UNIVERSITÄT ZU LÜBECK

From the Institute of Information Systems
of the University of Lübeck
Director: Prof. Dr. rer. nat. habil. Ralf Möller

Conceptual Orthospaces — An Embedding Framework Accounting for Negation Operators and Convexity Constraints

Dissertation
for Fullfillment of
Requirements
for the Doctoral Degree
of the University of Lübeck

from the Department of Computer Sciences and Technical Engineering

Submitted by
Mena Leemhuis
from Aurich (Ostfriesland)

Lübeck 2023

First referee	PD. Dr. Özgür L. Özçep
Second referee	Prof. Dr. Laura Giordano
Date of oral examination	18.04.2024

Abstract

Neural networks and in general subsymbolic learning approaches perform well on usual learning tasks, but they are black boxes lacking desired properties such as explainability or trustworthiness. Fully integrated symbolic/subsymbolic systems account for these properties, e.g., through the injection of symbolic information into a subsymbolic learning framework. Systems relying on embeddings can be considered as specific examples of such systems where facts, expressed in some logic, are embedded into a continuous space. Instances mentioned in the facts are mapped to points and the concepts mentioned in the facts are mapped to regions in that space. This allows to exploit geometric regularities in order to tackle typical learning tasks such as link prediction. Embeddings provide the first step towards filling the gap between qualitative, Tarskian style semantics which is used for deductive reasoning over the facts, and quantitative structures which are used for representing objects, relations, and concepts for learning purposes. However, to enable meaningful reasoning, embeddings of relations and concepts are not allowed to be shaped arbitrarily. Especially, convex sets turned out to be appropriate due to their computational advantages and due to their foundation in cognition. Convexity can be defined with a ternary betweenness relation and concepts can be defined as betweenness-closed (= convex) sets. Though many interesting phenomena of cognitive reasoning can be explained in such a framework, it is at least not obvious how to use betweenness for other, more logico-formal aspects of reasoning that, e.g., require defining logical operators. However, convexity is not the only relevant notion for embedding approaches. Next to it, also similarity is a notion vital for inference and allowing for modeling negation in form of dissimilarity. This leads to the possibility of modeling negation via the orthoframes of Goldblatt (1974). This dissertation tackles the question of how these two fundamental aspects of embedding approaches can be combined.

In this dissertation, I provide results for connecting betweenness and similarity, thus convexity and orthoframes. In particular, I investigate the construction of a space of convex concepts equipped with an orthogonality relation. I give a universal construction of a betweenness relation over an orthoframe and show that based on Euclidean betweenness (thus using the classical notion of convexity) and under some natural restrictions, convex cones (and a relaxation thereof) are the only structures capable of modeling such a space. Next to the propositional case, also the extension to modeling relations in an expressive way is considered, on the one hand based on reification allowing for modeling expressive partial relations, on the other hand focused on modeling an arbitrary quantifier depth. For increasing explainability and trustworthiness, it is not only necessary to have an interpretable approach but also to be able to determine its expressivity. Therefore, logical commitments are introduced as a way to determine the expressivity of embedding approaches.

Kurzfassung

Neuronale Netze und subsymbolische Lernverfahren im Allgemeinen liefern gute Ergebnisse für klassische Lernaufgaben, ihnen mangelt es aber aufgrund ihres Black-Box-Verhaltens an Erklärbarkeit und Vertrauenswürdigkeit. Systeme, die symbolische und subsymbolische Verfahren miteinander kombinieren, erlangen diese Eigenschaften, z.B. durch das Hinzufügen symbolischer Informationen in ein subsymbolisches Lernverfahren. Systeme, die auf Einbettungen basieren, sind ein Beispiel hierfür: Fakten, dargestellt mithilfe einer Logik, werden in einen kontinuierlichen Raum eingebettet. Instanzen, die in diesen Fakten genannt werden, werden auf Punkte in diesem Raum abgebildet, die genannten Konzepte auf eine Region in diesem Raum. Das erlaubt, geometrisches Verhalten zu nutzen, um typische Lernaufgaben wie die Vorhersage von neuen Tripeln zu erfüllen. Einbettungen sind der erste Schritt, um die Lücke zwischen qualitativer Semantik im Stil von Tarski, die zum deduktiven Schließen über Fakten genutzt wird, und quantitativen Strukturen, die genutzt werden, um Objekte, Relationen und Konzepte zum Lernen darzustellen, zu schließen. Aber um ein aussagekräftiges Schließen zu ermöglichen, dürfen die Konzepte nicht beliebig geformt sein. Konvexe Mengen haben sich als besonders nützlich herausgestellt, einerseits aufgrund ihrer Vorteile in der Rechenkomplexität, andererseits durch ihre Fundierung in der Kognition. Konvexität kann mithilfe einer ternären Dazwischen-Relation definiert werden, sodass Konzepte als Mengen dargestellt werden, die abgeschlossen gegenüber der Dazwischen-Relation und somit konvex sind. Obwohl so ein System die Erklärung vieler interessanter Phänomene des kognitiven Schließens ermöglicht, ist es nicht offensichtlich, wie man diese Dazwischen-Relation für andere, formal-logischere Aspekte des Schließens nutzen kann, z.B. solche, die die Definition logischer Operatoren notwendig machen. Konvexität ist nicht das einzige relevante Konzept für Einbettungen, auch Ähnlichkeit ist wichtig für Herleitungen und erlaubt es, Negation als Unähnlichkeit zu modellieren. Das ermöglicht es, die Negation mithilfe der Orthorahmen von Goldblatt (1974) zu modellieren. Diese Dissertation beschäftigt sich mit der Frage, wie diese beiden grundlegenden Aspekte der Einbettung kombiniert werden können.

In dieser Dissertation stelle ich Ergebnisse zur Verbindung zwischen der Dazwischen-Relation und der Ähnlichkeit vor, d.h. zwischen Konvexität und Orthorahmen. Insbesondere untersuche ich die Konstruktion konvexer Konzepte, die mit einer Orthogonalitätsrelation ausgestattet sind. Ich gebe eine universelle Konstruktion einer Dazwischen-Relation über einem Orthorahmen und zeige, dass, basierend auf einer euklidischen Dazwischen-Relation (d.h. unter Nutzung der klassischen Interpretation der Konvexität) und bei Erfüllung einiger kleinerer Einschränkungen, konvexe Kegel (und eine Verallgemeinerung dieser) die einzigen Strukturen sind, die so eine Modellierung erlauben. Zusätzlich zu dem propositionalen Fall habe ich auch die aussagekäftige Modellierung von Rollen betrachtet, einerseits basierend auf Reifikation um partielle Relationen abzubilden, und andererseits mit dem Fokus der Modellierung einer beliebigen Quantorentiefe. Um Erklärbarkeit und Vertrauenswürdigkeit zu erhöhen, ist es nicht nur wichtig, einen interpretierbaren Ansatz zu haben, sondern auch die Möglichkeit zu haben, die Aussagekraft des Verfahrens zu bestimmen. Ich habe den Begriff der logischen Garantien (englisch: logical commitments) eingeführt, um diese Aussagekraft formal zu bestimmen.

Acknowledgments

Writing my dissertation has been a long process that started with me attending a master's course on "Foundations of ontologies and databases for information systems". This lecture piqued my (from there on ever increasing) interest in "the world behind the scenes", the underlying logic and structure of established approaches, and resulted in me doing an internship and ultimately writing my master's thesis under the supervision of Özgür on the topic of embeddings and cones. However, after the successful completion of the masters' degree, the topic continued to fascinate me, so I had no choice but to write my dissertation on this topic as well. Throughout this period, the cones evolved further and this opened new possibilities, from al-cones to convex cones to pseudo-cones and, as is part of the title of this dissertation, to Conceptual Orthospaces.

First and foremost, I would like to express my special thanks to Özgür, whose lecture on "foundations" was the foundation for the enthusiasm this topic prompts in me to this day, and who accompanied me from internship to master and dissertation with many helpful discussions and advice. I would like to also thank Angela for her help with seemingly insurmountable bureaucratic obstacles compared to which working with the conceptual orthospaces was pure relaxation. On the technical side, I owe a thank you to LaTeX for allowing me to spend half the night working on my topic (and not on formatting footnotes). When I was stuck in my work, the discussion with my colleagues at lunch helped me get a mental break from it, as distraction is often the best help. In particular, the "word of the day" was always helpful, especially "da machste nichts dran"(nothing can be done about that). Last, but not least, I am very thankful to my parents and to Simon for their wholehearted support and help during this time.

Mena Leemhuis
Lübeck, November 2023

Contents

List of Abbreviations

abox assertional box.

COS Conceptual Orthospace.

DL Description Logic.

KGE Knowledge Graph Embedding.

Ne-Sy AI Neuro-Symbolic AI.

pl projective law.

SVM Support Vector Machine.

SVR Support Vector Regressor.

tbox terminological box.

XAI explainable AI.

ZSL Zero-Shot Learning.

List of Symbols

$\angle$ Angle between points in the space.

Pow Powerset.

GL_n General linear group of degree n.

$conv$ Convexity in the sense of betweenness-closure.

$conH$ Convex hull (based on Euclidean convexity).

cl Closure operator.

$cl_\perp$ Closure operator based on an orthogonality relation.

$\mathcal{C}^n$ Set of convex sets in the n-dimensional space.

$\mathcal{C}^n_0$ Set of closed convex sets in the n-dimensional space where each set contains $\{0\}$ at boundary or interior.

$\mathcal{C}^n_{cone}$ Set of closed convex cones in the n-dimensional space.

$\mathcal{PC}^n$ Set of closed pseudo-cones in the n-dimensional space.

$\mathcal{PPC}^n$ Set of closed relaxed pseudo-cones in the n-dimensional space.

$\mathcal{ALC}$ General description logic incorporating full negation.

$\mathcal{EL}^{++}$ Restricted description logic.

$\mathcal{EL}$ Restricted description logic.

$\overline{0}$ Bottom concept in DL.

$\overline{1}$ Top concept in DL.

$\sqcap$ Conjunction in DL.

$\sqcup$ Union in DL.

$\bigsqcap$ Infinitary conjunction in DL.

$\bigsqcup$ Infinitary disjunction in DL.

$\neg$ Negation in DL.

$\sqsubseteq$ Subsumption in DL.

$\equiv$ Equality in DL.

$\Delta^{\mathcal{I}}$ Domain of Interpretation in DL.

$\mathcal{I}$ Interpretation in DL.

$\mathcal{T}$ Tbox.

$\mathcal{A}$ Abox.

$\mathcal{O}$ Ontology.

N_c Set of constants in DL-vocabulary.

N_C Set of concept symbols in DL-vocabulary.

N_R Set of role symbols in DL-vocabulary.

qr Quantifier rank.

srank Semantic quantifier rank.

$\mathbf{srank}_c$ Cyclic semantic rank.

srank' Cyclic semantic rank, part 1.

$\mathbf{srank}_{\mathcal{T}}$ Tbox-specific semantic rank.

$\leq$ Partial order relation.

$\wedge$ Lattice-meet.

$\vee$ Lattice-join.

$<:$ Covering in lattices.

$\mathbb{1}$ largest lattice element.

$\mathbb{0}$ smallest lattice element.

$'$ orthocomplement.

Omin minimal orthologic.

& Conjunction.

$\bigvee$ Disjunction.

$\neg$ Negation.

$-$ Galois negation.

$\sim$ Similarity relation.

$\perp$ Orthogonality relation.

$\perp_g$ Orthogonality relation based on Goodman's matching operator.

$\perp_p$ Orthogonality relation based on polarity.

$\perp_{ps}$ Orthogonality relation based on the polarity operator on pseudo-cones.

$\perp_m$ Orthogonality relation based on the Minkowski-distance.

$\perp_u$ Orthogonality relation based on the uniform distance.

$\circ$ Polarity operator.

X Domain of a conceptual orthospace.

Y Set of admissible sets in a conceptual orthospace.

$\mathcal{L}_Y$ Ortholattice based on Y.

d Distance metric.

$\mathcal{B}$ Betweenness relation.

$\mathcal{B}_g$ Induced betweenness relation by Goodman.

$\mathcal{B}_E$ Euclidean betweenness relation.

$\mathcal{B}_a$ Betweenness based on the angle.

$\mathcal{B}_\perp$ Orthonegation-induced betweenness.

$\mathcal{B}'_\perp$ Orthonegation-induced betweenness based on an arbitrary orthogonality relation.

1 Introduction

After a motivation in Section 1.1, related work is presented in Section 1.2. The chapter ends with a presentation of the contributions in Section 1.3 and an outline of the structure of the dissertation in Section 1.4.

1.1 Motivation

Subsymbolic[1] learning approaches have gained importance in the last decades, in particular with the increasing relevance of neural networks. In the last years, however, awareness has arisen that not only the result quality in terms of quantitative criteria, such as precision and recall, is of importance. Instead, in order to garner further acceptance for subsymbolic learning approaches, it is important to increase the explainability and trustworthiness of the systems [Garcez and Lamb, 2023]. The aim of this dissertation is to shine a light on the black box of subsymbolic learning approaches.

There are different ways of thought to increase the explainability of a system: Post-hoc explainability of deep learning systems considers the explanation as one step at the end of the prediction process by determining the reasons for the classification result [Arrieta et al., 2020]. An example is class activation mapping [Zhou et al., 2016], highlighting the part of the input-image responsible for a specific output. Though such approaches lead to an increase in explainability, they still do not change the black box behavior, as they try to disclose the secrets of the black box only after it was trained. This leads to the idea of pre-explainability of a system in the sense of interpretability, namely changing the structure of the learning approach and thus circumventing the black box behavior

[1] Here, the widely used term *subsymbolic* should emphasize the contrast to symbolic methods and thus should depict the learning approaches based on instance-based information, e.g., similarity of instances. The term subsymbolic was introduced by Smolensky [1988]. Beside this term, also the distinction between *subconceptual* and *conceptual* and the terms *associatist* and *connectionist* are used (see, e.g., the work of Gärdenfors [2000, pp. 40ff.]). These terms have slightly different meanings. As the emphasis here is to clearly show the contrast to the symbolic paradigm, in this dissertation, the term *subsymbolic* is used as a general term, denoting approaches not classically classifiable as symbolic approaches.

directly in the design phase [Arrieta et al., 2020]. Thus, the aim is to design controlled, well-founded systems that follow mathematically and logically grounded principles.

An example for systems solely based on such principles are symbolic approaches. They rely on symbol manipulation based on rules and are thus more explainable and trustworthier than subsymbolic approaches. However, they have problems that subsymbolic approaches such as neural networks are able to circumvent. The most important problems are the *symbol-grounding problem* [Harnad, 1990], the impossibility of deriving the meaning/semantics of a symbol, and the *frame problem* [McCarthy and Hayes, 1969], the issue of choosing the relevant rules for the derivation of a desired result.

Therefore, the aim is to combine these two mindsets in order to get the best of both worlds. Fields of study dedicated to tackling this are *Hybrid-AI* [Marcus, 2020] and *Neuro-Symbolic AI (Ne-Sy AI)* [Garcez and Lamb, 2023]. Their aim is to create a combined symbolic-subsymbolic learning approach, on the one hand allowing for a sufficiently structured logically grounded system, on the other hand incorporating the advantages of subsymbolic approaches such as inconsistency-tolerance and groundings of symbols. These areas give only a broad context for the approaches and do not hint at how this combination should be achieved. There are several methods, ranging from creating an interface between a symbolic and a subsymbolic level to defining a complete, coherent approach unifying both paradigms. One especially successful idea for tackling the problem of combining subsymbolic and symbolic information is to model the subsymbolic information (thus instances and their similarities) geometrically as points in some space, e.g., but not limited to, a vector space, and the symbolic information by interpreting the symbols (concepts) as regions containing these points. This approach can be cognitively motivated, as has been done, e.g., by Gärdenfors in his work on *conceptual spaces* [Gärdenfors, 2000].

Geometric representations turned out to be useful also in practice, e.g., in the area of *Knowledge Graph Embedding (KGE)*. In KGE, knowledge graphs, i.e., sets of (subject, predicate, object)-triples, e.g., (“alice”,“loves”,“bob”), are embedded into some low-dimensional vector space (for an overview of KGE, see for example the survey of Q. Wang et al. [2017]). In this context, approaches such as *TransE* [Bordes et al., 2013] allow for link prediction with an accuracy that was impossible to achieve previously. These approaches benefit from the ability to model similarity of objects geometrically and the ability to do link prediction based on this similarity — and not only based on known rules about the relations. Whereas TransE and related approaches cannot be considered an instance of Ne-Sy AI, as they do not incorporate any symbolic information, the pioneering task of TransE paved the way for enhancements of KGE, extending learning with some form of additional information. This information ranges from using natural language descriptions to considering role paths [Q. Wang et al., 2017]. One possibility, hinting towards symbolic approaches, is to incorporate background information in form of axioms about concepts and their relations geometrically in the embedding. Concept subsumption can then be modeled as subset-relation and logical operators such as conjunction with the help of geometrical operators — as cognitively motivated by Gärdenfors [2000]. This allows for interpretability in the sense that a result is not only based on similarity to other results but also on background knowledge. Accordingly, such an approach increases the real-world plausibility. This is in contrast to the symbolic approach — while a Ne-Sy AI-approach does not provide a full explanation, it does give a strong indication in the direction of correct and thus trustworthy results, explainable based on background knowledge while circumventing the drawbacks of a purely symbolic approach. In an approach not incorporating background knowledge it would be, e.g., possible to state that an object is both a cat and a dog, whereas when the background knowledge that cats are disjoint from dogs is incorporated, this is not possible. This is not only useful in such simple examples but also for real-world problems, e.g., in a medical context, in which background information is required to ensure that two prescribed drugs are not contraindicative.

The aim, however, is not only to enable modeling structured information in some form but to enable modeling a background logic of sufficient strength, to avoid biases in the learning approach if, e.g., specific background information axioms cannot be modeled based on the given geometric structure. This is particularly true for the representation of negation and disjunction. However, there is a need for modeling negation in embedding approaches. Negation in learning approaches is a topic that is becoming increasingly popular, especially in learning for language models. Whereas in the past the word "not" was considered a stop-word and removed before training, it is now considered a vital part of the language model and its treatment is the subject of many research projects (see, for example, the survey of Morante and Blanco [2020]). It is also gaining importance in other areas, e.g., in KGE, as more and more KGE-approaches are able to model some (weak) form of negation (e.g., by Garg et al. [2019], Kulmanov et al. [2019] and Zhang et al. [2021]). Not only positive information but also information stating that something is not the case is of importance. One prominent example is again from the medical realm in which it is not only necessary to be able to model that a patient has a specific disease but also to model that a patient definitely does not have this disease. One easily modeled type of negation is negation as set-complement, where "e is not C" is true if it cannot be proven that "e is C" is true. This is not appropriate, as a proper model of negation is needed that enables expressing the difference between remaining silent on the question of whether entity e belongs to a concept C and stating that e belongs to the negation of C. Additionally, when considering this type of negation in the context of embeddings, the negated concept is mostly less structured than the concept itself, e.g., it is possible that the negation as set-complement of a convex representation of a concept is not convex anymore. At this point, the first technical question tackled in this dissertation arises:

Question 1. *What should a negation suitable for expressive embedding approaches look like?*

The area of negation is a widely researched topic and many different types of negation for many different use cases were proposed, both on philosophical and logical grounds (for an overview see, e.g., the book of Horn [2001]) but not all of them are both of sufficient strength and sufficiently well representable in a geometric setting.

The definition of a suitable negation is not the only thing to consider for a good learning approach. It is not only necessary to model logical operators but also to determine their strength and know their expressivity. This is in particular important for increasing the explainability and interpretability of a system. These attributes, together with others, make up, what I call, the *logical commitments* of an embedding approach. Logical commitments of a learning approach are its expressivity on different levels (not only the expressivity of logical operators), e.g., which relations are representable, which logic can be modeled, the question on the used inference service, i.e., whether a result is achieved via induction or deduction etc. The knowledge concerning the commitment of a learning approach is vital for a good result, as it is necessary to recognize biases resulting from a lack of representability. If a dataset contains a lot of transitive tuples but the embedding approach is not able to embed relations allowing for transitivity, then the approach is not adequate, independent of the results of quantitative quality criteria. This enforces the necessity to exactly determine in advance which logical commitments are needed for a given task and to determine which learning approach is able to fulfill these commitments and is thus able to model the given data in a sufficient way to enable for correct, explainable and bias-free learning. This leads to the second question handled in this dissertation:

Question 2. *What are the logical commitments of learning approaches and how are they determined?*

Here, the main observation towards answering this question is that the aim is not to model a perfect representation of reality but to find an approach of sufficient strength to model the desired task as simply as possible to allow for efficient learning and reasoning. Choosing an embedding approach based on the desired logical commitment sufficient for the given task changes the way of defining an embedding approach. Whereas actual embedding approaches are mainly based on the choice of some geometric object with appropriate properties as a basis and a determination of its logical commitments only afterwards (if it is determined at all), the aim should be to define the geometric structure based on some low-level attributes of the data and thus based on the data-oriented information need. But what are the main ingredients of embedding approaches and, thus, their foundational attributes? What attributes are at most necessary to model a sufficient embedding approach?

One basic ingredient is the *similarity* of objects, as it is the foundation of subsymbolic learning approaches and thus a vital part of a Ne-Sy AI-approach. It is not only important to gather information on the similarity of instances but also to determine beforehand which type of similarity should be used. Whereas the classical similarity is distance-based (e.g., based on the Euclidean distance), this is not the only one and especially not the best fitting one for all circumstances and all datasets [Tversky, 1977].

Though similarity is a vital notion for subsymbolic approaches, it is not sufficient to rely on similarity alone, in particular, when concepts and their interplay are considered. It is necessary to analyze the structure of the concepts and to restrict their structure to expressive representations. This can be done with the help of *betweenness*. Betweenness between three elements a, b and c states that b is in between a and c. This thus allows for some notion of order of the elements, e.g., by enforcing that if two elements are in a concept, then the element in between these elements is also in this concept. This leads to a more natural representation of concepts by enforcing convexity (betweenness-closure).

Convexity is vital for embedding approaches, as, based on the principles of cognitive economy [Gärdenfors, 2000, p. 70], a simple structure increases learning and reasoning capabilities. This is underlined by practical aspects such as the computational advantages of convex optimization and also the fact that actual embedding approaches mostly rely on convex regions, e.g., spheres [Kulmanov et al., 2019], boxes [Abboud et al., 2020] or arbitrary convex structures [Gutiérrez-Basulto and Schockaert, 2018]. This is not restricted to the standard Euclidean betweenness in terms of points on a line (just like the similarity is not restricted to be based on the Euclidean distance) but allows for betweenness (resp. similarity) of different strengths, depending on the use case.

This directly leads to the third question:

Question 3. *How does a general framework based on similarity and betweenness look like that can serve as basis for embedding approaches of different strengths?*

To get an embedding approach based on such a framework, some additional steps are necessary, e.g., defining a training objective or architecture. However, such a framework can serve as a basis, giving, based on some basic notions of similarity and betweenness, a possibility for modeling an expressive embedding. Using a geometrical structure as a basis for the creation of an embedding approach, in contrast, leads to a more restricted approach. This construction also helps in finding adequate structures for modeling negation, whereas many geometric approaches are unable to do so. Thus, starting with an arbitrary concept shape lowers the chance to model a sufficient negation based on this shape. Starting on the level of similarity and betweenness and using a general framework accordingly leads to an approach incorporating negation.

To revisit the statement of the beginning: The black box can be enlightened by, first, introducing symbolic information into the learning approach, second, determining the logical commitment of this approach and, third, defining the approach based on low-level properties of the underlying data.

Thus, to combine the three questions, the main question answered in this work is:

Question 4. *How does a general framework for embedding look like, able to model an expressive background logic and being as adaptable to specific needs as possible?*

These three questions lead to the main hypothesis which will be proven in the course of the dissertation:

Hypothesis. *It is possible to define a framework based on information about similarity and betweenness modeling logical operators, especially negation, geometrically.*

1.2 Related Work

Ne-Sy AI is a widely researched topic that led to many successful learning approaches. The focus in this section lies on embedding methods and in general on geometric learning. There are many KGE-approaches that incorporate background knowledge in many different forms and thus can be considered as Ne-Sy AI-approach. One variant is to model concepts directly as geometric objects, e.g., spheres [Kulmanov et al., 2019], boxes [H. Ren et al., 2020; Jackermeier et al., 2023; Xiong et al., 2022; Abboud et al., 2020; Peng et al., 2022], cones [Bai et al., 2021; Zhang et al., 2021; Özçep et al., 2020], subspaces of a Hilbert space [Garg et al., 2019], general convex regions [Gutiérrez-Basulto and Schockaert, 2018] and others. From these approaches, only those of Garg et al. [2019] and Özçep et al. [2020] are able to model negation sufficiently. Others, such as that of Zhang et al. [2021] can model negation but only in its weaker form of negation as set-complement. However, the approaches also differ in their way of handling logical operations in general. Whereas, e.g., Garg et al. [2019] handle logical operations as geometrical operations being able to model the properties of the logical operations exactly (e.g., an object x is in A and B iff it is in A and in B), this is not the case in all approaches. In contrast, in the approach of Zhang et al. [2021] conjunction is learned based on a neural network and thus models only approximately the correct behavior. Accordingly, the logical commitments are not exactly determinable.

Other approaches of KGE are, e.g., TransE [Bordes et al., 2013] or RotatE [Sun et al., 2019] which are simple models, embedding instances as points and relations (thus predicates) as translation resp. rotation. They do not incorporate any background information and, thus, are not able to model concepts sufficiently. Besides the approaches only modeling instances and their relations, there are several approaches incorporating some sort of additional information. This information ranges from textual descriptions to information on the concept membership and rules [Q. Wang et al., 2017]. Here, the focus lies on embedding background knowledge in form of concepts of a specific shape and modeling logical operators as geometric operators. An example for an approach, different from the explicit geometric approaches mentioned above, is, e.g., *semantically smooth embedding (SSE)* [Guo et al., 2015], embedding entities belonging to the same concept close to each other. Due to its simplicity it is not able to model subconcept relations and does not allow for an instance having several concepts. Other approaches, such as that of Xie et al. [2016], handle these restrictions but do not model the concepts as intersecting regions. Instead, for an object having multiple concepts, an embedding for each concept is learned separately. Then for link prediction only the embedding of the concept relevant for the considered role is used.

The usage of geometric models is not restricted to KGE, as it can be also used, e.g., in the context of multi-label learning. There, general approaches such as Support Vector Machines (SVMs) can be considered as modeling concepts as geometric regions. However, SVMs are mostly not able to model logical operators geometrically. Considering, e.g., two simple SVMs without a kernel having learned a discrimination between concept-membership and non-membership with one SVM for each concept. Then, the area depicting concept membership is a half space each, however, the conjunction of these two concepts is mostly not a half space but a conjunction of half spaces and thus not representable as a learned concept.

Geometric models are also used in their more abstract form in conceptual spaces [Gärdenfors, 2000]. The theory of conceptual spaces is focused on the information processing in the human brain and proposes that information is modeled geometrically in the brain by representing concepts as regions and instances as points in a space. In contrast to the embedding approach mentioned above, the focus here lies not on learnability but on a semantically meaningful representation of the knowledge in an embedding, thus, e.g., considerations on domains and meaningful dimensions. Conjunction should be, e.g., modeled not as set-intersection but as complex interplay of information of different domains [Gärdenfors, 2000, pp. 114 ff.]. Hereby, negation is described as being particularly difficult [Gärdenfors, 2014, p. 238]. In contrast, the focus in this dissertation does not lie on a representation having an exact semantics but on a representation enabling for learning and reasoning.

Though the above-mentioned embedding approaches lead to a more or less good result quality, they are, with two exceptions, not able to model a sufficient negation. These two approaches are based on a specific definition of the shape of the concept, namely subspaces of a Hilbert space and convex cones. Accordingly, the aim is to design a more general approach based on similarity and betweenness. Though there are some approaches that take into consideration the connection between similarity and betweenness, e.g., Sheremet et al. [2007] propose a description logic incorporating similarity information and Ibáñez-García et al. [2020] consider similarity and betweenness for concept interpolation, none of these approaches consider negation.

Based on these works, I identified the need for the definition of a general framework allowing for modeling expressive logical operators geometrically. Therefore, I introduced a framework combining betweenness and similarity as the basis for embedding approaches [Leemhuis and Özçep, 2023]. The case of convex cones as a special case has been presented by Özçep et al. [2020] for the propositional case (including an extension for an approximated representation of roles) and I extended it for the incorporation of relations [Leemhuis, Özçep, and Wolter, 2022a]. I also presented the application of the framework (again for the special case of cones) [Leemhuis, Özçep, and Wolter, 2020; Leemhuis, Özçep, and Wolter, 2022b].

1.3 Contributions

The aim of this dissertation is to define an abstract embedding framework based on betweenness and similarity, able to model negation and acting as a basis for embedding approaches. The contributions are as follows:

(1) Identifying orthonegation as suitable negation The modeling abilities of the proposed framework are highly based on the underlying definition of negation. I propose to use the orthogonality relation, as it is a symmetric and irreflexive relation interpretable as incompatibility and is thus strongly connected to the similarity being the basis of subsymbolic learning approaches. Next to the negation, also the type of conjunction and disjunction modeled is discussed. It turns out that such a modeling is both logically and psychologically justifiable.

(2) Betweenness and similarity as basis for an embedding approach I identify two basic notions, vital for embedding approaches, namely similarity and betweenness. This allows defining the negation based on dissimilarity and leads to an expressive and well-structured embedding with, due to convexity (betweenness-closure), good computational behavior.

(3) Conceptual Orthospaces Based on the orthogonality relation, betweenness and similarity, I create the framework of Conceptual Orthospaces (COSs). This framework is situated in the area of Ne-Sy AI and extends existing approaches by (i) enabling modeling negation, (ii) giving an unified underlying framework and, due to its high expressivity, enabling to act as basis for many different embedding approaches of different strength and (iii) allowing for theoretical considerations on the possible expressivity of embedding approaches and their limits.

(3a) Orthonegation-induced betweenness I propose several options for creating a COS. One option is to have an orthogonality relation given and to induce a betweenness relation. I examine two variants, enabling, on the one hand, to induce a strong betweenness relation based on a restricted orthogonality relation and, on the other hand, to induce a betweenness for an arbitrary orthogonality relation — to the cost of having a weaker resulting betweenness. Thus, I show that a COS can be created given an orthogonality relation.

(3b) Betweenness-induced orthogonality It is also possible to induce an orthogonality relation based on a given betweenness relation. This is done exemplarily for the case of Euclidean betweenness.

(4) COS based on Euclidean betweenness I consider Euclidean betweenness and show that (under some natural restrictions), the resulting COS based on Euclidean betweenness is not arbitrary, in fact, it is restricted to closed convex cones and a relaxation thereof and the orthogonality needs to be defined as polarity. Thus, it is possible to use COSs based on Euclidean betweenness — and thus possible to model an expressive negation in an embedding situated in the Euclidean space. However, it turns out that cones are vital for such constructions. This is underlined also when considering roles, as a cone-based COS enables to model roles expressively.

(5a) COS with roles COSs are not only usable in the propositional case but can also be extended with roles. For the case of a cone-based COS, I propose an interpretation of roles based on reification, enabling to model an expressive embedding, allowing, e.g., for partial interpretations and in particular for the interpretation of all knowledge bases in full $\mathcal{ALC}$.

(5b) COS for faithful embeddings Next to the reification-based approach for roles in a cone-based COS, I consider a different approach, also based on cones, enabling modeling ontologies with arbitrary quantifier rank faithfully and thus allowing for an high expressivity of the resulting COS.

(6) COS as basis for embedding approaches COS are not only usable on theoretical grounds but can also be brought to practice. I demonstrate the usability of COS-based embedding approaches with the help of two use-cases: Using COSs for Zero-Shot-Learning increases the interpretability and trustworthiness compared to a reference approach. Additionally, I underline the significance of the proposed COS by showing that an existing embedding approach is a special case.

(7) Logical commitments of learning approaches COSs increase the interpretability of learning approaches. However, to follow the lines of Ne-Sy AI, it is also important to consider the post-hoc explainability. The explainability has been considered only for subareas of embedding approaches, e.g., regarding properties of roles such as transitivity or reflexivity before. I propose the consideration of *logical commitments* allowing for a thorough view on the post-hoc explainability of embedding approaches, ranging from considerations on concept representations to the expressivity of the modeled logic to the inference service used.

1.4 Structure of the Dissertation

After the introduction, the dissertation starts with laying out the foundation for COSs. First, in Chapter 2 the general problem is stated, namely enhancing embedding approaches with background information. Therefore, ontologies are introduced as a form of modeling this background information and the combination of symbolic and subsymbolic information in general is discussed. Before coming to actual embeddings, I consider in Chapter 3 in detail which logical operators should be represented. Thereafter, the main part of this dissertation is divided into two parts:

- Part I introduces COSs as a framework for expressive embeddings.
  - Chapter 4 identifies the vital parts of embedding approaches for concept formation, namely betweenness and similarity.

- Chapter 5 gives a definition of COSs.
- Chapter 6 considers the construction of COSs based on a given similarity resp. betweenness relation

These three chapters are an extension of the publications

> Mena Leemhuis and Özgür L. Özçep (2022). "A Goodman-style Betweenness Relation on Orthoframes". In: *Proceedings of the 8th Workshop on Formal and Cognitive Reasoning, co-located with the 45th German Conference on Artificial Intelligence (KI 2022)*. Ed. by Christoph Beierle, Marco Ragni, Kai Sauerwald, Frieder Stolzenburg, and Matthias Thimm

> Mena Leemhuis and Özgür L. Özçep (Nov. 2023). "Conceptual Orthospaces—Convexity Meets Negation". In: *International Journal of Approximate Reasoning* 162, p. 109013. DOI: `10.1016/j.ijar.2023.109013`

- Part II considers applications of these COSs.
  - Chapter 7 considers the special case of a COS based on Euclidean betweenness. This is an extension of parts of

    > Mena Leemhuis and Özgür L. Özçep (Nov. 2023). "Conceptual Orthospaces—Convexity Meets Negation". In: *International Journal of Approximate Reasoning* 162, p. 109013. DOI: `10.1016/j.ijar.2023.109013`

  - Chapter 8 regards the extension of COS with roles. This chapter is based on the following two publications, whereas from the second one only one independent part was used.

    > Mena Leemhuis, Özgür L. Özçep, and Diedrich Wolter (2022a). "Knowledge Graph Embeddings with Ontologies: Reification for Representing Arbitrary Relations". In: *German Conference on Artificial Intelligence (Künstliche Intelligenz)*. Ed. by R. Bergmann, L. Malburg, S.C. Rodermund, and I.J. Timm. Lecture Notes in Computer Science 13404. Springer International Publishing, pp. 146–159. DOI: `10.1007/978-3-031-15791-2_13`

    > Özgür L. Özçep, Mena Leemhuis, and Diedrich Wolter (Oct. 2023). "Embedding Ontologies in the description logic ALC by Axis-Aligned

Cones”. In: *Journal of Artificial Intelligence Research* 78, pp. 217–267. DOI: 10.1613/jair.1.13939

- Chapter 9 showcases the use of COSs in embedding approaches. Parts of this chapter are based on the publication

  Mena Leemhuis, Özgür L. Özçep, and Diedrich Wolter (Oct. 2022b). “Learning with cone-based geometric models and orthologics”. In: *Annals of Mathematics and Artificial Intelligence* 90, pp. 1–37. DOI: 10.1007/s10472-022-09806-1

- Chapter 10 considers the logical commitment of embedding approaches in general, with a special focus on COSs.

The dissertation ends with a conclusion and a discussion of future work in Chapter 11.

2 Enhancing Learning Approaches with Background Knowledge

Multi-label learning as a classical learning approach aims at the determination of (possibly multiple) classes of a new instance based on training data in the form of instances and their classes. Subsymbolic solution strategies are based on, e.g., a collection of binary classifiers or on straightforward neural networks. They either do not incorporate the interrelation between classes or they incorporate it only implicitly based on correlation [Gibaja and Ventura, 2014]. Using a symbolic approach for solving such a problem can lead to problems when handling inconsistent data and to high complexity of the determination of the result. The same problem occurs in KGE where in standard subsymbolic approaches such as TransE [Bordes et al., 2013] the concept membership of instances and rules restricting the relations are not considered, whereas in symbolic approaches on knowledge graph completion, the huge size and low quality of knowledge graphs can lead to problems [Hogan et al., 2021].

There are many ideas combining subsymbolic and symbolic approaches to utilize the advantages of both. As the main focus of my work lies in employing these ideas, in the following, the necessary technical basis is introduced. First, in Section 2.1, the representation of symbolic knowledge with the help of ontologies and Description Logics is considered. Ontologies serve as a suitable representation of symbolic knowledge as a basis for a combined symbolic/subsymbolic framework. After that, in Section 2.2, the advantages and disadvantages of both the symbolic and the subsymbolic paradigms are discussed and Ne-Sy AI is explained as a general framework for a symbolic/subsymbolic combination. There are different variants for such a combination, ranging from a loosely connected, hybrid approach to a fully integrated approach, the focus in the following lies on an integration based on embedding. This is explained in Section 2.3 and in the following chapters.

2.1 Ontologies and Description Logic

Symbolic knowledge comes in many different forms, ranging from automata to representations in predicate logic [Bader and Hitzler, 2005]. One way of representing structured knowledge is via an ontology. An ontology is defined as “a formal explicit specification of a shared conceptualization for a domain of interest” [Staab and Studer, 2009, p.vii] (based on work of T. R. Gruber [1993]). In general, an ontology consists of *concepts* representing groups of individuals (e.g., the concept “human”) and connections (*relations*) between concepts resp. individuals. The basis of an ontology are generalization/specialization hierarchies [Guarino et al., 2009]. To define these, some structured formal language is needed. One example for such a language are Description Logics which are considered in detail later on in this section.

Ontologies are made for quite different purposes. It is, e.g., possible to differentiate between *top-level-ontologies* giving an overview of a topic and *core-ontologies* representing the backbone of a domain [Guarino et al., 2009]. Ontologies are highly individualized (at least the mid-level and core-ontologies) meaning that the modeled concepts are dependent on the interest/focus (thus the *domain of discourse*) of the ontology engineer. In fact, often several ontologies for the same domain exist, incorporating different concepts and different relations between them. But the ontologies should also represent some shared view (at least shared by a subgroup) [Guarino et al., 2009].

The choice of a suitable ontology expressive enough for the desired problem is thus vital. Not all datasets/use cases need the most expressive underlying ontology. Instead, the used ontology is chosen based on the information need. Accordingly, the most complex ontology is not necessarily the best fitting one. This fact will become relevant later on, when discussing the framework of Conceptual Orthospaces. There, the aim is not to be able to model the most expressive ontology but to model the required strength.

The basic idea of ontologies is conceptualization, thus making up an abstract simplified view of the world based on concepts [Guarino et al., 2009]. There, intensional knowledge, thus knowledge being true in all possible worlds (in contrast to extensional facts, only true in one specific world) is represented. To specify the semantics of concepts in an ontology, axioms can be used.

One structured formal language tailored towards representing ontologies are Description Logics (DLs). They have a logic-based semantics and are especially useful for reasoning [Baader, Calvanese, et al., 2007, p. 47]. The basis of DL is the distinction between instance- and concept-level. The basic building blocks are atomic concepts (unary predicates), atomic roles (binary predicates) and individuals (constants). A DL-ontology is defined as a pair of terminological box (tbox) and assertional box (abox); some researchers use it in a narrower sense considering only the tbox as ontology. The abox $\mathcal{A}$ contains extensional knowledge in form of assertions about the actual world, e.g., $human(alice)$ stating that “alice” is an instance of the concept “human”. The tbox $\mathcal{T}$ contains intensional knowledge being true in every possible world, e.g., $woman \sqsubseteq human$ stating that women are humans. DLs support classification, both of concepts and of instances [Baader, Calvanese, et al., 2007, p.48].

The tbox- and abox-statements can be represented by formulas in a fragment of first order predicate logic [Baader, Calvanese, et al., 2007, p. 54]. DLs are also closely related to modal languages, as it is possible to interpret roles as accessibility relations [Baader, Calvanese, et al., 2007, p. 42].

In the following, the basic formalism is presented. There are different types of DLs with different complexity. Here, $\mathcal{ALC}$ [Baader, Calvanese, et al., 2007] is considered in more detail as it is a widely used DL allowing for modeling full negation which will be the focus later on.

The vocabulary of a DL is given by a set of constants N_c, a set of role names N_R and a set of concept names N_C. The $\mathcal{ALC}$ concepts (concept descriptions) over $N_C \cup N_R$ are described by the grammar

$$C \longrightarrow A \mid \bar{0} \mid \bar{1} \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \exists R.C \mid \forall R.C$$

where $A \in N_C$ is an atomic concept, $R \in N_R$ is a role symbol, and C is an arbitrary concept. As the focus lies in this dissertation on the DL-reading of KGE-problems, the triple (subject, R, object) is denoted as $R(subject, object)$ and thus in the following, the term *role* is used to describe it. For readability, $\exists R^{n+1}.C = \exists R.\exists R^n.C$, where $\exists R^1.C = \exists R.C$.

An $\mathcal{ALC}$ *interpretation* is a tuple $(\Delta^{\mathcal{I}}, (\cdot)^{\mathcal{I}})$ with a set $\Delta^{\mathcal{I}}$, called the *domain*, and an *interpretation function* $(\cdot)^{\mathcal{I}}$ which assigns to each constant an element of $\Delta^{\mathcal{I}}$, to each concept name a subset of $\Delta^{\mathcal{I}}$, and to each role name a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The extension of an interpretation function $\mathcal{I}$ onto arbitrary concept descriptions is given in Table 2.1. Though, this interpretation function is the standard one, it is not the only possible one as can be seen, e.g., for inconsistency-tolerant Description Logics [Odintsov and Wansing, 2003]. Details on the definition of a different interpretation function, especially regarding the negation, can be found in Sections 3.3 and 5.1.

As stated before, an ontology $\mathcal{O}$ is defined as combination of $\mathcal{T}$ and $\mathcal{A}$, thus $\mathcal{O} = (\mathcal{T}, \mathcal{A})$. A tbox $\mathcal{T}$ consists of *general concept inclusions (GCIs)* $C \sqsubseteq D$ (“C is subsumed by D”) for concept descriptions C, D. The equality $C \equiv D$ can be defined as an abbreviation of $\{C \sqsubseteq D, D \sqsubseteq C\}$. An abox $\mathcal{A}$ consists of a finite set of *assertions*, i.e., facts of the form $C(a)$ or of the form $R(a, b)$ for $a, b \in N_c$ and $R \in N_R$. An interpretation $\mathcal{I}$ *models* a GCI $C \sqsubseteq D$, for short $\mathcal{I} \models C \sqsubseteq D$, iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. An interpretation $\mathcal{I}$ *models*

Name	Syntax	Semantics
top	$\overline{1}$	$\Delta^{\mathcal{I}}$
bottom	$\overline{0}$	$\emptyset$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
universal quantifier	$\forall R.C$	$\{x \in \Delta^{\mathcal{I}} \mid$ For all $y \in \Delta^{\mathcal{I}}$: if $(x, y) \in R^{\mathcal{I}}$ then $y \in C^{\mathcal{I}}\}$
existential quantifier	$\exists R.C$	$\{x \in \Delta^{\mathcal{I}} \mid$ There is $y \in \Delta^{\mathcal{I}}$ s.t. $(x, y) \in R^{\mathcal{I}}$ and $y \in C^{\mathcal{I}}\}$

Table 2.1: Syntax and semantics for the DL $\mathcal{ALC}$

an abox axiom $C(a)$, for short $\mathcal{I} \models C(a)$, iff $a^{\mathcal{I}} \in C^{\mathcal{I}}$ and it models an abox axiom of the form $R(a, b)$ iff $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$. An interpretation is a *model of an ontology* $(\mathcal{T}, \mathcal{A})$ iff it models all axioms appearing in $\mathcal{T} \cup \mathcal{A}$. An ontology $\mathcal{O}$ *entails* a (tbox or abox) axiom ax, for short $\mathcal{O} \models ax$, iff all models of $\mathcal{O}$ are also models of ax. A concept C is the *most specific concept* (msc) for an abox-element a and an ontology $\mathcal{O}$, if $\mathcal{O} \vDash C(a)$ and for all concepts C', $\mathcal{O} \vDash C'(a)$ entails $\mathcal{O} \vDash C \sqsubseteq C'$. The most specific concept does not necessarily exist.

An equality is a *definition* if its left-hand site is an atomic symbol. Thus, definitions can be used to introduce names to complex descriptions [Baader, Calvanese, et al., 2007, p. 55]. A set of definitions is called a tbox if no name is defined more than once. A tbox is *definitorial* if every interpretation of the base symbols only appearing on the right-hand site of definitions has exactly one extension that is a model of the tbox.

In the rest of this dissertation, I understand the term "tbox" in the wider sense, allowing for including general concept inclusions. Only in Section 8.2, it is necessary to consider tboxes in a stronger sense. It is possible that the tbox contains a cycle, in the sense that either directly or indirectly some concept uses itself for its definition. Therefore, the notion of *circular relationship* is introduced, stating that a concept is part of a circular relationship. This can be defined as a slight adaptation of the definition of Baader, Calvanese, et al. [2007, p. 56]: For two atomic concepts A and B, A *directly uses* B if it appears on the right hand side of the definition of A. A concept is part of a circular relationship if it *uses* itself (where *uses* is the transitive closure of *directly uses*). If a tbox is acyclic, then it is definitorial, as it allows for iteratively expanding the definitions by replacing the concept symbols on the right hand site of a definition by the symbols it stands for [Baader, Calvanese, et al., 2007, p. 57]. However, this is not possible for cyclic tboxes, thus, cyclic tboxes are not definitorial.

A special case are inclusion axioms, as a terminology including inclusion axioms is not definitorial. These are not problematic, as they can be reformulated into their *normalization* — assuming the DL allows for equality and conjunction — by replacing $A \sqsubseteq C$ by $A \equiv C \sqcap \overline{A}$ with a new base symbol $\overline{A}$ [Baader, Calvanese, et al., 2007, p. 64].

The *quantifier rank qr* for arbitrary concepts C depicts the maximal nesting depth of quantifiers in it. Formally, for an $\mathcal{ALC}$ concept the quantifier rank is defined by recursion as $qr(A) = 0$, $qr(\neg C) = qr(C)$, $qr(C_1 \sqcap C_2) = qr(C_1 \sqcup C_2) = \max\{qr(C_1), qr(C_2)\}$ and $qr(\exists R.C) = qr(\forall R.C) = qr(C) + 1$.

Basic, more restricted Description Logics are $\mathcal{EL}$, having the grammar

$$C \longrightarrow \overline{1} \mid A \mid C \sqcap C \mid \exists R.C,$$

and $\mathcal{EL}_{\overline{0}}$[1], extending $\mathcal{EL}$ with the $\overline{0}$-concept [Baader, Calvanese, et al., 2007, p. 535]. Another more expressive but still tractable DL is $\mathcal{EL}^{++}$, an extension of $\mathcal{EL}_{\overline{0}}$, allowing for nominals, a restricted form of concrete domains and a restricted form of role-value maps [Baader, Brandt, et al., 2005]. Though negation can not be modeled in $\mathcal{EL}^{++}$, it is at least possible to model disjointness via the bottom concept, thus two concepts C, D are disjoint iff $C \sqcap D = \overline{0}$ [Baader, Brandt, et al., 2005]. Additionally, it is possible to restrict the types of roles modeled, e.g., to transitive, reflexive or symmetric roles [Krötzsch et al., 2012]. A detailed introduction to Description Logics can be found in the book of Baader, Calvanese, et al. [2007].

We first consider theories based on propositional $\mathcal{ALC}$ (also called *Boolean* $\mathcal{ALC}$ in the following, in contrast to *full* $\mathcal{ALC}$ including roles), thus $\mathcal{ALC}$ without roles.

Each propositional tbox $\mathcal{T}$ generates a Boolean algebra (for details, see Appendix A), the so-called *Lindenbaum-Tarski algebra* [Tarski, 1935]. It is considered, as in order to apply the theory of ortholattices (see Chapter 3 and appendix A), the ontology need to be algebraized. The Lindenbaum-Tarski algebra can be defined for any theory in any logic. It is shown here how it can be defined for an $\mathcal{ALC}$ tbox $\mathcal{T}$. For concepts C, D let $\sim$ be the relation defined by $C \sim D$ iff $\mathcal{T} \models C \sqsubseteq D$ and $\mathcal{T} \models D \sqsubseteq C$. Relation $\sim$ is an equivalence relation inducing for each concept C an equivalence class $[C]$. Define operations $\wedge$, $\vee$, $\cdot'$ on the equivalence classes by setting $[C] \wedge [D] = [C \sqcap D]$, $[C] \vee [D] = [C \sqcup D]$ and $[C]' = [\neg C]$. As the relation $\sim$ is not only an equivalence relation but a congruence relation w.r.t. $\wedge$, $\vee$, $\cdot'$, the given equalities are indeed well-defined and one can show that the equivalence classes fulfill the axioms of a Boolean algebra. Just for completeness, we show here for the case of negation that indeed the logical operations on the equivalence classes are well-defined: $[C]'$ is defined by $[\neg C]$. Taken any other

[1] Originally, it is termed $\mathcal{EL}_\perp$, here it is termed $\mathcal{EL}_{\overline{0}}$ to avoid overloading of $\perp$.

representative D of the class $[C]$, i.e. $[D] = [C]$ or, equally, $C \sim D$. By definition this means $\mathcal{T} \models C \equiv D$. But then by definition of the entailment relation and the definition of $\neg$ for concepts we also have $\mathcal{T} \models \neg C \equiv \neg D$, in other words $[\neg C] = [\neg D]$. Hence, indeed $\cdot'$ on the equivalence classes is well defined.

In the following, also weaker logics than $\mathcal{ALC}$ are used to represent background knowledge. The strength of the DL and its interpretation is then clear out of context and (mostly) not explicitly stated.

2.2 Neuro-Symbolic Artificial Intelligence

For problem solving, e.g., link-prediction in KGE, two different paradigms have emerged: The symbolic and the subsymbolic paradigms.

Symbolic approaches are based on symbol manipulation and their aim is solely deductive. Prominent examples are logic programs [Gelfond and Leone, 2002]. A symbolic approach allows for exact and explainable reasoning. It does, however, have some drawbacks, such as the *symbol-grounding problem* [Harnad, 1990], meaning that the symbols have no grounding, e.g., in form of feature values, thus that a symbol is solely syntactical without any underlying semantic meaning. Without a grounding, it is not possible to define a new symbol based on new features, thus based on the underlying semantics, but only as a combination of already known symbols.

Another problem which is not restricted to symbolic approaches, however, is there particularly problematic, is the *frame problem*, the problem of finding the right rule to apply in a specific context. This is in particular difficult, as it is not clear which of the rules are actually relevant and help push the reasoning into the right direction. This problem gets in particular relevant at increasing size of the knowledge base [McCarthy and Hayes, 1969].

Supporters of a hybrid subsymbolic/symbolic framework, such as Gärdenfors, additionally state that symbolism can be considered as being incompatible with nature, as it does not allow for creative inductions or genuinely new knowledge [Gärdenfors, 2000, p.40].

In recent years, in particular with the third wave of AI [Kautz, 2022], the subsymbolic paradigm has gained importance, e.g., in the form of neural networks. Most subsymbolic approaches exploit similarity between instances to infer information about new instances based on their similarity to given ones. They circumvent some problems of the symbolic approach, especially the symbol grounding problem, increase the inconsistency-tolerance, and lead to a reasonable result quality. One well known drawback of neural network approaches (and also of most subsymbolic learning approaches in general) is their lack of explainability and interpretability. Research on explainability is, particularly due to the lack of explainability of neural networks, a topic of increasing importance and makes up a whole research area, the area of *explainable AI (XAI)* [Arrieta et al., 2020]. This is also considered in the area of KGE [Bianchi et al., 2020]. Whereas in symbolic reasoning, the way of getting a result is perfectly comprehensible, a neural network is mostly a black box and has no formally defined computational semantics [Garcez and Lamb, 2023]. This opens up the opportunity for adversarial attacks [K. Ren et al., 2020] and reduces explainability and trustworthiness. Other widely known problems are the generalization from training data and the vast amount of training data, time and energy needed [Garcez and Lamb, 2023].

Although there are efforts to tackle the problems of symbolic and subsymbolic approaches inside their own field, a widespread solution strategy is to combine the two paradigms. This combination can be done in several ways, e.g., as a neural subroutine in a symbolic problem solver or a neural network that converts a non-symbolic input in a symbolic output which then can be used as input for a symbolic problem solver [Kautz, 2022]. The need for neuro-symbolic approaches is illustrated in the following example.

Example 2.1. *Consider the problem of the calculation of a term given in handwritten digits (the example can be found in detail in the paper of Raedt et al. [2019]). Using a symbolic approach seems natural, as the calculation can then be easily done. However, the determination of the handwritten digits is difficult, as the digits are imprecisely written. On the other hand, the determination of digits is a relatively simple task for a neural network, whereas doing a calculation is complicated, error-prone and it is definitely unnecessary to train a neural network when a calculator can easily solve the problem. Therefore, this is a prime example for the use of a combination of symbolic and subsymbolic approaches. In this case, it is a* hybrid AI*-approach, combining the subsymbolic and symbolic parts via an interface (for details on the distinction between hybrid-AI and Ne-Sy AI, see below).*

Another example is link prediction in knowledge graphs. It is possible to use an embedding approach such as TransE [Bordes et al., 2013] on the subsymbolic level to embed instances as points and relations as translations into a space. With the help of a loss function an embedding of instances as points and roles as translation is learned by minimizing the distance $d(s+p, o)$ *for each triple* (s, p, o) *in the training set. Then, a new link, thus a new tuple* (s', p', o') *is predicted if* $d(s' + p', o') < \delta$ *for some fixed threshold* δ*, thus if two points are connected via a translation vector representing a role. This results in high quantitative result measures but is not able to incorporate known symbolic information, e.g., in form of rules excluding some inferences from being valid. However, this symbolic information in form of rules should be incorporated into the learning to* *enhance correctness and explainability of the approach. Therefore, Ne-Sy AI is needed.*

The general field of research looking into enhancing neural networks with symbolic information is called Neuro-Symbolic AI (Ne-Sy AI) [Garcez and Lamb, 2023]. The field of application of the principles of Ne-Sy AI is not restricted to neural networks but can also include other subsymbolic approaches due to its generic ideas. Ne-Sy AI tries to increase the explainability of neural networks by combining symbolic and subsymbolic information, as the main statement of the advocates of Ne-Sy AI is that a semantically sound, trustworthy and explainable system can only be achieved by not only considering subsymbolic but also symbolic layers (reasoning layers) [Garcez and Lamb, 2023]. On the one hand, this enables learning from experience and, on the other hand, it enables to draw reasons from what has been learned [Valiant, 2003]. This pattern can also be found in the human brain, as the brain can be seen as consisting of two systems: one fast, intuitive and automatic system and one system for deliberate and complex decisions [Kahneman, 2011].

In Ne-Sy AI, several different approaches are possible, roughly sortable along three axes [Bader and Hitzler, 2005]: The first axis describes the interrelation of symbolic and subsymbolic and thus whether the system is hybrid (as in Example 2.1) or integrated, namely tightly connected. The second axis describes the language used on the symbolic level and thus whether it is actually symbolic (as in automata) or logical. The third axis describes the usage for learning or reasoning. Although the focus in this dissertation lies on the theoretical foundations of such symbolic/subsymbolic frameworks and is especially not restricted to the use neural networks as a basis, it can be considered as following the paradigm of Ne-Sy AI.

Well-known examples of Ne-Sy AI are logic tensor networks [Badreddine et al., 2022] which enhance neural networks with a symbolic layer as a last layer, and also some approaches in KGE where background information is modeled geometrically in the embedding space (e.g., as geometric models by Gutiérrez-Basulto and Schockaert [2018]). These ideas are also common in the area of cognitive reasoning. A key advocate of combining symbolic and subsymbolic information is Gärdenfors [2000]. In fact, we will make use of ideas proposed by Gärdenfors and others in the area of cognitive reasoning in order to lay a foundation for a general framework of a fully integrated symbolic/subsymbolic system.

As the symbolic/subsymbolic approach needed is highly dependent on the problem to be solved, the problem to be solved in this dissertation is stated formally as follows: We consider the general task of enhancing learning and reasoning by incorporating background knowledge information into a learning problem. Given an abox, a subsymbolic learning approach based on this abox is improved by incorporating tbox information.

Such a symbolic/subsymbolic framework enables us to improve correctness, trustworthiness and interpretability by enforcing the prediction of ontologically correct results. One instance of this problem is depicted in the second example of Example 2.1.

2.3 Embeddings for Incorporating Background Knowledge

In Ne-Sy AI, the idea is to combine the neural learning and the symbolic reasoning parts in a way that allows utilizing the best aspects of both. But how can this be done? What can an Ne-Sy AI-system look like that is able to model sufficiently expressive symbolic information, e.g., in form of a tbox, especially information incorporating negation? When, e.g., considering link prediction in knowledge graphs, how can background knowledge be incorporated into the learning process?

One solution, proposed by researchers such as Smolensky [1988] and Gärdenfors [2000] states that enriching the symbolic space with some geometrical structure would be a solution for the symbol grounding problem. Numerical information enables defining some distance measure and thus enables modeling similarity both of concepts and of instances. The term "geometrical" is used in a broader sense by considering spaces with a betweenness relation. When having sets of instances, it can be given some form of structure to them by introducing betweenness relations between them (see Section 4.3). Such a geometric representation can be realized by embedding instances as points and concepts as regions in a geometric space. It enables modeling underlying symbolic information in form of an ontology by enforcing that the concepts fulfill logical constraints which can be modeled in a geometrical way, e.g., a subconcept relation via a subset relation. This representation allows for a comprehensive representation of logical knowledge in geometrical terms and thus enforces the classification result to adhere to the underlying symbolic knowledge.

It is widely argued that this approach of embedding symbols into some space works as model of brain processes. This is done, e.g., in the paradigm of *conceptual spaces* [Gärdenfors, 2000]. Cognition is described as a process involving three different levels, the symbolic, conceptual and subconceptual representation, all considering the problem in different levels of detail. Whereas the subconceptual (or associatist) representation is the most fine-grained level, comparable with the granularity a neural network works on, the conceptual level allows for seeing the bigger picture and recognizing non-local dependencies and makes up the connection between the numerical part and the symbolic part by representing symbols as regions in a space. The symbolic level depicts the most abstract level, enabling abstract reasoning [Gärdenfors, 2000, p. 33 ff.]. Though, as mentioned by Bouraoui et al. [2022], KGE and conceptual spaces differ substantially, e.g., regarding the modeling of explicit domains and quality dimensions, the basic idea of conceptual spaces, namely the geometric representation of symbolic knowledge based on convex regions, can also be applied to KGE.

As here the approach of embedding concepts as regions is considered, the question arises how these regions, thus these geometric structures, should look like. The basic idea of Gärdenfors and others is that the concepts and relations/logical operators between concepts should be modeled as explicit geometric structures resp. operations to directly combine symbolic and numerical data in a formal and clearly defined way. Gärdenfors states that the concepts need to be convex but does not explicitly say how the logical operators should be modeled geometrically.

There are many different embedding approaches modeling background information of different strengths, partly able to model some sort of negation (see Sections 3.3 and 10.2). In the following, however, it is not the aim to create another embedding approach similar to the others but based on a different underlying structure, but to give a general construction principle, pointing out which geometric (or even weaker, somehow structured) objects can be used as a basis for a learning approach and how logical operators (especially negation) can be modeled.

To achieve this, the first question is how the logical operators should be modeled. This question comes in two parts: First, it is necessary to determine the strength and expressivity of the logical operators, independent of their representation in geometrical terms. This is done, with a special focus on negation, in Chapter 3. Second, it needs to be determined how these logical operators are modeled in geometrical terms. This is in the focus of Part I.

3 The Basis: Choosing Logical Operators for an Expressive Embedding

Negation in learning approaches is a topic that is becoming increasingly popular, especially in learning for language models, as argued in Chapter 1. However, also in KGE, the importance of modeling negation has increased, as more and more KGE-approaches are able to model some (weak) form of negation (e.g., modeled via polarity by Garg et al. [2019], non-intersection by Kulmanov et al. [2019] and set-complement by Zhang et al. [2021]). As negation is an ambiguous notion with a wide range of interpretations, the aim of this chapter is to identify the type of negation suitable for the framework that will be presented in Part I.

The interpretation of negation has been widely discussed for centuries. One main observation, introduced by Aristotle [Horn, 2001, pp. 6 ff.], is the differentiation between contrary and contradictory negation. Whereas a contrary depicts non-intersection between statements, e.g., “the apple is red” vs. “the apple is green”, a contradiction depicts the literal negation, thus, e.g., “the apple is non-red”. Many embedding approaches are restricted to contrary negation, thus are able to model non-intersection of concepts but do not model the negation of a concept explicitly. An example for such an approach is, e.g., ELEm [Kulmanov et al., 2019], where concepts are modeled as spheres in $\mathbb{R}^n$. Though this negation is simple to model, it is not sufficient for all use cases, as is reflected upon in detail in Section 3.3.

Although there are many different interpretations of negation, in the following, I focus on a specific one, the orthonegation, as it is philosophically justifiable and has logical properties suitable for the use as the basis of an embedding approach.

The orthonegation is related to an observation, made by researchers including Demos [1917]. He states that a negative statement needs to be derived from a positive perception, e.g., the statement of someone being not at home is based on the (invalid) perception that someone is at home. Thus, the reference is the world of positive facts. The term “not” is applied to the whole statement and leads to the fact that the “not”-

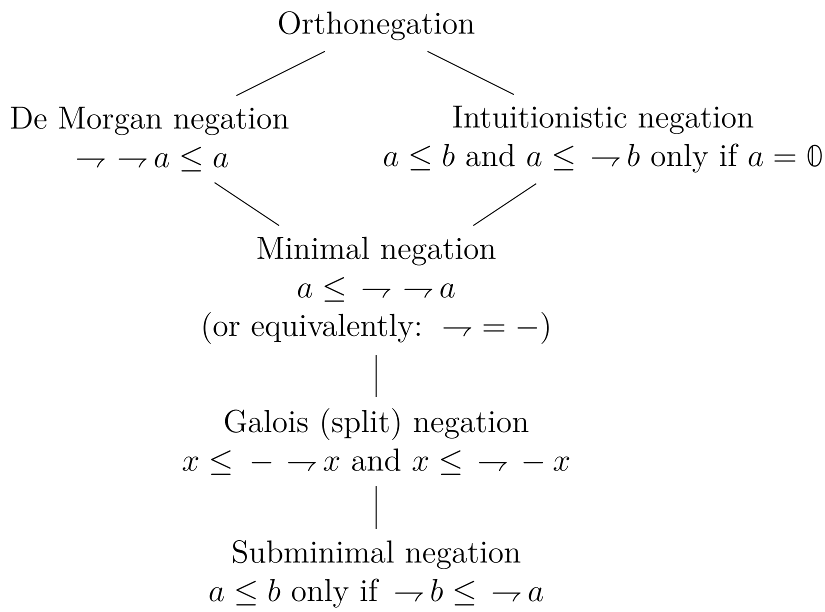

Figure 3.1: Dunn's Kite of Negation (adapted from [Dunn, 1996])

relation can be considered as notion of inconsistency or opposition.

The Birkhoff-von Neumann-Goldblatt interpretation of negation [Dunn, 1996; Birkhoff and Von Neumann, 1975; Goldblatt, 1974] follows this line of thought. Its main ingredient is that $\rightharpoondown A$ is true at a state of information χ if all states of information α in which A is true are incompatible with χ. Informally, this means that something is false if there is no possibility to make it true. Thus, a negation can be stated as follows [Dunn, 1996]:

$$\chi \models \rightharpoondown A \text{ iff for all } \alpha(\alpha \models A \text{ entails } \chi I \alpha),$$

where I depicts an incompatibility relation.

The following considerations on the type of negation are based on a poset $\mathfrak{P} = (P, \leq)$[1] with bounds $\mathbb{0}$ and $\mathbb{1}$. Elements of P are propositions and $a \leq b$ depicts "proposition a implies proposition b". Then, the negation $\rightharpoondown a$ of a proposition a can be seen as the weakest proposition incompatible with a [Dunn, 1996]:

$$x \leq \rightharpoondown a \text{ iff } x I a.$$

The properties of the resulting negation are specified by the properties of this incompatibility relation, especially by the question of whether the incompatibility should be

[1] For details on posets, see Appendix A.

irreflexive and symmetric. In Figure 3.1, Dunn's so-called Kite of negation can be seen, showcasing the types of negation based on different strengths of the incompatibility relation. Subminimal and Galois negations are based on a non-symmetric incompatibility and thus on two different notions of negation, $\neg$ and $-$, thus

$$x \leq -a \text{ iff } aIx.$$

However, as symmetry of the incompatibility relation is justifiable, both negations can be identified [Dunn, 1996, p. 10]. This symmetry leads to minimal negation, irreflexivity and symmetry of the incompatibility to intuitionistic negation. Additionally enforcing double negation elimination, thus $\neg\neg a = a$, and symmetry of the incompatibility leads to de Morgan negation and symmetry and irreflexivity to orthonegation [Dunn, 1996, p. 15].

Although it is possible to argue for an incompatibility relation that is not irreflexive and symmetric, in practice both properties seem to be reasonable and thus are assumed to be valid in the following. Double negation elimination is considered a basic property of negation [Horn, 2001, pp. 22 ff.] and, from a practical point of view, eases the handling of the negation. In particular, learning problems can be approached easier, as, e.g., overfitting is circumvented. However, these considerations can also be used as a starting point for more complex constructions, e.g., without assuming double negation elimination. Thus out of the possible negations (when considering Dunn's Kite) we are left with orthonegation.

In the following, the definition of orthonegation and the resulting orthologic are considered in more detail.

3.1 Orthologic and the Orthogonality Relation

Let $P = \{P_i \mid i \in \mathbb{N}\}$ be a set of proposition symbols and have logical symbols for binary conjunction $\&$, unary negation $\neg$, and define binary disjunction $\bigvee$ as $\neg(\neg A \,\&\, \neg B)$. The set of propositional formulae Fml is defined as usual. $A, B, C, \ldots$ stand for propositional formulae in Fml.

Consider natural deduction calculi with a derivability relation $\vdash$. Let $A \dashv\vdash B$ be the short notation for $A \vdash B$ and $B \vdash A$ and for a finite set of formulae $\Gamma = \{B_1, \ldots, B_n\}$ the notation $\Gamma \vdash A$ is a shorthand for $((B_1 \,\&\, \ldots) \,\&\, B_n) \vdash A$.

Definition 3.1. *[Goldblatt, 1974, Definition 1.1] An* orthologic *is a logic L such that for all $A, B, C \in Fml$*

- $A \vdash A$
- $A \mathbin{\&} B \vdash A$
- $A \mathbin{\&} B \vdash B$
- $A \vdash \neg\neg A$
- $\neg\neg A \vdash A$
- $A \mathbin{\&} \neg A \vdash B$
- *if* $A \vdash B$ *and* $B \vdash C$, *then* $A \vdash C$
- *if* $A \vdash B$ *and* $A \vdash C$, *then* $A \vdash B \mathbin{\&} C$
- *if* $A \vdash B$, *then* $\neg B \vdash \neg A$

The smallest logic satisfying these axioms is called *minimal orthologic* Omin. Omin is characterized by an ortholattice by defining $\&$ as lattice meet $\wedge$ and $\neg$ as orthocomplement $\cdot^{'}$ (for details on lattices in general, ortholattices and the restrictions of ortholattices mentioned below, see Appendix A).

The properties of an ortholattice are the following:

- $a \leq b$ entails $b^{'} \leq a^{'}$ (antitonicity)
- $a^{''} = a$ (double negation elimination)
- $\mathbb{0} = a \wedge a^{''}$ (intuitionistic absurdity)

An ortholattice does not necessarily fulfill distributivity, nor modularity or orthomodularity. Thus, it allows on the one hand for modeling very weak lattices, on the other hand it has the Boolean algebra as a special case (an ortholattice fulfilling distributivity is a Boolean algebra [Padmanabhan and Rudeanu, 2008, Proposition 4.10.1.]). Thus the class of ortholattices allows for defining logics of different strength, depending on the use case.

Goldblatt considers the orthonegation and the corresponding orthologic as a basis for defining an *orthoframe* to provide a basis for an intensional model theory of quantum logic. His main point is a representation theorem stating that every ortholattice can be represented by an orthoframe, a space X equipped with a symmetric and irreflexive orthogonality-relation $\perp$ [Goldblatt, 1974].

Definition 3.2. *[Goldblatt, 1974, Definition 2.1 and 2.2]*

$F = (X, \perp)$ *is an* orthoframe *iff* X *is a non-empty set and* $\perp$ *is an* orthogonality relation *on* X, *i.e.,* $\perp \subseteq X \times X$ *is irreflexive and symmetric.*

$M = (X, \perp, V)$ *is an* orthomodel *on the frame* $(X, \perp)$ *iff* V *is a function assigning to each propositional variable* p_i *a* $\perp$*-closed subset* $V(p_i)$ *of* X*. The truth of a Fml-formula* A *at* x *in* M *is defined recursively as follows (where* $M \Vdash_x A$ *depicts that* A *is true at* x *in model* M*).*

- $M \Vdash_x p_i$ *iff* $x \in V(p_i)$
- $M \Vdash_x A \,\&\, B$ *iff* $M \Vdash_x A$ *and* $M \Vdash_x B$
- $M \Vdash_x \neg A$ *iff for all* y*,* $M \Vdash_y A$ *only if* $x \perp y$*.*

A set $A \subseteq X$ is $\perp$-closed iff for all $x \in X, x \notin A$ only if there is $a \in X$ such that $a \perp A$ and $x \not\perp a$ (where $a \perp A$ is a shortcut for $a \perp a'$ for all $a' \in A$) [Goldblatt, 1974]. Denote $A^{\perp} = \{x \in X | x \perp a$ for all $a \in A\}$. This enables to define an equivalent definition of $\perp$-closure: A set $A \subseteq X$ is $\perp$-closed iff $A = A^{\perp\perp}$. Later on, several different $\perp$-relations $\perp_i$ are considered. For readability, the term $\perp$-closure is used when it is clear from context that it refers to $\perp_i$-closure.

These orthomodels are a structure of relevance, as Goldblatt shows that the logic Omin is strongly determined by the class Θ of all orthoframes, thus, that each orthoframe makes up an ortholattice and that each ortholattice is representable as a subortholattice of an orthoframe.

Theorem 3.3. *[Goldblatt, 1974, Corollary 3.6]* $\Gamma \vdash_L A$ *iff* $M : \Gamma \Vdash A$*.*

This connection between lattice and algebraic representation gives a hint on the representability of this negation in the context of an embedding.

In this construction, it is possible for an element a that neither $a \in A$ nor $a \in \neg A$, and it thus being "in-between". This hints to two different types of negation, namely: something is *implicitly false by default*, thus is not verified by any intended model, and something is *explicitly false by virtue of a direct proof*, thus is falsified in any intended model. These two negations can be differentiated as *weak and strong negations* [Herre et al., 1999, p. 121] (for details, see Section 3.3). This distinction is especially important in knowledge-based reasoning when relying on the open-world assumption. With the closed world assumption, both types of negation are identical [Herre et al., 1999, p. 154].

These considerations motivate considering in detail this "in-between", thus the elements a not modeling A but also not being incompatible to all elements modeling A. This is the topic of the next section.

3.2 Orthomodels Induce Partial Interpretations

These in-between states induce a form of partial interpretation allowing for incorporation of partial information. In the context of ontologies, the ability to account for partial information gets important if, e.g., the abox does not contain information about all concept memberships for all instances. For example, consider the abox $\mathcal{A} = \{A(a), B(a), A(b)\}$

over the signature $\sigma = \{A, B, a, b\}$, where no information is given whether $B(b)$ or $\neg B(b)$. In search for an adequate epistemic semantics a partial interpretation may help: it does not enforce to infer each missing concept membership either positively or negatively but enables to model this missing information as it is, namely as missing information. As stated by Özçep et al. [2023], this allows for modeling not only one possible, classical model of an ontology but all possible models and enables a learner to use this incomplete information without enforcing it to give statements about the missing concept membership assertions. A model having this property is termed *faithful model*. It is a model allowing for modeling the incomplete information regarding each concept of the ontology.

Definition 3.4. *[Özçep et al., 2023, Definition 3] Let $\mathcal{O}$ be a classically consistent (DL) ontology (or any other representation allowing the distinction between abox and tbox). For a (not necessarily classical[2]) interpretation $\mathcal{I}$ we have the following notions of being a faithful model of $\mathcal{O}$:*

- $\mathcal{I}$ *is a* strongly concept-faithful *model of $\mathcal{O}$ iff for each concept C and each constant b the following holds: if $b^\mathcal{I} \in C^\mathcal{I}$, then $\mathcal{O} \models C(b)$;*
- $\mathcal{I}$ *is a* weakly concept-faithful *model of $\mathcal{O}$ iff for each concept C and each constant b the following holds: if $b^\mathcal{I} \in C^\mathcal{I}$, then $\mathcal{O} \cup \{C(b)\}$ is classically satisfiable;*
- $\mathcal{I}$ *is a* strongly (weakly) abox-faithful *model of $\mathcal{O}$ iff it is strongly (weakly) concept-faithful and for each role R and constants a, b the following holds: if $(a^\mathcal{I}, b^\mathcal{I}) \in R^\mathcal{I}$, then $\mathcal{O} \models R(a, b)$ (resp. $\mathcal{O} \cup \{R(a, b)\}$ is classically satisfiable);*
- $\mathcal{I}$ *is a* strongly (weakly) tbox-faithful *model iff for all tbox axioms $\tau = C \sqsubseteq D$ the following holds: if $\mathcal{I} \models \tau$, then $\mathcal{O} \models \tau$ (resp. $\mathcal{O} \cup \{\tau\}$ is satisfiable) and if there is $x \in \Delta^\mathcal{I}$ with $x \in (C \sqcap \neg D)^\mathcal{I}$, then $\mathcal{O} \cup \{(C \sqcap \neg D)(a)\}$ is satisfiable for a fresh constant symbol a (resp. $\mathcal{O} \not\models \tau$).*

Because of their ability to model partial information, orthomodels are able to act as a concept-faithful model of an ontology (see, e.g., the considerations on concept-faithfulness of al-cones in Proposition 5.19).

As argued by Özçep et al. [2023], such a partial interpretation must be differentiated from partial models in the context of general logic programs, notably in the context of

[2] For details on alternatives for the interpretation function, see paragraph “Negation in Description Logic” of Section 3.3 and the interpretation based on COSs in Section 5.1. In Definition 3 of the paper by Özçep et al. [2023], geometric interpretations in form of some special type of cones have been considered, the notion is, however, also useful for arbitrary interpretation functions.

well-founded semantics [Van Gelder et al., 1991]. These differ from partial interpretations since those partial models are meant to treat $p \leftarrow \neg q$ and its contraposition $q \leftarrow \neg p$ differently, whereas in this case the contraposition rule (and double elimination) holds. In fact, as Van Gelder et al. [1991] note, there is a 3-valued logic interpretation for partial models in the well-founded semantics whereas for our partial models this kind of extensional semantics based on truth values is not possible: Consider the truth value of $p \vee q$. Assume that the ontology does not entail p (so p gets assigned the third truth value different from true and false) and q is not entailed (so q gets assigned the third truth value different from true and false). These two assignments do not determine the truth value of $p \vee q$: because the truth value of $p \vee q$ depends on whether the ontology entails the disjunction or not.

Now that I have introduced the technical notions of faithful models, the philosophical question on the type of modeled partiality might arise, in particular, one might ask whether it allows to model vagueness, uncertainty or truthlikeness. All three notions are highly discussed (e.g., by Williamson [1994] for vagueness and in Chapter 6 of the book of Yablo [2014] for truthlikeness (there called *verisimilitude*)) and a detailed philosophical discussion of the best fitting interpretation of the partiality in context of the Conceptual Orthospaces is beyond the scope of this work. Nonetheless, the notion of an orthomodel is that general that it should be possible to interpret it or an adaptation of it in each of the three types of partiality. It is, e.g., possible to consider *scaled Boolean Algebras* [Hardy, 2002] to model degrees of belief in uncertain propositions.

All three depict, in different variants, the "space in between" positive and negative instances of a concept. These three types of uncertain information have to be differentiated in order to (i) exactly determine the interpretation of the in-between in light of orthonegation and (ii) to include this notion in the discussion on the logical commitments of learning approaches in general (see Chapter 10).

Here we will focus on an interpretation of vagueness, uncertainty and truthlikeness by Godo and Rodríguez [2008], as they argue in the context of similarity-based reasoning which is related to the framework of this dissertation. *Vagueness*, also known from fuzzy set theory [Zadeh, 1996], describes the degree of truth of a statement when having vague information, e.g., "the mountain is high". In contrast, *uncertainty* is not based on a degree of truth but on a degree of confidence or belief, thus relying on incomplete information. Both notions are strongly connected, as having a vague premise might lead to an uncertain conclusion. *Truthlikeness* is based on similarity of statements, having a definitely false statement being similar to a true statement, e.g. the statement "Mount Everest is 8800m high". As Mount Everest is 8848m high, this statement is definitely false (and thus neither uncertain nor vague) but closer to the truth than, e.g., the statement "Mount Everest is 80m high".

This in-between also influences the representation of disjunction. Whereas for the conjunction "x is both in A and in B" it is naturally the case in an orthomodel that x is in A and x is in B, this is not directly adaptable to the case of disjunction. Although the law of excluded middle is valid in orthologic, it is not the case that from "x is in A or B" it follows that x is in A or x is in B. When considering a two-valued logic, this leads to the fact that some form of vagueness, uncertainty or truthlikeness is modeled.

3.3 Negation in Embedding Approaches

As mentioned above, negation in learning approaches is a topic of increasing popularity, also for embeddings.

There are at least two embedding approaches using an orthogonality relation as basis: Garg et al. [2019] use the classical interpretation of an orthomodel by interpreting concepts as linear subspaces of a Hilbert space and the negation as orthogonality relation in form of a polarity. Özçep et al. [2020] uses also negation as polarity but a restricted form of closed convex cones as concept interpretation.

But how is the negation in actual embedding approaches interpreted in general? It turns out that besides the use of the orthogonality relation, there are two different interpretations of the negation: Negation as non-intersection (not a negation in the stronger sense, leading to contrary instead of contradictory elements) and weak negation. Additionally, negative sampling is considered. The last is, in contrast to the first two, not a type of negation but a strategy for choosing negative elements enabling to learn the borders of a positive concept. However, as it is based on negative elements, it can be considered a form of negation strategy (or, at least, enforces to have some thought of the nature of negative elements).

Types of Used Negation

In the following, the two most common types of negation modeled in actual embedding approaches are introduced. A third type, appearing at least in the approaches of Garg et al. [2019] and Özçep et al. [2020] is the orthonegation described in Section 3.1.

Negation as contrary: non-intersection As mentioned at the beginning of this chapter, it is possible to distinguish between contrary and contradictory. When, e.g., considering the context of DLs, full negation can be modeled in the DL $\mathcal{ALC}$ or stronger DLs, whereas non-intersection of concepts can be modeled in weaker DLs like $\mathcal{EL}_{\bar{0}}$ or

$\mathcal{EL}^{++}$. Non-intersection is widely used, as many embedding approaches aim at embedding $\mathcal{EL}^{++}$. Examples are by Jackermeier et al. [2023], Kulmanov et al. [2019] and many more. Details on these approaches can be found in Section 10.2. The advantage of such an interpretation is its simplicity. It does not require to model a negated concept and to think about its shape at all but only to satisfy non-intersection. This, however, also means that it is not possible to state actual negation in the usual, stronger sense. Thus, though it has a simple structure, it is not useful for our aim of modeling negation.

Weak negation There is a distinction between strong and weak negation. For example in the context of knowledge bases, strong negation denotes explicit falsity, thus $\neg p$ is true iff p is false in all possible models , weak negation denotes non-truth, meaning $\neg p$ is true iff p cannot be verified in any model of the knowledge base [Herre et al., 1999, p. 121]. Weak negation can be interpreted as the closed world assumption being valid [Herre et al., 1999]. In contrast, the interpretation, e.g., with the help of the orthonegation allows for differentiating between strong negation ($M \models_x \neg A$) and weak negation ($M \not\models_x A$).

Due to its simplicity, it is a negation widely used in embedding approaches. Stating that everything not being true is false allows for dismissing the definition and learning of negative concepts. An example is ConE [Zhang et al., 2021] which essentially embeds concepts as cones and negated concepts as the set-complement of these cones. One main disadvantage of this interpretation of negation is that there is no distinction between not-knowing whether something is true and knowing whether something is false. This is especially problematic, e.g., in the area of medicine where it is vital to state the discrimination of whether an examination has not been done or an examination has been done but has a negative output. Weak negation also leads to the fact that the positive and negative concepts are not of the same strength. The positive is determined and the negative is interpreted as "everything else". Thus, the negative concept does not have any intended meaning. Accordingly, this type of negation is unsuitable for the embedding approach which is in the focus of the next chapters.

Negative Sampling

Negative sampling was introduced for the embedding approach word2vec by Mikolov et al. [2013] and has been widely examined and improved since then.

When considering a standard KGE-approach such as TransE [Bordes et al., 2013] with the training objective of maximizing the probability that each triple of the training-data is classified as valid triple, then this would result in a trivial classifier, setting the probability for each triple to one. As this is not appropriate, it is necessary to introduce corrupted triples into the training process. Then, learning is possible by maximizing the probability that triples of the data are classified as belonging to the data and corrupt ones are classified as not belonging to the data. This is called *negative sampling*. For details, see, e.g., the considerations on negative sampling in the context of word2vec by Goldberg and Levy [2014].

Whereas the basic negative sampling idea is based on drawing arbitrary corrupted tuples, the idea has evolved to incorporate more complicated techniques for identifying suitable negative tuples, as these tuples should, on the one hand, not be too restrictive – leading to too many false negatives – and, on the other hand, not too broad – resulting in many false positives.

In the area of KGE, the most-widely used technique for negative sampling is *uniform sampling*, where in a correct (subject, predicate, object)-triple the object is exchanged with an arbitrary other one [Qian et al., 2021]. As this mostly leads to easily distinguishable negative samples, several enhancements have been proposed, e.g., by using generative adversarial networks (GANs) to do adversarial training or by restricting the domain of which the objects are drawn (*custom cluster-based sampling*) [Qian et al., 2021]. Though many of these approaches lead to a high result quality, they do not enable modeling an expressive negation. The aim of negative sampling is to find samples as close to the positive ones as possible to determine the border of the positive region as exactly as possible. Negative sampling has the sole aim of modeling the positive correctly and has no incentive to model the negative. Thus, the negative samples do not have any inherent meaning.

Negation in Description Logic

As mentioned in Table 2.1, the usual interpretation of negation in a Description Logic setting is based on set complement ($(\neg A)^{\mathcal{I}} = \Delta^{\mathcal{I}} \backslash A^{\mathcal{I}}$). But DLs are not restricted to this interpretation. An example is work on *inconsistency-tolerant DL*, introduced, e.g., by Odintsov and Wansing [2003], where classical negation is replaced by constructive negation and thus $(\neg A)^{\mathcal{I}}$ is interpreted as $\sim (A)^{\mathcal{I}}$ where $\sim$ denotes the negation.

When considering an orthomodel as an interpretation of a DL, the problem occurs that the interpretation of negation as set-complement (as done in $\mathcal{ALC}$) is too strong, as partial information can then not be modeled. Thus, especially in the context of faithful models, it is necessary to change the interpretation. Therefore, for interpreting the orthomodel as DL-interpretation, there are two possibilities: If the underlying orthomodel fulfills distributivity, it is possible to interpret each element neither in a concept nor its negation as either belonging to the positive or negative concept, thus resolving the problem by approximating the interpretation in $\mathcal{ALC}$. This has been done, e.g., by Özçep et al. [2020] and is discussed in detail in Theorem 5.18. The other possibility is to change the interpretation of negation. This is discussed in detail for an interpretation based on COSs in Definition 5.3. In the following, an interpretation is used where the negated concept $(\neg A)^{\mathcal{I}}$ is interpreted as $(A^{\mathcal{I}})^{\perp}$, thus the interpretation of the negation is defined as the concept orthogonal to A.

3.4 Resulting Options for Modeling Conjunction and Disjunction

The last sections underline the advantages of using the orthogonality relation as basis for a negation. Though the question of the type of negation has been the most urgent one, it is also of importance to consider how conjunction and disjunction should be modeled. Whereas this seems to be straightforward at first sight, in fact, it is not.

For modeling disjunction, as seen in Section 3.2, it is not sufficient to model it based on the excluded middle in the sense that "x is in A or B" iff "x in A" or "x in B"[3]. This phenomenon is also called *overextension* and, together with its opposite *underextension* (an element belongs to concept A but not to A or B), is a widely researched topic both for disjunction and for conjunction, as it occurs both in human reasoning [Hampton, 1988a] and in logic in form of a *truth-value gap* [Priest, 2008, pp. 127 ff]. Overextension occurs when an element is in neither of the two concepts but at least in some way close to both and in between the two concepts [Hampton, 1988a]. It can occur in human reasoning when the disjunction of two concepts leads to interpreting the disjunction as a new super concept. One prominent example would be the disjunction "fruits or vegetables" possibly leading to the idea of "green groceries" and thus interpreting "mushrooms" as "fruits or vegetables" even though they are clearly neither fruits nor vegetables [Hampton, 1988a]. These considerations thus justify the use of the orthonegation and the disjunction based on it. In contrast, underextension is not represented by orthomodels, as this would interfere with its definition.

The representation of conjunction is also a highly researched topic. Whereas the simplest way would be to model conjunction as set-intersection, it is mainly argued that this is not sufficient. Resulting problems are, e.g., the so-called *pet-fish problem*, depicting that, e.g., a guppy is more prototypical to be a pet-fish than to be a pet or to be a fish [Osherson and Smith, 1981]. A more catchy example would be something "out of stone an being a lion" which relates to the concept of "stone lion" and is definitely not a lion in the usual sense. Another problem is *borderline contradiction*, e.g., stating

[3] Note that it is possible to model an orthoframe such that "x is in A or B" iff "x in A" or "x in B" for all $\perp$-closed sets A, B, e.g., for a space with two elements. These cases are, however, rare.

that someone is tall and short at the same time. This could be true if he is of normal size, thus in the borderline of both concepts [Bonini et al., 1999]. Additionally, the overextension and underextension can appear for conjunction in the same way as for disjunction [Hampton, 1988b].

These considerations lead to many different ideas of interpreting conjunction differently from set-intersection (see, e.g., the discussions in context of conceptual spaces [Gärdenfors, 2000]). At least the borderline contradiction-problem can be circumvented with the proposed interpretation based on the orthogonality relation, as based on the in-between, it is possible to model that a person is in between of being short and tall. For the other problems, however, it is necessary to find a trade-off between representability, computability and cognitive correctness. Here, the focus lies on a feasible model, usable in actual embedding contexts and therefore, the representation of conjunction as set-intersection is chosen, enabling for easier learning and the representation of a classical logic.

To wrap up the considerations in this and the last chapter, the aim is to create an embedding approach, interpreting negation as orthonegation, conjunction as set-intersection and disjunction as $\perp$-closure. Whereas there are some embedding approaches operating like this (e.g., by using cones [Özçep et al., 2020] or closed subspaces [Garg et al., 2019]), the aim here is not to find another approach but to find a general framework that can serve as the basis of several embedding approaches. The aim, in particular, is to exploit the properties of the given data to define a suitable embedding. Thus, the question, to be answered in the next section, is what basic, mandatory properties an embedding approach should have and which properties are adjustable to specific needs of the data on a low level of the embedding approach.

The aim is explicitly not to find a totally correct and psychologically grounded representation of knowledge but an approximation in the sense of Ne-Sy AI to allow for learning and reasoning.

Part I

Foundations of Conceptual Orthospaces

4 Concept Formation Based on Betweenness and Similarity

In Chapter 2, I argued that a concept-based view on learning and, in particular, a combination of symbolic and subsymbolic learning enhances correctness, explainability and quality of the results and one way of achieving this is by using an embedding approach. In Chapter 3, an opportunity for the interpretation of logical operators as geometric ones was identified: negation can be modeled with the help of an orthogonality relation, conjunction as set-intersection and disjunction as $\perp$-closure. But how can an embedding approach be defined that is, (i) as variable as possible and (ii) adaptable to as many datasets and their needs as possible? The solution is to identify the most basic properties of embedding approaches, adapting them to fit the properties of the dataset and providing a framework that allows to define an embedding approach based on these.

When considering standard approaches of KGE such as TransE, it turns out that similarity is a vital notion. An instance is considered being the object of an (subject, predicate, object)-triple (where subject and predicate are given) if it is sufficiently similar to the expected, calculated object [Bordes et al., 2013]. When considering also information about concept membership in a KGE-approach, an instance can only be member of the concept if it is sufficiently similar to at least some of the other members of that concept.

Other approaches (not being based on embeddings but geometric information in general) such as k-nearest-neighbor assign to a new instance the class label that most of its k neighbors (k most similar instances) have [Bishop, 2006, pp. 124 ff.].

There are many options of modeling such a similarity, e.g., based on the Euclidean distance. Choosing a suitable similarity relation is a design decision influencing the embedding in general and, as the orthogonality relation can be interpreted as dissimilarity, also the interpretation of the negation. Therefore, I identify the similarity as the first vital property of embedding approaches.

Although similarity is highly relevant, decision-making is not solely based on it, as conceptual information, e.g., in form of theories or rules may also be used (as argued in Section 2.2, e.g., based on the two systems of Kahneman [2011]). Therefore, a sufficient representation of concepts is needed. Solely based on similarity, this is, e.g., possible via prototype theory. Here, each concept is represented via a prototype and instances sufficiently similar to this prototype are classified as belonging to the concepts represented by the prototype [Geeraerts, 2006]. When defining concepts based on similarity with the help of prototype theory, the definition of conjunction as set-intersection is counterintuitive and not reasonable, as it leads to a new region based on a new prototype. Thus, prototypes are only appropriate when considering contrast sets, i.e., sets of jointly exhaustive and pairwise disjoint concepts [Schockaert and Gutiérrez-Basulto, 2021].

A different possibility is to represent concepts based on similarity, by interpreting the dissimilarity as orthogonality relation (for details, see Chapter 5) and define concepts as $\perp$-closed sets. An element is in the $\perp$-closure of a set iff it is orthogonal to all elements which are orthogonal to the set. Thus, the concept membership is not determined by considering the similarity to the instances of the positive concept but by considering the dissimilarity to instances of the negative concept.

One way to get a more detailed understanding of the structure of concepts is using betweenness. The notion of betweenness enables stating construction principles of the form "everything having a specific attribute value ranging from x to y", and, thus, states that all elements sharing the attributes of the instances of a concept are also members of this concept. Enforcing betweenness relations enhances the concepts with a structure. In particular, betweenness allows for statements about convexity of sets (thus, betweenness-closure) and enables restricting oneself to such convex sets. Convexity is both cognitively motivated [Gärdenfors, 2000] and motivated by practice, as many embedding approaches rely on convex sets as basis, in particular because of their computational advantages (for details, see Section 10.2).

Thus, in the following, these structures (similarity and betweenness) are introduced, starting with similarity and its properties in Section 4.1. Then, Section 4.2 explains in more detail why similarity alone is not sufficient and why, therefore, convexity is needed. The notion of convexity leads directly to a notion of betweenness, introduced in Section 4.3. At the end of this chapter, in Section 4.4, other existing approaches using betweenness (in particular the ones using it in combination with similarity) are discussed to underline the plausibility of the decision to use these two properties as a basis. Although there are some approaches incorporating similarity and betweenness, they either do not incorporate negation at all or only in a quite restricted form. Thus, this discussion highlights the need for the framework of Conceptual Orthospaces, combining betweenness and similarity and acting as the basis for expressive embedding approaches. This will be the topic of Chapter 5.

4.1 Similarity

The notion of similarity plays a fundamental role in theories of knowledge and behavior. It is an organizing principle and helps for classification, concept formation and generalization of learned results [Tversky, 1977]. This is long-established knowledge, voiced, e.g., by the philosopher Hume [1894, p. 104] in the context of its role in analogy:

> "ALL our reasonings concerning matter of fact are founded on a species of Analogy, which leads us to expect from any cause the same events, which we have observed to result from similar causes."

Regardless of whether some conceptual information is considered or not, the notion of similarity is vital for numerical learning approaches in general. Neural networks, in particular, are based on similarity in such a way that they classify elements as belonging to the same class if they are similar to each other. Similarity can be measured in different ways and its definition is highly domain-dependent and hence there is no similarity universally applicable for every use case [Goldstone and Son, 2005, pp. 30 ff.].

Thus, it is worth considering different types of similarity relations. One approach commonly used is modeling similarity as a metric, thus stating that objects a and b are similar if $d(a, b) \leq \tau$ for a distance d and a constant τ. Using similarity relations based on distance functions is not always appropriate [Tversky, 1977]. Tversky especially argues against the three attributes of a metric (the distance from a point to itself being zero, symmetry and the triangle inequality). First, he makes the point that enforcing an instance having a non-zero distance to itself is possible. This is based on, e.g., an experimental finding of Podgorny and Garner [1979] regarding the perception of similarity between different and also equal letters. People tend to find, e.g., letter C more similar to letter O than letter W to another occurrence of letter W. Symmetry is also not always given. As an example, Korea is more similar to China, than China is to Korea, as China is more salient than Korea. The triangle inequality is also not always fulfilled: Tversky gives the counterexample of Jamaica being similar to Cuba (geographically) and Cuba being similar to Russia (ideologically) but Jamaica not being similar to Russia.[1] Therefore, other, more general similarity relations are more useful in some cases, e.g., similarity based on sets and their interplay.

It is also discussed whether similarity is domain-dependent in the sense that it is only meaningful to measure similarity of two objects regarding one domain, Tversky [1977],

[1] All three counter-examples are of course highly debatable, in particular the first one. I will not go into detail here, as the aim of this considerations is only to point out that there are different viewpoints regarding properties of similarity.

e.g., argues that the similarity of an apple and a traffic light needs to be considered based on a domain, e.g., their color. On the other hand, e.g., Goldstone and Son [2005, p. 30] argue that similarity is particularly necessary and helpful in cases in which no explicit knowledge over the responsibilities of the domains is given or an assignment to a domain is not possible. To conclude, the basic properties of similarity are widely discussed and their validity is often questioned. As these decisions are highly dependent on the specific use case, in this dissertation, I focus on the classical and general form, stating that a similarity relation is a reflexive and symmetric binary relation and assume that similarity can be an arbitrary relation fulfilling these properties — without considering domain-dependence explicitly — in order to gain an approach that is as broad as possible. The focus lies on the distinction similar vs. non-similar and not on a graded similarity. This allows for a direct connection between similarity and the orthogonality relation, as orthogonality is irreflexive and symmetric and thus depicts a notion of dissimilarity. As the orthogonality relation is a general notion, it is possible to interpret it with the help of an arbitrary (dis-)similarity relation. Thus, the negation of an instance is interpreted as all elements which are sufficiently dissimilar to this instance. This goes perfectly in line with the considerations of Dunn [1996] and Goldblatt [1974] on incompatibility mentioned in Chapter 3.

Similarity as a Basis for Learning

Similarity is the basis for numerous learning approaches such as SVMs and those approaches based on neural networks. In the following, some standard approaches based on similarity are presented and it is argued why these are not sufficient in targeting the goal of defining an embedding with an expressive incorporated background knowledge.

One prominent example for the use of similarity is *multi-dimensional scaling* [Goldstone and Son, 2005], the embedding of objects and their similarity relations into a geometrical space by modeling their distance related to their similarity. This enables finding the underlying features representing the objects by interpreting the axes semantically and allows for dimension reduction in the sense that only dimensions, thus features, relevant for the similarity values are modeled [Goldstone and Son, 2005]. There are several different applicable similarity-relations [Saeed et al., 2018]. The validity of the approach is solely based on the chosen similarity. These approaches consider neither the shape of the resulting concepts (in fact, they do not consider concepts explicitly at all) nor background information in form of underlying rules. Thus, these approaches are not suitable for an expressive embedding approach incorporating negation.

Another prominent example are *Support Vector Machines (SVMs)*. Their main idea is to do binary classification by finding a discriminator maximizing the distance to instances

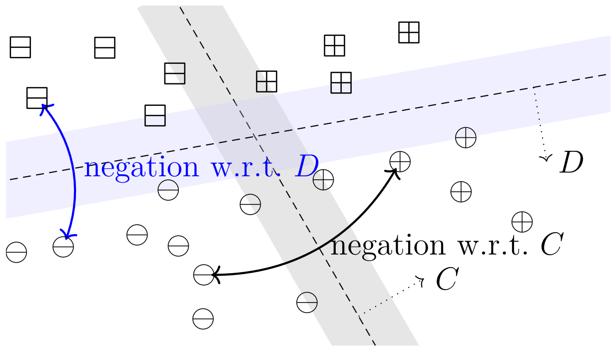

Figure 4.1: SVM classification as an example of missing homogeneity in classifiers

of both classes. Instead of using a linear classifier, it is possible to use the so-called *Kernel-Trick* to train, e.g., a polynomial classifier (for details, see, e.g., [Bishop, 2006]): Though SVMs are able to model some form of negation and also concepts, they are not able to model an expressive logic, as has been argued by Leemhuis, Özçep, and Wolter [2022] and is restated here. Let us consider a simple example depicted in Figure 4.1. For ease of presentation, we assume the classifier is a SVM without kernel (thus leading to linear classification). The actual problem also arises with complex kernels as well as with other classifiers. Figure 4.1 shows two concepts, C and D, with the lines determined by two SVMs (dashed lines) that separate the positive and negative training data for the individual concepts. Data points labeled with + belong to C, labels enclosed by a circle belong to D. Thus, a point labeled ⊞ is within C and outside D. Consider the conjunction of concepts C and D which is naturally represented by the set intersection of the concepts C and D. The geometrical structure of half spaces as used by linear kernel-SVMs lacks homogeneity in the sense that it is not closed under logic operations, i.e., it cannot represent $C \wedge D$. Clearly, we could construct a specific kernel to single out $C \wedge D$, but this will only shift the problem towards other logic combinations which cannot be represented. In other words, classifiers whose geometric structures lack homogeneity show shortcomings when trying to learn multiple concepts which are interrelated by an underlying logic structure.

Note that in order to support arbitrary concept disjunctions, full negation is required as both are tied together by De Morgan's law. The problem discussed above applies to any form of classifier in which the expressiveness of the geometric and logic structures is not balanced.

Other approaches, such as prototype theory, encounter issues with the representation of logical operators geometrically (as argued above) [Schockaert and Gutiérrez-Basulto, 2021] and an interpretation solely based on the orthogonality relation, though it allows for the representation of conjunction and disjunction, does not incorporate explicitly any geometrical structure and relies for concept definition on $\perp$-closure, thus on negative not on positive information.

There are approaches considering the geometrical structure appearing in the activation space of a neural network (thus the set of vectors being the output of each layer in a fully connected network) and use the convex hull to increase the result quality [Jia et al., 2022]. However, they neither explicitly incorporate background knowledge in form of rules nor consider negation (nor logical operators in general) in a geometrical sense. There are other approaches in this direction. The focus of these, however, lies on post-explainability of the resulting structure and they are thus not helpful for our problem. One possible solution along this line is to include the symbolic information in an additional step, as has been done, e.g., in logic tensor networks [Badreddine et al., 2022], where a neural network is trained of which in the last layer, the result is interpreted logically. Although this leads to good results, it results in a separation of the geometrical and the symbolic representation.

These considerations give a small overview and show that, though these approaches use similarity and some form of geometric structure, they are only distantly related to our aim of modeling logical operators geometrically and modeling concepts explicitly by specific and expressive geometric structures.

4.2 Concepts and Convexity

Though the notion of similarity enables the creation of learning approaches with an acceptable performance, the considerations on learning with similarity in the last section show a severe drawback: it is problematic when the expressivity of geometric and logic structures is not balanced.

Therefore, the problem must be approached from a different direction, based on the goal of balancing the geometrical and logical structures. The solution is loosely based on the conceptual spaces of Gärdenfors [2000] as discussed in Section 2.3: The concepts are embedded as regions into the geometric space on which the similarity relation is defined. Thus, the concepts are regions in a space X and a representation of an element e lying in the representation of a concept C is interpreted as e being of concept C. Having this as a basis, the first question is how the geometric representation of such concepts should be defined. To balance geometric and logic structures, these representations are not allowed to be arbitrarily shaped. When considering actual embedding approaches allowing for the geometric representation of concepts, representations are, e.g., spheres [Kulmanov et al., 2019], boxes [Abboud et al., 2020], cones [Bai et al., 2021], subspaces of a Hilbert space [Garg et al., 2019], general convex regions [Gutiérrez-Basulto and Schockaert, 2018], thus are mostly convex structures. This practice justifies the theoretical considerations about the specific shape of concepts done, e.g., by Gärdenfors [2000]. Concretely, he argues that a suitable shape of concepts is induced by the convexity property. This is explained as following a principle of cognitive economy [Gärdenfors, 2000, p. 70] (which is underlined by the computational advantages of convexity, e.g., in convex optimization). Moreover, convexity also enables a simple type of prediction, namely when instances a and c are in concept C, then also everything in between a and c needs to be in concept C. Although missing convexity does not necessarily lead to overfitting, enforcing a simple structure (as convexity) leads to greater induction capabilities (see, e.g., considerations on SVMs where having more support vectors leads to a more complex structure and thus to overfitting [Cristianini and Shawe-Taylor, 2000]).

Though Gärdenfors' considerations mainly refer to Euclidean convexity, his considerations are not restricted to it. He proposes that concepts should be convex regarding some betweenness relation $\mathcal{B}(a, b, c)$ in the following way: a set $C \subseteq X$, where X is an arbitrary domain, is *convex*, for short $isConv(C)$ iff for all $a, c \in C$ with $\mathcal{B}(a, b, c)$ (reading "b is in between of a and c") it follows that $b \in C$ [Gärdenfors, 2000, p. 69]. This can be formulated with a closure operator $conv(C)$ via $C = conv(C)$. The set $conv(C)$ is sometimes called the *convex hull* of C and can be defined as the smallest superset of C that is convex:

$$conv(C) = \bigcap \{A \supseteq C \mid isConv(A)\}$$

Convexity can be seen as a special type of closure operator according to the usual definition which we restate here.

Definition 4.1. *[Faure and Frölicher, 2000, Definition 1.2.5] A* closure operator *on a set X is a map $cl : Pow(X) \to Pow(X)$ satisfying the following two axioms:*

1. *$A \subseteq cl(A)$ for every $A \subseteq X$*
2. *$A \subseteq cl(B)$ entails $cl(A) \subseteq cl(B)$ for $A, B \subseteq X$*

Convexity is strongly dependent on the choice of the betweenness relation. Whereas a Euclidean betweenness leads to the well-known Euclidean convexity, there are other quite different forms resulting from non-Euclidean betweenness. There are many different betweenness relations but they all need to fulfill some basic axioms. These axioms and some examples of betweenness relations are discussed in the next section.

4.3 Betweenness

Betweenness plays a major role in concept formation and is a fundamental notion in geometry in general. Broadly spoken, betweenness is a ternary relation stating that some element is *in between* two other elements. In the Euclidean sense, betweenness depicts one element lying in between the other two on a line. Together with the notion of equidistance it allows for developing the whole elementary geometry axiomatically and, thus, can be considered as a basic notion of elementary geometry [Balbiani et al., 2007, p. 345-346]. However, betweenness is not restricted to Euclidean spaces. It is, e.g., possible to define lattice betweenness and to use the properties of this betweenness to conclude whether a lattice is distributive or modular [Pitcher and Smiley, 1942, p. 13]. Thus, betweenness is in general a way of enriching the underlying structure of a space. One main benefit of betweenness is that it allows for defining convex sets, namely as betweenness-closed sets, and enables utilizing the advantages of convex optimization (especially in but not limited to) Euclidean spaces.

The mathematically formal treatment of betweenness can be dated back to Huntington and Kline [1917] and is based on an axiomatization of geometries by Hilbert [1913]. There are several different axiomatizations of betweenness, focusing on different usage areas. In the following, the most basic properties of betweenness are given which are universal in the sense that they can be used (mostly) independently of the underlying space, stated, e.g., by Pitcher and Smiley [1942] and Gärdenfors [2000] and introduced by Huntington and Kline [1917].

(A1) $\mathcal{B}(a, b, c)$ iff $\mathcal{B}(c, b, a)$ (symmetry in the end points)

(A2) If $a \neq b \neq c$[2], then not both $\mathcal{B}(a, b, c)$ and $\mathcal{B}(a, c, b)$ (closure)

Axiom (A1) expresses symmetry, i.e., if a point b is in between a and c, then b is also between c and a. Axiom (A2) states that if a point is between two points it cannot have one of the points in between itself and the other point.

Next to these axioms, the following two axioms are the only remaining axioms compatible with (A1) and (A2) including three elements [Huntington and Kline, 1917]. One axiom depicts that each three points need to be connected via a betweenness relation (which seems to be unreasonable in our context, as it is in many cases not the case, e.g., when interpreting betweenness in the Euclidean sense [Pitcher and Smiley, 1942], and here the focus lies betweenness relations as general as possible). The other axiom is

(A3) If $\mathcal{B}(a, b, c)$ then a, b and c are distinct

[2] Here and in the following, $a \neq b \neq c$ is used as a shortcut for $a \neq b$, $b \neq c$ and $a \neq c$.

A betweenness fulfilling (A3) is called in the following *open*, one not fulfilling it is called *non-open*. These three axioms can be interpreted as the basics for a betweenness-relation. The validity of (A3) is a matter of definition whether an element can be in between itself or not.

Considering axioms containing four elements, there are eight axioms [Huntington and Kline, 1917] of which not all are appropriate for the general case. Huntington's and Kline's [1917] axioms 4. to 8. are based on the assumption that a betweenness relation makes up a line (in the usual, Euclidean sense) and that each two points are connected via a single line. These line-like properties are considered in detail in the next subsection, as they are too restrictive for the general case. From the remaining three axioms, the third one can be derived by using (A1) and the first two axioms, thus, only the first two axioms need to be considered:

(A4) If $\mathcal{B}(a, b, c)$ and $\mathcal{B}(b, c, d)$, then $\mathcal{B}(a, b, d)$

(A5) If $\mathcal{B}(a, b, d)$ and $\mathcal{B}(b, c, d)$, then $\mathcal{B}(a, b, c)$

Axioms (A4) and (A5) are sometimes called *outer transitivity* and *inner transitivity*. The list of betweenness axioms can be extended to include axioms with five or more points but these axioms are due to their complexity not further considered here.

Betweenness occurs in many different contexts and plays a major role in concept formation, in particular when using the idea of conceptual spaces of Gärdenfors [2000]. There, betweenness allows for defining convexity and thus enables establishing a psychologically grounded definition of concepts and properties in a geometric space.

Gärdenfors introduced the following betweenness axioms from which (B0)–(B4) are equivalent to the basic axioms mentioned before (including (A3)) and are mentioned here for completeness. In the following, the tags (B0)–(B4) will be used when referring to the betweenness axioms:

Definition 4.2. *Betweenness axioms according to Gärdenfors [2000, pp. 15 ff] are the following, based on a set X and $a, b, c, d \in X$*

(B0) *If $\mathcal{B}(a, b, c)$, then a, b, c are distinct*

(B1) *If $\mathcal{B}(a, b, c)$, then $\mathcal{B}(c, b, a)$*

(B2) *If $\mathcal{B}(a, b, c)$, then not $\mathcal{B}(b, a, c)$*

(B3) *If $\mathcal{B}(a, b, c)$ and $\mathcal{B}(b, c, d)$, then $\mathcal{B}(a, b, d)$*

(B4) *If $\mathcal{B}(a, b, d)$ and $\mathcal{B}(b, c, d)$, then $\mathcal{B}(a, b, c)$*

and as additional (possibly valid) axiom

(B5) *For any two points a and c in X, there is some point b such that $B(a, b, c)$.*

(B5) is an axiom not mentioned before and ensures density, making sure that between all two elements lies at least one element. This axiom is not always appropriate, as it is not applicable in discrete spaces, but is helpful as, unlike (B0)–(B4), it is a constructive axiom: whereas the other axioms can be fulfilled with an empty betweenness relation, this one can not. This axiom and some variants of it are considered further at the end of this section, because an axiom circumventing trivial betweenness relations will be of importance for the definition of COSs. (When considering an empty betweenness relation there is no sense in reasoning about convexity).

Having the betweenness relation, it is possible to define some simple notions based on it, especially the notion of half line and line. Let $a, b \in X$ be arbitrary, then define the line $L(a, b)$ through a, b (as has been done by Aboulker et al. [2016] for pseudo-metric betweenness) as

$$L(a, b) := \{x \mid x = a \text{ or } x = b \text{ or } \mathcal{B}(a, x, b) \text{ or } \mathcal{B}(x, a, b) \text{ or } \mathcal{B}(a, b, x)\}$$

Given a line L and points a, b on L, one has two disjoint open half lines $HL(a, b)$ and $HL^*(a, b)$ starting at a such that $L = \{a\} \cup HL(a, b) \cup HL^*(a, b)$. These can be defined as follows:

$$\begin{aligned} HL(a, b) &= \{x \mid \mathcal{B}(a, x, b) \text{ or } \mathcal{B}(a, b, x) \text{ or } x = b\} \\ HL^*(a, b) &= \{x \mid \mathcal{B}(x, a, b)\} \end{aligned}$$

Observe here that this definition does not necessarily lead to a line in the usual (Euclidean) sense, as can be seen in Figure 4.2 where all three examples represent valid lines but only the left one represents a line in the Euclidean sense. The middle example contains two elements b and b' both in between a and c but not interfering with each other. In the right example, the line has two distinct endpoints. Thus, in order to define a line in the Euclidean sense, it is necessary to define additional axioms for betweenness which are later given as (C2) and (C3).

Obviously, similarity (introduced in Section 4.1) and betweenness are highly related, or, to state it differently, similarity can be used to define a betweenness-relation. However, an element that is in between two others is not necessarily similar to both (or even to one of them) and vice versa.

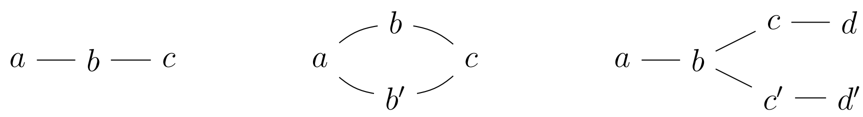

Figure 4.2: Three examples of lines $L(a, b)$ (where the edge indicates that from left to right a betweenness relation is given), the left one also being a line in the Euclidean sense.

Examples for Betweenness Relations

Betweenness relations are used in many fields, e.g., for antimatroids [Chvátal, 2009] and for lattices [Pitcher and Smiley, 1942, p. 13].

Three of the betweenness relations are of greater importance for the construction of COSs and thus are discussed in detail in the following: metric betweenness, intersection betweenness and betweenness based on Goodman's matching.

Metric betweenness The widest known betweenness-relations are in the group of *metric betweennesses* $\mathcal{B}_m$:

$$\mathcal{B}_m(a, b, c) \text{ iff } d(a, b) + d(b, c) = d(a, c) \text{ for } a, b, c \in X \text{ with } a \neq b \neq c,$$

where X is the domain and d is an arbitrary metric, e.g., the Euclidean, manhattan or polar distance [Gärdenfors, 2000, p. 72] in $\mathbb{R}^n$ or the angle where X is a n-ball. Whereas betweenness based on the Euclidean metric fulfills the axioms (B0)–(B5), the one based on an angle does not, as it does not fulfill (B3). This underlines the fact that the axioms (B0)–(B2) belong to the minimal set of axioms to be fulfilled by a betweenness relation. Metric betweenness based on an angle is in the following termed $\mathcal{B}_a$ (and used both for the case where X is a sphere and for the case where $X = \mathbb{R}^n$, though it is in this case not a metric betweenness).

The Euclidean metric and thus the Euclidean betweenness $\mathcal{B}_E$ plays a special role, as in the area of Euclidean geometry, betweenness has great importance, e.g., for the axiomatization of Tarski mentioned above. This also applies to COSs, as the Euclidean space with the Euclidean metric is a natural and reasonable research domain. Chapter 7 relates to COSs in the context of Euclidean betweenness. This betweenness fulfills (B0)–(B5) as well as other axioms, some of which will be discussed here, as they lead to a radical simplification of the resulting set of possible betweenness relations because of their strict constraints. One axiom [Borsuk and Szmielew, 2018, Axiom O5], complementing (B5) is

(C1) For any two (distinct) points a, b in X, there is some point c in X such that $\mathcal{B}(a, b, c)$.

Whereas (B5) ensures density in between two points, (C1) ensures density on the outside of two points. Considering the following axioms [Huntington and Kline, 1917, axioms 4. and 6.] leads to a restriction of the lines considered above. These axioms in particular enforce that the middle and right figure of Figure 4.2 do not represent valid betweenness relations anymore.

(C2) If $\mathcal{B}(a, b, c)$ and $\mathcal{B}(a, b', c)$, then either $\mathcal{B}(a, b', b)$ or $\mathcal{B}(a, b, b')$

(C3) If $\mathcal{B}(a, b, c)$ and $\mathcal{B}(a', b, c)$, then either $\mathcal{B}(a, a', c)$ or $\mathcal{B}(a', a, c)$

It can be proven that these axioms are valid for Euclidean betweenness [Borsuk and Szmielew, 2018, Theorem 18].

Intersection betweenness A betweenness relation which will be of great importance for some considerations on COSs is the *intersection betweenness*. The intersection betweenness enables defining betweenness on sets and is grounded on psychological considerations [Restle, 1959]. It is based on the idea that objects can be compared based on sets of their features. An object b is said to be in between a and c if it includes all the features common to a and c (but is allowed to have features which neither a nor c have) [Tversky and Gati, 1982]. Intersection betweenness can also be abstractly viewed as a betweenness relation on sets. When considering sets, some standard betweenness relations such as metric betweenness cannot be defined, as there is only qualitative information about the sets given. To break up this qualitative definition, a notion of order of sets based on betweenness is given. This notion was introduced by Restle [1959] and is also mentioned, e.g., by Tversky and Gati [1982], Burigana [2009] and Rautenbach et al. [2011]:

Definition 4.3. *[Rautenbach et al., 2011] For a set system $\mathcal{M} = (M_v)_{v \in V}$ indexed by the elements of a finite set V, the* intersection betweenness $\mathcal{B}_i$ *induced by $\mathcal{M}$ is defined as follows:*

$$\mathcal{B}_i(a, b, c) \text{ iff } M_a \cap M_c \subseteq M_b$$

The intersection betweenness is called *strict* if a, b and c are distinct.

Proposition 4.4 ([Restle, 1959, Theorem 3], [Burigana, 2009])**.** *Strict intersection betweenness $\mathcal{B}_i$ fulfills (B0)–(B1) but neither of (B2)–(B4).*

Though (B2)–(B4) are, in general, not fulfilled, it is possible to find instances of intersection betweenness also fulfilling (B2)–(B4). The restrictions necessary for fulfillment are mentioned by Burigana [2009] and Restle [1959] and are considered in some detail in the context of an orthonegation-induced betweenness relation in Section 6.1.

An example for the use of intersection betweenness in the context of DL is given by Schockaert, Ibáñez-García, et al. [2021]. Their aim is to model interpolation in DL, thus modeling that concept B is in between concepts A_1 and A_2. Therefore, they used a feature-based semantics and, thus, map each instance on a set of features with a mapping π. Betweenness is defined via intersection betweenness as

$$\mathcal{B}_i(a, b, c) \text{ iff } \pi(a) \cap \pi(c) \subseteq \pi(b).$$

Goodman's matching A betweenness-relation situated in a quite different context is the notion of *matching* introduced by Goodman [1977] and discussed in my publication [Leemhuis and Özçep, 2022]. In note 20 of his book on conceptual spaces [Gärdenfors, 2000], Gärdenfors remarks that Nelson Goodman in his work on the "Structures of appearance" [Goodman, 1977] established a notion of betweenness based on the notion of matching. Matching is a binary, symmetric and reflexive relation and thus can be interpreted as similarity relation and its negation as a dissimilarity, thus an orthogonality relation can be defined. Gärdenfors notes that Goodman meant to show that he can establish many properties of the betweenness relations, which suggests that not all of (B0)–(B4) are fulfilled. In fact, Goodman's book includes only proofs showing that (B0)–(B2) are fulfilled.

The ideas of Nelson Goodman concern a betweenness relation of so-called qualia, abstract phenomenal qualities, describing a sensory impression of a property under a specific condition, e.g., an impression that something appears to be green based on a given illumination. This restriction to qualia is, in fact, not a problem, as his general construction can also be used for non-qualia (as stated by himself and as shown in Section 6.1.2 for the adaptation). The theory of qualia is developed in a mereological system (a theory of parts and wholes) with a mereological sum operator (see, e.g., [Simons, 1987] for a modern, book-length treatment). I adapt Goodman's approach to a set theoretical setting (replacing sum by union). Though mereology has its merits and a long tradition in epistemology, metaphysics, and ontology research there are technical reasons why it has not found the same acceptance as set theory — at least by mathematicians [Hamkins and Kikuchi, 2016].

Goodman defined his notion of betweenness only for finite domains. This is the other point in which the adaptation of Goodman's construction deviates from the original approach: the adaptation considers also infinite (and even more: dense and continuous) spaces as they are used in general by embedding approaches.

The relevant betweenness relation of Goodman is called “betwixtness” (an old English word for betweenness) and reads as follows [Goodman, 1977, Definition D10.02]:

$$x/y/z \quad = \quad M(x,y) \;\&\; M(y,z) \;\&\; M(x,z) \;\& \\ G(x\dagger z, x\dagger y) \;\&\; G(x\dagger z, y\dagger z)$$

$M(x,y)$ stands for Goodman’s reflexive and symmetric notion of matching. This notion seems to be a primitive notion, i.e., to be undefined. The definition Goodman gives for matching [Goodman, 1977, Definition DA-2] relies on the notion of allying “A”. This notion in turn is defined on the basis of M. M is given only “implicitly” defined by some axioms [Goodman, 1977, pp. 209 ff]. Goodman has a mereological notion of symmetric difference denoted by the dagger † (p. 211) and a binary notion of aggregatively greater than $G(x,y)$. [Goodman, 1977, Definition D8.021] . This one is based on a primitive notion of “is of equal aggregate size” Z. (Translated to set theory Z just means: “has the same cardinality”). Definition D8.021 reads: $G(x,y)$ iff there is a proper part t of x such that t is equal in aggregate size to y. This seems to be the main place where Goodman assumes a finite domain: Because, if x were infinite, then one could find a proper part (a proper subset of x) that is of equal size as x. In that case, however, x would be greater than itself.

According to Goodman [1977, p. 219], his betweenness relation fulfills the following axioms (with the corresponding enumeration in Goodman’s book given in brackets): (B0) (Goodman: 10.25), (B1) (Goodman: 10.26), (B2) (Goodman: 10.27). There is no proof of (B3), which this adaptation does not fulfill either, and of (B4), which, in case of this adaptation, is provable.

(B3) is problematic due to several reasons. One problem is that each pair of x, y, z must match. This cannot be guaranteed: Whereas each pair in $\{a, b, c\}$ may match and each pair in $\{b, c, d\}$ may match, a and d may not match anymore. However, even when matching of each pair is assumed, (B3) is not necessarily fulfilled, as, e.g., the creation of circles in matching (meaning $M(a,b), M(b,c), M(c,d)$ and $M(d,a)$) could lead to the case that $\mathcal{B}(a,b,c)$ and $\mathcal{B}(b,c,d)$, but $\mathcal{B}(d,a,b)$ instead of $\mathcal{B}(a,b,d)$. In this paragraph, the original definition of betweenness by Goodman was presented. Its aforementioned adaptation is given in the context of orthonegation-induced betweenness in Section 6.1.2.

An Axiom of Density

This section ends with some considerations on the axiom of density (B5). This axiom (and also (C1)) is different from the other axioms. Whereas the others can be satisfied by considering the empty betweenness relation, this is not possible for these two axioms. Even when not considering the trivial betweenness relation, it is still possible to circumvent problems, e.g., with (B3) and (B4) by ensuring that each combination of two elements occurs only in one betweenness relation. Thereby, the premise would not be fulfilled and (B3) and (B4) would be trivially true.

Adding (B5) dismisses this possibility, as for each a and c in X, there must be some point b in X such that $\mathcal{B}(a, b, c)$ is valid. Thus, there must also be an d with $\mathcal{B}(a, d, b)$ etc. and the premises of the axioms (B3) and (B4) are fulfilled and the axioms need to be non-trivially fulfilled.

Thus, besides density, (B5) and (C1) additionally ensure some sort of connectedness: It is not possible to have different clusters of independent betweenness relations.

In particular when arguing about convexity, it is not desired to have trivial (or nearly trivial) betweenness relations or different clusters. The notion of convexity would lose its expressivity, because in that case it would be usable only locally. But requiring (B5) or (C1) or both to be valid for a betweenness relation is also not always reasonable as it requires "inner" resp. "outer" density which is not possible in discrete spaces.

To keep the advantages of this axiom without requiring the considered sets to be non-discrete, here a new axiom is presented combining the axioms (B5) and (C1):

(B5') For any two points a and c in X, there is some point b such that $B(a, b, c)$ or $\mathcal{B}(b, a, c)$ or $\mathcal{B}(a, c, b)$.

This is a relaxation of axiom I3 [Borsuk and Szmielew, 2018, p. 21], based on Euclidean betweenness, where it is stated that for all two distinct points there need to be at least one line through both of them. As for a general betweenness, a line can not be defined as in the usual Euclidean sense, here this relaxation is used.

4.4 Approaches Based on Similarity and Betweenness

The last sections showed that both similarity and betweenness are vital for an embedding approach with sufficient capacity of representing logical information geometrically. The next natural question is what such a combination could look like. Before proposing the framework of Conceptual Orthospaces (COSs) in Chapter 5, in the following, some approaches are surveyed which are related to an embedding approach based on similarity and betweenness. As shown in Sections 4.1 and 4.3, similarity and betweenness are both extensively researched areas. Therefore, in the following, no complete survey of related approaches is given. Instead, only a few examples are showcased which are either highly popular or closely related to the approach that will be proposed in Chapter 5.

One paradigm, incorporating similarity and betweenness as the basis are conceptual spaces as mentioned in Section 2.3. Although the main theory of conceptual spaces is more of a paradigmatic nature, there are also several practical approaches using conceptual spaces for learning and reasoning. For example, Bechberger and Kühnberger [2017] propose an approach based on star-shaped regions, able to model conjunction and disjunction and having the usage areas of concept combination and clustering. The framework, however, is not able to model conjunction and disjunction correctly in a logical sense, as the star-shaped regions they consider are not closed under conjunction and disjunction. Therefore, they need a repair mechanism and thus lose the direct interpretability of logical operators as geometric ones. Additionally, they argue that, depending on the type of convexity chosen, convexity is not necessarily the best option for modeling conceptual spaces, but they do not argue against convexity in general. Though their approach has the advantage of being simple and psychologically grounded, they are not able to reach our goals, namely convexity and modeling the logical operators as geometrical ones.

A different approach, also enabling the use of conceptual spaces in a practical context, was proposed by Jameel and Schockaert [2016]. There, an embedding into a vector space is considered where semantic types (thus concepts) are modeled as subspaces of this vector space. This allows for modeling type hierarchies and relations between entities of the corresponding types. Up to this point, the approach is quite similar to that of Garg et al. [2019] in which subspaces are also embedded. But in contrast to the work of Garg et al. [2019], here the dimensions can be considered as quality dimensions in the sense of conceptual spaces, thus having a semantic meaning. Though, this approach models concepts as convex regions, it does not incorporate any negation. For computational simplicity (and to allow for modeling negation as orthonegation), I neither consider explicit domains nor quality dimensions with a semantic meaning, as both add to the complexity of the learning approach [Bouraoui et al., 2022]. This differentiates COS from conceptual spaces. The aim of the COS is not psychological plausibility but rather learnability and simplicity of the representation. Therefore, the choice of the similarity should not be based on psychological grounds but based on simplicity and learnability.

Many approaches in the area of KGE that incorporate background knowledge geometrically use convexity (for a thorough discussion, see Section 10.2) and of course embedding is based on similarity. They, however, are focused on Euclidean convexity, do not consider the attributes of convexity and similarity in detail and are constrained to a specific geometric structure, whereas here the aim is to find a general form. Furthermore they mainly do not incorporate negation.

A different approach based on convexity is presented by Gutiérrez-Basulto and Schockaert [2018]. There, the same goal as here was set, namely to enable for an embedding approach that is on the one hand as expressive as possible and on the other hand enables defining the ontological commitment. They also identify convex regions as the basic geometrical structure both for the modeling of concepts and relations. They showed that an ontology can be embedded based on their convex regions when it fulfills the so-called *quasi-chainedness property*, but they do not consider disjunction and negation and restrict themselves to Euclidean spaces. This approach underlines the usability of convex regions, however, does not consider similarity (and thus in particular not dissimilarity) directly.

An area in which similarity in context of concept formation plays a vital role is the area of *similarity description logic* [Sheremet et al., 2007] and related areas, e.g., *fuzzy DL* [Borgwardt and Peñaloza, 2017] and rough DL [Peñaloza and Zou, 2013]. The similarity DL $\mathcal{SL}$ was introduced by Sheremet et al. [2007]. It enables enhancing the usual DL-axioms with information and reasoning regarding similarity of concepts. They introduce the quantifiers $\exists_s^{<a} C$, denoting all instances with a similarity of at least a to C and $C \leftleftarrows_s D$, including all instances being more similar to C than to D based on some similarity space. Thus, the similarity is considered based on different aspects (resp. different domains). Though this combines conceptual reasoning with a notion of similarity and enables for reasoning based on similarity (and allows also for the notion of closure and interior [Alenda and Olivetti, 2012]), it does not consider the shape of the concepts and thus does not enforce a geometric structuredness like convexity. Additionally, it only considers negation as set-complement. Though similarity is considered, it does not influence concept formation.

The approach most similar to the framework of COSs was introduced by Schockaert, Ibáñez-García, et al. [2021]. They pursued a different goal but came to results leading in a similar direction. Their approach is again situated in the area of similarity DL and they consider the enhancement of deductive reasoning based on the description logic $\mathcal{EL}_{\bar{0}}$ via the consideration of similarity information between different concepts. To accomplish this, they use an intersection betweenness on the features of these concepts. Convexity of concepts is enforced, however, the description logic $\mathcal{EL}_{\bar{0}}$ is considered as basis, and, thus, the problems regarding negation are not considered.

These approaches underline the relevance of betweenness in embedding approaches, e.g., in the context of conceptual spaces, not only by being the basis of convex concepts but also to allow for identifying concepts being in between other concepts [Derrac and Schockaert, 2015]. All these notions, though they incorporate some information about similarity, do not address all problems of the geometric regularities needed to ease reasoning and are not able to incorporate negation sufficiently.

One area in which negation can be modeled in a geometric (or at least topological) space, is spatial logic. Smyth and Webster [2007] present an approach based on closure spaces which is able to model negation. However, it considers closure spaces in general and does not consider convexity. The same problem appears for orthomodels. Though, they do not incorporate similarity information directly, it is possible to interpret incompatibility as inverse similarity. They can be also geometrically interpreted but do not contain any information about the shape of the concepts in form of convexity.

5 Conceptual Orthospaces as a Foundational Framework for Embeddings

As pointed out in the last chapters, Ne-Sy AI can be used to combine symbolic and subsymbolic learning approaches and, in this context, embeddings are widely used and work well. To enhance these embeddings with background information, it is necessary to model the concepts and the relations between concepts incorporated in the background knowledge geometrically. In Chapter 3, I decided to model negation with the help of the orthogonality relation and in Chapter 4 it was shown that similarity and convexity are basic notions for embedding approaches. These ideas are now combined to approach the main framework of this dissertation, the *Conceptual Orthospaces*.

At first, the definition of orthogonality as irreflexive and symmetric relation depicting incompatibility comes into play. As similarity is a reflexive and symmetric relation (see Section 4.1), the orthogonality relation can be considered as dissimilarity. This not only matches syntactically but also semantically, as instances are orthogonal to each other when they are sufficiently dissimilar. This enables modeling the orthogonality relation based on the similarity relation and thus combines the two notions.

By now, the framework that is to be created has the building blocks of a similarity relation (enabling the representation of the orthogonality relation), modeling conjunction as set-intersection and disjunction as $\perp$-closure. The remaining and most important question is how to represent concepts in this framework. As discussed in Section 4.2, they should be convex. However, this restriction is not enough. It is also necessary to balance convexity and similarity to allow for a suitable definition that enables modeling concepts that are both convex and $\perp$-closed. The next example depicts a possible use case.

Example 5.1. *Assume a user would like to have concepts that are convex in the sense that they fulfill Euclidean betweenness as the basis for his embedding approach. Now, he has to think about a suitable definition of negation. Figure 5.1(a) depicts an example, showing two concepts A and B in $\mathbb{R}^2$ that are convex w.r.t. Euclidean betweenness.*

Now, assume that the similarity relation has been chosen in a way leading to $A^\perp$ *and* $B^\perp$ *as shown in the figure. This embedding has several drawbacks: If we assume that the concepts are a result of a learning approach, then, as* $B^\perp$ *has in particular a complicated structure, it is prone to overfitting and thus not suitable for inferring new information. Though missing convexity does not necessarily lead to overfitting, enforcing a simple structure (as convexity) leads to greater induction capabilities (see, e.g., considerations on SVMs where having more support vectors leads to a more complex structure and thus to overfitting [Cristianini and Shawe-Taylor, 2000]).*

It can be argued that this problem is based on an artificial choice of a similarity relation. However, though the choice of a suitable similarity relation is important, it does not completely solve the problem. For illustration consider Figure 5.1(b): The similarity measure is chosen as quite a natural one, namely the Euclidean distance, and the negation is defined based on a threshold (thus, $a \perp b$ *iff* $d_E(a, b) \geq \tau$ *for* $a, b \in \mathbb{R}^n$*, a fixed* $\tau > 0$ *and* d_E *being the Euclidean distance). But even then it turns out that, despite* A *having a suitable structure,* $A^\perp$ *is not convex. This leads to a situation in which* A *and* $A^\perp$ *are not treated equally, which in turn means that, for a learning approach, a mapping to* A *is learned based on a well-formed convex structure, whereas the mapping to* $A^\perp$ *leads to a non-convex structure representing the absence of information on* A*.*

To sum up, the choice of an orthogonality relation is not arbitrary but dependent on the goals of the modeling and the chosen betweenness relation. Therefore, a framework that allows to define a suitable orthogonality relation for a given betweenness relation is needed. However, even then, it is not possible to consider all (or arbitrary) concepts, convex based on this betweenness relation. In fact, the definition of an expressive orthogonality relation could restrict the set of possible convex concepts. For example, Chapter 7 will show that for Euclidean betweenness and an expressive orthogonality relation, the only possible concept interpretations are some variants of convex cones.

Whereas in the example, the case of a given betweenness relation and the induced orthogonality relation was considered, another realistic scenario is that a similarity relation is chosen beforehand, e.g., an orthomodel is given. Though orthomodels are expressive structures to model concepts in a space equipped with negation, these concepts, though they are $\perp$-closed, are not necessarily well-shaped, e.g., convex in the Euclidean sense. For example, when creating an orthomodel based on $X = \mathbb{R}^n$, the concepts could be arbitrarily shaped in a geometric sense, even if they are $\perp$-closed as can be seen in Figure 5.1(a). What becomes relevant here is that the betweenness relation is not necessarily defined in the classical sense, thus as Euclidean betweenness. This enables determining a (weak) betweenness relation such that all $\perp$-closed sets are convex. Although this does not lead to the same expressivity in the sense of reduction of overfitting and com-

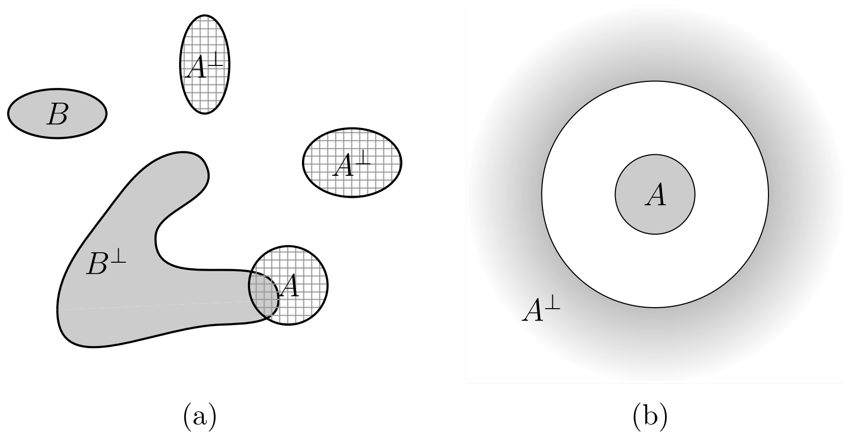

Figure 5.1: (a) Example of an orthomodel in $\mathbb{R}^2$ with two concepts A and B and an orthogonality relation that is not distance-based; (b) an example of an orthomodel in $\mathbb{R}^2$ with a concept A and a negation based on the Euclidean distance.

putational complexity, it at least eases the learning, as the structure of the concepts is examined and thus can be used, e.g., for optimization.

Thus, in general, there are two possibilities to model the connection between similarity and betweenness. The case beginning with similarity and the one starting with betweenness. The aim is not to define a framework modeling a betweenness leading to the strongest possible convexity and also not to determine the perfect orthogonality relation, but to give an underlying framework that enables designing an embedding approach accounting for logical operators, especially negation, based on either a given similarity or betweenness.

But before determining the details of the construction, in the following the general framework of the combination of similarity and betweenness is presented. This is then used as basis for the construction of orthonegation-induced betweenness in Section 6.1 and betweenness-induced orthogonality in Section 6.2 and throughout the rest of the dissertation.

The basic definition of the framework of Conceptual Orthospaces is given in Section 5.1. It will turn out that for different use cases different restrictions are necessary to gain sufficient expressivity. This is discussed in Section 5.2. The section is concluded by general considerations on the properties of COSs. This chapter is an extended version of my work presented in a journal publication [Leemhuis and Özçep, 2023].

5.1 Formalization

The main ingredients of COSs are the convexity of concepts based on a betweenness relation (not necessarily fulfilling all betweenness axioms) and an orthogonality relation.

Formally, assume that an arbitrary domain X containing elements/instances, a ternary betweenness relation $\mathcal{B}$ over X and a binary orthogonality relation $\perp$ over X are given. Additionally, there is a set Y of convex and $\perp$-closed sets based on X, with $Y \subseteq Pow(X)$, where $Pow(\cdot)$ is the powerset of X. The set Y can be seen as containing the subsets of X that make up the concepts, thus those entities for which the operators are defined. Whereas for an orthomodel all $\perp$-closed sets can be considered, here it is necessary to consider sets that are both $\perp$-closed and convex.

Definition 5.2. *A* Conceptual Orthospace (COS) *is a structure $S = (X, Y, \mathcal{B}, \perp)$ with the following components:*

- *a domain X,*
- *a ternary betweenness-relation $\mathcal{B}$ on X (not necessarily fulfilling all betweenness axioms),*
- *an orthogonality relation $\perp$ on X and*
- *a set $Y \subseteq Pow(X)$ where each $A \in Y$ is $\perp$-closed (based on $\perp$, for short $cl_\perp(A) = A$) and convex (based on $\mathcal{B}$, for short $conv(A) = A$) and which is closed under conjunction and negation, where for $A, B \in Y$*
  - *conjunction of A and B is defined as set-intersection $A \cap B$ and*
  - *negation of A is defined as $A^\perp$.*

The set Y is called complete *iff $Y = \{conv(A) \mid A \subseteq X\} = \{cl_\perp(A) \mid A \subseteq X\}$, i.e., iff each $\perp$-closed set in X is convex and vice versa.*

It is possible to interpret COS as DL interpretation, thus the elements of Y can be interpreted as representation of concepts and X as domain $\Delta^\mathcal{I}$. Using a COS as basis for an interpretation enforces the concept representations to be convex and $\perp$-closed and thus have the structure which was identified as suitable in the last chapters.

Definition 5.3. *Given a DL-vocabulary $\mathcal{V} = N_c \cup N_C$ of constants and concept symbols, the* COS-interpretation *(also called* geometric interpretation *in the following) based on a COS $S = (X, Y, \mathcal{B}, \perp)$ is a structure $(\Delta^\mathcal{I}, \cdot^\mathcal{I})$ where $\Delta^\mathcal{I} = X$, $\cdot^\mathcal{I}$ is the* denotation function *defined for all* $b \in N_c, A \in N_C$ *and concepts* C, D *over* $\mathcal{V}$ *such that the following conditions are fulfilled:*

$b^\mathcal{I} \in \Delta^\mathcal{I}$, $\quad (C \sqcap D)^\mathcal{I} = C^\mathcal{I} \cap D^\mathcal{I}$,
$A^\mathcal{I} \in Y$, $\quad (\neg C)^\mathcal{I} = (C^\mathcal{I})^\perp$,
$\overline{1}^\mathcal{I} = \Delta^\mathcal{I}$, $\quad (C \sqcup D)^\mathcal{I} = (\neg(\neg C \sqcap \neg D))^\mathcal{I}$

The interpretation of $\overline{0}^\mathcal{I}$ *depends on the use-case. Possibilities are, e.g.,* $\overline{0}^\mathcal{I} = \emptyset$ *or* $\overline{0}^\mathcal{I} = \{0\}$ *for some element* $0 \in \Delta^\mathcal{I}$*. Details on this can be found in Section 5.2.1.*

The notion of a COS-interpretation being a model (of an abox, tbox, ontology) *is defined in the same way as for classical interpretations according to Section 2.1.*

In the following are concept symbol and concept representation named equally when the meaning is clear from the context. The set Y is intended to contain the concepts which provide the basis for an embedding approach. To construct a COS, it is necessary to follow some construction principles, as Y needs to be constructed in a way that it is closed under logical operators and contains sets that are both $\perp$-closed and convex. Therefore, it is not possible to choose the betweenness and similarity arbitrarily.

First, we tackle the question of how to define the the Boolean logical operators conjunction, disjunction and negation over Y. Negation and conjunction are already given in the definition. Disjunction of $A, B \in Y$, thus in DL-notation $(A \sqcup B)^\mathcal{I}$ can be defined via de Morgan as $(A \sqcup B)^\mathcal{I} = cl_\perp(A \cup B)$.

Corollary 5.4. $(A \sqcup B)^\mathcal{I} = cl_\perp(A \cup B)$.

Proof. By definition, it follows that $(A \sqcup B)^\mathcal{I} = (A^\perp \cap B^\perp)^\perp$. And as $A^\perp \cap B^\perp = \{x \mid \forall y \in A : x \perp y \text{ and } \forall z \in B : x \perp z\} = \{x \mid \forall y \in (A \cup B) : x \perp y\} = (A \cup B)^\perp$, it follows that $(A^\perp \cap B^\perp)^\perp = (A \cup B)^{\perp\perp}$ and by definition $(A \cup B)^{\perp\perp} = cl_\perp(A \cup B)$. □

Interpreting disjunction as the $\perp$-closure leads to the fact that Y, ordered by inclusion, makes up a lattice interpreting lattice meet as set-intersection and lattice-join as $\perp$-closure [Faure and Frölicher, 2000, Proposition 1.2.7].

The second closure operator needed is the convex closure *conv*, as all sets in Y need to be convex. The equality $conv(\cdot) = cl_\perp(\cdot)$ holds only for complete COS.

In the following, an example for a COS is given based on closed convex cones. Therefore, first, closed convex cones are formally introduced:

Definition 5.5 ([Schneider, 2022, p. 7])**.** *A* convex cone *is a subset* $C \subseteq \mathbb{R}^n$ *with* $x+y \in C$ *and* $\lambda x \in C$ *for every* $x, y \in C$ *and* $\lambda \geq 0$*. A* closed convex cone *is a nonempty convex cone which is topologically closed (w.r.t. the Euclidean norm). "Convex" in the context of cones is defined in the classical way, i.e., based on Euclidean distance.*

Example 5.6. *As an example of a COS, the set of closed convex cones in* $\mathbb{R}^n$ *(*$\mathcal{C}^n_{cone}$*) can be considered. Let the domain* $X = \mathbb{R}^n$*, the betweenness be the standard Euclidean betweenness* $\mathcal{B}_E$ *and the orthogonality relation be defined as* $x \perp_p y$ *iff* $\angle(x, y) \geq 90°$*. This can be reformulated by using the standard definition of the scalar product* $\langle \cdot, \cdot \rangle$ *as* $x \perp_p y$ *iff* $\langle x, y \rangle \leq 0$*. To determine whether* $S = (\mathbb{R}^n, \mathcal{C}^n_{cone}, \mathcal{B}_E, \perp_p)$ *makes up a COS, it is necessary to prove (i) that* $\perp_p$ *is actually an orthogonality relation, (ii) that each set in* $Y = \mathcal{C}^n_{cone}$ *is* $\perp$*-closed and convex and (iii) that the set of closed convex cones is closed under intersection and polarity. Regarding (i),* $x \not\perp_p x$ *for* $x \in \mathbb{R}^n$ *and* $x \perp_p y = y \perp_p x$*, thus* $\perp_p$ *is irreflexive and symmetric, thus, an orthogonality relation. (ii) is the case, as* Y *contains the set of closed convex cones which are* $\perp$*-closed and obviously convex. (iii) is, e.g., stated by Schneider [2008]. Thus,* $S = (\mathbb{R}^n, \mathcal{C}^n_{cone}, \mathcal{B}_E, \perp_p)$ *is a COS. An example for such a COS with* $X = \mathbb{R}^n$ *and* $Y = \{A, B\}$ *can be seen in Figure 5.2(a).*

Thus, the main point of finding a COS is to find the two closure operators *conv* and $cl_\perp$, where the former relates to the betweenness relation and the latter to the orthogonality relation. To ease this task, we consider the connection between *conv* and $cl_\perp$ in order to find some restrictions which both closure operators need to rely on. For an arbitrary set in Y, the convex closure of that set is always a subset of (or equal to) its $\perp$-closure. When Y is complete, then $\perp$-closure and convexity are equivalent. The other direction does not follow directly, as it is possible to define Y as a subset of a complete Y'. In this case, $\perp$-closure and convexity are equivalent, but Y is not complete.

Proposition 5.7. *For a given COS* $S = (X, Y, \mathcal{B}, \perp)$ *and* $A, B \in Y$*, the following relations between* $cl_\perp$ *and conv hold:*

1. $cl_\perp(A) = conv(A) = A$*, i.e.,* $A = A^{\perp\perp}$ *and for all* $a, c \in A$ *with* $\mathcal{B}(a, b, c)$ *also* $b \in A$
2. $conv(A \cup B) \subseteq cl_\perp(A \cup B)$
3. *If* Y *is complete, then for all* $C \subseteq X$*:* $cl_\perp(C) = conv(C)$
4. *If for all* $C \subseteq X$*:* $cl_\perp(C) = conv(C)$*, then there is a complete* Y' *with* $Y \subseteq Y'$

Proof. Let $S = (X, Y, \mathcal{B}, \perp)$ be given.

1. This follows directly out of Definition 5.2.
2. Let $A, B \in Y$ be arbitrary. As Y is closed under disjunction, $cl_\perp(A \cup B) \in Y$, thus, as each set in Y is convex, $conv(cl_\perp(A \cup B)) = cl_\perp(A \cup B)$. With the conditions of a closure operator in Definition 4.1, it follows that $A \cup B \subseteq conv(cl_\perp(A \cup B))$ and thus $conv(A \cup B) \subseteq conv(cl_\perp(A \cup B))$ and thus $conv(A \cup B) \subseteq cl_\perp(A \cup B)$.

3. Let $C \subseteq X$ be arbitrary and Y be complete. Then, as Y contains all $\perp$-closed sets, $cl_\perp(C) \in Y$ and with 2. it follows that $cl_\perp(C) \supseteq conv(C)$. Thus, it remains to show that $cl_\perp(C) \subseteq conv(C)$. As Y is complete, $conv(C) \in Y$. Analogously to 2. it follows that $cl_\perp(C) \subseteq conv(C)$.

4. This follows directly out of Definition 5.2.

□

In particular, the proposition entails that $\perp$-closure and convex-closure need to be chosen in a way that the $\perp$-closure of a set results in a set equal to its convex-closure or in a superset of the convex-closure. $\perp$-closure thus needs to be a more inclusive notion than convex closure.

It is, in fact, possible that the subset relation in Proposition 5.7.2 is strict so that $conv(A \cup B) \subsetneq cl_\perp(A \cup B)$ for $A, B \in Y$.

Example 5.8. *For the COS in example 5.6, $\perp$- and convex-closure of closed convex cones can be defined both as convex hull (in the Euclidean sense), thus $cl_\perp(A \cup B) = conv(A \cup B)$ for $A, B \in Y$. However, there are also cases where this equality of convexity and $\perp$-closure does not hold. Consider again $X = \mathbb{R}^n$ and so-called* convex pseudo-cones. *A non-empty convex set $K \subseteq \mathbb{R}^n$ not containing the origin is a convex pseudo-cone if $\lambda x \in K$ for $x \in K$ and $\lambda \geq 1$. Betweenness is defined as Euclidean betweenness $\mathcal{B}_E$ and the orthogonality relation as $x \perp_{ps} y$ iff $\langle x, y \rangle \leq -1$ for $x, y \in X$ based on the usual scalar product $\langle \cdot, \cdot \rangle$ [Xu et al., 2023]. Thus, a possible COS is $S = (\mathbb{R}^n, \mathcal{PC}^n, \mathcal{B}_E, \perp_{ps})$ where $\mathcal{PC}^n$ is the set of all convex pseudo-cones in X. $\mathcal{PC}^n$ is closed under conjunction and negation but not under disjunction if $(A \sqcup B)^\mathcal{I}$ is defined as $conv(A \cup B)$ for $A, B \in \mathcal{PC}^n$. Therefore, $(A \sqcup B)^\mathcal{I}$ is rather defined as*

$$(A \sqcup B)^\mathcal{I} = \begin{cases} conv(A \cup B) & \text{if } \vec{0} \notin conv(A \cup B) \\ \mathbb{R}^n & \text{if } \vec{0} \in conv(A \cup B). \end{cases}$$

Details can be found in Section 7.3.

Both, the set of $\perp$-closed sets and the set of convex sets can be interpreted as a lattice each by interpreting lattice-meet as conjunction, i.e., set-intersection and lattice-join as $\perp$- or convex-closure resp. The lattice of Y with lattice meet as set-intersection and lattice-join as $\perp$-closure makes up a sublattice of the lattice of $\perp$-closed sets and thus the next proposition follows trivially.

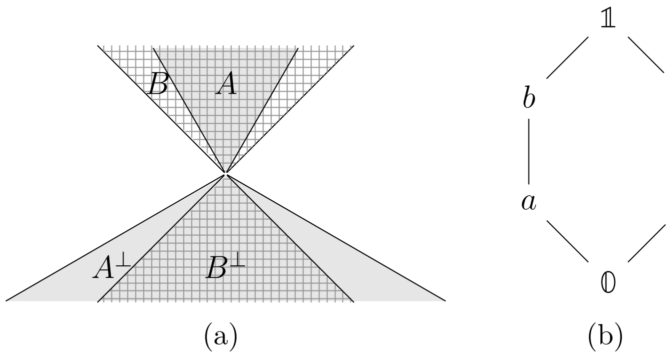

Figure 5.2: (a) Example of a COS based on closed convex cones with two concepts A and B; (b) lattice representation of this COS

Proposition 5.9. $(Y; \vee; \wedge)$ *with* Y *being the component* Y *of a COS with lattice meet interpreted as set-intersection and lattice join interpreted as* $\perp$*-closure makes up an ortholattice.*

The ortholattice based on Y is denoted as $\mathcal{L}_Y$ in the following. So the proposition states that $\mathcal{L}_Y$ is an ortholattice. The other direction is more interesting: How expressive is the resulting lattice $\mathcal{L}_Y$? Is it actually the case that it can represent every ortholattice? Or to state it differently: can each orthomodel be represented by a COS?

This problem is approached from two directions. First, an example is given, showing that COSs are able to model weak logics, especially those for which many of the well-known properties such as orthomodularity and distributivity are not fulfilled. Towards this end, the construction based on closed convex cones introduced in Example 5.6 is used, especially a COS $S = (\mathbb{R}^2, \{A, B\}, \mathcal{B}_E, \perp_p)$ depicted in Figure 5.2(a). There, Y makes up the lattice of Figure 5.2 (b). For this lattice it can be directly seen that it has neither the property of distributivity nor orthomodularity (for details, see Appendix A). It is also possible to model for some other, weaker rules counterexamples for this COS. This topic is considered in detail in Section 10.3, when determining the logical commitment of COSs based on cones. This example shows that COSs, in general, are able to model quite weak logics, especially ones not fulfilling orthomodularity. Are they, however, able to model the whole orthologic? Is every orthomodel thus representable by a COS?

As the betweenness relation is, by definition, not enforced to fulfill any betweenness axiom, it is, of course, possible to choose the trivial, thus empty, betweenness relation. Accordingly, every orthomodel can be interpreted as COS. However, the next theorem shows that there are COSs, even with a non-trivial betweenness relation, that are able to model all ortholattices.

In the following theorem, the orthonegation-induced betweenness is used which will be introduced in Section 6.1, as it enables defining a betweenness relation based on an orthomodel (and thus an ortholattice). Details can be found in Section 6.1 and are introduced in the proof whenever necessary.

Theorem 5.10. *Each ortholattice can be represented by a COS $S = (X, Y, \mathcal{B}, \perp)$ with non-trivial (i.e., non-empty) $\mathcal{B}$ fulfilling (B0)–(B4) (see page 47).*

Proof. Let $\mathcal{L}$ be an arbitrary ortholattice. By Theorem 3.3, it is possible to construct an orthomodel representing this ortholattice and thus $\perp$ is given. Define $S = (X, Y, \mathcal{B}_\perp, \perp)$ with the given orthogonality relation and $Y = \{cl_\perp(A) \mid A \subseteq X\}$, thus Y containing all $\perp$-closed sets (but Y not necessarily being complete). The betweenness relation is defined as follows: $\mathcal{B}_\perp(a, b, c)$ iff $a^\perp \cap c^\perp \subseteq b^\perp$ and $b \notin \{a, c\}$ for $a, b, c \in X$. This is, in fact, the orthonegation-induced betweenness which will be introduced in Section 6.1. As will be shown in Proposition 6.15 it is possible to extend this construction to get a betweenness relation fulfilling (B0)–(B2) and (B4) and it will be shown in Proposition 6.16 that based on this definition each $\perp$-closed set is convex. However, $\mathcal{B}_\perp$ does not necessarily fulfill (B3). As will be discussed in Section 6.1, the fulfillment of (B3) turns out to be difficult. However, here a simple workaround can be used: To gain (B3), the observation is helpful that convexity is kept when deleting betweenness relations (e.g., if $\{a, b, c, d\}$ is convex based on $\mathcal{B}_\perp(a, b, c), \mathcal{B}_\perp(b, c, d)$, then it is also convex solely based on $\mathcal{B}_\perp(a, b, c)$). This opens up the possibility to accomplish (B3) by deleting betweenness relations interfering with it. The resulting betweenness relation is not trivial, as to satisfy (B3) it is sufficient that the precondition is not fulfilled, thus that no $a, b \in X$ exists with $\mathcal{B}_\perp(a, b, \cdot)$ and $\mathcal{B}_\perp(a, \cdot, b)$ or $\mathcal{B}_\perp(\cdot, a, b)$. More details regarding the construction of the induced betweenness can be found in Section 6.1. □

Although the betweenness relation used is not trivial, it is still quite weak and thus does not answer the question of whether it is possible for an (arbitrary) strong betweenness relation to model all ortholattices. Though this theorem states a positive result regarding the expressivity of COSs, it also points out a drawback of the betweenness relation: when defining Y based on $\perp$-closed sets, it is always possible to define a COS by modeling an arbitrary weak betweenness relation (especially the trivial one). The weaker the betweenness relation, the more sets are convex — this is based on the fact that there are less betweenness relations preventing a set from being convex. To gain a meaningful betweenness relation, allowing for sufficiently structuring the space and for convex optimization strategies, it must be as strong as possible. However, when oneself

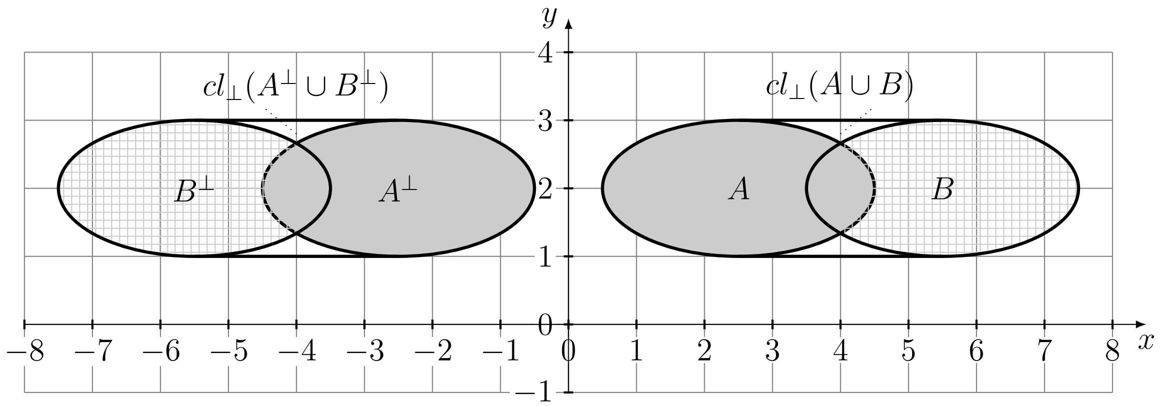

Figure 5.3: Example for a COS based on $X = \mathbb{R}^2$ and Euclidean betweenness. Though, it is a COS, the representation is not expressive.

is restricted to a given similarity relation, then having a weak betweenness relation is still better than not having one.

5.2 Special Forms of Conceptual Orthospaces

At this state, the definition of the general framework of COSs has been given. The simple structure of the COS as defined in Definition 5.2 and the discussion on the expressivity of the betweenness in Theorem 5.10 directly leads to the following question:

Question 5. *Is it sufficient to model betweenness and orthonegation as arbitrary closure operators (as long as they fulfill the conditions of Definition 5.2) to gain an expressive COS?*

The answer is "No!". Though, as discussed above, it is possible to represent both convexity and $\perp$-closure as (nearly arbitrary) closure operators, this does not necessarily lead to a sufficient level of expressivity. Consider the following example:

Example 5.11. *Consider Figure 5.3, with the domain $X = \mathbb{R}^2$ and let the betweenness relation be Euclidean betweenness. Y consists of A and B and their closure under conjunction, disjunction and negation.*

Based on the negated concepts denoted in the figure, all axioms of orthonegation are fulfilled, each set is convex, Y is closed for conjunction and disjunction (as disjunction is defined as $\perp$-closure), thus this example matches the definition of a COS.

Although it is a genuine COS and use cases are imaginable, such a COS brings up some peculiarities which are not suitable for all embedding scenarios. One peculiarity is the contrast in expressivity between the $\perp$-closure and the convex-closure. Although exactly equal strength of both is only possible for complete COS (and subsets of it), a strong connection between convexity and $\perp$-closure could mean, e.g., that they are equivalent for all sets in Y or at least are equivalent in the sense of $conv(A \cup A^\perp) = cl_\perp(A \cup A^\perp)$ for $A \in Y$. In this case, however, neither of the two are fulfilled and the connection between the two notions is quite weak. In particular, it is not possible to infer the negation of new concepts geometrically. Instead, it is necessary to define it for each new concept. Additionally, not all instances of the domain have a negation at all (thus, $x^\perp = \emptyset$ for some $x \in X$). Thus, another possible restriction of COSs would be to enforce all elements to have a negation (not being the empty set). Another argument, discussed in detail in Section 7.2, is that a COS is more expressive and can also be used better for learning if some form of invariance of the concepts is given, and, thus, that rotating or translating the concepts still lead to a valid COS. This is not the case in this example.

This topic of controlling and enhancing the expressivity of COSs, as pointed out in the example above, is handled in the next sections. In Section 5.2.1, some axioms are presented, handling the peculiarities mentioned in the previous example. In Section 5.2.2, a different way of controlling the expressivity is chosen by restricting $\mathcal{L}_Y$ to be distributive and in Section 5.2.3 faithful COSs are considered.

5.2.1 Conceptual Orthospaces with Enhanced Expressivity

To increase the strength of a COS without enforcing completeness, the following subsections present different axioms that increase the expressivity of a COS. This makes COSs more suitable for learning and reasoning tasks, as argued in Example 5.1 and Section 4.2, a structured embedding is less prone to overfitting than an unstructured one. Therefore, using a strong convexity as the Euclidean one leads to a structured space and, thus, to less overfitting. However, relying on Euclidean convexity is not an all-in-one solution and is in particular not helpful when based on it a weak orthogonality relation is defined as in Example 5.11. Therefore, in the following subsection, restrictions of the COS of different strengths are presented, allowing for the selection of the one appropriate for a specific use case.

A natural restriction is the axiom of density (B5), thus to state that all elements should have an element in between them. This enforces a form of connectivity of the elements of a concept C. Every element $a \in C$ has to be connected to every other element $c \in C$ via an element b with $\mathcal{B}(a, b, c)$. Thus, there aren't any unconnected clusters in C. For clarity, (B5) is repeated here and given as first restriction (R1).

(R1) For any two points a and c in X, there is some point b such that $\mathcal{B}(a, b, c)$.

This restriction does not hold for discrete sets. In order to widen applicability, (R1) can be relaxed to (R2) that states that all elements need to be related via a betweenness relation. In (R2), there is, however, no requirement for an element in between. This depicts the betweenness axiom (B5') and is repeated here for clarity.

(R2) For any two points a and c in X, there is some point b such that $\mathcal{B}(a, b, c)$ or $\mathcal{B}(b, a, c)$ or $\mathcal{B}(a, c, b)$.

When considering a non-complete Y, it is possible to relax this restriction even further, because elements in the $\perp$-closure of an element $x \in X$ (i.e., elements in the concept $C = cl_{\perp}(\{x\})$) can be considered as equal (or at least sufficiently similar). Therefore, it can be argued that a betweenness relation between elements in C interferes with the principle behind (B0), as the elements are not distinctive enough. To state it differently: if two elements are equal, they also need to behave equally and, thus, need to have the same properties regarding betweenness. This restriction is only relevant if Y is not complete, as otherwise $\{x\} = cl_{\perp}(\{x\})$, and thus the restriction would not be applicable.

(R3) For $a, c \in X$, (R1) (resp. (R2)) is not valid iff ($cl_{\perp}(\{a\}) = cl_{\perp}(\{c\})$ and for all $d, e \in X$: $\mathcal{B}(d, e, a)$ iff $\mathcal{B}(d, e, c)$ and $\mathcal{B}(d, a, e)$ iff $\mathcal{B}(d, c, e)$).

As argued in Example 5.11, convex hull and $\perp$-closure should have expressivities as similar as possible. (R4) enforces equality of convexity and $\perp$-closure regarding all elements of Y. As discussed in Example 5.8, this is not always achievable, therefore, as a second, weaker restriction, (R5) is introduced. It only enforces equality of closures for a concept and its negation.

(R4) For all $A, B \in Y$, $cl_{\perp}(A \cup B) = conv(A \cup B)$

(R5) For all $A \in Y$, $cl_{\perp}(A \cup A^{\perp}) = conv(A \cup A^{\perp})$

Note that $cl_{\perp}(A \cup A^{\perp}) = X$ is always the case, thus here the point is that $conv(A \cup A^{\perp})$ is enforced to be X. The next restriction is not based on concepts in Y but on elements in X and states that each element should have at least a negation at all, enforcing that the $\perp$-closure should, at least, not encompass the whole space.

(R6) For all $x \in X$: $cl_{\perp}(\{x\}) \subsetneq X$

The next restriction handles the representation of the $\overline{0}$-concept. Whereas one option is to state that $A \cap A^{\perp} = \emptyset$ for all $A \in Y$, another option is to model $A \cap A^{\perp} = \{0\}$ for all $A \in Y$ and some arbitrary but fixed element $0 \in X$. This means that $X^{\perp} = \{0\}$, thus $\{0\}$ depicts the contradictory concept. This allows for modeling inconsistencies in the data, and, thus, instances with contradicting labels in the space by explicitly stating their inconsistency. Therefore, a differentiation is possible between inconsistent information on an instance and a never-mentioned instance. The next axiom enforces the existence of such an 0-element.

(R7) $\emptyset \notin Y$ and there is exactly one $0 \in X$ with for all $A \in X$: $0 \subseteq A$.

In contrast, it is of course also possible to enforce that such an 0-element does not exists:

(R8) $\emptyset \in Y$ (and thus $X^{\perp} = \emptyset$).

If $\emptyset \in Y$, then $\emptyset \subseteq A$ for all $A \in Y$ and thus $\overline{0}^{\mathcal{I}} = \emptyset$. An example for a COS with restriction (R7) is the COS based on closed convex cones as introduced in Example 5.6, an example for a COS with restriction (R8) is the COS based on pseudo-cones as introduced in Example 5.8.

In the following, some COSs are exemplarily discussed based on the restrictions they fulfill. The first one considers the restrictions the example based on cones introduced in Example 5.6 fulfills.

Example 5.12. *Consider first the example based on closed convex cones mentioned in Example 5.6. There, all betweenness axioms are fulfilled (including (B5), thus (R1)). Additionally, the restriction (R4) and, as each convex cone contains* 0*, also (R7) is fulfilled. However,* Y *is not complete, as for all* $x \in X$*:* $cl_{\perp}(\{x\}) = \{\lambda x \mid \lambda \geq 0\}$.

This COS can be adapted to be based not on $X = \mathbb{R}^n$ but on the unit sphere.

Example 5.13. *Consider an* n*-dimensional sphere, where the orthogonality relation is defined as polarity as in Example 5.6. Then, for each point* x *on the sphere,* $cl_{\perp}(\{x\}) = \{x\}$ *and the* $\perp$*-closure is defined as convex hull and thus the COS is complete (or* $Y \subseteq Y'$ *for a complete* Y'*). The betweenness based on the angle (*$\mathcal{B}_a$*) is used. Then* Y *is complete, but the betweenness relation fulfills only (B0)–(B2) and (B4)–(B5) (and thus (R1)) but not (B3), as the definition of betweenness based on an angle introduces loops of the betweenness in the form of* $\mathcal{B}(a, b, c), \mathcal{B}(b, c, d)$ *and* $\mathcal{B}(c, d, a)$ *for some* $a, b, c, d \in \mathbb{R}^n$*. When considering the same setup but based on closed convex cones in* $\mathbb{R}^n$ *and not on convex regions of a sphere, (R1) is not fulfilled, however, (R3) is, as the only elements* *not having a betweenness relation are the elements in the cone* $C = \{\lambda x \mid \lambda \geq 0\}$ *for all* $x \in \mathbb{R}^n$*. (Note that* C *is actually a special cone called a ray, namely, the set of points on the line segment starting at the origin and lying in the direction of a vector* x*). Additionally, both variants fulfill (R4).*

The third example, a different direction is considered. A betweenness relation is given and an orthogonality relation is constructed based on it.

Example 5.14. *Consider a COS with an arbitrary but given betweenness relation fulfilling (B0)–(B5) and define the orthogonality relation as follows: Take an arbitrary and henceforth fixed element* $0 \in X$*. Then define an orthogonality relation for all* $x, y \neq 0$ *by:* $x \perp_0 y$ *iff* $\mathcal{B}(x, 0, y)$*. As intuition for this negation, a two-dimensional space with a betweenness relation* B_d *based on the Euclidean distance can be considered.* 0 *would be the point of origin and the negation of a point would be a half line starting at the point of origin. The relation is indeed an orthogonality relation: it is irreflexive due to (B0) and symmetric due to (B1). All* $\perp$*-closed concepts are convex, however, not all convex concepts are* $\perp$*-closed, thus* Y *is not complete. Consider, e.g., a convex set consisting of one element* x*. Thus, in* $cl_\perp(\{x\})$*, there are, amongst others, all elements* z *with* $\mathcal{B}(0, z, x)$*.*

The three examples showcase the usability of the above mentioned restrictions. These restrictions are of interest in particular in Chapter 7 when considering the induced orthogonality relation based on the Euclidean betweenness. There, exactly the motivational problem of this section occurs: without any restrictions, it is possible to create a variety of COSs based on Euclidean betweenness which are, however, quite restricted in their expressivity. When incorporating some of the above mentioned restrictions, it is possible to get an expressive induced orthogonality relation based on the polarity of convex cones (or pseudo-cones).

5.2.2 Distributive Conceptual Orthospaces

In the last subsection, some restrictions were presented to enhance the general expressivity of the COS, independent of thoughts about its logical commitments. In this section, the most important (or at least, often desired) logical commitment to the logical operations disjunction and conjunction, the distributivity over each other, is considered. As shown in Section 5.1, COSs have in general not the property of distributivity. I show how a restriction of a COS (resp. the lattice based on Y) can be constructed such that $\mathcal{L}_Y$ is distributive. This is possible by applying an alternative definition of distributivity which is not generally applicable, as an complement-operator is needed for the definition.

However, as $\mathcal{L}_Y$ makes up an ortholattice, this definition can be used (for details, see Appendix A and the proof of the following proposition).

(wLLJ) $a \wedge b \leq \mathbb{0}$, then $a \leq b'$.

The rule is based on MacNeille's axiomatization of Boolean algebras (according to [Padmanabhan and Rudeanu, 2008, axiom system B67, p. 114]) and is called (wLLJ) as it is a weakening of Johansson's constructive contraposition rule [Hartonas, 2016]. Thus, to regain distributivity, either two sets $A, B \in Y$ must not intersect or each element in A must be orthogonal to each element in B. This directly leads to the following proposition which is a generalization of my work in [Özçep et al., 2023, Proposition 5]

Proposition 5.15. *$\mathcal{L}_Y$ of a COS fulfills distributivity iff for each combination of two sets A and B (with $A, B \in Y$) it holds that $A \cap B \not\subseteq (\overline{0})^{\mathcal{I}}$ or for each $a \in A, b \in B$: $a \perp b$.*

Proof. We use MacNeille's axiomatization of Boolean algebras (according to [Padmanabhan and Rudeanu, 2008, axiom system B67, p. 114]): He states that a set L equipped with partial order $\leq$ and complement $\cdot'$ characterizes a Boolean algebra if the following axioms are fulfilled.

1. $\forall a \forall b(a \leq b \; \& \; b \leq a \Rightarrow a = b)$
2. $\forall a \forall b \forall c(a \leq b \; \& \; b \leq c \Rightarrow a \leq c)$
3. $\forall a \forall b[(a \wedge b) \leq a \; \& \; (a \wedge b) \leq b \; \& \; \forall c((c \leq a \; \& \; c \leq b) \Rightarrow c \leq (a \wedge b))]$
4. $\forall a \forall b(a \wedge a' \leq b)$
5. $(\forall c(a \wedge b \leq c)) \Rightarrow a \leq b'$
6. $\forall a \forall b(a \leq b \Rightarrow b' \leq a')$

All but one axiom, namely the 5$^{\text{th}}$ axiom, of this axiom system follow trivially from the definition of an ortholattice given in Chapter 3 and the geometric interpretation of COS and are fulfilled by all $\mathcal{L}_Y$ of COSs. Thus, $\cdot'$ can be interpreted as the orthocomplement. This reduces the proof to show the validity of the 5$^{\text{th}}$ axiom above. This axiom corresponds to the rule (wLLJ) which states that if $a \wedge b \leq \mathbb{0}$, then $a \leq b'$. (Note for the correspondence that $a \wedge b$ is smaller than all c iff $a \wedge b$ is smaller than $\mathbb{0}$).

"$\rightarrow$": (wLLJ) can be interpreted on a geometric level as follows: when A and B are disjoint, then A must be a subset of $B^{\perp}$. By the definition of negation as orthogonality operator this is possible only if for each $a \in A$ and each $b \in B$ it holds that $a \perp b$. When A and B intersect, then the condition is fulfilled trivially.

"$\leftarrow$": Follows from (wLLJ) and the definition of polarity. □

An application of this construction principle can be seen in the following example.

Example 5.16. *Again, the COS based on closed convex cones of Example 5.6 is considered. Consider the two different COSs in Figure 5.4.* $\mathcal{L}_Y$ *of the left one is not distributive: when using the standard definition of distributivity (see Appendix A) then* $b \wedge (a \vee a^{\perp}) = b$ *but on the other hand* $(b \wedge a) \vee (b \wedge a^{\perp}) = \overline{0}$. $\mathcal{L}_Y$ *of the right COS is distributive and each two concepts either intersect or are orthogonal to each other.*

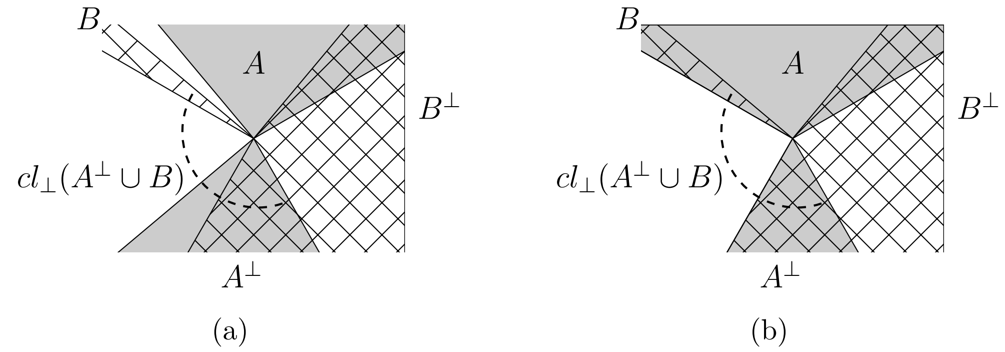

Figure 5.4: Model for a COS based on convex cones and (a) non-distributive and (b) distributive $\mathcal{L}_Y$

This example shows that it is possible to use a non-distributive $\mathcal{L}_Y$ of a COS as a basis (in this case the set $\mathcal{C}^n_{cone}$ of all closed convex cones) and restrict it by choosing a distributive subset.

Instead of restricting an existing COS such that $\mathcal{L}_Y$ is distributive, it is also possible to create the COS based on a structure that already enforces $\mathcal{L}_Y$ to be distributive. One example of this is presented in the following in form of al-cones.

Example 5.17 (based on [Özçep et al., 2023])**.** *Consider again the example based on closed convex cones presented in Example 5.6 and restrict the class of cones to so called* axis-aligned cones *defined as follows:*

$$C \text{ is an al-cone iff } C = (C_i)_{1 \le i \le n} = C_1 \times \cdots \times C_n, \text{ where } C_i \in \{\mathbb{R}, \mathbb{R}_+, \mathbb{R}_-, \{0\}\}.$$

It is still possible to consider the Euclidean betweenness and an orthogonality relation based on polarity. As shown by Özçep et al. [2023, Corollary 1], the set of axis-aligned cones is closed under conjunction, disjunction and negation and thus, $S = (\mathbb{R}^n, Y, \mathcal{B}_E, \perp_p)$ *with* Y *being the set of al-cones is a COS. According to Proposition 5.15, the lattice* $\mathcal{L}_Y$ *based on al-cones is distributive. An example al-cone-interpretation for the domain* $X = \mathbb{R}^2$ *and an empty tbox can be seen in Figure 5.5.*

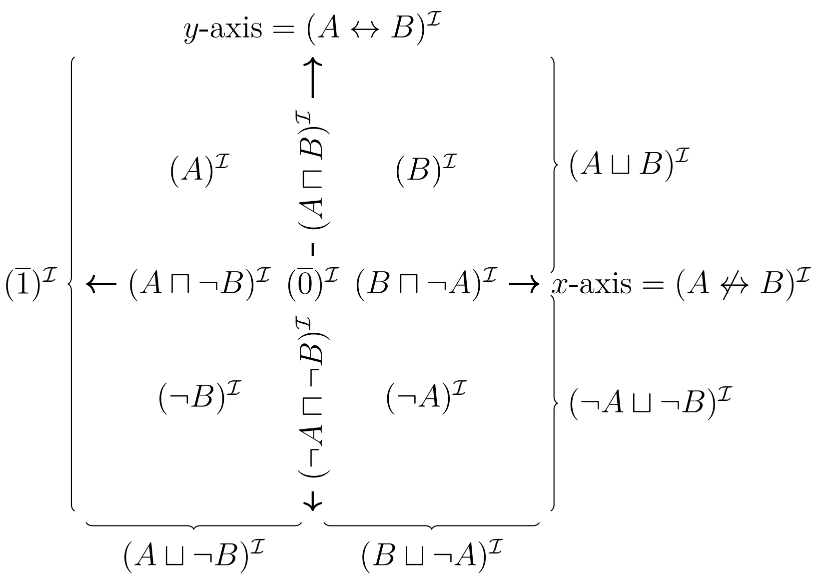

Figure 5.5: Representation of Y based on al-cones for an empty tbox.

Now, it is shown that Proposition 5.15, in fact, depicts a reasonable restriction, as it enables restricting the COS in such a way that Boolean $\mathcal{ALC}$-ontologies are representable. It is easy to show that a classically satisfiable Boolean $\mathcal{ALC}$-ontology can be represented via a COS, as, similar as in Theorem 5.10, an arbitrary weak betweenness relation can be used. Therefore, in the following, this direction is shown for an instance of the COS, the al-cones, presented in Example 5.17. These are based on Euclidean betweenness and thus represent a strong and expressive variant of COSs and thus show that COSs are, in fact, also able to model Boolean $\mathcal{ALC}$-ontologies in realistic circumstances. The proof is mainly based on the proof of Proposition 7 by Özçep et al. [2023].

Theorem 5.18. *Boolean $\mathcal{ALC}$-ontologies are classically satisfiable iff they are by a geometric model based on a COS $S = (X, Y, \mathcal{B}, \perp)$ fulfilling the restrictions of Proposition 5.15.*

Proof. "→": The proof is based on the fact that with a construction based on al-cones each $\mathcal{ALC}$-ontology is representable. As al-cone models are a special case of COSs, it follows that each Boolean $\mathcal{ALC}$-ontology is representable by a COS. The rest of the proof follows directly with the proof of Proposition 7 by Özçep et al. [2023].

"←": The following proof is a slight adaptation of the proof of Proposition 7 by Özçep et al. [2023]. Though the original theorem is proved for al-cones, it is also applicable to general COS. Assume that a COS-interpretation $\mathcal{I}_g$ of the ontology $\mathcal{O} = \mathcal{T} \cup \mathcal{A}$

exists. We have to construct a classical model $\mathcal{I} \models \mathcal{T} \cup \mathcal{A}$. For the specification of $\mathcal{I}$ it is sufficient to specify the denotations of the constants occurring in $\mathcal{A}$, denoted by N_c in the following, and for the concept symbols occurring in $\mathcal{O}$, denoted by $N_C = \{A_1, \dots, A_n\}$ in the following. We may further assume that each abox axiom is of the form $A_i(c)$. Otherwise, if the axiom is of the form $C(c)$ we would add a new atomic symbol $A_{n+j}$to N_C and replace $C(x)$ by $A_{n+j}(x), A_{n+j} \equiv C$.

Let the domain $\Delta^{\mathcal{I}}$ be just the set of constants N_c. The constants are interpreted by themselves, i.e., for any constant $c \in N_c$ set $c^{\mathcal{I}} = c$. We are left with specifying the denotations of each $A_i \in N_C$. For each constant $c \in N_c$ let $Z \subseteq \Delta^{\mathcal{I}_g}$ be the set representing the most specific concept including $c^{\mathcal{I}_g}$, denoted by concept C. Now consider the following set of algebraic atoms[1] compatible with C, i.e., consider the set

$$atc(c) = \{L = L_1 \sqcap \ldots \sqcap L_n \mid L_i \in \{A_i, \neg A_i\} \text{ and } L^{\mathcal{I}_g} \subseteq Z\}$$

Intuitively, $atc(c)$ describes the possible "groundings" of the concept C, as either C is an algebraic atom and thus $atc(c) = \{C\}$ or C is not. If C is not an algebraic atom, then the ability of the geometric model of modeling partial knowledge is used. Thus, e.g., $c^{\mathcal{I}_g} \in (C_1 \sqcup C_2)^{\mathcal{I}_g}$, but neither $c^{\mathcal{I}_g} \in (C_1)^{\mathcal{I}_g}$ nor $c^{\mathcal{I}_g} \in (C_2)^{\mathcal{I}_g}$. However, this ability is not given in a classical interpretation. Thus, an arbitrary element denoted $catc(c) \in atc(c)$ is chosen. Let $catc(c) = L_1 \sqcap \ldots \sqcap L_n$.

Now we can define $A_i^{\mathcal{I}} = \{c \in N_c \mid A_i = L_j$ for some conjunct L_j of $catc(c)\}$. Indeed, $\mathcal{I}$ is a model of the ontology. It makes all abox axioms $A_i(c)$ true, because $\mathcal{I}_g \models A_i(c)$, hence $catc(c)$ contains A_i and does not contain $\neg A_i$. The constructed model makes also every tbox axiom $C \sqsubseteq D$ true. Because, assume $c \in C^{\mathcal{I}}$. As c is completely specified w.r.t. each symbol A_i we also have $c^{\mathcal{I}_g} \in C^{\mathcal{I}_g}$, hence $c^{\mathcal{I}_g} \in D^{\mathcal{I}_g}$ and so $c \in D^{\mathcal{I}}$. □

5.2.3 Faithful Conceptual Orthospaces

As discussed in Section 3.2, faithfulness could be a desirable property. Therefore, the question arises, whether in COSs it is not only possible to model partiality but to model faithfulness. This question can be answered positively. Due to the variety of possible COSs, here I focus on a COS based on al-cones as defined in Example 5.17 and show that such a COS can be defined such that it allows for a strongly concept-faithful and tbox-faithful interpretation.

The following proposition is a slight adaptation of my work [Özçep et al., 2023, Proposition 8].

[1] Usually called "atom", here the term "algebraic" is added in order to prevent clashes with the atomic concepts in description logics. For details, see Appendix A.

Proposition 5.19. *For classically satisfiable Boolean $\mathcal{ALC}$-ontologies there is a strongly concept-faithful and tbox-faithful COS-model on some finite $\mathbb{R}^{2n}$ based on al-cones of the form $b_1 \times \cdots \times b_{2n}$ with $b_{2i} \in \{0, \mathbb{R}_+, \mathbb{R}_-, \mathbb{R}\}$ and $b_{2i+1} = b_{2i}$. Here n is the number of atomic elements in the Boolean algebra generated by the tbox of the ontology.*

Proof. First a (non-faithful) geometric model of the given tbox is generated in size n, as done in the proof of [Özçep et al., 2023, Prop. 7] and restated here: Like the $\mathcal{ALC}$-ontology $\mathcal{O}$, the induced Boolean algebra is finite, say containing no more than 2^k elements, with a finite number of algebraic atoms m. We choose a dimension n such that $4^n \geq 2^k$ and $n \geq m/2$, the idea being that we have to represent all of the 2^k Boolean concepts by embedding the algebraic atoms of the Boolean algebra directly onto half-axes of $\mathbb{R}$. Now one can see that O_4^n is the n-wise Cartesian product of the diamond shaped Boolean algebra $O_4 = \{\{0\}, \mathbb{R}_+, \mathbb{R}_-, \mathbb{R}\}$ with set inclusion as order: $\{0\} \subsetneq \mathbb{R}_+, \mathbb{R}_- \subsetneq \mathbb{R}$, where $\sqcap$, $\sqcup$, $\neg$ are defined component wise based on the $\sqcap$, $\sqcup$, $\neg$ defined on O_4. But in this way $\sqcap$ over O_4 is nothing else than set intersection, $\neg$ is polarity, and $\sqcup$ is defined by de Morgan. So, if $\mathcal{O}$ is satisfiable, then we can construct a model.

By definition, this model is tbox-faithful, thus only strong concept-faithfulness needs to be shown. In the resulting model every dimension is doubled, $(A \sqcup B)^{\mathcal{I}} = \mathbb{R}_+ \times \mathbb{R}_- \times \mathbb{R}$ for instance becomes $(A \sqcup B)^{\mathcal{I}} = \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_- \times \mathbb{R}_- \times \mathbb{R} \times \mathbb{R}$. To ensure strong concept-faithfulness it needs to be the case that for any object a in the abox if $a^{\mathcal{I}} \in C^{\mathcal{I}}$, then $O \models C(a)$. This can be achieved by embedding a into its most specific concept (msc) $M^{\mathcal{I}}$, for which not at the same time $a \in A^{\mathcal{I}}$ with $A \sqsubset M$ holds, i.e., not embedding it into a concept properly subsumed by $M^{\mathcal{I}}$. Let M be the msc of some object x. Consider the concepts C_i for $i = 1, \ldots, n$ properly subsumed by M. To ensure concept-faithfulness, $x^{\mathcal{I}} \in M^{\mathcal{I}}$ must be valid, however $x^{\mathcal{I}} \notin C_i^{\mathcal{I}}$ for all $i = 1, \ldots, n$. Each al-cone $C_i^{\mathcal{I}}$ must cover at least one dimension d where $(C_i^{\mathcal{I}})_d \subsetneq (M^{\mathcal{I}})_d$, as otherwise the subsumption would not be proper. In this dimension d it is either the case that $(M^{\mathcal{I}})_d = \mathbb{R}$ and $(C_i^{\mathcal{I}})_d \in \{\mathbb{R}_+, \mathbb{R}_-, \{0\}\}$ or $(M^{\mathcal{I}})_d = \mathbb{R}_+$ (analog for $\mathbb{R}_-$) and $(C_i^{\mathcal{I}})_d = \{0\}$. Consider first the second case, there $x^{\mathcal{I}}$ can be placed in this dimension in $\mathbb{R}_+ \backslash \{0\}$ and thus is not contained in $(C_i^{\mathcal{I}})_d$. For all other $C_j^{\mathcal{I}}$ (with $j \neq i$), either $(C_j^{\mathcal{I}})_d = (C_i^{\mathcal{I}})_d$ and thus $x^{\mathcal{I}} \notin C_j^{\mathcal{I}}$ or $(C_j^{\mathcal{I}})_d = \mathbb{R}_+$, then there must be a dimension $k \neq d$ ensuring that $C_j^{\mathcal{I}} \subsetneq M^{\mathcal{I}}$. In the first case this is not that easy. When considering the model without doubled dimensions, it is possible to place $(x^{\mathcal{I}})_d$ in $\mathbb{R}_-$, when $(C_i^{\mathcal{I}})_d = \mathbb{R}_+$, however, it is possible that there is an $C_j^{\mathcal{I}}$ with $(C_j^{\mathcal{I}})_d = \mathbb{R}_-$ and $(C_j^{\mathcal{I}})_k = (M^{\mathcal{I}})_k$ for $k = 1, \ldots, d-1, d+1, \ldots, n$, thus is a proper subsumption of M only because of the difference in dimension d. Then placing $(x^{\mathcal{I}})_d$ in $\mathbb{R}_-$ would result in $x^{\mathcal{I}} \in C_j^{\mathcal{I}}$, a contradiction. Thus, the doubled dimension is needed. Therefore, $(M^{\mathcal{I}})_{d,d+1} = \mathbb{R} \times \mathbb{R}$, $(C_i^{\mathcal{I}})_{d,d+1} = \mathbb{R}_+ \times \mathbb{R}_+$ and $(C_j^{\mathcal{I}})_{d,d+1} = \mathbb{R}_- \times \mathbb{R}_-$. Then, $(x^{\mathcal{I}})_{d,d+1}$ can be placed in $\mathbb{R}_+ \times \mathbb{R}_-$, being neither in $C_i^{\mathcal{I}}$ nor in $C_j^{\mathcal{I}}$. This follows analogously for the other cases and the other C_i. Thus, x is only contained in the msc and not in any concept properly subsuming it and therefore concept-faithfulness is given. □

5.3 Properties of Conceptual Orthospaces

As the framework of COSs is quite general and especially does not state many restrictions on the construction (as has been seen in the last section), many properties of COS are highly dependent on the specialties of the chosen COS and do not hold in general. Therefore, the focus in this section lies on the most basic property of COS, the completeness. To recapitulate, a set Y of a COS is complete iff $Y = \{conv(A) \mid A \subseteq X\} = \{cl_\perp(A) \mid A \subseteq X\}$. Intuitively, this means that each instance can be considered individually regarding its negation. This property can also be justified by looking at properties of closure operators, as a closure operator cl is called *simple* if it satisfies $cl(\{\emptyset\}) = \{\emptyset\}$ and $cl(\{x\}) = \{x\}$ for every $x \in X$ [Faure and Frölicher, 2000, p. 58]. This can be considered as a weaker form of completeness. In the further course of this section, it is shown that completeness is hard to gain but that it is possible to find a slightly weaker notion that is easier to approach.

5.3.1 Completeness

Completeness is not easily satisfiable, especially, when all betweenness axioms have to be fulfilled. This was pointed out in Example 5.12 for a COS based on closed convex cones and in Example 5.14 for a COS based on a given betweenness. This intuition is proven in the rest of this section. One main problem of completeness is that convexity relies on the ternary betweenness relation, meaning that one element is added based on two others. In contrast, $\perp$-closure can depend on an arbitrary number of elements. This can be seen in the following example:

Example 5.20. *Consider the example in Figure 5.6, where points connected by lines are orthogonal (i.e., $\perp$-related). The following betweenness relation*[2] *is used: Let $\mathcal{B}(a, b, c)$ iff $a^\perp \cap c^\perp \subseteq b^\perp$ and $b \notin \{a, c\}$ for all $a, b, c \in X$. Then the set $\{x, y, z\}$ is convex. As $x^\perp = \{u, u'', z, w, y\}$, $y^\perp = \{u', u'', x, w\}$ and $z^\perp = \{u, u', x, w\}$, it turns out that $cl_\perp(\{x, y\}) = (x^\perp \cap y^\perp)^\perp = (\{u'', w\})^\perp = \{x, y\}$. This follows analogously for the other $\perp$-closures of two elements out of $\{x, y, z\}$. However, $cl_\perp(\{x, y, z\}) = \{x, y, z, v\}$, as $x^\perp \cap y^\perp \cap z^\perp = \{w\}$ and $v \perp w$. Here, the interplay between more than two elements*

[2] This is again the orthonegation-induced betweenness discussed in detail in Section 6.1.

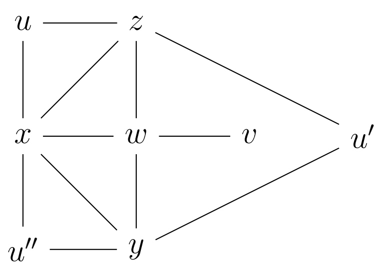

Figure 5.6: Example for a structure with a convex but not ⊥-closed set. Points connected by lines are orthogonal (i.e., ⊥-related).

is necessary to add elements to the ⊥-closure. This cannot be captured by the convex closure: thus, $\{x, y, z\}$ is convex but not ⊥-closed.

To enforce completeness, the ⊥-closure has to fulfill some properties to make it reducible to ternary relations. The next proposition formalizes the intuition of the last example: when an element is added to the ⊥-closure, then it must be possible to add it based on the ⊥-closure of two elements.

Proposition 5.21. *Let $S = (X, Y, \mathcal{B}, \perp)$. If Y is complete, then $cl_\perp$ fulfills the following property for $B \subseteq X$:*

$$\text{if } b \in cl_\perp(B)\backslash B, \text{ then there are } b_1, b_2 \in cl_\perp(B), b \neq b_1 \neq b_2 \text{ with } b \in cl_\perp(\{b_1, b_2\}) \text{ and there is no } b_3 \in cl_\perp(B) \text{ with } b \neq b_3 \text{ and } b \in cl_\perp(\{b_3\}).$$

Proof. Let Y be complete, $B \subseteq X$ and $b \in cl_\perp(B)\backslash B$. For the sake of proof by contradiction, assume that the condition is not fulfilled. Thus, there could be $b_3 \in cl_\perp(B)$ with $b \neq b_3$ and $b \in cl_\perp(\{b_3\})$, then Y is clearly not complete. Thus, assume there are no $b_1, b_2 \in cl_\perp(B)$ with $b \in cl_\perp(\{b_1, b_2\})$ with $b \neq b_1 \neq b_2$. According to Proposition 5.7.3, $cl_\perp(B) = conv(B)$, hence it follows that there are no $a, c \in conv(B)$ with $\mathcal{B}(a, b, c)$ (otherwise we would have $b_1 := a, b_2 := c$ with $b \in conv(\{b_1, b_2\}) = cl_\perp(\{b_1, b_2\})$). Thus b can't be added via convexity. Thus, $cl_\perp(B) \supsetneq conv(B)$, a contradiction to completeness. □

This can be considered as a weak form of the projective law (see below) which intuitively states that for two non-empty sets A and B the closure contains only elements directly between elements of A and B (thus, on a line-like structure). However, if the ⊥-closure fulfills the projective law and has a sufficiently strong betweenness, then the resulting COS is not complete.

Definition 5.22. *[Faure and Frölicher, 2000, Definition 3.3.1] A closure operator cl : $Pow(X) \to Pow(X)$ satisfies the* projective law (pl) *if for any non-empty subsets $A, B \subseteq X$ one has*

$$cl(A \cup B) = \bigcup\{cl(x \cup y) \mid x \in cl(A) \text{ and } y \in cl(B)\}$$

As common betweenness relations such as Euclidean betweenness fulfill the projective law they do not allow for a definition of a COS with a complete Y (if the betweenness relation fulfills (B0)–(B4)).

Theorem 5.23. *For a COS S as defined in Definition 5.2 the following implication holds: if $\mathcal{B}$ fulfills (B0)–(B4) and is not trivial (thus not empty) and $cl_\perp$ fulfills the projective law (pl), then Y is not complete.*

Proof. In the course of this proof, (B0) and (B1) are used without further notice. Let $S = (X, Y, \mathcal{B}, \perp)$ be a COS with $\mathcal{B}$ fulfilling (B0)–(B4) and $cl_\perp$ fulfilling the projective law and, for the sake of proof by contradiction, assume that Y is complete. Then, in particular due to completeness of Y it holds that $cl_\perp(\{a\}) = \{a\}$ for $a \in X$. Because of intuitionistic absurdity (see page 26) it follows that $cl_\perp(\{a\} \cup a^\perp) = X = conv(\{a\} \cup a^\perp)$ and with (pl) it follows that $cl_\perp(\{a\} \cup a^\perp) = \bigcup\{cl_\perp(\{a, y\}) \mid y \in a^\perp\}$. Because of completeness it holds that $cl_\perp(C) = conv(C)$ for $C \subseteq X$ (see Proposition 5.7.3). Thus, $cl_\perp(\{a\} \cup a^\perp) = \bigcup\{x \mid \mathcal{B}(a, x, y)$ with $y \in a^\perp\} \cup \{a\} \cup a^\perp$. Then, there are two possibilities: First, there are $b, c \in X$ with $\mathcal{B}(b, a, c)$ (thus, a is in between some other elements). If $b, c \in a^\perp$, then by convexity, $a \in a^\perp$, a contradiction. If $b \notin a^\perp$, then there is a $\mathcal{B}(a, b, y)$ for some $y \in a^\perp$, as $b \in cl_\perp(\{a\} \cup a^\perp)$ and $b \notin \{a\} \cup a^\perp$ and with (B3), it follows $\mathcal{B}(y, a, c)$. If $c \in a^\perp$, then $a \in a^\perp$, a contradiction. If $c \notin a^\perp$, then as for b, there is $\mathcal{B}(a, c, z)$ for $z \in a^\perp$ and with (B3) and $\mathcal{B}(b, a, c)$, it follows $\mathcal{B}(b, a, z)$ and based on this follows with (B3) and $\mathcal{B}(a, b, y)$ that $\mathcal{B}(y, a, z)$ and thus $a \in a^\perp$, a contradiction. Due to (B1), this is analogue for c.

The second case is that for a, there are no $b, c \in X$ with $\mathcal{B}(b, a, c)$. Assume $cl_\perp(\{a\} \cup a^\perp) \supsetneq \{a\} \cup a^\perp$. Then, there is $x \in X$ with $\mathcal{B}(a, x, y)$ for $y \in a^\perp$. Thus, for x and $x^\perp$ the same argument as in the first case leads to a contradiction. Thus, such an x does not exist, meaning $cl_\perp(\{a\} \cup a^\perp) = \{a\} \cup a^\perp$ for all $a \in X$, thus $a^\perp = X \backslash a$. Then, because of convexity, the betweenness relation must be empty, thus trivial. □

As can be seen in the proof, the problem is based, in particular, on the axiom (B3). When the betweenness relation fulfills only (B0) and commutativity (B1) (as natural restrictions) and (B3), then a construction is still not possible when the projective law is fulfilled. However, if (B3) is not fulfilled, then it is possible to create a complete COS.

Corollary 5.24.

1. *If $\mathcal{B}$ fulfills (B0)–(B1) and (B3) and is not trivial and $cl_\perp$ fulfills the projective law, then Y is not complete.*

2. *There is a COS $S = (X, Y, \mathcal{B}, \perp)$ with $cl_\perp$ fulfilling the projective law and $\mathcal{B}$ fulfilling (B0)–(B2) and (B4) with Y being complete.*

Proof.

1. The proof of Theorem 5.23 does not depend on (B2) and (B4).

2. This is shown Example 5.13 (where X is an n-dimensional sphere).

□

5.3.2 Non-completeness

Though completeness has its merits, as it enables having notions of convexity and $\perp$-closure of the same strength, it is not always achievable, as pointed out above. Besides these practical problems, it is also not always necessary or desired to call for completeness of Y. On the one hand, this can be the case if specific geometric objects are asked for as sets in Y because they have some desirable features, e.g., using convex cones to improve the optimization capabilities. On the other hand, there could be some objects that are not equal but should behave equal in a specific context, thus should have the same negation. This can only be modeled in an incomplete setting, as for a complete Y, these instances need to be identical.

Thus, it could be especially desirable to relax Proposition 5.21 such that if $b \in cl_\perp(B)\backslash B$, then there is $b_1, b_2 \in cl_\perp(B)$ with $b_1, b_2 \neq b$ and $b \in cl_\perp(\{b_1, b_2\})$, but, in contrast to Proposition 5.21, not enforcing $b_1 \neq b_2$. Thus, it should be allowed that there are elements $x \in X$ with $cl_\perp(\{x\}) \supsetneq \{x\}$ but to keep the restriction that not more than two elements should influence the addition of a new element to a set.

Definition 5.25. *For a COS S with $Y = \{cl_\perp(C) \mid C \subseteq X\}$, Y is called* restricted complete *if*

$$Y = \{conv(\psi(C)) \mid C \subseteq X\} \text{ with}$$
$$\psi(C) = \bigcup\{cl_\perp(c) \mid c \in C\}$$

For such a COS, it holds that $cl_\perp(A \cup B) = conv(A \cup B)$ for $A, B \in Y$ as the following proposition shows.

Proposition 5.26. *For a COS S with $Y = \{cl_\perp(C) \mid C \subseteq X\}$: if Y is restricted complete, then $cl_\perp(A \cup B) = conv(A \cup B)$ for $A, B \in Y$, thus (R4) is fulfilled.*

Proof. Let S be given with $Y = \{cl_\perp(C) \mid C \subseteq X\}$ and Y restricted complete. Let $A, B \in Y$ be arbitrary. Then, with the closure rules it follows that $A = \bigcup\{cl_\perp(a) \mid a \in A\}$ (analogously for B). With the condition of restricted completeness, $conv(\psi(A \cup B)) \in Y$ and thus $conv(A \cup B) \in Y$. With the same argument as in Proposition 5.7 it follows $cl_\perp(A \cup B) = conv(A \cup B)$ for $A, B \in Y$. ☐

This restricted completeness stays in contrast to (R3), for which the focus lies on the behavior of elements regarding the betweenness relation, whereas here the connection between $\perp$-closure and convexity is looked at.

Now, the statement from the beginning of this paragraph can be proven, namely that if Y is restricted complete, then either no more than two elements should influence the addition of a new element to a set or an element should be added based on the $\perp$-closure of one element alone.

Proposition 5.27. *Let $S = (X, Y, \mathcal{B}, \perp)$. If Y is restricted complete, then $cl_\perp$ fulfills the following property for $B \subseteq X$:*

$$\textit{if } b \in cl_\perp(B) \backslash B, \textit{ then there is } b_1, b_2 \in cl_\perp(B), b \neq b_1 \neq b_2 \textit{ with } b \in cl_\perp(\{b_1, b_2\}).$$

Proof. Let Y be restricted complete, $B \subseteq X$ and $b \in cl_\perp(B) \backslash B$. For the sake of proof by contradiction, assume that the condition is not fulfilled. Thus there are no $b_1, b_2 \in cl_\perp(B)$ with $b \in cl_\perp(\{b_1, b_2\})$. Due to Definition 5.25, it follows that $cl_\perp(B) = conv(\bigcup\{cl_\perp(\{c\}) \mid c \in B\})$. Thus, especially $b \notin cl_\perp(\{c\})$ for $c \in B$, as otherwise $b_1 = b_2 = c$ could be chosen. Thus $b \in conv(\bigcup\{cl_\perp(\{c\}) \mid c \in B\}) \backslash \bigcup\{cl_\perp(\{c\}) \mid c \in B\}$, thus b is added via convexity, thus there are $b_1, b_2 \in cl_\perp(B)$ with $\mathcal{B}(b_1, b, b_2)$ and thus, due to Proposition 5.7.2, $b \in cl_\perp(\{b_1, b_2\})$. A contradiction. ☐

An example is given by the convex cones of Example 5.12:

Example 5.28 (Example 5.12 continued)**.** *To repeat, let $S = (\mathbb{R}^n, \mathcal{C}^n_{cone}, \mathcal{B}_a, \perp_p)$ be a COS where $\mathcal{C}^n_{cone}$ is the set of all convex cones in $\mathbb{R}^n$, $\mathcal{B}_a$ a betweenness based on an angle and $\perp_p$ the polarity. Then $cl_\perp(\{x\}) = \{\lambda x \mid \lambda \geq 0\}$ for $x \in \mathbb{R}^n$, thus Y is not complete. But Y contains all $\perp$-closed concepts and as $\perp$-closure and the convex closure over cones can both be interpreted as conic hull, Y is restricted complete.*

The notion of restricted completeness thus bypasses the problems of completeness while keeping the advantages, especially the strong definition of convexity.

6 Induced Conceptual Orthospaces

There are two different possibilities for constructing a COS. Mostly, the structure of the given data is known to the extent that either a betweenness or a similarity relation are given or, at the very least, the desired structure of one of them is known. Then, two natural possibilities to construct a COS arise: First, given a similarity relation (and hence automatically an orthogonality relation), a COS is constructed by designing a betweenness relation. This possibility is considered in Section 6.1. Second, given a betweenness relation, a COS is constructed by designing a similarity relation. This is discussed for the general case in Section 6.2 and, later, for the special case of Euclidean betweenness in Chapter 7.

6.1 Orthonegation-induced Betweenness

Assume that a COS has to be created based on a given similarity relation. Although the simplest approach is to use the empty betweenness relation to model all $\perp$-closed sets as concepts (= betweenness-closed sets), this is obviously not suitable, as it does not contain any useful information about convexity. Thus, the aim is to define a betweenness relation relying on the similarity information. This would mean equipping the space with a betweenness relation that is as strong as possible. In the ideal case, one would define a betweenness relation such that Y of the COS is complete, meaning that closure with respect to the betweenness relation has exactly the same expressivity as the $\perp$-closure. But, as stated in the last section, completeness is not always achievable and appropriate. Thus, the aim is to come as close as possible, and, thus, to define a betweenness relation between as many elements as possible.

I illustrate the difference between an induced betweenness relation and an independently defined betweenness relation with the following example. It depicts that there is more than one possibility of defining a COS, even when the structure (in this case closed convex cones) and the orthogonality relation as polarity are given. On the one hand, it is possible to define an induced betweenness relation definable via an orthogonality relation, on the other hand, it is also possible to define a betweenness relation independent of the orthogonality relation. Both COSs look quite similar, however, have different properties regarding their betweenness.

Example 6.1. *Consider two COS* $S_i = (\mathbb{R}^n, \mathcal{C}^n_{cone}, \mathcal{B}_i, \perp_i)$*,* $i \in \{1, 2\}$*, where* $\mathcal{C}^n_{cone}$ *is the set of closed convex cones in* $\mathbb{R}^n$*. For* S_1*,* $\mathcal{B}_1$ *is defined as the Euclidean betweenness and* $\perp_1$ *as* $x \perp_1 y$ *iff* $\angle(x, y) \geq 90°$*. For* S_2*,* $\mathcal{B}_2$ *is defined on the basis of the angle operator* $\angle(\cdot, \cdot)$*, i.e., via* $\mathcal{B}_a$*, and* $\perp_2 = \perp_1$*. For both variants, the* $\perp$*-closure is defined as convex hull and both behave exactly the same regarding* Y*. For* S_2*, the betweenness can be directly defined based on the* $\perp$*-closure, as both are based on the angle. For* S_1*, however, this is not the case, as elements behaving the same way regarding the negation could behave differently regarding a betweenness relation.*

Defining a betweenness relation on the basis of an orthogonality relation proves to be challenging, as the definition is not guaranteed to result in a betweenness relation that fulfills all betweenness axioms (see Section 4.3) or some intended subset of them. To regain these axioms, one may, in principle, follow two directions. Either the orthogonality relation is restricted so that the induced betweenness fulfills all betweenness relations and is complete (or at least as complete as possible) or an arbitrary orthogonality relation is considered and the induced betweenness is restricted. Then, a wider variety of orthogonality relations can be used. However, one has to cope with a weaker betweenness relation. It turns out that a restriction of the orthogonality relation is easily transferable to a restriction of the betweenness. Hence, we follow the first direction and adapt it to the other one afterwards.

In the rest of this section, I define an induced betweenness, then consider case-by-case each betweenness axiom (as defined in Definition 4.2) — identifying those that are fulfilled — and propose restrictions of the orthogonality relation to fulfill the remaining ones. After that, these restrictions are adapted as restrictions of the betweenness induced by an arbitrary orthogonality relation. The task we tackle in the rest of this section can be formally stated as follows: Consider the structure S as defined in Definition 5.2. Let $\perp$ be given. The question is (i) how $\mathcal{B}$ is defined to be a reasonable induced betweenness and (ii) how $\perp$ is defined such that the induced betweenness relation fulfills the betweenness axioms.

Section 6.1.1 considers the definition of the induced betweenness and Section 6.1.2 discusses the betweenness of Goodman (see Section 4.3) as a special case. Sections 6.1.1 and 6.1.2 are mainly based on one of my publications, co-authored with Özgür Özçep [Leemhuis and Özçep, 2023]. Practical aspects of this induced betweenness, including its applicability to specific similarity relations, are discussed in Section 6.1.3.

6.1.1 A Definition of an Induced Betweenness

We aim at modeling a betweenness relation via an orthogonality relation. As the focus lies here on the orthogonality relation, it should especially be the case that all $\perp$-closed sets are convex, so it should be possible to include all $\perp$-closed sets in Y and, furthermore, the betweenness should mimic the properties and the structure of the orthonegation. As betweenness is a ternary relation and $\perp$-closure can also be interpreted as a ternary relation (at least in special cases, compare Proposition 5.21), it is possible to directly define betweenness based on orthogonality.

$$\text{For } a, b, c \in X : \mathcal{B}(a, b, c) \text{ iff } b \in cl_\perp(\{a, c\}) \quad \text{(or, equivalently, iff } a^\perp \cap c^\perp \subseteq b^\perp). \tag{6.1}$$

Equivalence is valid, as $cl_\perp(\{a, c\}) = (a^\perp \cap c^\perp)^\perp$, thus $b \in cl_\perp(\{a, c\})$ iff $b \in (a^\perp \cap c^\perp)^\perp$ iff $(a^\perp \cap c^\perp)^{\perp\perp} \subseteq b^\perp$ and thus $a^\perp \cap c^\perp \subseteq b^\perp$. This definition is preliminary and will be extended in Definition 6.2. This means that an element is between two others if this element is added to the $\perp$-closure of the other two elements.

It turns out that this definition of betweenness is an intersection betweenness as introduced in Definition 4.3. For the convenience of the reader, I shortly repeat the definition [Rautenbach et al., 2011]: For a set system $\mathcal{M} = (M_v)_{v \in V}$ indexed by the elements of a (here not necessarily finite) set V, the intersection betweenness $\mathcal{B}_i$ induced by $\mathcal{M}$ is defined as follows:

$$\mathcal{B}_i(a, b, c) \text{ iff } M_a \cap M_c \subseteq M_b.$$

As discussed in Section 4.3, intersection betweenness is originally motivated by the consideration of shared attributes of some elements. One element is in between two others if it shares the attributes which the others have in common. The same principle can be applied here to underline the validity of this definition. Instead of considering the attributes of the elements, their negations are considered. If it is assumed that an element is defined by its negation, then an element is in between two other elements, if the negation of the element shares the negated objects which the others have in common. Note the fundamental difference to the intersection betweenness and also to the approach of Schockaert, Ibáñez-García, et al. [2021]: whereas intersection betweenness is based on the idea that two sets are given, the set of elements and the set of attributes, in our case both sets are the same.

The definition of the betweenness relation induced by the orthogonality relation is stated in the following. This definition is in essence that given by Equation (6.1) but it has the additional constraint $b \notin \{a, c\}$. The rationale behind this to ensure that

$\mathcal{B}_\perp(a,b,c)$ should be valid only if b is actually added to the $\perp$-closure and not if it is already contained in it.

Definition 6.2. *Let* $S = (X, Y, \mathcal{B}_\perp, \perp)$ *be a conceptual orthospace. The* betweenness relation induced by $\perp$ on X *(for short $\mathcal{B}_\perp$) is defined as*

$$\mathcal{B}_\perp(a,b,c) \text{ iff } a^\perp \cap c^\perp \subseteq b^\perp \text{ and } b \notin \{a,c\}.$$

In the rest of this section, we discuss properties that an orthogonality relation needs to fulfill in order to establish a given (sub-)set of betweenness axioms as given on page 47. For each of the betweenness axioms, first a counterexample is given and then a possible restriction is presented. An exception is (B1) which is already fulfilled by $\mathcal{B}_\perp$.

Proposition 6.3. *Let S be a conceptual orthospace according to Definition 6.2, then $\mathcal{B}_\perp$ fulfills (B1).*

Proof. $\mathcal{B}_\perp(a,b,c) \Leftrightarrow a^\perp \cap c^\perp \subseteq b^\perp \Leftrightarrow c^\perp \cap a^\perp \subseteq b^\perp$, and as $b \notin \{a,c\}$, it follows $\mathcal{B}_\perp(c,b,a)$. □

The Axiom of Open Betweenness

First, the most basic axiom of betweenness, (B0), is established by stating that an element b different from a and c can only be in $cl_\perp(\{a,c\})$ if a and c are different. Thus, a betweenness relation should only be establishable when both elements a and c are responsible for the betweenness. This circumvents statements of the form $\mathcal{B}_\perp(a,b,x)$ for arbitrary $x \in X \backslash \{b\}$.

Proposition 6.4. *Let S be a conceptual orthospace according to Definition 6.2.*

1. *$\mathcal{B}_\perp$ does not necessarily fulfill (B0).*
2. *If $\perp$ fulfills the conditional*

$$\text{if } a^\perp \cap c^\perp \subseteq b^\perp \text{ and } b \notin \{a,c\}, \text{ then } a \neq c,$$

then $\mathcal{B}_\perp$ fulfills (B0) and (B1).

Proof.

1. Let $X = \{a,b,c\}$ and $a \perp b, a \perp c, b \not\perp c$. Then $b^\perp = c^\perp = \{a\}$ and $c^\perp \cap c^\perp \subseteq b^\perp$ and $b \notin \{c,c\}$ and thus $\mathcal{B}(c,b,c)$, a contradiction to (B0).

2. (B1) is obviously still fulfilled. Ad (B0): Let $\mathcal{B}_\perp(a,b,c)$. Then $a^\perp \cap c^\perp \subseteq b^\perp$ and $b \notin \{a,c\}$ and thus $a \neq c$, therefore, (B0) is valid.

□

Up to this point, $\mathcal{B}_\perp$ does only fulfill (B0) and (B1). We now consider one by one the remaining axioms.

The Axiom of Non-Exchangeability

Next, a restriction for the orthogonality relation to fulfill (B2) is defined, another basic betweenness axiom. The motivation of this restriction is similar to that of (B0), namely, that if $\mathcal{B}_\perp(a,b,c)$ holds then both a and c should influence the betweenness of b. Similar restrictions have also been considered for intersection betweenness [Burigana, 2009]. This restriction includes the restriction of Proposition 6.4 towards fulfilling (B0).

Proposition 6.5. *Let S be a conceptual orthospace according to Definition 6.2.*

1. *$\mathcal{B}_\perp$ does not necessarily fulfill (B2) (even if $\perp$ fulfills the restriction of Proposition 6.4).*

2. *$\mathcal{B}_\perp$ fulfills (B0)–(B2) if $\perp$ has the following property:*

$$\text{if } b \in cl_\perp(\{a,c\}) \text{ and } b \notin \{a,c\} \text{ then}$$
$$\text{there are } x \in a^\perp \backslash c^\perp \text{ and } y \in c^\perp \backslash a^\perp \text{ such that } b \perp x \text{ and } b \perp y.$$

Proof.

1. As a counterexample, assume $X = \{a,b,c,d\}$ and $a^\perp = \{b\}$, $c^\perp = \{d\}$ (and thus $b^\perp = \{a\}$, $d^\perp = \{c\}$). Then, $\mathcal{B}_\perp(x,y,z)$ for distinct $x,y,z \in X$ and especially $\mathcal{B}_\perp(a,b,c)$ and $\mathcal{B}_\perp(b,a,c)$.

2. (B1) is obviously still fulfilled. Ad (B0): Let $\mathcal{B}_\perp(a,b,c)$ and let (B0) be not fulfilled. As $b \notin \{a,c\}$, it must be the case that $a = c$. Then there is no $x \in a^\perp \backslash a^\perp$, a contradiction to the assumption stated in Point 2.

   Ad (B2): Let $\mathcal{B}_\perp(a,b,c)$ for arbitrary $a,b,c \in X$. Then $\mathcal{B}_\perp(b,a,c)$ and $\mathcal{B}_\perp(a,c,b)$ must not be valid to fulfill (B2). Thus, $a^\perp \cap c^\perp \subseteq b^\perp$ but $b^\perp \cap c^\perp \not\subseteq a^\perp$ and $a^\perp \cap b^\perp \not\subseteq c^\perp$. This follows directly based on the assumption stated in Point 2.

□

When $\mathcal{B}_\perp$ fulfills the property of Proposition 6.5, the definition of betweenness $a^\perp \cap c^\perp \subseteq b^\perp$ can be strengthened to $a^\perp \cap c^\perp \subsetneq b^\perp$. This is due to the fact that when $a^\perp \cap c^\perp = b^\perp$, then there is no $x \in a^\perp \backslash c^\perp$ with $x \in b^\perp$. This may evoke the impression that the restriction of $\mathcal{B}_\perp$ is too strong, as it prevents from defining betweenness relations with $a^\perp \cap c^\perp = b^\perp$. This is however not the case. When $a^\perp \cap c^\perp = b^\perp$, then $b^\perp \subseteq a^\perp$ and $b^\perp \subseteq c^\perp$. But this can't be the case based on the condition for (B0) in Proposition 6.4. With the antitonicity rule, it follows that $a^{\perp\perp} \subseteq b^{\perp\perp}$. But as $cl_\perp(\{x\}) = \{x\}$ for $x \in X$ (as $b \in cl_\perp(\{a, c\})$ only if $a \neq c$), it follows $a \subseteq b$ and thus $a = b$, a contradiction to $b \notin \{a, c\}$.

The Axiom of Inner-Transitivity

The axiom of inner-transitive betweenness (B4) can be established through the restriction that for $a, b, c \in X, \mathcal{B}_\perp(a, b, c)$ can only be the case if each element in $b^\perp$ is either in $a^\perp$ or $c^\perp$, i.e., that $b^\perp$ has no members being unique in the sense that they are neither in one of the other two sets, $a^\perp$ and $c^\perp$. This restriction has also been considered for intersection betweenness [Restle, 1959, p. 210]. If this restriction is valid, then the restriction of Proposition 6.5.2 can be simplified. Thus, this restriction also leads towards a fulfillment of (B0) and (B2).

Proposition 6.6. *Let S be a conceptual orthospace according to Definition 6.2.*

1. *$\mathcal{B}_\perp$ does not necessarily fulfill (B4) (even if $\perp$ fulfills the restriction of Proposition 6.5).*
2. *$\mathcal{B}_\perp$ fulfills (B4) if $\perp$ has the following property:*

$$\text{if } b \in cl_\perp(\{a, c\}) \text{ and } b \notin \{a, c\}, \text{ then}$$
$$\text{for all } x \perp b \text{ it follows that } x \perp a \text{ or } x \perp c \text{ (or equivalently } b^\perp \subseteq a^\perp \cup c^\perp).$$

3. *$\mathcal{B}_\perp$ fulfills (B0)–(B2) and (B4) if $\perp$ has the property of 2. and, additionally, $a \neq c$ is enforced.*

Proof.

1. Let $a^\perp, b^\perp, c^\perp$ and $d^\perp$ such that $\mathcal{B}(a, b, d)$ and $\mathcal{B}(b, c, d)$ is fulfilled and let $c^\perp \cap (a^\perp \backslash b^\perp) \backslash d^\perp \neq \emptyset$. Then $a^\perp \cap c^\perp \not\subseteq b^\perp$.
2. Let the antecedent of (B4) be fulfilled, thus $\mathcal{B}_\perp(a, b, d)$ and $\mathcal{B}_\perp(b, c, d)$ are valid. It is shown that $\mathcal{B}_\perp(a, b, c)$ is valid, thus that $a^\perp \cap c^\perp \subseteq b^\perp$. With $\mathcal{B}_\perp(b, c, d)$, it follows that $c^\perp \subseteq b^\perp \cup d^\perp$. Thus, with $a^\perp \cap d^\perp \subseteq b^\perp$, $a^\perp \cap c^\perp \subseteq a^\perp \cap (b^\perp \cup d^\perp) = (a^\perp \cap b^\perp) \cup (a^\perp \cap d^\perp) \subseteq b^\perp$. Thus, (B4) is valid.

3. (B0) and (B1) are trivially fulfilled and (B4) is still fulfilled. Ad (B2): For the sake of contradiction, let $\mathcal{B}_\perp(a,b,c)$ and $\mathcal{B}_\perp(b,a,c)$ be valid. As $b^\perp \subseteq a^\perp \cup c^\perp$, either $b^\perp \subseteq a^\perp$ or there is $x \in b^\perp \cap c^\perp$ with $x \notin a^\perp$. The second is a contradiction to $b^\perp \cap c^\perp \subseteq a^\perp$. Thus, it must be the case that $b^\perp \subseteq a^\perp$. With the antitonicity rule for $\perp$, it follows that $a^{\perp\perp} \subseteq b^{\perp\perp}$. But as $cl_\perp(\{x\}) = \{x\}$ for $x \in X$ (as $b \in cl_\perp(\{a,c\})$ only if $a \neq c$), it follows $a \subseteq b$ and thus $a = b$, a contradiction to the condition.

□

Before we deal with the axiom of outer-transitive betweenness (B3), we provide an example that illustrates how, for a given space and orthogonality relation, a betweenness can be defined and how it is proven that the orthogonality relation actually fulfills the rules leading to a sufficient definition of betweenness.

Example 6.7 (Example 5.13 continued)**.** *Consider a COS $S = (X, Y, \mathcal{B}_\perp, \perp_p)$ as defined in Definition 6.2. Let X be an n-dimensional unit sphere with the point of origin as center (as defined in Example 5.13). The orthogonality relation is defined via polarity with $\perp_p$. As discussed in Example 5.6, this leads to Y being a subset of the intersection of closed convex cones in $\mathbb{R}^n$ with the unit sphere. Thus, $\mathcal{B}_\perp(x,z,y)$ iff $\langle z,a\rangle \leq 0$ for all $a \in X$ with $\langle a,x\rangle \leq 0$ and $\langle a,y\rangle \leq 0$. Because of the bilinearity of the scalar-product $\mathcal{B}_\perp(x,z,y)$ holds for z being the linear combination of x and y: $z = \lambda x + \mu y$ with $\lambda, \mu > 0$, as then $\langle z,a\rangle = \langle \lambda x + \mu y, a\rangle = \lambda\langle x,a\rangle + \mu\langle y,a\rangle \leq 0$. For the case of z being not a linear combination we make use of Farkas' Lemma.*

Lemma 6.8 (Farkas' Lemma [Farkas, 1902])**.** *Let C be the convex cone (convex in the usual Euclidean sense) generated by vectors $v_1, \ldots, v_m \in \mathbb{R}^n$ (i.e., C is the smallest convex cone containing all v_i, thus the conic hull $conH(\{v_1, \ldots, v_m\})$) and let $z \in \mathbb{R}^n$. Then either $z \in C$ or there is $a \in \mathbb{R}^n$ such that $\langle a, v_i\rangle \leq 0$ for all i, and $\langle a,z\rangle > 0$.*

Essentially Farkas' Lemma states that the element z is either between x and y (based on a linear combination) or $\langle z,a\rangle > 0$ for some $a \in X$ with $\langle a,x\rangle \leq 0$ and $\langle a,y\rangle \leq 0$ and thus not $\mathcal{B}_\perp(x,z,y)$. With Farkas' Lemma, it follows that the induced betweenness is exactly $\mathcal{B}_a$, thus a metric betweenness based on the angle as introduced in Section 4.3. For each point $x \in X$ by definition $\{x\} = cl_\perp(\{x\})$. The orthogonality relation fulfills the constraints of Proposition 6.6.3 and thus (B0)–(B2) and (B4), but, as argued in Example 5.13, not (B3).

The Axiom of Outer-Transitivity

Now, after having established (B0)–(B2) and (B4), the question arises whether the axiom of outer transitivity (B3) can be established.

Proposition 6.9. *The betweenness relation $\mathcal{B}_\perp$ of a conceptual orthospace S with $\perp$ fulfilling the properties mentioned in Proposition 6.6.3 does not necessarily fulfill (B3).*

Proof. As a counterexample consider a COS based on a two-dimensional sphere as domain as defined in Example 6.7. (B3) is not fulfilled: Define some a as being at an angle of 0°, b at 90°, c at 170° and d at 200°. The premise of (B3) is fulfilled, however, not $\mathcal{B}_\perp(a, b, d)$. □

In general, when (B3) is not fulfilled it is possible to get some loop such that $\mathcal{B}_\perp(a, b, c)$ and $\mathcal{B}_\perp(b, c, d)$ but $\mathcal{B}_\perp(d, a, b)$ instead of $\mathcal{B}_\perp(a, b, d)$. Furthermore, it is possible to get a chain of betweenness relations, where in between the both endpoints of the chain are elements that are dissimilar to both endpoints. Restle [1959] gives a necessary condition for making (B3) true based on a definition of intersection betweenness fulfilling the restriction of (B4) mentioned in Proposition 6.6.3. We adapt his condition, which he states for finite sets, to arbitrary sets and to the context of orthonegation in the following proposition. The proposition rests on the notion of a linear array. A *linear array* is a sequence of sets $R = \langle {r_1}^\perp, {r_2}^\perp, \dots \rangle$ such that ${r_i}^\perp = A_i \cup B_i \cup C$, $A_i, B_i, C \neq \emptyset$ with $C = {r_1}^\perp \cap {r_2}^\perp \cap \dots$, $A_i \supsetneq A_{i-1}$ and $B_i \subsetneq B_{i-1}$ [Restle, 1959].

Proposition 6.10. *(based on [Restle, 1959, Theorem 5 and 6]). Let $R^* = \langle {r_1}^\perp, {r_2}^\perp, \dots \rangle$ be a sequence of sets representing the negations of $r_1, r_2, \dots$. Then the following equivalence holds for $\perp$ fulfilling the restrictions of Proposition 6.6.3:*

$$\mathcal{B}_\perp(r_i, r_j, r_k) \text{ for all } i < j < k \text{ iff } R^* \text{ is a linear array of sets.}$$

Proof. Restle [1959] proved his Theorems 5 and 6 for finite arrays and states that it could be trivially adapted to infinite ones. □

The difference between this condition and the others is that all others are based on local statements regarding the interplay of at most four elements, whereas this condition considers several betweenness relations making up a chain.

Proposition 6.11. *Let S be a conceptual orthospace. Assume $\perp$ fulfills the condition of Proposition 6.6.3 and the following condition:*

$$\text{if } b \in cl_\perp(\{a, c\}), c \in cl_\perp(\{b, d\}), c \notin \{b, d\} \text{ and } b \notin \{a, c\},$$
$$\text{then } R = \langle a^\perp, b^\perp, c^\perp, d^\perp \rangle \text{ is a linear array.}$$

Then $\mathcal{B}_\perp$ of S fulfills (B3).

Though this depicts a possible restriction, it is — because of the dependence on more than three elements — difficult to restrict the orthogonality relation sufficiently to fulfill these restrictions. Especially, it is not possible to fulfill (B3) if the projective law is valid.

Proposition 6.12. *Let $S = (X, Y, \mathcal{B}_\perp, \perp)$ be a COS defined as in Definition 6.2.3, where $Y = \{cl_\perp(C) \mid C \subseteq X\}$. If $cl_\perp$ fulfills the projective law (pl) (see page 78), then (B3) is not fulfilled or the orthogonality relation is trivial in the sense that $cl_\perp(A) = A$ for all $A \subseteq X$.*

Proof. Let $S = (X, Y, \mathcal{B}_\perp, \perp)$ be a COS defined as in Proposition 6.6.3, where $Y = \{cl_\perp(C) \mid C \subseteq X\}$. For all $a \in X$: $cl_\perp(\{a\}) = \{a\}$, as for $b \in X \backslash \{a\}$ by Proposition 6.6.3, $b \in cl_\perp(\{a\})$ is only the case, when $a^\perp \cap c^\perp \subseteq b^\perp$ with $a \neq c$, a contradiction, as $a = c$. Because of intuitionistic absurdity, it follows that $cl_\perp(\{a\} \cup a^\perp) = X$. With (pl) it follows $cl_\perp(\{a\} \cup a^\perp) = \{cl_\perp(\{a\} \cup y) \mid y \in a^\perp\}$, thus for $b \in cl_\perp(\{a\} \cup a^\perp)$, $b \notin \{a\} \cup a^\perp$ there is $y \in a^\perp$ with $a^\perp \cap y^\perp \subseteq b^\perp$, therefore by definition $\mathcal{B}_\perp(a, b, y)$. Hence, $cl_\perp(\{a\} \cup a^\perp) = \bigcup \{x \mid \mathcal{B}_\perp(a, x, y) \text{ with } y \in a^\perp\} \cup \{a\} \cup a^\perp$. As (B0)–(B2) and (B4) are valid, the rest of the proof can be trivially adapted from the proof of Theorem 5.23. There, it has been shown that for all $a \in X$: $cl_\perp(\{a\}) = \{a\}$ can not be accomplished, when (B0) – (B4) are valid, with a contradiction solely based on the use of axiom (B3). □

There are simple but radical means to restrict the orthogonality relation to fulfill (B3): The orthogonality relation could be constrained to not fulfill the premise of (B3), the space could be constrained in a way that only such a small subset is considered that the negation of a always intersect with the negation of every possible d correctly or the space is separated into subspaces where this is the case. However, these are severe restrictions on the orthogonality relation and do not capture the basic principle of (B3). Finding an orthogonality relation fulfilling (B3) non-trivially is difficult, as then (B3) excludes the existence of chains of arbitrary length containing a point between two points a and d such that this point is in the negation of both. At least for distance-based negations this is not possible.

The Axiom of Density

The betweenness axiom (B5) is not necessarily fulfilled and it is also not always possible to find a restriction such that (B5) is fulfilled, as it restricts the domain to dense ones. However, the orthogonality relation can be restricted such that the induced betweenness fulfills (B5). Axiom (B5) can be directly translated to a condition on the orthogonality relation stating that each $\perp$-closure of two elements must contain at least one additional element.

Proposition 6.13. *Let S be a COS according to Definition 6.2. If $\{a, c\} \subsetneq cl_\perp(\{a, c\})$ (for all $a, c \in X$ with $a \neq c$) then $\mathcal{B}_\perp$ of S fulfills (B5).*

(B5) is strongly connected to (B3). If (B5) is not fulfilled for some elements $a, d \in X$, it is not possible to fulfill (B3) (if (B3) is not trivially fulfilled), because then, e.g., $\mathcal{B}(a, b, c)$ and $\mathcal{B}(b, c, d)$ are valid, but there is no element $x \in X$ such that $\mathcal{B}(a, x, d)$. However, the fulfillment of (B5) obviously does not lead to the fact that (B3) is fulfilled.

Concluding the previous definitions for the satisfaction of the betweenness axioms, we identify in the following a restriction for the orthogonality relation summing up the previous results. As a consequence, if an orthogonality relation is found which fulfills these restrictions, then it can be used to define a conceptual orthospace with an induced betweenness relation fulfilling (B0)–(B2) and (B4).

Proposition 6.14. *Let $S = (X, Y, \mathcal{B}_\perp, \perp)$ be a COS defined as in Definition 6.2. $\mathcal{B}_\perp$ fulfills (B0)–(B2) and (B4) if for all $b \in cl_\perp(\{a, c\})$ with $b \notin \{a, c\}$ the following conditions are fulfilled:*

- *$a \neq c$ and*
- *for all $x \perp b$ it holds that $x \perp a$ or $x \perp c$.*

Induced Betweenness Based on Arbitrary Orthogonality Relations

As stated at the beginning of this section, there are two possibilities of inferring an orthonegation-induced betweenness: Either the orthogonality relation is restricted (as done above) and a betweenness relation is established containing as many triples as possible. Or the betweenness relation is restricted and the induced betweenness can be established for an arbitrary orthogonality relation but at the price of having only a subset of the betweenness relations.

In some applications, it might be desirable to allow for arbitrary orthogonality relations. Hence, in the following proposition, an adaptation of the definition in Proposition 6.14 to the case of restricting the betweenness is presented. Whereas the first and third condition are directly taken from the condition on the orthogonality relation (but interpreted as condition on the betweenness relation), the second condition regarding inequality of elements has to be added. The reason is that in contrast to the restricted orthogonality relation, the condition of $a \neq b \neq c$ does not enforce that $cl_\perp(\{x\}) = \{x\}$ for $x \in X$. Therefore, two different elements could have the same negation and, thus, as an additional constraint, $a^\perp \neq b^\perp \neq c^\perp$ must be added.

Proposition 6.15. *Let* $S = (X, Y, \mathcal{B}'_\perp, \perp)$ *for an arbitrary orthogonality relation* $\perp$. $\mathcal{B}'_\perp$ *fulfills (B0)–(B2) and (B4) if for all* $a, b, c \in X : \mathcal{B}'_\perp(a, b, c)$ *iff*

1. $a^\perp \cap c^\perp \subseteq b^\perp$ *and*
2. $a^\perp \neq b^\perp \neq c^\perp$ *(and thus* $a \neq b \neq c$*) and*
3. *for all* $x \perp b$ *it holds that* $x \perp a$ *or* $x \perp c$ *(equivalently* $b^\perp \subseteq a^\perp \cup c^\perp$*).*

Proof. Ad (B0): Follows directly out of condition 2.

Ad (B1): Trivial extension of the proof of Proposition 6.3.

Ad (B2): For the sake of contradiction, let $\mathcal{B}_\perp(a, b, c)$ and $\mathcal{B}_\perp(b, a, c)$ be valid. As $b^\perp \subseteq a^\perp \cup c^\perp$, either $b^\perp \subseteq a^\perp$ or there is $x \in b^\perp \cap c^\perp$ with $x \notin a^\perp$. The second is a contradiction to $b^\perp \cap c^\perp \subseteq a^\perp$. Thus, it must be the case that $b^\perp \subseteq a^\perp$. As $a^\perp \subseteq b^\perp \cup c^\perp$ and $b^\perp \subsetneq a^\perp$ (because of condition 2.), there is $x \in a^\perp \cap c^\perp$ with $x \notin b^\perp$, a contradiction to $a^\perp \cap c^\perp \subseteq b^\perp$.

Ad (B4): Assume that $\mathcal{B}'_\perp(a, b, d)$ and $\mathcal{B}'_\perp(b, c, d)$ is valid. The conditions for $\mathcal{B}'_\perp(a, b, c)$ are considered case by case. First, $a^\perp \cap c^\perp \subseteq b^\perp$, as $a^\perp \cap d^\perp \subseteq b^\perp$ and $c^\perp \subseteq b^\perp \cup d^\perp$ and thus $a^\perp \cap c^\perp \subseteq a^\perp \cap (b^\perp \cup d^\perp) = (a^\perp \cap b^\perp) \cup (a^\perp \cap d^\perp) \subseteq b^\perp$. Second, $a^\perp \neq b^\perp \neq c^\perp$: $a^\perp \neq b^\perp$ and $b^\perp \neq c^\perp$ follows by assumption. Thus, assume for the sake of contradiction that $a^\perp = c^\perp$. $c^\perp \subseteq b^\perp \cup d^\perp$ and $a^\perp \cap d^\perp \subseteq b^\perp$ (and thus $c^\perp \cap d^\perp \subseteq b^\perp$). Assume $c^\perp \nsubseteq b^\perp$, thus there is $x \in c^\perp$ with $x \in d^\perp \setminus b^\perp$, a contradiction to $c^\perp \cap d^\perp \subseteq b^\perp$. Thus, $c^\perp \subseteq b^\perp$ and with $b^\perp \neq c^\perp$, it follows $c^\perp \subsetneq b^\perp$ (and thus $a^\perp \subsetneq b^\perp$). With $b^\perp \subseteq a^\perp \cup d^\perp$, there is $x \in b^\perp \cap d^\perp$ with $x \notin a^\perp$ (and thus $x \notin c^\perp$), a contradiction to $b^\perp \cap d^\perp \subseteq c^\perp$. Third, assume for the sake of contradiction, $b^\perp \nsubseteq a^\perp \cup c^\perp$, thus there is $x \in b^\perp$ with $x \notin a^\perp, x \notin c^\perp$. As $b^\perp \subseteq a^\perp \cup d^\perp$ and $x \notin a^\perp$, it follows that $x \in d^\perp$. But as $b^\perp \cap d^\perp \subseteq c^\perp$ and $x \in b^\perp \cap d^\perp$, it follows $x \in c^\perp$, a contradiction. □

The constraint for (B3) is not adapted in the same manner here, as the restriction mentioned in Proposition 6.11 is based on four (or more) elements. If (B3) is not fulfilled directly by the induced betweenness relation, then it is not at all clear which betweenness relations should be deleted to fulfill (B3). The other conditions lead to a unique result, whereas this one does not and thus is omitted here.

Properties of a Conceptual Orthospace with Induced Betweenness

After regaining the betweenness axioms (B0)–(B2) and (B4) (and to some extent also (B3) and (B5)), here some notes regarding the properties of the resulting COS are stated. The observation is that the induced betweenness ensures that all $\perp$-closed sets are convex, but not vice versa that all convex sets are $\perp$-closed.

Proposition 6.16. *1. Let S be a conceptual orthospace according to Definition 6.2 or Proposition 6.14 or Proposition 6.15. Then, for $C \subseteq X$: If C is $\perp$-closed, then it is convex.*

2. Let S be a conceptual orthospace according to Definition 6.2 or Proposition 6.15. Then, for $C \subseteq X$: If C is convex, then it is not necessarily $\perp$-closed.

Proof.

1. Let $C \subseteq X$ be $\perp$-closed and assume C is not convex. Then there is $b \in X \backslash C$ and $a, c \in C$ with $\mathcal{B}_\perp(a, b, c)$ resp. $\mathcal{B}'_\perp(a, b, c)$, but then by definition of $\mathcal{B}_\perp$, $a^\perp \cap c^\perp \subseteq b^\perp$ and thus $b \in cl_\perp(\{a, c\})$. By the rules of closure operators (see Definition 4.1.2), if $\{a, c\} \subseteq cl_\perp(C)$, then $cl_\perp(\{a, c\}) \subseteq cl_\perp(C)$ and thus $b \in C$, a contradiction.
2. As a counterexample, Example 5.20 can be used. There, the betweenness is exactly defined as orthonegation-induced betweenness and it contains a convex but not $\perp$-closed set.

□

6.1.2 The Induced Betweenness of Goodman as Special Case

An induced betweenness relation situated in a quite different setting is that of Goodman [1977], introduced in Section 4.3. In the following, I show that it can be adapted to be the basis of a COS and that it can be considered as a subset of the orthonegation-induced betweenness defined in Proposition 6.15.

Let $\perp_g$ be defined based on the similarity relation $\sim$ (called "matching relation" by Goodman) as $x \not\perp_g y$ iff $x \sim y$ for $x, y \in X$. Let $A \triangle B = A \setminus B \cup B \setminus A$ denote the symmetric difference of two sets A, B. A corresponding notion of symmetric difference $x \oplus y$ for elements $x, y \in X$ can be defined as follows:

$$x \oplus y = \{z \in X \mid z \sim x\} \triangle \{z \in X \mid z \sim y\}$$

Thus, $x \oplus y$ denotes the set of elements z which are similar to one of x, y but not the other. Now, the adaptation of Goodman's betweenness relation is given in the following definition.

Definition 6.17. *The* Goodman-style betweenness relation based on subset-inclusion $\mathcal{B}_g$ *is defined as follows:*

$$\mathcal{B}_g(x, y, z) \quad \textit{iff} \quad x \sim y \;\&\; y \sim z \;\&\; x \sim z \;\&\; x \oplus y \subsetneq x \oplus z \;\&\; y \oplus z \subsetneq x \oplus z \quad (6.2)$$

This mimics Goodman's construction (see Definition 10.02 of Goodman [1977] and Section 4.3) but replaces the greater-than relation G with proper set inclusion. The reason is that we — in contrast to Goodman — do not assume that X is finite.

Which of the betweenness axioms (Bi) are fulfilled? We show that (B0), (B1), (B2), and (B4) are fulfilled and construct a counterexample for (B3).

In the proof of (B4) we need the following lemma.

Lemma 6.18. *If $x \oplus y \subseteq x \oplus z$ holds, then: $x \oplus y \subsetneq x \oplus z$ is the case iff $y \oplus z \neq \emptyset$.*

Proof. Assume $x \oplus y \subsetneq x \oplus z$, then there is $v \in x \oplus z$ and $v \notin x \oplus y$. If $v \sim x$ and not $v \sim z$, then we must have $v \sim y$ and hence $v \in z \oplus y$. If not $v \sim x$ and $v \sim z$, then $v \sim y$ cannot be the case and we again have $v \in z \oplus y$. Assume now that $y \oplus z \neq \emptyset$, say with $v \in z \oplus y$, i.e., $v \sim z$ iff not $v \sim y$. We have to show that $x \oplus y \subsetneq x \oplus z$ holds. If not, then we have $x \oplus y = x \oplus z$. On the one hand $v \in x \oplus z$ iff ($v \sim x$ iff not $v \sim z$). On the other hand $v \in x \oplus y$ iff ($v \sim x$ iff not $v \sim$). So we have ($v \sim x$ iff not $v \sim z$) iff ($v \sim x$ iff not $v \sim y$). Contradiction (with $v \sim z$ iff not $v \sim y$). □

Proposition 6.19. *The Goodman-style betweenness relation according to Equation* (6.2) *fulfills (B0), (B1), (B2), and (B4) but in general not (B3).*

Proof. Ad (B0): If $\mathcal{B}_g(x, y, z)$, then x, y, z are distinct. The definition guarantees that x and z must be different. Because otherwise $x \oplus z = \emptyset$ and the empty set cannot have a proper subset. If $x = y$ were the case, then substituting in the definition of $\mathcal{B}_g(x, y, z)$ the y with x wold lead to $x \oplus z \subsetneq x \oplus z$, a contradiction. Analogously the assumption $y = z$ leads to a contradiction.

Ad (B1): If $\mathcal{B}_g(x, y, z)$, then $\mathcal{B}_g(z, y, x)$. This follows from the fact that $\oplus$ is symmetric: let $\mathcal{B}_g(x, y, z)$, i.e., $x \sim y \;\&\; y \sim z \;\&\; x \sim z \;\&\; x \oplus y \subsetneq x \oplus z \;\&\; y \oplus z \subsetneq x \oplus z$. Then $x \sim y \;\&\; y \sim z \;\&\; x \sim z \;\&\; z \oplus y \subsetneq z \oplus x \;\&\; y \oplus x \subsetneq z \oplus x$, i.e., $\mathcal{B}_g(z, y, x)$.

Ad (B2): If $\mathcal{B}_g(x, y, z)$, then not $\mathcal{B}_g(y, x, z)$. Let $\mathcal{B}_g(x, y, z)$, i.e., $x \sim y \;\&\; y \sim z \;\&\; x \sim z \;\&\; x \oplus y \subsetneq x \oplus z \;\&\; y \oplus z \subsetneq x \oplus z$ and assume $\mathcal{B}_g(y, x, z)$. Then the former implies $y \oplus z \subsetneq x \oplus z$, whereas the latter implies $x \oplus z \subsetneq y \oplus z$, a contradiction.

Ad (B3): Not necessarily: if $\mathcal{B}_g(x, y, z)$ and $\mathcal{B}_g(y, z, w)$, then $\mathcal{B}_g(x, y, w)$. Consider the following counterexample: Let $X = \{a, b, c, d, e\}$ and $\sim$ with $a \sim b, a \sim c, b \sim c, b \sim d, c \sim d, c \sim e$ and $d \sim e$ and assertions $v \sim v$ and if $v \sim w$, then also $w \sim v$ for all $v, w \in X$. Then: $a \oplus c = \{d, e\}, a \oplus b = \{d\}, b \oplus c = \{e\}$, hence $\mathcal{B}_g(a, b, c)$; moreover, $b \oplus d = \{a, e\}, c \oplus d = \{a\}$, hence $\mathcal{B}_g(b, c, d)$. But $a \sim d$ does not hold, hence $\mathcal{B}_g(a, b, d)$ does not hold.

Ad (B4): If $\mathcal{B}_g(x, y, z)$ and $\mathcal{B}_g(y, w, z)$, then $\mathcal{B}_g(x, y, w)$. Let $\mathcal{B}_g(x, y, z)$, i.e.,

$$x \sim y \;\&\; y \sim z \;\&\; x \sim z \;\& \tag{6.3}$$
$$x \oplus y \subsetneq x \oplus z \;\& \tag{6.4}$$
$$y \oplus z \subsetneq x \oplus z \tag{6.5}$$

and assume that $\mathcal{B}_g(y, w, z)$ holds, i.e.,

$$w \sim y \;\&\; y \sim z \;\&\; w \sim z \;\& \tag{6.6}$$
$$y \oplus w \subsetneq y \oplus z \;\& \tag{6.7}$$
$$w \oplus z \subsetneq y \oplus z \tag{6.8}$$

We have to show: $\mathcal{B}_g(x, y, w)$, i.e., that the following conditions hold:

$$x \sim y \;\&\; y \sim w \;\&\; x \sim w \;\& \tag{6.9}$$
$$x \oplus y \subsetneq x \oplus w \;\& \tag{6.10}$$
$$y \oplus w \subsetneq x \oplus w \tag{6.11}$$

Proof of (6.9): We have trivially $x \sim y$ and $w \sim y$. But we also have $x \sim w$ due to the following: $y \oplus w \subsetneq y \oplus z$ hold. Now, as $y \sim x$ and $z \sim x$ it follows that $x \notin y \oplus z$. Assume for sake of contradiction that not $x \sim w$. As $y \sim x$, this would give us $x \in y \oplus w$, but then x would have to be in $y \oplus z$, leading to a contradiction.

Proof of (6.10): $x \oplus y \subsetneq x \oplus w$. We show first that the $\subseteq$ relation holds: Assume that v exists with $v \in x \oplus y$ and thus $v \in x \oplus z$. Case 1: $v \sim x$ and not $v \sim y$. Thus, not $v \sim z$. We have to show that $v \in x \oplus w$, i.e., not $v \sim w$. If $v \sim w$ were the case, then $v \in w \oplus z \subsetneq y \oplus z$, thus $v \sim y$, a contradiction. Case 2 is analog to case 1: not $v \sim x$ and $v \sim y$. Thus, $v \sim z$. We have to show that $v \in x \oplus w$, i.e., $v \sim w$. If not $v \sim w$ were the case, then $v \in w \oplus z \subsetneq y \oplus z$, thus not $v \sim y$, a contradiction.

Now we are left with showing that $x \oplus y$ is a proper subset of $x \oplus w$ or in other words that $x \oplus w$ is not a subset of $x \oplus y$. Due to Lemma 6.18, in order to show $x \oplus y \subsetneq x \oplus w$, we have to show that $y \oplus w$ is not empty. But this is due to Lemma 6.18 and (6.8).

Proof of (6.11): We have to show $y \oplus w \subsetneq x \oplus w$. We start again by showing $\subseteq$: Let $v \in y \oplus w$. Case 1: $v \sim y$ and not $v \sim w$. We have to show $v \in x \oplus w$ and thus $v \sim x$. Assume not $v \sim x$. But then (as $v \sim y$) $v \in x \oplus y$. Due to (6.10) we then have $v \in x \oplus w$, giving us a contradiction. Similar for case 2 having not $v \sim y$ and $v \sim w$. We are left with showing that $x \oplus w$ is not a subset of $y \oplus w$. Due to Lemma 6.18 we have to show that $y \oplus x$ is not empty which is due to Lemma 6.18 and Section 6.1.2. □

As Goodman's construction is intended for sensory impressions, the matching operation is assumed to be a similarity relation which is usually not transitive. And in fact, the betweenness construction according to Goodman does not work out for arbitrary orthogonality relations. Consider an orthogonality relation $\perp_g$ such that the following holds:

$$\text{For all } x, y, z \in X\text{: if } x \perp_g z\text{, then } x \perp_g y \text{ or } y \perp_g z. \qquad \text{(Trans*)}$$

This condition is equivalent to the condition stating that $\sim$ is transitive[1], i.e., as $\sim$ already is assumed to be symmetric and reflexive, that $\sim$ is an equivalence relation. In this case the betweenness relation becomes trivial: if $\perp_g$ fulfills (Trans*) then no triple of elements stands in $\mathcal{B}_g$-relation. The reason is that as $\sim$ becomes transitive, $x \sim y$ and $y \sim z$ and $y \sim z$ would mean that, e.g., for all $v \in X$ with $y \sim v$, it follows $z \sim v$, therefore $x \oplus y = y \oplus z = x \oplus z = \emptyset$. But then $x \oplus z$ cannot have any proper subset.

Having an arbitrary orthogonality relation and a betweenness-relation based on it according to Definition 6.17, it is possible to show that the $\perp$-closed sets are actually convex.

Proposition 6.20. *If a set $A \subseteq X$ is $\perp$-closed, then it is convex.*

Proof. Assume $A \subseteq X$ to be $\perp$-closed and let $a, c \in A$. Assume that A is not convex. Thus, there is a $b \notin A$ with $\mathcal{B}_g(a, b, c)$. $b \notin A$ means that there is a $d \in A^\perp$ with $d \sim b$. By definition of $\mathcal{B}_g$, $a \oplus b \subsetneq a \oplus c$. As $a \perp_g d$, $d \in a \oplus b$, thus $d \in a \oplus c$, thus $d \sim c$. A contradiction, as then $d \notin A^\perp$. □

Thus, the betweenness relation of Goodman makes up a conceptual orthospace, when, e.g., identifying Y with the set of all $\perp$-closed sets. Thus, $S = (X, Y, \mathcal{B}_g, \perp_g)$ is a COS for $Y \subseteq \{cl_\perp(A) \mid A \subseteq X\}$ closed under intersection and orthonegation and an arbitrary domain X. In the following, it is shown that this betweenness relation can be considered as an instance of the orthonegation-induced betweenness defined in Proposition 6.15.

Proposition 6.21. *Let $S_1 = (X, Y_1, \mathcal{B}'_\perp, \perp_g)$ and $S_2 = (X, Y_2, \mathcal{B}_g, \perp_g)$ with $\mathcal{B}'_\perp$ as defined in Proposition 6.15. For all $a, b, c \in X$:*

$$\textit{if } \mathcal{B}_g(a, b, c)\textit{, then } \mathcal{B}'_\perp(a, b, c)$$

Proof. Let for $a, b, c \in X$, $\mathcal{B}_g(a, b, c)$ be valid. The conditions of Proposition 6.15 are considered one by one:

[1] With "if $x \perp_g z$, then $x \perp_g y$ or $y \perp_g z$", it follows that "if $x \sim y$ and $y \sim z$, then $x \sim z$", thus transitivity.

1. $a^\perp \cap c^\perp \subseteq b^\perp$: Let $a^\perp \cap c^\perp \not\subseteq b^\perp$, therefore there is an $x \in a^\perp \cap c^\perp$ with $x \notin b^\perp$. Then, $x \in a \oplus b$ but $x \notin a \oplus c$, a contradiction, as $\mathcal{B}_g(a, b, c)$ is only the case if $a \oplus b \subsetneq a \oplus c$.

2. $a^\perp \neq b^\perp \neq c^\perp$: With $a \oplus b \subsetneq a \oplus c$, it follows that $a \oplus c \neq \emptyset$ and thus $a^\perp \neq c^\perp$. Assume for the sake of contradiction that $a^\perp = b^\perp$, then $b \oplus c = a \oplus c$, a contradiction to $b \oplus c \subsetneq a \oplus c$ (analogously for $b^\perp = c^\perp$).

3. $b^\perp \subseteq a^\perp \cup c^\perp$: Let $b^\perp \not\subseteq a^\perp \cup c^\perp$, thus there is an $x \in b^\perp$ with $x \notin a^\perp$ and $x \notin c^\perp$. Then $x \in a \oplus b$, but $x \sim a$, $x \sim c$, thus $x \notin a \oplus c$, a contradiction.

□

6.1.3 Application of an Orthonegation-induced Betweenness

Whereas in the previous subsection two definitions of an orthonegation-induced betweenness relation were given, in this section I deal with questions of applying them. Concrete, I deal with the questions whether some specific similarity relations are usable as a basis for the definition of an orthonegation-induced betweenness in the sense of whether they fulfill the restrictions of Proposition 6.14 and whether it is possible to gain an induced betweenness relation in the style of Proposition 6.15. In the following, the term "orthonegation-induced betweenness" relates to the definition of it in the last section and does not cover orthonegation-induced betweenness relations in general.

Before answering the questions, I want to discuss why the identification of an induced betweenness relation is actually helpful. As explained in Chapters 4 and 5, similarity and betweenness are both vital for a good embedding approach. However, when starting with a betweenness relation, it is possible to enforce a specific strength of the betweenness. This is not possible when having a similarity relation given. In this case, the determination of the betweenness relation is passive in the sense that it is enforced by the similarity relation to have a specific strength. One case where it is not possible to rely on a given similarity relation is when background knowledge is given and the resulting COS based on an orthonegation-induced betweenness interferes with this background information. Therefore, the induced betweenness can only be used when either the given background knowledge matches the resulting COS or the background knowledge is weak enough for not interfering or when only the type of similarity relation is given as basis, however, the similarity of instances is determined via an embedding approach.

It could be the case that the concepts are convex by definition of the induced betweenness relation, but may, however, not be convex based on Euclidean betweenness or another expressive notion of betweenness. In such cases, the arguments for the advantages of convex concepts cannot be applied directly, e.g., they have not necessarily obvious computational benefits. The dissimilarity relation itself already results in an ortholattice. Ortholattices can represent negation, conjunction and disjunction sufficiently — so why should betweenness be specified?

The first reason is that the determination of the betweenness allows for the examination of the underlying structure of the space. This structure is not necessarily recognizable solely based on the orthogonality relation given and, thus, considering this structure in detail could help to use the advantages of convexity not visible beforehand.

In other cases, when one recognizes that the structure is weak, an option would be to dismiss this type of orthogonality relation and use a different one which induces a more expressive betweenness relation. It is also possible to rate the orthogonality relations based on their ability to allow for expressive induced betweenness relations. As will turn out later in this section where some examples of induced betweenness are considered, an induced betweenness based on the Euclidean metric leads to a less expressive betweenness as the one induced by the uniform norm. Though both are less expressive than, e.g., the Euclidean betweenness, the latter one does at least contain betweenness relations at all.

This directly leads to the second option: If the orthogonality relation has to be kept, it is at least possible to specify the structure as much as possible. It is, e.g., possible that the concepts are locally stronger structured, thus, that, e.g., (B5) is valid, when only one region and not the whole space is considered and, thus, the advantages of convexity can be used locally. It is also possible to understand the structure of the data better by recognizing which elements are influential in the sense that they are part of a betweenness relation. The resulting COS need to be considered with care, though, as it is mostly only weakly structured.

In the following paragraphs, the induced betweenness is determined for some well-known similarity relations.

Minkowski Distance-based Orthogonality Relation

As mentioned in Section 4.1, distance-based similarity relations are widely used in many contexts. Here, $\perp_m$ is defined as an orthogonality relation based on the Minkowski distance $d_m(x,y) = \sum_{i=1}^{m}(|x_i - y_i|^p)^{\frac{1}{p}}$ for $p > 1$ and $p < \infty$ ($p = 1$ and $p = \infty$ are considered as special cases later on), by using a threshold $\tau \geq 0$, resulting in $x \perp_m y$ iff $d_m(x,y) > \tau$. However, it turns out that the induced betweenness relation is empty. An illustration of the problem for the special case of a Euclidean metric and arbitrary $\tau \geq 0$ can be seen in Figure 6.1(a).

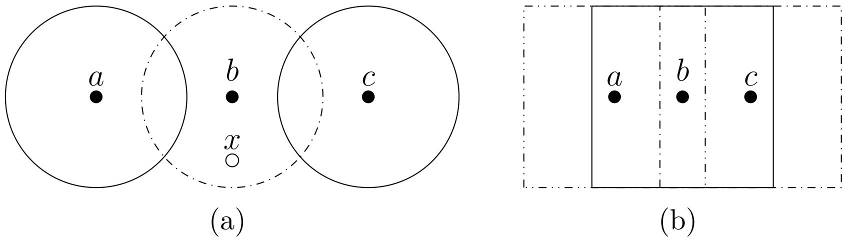

Figure 6.1: (a) Example for an induced betweenness based on Euclidean distance: For each of the three points the respective $\cdot^\perp$ is outside of the respective circle. There is an $x \in a^\perp \cap c^\perp$ with $x \sim b$, thus $x \notin b^\perp$ and thus not $\mathcal{B}'_\perp(a, b, c)$; (b) Example for an induced betweenness based on uniform distance (where again the space outside the respective square depicts the resp. negated concept): For the orthogonality relation interpreted as uniform norm, these a, b, c lead to $\mathcal{B}'_\perp(a, b, c)$.

Proposition 6.22. *For $S = (\mathbb{R}^n, Y, \mathcal{B}'_\perp, \perp_m)$, the induced betweenness relation $\mathcal{B}'_\perp$ is empty (independent of the choice of τ).*

Proof. Assume for the sake of contradiction that there are $a, b, c \in \mathbb{R}^n$ with $\mathcal{B}'_\perp(a, b, c)$. With Proposition 6.15, it follows that $a \neq b \neq c$. Thus, $d_m(a, b) > 0$. Choose x_1 and x_2 orthogonal to the line ab such that $d_m(x_1, b) = \tau$ and $d_m(x_2, b) = \tau$ with $d_m(x_1, b) + d_m(x_2, b) = d_m(x_1, x_2)$, thus b lying in between x_1 and x_2 in the usual sense. Then, $d_m(a, x_1) > \tau$ and $d_m(a, x_2) > \tau$, thus $x_1, x_2 \in a^\perp$. As $x_1, x_2 \notin b^\perp$ it must be the case that $x_1, x_2 \notin c^\perp$ to allow for $\mathcal{B}'_\perp(a, b, c)$. Based on the definition of the Minkowski-distance, this is only possible if $b = c$, a contradiction. Thus, it is not possible to choose $a, b, c \in \mathbb{R}^n$ such that $\mathcal{B}'_\perp(a, b, c)$ and thus the betweenness is empty. □

Uniform Distance-based Orthogonality Relation

Next, the uniform distance is considered, which is defined for $x, y \in \mathbb{R}^n$ by

$$d_u(x, y) = \max_{i=1}^{n} |x_i - y_i|.$$

The orthogonality relation is again defined as above as $x \perp_u y$ iff $d_u(x, y) > \tau$ for arbitrary $\tau \geq 0$. An example can be seen in Figure 6.1(b). It turns out that the problem as for the Euclidean distance does not occur. In particular, instances are in between each other if they are on a line parallel to the axes and restricted to a distance between the outer points smaller than or equal 2τ. When looking at Figure 6.1(b) the negated concept is still not convex in the classical Euclidean sense. But in contrast to the Euclidean distance-based negation of Figure 6.1(a), it is at least possible to define betweenness locally, thus the advantages of convexity can be used at least locally. Of course, for many applications, it is desirable to regain stronger convexity. This is, e.g., possible when considering angle-based similarity as done in the next paragraph.

The next proposition identifies the possible betweenness relations induced by a uniform distance.

Proposition 6.23. *For $S = (\mathbb{R}^n, Y, \mathcal{B}'_\perp, \perp_u)$, the induced betweenness relation is*

$$\mathcal{B}'_\perp(a,b,c) \text{ iff } d_u(a,b) + d_u(b,c) = d_u(a,c) \;\&\; d_u(a,c) \leq 2\tau \;\& \\ \text{the line through } a,b,c \text{ is parallel to an axis.}$$

Proof. Follows trivially out of the definition of the uniform distance. □

The Manhattan-distance, defined for $x, y \in \mathbb{R}^n$ by

$$d_{ma}(x,y) = \sum_{i=1}^{n} |x_i - y_i|$$

behaves similarly, instead of allowing for betweenness parallel to the axes allowing for betweenness diagonal to the axes.

Angle-based Orthogonality Relation

The application of the orthogonality based on the angle between instances has already been considered in Example 6.7 for the special case of an angle of 90°. There, it turned out that the induced betweenness is the betweenness based on the angle. Thus, an orthogonality relation based on the angle is a well-working orthogonality relation leading to a strong betweenness relation. As shown in Example 6.7, in particular choosing the domain X as a sphere leads to an expressive induced betweenness relation, as $\perp_p$ then fulfills the restrictions of Proposition 6.14. Even when considering $X = \mathbb{R}^n$, the induced betweenness in form of the betweenness based on the angle is a quite strong betweenness.

This does not work out if an angle different from 90° is chosen, i.e., setting

$$x \perp_p' y \text{ iff } \angle(x,y) \geq \tau$$

for $x, y \in \mathbb{R}^n$ and $\tau \neq 90°$. In this case, the same problem occurs as for the Minkowski distance: the induced betweenness relation is empty for $n > 2$. Only when considering the two-dimensional case, the induced betweenness is non-trivial and usable.

Proposition 6.24. *For* $S = (\mathbb{R}^n, Y, \mathcal{B}'_\perp, {\perp_p}')$*, the induced betweenness relation is empty if* $\tau \neq 90°$ *and* $n > 2$*.*

Proof. Let a $S = (\mathbb{R}^n, Y, \mathcal{B}'_\perp, {\perp_p}')$ be given with $n > 2$. For the case of $\tau < 90°$, based on the same argument as used in the proof of Proposition 6.22, it follows that the induced betweenness relation is trivial.

For the case of $\tau > 90°$, it is possible to fulfill the statements of Proposition 6.15.1 and Proposition 6.15.2 but not Proposition 6.15.3 which can be again argued for based on the proof of Proposition 6.22. □

6.2 Betweenness-induced Orthogonality

In the last section, one variant of an orthonegation-induced betweenness relation was presented. In the discussion of the application of such relations in Section 6.1.3, it turned out that, although defining such an induced betweenness could be helpful, it is, in most cases, impossible to regain a strong betweenness relation. There are many use cases imaginable in which a strong betweenness is needed and the user hence would like to define a COS based on the desired betweenness relation and choose a suitable orthogonality relation based on it.

I tackle this problem in several ways: On the one hand, the special case of Euclidean betweenness is discussed in Chapter 7; on the other hand, the induced orthogonality relation based on arbitrary betweenness relations is showcased in this section. In contrast to the case of orthonegation-induced betweenness, though, it is, in this case, not possible to use a general construction principle for a betweenness-induced orthogonality relation, as argued in the following and illustrated in Example 6.25.

The main problem is the different influence of the two closure operators on the resulting COS. Whereas the induced betweenness for an arbitrary orthogonality relation could be empty — resulting in a correct (but not useful) COS — the orthogonality relation is not allowed to be trivial. In particular, for all $A, B \in Y$, $conv(A \cup B) \subseteq cl_\perp(A \cup B)$, thus for a non-trivial betweenness, the $\perp$-closure is also not allowed to be trivial: If $\perp$ is trivial in the sense that for all $x, y \in X, x \perp y$, then $cl_\perp(\{x, y\}) = \{x, y\}$ for all $x, y \in X$ and thus for a non-trivial betweenness $cl_\perp(\{x, y\}) \subset conv(\{x, y\})$ for some $x, y \in X$. If $\perp$ is trivial in the sense that for all $x, y \in X, x \not\perp y$, then $x^\perp = \{\emptyset\}$ and $cl_\perp(\{x\}) = X$, thus there are no $\perp$-closed sets beside $\emptyset$ and X and thus the resulting COS is trivial. A naive solution for this problem would be to define $cl_\perp(A \cup B) = conv(A \cup B)$ for all $A, B \in Y$ with $Y \subseteq \{conv(C) \mid C \subseteq X\}$: this definition does not lead directly to a definition of the orthogonality relation. The only information is that the rules of an orthogonality relation need to be fulfilled.

This leads to the observation that the definition of a betweenness-induced orthogonality relation is highly dependent on the chosen betweenness, the chosen domain X and also on the desired Y. For orthonegation-induced betweenness, the orthoframe $(X, \perp)$ makes up an ortholattice which is preserved in the constructed COS (as Y contains all $\perp$-closed sets). However, if in addition to an ortholattice, a domain X and a betweenness relation are given, this ortholattice may not be representable anymore — even based on an arbitrary orthogonality relation. This is illustrated in the following example.

Example 6.25. *Let $X = \mathbb{R}$ and let the betweenness be the Euclidean one. The problem of defining an orthogonality relation based on a given betweenness becomes relevant if*

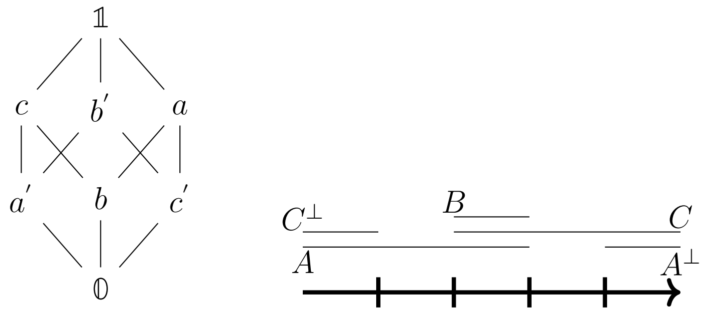

(a) (b)

Figure 6.2: Illustration of the problem of inducing an orthogonality relation mentioned in Example 6.25. The lattice in (a) should be represented, (b) shows one possible construction. For illustration purposes, here the concepts are represented below each other but are all placed in $\mathbb{R}$.

complex lattices are needed to be modeled. Therefore, assume that $\mathcal{L}_Y$ is the lattice represented in Figure 6.2(a), a simple, distributive lattice enforcing some non-trivial conjunctions and disjunctions for the resulting COS. Let $A, B, C \subseteq X$ with A, B, C be convex such that $A \cap C = B$. Thus, A, B, C are all intervals of X. An example for these three sets can be seen in Figure 6.2(b). Now the question arises as to how the orthogonality relation is modeled, in particular, what $A^\perp, B^\perp$ and $C^\perp$ need to look like. As $a^{'} \leq c$ and $c^{'} \leq a$, the construction has to look like in Figure 6.2(b) or similar. But independent of the actual construction, it is the case that $B^\perp = cl_\perp(A^\perp \cup C^\perp)$ and as $conv(A^\perp \cup C^\perp) \subseteq cl_\perp(A^\perp \cup C^\perp)$, it follows that $B \subseteq B^\perp$, a contradiction. Thus, it is not possible to construct a COS modeling the proposed lattice based on the Euclidean betweenness. However, when dismissing the betweenness, it is possible to model the lattice in the given domain, and, thus, to find a suitable orthoframe (in fact, the exact same one as seen in Figure 6.2(b) can be used when considering a weaker betweenness).

This illustrates that given a domain and a betweenness relation it is not always possible to construct a COS of sufficient expressivity.

Another related question regarding the process of constructing an induced orthogonality relation is how Y should be defined. Following the construction of a COS based on an induced betweenness, where $Y = \{cl_\perp(C) \mid C \subseteq X\}$, the aim is to incorporate the betweenness in the general construction as much as possible, thus enforcing the orthogonality relation to be of the same strength as the betweenness relation. The first idea would be to require completeness. This, however, is difficult to gain (as has been discussed in Section 5.3.1 and will turn out for the case of Euclidean betweenness in

Chapter 7). Therefore, the relaxation in form of (R4) can be required, enabling for as many betweenness relations as possible to be used as the basis for an induced orthogonality relation:

$$cl_{\perp}(A \cup B) = conv(A \cup B) \text{ for } A, B \in Y.$$

However, even in this case, it is, as discussed in Example 6.25, sometimes impossible to construct an expressive COS at all. Even when an expressive COS can be constructed, it is possible that only a small subset of the sets, convex based on the betweenness relation, can be modeled. This is illustrated in Chapter 7 in which I show for the Euclidean betweenness that based on some natural restrictions, the orthogonality relation needs to be the polarity and the set of concepts considered needs to be a subset of closed convex cones.

As illustrated in Chapter 7, the problem of finding an induced orthogonality relation by no means leads to the fact that COSs based on a given betweenness are problematic. In contrast, it states that due to the difference in expressivity of different betweenness relations, it it necessary to construct a COS for each betweenness individually to get the best solution.

Part II

Applications of Conceptual Orthospaces

7 Euclidean Convexity as Basis of Conceptual Orthospaces

The discussion on betweenness-induced orthogonality relations in Section 6.2 lead to the observation that such an induced orthogonality relation is highly dependent on the chosen betweenness relation. Therefore, here we focus on Euclidean betweenness, because it is a widely used betweenness relation and helps to substantiate practically the abstract definition of COSs. Moreover, Euclidean betweenness leads to a natural form of convexity, facilitates calculation techniques such as convex optimization and is in general a widely researched topic. The aim of this chapter is, on the one hand, to exemplify the determination of an betweenness-induced orthogonality relation and, on the other hand, to examine COSs based on Euclidean betweenness in detail. Thus, in the following, the term "convexity" is used in its usual meaning as Euclidean convexity.

Euclidean betweenness was introduced in Section 4.3 and fulfills many betweenness axioms, in particular, (B0)–(B5) and (C2)–(C3) (see pages 47 and 50).

The aim of this section is not only to find one COS based on Euclidean betweenness but to examine the possible COSs based on Euclidean betweenness in general. The overall result of this examination is a characterization result: Using theorems out of the area of duality theory [Artstein-Avidan and Milman, 2008], I am able to characterize the class of COSs induced by Euclidean betweenness as that class of COSs that rely on different types of cones (as members of Y).

Sections 7.1 and 7.2 are mainly based on a joint journal publication [Leemhuis and Özçep, 2023], whereas Section 7.3 extends this result by relaxing the constraints of Sections 7.1 and 7.2.

7.1 Euclidean Convex Structures as Conceptual Orthospaces

Let $X = \mathbb{R}^n$ be the domain and let $\mathcal{B}_E$ denote Euclidean betweenness, i.e.,

$$\mathcal{B}_E(a, b, c) \text{ iff } d(a, b) + d(b, c) = d(a, c) \text{ and } a \neq b \neq c,$$

with $d(x,y) = \sqrt{\Sigma_{1 \le i \le n}(x_i - y_i)^2}$ being the standard Euclidean metric.

This betweenness can be intuitively described as follows: an element b is between a and c if it lies on a line segment through a and c. This can be reformulated to the standard definition of Euclidean convexity: In $\mathbb{R}^n$, a set $C \subseteq \mathbb{R}^n$ is convex iff for all $a, b \in C$, $\lambda a + (1-\lambda)b \in C$ for $0 \le \lambda \le 1$. Let $\mathcal{C}^n$ be the set of all convex sets in $\mathbb{R}^n$ (including the empty set).

$$\mathcal{C}^n := \{C \subseteq \mathbb{R}^n \mid C \text{ is convex}\}$$

The convex hull of some $C \subseteq \mathbb{R}^n$ (for short $conH(C)$) can be defined as the smallest convex set containing C, concretely, by constructing it as the intersection of convex sets containing C:

$$conH(C) = \bigcap\{A \mid A \supseteq C \text{ and } A \text{ is convex}\}.$$

As the domain itself is a convex set containing C, the intersection is not empty, unless C is empty. Here the term "convex" is used for "convex w.r.t. Euclidean betweenness".

Thus, it seems natural to set up a COS S over the domain of $\mathbb{R}^n$ based on the Euclidean betweenness, thus of the form $S = (\mathbb{R}^n, Y, \mathcal{B}_E, \perp)$ with $Y \subseteq \mathcal{C}^n$. But now two questions arise:

Question 6.

a) How should $\perp$ be defined so that S is a COS?

b) Can we choose $Y = \mathcal{C}^n$, i.e., can $\perp$ be defined such that Y is complete?

Before tackling the first question by defining a suitable orthogonality relation, we deal with the disjunction operator. Disjunction is a closure operator cl and as proven in Proposition 5.7, the disjunction needs to be interpreted as convex hull, when Y is complete. However, also if Y is not complete, it seems to be natural to interpret the disjunction of elements of Y as convex hull, both based on the interpretation as a conceptual space in the sense of Gärdenfors [2000] and based on the discussion in Section 6.2. Thus, in the following, the constraint (R4) (stating that for all $A, B \in Y$, $cl_\perp(A \cup B) = conv(A \cup B)$, see page 68 for details) is assumed to be valid.

This definition of conjunction and disjunction leads to the fact that $(\mathcal{C}^n, \wedge, \vee)$ makes up a bounded lattice, where $c \wedge d$ is interpreted as $C \cap D$ and $c \vee d$ is interpreted as $conH(C \cup D)$ for $C, D \in \mathcal{C}^n$, $\mathbb{1}$ is interpreted as $\mathbb{R}^n$ and $\mathbb{0}$ as $\emptyset$.

Unfortunately, this result does not say that we can immediately construct an orthocomplement on the lattice. In fact, the second question has to be answered negatively, as it is not possible to define a COS where each convex set is $\perp$-closed. The next proposition is, in fact, not dependent on (R4): even the weaker restriction (R5) (stating that for all $A, A^\perp \in Y$ it must be the case that $cl_\perp(A \cup A^\perp) = conv(A \cup A^\perp)$), prevents defining a COS, as shown in the following proposition.

Proposition 7.1. *$S = (\mathbb{R}^n, \mathcal{C}^n, \mathcal{B}_E, \perp)$ is not a COS for an arbitrary orthogonality relation $\perp$ fulfilling either (R4) or (R5).*

Proof. For $S = (\mathbb{R}^n, \mathcal{C}^n, \mathcal{B}_E, \perp)$ to be a COS, it must be the case that $A, A^\perp \in \mathcal{C}^n$ can be defined to satisfy the rules of the orthonegation, in particular, $cl_\perp(A \cup A^\perp) = \mathbb{R}^n$. With (R4) resp. (R5), it follows that $conH(A \cup A^\perp) = \mathbb{R}^n$.

Assume for the sake of contradiction that S is a COS. Let $A = \{x\}$ with $x \in \mathbb{R}^n$. A is convex and thus $A \in \mathcal{C}^n$ and thus it is by definition also $\perp$-closed. As $A^\perp \in Y$ needs to be the case and thus $conv(A^\perp) = A^\perp$, due to the rules of convexity, $A^\perp$ can be at most a half space (thus one of the parts in which an arbitrary hyperplane (an $n-1$ dimensional linear subspace) divides the space $\mathbb{R}^n$). Assume $A^\perp$ is a half space, then $conH(A \cup A^\perp) \neq \mathbb{R}^n$ and, thus, this is also not the case for $A^\perp$ being smaller. □

This goes perfectly in line with Theorem 5.23. $\mathcal{B}_E$ fulfills (B0)–(B4) and $cl_\perp$ fulfills the projective law, thus, Y can not be complete.

As a consequence of the negative answer, we change the second question to the following question: "Which subsets of $\mathcal{C}^n$ are potential candidates for Y in the COS?" Note that this question does not just ask to find some subset but to give a criterion on subsets identifying them as appropriate candidates for Y.

In the following, we assume that restriction (R7) applies, thus for $\vec{0} \in X$ it follows that $cl_\perp(\{\vec{0}\}) = \{\vec{0}\}$ and $\{\vec{0}\}^\perp = X = \mathbb{R}^n$ (where of course instead of $\{\vec{0}\}$ an arbitrary other singleton set can be chosen) and focus on topologically closed sets. The case where (R7) is not fulfilled and instead $X^\perp = \emptyset$ is considered in Section 7.3. Thus, the largest possible candidate for Y is the set $\mathcal{C}_0^n$ consisting of all (topologically) closed convex sets containing 0.

$$\mathcal{C}_0^n := \{C \mid C \in \mathcal{C}^n \text{ and } 0 \in C \text{ and } C \text{ is topologically closed }\}$$

To construct an orthogonality relation, a mapping $\phi : \mathcal{C}_0^n \to \mathcal{C}_0^n$ is needed which maps each $C \in \mathcal{C}_0^n$ into some $D \in \mathcal{C}_0^n$ such that ϕ fulfills the properties of the orthonegation. In particular, ϕ has to fulfill $\phi(\phi(C)) = C$ for all $C \in \mathcal{C}_0^n$ and $C \subseteq D \to \phi(C) \supseteq \phi(D)$ for all $C, D \in \mathcal{C}_0^n$ (double negation elimination and contraposition, if ϕ is interpreted as negation operator). For these restrictions, there are corresponding restrictions in projective geometry, especially in the concept of *duality*. Duality in a form adapted to this setting is defined as follows:

Definition 7.2 ([Artstein-Avidan and Milman, 2008, Definition 2])**.** *A map* $\phi : \mathcal{C}_0^n \to \mathcal{C}_0^n$ *generates a* duality *in* $\mathbb{R}^n$ *if the following two properties are satisfied:*

1. *For any* $C \in \mathcal{C}_0^n$ *it follows that* $\phi(\phi(C)) = C$.
2. *For any two* $C, D \in \mathcal{C}_0^n$ *satisfying* $C \subseteq D$ *it follows that* $\phi(C) \supseteq \phi(D)$

Thus, we see that the map we are looking for has to be a duality map. Based on this duality, some characterization results are given, leading to the definition of a unique mapping representing the orthonegation. For the class $\mathcal{C}_0^n$ there are results examining the properties of such dualities and one can prove that they are essentially unique.

Theorem 7.3. *[Artstein-Avidan and Milman, 2008, Theorem 10] Let* $n \geq 2$ *and* $\phi : \mathcal{C}_0^n \to \mathcal{C}_0^n$ *be a duality as defined in Definition 7.2. Then there is a symmetric* $g \in GL_n$ *such that for all* $C \in \mathcal{C}_0^n$

$$\phi(C) = (gC)^{\circ'},$$

where $(gC)^{\circ'}$ *is the polar body*

$$(gC)^{\circ'} = \{x \in \mathbb{R}^n | \langle x, gy \rangle \leq 1 \text{ for all } y \in C\}. \tag{7.1}$$

Here we use the usual notation GL_n *for the general linear group of degree* n*, i.e., the group over the set of invertible* $n \times n$*-matrices with entries over the reals.*

This states that a mapping having the properties of duality is nothing else than the polar-body mapping (up to multiplication with an invertible matrix g).

As the similarity between the definitions of duality and orthonegation is obvious, it might seem to be straightforward to use the polar body defined in Theorem 7.3 as a negation. For this interpretation, it would be necessary to prove that the third rule of orthonegation, intuitionistic absurdity, is valid. However, as Corollary 7.4 below shows, this is not the case and thus, this restricted set $\mathcal{C}_0^n$ is also not an appropriate candidate for a COS. But now, as stated in Theorem 7.3, this construction is the only one fulfilling the first two rules of orthonegation based on the set $\mathcal{C}_0^n$, further restrictions on the type of convex bodies considered are necessary.

Corollary 7.4. *For each map* ϕ *there is a* $C \in \mathcal{C}_0^n$ *such that* $C \cap \phi(C) \neq \{\vec{0}\}$.

Proof. Let $n \geq 2$ and choose a $g \in GL_n$ according to Theorem 7.3. For the sake of proof by contradiction, assume that $C \cap \phi(C) = \{\vec{0}\}$ for all $C \in \mathcal{C}_0^n$. Then, the following would hold:

(*) For all C there is no $x \in C$ with $x \in \phi(C) = (gC)^{\circ'}$.

We now construct C, x that falsify (*). Choose $x \in \mathbb{R}^n \backslash \{\vec{0}\}$ such that $\langle x, gx \rangle \leq 1$. Such an x always exists, as for arbitrary $z \in \mathbb{R}^n$, $\langle z, gz \rangle = \mu$ for some $\mu \in \mathbb{R}$. If $\mu \leq 1$, then let $x = z$ otherwise let $x = \frac{1}{\mu} z$, then $\langle x, gx \rangle = \frac{1}{\mu^2} \langle z, gz \rangle \leq 1$. Let $C = conH(\{0, x\}) = \{\lambda x \mid 0 \leq \lambda \leq 1\}$. As C is defined as convex hull we must have $x \in C$. But now, as for all $y \in C$, $y = \lambda x$ for $0 \leq \lambda \leq 1$, $\langle x, gy \rangle = \langle x, g(\lambda x) \rangle = \lambda \langle x, gx \rangle \leq 1$, it follows that $x \in \phi(C)$. But $x \in C$ holds by definition. So we reach a contradiction to (*). ☐

Note that $C \cap \phi(C) = \{\vec{0}\}$ for $C \in \mathcal{C}_0^n$ means $c \wedge c' = \mathbb{0}$, hence this corollary indeed shows that intuitionistic absurdity is falsified.

Thus, even the restriction to $\mathcal{C}_0^n$ is not enough to allow for the definition of an orthogonality relation. As this construction is based on necessary conditions for the establishment of duality, it is not possible in general to construct an orthogonality relation for general convex bodies including $\{\vec{0}\}$. An exception would be of course to define conjunction or disjunction in a different way which is not further considered here, because defining conjunction as set intersection and disjunction as convex hull seems to be natural and also psychologically grounded.

Now, as I will show in the following, it is possible to restrict the set of convex bodies further to allow for the definition of an orthogonality relation. There are several variants of the duality result of Theorem 7.3 in the literature, e.g., for convex structures including the point of origin in its interior [P. M. Gruber, 1991] or for ellipsoids [Artstein-Avidan and Slomka, 2011], but it turns out that the only restriction of $\mathcal{C}_0^n$ fulfilling intuitionistic absurdity is the set of closed convex cones. The following definition has already been given in Definition 5.5 and is repeated here for clarity.

Definition 7.5 ([Schneider, 2022, p. 7])**.** *A* convex cone *is a subset $C \subseteq \mathbb{R}^n$ with $x+y \in C$ and $\lambda x \in C$ for every $x, y \in C$ and $\lambda \geq 0$. A* closed convex cone *is a nonempty convex cone which is topologically closed (w.r.t. the Euclidean norm). The set of closed convex cones in $\mathbb{R}^n$ is named $\mathcal{C}^n_{cone}$.*

With this definition, the next proposition follows:

Proposition 7.6. *A set $\mathcal{D} \subseteq \mathcal{C}_0^n$ closed under the map ϕ of Theorem 7.3 and homogeneous in the sense that it is invariant under any rotation and reflection (i.e., application of an orthogonal matrix) fulfills $D \cap \phi(D) = \{\vec{0}\}$ for all $D \in \mathcal{D}$ only if $\mathcal{D} \subseteq \mathcal{C}^n_{cone}$.*

Proof. Let $n \geq 2$ and $g \in GL_n$ be arbitrary fulfilling the condition of Theorem 7.3 and let $D \cap \phi(D) = \{\vec{0}\}$ for all $D \in \mathcal{D}$. Let $D \in \mathcal{C}_0^n$ be arbitrary and $x \in D$ with $x \neq \vec{0}$. To fulfill $D \cap \phi(D) = \{\vec{0}\}$, there must be $y \in D$ with $\langle x, gy \rangle > 1$, as otherwise $x \in \phi(D)$.

As D is convex, and $x \in D$, also $\lambda x \in D$ for $0 \leq \lambda \leq 1$. Let $\langle x, gy \rangle = \gamma$ and $0 < \lambda < \frac{1}{\gamma}$. Then $\langle \lambda x, gy \rangle < 1$ and thus $\lambda x \in \phi(D)$ would be the case, a contradiction to $D \cap \phi(D) = \{\vec{0}\}$. Thus, it must be the case that $\frac{1}{\lambda} y \in D$ (or a different element z with $\langle x, gz \rangle > 1$). As $\lambda > 0$ can be arbitrary small, the line $\{\frac{1}{\lambda} z \mid \lambda \geq 0\}$, for some $z \in \mathbb{R}^n$, is contained in D.

Thus, $\mathcal{D}$ could contain closed convex cones (including linear subspaces and half spaces) and the convex hulls over convex cones and convex bodies. Though, when considering only a few $\perp$-closed sets in Y, both constructions are possible, the construction based on convex hulls over convex cones and convex bodies get problematic, when each element in $\mathbb{R}^n$ should be contained in a concept in Y (and not only in the $\overline{1}$-concept). Then, modeling concepts in Y depends on the other concepts already contained in Y, as the set of all convex hulls over convex cones and convex bodies is not closed under conjunction. Thus, the resulting map ϕ is especially not invariant under rotation and reflection, whereas the set of closed convex cones is closed under conjunction, negation and disjunction and for each convex body, it can be defined a unique convex cone (the conic hull) which is the minimal cone containing the convex body. □

Note that we defined a homogeneity principle based on rotation applications because this principle is compatible with the constraint $D \cap \phi(D) = \{\vec{0}\}$, whereas, e.g., invariance under translations would not.

Considering the set of convex cones $\mathcal{C}^n_{cone}$, it turns out that the polarity operator can be simplified to

$$x \perp_p y \text{ iff } \langle x, y \rangle \leq 0 \text{ for } x, y \in \mathbb{R}^n \tag{7.2}$$

and thus

$$C^\circ = \{y \in \mathbb{R}^n | \langle y, x \rangle \leq 0 \text{ for all } x \in C\} \tag{7.3}$$

for $C \in \mathcal{C}^n_{cone}$. We have to show $C^\circ = C^{\circ\prime}$. "$\subseteq$" is clear. For the other direction we show that if $y \notin C^\circ$, then $y \notin C^{\circ\prime}$. Let $y \notin C^\circ$, hence there is $x \in C$ with $\mu := \langle x, y \rangle > 0$. As C is a convex cone one can choose $\lambda > \frac{1}{\mu}$ and $\lambda x \in C$. But now $\langle y, \lambda x \rangle > 1$. So $y \notin C^{\circ\prime}$.

7.2 Concepts in Euclidean Conceptual Orthospaces are Convex Cones

When restricting the space of concepts to closed convex cones, obviously, the duality based on polarity can still be used. However, the question arises whether this representation is also unique for cones or whether it is possible to find a different mapping fulfilling the properties of the duality without relying on the polarity operator. The answer is twofold: On the one hand, in $\mathbb{R}^2$, the uniqueness is no longer given, but on the other hand, more importantly, the same theorem as for the convex bodies considered beforehand can also be stated for closed convex cones in $\mathbb{R}^n$ for $n \geq 3$.

Theorem 7.7. *[Schneider, 2008, Corollary 1] Let $n \geq 3$ and $\phi : \mathcal{C}^n_{cone} \to \mathcal{C}^n_{cone}$ be a duality as defined in Definition 7.2.*

Then there is a symmetric $g \in GL_n$ such that

$$\phi(C) = gC^\circ$$

for all $C \in \mathcal{C}^n_{cone}$ where C° is the polar body as defined in Equation (7.3) and GL_n is the general linear group of degree n. Note that here, in contrast to Theorem 7.3, $\phi(C) = gC^\circ = \{gx \in \mathbb{R}^n \mid \langle x, y\rangle \leq 0 \text{ for all } y \in C\}$[1].

This rule can be further restricted, as it turns out that $g \in GL_n$ can not be arbitrary and especially has to be a positive diagonal matrix.

Proposition 7.8. *$C \cap gC^\circ = \{\vec{0}\}$ for all $C \in \mathcal{C}^n_{cone}$ is valid iff g is a positive diagonal matrix.*

Proof.
"←": Let g be a positive diagonal matrix. Let $C \cap gC^\circ \neq \{\vec{0}\}$. Then there is $x \in C$ with $x = gy$ for some $y \in C^\circ$. As $y \in C^\circ$, especially $\langle x, y\rangle \leq 0$ and, as $x = gy$, $\langle gy, y\rangle \leq 0$. As g is positive and diagonal, this is only the case if $x = \vec{0}$. Thus, $C \cap gC^\circ = \{\vec{0}\}$, a contradiction to the assumption.

"→": Let g not be a diagonal matrix (or a diagonal matrix containing a non-positive element). Then there is $y \in \mathbb{R}^n \backslash \{\vec{0}\}$ such that $\langle y, gy\rangle \leq 0$. Let $C = \{\lambda gy \mid \lambda \geq 0\}$ and thus $y \in C^\circ$, as for $x \in C$ with $x = \lambda gy$: $\langle x, y\rangle = \langle \lambda gy, y\rangle = \lambda\langle gy, y\rangle \leq 0$. As $\phi(C) = gC^\circ$, it follows $gy \in \phi(C)$. As $gy \in C$ by definition, this leads to a contradiction. □

Thus, g can, in particular, be interpreted as identity and this leads to the following conceptual orthospace:

Proposition 7.9. *The structure $S = (\mathbb{R}^n, \mathcal{C}^n_{cone}, \mathcal{B}_E, \perp_p)$ with the set of closed convex cones $\mathcal{C}^n_{cone}$ in $\mathbb{R}^n$, the Euclidean betweenness $\mathcal{B}_E$ and the polarity $\perp_p$ (thus $x \perp_p y$ iff $\langle x, y\rangle \leq 0$ for $x, y \in \mathbb{R}^n$) is a COS.*

[1] This is done only for simplicity here, Schneider [2008] shows that $(gC)^\circ = g^{-1}C^\circ$ for $C \in \mathcal{C}^n_{cone}$, thus both notions are interchangeably usable.

Proof. $\mathcal{B}_E$ obviously fulfills the betweenness axioms and $\perp_p$ is an orthogonality relation, as proven in Theorem 7.7 and Proposition 7.8, each closed convex cone is $\perp$-closed based on polarity and the set $\mathcal{C}^n_{cone}$ is closed under conjunction, intersection and negation and makes up an ortholattice [Schneider, 2008]. □

As shown by the derivation of this COS in this section, this is not only one of many possible COS but it is also the only possible one in Euclidean space based on some natural assumptions.

Theorem 7.10. *For a domain $\mathbb{R}^n$ with $n \geq 3$ with a betweenness relation defined as Euclidean betweenness $\mathcal{B}_E$, $\perp$-closure of elements of Y defined as convex hull and $X^\perp = \{\vec{0}\}$ (thus (R4) and (R7) are fulfilled), a COS has to be of the form of Proposition 7.9 where $Y = \mathcal{C}^n_{cone}$ or a sublattice of it and $\perp_p$ is possibly scaled by a positive diagonal matrix.*

Notice that this leads to the COS as defined in Example 5.6. Of course, when considering a domain $X = \mathbb{R}^n$ with $n \leq 2$, also other COSs could be possible. However, such a construction in two dimensions or less is not expressive enough for most cases (consider, e.g., the application of Helly's theorem mentioned by Gutiérrez-Basulto and Schockaert [2018]).

7.3 Generalization to Pseudo-Cones

In the last section, it turned out that for a COS based on Euclidean betweenness and some further constraints closed convex cones need to be considered as underlying structure. This has been established based on the assumptions that $\perp$-closure of elements of Y is defined as convex hull and that $X^\perp = \{\vec{0}\}$. Though, both restrictions are well justifiable, the question comes to mind whether it is possible to find such a necessary structure even when these restrictions are omitted resp. relaxed. Without any restriction on the resulting COS, it is not possible to state much about the structure of the resulting concepts. This was discussed in Section 5.2 and, in particular, regarding Euclidean betweenness, in Example 5.11.

Therefore, even when relaxing the constraints used in the last section it is necessary to define some constraints. In particular, instead of (R7), (R8) is used, therefore $X^\perp = \emptyset$. As argued in the last section, having (R4) would be beneficial, but due to reasons that will turn out later on in this section, this is not always possible. For the following propositions, (R5) is sufficient and, later on, the discussion on (R4) will be picked up again.

These constraints lead to the following question:

Question 7. *How should $\perp$ and Y be defined for a COS $S = (\mathbb{R}^n, Y, \mathcal{B}_E, \perp)$ such that*

- *For $A, A^\perp \in Y$, $cl_\perp(A \cup A^\perp) = conv(A \cup A^\perp)$* *(R5)*
- *$X^\perp = \emptyset$* *(R8)*

is valid?

Here and in the following, the focus lies again on topologically closed sets. First, remember that Proposition 7.1 showed that Y cannot be an arbitrary set of convex structures and especially not the set of all convex sets $\mathcal{C}^n$.

For the determination of a suitable Y, in the following *pseudo-cones* known from the literature and a relaxation of it defined by me are considered. As a first step, it is shown that Y must be at least the set of *relaxed pseudo-cones* and later on, this result can be restricted to the stronger class of pseudo-cones.

Definition 7.11.

- *A non-empty closed convex set C not containing the origin is called* pseudo-cone *if $\lambda x \in C$ for all $x \in C$ and $\lambda \geq 1$ [Xu et al., 2023].*
- *A non-empty closed convex set C is called* relaxed pseudo-cone *if for all $x \in C$, there is $y \in \mathbb{R}^n$ with $y \neq x$ or $y = \vec{0}$ such that for all $\lambda \geq 0 : x + \lambda(x - y) \in C$.*

Let $\mathcal{PC}^n$ be the set of all closed pseudo-cones in $\mathbb{R}^n$ and $\mathcal{PPC}^n$ the set of all closed relaxed pseudo-cones in $\mathbb{R}^n$. Relaxed pseudo-cones are relaxed in the sense that each convex cone and each pseudo-cone are also relaxed pseudo-cones.

Lemma 7.12. *Consider the set set of all pseudo-cones in $\mathbb{R}^n$ ($\mathcal{PC}^n$), the set of all relaxed pseudo-cones in $\mathbb{R}^n$ ($\mathcal{PPC}^n$) and the set of all convex cones in $\mathbb{R}^n$ ($\mathcal{C}^n_{cone}$):*

1. *$\mathcal{C}^n_{cone} \subseteq \mathcal{PPC}^n$*
2. *$\mathcal{PC}^n \subseteq \mathcal{PPC}^n$*

Proof. Consider the definition of relaxed pseudo-cones in Definition 7.11 and choose $y = \vec{0}$. Then, the definition reduces to $\forall x \in C : \forall \lambda \geq 0 : x + \lambda x \in C$. Based on this definition, both $\mathcal{C}^n_{cone}$ and $\mathcal{PC}^n$ are special cases. □

An example for a pseudo-cone can be seen in Figure 7.1(a), an example for a relaxed pseudo-cone in Figure 7.1(b).

The next proposition shows that an orthogonality relation can only be constructed if the class of relaxed pseudo-cones or a subset of it is considered.

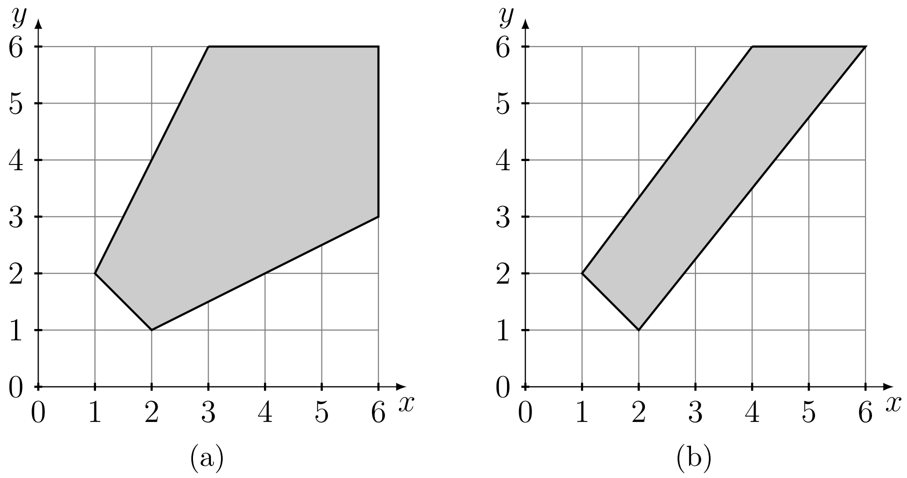

Figure 7.1: (a) Example for a pseudo-cone in $\mathbb{R}^2$; (b) Example for a relaxed pseudo-cone in $\mathbb{R}^2$

Proposition 7.13. *If for a set $\mathcal{D} \subseteq \mathcal{C}_c^n$ (where $\mathcal{C}_c^n$ is the set of all closed convex sets in $\mathbb{R}^n$) an arbitrary orthogonality relation $\perp$ (fulfilling (R5)) can be constructed, then $\mathcal{D} \subseteq \mathcal{PPC}^n$.*

Proof. Due to (R5) it must be the case that for $A \in \mathcal{D}$, $cl_\perp(A \cup A^\perp) = conv(A \cup A^\perp) = \mathbb{R}^n$. Due to the definition of Euclidean betweenness and thus of convexity, $A^\perp$ can be at most a half space. Therefore, A must contain at least a half line to satisfy $conv(A \cup A^\perp) = \mathbb{R}^n$. This is also the case for a smaller $A^\perp$.

As each element in $\mathcal{D}$ is convex, it must be the case that $conv(A) = A$ and therefore, A must be a relaxed pseudo-cone. A special case is $A = \{\vec{0}\}$. However, $\mathbb{R}^n \in \mathcal{PPC}^n$ and $conv(\{\vec{0}\} \cup \mathbb{R}^n) = \mathbb{R}^n$. □

Despite this positive result, it is still not sufficient to restrict Y to $\mathcal{PPC}^n$. The problem is especially that $\mathcal{PPC}^n$ is not closed under conjunction (consider, e.g., for some $a \in \mathbb{R}^n \backslash \{\vec{0}\}$, $A = \{\lambda a \mid \lambda \geq 1\}$ and $B = \{\lambda a \mid \lambda \leq 1\}$, $A, B \in \mathcal{PPC}^n$, but $A \cap B = \{a\} \notin \mathcal{PPC}^n$). Therefore, a restricted set $\mathcal{D} \subseteq \mathcal{PPC}^n$ is searched for such that it is closed under conjunction and disjunction.

There are two options: First, consider the relaxed pseudo-cone $\{\vec{0}\} \in \mathcal{PPC}^n$. If it is contained, then, as shown by the next proposition, it has to be the case that $X^\perp = \{\vec{0}\}$. This leads to the fact that, based on a minor restriction, again closed convex cones need to be considered.

Proposition 7.14. *Let a COS $S = (\mathbb{R}^n, \mathcal{D}, \mathcal{B}_E, \perp)$ for $\perp$ being an arbitrary orthogonality relation (based on (R4)) and $\mathcal{D} \subseteq \mathcal{PPC}^n$.*

If $\{\vec{0}\} \in \mathcal{D}$, then $\mathcal{D} \subseteq \mathcal{C}^n_{cone}$ and $\perp$ is the polarity $\perp_p$ as defined in Equation (7.3) *possibly scaled by a positive diagonal matrix.*

Proof. There are two possibilities interpreting $\{\vec{0}\} \in \mathcal{D}$: either $\{\vec{0}\}^\perp = \mathbb{R}^n$ or $\{\vec{0}\}^\perp \neq \mathbb{R}^n$. First, let $\{\vec{0}\}^\perp = \mathbb{R}^n$. Thus, for each $A \in \mathcal{D}$, $\{\vec{0}\} \in A$. Due to Lemma 7.12, each closed convex cone is a relaxed pseudo-cone and by definition $\{\vec{0}\}^\perp = \mathbb{R}^n$, thus, as (R4) is valid, with Theorem 7.10 the statement follows.

The remaining case is that $\{\vec{0}\}^\perp \neq \mathbb{R}^n$. But as argued in Proposition 7.13, there is no $A \in \mathcal{PPC}^n$ with $A \subsetneq \mathbb{R}^n$ such that $conv(A \cup \{\vec{0}\}) = \mathbb{R}^n$. This contradicts the definition of a COS, as then $\mathcal{D}$ is not closed under negation. □

This proof can also be used as an alternative proof for the need of closed convex cones in Section 7.1, independent of duality and thus underlines that result further.

The remaining option is that $\mathcal{D} \subseteq \mathcal{PPC}^n$ is considered with $\{\vec{0}\} \notin \mathcal{D}$ and either $\vec{0} \notin A$ for all $A \in \mathcal{D}$ or $\vec{0} \in A$ only for some $A \in \mathcal{D}$. As in the latter case again the problem occurs that $\mathcal{D}$ is not closed under intersection, here the former case is considered.

Here, the same line of argument can be followed as for Proposition 7.6 regarding closed convex cones. Though, in contrast to the case of Proposition 7.6, here, the representation of the disjunction is not yet determined, it is sufficient to consider the case of closure based on intersection of sets.

Proposition 7.15. *A set $\mathcal{D} \subseteq \mathcal{PPC}^n$ where $\{\vec{0}\} \notin A$ for $A \in \mathcal{D}$ is closed under intersection and invariant under any rotation and reflection (i.e. application of an orthogonal matrix) only if $\mathcal{D} \subseteq \mathcal{PC}^{n\prime}$ (where $\mathcal{PC}^{n\prime}$ consists of the elements of $\mathcal{PC}^n$ translated by an arbitrary but fixed $t \in \mathbb{R}^n$ (thus $\mathcal{PC}^{n\prime} = \mathcal{PC}^n$ for $t = \vec{0}$))*

Proof. Follows analogously to the second half of the proof of Proposition 7.6. □

The next question is how the $\perp$-closure should be defined. As argued, the best case would be to satisfy (R4), thus achieving equivalence of convex hull and $\perp$-closure regarding elements of Y. However, the problem is that pseudo-cones are not closed under conic hull. A problem occurs when $\vec{0} \in conH(A \cup B)$ for $A, B \in \mathcal{PC}^n$, as illustrated in Figure 7.2. Therefore, the constraint is relaxed as little as possible, resulting in the following:

Proposition 7.16. *[Xu et al., 2023, Lemma 3.11] The class of closed pseudo-cones $\mathcal{PC}^n$ including $\{\vec{0}\}$ and $\mathbb{R}^n$ makes up a lattice where conjunction is again defined as*

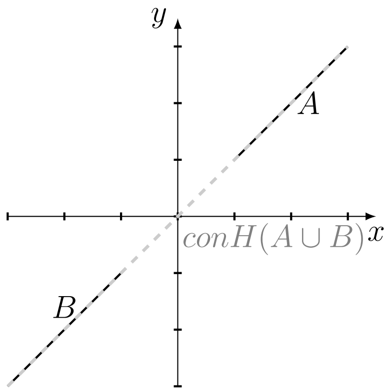

Figure 7.2: Example for a convex hull $conH(A \cup B)$ of two pseudo-cones A, B that is not a pseudo-cone.

set-intersection, disjunction as

$$(A \sqcup B)^{\mathcal{I}} = \begin{cases} conH(A, B) & \text{if } \vec{0} \notin conH(A \cup B) \\ \mathbb{R}^n & \text{if } \vec{0} \in conH(A \cup B) \end{cases} \tag{7.4}$$

for $A, B \in \mathcal{PC}^n$, where $\mathbb{0}$ is represented as $\emptyset$ and $\mathbb{1}$ is represented as $\mathbb{R}^n$.

Thus, the relaxed version of (R4) for the class of pseudo-cones can be stated as follows:

(R4)$_{\textbf{pseudo-cone}}$ For all $A, B \in Y$: If $\vec{0} \notin conv(A \cup B)$, then $cl_\perp(A \cup B) = conv(A \cup B)$.

Thus, the only possible COS to consider based on the Euclidean betweenness and the constraints mentioned above is $S = (\mathbb{R}^n, \mathcal{PC}^n, \mathcal{B}_E, \perp)$. The question remains how $\perp$ should be defined and whether there is a unique definition for $\perp$ given. It turns out that the results for closed convex cones are adaptable to the case of pseudo-cones and that $\perp$ can be defined based on the following polarity operator: The polarity C^* [Xu et al., 2023] is defined as

$$C^* = \{y \in \mathbb{R}^n | \langle x, y \rangle \leq -1 \text{ for all } x \in C\} \tag{7.5}$$

Thus, the polarity operator $\perp_{ps}$ is defined as

$$x \perp_{ps} y \text{ iff } \langle x, y \rangle \leq -1 \text{ for } x, y \in \mathbb{R}^n \tag{7.6}$$

Each polar of a pseudo-cone is again a pseudo-cone [Xu et al., 2023].

Based on the same line of argument as done for the case of convex cones, the following is stated:

Theorem 7.17. *[Xu et al., 2023, Corollary 1.2] Let $n \geq 2$ and $\phi : \mathcal{PC}^n \to \mathcal{PC}^n$ be a duality as defined in Definition 7.2. Then there is a symmetric $g \in GL_n$ such that*

$$\phi(C) = gC^*$$

for all $C \in \mathcal{PC}^n$ where C^* is the polar body as defined in Equation (7.5) and GL_n is the general linear group of degree n.

It is possible to prove that intuitionistic absurdity is valid for g being a positive diagonal matrix. The proof proceeds analogously to the one for closed convex cones in Proposition 7.8.

Proposition 7.18. *The duality as defined in Equation (7.5) fulfills*

$$C \cap \phi(C) = C \cap gC^* = \emptyset$$

if and only if g is a positive diagonal matrix.

Proof.
"$\leftarrow$": Let g be a positive diagonal matrix. Assume $x \in C \cap gC^* \neq \emptyset$ for some $C \in \mathcal{PC}^n$. Thus, by definition, $x = gy$ for some $y \in C^*$ and $\langle x, y \rangle \leq -1$ and thus $\langle gy, y \rangle \leq -1$. As g is positive and diagonal, this is not possible.

"$\rightarrow$": Let g not be a diagonal matrix (or a diagonal matrix containing a non-positive element). Then there is $y \in \mathbb{R}^n \backslash \{\vec{0}\}$ such that $\langle y, gy \rangle \leq -1$. Let $C = \{\lambda gy \mid \lambda \geq 0\}$ and thus $y \in C^*$, as for $x \in C$ with $x = \lambda gy$: $\langle x, y \rangle = \langle \lambda gy, y \rangle = \lambda \langle gy, y \rangle \leq -1$. As $\phi(C) = gC^*$, it follows $gy \in \phi(C)$. As $gy \in C$ by definition, this leads to a contradiction. □

This leads to the following theorem, combining the results of this and the last section:

Theorem 7.19. *For a domain $\mathbb{R}^n$ with $n \geq 3$ and a betweenness relation defined as Euclidean betweenness $\mathcal{B}_E$, a COS $S = (\mathbb{R}^n, Y, \mathcal{B}_E, \perp)$ fulfills (R4) (resp. (R4)$_{pseudo\text{-}cone}$) iff*

- *$Y \subseteq \mathcal{C}^n_{cone}$ and $\perp = \perp_p$ is the polarity operator as defined in Equation (7.2) — possibly scaled by a positive diagonal matrix or*
- *$Y \subseteq \mathcal{PC}^n$ and $\perp = \perp_{ps}$ is the polarity operator as defined in Equation (7.6) — possibly scaled by a positive diagonal matrix.*

These results show the importance of (pseudo)-cones for COSs based on Euclidean betweenness and underline the intuition that convex cones are suitable structures for representing negation, as, e.g., stated by Özçep et al. [2020] regarding al-cones and by Garg et al. [2019] regarding axis-parallel linear subspaces.

8 Conceptual Orthospaces with Roles (Binary Relations)

In the last chapters, the framework of COSs has been presented in a propositional context, i.e., without including roles. To strengthen its applicability in the context of classical KGE-approaches, in the following, the extension of COS with roles is considered, in particular in the context of cone-COS.

Roles can be represented with the help of many different geometric operators, ranging from translation — as, e.g., used in TransE [Bordes et al., 2013] and illustrated in Figure 8.1(a) — to rotation as in RotatE [Sun et al., 2019]. In conceptual spaces, roles are considered as being represented in a higher dimensional space than the instances they rely on. For example, "longer than" for entities in an one-dimensional space can be represented as a two-dimensional space containing a region representing "longer than" where the value of one axis is greater than on the other axis [Gärdenfors, 2000, p. 92]. This is illustrated in Figure 8.1(b). This idea is also widely used in practice, e.g., for the embedding of closed subspaces of a Hilbert space [Garg et al., 2019], cones [Özçep et al., 2020] and convex sets in general [Gutiérrez-Basulto and Schockaert, 2018].

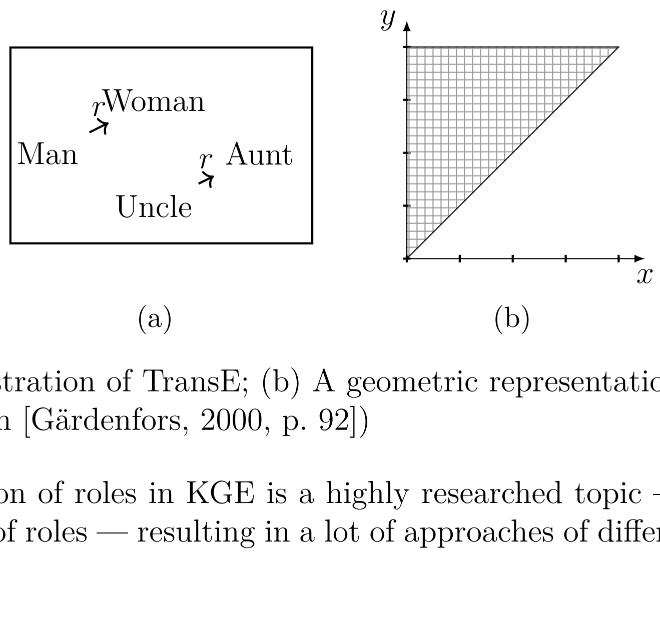

Figure 8.1: (a) Illustration of TransE; (b) A geometric representation of "y longer than x" (from [Gärdenfors, 2000, p. 92])

The representation of roles in KGE is a highly researched topic — as KGE is about the representation of roles — resulting in a lot of approaches of different complexity and expressivity [Q. Wang et al., 2017].

TransE [Bordes et al., 2013] continues to be the classical reference for a knowledge graph embedding of roles, drawing its charm from a simple geometrical representation of roles that can be learned efficiently. Indeed, TransE represents roles as vector translations, and hence embedding a triple (s, R, o) into a continuous space is easily integrated with a loss function used for learning the embedding. The downside of this simple representation is TransE's limited expressivity [Kazemi and Poole, 2018]: only binary relations that are functional in their first argument can be modeled. Such considerations on expressivity have lead to many other embedding approaches that rely on more complex representations of roles. Most of them are in the tradition of TransE — some more such as TransH [Z. Wang et al., 2014], TransR [Lin et al., 2015] that rely on representing roles as translation, and some less, relying on other, more involved geometrical operations such as SimplE [Kazemi and Poole, 2018] or Rescal [Nickel et al., 2011].

However, it is not only necessary to model roles for contingent facts of the form $R(a, b)$ (in DL speak: on the abox-level) but also for terminological assertions of the form $\exists R.C \sqsubseteq D$, say (in DL speak: on the tbox-level).

Many KGE-approaches with incorporation of background knowledge suffer from limitations in expressivity, both regarding expressivity on abox- and on tbox-level. This is discussed in detail in Section 10.2. In particular, important features such as partiality and non-functionality of roles can mostly not be modeled correctly. By Gutiérrez-Basulto and Schockaert [2018] subjects and objects of triples are represented as vectors and arbitrary n-ary relations as subsets of the n-Cartesian product over the embedding space. An obvious downside of that approach is the high increase in dimensionality required with adverse effects on learning. Here, I take a middle-road: I insist on representing concepts (unary relations) as sets of vectors but allow representations of arbitrary binary relations in mathematically well-behaved operations.

Thus, the aim is to enhance the COS with an interpretation of roles enabling to model both abox- and tbox-axioms, ideally of infinite quantifier rank or at least of an arbitrary but fixed quantifier rank and being able to model full negation also for axioms including roles. Finding an adequate representation of roles highly depends on the specific COS. Therefore, here, mainly the cone-COS as introduced in Example 5.6 is considered. However, the general idea of reification introduced in the following is a notion independent of cone-COS.

To tackle the problem of finding a cone-COS with roles, first, the concept of *reification* is introduced as a general approach of handling roles. I adopt the idea of relying on matrix multiplication to represent roles previously used in KGEs, but I rewrite roles into equivalent structures which allows to model arbitrary roles, including partial and non-functional ones. The idea of reification is to represent roles as objects in the embedding space. In my case, e.g., a triple (a, R, b) would be represented by an object $c_{R(a,b)}$ and two relations stating that its "arguments" are a and b: $\pi_{1,R}(c_{R(a,b)}) = a$, $\pi_{2,R}(c_{R(a,b)}) = b$. The only relations $\pi_{i,R}$ that have to be represented and learnt, then, are functional relations, namely projections of triples to its subject and its object. This construction also enables to model partiality, as not each element needs to have a connection to a reified relation. While reification is a well-known approach in logic modeling, the technical challenge tackled here is to apply this approach to KGE and especially to the DL $\mathcal{ALC}$.

In the following, reification is applied to the cone-COS in Section 8.1. In this context, a restriction is presented considering the distributive case, enabling to show that each $\mathcal{ALC}$-ontology can be represented by a cone-COS with roles. Section 8.1 is based on my joint work work with Özgür Özçep and Diedrich Wolter [Leemhuis, Özçep, and Wolter, 2022a].

The chapter ends in Section 8.2 with a consideration on faithful embeddings based on cone-COS, not only for concepts, but also for roles, in particular regarding the representation of arbitrary quantifier ranks, based on my previously published work [Özçep et al., 2023] co-authored with Özgür Özçep and Diedrich Wolter.

8.1 Cone Embedding with Roles

The aim of this section is to present an extension of the framework of COS based on convex cones to be able to model roles. These roles should not only be modeled on instance level (thus modeling $R(a, b)$) but also on concept level (thus modeling, e.g., $A \sqsubseteq \exists R.B$).

We define H_m as the m-dimensional hyperoctant cone $H_m \subset \mathbb{R}^m$ generated by m vectors $\{(1\,0\,\cdots 0)^T, (0\,1\,0\,\cdots 0)^T, (0\,0\,1\,0\,\cdots 0)^T, \ldots, (0 \cdots 0\,1)^T\}$.

We now define reification as illustrated in Figure 8.2(a) to relate two concepts C, D. Roles are represented like concepts, i.e., by convex sets of specific geometric shape, and projections π_1, π_2 are introduced that link the embedded roles with the corresponding concepts. The main advantage of this over previous attempts is that the use of projections allows non-functional and partial relations to be represented. Approaches representing roles by geometric transformations in the embedding space such as TransE [Bordes et al., 2013] are attractive as they do not require further dimensions to be introduced, yet at the severe cost of being only able to represent functional relations, i.e., any x is always related to exactly one y. Several work on remedying this severe limitation has been done, but no general cure is possible when relying on a single geometric transformation function.

Definition 8.1. *We say a* reification *of a binary relation R between two cones C, D*

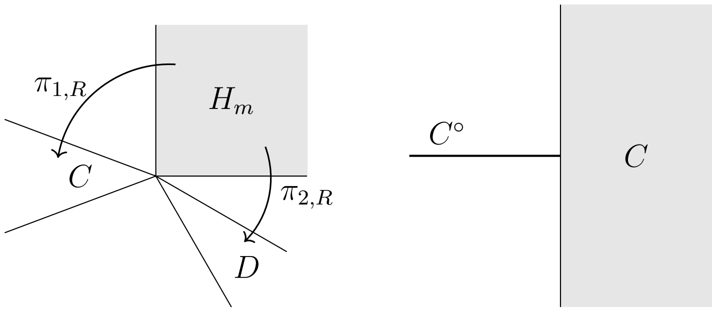

Figure 8.2: (a) Reification of role R is based on linear functions $\pi_{1,R}$, $\pi_{2,R}$ that project role concept H_m to its arguments. (b) Illustration for reification requiring H_m with $m > n$.

is given by a hyperoctant H_m and two linear functions $\pi_{1,R}$ and $\pi_{2,R}$ given as matrices $M \in \mathbb{R}^{n \times m}$.

Let us now extend the cone-COS with roles. A cone interpretation for $\mathcal{ALC}$ maps symbols and formulae to cones in $\mathbb{R}^n$. The definition is as usual, except that we exclude the origin $\vec{0}$ from the domain. This creates a convenient way for projections in the reification of roles to map subspaces to a well-defined "nirvana" $\vec{0}$ whenever a mapping to the empty set is required. For example, the formula $\forall R.\overline{1} \sqsubseteq \overline{0}$ saying that nothing is reachable by role R can elegantly be represented by setting $\pi_{2,R}$ to $\mathbf{0} \in \mathbb{R}^{n \times n}$, the projection to $\vec{0}$.

Definition 8.2. *A* cone-COS with binary relations *for a given $\mathcal{ALC}$ vocabulary $\mathcal{V} = N_c \cup N_C \cup N_R$ of constants, concept and role symbols is a structure $S = (\mathbb{R}^n \backslash \{\vec{0}\}, \mathcal{C}^n_{cone}, \mathcal{B}_E, \perp_p)$, based on the set of closed convex cones $\mathcal{C}^n_{cone}$, Euclidean betweenness $\mathcal{B}_E$ and the orthogonality relation $\perp_p$ for some $n \in \mathbb{N}$ and a* denotation function $\cdot^{\mathcal{I}}$ *defined for all $b \in N_c, A \in N_C, R \in N_R$ and concepts C, D over $\mathcal{V}$ such that the following conditions are (additionally to the ones in Definition 5.2) fulfilled:*

$R^{\mathcal{I}} \subseteq X \times X$, $\quad (\exists R.C)^{\mathcal{I}} = \pi_{1,R}\left(\pi_{2,R}^{-1}(C^{\mathcal{I}}) \cap H_m\right)$

$\overline{1}^{\mathcal{I}} = X$, $\quad$ *where H_m, $\pi_{1,R}$, $\pi_{2,R}$ are a reification of role R*

$\overline{0}^{\mathcal{I}} = \emptyset$, $\quad$ *and $\pi_{1,R}\left(\pi_{2,R}^{-1}(C^{\mathcal{I}}) \cap H_m\right)$ is a convex cone.*

$(\forall R.C)^{\mathcal{I}} = (\neg\exists R.\neg C)^{\mathcal{I}}$

In the following, the term "cone-COS" depicts cone-COS with binary relations. We now show that arbitrary $\mathcal{ALC}$ ontologies $\mathcal{O} = (\mathcal{A}, \mathcal{T})$ consisting out of abox $\mathcal{A}$ and tbox

$\mathcal{T}$ are representable by a cone-COS. First, we define how roles on the abox level are modeled.

Definition 8.3. *Given an $\mathcal{ALC}$ vocabulary with concept symbols N_C, constant symbols N_c, and role symbols N_R, and an $\mathcal{ALC}$ ontology $\mathcal{O}$, we say that the roles in $\mathcal{O}$ are representable if there is a cone-COS that is a model of $\mathcal{O}$ and $\mathcal{O} \models R(a,b)$ if and only if $b \in \pi_{1,R}(\pi_{2,R}^{-1}(a) \cap H_m)$ for some hyperoctant H_m.*

Proposition 8.4. *A cone-COS maps all concept descriptions to closed convex cones.*

Proof. This is clear for atomic concepts, intersection, disjunction and for the polar operator as shown in Example 5.6. Linear mappings also preserve cones, as they distribute over arbitrary linear combinations (not only those with positive scalars). For the existential, being a convex cone is enforced directly by the definition. □

Note that enforcing closed convex cones for the embedding of existentials is not a strong constraint. Taking the null vector into account one can show that the inverse preserves cones: Let C be a cone and M be a linear mapping as used in reification. Let $v \in M^{-1}[C]$, then $M(v) = w \in C$. Then for $\lambda > 0$ due to linearity $M(\lambda v) = \lambda M(v) = \lambda w$ and $\lambda w \in C$ due to the fact that C is a cone. Let $v_1, v_2 \in M^{-1}[C]$, then $M(v_1) = w \in C, M(v_2) = w_2 \in C$ (for some w_1, w_2). Now $M(v_1+v_2) = M(v_1)+M(v_2) = w_1+w_2 \in C$, so $v_1 + v_2 \in M^{-1}[C]$.

Reification employs matrix multiplications like several previous approaches, but it employs an 'in-between stop' at H_m which is the central trick to represent 1-to-k relations by making π_1 a k-to-1 mapping. Let us consider a simple example shown in Figure 8.2 to see that a stop H_m is necessary and that even $m > n$ may be necessary.

Example 8.5. *We consider the cone C generated by vectors $\{(0\,1)^T, (0\,-1)^T, (1\,0)^T\}$ in $\mathbb{R}^2$ shown in Figure 8.2 (b). Its negation given by the polarity operator C° is the cone generated by $\{(-1\,0)^T\}$. Now consider background knowledge $\exists R.C = \overline{1}$ saying that any entity is reachable from C by means of role R. $\overline{1}$ is interpreted as $\mathbb{R}^2 \setminus \{\vec{0}\}$ and it requires four independent rays $\lambda_i c_i$, $\lambda_i > 0$, $c_i \in C$ to span $\mathbb{R}^2$, more than offered by C. It requires at least H_4 to achieve the desired mapping:*

$$\pi_{1,R} := \begin{pmatrix} 0 & 0 & 1 & -1 \\ 1 & -1 & 0 & 0 \end{pmatrix}, \pi_{2,R} := \begin{pmatrix} 0 & 0 & 1 & 0 \\ 1 & -1 & 0 & 0 \end{pmatrix}$$

For any point $\vec{x} = (x_1 \cdots x_4)^T$ in H_4 we have $\pi_2(\vec{x}) = x_1 \begin{pmatrix} 0 \\ 1 \end{pmatrix} + x_2 \begin{pmatrix} 0 \\ -1 \end{pmatrix} + x_3 \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and since $x_1, \ldots, x_4 \geq 0$ we have $\pi_2(\vec{x}) \in C$ and $\pi_1(\vec{x})$ covers $\mathbb{R}^2$ for $\vec{x} \in H_4$. In general, H_m with $m > n$ is required when concept C is a sub-space of lower dimensionality than the concept it is related to.

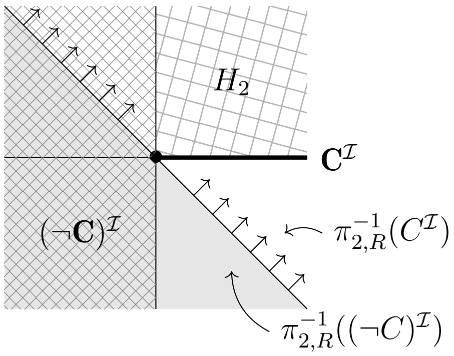

$$\pi_{1,R} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

$$\pi_{2,R} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}$$

Figure 8.3: Example for the construction of a geometric model of a tbox consisting of $\exists R.C = C$ and $\exists R.\neg C = \overline{0}$

Let us discuss a more involved example showcasing the ability to model complex relationships.

Example 8.6. *Consider concept $C^{\mathcal{I}}$ represented as the positive x-axis, the complement $(\neg C)^{\mathcal{I}}$ is the negative half space in $\mathbb{R}^2$ shown in Figure 8.3. This model fulfills the two tbox-axioms $\exists R.C = C$ and $\exists R.\neg C = \overline{0}$. The first axiom is fulfilled, as $\pi^{-1}_{2,R}(C^{\mathcal{I}})$ is the region marked with arrows in the figure. The intersection with the region of possible relational facts H_2 results in H_2, the upper right quadrant. This is mapped to $C^{\mathcal{I}}$ by $\pi_{1,R}$. As $\pi^{-1}_{2,R}(\neg C)^{\mathcal{I}}$ does not intersect with H_2, $\exists R.\neg C = \overline{0}$ is valid.*

Therefore, partiality is obviously given. To show non-functionality, the instances need to be considered. Assume $c^{\mathcal{I}} = (\lambda\, 0)^T$ for an arbitrary $\lambda > 0$. $\pi^{-1}_{1,R}(c^{\mathcal{I}}) \cap H_2 = \{(\lambda\, \mu)^T \mid \mu \geq 0\}$ and $\pi_{2,R}\pi^{-1}_{1,R}(c^{\mathcal{I}}) = \{(\lambda + \mu\, 0)^T \mid \mu \geq 0\}$. Thus, a $c^{\mathcal{I}} = (1\, 0)^T$ has a relation to all $b^{\mathcal{I}} = (\gamma\, 0)^T$ for $\gamma \geq 1$.

A property of this approach, besides its non-distributivity of $\sqcap$ over $\sqcup$, is its non-distributivity of $\exists$ over $\sqcup$, meaning $(\exists R.(C \sqcup D))^{\mathcal{I}} \neq (\exists R.C \sqcup \exists R.D)^{\mathcal{I}}$. Despite this is not classical, e.g., different from $\mathcal{ALC}$-semantics, it may be quite useful in modeling: Assume a binary relation E is introduced to model whether a person is examined to have a specific disease. Thus, by asserting $E(\mathit{person}, \mathit{disease})$ medical knowledge about a person at a specific point in time is reported. Now, it might be the case that an examination is not exact and thus results in the knowledge that the person could have disease A or disease B. However, assuming $\exists E.(A \sqcup B) = \exists E.A \sqcup \exists E.B$ would result in the conclusion that at this stage of examination it is already known which exact disease the person has. However, this exact specification was presumed not to be possible, and therefore, the instance representing the person should be placed in the embedding of $\exists E.(A \sqcup B)$ but neither in the embedding of $\exists E.A$ nor in the embedding of $\exists E.B$ because for both of them there is no justification in the examination. Thus, a lack of distributivity can be helpful to bridge gaps in semantics.

Distributive Embedding

The approach described so far leads to a possibility of expressing relational knowledge in general orthologics which may be relevant for some applications as we have argued above. However, many knowledge bases consider stronger orthologics, i.e., expect distributivity to hold, which include all classical logics.

One prominent example is $\mathcal{ALC}$ with classical semantics. Here one requires the ortholattice also to be a Boolean algebra. In fact, classical $\mathcal{ALC}$-tboxes are characterized according to Schild [1991] by the fact that the existential is a strong operator, i.e., the existential quantifier with classical $\mathcal{ALC}$ semantics satisfies the following two properties: $(\exists R.\overline{0})^{\mathcal{I}} = (\overline{0})^{\mathcal{I}}$ and $(\exists R.(C \sqcup D))^{\mathcal{I}} = (\exists R.C \sqcup \exists R.D)^{\mathcal{I}}$.

Therefore, to adapt the approach to $\mathcal{ALC}$, distributivity of $\sqcap$ over $\sqcup$ and distributivity of $\exists$ over $\sqcup$ must be achieved. The first property can be met by restricting cones to al-cones as introduced by Özçep et al. [2020] and explained in Example 5.17 since geometric models based on al-cones fulfill distributivity. Al-cones have a finite basis and their generating vectors only consists out of components $+1$, -1, and 0. The second property can be met by restricting the modeling of roles in form of the role distributivity property which is motivated by the intuition that if there is a relation for an element in $(C \sqcup D)^{\mathcal{I}}$ which is neither in $C^{\mathcal{I}}$ nor in $D^{\mathcal{I}}$, then this relation must be based on relations of elements of C and of D.

Definition 8.7. *Role distributivity property RDP: if $x \in (C \sqcup D)^{\mathcal{I}} \backslash C^{\mathcal{I}} \backslash D^{\mathcal{I}}$ and $\pi_{2,R}^{-1}(x) \cap H_m \neq \overline{0}^{\mathcal{I}}$, then it exists $x_c \in (C)^{\mathcal{I}}$ and $x_d \in (D)^{\mathcal{I}}$ with $x = x_c + x_d$ and $\pi_{2,R}^{-1}(x_c) \subseteq H_m$ and $\pi_{2,R}^{-1}(x_d) \subseteq H_m$.*

Each two concepts C and D must fulfill the RDP to regain distributivity of the $\exists$-operator.

Proposition 8.8. *If RDP is fulfilled, then $\exists R.(C \sqcup D) = \exists R.C \sqcup \exists R.D$ is valid.*

Proof. $\exists R.C \sqcup \exists R.D \sqsubseteq \exists R.(C \sqcup D)$ holds in any case. Therefore, it is sufficient to show $\exists R.(C \sqcup D) \sqsubseteq \exists R.C \sqcup \exists R.D$. Therefore, for all $x \in (\exists R.(C \sqcup D))^{\mathcal{I}}$ it needs to follow that $x \in (\exists R.C \sqcup \exists R.D)^{\mathcal{I}}$. Let $x \notin (\exists R.C)^{\mathcal{I}}$, $x \notin (\exists R.D)^{\mathcal{I}}$, as trivial in the other cases. Therefore with RDP follows that $x = \pi_{1,R}(\pi_{2,R}^{-1}(x_c + x_d) \cap H_m)$ for a $x_c \in C$ and a $x_d \in D$. With linearity of $\pi_{2,R}$ it follows that $x = \pi_{1,R}((\pi_{2,R}^{-1}(x_c) + \pi_{2,R}^{-1}(x_d)) \cap H_m)$ and

$\pi_{2,R}^{-1}(x_c) \subseteq H_m$ and $\pi_{2,R}^{-1}(x_d) \subseteq H_m$ means $x = \pi_{1,R}((\pi_{2,R}^{-1}(x_c) \cap H_m) + (\pi_{2,R}^{-1}(x_d) \cap H_m))$. With linearity of $\pi_{1,R}$ follows equality. □

Having this property, it is possible to show that all $\mathcal{ALC}$ knowledge bases are representable by a cone-COS.

Proposition 8.9. *All $\mathcal{ALC}$ knowledge bases are representable by a cone-COS based on al-cones.*

Proof. We show that an al-cone-COS of an $\mathcal{ALC}$ knowledge base without roles, i.e., only considering the Boolean part, can be extended to a model for roles as well.

Since $\mathcal{ALC}$ features the finite model property we may assume that the geometric model is finite and represents all facts following from a given ontology $\mathcal{O}$. Hence assume all concepts to be represented by cones and all constants by vectors in $\mathbb{R}^n$. We write $a^{\mathcal{I}_B}$ to refer to the vector obtained for constant a in the Boolean embedding and we write $a^{\mathcal{I}}$ for its embedding we seek to construct.

We iteratively construct a geometric model from a Boolean geometric model based on al-cones and a corresponding $\mathcal{ALC}$ model by processing role after role in a two-step process. Initially, we initialize $\cdot^{\mathcal{I}}$ by setting $\cdot^{\mathcal{I}}$ to the Boolean-only model $\cdot^{\mathcal{I}_B}$. In the first step, we consider a role R with $|\{(a,b)|\mathcal{O} \models R(a,b)\}| = m$, where $|\cdot|$ denotes the cardinality, and assume $R = \{(a_1,b_1),\ldots,(a_m,b_m)\}$ in the finite model. We extend the dimensions of our model from n to $n(m+1)$ by cloning the components of all vectors $\vec{x}$ that generate some concept. Let 0_k denote k consecutive zero components in a vector, then we can describe the modification of the embedding $c^{\mathcal{I}}$ for any constant c as follows:

$$\phi(c) \quad = \sum_{i=1,\ldots,m,c=a_i \vee c=b_i} (0_{n\cdot i}\,(c^{\mathcal{I}})^T\,0_{m-i})^T \tag{8.1}$$

$$c^{\mathcal{I}} \quad \leftarrow \begin{cases} \phi(c) & \phi(c) \neq 0_{n(m+1)} \\ ((c^{\mathcal{I}})^T\,0_{n\cdot m})^T & \text{otherwise} \end{cases} \tag{8.2}$$

Note that $\phi(c) \neq 0_{n(m+1)}$ occurs exactly if there is at least one a_i or b_i with $c = a_i$ or $c = b_i$. Doing so, we separate all entities in $\mathrm{dom}(R) \cup \mathrm{Img}(R)$ that occur in the model. In particular, we achieve that $\lambda a_i^{\mathcal{I}} \in \mathrm{dom}(R)$, $\lambda > 0$ if and only if $a_i^{\mathcal{I}} \in \mathrm{dom}(R)$ and likewise for $\mathrm{Img}(R)$. We repeat the process for all roles.

In the second step, we need to construct the reification of any role R which can be done as follows. Assume again $R = \{(a_1,b_1),\ldots,(a_m,b_m)\}$ and then define a reification based on hyperoctant H_l embedded in the model using projections

$$\pi_1 = \left[a_1^{\mathcal{I}} \cdots a_m^{\mathcal{I}}\right], \pi_2 = \left[b_1^{\mathcal{I}} \cdots b_m^{\mathcal{I}}\right],$$

where $[\cdots]$ represents a matrix composed out of column vectors. By construction of $a_i^{\mathcal{I}}$, $b_i^{\mathcal{I}}$, we have $c_{R(a,b)} = (0_{i-1}\, 1\, 0_{n-1})^T \in H_l$ which corresponds to $R(a_i, b_i)$ since $\pi_1(c_{R(a,b)}) = a_i^{\mathcal{I}}$ and $\pi_2(c_{R(a,b)}) = b_i^{\mathcal{I}}$. It thus follows $\pi_1(H_l) \supseteq \mathrm{dom}(R)$ and $\pi_2(H_l) \supseteq \mathrm{Img}(R)$, respectively. Also by construction, for any $c \notin \mathrm{dom}(R)$ we have $c^{\mathcal{I}} \notin \pi_1(H_l)$ since $a_i^{\mathcal{I}}$ and $c^{\mathcal{I}}$ reside then in mutually exclusive sub-spaces according to Equation (8.2). □

We note that this proof, albeit constructive, is of theoretical nature since it exploits a large amount of dimensions for H_m to ease the construction.

It is not only possible to represent each $\mathcal{ALC}$ knowledge base with such a geometric interpretation, it is also possible to interpret each geometric model based on the axis-aligned cones introduced in the above proposition as an $\mathcal{ALC}$-ontology.

Proposition 8.10. *A geometric interpretation based on al-cones fulfilling RDP, where $\pi_{1,R}(\pi_{2,R}(C) \cap H_m)$ maps to an al-cone for each half-axis C, represents an $\mathcal{ALC}$ knowledge base.*

Proof. A geometric interpretation without considering roles is shown by Özçep et al. [2023, Proposition 7]. Therefore, it is sufficient to show that the relational part also fulfills the restrictions of $\mathcal{ALC}$. $\exists R.\overline{0} = \overline{0}$ is fulfilled by construction. The distributivity of $\exists$ over $\sqcup$ is ensured by RDP, as shown in Proposition 8.8. As each half-axis is mapped to an al-cone, because of linearity of π, each concept is mapped to a union of al-cones, which is still an al-cone. Also because of linearity, it is ensured that the properties needed for roles, e.g. $\exists R.C \sqsubseteq \exists R.\overline{1}$ are fulfilled. The negation of $\exists R.C$ is given by polarity (as it would not be a geometric model otherwise). Therefore, the resulting geometric model represents an $\mathcal{ALC}$ knowledge base. □

8.2 Cone Embedding for Ontologies with Arbitrary Quantifier Rank

In the last section, reification turned out to be a useful approach for the incorporation of roles into cone-COSs. As stated in Proposition 8.9, it allows for modeling all $\mathcal{ALC}$-knowledge bases. In this section, the focus switches from solely modeling knowledge bases to modeling them faithfully. For Boolean $\mathcal{ALC}$-ontologies, constructing a faithful COS is possible as shown in Proposition 5.19 (for the definition of faithfulness, see Definition 3.4). This result is now extended to full $\mathcal{ALC}$.

For a given tbox, faithfulness can be gained for each possible abox. Obviously, the complexity of the resulting model correlates with the complexity of the abox. In the proof of the construction of faithfulness for an al-cone-COS for Boolean $\mathcal{ALC}$-ontologies in Proposition 5.19, the construction is done for an abox based on the atomic elements in the Boolean algebra generated by the tbox of the ontology, thus for the most general abox.

Thus, in order to gain faithfulness, it is necessary to look at each possible abox. For Boolean-$\mathcal{ALC}$ al-cone-COSs, it was possible to model each atomic element on a half-axis. However, an extension of this construction to full $\mathcal{ALC}$ is not straightforwardly possible. Problems occur in particular because of the occurrence of chains of roles with arbitrary length in the abox. In the rest of this section, I consider these problems in detail and present a solution, based on al-cone-COSs but different to the reification-based approach of the last section. The main result of this section is the proof that it is possible to model each full $\mathcal{ALC}$-ontology faithfully with the help of al-cone-COSs.

Remember: An al-cone-COS $S = (\mathbb{R}^n, Y, \mathcal{B}_E, \perp_p)$ is based on the Euclidean betweenness $\mathcal{B}_E$, the orthogonality relation $\perp_p$ based on polarity, and Y consisting of axis-aligned cones. An axis-aligned cone is defined as

$$C \text{ is an al-cone iff } C = (C_i)_{1 \leq i \leq n} = C_1 \times \cdots \times C_n, \text{ where } C_i \in \{\mathbb{R}, \mathbb{R}_+, \mathbb{R}_-, \{0\}\}.$$

For details, see, e.g., Example 5.17. The definition immediately entails the fact that al-cones in $\mathbb{R}^n$ are characterized by the sets of half-axes contained in it. So I introduce an operator $\mathit{halfAxes}(\cdot)$ denoting the half-axes of an al-cone X:

$$\begin{aligned} \mathit{halfAxes}(X) \;\; = \;\; & \{H \mid H = \{0\} \times \cdots \times \{0\} \times X_i \times \{0\} \times \cdots \times \{0\} \in X \\ & \text{for } X_i \in \{\mathbb{R}_+, \mathbb{R}_-\} \text{ and } i = 1, \ldots, n\} \end{aligned}$$

In the following, for better readability, the al-cone-COS-interpretation is abbreviated by the term *geometric/al-cone interpretation* $(\Delta^\mathcal{I}, \mathcal{I})$.

Our main task is to define semantics for roles and quantifiers. In contrast to established embedding approaches such as TransE [Bordes et al., 2013] and to the last section the aim here is not to model the role $R(a, b)$ between two specific instances a and b directly, but to model it on the conceptual level by proving assertions such as $\exists R.C(a)$. This is done by defining an interpretation $(R)^\mathcal{I}$ of role R as matrix $\mathcal{R}$, such that $R(a, b)$ is represented by the statement that $a^\mathcal{I}$ is in the al-cone generated by $\mathcal{R}b^\mathcal{I}$. Likewise, al-cones can be transformed by matrix multiplication. We will later see that $\mathcal{R}$ can be chosen such that the transformation always yields an al-cone. Then, for an instance $a^\mathcal{I}$ the interpretations of the type $(\exists R.C)^\mathcal{I}$ with $a^\mathcal{I} \in (\exists R.C)^\mathcal{I}$ can be determined by applying the matrix $\mathcal{R}$ to the interpretation $(C)^\mathcal{I}$ (where C itself can also contain an existential) and checking whether $a^\mathcal{I}$ is contained. Therefore, abox-faithfulness in its original form cannot be ensured. This distinguishes this approach from the reification-based on of the last section: that one discusses expressivity on abox- and tbox-level, whereas this one is focused solely on the tbox and the faithful representation thereof.

In technical terms, modeling (also roles) on the conceptual level amounts to showing that the Lindenbaum-Tarski algebra of an $\mathcal{ALC}$ tbox makes up a Boolean algebra with operators (BAOs) [Jónsson and Tarski, 1951; Jónsson and Tarski, 1952]. This in turn amounts to showing that on top of having a Boolean algebra the semantics of $\exists R.(\cdot)$ in a geometric model $\mathcal{I}$ gives an operator with the following properties: $(\exists R.\overline{0})^{\mathcal{I}} = (\overline{0})^{\mathcal{I}}$ and $(\exists R.(C \sqcup D))^{\mathcal{I}} = (\exists R.C)^{\mathcal{I}} \sqcup (\exists R.D)^{\mathcal{I}}$.

This enables negation to be generalized for concepts that contain roles. So we may express also negated concepts with roles such as $(\neg\exists R.C)(a)$. (Of course we do not model explicitly negation for roles as there is no negation operator for roles in $\mathcal{ALC}$).

The aim is now to adapt the notion of faithfulness to roles and to find a faithful geometric embedding based on al-cones including roles. Concept-faithfulness can be directly adapted to full $\mathcal{ALC}$, on the one hand restricted up to a specific quantifier rank, on the other hand for an unrestricted ontology.

The next definition adapts the geometric faithful models from Definition 3.4 to full $\mathcal{ALC}$.

Definition 8.11. *Let $\mathcal{O}$ be a classically consistent (DL) ontology (or any other representation allowing the distinction between abox and tbox). For a (not necessarily classical) interpretation $\mathcal{I}$ we have the following notions of being a faithful model of $\mathcal{O}$:*

1. *$\mathcal{I}$ is an m-*(quantifier)-rank concept-faithful *model of $\mathcal{O}$ iff for each concept C with rank at most m and each constant b the following holds: if $b^{\mathcal{I}} \in C^{\mathcal{I}}$, then $\mathcal{O} \models C(b)$*
2. *$\mathcal{I}$ is a* concept-faithful *model of $\mathcal{O}$ iff for each concept C and each constant b the following holds: if $b^{\mathcal{I}} \in C^{\mathcal{I}}$, then $\mathcal{O} \models C(b)$.*

First, I consider the problem only on a conceptual level, by which I mean that I do not specify a geometric interpretation of roles but only geometric models of concepts including role information. An idea for the construction would be to consider algebraic atoms including roles and to place them on half-axes as for the Boolean case. This is not possible due to the fact that the algebra of $\mathcal{ALC}$ concepts is not an atomic lattice as in the Boolean case.

Example 8.12. *Consider as example an abox containing* loves(narcissus, narcissus) *and* Vain(narcissus) *described by Baader and Küsters [2006]. The instance* narcissus *fulfills all concepts of the form* $\exists \text{loves}^n.\text{Vain}$, $n = 1, 2, \dots$ *which, connected conjunctively, give a chain of concepts C_i that become infinitely narrower and narrower. The C_is are defined by:*

$$C_0 := \exists R.\text{Vain} \textit{ and } C_i := C_{i-1} \sqcap \exists \text{loves}^{i+1}.\text{Vain}$$

Thus, narcissus *would be in a concept of each quantifier rank.*

Even without a circular relationship in the abox, it is possible to have concepts of arbitrary quantifier rank in form of a chain of roles of arbitrary depth, e.g. $R(a_1, a_2)$, $R(a_2, a_3)$, ..., $R(a_{i-1}, a_i)$ for an arbitrary (unknown) $i \in \mathbb{N}$. Depth i cannot be determined because, due to the open world assumption, it is possible that roles exist for which the extension is not stated in the abox. Thus we see that the Lindenbaum-Tarski algebra of $\mathcal{ALC}$ tboxes (in this case) may not be atomic and an abox may contain an object for which there is no most specific concept as exemplified in Example 8.12.

Sketch of the solution. Let me sketch my solution to tackle this problem of non-atomicity before we embark on technicalities. I am going to show that I can embed the Lindenbaum-Tarski algebra of an $\mathcal{ALC}$-tbox into an extended algebra with elements that play the role of atomic elements. These elements are constructed as infinite conjunctions of algebraic elements of the Lindenbaum-Tarski Algebra of a special kind (elements of $\mathcal{X}_i$ in Lemma 8.28 below). So, actually, I am doing more than embedding just $\mathcal{ALC}$-concepts but I embed carefully chosen infinite versions of $\mathcal{ALC}$ concepts.

To be more concrete, the infinitary logic $\mathcal{ALC}^{inf}$ I refer to is defined as follows: The logical symbols formulae of $\mathcal{ALC}^{inf}$ consist of strings that make up trees that can infinitely branch and can have infinitely long branches. As for finite trees, for these trees we have a notion of a level, level 0 being the level of the root. The branches of the trees are infinite words $\omega : \mathbb{N} \to \Sigma$ where

$$\Sigma = N_R \cup N_C \cup \{\exists, \sqcap, \sqcup, \neg, \bigsqcap, \bigsqcup, (,), .\}$$

is the alphabet consisting of a vocabulary, the logical symbols of $\mathcal{ALC}$ and two additional logical symbols $\bigsqcap, \bigsqcup$. Example words are of the form $\exists R.\exists R.\dots$ or $\exists R.\exists S \sqcap C.\exists R\dots$. Hence, if one wants to define specific subsets of those formulae (as done, e.g., in Lemma 8.33), one has to rely on coinduction [Rutten, 2005]: Operations on infinite branches ω (of a tree) can be defined by defining the operation on the head $\omega(0)$ and on the derivative ω' of ω defined as $\omega'(n) = \omega(n+1)$. The two additional logical symbols of $\mathcal{ALC}^{inf}$ mentioned above, namely $\bigsqcap, \bigsqcup$, are intended to model infinitary conjunction and disjunction (these lead to the infinite branchings). The definition of an $\mathcal{ALC}^{inf}$ concept extends the definition of an $\mathcal{ALC}$ concept by two rules: If X is a (possibly infinite) set of $\mathcal{ALC}^{inf}$ concepts then so are $\bigsqcap X$ and $\bigsqcup X$. The semantics is as expected: $(\bigsqcap X)^{\mathcal{I}} = \bigcap\{C \mid C \in X\}$ and $(\bigsqcup X)^{\mathcal{I}} = \neg \bigcap\{\neg C \mid C \in X\}$.

Intuitively, the extended algebra is a natural limit of arbitrary Lindenbaum-Tarski algebras for $\mathcal{ALC}$ tboxes. The notion of limit can be made precise by a (new) rank notion. This notion is also necessary to account for the fact that with our construction there might be infinitely many algebraic atoms in our extended algebra — due to the fact that they are based on a possibly infinite number of conjuncts.

I am going to consider first the simple case of an empty tbox in the non-cyclic and in the cyclic case and then treat arbitrary tboxes.

Handling the Non-Cyclic Case

Now let us fill the sketch with details. How do the algebraic atoms in the extended algebra look like? They are defined in the Lindenbaum-Tarski algebra of a tbox in $\mathcal{ALC}$ considering the equivalence of $\mathcal{ALC}^{inf}$ concepts.

The idea is to construct them by an adequate choice of its conjuncts by using, e.g., only one positive existential and infinitely many negative ones in it. One first approach by Özçep et al. [2020] is to restrict $\mathcal{ALC}$-ontologies to contain concepts up to a specific quantifier rank to get a m-(quantifier)-rank concept-faithful model. But this restricts the expressivity of the model, as faithfulness for higher ranks is lost.

Independent of the used construction, a drawback of a faithful interpretation is that it results in an infinite-dimensional cone model because of its infinitely many algebraic atoms (except for highly restricted ontologies). Because of the infinite-dimensional model, it is not possible to create the model as a whole in a suitable way, therefore, it is necessary to be able to extend the model iteratively whenever necessary without influencing the existing one. This also enables us to stick to the half-axis based construction principle of the Boolean case.

Thus, it is necessary to have a creation principle for algebraic atoms of the extended algebra which is able to handle their (possibly) infinite number. One idea is to use some notion of rank to model concepts of at most a specific rank, and create out of those rank-restricted concepts algebraic atoms based on the extended algebra. Then the modeling known from the Boolean case can be used. Therefore, a notion of rank is searched for, for which rank-restricted algebraic atoms are also algebraic atoms (in the extended algebra) in the non-restricted case and thus enables for iterative extension without having to change the original model.

Therefore, first a definition of rank-restricted interpretations is needed.

Definition 8.13. *An interpretation $\mathcal{I}$ is an m-*rank model of an ontology $(\mathcal{T}, \mathcal{A})$ *with an empty tbox $\mathcal{T}$ iff for all constants a and concepts C such that $\mathcal{A} \models C(a)$ and the rank of C is at most m: $\mathcal{I} \models C(a)$.*

This definition is independent of the type of rank considered. To demonstrate the necessary properties, we consider first the standard (syntactical) notion of quantifier rank as introduced in Section 2.1 and show its inappropriateness.

Example 8.14. *Consider a simple ontology with an empty tbox, a concept symbol A and a role symbol R. For the creation of a* 0*-quantifier-rank model, only the algebraic* *atom A and $\neg A$ need to be considered, as the rank $qr(\exists R.\overline{1}) > 0$, $qr(\neg\exists R.\overline{1}) > 0$. Thus, a one-dimensional space is needed. Let $\mathbb{R}_+$ represent $A^{\mathcal{I}}$, then $\mathbb{R}_-$ represents $(\neg A)^{\mathcal{I}}$ (or vice versa). A higher-ranked concept, e.g., $A \sqcap \exists R.A$ can be approximated with concept A, as $A \sqcap \exists R.A \sqsubseteq A$. Therefore, all instances of $(A \sqcap \exists R.A)^{\mathcal{I}}$ can be placed in $\mathbb{R}_+$. The same holds for instances in $(A \sqcap \exists R.\neg A)^{\mathcal{I}}$. Therefore, each k-quantifier-rank model (for arbitrary $k \geq 0$) can represent all instances of the abox correctly, however, not necessary as specific as it would be possible using a higher rank.*

Therefore, algebraic atoms of quantifier-rank restricted models are not algebraic atoms of unrestricted models. If the model is extended to an 1-quantifier-rank model, A is not an algebraic atom anymore, as it contains, e.g., $A \sqcap \exists R.A$ and $A \sqcap \neg\exists R.A$. Therefore, the instances placed in $\mathbb{R}_+$ in the smaller model cannot be placed in $\mathbb{R}_+ \times \{0\} \times \cdots \times \{0\}$ in the bigger model as this would contradict $\exists R.A \sqcap \neg\exists R.A = \overline{0}$. Therefore, increasing the represented quantifier rank leads to the necessity of constructing a completely new model and thus, the quantifier rank is not suitable for the creation of an iteratively extendable model based on the rank.

To circumvent the restriction mentioned in the example, it is necessary to define a new notion of rank, which I call the *semantic quantifier rank*. This leads to the possibility of creating a semantic-rank-restricted model which depicts a subspace of the model of the whole ontology, meaning that if an algebraic atom is modeled on a specific half-axis in the model with rank-restricted concepts, this algebraic atom is also placed on exactly this half-axis in the model of the whole ontology. Thus, algebraic atoms of a smaller rank remain algebraic atoms in models of higher rank and the model is iteratively extendable without influencing the underlying smaller sized model. To gain this, concepts of smaller ranks cannot be interpreted as approximations of concepts of higher rank as done in Example 8.14. If a concept is either of the represented rank or smaller, then it can be represented correctly and accurate, if it has a higher rank, then it can not be represented or approximated at all (meaning higher ranked concepts are modeled as $\overline{0}$ in a model representing lower ranked concepts). By circumventing the approximation in this way, it is not necessary to change the existing model when adding concepts with a higher rank and thus adding new dimensions to the model.

Intuitively, the semantic quantifier rank describes on which depth it is possible to model a concept having an actual extension (not being $\overline{0}$). It therefore represents the necessary size of the model to be able to model a specific concept. The definition tests in the beginning whether the concept or a part of the concept is equivalent to $\overline{1}$ or $\overline{0}$ and then proceeds by induction on the structure of the concept. The semantic rank gives either a natural number or ∞. Hence we assume in the following definition that ∞ is greater than all natural numbers and that $\min, \max, +$ handle ∞ accordingly.

Both, srank and *reduce* are defined for existentials only, thus universal quantifiers are interpreted as existentials by setting $\forall R.C = \neg\exists R.\neg C$. $reduce(D)$ is defined inductively as $reduce(D) = D$ for D being an atomic symbol or $reduce(D) = \overline{1}$ when D is equivalent to $\overline{1}$ or $reduce(D) = \overline{0}$ when D is equivalent to $\overline{0}$. $reduce(C \sqcap D) = reduce(C)$ if $reduce(D)$ is equivalent to $\overline{1}$, $reduce(C \sqcap D) = reduce(D)$ if $reduce(C)$ is equivalent to $\overline{1}$ and $reduce(C \sqcap D) = reduce(C) \sqcap reduce(D)$ otherwise; similarly for $\sqcup$ and $\neg$. $reduce(\exists R.C) = \overline{0}$ if $reduce(C) = \overline{0}$, else $\exists R.reduce(C)$.

Definition 8.15. *The* semantic quantifier rank srank *is defined for each $\mathcal{ALC}$ concept $D = reduce(D')$ as follows:*

- $srank(D) = \infty$ *if D is equivalent to $\overline{0}$*
- $srank(D) = 0$ *if D is equivalent to $\overline{1}$*
- $srank(D) = 0$ *if D is a concept symbol and different from $\overline{0}$ and $\overline{1}$*
- $srank(\neg D) = 0$ *if D is a concept symbol and different from $\overline{0}$ and $\overline{1}$*
- $srank(C \sqcap D) = \max(\{srank(C), srank(D)\})$
- $srank(C \sqcup D) = \min(\{srank(C), srank(D)\})$
- $srank(\exists R.D) = srank(D) + 1$
- $srank(\neg\exists R.D) = 0$
- $srank(\neg(C \sqcap D)) = \min(\{srank(\neg C), srank(\neg D)\})$
- $srank(\neg(C \sqcup D)) = \max(\{srank(\neg C), srank(\neg D)\})$

Note that the case of the universal quantifier is captured by its equivalent description using the existential quantifier: $\forall R.C = \neg\exists R.\neg C$ and thus $srank(\forall R.C) = 0$.

The following example illustrates the calculation of sranks.

Example 8.16. *Consider concepts $\exists R.C$ and $\exists R^2.C$ for concept symbol C:*

$$srank(\exists R.C) = 1, srank(\neg\exists R.C) = 0,$$

and

$$srank(\exists R^2.C) = 2, srank(\neg\exists R^2.C) = 0.$$

The definition can be interpreted as follows. Concept symbols can be represented in each model, for these it is not necessary to model roles. Having a conjunction, it is necessary to model both conjuncts, therefore, the conjunct with maximal srank must be considered. Having a union, it is sufficient to model the part with the smaller srank to have a union which is not bottom. The srank is increased when an existential is used, as it increases the depth of the path defined by the roles. A negated existential has a srank of zero, as there is no role necessary to model it (only the nonexistence of a role needs to be modeled).

It is necessary to consider concepts and parts of concept D equivalent to $\overline{0}$ and $\overline{1}$ separately. For a concept equivalent to $\overline{1}$ there is an extension independent of the srank considered and an arbitrary srank can be used to model the extension. The srank of a concept equivalent to $\overline{0}$ is ∞, as there can't be any concept extension (different from the point of origin) representing the $\overline{0}$-concept. Therefore, concepts being equivalent to $\overline{0}$ or $\overline{1}$ are ignored for creating the srank. Thus, having a concept for which the srank should be determined, it has to be transformed, e.g., by using de Morgan or the distributivity rule to determine all parts representing the $\overline{0}$-concept. After that, for the remaining terms the srank can be determined. One can show that the defined rank notion is indeed semantical in the sense that equivalent concepts (based on an empty tbox) have the same quantifier rank. Later on, this approach will be extended to ontologies restricted by a non-empty tbox.

Proposition 8.17. *If $C \sqsubseteq D$, then $srank(C) \geq srank(D)$.*

Proof. First, it is shown that the srank of a concept is independent of its syntactical form (except of subconcepts being equivalent to bottom and top as mentioned in Definition 8.15). The rules of de Morgan preserve the srank, as $srank(\neg(C \sqcap D)) = srank(\neg C \sqcup \neg D) = \min(\{srank(\neg C), srank(\neg D)\}) = srank(\neg C \sqcup \neg D)$. Distributivity is also preserved by srank as for example

$$\begin{aligned}
&srank((A \sqcap B) \sqcup C) \\
&= \min\big(\{\max(\{srank(A), srank(B)\}), srank(C)\}\big) \\
&= \max\big(\{\min(\{srank(A), srank(C)\}), \min(\{srank(B), srank(C)\})\}\big) \\
&= srank\ \big((A \sqcup B) \sqcap (B \sqcup C)\big).
\end{aligned}$$

This follows in the same way for the other cases of de Morgan and distributivity.

Let $C \sqsubseteq D$ and $srank(C) = i$. As $C \sqsubseteq D$, $D = C \sqcup E$ for some E. Thus $srank(D) = \min(\{srank(C), srank(E)\}) \leq srank(C) = i$. This can be done, as the srank of a concept is independent of its syntactical form as shown above. Thus, $srank(D) \leq srank(C)$. □

The following lemma shows the necessary properties mentioned in the motivation. Let a geometric model M_i be defined as an i-srank-model.

Lemma 8.18. *Model M_i and model M_{i+1} of an ontology $\mathcal{O}$ with an empty tbox have the following properties:*

1. *An algebraic atom C of M_i with $srank(C) = i$ is an algebraic atom of the unrestricted geometric model of $\mathcal{O}$.*
2. *M_i and M_{i+1} can be chosen such that M_i is a subspace of M_{i+1}.*

Proof.

1. Let C be an algebraic atom with $srank(C) = i$ and assume C would not be an algebraic atom of the unrestricted geometric model. Therefore, there is a concept D with $C \sqcap D \neq \overline{0}$ with $srank(D) = j > i$. As $srank(D) = j$, D must contain at least one conjunct D' with $srank(D') = j$ and thus $srank(\neg D) = 0$. Now, consider $C \sqcap \neg D$: If $C \sqcap \neg D \neq \overline{0}$, then either C is not an algebraic atom in M_i or $C \sqsubseteq \neg D$, but then $C \sqcap D = \overline{0}$, a contradiction. Thus, assume $C \sqcap \neg D = \overline{0}$. Then, $C \sqsubseteq D$ (due to (wLLJ), see Appendix A for details) but with Proposition 8.17 $srank(C) \geq srank(D)$, a contradiction, as $srank(C) = i < j = srank(D)$.
2. Let M_i be represented in some $\mathbb{R}^n$ where the algebraic atoms are interpreted as half-axes. Let $\mathcal{Y}_i$ be the set of algebraic atoms of M_i and $\mathcal{Y}_{i+1}$ the set of algebraic atoms of M_{i+1}. With 1. it follows that $\mathcal{Y}_i \subseteq \mathcal{Y}_{i+1}$. Algebraic atoms of $\mathcal{Y}_{i+1} \backslash \mathcal{Y}_i$ can thus be represented in some $\mathbb{R}^m$ independently of the $\mathbb{R}^n$ of M_i. Both spaces can be concatenated to a space $\mathbb{R}^{m+n}$ representing M_{i+1}.

□

Example 8.19 (Example 8.14 continued)**.** *Consider again the simple ontology with an empty tbox, a concept symbol A and a role symbol R. In the following, it is shown that the srank is suitable for the iterative creation of a geometric model. Consider a 0-srank-model. It is as for the 0-quantifier-rank model a one-dimensional space, but with the difference that not A and $\neg A$ are represented but the algebraic atoms of srank 0, thus, $A \sqcap \neg \exists R.\overline{1}$ and $\neg A \sqcap \neg \exists R.\overline{1}$ on $\mathbb{R}_+$ resp. $\mathbb{R}_-$. Thus, only instances having no relation at all can be represented and instances of, e.g., $(A \sqcap \exists R.A)^{\mathcal{I}}$ cannot be placed in this model. This circumvents the problem which appeared in Example 8.14 when extending the dimension. Here, the 1-srank-model consists of the algebraic atoms $A \sqcap \neg \exists R.\overline{1}$ and $\neg A \sqcap \neg \exists R.\overline{1}$ on the one hand and $V \sqcap \exists R.(W \sqcap \neg \exists R.\overline{1}) \sqcap \neg \exists R.(\exists R.\overline{1})$ for $V, W \in \{A, \neg A\}$ on the other hand. As there is no algebraic atom in the 1-srank-model which is a proper subsumer* *of an algebraic atom in the* 0*-srank-model, it is possible to keep the representation of* $(A \sqcap \neg\exists R.\overline{1})^{\mathcal{I}}$ *by extending it from* $\mathbb{R}_+$ *to* $\mathbb{R}_+ \times \{0\} \times \{0\}$ *(analogously for* $(\neg A \sqcap \neg\exists R.\overline{1})^{\mathcal{I}}$*). The algebraic atoms of the form* $V \sqcap \exists R.(W \sqcap \neg\exists R.\overline{1}) \sqcap \neg\exists R.(\exists R.\overline{1})$ *for* $V, W \in \{A, \neg A\}$ *can then be placed on the half-axes of the second and third dimension and thus do not interfere with the lower dimensional model.*

Having this notion of a semantic rank, it is possible to create a srank-restricted model and extend it iteratively. However, this construction treats roles completely on a conceptual level in the sense that there is no geometric operation for the representation of the role. To mitigate this problem, we observe that it should be possible to describe the changes in the semantic rank caused by a role in a geometrical way. For example, if $a^{\mathcal{I}} \in (\exists R^2.\overline{1} \sqcap \neg\exists R^3.\overline{1})^{\mathcal{I}}$ and $R(a,b)$ is valid, applying an interpretation of R one time should result in a $b^{\mathcal{I}} \in (\exists R.\overline{1} \sqcap \neg\exists R^2.\overline{1})^{\mathcal{I}}$. Thus, a representation $R^{\mathcal{I}}$ of role R is needed.

The role R is represented as an incidence matrix $\mathcal{R}$ that maps each half-axis to arbitrary many half-axes of the model. Interpreting R as incidence matrix allows for extending it iteratively while increasing dimensions.

Definition 8.20. *An* al-cone interpretation $\mathcal{I}$ *is a Boolean al-cone interpretation including additionally matrices* $\mathcal{R}$ *representing roles* R*. An* al-cone interpretation *of* $\mathcal{ALC}$ *concepts is defined recursively as for Boolean concepts and defining the concepts of the form* $\exists R.C$ *as al-cone, as* $(\exists R.C)^{\mathcal{I}} = conH(\{\mathcal{R}^T y \mid y \in C^{\mathcal{I}}\})$ *with* R *interpreted as incidence matrix* $\mathcal{R}$*. The definition of the all quantifier is given by de Morgan, i.e.,* $(\forall R.C)^{\mathcal{I}} = (\neg\exists R.\neg C)^{\mathcal{I}}$*.*

As $\mathcal{R}$ is an incidence matrix of size $\mathbb{R}^{2n \times 2n}$, where n is the size of the geometric model, the above mentioned definition contains a slight simplification. More concretely, it is based on the incidence vectors $y' \in \mathbb{R}^{2n}$ of $y \in R^{\mathcal{I}}$ where in even dimensions of y' the positive half-axes and in odd dimensions the negative half-axes are represented. Multiplication gives the incidence vector $x' = \mathcal{R}^T y'$ which needs to be transformed to al-cones by splitting it into half-axes. Therefore, $\exists R.C = conH(\{halfAxes(x') \mid x' = \mathcal{R}^T y' \text{ with } y' = \text{incidenceVector}(y) \text{ with } y \in C^{\mathcal{I}}\})$. For simplicity this is abbreviated with $(\exists R.C)^{\mathcal{I}} = conH(\{\mathcal{R}^T y \mid y \in C^{\mathcal{I}}\})$ in the following.

Example 8.21. *Consider the example of Narcissus mentioned in Example 8.12. Assume an empty tbox, one concept* Vain *and one role* loves*, for short* V, R*.*

Let the al-cone interpretation be in $\mathbb{R}^2$*,* $V^{\mathcal{I}} = \mathbb{R}_+ \times \mathbb{R}_-$ *and*

$$\mathcal{R}^T = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

The incidence vector of instance $y^{\mathcal{I}} = (1\ -1)^T \in V^{\mathcal{I}}$ *would be* $y' = (1\,0\,0\,1)^T$*.* $x' = \mathcal{R}^T y' = (1\,0\,1\,0)^T$ *and thus* $x^{\mathcal{I}} = (1\,1)^T$*. This can be split into the half-axes* $\mathbb{R}_+ \times \{0\}$ *and* $\{0\} \times \mathbb{R}_+$*. Thus, when considering all instances of* $V^{\mathcal{I}}$*, then it results in* $(\exists R.V)^{\mathcal{I}} = conH(\{(1\,0)^T, (0\,1)^T\}) = \mathbb{R}_+ \times \mathbb{R}_+$*.*

$\mathcal{R}$ can be used to determine possible roles for a specific instance x, e.g., it can be checked, whether $(\mathcal{R}x^{\mathcal{I}}) \cap C^{\mathcal{I}} \neq \overline{0}^{\mathcal{I}}$, to state that $x^{\mathcal{I}} \in (\exists R.C)^{\mathcal{I}}$. As the concept $\exists R.C$ is created based on the instances of C, to answer such questions, the incidence matrix needs to be used in non-transposed form $\mathcal{R}$ here.

As for the Boolean case, it is also possible in the non-Boolean case to model partial knowledge, thus considering faithfulness, e.g., $x \in (\exists R.\overline{1} \sqcap (\exists S.\overline{1} \sqcup \neg\exists S.\overline{1}))^{\mathcal{I}}$ (where the second part does not add any information and is only added for demonstration purposes), meaning it is known that a relation with role R exists, however, nothing is known about a role S. This can be modeled in the same way as in the Boolean case, by placing instance $x^{\mathcal{I}}$ in $(\exists R.\overline{1})^{\mathcal{I}}$, yet between $(\exists S.\overline{1})^{\mathcal{I}}$ and $(\neg\exists S.\overline{1})^{\mathcal{I}}$. Reasoning about partial knowledge works as follows: Consider the question whether $x^{\mathcal{I}} \in (\exists S.\overline{1})^{\mathcal{I}}$, $\mathcal{S}x^{\mathcal{I}}$ would not result in $\overline{0}^{\mathcal{I}}$, but this does not mean that $x^{\mathcal{I}} \in (\exists S.\overline{1})^{\mathcal{I}}$ is necessarily valid. Therefore, as a second step, it is necessary to check the other way around, namely, to check whether $x^{\mathcal{I}} \in conH(\{\mathcal{S}^T c \mid c \in \overline{1}^{\mathcal{I}}\}))$. The combination of the two steps results in determining whether x is a positive instance of a concept, or a negative instance or whether only partial knowledge is given. Thus, x has possibly but not necessarily a role $\mathcal{S}$.

In fact, representation $\mathcal{R}$ of role R behaves as desired with existential (and hence also with universal) quantifiers: it makes $\exists R$ indeed to a normal, additive operator f in a Boolean algebra with operators, i.e. f fulfills $f(\mathbb{0}) = \mathbb{0}$ (where $\mathbb{0}$ is the smallest element in the algebra; see Appendix A) and $f(a \vee b) = f(a) \vee f(b)$.

Proposition 8.22.

1. *Given an al-cone interpretation* $\mathcal{I}$ *and an arbitrary* $\mathcal{R}$*, then* $(\exists R.C)^{\mathcal{I}}$ *yields for arbitrary al-cones* $C^{\mathcal{I}}$ *an al-cone.*
2. $\exists R.C$ *interpreted as* $\mathcal{R}^T C^{\mathcal{I}}$ *maps the cone representing the bottom concept onto itself:* $(\exists R.\overline{0})^{\mathcal{I}} = \overline{0}^{\mathcal{I}}$ *and distributes over cone disjunction:* $\mathcal{R}^T(C^{\mathcal{I}} \sqcup D^{\mathcal{I}}) = conH(\mathcal{R}^T C^{\mathcal{I}} \cup \mathcal{R}^T D^{\mathcal{I}})$*, i.e.,* $\exists R.(C \sqcup D) = \exists R.C \sqcup \exists R.D$*.*

For the proof of this proposition, the following lemma is needed:

Lemma 8.23. *[Özçep et al., 2023, Proposition 3] For all al-cones* X, Y*:*

$$halfAxes(conH(X \cup Y)) = halfAxes(X) \cup halfAxes(Y)$$

Proof of Proposition 8.22.

1. The incidence matrix $\mathcal{R}$ only consists of ones and zeros, i.e., $(\mathcal{R})_{i,j} \in \{0, 1\}$. Therefore, multiplication with arbitrary half-axes according to the aforementioned calculation procedure results in a convex hull over half-axes, i.e., an al-cone.

2. $(\exists R.\overline{0})^{\mathcal{I}} = \overline{0}^{\mathcal{I}}$, as $(\exists R.\overline{0})^{\mathcal{I}} = conH(\{\mathcal{R}^T y \mid y \in \{\vec{0}\}\} = \{\vec{0}\} = \overline{0}^{\mathcal{I}}$.

   Due to Lemma 8.23, for $(C \sqcup D)^{\mathcal{I}}$ there is no half-axis (thus no al-cone) which is not contained in either $C^{\mathcal{I}}$ or $D^{\mathcal{I}}$, yet contained in $(C \sqcup D)^{\mathcal{I}}$. Having this insight and the above mentioned calculation rules for $(\exists R.C)^{\mathcal{I}}$, it follows that

$$\begin{aligned}
(\exists R.(C \sqcup D))^{\mathcal{I}} &= conH(\{\mathcal{R}^T y \mid y \in (C \sqcup D)^{\mathcal{I}}\}) \\
&= conH(\{\mathcal{R}^T y \mid y \in halfAxes((C \sqcup D)^{\mathcal{I}})\}) \\
&= conH(\{\mathcal{R}^T y, \mathcal{R}^T z \mid y \in halfAxes(C^{\mathcal{I}}), z \in halfAxes(D^{\mathcal{I}})\} \\
&= conH(\{conH(\{\mathcal{R}^T y \mid y \in C^{\mathcal{I}}\}), conH(\{\mathcal{R}^T z \mid z \in D^{\mathcal{I}}\})\}) \\
&= (\exists R.C \sqcup \exists R.D)^{\mathcal{I}}
\end{aligned}$$

□

Now, our first aim is to find a satisfiable geometric model of a given ontology based on the interpretation of roles as incidence matrices on a conceptual level, meaning, if $\mathcal{O} \models C(b)$, then $b^{\mathcal{I}} \in C^{\mathcal{I}}$ for each concept C and each constant b and, if $\mathcal{O} \models R(a, b)$ and $\mathcal{O} \models C(b)$, then $a^{\mathcal{I}} \in (\exists R.C)^{\mathcal{I}}$. In the first step, the focus does not lie on the correct representation of roles, i.e., guaranteeing that if $\mathcal{O} \models R(a, b)$, then $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$. Thus, "role-satisfiability" is considered, i.e. the mapping into correct regions in the sense that if $\mathcal{O} \models R(a, b)$, then $b^{\mathcal{I}} \in \mathcal{R}a^{\mathcal{I}}$.

Theorem 8.24. *$\mathcal{ALC}$-ontologies are classically satisfiable iff they are satisfiable by a concept faithful geometric model on some (possibly infinite) $\mathbb{R}^n$ using sets of the form $b_1 \times \cdots \times b_n$ with $b_i \in \{\{0\}, \mathbb{R}_+, \mathbb{R}_-, \mathbb{R}\}$ and incidence matrices in $\mathbb{R}^{2n \times 2n}$.*

In order to prove this theorem, we reduce it to the following subproblems. First, concept-faithfulness is proven for an empty tbox with an acyclic abox. After that, it is extended to an arbitrary tbox with an acyclic abox. At the bottom line this leads to the case of an arbitrary tbox and a cyclic abox mentioned in Theorem 8.24 and thus to the conclusion that it is possible to represent arbitrary $\mathcal{ALC}$-ontologies using al-cones.

The algebraic atom that a half-axis represents should be determined by the applicability of role R, thus by exploring the possible paths starting on this half-axis by applying the incidence matrix R. The role matrix is as the geometric model dependent on the srank and iteratively extendable.

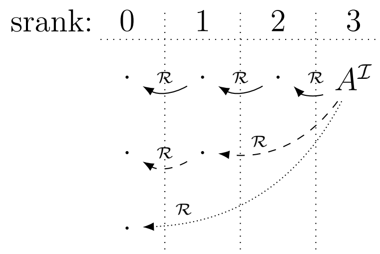

Figure 8.4: Visualization of the assessment of concepts to the al-cone $A^{\mathcal{I}}$. Starting with A, the arrows depict the application of the role R (thus incidence matrix $\mathcal{R}$). The numbers indicate the srank of the resulting concepts, thus, role R can be applied until an srank of 0 is reached.

Example 8.25. *An example can be seen in Figure 8.4. Assume only one role symbol R and no concept symbols are given. There are three different connections starting at the half-axis $A^{\mathcal{I}}$ with R, the dashed, dotted and solid one. Following the possible paths leads to the fact that $A \sqsubseteq \exists R^3.\neg\exists R.\bar{1} \sqcap \exists R^2.\neg\exists R.\bar{1} \sqcap \exists R.\neg\exists R.\bar{1}$ and additionally to the fact that all other roles are not possible, thus $A \sqsubseteq \neg\exists R^4.\bar{1}$, therefore, $A = \exists R^3.\neg\exists R.\bar{1} \sqcap \exists R^2.\neg\exists R.\bar{1} \sqcap \exists R.\neg\exists R.\bar{1} \sqcap \neg\exists R^4.\bar{1}$.*

Additionally, it is necessary to define an underlying propositional model to determine the intersection of a concept with the propositional concepts. This intuition can be represented as incidence matrix.

Definition 8.26. *Given roles $R, S, T, ...$, the* incidence matrices $\mathcal{R}^*, \mathcal{S}^*, ...$ of $R, S, T, \ldots$ *are defined as follows:*

$$\mathcal{R}^* = \begin{pmatrix} \mathcal{R}'_{0\to 0} & \mathcal{R}'_{1\to 0} & \mathcal{R}'_{2\to 0} & \mathcal{R}'_{3\to 0} & \cdots \\ 0 & 0 & \mathcal{R}'_{2\to 1} & \mathcal{R}'_{3\to 1} & \cdots \\ 0 & 0 & 0 & \mathcal{R}'_{3\to 2} & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{8.3}$$

with submatrices $\mathcal{R}'_{i\to j}$, where for all $i, j, k \in \mathbb{N}$ with $j < i$ and $k < j$ there are $m, n, o \in \mathbb{N}$ such that $\mathcal{R}'_{i\to j} \in \mathbb{R}^{m\times n}$ and $\mathcal{R}'_{j\to k} \in \mathbb{R}^{o\times m}$, $\mathcal{R}'_{0\to 0}$ is a zero matrix and the submatrices have a fixed size dependent on the number of roles and concepts given.

Intuitively, the submatrix $\mathcal{R}'_{i\to j}$ depicts the application of role R to find relations between instances of concepts with srank i and srank j and thus, submatrix $\mathcal{R}'_{i\to j}$ influences the region of the geometric model containing algebraic atoms of srank i.

Example 8.27. *Consider an empty tbox, with no concept symbol and one role symbol R. First, it is necessary to determine the number of algebraic atoms added to a i-srank-model when considering an $i+1$-srank-model. Whereas this is straightforward for $i \in \{0,1\}$ with $(\neg\exists R.\bar{1})^{\mathcal{I}}$ for the 0-srank-model and $\exists R.\neg\exists R.\bar{1} \sqcap \neg\exists R^2.\bar{1}$ for the 1-srank-model, there are more algebraic atoms for the 2-srank-model: $\exists R^2.(\neg\exists R.\bar{1}) \sqcap \exists R.(\neg\exists R.\bar{1}) \sqcap \neg\exists R^3.\bar{1}$ and $\exists R^2.(\neg\exists R.\bar{1}) \sqcap \neg\exists R.(\neg\exists R.\bar{1}) \sqcap \neg\exists R^3.\bar{1}$. Based on this, the following incidence matrix $\mathcal{R}^*$ can be created:*

$$\mathcal{R}^* = \begin{pmatrix} \begin{pmatrix}0\end{pmatrix} & \begin{pmatrix}1\end{pmatrix} & \begin{pmatrix}1 & 0\end{pmatrix} & \dots \\ 0 & 0 & \begin{pmatrix}1 & 1\end{pmatrix} & \dots \\ 0 & 0 & 0 \;\; 0 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

where the submatrices (in brackets) in the first row are from left to right $\mathcal{R}'_{0\to 0}, \mathcal{R}'_{1\to 0}$ and $\mathcal{R}'_{2\to 0}$ and in the second row $\mathcal{R}'_{2\to 1}$. So, how was this matrix created and how can it be used for determining the concept of some incidence vector x? In the incidence matrix presented here, the roles up to a srank of two are depicted. Thus, all algebraic atoms mentioned above are representable by the incidence matrix. In the following, it is shown for the half-axes of the geometric model which algebraic atom they depict (where $x^{\mathcal{I}}, \dots, z^{\mathcal{I}}$ represent incidence vectors of half-axes). Considering $w^{\mathcal{I}} = (1\ 0\ \dots)^T$: as $\mathcal{R}^* w^{\mathcal{I}} = \{\vec{0}\}$, there is no role starting in w and thus $w \in \neg\exists R.\bar{1}$. For $x^{\mathcal{I}} = (0\ 1\ 0\ \dots)^T$, $\mathcal{R}^* x^{\mathcal{I}} = (1\ 0\ \dots)^T = w^{\mathcal{I}}$. Thus, for x there is a role R ending in w, such ending in a region where no role is possible anymore. Thus, $x^{\mathcal{I}} \in \exists R.(\neg\exists R.\bar{1}) \sqcap \neg\exists R^2.\bar{1}$. For $y^{\mathcal{I}} = (0\ 0\ 1\ 0\ \dots)^T$, $\mathcal{R}^* y^{\mathcal{I}} = \{(0\ 1\ 0\ \dots)^T = x^{\mathcal{I}}, (1\ 0\ 0\ \dots)^T = w^{\mathcal{I}}\}$. Hence it is, on the one hand, possible to apply the role two times without reaching $\vec{0}$ (from $y^{\mathcal{I}}$ over $x^{\mathcal{I}}$ to $w^{\mathcal{I}}$ and, on the other hand, to directly reach $w^{\mathcal{I}}$, i.e., applying the role only one time. Thus $y^{\mathcal{I}} \in \exists R^2.(\neg\exists R.\bar{1}) \sqcap \exists R.(\neg\exists R.\bar{1}) \sqcap \neg\exists R^3.\bar{1}$. For $z^{\mathcal{I}} = (0\ 0\ 0\ 1\ 0\ \dots)^T$, this is not possible, as $z^{\mathcal{I}}$ is mapped to $x^{\mathcal{I}}$, however not to $w^{\mathcal{I}}$ directly. Thus, $z^{\mathcal{I}} \in \exists R^2.(\neg\exists R.\bar{1}) \sqcap \neg\exists R.(\neg\exists R.\bar{1}) \sqcap \neg\exists R^3.\bar{1}$.

In each dimension k, the two half-axes of the geometric model represent one algebraic atom each. Interpreting half-axes as incidence vectors leads to the case that the positive half-axis in dimension k is represented by a 1 in the incidence vector at position $2k$ and that the negative half-axis in dimension k is represented by a 1 in the incidence vector at position $2k-1$. The algebraic atom can be determined by considering the matrices $\mathcal{R}^*, \mathcal{S}^*, \dots$ at column $2k$. Thus, to represent all concepts of a specific srank i, in each submatrix containing the submatrices $\mathcal{R}'_{i\to j}$, where $j = \{0, \dots, i-1\}$, each possible concept of srank i has to be created. Before stating the main theorem of reaching concept faithfulness, first, the form and creation of srank-algebraic atoms is considered.

Lemma 8.28. *Let $N_R \cup N_C \cup N_c$ be the signature of the ontology under consideration. For a srank of* 0*, the set*

$$\mathcal{X}_0 = \left\{ Z \sqcap \left(\bigsqcap_{R \in N_R} \neg \exists R.\overline{1} \right) \middle| Z \in \mathcal{M}_0 \right\},$$

depicts all algebraic atoms of srank 0*. Here* $\mathcal{M}_0$ *depicts the set of algebraic atoms given based only on the concept symbols.*

For a given srank $i > 0$*, the set of algebraic atoms of srank* i *can be created based on the algebraic atoms of srank* $0, \dots, i-1$

$$\mathcal{X}_i = \left\{ Z \sqcap \left(\bigsqcap_{R \in N_R} \bigsqcap_{B \in \mathcal{X}_{0,\dots,i-1}} \varphi_{R,B} \right) \sqcap \left(\bigsqcap_{R_j \in N_R \text{ for } j=1,\dots,i+1} \neg \exists R_1.(\exists R_2.(\dots \exists R_{i+1}.\overline{1})) \right) \middle| Z \in \mathcal{M}_0, \varphi_{R,B} \in \{\exists R.B, \neg \exists R.B\} \right\},$$

such that for each $A \in \mathcal{X}_i$*,* $A \sqsubseteq \exists R.B$ *for an* $R \in N_R$ *and* $B \in \mathcal{X}_{i-1}$*.*

Proof. Each concept $A \in \mathcal{X}_i$ has srank i, as $A \sqsubseteq \exists R.B$ for a $R \in N_R$ and $B \in \mathcal{X}_{i-1}$. It has to be shown that it is an algebraic atom. If this is not the case then there has to be an algebraic atom A' for which $\overline{0} \neq A' \sqcap A$ and $A' \sqsubset A$. Assume $srank(A') = j > i$, thus, $A' \sqsubseteq C$ where $C = \exists R_1.(\exists R_2.(\dots \exists R_j.\overline{1})$ for $R_1, ..., R_j \in N_R$, thus a chain of j positive existentials must exists. Because of the second part of the definition of $\mathcal{X}_i$, $A \sqcap C = \overline{0}$ and thus $A \sqcap A' = \overline{0}$, a contradiction. Thus assume $srank(A') \leq i$: By definition, each conjunct containing positive existentials up to srank i is either positively or negatively contained in A (because of the first part of the definition of $\mathcal{X}_i$). Thus, it remains to show that there is no conjunct of the type $D = \neg \exists R.E$ of A' for arbitrary $R \in N_R$ and arbitrary concept E such that $A \sqcap D \neq \overline{0}$ and $D \neq A$. Because of the second part of the definition, all D including one negative and then i or more positive existentials are already included. Thus, assume D has one negative and $k < i$ positive existentials (with $k \geq 0$), followed by a negative one and arbitrary existentials afterwards. Then, $D = \neg \exists R.B$, where $B \in \mathcal{X}_{k-1}$ and thus already positively or negatively contained in A, a contradiction. Thus, A is an algebraic atom of srank i. As all combinations of algebraic atoms of a lower srank are considered, $\mathcal{X}_i$ contains all algebraic atoms of srank i. □

To clarify the construction of the algebraic atoms, a small example is given.

Example 8.29. *Consider an ontology with one concept symbol A and one role symbol R and an empty tbox. The algebraic atoms of the Boolean part of the ontology comprise $\mathcal{M}_0 = \{A, \neg A\}$. The construction of $\mathcal{X}_0$ results in $\mathcal{X}_0 = \{A \sqcap \neg\exists R.\bar{1}, \neg A \sqcap \neg\exists R.\bar{1}\}$ Then $\mathcal{X}_1$ is constructed as follows:*

$$\begin{aligned}\mathcal{X}_1 &= \left\{ Z \sqcap \left(\prod\nolimits_{B \in \mathcal{X}_0} \varphi_{R,B} \right) \sqcap \left(\neg\exists R^2.\bar{1} \right) \middle| Z \in \{A, \neg A\}, \varphi_{R,B} \in \{\exists R.B, \neg\exists R.B\} \right\} \\ &= \{Z \sqcap \exists R.(A \sqcap \neg\exists R.\bar{1}) \sqcap \exists R.(\neg A \sqcap \neg\exists R.\bar{1}) \sqcap \neg\exists R^2.\bar{1}, \\ &\quad Z \sqcap \exists R.(A \sqcap \neg\exists R.\bar{1}) \sqcap \neg\exists R.(\neg A \sqcap \neg\exists R.\bar{1}) \sqcap \neg\exists R^2.\bar{1}, \\ &\quad Z \sqcap \neg\exists R.(A \sqcap \neg\exists R.\bar{1}) \sqcap \exists R.(\neg A \sqcap \neg\exists R.\bar{1}) \sqcap \neg\exists R^2.\bar{1} \mid Z \in \{A, \neg A\}\}\end{aligned}$$

The second part, denoting the negative existentials, can be interpreted as all concepts that are not accessible through the role and, thus, not accessible through the incidence matrix.

Algebraic atoms are restricted so that there must be at least one $B \in \mathcal{X}_{i-1}$ of $srank(B) = i - 1$ with a positive existential due to the fact that algebraic atoms of srank i are considered. Without this restriction, $\mathcal{X}_i$ could contain algebraic atoms of a lower srank.

Based on Lemma 8.28 it is possible to prove the next proposition that states that interpreting the ontology based on the incidence matrices $\mathcal{R}^*$ is a suitable concept-faithful interpretation. The idea is to have an underlying geometric model which only contains information about propositional concepts and having the incidence matrices $\mathcal{R}^*, \mathcal{S}^*, \ldots$ to define the algebraic atoms based on roles.

Proposition 8.30. *Let $\mathcal{O}$ be an ontology with an empty tbox. Let $\mathcal{I}$ be a geometric interpretation of $\mathcal{O}$ constructed based on an interpretation $\mathcal{M}$ which is $\mathcal{M} = \mathcal{M}_0 \times \mathcal{M}_0 \times \ldots$, thus an infinite direct product of one 0-(quantifier)-rank-concept-faithful geometric interpretation $\mathcal{M}_0$. Furthermore, let each role R be interpreted as $\mathcal{R}^*$ as in Definition 8.26 with $\mathcal{R}'_{0\to 0}$ in $\mathbb{R}^{m \times m}$ and $m = 2|\mathcal{M}_0|$. Then: It is possible to construct $\mathcal{R}^*$ for each R in a way that $\mathcal{I}$ is concept-faithful (w.r.t. $\mathcal{O}$), and for each constants a, b and role R, if $\mathcal{O} \models R(a, b)$, then there is a representation $b^{\mathcal{I}}$ and $a^{\mathcal{I}}$ so that $b^{\mathcal{I}} \in \mathcal{R}^* a^{\mathcal{I}}$, if $R(a, b)$ is not part of a circular relationship.*

In the proposition above we rely on the usual "direct product" operation on structures that is known in model theory and that can be applied to a possibly infinite number of input structures [Chang and Keisler, 1990, p. 224]: the domain is the Cartesian product of the domains of the input structures and the interpretations of the non-logical symbols are given component-wise.

Proof. First, it is shown that the construction is consistent, meaning no contradictions are induced. The construction of the 0-rank-faithful model can be done as shown in Proposition 5.19 and is consistent, as it consists of Boolean $\mathcal{ALC}$. As $\mathcal{R}^*$ only influences positive existential quantification, negation of an existential is defined via polarity. It is by definition not possible that a concept and its negation intersect. As shown in Proposition 8.22, the role operator results in an al-cone and fulfills $(\exists R.\overline{0})^\mathcal{I} = \overline{0}^\mathcal{I}$ and $\mathcal{R}^T(A^\mathcal{I} \sqcup B^\mathcal{I}) = conH(\mathcal{R}^T A^\mathcal{I} \cup \mathcal{R}^T B^\mathcal{I})$ for concept symbols A, B. To show concept-faithfulness, it is sufficient to show satisfiability, as the construction principle of Proposition 5.19 can be used similarly as in the Boolean case to construct a concept-faithful model. Thus, it is shown that all algebraic atoms induced by the tbox can be represented. Therefore, in the following, only half-axes (as representatives of algebraic atoms) are considered.

The proof is done based on induction over the srank. Therefore, it is shown that each algebraic atom of a specific srank can be represented and therefore, for the geometric interpretation of an arbitrary srank, faithfulness is reached. This is done by showing that the construction of $\mathcal{R}^*$, $\mathcal{S}^*, \dots$ exactly leads to the atomic concepts represented in Lemma 8.28. First, it is shown that the subspace $\mathbb{R}^m$ of the first m dimensions of the geometric model is a 0-srank-model and incorporates all algebraic atoms of srank 0. The subspace contains $\mathcal{M}_0$, which is a 0-(quantifier)-rank-faithful geometric model. As $\mathcal{R}'_{0\to 0}$ equals zero for all roles R, each application of the role operator (in form of $\mathcal{R}^* d$ for a half-axis d in $\mathbb{R}^m$) results in $\{\vec{0}\}$, thus is not possible. Thus, it exactly matches the definition of Lemma 8.28. Next, assume that the underlying geometric model combined with the submatrix including all submatrices $\mathcal{R}'$ up to $\mathcal{R}'_{i-1\to i-2}$ represents all atomic concepts up to srank $i-1$. It is shown that the submatrices $\mathcal{R}'_{i\to j}$ for $j \in \{0, ..., i-1\}$ for each role R represent in combination with the geometric model all algebraic atoms of srank i and do not lead to any inconsistencies. The submatrix

$$\mathcal{R}^*_i = \begin{pmatrix} \mathcal{R}'_{i\to 0} \\ \mathcal{R}'_{i\to 1} \\ \vdots \\ \mathcal{R}'_{i\to i-1} \\ 0 \\ \vdots \end{pmatrix} \tag{8.4}$$

for role R combined with the respective submatrices for the other roles represents in each column one algebraic atom in form of a mapping to lower sranks. Consider, e.g., the incidence vector x' of an algebraic atom X which contains only one non-zero element. $\mathcal{R}^*x'$ can thus be reduced to a $r \cdot 1 = \mathcal{R}^*x'$ where r is one column of $\mathcal{R}^*$.

The exact content of each column r is not considered here, it is adequate to show that each possible algebraic atom could be represented in such a column as each column represents independent of the others one possible half-axis (algebraic atom). $\mathcal{X}_i$ as defined in Lemma 8.28 is considered. The first conjunct of each algebraic atom, the Boolean concept A, is contained, as the geometric model consists of an infinite concatenation of the 0-rank geometric models. Therefore, also in the area influenced by the submatrix considered, there is an intersection with each of the Boolean concepts A possible. Next, it has to be shown that the submatrix $\mathcal{R}'_{i \to j}$ maps an algebraic atom of srank i really to one (ore more) algebraic atoms of srank j and only to them. This is satisfied by the construction principle of the matrices mentioned in Definition 8.26. It is not possible that the algebraic atom has a higher srank, as the space under the diagonal of $\mathcal{R}^*$ is not populated and therefore, only a reduction of the rank is possible and thus $(\bigsqcap_{R \in N_R} \bigsqcap_{C \in \mathcal{X}_{i+1,i+2,\ldots}} \neg \exists R.C)$ is fulfilled. As $\mathcal{R}^*_i$ can contain arbitrary many ones in a column, it is possible to model an existential for arbitrary $B \in \mathcal{X}_{0,\ldots,i-1}$ and each combination of roles, as $\mathcal{S}^*, \mathcal{T}^*, \ldots$ also can influence the column considered. A column where all role-matrices $\mathcal{R}'_{i \to i-1}$ have only zero entries is not allowed as it would interfere with the restriction of Lemma 8.28 for having at least one role to an algebraic atom with srank $i-1$.

When each column-tuple of $\mathcal{R}^*, \mathcal{S}^*, \ldots$ is unique, then each algebraic atom is on a unique half-axis. Therefore, for non-cyclic aboxes, instances $a^\mathcal{I}$ and $b^\mathcal{I}$ have a unique al-cone each. Thus, when $R(a,b)$ is valid, then $b^\mathcal{I} \in \mathcal{R}^*(a^\mathcal{I})$. □

Handling the Cyclic Case

Now we proceed with the problem of handling cyclicity. We observe that cyclicity cannot be represented correctly in the construction above. Assume $R(a,b)$ and $R(b,a)$ is given in the abox. Then, applying $\mathcal{R}$ on a concept A with $a \in A$ would not reduce its srank, as afterwards it is still possible to apply $\mathcal{R}$ infinitely many times. Thus, it is necessary to represent concepts with an infinite srank.

The set of possible algebraic atoms has been determined in Lemma 8.28. The incidence matrix $\mathcal{R}^*$ is created based on the idea of iteratively increasing the srank and modeling all possible atomic concepts of this srank. Now, it is necessary to consider concepts $X \in \mathcal{X}_\infty$, thus, circular relationships. These concepts are not considered in such an iterative approach. To use the same construction principle of iterative extension it is necessary to define a new notion of rank which is a combination of srank and the cycle depth of the concepts. This rank then enables for iterative extension.

This is done based on the idea that each application of a role is either part of a cycle or not. For each non-cyclic role, the srank can be determined. Out of this, the maximum is chosen. For non-circular relationships, the srank does not change. The motivation behind this is that having a circular relationship, it is also possible to have non-cyclic behavior in parts, e.g., Narcissus could love himself (a cycle) but could also love a person not loving anyone (a non-cyclic behavior). This has to be determined for each part of the cycle, as in the incidence matrix extended for circular relationships $\mathcal{R}^{**}$, which will be considered in more detail in the proof of Proposition 8.34 below, it is necessary to place the concept representing a cycle at a srank where a mapping of the non-cyclic parts of the concept is possible. On the other hand, the maximum depth of a cycle can be considered. A conjunct with a srank of infinity needs to contain a cycle. This means that at some point, applying role R leads to an algebraic atom which has been visited before. We introduce *cycle_depth* to represent the depth of this cycle. Regarding the Narcissus example, cycle depth would be one.

Definition 8.31. *The* cyclic semantic rank srank$_c$ ($srank'$,$cycle\_depth$) *of a concept C given by its defining formula is a pair determined as follows:*

- *If $C = \overline{0}$, then $srank'(C) = srank(C) = \infty$ and cycle_depth undefined.*
- *Let C_t be the (possibly infinite) computation tree that would unfold when computing $srank(C)$, i.e., $srank(C) = srank(\mathrm{root}(C_t))$.*
- *If $srank(C) = r < \infty$, then $srank_c = (r, 0)$;*
- *else transform C_t to C_t' by copying C_t and replacing any node $N \in C_t$ if a node representing the same concept term as N occurs in the subtree rooted at N. N is replaced by a new concept symbol N', which yields $srank(N') = 0$ according to Definition 8.15; the subtree rooted at N is removed. This makes $srank' = srank(C_t')$ evaluate to $r < \infty$. Let cycle_depth then be the minimum number of occurrences of $\exists$ between re-appearances of N in the subtree in C_t, maximized over all nodes N that were replaced.*

Intuitively, the definition extends *srank* to infinite trees by cutting off re-appearing nodes N and recording the maximum length of cycles cut off.

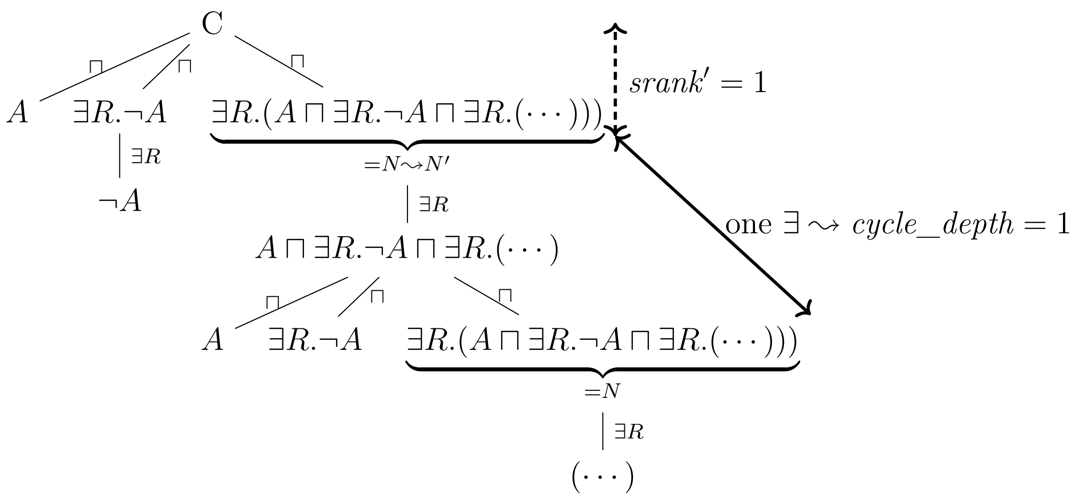

Figure 8.5: Computation tree C_t of the concept $C = A \sqcap \exists R.\neg A \sqcap \exists R.(A \sqcap \exists R.\neg A \sqcap \exists R.(\dots))$.

Example 8.32. *Consider again an empty tbox and an ontology with a role symbol R and a concept symbol A as arising in Example 8.12 (narcissus-vain). Consider $srank_c$ of concept C given by the infinite term $C = A \sqcap \exists R.\neg A \sqcap \exists R.(A \sqcap \exists R.\neg A \sqcap \exists R.(\dots))$. First, the computation tree C_t as shown in Figure 8.5 is created. The calculation of the srank leads to $srank(C) = \infty$, as a cycle is included. As marked in the figure with N, the node $\exists R.(A \sqcap \exists R.\neg A \sqcap \exists R.(\dots))$ occurs several times and its first occurrence is thus replaced with a new symbol N' and is given a $srank(N') = 0$. Since $srank(\exists R.\neg A) = 1$, w have $srank'(C) = 1$ and $cycle\_depth(C) = 1$, as both occurrences of N are connected via one existential role quantification, thus every step the origin of the cycle is reached again. Thus, $srank_c(C) = (1, 1)$.*

It is, analogously to Lemma 8.28 for the non-cyclic case, possible to define a srank$_c$-dependent construction principle for algebraic atoms.

Lemma 8.33. *Let $N_R \cup N_C \cup N_c$ be the signature of the ontology under consideration. Then $\mathcal{X}_{i,0}$ as defined below depicts all algebraic atoms of srank$_c$(i,0), and $\mathcal{X}_{0,1}$, which is* *defined by coinduction [Rutten, 2005] below, depicts all algebraic atoms of* $srank_c(0,1)$*.*

$$\mathcal{X}_{i,0} = \mathcal{X}_i$$

$$\mathcal{X}_{0,1} = \left\{ X^k_{0,1} \;\middle|\; \begin{array}{l} \text{There is } Z \in \mathcal{M}_0, \psi_{R,X^k_{0,1}} \in \{\exists R.(X^k_{0,1})', \neg\exists R.(X^k_{0,1})'\} \text{ such that} \\ X^k_{0,1} = Z \sqcap \left(\bigsqcap_{R\in N_R} \bigsqcap_{B\in \mathcal{X}_0,\mathcal{X}_1,\ldots} \neg\exists R.B \right) \\ \sqcap \left(\bigsqcap_{R\in N_R} \psi_{R,(X^k_{0,1})'} \right) \sqcap \left(\bigsqcap_{R\in N_R} \neg\exists R.\neg(X^k_{0,1})' \right) \end{array} \right\}$$

where $\mathcal{M}_0$ *depicts the set of algebraic atoms given based only on the concept symbols and for each* $X^k_{0,1}$ *at least one* $\psi^k_{R,X_{0,1}}$ *must appear positive.* $(X^k_{0,1})'$ *is the derivative of* $X^k_{0,1}$ *required for the coinductive definition.*

Proof. Having a cycle-depth of zero, the construction reduces to the construction of $\mathcal{X}_i$, depicted in Lemma 8.28, thus $\mathcal{X}_{i,0} = \mathcal{X}_i$, as no cycles are contained. Now, consider $X^k_{0,1} \in \mathcal{X}_{0,1}$. $srank'(X^k_{0,1}) = 0$, as all concepts $B \in \mathcal{X}_0, \mathcal{X}_1, \ldots$ are only reachable by negated roles. The cycle-depth of $X^k_{0,1}$ is 1, as at least one $\psi^k_{R,X_{0,1}}$ must appear positive, thus a cycle is included. It remains to show that $X^k_{0,1}$ actually is an algebraic atom. Assume this is not the case, thus an algebraic atom $A \sqsubset X^k_{0,1}$ must exists such that $X^k_{0,1} \sqcap A \neq \overline{0}$. First, assume $srank'(A) > 0$. A contradiction, as all concepts of a higher $srank'$ appear negated. Second, assume $srank_c(A) = (0,0)$. This cannot be the case, as all existentials not part of a circular relationship are negated and at least one cycle appears positive in $X^k_{0,1}$ and would interfere with the assumption of a cycle depth of 0. Third, assume $srank_c(A) = (0,1)$. However, as in $X^k_{0,1}$ all concepts with a cycle-depth greater 1 are negated and each possible cycle of length 1 either appears positive or negative, this is not possible. Fourth, assume $srank_c(A) = (i,j)$ with $i \geq 0, j > 1$, thus there must exists a cycle of length at least two. However, as $X^k_{0,1}$ contains a conjunction with $\bigsqcap_{R\in N_R} \neg\exists R.\neg X^k_{0,1}$, it is not possible to have a cycle not ending in $X^k_{0,1}$ thus it is not possible to have a cycle depth greater one. Thus, $X^k_{0,1}$ is an algebraic atom with $srank_c(X^k_{0,1}) = (0,1)$. □

The sets of algebraic atoms $\mathcal{X}_{i,j}$ can be derived based on the above two but are omitted here for reason of readability.

Now, based on srank_c, an iterative creation of the geometric model is possible. For each of these tuples, a submatrix, thus, a specific region in the geometric model can be created to depict the concepts having this configuration. The matrix where all submatrices representing a cycle depth greater than zero (thus containing a cycle) are set to zero contains exactly the submatrices contained in $\mathcal{R}^*$ and has its behavior (except including more zero matrices). The basic idea for the submatrix $\mathcal{R}'_{i,j\to j}$ depicting a specific srank$_c$ i and a cycle-depth j is to place it at position $\mathcal{R}^{**}_{k...l,k...l}$ for some k and l in the new incidence matrix $\mathcal{R}^{**}$, thus, to enable to model connections trough roles between arbitrary columns of the matrix, thus, creating arbitrary circles. In the same column, it is also possible to place some $\mathcal{R}'_{i,j\to h}$ for $h < j$ and j is the srank' of srank$_c$ i.

Thus, Proposition 8.30 can be extended to a proposition covering cyclic aboxes.

Proposition 8.34. *Let $\mathcal{O}$ be an ontology with an empty tbox. Let $\mathcal{I}$ be a geometric interpretation of $\mathcal{O}$ constructed based on an interpretation $\mathcal{M}$ which is $\mathcal{M} = \mathcal{M}_0 \times \mathcal{M}_0 \times \ldots$, thus an infinite direct product of one 0-(quantifier)-rank-concept-faithful geometric interpretation $\mathcal{M}_0$. Let further $\mathcal{I}$ interpret each role R as some incidence matrix. Then: $\mathcal{I}$ is concept faithful (w.r.t. $\mathcal{O}$), and for each constants a, b and role R, if $\mathcal{O} \models R(a,b)$, then $b^{\mathcal{I}} \in \mathcal{R}^{**}a^{\mathcal{I}}$.*

Proof. In the following, it is shown how concepts having a cycle depth greater zero and a srank' of zero and concepts having a srank' greater zero and a cycle-depth of one can be modeled in an incidence matrix $\mathcal{R}^{**}$. The proof of the other concepts follows analogously.

Assume a cycle depth of i and an srank' of zero. It is shown that a submatrix $\mathcal{R}'_{i,j\to j}$ of $\mathcal{R}^{**}_{k...l,k...l}$ can be chosen which models these dependencies. As the *srank'* is zero, all other elements in columns $k, \ldots, l$ are zero. Thus, all possible combinations of conjunctions of zero *srank'* concepts with roles have to be considered. These are only finitely many and could be represented in this submatrix.

The incidence matrix when considering only elements having a cycle depth of zero or one has the following form:

$$\mathcal{R}^{**} = \begin{pmatrix} \mathcal{R}'_{0,0\to 0} & 0 & \mathcal{R}'_{2,1\to 0} & \mathcal{R}'_{3,1\to 0} & \cdots \\ 0 & \mathcal{R}'_{1,1\to 1} & \mathcal{R}'_{2,1,1\to 0} & 0 & \cdots \\ 0 & 0 & \mathcal{R}'_{2,1\to 1} & 0 & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \tag{8.5}$$

where the first column with submatrices depicts srank$_c$ of $(0,0)$, the second column with submatrices the srank$_c$ $(0,1)$, the third $(1,1)$ and the fourth $(1,0)$. $\mathcal{R}^{**}$ contains the submatrices of $\mathcal{R}^*$ as submatrices, as there are concepts without any circular relationshipy possible which are represented by interaction with the submatrices of $\mathcal{R}^*$. The circular relationship is represented by the additional submatrices $\mathcal{R}'_{j,i\to i}$ which are on the diagonal and enable to cover arbitrary connections between algebraic atoms. To model cycles of length two, we have to ensure that when the submatrix is one at position $\{k,l\}$ it needs to be one at position $\{l,k\}$. As for each srank_c there are only finitely many combinations, it is possible to model the incidence matrix iteratively in this way. Extension to more than one role and a higher srank_c is achieved analogously. □

Non-Empty Tboxes

As in general it cannot be assumed and is not suitable to have an empty tbox, in the following it is shown how the geometric model and the incidence matrix can be restricted to fulfill the tbox axioms.

The same iterative modeling approach from above is going to be used here. But now, the information of each tbox-axiom has to be incorporated, even if the model is restricted to an srank smaller than the srank of a part of the axiom. Therefore, it is necessary to know the relevance of specific tbox-axioms to the srank-models. It is not possible to use the definition of a srank for the empty tbox without change, as, e.g., having the axiom $A \sqsubseteq \exists R.C$, then the srank would be $srank(A) = 0$, $srank(\exists R.C) = 1$. However, A is known to be a subconcept of $\exists R.C$, this means, each for each $a \in A$, there must be a $b \in C$ with $R(a,b)$. Therefore, it is not possible to place A in a 0-srank-model, as then, it would be possible that an element in A does not have any role, what conflicts with the axiom. Having a concept, e.g., $\exists R.C$ with $srank(\exists R.C) = 3$, this influences also concepts based on the srank of this concept, e.g., $\exists R^2.C$, thus, this concept would have $srank(\exists R^2.C) = 4$, as applying a role increases the rank by one.

To model this, an extension of the semantic rank accounting for tboxes is defined.

Definition 8.35. *The* tbox-specific semantic rank $srank_{\mathcal{T}}$ (tbox-srank for short) *is defined as follows:* $srank_{\mathcal{T}}(C) = srank(C) = \infty$ *if* C *is part of a circular relationship (see Section 2.1 for details) and else is defined based on the same defining rules as for the srank (see Definition 8.15), with one added rule:*

$$srank_{\mathcal{T}}(C) = max(\{srank(C)\} \cup \{srank(D) \mid \mathcal{T} \vDash C \sqsubseteq D\})$$

where the other rules are changed to

$$\begin{aligned}
srank(C \sqcap D) &= \max(\{srank_{\mathcal{T}}(C), srank_{\mathcal{T}}(D)\})\\
srank(C \sqcup D) &= \min(\{srank_{\mathcal{T}}(C), srank_{\mathcal{T}}(D)\})\\
srank(\neg(C \sqcap D)) &= \min(\{srank_{\mathcal{T}}(\neg C), srank_{\mathcal{T}}(\neg D)\})\\
srank(\neg(C \sqcup D)) &= \max(\{srank_{\mathcal{T}}(\neg C), srank_{\mathcal{T}}(\neg D)\})
\end{aligned}$$

We extend the arithmetics for ∞ in the usual way, setting $\infty + 1 = \infty$.

Thus, the calculation of $\text{srank}_\mathcal{T}$ is influenced on the one hand by the srank of the concept but on the other hand on its subsumption-relation to other concepts.

Lemma 8.36. *If a tbox contains a concept $\neg\exists R.C$ for arbitrary role R and concept C and $srank_\mathcal{T}(\neg\exists R.C) > 0$ or contains a concept D with $srank_\mathcal{T}(D) = \infty$, then the tbox contains a circular relationship.*

Proof. The lemma follows trivially for infinite srank. Assume a concept $\neg\exists R.C$ for arbitrary R and a propositional concept C and a tbox-axiom $\neg\exists R.C \sqsubseteq \exists S^i.D$ for arbitrary S, i and propositional concept D. Then, we get $srank_\mathcal{T}(\neg\exists R.C) = i$. $\exists R.C$ has a $srank_\mathcal{T}(\exists R.C) = 1$, as C is a propositional concept. Thus, one disjunct of $\exists R.C$ is $\exists R.(C \sqcap \neg\exists R.\overline{1} \sqcap \neg\exists S.\overline{1})$. This is a contradiction, as $\neg\exists R.\overline{1} \sqcap \neg\exists S.\overline{1} = \overline{0}$. Therefore, an extension is necessary and thus $\exists R.C = \exists R.(C \sqcap \neg\exists R.\overline{1} \sqcap \exists S^i.D)$. This can be done in the same way for D and thus leads to an infinite extension of the concept and thus to an infinite $srank_\mathcal{T}$. Therefore, a circular relationship exists. This can trivially be adapted to general axioms. □

Example 8.37 (Example 8.16 continued)**.** *Consider a tbox with $\exists R.C \equiv \exists R^2.C$ (and therefore $\neg\exists R.\exists R.C \equiv \neg\exists R.C$). Based on the sranks calculated above*

$$srank(\exists R.C) = 1, srank(\neg\exists R.C) = 0, srank(\exists R^2.C) = 2, srank(\neg\exists R.\exists R.C) = 0,$$

it is possible to calculate the tbox-sranks:

$$srank_\mathcal{T}(\neg\exists R.\exists R.C) = \max(\{srank(\neg\exists R.\exists R.C), srank(\neg\exists R.C)\}) = 0$$

For the first axiom, the calculation is more complex:

$$\begin{aligned} srank_\mathcal{T}(\exists R.C) &= \max(\{srank(\exists R.C)\} \cup \{srank(\exists R^2.C)\}) \\ &= \max(\{1\}, \{srank_\mathcal{T}(\exists R.C) + 1\}) \end{aligned}$$

Therefore, $srank_\mathcal{T}(\exists R.C) = srank_\mathcal{T}(\exists R^2.C) = \infty$.

We first focus on acyclic tboxes and aboxes, this means that each concept has a finite tbox-srank and each concept of the form $\neg\exists R.C$ for arbitrary R, C has an tbox-srank of 0. Then, the creation of a geometric model can be done as described in Definition 8.26 and Proposition 8.22, except that the axioms have to be considered. Thus, the first appearance of a concept is at its tbox-srank, before that, it appears only negative. Therefore, the 0-quantifier-rank concept-faithful model for the representation of the

0-srank-model is only created with the propositional concepts which have a tbox-srank of 0 and is extended for each tbox-srank with the propositional concepts having this rank. Submatrices $R'_{i \to j}$ are not allowed to perform any mappings that contradict the axioms.

Definition 8.38. *A geometric interpretation for an arbitrary non-cyclic tbox and non-cyclic abox is given as*

- *an incidence matrix $\mathcal{R}^*$ as defined in Definition 8.26 for each role R;*
- *each $\mathcal{R}'_{i \to 0,\dots,i-1}$ is based on the part of the geometric model representing a product of arbitrary many 0-quantifier-rank geometric models of the propositional concepts having a tbox-srank of at least i.*

Proposition 8.39. *Let $\mathcal{O}$ be an ontology with an arbitrary non-cyclic tbox and non-cyclic abox. Let $\mathcal{I}$ be a geometric interpretation of $\mathcal{O}$ as defined in Definition 8.38. Then: $\mathcal{I}$ is concept-faithful (w.r.t. $\mathcal{O}$), and for each constants a, b and role R, if $\mathcal{O} \models R(a,b)$, then there is a representation $b^{\mathcal{I}}$ and $a^{\mathcal{I}}$ so that $b^{\mathcal{I}} \in \mathcal{R}^* a^{\mathcal{I}}$, if $R(a,b)$ is not part of a circular relationship.*

Proof. As the tbox is acyclic, for each concept C, either $srank_{\mathcal{T}}(C) = 0$ or $srank_{\mathcal{T}}(\neg C) = 0$. Thus, it is possible to model a k-srank-model for arbitrary $k \geq 0$ without having both C and $\neg C$ to be contradictory. As a concept C is only different from $\overline{0}^{\mathcal{I}}$ in a k-$srank_{\mathcal{T}}$-model, if $srank_{\mathcal{T}}(C) \leq k$, it has to be ensured that a concept does not appear at a lower srank and that only non-contradictory algebraic atoms are modeled. Non-contradiction is ensured for propositional concepts because of Definition 8.38.

As proven in Proposition 8.30, each algebraic atom can be represented. Based on this, only columns are considered which represent the desired algebraic atoms and thus, the restrictions are satisfied. □

This leads to the proof of Theorem 8.24 stating that it is possible to model a faithful geometric model for ontologies over full $\mathcal{ALC}$.

Proof of Theorem 8.24. Considering tboxes with finite ranks, the proof follows directly from Proposition 8.39.

Now consider tboxes with infinite ranks, e.g., a tbox containing the axiom $A = \exists R.A$. Then $srank_{\mathcal{T}}(A) = \infty$ and $srank_{\mathcal{T}}(\neg A) = \infty$. Or consider a tbox containing a concept $\neg\exists R.C$ with tbox-srank greater than 0, as described in Lemma 8.36. Therefore, it is necessary to consider the construction introduced in Proposition 8.34. There, it is allowed to model cycles in the abox, and therefore, infinite sranks. Based on this, the construction principle depicted in Propositions 8.34 and 8.39 can be used. □

The approach presented above enables us to create a faithful geometric model of a given ontology. It opens up the possibility to restrict the model to a given srank without affecting the expressivity of the given concepts. However, even for a restricted rank, the faithful geometric model could grow exponentially (depending on the tbox) and is therefore possibly not practical because of its size. Therefore, on the one hand, it is possible to extend the tbox with axioms which model a known bias in the data to circumvent that this bias is learned. This incorporates helpful information and reduces the size of the model. On the other hand, it is possible to focus on specific subparts of the model and model them faithfully. Thus, being able to model an ontology faithfully is not only helpful when full faithfulness is needed but also helpful as it is then known that each desired subproblem can be modeled correctly.

Example 8.40. *Consider again the example of Narcissus first mentioned in Example 8.12. Assume an empty tbox, one concept* Vain *and one role* loves*, for short V, R.*

It is possible to create a faithful representation, yet quite complex and of infinite dimension. Therefore, it is necessary to make suitable restrictions on the faithfulness. Here, the focus lies on reflexivity of the role loves regarding Narcissus. *Therefore, reflexivity has to be modeled faithfully. This still results in an infinite model, however, of a simpler structure.*

One example for modeling is

$$\begin{aligned} V^{\mathcal{I}} &= \mathbb{R}_+ \times \mathbb{R}_+ \times \dots \\ (\neg V)^{\mathcal{I}} &= \mathbb{R}_- \times \mathbb{R}_- \times \dots . \end{aligned}$$

Then the role R can be modeled as

$$\mathcal{R} = \begin{pmatrix} 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & \dots \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 1 & \dots \\ & & 1 & 0 & & & & & & & \\ & & 0 & 1 & & & & & & & \\ & & 0 & 0 & & & & & & & \\ & & \vdots & \vdots & & & & & & & \end{pmatrix},$$

where the empty regions contain zeros.

The geometric model is based on only one concept. Thus, each odd column of $\mathcal{R}$ represents a conjunction with $V^{\mathcal{I}}$, each even column a conjunction with $(\neg V)^{\mathcal{I}}$. The first two columns represent the area of the geometric model where no role is possible. Therefore, a person Charlie represented by the incidence vector $c^{\mathcal{I}} = (1\,0\,0\,\dots)^T$ would be vain (as $c^{\mathcal{I}} \in V^{\mathcal{I}}$) but would not love any person (as $\mathcal{R}c^{\mathcal{I}} = \{\vec{0}\} = \vec{0}^{\mathcal{I}}$).

The fifth to the tenth column represent the different algebraic atoms of srank 1. *A person Bob represented by* $b^{\mathcal{I}} = (0\,0\,0\,0\,1\,0\,\ldots)^T$ *is vain, as* $b^{\mathcal{I}} \in V^{\mathcal{I}}$. $\mathcal{R}b^{\mathcal{I}}$ *leads to the incidence vectors* $(1\,0\,\ldots)^T$ *and* $(0\,1\,0\,\ldots)^T$*, thus,* $b \in \exists R.(V \sqcap \neg\exists R.\overline{1}) \sqcap \exists R.(\neg V \sqcap \neg\exists R.\overline{1})$. *Thus, Bob is vain and loves one person being vain and one person not being vain, whereas both the beloved persons do not love anyone.*

The third and fourth column represent the reflexivity and thus the property relevant for expressing the narcissism of Narcissus. As Narcissus is vain, $n^{\mathcal{I}}$ *has to be modeled in an odd dimension of the model. Thus,* $n^{\mathcal{I}} = (0\,0\,1\,0\,\ldots)^T$. $\mathcal{R}n^{\mathcal{I}}$ *leads to* $(0\,0\,1\,0\,\ldots)^T$*, thus it is mapped onto itself. Therefore,* $n \in \exists R.(V \sqcap \exists R.(V \sqcap \ldots))$ *and reflexivity is modeled. Additionally, it leads to* $(1\,0\,\ldots)^T$ *and* $n \in \exists R.(V \sqcap \neg\exists R.\overline{1})$. *Thus, Narcissus loves himself and he loves a vain person which does not love anyone.*

9 Conceptual Orthospaces in Practice

I proposed COSs as answer to the question stated in the introduction: what does an expressive framework modeling background knowledge have to look like? COSs are a framework enabling to define a non-classical interpretation of an ontology. However, COSs do not implement any concrete embedding approach, but are a foundation of it. Therefore, in this chapter the bracket opened up in the introduction is closed by illustrating the use of COSs.

9.1 Conceptual Orthospaces as Basis for Embedding Approaches

Assume some dataset is given, containing instances and their (possible multiple) labels/classes. Assume additionally that some background knowledge in form of an ontology is given. If roles are contained, than this dataset could be, e.g., a knowledge graph.

The first distinction must be made whether dissimilarity or betweenness should be used as starting point. The fact that similarity information is given does not directly qualify for defining an orthonegation-induced betweenness. As discussed in Section 6.1.3, many similarity relations are not suitable for the creation of an expressive embedding. It could be the case that even with a given similarity relation it is more useful to choose a betweenness as basis for training an embedding.

It is also necessary to think about desired restrictions of the COS, e.g., the strength of the connection between similarity and betweenness as argued in Section 5.2. Based on these considerations, a COS can be created with the help of the construction principles of orthonegation-induced betweenness (see Proposition 6.14 or Proposition 6.15), with the help of cones in case of Euclidean betweenness (see Chapter 7) or by creating a COS from scratch.

However, having a COS is only the first step towards a learning approach. A concrete, given COS opens up many possibilities to actually model a given ontology. In particular, for a COS Y is abstractly defined and it is necessary to choose relevant sets of Y to model actual concepts resp. the actual Y needs to be constructed. It is also necessary to define the amount of faithfulness needed. Whereas a strongly faithful model allows for modeling incomplete information correctly, it hinders inferring new labels for given instances and thus increases overfitting. Additionally, it is possible that a faithful model leads to a very high dimension needed for the representation. This is, of course, dependent on the COS used but is, e.g., a problem for a COS based on convex cones. On the other hand, a too severe reduction of the faithfulness is also not helpful, as it leads to a reduction of the background knowledge information used and prevents from modeling partial information. Next to the faithfulness of concepts, it is also necessary to consider faithfulness of roles. The first question is whether roles are contained at all. If they are contained, then a COS needs to be enhanced with roles, e.g., by, but not limited to, using reification as introduced in Chapter 8.

The next question concerns the inference service used. The exact aim of the approach needs to be determined, e.g., link prediction. This also influences the determination of Y: is it the case that Y is exactly determined, thus the concept representation is already given and the embedding is solely learned for the instances or is it necessary to learn both, the concept and the instance representation? In the following, Section 9.2 exemplifies the former and Example 9.1 the latter case. A third case is that the instances are already given in the domain X and only the concept representation needs to be learned.

In the rest of this chapter, the procedure of creating an embedding approach based on COSs is exemplarily depicted based on two different use cases. On the one hand, in Section 9.2 an approach based on zero-shot learning with an al-cone-COS based on my joint work with Özgür Özçep and Diedrich Wolter [Leemhuis, Özçep, and Wolter, 2022b] is presented and, on the other hand, in Example 9.1 the quantum embedding [Garg et al., 2019], situated in the area of KGE is presented which, though it was invented independently, can be considered as being a special case of a COS.

Example 9.1. *It is not only possible to create new approaches based on COSs, it is also possible to underline the importance of COSs for KGE by showing that existing approaches can be interpreted following the basic ideas of COSs. This is in particular the case for the Quantum Embeddings of Garg et al. [2019].*

The aim of Garg et al. [2019] is to do both link prediction and property testing, thus examining whether an entity belongs to a concept. The idea of the Quantum Embeddings is based on quantum logic [Birkhoff and Neumann, 1936] and each unary and binary relation (thus concepts and roles) is represented as linear subspace of a vector space. Conjunction of concepts is represented as set-intersection, disjunction as convex hull. The linear subspaces are restricted to axis-parallel ones to enforce the resulting lattice (considering conjunction as lattice meet and disjunction as lattice join) having the property of distributivity (though, the lattice of linear subspaces is orthomodular, *here it is distributive, as only axis-parallel linear subspaces are considered). Negation is defined via the polarity operator. Instances are mapped to unit-length-vectors in* $\mathbb{R}^d$. *The embedding space can be considered as a restricted form of an al-cone-COS. Instead of* Y *containing al-cones in general,* Y *is restricted to axis-parallel linear subspaces of* $\mathbb{R}^n$. *Roles are represented as axis-parallel subspaces of* $\mathbb{R}^{2d}$. *Then* $R(a, b)$ *is mapped to the vector* $(x_a\, x_b)^T$ *where* $x_a, x_b \in \mathbb{R}^d$ *represent* a, b *resp. and* $(x_a\, x_b)^T \in S_R$*, where* S_R *is the linear subspace representing* R. *To reduce the computational complexity, each axis-parallel linear subspace is represented via a so-called indicator-vector* $y_C \in \mathbb{R}^d$ *for a concept* C *which is* 1 *in dimension* i *if the* ith *axis is part of the subspace representing* C. *These indicator vectors then allow for modeling loss functions, e.g., for concept-membership or fulfillment of tbox-axioms. For details, see the work of Garg et al. [2019]. Thus Quantum Embeddings are a special case of COSs and demonstrate that COSs can be used as basis of actual embedding approaches.*

9.2 A Conceptual Orthospace for Zero-Shot Learning

In the following, a COS-based learning approach for zero-shot-learning is presented. First, zero-shot-learning is formally introduced and one established solution strategy is discussed which will be used as basis for the proposed approach. After that, the construction of the COS is shown, followed by the experiments done and their results.

The Dataset and the Zero-Shot Learning-Problem

Zero-Shot Learning (ZSL) is a multi-class learning task in which each instance has to be assigned exactly one label. The distinct feature of ZSL is that new instances are classified with labels that haven't been seen while training [Xian et al., 2019]. To be able to label previously unseen classes, some auxiliary information is needed. This information is, e.g., given in form of per-class-attribute information, meaning for each class a set of attributes is given. For example, a ZSL task may require a previously unseen class 'zebra' to be identified, provided the information that zebras exhibit the feature 'striped'. A ZSL learner would need to identify an instance of the zebra class, robustly discriminating it from horses (which share many attributes of zebras but being striped) and tigers (which are striped but unlike zebras otherwise). Formally, $\{(x_i, y_i)|i = 1 \ldots n\}$ are the training instances x_i denoted with labels $y_i \in \mathcal{Y}^{tr}$ that belong to the training classes. Test instances $\{(x_j, y_j)|j = n+1 \ldots n+m\}$ with labels $y_j \in \mathcal{Y}^{ts}$ constitute test classes, where $\mathcal{Y}^{tr}$ and $\mathcal{Y}^{ts}$ are disjoint. Additionally, for each class both of the training and test set, the information about its positive and negative attributes is given.

One prominent approach for solving the ZSL problem is EXEM [Changpinyo et al., 2017] which is based on the insight that instances of unseen labels will cluster around the semantic embedding of that class, thus, so to say, the instance representing these attribute combination prototypically. Therefore, based on the seen classes, a predictive function is learned to predict the semantic embeddings (called *exemplars*) of the unseen classes. The approach works in detail as follows: A transformation function ϕ is learned which is able to map for each class c its attribute vector a_c to the exemplar in the feature space v_c, therefore, $\phi(a_c) \approx v_c$. Vector v_c of a known class is the mean of all instances of that class obtained after performing a class-unspecific PCA projection computed based on training data of seen classes. Then, d Support Vector Regressors (SVRs) (see, e.g., [Awad and Khanna, 2015] for details) with rbf-Kernel are trained, where d is the dimension of the PCA. Using these regressors, the exemplars of unseen classes can be predicted based on their attribute vectors. For a test instance, 1-nearest neighbor (1NN) is used to select the nearest visual exemplar based on (standardized) Euclidean distance. For more details, see the work of Changpinyo et al. [2017].

For evaluation, we consider a standard dataset for ZSL, the Animals with Attributes (AWA2) dataset proposed by Xian et al. [2019]. AWA2 comprises 40 classes for training and ten for testing, each class having 85 either positive or negative attributes. In total $37,322$ images of animals are contained. The classes have 746 images on average, the least populated class 100 and the most populated class 1645 images [Xian et al., 2019]. As feature space the embedding of the images as 2048-dimensional top-layer pooling unit of the ResNet is used [Xian et al., 2019]. As train/test split, the standard split of this dataset is used.

The Resulting Conceptual Orthospace

EXEM is based on classifying an instance based on its nearest neighbor and uses background knowledge in form of attributes. The aim is now to extend this approach by using COSs. Thus, I am not relying on the nearest neighbor for classification but I am using geometric information. Based on these preconditions, first a COS is created and then, based on it, a learning approach is proposed. Here, instances and their numerical features are given. Thus, one possibility would be to choose some notion of similarity based on these numerical features and then create an orthonegation-induced betweenness relation. This leads, however, as argued in Section 6.1.3, often to a weak structure and does not necessarily mimic the underlying background knowledge. Therefore, I choose here the other direction, having a betweenness relation as basis and defining the orthogonality relation based on it. Then, the first question is which concepts should be modeled. As zero-shot-learning is modeled, there isn't any instance in the intersection of label names (nothing is a tiger and a zebra at the same time). Therefore, the COS is modeled based on the attribute space: The exemplar of each known label can be placed at the conjunction of its attributes (and negated attributes). However, how should this space be defined? As distributivity must be valid and Euclidean betweenness seems to be a good choice, the only possible construction (as shown in Chapter 7) is a cone-COS and due to the property of distributivity, an al-cone-COS is chosen. This COS can be directly created without training: Each unseen label is placed on a unique half-axis and the concepts representing the regions are determined based on the attributes the labels have. Thus, the resulting al-cone-COS is of dimension $|\mathcal{Y}^{te}|/2$. The exact placement can be used as a hyperparameter to be tuned, here, the hamming distance is used: labels which attribute vectors have a high hamming distance are placed more distant to each other than labels with a smaller hamming distance, as they can be considered as sufficiently similar. Due to the construction, the betweenness is not relevant here, however, lead to the fact that a well-structured COS can be considered.

Now, this al-cone-COS can be used as the basis for an embedding approach. The approach of Changpinyo et al. [2017] is focused on the distance of a new instance to the predicted exemplars of unseen classes. Using the al-cone-COS allows for extending this approach such that it is not only distance- but geometry-based and, thus, is able to incorporate partial information. The overall approach is similar to that of Changpinyo et al. [2017]: First, the approach of Changpinyo et al. [2017] is used to learn the mapping $\phi(a_c) \approx v_c$. However, now, the label of a new instance is not determined based on the nearest exemplar of a label but the al-cone-COS is incorporated: The mapping ϕ is used to train for each dimension of the al-cone-COS a SVM, discriminating instances with a positive value in this dimension from instances with a negative value in this dimension. The SVM is trained based on the mapping of these elements (thus based on the predicted exemplars). This technique would suffer from lack of data, as on each half-axis only one element is placed, namely the exemplar representing the unseen label. This would lead to the fact that for the training of each classifier only two elements were present. As in this case the advantages of the geometric model can't be used (and the SVM can't be trained suitably), some additional partial information is added. This is done by predicting not only the visual exemplars of the unseen classes but also visual exemplars of unseen superclasses, e.g., *A or B*. Thus, an attribute vector can be determined which could be A or B with the same probability. This vector would have positive (resp. negative) attributes, when both A and B have this attribute (positive, resp. negative). When A and B differ in an attribute, this would be set to zero. Introducing superclasses for training leads on the one hand to the advantage of increasing the number of training instances. On the other hand, it enables to consider the certainty of the classifier's decisions (I discuss this features in the result section further below). A new instance is then classified based on the SVMs for each dimension of the al-cone-COS. In contrast to EXEM, not necessarily one label is proposed, it is also possible that partial information is returned.

Implementation

For the baseline approach, we use the implementation[1] provided by Changpinyo et al. [2017] for their approach EXEM. For adaptation to AWA2 we changed the kernel used to a polynomial kernel because of better classification properties on this dataset. The approach is evaluated by using a class-based five-fold cross-validation (thus, all instances of one class are either completely in the training or completely in the validation set to ensure ZSL also for the cross validation) to determine the best hyperparameters and tested on the proposed test split. The hyperparameters of the regressor have been taken based on the minimization of the Euclidean distance, as done by Changpinyo et al. [2017]. Based on this, the hyperparameters of the SVM have been tuned as a trade-off between overall recall and recall for returning one label, as the result should be as specific as possible by being able to return labels to as many instances as possible. Execution time of the training is very short (few seconds on standard laptop), due to consideration of visual exemplars as representants of the class. The implementation can be found on GitHub[2].

Result and Discussion

Figure 9.1 presents precision and recall obtained for our approach in comparison to the baseline and is accompanied by Table 9.1 which presents the numerical values and their variance. Individual marks in the plot show precision and recall obtained for a classifier output containing the nearest, the two nearest,… exemplars of class labels in case of EXEM. For my approach, for the case of one label only those instances are considered for which a mapping on a half-axis was possible, thus those instances which can be exactly classified. For the case of two labels, all instances are returned containing some partial information, for which, however, only two labels are possible outcomes. This is done analogously for the other cases.

For example, the best class label obtained in the baseline approach achieves about 0.71 precision and 0.71 recall, whereas a single class label obtained by our approach achieves 0.895 precision and 0.094 recall. As can be seen, the baseline outperforms our approach in terms of recall as it can potentially generate any class label, even those not agreeing

[1] See `https://github.com/pujols/Zero-shot-learning-journal`
[2] See `https://github.com/mleemhuis/AMAI_conelearning`

Table 9.1: Precision and recall with respect to the amount of labels returned for classification

# returned labels	precision		recall	
	our approach	EXEM	our approach	EXEM
1	0.895 ± 0.059	0.710 ± 0.048	0.094 ± 0.041	0.710 ± 0.048
2	0.465 ± 0.043	0.443 ± 0.014	0.549 ± 0.092	0.885 ± 0.028
3	0.399 ± 0.039	0.317 ± 0.005	0.816 ± 0.112	0.950 ± 0.016
4	0.378 ± 0.045	0.243 ± 0.003	0.926 ± 0.064	0.973 ± 0.012
5	0.373 ± 0.047	0.20 ± 0.001	0.955 ± 0.031	0.987 ± 0.007
6	0.372 ± 0.048	0.166 ± 0	0.959 ± 0.015	0.995 ± 0.002
7	0.372 ± 0.048	0.1425 ± 0	0.960 ± 0.013	0.998 ± 0
8	0.372 ± 0.048	0.125 ± 0	0.960 ± 0.013	0.999 ± 0
9	0.372 ± 0.048	0.111 ± 0	0.960 ± 0.013	0.999 ± 0
10	0.372 ± 0.048	0.100 ± 0	0.960 ± 0.013	1 ± 0

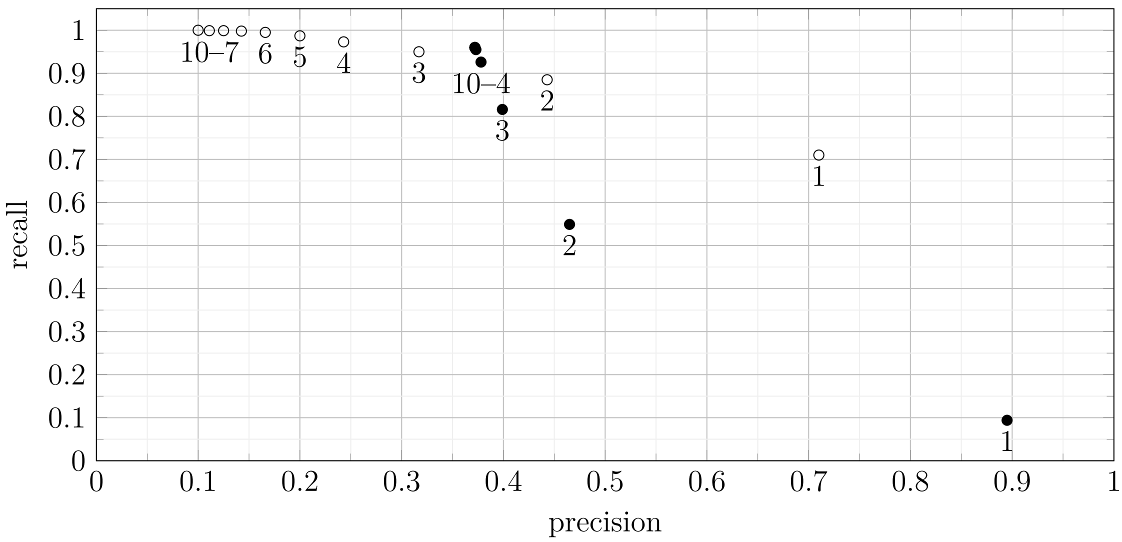

Figure 9.1: Recall and precision for our approach (black) and the other approach (white), based on the proposed test set. Numbers represent the number of labels returned.

Table 9.2: Uncertainty in classification of maliciously crafted AWA2 variant: percentage of instances labeled with exactly i labels for our approach

number of returned labels	instances with i labels (in %)
0	0
1	0
2	0
3	0
4	0
5	3.2
6	23.6
7	44.1
8	25.8
9	3.2
10	0.01

with the auxiliary information. In contrast, our approach is restricted to labels agreeable with auxiliary information given, which in case of sparse and noise training data may inhibit perfect recall. With respect to precision our approach outperforms the baseline, both in terms of the maximal precision achievable and in terms of the precision relative to the number of class labels generated. In other words, the proposed method is well-suited to reject classifier decisions that are not agreeable with background information. Such a feature can be important in applications where wrong decisions must be avoided. This is accomplished by an approach capable of respecting the inherent structure of the feature space.

The results with respect to the proposed approach can be explained based on the fact that it splits the feature space into subspaces representing the areas of the geometric model. The region close to the visual exemplar is represented by a half-axis of the geometric model, further away from the visual exemplars (near the partial information visual exemplar), the union of several half-axes can be found. Thus, is is possible to determine the certainty of a classification result simply by counting the number of possible labels. Considering the baseline, it is possible to increase the recall rate simply by not considering the nearest neighbor but returning the labels of the two, three, four,... nearest exemplars. Such an approach does however provide no information about which distance is still agreeable with the auxiliary information. In contrast, our approach has the ability to state which labels are definitely incorrect for a given instance.

To demonstrate the ability of rejecting labels that are not agreeable, I conduct an additional experiment involving incorrect attributes. I negate the attribute representation of the test set (meaning $a'_c = -a_c$). Based on this, 1NN is executed, resp. the SVMs of our geometric model are trained. The change has the effect that the stated exemplar is not associated with the original exemplar. As expected, both approaches show a recall and precision near zero. However, with our approach one is able to recognize the problem as it returns for no element less than four labels, as can be seen in Table 9.2.

As explained above, the amount of class labels is indicative for the uncertainty of our classifier. Using the capabilities of modeling negation, it can be stated which labels are definitely not correct for an instance. Therefore, the ability of the geometric model to model negation and disjunction is used.

This approach on ZSL and the approach for KGE of Example 9.1 show that the framework of COSs is not only usable in theory but can also be used as a basis for actual learning approaches.

10 Logical Commitments of Learning Approaches

Embedding approaches such as TransE [Bordes et al., 2013] are simple methods to enable link prediction with a high accuracy. However, accuracy is not the only attribute that needs to be considered, other important properties are explainability and trustworthiness. The need for having explainable and trustworthy systems led to the research area of *explainable AI (XAI)* [Arrieta et al., 2020]. XAI is a widely researched area incorporating many different forms of explanations and many ways of increasing the explainability of systems, e.g., by textual or feature relevance explanations. In the last chapters, I discussed a framework able to incorporate information in the form of axioms and rules into the embedding process, thus increasing the interpretability of the embedding. However, for such an approach, it is also necessary to consider its post-hoc explainability. In KGE-approaches the incorporation of logical rules increases not only the interpretability but also the explainability of a system [Bianchi et al., 2020]. However, the embedding of these rules introduces an aspect into the embedding approach which needs explanations by itself because rules are expressed in a logic based on some assumptions. We restate these assumptions under the term "commitments". And, thus, the question arises: which *logical commitments* are actually representable? Here, I focus on the logical commitments of embedding approaches and apply them to COSs.

The main point is that even simple embedding approaches are able to model (implicitly) some sort of logical information (the term "logical" is seen here in the widest possible sense as some form of language of structured representation of knowledge). Examples are properties of roles, e.g., whether a role is transitive or functional, or the behavior of concepts, e.g., whether they are closed under intersection.

This is particularly useful for determining whether a given approach is applicable to a given dataset. When, e.g., the dataset contains a lot of tuples having transitive or symmetric behavior and TransE is used, then even if the result has reasonably good scores, it will be biased towards representing only the non-transitive, non-symmetric parts of the data. However, when the dataset has a mostly simple structure, TransE could be appropriate due to its feasibility and generally good performance. This problem is not restricted to the properties of roles. It also occurs, e.g., when planning to model background knowledge in $\mathcal{ALC}$ with an embedding approach only able to model $\mathcal{EL}^{++}$ or when one has a dataset containing a lot of partial information and uses an embedding approach unable to model this partiality.

The main problem here is not that the embedding should not be used because of its restricted expressivity — TransE, e.g., leads to a good result quality despite its restrictions — but to use the embedding approach only in a context where it is appropriate based on its expressivity. Although there are empirical proofs that the use of rules for learning can improve the result [Gutiérrez-Basulto and Schockaert, 2018], it is not necessarily the best option to model the most expressive embedding possible. The reason is that an embedding technique can be considered as a special type of dimension-reduction strategy as a smaller dimension enforces stronger regularities and thus enables induction [Gutiérrez-Basulto and Schockaert, 2018].

Thus, the aim of this chapter is to give a framework to determine the expressivity of a given embedding approach. Knowledge about the expressivity of an approach can also help to increase the trust in the system by increasing its explainability. When one knows that a system can only learn symmetric relations, it would be absurd to deduce from the embedding that a relation has to be symmetric.

How can this "expressivity in logical terms" be quantified? Kazemi and Poole [2018] introduce the notion of *full expressivity*. They call an approach *fully expressive* if, for a pair of disjoint sets of positive triples P and negative triples N, an embedding (with arbitrary high dimension) can be found such that each triple in P is correctly represented whereas no triple of N is represented. Although this rule is a good basis for considering the expressivity, it is not sufficient, in particular as it is based on an arbitrary number of dimensions. Morevover, although it represents expressivity on data level, it does not do so on ontology level [Gutiérrez-Basulto and Schockaert, 2018]. There are other approaches for defining the expressivity of KGE-approaches such as the *quasi-chainedness-property*[1] or general considerations on the expressivity of learning approaches by Schockaert [2021] and Kazemi and Poole [2018]. These, however, are either not generally applicable or consider only a subset of possible expressivity features. Therefore, I propose the concept of *logical commitments* to enable for a thorough determination of the expressivity of embedding approaches. This term is inspired by the term *ontological commitment*, used for example by Davis et al. [1993] and T. R. Gruber [1995]. Davis et al. [1993] argue that a main part of knowledge representation is the determination of ontological commitments. Although an ontology aims at representing all facts true in every possible world, this representation is not possible in practice due

[1] An existential rule (the datalog rule $B_1 \wedge \cdots \wedge B_n \rightarrow \exists X_1, \ldots, X_j.H$) is *quasi-chained* if for all $1 \leq i \leq n$: $|(vars(B_1) \cup \cdots \cup vars(B_{i-1})) \cap vars(B_i)| \leq 1$ [Gutiérrez-Basulto and Schockaert, 2018].

to its complexity. Thus, for an ontology engineer, it is important to focus on a specific subarea, adapted to his problem; thus, he makes commitments to the modeled ontology. Although an ontological commitment is a restriction of the modeled reality, it gives, as Davis et al. [1993] state, guidance for both the human user and the reasoning machine to find the relevant parts of the problem considered which would not be visible in an overwhelmingly complex ontology. These ontological commitments come into play for different levels of the ontology design process, starting with the most basic choices of the representation system of the ontology and continuing with the choice of concepts, relations and axioms.

Considering a KGE-problem incorporating background knowledge, the ontological commitments of the background logic ontology play a vital role in the perception of the problem and its solution. However, when considering the overall approach, not only the ontology design needs to be considered but also its representation, e.g., the geometric representation of concepts. Additionally, whereas the ontological commitments are part of the ontology design process, here the focus lies on post-hoc explainability. Thus, here *logical commitments* are considered. As we take a concept-centered approach, these logical commtiments are determined on several levels reflecting relations between concepts and operators on them: concept representations, representations of operators on concept level, the properties of roles, role operators and quantifiers and the inference services.

First, in Section 10.1 the different levels of logical commitments are introduced and explained in detail. Their usefulness is demonstrated by applying them to several embedding approaches in Section 10.2. Last but not least, in Section 10.3, I come back to COSs and determine their logical commitments.

10.1 Types of Logical Commitments

There are several different types of logical commitments, discriminable into four levels. The most basic decision concerns the representation of concepts, e.g., probabilistic or fuzzy, and their negation, representing a weak or strong negation. A geometric representation of concepts is not necessary (in this case, the logical commitment is only considered based on level three and four), but when a geometric representation is given, then it is necessary to consider it in detail, as the representation influences the strength of the background knowledge representable. The second level is highly connected to the first level and concerns the operators on concept level by considering their expressivity, e.g., whether they make up a Boolean algebra. The third level depicts the expressivity of role operators and quantifiers, on the one hand the relations between individuals on

inference services	...	deduction	link prediction			...	4. level
role operators and quantifiers	...	$\exists$	$\exists R.$	$\forall$	$R(a,b)$	...	3. level
operators on concept level	...	$\langle \mathbb{1}, \mathbb{0}, \wedge, \vee, \cdot' \rangle$ ortholattice	$\langle \mathbb{1}, \mathbb{0}, \wedge, \vee, \cdot' \rangle$ Boolean alg.			...	2. level
concept representation	...	probabilistic	binary partial	binary total		...	1. level

Figure 10.1: Different levels of logical commitments, implicit or explicit choices in the design of a knowledge representation. The arrow indicates a possible combination of ingredients.

instance level, e.g., symmetry or reflexivity and, on the other hand, general assumptions about roles and their axioms, e.g., existentials or modal operators. The fourth level depicts the calculus used, e.g., (types of) induction or deduction. These levels can be seen in Figure 10.1 and will be discussed in more detail in the following subsections.

Logical commitments are either explicitly or implicitly given. An explicit commitment is, e.g., defined by Kulmanov et al. [2019], where concepts are represented as spheres to allow for the representation of $\mathcal{EL}^{++}$. In the same approach, the relations are represented as translations. This leads, by design of the embedding approach, to the implicit commitment that only intransitive relations can be represented.

10.1.1 Level 1: Concepts

The lowest level concentrates on the representation of concepts and their interaction in geometric space. The first question concerns the attributes of the underlying embedding space: are there explicit *quality dimensions* [Gärdenfors, 2000, p. 5], thus do the dimensions have a semantic meaning? An example is the representation of the color domain, given by Gärdenfors [2000, p. 11]: it defines the colors based on brightness, chromaticness and hue as quality dimensions. It is also possible that the embedding approach results in implicit quality dimensions, and, thus, that it is possible to give dimensions a semantic meaning after training, as done, e.g., by Jameel and Schockaert [2016]. A related question is whether domain information is incorporated, e.g., in the form of explicit subspaces, depicting specific domains. The difference between using dimensions with and without a semantic meaning is illustrated in the following example.

Example 10.1. *Again, the embedding of colors is considered. When interpreting the col-*

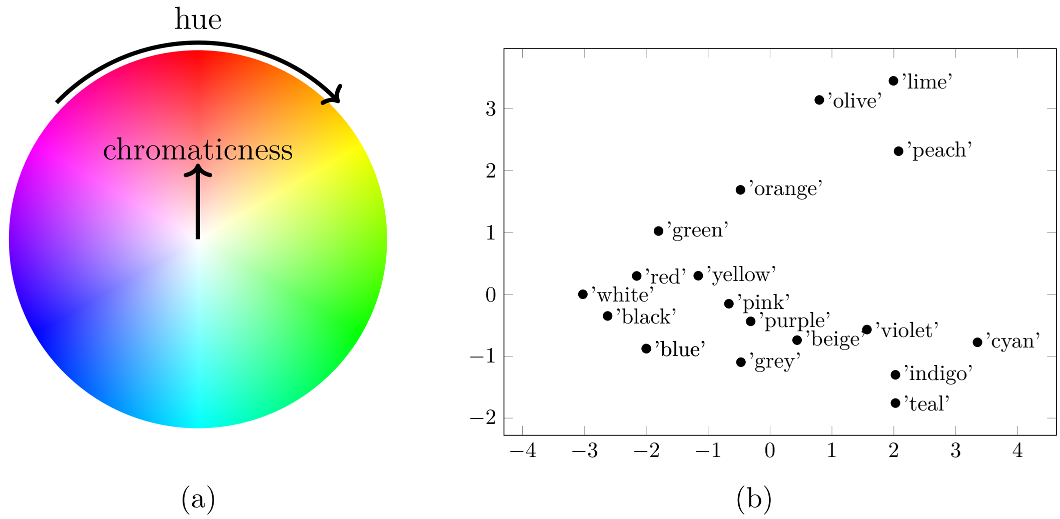

Figure 10.2: (a) Example of a color spindle (third axis of brightness is omitted); (b) Embedding of some well-known color terms with the help of word2vec.

ors as being attributes solely in the color domain and when using quality dimensions, this results in the color spindle (as can be seen in Figure 10.2(a)) with the two quality dimensions chromaticness and hue (brightness is omitted for simplicity). In Figure 10.2(b), an embedding of some colors with the help of word2vec [Mikolov et al., 2013] is presented.[2] *This representation is not solely focused on the color domain (it seems some of the colors are also interpreted as fruits (e.g., peach, olive) and grouped together) and there is no quality dimension recognizable. Note that it is not necessarily the case that an embedding like word2vec does not result in quality dimensions with a semantic meaning (they are only not enforced to have a meaning) and that here the original embedding space is high-dimensional, such that a reduction to two dimensions is a severe reduction and is done here only for illustration purposes.*

The most important question on the first level is how concepts are modeled. In approaches modeling some background knowledge, this is done explicitly, e.g., as n-dimensional balls [Kulmanov et al., 2019] or as convex regions [Gutiérrez-Basulto and Schockaert, 2018]. In this case, different design decisions for the concept extension can be made. One basic discrimination is based on whether a concept is embedded as a

[2] Word2vec embeds words into a vector space based on a text corpus such that words appearing in similar contexts are mapped closer together. Here, I used the pretrained model *glove-wiki-gigaword-100* (`https://huggingface.co/fse/glove-wiki-gigaword-100`, accessed 18.10.2023). For illustration purposes, I reduced the dimensions with the help of principal component analysis.

primitive notion, as has been done, e.g., for spheres [Kulmanov et al., 2019] or whether it is incorporated as a collection of attributes, e.g., representing an instance of the concept “apple” based on attributes in the color, the shape and the taste domain, as proposed, e.g., by Gärdenfors [2000]. On the other hand, it is also possible to model concepts reified in form of points in the space, as, e.g., done in SimplE [Kazemi and Poole, 2018].

Another discrimination is whether negation is represented, and, thus, whether the concept has not only an extension but also an anti-extension. Not all approaches declare an anti-extension, as, e.g., by Kulmanov et al. [2019] the DL $\mathcal{EL}^{++}$ is modeled which only incorporates non-intersection of concepts but no anti-extension. Approaches modeling an anti-extension are able to model an expressive negation, thus modeling an anti-extension opens up the possibility for modeling a wide range of ontologies/an expressive background knowledge.

When including the anti-extension, there are several possibilities, e.g., modeling weak negation, thus partitioning the space into extension and anti-extension (called binary total in Figure 10.1), or modeling strong negation where extension and anti-extension do not cover the whole space (thus resulting in a binary partial definition). Related to modeling anti-extension — but also applicable when no anti-extension is modeled, but, e.g., for disjunction overextension occurs — it is necessary to consider the space “in-between”, namely, whether uncertainty, vagueness or truthlikeness or none of them are modeled (for details, see Section 3.2). Besides this, there are other possible variants, e.g., modeling with a three-valued logic.

In addition to having tight borders such as for spheres [Kulmanov et al., 2019], it is also possible to define fuzzy or probabilistic concept extensions.

Another important point is whether inconsistencies can be modeled. On the one hand, modeling inconsistencies can be achieved based on the loss function. In KGE, the embedding is trained by minimizing a loss function. This loss function depends on the proposed approach, however, in general penalizes embeddings interfering with the background knowledge. If the training is continued until the loss equals zero, then a consistent embedding is learned. However, not enforcing the loss to be zero allows for modeling inconsistencies. This is, however, not comprehensible as the user is not able to know whether a result is based on an inconsistent input. Thus, it is a logical commitment to be able to model inconsistencies explicitly, as then it is possible to keep the inconsistencies and analyze them, e.g., as discussed in Section 5.2.1 with the help of an element 0 which represents $\overline{0}^{\mathcal{I}}$. Another commitment is whether inconsistency of the background knowledge can be tackled. Although this is not discussed further here, approaches can be found, e.g., in the work of Flouris et al. [2006].

After analyzing the representation of concepts, the next question regards the modeling of logical operators geometrically. One needs to consider which operators are representable at all and how they are represented. Whereas, e.g., many approaches model conjunction as set-intersection, doing so is not a necessary condition. For the cone-based embedding ConE [Zhang et al., 2021], for example, the conjunction is approximated based on weight vectors. Additionally, properties like over- and underextension, both for conjunction and disjunction, need to be considered (for details, see Section 3.4).

10.1.2 Level 2: Operators on Concept Level

Based on the representation of concepts, their operators and their expressivity are considered. A binary total interpretation, e.g., will lead to a Boolean algebra if conjunction and disjunction are interpreted as set-intersection and set union [Dunn, 1996, p.8]. This level is strongly related to the one before, as the geometric representation of concepts highly influences the logical expressivity.

At this level, in particular, the determination of the logical commitments is difficult. Although it is relatively easy to determine whether, e.g., the represented lattice is restricted to a Boolean algebra or an ortholattice, it is difficult to determine the exact strength, thus, e.g., whether actually each Boolean algebra resp. each ortholattice can be modeled.

Even when allowing for an arbitrary number of dimensions, such a characterization result is non-trivial. For example, for the case of embedding concepts as linear subspaces in an n-dimensional space as done in a practical embedding setting by Garg et al. [2019] and in a theoretical setting in the area of quantum logic, considerations on representatbility made up a whole research area [Birkhoff and Neumann, 1936]. Even if the strength has been determined for the general approach, it is additionally necessary to consider differences in expressivity when a smaller number of dimensions is used. This can be seen in the case of Example 6.25 in which I showed that some lattice was not representable based on $X = \mathbb{R}$. However, when increasing the size of the domain to $X = \mathbb{R}^2$, it is possible to construct this lattice (an example relying on cones can be seen in Figure 5.4(b)). Thus, even when finding an axiomatization, the number of dimensions used for the actual learning problem has to be considered to determine whether the data in this dimension is actually representable. In particular, an approach with a high expressivity in an infinite (or quite big) dimension could be less usable than an approach having an overall lower expressivity, which is, however, higher in a lower dimension.

Besides the Boolean algebra and the ortholattice, DLs are often used in embedding approaches to define the strength of the underlying background logic. There, mostly $\mathcal{EL}$ and its extensions are used (see Section 2.1), which can be equivalently represented in first order logic. When modeling the expressivity of the concept operators with the help of DL, this also influences the next level of logical commitments, the role operators and quantifiers, as they are incorporated in the definition of the DL.

After determination of the concept operators and the underlying logic, it is possible to determine whether the given ontology can be modeled faithfully (regarding the concept level) with the help of the given approach. As argued in Section 3.2, this faithfulness is not necessarily a desirable property. Faithfulness has the same drawback as full expressivity in the sense that it does not state anything about the number of dimensions needed. The general possibility of creating a faithful model thus does not state that it is actually possible to create one for an actual learning approach. However, in contrast to fully-expressivity, it incorporates the ontological information and not only the instance information.

10.1.3 Level 3: Roles, Role Operators and Quantifiers

The third level is relevant for all approaches, even if concepts are not considered, as roles are of vital importance, in particular for KGE. For approaches not considering concepts, level one and two are mostly not applicable, thus, for those approaches, the logical commitments are determined based on levels three and four.

The ability to induce new triples in knowledge graphs is dependent on the ability to model expressive role operators. This is a widely researched area. One embedding approach for which its expressivity was widely discussed is TransE, whose relations are either trivial or restricted to anti-symmetric, irreflexive and intransitive relations [Kazemi and Poole, 2018]. There are many extensions of TransE able to circumvent at least some of these restrictions (for more details, see the overview by Q. Wang et al. [2017] and for their expressivity, see the overview by Y. Wang et al. [2018]).

The most basic properties of roles are (anti-)symmetry, (ir-)reflevity, (in-)transitivity and (non-)functionality. These can be stated explicitly in a learning approach, e.g., when modeling the embedding of a specific description logic, or can be enforced implicitly by construction of the embedding of the relation, e.g., by using translation in TransE. Additionally, there are other properties that can be considered, e.g., *quasi-chainedness*, as introduced by Gutiérrez-Basulto and Schockaert [2018].

When considering not only instances but also concepts, it is possible to determine the expressivity of quantifiers, e.g., whether exists- and forall-quantifier are representable at all, whether they are dependent on each other or whether they are distributive (in the sense that $\exists R.(A \sqcup B) = \exists R.A \sqcup \exists R.B$ or $\forall R.(A \sqcap B) = \forall R.A \sqcap \forall R.B$, resp.). Whereas the propositional expressivity was considered in the last levels, in this level the expressivity of the roles is considered, thus, e.g., whether only $\mathcal{EL}^{++}$ is representable (where $\forall$-constraints are not representable), or whether only $\mathcal{DL}_{Lite}$ is representable which does not allow for qualified existentials on the left-hand side [Artale et al., 2009].

This is again dependent on the first and second level, as, e.g., the definition of a full negation leads to the fact that the all-quantor can be defined as a negated exist-quantor.

Strongly related to the considerations on properties of roles and tbox-axioms containing roles are considerations on the quantifier rank of concepts and the chain length of abox-assertions that can be modeled. Transitivity, e.g., can not be modeled if it is not possible to model the assertions $R(a, b), R(b, c)$ at all, thus no chains of length 2.

Example 10.2. *Consider a vocabulary containing arbitrary many constants* $x_i, 0 < i \leq n$ *for* $n \in \mathbb{N}$ *and let the abox contain* $R(x_i, x_{i+1})$ *for all* i *with* $0 < i < n$. *Then, the abox contains a chain of length* $n - 1$. *If* $R(x_n, x_1)$ *is contained, then it contains a cycle. TransE, for example, is not able to model cycles.*

Thus, it is necessary to determine whether an embedding approach is able to model such arbitrary long chains or even cycles. This commitment is not only needed on abox-level but also on tbox-level. On the one hand, tbox-axioms could have an arbitrary quantifier rank or could contain cycles. On the other hand, it should be possible to model arbitrary quantifier ranks to construct a strongly concept-faithful model (see Definition 8.11). The problem of arbitrary quantifier rank was considered in detail in Section 8.2, where a solution strategy for the case of an embedding based on convex axis-aligned cones was presented.

Next to roles, on this level, it is also possible to consider (other) modal operators and their expressivity (note that DLs are closely related to modal logic when interpreting roles as accesibility relations [Baader, Calvanese, et al., 2007, p.41-42]).

10.1.4 Level 4: Inference Services

The fourth level, the inference service, is the only one that is mainly based on an explicit design decision. Here, the inference service is chosen based on the goal of the approach considered. The goal can be, e.g., link prediction as in TransE and other classical KGE-approaches or query answering, thus embedding FOL queries and answering a query by following the corresponding computation graph and returning matching instances, as done, e.g. by Zhang et al. [2021]. Other options are to do property testing, thus examining whether an entity belongs to a concept, and the identification of the most specific concept of an instance. It is also possible, as done, e.g., by Xiong et al. [2022] to do non-classical plausible reasoning in form of predicting, e.g., $C \sqsubseteq D$ with a probability of 0.9, thus, inferring a new subsumption relation which is not entailed by the input knowledge base. This can, on the one hand, be interpreted as the advantage of doing non-classical plausible reasoning, on the other hand, it can be interpreted as problem of missing expressivity, thus, e.g., the choice of a dimension size which is too small for the faithful representation of the knowledge base.

10.2 Determination of Logical Commitments Based on KGE-approaches

After introducing the logical commitments and categorizing them into different levels, in the following two sections, the level-based classification is applied to a variety of different embedding approaches. To illustrate the applicability of logical commitments, three different approaches, resp. groups of approaches are considered. First, the logical commitments of approaches not incorporating background knowledge are considered. For this, I concentrate on TransE, a well-known member of this group. Next, SimplE [Kazemi and Poole, 2018] is considered as an approach not able to model concepts explicitly, however, can model them implicitly. After that, the logical commitments of approaches incorporating information on concepts (thus, either explicit background information or implicit information about concept relations incorporated into the training data) are considered. Therefore, first, some logical commitments of such embedding approaches are exemplarily examined and then, in Section 10.3, the procedure of determining the logical commitments is exemplified using COSs.

TransE

Although TransE [Bordes et al., 2013] is a simple method, its application leads to surprisingly good results. TransE embeds (subject, predicate, object)-triples into a low-dimensional vector space. The subject and object are embedded as points in the space, whereas the predicate is embedded as a translation. If $||s + p - o|| < \delta$ for some fixed treshold $\delta > 0$ (where $|| \cdot ||$ is the standard Euclidean metric), then (s,p,o) is assumed to be valid. With the help of negative sampling, corrupted triples are defined and both correct and corrupted triples are used as input for a loss function. TransE is solely based on abox-information and thus does not incorporate any background information.

Now, the logical commitments are considered per level: Starting with the first level, no quality dimensions or domains are modeled, neither explicitly nor implicitly. Concepts are not modeled explicitly. Concepts could be modeled implicitly, as the abox-statement $C(a)$ can be represented as triple $(a, \text{ instance-of}, C)$, thus interpreting C as reified instance. However, as only functional relations can be modeled — in fact, relations (s, p, o_i) where all o_i are inside an n-ball of radius δ — this leads to a scenario in which concept C is only allowed to incorporate instances that are inside a small margin δ around a.

It is possible to model $\exists R.\overline{1}$ as a concept, however, due to missing partiality, it leads to $\exists R.\overline{1} = \overline{1}$ for all $R \in N_R$. Due to the missing representability of concepts in level 1, commitments of level 2 are not applicable.

Level 3 and 4 stand in contrast to the first two levels and are also well applicable for approaches not incorporating background information. Relations in TransE are, e.g., functional (or at least nearly functional in the sense that two triples (s, p, o_1) and (s, p, o_2) are only possible if o_1 and o_2 are close together (having a distance of at least δ)). It is not possible either to model partial information explicitly with TransE. A subject-vector translated with a predicate vector needs to end up at some point in the space. The only possibility to model partiality is if this point is not populated by any known instance. This, however, can not be guaranteed. Other restrictions of TransE are shown by Kazemi and Poole [2018].

An advantage of TransE is that it is possible to model an arbitrarily long chain in the abox, as relations are modeled as translation in a vector space. However, the chains are quite restricted: particularly, when more than one chain is modeled. Consider, e.g., the two chains $R(a, b), R(b, c), R(c, d)$ and $S(d, b), S(b, c), S(c, a)$. It is not possible to model both based on relations represented as translations, except $a = b = c = d$. The logical commitment of the fourth level is link prediction.

SimplE

I consider SimplE [Kazemi and Poole, 2018] as a second example of an embedding approach not relying (explicitly) on the embedding of concepts based on some geometric structure. It is a widely known bilinear method, in which two vectors $h_e, t_e \in \mathbb{R}^d$ represent each instance and two vectors $v_r, v_{r^{-1}} \in \mathbb{R}^d$ represent each role. The loss function is then defined as $0.5 \cdot (\langle h_{e_i}, v_r, t_{e_j} \rangle + \langle h_{e_j}, v_{r^{-1}}, t_{e_i} \rangle)$, where $\langle v, w, x \rangle = \sum_{i=1}^{n} v_i \cdot w_i \cdot x_i$. This can also be written as $(h_{e_i}\, t_{e_i})^T M_r (h_{e_j}\, t_{e_j})^T$ with $M_r \in \mathbb{R}^{d \times d}$ (for details on the structure of M_r, see [Kazemi and Poole, 2018]), thus SimplE is a bilinear model. Learning is done based on stochastic gradient descent.

In the same manner as for TransE, concepts are not explicitly incorporated. However, in contrast to TransE, concepts can be implicitly defined as convex regions. For example, consider the two triples $(a_1,\ \text{instance-of}, C)$ and $(a_2,\ \text{instance-of}, C)$ and assume they are both correctly modeled. This is possible, as SimplE is not functional (as it is, in fact, fully expressive (see the introduction of this chapter or the work of Kazemi and Poole [2018] for details)). Thus, it needs to be the case that $0.5 \cdot (\langle h_{a_i}, v_{\text{instance-of}}, t_C \rangle + \langle h_C, v_{\text{instance-of}^{-1}}, t_{a_i} \rangle) \leq \tau$ for $i \in \{1, 2\}$. Now, assume an instance a_3 is given with $a_3 = \lambda a_1 + (1-\lambda) a_2$, where $a_i = (h_{a_i}\, t_{a_i})^T$ for $i \in \{1, 2\}$. This leads to $0.5 \cdot (\langle h_{a_3}, v_{\text{instance-of}}, t_C \rangle + \langle h_C, v_{\text{instance-of}^{-1}}, t_{a_3} \rangle) \leq \tau$ and thus leads to the fact that, although C is represented as a vector $(h_C\, t_C)^T$, the set of instances in the extension of C is convex.

Although it is possible to define conjunction and subset-relations straightforwardly, SimplE does not model negation. Though, negative sampling is used for learning, negative information is not modeled in the system. Therefore, it is possible to embed knowledge expressed in $\mathcal{EL}^{++}$ or a subset of it. Even for such a case, one needs to examine whether the representation of the "instance-of"-relation is strong enough to enable modeling arbitrary forms of conjunction and subset relations. This is, at least based on my interpretation mentioned above, not possible. This is reflected in the work of Gutiérrez-Basulto and Schockaert [2018], who argue that it is not possible to model relations of the type "$(a,\ \text{instance-of}, C)$ and $(C,\ \text{is-a}, D)$ entails $(a,\ \text{instance-of}, D)$" by bilinear methods without enforcing triviality [Gutiérrez-Basulto and Schockaert, 2018, Proposition 2].

The Logical Commitments of Approaches with Explicit Representation of Concepts

Of the last two approaches, TransE did not contain conceptual information at all, whereas for SimplE, it was at least implicitly contained. Now, the approaches incorporating explicit geometric concept representations are considered. On the one hand, there are query embedding approaches [H. Ren et al., 2020; Zhang et al., 2021]. On the other hand, there are approaches allowing for KGE with the incorporation of background knowledge information and, regarding a slightly different topic, approaches which handle the embedding in general without being explicitly categorizable [Jameel and Schockaert, 2016; Gutiérrez-Basulto and Schockaert, 2018]. These approaches rely on the representation of concepts as boxes, spheres, several types of cones (including linear subspaces) or general convex sets. These approaches have much in common, in particular when based on the same concept representation. The approaches considered are: representing concepts as boxes by Query2Box [H. Ren et al., 2020], Box^{2}EL [Jackermeier et al., 2023], Faithful[3] Embeddings in $\mathcal{EL}^{++}$ [Xiong et al., 2022] and ELBE [Peng et al., 2022], as cones by ConE [Zhang et al., 2021], Embed2Reason [Garg et al., 2019], conceptual subspaces [Jameel and Schockaert, 2016], al-cones [Özçep et al., 2020; Özçep et al., 2023] and cones with roles [Leemhuis, Özçep, and Wolter, 2022a], as spheres by ELEm [Kulmanov et al., 2019] and EmEL^{++} [Mondal et al., 2021] and as general convex regions by geometric models [Gutiérrez-Basulto and Schockaert, 2018]. As al-cones and cones with roles can be considered as COS, these are not incorporated in the discussion here but considered in the next section when determining the logical commitments of COSs.

[3] The term faithful refers here to the possibility of representing conjunction faithfully as set-intersection and not to the faithfulness mentioned in Definition 3.4.

As the identification of the logical commitments of all these approaches is beyond the scope of this section, here, for each of the four levels, the logical commitments of these approaches are sketched with a special focus on outstanding properties.

This collection is by no means complete. In particular, logical commitments can also be determined for embeddings not based on explicit geometric structures but nonetheless incorporating background knowledge, e.g., a cluster-based architecture by Aspis et al. [2022] or *Logic Tensor Networks* [Badreddine et al., 2022], which incorporate a logic layer as the last layer of a neural network.

Level 1: Concepts All of the above-mentioned approaches represent concepts as geometric objects. In contrast to the other approaches, only Jameel and Schockaert [2016] consider an embedding approach in which quality dimensions and domain information are learned.

Most of the approaches are restricted to $\mathcal{EL}^{++}$, and, thus do not incorporate any information on negation. Garg et al. [2019] define concepts as axis-parallel linear subspaces and their negation based on polarity. Therefore the concept representation is binary partial. Zhang et al. [2021] consider cones and their negation as set-complement, thus, a binary total representation. All other approaches mentioned above do not consider negative information at all — except for training based on negative sampling (see Section 3.3 for details).

All the approaches allow for some form of conjunction, but only a few allow for disjunction. Conjunction is mostly modeled as set-intersection, with four exceptions: H. Ren et al. [2020] and Zhang et al. [2021] represent conjunction as multi-layer perceptron to approximate set-intersection, Kulmanov et al. [2019] and Mondal et al. [2021] consider spheres which are not closed under intersection and thus are only able to model axioms of the form $A \sqcap B \sqsubseteq C$ by approximating it by enforcing $C^{\mathcal{I}}$ to be strictly greater than $(A \sqcap B)^{\mathcal{I}}$. Thus, this approach leads to overextension of the conjunction (for details on overextension, see Section 3.4), whereas the multi-layer perceptron leads to a non-determinable different interpretation of conjunction — and thus decreases the explainability. Disjunction is only modeled for three approaches and of them only by one approach geometrically: H. Ren et al. [2020] and Zhang et al. [2021] focus on query embedding and propose rewriting the query into disjunctive normal form, then determining the results for each part, doing a disjunction on symbolic level afterwards. Therefore, the representation of disjunction misses the geometric information incorporated, e.g., in the representation of conjunction as set-intersection. The only approach modeling disjunction geometrically is by Garg et al. [2019], where it is modeled as convex closure of axis-parallel linear subspaces. Thus, there the disjunction leads to an overextension.

Level 2: Operators on concept level Most of the approaches model the DL $\mathcal{EL}^{++}$, as $\mathcal{EL}^{++}$ is a DL that is widely usable in many domains due to its simplicity [Baader, Brandt, et al., 2005]. Exceptions are the work of Garg et al. [2019] considering $\mathcal{ALC}$ (in fact, the construction would be able to also model orthomodular lattices) and, situated in the area of query embedding, ConE by Zhang et al. [2021] who propose to be able to model FOL-queries. However, they all need to be examined for their ability to model each $\mathcal{EL}^{++}$- resp. $\mathcal{ALC}$-ontology. This is definitely not possible for a fixed dimension, as shown by Gutiérrez-Basulto and Schockaert [2018] based on Helly's theorem. However, even for an infinite dimension, it is not necessary the case that all $\mathcal{EL}^{++}$- resp. $\mathcal{ALC}$-ontologies can be modeled. As discussed in Section 10.1.2 for quantum logic, the exact determination of the expressivity, e.g., in the form of a representation theorem, is complicated and is therefore not discussed in detail here. However, it is possible to at least find a lower bound on the expressivity by examining counterexamples for well-known rules. This is done exemplarily for a COS based on closed convex cones in Section 10.3. As DLs consider not only concepts but also roles, there is a tight connection to the considerations on role operators of the next paragraph.

Level 3: Roles, role operators and quantifiers The expressivity of embedding approaches regarding properties such as reflexivity or functionality is frequently discussed for many of the approaches considered. This can be seen, e.g., in the considerations of Kazemi and Poole [2018]. As many KGE-approaches are mainly focused on link prediction (see next paragraph on inference services), the representation of roles is that part of the embedding approaches in which they deviate the most from each other to overcome the restrictions in expressivity in terms of the logical commitments. In cases in which the representation of roles is based on an established representation and only extended to concepts, it is possible to infer the properties based on the properties of the underlying approach. H. Ren et al. [2020], Kulmanov et al. [2019], Peng et al. [2022] and Mondal et al. [2021], e.g., rely on the representation of relations as translation and thus have to cope with most of the restrictions of TransE. As these properties have been widely discussed for many of these approaches in the literature (see, e.g., the discussions by Kazemi and Poole [2018]), they are not discussed in detail here.

One commitment particularly important for approaches incorporating conceptual information is the expressivity of roles on concept level, thus, whether, e.g., $\exists R.C \sqsubseteq B$ is representable. This has been discussed in the last paragraph in the context of the DL represented.

Based on these considerations, one main point is to determine whether it is possible to represent concepts with an arbitrary quantifier rank and relations with an arbitrary chain length. The relation-based approaches by H. Ren et al. [2020], Kulmanov et al. [2019], Peng et al. [2022] and Mondal et al. [2021] are able to model arbitrary long chains of instances (thus $R(x_i, x_{i+1})$ for $0 < i \leq n$ for some fixed but arbitrary $n \in \mathbb{N}$) and also assertions of the form $\exists R.\exists R. \ldots \overline{1}$. However, they cope with the same restrictions as TransE and are, thus, only able to model restricted chains (for details, see the considerations on TransE earlier on in this section). Xiong et al. [2022] use affine transformations and thus circumvent partly the problems of TransE. In their approach, it is, however, still problematic, e.g., to model cycles. Garg et al. [2019] rely on representing the relations as $2d$-dimensional axis-parallel linear subspaces for d-dimensional representations of instances. The dimension is, thus, doubled to represent $R(a, b)$ as vector $(a^{\mathcal{I}}\, b^{\mathcal{I}})^T$ lying in axis-parallel linear subspaces R. Although an arbitrary chain length is possible, the quantifier depth is restricted, as the number of axis-parallel linear subspaces is restricted for a given dimension.

Level 4: Inference services Most approaches do link prediction and subsumption reasoning. As described in Section 10.1.4, it is also possible to do non-classical plausible reasoning as done, e.g., by Xiong et al. [2022]. As stated above, the aim of H. Ren et al. [2020] and Zhang et al. [2021] is query embedding and thus query answering.

10.3 The Logical Commitments of Conceptual Orthospaces

COSs enable incorporating an expressive background logic into an embedding approach, in particular by allowing to model negation. However, what are the exact logical commitments of such approaches, and, thus, what can actually be modeled? Parts of this question have already been answered in the last chapters and are now combined in this section to explore the overall logical commitment of COSs. In this chapter, I partly determine the logical commitments of COS and, whenever an exact determination of the commitments is not possible, discuss the problems in detail. For details on the commitments mentioned, the respective chapters are referenced.

Level 1: Concepts

As one of the main motivations for defining a COS was to allow for the representation of negation, the concept extension is binary. As discussed in Section 3.2, partiality can be modeled, thus, it can even be considered as binary partial, where the exact interpretation of the partiality is dependent on the specific COS and the use case. The concepts have a sharp boundary, and are thus neither probabilistic nor fuzzy. However, it would be worth considering extensions of COSs allowing for it. Quality dimensions and explicit domains are not considered. Whether they are implicitly contained needs to be determined based on the actual embedding approach. Conjunction is interpreted as set-intersection, disjunction as $\perp$-closure. Therefore, overextension of the disjunction could occur (this is again dependent on the specific COS, as it could be the case that the $\perp$-closure is equivalent to the set-union; for details, see Section 3.4). A nice property of COSs is that, with the help of an element $0 \in X$ which represents $\overline{0}^{\mathcal{I}}$, it is possible to model inconsistent instances explicitly (for details, see Section 5.2.1).

Level 2: Operators on Concept Level

Lattice $\mathcal{L}_Y$ of a COS is an ortholattice (as shown in Theorem 5.10). However, as can be seen, e.g., in Example 6.25, the actually representable lattice is highly dependent on the specific COS considered.

Therefore, in the following, the determination of the second level logical commitments is illustrated with a COS based on closed convex cones, where the domain X is $\mathbb{R}^n$, Y is a subset of the set of closed convex in $\mathbb{R}^n$, standard Euclidean betweenness $\mathcal{B}_E$ is used and the orthogonality relation is defined based on the polarity (in the following, termed “cone-COS”; for details, see Example 5.6).

As a COS is considered, the rules of orthologic are the lower bound of expressivity of the cone-COS. The remaining question is whether each ortholattice can be represented as induced lattice $\mathcal{L}_Y$ of a cone-COS. As far as I know, the answer to this question is open. It is at least possible to construct a number of counterexamples for many commonly used rules, e.g., distributivity and orthomodularity, thus showing that cone-COSs are able to model a wide variety of lattices. These rules are considered in detail below, to give an impression of the expressivity of cone-COSs.

Table 10.1 summarizes some prominent rules for which counterexamples in a cone-COS exists. The rules above the double line are rules in a propositional calculus that are either discussed directly in the logic literature or that are obvious translations of rules discussed in the context of lattice theory. The rules below the double line are defined without a reference to a calculus but directly with lattice-theoretic notions because these are not expressible in a propositional calculus without propositional quantifiers. The calculus of Goldblatt [1974] considered in Section 3.1 does not need quantifiers. Counterexamples for (D), (MSD), (wLLJ), (LLJ), (JSD) and (W) are shown in Figure 10.3. The instance on the top right of Figure 10.3 is a counterexample to orthomodularity. Thus, it is also a counterexample for (M) and (D). Moreover, the same instance serves as a counterexample for the condition of Birkhoff (Bi) and hence also of (MS), (Mac1), (SM):

Table 10.1: Well-known rules excluded from a logic of cones.

Name	**Propositional Calculus Rule**	**Comment**
Distributivity (D)	$A \mathbin{\&} (B \bigvee C) \dashv\vdash (A \mathbin{\&} B) \bigvee (A \mathbin{\&} C)$	Adding (D) to ortholattices gives Boolean logic
Meet-Semi-Distributivity (MSD)	$\dfrac{A \mathbin{\&} C \dashv\vdash B \mathbin{\&} C}{A \mathbin{\&} C \dashv\vdash (A \bigvee B) \mathbin{\&} C}$	Weakening of (D)
Join-Semi-Distributivity (JSD)	$\dfrac{A \bigvee C \dashv\vdash B \bigvee C}{A \bigvee C \dashv\vdash (A \mathbin{\&} B) \bigvee C}$	Weakening of (D)
Modularity (M)	$\dfrac{C \vdash A}{A \mathbin{\&} (B \bigvee C) \vdash (A \mathbin{\&} B) \bigvee C}$	Weakening of (D)
Orthomodularity (OMR)	$\dfrac{A \vdash B,\ \neg A \vdash C}{A \bigvee (B \mathbin{\&} C) \dashv\vdash (A \bigvee B) \mathbin{\&} (A \bigvee C)}$	Weakening of (M)
Johansson's minimal negation (LLJ)	$\dfrac{A \mathbin{\&} B \vdash C}{A \mathbin{\&} \neg C \vdash \neg B}$	Only rule in Dunn's kite falsified by cones.[4]
Weak J.'s negation (wLLJ)	$\dfrac{A \mathbin{\&} B \vdash D \mathbin{\&} \neg D}{A \vdash \neg B}$	Weakening of (LLJ)
Name	**Lattice rule**	**Comment**
M-symmetry [Grätzer, 2011, p. 336]	If $M(a,b)$ then $M(b,a)$	Weakening of (M)
Mac Lane's Cond. (Mac_1) [Stern, 1991, p. 111]	If $b \wedge c < a < c < b \vee a$ then there is a d with $b \wedge c < d \leq b$ and $a = (a \vee d) \wedge c$	Weakening of (M)
Semimodularity (SM) [Stern, 1991, p. 2]	If $a \wedge b <: a$ then $b <: a \vee b$ ($<:$ denotes the covering relation)	Weakening of (MS) and (Mac_1) equivalent in the finite
Birkhoffs's Covering (Bi) [Stern, 1991, p. 3]	If $a \wedge b <: a, b$ then $a, b <: a \vee b$	Weakening of (SM) and (Mac_1)
Whitman Cond. (W) [Grätzer, 2011, p.479]	If $a \wedge b \leq c \vee d$ then $a \leq c \vee d$ or $b \leq c \vee d$ or $a \wedge b \leq c$ or $a \wedge b \leq d$	Used in discussion of free lattices

Let a and b of the rule (Bi) be instantiated by b and a' of the instance on the top right of Figure 10.3. Then the precondition is fulfilled as $\mathbb{0} = b \wedge a' <: b$ and $\mathbb{0} = b \wedge a' <: a'$. On the other hand we do not have $b <: b \vee a' = \mathbb{1}$ nor $a' <: b \vee a' = \mathbb{1}$.

As we can find cone-COSs not fulfilling the distributivity law (D) of $\wedge$ over $\vee$, the logic generating cones is indeed a non-classical propositional logic. Moreover, some intuitive rewritings known from classical logics are not possible. For example, Johansson's constructive contraposition (LLJ) and its weakening (wLLJ) do not hold. Closed convex

[4] Based on the version of Dunn's kite introduced, e.g., by Hartonas [2016], not mentioned in the version discussed in Chapter 3

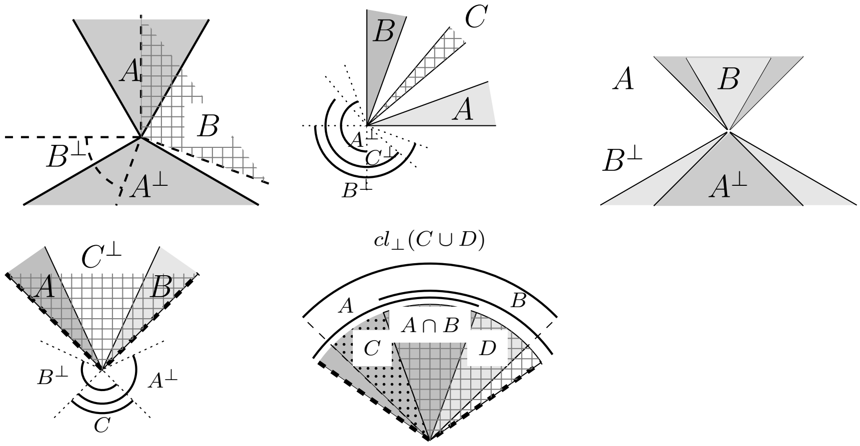

Figure 10.3: Counterexamples for rules falsified by cones. Top left: excluded middle. Top middle: distributivity. Top right: orthomodularity. Bottom left: Join-Semi-Distributivity. Bottom right: Whitman Condition

cones falsifying (wLLJ) show that there are two different notions of complement. The first says that b is a complement of a iff $a \wedge b \leq \mathbb{0}$, the second iff $a \leq b'$. In fact, it is known that any ortholattice which has a unique complement must in fact have the distributivity property and hence must be a Boolean algebra (though this is not the case for arbitrary lattices [Dilworth, 1945]).

As orthomodularity is not fulfilled, the kind of logic induced by cone-COSs cannot be that of quantum logics [Chiara and Giuntini, 2002] (in fact, the main underlying geometric structure for quantum logics is not that of a cone, at least in the pioneering work of Birkhoff and Von Neumann [1975], but that of a closed subspace of a Hilbert space).

The Whitman condition is discussed in the context of free product lattices. A counterexample in $\mathbb{R}^2$ is portrayed in the bottom right of Figure 10.3.

Though the considerations above give a reasonable lower bound for the representability of cone-COSs, they are not sufficient for a representation of the exact logic representable. To give such a representation theorem, it would be necessary to state, on the one hand, the general case for infinite dimensions and, on the other hand, the (more realistic) case of the representability when the dimension is restricted. However, such representation theorems are hard to establish and out of scope of this dissertation. Finding such a representation theorem is in fact possible for a further restriction of the COS by considering only axis-aligned cones. As discussed in Theorem 5.18, this leads to the fact that each and only Boolean algebras can be represented.

An advantage of cone-COSs (and COSs in general) is their binary partial interpretation. Therefore, representing different levels of faithfulness is possible, e.g., for COSs based on axis-aligned cones, I showed that for classically satisfiable Boolean $\mathcal{ALC}$-ontologies there is a strongly concept-faithful and tbox-faithful al-cone-COS (see Proposition 5.19).

Level 3: Roles, Role Operators and Quantifiers

So far, I have considered COSs only in the propositional case, thus without incorporation of any relations. COSs with roles were tackled in Chapter 8. These are again highly dependent on the specific COS considered and therefore are discussed in Chapter 8 mainly based on cone-COSs. One possibility is using reification (for details, see Chapter 8) which lead to an embedding approach able to model expressive relations, fulfilling, e.g., partiality. This approach does not only allow for incorporation of relations in the abox but also in the tbox by allowing for statements of the form $\exists R.A \sqsubseteq B$, in fact, is able to model all $\mathcal{ALC}$ knowledge bases (see Proposition 8.9). When considering the cone-COS based on al-cones and a different technique of modeling relations, it is possible to model (theoretically) an arbitrary quantifier rank and even cycles. This was discussed in detail in Section 8.2. On the other hand, it is also possible to model non-distributive behavior, thus, e.g., $(\exists R.(C \sqcup D))^{\mathcal{I}} \neq (\exists R.C \sqcup \exists R.D)^{\mathcal{I}}$.

Level 4: Inference Services

The calculus is not only dependent on a specific COS but also on a specific embedding approach which is based on this COS. Therefore, here no statement is possible. However, as discussed in the last section, induction and deduction are conceivable.

11 Conclusion and Outlook

In the course of this dissertation, I showed that COSs are an option to answer the questions of the introduction positively. Or, to restate the working hypothesis: It is in fact possible to define a framework based on information about similarity and betweenness able to be the basis of an embedding approach modeling logical operators, especially negation, geometrically, namely the Conceptual Orthospaces.

11.1 Summary of Contributions

The contributions of this dissertation are summarized as follows:

The Conceptual Orthospace as foundational framework COSs open up an expressive opportunity for embeddings incorporating negation. They are able to represent the DL $\mathcal{ALC}$ and also non-distributive lattices. They are both usable for propositional embeddings and for embedding of roles and thus full $\mathcal{ALC}$ can be modeled. As COSs are based on a simple structure, they allow for general examinations on the expressivity of embedding approaches, e.g., I showed that an ortholattice is representable based on a non-trivial betweenness relation.

The importance of cones for embeddings in the Euclidean space Cone-based COSs occurred throughout the dissertation. In Chapter 7, I showed that this is not by chance but based on the result that cones are vital for embedding approaches that rely on the Euclidean space, are convex based on Euclidean betweenness and interpret negation based on the orthogonality relation. In fact, they are not only vital, but the only option.

Applications of Conceptual Orthospaces Though, COSs are a theoretical framework, it is possible to bring them into practice, as demonstrated for Zero-Shot Learning. Along this dissertation, many different COSs of different strength were proposed, not limited to cones. Thus, though many embedding approaches are based on Euclidean spaces and Euclidean convexity, this shows that there are more options to create an embedding approach.

Modeling relations with arbitrary quantifier depth I constructed an embedding approach able to model arbitrary quantifier depths faithfully. Though this needs an infinite dimensions, it enables to increase the size of the geometric model iteratively when needed and thus does not enforce retraining of the whole model when new data is given.

11.2 Future Work

Using the framework of COSs as basis, there are many possibilities, both theoretically and practically oriented, to advance further. Some of them are presented in the following.

Extending research in actual learning approaches Although this dissertation focused on the foundational and theoretical aspects of COSs, the short consideration on learning approaches based on COSs in Chapter 9 showed promising results and thus it seems to be worth to be investigated further, both based on different datasets but also based on other COSs or other/adapted learning approaches based on the COSs already defined.

Determining the logical commitments of learning approaches In this dissertation, logical commitments are defined and applied to some existing learning approaches. This can be extended to other approaches and extended further on the ones examined. One open question is for example the exact expressivity of the lattice based on closed convex cones. It is also possible to dive into more detail on the logical commitment of COSs as such.

The interpretation of the partial interpretation Another interesting topic is the exact identification of the partiality. In Section 3.2 some possible interpretations have been pointed out but a thorough discussion — in particular with discussion of approaches in (analytical) philosophy — of these and possibly adaptations of the COS to fit to the different interpretations is a topic of future work.

Identifying instances of COSs Perhaps the most important future work is to identify possible other instances of COSs. Though, several different COSs have been showcased, there is still a lot of potential, especially for not-distance based similarity relations or in general COSs not based on Euclidean spaces.

Changing the preconditions At last, it would be also interesting to extend the approach to other variants of negations, such as those making up Dunn's Kite of negation, e.g., dismissing double negation elimination (such as in intuitionistic negation).

Part III

Appendix

A Lattices

Lattices enable representing logical structures in an algebraic way. They can be used, e.g., to increase the understanding of logic, set theory, and projective geometry [Birkhoff, 1938]. Lattices represent an order structure based on joins and meets of objects. The interpretation of join and meet is context dependent and can be defined, e.g., as the least common multiple and the greatest common divisor for real numbers or based on set operations [Birkhoff, 1938].

Lattices are defined as follows: A *poset* $(L, \leq)$ is a pair of a domain L and a (partial) order relation $\leq$. A poset is a *lattice*, if every two elements have a supremum (a smallest upper bound, written $\sup(a, b) = a \vee b$), called *join*, and an infimum (a greatest lower bound, written $\inf(a, b) = a \wedge b$), called *meet* [Grätzer, 2011, p. 9].

A lattice is *bounded* iff it contains a smallest and a largest element (denoted by $\mathbb{0}$ and $\mathbb{1}$ respectively) [Grätzer, 2011, p. 10].

An element a is *covered by* b (written $a <: b$), if $a \leq b$, but there is no c with $c \neq a$, $c \neq b$ and $a \leq c \leq b$ [Grätzer, 2011, p. 6]. $a = b$ iff $a \leq b$ and $b \leq a$.

An *algebraic atom*[1] b in such a lattice is an element that covers $\mathbb{0}$, i.e., $\mathbb{0} < b$ holds and for all a with $\mathbb{0} \leq a \leq b$ either $a = \mathbb{0}$ or $a = b$. If the lattice is a Lindenbaum-Tarski algebra, we call "algebraic atom" also a representative of the equivalence class that is a algebraic atom in the lattice. The context makes clear whether we mean the equivalence class or a representative. A lattice is called *algebraically atomic* iff each element has an algebraic atom below it. Intuitively, in algebraic atomic lattices one excludes the possibility of having an infinitely descending chain whose infimum is $\mathbb{0}$ [Grätzer, 2011, p. 101].

Two elements $a, b \in L$ are *perspective* to an element $c \in L$ iff $a \vee c = b \vee c$ and $a \wedge c = b \wedge c = \mathbb{0}$ [Grätzer, 2011, p. 239].

Beside this general form of lattices, there are some special lattices, defined by the rules they fulfill:

The most common rule is distributivity. A lattice is called *distributive* iff for all $a, b, c \in L : a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$ [Padmanabhan and Rudeanu, 2008, p. 53].

[1] Note that in contrast to the usual wording, we add the specification "algebraic" in order to prevent clashes with the atomic concepts in DL.

An element a' is called the *complement* of a iff $a' \wedge a = \mathbb{0}$ and $a' \vee a = \mathbb{1}$ [Grätzer, 2011, p. 97]. Let $a \in [b, c]$. x is a *relative complement of a in* $[b, c]$ if $a \vee x = c$ and $a \wedge x = b$ [Grätzer, 2011, p. 97]. A bounded lattice L is called *complemented* if each element $a \in L$ has a complement and a bounded lattice L is called *relatively complemented* if each element $a \in L$ has a relative complement for each interval containing it [Grätzer, 2011, p. 98]. Each relatively complemented lattice is also complemented [Faure and Frölicher, 2000, p. 17]. The (relative) complement is not necessarily unique. A distributive, complemented lattice makes up a *Boolean algebra* [Padmanabhan and Rudeanu, 2008, p. 71].

A weaker notion is modularity. There are several ways of determining modularity. One is the following: A pair od elements (a, c) of a lattice L is called *modular* (written $M(a, c)$) iff $c \leq a$ implies that for all $b \in L$, $a \wedge (b \vee c) = (a \wedge b) \vee c$ [Grätzer, 2011, p. 335]. A lattice is called *M-symmetric* if $M(a, c)$ implies that $M(c, a)$ for every $a, c \in L$. A lattice is *modular* iff $M(a, c)$ for all $a, c \in L$ [Grätzer, 2011, p. 336].

An *orthomodular* lattice is defined as follows [Rédei, 1998, pp. 35–36]: If $a \leq b$ and $a^{'} \leq c$ then $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$. Here and in the following, $\cdot^{'}$ denotes the orthocomplement. The above rules were presented in decreasing strength, meaning for example that every distributive lattice is also modular and orthomodular.

A rule weaker than the ones introduced above is the ortholattice rule. An ortholattice is the lattice-equivalent to orthologic which is introduced in Chapter 3. A lattice is called an ortholattice iff an orthocomplement exists for which the following three conditions hold for all $a, b \in L$ [Chiara and Giuntini, 2002, pp. 11–12]:

- $a \leq b$ entails $b^{'} \leq a^{'}$ (antitonicity)
- $a^{''} = a$ (double negation elimination)
- $\mathbb{0} = a \wedge a^{'}$ (intuitionistic absurdity)

Due to these rules, an ortholattice satisfies De Morgan's laws, i.e., for any $a, b \in L$ it holds that $(a \wedge b)^{'} = a^{'} \vee b^{'}$ (and dually: $(a \vee b)^{'} = a^{'} \wedge b^{'}$). An orthomodular lattice is an ortholattice with some additional conditions. An ortholattice is distributive iff it fulfills the following rule (wLLJ): $a \wedge b \leq \mathbb{0}$, then $a \leq b^{'}$ [Padmanabhan and Rudeanu, 2008, axiom system B67, p. 114]. The rule is called (wLLJ) as it is a weakening of Johansson's constructive contraposition rule [Hartonas, 2016]. Next to this lattice rules, there are several others, they are introduced when needed, especially in Table 10.1.

Next to the order-based notion of an ortholattice there is an algebraic notion of an ortholattice. These notions are equivalent in the sense that each ortholattice induces an algebraic ortholattice and vice versa. A main difference comes with w.r.t. building substructures. As we need the notion of a subortholattice in the algebraic sense we give here the algebraic notions of ortholattice and subortholattice. An *(algebraic) ortholattice* is a set L with functions defined on it, namely a structure $(L, \wedge, \vee, \cdot', \mathbb{0}, \mathbb{1})$ fulfilling the following properties:

- $a \vee a = a$, $a \wedge a = a$. (idempotence)
- $a \vee b = b \vee a$, $a \wedge b = b \wedge a$. (commutativity)
- $(a \vee b) \vee c = a \vee (b \vee c)$, $(a \wedge b) \wedge c = a \wedge (b \wedge c)$ (associativity)
- $a \vee (a \wedge b) = a$, $a \wedge (a \vee b) = a$ (absorption)
- $a \wedge \mathbb{0} = \mathbb{0}$, $a \vee \mathbb{0} = a$, $a \vee \mathbb{1} = \mathbb{1}$, $a \wedge \mathbb{1} = a$
- $a'' = a$ (double negation elimination)
- $\mathbb{0} = a \wedge a'$ (intuitionistic absurdity)
- $(a \vee b)' = a' \wedge b'$, $(a \wedge b)' = a' \vee b'$ (De Morgan)

Given an ortholattice, its algebraic variant is defined as usual, e.g., defining $a \wedge b$ as the largest lower bound of a, b. And vice versa, given an algebraic ortholattice, the order of an ortholattice can be defined by setting $a \leq b$ iff $a \wedge b = a$. An ortholattice in an algebraic sense $(K, \wedge^K, \vee^K, \cdot'^K, \mathbb{0}^K, \mathbb{1}^K)$ is a *subortholattice* of the lattice $(L, \wedge^L, \vee^L, \cdot'^L, \mathbb{0}^L, \mathbb{1}^L)$ iff $K \subseteq L$ and $\mathbb{0}^K = \mathbb{0}^L$, $\mathbb{1}^K = \mathbb{1}^L$ and for all $a, b \in K$: $a \wedge^K b = a \wedge^L b \in K$ and $a \vee^K b = a \vee^L b \in K$ and $a'^K = a'^L \in K$.

A lattice can be converted into a logic an back by relying on the Lindenbaum-Tarski-Algebra. The elements of the lattice are interpreted as equivalence classes of formula and the ordering is based on derivability. Each letter representing a symbol of the logic is denoted with the respective lower case letter in the lattice.

Definition A.1. *[Restall, 2002, p. 183] The Lindenbaum-Tarski Algebra of a system is defined as follows:*

- *The elements are the sets* $[A] = \{B : A \dashv\vdash B\}$ *of equivalence classes of provably equivalent formulae.*
- *The ordering* $\leq$ *is defined by setting* $[A] \leq [B]$ *iff* $A \vdash B$.
- $[A] \vee [B] = [A \bigvee B]$, $[A] \wedge [B] = [A \,\&\, B]$ *and* $[A]' = [\rightharpoondown A]$.

Bibliography

Abboud, Ralph, Ismail Ceylan, Thomas Lukasiewicz, and Tommaso Salvatori (Dec. 2020). “BoxE: A box embedding model for knowledge base completion”. In: *NIPS’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, pp. 9649–9661.

Aboulker, Pierre, Xiaomin Chen, Guangda Huzhang, Rohan Kapadia, and Cathryn Supko (July 2016). “Lines, Betweenness and Metric Spaces”. In: *Discrete & Computational Geometry* 56.2, pp. 427–448. DOI: 10.1007/s00454-016-9806-2.

Alenda, Régis and Nicola Olivetti (2012). “Preferential Semantics for the Logic of Comparative Similarity over Triangular and Metric Models”. In: *European Conference on Logics in Artificial Intelligence*. Springer Berlin Heidelberg, pp. 1–13. DOI: https://doi.org/10.1007/978-3-642-33353-8_1.

Arrieta, Alejandro Barredo, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera (June 2020). “Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI”. In: *Information Fusion* 58, pp. 82–115. DOI: 10.1016/j.inffus.2019.12.012.

Artale, Alessandro, Diego Calvanese, Roman Kontchakov, and Michael Zakharyaschev (Oct. 2009). “The DL-Lite Family and Relations”. In: *Journal of artificial intelligence research* 36, pp. 1–69. DOI: 10.1613/jair.2820.

Artstein-Avidan, Shiri and Vitali Milman (May 2008). “The Concept of Duality for Measure Projections of Convex Bodies”. In: *Journal of Functional Analysis* 254.10, pp. 2648–2666. DOI: 10.1016/j.jfa.2007.11.008.

Artstein-Avidan, Shiri and Boaz A. Slomka (Sept. 2011). “Order Isomorphisms in Cones and a Characterization of Duality for Ellipsoids”. In: *Selecta Mathematica* 18.2, pp. 391–415. DOI: 10.1007/s00029-011-0069-8.

Aspis, Yaniv, Krysia Broda, Jorge Lobo, and Alessandra Russo (July 2022). “Embed2Sym - Scalable Neuro-Symbolic Reasoning via Clustered Embeddings”. In: *Proceedings of the Nineteenth International Conference on Principles of Knowledge Representation and Reasoning*. Ed. by Gabriele Kern-Isberner, Gerhard Lakemeyer, and

Thomas Meyer. Vol. 19. 1. International Joint Conferences on Artificial Intelligence Organization, pp. 421–431. DOI: 10.24963/kr.2022/44.

Awad, Mariette and Rahul Khanna (2015). "Support Vector Regression". In: *Efficient Learning Machines*. Berkeley, CA: Apress, pp. 67–80. DOI: 10.1007/978-1-4302-5990-9_4.

Baader, Franz, Sebastian Brandt, and Carsten Lutz (2005). "Pushing the EL Envelope". In: *IJCAI'05: Proceedings of the 19th international joint conference on Artificial intelligence*, pp. 364–369.

Baader, Franz, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, eds. (Aug. 2007). *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press. DOI: 10.1017/CBO9780511711787.

Baader, Franz and Ralf Küsters (2006). "Nonstandard Inferences in Description Logics: The Story So Far". In: *Mathematical Problems from Applied Logic I*. Ed. by D.M. Gabbay, S.S. Goncharov, and M. Zakharyaschev. Vol. 4. International Mathematical Series. Springer New York, pp. 1–75. DOI: 10.1007/0-387-31072-X_1.

Bader, Sebastian and Pascal Hitzler (2005). "Dimensions of Neural-symbolic Integration - A Structured Survey". In: *ArXiv* abs/cs/0511042. DOI: 10.48550/arXiv.cs/0511042.

Badreddine, Samy, Artur S. d'Avila Garcez, Luciano Serafini, and Mike Spranger (Feb. 2022). "Logic Tensor Networks". In: *Artificial Intelligence* 303, p. 103649. DOI: 10.1016/j.artint.2021.103649.

Bai, Yushi, Zhitao Ying, Hongyu Ren, and Jure Leskovec (2021). "Modeling heterogeneous hierarchies with relation-specific hyperbolic cones". In: *Advances in Neural Information Processing Systems* 34, pp. 12316–12327.

Balbiani, Philippe, Valentin Goranko, Ruaan Kellerman, and Dimiter Vakarelov (2007). "Logical Theories for Fragments of Elementary Geometry". In: *Handbook of Spatial Logics*. Ed. by Marco Aiello, Ian Pratt-Hartmann, and Johan Van Benthem. Dordrecht: Springer Netherlands, pp. 343–428. DOI: 10.1007/978-1-4020-5587-4_7.

Bechberger, Lucas and Kai-Uwe Kühnberger (2017). "A Thorough Formalization of Conceptual Spaces". In: *KI 2017: Advances in Artificial Intelligence*. Springer International Publishing, pp. 58–71. DOI: 10.1007/978-3-319-67190-1_5.

Bianchi, Federico, Gaetano Rossiello, Luca Costabello, Matteo Palmonari, and Pasquale Minervini (2020). "Knowledge Graph Embeddings and Explainable AI". In: *Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges*. Ed. by Ilaria Tiddi, Freddy Lécué, and Pascal Hitzler. Vol. 47. Studies on the Semantic Web. IOS Press, pp. 49–72. DOI: 10.3233/SSW200011.

Birkhoff, Garrett (1938). “Lattices and their Applications”. In: *Bulletin of the American Mathematical Society* 44.12, pp. 793–800. DOI: 10.1090/S0002-9904-1938-06866-8.

Birkhoff, Garrett and John Neumann (1936). “The Logic of Quantum Mechanics”. In: *Annals of Mathematics* 37, pp. 823–843.

Birkhoff, Garrett and John Von Neumann (1975). “The Logic of Quantum Mechanics”. In: *The Logico-Algebraic Approach to Quantum Mechanics: Volume I: Historical Evolution*. Ed. by C. A. Hooker. Dordrecht: Springer Netherlands, pp. 1–26. DOI: 10.1007/978-94-010-1795-4_1.

Bishop, Christopher M. (2006). *Pattern Recognition and Machine Learning*. 1st ed. Information Science and Statistics. Springer New York, NY. ISBN: 978-0-387-31073-2.

Bonini, Nicolao, Daniel N. Osherson, Riccardo Viale, and Timothy Luke Williamson (Dec. 1999). “On the Psychology of Vague Predicates”. In: *Mind & Language* 14.4, pp. 377–393. DOI: 10.1111/1468-0017.00117.

Bordes, Antoine, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko (Dec. 2013). “Translating Embeddings for Modeling Multi-relational Data”. In: *NIPS’13: Proceedings of the 26th International Conference on Neural Information Processing Systems*, pp. 2787–2795.

Borgwardt, Stefan and Rafael Peñaloza (2017). “Fuzzy Description Logics – A Survey”. In: *Scalable Uncertainty Management*. Ed. by Serafín Moral, Olivier Pivert, Daniel Sánchez, and Nicolás Marín. Lecture Notes in Computer Science 10564. Cham: Springer International Publishing, pp. 31–45. DOI: 10.1007/978-3-319-67582-4_3.

Borsuk, Karol and Wanda Szmielew (2018). *Foundations of Geometry: Euclidean, Bolyai-Lobachevskian, and Projective Geometry*. Dover Books on Mathematics. Mineola, New York: Dover Publications.

Bouraoui, Zied, Víctor Gutiérrez-Basulto, and Steven Schockaert (2022). “Integrating Ontologies and Vector Space Embeddings Using Conceptual Spaces”. In: *International Research School in Artificial Intelligence in Bergen (AIB 2022)*. Ed. by Camille Bourgaux, Ana Ozaki, and Rafael Peñaloza. Vol. 99. Open Access Series in Informatics (OASIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 3:1–3:30. DOI: 10.4230/OASIcs.AIB.2022.3.

Burigana, Luigi (2009). “Tree Representations of Betweenness relations Defined by Intersection and Inclusion”. In: *Mathematics and Social Sciences* 185.1, pp. 5–36.

Chang, Chen Chung and Howard Jerome Keisler (1990). *Model Theory*. Vol. 73. Studies in Logic and the Foundations of Mathematics. Elsevier Science. ISBN: 9780080880075.

Changpinyo, Soravit, Wei-Lun Chao, and Fei Sha (Oct. 2017). “Predicting Visual Exemplars of Unseen Classes for Zero-Shot Learning”. In: *2017 IEEE International*

Conference on Computer Vision (ICCV). IEEE, pp. 3496–3505. DOI: 10.1109/ICCV.2017.376.

Chiara, Maria Luisa Dalla and Roberto Giuntini (2002). “Quantum Logics”. In: *Handbook of Philosophical Logic*. Ed. by Dov M. Gabbay and F. Guenthner. Vol. 6. Dordrecht: Springer Netherlands, pp. 129–228. DOI: 10.1007/978-94-017-0460-1_2.

Chvátal, Vašek (2009). “Antimatroids, Betweenness, Convexity”. In: *Research Trends in Combinatorial Optimization*. Ed. by William Cook, László Lovász, and Jens Vygen. Springer Berlin Heidelberg, pp. 57–64. DOI: 10.1007/978-3-540-76796-1_3.

Cristianini, Nello and John Shawe-Taylor (Mar. 2000). *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press. DOI: 10.1017/CBO9780511801389.

Davis, Randall, Howard Shrobe, and Peter Szolovits (1993). “What is a Knowledge Representation?” In: *AI Magazine* 14.1, pp. 17–33. DOI: 10.1609/aimag.v14i1.1029.

Demos, Raphael (Apr. 1917). “A Discussion of a Certain Type of Negative Proposition”. In: *Mind* 26.102, pp. 188–196. DOI: 10.1093/mind/xxvi.1.188.

Derrac, Joaquín and Steven Schockaert (Nov. 2015). “Inducing semantic relations from conceptual spaces: a data-driven approach to plausible reasoning”. In: *Artificial Intelligence* 228, pp. 66–94. DOI: 10.1016/j.artint.2015.07.002.

Dilworth, R. P. (1945). “Lattices With Unique Complements”. In: *Transactions of the American Mathematical Society* 57.1, pp. 123–154.

Dunn, Jon Michael (Mar. 1996). “Generalized Ortho Negation”. In: *Negation: a notion in focus*. Ed. by Heinrich Wansing. Vol. 7. Perspektiven der Analytischen Philosophie / Perspectives in Analytical Philosophy. Walter de Gruyter Berlin, pp. 3–26. DOI: 10.1515/9783110876802.3.

Farkas, Julius (Jan. 1902). “Theorie der einfachen Ungleichungen.” In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 124, pp. 1–27. DOI: 10.1515/crll.1902.124.1.

Faure, Claude-Alain and Alfred Frölicher (2000). *Modern Projective Geometry*. Vol. 521. Mathematics and Its Applications. Springer Dordrecht. DOI: 10.1007/978-94-015-9590-2.

Flouris, Giorgos, Zhisheng Huang, Jeff Z. Pan, Dimitris Plexousakis, and Holger Wache (2006). “Inconsistencies, Negations and Changes in Ontologies”. In: *AAAI’06: proceedings of the 21st national conference on Artificial intelligence* 21.2, pp. 1295–1300.

Garcez, Artur d’Avila and Luís C. Lamb (Mar. 2023). “Neurosymbolic AI: the 3rd wave”. In: *Artificial Intelligence Review* 56.11, pp. 12387–12406. DOI: 10.1007/s10462-023-10448-w.

Gärdenfors, Peter (2000). *Conceptual Spaces: The Geometry of Thought*. Cambridge, Massachusetts: The MIT Press. DOI: 10.7551/mitpress/2076.001.0001.

— (2014). *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. The MIT Press. DOI: 10.7551/mitpress/9629.001.0001.

Garg, Dinesh, Shajith Ikbal, Santosh K. Srivastava, Harit Vishwakarma, Hima Karanam, and L Venkata Subramaniam (2019). "Quantum Embedding of Knowledge for Reasoning". In: *NIPS'19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, pp. 5594–5604.

Geeraerts, Dirk (Sept. 2006). "Prototype theory". In: *Cognitive Linguistics: Basic Readings*. Mouton de Gruyter. Chap. 4, pp. 141–165. DOI: 10.1515/9783110199901.141.

Gelfond, Michael and Nicola Leone (June 2002). "Logic programming and knowledge representation—The A-Prolog perspective". In: *Artificial Intelligence* 138.1-2. Knowledge Representation and Logic Programming, pp. 3–38. DOI: 10.1016/S0004-3702(02)00207-2.

Gibaja, Eva and Sebastián Ventura (Nov. 2014). "Multi-label learning: a review of the state of the art and ongoing research". In: *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 4.6, pp. 411–444. DOI: 10.1002/widm.1139.

Godo, Lluís and Ricardo O. Rodríguez (2008). "Logical approaches to fuzzy similarity-based reasoning: an overview". In: *Preferences and Similarities*. Ed. by G. Della Riccia, D. Dubois, R. Kruse, and HJ. Lenz. Vol. 504. Springer Vienna, pp. 75–128. DOI: 10.1007/978-3-211-85432-7_4.

Goldberg, Yoav and Omer Levy (2014). "Word2vec Explained: Deriving Mikolov et al.'s Negative-sampling Word-embedding Method". In: *ArXiv* abs/1402.3722. DOI: 10.48550/arXiv.1402.3722.

Goldblatt, R. I. (1974). "Semantic analysis of orthologic". In: *Journal of Philosophical Logic* 3.1-2, pp. 19–35. DOI: 10.1007/bf00652069.

Goldstone, Robert L. and Ji Yun Son (2005). "Similarity". In: *Cambridge Handbook of Thinking and Reasoning*. Ed. by Keith J. Holyoak and Robert G. Morrison. Cambridge University Press, pp. 13–36.

Goodman, Nelson (1977). *The Structure of Appearance*. Boston Studies in the Philosophy of Science. Springer Netherlands. DOI: 10.1007/978-94-010-1184-6.

Grätzer, George (2011). *Lattice Theory: Foundation*. Springer Basel. DOI: 10.1007/978-3-0348-0018-1.

Gruber, Peter M. (Dec. 1991). "The endomorphisms of the lattice of convex bodies". In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 61.1, pp. 121–130. DOI: 10.1007/BF02950756.

Gruber, Thomas R. (June 1993). "A translation approach to portable ontology specifications". In: *Knowledge Acquisition* 5.2, pp. 199–220. DOI: 10.1006/KNAC.1993.1008.

— (Nov. 1995). "Toward Principles for the Design of Ontologies Used for Knowledge Sharing?" In: *International Journal of Human-Computer Studies* 43.5, pp. 907–928. DOI: 10.1006/ijhc.1995.1081.

Guarino, Nicola, Daniel Oberle, and Steffen Staab (2009). "What Is an Ontology?" In: *Handbook on Ontologies*. Ed. by Steffen Staab and Rudi Studer. Springer Berlin Heidelberg, pp. 1–17. DOI: 10.1007/978-3-540-92673-3_0.

Guo, Shu, Quan Wang, Bin Wang, Lihong Wang, and Li Guo (2015). "Semantically Smooth Knowledge Graph Embedding". In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, pp. 84–94. DOI: 10.3115/v1/P15-1009.

Gutiérrez-Basulto, Víctor and Steven Schockaert (2018). "From Knowledge Graph Embedding to Ontology Embedding? An Analysis of the Compatibility between Vector Space Representations and Rules". In: *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October – 2 November 2018*. Ed. by Michael Thielscher, Francesca Toni, and Frank Wolter. AAAI Press, pp. 379–388.

Hamkins, Joel David and Makoto Kikuchi (May 2016). "Set-theoretic mereology". In: *Logic and Logical Philosophy* 25.3. Ed. by A. C. Varzi and R. Gruszczyński, pp. 285–308. DOI: 10.12775/llp.2016.007.

Hampton, James A. (Nov. 1988a). "Disjunction of natural concepts". In: *Memory & Cognition* 16.6, pp. 579–591. DOI: 10.3758/BF03197059.

— (Jan. 1988b). "Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion". In: *Journal of Experimental Psychology: Learning, Memory and Cognition* 14.1, pp. 12–32. DOI: 10.1037/0278-7393.14.1.12.

Hardy, Michael (Aug. 2002). "Scaled Boolean algebras". In: *Advances in Applied Mathematics* 29.2, pp. 243–292. DOI: 10.1016/S0196-8858(02)00011-8.

Harnad, Stevan (June 1990). "The symbol grounding problem". In: *Physica D: Nonlinear Phenomena* 42.1-3, pp. 335–346. DOI: 10.1016/0167-2789(90)90087-6.

Hartonas, Chrysafis (2016). "Reasoning with Incomplete Information in Generalized Galois Logics Without Distribution: The Case of Negation and Modal Operators". In: *J. Michael Dunn on Information Based Logics*. Ed. by Katalin Bimbó. Outstanding Contributions to Logic 8. Springer International Publishing, pp. 279–312. DOI: 10.1007/978-3-319-29300-4_14.

Herre, Heinrich, Jan Jaspars, and Gerd Wagner (1999). "Partial Logics with Two Kinds of Negation as a Foundation for Knowledge-Based Reasoning". In: *What is negation?* Ed. by Dov Gabbay and Heinrich Wansing. Applied Logic Series 13. Springer Netherlands, pp. 121–159. DOI: 10.1007/978-94-015-9309-0_7.

Hilbert, David (1913). *Grundlagen der Geometrie*. Wissenschaft und Hypothese 7. Leipzig, Berlin: B. G. Teubner.

Hogan, Aidan, Eva Blomqvist, Michael Cochez, Claudia D'Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann (July 2021). "Knowledge Graphs". In: *ACM Computing Surveys (CSUR)* 54.4, pp. 1–37. DOI: 10.1145/3447772.

Horn, Laurence R. (2001). *A natural history of negation*. The David Hume Series. CSLI Publications. ISBN: 1-57586-336-7.

Hume, David (1894). *An Enquiry Concerning the Human Understanding, and an Enquiry Concerning the Principles of Morals*. Ed. by L.A. Selby-Bigge. Oxford Clarendon Press.

Huntington, Edward V. and J. Robert Kline (1917). "Sets of independent postulates for betweenness". In: *Transactions of the American Mathematical Society* 18.3, pp. 301–325. DOI: 10.2307/1988957.

Ibáñez-García, Yazmín Angélica, Víctor Gutiérrez-Basulto, and Steven Schockaert (July 2020). "Plausible Reasoning about EL-Ontologies using Concept Interpolation". In: *Proceedings of the Seventeenth International Conference on Principles of Knowledge Representation and Reasoning*. Ed. by Diego Calvanese, Esra Erdem, and Michael Thielscher. International Joint Conferences on Artificial Intelligence Organization, pp. 506–516. DOI: 10.24963/kr.2020/51.

Jackermeier, Mathias, Jiaoyan Chen, and Ian Horrocks (2023). "Box2EL: Concept and Role Box Embeddings for the Description Logic EL++". In: *ArXiv*. DOI: 10.48550/ARXIV.2301.11118.

Jameel, Shoaib and Steven Schockaert (2016). "Entity Embeddings with Conceptual Subspaces as a Basis for Plausible Reasoning". In: *Proceedings of the Twenty-Second European Conference on Artificial Intelligence*. ECAI'16. The Hague, The Netherlands: IOS Press, pp. 1353–1361. DOI: 10.3233/978-1-61499-672-9-1353.

Jia, Yuting, Shao Zhang, Haiwen Wang, Ying Wen, Luoyi Fu, Huan Long, Xinbing Wang, and Chenghu Zhou (Aug. 2022). "Investigating the geometric structure of neural activation spaces with convex hull approximations". In: *Neurocomputing*. Vol. 499. Elsevier BV, pp. 93–105. DOI: 10.1016/j.neucom.2022.05.019.

Jónsson, Bjarni and Alfred Tarski (Oct. 1951). "Boolean Algebras with Operators. Part I". In: *American Journal of Mathematics* 73.4, pp. 891–939. DOI: 10.2307/2372123.

— (Jan. 1952). "Boolean Algebras with Operators. Part II". In: *American Journal of Mathematics* 74.1, pp. 127–162. DOI: 10.2307/2372074.

Kahneman, Daniel (2011). *Thinking, Fast and Slow.* Farrar, Straus and Giroux. ISBN: 978-0374275631.

Kautz, Henry A. (Mar. 2022). "The third AI summer: AAAI Robert S. Engelmore Memorial Lecture". In: *AI Magazine* 43.1, pp. 105–125. DOI: 10.1002/aaai.12036.

Kazemi, Seyed Mehran and David Poole (2018). "SimplE Embedding for Link Prediction in Knowledge Graphs". In: *NIPS'18: Proceedings of the 32nd International Conference on Neural Information Processing Systems.* Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Vol. 31. Curran Associates, Inc., pp. 4289–4300.

Krötzsch, Markus, František Simančík, and Ian Horrocks (2012). "A Description Logic Primer". In: *arXiv* abs/1201.4089. DOI: 10.48550/arXiv.1201.4089.

Kulmanov, Maxat, Wang Liu-Wei, Yuan Yan, and Robert Hoehndorf (Aug. 2019). "EL Embeddings: Geometric Construction of Models for the Description Logic EL++". In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19.* International Joint Conferences on Artificial Intelligence Organization, pp. 6103–6109. DOI: 10.24963/ijcai.2019/845.

Leemhuis, Mena and Özgür L. Özçep (2022). "A Goodman-style Betweenness Relation on Orthoframes". In: *Proceedings of the 8th Workshop on Formal and Cognitive Reasoning, co-located with the 45th German Conference on Artificial Intelligence (KI 2022).* Ed. by Christoph Beierle, Marco Ragni, Kai Sauerwald, Frieder Stolzenburg, and Matthias Thimm.

— (Nov. 2023). "Conceptual Orthospaces—Convexity Meets Negation". In: *International Journal of Approximate Reasoning* 162, p. 109013. DOI: 10.1016/j.ijar.2023.109013.

Leemhuis, Mena, Özgür L. Özçep, and Diedrich Wolter (2020). "Multi-label Learning with a Cone-Based Geometric Model". In: *Proceedings of the 25th International Conference on Conceptual Structures (ICCS 2020).* Springer International Publishing, pp. 177–185. DOI: 10.1007/978-3-030-57855-8_13.

— (2022a). "Knowledge Graph Embeddings with Ontologies: Reification for Representing Arbitrary Relations". In: *German Conference on Artificial Intelligence (Künstliche Intelligenz).* Ed. by R. Bergmann, L. Malburg, S.C. Rodermund, and I.J. Timm. Lecture Notes in Computer Science 13404. Springer International Publishing, pp. 146–159. DOI: 10.1007/978-3-031-15791-2_13.

— (Oct. 2022b). “Learning with cone-based geometric models and orthologics”. In: *Annals of Mathematics and Artificial Intelligence* 90, pp. 1–37. DOI: 10.1007/s10472-022-09806-1.

Lin, Yankai, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu (2015). “Learning Entity and Relation Embeddings for Knowledge Graph Completion”. In: *AAAI'15: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Austin, Texas: AAAI Press, pp. 2181–2187.

Marcus, Gary (2020). “The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence”. In: *arXiv*. DOI: 10.48550/arXiv.2002.06177.

McCarthy, John and Patrick J. Hayes (1969). “Some Philosophical Problems from the Standpoint of Artificial Intelligence”. In: *Machine Intelligence 4*. Ed. by B. Meltzer and D. Michie. Edinburgh University Press, pp. 463–502.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean (2013). “Distributed Representations of Words and Phrases and Their Compositionality”. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems*. NIPS'13 2, pp. 3111–3119.

Mondal, Sutapa, Sumit Kaur Bhatia, and Raghava Mutharaju (2021). “EmEL++: Embeddings for EL++ Description Logic”. In: *Proceedings of the AAAI 2021 Spring Symposium on Combining Machine Learning and Knowledge Engineering (AAAI-MAKE 2021)*. Ed. by Andreas Martin, Knut Hinkelmann, Hans-Georg Fill, Aurona Gerber, Doug Lenat, Reinhard Stolle, and Frank van Harmelen. Vol. 2846. CEUR Workshop Procedings.

Morante, Roser and Eduardo Blanco (Dec. 2020). “Recent advances in processing negation”. In: *Natural Language Engineering* 27.2, pp. 121–130. DOI: 10.1017/S1351324920000534.

Nickel, Maximilian, Volker Tresp, and Hans-Peter Kriegel (2011). “A Three-way Model for Collective Learning on Multi-relational Data”. In: *Proceedings of the 28th International Conference on International Conference on Machine Learning*. Ed. by Lise Getoor and Tobias Scheffer. ICML'11. Bellevue, Washington, USA: Omnipress, pp. 809–816. ISBN: 978-1-4503-0619-5.

Odintsov, Serguei P. and Heinrich Wansing (2003). “Inconsistency-tolerant Description Logic: Motivation and Basic Systems”. In: *Trends in Logic*. Ed. by Vincent F. Hendricks and Jacek Malinowski. Springer Netherlands, pp. 301–335. DOI: 10.1007/978-94-017-3598-8_11.

Osherson, Daniel N and Edward E Smith (Jan. 1981). “On the adequacy of prototype theory as a theory of concepts”. In: *Cognition* 9.1, pp. 35–58. DOI: 10.1016/0010-0277(81)90013-5.

Özçep, Özgür L., Mena Leemhuis, and Diedrich Wolter (July 2020). "Cone Semantics for Logics with Negation". In: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*. International Joint Conferences on Artificial Intelligence Organization, pp. 1820–1826. DOI: 10.24963/ijcai.2020/252.

— (Oct. 2023). "Embedding Ontologies in the description logic ALC by Axis-Aligned Cones". In: *Journal of Artificial Intelligence Research* 78, pp. 217–267. DOI: 10.1613/jair.1.13939.

Padmanabhan, R and S Rudeanu (Aug. 2008). *Axioms for Lattices and Boolean Algebras*. World Scientific. DOI: 10.1142/7007.

Peñaloza, Rafael and Tingting Zou (2013). "Roughening the EL Envelope". In: *Frontiers of Combining Systems*. Ed. by P. Fontaine, C. Ringeissen, and R. A. Schmidt. Lecture Notes in Computer Science 8152. Springer Berlin Heidelberg, pp. 71–86. DOI: 10.1007/978-3-642-40885-4_6.

Peng, Xi, Zhenwei Tang, Maxat Kulmanov, Kexin Niu, and Robert Hoehndorf (2022). "Description Logic EL++ Embeddings with Intersectional Closure". In: *ArXiv* abs/2202.14018. DOI: 10.48550/arXiv.2202.14018.

Pitcher, Everett and Malcolm F. Smiley (July 1942). "Transitivities of Betweenness". In: *Transactions of the American Mathematical Society* 52.1, pp. 95–114. DOI: 10.2307/1990155.

Podgorny, Peter and W. R. Garner (Jan. 1979). "Reaction time as a measure of inter- and intraobject visual similarity: Letters of the alphabet". In: *Perception & Psychophysics* 26.1, pp. 37–52. DOI: 10.3758/BF03199860.

Priest, Graham (Apr. 2008). *An Introduction to Non-Classical Logic*. Cambridge University Press. DOI: 10.1017/cbo9780511801174.

Qian, Jing, Gangmin Li, Katie Atkinson, and Yong Yue (Jan. 2021). "Understanding Negative Sampling in Knowledge Graph Embedding". In: *International Journal of Artificial Intelligence & Applications* 12.1, pp. 71–81. DOI: 10.5121/ijaia.2021.12105.

Raedt, Luc De, Robin Manhaeve, Sebastijan Dumancic, Thomas Demeester, and Angelika Kimmig (Aug. 2019). "Neuro-Symbolic = Neural + Logical + Probabilistic". In: *Proceedings of the 2019 International Workshop on Neural-Symbolic Learning and Reasoning (NeSy 2019), Annual workshop of the Neural-Symbolic Learning and Reasoning Association, Macao, China*. Ed. by Derek Doran, Artur S. d'Avila Garcez, and Freddy Lécué.

Rautenbach, Dieter, Vinícius Fernandes dos Santos, Philipp M. Schäfer, and Jayme L. Szwarcfiter (Mar. 2011). "Characterization and representation problems for intersection betweennesses". In: *Discrete Applied Mathematics* 159.5, pp. 389–395. DOI: 10.1016/j.dam.2010.12.009.

Rédei, Miklós (1998). *Quantum Logic in Algebraic Approach*. Fundamental Theories of Physics. Springer Dordrecht. DOI: 10.1007/978-94-015-9026-6.

Ren, Hongyu, Weihua Hu, and Jure Leskovec (2020). "Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings". In: *arXiv*. DOI: 10.48550/arXiv.2002.05969.

Ren, Kui, Tianhang Zheng, Zhan Qin, and Xue Liu (Mar. 2020). "Adversarial Attacks and Defenses in Deep Learning". In: *Engineering* 6.3, pp. 346–360. DOI: 10.1016/j.eng.2019.12.012.

Restall, Greg (2002). *An Introduction to Substructural Logics*. Routledge.

Restle, Frank (Sept. 1959). "A metric and an ordering on sets". In: *Psychometrika* 24.3, pp. 207–220. DOI: 10.1007/BF02289843.

Rutten, J. J. M. M. (Feb. 2005). "A coinductive calculus of streams". In: *Mathematical Structures in Computer Science* 15.1, pp. 93–147. DOI: 10.1017/S0960129504004517.

Saeed, Nasir, Haewoon Nam, Mian Imtiaz Ul Haq, and Dost Bhatti Muhammad Saqib (May 2018). "A Survey on Multidimensional Scaling". In: *ACM Computing Surveys* 51.3, pp. 1–25. DOI: 10.1145/3178155.

Schild, Klaus (Aug. 1991). "A Correspondence Theory for Terminological Logics: Preliminary Report". In: *IJCAI'91: Proceedings of the 12th international joint conference on Artificial intelligence*. Ed. by John Mylopoulos and Ray Reiter. Vol. 1, pp. 466–471.

Schneider, Rolf (2008). "The endomorphisms of the lattice of closed convex cones". In: *Contributions to Algebra and Geometry* 49.2, pp. 541–547.

— (2022). *Convex Cones - Geometry and Probability*. Vol. 2319. Lecture Notes in Mathematics. Springer Cham. DOI: 10.1007/978-3-031-15127-9.

Schockaert, Steven (2021). "Modelling Monotonic and Non-Monotonic Attribute Dependencies with Embeddings: A Theoretical Analysis". In: *3rd Conference on Automated Knowledge Base Construction*. Automated Knowledge Base Construction (AKBC). DOI: 10.24432/C5GW2Z.

Schockaert, Steven and Víctor Gutiérrez-Basulto (2021). "Modelling Symbolic Knowledge Using Neural Representations". In: *Reasoning Web. Declarative Artificial Intelligence : 17th International Summer School 2021, Leuven, Belgium*. Springer, pp. 59–75. DOI: 10.1007/978-3-030-95481-9_3.

Schockaert, Steven, Yazmín Angélica Ibáñez-García, and Víctor Gutiérrez-Basulto (2021). "Modelling Concept Interpolation in Description Logics using Abstract Betweenness Relations". In: *Proceedings of the 34th International Workshop on De-*

scription Logics (DL 2021) part of Bratislava Knowledge September (BAKS 2021), Bratislava, Slovakia, September 19th to 22nd, 2021. Ed. by Martin Homola, Vladislav Ryzhikov, and Renate A. Schmidt. Vol. 2954. CEUR Workshop Proceedings.

Sheremet, Mikhail, Dmitry Tishkovsky, Frank Wolter, and Michael Zakharyaschev (Mar. 2007). "A Logic for Concepts and Similarity". In: *Journal of Logic and Computation* 17.3, pp. 415–452. DOI: 10.1093/logcom/exm007.

Simons, Peter M. (1987). *Parts: A Study in Ontology*. Oxford University Press. ISBN: 9780199241460.

Smolensky, Paul (Mar. 1988). "On the proper treatment of connectionism". In: *Behavioral and Brain Sciences* 11.1, pp. 1–74. DOI: 10.1017/S0140525X00052432.

Smyth, Michael B. and Julian Webster (2007). "Discrete Spatial Models". In: *Handbook of Spatial Logics*. Ed. by Marco Aiello, Ian Pratt-Hartmann, and Johan Van Benthem. Springer Netherlands. Chap. 12, pp. 713–798. DOI: 10.1007/978-1-4020-5587-4_12.

Staab, Steffen and Rudi Studer, eds. (2009). *Handbook on Ontologies*. Springer Berlin Heidelberg. DOI: 10.1007/978-3-540-92673-3.

Stern, Manfred (1991). *Semimodular Lattices*. Vieweg+Teubner Verlag. DOI: 10.1007/978-3-663-12478-8.

Sun, Zhiqing, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang (2019). "RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space". In: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.

Tarski, Alfred (1935). "Grundzüge des Systemenkalküls. Erster Teil". In: *Fundamenta Mathematicae* 25, pp. 503–526. DOI: 10.4064/FM-25-1-503-526.

Tversky, Amos (July 1977). "Features of similarity". In: *Psychological Review* 84.4, pp. 327–352. DOI: 10.1037/0033-295x.84.4.327.

Tversky, Amos and Itamar Gati (1982). "Similarity, separability, and the triangle inequality." In: *Psychological Review* 89.2, pp. 123–154. DOI: 10.1037/0033-295X.89.2.123.

Valiant, Leslie G. (Jan. 2003). "Three Problems in Computer Science". In: *Journal of the ACM* 50.1, pp. 96–99. DOI: 10.1145/602382.602410.

Van Gelder, Allen, Kenneth A Ross, and John S Schlipf (July 1991). "The well-founded semantics for general logic programs". In: *Journal of the ACM (JACM)* 38.3, pp. 619–649. DOI: 10.1145/116825.116838.

Wang, Quan, Zhendong Mao, Bin Wang, and Li Guo (Dec. 2017). "Knowledge Graph Embedding: A Survey of Approaches and Applications". In: *IEEE Transactions on Knowledge and Data Engineering* 29.12, pp. 2724–2743. DOI: 10.1109/tkde.2017.2754499.

Wang, Yanjie, Rainer Gemulla, and Hui Li (Feb. 2018). "On Multi-Relational Link Prediction with Bilinear Models". In: *AAAI'18/IAAI'18/EAAI'18: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*. New Orleans, Louisiana, USA: AAAI Press, pp. 4227–4234.

Wang, Zhen, Jianwen Zhang, Jianlin Feng, and Zheng Chen (2014). "Knowledge Graph Embedding by Translating on Hyperplanes". In: *AAAI'14: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. AAAI Press, pp. 1112–1119.

Williamson, Timothy (1994). *Vagueness*. London: Routledge. DOI: 10.4324/9780203014264.

Xian, Yongqin, Christoph H. Lampert, Bernt Schiele, and Zeynep Akata (Sept. 2019). "Zero-Shot Learning—A Comprehensive Evaluation of the Good, the Bad and the Ugly". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41.9, pp. 2251–2265. DOI: 10.1109/TPAMI.2018.2857768.

Xie, Ruobing, Zhiyuan Liu, and Maosong Sun (July 2016). "Representation Learning of Knowledge Graphs with Hierarchical Types". In: *IJCAI'16: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pp. 2965–2971.

Xiong, Bo, Nico Potyka, Trung-Kien Tran, Mojtaba Nayyeri, and Steffen Staab (2022). "Faithful Embeddings for EL++ Knowledge Bases". In: *The Semantic Web – ISWC 2022*. Springer International Publishing, pp. 22–38. DOI: 10.1007/978-3-031-19433-7_2.

Xu, Yun, Jin Li, and Gangsong Leng (2023). "Dualities and endomorphisms of pseudo-cones". In: *Advances in Applied Mathematics* 142, p. 102434. DOI: 10.1016/j.aam.2022.102434.

Yablo, Stephen (2014). *Aboutness*. Carl G. Hempel Lecture Series. Princeton University Press. ISBN: 9780691144955.

Zadeh, Lotfi A. (May 1996). "Fuzzy Sets". In: *Advances in Fuzzy Systems - Applications and Theory*. WORLD SCIENTIFIC, pp. 19–34. DOI: 10.1142/9789814261302_0001.

Zhang, Zhanqiu, Jie Wang, Jiajun Chen, Shuiwang Ji, and Feng Wu (2021). "ConE: Cone Embeddings for Multi-Hop Reasoning over Knowledge Graphs". In: *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Curran Associates, Inc., pp. 19172–19183.

Zhou, Bolei, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba (June 2016). “Learning Deep Features for Discriminative Localization”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2921–2929. DOI: 10.1109/CVPR.2016.319.