UNIVERSITÄT ZU LÜBECK

From the Institute of Information Systems
of the University of Lübeck
Director: Prof. Dr.-Ing. Nele Rußwinkel

Hierarchies of Action Control: Modeling Sensorimotor and Cognitive Mechanisms of Human Adaptation

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
Nils Wendel Heinrich
from Bonn

Lübeck 2026

First referee: Prof. Dr. Nele Rußwinkel
Second referee: Prof. Dr. Eva Wiese
Date of oral examination: 17.04.2026
Approved for printing. Lübeck, 22.04.2026

Abstract

Humans act in environments that are uncertain, dynamically changing, and only partially controllable. Successful behavior in such settings requires more than reactive feedback correction or rule-based strategies. Agents must continuously regulate how precisely action goals are specified, how strongly predictions guide action, and how they adapt when control deteriorates. Despite extensive work on motor control, decision-making, and learning, existing accounts lack a unified explanation of how behavioral strategies, subjective control beliefs, and computational learning mechanisms jointly support adaptive action control in these dynamic environments that are characterized by uncertainty and action–effect contingencies that change over time.

This dissertation develops and empirically grounds a multi-level framework of action control that integrates hierarchical accounts of intention, graded degrees of control, and belief-based regulation of agency. Action control emerges from the dynamic interaction between behavioral markers of ongoing regulation, cognitive forward models supporting both reactive and proactive action selection, and computational mechanisms driving parameter adjustments of these internal forward models.

Across four studies that share the same continuous task environment, action control is investigated at behavioral, cognitive, and computational levels. Studies 1 and 2 show that gaze behavior implements a hierarchical organization of action goals over different temporal horizons. Close fixations support immediate control and state-dependent regulation, whereas distant fixations anchor attention to future task-relevant locations, supporting proactive planning. These complementary fixation types adapt flexibly to changes in environmental dynamics and action–effect contingencies, revealing how perceptual control implements both reactive and anticipatory action.

Study 3 treats the Sense of Control, the subjective feeling of being in control, as a latent belief that is updated via Bayesian integration. Participants provided control ratings after each trial, serving as observable indicators of their evolving beliefs. Analysis showed that participants derived their ratings by integrating prior expectations with observed performance outcomes. Control ratings decreased when environmental outcomes violated participants' expectations, and individuals differed in how strongly they weighted performance evidence, indicating variability in belief updating. These results suggest that the Sense of Control reflects the predictive accuracy of internal forward models.

Study 4 translated these insights into a cognitive architecture combining forward predictions with error-driven learning. The Sense of Control from Study 3, modeled using the Bayesian framework, was reinterpreted as uncertainty associated with an internal forward model of the task's environmental dynamics. This uncertainty modulated whether forward predictions were used during action selection. Simulations showed that when high uncertainty led the architecture to suspend forward prediction, learning was severely impaired. Adaptation was slow, performance remained below that of human participants, and internal representations failed to improve because no informative prediction errors were generated. When the architecture relied on predictions despite high uncertainty, corrective error signals enabled efficient parameter updating, resulting in learning trajectories and performance comparable to human participants. This demonstrates that engaging predictive mechanisms early, even under uncertainty, is essential for effective adaptation to dynamic environments.

Together, these studies establish a process-based framework in which adaptive action control is neither purely reactive nor rigidly predictive. Instead, it emerges from coordinated mechanisms operating across multiple levels and timescales, with uncertainty guiding when internal models are updated incrementally and when they must be fundamentally restructured. By linking gaze behavior, subjective control beliefs, and predictive learning, this dissertation provides a foundation for investigating structural learning, proactive control, and failures of agency in both human and artificial systems.

Kurzfassung

Menschen handeln in Umgebungen, die unsicher, dynamisch und nur teilweise kontrollierbar sind. Erfolgreiches Verhalten unter solchen Bedingungen erfordert mehr als reine Feedback-Korrektur oder regelbasierte Strategien. Handelnde Systeme müssen fortlaufend regulieren, wie präzise Handlungsziele spezifiziert sind, wie stark Vorhersagen das Handeln leiten und wie sie ihr Verhalten anpassen, wenn Kontrolle verloren geht. Trotz umfangreicher Forschung in den Bereichen der motorischen Kontrolle, Entscheidungsfindung und des Lernens fehlt bislang eine einheitliche Erklärung dafür, wie Verhaltensstrategien, subjektive Kontrollüberzeugungen und komputationale Lernmechanismen gemeinsam adaptive Handlungskontrolle in dynamischen Umgebungen ermöglichen, die durch Unsicherheit und sich verändernden Handlungs-Wirkungs-Zusammenhängen gekennzeichnet sind.

Diese Dissertation entwickelt und untersucht empirisch einen hierarchisch organisierten theoretischen Rahmen der Handlungskontrolle, der intentionale Strukturen, graduelle Ausprägungen von Kontrolle und metakognitive Regulation von Handlungseinfluss integriert. Handlungskontrolle entsteht dabei aus dem dynamischen Zusammenspiel von Verhaltensindikatoren fortlaufender Regulierung, kognitiven Vorhersagemodellen, die sowohl reaktive als auch proaktive Handlungswahl unterstützen, sowie komputationalen Mechanismen, die die Parameter dieser internen Vorhersagemodelle anpassen.

In vier Studien, die in derselben kontinuierlichen Aufgabenumgebung durchgeführt wurden, wird Handlungskontrolle auf der Verhaltensebene, der kognitiven Ebene und der Ebene der Modellierung untersucht. Studien 1 und 2 zeigen, dass Blickverhalten eine hierarchische Organisation von Handlungszielen über unterschiedliche zeitliche Bezugsebenen hinweg realisiert. Close fixations unterstützen unmittelbare Kontrolle und zustandsabhängige Regulation, während distant fixations die Aufmerksamkeit auf zukünftig relevante Positionen lenken und so proaktive Planung unterstützen. Beide Fixationstypen passen sich flexibel an Veränderungen der Umgebungsdynamik und der Handlungs-Wirkungs-Zusammenhänge an und verdeutlichen, wie visuelle Kontrolle reaktive und antizipierende Aspekte der Handlungskontrolle miteinander verbindet.

Studie 3 konzeptualisiert den Sense of Control, das subjektive Gefühl in Kontrolle zu sein, als latente Überzeugung, die über Bayes'sche Integration aktualisiert wird. Teilnehmende gaben nach jedem Durchgang Einschätzungen ihres Kontrollgefühls ab, die

als beobachtbare Indikatoren dieser sich entwickelnden Überzeugung dienten. Analysen zeigten, dass diese Kontrollbewertungen aus der Integration vorheriger Erwartungen mit beobachteten Handlungsauswirkungen hervorgingen. Das Kontrollgefühl nahm ab, wenn Handlungsauswirkungen von den Erwartungen abwichen, wobei sich deutliche individuelle Unterschiede in der Gewichtung leistungsbezogener Evidenz zeigten. Die Ergebnisse sprechen dafür, dass der Sense of Control die prädiktive Genauigkeit interner Vorhersagemodelle widerspiegelt.

Studie 4 überführte diese Befunde in eine kognitive Architektur, die kognitive Vorhersagen mit fehlergetriebenem Lernen kombiniert. Der in Studie 3 innerhalb eines Bayes'schen Modellierungsansatzes beschriebene Sense of Control wurde dabei als Unsicherheit eines internen Vorhersagemodells der Umgebungsdynamik interpretiert. Diese Unsicherheit bestimmte, ob kognitive Vorhersagen in die Handlungswahl einbezogen wurden. Simulationen zeigten, dass das Unterbinden von Vorhersagen bei hoher Unsicherheit zu stark beeinträchtigtem Lernen führte. Die Handlungsanpassung verlief langsam, die Leistung blieb unter der menschlicher Teilnehmender, und interne Repräsentationen verbesserten sich nicht, da keine informativen Vorhersagefehler generiert wurden. Wurde kognitives Vorhersagen hingegen trotz hoher Unsicherheit genutzt, ermöglichten korrigierende Fehlersignale effiziente Parameter-Updates, was zu Lernverläufen und Leistungsniveaus führte, die mit denen menschlicher Teilnehmender vergleichbar waren. Dies unterstreicht, dass die frühe Einbindung prädiktiver Mechanismen, selbst unter Unsicherheit, eine zentrale Voraussetzung für erfolgreiche Anpassung in dynamischen Umgebungen ist.

Zusammen begründen diese Studien einen theoretischen Rahmen, in dem adaptive Handlungskontrolle weder als rein reaktiv noch als starr prädiktiv verstanden wird. Stattdessen entsteht sie aus koordinierten Mechanismen, die über mehrere Ebenen und Zeitskalen hinweg interagieren, wobei Unsicherheit steuert, wann interne Modelle schrittweise angepasst und wann sie grundlegend reorganisiert werden müssen. Durch die Verknüpfung von Blickverhalten, Sense of Control und prädiktivem Lernen schafft diese Dissertation eine Grundlage für die Untersuchung strukturellen Lernens, proaktiver Kontrolle und von Störungen der Handlungskontrolle in menschlichen wie auch künstlichen Systemen.

Acknowledgments

Nele — Thank you for giving me the freedom to find my own path, and for nudging me, sometimes gently and sometimes a bit more firmly, when I started drifting in the wrong direction. I couldn't have ended up with a better supervisor. Watching you receive the call to Lübeck and navigate that turbulent time showed me that no one is immune to uncertainty and that confidence is something we learn to hold on to. Seeing how you built your team, recognized strengths in others, and helped people grow has left a deep impression on me.

Stefan — I have rarely met anyone as quietly competent as you. You always knew how to ask the right questions. You balanced real curiosity for ideas with a pragmatic instinct for what is worth focusing on. I'm grateful you were the second PI on this project. You and Nele complement each other perfectly; you should consider leading more projects together.

Eva — I truly enjoyed our conversations. From you, I learned more about the realities of research and university life than from anyone else. You never hid how demanding it all can be, and through that honesty, you became a kind of moral anchor for me. Thank you for letting me stay, and for giving me a place where I felt welcome.

To the members of the KKE group — Thank you for giving me not only a place to work but a team to belong to. We often approached things differently methodologically, but I genuinely enjoyed discussing your research and sharing mine. Winnie, I'm looking forward to all our future sleepovers.

Annika — You were my partner in struggling through things together. Our shared sessions of complaining, but also actually figuring things out, kept me going. Your feedback on the chapters of this thesis meant a lot to me.

Maren — Of all the PhD students I know, you deserve to be a PhD the most. Your rigorous way of doing research has influenced my own work more than you probably realize. I really enjoyed working with you, and I'm grateful for your feedback on chapters of this thesis.

To my family — My dreams have taken me to another city, but you have always stayed close, even from afar.

To my siblings — Nico, when I think of determination, I think of you. Alice, when I think of strength, I think of you. Annika, when I think of love, I think of you. I cannot imagine not being part of us. Thank you for being there.

Mama and Papa — You have always made me feel that no matter what I do, you stand behind me, whether I fail or succeed.

Hasret — Everything I achieve, I achieve through you.

Nils Wendel Heinrich, Lübeck, January 2026

Contents

1 Introduction

Humans routinely act in environments that are uncertain, dynamically changing, and in which outcomes are not fully determined by their actions. In such settings, successful behavior requires more than reacting to immediate sensory input or following a fixed policy. Instead, humans continuously regulate how precisely they specify the goals guiding their actions, how strongly they act toward anticipated outcomes, and how they adapt their strategies when control deteriorates or is lost. These processes of human action control are studied at multiple levels: behavioral markers reveal how control unfolds in action, cognitive mechanisms govern how feedback and beliefs guide ongoing control, and computational mechanisms formalize the strategies and constraints that support flexible action under uncertainty. But understanding the mechanisms that support human action control remains a central challenge in the field of cognitive science.

This challenge has become increasingly important as artificial systems act autonomously in real-world environments where mistakes have serious consequences. In many applied domains, autonomous agents are expected to operate reliably under uncertainty, interact with complex environments, and adapt to changing conditions. However, current agentic systems, which are often based on large-scale statistical learning, typically lack an explicit representation of their own control capabilities or of uncertainty about action outcomes. As a result, they struggle to regulate behavior when predictions become unreliable or when effective control breaks down (Uesato et al., 2018; Manchingal et al., 2025).

Existing approaches to autonomous behavior emphasize different aspects of this problem. Large-scale learning systems derive control policies from optimization over extensive datasets (LeCun, Bengio, and Hinton, 2015; Sutton and Barto, 2018). Classical symbolic systems rely on explicit reasoning, rule systems, and transparent planning (Newell and Simon, 1976; Russell and Norvig, 2020). Other approaches draw on control theory and robotics, using feedback control, state estimation, and model-predictive planning to ensure stability and performance in dynamic environments (Siciliano and Khatib, 2016; Camacho and Bordons, 2007). Hybrid systems combine learned perception modules with symbolic or optimization-based planners (Kaelbling and Lozano-Pérez, 2013; Levine et al., 2016). While each of these approaches captures important aspects of action, none directly addresses how human and artificial agents monitor and regulate their own degree of control while acting.

This shortcoming in capturing how agents monitor and regulate their own control places specific demands on empirical and computational research. Studying action under continuous, dynamically changing, and uncertain conditions requires task environments that support sustained regulation over time, behavioral measures that reflect ongoing control rather than isolated decisions, and models that link these measures to internal beliefs and computational mechanisms. However, much of the existing literature has not met these requirements.

First, most empirical and computational approaches lack *task continuity*, relying on discrete and simplified task structures that abstract away the dynamics of real-world action. Such tasks limit the formation of higher-level goals and their translation into ongoing action regulation. Second, there is little *cross-level integration* between behavioral markers, cognitive beliefs, and computational mechanisms, leaving unclear how these levels jointly produce adaptive control. Both of these research gaps are discussed in more detail in Section 1.2.9.

Finally, existing models rarely make explicit how agents represent, update, and use estimates of their own influence over outcomes during ongoing behavior. As a result, control *self-monitoring* is typically treated implicitly or descriptively, rather than being specified as a mechanistic component of action control.

Against this background, the present work approaches action control from the human perspective. Unlike current artificial systems, humans possess a sense of control that reflects their ongoing ability to influence outcomes (for an overview, see Grünbaum and Christensen, 2020). The subjective control experience is supported by the integration of behavioral feedback, subjective experience, and internal models, and allows humans to sustain effective action under uncertainty. This perspective provides a theoretical and empirical framework for understanding robustness and adaptability when environments are unpredictable and agents are not fully in control.

Despite decades of research on motor control, decision-making, and learning, relatively little is known about how humans represent and update beliefs about their own ability to influence outcomes at the level of subjective experience and action regulation in uncertain and dynamically changing environments. Yet such beliefs are critical for adaptive behavior, as they determine whether human agents persist, explore alternatives, or disengage when control deteriorates (Abramson, Seligman, and Teasdale, 1978). Contemporary accounts of agency emphasize that judgments of control are derived from integrating multiple sources of information, including sensorimotor and cognitive information (Giersiepen et al., 2025; Synofzik, Vosgerau, and Newen, 2008). Recent work across psychology, neuroscience, and human–machine interaction has highlighted that failures of action are often not due to insufficient skill or information, but to miscalibrated beliefs about control (Hyland, 2020; Lee and See, 2004).

At the same time, advances in experimental methods and computational modeling make it possible to study action control at multiple levels, from fine-grained behavior to cognitive processes and explicit computational mechanisms. Because action control is ultimately expressed in behavior, studying behavioral markers provides a crucial foundation for understanding the cognitive and computational processes that generate them. Computational models, in turn, allow us to formalize the mechanisms by which humans integrate feedback, update beliefs, and adjust their actions under uncertainty. This layered approach creates an opportunity to systematically investigate the mechanisms that support effective action control and naturally leads to the central research question:

> **What behavioral, cognitive, and computational mechanisms enable humans (and artificial agents modeled after them) to exert effective action control across uncertain and dynamically changing environments?**

To address this research question, I decompose action control into interacting behavioral, cognitive, and computational levels. Behavior provides observable markers through which action control can be measured. At the cognitive level, control is explained in terms of inferred processes by which agents monitor, evaluate, and regulate their own influence over action outcomes. At the computational level, these processes are put in explicit models that specify how feedback is integrated, beliefs are updated, and actions are selected over time. Together, this yields a coherent account of action control that is empirically grounded and mechanistically explicit.

All studies in this dissertation are conducted in a dynamic task environment that preserves task continuity. I begin with an empirical analysis of fine-grained behavior in this environment (Study 1, Section 2.2.1). Using a data-driven approach, I identify two distinct types of fixational eye movements, which provide high-resolution behavioral indicators of how humans monitor and adjust their actions in real time. Building on this, Study 2 (Section 2.2.2) investigates how these eye movements adapt to environmental predictability, revealing both reactive and proactive modes of control. In this study, I also collect participants' subjective experiences of control, establishing a first link between observed behavior and how humans experience their own agency. Together, these first two studies establish a behavioral foundation for understanding the cognitive mechanisms that support adaptive control.

In Study 3 (Section 3.2), I use a sequential Bayesian integration framework to model participants' subjective experiences of control and address control self-monitoring at a mechanistic level. The Bayesian model formalizes how humans update beliefs about their own ability to influence outcomes based on performance and expectations. Lastly, Study 4 (Section 4.2) examines how these processes support learning and adaptive behavior through cross-level integration. I present a computational cognitive model built in the ACT-R architecture (J. R. Anderson, Bothell, et al., 2004), which enables mechanistic modeling of learning and adaptive strategies by proceduralizing cognition through the integration of perceptual, motor, and declarative modules. The cognitive model learns flexible strategies through interaction with a changing task environment. It incorporates the subjective control responses from Study 3 and condenses them into an uncertainty parameter linked to the task representation. This uncertainty parameter influences predictions, guides action plans, and shapes behavior. Both the task representation and its associated uncertainty are updated continuously, depending on whether the predictions derived from the representation are confirmed, closing the loop between perception, cognition, and action.

These studies provide a comprehensive, layered account of human action control, from observable behavior to cognitive and computational mechanisms, and reveal how humans monitor control, adapt to change, and form beliefs that support effective action.

1.1 Theoretical Foundations of Human Action Control

Exercising control over one's actions is a deceptively simple notion. In ordinary language, we speak as though control is something an agent either possesses or lacks: one is in control of a car, one loses control of one's temper, one regains control after a slip. Yet in practice, control is rarely clear-cut. Agents may exercise varying degrees of control, depending on how actively they regulate and monitor their behavior in pursuit of a goal (Shepherd, 2014). The philosophically inclined distinction between being in control and exercising control is therefore not merely semantic, but conceptual. Being in control describes a relatively stable state of successful regulation, whereas exercising control points to the process of regulating, adjusting, and coordinating one's actions in real time, especially under conditions of uncertainty or change (Rotter, 1966; Wennerhold and Friese, 2023). Exercising control therefore lies along a continuum: from fluent, automatic actions that are executed with little conscious oversight, to deliberate, effortful regulation when circumstances demand active intervention.

This continuum raises foundational questions about what ”control” consists in and how it can be understood scientifically. Is control primarily a matter of physical feedback and stability, or of cognitive representation and intention? Do simple systems that maintain homeostasis, like thermostats or bacterial chemotaxis, already exercise control, or is true control a hallmark of cognitive systems capable of goal-directed reasoning?

These questions motivate the need for a more formal characterization of control that can accommodate its graded nature and its embedding within the cognitive hierarchy.

1.1.1 Shepherd's Contours of Control

Shepherd's (2014) Contours of Control offers an influential attempt to define the structure of control in a way that bridges philosophical and empirical perspectives. His central insight is that control is not a single, indivisible trait or ability but rather a set of interacting mechanisms and processes work together to achieve control. Control is distributed across processes like sensing, predicting, and acting. He emphasizes that control should be analyzed in terms of how the system functions rather than as a single static property an agent possesses. The focus should be on how different components contribute to goal-directed regulation and adaptability. In this respect, Shepherd identifies two key features that together characterize a control system: *guidance* and *flexible repeatability*.

First, a system must guide action appropriately in the service of a goal. According to Shepherd, a process counts as control only if it produces outcomes in a way that is sensitive to goal-relevant information. This introduces a normative dimension that has been explored extensively in the philosophy of action, where Fischer and Ravizza (1998) argue that guidance control requires mechanisms to be appropriately responsive to reasons. This implies that control involves more than mere causation: actions can be better or worse guided relative to the goal they are meant to serve. For example, a navigator reaching their destination by first walking in the wrong direction is clearly worse than taking a direct path, even if both eventually succeed. Guidance therefore requires a structured relation between the system's internal organization and the external conditions under which the goal is pursued (as discussed by Bratman, 1987). The system's structure must be such that it can respond appropriately to features of the environment that are relevant to the goal (Hommel et al., 2001; Carver and Scheier, 2013).

Importantly, this does not in itself require the system to internally represent its goals. Rather, the system's dynamics can instantiate guidance by being appropriately coupled to environmental feedback. In the example of navigating towards a destination, a system might maintain the correct direction not by consulting an internal map, but by continuously adjusting its trajectory in response to perceptual cues that signal proximity to the destination. In this sense, control can emerge from the coordination between perception and action, where the system maintains its behavior within the bounds of goal-relevant success conditions.

Second, control requires flexible repeatability: the ability to achieve the same goal across varying contexts and perturbations. A thermostat that maintains temperature under changing external conditions, a driver who steers successfully on both dry and wet roads, or, revisiting the earlier example, a navigator who reaches their destination from various starting positions, all demonstrate this capacity for flexible adjustment. Shepherd calls this *wide* control (in contrast to narrow control): the ability to coordinate and integrate multiple sub-processes across different temporal and situational scales. Wide control enables agents to adjust strategies, balance competing constraints, and maintain goal-directed action under uncertainty. This capacity for flexible repeatability distinguishes controlled behavior from mere reflexes or rigidly preprogrammed routines.

Importantly, Shepherd's framework remains nonrepresentational in spirit. He characterizes control as a relational property between an agent, its actions, and its environment, rather than as a process that necessarily depends on internal models or symbolic representations. This emphasis on functional organization allows his account to scale across biological, cognitive, and social levels of explanation, as long as these systems exhibit guidance and flexible repeatability. Shepherd thereby offers a principled answer to the question of what it means to exercise control: to actively maintain goal-directed regulation in the face of potential deviation, and to do so with sufficient flexibility to handle changing circumstances. However, while this captures the structure of control, it leaves open an important aspect: how can we distinguish between systems that merely exhibit goal-directed dynamics and those that understand or represent what they are controlling? This issue is central to Gentry's (2021) extension of Shepherd's framework.

1.1.2 Gentry's Semantic Expansion of Control

Gentry (2021) agrees with Shepherd that control involves guiding action in the service of a goal and doing so flexibly. But he argues that these conditions alone are not sufficient for what he calls genuine control. A system could satisfy both features (for example, a well-designed autopilot or biological homeostat) without possessing any understanding of what it is controlling. Such systems, Gentry argues, lack the capacity for normative error: they cannot be wrong about what they are doing, they can only malfunction.

To capture what distinguishes genuine control from mere regulation, Gentry introduces a third requirement: in addition to exhibiting guidance and flexible repeatability, the system must be capable of semantic information processing. This means that the system's internal states represent aspects of the world in ways that can be true or false, accurate or mistaken. Semantic information processing allows a system to interpret sensory inputs and internal signals as information about the world (e.g. the position of an object, the speed of movement, the distance to a goal). Without this semantic dimension, control lacks the normative grounding necessary for evaluation in terms of success, error, or misrepresentation.

This addition is crucial because it explains how control systems can be sensitive to what should happen rather than merely what does happen. A thermostat can fail if its sensor breaks, but it cannot misrepresent the temperature. The poor lost navigator who wants to reach their destination, by contrast, can plan incorrectly because they believe they are heading north when they are actually heading east. The difference lies in representational capacity: the navigator's control depends on internal models that can be more or less accurate, allowing for prediction and correction based on semantic information about the world.

The extended control systems theory also fits within embodied cognition, according to which cognitive processes are partly realized through the interaction of brain, body, and environment (Shapiro, 2007). Gentry frames semantic information processing as part of an extended system and argues that representational states are not confined to the brain but include contributions from the body and environment. The navigator finding their way, for example, depends on the coordinated processing of visual landmarks, spatial memory, and movement actions. Together these components form a distributed system that integrates information about position, orientation, and speed.

1.1.3 Pacherie's Hierarchical Structure of Intention and the Sense of Control

Taking Shepherd's functional analysis and Gentry's emphasis on information-processing capacities, Pacherie's (2008) framework provides a complementary, representational perspective on how control is organized within the agent. Where Shepherd and Gentry characterize control in terms of guidance, flexibility, and informational coupling with the environment, Pacherie examines how these capacities are reflected in a hierarchical architecture of intention. She discusses how agents internally represent and progressively refine their goals across multiple temporal and cognitive scales.

In her dynamic theory of intentions (Pacherie, 2006; Pacherie, 2008), the goal of an action can be specified at distinct levels of intention. Distal intentions (D-intentions) encode abstract, temporally extended goals and general desires for action, such as writing a dissertation or getting in shape. They determine what the agent aims to achieve and structure longer sequences of behavior. Pacherie aligns D-intentions with the agent's beliefs and desires, highlighting the relative indeterminacy of this level. Proximal intentions (P-intentions) translate these abstract goals into context-specific action plans, incorporating environmental constraints and current opportunities for action. For example, the distal intention to write a dissertation is translated into the proximal intention to open the document and draft a paragraph. Motor intentions (M-intentions) operate at the execution level. They specify the detailed motor parameters needed to realize proximal plans and dynamically adjust movements based on sensory feedback, such as the fine-grained muscle activations required to type a sentence on a keyboard (for a schematic of Pacherie's hierarchical framework, see Figure 1.1).

This hierarchy bridges abstract planning and embodied motor control. It also captures the dynamic interplay between top-down guidance and bottom-up feedback: higher-level intentions constrain lower-level execution, while sensory feedback can prompt adjustments at both motor and proximal levels. Importantly, Pacherie connects this functional organization to the Sense of Control (SoC), the subjective feeling of being in control or having to exert control to maintain a course of action. She proposes that the SoC arises from the successful integration of predictive and feedback signals across the different levels of the intentional hierarchy. When motor outcomes match predicted results, the agent experiences a strong SoC; when mismatches occur, such as in cases of involuntary movement or distorted tool feedback, the SoC diminishes. In this way, control involves a continuous cycle of descending specification and ascending correction: intentions guide action down the hierarchy, while sensory and performance feedback inform adjustments upward.

This phenomenological dimension complements Shepherd's and Gentry's theories in several ways. First, Pacherie provides a temporal articulation of control: the capacity to exercise control depends not only on having goals and semantic representations but also on coordinating them across distinct timescales of intention. Second, her model grounds the subjective experience of exercising control in mechanisms of prediction and comparison, aligning closely with computational frameworks such as predictive processing and active inference (K. Friston, 2005; K. Friston, 2010). In these models, the system minimizes prediction errors across hierarchical levels, effectively exercising control by keeping sensory input within expected bounds. Pacherie's intention cascade can therefore be interpreted as a cognitive architecture that implements both Shepherd's goal-guided flexibility and Gentry's semantic normativity within a hierarchical predictive framework.

Taken together, these three accounts offer a comprehensive account of action control. Shepherd identifies its core functional requirements: goal-directed regulation and flexible repeatability. Gentry adds that genuine control depends on semantic information processing, which allows systems to represent, misrepresent, and correct their internal models and behavior. Pacherie discusses control as being organized hierarchically through intentions that connect abstract goals with concrete motor actions, and that this organization gives rise to the phenomenological SoC through predictive processing.

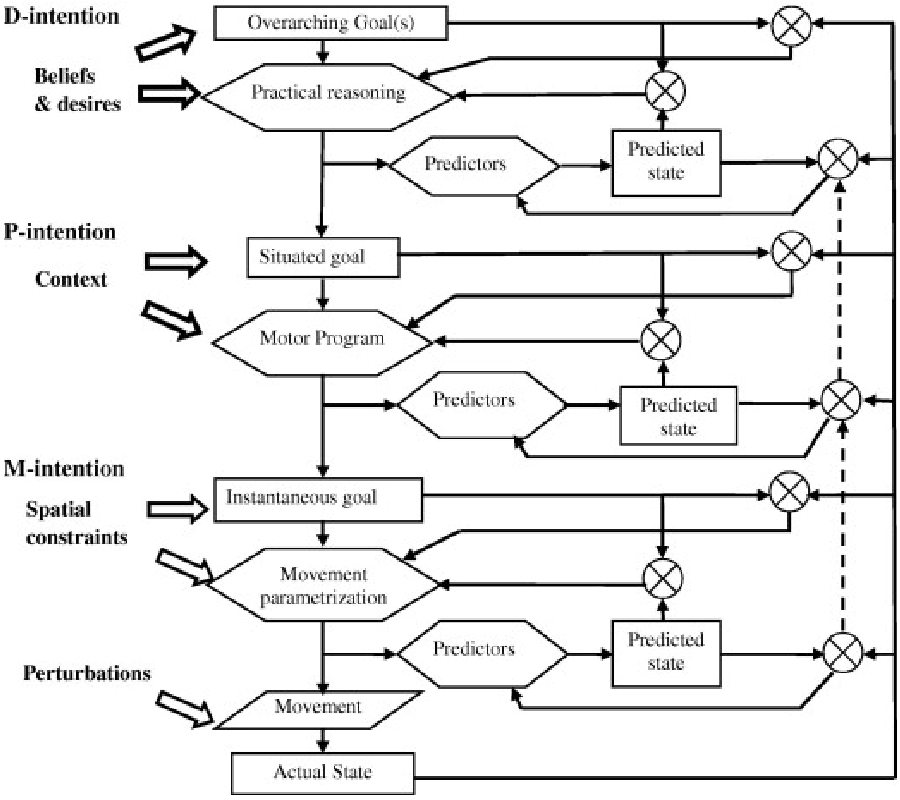

Figure 1.1 Illustration of the hierarchical model of action specification by Pacherie (2008). At the highest level, distal intentions (D-intentions) arise from beliefs and desires and generate overarching goals through practical reasoning. These are translated into proximal intentions (P-intentions) by incorporating contextual information to form situated goals. P-intentions are then further specified into motor intentions (M-intentions) by considering spatial constraints, resulting in instantaneous goals that guide movement parametrization. Each hierarchical level contains its own internal forward model (predictors) that generates predicted states. The actual state produced by motor execution is compared against predicted states across all three levels of the hierarchy (indicated by the crossed circles), with discrepancies at each level reflecting different aspects of prediction error. This multi-level comparison allows the hierarchy to detect control failures at varying degrees of abstraction, from high-level goal violations to low-level motor deviations.

1.2 Empirical Foundations of Human Action Control

The philosophical accounts above provide a conceptual framework for understanding what control is and how it is structured. To bridge these theoretical insights with empirical investigation, I will review experimental research on human action control that has established fundamental principles about how actions are represented, selected, organized, and executed.

1.2.1 Perception-Action Integration

The study of human action control has its conceptual roots in William James's ideomotor principle (1890), which proposed that actions are initiated and controlled through the anticipation of their sensory effects. James suggested that the mental representation of an action's outcome is sufficient to trigger the action itself, establishing a bidirectional link between perception and action. This principle remained largely dormant in experimental psychology until its revival in contemporary cognitive research.

Hommel, Müsseler, Aschersleben, and Prinz (2001) elaborated James's original insight into a comprehensive Theory of Event Coding (TEC), which fundamentally reconceptualized how perception and action are represented in the cognitive system. Rather than treating perception and action as separate domains with distinct codes, TEC proposed that both are coded in a common representational medium using distal event codes. According to this framework, perceived events and planned actions share the same representational format, which explains a wide range of perception-action interaction effects. The theory accounts for phenomena such as ideomotor compatibility effects, where the similarity between perceived and to-be-executed actions facilitates or interferes with performance (Prinz, 1997). Hommel and colleagues demonstrated that feature overlap between perceived stimuli and action plans can lead to both facilitation when features match intended actions and interference when they mismatch, providing strong empirical support for common coding principles.

Elsner and Hommel (2001) extended this theoretical framework by demonstrating how action-effect associations are acquired through experience. Their research showed that even arbitrary relationships between actions and their sensory consequences become encoded such that later perception of the effect can prime or trigger the associated action. This work provided crucial empirical evidence for the bidirectional nature of action-effect links and established learning mechanisms through which the ideomotor principle operates. The ideomotor framework and TEC empirically ground Gentry's notion of semantic information processing: actions are controlled not merely through mechanical feedback but through representational states that link movements to their perceptual consequences. This bidirectional coupling between perception and action constitutes a form of semantic control, where internal representations can be validated or contradicted by sensory outcomes.

1.2.2 Hierarchical Control Mechanisms

A central challenge in understanding human action control concerns how complex, goal-directed behavior is organized across multiple levels of representation and control. Norman and Shallice (1986) proposed an influential framework distinguishing between routine action control through contention scheduling and deliberate control through a supervisory attentional system. In their model, contention scheduling manages well-learned action sequences through a competitive process among action schemas, with the most highly activated schema winning control. The supervisory attentional system intervenes when routine control is insufficient, such as in novel situations, during planning and decision-making, when overcoming habitual responses, or in error correction. This framework provided a theoretical basis for understanding both automatic and controlled aspects of action and has been particularly influential in explaining action slips (Reason, 1984), perseverative errors (Luria, 1966), and deficits following frontal lobe damage (Shallice and Burgess, 1991).

Cooper and Shallice (2000) developed computational implementations of these ideas. The authors specified mechanisms through which hierarchical goals activate and constrain lower-level action schemas. Their models demonstrated how the interaction between routine and supervisory control could account for complex patterns of both normal and pathological action control, including the temporal organization of everyday action sequences.

This empirical work on hierarchical control mechanisms directly instantiates Pacherie's distinction between distal, proximal, and motor intentions. The supervisory attentional system operates at the level of distal and proximal intentions, selecting among competing action plans and contextualizing them appropriately, while contention scheduling implements the more automatic execution corresponding to motor intentions. The framework also embodies Shepherd's notion of flexible repeatability: the supervisory attentional system enables wide control by intervening when routine processes are insufficient, allowing the same goal to be achieved across varying contexts.

1.2.3 Action Selection and Inhibitory Control

Flexible action control requires not only the activation of appropriate responses but also the inhibition of inappropriate ones. The ability to stop or inhibit planned or ongoing actions is a crucial component of adaptive behavior. Logan and Cowan (1984) developed the stop-signal paradigm and the associated race model to measure the latent process of response inhibition. In this task, participants perform a primary choice reaction time task but occasionally receive a stop signal instructing them to withhold their response. By varying the delay between the go stimulus and stop signal, the model estimates the stop-signal reaction time (the latent time required to inhibit a response). This paradigm established inhibitory control as a measurable cognitive process with characteristic temporal dynamics and demonstrated that stopping is slower than going, revealing a fundamental asymmetry in action control.

Verbruggen and Logan (2008) extended this framework by distinguishing between reactive and proactive inhibition. Reactive inhibition refers to stopping a response after a stop signal appears, while proactive inhibition involves the advance preparation to potentially stop, which slows overall responding but improves stopping ability. Their work demonstrated that inhibitory control is not a singular process but involves multiple modes that can be strategically deployed depending on task demands. This distinction has been important for understanding how humans adaptively adjust their control modes based on environmental demands.

The distinction between reactive and proactive control modes again reflects Shepherd's emphasis on flexible repeatability as a fundamental cognitive capacity. Effective control requires not only the ability to execute goal-directed actions but also the capacity to inhibit them when circumstances change. Proactive inhibition, in particular, exemplifies wide control by integrating information about task context and potential future demands to adjust control settings in advance.

1.2.4 Sequential Action and Motor Programming

The organization of sequential actions poses the classical ”serial order problem” articulated by Lashley (1951), who noted that action sequences often unfold too rapidly for each element to be triggered by sensory feedback from the previous element. This implies that sequences must be represented and controlled through internal programs or plans rather than simple stimulus-response chains. Lashley's insights suggest that understanding action control requires explaining how extended action sequences are mentally represented and executed.

Rosenbaum (1980) discussed the concept of motor programming to explain how action sequences are prepared before execution. His research on movement pre-cueing demonstrated that various parameters of upcoming movements (such as direction, extent, and force) can be specified in advance. Each parameter contributes independently to preparation time: providing advance information about movement direction reduces preparation time by a certain amount, and providing information about extent or force produces additional independent reductions. This parametric independence established that motor programs have a compositional structure in which movement features are separately specified and assembled.

This compositional structure of motor programs provides a mechanistic account of how Pacherie's motor intentions (M-intentions) are implemented. The separate specification of movement parameters corresponds to the progressive refinement of action plans as they descend through the intentional hierarchy, with each level adding more concrete detail to the abstract goal representation.

1.2.5 Goal-Directed and Habitual Control Systems

A fundamental distinction in action control concerns whether actions are selected based on their anticipated outcomes (goal-directed) or triggered automatically by environmental stimuli (habitual). Dickinson and Balleine (1994) developed a framework based on instrumental learning research. The authors proposed that goal-directed actions are mediated by action-outcome associations and are sensitive to changes in outcome value. Habitual actions, on the other hand, are mediated by stimulus-response associations and persist despite outcome devaluation. Though initially developed through animal research, this framework has become foundational for understanding human action control.

De Wit and Dickinson (2009) adapted outcome devaluation procedures for human participants, demonstrating that well-trained actions can become insensitive to changes in outcome value, indicating a transition from goal-directed to habitual control. Their research established that human action control, like that of other animals, involves multiple systems with different computational properties, and that extensive training can shift control from a flexible but computationally demanding goal-directed system to an efficient but inflexible habitual system.

This dual-system architecture can be mapped onto the distinction between Shepherd's narrow and wide control. Habitual control represents narrow control: efficient and reliable within a restricted set of circumstances but inflexible when conditions change. Goal-directed control, by contrast, embodies wide control: more resource-intensive but capable of flexibly adjusting to new circumstances by computing action-outcome contingencies. The transition between these systems represents a fundamental trade-off in action control between computational efficiency and behavioral flexibility.

1.2.6 Motor Control and Optimal Feedback Processing

Wolpert and colleagues (1995) developed computational theories of optimal feedback control that explain how the motor system plans and executes movements under uncertainty (see also Todorov and Jordan, 2002). Their framework proposes that the brain maintains forward models that predict the sensory consequences of motor commands and inverse models that determine motor commands needed to achieve desired outcomes. By combining predicted and actual sensory feedback, the system can rapidly detect prediction errors, correct them, and adapt to changing conditions. The optimal control framework has been influential in explaining both reaching movements and more complex actions and offers a principled account of how the brain exercises motor control.

This framework offers a computational implementation of the prediction-comparison mechanisms that Pacherie identifies as central to the SoC. The forward model generates predictions about sensory consequences, and these predictions are continuously compared with actual feedback. When predictions match outcomes, control is experienced as fluent and successful; when mismatches occur, they trigger corrective adjustments and may diminish the SoC. Wolpert's framework thus bridges computational modeling with phenomenological experience; it explains the underlying computational architecture from which the subjective feeling of control emerges.

1.2.7 Cognitive Control and Conflict Monitoring

Extending beyond specific action control mechanisms, researchers have investigated domain-general cognitive control processes that regulate action selection. Botvinick, Braver, Barch, Carter, and Cohen (2001) proposed the conflict monitoring theory, suggesting that the anterior cingulate cortex monitors for conflicts between competing response tendencies and signals the need for increased cognitive control. When conflict is detected, control is upregulated to improve subsequent performance. This phenomenon is known as the conflict adaptation or Gratton effect (Gratton, Coles, and Donchin, 1992). The conflict monitoring theory provided a unified account of performance adjustments across multiple tasks and established conflict detection as a key mechanism in action control.

Braver (2012) proposed the dual mechanisms of control framework, distinguishing between proactive control (sustained, anticipatory maintenance of goal-relevant information) and reactive control (transient, stimulus-driven reactivation of goals). Proactive control involves actively maintaining task goals and is resource-intensive but enables optimal performance, while reactive control is more efficient but leads to occasional errors when goals are not accessible. Individual differences in the balance between these control modes have been linked to working memory capacity, aging, and clinical conditions. The dual mechanisms of control framework has been influential in explaining how control is deployed flexibly depending on both individual characteristics and situational demands.

Braver's distinction between proactive and reactive control modes directly parallels Verbruggen and Logan's distinction in the inhibitory control domain, suggesting a general principle: control systems can operate in either an anticipatory mode that maintains readiness at a cost or a reactive mode that responds to immediate demands more efficiently but less reliably. This trade-off between sustained preparedness and reactive adjustment reflects the tension between different strategies for achieving Shepherd's flexible repeatability under varying environmental demands.

1.2.8 Dynamic Approaches to Action Control

Traditional reaction time paradigms treat action control as a series of discrete processing stages culminating in a ballistic motor output (Sternberg, 1969). However, this discrete stage model was challenged. Analyzing continuous motor trajectories in reaching tasks revealed that cognitive processing influences motor output continuously throughout movement execution (Spivey, Grosjean, and Knoblich, 2005; Dale, Kehoe, and Spivey, 2007). For instance, when participants face competing response options, their hand trajectories curve toward the non-selected alternative before settling on the final choice. This finding revealed the real-time dynamics of competition resolution. The work of these authors established that examining continuous motor output gives insight into cognitive processes that are obscured by discrete dependent measures.

This dynamic perspective fundamentally reconceptualizes action control as a continuous process of real-time adjustment rather than a sequence of discrete decisions followed by ballistic execution. It empirically supports the view that control emerges from ongoing coupling between perception and action, consistent with both Shepherd's emphasis on guidance as a continuous relational property and Gentry's notion that semantic information processing involves continuous interpretation of sensory signals to adjust behavior. The continuous trajectories reveal not just what action was selected but how the selection process unfolded over time, making visible the competition and resolution of alternative action plans.

1.2.9 Research Gaps and Thesis Contributions on Human Action Control

The research reviewed above has established several fundamental principles of human action control. Actions are represented through their anticipated effects and share common codes with perception, providing a foundation for semantic information processing. Control is hierarchically organized, with higher-level goals constraining lower-level action selection through cascading intentional structures. Multiple mechanisms enable flexible action selection, including inhibitory control, conflict monitoring, and the balance between proactive and reactive control strategies. Action control involves multiple systems, including goal-directed and habitual mechanisms that dominate under different conditions and reflect trade-offs between flexibility and efficiency. Motor execution reflects sophisticated computational principles, particularly predictive processing and optimal feedback control, that solve coordination problems under uncertainty.

Despite these advances, significant gaps remain in our understanding of human action control. First, most empirical research lacks task continuity, relying on highly simplified laboratory paradigms with discrete trials, minimal environmental complexity, and actions with no meaningful consequences. Such task structures abstract away the dynamics of real-world action and limit the formation and translation of higher-level goals into ongoing action regulation.

Second, cross-level integration within the control hierarchy remains empirically underspecified. In particular, it is unclear how high-level goals constrain and coordinate lower-level perceptual and motor processes in real time. Few studies have examined action control in truly dynamic environments where goals, affordances, and optimal actions change continuously and unpredictably.

Ultimately, the field has developed a sophisticated understanding of isolated mechanisms, but there is no consensus on how these mechanisms operate together in the integrated system that produces adaptive behavior in realistic contexts.

This dissertation addresses these gaps by investigating action control in a dynamic environment that requires continuous sensorimotor regulation under varying levels of uncertainty. Rather than isolating individual mechanisms through simplified laboratory tasks, I examine how multiple control processes coordinate to maintain goal-directed behavior when environmental demands change unpredictably.

1.3 Bridging Theory and Measurement: The Sensorimotor Interface

Building on the conceptual frameworks of Shepherd, Gentry, and Pacherie, and informed by the empirical findings reviewed above, I propose that control emerges from the coordination of structural capacities, representational processes, and intentional organization. Importantly, this view suggests that control is not limited to internal mental states but extends through the dynamic interaction between agent and environment. If control relies on flexible regulation that is guided by semantic information, then we should expect to observe these regulation processes most clearly at the sensorimotor system, the interface where agents directly engage with their environment. Here, semantic information is ultimately translated into concrete motor adjustments and updated through sensory feedback. This involves the integration of sensory information for motor planning, the coordination of motor output, and the ongoing refinement of movements to maintain precision and stability toward a goal.

This insight motivates the empirical approach of this dissertation. Rather than treating control as a purely internal cognitive process, I examine how it operates within the sensorimotor dynamics that link perception and action. The dynamic approaches reviewed above (see Spivey, Grosjean, and Knoblich, 2005; Dale, Kehoe, and Spivey, 2007) demonstrated that continuous motor trajectories reveal ongoing cognitive processes. I extend this logic to eye movements, which represent a particularly rich window into action control. Eye movements are both perceptual (directing gaze to gather information) and motor (precisely controlled ballistic movements), making them an ideal domain for investigating the sensorimotor coupling that underlies control.

Moreover, subtle patterns of eye movement control may reflect fine-grained control processes operating at multiple levels of the intentional hierarchy. If the hierarchical control architecture proposed by Pacherie and empirically supported by Norman and Shallice (1986) and Cooper and Shallice (2000) extends to oculomotor control, we should expect different patterns of eye movements to reflect different levels and modes of control.

To investigate these processes, all studies reported in this work use a custom experimental environment: the Dodge Asteroids task. In this computer-based task, participants use keys on a keyboard to steer a small spaceship through a corridor filled with obstacles. The environment allows precise measurement of motor behavior (steering trajectories) and, when combined with eye tracking, also captures the sensory processes involved in control (gaze patterns). The task requires continuous sensorimotor coordination to maintain goal-directed behavior under uncertainty. This makes it an ideal testbed for investigating the mechanisms of action control in dynamic environments.

Two key manipulations modulate environmental dynamics within the Dodge Asteroids task. First, input noise reflects motor variability, introducing uncertainty into steering by sampling each movement step from a normal distribution with varying standard deviation (higher deviation means higher uncertainty in the mapping between intended and actual movement). Second, drift displaces the spaceship in a predictable direction without user input, indicated by a visual cue. These manipulations allow systematic investigation of how control processes adapt to different forms of environmental uncertainty: unpredictable variability (noise) versus predictable systematic perturbations (drift).

The Dodge Asteroids environment thus operationalizes several key concepts from the theoretical framework. Shepherd's guidance is reflected in how well participants maintain position within the corridor and avoid obstacles. Flexible repeatability is tested by varying noise and drift across conditions: successful control requires adjusting strategies to accommodate these changing dynamics. Gentry's semantic information processing is evident in how participants use visual information (obstacle positions, drift cues, corridor boundaries) to predict future states and plan corrective actions. Pacherie's intentional hierarchy is reflected in the coordination between high-level goals (reach the end safely), proximal intentions (maintain safe position, avoid upcoming obstacle), and motor intentions (execute specific steering adjustments).

The first two studies introduced in Chapter 2 use this environment in combination with eye tracking to explore the sensorimotor functions underlying human action control. Study 1 investigates whether eye-movement control reflects goal-directed behavior under varying levels of uncertainty. I identify two types of fixational eye movements with subtle distinctions in oculomotor patterns that may reflect different modes of action control. Building on these findings, Study 2 examines how these two types of fixational eye-movement control systematically change with the predictability of environmental dynamics and how these changes relate to the ability to exercise control (as framed by Shepherd) and to guide actions appropriately based on semantic information (as framed by Gentry). In Study 3 covered in Chapter 3 , I model the subjective experience of control using computational frameworks that parallel the integration across multiple levels of the control hierarchy identified by Pacherie. Finally, in Study 4 (Chapter 4), I integrate these findings into a unified cognitive architecture that learns and adapts through interaction with the environment. The cognitive model implements the hierarchical, predictive, and adaptive principles identified in both the philosophical and empirical literature.

By moving systematically from behavioral observation through computational modeling to cognitive architecture, this dissertation bridges philosophical accounts of what control is and empirical understanding of how it works. In doing so, it aims to contribute to both a deeper theoretical understanding of human action control and the development of artificial agents with similar capabilities.

2 Extending Control Through the Sensorimotor Interface

The sensorimotor system provides a natural site for investigating extended control empirically. Eye movements, in particular, offer a high-resolution window into how control adapts in real time. They continuously align sensory sampling with task demands and reflect the interplay of prediction, attention, and motor planning.

This chapter examines eye-movement control as a specific instance of sensorimotor coordination. I focus on the adaptation of gaze behavior under uncertainty and what this reveals about the mechanisms supporting flexible regulation in dynamic environments.

2.1 Eye Movements as Indicators of Predictive Sensorimotor Control

At the sensorimotor interface, internal predictions meet external dynamics. The brain continuously generates expectations about sensory input and updates them through feedback, minimizing the difference between predicted and actual input (K. Friston, 2010; Pezzulo, Rigoli, and K. J. Friston, 2018). Within this predictive framework, control operates through anticipation rather than reaction and can be assessed by the extent to which internal models capture environmental dynamics and guide action accordingly. This way, cognition, perception, and movement form integrated components of an extended control system whose interplay can be examined through the sensorimotor interface (Engel et al., 2013). Among the various sensorimotor functions, eye movements provide a particularly direct view of this predictive coupling. Because gaze behavior unfolds on a scale of milliseconds, it may serve as a sensitive indicator of adaptive control under changing environmental conditions.

In natural behavior, eye movements do not merely respond to stimuli, rather they anticipate them. They guide the acquisition of visual information that supports ongoing and future actions (Land and B. Tatler, 2009; B. W. Tatler et al., 2011). The spatial and temporal patterns of gaze reveal shifts in control strategies as uncertainty rises and predictive demands grow (Wolpert and Landy, 2012). Examining eye-movement control thus makes the coordination of internal models and sensory feedback observable and empirically testable.

2.2 Investigating Eye-Movement Control in Dynamic Environments

2.2.1 Study 1: Unsupervised Clustering Uncovers Two Distinct Types of Fixational Eye-Movements in Dynamic Environments

Study 1 examined whether eye-movement control reflects goal-directed adaptation under varying input noise in the Dodge Asteroids environment. Using unsupervised clustering, two distinct types of fixational eye movements were identified, called Type 0 and Type 1 fixations, which differed systematically in spatial allocation and temporal dynamics. Type 0 fixations occurred less frequently but were generally longer in duration and located closer to the spaceship, whereas Type 1 fixations were shorter and tended to occur farther from both the spaceship and the closest obstacles.

As input noise increased, the number of fixations of both types decreased, but their relative proportion remained stable. This suggests that both fixation types are essential and maintained in balance for efficient control. Type 0 fixations became shorter in duration and more narrowly centered around the spaceship, indicating increasingly focused monitoring of the immediate control region when steering uncertainty rose. In contrast, Type 1 fixations became longer in duration and shifted farther from nearby obstacles, often toward open regions in the environment. This spatial redistribution suggests that under higher motor uncertainty, participants used longer fixations to support predictive tracking of the target locations they steered toward, consistent with the notion of predictive sensorimotor control described by Wolpert & Landy (2012).

Together, these patterns indicate that gaze behavior adapts flexibly to uncertainty by adjusting both fixation duration and spatial deployment. The identification of two distinct fixation types further reveals that gaze control is composed of complementary mechanisms that fulfill different functional roles. Type 0 fixations appear to support immediate, reactive monitoring near the controlled spaceship, whereas Type 1 fixations enable proactive tracking of anticipated target locations. This division suggests that adaptive control involves a dynamic balance between reactive stabilization and proactive planning, with eye movements actively coordinating both modes of adaptation.

These findings establish that gaze behavior provides a sensitive indicator of sensorimotor control and reveal distinct modes of adaptation, reactive and proactive, that may generalize beyond oculomotor behavior. This raises the question of how these modes systematically adapt to different types of environmental dynamics beyond motor variability alone.

Contribution to Study 1

I conceptualized the study idea and design. I programmed the experiment and contributed majorly to data collection (participants were sourced through a local participant pool for students at the University of Potsdam). I analyzed and interpreted the data and created all visualizations. I wrote the manuscript.

Study 1

Heinrich, N. W., Jahanian Najafabadi, A., & Pérez-Osorio, J. (2025).
Unsupervised Clustering Uncovers Two Distinct Types of
Fixational Eye-Movements in Dynamic Environments.
2025 IEEE Signal Processing in Medicine and Biology Symposium
(SPMB).

Unsupervised Clustering Uncovers Two Distinct Types of Fixational Eye-Movements in Dynamic Environments

N. W. Heinrich[1], A. Jahanian Najafabadi[2] and J. Perez-Osorio[3]

1. Dept. of Computer Science & Engineering, Universität zu Lübeck, Lübeck, Germany
2. Dept. Cognitive Neuroscience, Bielefeld University, Bielefeld, Germany
3. Cognitive Psychology & Ergonomics, Technische Universität Berlin, Berlin, Germany
nils.heinrich@uni-luebeck.de, amir.jahanian@uni-bielefeld.de, j.perez-osorio@tu-berlin.de

Abstract— **The present study extends previous work investigating eye-movement control in a dynamic, time-constrained task environment. Earlier studies showed that the predictability of environmental dynamics influenced fixation allocation and the initiation sites of smooth pursuits. In those experiments, however, eye movements were clustered solely based on whether a reference point fell within foveal or peripheral vision, which may have confounded the analyses. This study introduces a bottom-up, data-driven clustering approach to identify fixation types across multiple spatial and temporal dimensions. Six participants steered a spaceship to avoid obstacles under varying levels of motor control uncertainty. We identified two fixation types: Type 0 fixations, which were longer and centrally located near the spaceship, and Type 1 fixations, which occurred farther from the agent and were directed toward more open areas of the screen. Linear mixed modeling revealed that with increasing input noise, Type 0 fixations became shorter and more focused, while Type 1 fixations became longer and shifted farther from nearby obstacles. These patterns suggest adaptive gaze strategies, with Type 1 fixations potentially supporting predictive tracking under high control uncertainty. Our findings provide further evidence of how eye movements flexibly support action in complex, real-time environments.**

Keywords— Action control, Eye movements, Adaptive gaze strategy, Unsupervised clustering, Linear mixed modeling

I. INTRODUCTION

Action control refers to the cognitive and motor processes that allow humans to select, initiate, and adapt actions in response to goals and changing environments. It involves planning movements, predicting their sensory consequences, and making adjustments based on feedback [1]. In dynamic environments, this process is not linear or pre-scripted: actions are monitored in real time, and behavior is constantly shaped by how well-intended outcomes match actual results. Rather than following a rigid plan, humans adapt their actions on the fly, correct for errors, or abandon goals when needed.

A central component of action control is the visual system. Vision provides real-time information about the environment, allowing humans to plan and guide their actions with precision. During movement, the eyes actively support behavior by anticipating upcoming demands, selecting relevant targets, and monitoring progress [2]. Where we look often reflects our current goals and, in many cases, even predicts our next actions. This close coupling between gaze and motor control has been demonstrated across a wide range of tasks, from simple everyday activities like tea-making [3] to complex behaviors such as driving [4, 5] (for a review, see Hayhoe & Ballard, 2005 [6]). In such contexts, gaze not only signals intention but also enables rapid corrections and flexible adjustments, particularly when the environment changes unexpectedly. Moreover, beyond guiding our own actions, gaze plays a crucial role in predicting the actions of others [7].

Eye movements serve multiple roles during action control. At times, fixations support immediate motor guidance, for example, by locking onto a target that the hand or body is moving toward [8]. In other moments, they serve more exploratory functions, scanning ahead to anticipate what is coming next (e.g., an obstacle or a slippery surface) or checking task-relevant features in the periphery (e.g., road markings on a bike lane). Rather than sticking to a single function, gaze constantly switches between these roles depending on the demands of the task. In stable and predictable environments, visual guidance often dominates, with the gaze closely tracking the current goal. But as tasks become more challenging or when uncertainty increases, exploratory sampling becomes more prominent. In such cases, eye movements reflect not just ongoing actions but also internal processes like planning, monitoring, or even doubt. The balance between these roles potentially offers insights into how humans implement action control strategies to proactively respond to moment-to-moment changes of their environment.

In the present study, we use a continuous, dynamic task with systematically varied levels of predictability to investigate how gaze supports flexible action control in a complex yet controlled setting. Participants steer a spaceship through an environment featuring obstacles and unpredictable steering, input noise, requiring continuous adjustments in their action control. Eye movements were recorded at high sampling rates, tracking rapid shifts in gaze allocation. With this setup, we explored how gaze shifts between guiding immediate actions, anticipating upcoming demands, and adapting to uncertainty in real-time.

This work extends on previous studies using the same task environment. In Abalakin et al. (2024) [9], we manipulated the predictability of environmental drift, a passive displacement of the spaceship that was either visually indicated or not. We found that fixational eye-movements converged toward an optimal distance from the spaceship when the drift was visible (predictable), whereas invisible (unpredictable) drift prevented such adaptation. In Heinrich et al. (2024) [10], unpredictability in steering was varied on a smaller scale with less distinct levels than in the present study. There, we identified smooth pursuit eye-movements that moved through the environment and typically terminated closer to the spaceship than when they were initiated. These pursuits were initiated closer to the spaceship under mild steering uncertainty but farther away under moderate steering uncertainty. In both studies, fixational eye movements were clustered solely based on whether the spaceship fell within the foveal/parafoveal region or the peripheral field.

Building on these findings, the present study extends the manipulation of input noise and introduces a bottom-up, data-driven clustering approach to identify fixation types from multiple spatial and temporal dimensions. With this, we are aiming for a more comprehensive characterization of how gaze supports real-time action control under varying levels of motor uncertainty.

II. METHODS

This study was approved by the ethics committee of the Technische Universität Berlin (proposal KMDS-WS-01-20190814-E2). Six students at the University of Potsdam were recruited through SONA system. All participants had either normal or corrected-to-normal vision (with contact lenses) and no known history of neurological disorders. Before the experiment, informed consent for research and publication was obtained from all individuals. To prevent potential biases, participants were not briefed on the study's hypotheses but only about the task itself. Participants were compensated with 1.5 participation hours.

II-A. Experiment and Procedure

We used the same experimental environment as in our previous studies: **Dodge Asteroids**, a custom environment built in `Python` [11] using the `PyGame` library [**?**]. The environment runs at 60 FPS and builds on a simulation framework originally developed by Kahl et al. (2022) [12] to evaluate a computational model of action control. Dodge Asteroids provides a well-controlled yet dynamic setting, where the environment changes independently of the agent's behavior, while still allowing the agent to act freely and select sub-goals. This combination makes it particularly well suited for investigating the temporal dynamics of action control. Paired with high-frequency eye tracking, the environment allows us to examine how the statistical properties of action goals evolve in response to increasing uncertainty in the agent's world model [9] or loss of motor control [10, 13].

Obstacles are scattered throughout the environment, with their x- and y-positions drawn from a uniform distribution bound by the width and height of the environment.

Participants steered a spaceship that automatically *falls* through the environment but can be pushed in either horizontal direction using the keyboard (Y = left, M = right; QWERTZ layout). The task was to reach the bottom end without crashing into obstacles.

The Dodge Asteroids environment had a width of 720 pixels and a height of 13,500 pixels. Free fall was 6 pixels each frame. Spaceship and obstacles were 36 pixels in width and height. At any moment during gameplay, only a small part of the full environment, the observation space, was visible (see left side of Figure 1 a). A full example layout is shown on the right of Figure 1 a.

We manipulated the accuracy with which the spaceship is steered by means of *input noise*. Each frame, a key is down to steer the spaceship, its' horizontal displacement is drawn from a normal distribution centered above 6 pixels and with varying standard deviation. Increasing standard deviations impose increasing uncertainty in motor control. There were 5 levels of input noise given by the standard deviations of 0 (no noise, i.e., accurate control), 0.5, 1.0, 1.5, and 2.0.

The experiment was presented on a 28" ASUS PB277Q screen with a 1920x1080 resolution and a refresh rate of 60Hz (equal to the FPS of the Dodge Asteroids environment). Participants were seated with their heads stabilized on a chin rest positioned 80cm from the screen. Eye movements were recorded using a ViewPixx TRACKPixx eye-tracker (VPixx Technologies, Saint-Bruno, QC, Canada), which tracked both eyes at a sampling rate of 2,000Hz. The setup was mounted on a height-adjustable table, ensuring the chin rest and participants' faces remained at a consistent height relative to the screen. The spaceships' position on the screen was kept constant as a static reference point, fixed at coordinates x = 954 and y = 270 pixels (position of the upper left corner of the sprite). Consequently, any movement in the environment, whether due to free fall or steering, caused the surroundings to move relative to the spaceship at the screen center. New objects appear at the bottom of the screen and travel upwards until exiting at the top, taking approximately 2.45s from bottom to top. A gray bar, 270 pixels in height and spanning the full screen width, was displayed at the bottom to discourage gaze shifts beyond the screen (in

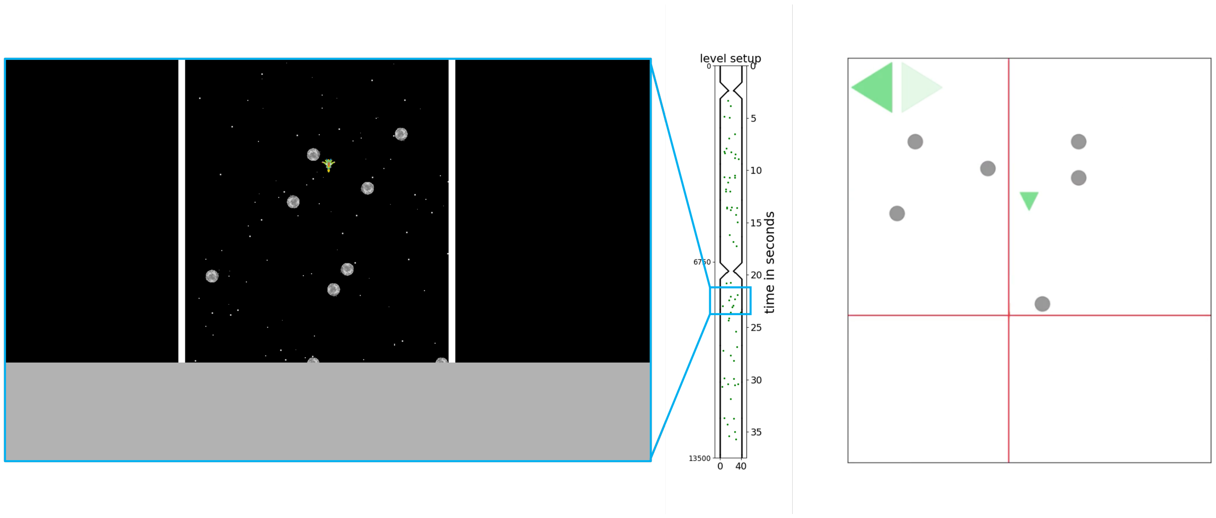

(a) **(b)**

Figure 1. (a) Instance within the Dodge Asteroids experimental environment. The green spaceship remains static at this position on the screen. Obstacles are spread throughout the instance. The full layout of this trial is shown on the right. (b) Simplified representation of a (different) instance within the Dodge Asteroids environment. Here, the participant is steering to the left, indicated by the bright green-colored triangle in the upper left corner. The red cross marks the gaze position.

the flight direction of the spaceship and where new objects appear).

Before the experiment was conducted, we generated 6 distinct environment layouts by randomly placing 62 obstacles (sampling x- and y-positions from a uniform distribution as described above). Participants completed each layout under all five input noise levels, resulting in 30 unique combinations of layout and noise condition.

Prior to every session, the eye-tracker was calibrated individually for each participant using a 9-point grid. Participants then completed a training level measuring 36,000 pixels in length to familiarize themselves with the task. After training, participants solved the 30 unique combinations of layout and input noise that were presented in a randomized order. If a participant crashed during a trial, the same combination was reintroduced later in the session. Each combination could be attempted up to three times. However, combinations that resulted in three consecutive crashes were removed from the participant's remaining sequence. This procedure was designed to prevent the learning of specific obstacle patterns through repetition. Prior to each new trial, a 5-point eye-tracker recalibration was performed, allowing participants to take short breaks in between trials during which they could disengage from the chin rest if needed. Both the 9-point and 5-point calibrations were only accepted if the mean error of both eyes during validation was below 50 pixels. Otherwise, the calibration was repeated. Participants completed the experiment in roughly 60 minutes.

For the present analyses, we considered only trials without crashes, since crash trials end prematurely, affecting fixation counts and durations, and fixations just before crashes may involve distinct action-related processes; these will be examined in a separate study.

III. DATA ANALYSIS

Fixation detection was performed as follows. We identified saccades in the dataset using an established velocity-based detection algorithm [14, 15]. A sample was classified as part of a saccade if the eye moved with a velocity of at least 0.5° for four or more consecutive samples (≥ 0.002s), with the velocity threshold determined using a multiplier of $\lambda = 6$. Periods in which either eye signal was lost were labeled as blinks. Blinks, as well as the events directly preceding or following a blink (i.e. fixations or saccades), were excluded from further analysis. Finally, fixations were defined as time intervals between consecutive saccades. In total, we identified 31,505 fixations on the basis of which we conducted the analyses described below.

Briefly going over the structure of the data, each row is a fixation which is linked to its distance to the spaceship in visual degrees, its fixation duration in seconds, its distance to the closest obstacle in visual degrees, and the input noise in the trial the fixation occurred.

All data processing and analyses were conducted using `Python` (v3.9.18), and relevant scientific libraries (`numpy` [16] and `pandas` [17]). For reproducibility of our results, we used a random seed(36)[1].

III-A. Principal Component Analysis

Principal Component Analysis (PCA) was applied using the `scikit-learn (sklearn)` library [18] to explore the structure of visual and spatial information related to participants' fixations during spaceship navigation. PCA is an unsupervised dimensionality reduction method that identifies orthogonal components maximizing the variance in the feature space, without reference to any target variable. We entered three features: distance to the spaceship, fixation duration, and distance to the closest obstacle. All features were z-standardized prior to the analysis.

The PCA yielded three orthogonal components that together explained 100% of the total variance in the projected data, with individual components accounting for 39.9%, 33.2%, and 26.9%, respectively. Since each component explained a substantial portion of the variance, all three components were kept for the clustering analysis described in the next section.

III-B. Quantile-Based Clustering

We applied quantile-based clustering [19] to the three identified features using the `QuClu` library. Here, data points were assigned to the closest quantile (or a weighted sum of quantiles in higher dimensions). We chose this clustering method over standard clustering techniques (e.g., k-means) because it is more robust to outliers and non-Gaussian or skewed distributions and allows for variable-wise scaling and quantile normalization [20]. This makes quantile-based clustering especially beneficial for fixation data, where feature distributions often exhibit heavy tails, unequal scaling, and non-linear effects.

Clustering was performed with B = 50 resamplings to ensure stability of the quantile estimates. The number of clusters K was varied from 2 to 10. For each k, we computed the silhouette score, a standard internal validity index that quantifies how well-separated and cohesive the resulting clusters are (Figure 2). The highest silhouette score was obtained for K = 2, suggesting that the fixation data consists of two types of fixations.

All fixation events were assigned to the distinct fixation types based on the resulting clusters (Type 0 and Type 1). Next, two fixation types were described based on their original feature values, highlighting the differences in spatial and temporal fixation patterns. Following the descriptive analysis, we investigated whether these types also show distinct sensitivity to input noise using linear mixed modeling.

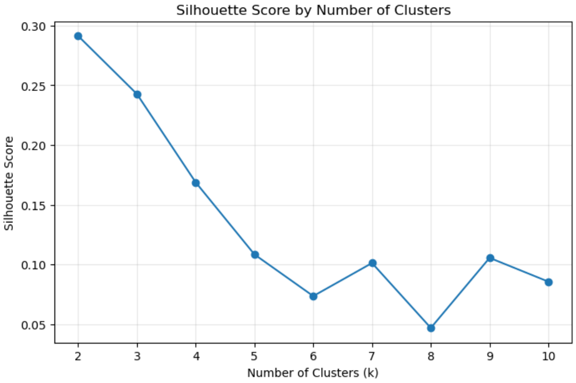

Figure 2. Silhouette scores (y-axis) for fixation data clustered into various number of clusters (x-axis).

III-C. Descriptives

There could be a general difference in how often fixations of the different types are initiated. Therefore, we calculated the total number of fixations per fixation type for each trial, resulting in one data point per trial and fixation type ($N_{\text{fix},0}$ and $N_{\text{fix},1}$). The mode of Type 0 fixations per trial, estimated using kernel density estimation (KDE), was 22.525. For Type 1 fixations, the mode was 29.253 fixations per trial.

Type 0 fixations featured a mean distance to the spaceship of 2.544° (1.489), a mean fixation duration of 0.240s (0.370), and a mean distance to the closest obstacle of 3.305° (1.920).

Compared to Type 0 fixations, Type 1 fixations featured a greater mean distance to the spaceship of 8.198° (2.473), a shorter mean fixation duration of 0.178s (0.363), and a greater mean distance to the closest obstacle of 4.241° (1.917).

Note that the distributions of all variables of interest are skewed to the right, with more probability density gathered in the right tail compared to the left tail (Figure 3). To mitigate potential violations of model assumptions, we explored appropriate data transformations before proceeding with linear mixed modeling.

III-D. Linear Mixed Modeling

Box-Cox distributional analyses were conducted [21] using the `scipy.stats` library [22] for data transformation. Applied transformations and corresponding lambda values are stated in the individual paragraphs. Linear mixed models were built using the `statsmodels.formula.api` library [23] and in every model, the participant ID was entered as a random

---

[1] Git repository containing experimental data and Jupyter Notebook (`Clustering&Analysis.ipynb`) with code for all analyses described in this paper.

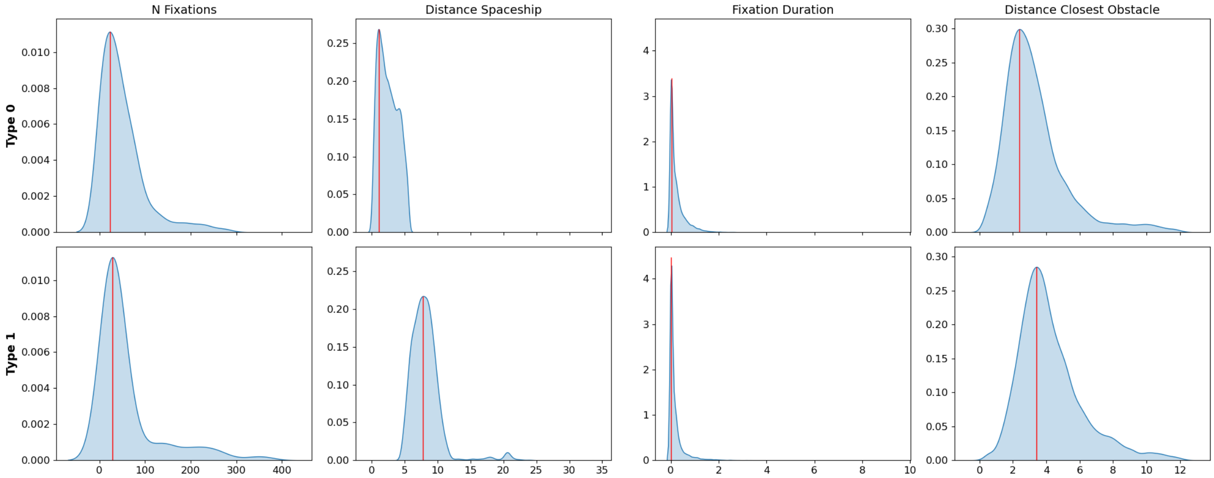

Figure 3. Kernel density estimates (KDEs) for 4 different metrics (columns) of Type 0 and Type 1 fixations (rows) identified through quantile-based clustering. The first column shows the number of fixations of the respective type in a given trial. Columns 2 to 4 are the metrics based on which the 2 fixation types were identified in the clustering process.

intercept effect and input noise as a numerical fixed effect.

First, we investigated whether the number of fixations of each type varies systematically with input noise, under the hypothesis that different fixation types i.e. their roles become more or less relevant as uncertainty in motor control increases. $N_{\text{fix},0}$ was log transformed (λ = 0.239). The number of Type 0 fixations decreased with increasing input noise (β = -0.181, σ = 0.079, 95% CI -0.335 - -0.027, p = .021). $N_{\text{fix},1}$ was also log transformed (λ = 0.042). We also found a significant decrease in the number of fixations with increasing input noise (β = -0.415, σ = 0.069, 95% CI -0.550 - -0.279, p $<$ .001).

We then turned to more detailed analyses, using individual fixations as data points for the respective cluster. Each fixation was linked to its specific characteristics (distances to objects on the screen and its‘ duration) and the input noise level of the trial during which it occurred.

Type 0 fixations‘ distance to spaceship in visual degrees was square root transformed (λ = 0.539). We found that the higher the input noise level was, the closer Type 0 fixations were allocated to the spaceship (β = -0.023, σ = 0.006, 95% CI -0.035 - -0.012, p $<$ .001).

A reciprocal transformation was applied to the distance to the spaceship in Type 1 fixations (λ = -0.909). There was no significant effect of input noise on distance to spaceship (β = -0.001, σ = 0.000, 95% CI -0.001 - 0.000, p = .248).

Fixation duration of Type 0 fixations was log-transformed (λ = 0.093). Input noise had a significant negative effect (β = -0.097, σ = 0.018, 95% CI -0.132 - -0.063, p $<$ .001), meaning that the higher the uncertainty in motor control was, the shorter Type 0 fixations lasted.

Type 1 fixations‘ duration was also log-transformed (λ = -0.017). Contrary to the effect of input noise in Type 0 fixation duration, increasing input noise was associated with longer fixation durations (β = 0.043, σ = 0.016, 95% CI 0.011 - 0.075, p $<$ .009).

We log-transformed the distance to the closest obstacle of Type 0 fixations (λ = 0.163). Input noise had no significant effect on the variable of interest (β = -0.001, σ = 0.007, 95% CI -0.014 - 0.013, p = .981).

For the distance to the closest obstacle in Type 1 fixations, the Box-Cox distributional analysis also indicated a log transformation (λ = 0.246). Increasing input noise was associated with longer distances to the closest obstacle (β = 0.034, σ = 0.005, 95% CI 0.023 - 0.044, p $<$ .001).

IV. DISCUSSION

With the experiment described here, we continue a series of experiments investigating the relationship between oculomotor control and action control. Participants navigated a spaceship through a dynamic but well-controlled environment, attempting to avoid collisions with obstacles while dealing with varying levels of motor noise that impaired their control. We recorded their eye movements throughout gameplay to investigate how gaze supports action control under uncertainty. Based on the assumption that fixations serve different functional roles during action control, we expected that distinct types of fixations could be distinguished statistically. We therefore applied quantile-based clustering and analyzed the obtained clusters using linear mixed modeling.

Type 0 fixations were initiated less often in a trial than Type 1 fixations. In both fixations types, the number of fixations decreased with increasing input noise.

Type 0 fixations were generally closer to the spaceship and longer in duration than Type 1 fixations. As input noise increased, they became even more tightly focused on the ship and decreased in duration. Their distance to the nearest obstacle, however, remained unaffected. These fixations may be used to monitor the regions in the immediate vicinity of the spaceship, and the effects of input noise indicate that monitoring becomes more rigorous with heightened demands on real-time motor control. Interestingly, the spatial pattern of Type 0 fixations resembles the center bias in fixations reported by Burlingham et al. (2024) [24]. In the context of our task, Type 0 fixations are allocated near a central reference point, potentially to optimize energetic efficiency, particularly under conditions of control loss or increased task difficulty.

Type 1 fixations, in contrast to Type 0, tended to occur farther from both the spaceship and the nearest obstacle and were shorter in duration. As input noise increased, these fixations were allocated even farther from obstacles and became longer in duration, while their average distance to the spaceship remained stable. Notably, they were often directed toward relatively empty regions of the screen, as indicated by their greater distance from nearby obstacles. This spatial pattern suggests that Type 1 fixations may serve to track anticipated target locations, the positions of which change over time due to the dynamic visualization of the task. In line with this interpretation, our findings appear to recover the two types of fixation roles described in Lisberger (2015) [8], with Type 0 fixations anchored near the spaceship and relying on peripheral vision to monitor environmental features, while Type 1 fixations are more directly locked onto target positions. Moreover, the increasing distance to obstacles under higher input noise, when control outcomes become more uncertain, may reflect an adaptive, forward-looking gaze strategy aimed at reducing risk by gathering information about future action-relevant locations.

If Type 1 fixations are directed toward target positions, and these target positions move because the environment shifts around a spaceship that remains centered on the screen, then they are not fixations in the strict sense. Rather, they likely represent smooth pursuit eye movements, with their velocity corresponding to the motion of anticipated target positions across the screen (similar to the *foveated action goals* described by Heinrich et al., 2024 [10]). Smooth pursuits can directly support motor behavior, as the pursuit signal itself may be sufficient to improve (hand) motor control during interactions with moving objects [25]. This characteristic could be another distinguishing feature between Type 1 and Type 0 fixations. Type 0 fixations are primarily directed at the spaceship itself, an object whose position remains stable on the screen, and thus would not elicit smooth pursuit. Investigating whether Type 1 fixations indeed exhibit smooth pursuit-like properties could be a valuable direction for future experiments.

Beyond their role in efficient visuomotor control, Type 0 and Type 1 fixations may reflect distinct modes of self–environment coupling. Stable, Type 0 fixations help maintain a coherent perceptual anchor and predictable sensory input, whereas Type 1, pursuit-like fixations support adaptive engagement with dynamically changing action goals. Smooth pursuits rely on predictive mechanisms, such as corollary discharge, to maintain accurate tracking once the fovea is stabilized on a moving target, and disrupted predictive signals can impair pursuit performance [26]. This functional distinction suggests that the balance between Type 0 and Type 1 fixations may be critical for efficient action control: Type 1 fixations allow participants to track anticipated target locations, especially when motor control is uncertain, while Type 0 fixations preserve stability when cognitive resources cannot be fully allocated to action selection. Notably, these same perceptual and oculomotor processes are often disrupted in clinical populations, such as patients with schizophrenia and borderline personality disorder, who exhibit unstable self-perception, altered smooth pursuit, and difficulties maintaining spatial anchoring. Our classification framework could thus provide a computational or behavioral marker of how efficiently individuals balance sensory stability and goal-directed engagement, a balance that appears perturbed in disorders involving self-disturbances.

This study may present a novel approach to studying eye movements in dynamic, time-constrained environments. It is inspired by scenarios such as natural driving but simplified in terms of complexity and participant behavior (e.g., head movements constrained by a chin rest) for better experimental control. The paradigm combines a screen-based task, where participants freely choose subgoals, with high-frequency eye tracking. With this setup, we explore the functional roles of eye movements in real-time action and how these roles adapt under varying levels of control uncertainty. We are continuously refining the paradigm to validate our findings, and welcome discussions and collaborations.

ACKNOWLEDGMENT

This research was funded by the German Research Foundation (DFG) Priority Program 2134 'The Active Self'. We thank the Potsdam EyeLab team, led by Prof. Ralf Engbert, for their support in data collection.

REFERENCES

[1] B. Elsner and B. Hommel, "Effect anticipation and action control," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 27, no. 1, pp. 229–246, 2001.

[2] M. F. Land, "Eye movements and the control of actions in everyday life," *Progress in Retinal and Eye Research*, vol. 25, no. 3, pp. 296–324, 2006.

[3] M. Land, N. Mennie, and J. Rusted, "The roles of vision and eye movements in the control of activities of daily living," *Perception*, vol. 28, no. 11, pp. 1311–1328, 1999.

[4] F. I. Kandil, A. Rotter, and M. Lappe, "Driving is smoother and more stable when using the tangent point," *Journal of vision*, vol. 9, no. 1, pp. 11–11, 2009.

[5] R. M. Wilkie and J. P. Wann, "Driving as night falls: The contribution of retinal flow and visual direction to the control of steering," *Current Biology*, vol. 12, no. 23, pp. 2014–2017, 2002.

[6] M. Hayhoe and D. Ballard, "Eye movements in natural behavior," *Trends in cognitive sciences*, vol. 9, no. 4, pp. 188–194, 2005.

[7] J. R. Flanagan and R. S. Johansson, "Action plans used in action observation," *Nature*, vol. 424, no. 6950, pp. 769–771, 2003.

[8] S. G. Lisberger, "Visual guidance of smooth pursuit eye movements," *Annual review of vision science*, vol. 1, no. 1, pp. 447–468, 2015.

[9] N. Abalakin, N. W. Heinrich, A. Österdiekhoff, S. Kopp, and N. Russwinkel, "Allocation of fixational eye movements in response to uncertainty in dynamic environments," *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 46, 2024.

[10] N. W. Heinrich, A. Österdiekhoff, S. Kopp, and N. Russwinkel, "Goal-directed allocation of gaze reflects situated action control in dynamic tasks," *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 46, 2024.

[11] G. Van Rossum and F. L. Drake, "Python 3 Reference Manual," Scotts Valley, CA, 2009.

[12] S. Kahl, S. Wiese, N. Russwinkel, and S. Kopp, "Towards autonomous artificial agents with an active self: Modeling sense of control in situated action," *Cognitive Systems Research*, vol. 72, pp. 50–62, 2022.

[13] A. Österdiekhoff, N. W. Heinrich, N. Russwinkel, and S. Kopp, "Sense of control in dynamic multitasking and its impact on voluntary task-switching behavior," *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 46, 2024.

[14] R. Engbert and R. Kliegl, "Microsaccades uncover the orientation of covert attention," *Vision Research*, vol. 43, no. 9, pp. 1035–1045, 2003.

[15] R. Engbert and K. Mergenthaler, "Microsaccades are triggered by low retinal image slip," *Proceedings of the National Academy of Sciences*, vol. 103, no. 18, pp. 7192–7197, 2006.

[16] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with NumPy," *Nature*, vol. 585, no. 7825, pp. 357–362, 2020.

[17] W. McKinney, "Data structures for statistical computing in python," *Proceedings of the 9th Python in Science Conference*, S. van der Walt and J. Millman, Eds., 2010, pp. 56–61.

[18] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[19] C. Hennig, C. Viroli, and L. Anderlucci, "Quantile-based clustering," *Electronic Journal of Statistics*, vol. 13, no. 2, pp. 4849–4883, 2019.

[20] L. A. García-Escudero, A. Gordaliza, C. Matrán, and A. Mayo-Iscar, "A review of robust clustering methods," *Advances in Data Analysis and Classification*, vol. 4, pp. 89–109, 2010.

[21] G. E. P. Box and D. R. Cox, "An analysis of transformations," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 26, no. 2, pp. 211–243, 1964, publisher: Oxford University Press.

[22] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.

[23] S. Seabold and J. Perktold, "Statsmodels: Econometric and statistical modeling with python," *Proceedings of the 9th Python in Science Conference*, vol. 57, 2010, pp. 92–96.

[24] C. S. Burlingham, N. Sendhilnathan, O. Komogortsev, T. S. Murdison, and M. J. Proulx, "Motor "laziness" constrains fixation selection in real-world tasks," *Proceedings of the National Academy of Sciences*, vol. 121, no. 12, p. e2302239121, 2024.

[25] O. Sinha, S. Madarshahian, A. Gomez-Granados, M. L. Paine, I. Kurtzer, and T. Singh, "Smooth pursuit eye movements contribute to anticipatory force control during mechanical stopping of moving objects," *Journal of neurophysiology*, vol. 129, no. 6, pp. 1293–1309, 2023.

[26] K. N. Thakkar and M. Rolfs, "Disrupted corollary discharge in schizophrenia: evidence from the oculomotor system," *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, vol. 4, no. 9, pp. 773–781, 2019.

2.2.2 Study 2: Using Eye Movements to Understand Sense of Control in Situated Action

The second study investigated how humans adjust their visuomotor control when confronted with different environmental dynamics that are predictable to various degrees. We reused the clustering method from Study 1 to distinguish between Close and Distant Fixations (as in close to or far from the spaceship, as this dimension of the data explained the highest variance). We examined (oculomotor) behavioral adaptation of both fixation types and subjective Sense of Control (SoC; termed Judgement of Control, JoC, in the referenced study) across three experiments manipulating environmental predictability through input noise and drift.

Gaze Adaptation to Environmental Dynamics

When assessing gaze behavior, we controlled for visual complexity by adding the number of obstacles on screen to all linear models used for analysis.

Input noise introduced unpredictability in motor execution and affected gaze allocation in Close Fixations. Moderately strong input noise decreased the distance to the spaceship in Close Fixations, indicating tighter monitoring of the controlled object. With increasing noise intensity up to extreme levels, Close Fixations' distance to the spaceship showed dynamic adjustments. The distance alternated between increasing and decreasing across different uncertainty levels, suggesting that participants adapted their gaze strategies flexibly depending on the degree of motor unpredictability they experienced. There was no consistent effect of input noise on Distant Fixations.

Drift sections imposed predictable but unavoidable lateral displacements and prompted proactive gaze adaptations in Distant Fixations. The number of drift sections on screen significantly increased the distance to the spaceship in Distant Fixations. This indicates that participants looked farther ahead to anticipate forced lateral movements. The spatial redistribution reflects proactive scanning in response to predictable external perturbations.

These patterns demonstrate that visuomotor behavior adapts to both the type and intensity of environmental dynamics. Input noise leads to closer monitoring of the controlled object through reactive gaze adaptation, while drift prompts anticipatory scanning of future positions through proactive adaptation.

Sense of Control and Environmental Predictability

Successfully completing trials consistently increased SoC across all experiments, confirming that performance outcomes do influence subjective control. However, participants' reported SoC also varied systematically with environmental predictability, revealing that subjective control is not entirely based on performance outcomes.

Input noise showed effects on SoC primarily at higher intensities. Lower noise levels that reduced actual performance did not affect SoC, while only the highest noise levels decreased both performance and subjective experience. This pattern indicates that noise needs to be quite severe before it affects SoC, even though it starts affecting performance at moderate-high levels. Participants retained confidence in their control despite declining success rates, suggesting that adaptive action control can compensate for moderate motor variability and preserve the subjective experience of control.

The presence of visual cues for drift sections did not affect subjective SoC ratings or actual performance. However, when drift sections had no visual cue, both SoC and performance decreased. This demonstrates that the ability to anticipate perturbations determines both subjective control and behavioral success. Critically, uninformative visual cues (when no drift sections were present, but the visual cue was displayed) did not affect SoC and did not affect performance, confirming that participants' responses tracked actual environmental dynamics rather than mere visual features.

Together, these findings suggest that SoC reflects the quality of one's internal model of environmental dynamics rather than performance outcomes alone. When environmental dynamics are predictable and thus allow for accurate predictions (such as with visible drift or moderate noise), individuals maintain a high SoC. Even when such predictable dynamics are challenging and performance declines, SoC remains intact because the loss in performance occurs against a background of predictable action possibilities that are represented in the internal model. Human agents can anticipate what actions are available to them in specific situations, even when those actions are difficult to execute successfully. In contrast, when dynamics cannot be anticipated (such as with invisible drift or extreme input noise), SoC decreases because predictive models fail to capture the environment's behavior. Without an accurate model of environmental dynamics, agents cannot maintain a repertoire of action possibilities to respond effectively, leading to decreases in both SoC and performance.

The results align with frameworks that treat SoC as a form of self-belief updating. Individuals continuously update their expectations about their own efficacy based on the quality of their predictive models and the action possibilities they can afford, rather than on performance outcomes. When the sensorimotor system can anticipate and predict environmental changes, both gaze behavior and subjective experience reflect this predictive capability.

Contribution to Study 2

I fully conceptualized the overall study as well as the ideas and designs for Experiments 1 and 2, and I contributed to the conceptual development of Experiment 3. I implemented the experimental programs for Experiments 1 and 2, and Experiment 3 was developed on the basis of my original script. Participants were sourced through the University of Potsdam's student participant pool, and data collection for all three experiments was primarily carried out by trained student assistants. I analyzed and interpreted all data, created all visualizations, and wrote the manuscript.

Study 2

Heinrich, N. W., Österdiekhoff, A., Kopp, S., & Rußwinkel, N. (2025).
Using Eye Movements to Understand Sense of Control in Situated Action.
Cognitive Science, 49(12), e70154.
`https://doi.org/10.1111/cogs.70154`

COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 49 (2025) e70154

ISSN: 1551-6709 online
DOI: 10.1111/cogs.70154

Using Eye Movements to Understand Sense of Control in Situated Action

Nils Wendel Heinrich,[a,b] Annika Österdiekhoff,[c] Stefan Kopp,[c]
Nele Russwinkel[a]

[a]*Institute of Information Systems (IFIS), Department of Computer Science and Engineering, Universität zu Lübeck*
[b]*Cognitive Psychology and Ergonomics, Technische Universität Berlin*
[c]*Social Cognitive Systems, Faculty of Technology, CITEC, Bielefeld University*

Received 6 March 2025; received in revised form 17 September 2025; accepted 24 November 2025

Abstract

This series of studies investigated the interplay between the Sense of Control, continuous action control, and eye-movement behavior in dynamic and uncertain environments. Across three experiments, we used a custom-designed environment combined with eye-tracking to examine how action goal pursuit and visual strategies were adapted to deal with motor perturbations of varying predictability. Participants steered a spaceship, avoiding walls and obstacles while contending with random input noise and predictable horizontal drift. We found that changes in fixation distances to a reference point, the spaceship, indicated the type of action control employed. Input noise was associated with decreasing distances in fixations already close to the spaceship, addressing immediate demands for maintaining the spaceship's trajectory. In contrast, fixations allocated within the outer vicinity of the spaceship featured even longer distances in response to drift, suggesting visual exploration and proactive planning. That is, reactive strategies of action control were characterized by immediate responses to unpredictable disturbances, whereas proactive strategies reflected anticipatory adjustments to predictable changes. Furthermore, judgments about the own Sense of Control were closely tied to participants' ability to anticipate and adapt to environmental features. Invisible perturbations led to control loss and reduced task performance, but predictable perturbations allowed participants to maintain a high Sense of Control and still successfully solve the task. These results highlight how

---

Correspondence should be sent to Nils Wendel Heinrich, Institute of Information Systems (IFIS), Department of Computer Science and Engineering, Universität zu Lübeck, Ratzeburger Allee 160, Building 64, 23562 Lübeck, Germany. E-mail: nils.heinrich@uniluebeck.de

cognitive processes and sensorimotor control interact to navigate uncertain environments by flexibly balancing reactive and proactive strategies of action control.

Keywords: Sense of Control; Situated action control; Eye-movement control; Dynamic environments; Statistical modeling; Embodied cognition

1. Introduction

Actions are executed with the intention of achieving a specific outcome, the action goal. At any given moment, an agent is aware of the various outcomes it can achieve with actions. With *action outcomes*, we refer to the changes in the environment the agent can bring about. All these possible actions constitute the agents action field (Kahl, Wiese, Russwinkel, & Kopp, 2022). The cognitive process of choosing among these possibilities within the action field is known as action selection. However, the action field can vary significantly even under similar conditions, partly because actions rarely occur in isolation.

Human agents act within dynamic environments, performing sequences of consecutive actions. Thus, individual actions cannot be considered completely separate from one another. Naturally, the selection of each new action likely depends on how effective preceding actions have been. The evaluation of actions is based on the degree to which they achieved the intended action goal and whether difficulties arose during execution. Said difficulties relate to the violation of the predicted sensory consequences. While we perform an action, a forward model generates predictions regarding the sensory feedback associated with each of the individual motor commands that constitute the action. The prediction is compared with the actual sensory feedback. Mismatches indicate prediction errors, signaling a deviation from the planned course of action. But even if many prediction errors were encountered, an action can still achieve its intended action goal. In fact, agents can use the feedback from prediction errors to adjust and correct actions in real time (Synofzik, Vosgerau, & Newen, 2008). This is similar to cycling in strong side winds that can unexpectedly push the handlebars causing us to deviate from our intended paths. But one can easily correct by steering straight again and mostly stay on course. In human agents, leveraging feedback is discussed to be based on a phenomenal experience, the Sense of Control (SoC; Jeannerod, 2003; Pacherie, 2008). It entails both the feeling of being in control of an action and the need to exert control to maintain the course of an action when encountering difficulties (Pacherie, 2008). Composing action fields, action goal selection, predicting sensory consequences, and correcting the course of actions are all components of action control, and experiencing control loss could lead to an adaptation of any of these components. There are theoretical frameworks attempting to explain the link between the SoC and adaptation of action control, but testing these models remains a challenge, especially with regard to action control within dynamic environments.

To address this challenge, we designed a dynamic experimental environment and introduced conditions known to influence participants' SoC (Österdiekhoff, Heinrich, Russwinkel, & Kopp, 2024). We conducted several experiments using our custom dynamic environment in combination with eye-tracking. We argue that it is possible to infer statistics of action goal properties from gaze patterns on the screen (Heinrich, Österdiekhoff, Kopp, & Russwinkel, 2024). Measuring changes in fixational eye movements in response to the different experimental conditions thus allows us to investigate how the SoC and action control are connected.

1.1. Sense of Control

The SoC is a crucial component of the broader Sense of Agency, which also includes a sense of being the cause of an action and a sense of having initiated the action in the first place (Pacherie, 2008). The Sense of Agency in its entirety would lead to an agent assuming ownership for the action, which enables differentiating between self-generated and externally caused events which fundamentally links to our perception of self.

In literature, it is mostly assumed that the SoC is based on the comparison between predicted sensory consequences of a motor command and the actual sensory feedback (Synofzik et al., 2008). If prediction and feedback match, the SoC is increased (in turn leading to a high Sense of Agency, as discussed by Blakemore, Wolpert, & Frith, 2002; Haggard, 2005; Tsakiris, Prabhu, & Haggard, 2006); conversely, it diminishes when prediction errors in sensory states occur.

But the comparison can go beyond mere sensory states. It can also assess the match between the intended action goal and the actual environmental state that was brought about by the action. This comparison is on higher conceptual levels than the one identifying prediction errors, as it requires the agent to interpret the state it is in. Comparing intended and achieved outcome refers to the *consistency* criterion of the postdictive account of agency (Wegner, 2003; Wegner & Wheatley, 1999), which assesses whether an action has successfully achieved its intended goal. The conclusion can, of course, only be drawn once the action is completed. This aspect of goal completion is likely what individuals recall when asked to retrospectively judge their own control (Chambon & Haggard, 2012). If the consistency criterion is not met, lower values for control judgments are expected. The distinction between prediction errors and the consistency criterion is crucial, because while both affect the SoC, it is quite possible that they trigger different types of action control adaptation.

1.2. Action control

Frameworks explaining action control often draw heavily from motor control accounts and almost all of them assume a forward model that predicts the outcome of the planned motor program, and additionally, an inverse model which infers the sequence of motor commands that will achieve a desired state (Cooper, 2010). This theory of model-based motor control, therefore, relies on mental representations of the environment to link motor commands to sensory states and desired outcomes. But it also explains mechanisms such as error detection, feedback processing, and sensory integration that underpin short-term motor adaptation and longer-term motor learning (Jordan, 1996). The feedback at the core of these mechanisms can be extrinsic (e.g., explicit guidance or information on performance), but in this work, we will focus on intrinsic feedback (specifically through proprioception and vision).

One more recent and comprehensive framework that links intrinsic feedback and adaptation in action control is proposed by Kahl et al. (2022). The authors describe action control as adapting through recurrent feedback loops in situations of reduced control. They refer exactly to the comparator mechanism of which the output is fed into the SoC of the agent and link it to action selection in particular. Kahl et al. distinguish between two types of SoC, each corresponding to a different level in the action control hierarchy. The comparison between sensory states is associated with a low-level SoC that is located at the sensorimotor control layer. In contrast, the comparison between intended and actual action outcome is linked to a high-level SoC that in turn is associated with higher cognitive control levels. If either comparison identifies a mismatch, the respective SoC is decreased. Additionally, a considerable drop in low-level SoC also causes a decline in high-level SoC. The authors discuss how this dynamic high-level SoC influences top-down (goal-directed) action control by restricting the composition of the action field in situations of control loss. Meaning that in challenging situations, the action field will only consist of action possibilities that are more likely to be accomplished given the constraints of diminished SoC (i.e., less likely to elicit prediction errors or fail to meet the consistency criterion). Because the action goal is selected from the action field, consequently, the selected action goal is bound to the SoC in the given moment. However, Kahl et al. did not specify what exactly this means for the properties of action possibilities that more likely constitute the action field and the selected action goal. Specifically, regarding the action goal properties, we intend to provide experimental results and, in addition, explore a new aspect of anticipation that has not yet been discussed.

Finally, Kahl et al. do assume that the sensorimotor control layer is highly capable of responding to drops in low-level SoC all on its own. As in the example of cycling in strong winds, small adaptations in motor control can already overcome the difficulties. With the experiments described in this work, we cannot reliably distinguish between action control adaptation that is elicited by lower or higher levels of the action control hierarchy. Therefore, we will not attempt to map the different ways of adapting on specific levels, but we will discuss different types of action control adaptation.

The mechanisms described above only cover one type of action control adaptation: *reactive* adaptation of action control as it is action control that is adapted in response to control loss (Meiran, Cole, & Braver, 2012). But an agent can already anticipate a challenging situation with an associated decline in SoC and prepare itself accordingly for the upcoming circumstances. Here, the agent engages in *proactive* adaptation of action control, yet this presumes that it has a high degree of control. An increased SoC means that the forward model is accurate that in turn generates accurate predictions. Thus, proactive adaptation of action control is based on the ability to correctly project the circumstances of an upcoming situation. The agent then controls its actions in a way that is directed toward goal completion even under more challenging conditions. This illustrates the bidirectionality of the system. The SoC arises from action control, but also feeds back into action control. Despite extensive research on the phenomenological experience of control, few studies have explored how this experience impacts action control. In this paper, we address the research question of how action control changes with different degrees of SoC. Nevertheless, this poses the challenge of how to measure shifts in action control.

1.3. Inferring action goals from eye movements

In dynamic environments, it is almost impossible to predict when agents will select an action goal and start pursuing it, as individual instances are not explicitly separated from each other. Likewise, action goals can be abandoned at any time to initiate a new action selection process. To better understand human behavior in these settings, we need a reliable way to measure action goals. Visuomotor control theories suggest that visual focus can effectively guide motor responses, resulting in a seamless integration of visual perception and motor actions (Hayhoe, 2017). This can be observed, for example, in many simple mundane activities (e.g., making tea; Land, Mennie, & Rusted, 1999), but is also heavily researched in natural driving, which features a much more complex and dynamic environment. Research shows that drivers tend to steer toward the point they are looking at (Cina & Rad, 2024; Wilkie & Wann, 2002). Steering is visually guided by focusing on journey points on the road ahead, and as the car reaches these journey points, the gaze shifts further ahead again (Land & Horwood, 1995). However, drivers do not keep their gaze fixed on journey points; they periodically sample from various parts of the road (Kandil, Rotter, & Lappe, 2009; Wilkie & Wann, 2002). Additionally, drivers' gaze is directed not straight at journey points but also at other targets important for the driving maneuver, such as other cars during overtaking or the inside of a curve while turning (Lappi, 2022; Land & Lee, 1994). The final gaze location, therefore, is a synthesis of the journey point and the other targets. Even though it may be difficult to directly infer action goals from gaze locations, the current action goal significantly influences where the gaze is directed (Heinrich et al., 2024). This means that tracking the eyes of participants in a navigation task, should enable us to statistically differentiate between action control adaptation that is reactive or proactive in nature. And if we can identify distinct types of action control adaptation, we can attempt to assign them to the different experimental conditions that we know have influence on the SoC.

For the purpose of this work, we have to make specific assumptions about the connection of reactive and proactive adaptation of action control and changes in eye-movement behavior. In the context of continuous action control, reactive adaptation could mean that action goals are selected with shorter time horizons; they target intended outcomes that are not as far ahead in time, as predicting the environmental state becomes increasingly uncertain with more time steps. This would mean that the eyes would be allocated closer to a reference point that is the start of a planned course of action. Proactive action control, on the other hand, would imply that the eyes fixate more strongly on certain areas in which challenging conditions are anticipated. We specifically assess fixational eye movements of participants controlling an agent in a custom-designed experimental environment. We investigate the extent to which the statistics of these fixations can be explained by unpredictable and predicable features of the environment while accounting for the complexity of the visual scene. The resulting differences should also align with self-assessments of control (Judgment of Control; JoC).

2. Methods

2.1. Setup

Participants were seated comfortably in a chair with a front rest with their heads placed in a chin rest at a height of 23.2 cm on a height-adjustable table. The task was presented using a 28'' ASUS PB277Q screen with a 60 Hz refresh rate and a resolution of 1920x1080 pixels. The monitor was placed on the table at 80 cm distance from the participants head, with the center of the screen slightly below eye level. A high-frequency eye-tracking camera was positioned below the screen in 26.4 cm distance to the eyes. For task input, a keyboard on which the relevant keys were marked was placed directly in front of the participants.

2.2. *Dodge Asteroids environment*

Inspired by the simulation environment of Kahl et al. (2022), we have designed a custom experimental environment called the *Dodge Asteroids* environment[1] (Abalakin, Heinrich, Österdiekhoff, Kopp, & Russwinkel, 2024; Heinrich, Russwinkel, Österdiekhoff, & Kopp, 2023; Österdiekhoff, 2023). It is implemented in Python (Van Rossum & Drake, 2009) using the PyGame package (Shinners & Pygame Community, 2000) and runs at 60 FPS. A spaceship of 36 pixels in width and length automatically traverses downward in an environment of 720 pixels width and length of up to 18,000 pixels. The free fall of the spaceship is 6 pixels every time frame. Participants can press either Y or M on the keyboard to move the spaceship 6 pixels to the left or right, respectively. There are walls on both sides of the environment. Obstacles of 36 pixels in width and length are randomly scattered throughout with their x and y coordinates sampled from a uniform distribution bounded by width and height of the environment. The goal in each trial is to steer the spaceship across the finish line at the bottom of the environment without crashing into walls or incoming obstacles. The finish line is indicated by a green line ranging horizontally from one wall to the other.

The Dodge Asteroids environment features two experimental manipulations.

Input noise can make it more difficult to control the spaceship. If a key is pressed to move the spaceship horizontally, in each game frame the key is down, the step is sampled from a normal distribution centered over the usual step size of 6 pixels and with varying standard deviation. This means that with increasing standard deviation, there is more uncertainty linked to the horizontal movement of the spaceship. This will evoke prediction errors in individual sensory states in the short term, but in the long term will also lead to a failure of meeting the consistency criterion when pursuing action goals.

Drift is indicated by red bars of 18 pixels width outside the walls, 45 pixels away. Drift sections span 270 pixels in vertical direction exactly the length of the red bar. As soon as the spaceship enters the drift section, it is moved by 3 pixels (exactly half the step size when a key is pressed) in a specific horizontal direction in each game frame. The direction of the horizontal movement is indicated by the side on which the red bar is displayed: If the red bar is displayed on the left-hand side outside the walls, the spaceship is pushed to the right and vice versa (comparable to wind, which emanates from the drift bar and acts on the spaceship). The positions of drift sections within the environment as well as drift directions are randomized. The point in time of the onset of drift, as well as its effect, is predictable. We use this manipulation to investigate the role of (missing) prediction errors while still affecting the steering of the spaceship.

For Experiment 1, we randomly generated six different layouts of the Dodge Asteroids environment, defined by where obstacles and drift sections are located. Three of the layouts

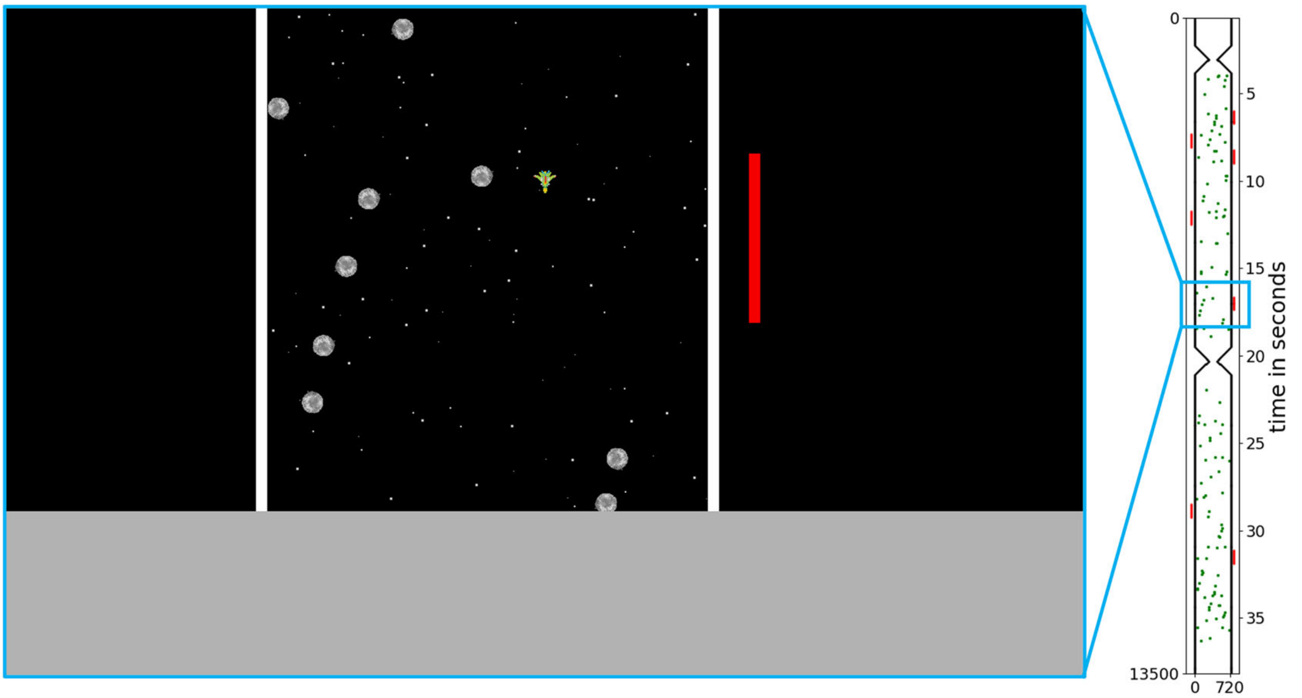

Fig. 1. Instance within the Dodge Asteroids environment (based on the figure by Heinrich et al., 2023). The green spaceship, fixed at the horizontal center of the screen, serves as a reference point for fixational eye movements. Obstacles are spread throughout the instance. The red bar indicates a drift section, in which the spaceship is pushed to the left. The full layout of this trial is shown on the right.

had a length of 9000 pixels (short), the other three have a length of 18,000 pixels (long). We introduced three different degrees of difficulty (easy vs. medium vs. hard) depending on the number of obstacles and drift sections. The number of obstacles in short layouts ranged from 12 (easy), over 34 (medium) to 68 (hard), whereas in long layouts, it ranged from 32 (easy), over 84 (medium) to 168 (hard). The amount of drift in short layouts ranged from four drift sections (easy & medium) to eight drift sections (hard), whereas in long layouts, it ranged from nine drift sections (easy & medium) to 18 drift sections (hard). An example layout is shown on the right-hand side of Fig. 1. Each layout was played with each of the input noise settings while drift could be on or off resulting in a 3x2 design: input noise (with standard deviations of 0, 3, and 6 pixels referred to as none, weak, and strong input noise, respectively) x drift (on vs. off). Therefore, each layout was played in a total of six different configurations, resulting in a total of 36 different configurations.

For Experiment 2, we randomly generated six layouts of the same length, 13,500 pixels. Three of difficulty medium (58–62 obstacles), and the other three of difficulty hard (116–124 obstacles). We did not include drift in Experiment 2. In return, we did introduce two further levels for input noise. The standard deviations of these additional levels were 9 and 12 (0, 3, 6, 9, and 12 pixels of standard deviation). Each layout was played with each of the five different input noise levels, resulting in a total of 30 different configurations.

For the latest iteration of the Dodge Asteroids environment, Experiment 3, we again random generated six layouts of 13,500 pixels length. Here, also three layouts were of medium difficulty (58–62 obstacles) and the other three were of hard difficulty (116–124 obstacles). Every layout regardless of difficulty contained eight drift sections at fixed position. We did not include input noise in this experiment, but rather introduced two new types of drift. In addition to *drift* (drift sections indicated by red bar and acts on spaceship), there was *invisible* drift (red bar was not shown, but drift acts on spaceship) and *fake* drift (red bar was shown, but drift did not act on the spaceship).

Experiment 3 was organized into three blocks, so that one new drift type was encountered at a time. We included a normal drift block as control condition that was always presented first. Within this block, half of the drift sections were normal drift and the other half had no drift. In the invisible block, only half of the drift sections in every trial were invisible, with the other being normal drift. In the fake block, again, only half the drift sections were fake and the other half were normal drift. The order in which the invisible and fake block were presented was counterbalanced across participants. Each layout was played at least twice within a block. Once with four random drift sections as no, invisible, or fake drift and the other four as normal drift. And the other time with the drift types reversed, so that the drift sections that were normal drift before were now no, invisible, or fake drift. This adds up to a total of 36 configurations. We included another training configuration with normal drift only in between blocks to establish a baseline.

During gameplay, participants saw only a small section of the entire environment drawn on the screen, the observation space (enlarged section on the left of Fig. 1). This is done to prevent longer-term planning and promote situated action. It also enables associating eye movements with a given instance.

The following procedure was used across all three experiments.

2.3. Procedure

Once the height of the table had been adjusted and participants had found a comfortable position, the eye tracker was calibrated using 9-point grid. Participants were instructed about the objective of the Dodge Asteroids environment and told that they will encounter difficulties (no explicit instruction of input noise or drift was given). Immediately afterward, participants entered training to familiarize themselves with controlling the spaceship. The random generated layout exclusively for training had a length of 18,000 pixels and featured the statistics of the easy difficulty. In Experiments 1 and 2, input noise of standard deviation 3.0 occurred halfway through the training layout. Likewise, drift was enabled from halfway through the training layout in Experiments 1 and 3. In all three experiments, participants had a total of three attempts to complete the training configuration without crashing. After completion of the training, the experiment started. A recalibration using a 5-point grid was carried out prior to each trial. Participants started the recalibration themselves by pressing the space bar, which was always immediately followed by the start of the next configuration. At the start, the spaceship flew in from the top of the horizontal center screen until it reached screen x- and y-coordinates 954 and 270 above the vertical center of the screen (coordinates refer to the top left corner of the spaceship sprite). From this point onward, participants were able to control the spaceship via key press. Key presses would not result in the spaceship moving across the screen, but the environment moving around the spaceship so that it remained static at all times. For reference, it took roughly 2.45 s for an object that appeared on the bottom to disappear again at the top of the screen. The static center position of the spaceship was chosen to be at eye level.

Regardless of completion or crash, a question assessing general control during steering was presented on the screen after each configuration (JoC; “In the most recent run, how strong was your feeling of control?”). Participants indicated their degree of control by pressing one of the number keys 1–7. A short break was possible after each trial allowing participants to disengage from the chin rest. The order in which the different configurations were played was fully randomized. If a configuration was completed, it was removed from the list of configurations to be played. Otherwise, if a crash occurred, the trial was abandoned and mixed back into the list of configurations to be played. For each configuration, participants were given only three attempts to prevent familiarization with specific instances within the layout. If all three attempts ended in a crash, the configuration was removed from the list of configurations to be played and not presented again.

On average, completing a layout of 9000 pixels length took 29.08 s (2.71), while a layout of 13,500 pixels length took 43.56 s (3.99), and a layout of 18,000 pixels length took 58.59 s (2.07). Deviations can occur because individual frames may last longer than $\frac{1}{60}$ s, depending on how many objects need to be drawn on the screen. In Experiment 1, the average total playing time across all participants was 1867.08 s (120.99). In Experiment 2, it was 2061.13 s (189.11). And, in Experiment 3, the average total playing time was 1549.71 s (92.68).

Once the list of configurations to be played was empty, the main part of the experiment ended, and participants were able to disengage from the chin rest for good. They were given a pen-and-paper questionnaire about video game expertise (Nolan, 2022) that was followed by a semi-structured interview. The interview questions related to whether a total loss of control was experienced while steering the spaceship and what specific difficulties were encountered that may have led to a loss of control. Afterward, participants were debriefed.

The whole procedure was meant to never exceed a duration of 1.5 h.

2.4. *Recruiting and exclusion criteria*

Participants were recruited from the pool of students at the University of Potsdam using the SONA platform. Participants had normal or corrected to normal vision. Participation was compensated by receiving credit points. Each participant provided written informed consent. Not completing the training configuration within three attempts was an exclusion criterion that resulted in no actual exclusion.

Experiment 1 was conducted in the period from January 24 to March 31, 2023, in which a total of 27 participants were recruited. The six participants in Experiment 2 were recruited in the period from May 21 to June 07, 2024. For Experiment 3, we gathered the data of a total of 26 participants in the period from May 30 to June 15, 2023. The average age of participants across all three experiments was 23.57 years (SD = 4.89 years). No participant was allowed to participate in multiple of the experiments.

All three studies were conducted in accordance with the Declaration of Helsinki.[2]

2.5. *Eye-tracking*

Eye movements were recorded binocularly using the ViewPixx TRACKPixx eye tracker (VPixx Technologies, Saint-Bruno, QC, Canada) with a sampling rate of 2000 Hz. The eye tracker recorded X- and Y-axes coordinates of the individual eyes of the participants. To obtain the final gaze position, we computed the Euclidean center point of both eyes. Binocular tracking allowed us to approximate the gaze position to the signal of the other available eye if the signal of one eye was lost.

While navigating the Dodge Asteroids environment (not during instructions or calibrations), a gray bar was displayed at the bottom of the screen, ranging the whole width of the screen with height 270 pixels. This prevented participants from looking outside the screen when positioning the gaze toward the bottom edge of the environment where new objects appeared.

Fixation detection within the data samples of the eye-tracker was done for every experiment individually and by using a velocity-based algorithm. Rows were annotated as fixation when both eyes traveled no more than 0.026° (1.25 pixels) between samples for a minimum of 25 consecutive samples (VPixx Technologies Inc., 2024). These may also have included smooth pursuit eye movements, since if the gaze was not aimed directly at the spaceship, the fixated point moved due to how we implemented movement on the screen visually. We verified the remaining data in a follow-up step by applying a velocity-based saccade detection algorithm (Engbert & Kliegl, 2003; Engbert & Mergenthaler, 2006). Here, rows were annotated as saccade when both eyes traveled at least 0.5° for a minimum of four consecutive samples ($2.00*10^{-3}$ ms). We used a multiplier $\lambda = 6$ to compute the velocity threshold. This way, we detected a total of 96,263, 31,505, and 137,190 fixations for Experiment 1, 2, and 3, respectively.

We had to bear in mind that fixations that foveate the spaceship on the one hand or the surroundings on the other hand have different functions. If the spaceship is within high-accuracy vision, the main focus lies probably on monitoring the individual movements of the spaceship in the immediate vicinity. On the contrary, if the gaze is shifted away from the spaceship and placed within the more distant environment, the main focus is most likely on visual exploration or the pursuit of a specific action goal, to which the gaze is primarily directed. We, therefore, hypothesized a distinction between two types of fixations based on the visual region in which the spaceship is located. *Close* Fixations keep the spaceship within the parafovea thus within a radius of roughly 5° of visual angle from the exact point of fixation at the time it is initiated. Fixations that are initiated further away than roughly 5° of visual angle from the spaceship (spaceship within peripheral vision) are labeled *Distant* Fixations (Fig. 2). We verified splitting the data using a cluster algorithm, the K-quantiles clustering algorithm (Hennig, Viroli, & Anderlucci, 2019), which is a special case of k-means. We applied K-quantiles by using the QuClu package (Hennig, Viroli, & Anderlucci, 2022) in R (R Core Team, 2022). We set k, the number of clusters to be identified = 2, b, the number the initialization process is repeated = 50, and used the default, unconstrained method VS (Variable-wise theta and Scaled variables). The only data dimension passed to the algorithm was the distance to the spaceship. For the data of Experiment 1, the algorithm identified a

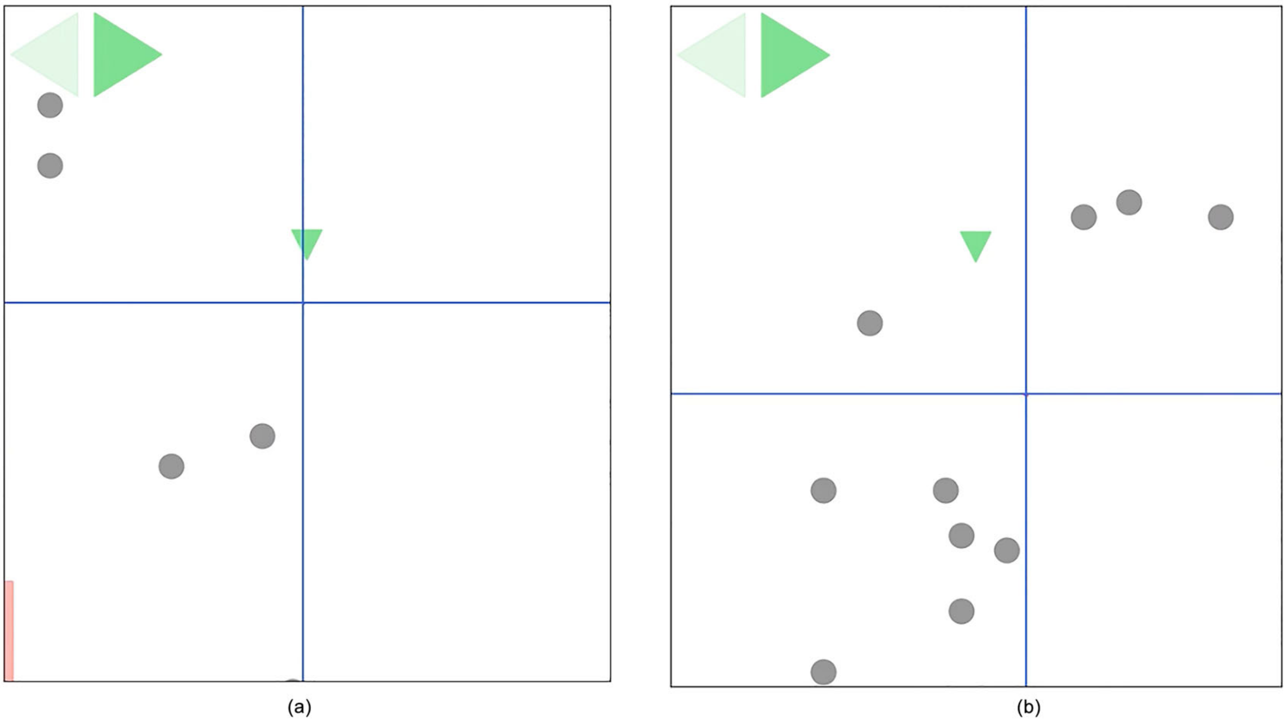

Fig. 2. Simplified illustration of two separate instances of the Dodge Asteroids game environment in Experiment 1. The spaceship is indicated by the small green triangle. The blue cross indicates the fixation location. (a) shows a Close Fixation being executed within a 5.70° radius around the spaceship. (b) shows a Distant Fixation being executed outside a 5.70° radius of the spaceship. In both cases, the participant is steering to the right indicated by the highlighted green arrow. Note that in (a), gaze is directed toward the immediate space in front of the spaceship, whereas in (b), gaze is located farther away in empty space, likely at a point along the planned trajectory of the spaceship within the environment.

threshold value at 5.70°. For Experiment 2, the threshold was identified at 5.65°. In the data of Experiment 3, the threshold was identified at 5.00°. The respective data set was split into Close and Distant Fixations on the basis of the clusters found by advanced K-quantiles.

Hypothesis testing was done individually for Close and Distant Fixations.

2.6. *Measures*

Completion is a variable with binary outcome that represents having completed the configuration. True in the case of managing to cross the finish line at the bottom of the environment and False in the case of a crash. For predicting completion, we included the features of the configuration as covariates (difficulty, input noise level, drift presence, drift type).

Judgment of control (JoC) is the term used for the responses to the general control question after each trial. Participants reported their judgments using a 7-point Likert scale. JoC responses were predicted by the features of the configuration (difficulty, input noise level, drift presence, and drift type). We included two additional covariates: first, the number of fixations executed during the current trial; and second, the number of consecutive crash-completions prior to the current trial, defined as the number of times a crash occurred followed by a completion of the very next trial.

Distance Measures. Distance to Spaceship & Distance to Closest Obstacle were derived from the fixation location at the point in time when the fixation was initiated. For both Close and Distant Fixations, the fixation location was specified regarding the distance to a reference point in degrees of visual angle. The reference point was either the center of the spaceship or the center of the closest obstacle. In models predicting the Distance to Spaceship or Distance to Closest Obstacle, we included input noise level or drift type covariates. We also included the number of obstacles and the number of drift sections visible on screen in the first frame the fixation was detected. With the last two variables, we accounted for the complexity of the visual scene.

2.7. Data analysis

We applied (generalized) linear mixed modeling using the MixedModels package (Bates, 2015) within the Julia programming language (Bezanson, Edelman, Karpinski, & Shah, 2017).

A Box–Cox distributional analysis (Box & Cox, 1964) with the MASS package (Venables & Ripley, 2002) in R indicated that several transformations are required. These were: a logarithmic transformation for distance to the closest obstacle in Distant Fixations in all three experiments; a square root transformation for distance to the spaceship in Close Fixations in Experiments 2 and 3; a reciprocal transformation for distance to the spaceship in Distant Fixations in Experiments 1 and 3; a reciprocal of the logarithmic transformed value for distance to the spaceship in Distant Fixations in Experiment 2. Lastly, due to the binary nature of the completion variable, we chose a Bernoulli link function when predicting the probability to complete a trial.

Maximal allowed random effects structures were explored for each model individually (Barr, Levy, Scheepers, & Tily, 2013). Including as many random effects as the experimental design justifies minimizes the risk of Type I errors by accounting for variance that might otherwise confound the variables of interest. This makes the effects found for variables of interest more robust against retesting. Model selection of the random effects structures was based on the Bayesian information criterion (Chakrabarti & Ghosh, 2011). Fixed effects structures were not subjected to a selection process and kept for hypothesis testing. All models were fitted using the maximum likelihood criterion. Hypothesis testing is based on z-scores applying a $|z| \geq 2.0$ significance criterion. In addition to z-scores, we report lower and upper bounds of the 95% highest density confidence interval. Confidence intervals were obtained using a parametric bootstrap with 10,000 replications based on the final selected model.[3]

2.8. Hypotheses

We hypothesize that in Experiment 1, increasing input noise will lead to fixations being initiated closer to the spaceship (H1.1). The reason for this is that input noise will elicit a growing number of prediction errors in visual sensory states and consequently not meeting the consistency criterion. This will prompt participants to select action goals with shorter time horizons that are also accounted for in gaze allocation. Additionally, in trials with higher input noise, the resulting rise in prediction errors and failure to meet the consistency criterion are expected to lead to control loss in participants, indicated by decreased JoCs (H1.2). Similarly, we expect that the probability to complete a layout (Completion Probability) is negatively impacted by increased input noise (H1.3).

Drift imposes unintended movement on the spaceship. However, given that drift sections are indicated by a red bar, we expect participants to anticipate these situations. Thus, they will allocate their gaze further away in front of the spaceship toward these situations (H1.4). We further hypothesize that there will be no significant difference in the JoC between trials with and without drift, as the anticipation of drift situations and the subsequent behavior, where participants can easily engage in proactive adaptation of action control, will help to prevent any control loss (H1.5). Furthermore, we do not expect any differences in Completion Probability between trials with and without drift (H1.6).

In Experiment 2, we expect that with increasing input noise, participants will initiate fixations at distances progressively closer to the spaceship (H2.1). Likewise, trials with increasing input noise will be associated with lower JoCs (H2.2) and decreased Completion Probability (H2.3).

In Experiment 3, we hypothesize that, compared to the normal block, the induced uncertainty in the invisible block will prompt participants to execute fixations at smaller distances to the spaceship (H3.1). This is based on adapted action selection and participants attending prediction errors occurring when drift suddenly shifts the spaceship. We further expect that trials within the invisible block are associated with lower JoCs, as prediction errors should occur in invisible drift situations, which could lead to the consistency criterion being challenged (H3.2). This will also lead to a decreased Completion Probability in trials within the invisible block compared to trials within the normal block (H3.3). The fake block as a control condition has visual indicators for drift sections, but these are not always reliable. We hypothesize that once this has been learned, there will be no loss of control, as it should be easy to adapt generated predictions which are then based solely on the own actions. Therefore, we expect that fixations during the fake block have distances to the spaceship similar to those in the normal block (H3.4: no significant difference in the distance to spaceship between fake and normal block). Likewise, we expect no differences in JoCs (H3.5) or Completion Probability (H3.6) between trials within the fake and normal block.

3. Results

To get a good balance in terms of JoCs and post-experiment interviews, we aimed for a 50% chance of completing a configuration of layout and experimental manipulations in the piloting phase. Across all experiments, on average, participants completed a configuration in 58.9% (9.4%) of times with a minimum completion rate of 41.9% and a maximum completion rate of 77.3%. Available interview data about the experience of total control loss and the factors to which a control loss was attributed (keeping head static in chin rest, accurate description of input noise), as well as gaming experience, were explored as possible random intercept effects using null models of the various predicted variables. None of the factors explained an adequate amount of variance and we resorted to including participant ID as random intercept effect in every model we built.

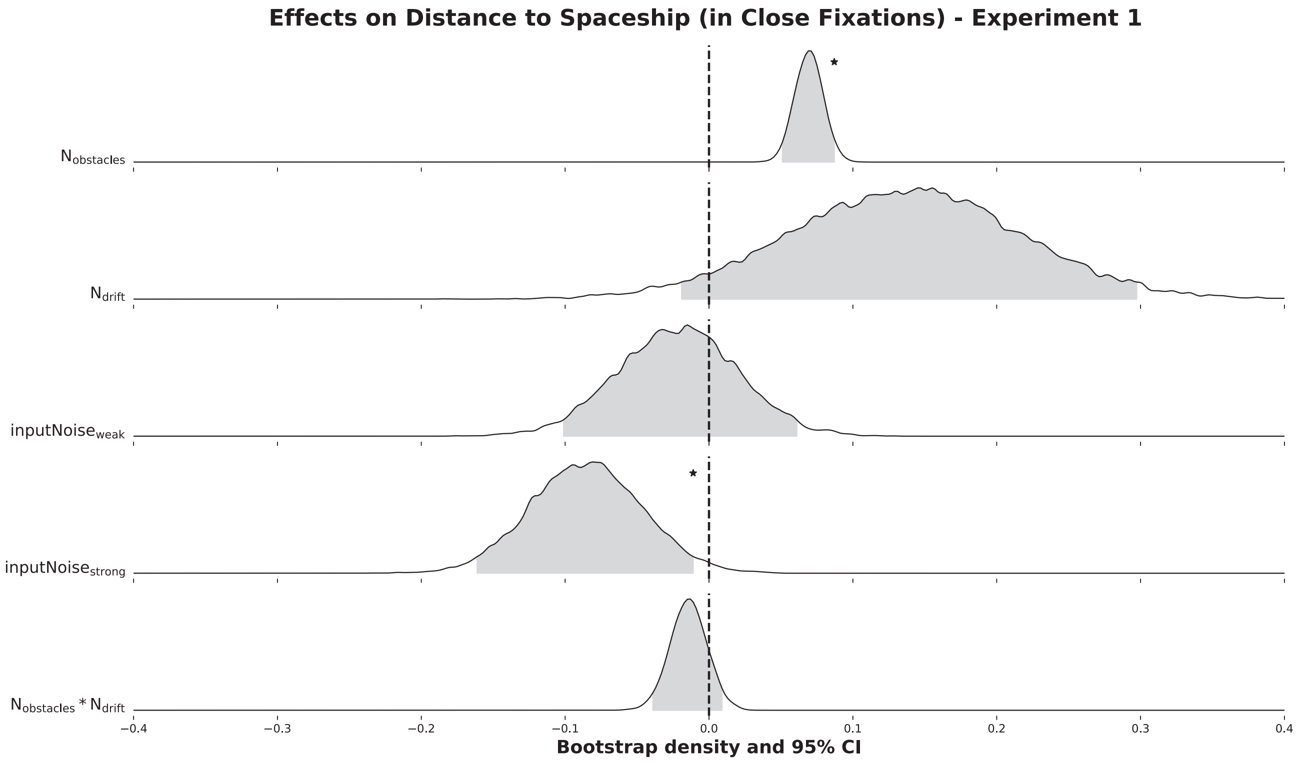

Fig. 3. Distributions of β values for all covariates obtained from a parametric bootstrap with 10,000 iterations of the final selected model. The model predicted the distance between Close Fixations of Experiment 1 and the spaceship in visual degrees (no transformation required for the predicted variable). The gray shaded area in the distributions indicates the 95% - highest density confidence interval. Significant effects are labeled with an asterisk.

3.1. Experiment 1

Due to technical errors, interview data of one of the participants in Experiment 1 was lost. In two out of 26 available interviews, participants hinted toward input noise as disturbance but were not able to accurately describe its effects. Given descriptions included time delay between key press and spaceships' horizontal movement and halting visualization. Two participants were able to accurately describe the effects of input noise ("varying steps," "varying displacement in every game tick"). In the remaining 21 interviews, input noise or factors relating to input noise were not mentioned as causes for control loss. In two interviews, solely the number of obstacles was stated to reduce control. Solely the presence of drift was mentioned 11 times. The combination of both were stated eight times to have caused a loss in control. Out of the four times input noise was mentioned or hinted toward, two times all three factors (many obstacles, drift, and input noise) were mentioned with one time input noise being accurately described and one time it being described as delays. The other two times input noise was mentioned as the sole factor causing control loss, again once being accurately described and the other time it being described as halting visualization.

Fig. 3 shows the distribution of β values obtained by a parametric bootstrap with 10,000 iterations when predicting the distance to the spaceship in Close Fixations. The number of obstacles on screen significantly increased the predicted variable ($\beta = 0.07$, $\sigma = 0.01$, 95%

CI 0.05 to 0.09, z = 7.30). In contrast, in comparison to weak input noise, strong input noise significantly decreased the distance to the spaceship ($\beta = -0.09$, $\sigma = 0.04$, 95% CI −0.16 to −0.01, z = −2.19). Neither the number of drift sections ($\beta = 1.14$, $\sigma = 0.08$, 95% CI −0.03 to 0.28, z = 1.70), nor weak input noise (in comparison to no input noise; $\beta = -0.02$, $\sigma = 0.04$, 95% CI −0.10 to 0.06, z = −0.53) influenced the predicted variable.

When predicting the distance to the spaceship in Distant Fixations, we also found a significant increase in the predicted variable with increasing number of obstacles ($\beta = 1.03*10^{-3}$, $\sigma = 1.51*10^{-4}$, 95% CI $7.52*10^{-4}$ to $1.35*10^{-3}$, z = 6.85). There was no significant main effect for the number of drift sections ($\beta = -3.60*10^{-4}$, $\sigma = 1.06*10^{-3}$, 95% CI $-2.41*10^{-3}$ to $1.74*10^{-3}$, z = −0.34). Similarly, either level of input noise did not affect the distance to the spaceship in Distant Fixations (weak: $\beta = 8.35*10^{-4}$, $\sigma = 1.10*10^{-3}$, 95% CI $-1.36*10^{-3}$ to $2.95*10^{-3}$, z = 0.76; strong: $\beta = 3.72*10^{-4}$, $\sigma = 7.47*10^{-4}$, 95% CI $-1.12*10^{-3}$ to $1.79*10^{-3}$, z = 0.50).

Having successfully completed the trial significantly increased JoCs ($\beta = 1.52$, $\sigma = 0.16$, 95% CI 1.19–1.82, z = 9.34). The following results stem from a model with Completion entered as random intercept effect. Increasing the difficulty of the layout to hard by increasing the total number of obstacles in the layout significantly decreased JoCs ($\beta = -0.44$, $\sigma = 0.10$, 95% CI −0.64 to −0.24, z = −4.30). Likewise, strong input noise significantly decreased JoCs ($\beta = -0.34$, $\sigma = 0.06$, 95% CI −0.46 to −0.21, z = −5.38). We found additional positive effects on JoC for the number of fixations executed throughout the trial ($\beta = 1.16*10^{-3}$, $\sigma = 3.61*10^{-4}$, 95% CI $4.58*10^{-4}$ to $1.87*10^{-3}$, z = 3.22), and the number of times the participant crashed prior to the current trial but immediately after completed a layout (consecutive crash-completion; $\beta = 0.03$, $\sigma = 7.18*10^{-3}$, 95% CI 0.01–0.04, z = 3.64). Also, hard layout difficulty and the presence of drift interacted negatively with each other ($\beta = -0.68$, $\sigma = 0.14$, 95% CI −0.95 to −0.41, z = −4.92). Ultimately, the presence of drift in a given trial did not affect JoCs ($\beta = -0.27$, $\sigma = 0.14$, 95% CI −0.38 to −0.03, z = −1.87).

The probability to complete a layout was negatively affected by increased difficulty of the layout (medium: $\beta = -2.17$, $\sigma = 0.64$, z = −3.39; hard: $\beta = -3.64$, $\sigma = 0.62$, z = −5.90) and strong input noise ($\beta = -0.84$, $\sigma = 0.19$, z = −4.42). Note that the presence of drift did not influence Completion Probability ($\beta = -0.28$, $\sigma = 0.79$, z = −0.35).

3.2. Experiment 2

This experiment featured a total of five levels of input noise. In the subsequent interviews, all six of the participants were able to describe input noise accurately. Four participants mentioned that slight stuttering of the game (most likely due to input noise levels up to 6) was initially mistaken for technical problems with the screen. However, when more severe inaccurate controls (input noise of 9 and 12) began to occur periodically during the experiment, they realized that it was part of the experiment.

Here, when predicting the distance to the spaceship in Close Fixations, we found that the number of obstacles on screen significantly increased the predicted variable ($\beta = 0.01$, $\sigma = 3.07*10^{-3}$, 95% CI $6.73*10^{-3}$ to 0.02, z = 4.25). For the input noise, we have chosen the

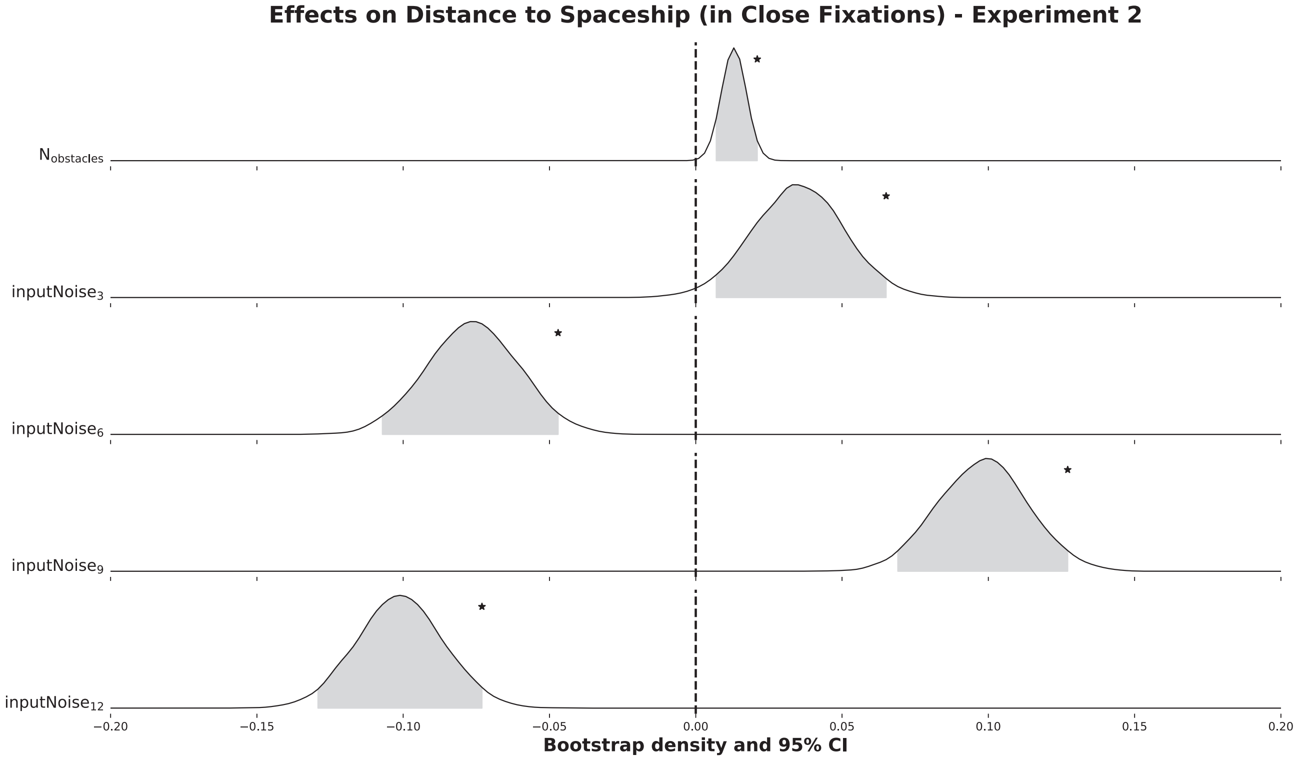

Fig. 4. Distributions of β values for all covariates obtained from a parametric bootstrap with 10,000 iterations of the final selected model. Here, the model predicted the distance between fixation and spaceship in Close Fixations of Experiment 2. The predicted variable was square root-transformed and is given in visual degrees. The gray shaded area in the distributions indicates the 95%-highest density confidence interval. Significant effects are labeled with an asterisk. The subscripts of the input noise covariates refer to the standard deviation in pixels of the normally distributed disturbance when steering the spaceship. Thus, larger subscripts mean increased induced uncertainty in sensorimotor control. The contrasts were chosen so that each level of input noise is compared against the previous one.

contrasts so that each level is compared with the previous level. We found nuanced effects for every individual level; refer to Fig. 4 for a better overview. Compared to the intercept of no input noise, input noise:3 increased the distance to the spaceship ($\beta = 0.04$, $\sigma = 0.02$, 95% CI $3.99*10^{-3}$ to 0.07, z = 2.24). But compared to input noise:3, input noise:6 decreased the distance to the spaceship ($\beta = -0.08$, $\sigma = 0.02$, 95% CI −0.11 to −0.05, z = −4.83). However, the comparison of input noise:9 against input noise:6 yielded an increase again ($\beta = 0.10$, $\sigma = 0.02$, 95% CI 0.07–0.13, z = 6.36). And finally, compared to input noise:9, input noise:12 significantly decreased the distance to the spaceship again ($\beta = -0.10$, $\sigma = 0.02$, 95% CI −0.13 to −0.07, z = −6.90).

Similar to the effect in Close Fixations, the number of obstacles increased the distance to the spaceship in Distant Fixations ($\beta = 1.02*10^{-3}$, $\sigma = 1.52*10^{-3}$, 95% CI $7.26*10^{-4}$ to $1.33*10^{-3}$, z = 6.72). Compared to no input noise, input noise:3 significantly decreased the predicted variable ($\beta = -2.45*10^{-3}$, $\sigma = 1.03*10^{-3}$, 95% CI $-4.48*10^{-3}$ to $-4.70*10^{-4}$, z = −2.38). The distance to the spaceship further decreased in response to input noise:6 ($\beta = -3.75*10^{-3}$, $\sigma = 1.09*10^{-3}$, 95% CI $-5.90*10^{-3}$ to $-1.64*10^{-3}$, z = −3.45).

Compared to input noise:6, input noise:9 then led to an increase in the predicted variable ($\beta = 3.64{*}10^{-3}$, $\sigma = 1.07{*}10^{-3}$, 95% CI $3.64{*}10^{-3}$ to $1.07{*}10^{-3}$, $z = 3.41$). Further increasing input noise to 12 yielded no significant difference to input noise:9 ($\beta = 1.01{*}10^{-3}$, $\sigma = 1.05{*}10^{-3}$, 95% CI $-1.05{*}10^{-3}$ to $1.05{*}10^{-4}$, $z = 0.97$).

Successfully completing a trial in Experiment 2 led to an increase in JoCs ($\beta = 1.74$, $\sigma = 0.30$, 95% CI 1.14–2.31, $z = 5.84$). Of the input noise levels, only input noise:9 led to a significant decrease in JoCs compared to input noise:6 ($\beta = -0.89$, $\sigma = 0.22$, 95% CI -1.33 to -0.45, $z = -3.98$). Although there is a negative trend across all input noise levels, no other comparison yielded significance. Interestingly, we also found no other effects for covariates that we explored in Experiment 1, such as increased difficulty of the layout ($\beta = 0.13$, $\sigma = 0.14$, 95% CI -0.15 to 0.41, $z = 0.93$), the total number of fixations layout ($\beta = 5.69{*}10^{-4}$, $\sigma = 8.59{*}10^{-4}$, 95% CI $-1.21{*}10^{-3}$ to $2.23{*}10^{-3}$, $z = 0.66$), or the number consecutive crash-completions ($\beta = 3.55{*}10^{-3}$, $\sigma = 0.02$, 95% CI -0.05 to 0.05, $z = 0.15$).

Completion Probability in Experiment 2 decreased when layout difficulty increased ($\beta = -1.13$, $\sigma = 0.32$, $z = -3.49$). Imposing weak (3 pixels SD) or (6 pixels SD) moderate input noise had no effect on Completion Probability (input noise:3, $\beta = -0.08$, $\sigma = 0.67$, $z = -0.12$; input noise:6, $\beta = -0.11$, $\sigma = 0.64$, $z = -0.17$). But in comparison to input noise:6, input noise:9 significantly decreased Completion Probability ($\beta = -1.46$, $\sigma = 0.53$, $z = -2.73$), with input noise:12 further decreasing Completion Probability ($\beta = -1.05$, $\sigma = 0.38$, $z = -2.75$).

3.3. Experiment 3

Experiment 3 featured new types of drift sections, namely, invisible and fake drift, which participants encountered in corresponding experimental blocks.

In the semi-structured interviews following the experiment, only four of the 26 participants noticed and described fake drift. And only three of these confirmed that they felt their control over the spacecraft was restricted by fake drift. All but one participant mentioned invisible drift as a factor limiting their control.

When predicting the distance to the spaceship in Close Fixations, the distance increased with every additional obstacle on screen ($\beta = 0.01$, $\sigma = 2.77{*}10^{-3}$, 95% CI 0.01–0.02, $z = 4.48$). No other significant main effects were found. Neither the number of drift sections on screen ($\beta = 1.11{*}10^{-3}$, $\sigma = 7.43{*}10^{-3}$, 95% CI -0.01 to 0.02, $z = 0.15$), nor the comparison of any of the two new drift types versus normal drift yielded significance (invisible block: $\beta = -0.03$, $\sigma = 0.02$, 95% CI -0.07 to 0.02, $z = -1.17$; fake block: $\beta = -0.02$, $\sigma = 0.03$, 95% CI -0.07 to 0.03, $z = -0.82$).

The model predicting the distance to the spaceship in Distant Fixations revealed a similar effect for the number of obstacles on screen as in Close Fixations. Here again, an increasing number of obstacles is associated with an increasing distance to the spaceship ($\beta = 7.99{*}10^{-4}$, $\sigma = 1.03{*}10^{-4}$, 95% CI $6.02{*}10^{-4}$ to $1.01{*}10^{-3}$, $z = 7.75$). However, in contrast to Close Fixations, the number of drift sections on screen significantly increased the distance to the spaceship in Distant Fixations ($\beta = 4.45{*}10^{-3}$, $\sigma = 7.65{*}10^{-4}$, 95% CI $2.94{*}10^{-4}$ to $5.93{*}10^{-3}$, $z = 5.82$). Yet, neither type of drift showed a significant effect on the predicted variable (invisible block: $\beta = 3.40*10^{-3}$, $\sigma = 2.58*10^{-3}$, 95% CI $-2.77*10^{-3}$ to $1.87*10^{-3}$, z = 1.32; fake block: $\beta = 1.22*10^{-3}$, $\sigma = 4.06*10^{-3}$, 95% CI $-7.04*10^{-3}$ to $8.95*10^{-3}$, z = 0.30).

Successful completion of trials also increased JoCs reported in Experiment 3 ($\beta = 1.47$, $\sigma = 0.02$, 95% CI 1.44–1.51, z = 81.65). Entering successful completion as random intercept effect into the model, we found that raising layout difficulty from medium to hard decreased JoCs ($\beta = -0.19$, $\sigma = 0.02$, 95% CI −0.22 to −0.16, z = −11.77). Further, JoCs in the invisible block were significantly lower than those in the normal block ($\beta = -0.69$, $\sigma = 0.02$, 95% CI −0.73 to −0.65, z = −34.54). There was no difference between JoCs in the fake and the normal block ($\beta = -0.03$, $\sigma = 0.02$, 95% CI −0.08 to $6.90*10^{-3}$, z = −1.66). Also here, in Experiment 3, we found positive effects for the number of fixations executed throughout a trial ($\beta = 6.57*10^{-4}$, $\sigma = 7.61*10^{-5}$, 95% CI $5.10*10^{-4}$ to $8.12*10^{-4}$, z = 8.63) and the number of consecutive crash-completions ($\beta = 0.03$, $\sigma = 1.97*10^{-3}$, 95% CI 0.03–0.04, z = 17.11). For a better overview and comparison of the effects on JoC in all three experiments, refer to Fig. 5.

Surprisingly, Completion Probability was not affected by increased layout difficulty ($\beta = 0.14$, $\sigma = 0.14$, z = 0.98). However, compared to the normal block, in the invisible block, Completion Probability was significantly decreased ($\beta = -0.91$, $\sigma = 0.17$, z = −5.46). Lastly, Completion Probability was not affected by the fake block compared to the normal block ($\beta = 0.12$, $\sigma = 0.18$, z = 0.63).

3.4. Additional analyses

In further exploratory analyses, we tested what affects the distance to the closest obstacle in Distant Fixations across experiments. If we assume that the locations of Distant Fixations reflect action goals that are pursued (Heinrich et al., 2024), it can be informative to explore how action selection is adapted in order to be more efficient under control loss—for instance, by maintaining greater distances from obstacles when there is input noise or drift.

In Experiment 1, the predicted distance increased the more drift sections were on screen ($\beta = 0.04$, $\sigma = 0.02$, 95% CI 0.01–0.07, z = 2.59). Neither weak input noise ($\beta = -3.69*10^{-3}$, $\sigma = 5.42*10^{-3}$, 95% CI −0.02 to $6.25*10^{-3}$, z = −0.68) nor strong input noise ($\beta = 7.63*10^{-3}$, $\sigma = 5.60*10^{-3}$, 95% CI $-2.83*10^{-3}$ to 0.02, z = 1.36) had an effect on the distance to the closest obstacle.

In Experiment 2, input noise:3 led to an increase in the distance to the closest obstacle in Distant Fixations ($\beta = 0.04$, $\sigma = 0.01$, 95% CI 0.01–0.07, z = 2.88). Further, increasing input noise yielded no significant difference in comparison.

Contrary to the findings in Experiment 1, in Experiment 3, we found that the number of drift sections on screen decreased the distance to the closest obstacle ($\beta = -0.11$, $\sigma = 0.03$, 95% CI −0.17 to −0.06, z = −4.31). In turn, compared to the normal block, fixations in the invisible and fake blocks featured longer distances to the closest obstacle (invisible block: $\beta = 0.03$, $\sigma = 0.01$, 95% CI $6.51*10^{-3}$ to 0.05, z = 2.50; fake block: $\beta = 0.04$, $\sigma = 0.01$, 95% CI $8.31*10^{-3}$ to 0.06, z = 2.68).

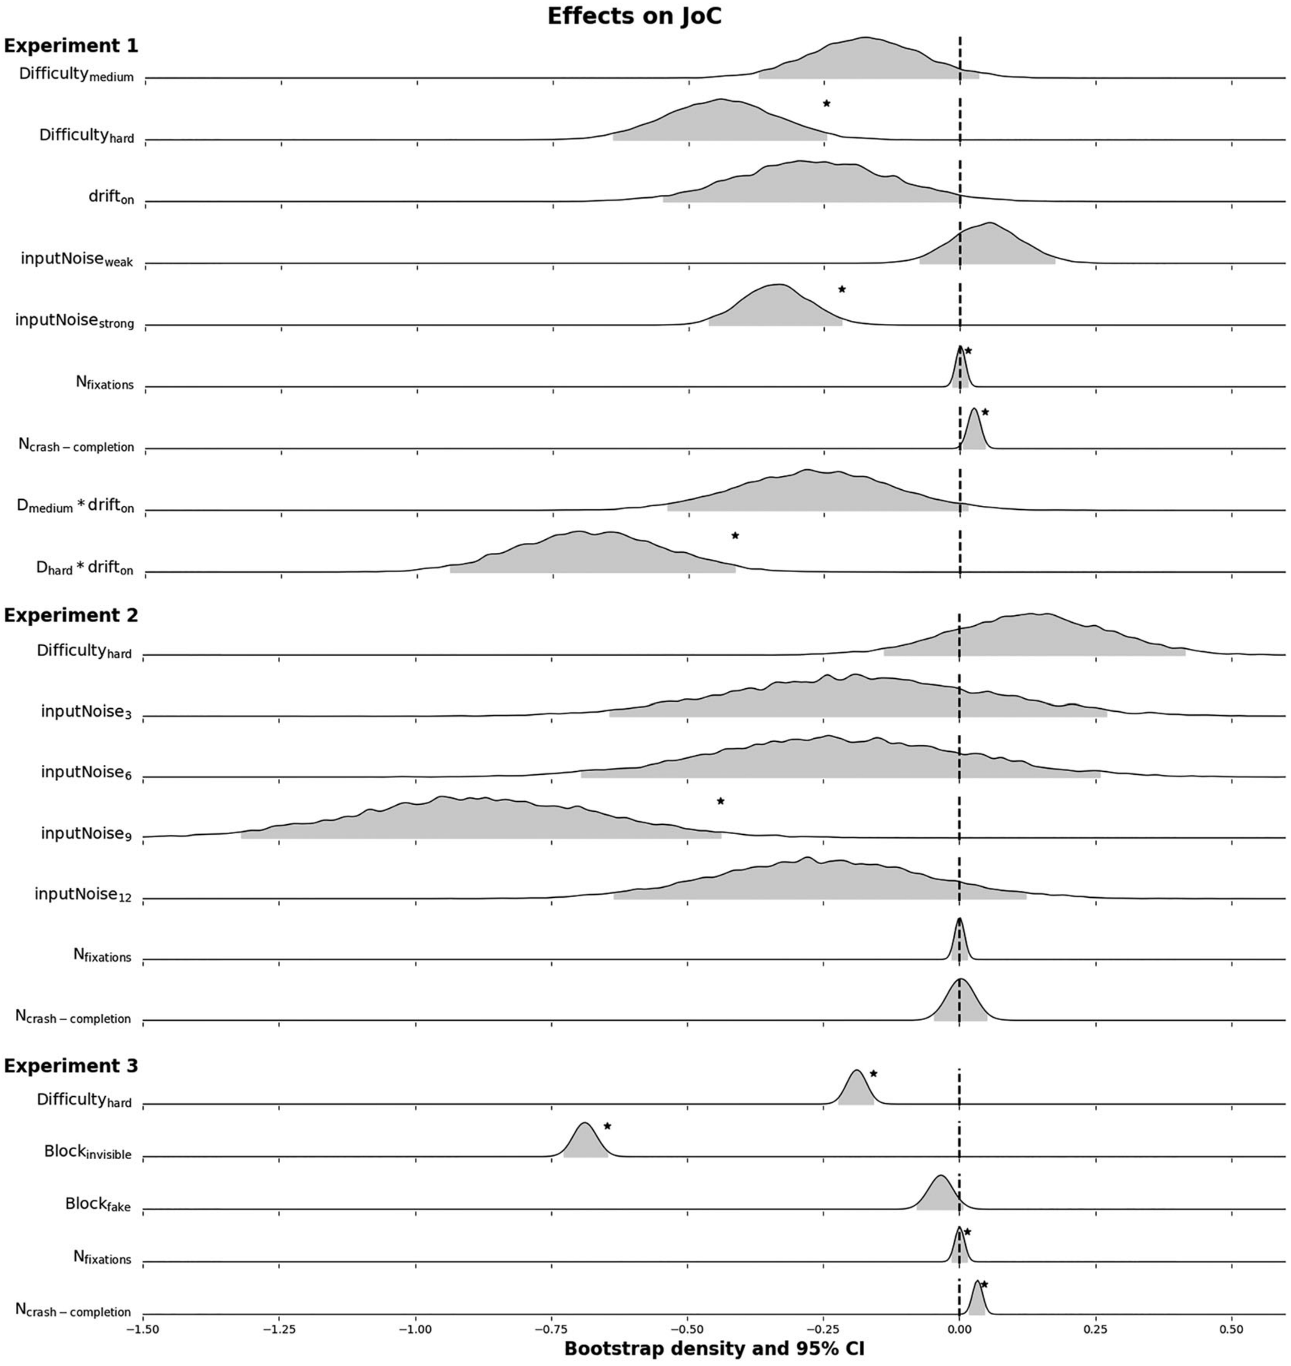

Fig. 5. Distributions of β values for all covariates obtained from a parametric bootstrap with 10,000 iterations. The parametric bootstraps were based on the final selected model in the respective experiment. Each of the models predicted the control judgment responses (JoC; possible responses ranged from 1 to 7) given by participants after every trial. No transformations were required, which simplifies the comparability of the effects across experiments. The gray shaded area in the distributions indicates the 95%-highest density confidence interval. Significant effects are labeled with an asterisk. Note that $N_{\text{fixations}}$ positively influenced JoCs in Experiments 1 and 2, but the kernel bandwidth used for the kernel density estimates in generating the distributions was likely too wide to accurately capture its small effect.

We conducted an additional analysis to examine how the proportion of Distant Fixations (and conversely, Close Fixations) within the total number of fixations changes as a function of the experimental manipulations.

In Experiment 1, we found no significant main effects. However, there was a significant interaction between high level difficulty and drift ($\beta = -0.08$, $\sigma = 0.03$, 95% CI -0.14 to -0.02, $z = -2.51$). When level difficulty increased from medium to high and drift was enabled, the proportion of Distant Fixations decreased, meaning that the proportion of Close Fixations increased.

This interaction could not be replicated in Experiment 3, and we found no effects for other types of drift. Likewise, in line with the results of Experiment 1, we observed no effects of input noise on the proportion of Distant Fixations in Experiment 2.

4. Discussion

In this series of studies, we investigated how motor perturbations of varying predictability—input noise and drift—impact performance, control judgments (JoC), and eye-movement behavior in a dynamic, interactive environment. We differentiated between reactive and proactive strategies of action control by examining changes in participants' behavior while steering a spaceship. But it was only when we distinguished between different types of fixations that the results contributed to our understanding of the dynamics of situated action control and their link to shifts in the SoC.

Our analysis of fixation distances to the spaceship in Experiment 1 revealed mixed support for H1.1. As assumed, Close Fixations showed that higher levels of input noise reduced the distance to the spaceship, reflecting cautious behavior in response to control loss. However, in Distant Fixations, input noise did not affect fixation distances.

In Experiment 2, input noise also did not lead to a monotonic decrease in fixations distance to the spaceship (H2.1). In Close Fixations instead, we observed an effect, with each level of input noise oscillating in the opposite direction from the level before. We cannot exclude the possibility that these results are partly due to subjects being aware of input noise as an experimental condition. Our initial understanding of these results is that participants were actively trying to visually ignore the stuttering of obstacles in the environment and focus on the spaceship. However, the interpretation of these effects remains mostly unclear. In Distant Fixations, we found a successive reduction of the distance through the first two levels of input noise. Then, however, we found a slight increase in distance for input noise:9. Note that this increase did not go beyond the distance of the intercept of no input noise. These results better fit a step-function than the assumed monotonic decreasing one: once uncertainty reached a critical threshold, participants adapted by shifting gaze closer to the spaceship, a reactive adaptation of action control, but further increases in noise did not exacerbate this effect. This threshold-like response highlights the limits of sensorimotor compensation under unpredictable conditions.

The increased distance to obstacles observed under input noise in Distant Fixations of Experiment 2 further supports the idea of reactive adaptation of action control. By prioritizing immediate safety and feasibility, participants demonstrated a shift in action goals under uncertainty, pursuing those more likely to succeed. This suggests that participants engaged in reactive adaptation of action control to cope with unpredictable disturbances, prioritizing immediate control over projecting their environment into the more distant future. The action selection process is adapted in terms of its top-down information processing component. Initially, agents were conservative in their top-down control, relying primarily on sensorimotor processing with action selection largely driven bottom-up. However, during control loss, top-down control—and consequently cognitive processing—increases. Planning becomes more incremental to maintain high performance.

The conclusion that reactive adaptation of action control is based on an awareness boundary, as Kahl et al. (2022) assume, is further supported by the fact that JoC did not decrease linearly with increasing input noise but instead followed a threshold function. Once participants experienced control loss, additional input noise did not further reduce their SoC (contrary to H1.2 and H2.2). These findings suggest that control judgments are influenced more by cognitive processes, such as explaining the environment, than by sensorimotor compensations for disturbances.

We could not confirm that gaze is allocated further away in front of the spaceship in drift situations in Experiment 1 (H1.4). Yet, particularly changes in the statistics of Distant Fixations in Experiment 3 indicated that participants learned and adapted their behavior in response to predictable disturbances, even in the absence of visible feedback. Here, the number of drift sections on the screen increased the distance to the spaceship in Distant Fixations, supporting H1.4, which we stated for Experiment 1. But we do have to reject H3.1, as the invisible block had no effect on distances between fixation and spaceship. We would have expected participants to allocate their gaze closer to the spaceship, as invisible drift can occur suddenly and unexpectedly and that would create immediate demands for maintaining the spaceships trajectory similar to high input noise conditions. Further exploratory analyses of Distant Fixations revealed that number of drift sections decreased distance to the closest obstacle in Experiment 3. Only within the invisible drift block, fixations featured increased distances to closest obstacles. This suggests that the presence of a mental model of environmental features like drift enables proactive strategies, allowing participants to anticipate and prepare for sudden perturbations. The results for drift in general provided support for proactive adaptation of action control. Note that fake drift had no effect on the distance to the spaceship (H3.4) or to the closest obstacle, suggesting that proactive adaptation of action control is used to mitigate conditions that are detrimental and arise when the agent does not have an accurate model of the environment.

As hypothesized, normal drift did not affect JoC (H1.5). This lack of impact suggests that proactive adaptation of action control prevented participants from experiencing a loss of control. Fake drift probably had no effect on JoC (H3.5), as this type of drift did not result in a loss of control over the spaceship. Only when drift became unpredictable and actually affected the spaceship, as in the invisible drift condition, did participants report control loss, aligning with H3.2. Taken together, this indicates that control loss arises not from the drift disturbance itself but from the inability to link the perceived disturbance to the mental model of the environment.

We explored other variables that could be integrated when judging the own SoC. We found that in Experiments 1 and 3, participants reported higher JoCs for every additional fixation they executed or for every time they successfully completed a trial after crashing (Fig. 5). However, this effect was observed only in experiments, which featured predictable drift and is missing in Experiment 2, which only featured unpredictable input noise. This may be due to the fact that the SoC can be improved by adapting one's behavior (oculomotor or steering) in response to predictable changes in the environment. This again supports that proactive adaptation of action control is based on anticipations and projections of the environment.

The inability to explain perceived disturbances based on a mental model of the environment was linked to poorer performance. For instance, random noise in control over the spaceship and invisible drift led to reduced Completion Probabilities, as predicted by H1.3, H2.3, and H3.3. This highlights the importance of predictability in supporting effective action control and maintaining performance. However, this effect is most likely because the situations that were caused by inaccurate models of the environment were detrimental, supported by the fact that fake drift did not affect performance (H3.6).

An additional analysis provided further insight into the relative use of Close versus Distant Fixations. Both types of fixational eye movements are crucial for action control, which may explain why we did not observe consistent differences in the proportion of Distant Fixations across experimental manipulations. This indicates that participants flexibly balance both fixation types depending on task demands, maintaining a robust allocation strategy even under varying levels of noise and drift. The significant interaction between hard level difficulty and drift in Experiment 1, however, shows that this balance can shift under particularly demanding conditions, most likely within drift sections where immediate control needs rise sharply, especially when many obstacles are present. The absence of this effect in Experiment 3 points to context-dependence and highlights boundary conditions (obstacle density) that may be worth investigating to determine under which circumstances changes in fixation proportions emerge.

This work highlights visual focus as an important measure for shifts between reactive and proactive strategies in situated action control, when people respond to environmental features of varying predictability. However, the role of fixation distances to a reference point in action goal pursuit under (un)certainty needs further study. Gaze metrics should be assessed not just as static measures. Examining how gaze transitions between regions could offer deeper insights into the link between the SoC and cognitive processes. Close Fixations seem to represent immediate control needs, while Distant Fixations may signal readiness for visual exploration and proactive adaptation of action control. This distinction aligns with the observed effects of experimental manipulations on these two types of fixations. Identifying especially the transitions between Close and Distant Fixations might be a key indicator to better understand action control holistically.

Our study has limitations in separating sensorimotor control from cognitive control processes. However, we believe that accounting for both levels and their interplay is essential to fully understand real-world behavior. To achieve this, it is important to allow flexibility not only in how movements are controlled but also in how action goals are selected. Future research could explore the connection between these two levels by combining eye-tracking with precise joystick tracking. For tasks involving reaching, devices that track hand or mouse movements can also provide valuable insights.

Additionally, future research should incorporate computational models to rigorously formalize how fixation distances relate to action control. We have already taken the first step in the direction of simulation-capable models by using a type of data analysis that provides mathematical models. Such models are continuously adapted and can simulate various fixation locations and strategies over time. These simulations display how visuomotor behavior evolves, adapting to predictable and unpredictable features of the environment but being dependent on prior behavior. We believe that approaching action control from the perspective of dynamical systems is crucial for a better understanding.

In this last part, we would like to take the discussion beyond the scope of the experiments described here. We have not yet addressed the concept of multisensory perception. Prediction errors can occur in a variety of modalities, such as an action not producing the expected sound. Where it is easy to infer how prediction errors in visual perception cause adjustments in action goals (here measured by eye movements), it is much harder to discuss how auditory prediction errors (or even haptic or olfactory ones) influence action plans. According to multilayered models of action control such as that of Kahl et al. (2022), prediction errors detected in the sensorimotor system are integrated into a holistic representation of the action at the cognitive level (which includes the SoC). However, it is quite possible that the different modalities are not represented in equal proportions. Additionally, prediction errors in the different sensors may lead to different adaptations in action control and not to a certain adaptation to different degrees. The concept of the SoC refers to the overall control experience. This means that all prediction errors, regardless of the sensory origin, are integrated into a low-level SoC (whether individual or same), which then translates into a single higher-level SoC. This raises the question of what effects a haptic prediction error can be expected to have on the SoC compared to an auditory prediction error, for example. Because haptic feedback is more integral to precise motor actions (e.g., holding, pressing, balancing), we think that in comparison to auditory prediction errors, haptic ones will lead to a greater loss in control. There are interesting interactions to be explored as well. Like what happens to the SoC and action control when prediction errors occur in all modalities except visual perception, the modality we rely on the most? Also, how does the SoC play into the rich representation of the action and how specifically does it elicit adaptations in action control? Regarding time perception, if expected changes in the environment are brought about, but at the wrong time, this will certainly lead to adaptations on the temporal dimension of subsequent actions and anticipated action feedback. Finally, it may even be that the sensory origin of prediction errors primarily drives whether reactive or proactive adaptation of action control is applied.

Under the assumption that the current action goal influences where visual attention is allocated, we used eye-tracking to implicitly measure action goals. As mentioned above (Section 1.3), eye-movement behavior—the allocation of visual attention—is driven by multiple factors, which can broadly be categorized as either top-down or bottom-up processing. Action goals fall under top-down (goal-directed) processing, while strictly visual features such as edges, colors, and movement are assigned to bottom-up (stimulus-based) processing. Behavior at any moment results from a combination of both types, but they do not contribute equally. We propose that the SoC, reflecting how effectively one's past action goals have been implemented, determines how top-down and bottom-up processing are weighted in shaping behavior. When SoC is high, top-down processing predominates because the pursued goal is well-defined, supported by an accurate cognitive representation of the environment. In contrast, when SoC is low, bottom-up processing has a stronger influence, meaning behavior is guided mainly by low-level perceptual cues. The action goal is then only vaguely represented, lacking the precision needed to effectively drive behavior. Primarily, bottom-up-driven attention may support exploration of the surroundings while still solving the task at hand, helping to gather new evidence and refine the mental model of the environment.

Lastly, a potential long-term goal might be to integrate SoC into autonomous systems. In this case, action selection would respond to real-time changes in environmental conditions but also factor in previous action selection and the feasibility during execution. This approach can improve the fluency and efficiency of action control, and increase the system's overall capability and performance.

Overallvv, the findings underscore the adaptability of situated action control, illustrating how humans adjust their strategies to navigate complex and continuous environments. What we found contributes to our understanding of the general interplay between vision, motor control, and cognitive processes. Beyond these insights, we welcome discussions on the broader implications of our research, particularly how it relates to other domains of cognition and behavior.

Acknowledgments

Open access funding enabled and organized by Projekt DEAL.

Funding

This research was funded by the German Research Foundation (DFG) Priority Program 2134 "The Active Self."

Notes

1 Git repository of the Dodge Asteroids experimental environment.
2 This research was reviewed and approved by the Ethics Committee of the Department of Psychology and Ergonomics (IPA) of the Technische Universität Berlin.
3 Data and analysis code, including model selection, is available in the following Git repositories: Experiment 1, Experiment 2, and Experiment 3.

References

Abalakin, N., Heinrich, N. W., Österdiekhoff, A., Kopp, S., & Russwinkel, N. (2024). Allocation of fixational eye movements in response to uncertainty in dynamic environments. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume *46*.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Bates, D. (2015). Mixedmodels: A julia package for fitting (statistical) mixed-effects models. https://github.com/dmbates/MixedModels.jl

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, *59*(1), 65–98. https://doi.org/10.1137/141000671

Blakemore, S.-J., Wolpert, D. M., & Frith, C. D. (2002). Abnormalities in the awareness of action. *Trends in Cognitive Sciences*, *6*(6), 237–242. https://doi.org/10.1016/S1364-6613(02)01907-1

Box, G. E. P., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *26*(2), 211–243. https://doi.org/10.1111/j.2517-6161.1964.tb00553.x

Chakrabarti, A., & Ghosh, J. K. (2011). AIC, BIC and recent advances in model selection. In P. S. Bandyopadhyay & M. R. Forster (Eds.), *Philosophy of statistics, volume 7 of Handbook of the philosophy of science* (pp. 583–605). Amsterdam: North-Holland. https://doi.org/10.1016/B978-0-444-51862-0.50018-6

Chambon, V., & Haggard, P. (2012). Sense of control depends on fluency of action selection, not motor performance. *Cognition*, *125*(3), 441–451. https://doi.org/10.1016/j.cognition.2012.07.011

Cina, M., & Rad, A. B. (2024). We steer where we stare. *Transportation Research Interdisciplinary Perspectives*, *25*, 101092. https://doi.org/10.1016/j.trip.2024.101092

Cooper, R. P. (2010). Forward and inverse models in motor control and cognitive control. In *Proceedings of the AISB 2010 Symposium on AI-Inspired Biology* (pp. 108–110).

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*(9), 1035–1045. https://doi.org/10.1016/S0042-6989(03)00084-1

Engbert, R., & Mergenthaler, K. (2006). Microsaccades are triggered by low retinal image slip. *Proceedings of the National Academy of Sciences*, *103*(18), 7192–7197. https://doi.org/10.1073/pnas.0509557103

Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences*, *9*(6), 290–295. https://doi.org/10.1016/j.tics.2005.04.012

Hayhoe, M. M. (2017). Vision and action. *Annual Review of Vision Science*, *3*(1), 389–413. https://doi.org/10.1146/annurev-vision-102016-061437

Heinrich, N. W., Österdiekhoff, A., Kopp, S., & Russwinkel, N. (2024). Goal-directed allocation of gaze reflects situated action control in dynamic tasks. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume *46*.

Heinrich, N. W., Russwinkel, N., Österdiekhoff, A., & Kopp, S. (2023). A straightforward implementation of sensorimotor abstraction in a two-layer architecture for dynamic decision-making. In *MathPsych/ICCM/EMPG 2023*.

Hennig, C., Viroli, C., & Anderlucci, L. (2019). Quantile-based clustering. *Electronic Journal of Statistics*, *13*(2), 4849–4883. https://doi.org/10.1214/19-EJS1640

Hennig, C., Viroli, C., & Anderlucci, L. (2022). *QuClu: Quantile-Based Clustering Algorithms*. R package version 1.0.1. https://cran.r-project.org/package=QuClu

Jeannerod, M. (2003). The mechanism of self-recognition in humans. *Behavioural Brain Research*, *142*(1–2), 1–15. https://doi.org/10.1016/S0166-4328(02)00384-4

Jordan, M. I. (1996). Computational aspects of motor control and motor learning. In H. Heuer & S. W. Keele (Eds.), *Handbook of perception and action*, volume *2* (pp. 71–120). Elsevier.

Kahl, S., Wiese, S., Russwinkel, N., & Kopp, S. (2022). Towards autonomous artificial agents with an active self: Modeling sense of control in situated action. *Cognitive Systems Research*, *72*, 50–62. https://doi.org/10.1016/j.cogsys.2021.11.005

Kandil, F. I., Rotter, A., & Lappe, M. (2009). Driving is smoother and more stable when using the tangent point. *Journal of Vision*, *9*(1), 11. https://doi.org/10.1167/9.1.11

Land, M., & Horwood, J. (1995). Which parts of the road guide steering? *Nature*, *377*(6547), 339–340. https://doi.org/10.1038/377339a0

Land, M., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, *28*(11), 1311–1328. https://doi.org/10.1068/p2935

Land, M. F., & Lee, D. N. (1994). Where we look when we steer. *Nature*, *369*(6483), 742–744. https://doi.org/10.1038/369742a0

Lappi, O. (2022). Gaze strategies in driving–An ecological approach. *Frontiers in Psychology*, *13*, 821440. https://doi.org/10.3389/fpsyg.2022.821440

Meiran, N., Cole, M. W., & Braver, T. S. (2012). When planning results in loss of control: Intention-based reflexivity and working-memory. *Frontiers in Human Neuroscience*, *6*, 104. https://doi.org/10.3389/fnhum.2012.00104

Nolan, D. (2022). Video survey. Accessed: 2022-12-08. https://www.stat.berkeley.edu/~nolan/surveys/videoSurvey.pdf

Österdiekhoff, A. (2023). Incorporating sense of control in dynamic multi-tasking problems: A model-based hierarchical reinforcement learning approach. In *Doctoral Consortium at the 15th International Conference on Agents and Artificial Intelligence (ICAART).*

Österdiekhoff, A., Heinrich, N. W., Russwinkel, N., & Kopp, S. (2024). Sense of control in dynamic multitasking and its impact on voluntary task-switching behavior. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume *46*.

Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, *107*(1), 179–217. https://doi.org/10.1016/j.cognition.2007.09.003

R Core Team. (2022),. *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.r-project.org/

Shinners, P., & Pygame Community. (2000). Pygame. Accessed: 2022-08-29. https://www.pygame.org/

Synofzik, M., Vosgerau, G., & Newen, A. (2008). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, *17*(1), 219–239. https://doi.org/10.1016/j.concog.2007.03.010

Tsakiris, M., Prabhu, G., & Haggard, P. (2006). Having a body versus moving your body: How agency structures body-ownership. *Consciousness and Cognition*, *15*(2), 423–432. https://doi.org/10.1016/j.concog.2005.09.004

Van Rossum, G., & Drake, F. L. (2009). Python 3 reference manual.

Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with S* (4th ed.). New York: Springer. https://www.stats.ox.ac.uk/pub/MASS4/

VPixx Technologies Inc. (2024),. VPixx Software Tool Update Guide 3R5. Accessed: 2025-01-20. https://vpixx.com/developer/VPixx_Software_Tool_Update_Guide_3R5.pdf

Wegner, D. M. (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences*, *7*(2), 65–69. https://doi.org/10.1016/S1364-6613(03)00002-0

Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation: Sources of the experience of will. *American Psychologist*, *54*(7), 480–492. https://doi.org/10.1037/0003-066X.54.7.480

Wilkie, R. M., & Wann, J. P. (2002). Driving as night falls: The contribution of retinal flow and visual direction to the control of steering. *Current Biology*, *12*(23), 2014–2017. https://doi.org/10.1016/S0960-9822(02)01259-4

2.3 Conclusion

These studies demonstrate that both gaze behavior and SoC depend on the match between internal predictive models and observed environmental dynamics. Adaptive control is grounded in the continuous coordination of prediction, action, and feedback at the sensorimotor interface. When this coordination fits the actual dynamics of the environment, behavior remains effective within the agent's action possibilities, and the experience of control is preserved.

Finally, the findings suggest that SoC operates as a form of self-belief updating, where expectations about efficacy are continuously updated. This raises the question of the computational principles that govern such updating.

3 A Bayesian Model of Sense of Control Updating

The previous chapter established that SoC reflects a form of self-belief updating, in which expectations about efficacy are continuously updated based on sensorimotor and performance feedback. These findings build on Pacherie's (2006; 2008) dynamic hierarchical model, in which control is established through the integration of predictions and feedback across multiple levels of action intention. Pacherie (2008) draws on the work of Moore and Haggard (2007) on intentional binding to propose that such experiences depend on Bayesian inference, weighting prior knowledge and sensory data according to their reliability. However, while Moore and Haggard speculated about Bayesian mechanisms for binding actions to their effects, and Pacherie endorsed this approach for understanding the sense of intentional causation, neither developed a formal computational model of how control beliefs are updated over time based on performance feedback.

3.1 Integrating Sensorimotor and Evaluative Feedback

Previous work by our group (Giersiepen et al., 2025) empirically investigated how sensorimotor cues and evaluative feedback combine to produce SoC in a complex, goal-directed task. Fifty participants had to solve a modified version of the Dodge Asteroids task. Here a spaceship had to be steered following a path as closely as possible. Performance was measured as the summed total distance between the spaceship and the path throughout the trial, and task difficulty was manipulated through input noise. After each trial, participants received feedback. Unbeknownst to them, the valence of this feedback was systematically varied. The results supported that SoC drew on both low-level sensorimotor information and high-level evaluative feedback. Higher levels of input noise reduced SoC. Positive feedback increased SoC, and negative feedback lowered it, with the effect of negative feedback being noticeably stronger. Individual differences in trait control beliefs, specifically locus of control and depressive symptoms, moderated these effects. These findings demonstrate that in dynamic environments where action-effect contingencies unfold over time, SoC reflects an integration of information drawn from multiple levels of Pacherie's hierarchical framework.

However, the computational principles governing how new information from either level (sensorimotor cues or evaluative feedback) updates control beliefs over time remain unspecified. The following section introduces a Bayesian model to formalize these updating dynamics.

3.2 Study 3: Towards a Bayesian Cognitive Model of Self-Belief Updating

I treat SoC as a latent belief that changes through Bayesian inference, combining prior expectations about personal efficacy with performance evidence. Bayesian modeling provides a mathematically grounded way to express this updating process and to explain individual differences in sensitivity to performance evidence, which is crucial for understanding the effects observed in Giersiepen et al. (2025).

To test this framework empirically, I developed a modified version of the Dodge Asteroids task that required participants to learn environmental dynamics across trials while continuously reporting their control beliefs. In this version of the Dodge Asteroids environment, participants completed trials that contained four drift sections, and obstacles appeared only within these sections. Steering was disabled during drift, and each section imposed a consistent lateral movement that was indicated visually. There was no input noise. Participants needed to learn the environmental drift dynamics and plan their initial steering so that the spaceship passed through the obstacles without crashing. Trial performance was defined by whether a crash occurred and by how many drift sections were completed, ranging from zero to four, with four indicating that the trial was completed without a crash[1]. At the beginning of each trial, participants rated their expected performance, and after the trial they reported their SoC. After completing the task, participants filled out a Big Five personality questionnaire.

I developed a Bayesian model that describes how prior expectations and performance evidence shaped the reported SoC across trials. The model treated the participant's SoC from the previous trial as the prior belief for the current trial. Only for the first trial, where no such information was available, the expectancy rating was used as the prior. Trial performance served as the evidence that updated this belief. The influence of performance evidence was determined by two participant-specific regression coefficients

---

[1] Although these measures may seem redundant, a model selection for linear mixed models showed that crash occurrence and the number of completed drift sections each explained unique variance in SoC ratings, which justified keeping both in the final model.

estimated from the data. These coefficients were applied to the performance indicators on each trial, producing a likelihood weight that varied from trial to trial depending on the available evidence. The updated belief was computed as a combination of the prior SoC and this weighted trial performance likelihood, yielding the model's predicted SoC for the trial. This posterior belief then served as the prior for the next trial.

The model reproduced overall trends in the data but also revealed substantial individual differences in how performance evidence influenced belief updating. Some participants relied moderately on performance outcomes, while others relied strongly on them. Contrary to our hypotheses, these differences were not explained by overall performance levels or Big Five personality traits. Although the model captured stable patterns of belief updating, it did not accurately predict rapid within-participant changes in evidence weighting. Participants often appeared to place greater weight on performance evidence after successfully completing multiple drift sections following a crash. The Bayesian model therefore requires revision with an additional term in the likelihood weighting that increases the influence of performance evidence according to the number of completed drift sections when a crash occurred in the previous trial.

Another limitation concerns how the model treats evidence that contradicts prior expectations. Research suggests that unexpected or surprising outcomes trigger adaptive responses. These responses often include increases in learning rate that support faster belief updating when predictions are violated (Gershman, 2015; Körding and Wolpert, 2004; Tenenbaum, Griffiths, and Kemp, 2006). The present model does not distinguish between expected and unexpected evidence (prediction errors). The influence of performance evidence depends only on its magnitude and the participant-specific coefficients, not on the discrepancy between expectations and outcomes. Incorporating sensitivity to prediction error magnitude may improve the model's ability to capture dynamics in belief updating, particularly when participants encounter outcomes that substantially deviate from their expectations.

3.2.1 Contribution to Study 3

I conceptualized the study idea and design, programmed the experimental task, and fully carried out data collection (participants were recruited from a local participant pool for students at the Technische Universität Berlin). I conducted the data analysis and created all visualizations, and I made major contributions to the interpretation of the results. I wrote the manuscript. The paper was presented at the MathPsych / ICCM 2025 by the second first author Rebecca von Engelhardt.

Study 3

Heinrich, N. W., von Engelhardt, R., Österdiekhoff, A., Kopp, S., & Russwinkel, N. (2025).

Towards a Bayesian Cognitive Model of Self-Belief Updating.

Proceedings of the International Conference on Cognitive Modeling (ICCM 2025).

Towards a Bayesian Cognitive Model of Self-Belief Updating

Nils Wendel Heinrich[1,2,*], **Rebecca von Engelhardt**[1,*], **Annika Österdiekhoff**[3], **Stefan Kopp**[3], **Nele Russwinkel**[1]
nils.heinrich@uni-luebeck.de
[1]Institute of Information Systems (IFIS), Department of Computer Science and Engineering, Universität zu Lübeck, Lübeck, Germany
[2]Cognitive Psychology and Ergonomics, Technische Universität Berlin, Berlin, Germany
[3]Social Cognitive Systems, Faculty of Technology, CITEC, Bielefeld University, Bielefeld, Germany
[*]These authors contributed equally to this work

Abstract

Bayesian inference is a powerful mathematical framework for modeling cognitive processes, and it has been widely applied in computational models of sensorimotor integration. However, higher-level cognitive functions, such as updating beliefs about one's own sense of control, may also rely on processes that resemble Bayesian integration. We conducted an experiment in which participants navigated a spaceship through sections characterized by temporary control loss. They had to use their knowledge about the environment to anticipate the ship's trajectory several steps ahead. Before each trial, participants rated how well they expected to perform; afterward, they rated how much control they felt they had. Through model selection, we found that both expectancy about one's own performance and objective measures of performance influenced participants' sense of control ratings. Additionally, the control rating from the previous trial predicted expectancy in the following trial, suggesting a sequential updating process. We implemented a sequential Bayesian updating model that integrates prior expectancy and new performance evidence to infer participants' control judgments. We discuss the model's performance and limitations in light of our empirical findings. Our goal is to integrate this framework into a cognitive model using the ACT-R architecture. This is ongoing work, and we share initial results.

Keywords: Action Control; ACT-R cognitive architecture; Bayesian modeling; Belief updating; Cognitive modeling

Introduction

Humans hold beliefs about their ability to act, that are shaped by prior and recent performance and individual weighting of new performance evidence. (Tuckman & Sexton, 1990). A reliable way to model this belief updating over time is Bayesian updating (Moutoussis, Fearon, El-Deredy, Dolan, & Friston, 2014), a framework that has been successfully applied to understand how the brain integrates information from various sources.

Behavior is inherently goal-directed: when a goal is achieved, the behavior is considered effective and associated with good performance; when a goal is not met, the behavior is seen as ineffective and consequently associated with poor performance. Humans generate behavior based on a mental model of the world, an internal representation of their environment. This model helps them interpret objects and dynamics around them, anticipate future states, and make decisions. By relying on this representation, individuals can derive goals they want to achieve and determine the actions needed to bring about the desired changes.

Self-belief about performance, or how capable an individual perceives itself to be, is a part of this mental representation. But the representation and self-belief are also interconnected in a bidirectional way: the representation guides behavior, which in turn affects self-belief, and this updated self-belief as part of the representation influences future behavior. This dynamical system of self-belief and behavior results in very efficient action control when we consider human behavior.

We aim to integrate the Bayesian framework for self-belief updating with a cognitive model of navigation. Our goal is to develop a model that explains behavior as arising from a mental representation of the environment, which includes the agent's self-belief. For this we use our custom-developed experimental environment, in which an agent steers a spaceship while trying to avoid crashing into obstacles. We explore how the mental model of the agent influences self-belief about performance and interacts with personality traits. To achieve this, we analyze and discuss how human experimental data fits a Bayesian updating process. In a next step, we introduce a cognitive model that we implemented in the ACT-R architecture. The ACT-R model features a situated state chunk that reflects the representation of the environment. Behavior is directly produced on the basis of the values held in the chunk slots. We would like to integrate these frameworks to give the model a sense of self-belief. This is ongoing work, and we welcome discussions on how to best combine both approaches.

Action Control

Action control is a broad concept describing how humans interact with their environment. It includes the preparation, execution, and regulation of actions. A core assumption is that human behavior is goal-directed, meaning that actions are executed with a desired state of the environment (or the agent within it) in mind (Frings et al., 2020). To do this, humans rely on a mental representation of the environment, allowing them to anticipate goal states and plan actions accordingly.

In real-world behavior, actions are rarely isolated. Instead, they form sequences of consecutive actions that are closely interconnected. This makes studying action control challenging. Despite these complexities, we focus on shedding light on the process of action preparation. Specifically, we investigate how people form self-beliefs about their own performance or control that is based on action control. We consider the self-belief to be part of their mental representation of the environment, or more precisely, of themselves within the environment. We propose that action preparation is largely driven by self-belief and that actual performance, in turn, feeds back into and updates this belief, creating a bidirectional relationship.

To model this evolving process of self-belief updating across trials, we plan to apply the Bayesian framework. But first, we examine how inter-individual differences, measured through the Big Five personality traits using the NEO-Five-Factor Inventory (NEO-FFI) (Borkenau & Ostendorf, 2008), are associated with this potential updating process. We present a first iteration of a cognitive model of action preparation, specifying the cognitive processes necessary for effective action planning within a custom-developed experimental environment. This model is implemented using the ACT-R cognitive architecture (Anderson et al., 2004). Ultimately, we aim on combining the Bayesian updating process and the ACT-R model of action preparation. With this, we want to contribute to the discussion on integrating mathematical models of cognition, such as Bayesian frameworks, with cognitive architectures that take a more symbolic and structured approach to modeling cognition.

Before an action is initiated, it must be prepared. A key part of action preparation is the process of selecting an action (Press, Gherri, Heyes, & Eimer, 2010). Often, multiple actions can achieve the same outcome or slightly different ones that still satisfy the goal state, and a choice must be made between the different options. However, humans are not optimal, but rather rational action controllers: we do not always choose the best possible option, but instead exhibit specific biases. Biases become especially pronounced under time constraints, which is common in everyday life. In these situations, there is insufficient time to weigh all possible actions. Affective beliefs, including the self-belief about one's own performance, are always present and influence action control. These beliefs, along with situational information, contribute to the action selection process through top-down processing. In this case, top-down processes help prioritize actions based on past experiences and internal goals.

Once an action is prepared, it is initiated and continuously adjusted based on sensory feedback (Synofzik, Vosgerau, & Newen, 2008). Prediction errors, differences between expected and actual sensory input while the action is executed, guide these adjustments in real time. The ability to correct the course of an action on the fly reduces the demand on the initial preparation/selection process. Action execution with the prediction and matching of sensory feedback, can be modeled using predictive coding (Rao & Ballard, 1999). One key aspect of this framework is that larger prediction errors require larger adjustments to maintain the intended course of action.

Once an action is completed, its outcome can be evaluated: whether or not the action goal was met. This evaluation may trigger adjustments, but unlike the real-time adjustments at the sensorimotor level that refine the course of an action, these occur at a higher cognitive level within the action control hierarchy. If the intended outcome is not achieved, for example, the mental model of the environment may need to be adjusted, as the failure implies that the current representation was inaccurate.

There are two different adjustments that can be made at this level. First, the representation of the environment may be refined, improving how the individual perceives and anticipates interactions within it. Second, the self-belief - a metacognitive construct reflecting, among other aspects, how the individual perceives itself as an efficient agent within the environment - may be updated, either strengthened or weakened depending on whether the action goal was achieved or not (Müller-Pinzler et al., 2022). These adaptations influence future action preparation, changing action selection in a way that enhances the likelihood of successfully achieving action goals; interacting more efficiently with the environment.

Self-Belief and Self-Belief Updating

Self-belief is shaped by both stable traits and dynamic, context-dependent confidence. This distinction was shown in fMRI studies, where global self-performance, a trait-like aspect of self-belief, was associated with activity in the ventral striatum, while local confidence, which varies with situational demands, was reflected in frontoparietal network activity (Rouault & Fleming, 2020).

Moreover, evidence suggests that self-belief is continuously updated through Bayesian inference, where prior experiences of meeting action goals are integrated with new evidence about the own performance (Moutoussis et al., 2014). The stable trait of global self-performance influences this updating process, regulating how much weight is given to new evidence that is either positive or negative.

In the following, we present an experiment designed to assess self-belief based on individual performance, with the goal of modeling it as a Bayesian updating process. We also examine all five personality traits of the Big Five, using the NEO-FFI questionnaire, under the assumption that these traits might influence how new evidence is weighted in the self-belief Bayesian updating process.

For the time being, we focus on deriving self-belief from action preparation (based on mental representations) and not from the ability to correct poorly prepared actions during execution. Therefore, we have designed our experimental paradigm in such a way that excludes the possibility to correct actions in real time. This presented a methodological challenge, which we address with an environmental feature we call drift (see experiment section).

Methods

Experiment

The main part of the experiment consists of a game in which participants have to steer a spaceship through a closed environment which is bounded by walls on both sides, left and right. The experimental environment is implemented in Python (Van Rossum, Guido & Drake, 2009) using the PyGame package (Shinners, 2011) and runs at 60 FPS[1]. The spaceship automatically moves downwards through the environment at constant speed. Participants are tasked to avoid crashing into walls or obstacles and make it to the finish line at the bottom of the environment. They can use the keyboard to move horizontally: [Y] for left and [M] for right. The spaceship is always displayed in the screen center and the environment will shift around the spaceship. Therefore, if participants want to move to the right, the whole environment is shifted to the left instead.

Participants encounter so-called *drift* sections, indicated by a red bar on either side. Once the spaceship enters a drift section with the red bar displayed for example on the right side, the spaceship will be pushed to the left for the duration of the drift section; the length of the red bar (if in contrast the red bar is displayed on the left, the spaceship is pushed to the right). An example of such a drift section can be seen in Figure 1. Within a drift section, it is not possible for the participants to steer, i.e. pressing the Y and M keys does not have any impact on the spaceship's movements. However, the horizontal push of the drift is constant and can therefore be anticipated once the participants have learned about it.

Drift is the key environmental feature that participants must mentally represent to accurately predict the spaceship's trajectory and to effectively prepare for drift sections.

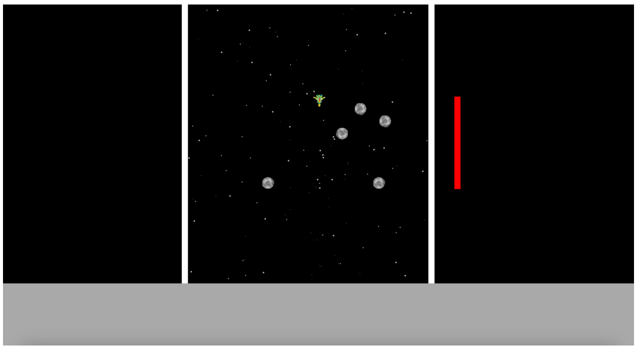

Figure 1: Depiction of an instance within the experimental environment. Shown here is the space of the environment that is drawn on the screen during the experiment. Obstacles are scattered between the two walls on the left and right. The red bar indicates a drift section with push to the left.

The study is divided into multiple levels. Every level consists of four sequential drift sections. Obstacles that need to be avoided are placed solely within drift sections. All drift sections have the same length, but the distance between two following sections and the side on which the red bar is displayed (its push direction) vary. Levels were created randomly including the position of the obstacles, but it was ensured that there is always a possible way for the spaceship to navigate through them.

If participants manage to pass all four drift sections, they reach a green line at the end which means the level is successfully completed. It took on average 18.64 seconds (3.53) to complete a level (deviations may occur due to rendering of the game environment). If the spaceship crashes into an obstacle or in one of the walls, the level ends prematurely and a display is shown with a notification about the crash.

Procedure

Before participants started with the actual trials of the experiment, a training level was presented to familiarize them with the environment and steering the spaceship. The training level consisted of multiple drift sections, first without and then with obstacles. If participants crashed during the training, the training level was presented again until they successfully navigated through it once to ensure the task was sufficiently understood. After that, the trials started. In total, there were 80 different levels. If a participant failed to complete a level, the same level was presented again - up to three times in total - later during the experiment (level was inserted at a random position within the sequence). Before every trial, participants were asked to rate the expectation of their own performance in the upcoming trial (”In the upcoming run, how good do you think you will perform?”) on a scale ranging from 1 (”very poorly”) to 7 (”excellent”) which we will refer to with the term *expectancy* from now on. After each trial, they had to rate their *sense of control* (SoC; ”In the most recent run, how strong was your feeling of control?”) on the same scale (1-7) which is here used to assess self-belief. Depending on the participants' performance, the main part of the experiment took approximately 40-55 minutes. After having completed the trials, participants were asked to fill out two questionnaires. The first one was the German version of the 30 items NEO-FFI which is commonly used to assess the Big Five personality traits: Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism. The second questionnaire was a short four item questionnaire asking for video game expertise such as the frequency and types of games played (Nolan, 2022). From briefing to debriefing and including time for potential questions by participants, the whole experiment took approximately one hour and 10 minutes.

Data Analysis

A total of 26 participants were collected between 09.02.2025 and 12.05.2025 and data collection was completed since submitting the first version of this paper. Below we describe the final model selection.

We used linear mixed modeling (Bates, 2015) in the R programming language (R Core Team, 2022) to analyze the factors influencing responses to both the SoC and expectancy question separately. A Box-Cox distributional analysis (Box & Cox, 1964) showed that the predicted variable SoC required a square root transformation. No transformation was required for expectancy. We included participant ID as a random intercept for both model selections.

---

[1] Git repository containing the experimental environment.

Either model selection was done applying the following procedure. We determined the fixed effects structure by comparing Bayesian Information Criterion (BIC) values (Chakrabarti & Ghosh, 2011). Since BIC penalizes models with more parameters while also accounting for the number of observations, this approach helps prevent overfitting and ensures a simpler final model. Starting with a full model that included all possible fixed effects, we removed individual variables step by step. After each removal, we compared the new model's BIC with the previous one. Once the fixed effects structure had been determined, we evaluated the random slopes by also referring to the BIC - starting with the full random effects structure and removing individual random slopes step by step. This way, we successively approached the best model structure for the present data set.

All models were fitted using the maximum likelihood criterion and we used a random seed (36) for reproducibility of our results[2].

Results

On average participants successfully solved 24.02% (0.07) of the trials.

For the effects in the linear models predicting SoC and expectancy, we report parameter estimates along with their standard deviations, confidence intervals (obtained from parametric bootstraps with 10,000 iterations), and p-values.

SoC. We predicted SoC responses on a trial-by-trial basis.

Referring to the intraclass correlation coefficent of a null model, participant ID accounted for 46.05% of the total variance in SoC responses, indicating a notable individual difference component.

The initial model included all five dimensions of the NEO-FFI. However, none of these factors remained in the final model after model selection. The final model featured the following predictors: expectancy (response to the expectancy question before the trial), N_{drift} (how far participants progressed in the trial, ranging from 0 to 4), and *crashed* (whether the participant crashed in the current trial).

The higher participants expected their performance to be, the higher they rated their own control during the trial ($\beta = 0.10$, $\sigma = 0.01$, 95% CI [0.07, 0.12], $p < .01$). Actual performance also influenced SoC. It increased with the number of successfully completed drift sections (N_{drift}; $\beta = 0.07$, $\sigma < 0.01$, 95% CI [0.06, 0.08], $p < .01$). Conversely, SoC decreased when participants crashed in the current trial ($\beta = -0.32$, $\sigma = 0.04$, 95% CI [-0.39, -0.24], $p < .01$).

Expectancy. We again predicted expectancy on a trial-by-trial basis.

In a null model, the intraclass correlation coefficent indicated that participant ID accounted for 58.90% of the total variance in expectancy responses, highlighting significant individual differences also in performance expectations.

The final model predicting expectancy included two predictors: $N_{crash-success}$ (the number of times a participant crashed but successfully completed the very next trial) and SoC_{n-1} (the SoC response from the previous trial).

Interestingly, after random slopes had been selected, $N_{crash-success}$ did not yield significance ($\beta < 0.01$, $\sigma < 0.01$, 95% CI [<-0.01, 0.02], $p = .153$). However, higher SoC ratings during the last trial were associated with significantly greater expectancy, increasing with each point on the 1–7 scale ($\beta = 0.34$, $\sigma < 0.01$, 95% CI [0.33, 0.36], $p < .01$).

Discussion of Experimental Data

SoC as self-belief was influenced by both prior expectancy and new evidence about performance. However, against our hypotheses, there were no significant effects for any of the NEO-FFI personality dimensions. We assume that SoC results from a sequential updating process that can be modeled in Bayesian terms, where expectancy serves as the prior and new evidence about performance acts as the likelihood. In this framework, we had expected personality traits to modulate how prior and likelihood are weighted during updating, but the lack of significant personality effects implies a more uniform updating process across individuals.

Expectancy was mainly explained by the last trial's SoC, suggesting that the posterior belief from the previous trial carries over into the next one as prior knowledge, consistent with our Bayesian interpretation. Based on these findings, we have now constructed a mechanistic Bayesian updating model of self-belief that defines the sequential updating process we assume to be underlying SoC.

Sequential Bayesian Updating of Self-Belief

We translated the findings of our model selection directly into the structure of a Bayesian updating model of self-belief. We introduce weighting parameters driving the relative influence of prior expectancy and new performance-related evidence. The prior weight is fixed at 1, a reference, while the weight on the likelihood is treated as a free parameter that the model estimates from the data. This parameterization allows us to examine how strongly participants rely on new evidence relative to their prior expectations when forming SoC judgments.

$\beta_{i,1}$ and $\beta_{i,2}$ denote the participant-specific regression coefficients. The likelihood components are $N_{i,t}^{\text{drift}}$, representing the numerical *drift* evidence on trial t for participant i, and $C_{i,t}$, representing the *crashed* evidence. The log-likelihood weight is then given by:

$$\log w_{i,t} = \beta_{i,1} \cdot N_{i,t}^{\text{drift}} + \beta_{i,2} \cdot C_{i,t} \tag{1}$$

The resulting likelihood weight is:

$$w_{i,t} = \exp\left(\log w_{i,t}\right) \tag{2}$$

The predicted SoC, denoted $\mu_{i,t}^{\text{SoC}}$, is computed as a precision-weighted average of prior and likelihood:

---

[2] Git repository containing data and R analysis script (and also the sequential Bayes model described later).

$$\mu_{i,t}^{\mathrm{SoC}} = \frac{Prior_{i,t} + w_{i,t} \cdot N_{i,t}^{\mathrm{drift}}}{1 + w_{i,t}} \tag{3}$$

Finally, the observed SoC judgments are normally distributed around the posterior mean:

$$\mathrm{SoC}_{i,t} \sim \mathcal{N}\left(\mu_{i,t}^{\mathrm{SoC}}, \sigma^2\right) \tag{4}$$

Dynamic Likelihood Weighting

In this paper, we cannot examine every participant, but we highlight a small selection as examples. Their individual likelihood weighting results are shown in Figure 2. We divided these participants into low and high performers, based on the idea that performance (i.e., the likelihood) would influence how strongly it is weighted when rating one's own SoC. However, we do not observe a clear difference in likelihood weighting based on performance level. While we expected high performers to give more weight to performance evidence (a), we also see that some high-performing individuals assign relatively low weight to it (though the likelihood weight is still higher than that of the prior which is indicated by the dashed line). Similarly, low-performing individuals (b) may also place a strong emphasis on performance evidence when forming their SoC. We had expected that, especially among low performers, personality traits might help explain these differences in weighting. However, since the personality effects were not significant, we cannot draw firm conclusions about the source of these individual differences.

To evaluate the overall fit of our Bayesian model and its ability to predict individual SoC ratings, we computed the root mean squared error (RMSE) between predicted and observed ratings. The model yielded an RMSE of 0.87 (SD = 0.22), indicating that, on average, predicted ratings deviate from actual ratings by nearly one point on the 7-point SoC scale. Moreover, the model's goodness of fit for individual participants is not explained by their overall success rate. That is, the model does not systematically predict SoC ratings more accurately for high performers than for low performers. While the RMSE shows that the model captures overall patterns in the data, it fails to account for more dynamic within-participant changes in weighting across trials.

Figure 2 suggests that the model underestimates abrupt shifts in how participants rely on performance evidence. For instance, in some cases, a crash followed by a successful trial seems to trigger a sharp increase in the weighting of the likelihood - an effect the model does not fully anticipate and thus may lead to systematic prediction errors in these situations. This is most likely due to the current model's estimated weighting parameter being based only on how many drift sections had been encountered ($N_{i,t}^{\mathrm{drift}}$) and whether there was a crash or not ($C_{i,t}$; Equation 1). These limitations highlight areas where the model could be refined and should be kept in mind when evaluating its predictive performance.

During model selection, we considered $C_{trial-1}$ as a potential fixed effect for SoC, but it was excluded from the final model. One possibility is that the model is currently missing a relevant interaction between $C_{trial-1}$ and N_{drift}, which may be necessary to improve predictive accuracy. A clear next step in our analysis is to revisit the model selection process with this interaction term included.

Once the Bayesian model reliably captures how participants form judgments about their sense of control (SoC) and the information they rely on, we can embed this mechanism into our cognitive model of action preparation. The aim is for the cognitive model to use the current SoC as a basis for action planning and to adapt its behavior dynamically, depending on how much in control it perceives itself to be.

We have already begun developing a model of action preparation that starts with a rough idea of drift (how long it lasts and how far it pushes) and refines its internal representation over time through acting within the DodgeAsteroids environment. In the following section, we provide a conceptual overview of this model.

A Computational Cognitive Model of Action Preparation

The ACT-R cognitive architecture (Anderson et al., 2004) separates declarative memory (facts, concepts) from procedural memory (how-to knowledge/rules). This makes it well-suited for our purposes, as the model can encode an initial impression of drift as a chunk in declarative memory, which is then turned into effective procedural strategies. But, using ACT-R for this kind of learning process also presents challenges. It requires careful design of how declarative knowledge, such as different representations of drift, is encoded and retrieved. The model must also handle increased complexity when the environment contains subtle or probabilistic cues, which can be difficult to represent and act on within a symbolic framework. Since the ACT-R model is not the primary focus of this paper, we provide only a brief overview.

Model Components and Functions

Our model of action preparation is implemented using a hybrid approach. Some of the core components are written in Python and executed within ACT-R running in Lisp, using the py4cl library for interoperability. The Python experimental environment and the ACT-R model communicate through ACT-R's RPC interface, managed by a socket connection.

Drift is encoded as a vector in the `situated-state` chunk within declarative memory. This chunk includes slot values for the predicted horizontal displacement (`drift-x`), the estimated duration of the drift (`drift-y`), and the direction from which the drift originates (`drift-direction`), based on the side where the red drift-indicating bar is perceived. In our current model, `drift-x` is initialized with a random value and `drift-y` is set to the true length of the drift section. The chunk also stores the spaceship's current horizontal and vertical position (`agent-x` and `agent-y`), providing a unified representation of both, the environmental dynamics and the agent's state. The chunk is retrieved when the model per-

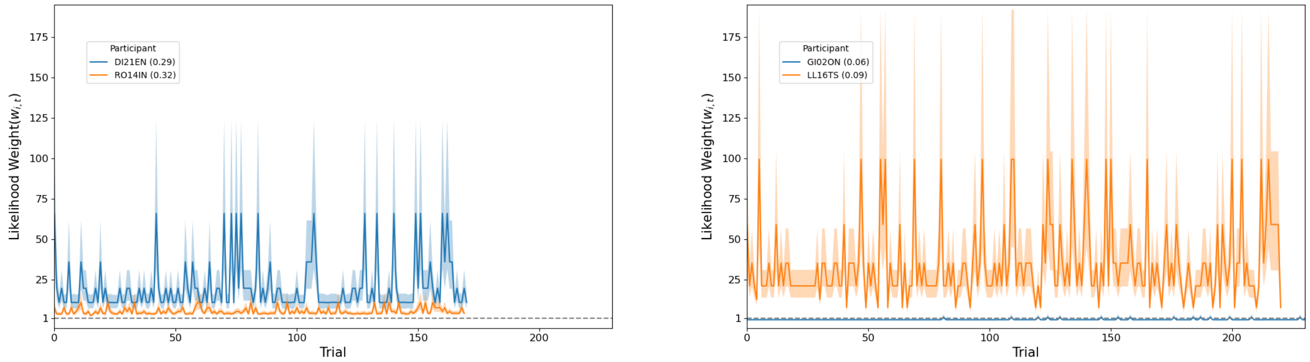

(a) High performers

(b) Low performers

Figure 2: Dynamic likelihood weighting across trials for four participants, grouped by (a) high and (b) low performance. Axes are scaled equally for better comparison. Dashed lines indicate the fixed prior weight of 1. The low-performing individuals required more trials to complete the experiment and are identified by lower success rates (in parentheses after participant IDs). There is no clear difference in likelihood weighting between high and low performers. In both groups, there are individuals who place considerably more weight on the likelihood, as well as those who weight prior and likelihood nearly equally.

ceives a red object just outside the environment's walls, either on the left or right, which serves as a retrieval cue.

To prepare an action, the model uses the values in the chunk to project a drift vector onto the upcoming drift section, ensuring that the vector does not pass through any of the obstacles. The drift vectors x-component corresponds to `drift-x`, while the y-component corresponds to `drift-y`. The sign of the x-component is determined by `drift-direction` (that is, whether the drift originates from the left or right side of the agent) so that the vector points in the correct horizontal direction when projected. The vector is anchored at the vertical start of the drift section. The model then selects a high-level goal: the position it aims to steer toward, based on the starting point of the projected trajectory. It must learn that not all positions can be reached from the current position of the spaceship, `agent-x` and `agent-y`. Depending on the spaceship's horizontal position, the vertical distance to the drift section might not allow enough time to steer horizontally to the target position given by `high-level goal-x`. This requires balancing between the safest available path and one that is physically achievable.

Motor control in the model is straightforward. Once a high-level goal is selected, the model compares the goal's x-coordinate with the current horizontal position of the spaceship (`agent-x`). If the goal lies to the left, it initiates a Y-key press; if to the right, it initiates an M-key press. In cases where no high-level goal is defined, the model defaults to steering toward the horizontal center of the environment.

It is important to note that we consider this vector-based computation a hack rather than an exact cognitive process. Humans likely do not engage in explicit vector calculus to plan movement. Instead, we assume they rely on heuristics based on spatial features. To better reflect this in upcoming iterations, we will define ACT-R production rules that capture rule-based action preparation.

Next Steps

Our next steps focus on two main objectives: reconsidering the parameterization of the Bayesian updating process underlying self-belief, and refining the computational cognitive model of action preparation. Through the exploration of additional interactions, for example between $C_{trial-1}$ and N_{drift}, we aim to capture more accurately how prior experiences and new performance outcomes dynamically influence belief updating. Additionally, we plan to expand the Bayesian inference mechanism to support probabilistic updating of hypotheses about the drift environmental feature, enabling the model to adapt more flexibly to uncertainty in task dynamics.

We will continue refining our computational cognitive model to align more closely with human behavior, improving its mechanisms for action selection and navigation. The model is supposed to successfully complete trials at a level comparable to human participants. Once this is achieved, we plan to integrate the Bayesian updating framework into the cognitive model, allowing it to explain both performance and self-belief updating in response to task outcomes.

A key challenge remains in effectively combining these two modeling frameworks. One potential method is embedding the mathematical model identified through model selection into the ACT-R cognitive architecture. This would allow the cognitive model to not only replicate human performance but also dynamically update self-belief based on task outcomes. However, we seek to explore alternative ways of integrating these frameworks beyond simply embedding a data-driven mathematical model within ACT-R. Future work will focus on different integration approaches to develop a unified framework that captures both action control and self-belief updating in a cognitively plausible manner.

Acknowledgments

This research was funded by the German Research Foundation (DFG) Priority Program 2134 'The Active Self'.

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Bates, D. (2015). *MixedModels: A Julia Package for Fitting Statistical Mixed-Effects Model.* Retrieved from `https://github.com/dmbates/MixedModels.jl`

Borkenau, P., & Ostendorf, F. (2008). *NEO-FFI : NEO-Fünf-Faktoren-Inventar nach Costa und McCrae, Manual.* Hogrefe. Retrieved from `https://pub.uni-bielefeld.de/record/1900471`

Box, G. E., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *26*(2), 211–243. Retrieved 2024-01-03, from `https://academic.oup.com/jrsssb/article-abstract/26/2/211/7028064` (Publisher: Oxford University Press)

Chakrabarti, A., & Ghosh, J. K. (2011, January). AIC, BIC and Recent Advances in Model Selection. In P. S. Bandyopadhyay & M. R. Forster (Eds.), *Philosophy of Statistics* (Vol. 7, pp. 583–605). Amsterdam: North-Holland. Retrieved 2024-01-03, from `https://www.sciencedirect.com/science/article/pii/B9780444518620500186` doi: 10.1016/B978-0-444-51862-0.50018-6

Frings, C., Hommel, B., Koch, I., Rothermund, K., Dignath, D., Giesen, C., ... others (2020). Binding and retrieval in action control (brac). *Trends in Cognitive Sciences*, *24*(5), 375–387.

Moutoussis, M., Fearon, P., El-Deredy, W., Dolan, R. J., & Friston, K. J. (2014). Bayesian inferences about the self (and others): A review. *Consciousness and cognition*, *25*, 67–76.

Müller-Pinzler, L., Czekalla, N., Mayer, A. V., Schröder, A., Stolz, D. S., Paulus, F. M., & Krach, S. (2022). Neurocomputational mechanisms of affected beliefs. *Communications biology*, *5*(1), 1241.

Nolan, D. (2022). *Video survey.* `https://www.stat.berkeley.edu/~nolan/surveys/videoSurvey.pdf`. Retrieved from `https://www.stat.berkeley.edu/~nolan/surveys/videoSurvey.pdf` (Accessed: 2022-12-08)

Press, C., Gherri, E., Heyes, C., & Eimer, M. (2010). Action preparation helps and hinders perception of action. *Journal of Cognitive Neuroscience*, *22*(10), 2198–2211.

R Core Team. (2022). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `https://www.R-project.org/` (Version 4.1.0)

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, *2*(1), 79–87.

Rouault, M., & Fleming, S. M. (2020). Formation of global self-beliefs in the human brain. *Proceedings of the National Academy of Sciences*, *117*(44), 27268–27276.

Shinners, P. (2011). *PyGame.* Retrieved 2022-08-29, from `http://pygame.org`

Synofzik, M., Vosgerau, G., & Newen, A. (2008, March). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, *17*(1), 219–239. Retrieved 2023-11-27, from `https://www.sciencedirect.com/science/article/pii/S1053810007000268` doi: 10.1016/j.concog.2007.03.010

Tuckman, B. W., & Sexton, T. L. (1990). The relation between self-beliefs and self-regulated performance. *Journal of Social Behavior and Personality*, *5*(5), 365.

Van Rossum, Guido, & Drake, F. L. (2009). *Python 3 Reference Manual.* Scotts Valley, CA: CreateSpace.

3.3 Conclusion

Despite its limitations, the Bayesian model establishes important groundwork for understanding the SoC as a process of belief updating. It formalizes how performance outcomes update expectations to produce subjective control judgments and provides a quantitative framework that can be tested and refined. However, the model does not specify the mechanisms that generate expectations or produce action plans based on these expectations. The next chapter addresses these gaps by embedding belief updating within a mechanistic architecture that simulates strategy learning and action control through explicit cognitive processes that factor in a form of self-belief.

4 Cognitive Modeling of Error-Driven Learning in Human Action Control

The present study builds on the empirical and theoretical insights developed throughout the previous chapters by embedding belief updating, uncertainty, and action planning into a mechanistic cognitive model. A mechanistic model in cognitive science is a model that explains how a cognitive system solves a task by specifying the internal components, processes, and interactions that produce observable behavior.

Studies 1 and 2 examined sensorimotor functions and inferred different modes of action control, demonstrating how humans adapt their behavior in response to environmental dynamics. Study 3 shifted focus to subjective control judgments and identified belief updating as a key mechanism underlying the SoC. However, these studies stopped short of specifying how such beliefs are formed, maintained, and used during action selection. Study 4 addresses this gap by developing a computational cognitive model that implements these mechanisms in an environment requiring adaptive action control.

4.1 Prediction, Uncertainty, and Adaptive Behavior

Action control in dynamic environments is fundamentally predictive. Agents must anticipate how the environment will evolve in response to their actions, update predictions based on prediction errors, and choose actions that minimize risk. This perspective aligns with broader theoretical frameworks such as predictive processing (Clark, 2013) and active inference (K. Friston, 2010), which position prediction and error minimization as core organizing principles of cognition and behavior.

Studies 1 and 2 examined the effects of input noise and drift on eye-movement control, revealing two distinct modes of adaptive behavior. Input noise led to closer monitoring of the controlled object through reactive gaze adaptation, while drift prompted anticipatory scanning of future positions through proactive adaptation. The key feature that distinguishes input noise and drift is the degree of uncertainty associated with the respective environmental dynamic. Input noise is inherently less predictable than drift, as it lacks consistency that can be learned. This distinction motivates the hypothesis that uncertainty plays a central role in action selection. When environmental dynamics are uncertain, agents may rely on conservative heuristics and less on outcome predictions; when dynamics become predictable (either through learning or inherent consistency) agents can adopt more refined, model-based strategies that incorporate predicting the outcomes of various actions currently available to the agent.

Crucially, environmental dynamics in the real world rarely consist of pure random noise (like input noise) but rather contain learnable structure and regularities. Human action controllers are particularly adept at rapidly extracting these consistencies and integrating them into predictive models that guide action selection (Clark, 2013; K. Friston, 2010; Schulz and Gershman, 2019). Capturing this adaptive transition from reactive to predictive control requires a framework that unifies learning, prediction, and action selection within a single operational architecture.

The ACT-R cognitive architecture provides such a framework (J. R. Anderson and Lebiere, 1998). It specifies mechanisms for declarative memory, procedural learning, and resource-limited perception and action, and therefore enables holistic modeling of human cognition. By embedding predictive, uncertainty-driven learning and action selection in ACT-R, I can implement the theoretical concepts of hierarchical control and predictive processing discussed in Chapter 1. The resulting model can then be evaluated by comparing simulated behavior against human performance.

ACT-R simulates learning as an interaction between declarative and procedural memory systems. The architecture builds up and maintains an episodic memory of past experiences that can be retrieved to inform future decisions. This capability is critical for the present study: the model must not only learn a mapping between drift indicators and outcomes but also store and recall past episodes of obstacle configurations to guide action selection in the present.

This approach builds on previous work from our group that examined adaptive processes in multitasking using reinforcement learning (RL) models (Österdiekhoff, Heinrich, Russwinkel, et al., 2026). RL is fundamentally designed to exploit learnable structure from environmental dynamics through interaction. In Österdiekhoff et al. (2026), we applied this method to find efficient task-switching strategies in a dual-task environment in which two tasks are concurrently running, while only one task can be seen and controlled at a time. The agent architecture consists of two subtask agents and a meta-agent that selects which task to control at each step. The meta-agent's task-switching decision incorporates a SoC computed by prediction modules that estimate action outcomes. The SoC is based on two signals: a prediction error measuring how the predicted next state differs from the actual one, and a need-for-control value reflecting how much better the agent would perform by acting optimally rather than doing nothing. These signals provide a simple estimate of how well the agent can currently manage the current task. The meta-agent uses the resulting SoC value to choose whether to stay in the current or switch to the other task. Across nearly all tested conditions—including easy and hard level layouts and different intensities of input noise in none, one, or both subtasks—agents equipped with this SoC mechanism outperformed agents without SoC, a switch-every-frame baseline, and even human participants.

However, that work was primarily inspired by cognitive mechanisms rather than designed to be cognitively realistic. While RL models can learn effective policies in dynamic environments, they lack native support for episodic memory retrieval. Although they can be engineered to incorporate predictive representations of environmental dynamics, this requires substantial additional architecture and does not necessarily reflect how humans learn and adapt.

The present study therefore adopts ACT-R as a more cognitively plausible modeling framework. The goal is theory-testing; to build a cognitive model that learns to control the spaceship in the Dodge Asteroids environment in a human-like manner. The model learns through error-driven updates of its internal representation of drift dynamics, using memory retrieval, predictions based on internally maintained drift parameters, and uncertainty-sensitive action selection to navigate drift sections and adopt efficient strategies. By comparing the model's performance to human learning trajectories from Study 3, I can evaluate whether its internal mechanisms reproduce human behavioral patterns. I specifically focus on the interplay between prediction errors, uncertainty, and memory retrieval to assess both the plausibility and the limitations of the underlying mechanisms.

The study makes several contributions. First, it demonstrates how adaptive action control can emerge from the interplay between memory, prediction, and uncertainty within a unified cognitive architecture. Second, it provides a mechanistic account of the belief-updating processes identified in Study 3, showing how internal models are constructed and refined through experience. Third, it serves as the culmination of the dissertation's progression from conceptual theory (Chapter 1) to empirical investigation (Chapter 2 and 3) to mechanistic modeling and bridges philosophical insights about predictive processing and hierarchical control with a concrete, testable computational implementation.

4.2 Study 4: Error-Driven Learning in Human Action Control: A Computational Cognitive Model

Study 4 refers to the exact same study reported in Chapter 3, but here I assessed different metrics in the data and used it for a different, more mechanistic type of modeling. To reiterate, the experimental task simulates a spaceship navigating a vertical corridor containing obstacles. Participants can freely steer the ship horizontally but only during non-drift segments. At several points in the corridor, the ship enters a drift section in which horizontal control is temporarily disabled, and the ship is displaced in a visually indicated direction. The objective is to align the ship's horizontal position before the drift so that the forced displacement will not result in a collision.

The task structure creates a clear separation between planning and execution. Because steering is disabled during drift, action selection must occur before the drift begins and relies on prediction rather than online correction. Early in learning, there are many severe prediction errors. The visual drift indicator must be associated with the actual displacement that follows. Over time, participants refine their expectations of drift magnitude and direction, reduce uncertainty, and adopt efficient strategies for entering drift sections. This combination of prediction, error feedback, and strategy adaptation makes the task well suited for cognitive modeling.

Participants learned the task rapidly. Although initial performance varies widely, most individuals converge on an efficient steering strategy within 9 trials. They become increasingly consistent in their choice of entry points and show reduced variance in drift estimation. I used three behavioral markers as empirical benchmarks for evaluating the cognitive model: (i) the duration of the early exploration phase, (ii) the drift-entry distance from the horizontal center, and (iii) overall performance.

4.2.1 Cognitive Model Implementation

The cognitive model is implemented in ACT-R reflects the core idea that action control emerges from the interaction between learned predictions and uncertainty-driven decision processes. Predictions are implemented through declarative memory, action selection through procedural rules, and uncertainty through a dynamic parameter integrated into the task representation.

The model consists of the following interacting components. The perceptual module extracts the drift indicator and relevant spatial information. Declarative memory stores both an internal representation of the drift dynamics and experiences of drift outcomes. The drift representation is implemented as a set of parameters that build a linear model that predicts the resulting lateral movement; these parameters are updated via prediction errors after each drift section. Each drift experience is encoded as a chunk containing a discretized kernel representation (the convolved visual input), the selected entry point into the drift section, and whether the section was completed successfully or ended in a crash. Procedural memory includes production rules that map perceptual cues and retrieved memories to action decisions.

When approaching a new drift section, the model retrieves the most relevant chunk containing a past drift experience using ACT-R's similarity-based retrieval. It also retrieves the drift linear model representation. The retrieved drift model provides a prediction of the upcoming displacement. At all times, the cognitive model holds an uncertainty value that reflects its confidence in its predictions.

If the retrieved past instance was successful and its entry point appears safe in the current situation, the model uses that entry point directly. However, when no past instance is retrieved, when the retrieved drift instance resulted in a crash, or when the cognitive model evaluates that the entry point would result in a crash in the current situation, it generates multiple candidate entry points. Each candidate is assessed according to its risk. The safest candidate entry point is selected as the target position to steer toward.

Learning occurs on multiple levels. Declarative memory accumulates past experiences of action selections for drift sections and improves the accuracy of retrieval-based predictions. Procedural learning strengthens production rules that lead to successful outcomes and enables more efficient and consistent action selection over time. Prediction errors drive updates in the drift representation. After each drift, the model compares the predicted displacement to the actual outcome, adjusts its internal parameters, and updates the associated uncertainty. This mechanism operationalizes the belief-updating framework introduced in Chapter 3.

The free parameters of the cognitive model were fitted to human data based on two criteria: first, the distribution of performance across trials, specifically the proportion of trials in which 0–4 drift sections were successfully solved, and second, the learning time, defined as the number of trials required to establish a stable steering strategy.

4.2.2 Results

The ACT-R model initially included a conservative fallback strategy. When uncertainty exceeded a threshold, the model selected an entry point without generating predictions or assessing risk. Instead, it computed a horizontal offset to the corridor center based on the uncertainty parameter and did not factor in the actual placement of obstacles. This was meant to reflect an uninformed, computationally cheap early learning strategy. The hypothesis was that this approach might capture cautious early human behavior.

However, the model with this fallback strategy learned much more slowly than human participants, settling on a stable steering strategy only after approximately 44 trials (compared to 9 trials in humans). Because uncertainty is typically high during early learning, the fallback heuristic was activated frequently in the initial phase. Since no predictions are generated under the fallback, it reduced exposure to prediction errors and consequently slowed adaptation. The model also deviated from human exploration behavior in that it initially sampled entry points with large distances from the horizontal center of the corridor, whereas human participants began with more conservative entry points closer to the center. Despite these differences, the model broadly captured overall performance patterns, showing high proportions of trials in which no drift sections were solved, few trials with 1–3 drift sections solved, and a subsequent increase in trials in which all four drift sections were successfully completed. The discrepancies motivated a revision of the model with the fallback mechanism removed.

The revised model forced the agent to rely on its uncertain predictions even in early trials. This increased exposure to prediction errors resulted in a steeper learning curve of 14 trials. The revised model also better matched other observed human behavioral patterns. Specifically, it initially sampled more conservative entry points with smaller distances to the corridor's horizontal center and gradually increased these distances over the course of the 14-trial exploration phase. Moreover, the revised model captured overall performance patterns even more accurately by reproducing the characteristic distribution of trials in which drift sections were solved. It showed higher proportions of trials with either 0 or all 4 drift sections completed, and lower proportions of trials in which only 1–3 sections were solved. This finding challenges the initial assumption that humans would adopt prediction-independent strategies under high uncertainty. Instead, they appear to rely on prediction earlier and more heavily than anticipated, using prediction errors to refine internal forward models of drift dynamics.

4.2.3 Discussion and Limitations

The success of the fallback-free model suggests that effective learning in dynamic environments depends on engaging with predictive mechanisms early, even when uncertainty is high, rather than deferring to prediction-independent heuristics until confidence is established. The revised model achieved performance comparable to human participants, and even the behavior during exploration resembled that of humans. However, humans' exploration phases are still shorter than the models.

Several limitations remain. The model implements drift estimation through a linear model, which, while potentially plausible if humans mentally represent drift as a directional vector, may oversimplify the actual cognitive representations involved. The model's uncertainty value modulates how many candidate trajectories are generated and evaluated, but predictions themselves may also adapt in sophistication as a function of uncertainty. For instance, humans under high uncertainty might not mentally simulate complete trajectories but instead rely on simpler heuristics within the prediction process itself. Further, the architectural assumption of retrieval-based drift predictions may not fully capture human generalization abilities. Future extensions could incorporate hierarchical Bayesian inference, as discussed in Chapter 3 for subjective SoC, for a more sophisticated modeling of uncertainty dynamics.

4.2.4 Contribution to Study 4

The study described here and Study 3 in Chapter 3 refer to the same experiment, and therefore the same contributions apply regarding the conceptualization of the study idea and design, programming of the experimental task, and data collection. In addition, I conducted the data analysis, implemented and evaluated the ACT-R cognitive model, created all visualizations, contributed substantially to the interpretation of the results, and wrote the manuscript.

Study 4

Heinrich, N. W., Österdiekhoff, A., Stuwe, M., Kopp, S., & Rußwinkel, N. (2026).

Error-Driven Learning in Human Action Control.

Manuscript submitted for publication.

Error-Driven Learning in Human Action Control: A Computational Cognitive Model

Nils Wendel Heinrich[1,2], **Annika Österdiekhoff**[3], **Martin Stuwe**[1], **Stefan Kopp**[3], **Nele Russwinkel**[1]
nilsheinrich91@gmail.com
[1]Institute of Information Systems (IFIS), Department of Computer Science and Engineering, Universität zu Lübeck, Lübeck, Germany
[2]Cognitive Psychology and Ergonomics, Technische Universität Berlin, Berlin, Germany
[3]Social Cognitive Systems, Faculty of Technology, CITEC, Bielefeld University, Bielefeld, Germany

Abstract

Action control in dynamic environments requires adaptive mechanisms that integrate prediction, memory retrieval, and uncertainty adjustments. This study presents a cognitive model of action control that is situated in a dynamic task environment. The task involves steering a spaceship through a corridor while avoiding obstacles. It features drift sections in which control is temporarily lost and the ship is pushed in a visually indicated direction. The cognitive model is implemented in ACT-R and combines declarative–procedural learning with an adaptive drift model that continuously updates its parameters and associated uncertainty through prediction errors. It uses its learned representation of drift dynamics to select actions that minimize collision risk. Model performance was evaluated against human behavioral data in terms of final performance levels and the number of trials required to develop an efficient steering strategy. The cognitive model initially relied on a heuristic fallback for action selection under high uncertainty but took significantly longer than humans to adopt an efficient strategy. A revised model without this fallback learned faster and more closely reproduced human-like learning patterns. The findings demonstrate that action control in continuous environments emerges from the interplay between memory, prediction, and uncertainty. Our framework offers a bridge between symbolic cognitive architectures and predictive, error-driven theories of learning.

Keywords: Action control; Dynamic environments; ACT-R cognitive architecture; Cognitive modeling; Error-driven learning; Predictive control

1 Introduction

Humans constantly interact with complex and unpredictable environments, from navigating crowded streets to operating tools or vehicles. Successfully achieving goals in such contexts requires more than simple reflexes: the brain must predict the consequences of actions, detect errors, and adapt behavior over time (Wolpert & Flanagan, 2001; Shadmehr, Smith, & Krakauer, 2010). Understanding how humans control actions in dynamic environments and building cognitive models of these processes not only advances cognitive science but also informs the development of artificial agents and decision systems that can act robustly under uncertainty (Laird, Lebiere, & Rosenbloom, 2017).

1.1 Action control

Successful action in everyday life requires translating goals into effective movements while coping with dynamic environments. Objects move, forces shift, and unexpected events can disrupt even well-prepared actions. The human brain relies on forward models to predict the outcomes of planned actions (Wolpert & Landy, 2012; Shadmehr et al., 2010). By comparing predictions with sensory feedback, the brain detects errors and adjusts actions in real time, maintaining goal-directed behavior despite variability and noise. Because the relationship between actions and outcomes can change across contexts, forward models are continuously refined: individuals update their internal models based on prediction errors, gradually improving their capacity to anticipate outcomes and execute actions proactively (Shadmehr et al., 2010; Wolpert, Diedrichsen, & Flanagan, 2011). With repeated experience, this adaptive process enhances flexibility and supports performance across a range of dynamic challenges.

Action control also draws on higher-level cognitive processes that extend beyond mere motor control. Goal maintenance keeps intended outcomes in focus, ensuring actions remain aligned with relevant objectives even when distractions or competing demands arise (Badre & Nee, 2018; Miller & Cohen, 2001). Decision-making draws on internal simulations, where forward models predict the consequences of one's actions and the environment's likely responses.(Pezzulo & Cisek, 2016). By predicting how different actions unfold, individuals evaluate alternatives and select the course of action most likely to achieve their goals. This process supports proactive planning that enables individuals to anticipate challenges and adjust strategies before errors occur (Friston, 2010). Flexible adaptation is the ability to revise plans in response to unexpected events or failed predictions, facilitating rapid adjustment to new conditions. These predictive control and error-driven learning processes construct internal models that link environmental cues to expected outcomes and simulate entire action sequences, helping humans maintain effective, goal-directed behavior in complex, unpredictable environments.

1.2 Error-driven learning

Although humans share common mechanisms of action control, learning takes place along individual trajectories. Each trajectory reflects the specific challenges encountered previously. Some environments foster rapid refinement of internal models, while others create difficulties that slow or redirect adaptation (Smith, Ghazizadeh, & Shadmehr, 2006; Diedrichsen, White, Newman, & Lally, 2010). As a result, skill acquisition is shaped by a history of successes, failures, and corrective adjustments, producing individual differences in predictive accuracy and flexibility (Krakauer & Mazzoni, 2011; Wu, Miyamoto, Castro, Ölveczky, & Smith, 2014).

To understand how humans learn to act effectively in dynamic environments, it is essential to model the underlying learning mechanisms. Our aim is therefore to formalize a computational cognitive model that captures general principles of individual learning in dynamic environments. Rather than reproducing each participant's individual learning trajectory, the cognitive model is supposed to reproduce the aggregate learning patterns observed across participants and to explain them through error-driven adaptation and uncertainty reduction. We test the cognitive model in an experimental environment that also supports controlled human studies and allows direct behavioral comparison.

2 Anonymized environment

We designed an experimental environment which we implemented as a video game in Python (Van Rossum, Guido & Drake, 2009) using the PyGame library (Shinners, 2011)[1]. Participants steer a spaceship that moves automatically downward through a space corridor bounded by walls. Horizontal movement is controlled via keyboard input: [Y] to move left, [M] to move right. The spaceship, however, remains centered on screen and the environment shifts around it to create the perception of continuous motion. We dubbed it the Dodge Asteroids environment (Heinrich, Österdiekhoff, Kopp, & Rußwinkel, 2025) and it runs at 60 FPS.

The key feature of this iteration of the Dodge Asteroids environment is the presence of drift sections, indicated by red bars on either side outside the walls. When the spaceship enters a drift section, it is pushed in the horizontal direction opposite to the indicator: a bar on the right causes leftward drift, and a bar on the left causes rightward drift. Participants cannot steer during these sections. Therefore, they are required to anticipate the drift when selecting an initial entry point into the drift section. This design emphasizes action planning under uncertainty and the ability to adjust strategies based on expected outcomes.

The environment further features obstacles, which are placed exclusively within drift sections. Participants' primary task is to avoid collisions with both obstacles and the corridor walls. Because horizontal drift cannot be corrected during drift sections, they need to anticipate the spaceship's trajectory and incorporate it into their action plan. This can then be used to select an appropriate entry point in order to navigate the drift section without crashing. Drift magnitude is constant within and across sections (corresponding to a horizontal displacement of 230 pixels over a fixed drift section length of 210 pixels), but the horizontal position of the indicator (the drift direction) and the locations of obstacles vary across drift sections. This creates a dynamic environment that challenges participants to constantly adjust their action plans.

By design, the Dodge Asteroids environment is well-suited to investigate continuous action control, specifically the integration of predictive planning, execution, and feedback-based updating. Participants must learn the relationship between visual indicators and environmental perturbations, continuously adapt entry points, and refine motor strategies based on the observed outcomes of previous actions. anonymous

While the task constraints are simple, the underlying challenge (anticipating consequences, adjusting plans, and refining strategies) reflects core principles of human action control. To formalize these processes and link them to empirical behavior, we implemented a cognitive model using the ACT-R architecture.

3 A cognitive model of human action control

We defined mechanisms that reflect short-term, error-based adaptation similar to human motor learning. We embedded the mechanisms within the higher-level cognitive architecture ACT-R (Anderson et al., 2004) to investigate how principles of motor adaptation, such as prediction error correction and uncertainty reduction, interact with the deliberate, goal-driven processes of human action control.

Cognitive architectures such as ACT-R provide a framework for implementing psychological theories as computational models that simulate task performance (Anderson & Lebiere, 2014). ACT-R is grounded in the Common Model of Cognition (Laird et al., 2017), which specifies core components of human cognition such as perception, memory, learning, and action. Within ACT-R, knowledge is represented symbolically as chunks in declarative memory and production rules in procedural memory. These symbolic structures govern task execution by specifying what the system knows and how it can act on that knowledge. At the same time, ACT-R uses subsymbolic mechanisms, variables that drive behavior internally. These include chunk activation, production utili-

---

[1] Git repository containing the experimental environment.

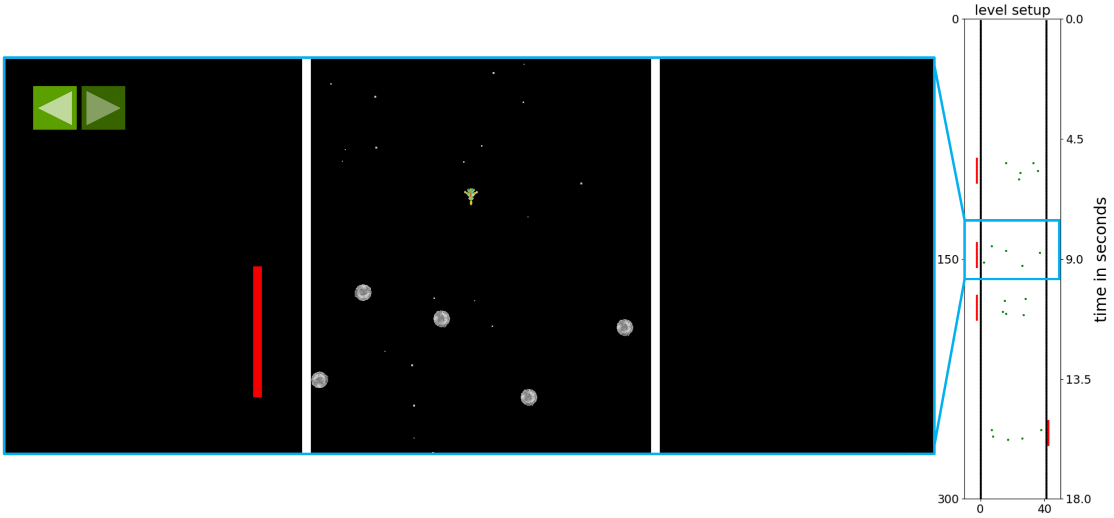

FIGURE 1 Instance within the Dodge Asteroids environment (based on the figure by Heinrich et al., 2023). The green spaceship is positioned at the horizontal center of the screen. Steering input is directed to the left, as indicated by the green arrow in the upper-left corner. The red bar indicates a drift section in which the spaceship is pushed to the right, with obstacles randomly placed within the section. The full layout of this trial is shown on the right.

ties, and associative strengths, and they are continuously updated during simulation. ACT-R further provides a set of adjustable parameters like retrieval latency factors, activation noise, or learning rates. These also determine how the model behaves in measurable terms but can be systematically varied. In a parameter estimation, the fit between model and human behavioral data can be assessed for each new set of parameter values. This way, we can identify a set of parameter values that best capture human performance while our model remains grounded in cognitive theory.

The computational cognitive model is designed to operate within the previously introduced Dodge Asteroids environment. To do so, it incorporates task-specific functions for perceiving the spaceship, corridor walls, obstacles, and drift indicators, as well as a drift-model stored as a declarative chunk within ACT-R. For clarity, the combination of ACT-R, the drift chunk, and associated perceptual mechanisms is referred to as the *cognitive model* throughout this work.

3.1 Model structure

The cognitive model solves the task in three stages: (1) perceiving the drift indicator and obstacle layout, (2) predicting outcome of possible actions and action selection, and (3) observing outcome of selected action and error-driven learning.

Perception of the drift indicator and obstacles The visual module encodes two aspects of the drift section: the drift indicator and the locations of obstacles.

The drift indicator's location (on which side of the corridor the red bar is displayed) is encoded as a signed value I, stored in the goal buffer. Positive values indicate the bar is on the right, negative on the left. Because drift always occurs in the opposite direction of the indicator, the model learns a negative slope parameter a.

Obstacle perception is implemented through a coarse spatial encoding of the drift section. The corridor below the spaceship is divided into $N_{\text{bins}} = 102$ kernels, each aggregating several pixels. For each bin, the model computes the proportion of obstacle pixels; if this proportion exceeds a threshold of 7%, the bin is marked as occupied. The threshold is deliberately chosen to be low enough to capture even small corners of obstacles that fall within a kernel, ensuring that potentially relevant features are not overlooked. This discretization reduces visual complexity while still preserving essential spatial information about free versus occupied regions of the corridor. At the same time, the bin-based representation aligns with ACT-R's principle of chunked, symbolic encoding: each bin can be mapped to a discrete symbol ("_" for free, "o" for occupied), allowing spatial layouts to be processed as symbolic sequences while remaining grounded in the underlying pixel environment. The resulting symbolic description of the corridor is used in the next stage to predict the outcomes of possible actions. The choice of $N_{\text{bins}} = 102$ provides sufficient granularity for representing the environment in this task. Finally, both N_{bins} and the threshold are fixed parameters in the present work and are not subject to model fitting.

Action outcome prediction and action selection Initially, the mapping between the visual drift indicator (red bar) and the resulting drift displacement is highly inaccurate. During the task, this relationship is gradually learned through repeated observation and prediction errors. The learned mapping is stored in the drift-model chunk as a simple relational schema (e.g., “indicator-left → drift-right”), which informs subsequent predictions ($\hat{d}$) and action selection.

When a drift indicator is detected in the visual buffer, a production matching on that visual input issues a retrieval request for the corresponding drift-model chunk from declarative memory. This chunk stores the current parameter estimates: slope a and intercept b, which are then used to compute the predicted drift displacement in pixels $\hat{d}$ in the goal buffer:

$$\hat{d} = a * I + b \tag{1}$$

where a captures the strength and direction of the observed drift indicator I, and b accounts for any baseline offset in cases in which there is no drift ($I = 0$).

Action selection proceeds via a forward-simulation mechanism. Upon detecting a drift indicator, the cognitive model first determines a set of candidate starting positions, or entry points $x_{\text{entry}}^{(1\to n)}$. For this, it attempts to retrieve a past instance from declarative memory that matches the current discretized obstacle layout. If no sufficiently similar instance is retrieved, or if the retrieved instance previously resulted in a crash, candidate entry points are sampled from a uniform distribution within the feasible range. This range is derived from the drift model: based on the predicted displacement across the section, regions near the wall opposite to the drift indicator are excluded, as these would likely lead to collisions given the model’s current understanding of drift dynamics.

If a successfully solved instance is retrieved, its stored entry point $x_{\text{entry,recalled}}$ serves as the mean of a normal distribution from which new candidates are drawn. The standard deviation of this distribution increases with the model’s current uncertainty σ (Eq. 6). The number of sampled candidates also scales with σ, ranging from $n = 10$ under low uncertainty to $n = 3$ under high uncertainty.

For each of the n candidate entry points, the cognitive model simulates the expected trajectory at the pixel scale and overlays it with the bin-based obstacle representation to detect potential collisions. Collision risk is defined as the number of occupied bins intersected by the predicted trajectory (each bin counting as a single collision regardless of how many pixels are intersected). The trajectory with the lowest predicted collision risk is selected.

This mechanism combines two forms of decision behavior: uninformed exploration through uniform sampling when the model cannot draw on successfully solved instances from the past, and uncertainty-driven variability through normal sampling when reliable experience is available. The latter is consistent with empirical findings showing that higher target uncertainty leads to greater movement variability in reaching and action-selection tasks (Krüger & Hermsdörfer, 2019; Enachescu, Schrater, Schaal, & Christopoulos, 2021; Wolpert & Landy, 2012).

However, if σ exceeds a certain threshold (e.g., $\sigma > 0.8$), the model falls back on a heuristic strategy that bypasses forward simulation and applies a conservative safety margin k to arrive at a final x_{entry}:

$$x_{entry} = -\hat{d} \pm margin, \qquad margin = k * \sqrt{\sigma} \tag{2}$$

Here, the minus sign before the predicted drift displacement $\hat{d}$ indicates that the cognitive model compensates in the opposite direction of the predicted drift. Because the prediction is highly uncertain ($\sigma > 0.8$), the cognitive model will add a safety margin that is based on the scaling parameter k and $\sqrt{\sigma}$, which is the standard deviation of the drift displacement in pixels. This way the margin grows proportionally to how variable the predictions are.

The fallback mechanism can be understood as a heuristic strategy for decision-making under uncertainty (Tversky & Kahneman, 1974; Mousavi & Gigerenzer, 2014, 2017). When predictive confidence is low, i.e. uncertainty is high ($\sigma > 0.8$), the model suspends fine-grained optimization and instead applies a simple spatial rule based on a fixed safety margin k. This reflects the human tendency to rely on fast and frugal heuristics when probabilistic information is unreliable, and favor easily computable actions over precise adjustments. Conceptually, the safety margin k functions as an *anchor* (Tversky & Kahneman, 1974), biasing entry-point selection toward the corridor center when k is small or toward the drift indicator when k is large.

Because the fallback is only triggered under high uncertainty, it primarily operates before a stable mapping between indicators and drift displacements has been learned. Although this initially slows precise adaptation, it captures the idea that heuristics can serve as adaptive tools for coping with uncertainty rather than as decision biases.

Once an entry point is selected, the cognitive model steers toward it by pressing the corresponding key to reduce the horizontal distance (Y for left, M for right). When no drift section is visible, the model maintains a central position by steering toward the horizontal center of the screen.

Action outcome observation and error-driven learning After the drift section (or in case of a crash), the model observes the actual drift displacement d_{obs} and computes the prediction error:

$$pe = d_{obs} - \hat{d} \tag{3}$$

If the trajectory is interrupted by a crash before the drift section is completed, $\hat{d}$ is proportionally scaled to the fraction of the drift section traversed at the point of impact. This ensures that the prediction error pe is still computed on a comparable basis.

Parameters of the drift-model are then updated using a local delta rule weighted by uncertainty:

$$a \leftarrow a + \eta * pe * I, \qquad b \leftarrow b + \eta * pe \tag{4}$$

with:

$$\eta = \frac{\eta_0}{1+\sigma} \tag{5}$$

By defining learning rate η this way, the base rate η_0 primarily drives the learning speed under low uncertainty, while increasing σ progressively slows adaptation so that uncertainty dominates when it is high. Finally, uncertainty σ is adjusted heuristically:

$$\sigma \leftarrow \max(\sigma_{min}, 0.9\sigma + 0.1pe^2) \tag{6}$$

Large prediction errors maintain or even increase uncertainty, keeping the model cautious and encouraging continued reliance on safety margins or conservative action selection. Consistently small errors, on the other hand, reduce σ, shrinking both the margin and the effective learning rate, which reflects confidence calibration during skill acquisition. The parameter σ_{min} sets a lower bound on uncertainty, preventing it from reaching zero and ensuring that the model always retains some minimal caution. For an overview of all parameters of the cognitive model and their meanings, see Table 1.

In addition to updating drift-model parameters, the model encodes each drift section instance as a chunk in ACT-R's declarative memory, storing the discretized kernel representation, the entry point, and whether the drift section was solved without a crash. These chunks are retrieved based on kernel similarity when the model encounters comparable drift sections. This instance-based mechanism complements the continuous drift model: error-driven learning provides a generalized mapping from indicators to displacements for predicting trajectories, while declarative memory stores specific past experiences with their associated confidence/uncertainty levels. The cognitive model requires both mechanisms to learn efficient action selection.

See Figure 2 for a schematic overview of the cognitive model.

3.2 ACT-R implementation details

The mapping between I and drift is stored as a single declarative chunk (drift-model) in declarative memory, containing slots for a, b, σ, and k. When perceiving a drift indicator, the cognitive model retrieves this chunk, performs computations in the goal buffer, and then creates a new drift-model chunk with the updated parameter values.

The area on the screen beneath the spaceship, bounded by the corridor walls, is converted into a symbolic representation by binning. The area spans 714 pixels horizontally and 270 pixels vertically. It is subdivided into 17 bins arranged in 6 rows, with each standard bin covering 42 pixels horizontally and 45 pixels vertically. Binning is computed dynamically: if the screen width cannot be divided evenly, the remaining pixels are grouped into an additional column spanning fewer pixels horizontally. Forward simulations and safety margin calculations are implemented as production rules operating on these chunk values. Perception of the indicator and obstacles uses the visual buffer, while decisions and predictions occur in the goal buffer via productions.

The ACT-R model, implemented in Lisp, was integrated with the Python-based experimental environment via the py4cl library (Dudson, 2023), using ACT-R's RPC interface over a socket connection for real-time communication.

4 Experiment

4.1 Setup

We conducted an experiment designed to collect behavioral data for comparison with predictions from our computational cognitive model implemented in ACT-R. The experimental setup consisted of multiple levels, each containing four consecutive drift sections. Obstacles appeared only within these sections, requiring participants to anticipate drift effects when planning their entry points. All drift sections had the same length, but the indicator side (signaling drift direction) and the spacing between sections varied across levels.

Levels were generated procedurally with randomized obstacle positions, while ensuring at least one viable path through each section. A level was successfully completed when participants navigated all four drift sections and crossed a green finish line at the bottom of the corridor. On average, a level took 18.64 seconds (SD = 3.53) to complete, with slight variations due to rendering performance. If the spaceship collided with an obstacle or a wall, the level ended immediately, and a crash notification was displayed.

4.2 Procedure

Before starting the main experiment, participants completed a training level to familiarize themselves with the game mechanics and steering controls. The training level included several drift sec-

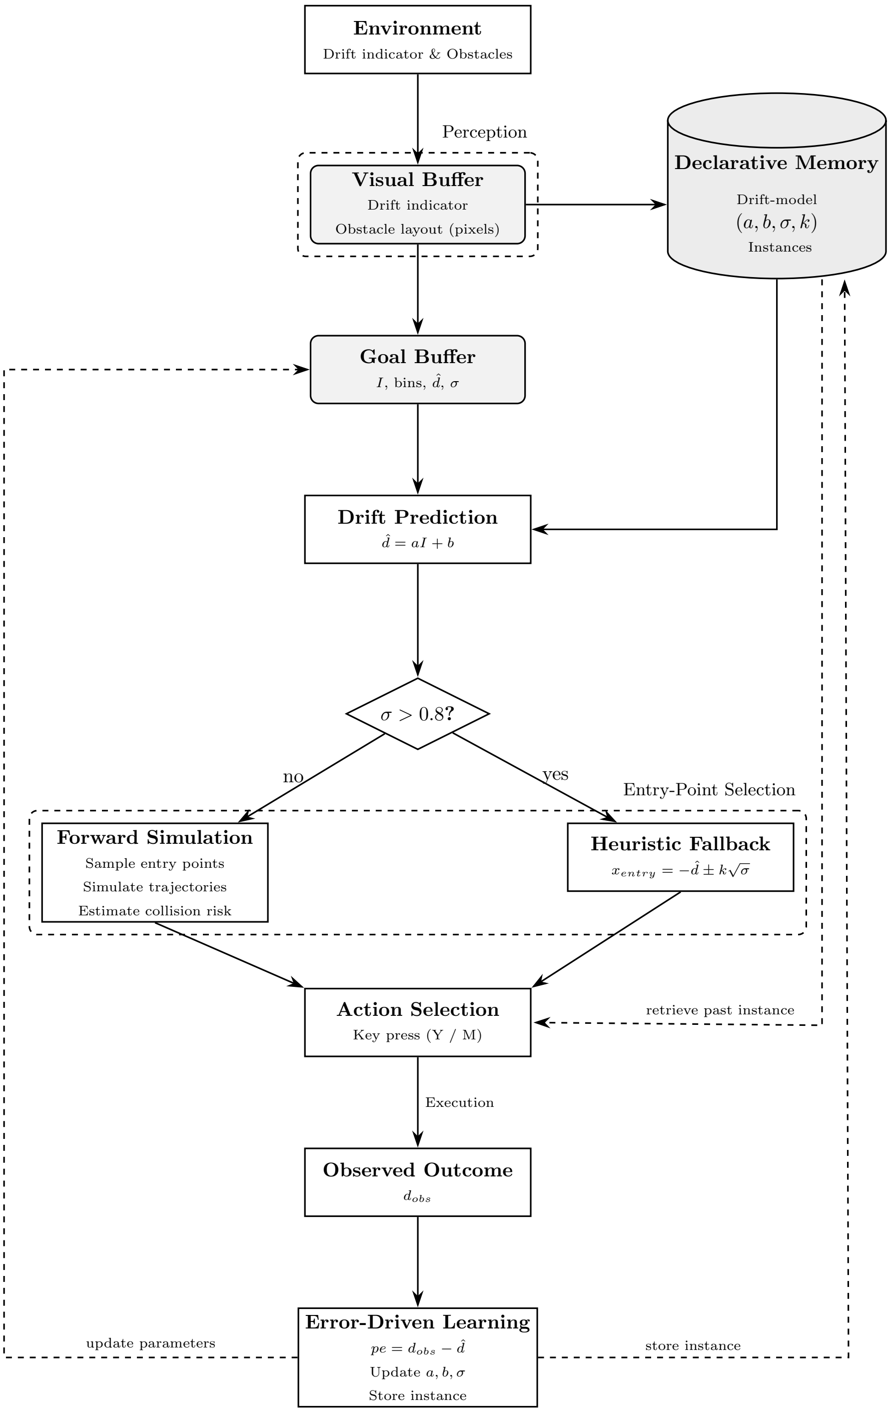

FIGURE 2 Schematic overview of the ACT-R cognitive model for the Dodge Asteroids task. ACT-R components are shown in light gray; solid arrows indicate the primary flow of online processing, and dashed arrows indicate learning- and memory-related processes. Visual information about the drift indicator and obstacles is encoded in the visual buffer and transformed into symbolic form in the goal buffer. Action selection proceeds via forward simulation of candidate trajectories or via a heuristic fallback under high predictive uncertainty. If a previously solved instance stored in declarative memory matches the current situation, its associated entry point is retrieved and used to guide entry-point selection. Following execution, the observed outcome is compared to the predicted displacement, generating a prediction error that drives parameter updates, uncertainty adjustments, and the storage of new instances.

Table 1 Model parameters and their functional roles in the ACT-R implementation.

Parameter	Symbol	Function
Slope	a	Encodes the strength and direction of the mapping between indicator position I and predicted drift displacement $\hat{d}$. Learned from prediction errors.
Intercept	b	Accounts for systematic baseline offsets in predicted drift when $I = 0$. Learned from prediction errors.
Indicator position	I	Horizontal position of the drift indicator (signed). Input to drift prediction (Eq. 1).
Safety margin scaling	k	Multiplies $\sqrt{\sigma}$ to determine conservative offset from predicted entry point (Eq. 2). Ensures risk-averse fallback under high uncertainty.
Entry point	x_{entry}	Horizontal starting position for drift section. Selected either via forward simulation of candidate trajectories (preferred) or by applying the safety margin fallback under high uncertainty.
Predicted displacement	$\hat{d}$	Model's prediction of drift displacement, based on a and b. Used to determine planned entry point for drift sections.
Observed displacement	d_{obs}	The actual displacement experienced in a drift section. Used to compute prediction error ($pe = d_{obs} - \hat{d}$).
Prediction error	pe	Difference between observed and predicted drift displacement (Eq. 3). Drives parameter updates and uncertainty adjustments.
Learning base rate	η_0	Sets the base speed of parameter updating. Effective learning rate η decreases with increasing σ, stabilizing learning under uncertainty.
Uncertainty	σ	Tracks reliability of predictions. High σ restricts forward simulations to fewer/narrower candidate entry points and widens safety margins. Low σ permits broader exploration and shrinks safety margins. Updated from squared prediction errors.

tions, initially without obstacles and later with obstacles added. If a participant crashed during training, the level was repeated until they successfully completed it once, ensuring they fully understood the task.

The main experiment consisted of 80 unique levels. If a participant failed a level, the same level was reintroduced later in the session with it being inserted at a random position in the level sequence. Each level could be attempted up to three times. If all three attempts ended in a crash, the level was removed from the sequence to prevent memorization of obstacle locations.

Before each level, participants rated their expected performance ("In the upcoming run, how good do you think you will perform?") on a 1–7 scale (1 = very poorly, 7 = excellent). We refer to this measure as expectancy. After completing a level, they rated their sense of control (SoC; "In the most recent run, how strong was your feeling of control?") using the same 1–7 scale. These ratings provided a subjective measure of performance confidence and perceived control.

The main part of the experiment took approximately 40–55 minutes, depending on individual performance. After completing all trials, participants filled out two questionnaires: (1) the German version of the NEO-FFI (30 items) to assess the Big Five personality traits (Openness, Conscientiousness, Extraversion, Agreeableness, Neuroticism), and (2) a short four-item survey on video game expertise, including play frequency and preferred game types (Nolan, 2022). From briefing to debriefing, the entire session lasted about 70 minutes, including time for questions.

Although participants provided expectancy, sense of control, and personality ratings, these measures are not analyzed here; however, we will discuss how they could help improve model fit or serve as a sanity check (see Section 8).

Prior to participation, all participants provided written informed consent and were informed about the procedures for data handling. Participation was voluntary, and participants could discontinue the experiment at any time without penalty and request withdrawal of their data. After completing the experiment, participants were debriefed regarding the purpose of the study and the comparison of their behavior with a computational cognitive model.

4.3 Recruiting and exclusion criteria

Data of 26 participants was collected from February 9, 2025, until May 12, 2025. Completing the training level within a reasonable number of attempts was an exclusion criterion, but all participants completed the training within a maximum of four attempts; therefore, no participants were excluded from the final analysis.

The training level first introduced steering the spaceship, followed by steering it past obstacles, then navigating a single drift section without obstacles, and finally a single drift section with only two obstacles to avoid. This design was intended to minimize participants' prior exposure to drift before the main experiment.

This study was conducted in accordance with the Declaration of Helsinki. It study was reviewed and approved by the Ethics Committee of the Department of Psychology and Ergonomics (IPA) of the Technische Universität Berlin.

4.4 Measures

We assessed the following summary statistics:

Trial success rate A binary variable ($\texttt{done} \in 0,1$) indicating whether a participant completed the trial without crashing. This measure was used to estimate overall success rates and to model the probability of trial completion across time.

Drift section success rate The number of drift sections successfully managed within a trial ($N_{\text{drift}} \in 0,1,2,3,4$). This measure captures partial progress within trials and provides a finer-grained metric compared to overall trial-based success rate.

Center tendency at drift onset For each drift section, we extracted the first frame in which drift was applied. We then computed the horizontal distance between the spaceship's position at the screen center ($x = 430$ pixels) and the corridor center. Distances were corrected for drift direction, such that negative values indicate deviations against the drift direction and positive values indicate deviations in the drift direction. This measure reflects participants' anticipatory positioning strategy when entering drift sections.

4.5 Experiment results

The present behavioral data has not been analyzed in this form or published elsewhere. However, subjective reports of expectancy and sense of control from the same experiment were analyzed in a separate study examining sequential Bayesian updating processes (Heinrich, von Engelhardt, Österdiekhoff, Kopp, & Russwinkel, 2025).

Across participants, the average maximum trial was 186.04 (SD = 16.69), with a range from 147 to 230 trials.

Participants varied in how many drift sections they successfully completed within a trial. Across all trials, success rates were distributed as follows: 0 sections completed - 29.2%, 1 section - 24.2%, 2 sections - 13.2%, 3 sections - 9.9%, and 4 sections - 23.4%.

We computed a logistic regression predicting the probability of successfully completing a trial ($\texttt{done} = 1$) from trial number. It indicated a significant improvement over time, $\beta = 0.0017, \sigma = 0.001, z = 2.738, p = .006$, suggesting that participants became more likely to complete entire trials as the experiment progressed. Similarly, an OLS regression of the number of successfully completed drift sections (N_{drift}) on trial number revealed a positive slope, $\beta = 0.0012, \sigma < 0.001, t = 3.001, p = .003$, confirming gradual performance improvements across trials even for individual drift sections.

To examine horizontal positioning at the onset of drift sections, we selected the first frame of each drift section, yielding one observation per drift. We then computed the corrected distance between the spaceship's entry point and the horizontal center of the corridor, such that negative values reflected deviations against the drift direction and positive values reflected deviations in the drift direction.

The mean corrected distance was highly similar for leftward drifts (M = –136.19 px, SD = 101.29) and rightward drifts (M = –135.38 px, SD = 100.78). A two-sample t-test confirmed no significant difference between the two drift directions, $t(12{,}178.71) = 0.09, p = 0.93$.

We next examined how deviations from the horizontal center evolved over the course of the experiment. For each drift section within a trial, we computed the absolute distance between the participant's horizontal position and the screen center at drift onset (up to four observations per trial). The trajectory of these values across trials was modeled using piecewise linear regression (implemented with the `pwlf` Python package)[2]. This method estimates linear slopes and change points (breakpoints) in the time course. A four-segment model (three breakpoints) was fitted. The resulting fit indicated breakpoints at trials 9.00, 139.63, and 142.00, suggesting distinct phases in participants' center tendency. Because the residuals may not be normally distributed due to the ordinal scale, we checked for robustness using a Box–Cox transformation of the distance data ($\lambda = 0.69$) and computed an additional piecewise linear regression of the transformed data. It yielded breakpoints at 9.00, 134.30, and 159.78, which we consider sufficiently similar; thus, all reported values refer to the original (untransformed) data breakpoints.

The first breakpoint at $t = 9$ reflects the initial

---

[2] We used a fixed random seed (36) throughout the work to ensure reproducibility.

learning phase, in which participants rapidly acquired an efficient center offset strategy. After approximately nine trials, performance reached an asymptote, consistent with fast early learning.

The second breakpoint at $t \approx 140$ likely reflects a qualitative change in behavior. Toward the end of the experiment, participants re-encountered unsolved drift constellations, which were presented for a third time. At this stage, participants may have adopted a more exploratory strategy, deliberately steering further from the center to test alternative approaches. The subsequent breakpoint at $t \approx 142$ supports this interpretation, marking the onset of a late exploratory phase rather than a decline in performance. It should be noted, however, that this third breakpoint may partly reflect the methodological choice to fit a four-segment piecewise linear model, which necessarily imposes additional breakpoints; the presence of this breakpoint could therefore be influenced by the fitting procedure rather than purely by behavioral changes.

Since our computational model is designed to capture initial learning but not explicit late-phase exploration, we restricted model fitting to human performance up to the second breakpoint (i.e., the first 140 trials). This ensures that the model reflects early acquisition of efficient strategies while excluding exploratory behaviors outside its scope.

5 Parameter inference

We inferred the values of several parameters of the model to align model performance with human success rates and learning time. We fitted ACT-R-specific parameters that influence decision-making, variability, and learning dynamics. First, we adjusted declarative memory parameters: activation noise (`ans`), which introduces variability in retrieving learned drift mappings, and the mismatch penalty (`mp`), which biases retrieval toward contextually appropriate chunks. These parameters determine how reliably the model recalls prior experiences when selecting an entry point. Second, we tuned production system parameters, including utility noise (`egs`) and the utility learning rate (`alpha`), which control the degree of exploratory versus conservative action selection and the speed at which successful strategies are reinforced. We further fitted cognitive parameters specific to the task at hand: the safety margin k, the error-based adaptation parameter η_0, and uncertainty parameters σ_0 and σ_{min}. k drives how conservatively the model positions the spaceship toward the corridor center under uncertainty. A larger safety margin reflects commitment-averse planning meant to increase feedback control. η_0 determines how fast internal drift mapping estimates (a and b) are updated after observing prediction errors (Equation 5). This parameter directly influences the model's learning curve and its ability to refine forward predictions across trials. Finally, σ_0 is the initial uncertainty linked to the drift-model, whereas σ_{min} is the floor bound of uncertainty σ (Equation 6). This selection of parameters captures both variability in action selection and the adaptive processes that allow the cognitive model to improve over time.

In order to evaluate how well the cognitive model accounts for human behavior, we fit model parameters θ (vector of values of all parameters-to-be-fitted in fixed order) such that simulated performance matches human data in two aspects: (1) the distribution of success outcomes across trials, and (2) the time course of learning, reflected in how quickly performance approaches its asymptote (reflected by the first breakpoint found in human data). These two targets were combined into a single (weighted) objective function that penalizes mismatches in both outcome distributions and learning-time estimates.

5.1 Data preprocessing

Trial index served as the temporal axis for estimating learning dynamics. Each trial t provided an outcome category $cat \in \{0,1,2,3,4\}$, reflecting the number of drift sections successfully managed. For modeling learning curves, we defined trial-level success values $success \in [0,4]$ directly from these outcomes, without temporal binning. Finally, only trials up to $t = 140$ (second breakpoint) were included in the estimation to focus on early learning dynamics, excluding later exploratory behavior.

5.2 Simulating model outcomes

For each candidate parameter vector θ, the model was simulated across a randomized trial sequence (analogous to the randomization across human participants), and this simulation was repeated $R = 50$ times to stabilize outcome estimates.

For each repetition, we recorded the number of drift sections successfully completed (categories $cat = 0,\dots,4$). The model-predicted probability of each outcome category was then estimated as the proportion of simulation runs in which that category occurred - yielding the *outcome distribution*:

$$p^{\text{mod}}_{cat}(\theta) \approx \frac{1}{R}\sum_{r=1}^{R} \mathbf{1}\{\text{outcome in run } r = cat\}. \quad (7)$$

To avoid numerical issues from zero or extremely small probabilities, a Laplace smoothing constant of $+0.5$ was applied to all counts before computing likelihoods.

5.3 Multinomial negative log-likelihood

To assess the fit of the model to the observed human outcome distributions, we computed the multinomial negative log-likelihood (*NLL*) of the aggregated human outcome distribution given the model's predicted outcome probabilities:

$$\mathcal{L}_{\mathrm{NLL}}(\theta) = -\sum_{cat=0}^{4} n_{cat}^{\mathrm{human}} \log\left(p_{cat}^{\mathrm{model}}(\theta)\right). \quad (8)$$

This formulation assesses how well the model reproduces the overall pattern of human performance across drift sections.

5.4 Learning-time mismatch

In addition to the outcome distributions, the cognitive model is evaluated on the basis of learning dynamics.

The time course of learning was modeled using an exponential approach-to-asymptote function meant to capture typical shapes of learning curves observed in skill acquisition (rapid early learning and slow improvement over time):

$$S(t) = L - G\exp(-t/\tau), \quad (9)$$

where L represents the asymptotic performance level, G the initial gap from the asymptote, and τ the learning time constant expressed in trial number t. This function was fit to the human trial-level success data using nonlinear least-squares optimization (`scipy.optimize.curve_fit`). Initial guesses and bounds were deliberately set as follows: $L_0 = \max(success)$ with bounds $L \in [0, 4.5]$, $G_0 = L_0 - \min(success)$ with bounds $G \in [0, 4.5]$, and $\tau_0 = 1$ with bounds $\tau \in [0.1, 50]$. The choice of $\tau_0 = 1$ reflects a conservative prior for rapid learning, while wide bounds allow the optimizer to capture slower learning if present. Before parameter inference, we estimated the human learning time constant $\hat{\tau}_{\mathrm{human}}$ with variability assessed by bootstrap resampling of participants.

The exponential function was fitted to the model's trial-level success curve for each different θ, yielding $\hat{\tau}_{\mathrm{model}}(\theta)$. We then computed a normalized penalty for learning-time mismatch:

$$\mathcal{L}_{\tau}(\theta) = \frac{\left(\hat{\tau}_{\mathrm{model}}(\theta) - \hat{\tau}_{\mathrm{human}}\right)^2}{\mathrm{Var}(\hat{\tau}_{\mathrm{human}})}. \quad (10)$$

5.5 Combined loss function

The overall objective function combined both terms in a weighted sum:

$$\mathcal{L}(\theta) = \mathcal{L}_{\mathrm{NLL}}(\theta) + w\mathcal{L}_{\tau}(\theta), \quad (11)$$

where w balances the relative contribution of the two components (with $w < 1$, weighting $\mathcal{L}_{\mathrm{NLL}}(\theta)$ more heavily). We set $w = 1$ by default, as in this work, we place equal emphasis on final performance and learning time.

5.6 Parameters and optimization

The parameter vector θ included ACT-R parameters affecting success and learning dynamics: activation noise (`ans`), mismatch penalty (`mp`), utility noise (`egs`), learning rate (`alpha`), and task-specific cognitive parameters: learning base-rate η_0, safety margin k, the initial uncertainty σ_0, and its minimum $\sigma_{\min}$. Bounds were deliberately set as follows: `ans` $\in [0.01, 0.5]$, `mp` $\in [0, 4]$, `egs` $\in [0.01, 1.0]$, `alpha` $\in [0.01, 0.5]$, $\eta_0 \in [0.01, 0.5]$, $k \in [0, 300]$, $\sigma_0 \in [0.1, 10.0]$, and $\sigma_{\min} \in [0.1, 10.0]$.

Optimization proceeded in two stages. First, we performed a global search using a space-filling design (Sobol sequences) to evaluate $\sim$1000 candidate parameter sets at reduced replication ($R = 10$) for efficiency. Second, the best 117 candidates were explored more thoroughly using a local optimizer (Bayesian optimization) with full replication of $R = 50$.

The cognitive model is supposed to reflect the learning trajectory up to the second breakpoint in the human data, as well as the resulting performance after the initial learning (corresponding to the first two components of the piecewise linear regression). We therefore excluded all trials beyond trial 140 before splitting the data into training and testing sets. Model behavior during parameter inference was simulated on a training set comprised of randomly sampled 75% of the truncated trial data across all levels and participants. The remaining 25% compose the testing set that is used to validate the final fit to human behavior.

6 Model results

The cognitive model was fitted to the training set of human behavioral data using maximum-likelihood estimation of key ACT-R and task-specific parameters. The best-fitting parameter values were as follows: activation noise (`ans`) = 0.09, mismatch penalty (`mp`) = 0.91, utility noise (`egs`) = 0.02, and utility learning rate (`alpha`) = 0.08. Regarding the task-specific parameters, the base learning rate was estimated as $\eta_0 = 0.36$, the safety margin as $k = 109$, initial uncertainty as $\sigma_0 = 2.0$, and minimum uncertainty as $\sigma_{\min} = 0.3$.

The model reproduced the overall variability in task performance observed in humans. The distribution of outcomes indicated that 29.6% of trials resulted in zero sections completed, 23.7% in one section, 16.4% in two sections, 12.6% in three sections, and 17.7% in four sections completed successfully.

A logistic regression predicting trial completion (i.e., success vs. failure) from trial number revealed a significant increase in success over time, $\beta = 0.0049$, $\sigma < 0.001$, $z = 7.486$, $p < 0.001$. This pattern indicates that the model's task performance improved with task progression, consistent with learning observed in the training set. The number of correctly solved drift sections also increased significantly over trials, $\beta = 0.0027$, $\sigma < 0.001$, $z = 7.404$, $p < 0.001$, confirming that the model gradually acquired the mapping between drift indicators and displacements.

The mean corrected horizontal position of the spaceship was comparable for leftward drifts (M = –103.31 px, SD = 69.82) and rightward drifts (M

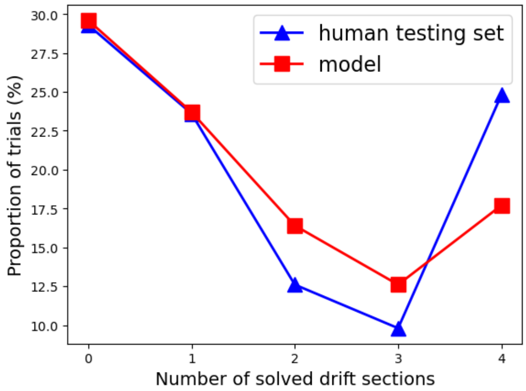

FIGURE 3 Proportion of trials (%) as a function of the number of solved drift sections for the cognitive model (red) and human testing set (blue). $D_{KL}(human||model) = 0.0300$ bits (with base=2).

= −102.74 px, SD = 69.94). A two-sample t-test confirmed no significant difference between drift directions, $t(19{,}120.83) = 0.56$, $p = 0.57$, indicating that the model did not learn different horizontal positioning strategies for the two drift directions.

We again used a piecewise linear regression to characterize the learning trajectory over the first 140 trials (the range used for model fitting). We computed a two-segment model with a single breakpoint reflecting reaching an asymptote in center tendency. The procedure identified a transition point at trial ≈ 44.

6.1 Model validation

Model predictions based on the fitted parameters were compared against the testing set of human data (25% of total trials). With this, we assessed the model's ability to generalize beyond training data.

The distribution of the number of completed drift-section in the testing data (0: 29.3%, 1: 23.6%, 2: 12.6%, 3: 9.8%, 4: 24.8%) was broadly similar in shape to that of the model on the training set. We computed the relative entropy (Kullback–Leibler divergence) using scipy.stats.entropy (base=2) to assess the overall agreement between the simulated data and the human test, which yielded a $D_{KL}(human||model)$ of 0.0300 bits. While the KL divergence reflects a high agreement, the model slightly overestimated the proportion of intermediate outcomes (one to three sections) and underestimated the frequency of fully successful trials (see Figure 3).

A logistic regression predicting trial completion from trial number in the testing data revealed a significant positive learning trend ($\beta = 0.0037$, $\sigma = 0.002$, $z = 1.981$, $p = 0.048$), indicating modest but reliable improvement over time. The corresponding effect for the number of successfully solved drift sections was also positive ($\beta = 0.0033$, $\sigma = 0.001$, $t = 2.643$, $p = 0.008$). These trends are consistent in direction with the model's learning behavior, but considerably weaker in magnitude, suggesting that the model may have needed to learn more efficiently than humans to solve the task, as also indicated by the comparison of the piecewise linear regression below.

The human testing set, like the model, showed no difference in horizontal positioning between drift directions ($t(2{,}292.53) = 0.92$, $p = 0.36$).

While the model was fitted based on learning-time mismatch in performance (performance over trials; see section 5.4), in the validation we assessed learning in terms of center tendency, i.e., when the agent settles on a sufficiently effective horizontal positioning strategy. With this, we aim to validate the model on a more behaviorally informative measure of learning.

The piecewise linear regression on the testing data identified an early breakpoint at ≈ 9 trials, in contrast to the model's later breakpoint at ≈ 44 trials. This indicates that humans required substantially fewer trials to reach an efficient strategy for horizontal positioning. Together with the outcome distribution, which already reflected relatively high success rates, this suggests that human participants entered the task with an approximately effective strategy and only fine-tuned their behavior over the course of 9 trials. More specifically, humans started with relatively small horizontal distances to the corridor center and gradually increased these distances until reaching a stable positioning strategy, whereas the model began with large offsets and learned to reduce them over many trials (most likely caused by k being fit to 109; see Figure 4).

The fitted parameter values indicate that the model required relatively low stochastic variability during retrieval and decision-making (activation noise = 0.03; utility noise = 0.02), suggesting stable memory access and consistent choice behavior across trials. The moderate mismatch penalty (`mp` = 0.91) implies a fairly narrow retrieval tolerance, promoting the selection of closely matching declarative chunks and thus more deterministic performance once learning stabilized. The estimated utility learning rate (`alpha` = 0.08) reflects gradual but steady adjustment of production utilities, consistent with incremental reinforcement-based adaptation rather than abrupt strategy shifts.

Within the task-specific learning module, the relatively high base learning rate ($\eta_0 = 0.36$) suggests that updates of the drift model following prediction errors were fairly rapid. Further, the initial uncertainty of $\sigma_0 = 2.0$ implies that the model began with an already considerable precision in its drift predictions. Over time, uncertainty decreased toward the lower bound ($\sigma_{\min} = 0.3$).

The combination of a comparatively low initial uncertainty, a high base learning rate, and a tight lower bound (relative to the fallback threshold of

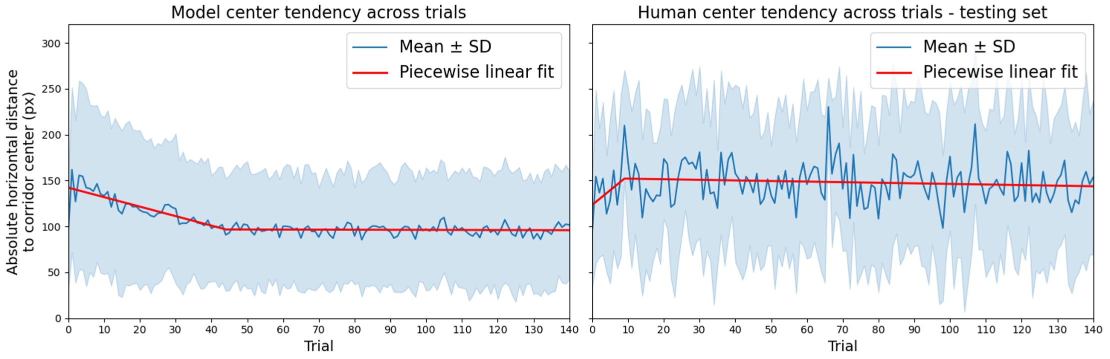

FIGURE 4 Mean absolute horizontal distance to the corridor center (px) across trials for the cognitive model (left) and the human testing set (right). Blue lines indicate the mean with shaded areas representing ±1 SD. Red lines show the two-segment piecewise linear regression fit.

$\sigma > 0.8$) likely produced the early learning curve observed here, possibly compensating for inefficient learning during the initial phase of the experiment when fallback activation was frequent and prevented error-driven adaptation.

Taken together, the validation results suggest that while the model captures the overall structure of human performance, its early learning dynamics differ substantially. One potential reason may be the model's fallback mechanism, which likely limited effective learning during early trials. In the next step, we will therefore fit and evaluate a revised model without this mechanism.

7 Model revision and reevaluation

The original model incorporated a safety-margin fallback intended to promote cautious entry-point selection under high uncertainty. However, inspection of the fitted parameters revealed that this mechanism did not lead to more conservative positioning. Instead, it produced entry points that were largely independent of the actual obstacle constellations, occasionally resulting in inefficient or counterproductive strategies. Because that fallback-driven entry points are not grounded in the specific instance, storing and recalling them from memory offered little benefit and likely hindered effective learning, particularly early in the task. We therefore removed the safety-margin mechanism and refitted the model to the human training data, excluding the safety-margin parameter k while keeping all other parameters and fitting procedures unchanged. The revised model now selects entry points only based on observed obstacle constellations and drift predictions.

The revised model was refitted on the human training data using the same optimization procedure. The best-fitting parameters were: activation noise (`ans`) = 0.03, mismatch penalty (`mp`) = 1.2, utility noise (`egs`) = 0.07, utility learning rate (`alpha`) = 0.23, $\eta_0 = 0.31$, $\sigma_0 = 1.8$, and $\sigma_{\min} = 0.2$, with the safety-margin parameter k omitted.

The revised version produced outcome distributions that more closely resemble those of human participants (0 sections: 36.6%, 1 section: 21.3%, 2 sections: 14.2%, 3 sections: 9.0%, 4 sections: 19.0%), yielding a slightly improved KL divergence of $D_{KL}(human||model_{without\,fallback})$ =0.0265 bits (compared to 0.0300 for the previous model). Although the model somewhat overfitted the proportion of trials with zero sections completed, it successfully reproduced the sharp decline across intermediate categories (1–3 sections) and the subsequent rebound for fully completed trials (4 sections; see Figure 5).

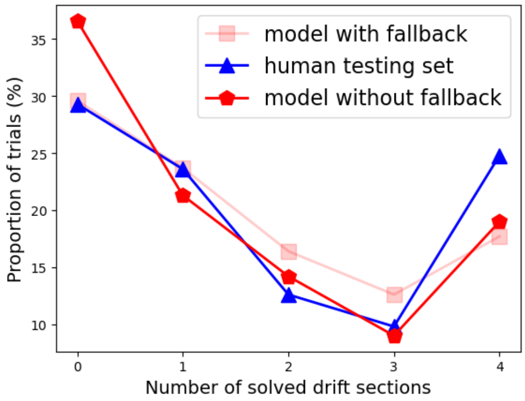

FIGURE 5 Proportion of trials (%) as a function of the number of solved drift sections for the *revised* cognitive model (red) compared with the original model (transparent red) and the human testing set (blue). $D_{KL}(human||model_{without\,fallback}) = 0.0265$ bits (with base=2).

Performance, again, improved significantly across trials, both in terms of binary success ($\beta = 0.0049$, $\sigma = 0.001$, $z = 7.40$, $p < .001$) and the number of successfully managed drift sections ($\beta = 0.0030$, $\sigma < 0.001$, $t = 7.88$, $p < .001$).

Also the revised model showed no difference in horizontal positioning between leftward and rightward drift sections ($t(17{,}362) = -0.10$, $p = 0.92$), replicating the symmetric spatial policy observed in human data. Piecewise linear regression identified a breakpoint at $t \approx 14$, indicating that the revised model reached a stable positioning strategy much earlier than in the original version (previously $t \approx 44$).

Importantly, the revised model now started with relatively small offsets from the corridor center and gradually increased them until reaching an efficient horizontal positioning strategy. The gradual-increase adaptation mirrored the human learning pattern (see Figure 6). This suggests that removing the fallback mechanism allowed the model to learn directly from task-relevant feedback, producing more instance-specific and efficient action selection.

8 Discussion

It is worth reflecting on the drift model and why we implemented it the way we did. It is a compact, low-dimensional summary of how the corridor moves over time. That the model required a correctly-specified linear representation of this dynamic to match human learning times is, in a sense, remarkable. As modelers, we pre-specified the degree of the polynomial for the drift model. Yet even with the pre-specified polynomial, learning was substantially slower than in humans. This suggests that the ability to form a parsimonious, actionable model of an environmental dynamic is central to the learning efficiency humans display. One might ask whether humans begin with a first-order approximation when encountering novel dynamics, increasing model complexity only when simpler representations fail to support adequate prediction. If so, the choice of model degree would itself constitute a cognitive heuristic: a prior over environmental structure that constrains and accelerates learning. How intelligent systems might acquire such priors autonomously, rather than inheriting them from their designers, remains an open challenge for cognitive modeling in dynamic environments.

The original model, while broadly capturing human performance trends, showed important discrepancies in its learning dynamics: (1) it required a substantially longer learning period, settling on a stable positioning strategy only after approximately 44 trials, and (2) it initially adapted in the opposite direction of human participants, starting with a large horizontal offset from the corridor center and reducing rather than increasing offsets. Both patterns can be traced to the heuristic fallback, which was designed to address the challenges of early learning where uncertainty about the environmental dynamic is highest and forward simulation is least reliable. By bypassing simulation under these conditions, it was intended to provide a safe default strategy. Instead, it introduced an action selection policy that was inefficient and poorly aligned with human behavior. The fallback applied an abstract safety margin independent of the actual spatial layout of obstacles, producing uninformative memory traces despite achieving reasonable overall performance in later stages. The behavioral data suggest that human participants do not operate this way. Even in the earliest trials, when uncertainty about the drift dynamic must have been substantial, participants' entry points reflect sensitivity to the observable structure of the environment, including the positions and arrangement of obstacles, rather than a feature-blind fallback. This is consistent with a broader body of evidence suggesting that human decision-making under uncertainty is guided by fast, ecologically grounded heuristics that exploit available perceptual information rather than suspending it (Gigerenzer & Todd, 1999; Gibson, 1979). The failure of the abstract fallback to capture this pattern reinforces the conclusion that uncertainty, in human cognition, modulates how the environment is processed, not whether it is processed at all.

Removing the fallback mechanism led to a clear improvement in model–human correspondence, particularly in the way the model adapted in early learning (now starting with small horizontal offsets and increasing over time) and the number of trials it required to settle on a positioning strategy (now 14 trials compared to 44 in the original model and 9 in humans.). This suggests that the heuristic safety margin, although conceptually motivated as a tool for decision-making under uncertainty (Mousavi & Gigerenzer, 2017), may have interfered with effective learning early in the task by generating uninformative memory traces. Without the fallback, recalled instances became more useful because they were grounded in the actual obstacle layout rather than in an abstract safety rule.

While the removal of the fallback mechanism improved model–human correspondence, it also revealed that the model may have been overdesigned in its initial form. The combination of forward simulation and heuristic safety margins with uncertainty-based gating created partially redundant control pathways. But instead of completely discarding the heuristic fallback, a more parsimonious account could be achieved by revising the candidate selection process. We envision a less discrete fallback rule. Basing the number and spatial spread of evaluated entry points on uncertainty already provides a natural way to represent heuristic decision making under uncertainty. Additionally, internal uncertainty could control how far into the future the cognitive simulation extends and thus reduce the scope of forward simulations in cases of high uncertainty. Much like humans committing to a decision that appears reasonable in the first few moments, but as the situation continues to unfold, a different choice would have served them better. At the same time, sampling biases other than the normal or uniform sampling policies de-

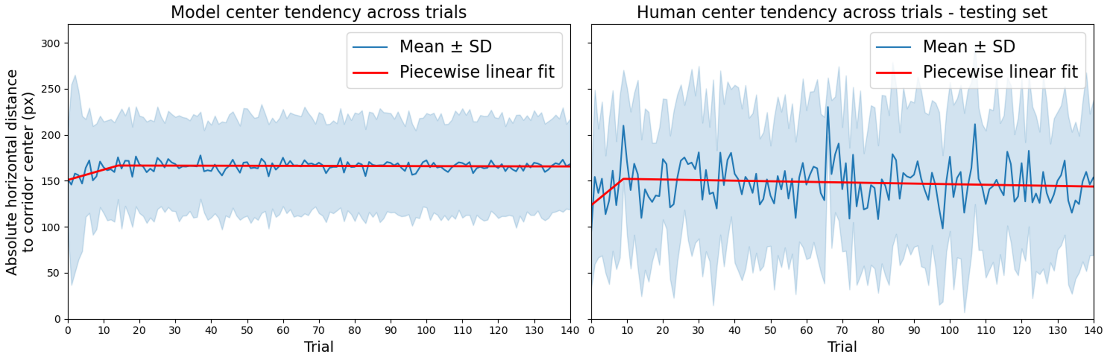

FIGURE 6 Mean absolute horizontal distance to the corridor center (px) across trials for the *revised* cognitive model (left) and the human testing set (right). Blue lines indicate the mean with shaded areas representing ±1 SD. Red lines show the two-segment piecewise linear regression fit.

scribed in section 3.1 could be explored. Both of these might lead to a simpler and cognitively more plausible way of dealing with uncertainty in early trials.

How can ACT-R models achieve human-like decision-making in dynamic environments without internal models of precisely the environmental dynamics they are meant to solve? The present model navigates this tension through a learned linear representation whose structure was prespecified by the modelers, which was necessary to match human learning times, as discussed above, but difficult to reconcile with a cognitively plausible account of how such representations arise. Predictive coding frameworks address this directly, proposing that the brain constructs and refines hierarchical generative models of the environment through continuous prediction error minimization (Friston, 2005; Clark, 2013). ACT-R's symbolic, modular architecture does not naturally accommodate this kind of graded, continuous updating (Anderson et al., 2004), yet the uncertainty-driven learning implemented here shares a functional resemblance with it. The tension between these accounts may be less about architecture and more about granularity: bounded rationality research suggests that humans often achieve adaptive behavior not through rich generative models but through fast, frugal heuristics that exploit environmental structure (Simon, 1955; Gigerenzer & Todd, 1999), the same principle that motivated the fallback, and whose failure here illustrates how difficult it is to specify such heuristics without grounding them in the very environmental knowledge they are meant to substitute.

Another point of discussion is that parameter inference in the present study relied on randomized trial sequences to approximate between-participant variability. While this approach captures average learning behavior, it may obscure the trial-specific adaptations observed in individual participants. A promising next step would be to fit the model directly to individual trial sequences. This would allow us to examine how parameter estimates reflect differences in learning rate, uncertainty reduction, and exploration strategies across participants. In addition, subjective measures such as expectancy and sense of control (SoC) could be compared to the model's internal variables to evaluate whether the model's continuous uncertainty estimates correspond to participants' self-reported experience of control.

A related limitation concerns the amount of prior exposure to the drift dynamic that participants brought to the experiment. Training was kept deliberately minimal, introducing drift only through a single section without obstacles and a single section with two obstacles (that may have been attempted several times to successfully solve the training level), to ensure that learning of the drift dynamic would unfold during the main experiment. However, perhaps participants did not have sufficient grounding in the environmental dynamic to develop the kind of compact, actionable representation the model assumes, or perhaps they had already formed a more accurate representation than the linear model captured at the beginning of the experiment. Future studies should therefore consider how training protocols in dynamic environments can be designed to give participants knowledge about the environmental dynamic without preempting the learning trajectory of interest. A simple verbal description of the drift mechanic prior to the experiment, for instance, might strike this balance by providing conceptual grounding without introducing motor or strategic experience.

9 Conclusion

We have designed a cognitive model of action control that learns to act in a custom dynamic task environment, the Dodge Asteroids environment. In this task, a spaceship is steered through a narrow corridor filled with obstacles that must be avoided (with the spaceship fixed at the screen center and the corridor moving around it). The environment features drift sections during which control over the spaceship is lost and it is pushed to one side, as indicated by a visual drift cue. Solving the task requires learning the drift dynamics and selecting entry points to drift sections that minimize the risk of collision. The cognitive model was implemented in ACT-R and evaluated against human participants. Model and human task behavior were compared in terms of final performance levels and the learning time required to develop an efficient strategy for selecting entry points.

In ACT-R, learning arises from the interaction between declarative and procedural memory. Declarative learning strengthens or weakens memory chunks based on retrieval success and frequency, while procedural learning adjusts production utilities and compiles frequently co-occurring rules into more efficient sequences that are selected based on expected gain rather than retrieved from memory, which speeds up the selection process and reduces its susceptibility to retrieval failures. In the present model, these mechanisms are complemented by the adaptive updating of drift-model parameters (a, b) and the associated uncertainty σ. The drift model is a linear abstraction of the environmental drift dynamic and the uncertainty term acts as a metacognitive moderator of learning. It controls how strongly prediction errors (when the spaceship's path deviates from what was anticipated) adjust the drift model and regulates exploration in action selection by modulating the number and spread of candidate entry points and whether a heuristic fallback strategy is used.

The heuristic fallback was implemented as a rule that applied an abstract safety margin to action selection when uncertainty about the drift dynamic was high, bypassing forward simulation and thus avoiding reliance on an imprecise drift representation. This rule, however, proved detrimental to early learning. By decoupling action selection from the actual spatial layout of the environment, it prevented the model from encountering meaningful prediction errors early on, causing it to take substantially longer to settle on an efficient entry point strategy than human participants. Removing the fallback rule improved model–human correspondence considerably. Without it, the model relied on forward simulation from the very first drift sections, and since its drift representation was initially imprecise, it produced prediction errors that drove rapid adjustment of the drift model parameters. These early errors proved essential for aligning the model's learning trajectory with that of human participants. This suggests that humans engage with the environmental dynamic from the earliest trials, rather than suspending it in favor of an abstract safety rule.

The revised cognitive model provides a more coherent account of early learning in the drift task. By integrating ACT-R's declarative–procedural learning mechanisms with a dynamic representation of uncertainty, it shows that human-like adaptation in continuous environments emerges not merely from rule-based learning, but from error-driven adjustments through the continuous interplay between memory retrieval, prediction, and uncertainty updating. These results imply a broader framework in which metacognitive control, implemented as uncertainty-guided exploration of action possibilities, may link symbolic cognitive architectures with predictive, error-based theories of learning.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used ChatGPT (OpenAI) and Claude (Anthropic) to assist with the drafting and paraphrasing of written content, as well as with the structuring and organization of the manuscript. These tools were not used to generate ideas or intellectual contributions. Claude Code (Anthropic) was used to support the implementation of the socket-based communication layer between the Python environment and ACT-R's RPC interface. All AI-assisted content was subsequently reviewed, thoroughly tested where applicable, and edited by the authors. The authors take full responsibility for the integrity and accuracy of the published work.

Acknowledgements

This research was funded by the German Research Foundation (DFG) Priority Program 2134 "The Active Self".

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Anderson, J. R., & Lebiere, C. J. (2014). *The atomic components of thought*. Psychology Press.

Badre, D., & Nee, D. E. (2018). Frontal cortex and the hierarchical control of behavior. *Trends in cognitive sciences*, *22*(2), 170–188.

Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204.

Diedrichsen, J., White, O., Newman, D., & Lally, N. (2010). Use-dependent and error-based learning of motor behaviors. *Journal of Neuroscience*, *30*(15), 5159–5166.

Dudson, B. (2023). *py4cl: Call python from common lisp.* (GitHub repository, `https://github.com/bendudson/py4cl`)

Enachescu, V., Schrater, P., Schaal, S., & Christopoulos, V. (2021). Action planning and control under uncertainty emerge through a desirability-driven competition between parallel encoding motor plans. *PLoS computational biology*, *17*(10), e1009429.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological sciences*, *360*(1456), 815–836.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, *11*(2), 127–138.

Gibson, J. J. (1979). *The ecological approach to visual perception* (Classic ed.). Psychology Press.

Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: the adaptive toolbox. In G. Gigerenzer & P. M. Todd (Eds.), *Simple heuristics that make us smart* (pp. 3–34). New York, NY: Oxford University Press.

Heinrich, N. W., Österdiekhoff, A., Kopp, S., & Rußwinkel, N. (2025). Using eye movements to understand sense of control in situated action. *Cognitive Science*, *49*(12), e70154. doi: 10.1111/cogs.70154

Heinrich, N. W., Russwinkel, N., Österdiekhoff, A., & Kopp, S. (2023). A Straightforward Implementation of Sensorimotor Abstraction in a Two-Layer Architecture for Dynamic Decision-Making. In *MathPsych/ICCM/EMPG 2023.*

Heinrich, N. W., von Engelhardt, R., Österdiekhoff, A., Kopp, S., & Russwinkel, N. (2025). Towards a bayesian cognitive model of self-belief updating. In *Proceedings of international conference on cognitive modeling (iccm 2025).*

Krakauer, J. W., & Mazzoni, P. (2011). Human sensorimotor learning: adaptation, skill, and beyond. *Current opinion in neurobiology*, *21*(4), 636–644.

Krüger, M., & Hermsdörfer, J. (2019). Target uncertainty during motor decision-making: The time course of movement variability reveals the effect of different sources of uncertainty on the control of reaching movements. *Frontiers in Psychology*, *10*, 41.

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *Ai Magazine*, *38*(4), 13–26.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, *24*(1), 167–202.

Mousavi, S., & Gigerenzer, G. (2014). Risk, uncertainty, and heuristics. *Journal of Business Research*, *67*(8), 1671–1678.

Mousavi, S., & Gigerenzer, G. (2017). Heuristics are tools for uncertainty. *Homo Oeconomicus*, *34*(4), 361–379.

Nolan, D. (2022). *Video survey.* `https://www.stat.berkeley.edu/~nolan/surveys/videoSurvey.pdf`. Retrieved from `https://www.stat.berkeley.edu/~nolan/surveys/videoSurvey.pdf` (Accessed: 2022-12-08)

Pezzulo, G., & Cisek, P. (2016). Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends in cognitive sciences*, *20*(6), 414–424.

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual review of neuroscience*, *33*(1), 89–108.

Shinners, P. (2011). *PyGame.* Retrieved 2022-08-29, from `http://pygame.org`

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118.

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS biology*, *4*(6), e179.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, *185*(4157), 1124–1131.

Van Rossum, Guido, & Drake, F. L. (2009). *Python 3 Reference Manual.* Scotts Valley, CA: CreateSpace.

Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature reviews neuroscience*, *12*(12), 739–751.

Wolpert, D. M., & Flanagan, J. R. (2001). Motor prediction. *Current biology*, *11*(18), R729–R732.

Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current opinion in neurobiology*, *22*(6), 996–1003.

Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P., & Smith, M. A. (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature neuroscience*, *17*(2), 312–321.

4.3 Conclusion

Study 4 demonstrates that error-driven learning mechanisms integrated within the ACT-R architecture can account for key aspects of human action control in dynamic environments. By linking prediction, memory, and uncertainty into a comprehensive cognitive model, it provides a mechanistic account of learning and action selection. This final study completes the dissertation's progression from behavioral description to computational modeling to architectural integration. It illustrates how a unified theory of action control can be grounded in both psychological and computational principles.

5 General Discussion

5.1 Theoretical Implications

This dissertation set out to examine the *behavioral, cognitive, and computational mechanisms that enable humans to exert effective action control across uncertain and dynamically changing environments.* I have presented the core findings in chapters 2-4, but their theoretical significance is less straightforward and requires unpacking from several angles.

5.1.1 Hierarchical Intentions Embodied in Fixation Dynamics

Study 1 identified two distinct types of fixational eye movements with complementary functions: Close/Type 0 fixations, which are shorter in duration and cluster around the spaceship, and Distant/Type 1 fixations, which are longer in duration and located farther ahead of the spaceship. Study 2 showed that these fixation types respond differently to environmental demands. With increasing input noise, Close fixations tightened around the ship, whereas during drift sections Distant fixations shifted even farther ahead. Taken together, these findings indicate that gaze control implements hierarchical intentions through the coordinated use of two fixation modes that support different levels of action control.

Both fixation types reflect the pursuit of Pacherie's (2006; 2008) category of proximal intentions (P-intentions). They guide the unfolding of action rather than specifying abstract goals (D-intentions) or controlling fine-grained motor execution (M-intentions). Yet the two fixation types differ in the degree to which they are grounded in concrete action plans and environmental structure.

Distant fixations are tightly linked to upcoming action demands. They are directed at specific future task-relevant locations, like open spaces between obstacles or within drift sections. Their longer durations suggest that attention remains anchored to these locations until the planned action is executed or task demands shift. This sustained engagement support proactive planning by extracting information that constrains near-future maneuvers. In this sense, Distant fixations correspond to highly specified proximal intentions, monitoring whether the planned maneuvers are unfolding as anticipated.

Close fixations, by contrast, are not tied to specific environmental features or landmarks. Instead, they provide a more diffuse monitoring of the immediate control region around the spaceship. Their functional role is best described as maintaining a safety envelope: avoid imminent collision or preserve response margins. These fixations support proximal intentions as well, just at a more abstract, non-object-specific level. Rather than specifying an exact future action, they maintain readiness to respond to whichever immediate demands arise.

The looser grounding of Close fixations is likely driven by time constraints inherent to the task. Generating and maintaining the concrete action plans that Distant fixations embody demands sustained cognitive resources, but these cannot be sustained indefinitely. Close fixations provide necessary intervals during which attention disengages from specific environmental targets, allowing the system to formulate the next concrete action plan. Under high time pressure, alternating between specific action targeting (via Distant Fixations) and diffuse vigilance (via Close Fixations) may be more efficient than attempting to maintain continuous detailed planning.

The finding of Study 2 that the proportion of Close and Distant fixations remained the same even as the allocation of the individual types adapted to task demands supports this claim. It implies that the two fixation types are not optional strategies that can be abandoned under difficulty; they are structurally necessary components of efficient gaze control. This adaptive pattern shows that the system maintains guidance in Shepherd's (2014) sense, in that visual attention remains organized around goal-relevant information that guides action control.

This interpretation extends Pacherie's framework by showing that hierarchical intentions are not merely abstract representational categories but emerge directly from the temporal and spatial organization of action. These findings offer concrete mechanistic grounding for the comprehensive account of action control introduced in Chapter 1 and demonstrate how the abstract structure of intentions manifests in the coordinated deployment of perceptual and motor systems. They reveal that effective action control depends on mechanisms that operate simultaneously at multiple levels of specificity, continuously balancing between immediate stability and forward planning.

5.1.2 Uncertainty as Productive Force in Learning and Control

High uncertainty does not always impair control and can, under certain conditions, facilitate learning. Study 2 demonstrated that moderate input noise, while reducing performance, did not diminish SoC. Participants maintained confidence in their ability to control the spaceship even as objective performance declined. In Study 4, the ACT-R model learned faster and more effectively when forced to generate predictions under high uncertainty rather than falling back to prediction-independent heuristics.

Together these patterns reveal that uncertainty serves dual, seemingly contradictory functions. On one hand, it degrades the precision of predictions and increases execution variability, directly impairing performance. On the other hand, it forces engagement with predictive mechanisms, generating prediction errors that help refine the forward model. The critical insight is that uncertainty is productive when it leads to prediction errors. Attempting to sidestep prediction under uncertainty (as the initial ACT-R with fallback did) prevents the accumulation of the very information needed to reduce uncertainty through learning.

This finding argues against simple models that do not represent uncertainty and favor purely habitual or heuristic responses. Instead, adaptive control appears to require ”leaning into” uncertainty by actively generating predictions, even when those predictions are initially poor. This principle aligns with the active inference framework (K. Friston, 2010), which proposes that agents reduce uncertainty by actively gathering information. In particular, it highlights that *making predictions under uncertainty* is the primary learning mechanism, not merely a consequence of reduced uncertainty.

There are limits to when uncertainty can be productive. Study 2 showed that SoC did not decrease gradually with increasing noise but dropped sharply once noise became severe, indicating a threshold beyond which uncertainty overwhelms the capacity for coherent prediction. Study 4 also revealed that the ACT-R model needed more trials than human participants to reach similar performance, suggesting that humans employ additional mechanisms for learning under uncertainty that the current model does not implement. Understanding how agents decide when to engage in predictive mechanisms, and when to fall back on non-predictive heuristics due to excessive uncertainty, remains an open question.

The central insight is that adaptive systems need uncertainty. They must be placed in conditions where prediction errors can arise and accumulate, as learning depends fundamentally on this engagement. Reducing uncertainty too aggressively may paradoxically limit opportunities for building the refined and flexible internal representations needed for robust control.

5.1.3 Model Coherence versus Model Precision

The previous section established that uncertainty drives learning when the agent recognizes prediction errors. This raises the question of what separates productive uncertainty from uncertainty that diminishes SoC.

In Study 2, invisible drift impaired both performance and SoC more severely than extreme input noise. The difference lies in how each manipulation affects internal forward models. Extreme noise degrades precision; participants understand what needs to be done but they cannot execute it reliably. Invisible drift, however, violates model coherence by introducing a hidden causal force that participant's representations fail to account for. Participants cannot build a coherent understanding of why their actions produce the outcomes they do, and no amount of effort compensates when the model itself misrepresents task dynamics.

This hints toward two distinct dimensions of internal model quality: *precision* and *coherence*. Model precision refers to the accuracy and reliability of predictions when the model's structure is correct. In predictive processing frameworks, precision reflects the weight or certainty assigned to prediction errors. High noise degrades precision, meaning that predictions become more variable and unreliable, but the underlying causal understanding remains intact (K. Friston, 2010). Model coherence, by contrast, refers to whether the model structure itself correctly represents the causal organization of the environment. When coherence breaks down, the model is not merely imprecise; it is fundamentally incorrect about the forces acting in the environment, regardless of the precision of its predictions (Clark, 2013).

Model coherence connects to Gentry's (2021) claim that genuine control requires semantic information processing. According to Gentry, a system must be able to detect when its representation or action selection has gone wrong (otherwise it simply malfunctions). This capacity marks the difference between a merely mechanistic process and a system that genuinely exercises control. The distinction between precision loss versus coherence violations is equal to the difference between noisy and wrong representations. A normative error occurs explicitly when agents recognize their internal model is fundamentally wrong.

This also explains why moderate noise in Study 2 did not reduce SoC even as performance declined. Precision loss alone does not diminish the subjective experience of control because the agent retains a valid causal model. Participants understood the relationship between actions and outcomes; uncertainty merely clouded execution. But violations of model coherence immediately resulted in a loss of SoC because participants recognized that their understanding was inadequate.

Distinguishing between model precision and model coherence clarifies the conditions under which uncertainty becomes productive. Uncertainty that challenges precision without violating coherence forces deeper engagement with predictive mechanisms, as argued in the previous section. The ACT-R model in Study 4 learned faster when generating predictions under high uncertainty (at the beginning of the experiment) precisely because the model structure was correct. Here, uncertainty drove parameter refinement without requiring structural revision. The model had to learn the strength of connections (e.g., between drift indicators and drift dynamics), but the relevant causal variables were already established (by me, the modeler).

Invisible drift presents a different challenge. Here, learning requires not only parameter adjustment but also recognizing that the model is missing critical causal structure. While Study 3 only featured visible drift, it provides relevant insight through the Bayesian model's inability to account for participants' rapid evidence reweighting after crashes. The results suggest that agents not merely update beliefs about probabilities but engage in meta-cognitive monitoring and actively track whether their causal model remains valid. Crashes may signal potential model inadequacy and trigger increased sensitivity to subsequent evidence which may lead to revising one's understanding of task dynamics much faster and much more easily. This meta-cognitive mechanism is particularly critical when facing invisible causal forces that violate model coherence entirely.

Formalizing a Bayesian framework that is capable of structural model revision would require going beyond standard parameter updates (which assume fixed likelihood functions). There needs to be a mechanism for recognizing when the current likelihood function cannot accommodate the evidence and revising the set of represented causal variables. But this remains an open challenge for computational models of agency and control.

The distinction between precision and coherence has clinical implications. Conditions like depression and anxiety are often characterized by reduced SoC, but the mechanisms may differ. Some individuals may maintain coherent causal models but underestimate their precision ("I understand what I need to do, but I can't execute it"). Others may fail to construct or maintain coherent models of how their actions relate to outcomes, experiencing a more fundamental breakdown in agency. Recent work in computational psychiatry supports this perspective. Gilbert, Wusinich, and Zarate (2022) propose a predictive coding framework for understanding major depression, highlighting how disruptions in generative models and aberrant prediction-error processing can impair the brain's internal representation of the environment. These findings suggest that, in some cases, reduced SoC reflects violations of coherence in internal models rather than low confidence in execution.

More broadly, these findings suggest that agents (biological or artificial) must balance two imperatives: maintaining coherent representations of causal structure while remaining open to revising those representations when evidence accumulates that they are wrong. Shielding agents from uncertainty prevents both parameter refinement (as discussed in Section 5.1.2) and the discovery of model inadequacies. Effective learning requires embracing uncertainty, but only when the system has mechanisms for detecting when precision loss signals that the internal forward model itself needs restructuring.

5.1.4 Representational Inertia and the Lag Between Performance and SoC

If precision and coherence indicate the quality of internal models, it raises the question of how internal models are revised. Why does SoC remain stable even when performance declines? Study 2 showed that whereas extreme noise reduced performance and subjective control, moderate noise only affected the former. This asymmetry suggests that internal models do not update in lockstep with environmental changes. Instead, they exhibit representational inertia, which is the tendency to preserve existing model structure until evidence strongly indicates that it is no longer adequate.

This inertia can be explained by the fact that agents can compensate a wide range of environmental dynamics through parametric adjustments while maintaining model structure. Moderate noise, for instance, may be handled by widening confidence intervals or adjusting gain parameters in sensorimotor control without altering the underlying causal representation. If the existing model structure can "explain" the environment, even imprecisely, SoC is maintained. Participants still feel in control because they understand what *should* happen, even if noise makes execution unreliable.

Only when environmental changes exceed the model's parametric flexibility does structural updating become necessary. Here, the agent must revise the causal organization itself, not just tune its parameters. This threshold aligns with the sharp SoC decline observed under extreme noise in Study 2. These threshold dynamics align with Study 3's Bayesian evidence-reweighting model, which suggested meta-cognitive monitoring of model adequacy. Here, SoC reflected not immediate performance outcomes but whether the agent believes their internal model still accounts for the evidence.

Kahl et al. (2022) propose a two-layered agent architecture in which each level is associated with its own SoC (as depicted in Figure 5.1). At the sensorimotor control layer, the low-level SoC remains stable as long as prediction errors can be resolved through continuous parametric adjustments. This layer operates largely automatically and maintains control when the environment stays within the internal model's representational bounds. When environmental changes exceed this flexibility, accumulated prediction errors push the low-level SoC below a critical threshold, which activates the cognitive control layer. At this second level, the agent uses partly symbolic representations to construct and update the internal model and generate longer-horizon predictions. Kahl et al. describe the high-level SoC as a self-assessment that integrates lower-level evaluations with judgments about goal progress that guides strategic adjustments. The implications of this thesis extend Kahl et al.'s view by proposing that the high-level SoC also enables more demanding forms of control, such as structural model revision, deliberate hypothesis testing, and flexible strategy selection, which cannot be achieved through

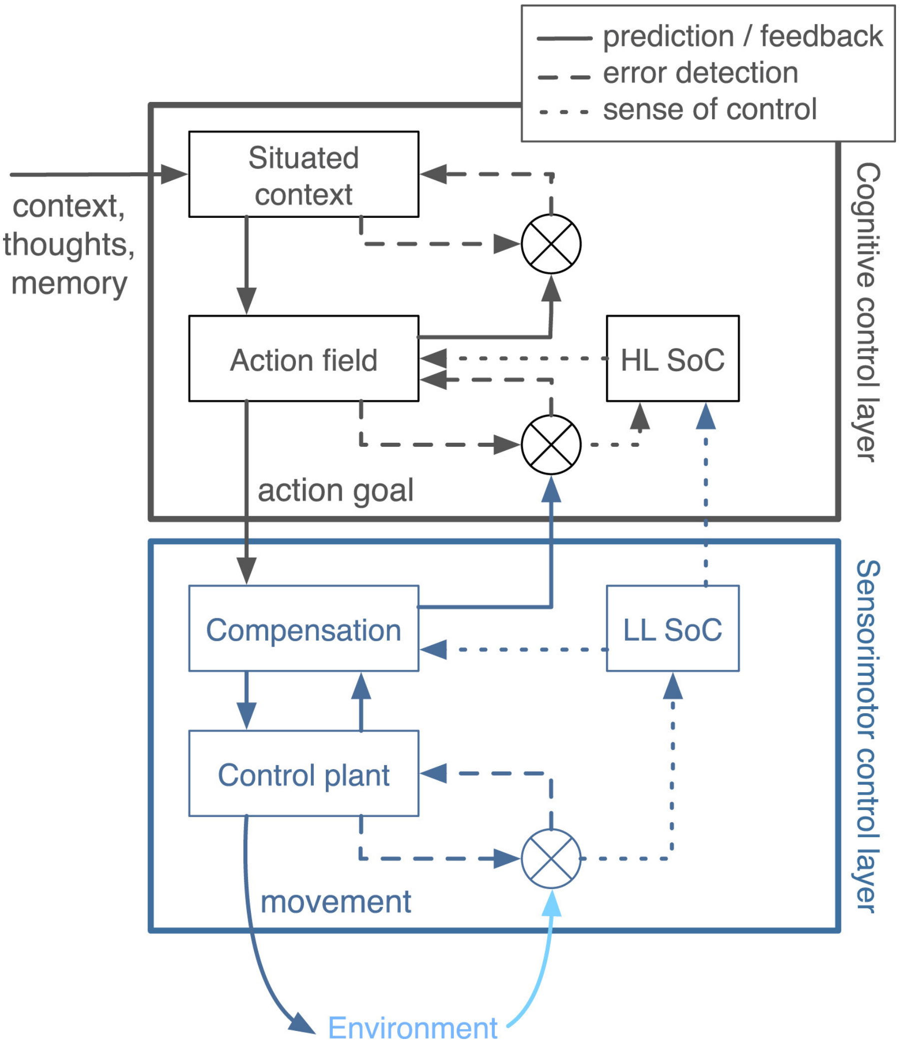

Figure 5.1 Illustration of the two-level hierarchical control architecture by Kahl et al. (2022). At the cognitive control layer, contextual information informs the situated context, which specifies an action field from which an action goal is selected. This goal is translated into motor commands at the sensorimotor control layer, where the control plant executes movements. Each level contains its own Sense of Control (SoC) module: a higher-level SoC (HL SoC) for cognitive control and a lower-level SoC (LL SoC) for sensorimotor control. Both modules compare predicted states against actual states (crossed circles) to detect prediction errors (dashed lines). When sensorimotor prediction errors exceed what parametric adjustments can resolve, the low-level SoC drops below a critical threshold, activating the cognitive control layer to update the situated context and generate a new action field.

parameter tuning alone. A decline in the high-level SoC marks the transition from an automatic, model-preserving regime to a reflective, model-revising one, mirroring the threshold dynamics observed in Study 2.

5.1.5 Two Pathways to Proactive Control

The final component of the theoretical framework explored in this dissertation concerns the mechanisms that support proactive control. Proactive control refers to preparing for future states rather than merely reacting to present ones. Across Studies 2 and 4, the results suggest that proactive control can emerge through two distinct mechanisms with different computational demands: externally-cued proactivity and internally-enabled proactivity.

Externally-cued proactivity arises when the environment provides reliable signals about upcoming perturbations. In Study 2, visible drift indicators prompted participants to increase their Distant Fixations, looking farther ahead to anticipate forced lateral movement. Externally-cued anticipatory behavior occurred independently of trial (how far participants have progressed within the experiment) and therefore independently of the accuracy or coherence of internal drift estimates. The presence of a signal indicating that drift was imminent appeared sufficient to elicit proactive adjustments. It is therefore reasonable to assume that such external cues can scaffold proactive control even when internal models are uncertain.

Internally-enabled proactivity, in contrast, refers to anticipatory control that is enabled by an agent's own predictive model rather than by external signals. In Study 4, the ACT-R model increasingly used forward simulation as its drift parameters were tuned by prediction errors, showing how learning can turn a generative model into a tool for anticipatory action. However, because only visible drift cues were present in that task, we cannot conclude that participants relied on internal prediction in the absence of cueing. Still, the model provides a clear mechanistic account of how internally-driven proactivity would arise: as model precision and coherence improve, forward simulation becomes reliable enough to guide entry-point selection even without external prompts. Testing this directly will require tasks that remove or mask external cues so we can observe whether learned internal models alone support proactive behavior.

These two pathways have distinct computational profiles and fail states:

Externally-cued proactivity is robust to internal uncertainty, since even a coarse generative model can benefit from an explicit signal about impending perturbations, but breaks down when cues are absent or misleading, as in the invisible-drift condition of Study 2.

Internally-enabled proactivity does not depend on external signals and can function even when environmental cues are unavailable, but only once the internal model is sufficiently precise, which takes time and experience.

Critically, proactive control requires at least one of these pathways. When both mechanisms are unavailable (when model quality is low and no external cue signals when to anticipate) behavior collapses into reactive control. This pattern was evident both in Study 2 (invisible drift) and in the early learning stages of the ACT-R model in Study 4, where agents behaved conservatively and remained close to the horizontal center when entering drift sections.

Flexible repeatability, as described by Shepherd (2014), depends on proactive control. Participants and cognitive models adapted successfully to moderate uncertainty and to variation in drift position or direction, showing that anticipatory policies can transfer as long as the underlying predictive structure remains stable. The study results also highlight the limits of flexible repeatability. Performance declines when internal forward models no longer match the environment and when action execution becomes too noisy. In these situations, proactive control is not absent in principle, but the predictive basis required for its implementation is no longer reliable and flexible repeatability breaks down.

The limits of proactive control have implications for artificial agents and human-support systems. Maintaining goal-directed behavior in uncertain environments requires either internalizing a sufficiently stable predictive model or receiving external support that guides anticipation. Systems that offer neither form of support, and that do not allow enough learning for internal models to improve, limit users and agents to purely reactive and less efficient control strategies.

5.2 Future Directions

Section 5.1 points to several open questions concerning model coherence, representational learning, and the conditions under which proactive control emerges. In this section, I outline concrete experimental and modeling directions that can address these questions.

5.2.1 Structural Learning

The implications of model coherence and representational inertia extend directly to the ACT-R model used in Study 4. In its current form, the model relies on a fixed linear forward model whose parameters are refined through prediction errors. This works well as long as the environment conforms to those linear assumptions, allowing the model to improve performance without altering its structure. But within the notion of representational-inertia, and in line with the argument of Kahl et al. (2022), this also highlights a limitation: the cognitive model cannot detect or compensate for structural mismatches in the environment. If the true drift dynamics were nonlinear or otherwise incompatible with the model's representational form, prediction errors would accumulate in ways that parameter updates could not resolve.

Future work should therefore investigate when agents must transition from parameter refinement to structural revision. This can be tested experimentally by introducing nonlinear or structurally altered drift dynamics in later blocks of the Dodge Asteroids task. One possible manipulation would be to make drift magnitude dependent on the spaceship's horizontal position. It would expose human participants as well as cognitive models to situations in which prediction errors cannot be corrected solely by adjusting parameters.

SoC can be conceptualized as an index of structural adequacy rather than moment-to-moment task success. As long as agents believe "my model is basically right", SoC is preserved even when performance suffers. When agents recognize that "my model cannot possibly be right", SoC diminishes. This explains nonlinear dynamics: SoC should remain stable across a broad range of perturbations, then decline sharply once representational inertia is overcome and structural updating becomes necessary. Cognitive models can be implemented that they produce behavior based on internal forward models. SoC can then guide updating these internal models (through parameter updates as in Study 4 or structural adjustments). Combining this modeling approach with identifying behavioral markers (e.g. in eye movements) near the transition threshold in humans would enable to close the loop and to determine precisely when agents recognize structural inadequacies and test how internal model adjustments correlate with observable behavior.

Extending the discussion of internal models and SoC, a related issue is representational inertia in clinical contexts. Depression and learned helplessness often involve stable, negative beliefs about personal efficacy that persist despite contradictory evidence (Würtz et al., 2024; Hyland, 2020; Abramson, Seligman, and Teasdale, 1978). The current framework suggests these beliefs may reflect not merely biased evidence weighting but structural representations that are difficult to revise. In this case, interventions might be more effective if they target model structure explicitly. Rather than only providing positive performance feedback that might drive parameter adjustments into existing negative models, interventions could help individuals build alternative causal representations (Schneider et al., 2023; Marian and Filimon, 2010).

5.2.2 Transitioning from Externally-cued to Internally-enabled Proactive Control

As stated above, proactive control can emerge through two pathways. The environment may provide reliable signals about when anticipatory control is required. Or proactive behavior is enabled through the precision and coherence of internally learned generative models that support forward simulation. This two-pathway account offers a mechanistic explanation of when agents can shift from reactive to proactive control.

A key direction for future research is to disentangle these pathways more explicitly. In Study 2, proactive adjustments were strongly driven by external cues; in Study 4, drift cues were visible throughout the experiment, meaning the ACT-R model's internally enabled proactivity cannot be isolated. Future experiments using the Dodge Asteroids or related experimental environments should remove or mask external drift indicators after initial learning of the task. This would allow to test how agents operate under conditions where anticipation is only guided by their internal forward models.

Such designs would also help identify the transition point between the two pathways. If internally enabled proactivity emerges only once a generative model reaches a certain precision threshold, this transition should have traceable behavioral signatures, possibly reflected in eye-movement patterns, riskier entry point selections, or reduced reliance on conservative strategies. ACT-R simulations could be used to map these thresholds and test how model precision interacts with environmental cue availability.

In ACT-R, the transition from deliberate, cue-dependent control to more autonomous, proactive behavior is typically modeled through mechanisms such as production compilation, which gradually proceduralizes repeated cognitive operations. However, for tasks that involve noisy, variable, or continuous dynamics, these mechanisms alone may not suffice. In such cases, combining ACT-R with reinforcement learning or Bayesian inference can provide the additional flexibility needed for robust, cue-independent proactivity and more precise internal predictions. This type of modeling should be explored to explain the cognitive processes in the Dodge Asteroids environment with the adjustments described above.

More broadly, this line of work raises further questions about how agents choose between cue-driven and model-driven anticipatory strategies. Do the two pathways operate additively, competitively, or in a hierarchical fashion where external cues dominate until internal models become sufficiently reliable? Investigating these interactions, both empirically and through cognitive modeling, would clarify how proactive control emerges and how agents allocate control under uncertainty.

5.3 General Conclusion

This dissertation set out with the question: What *behavioral*, *cognitive*, and *computational* mechanisms enable humans (and artificial agents modeled after them) to exert effective action control across uncertain and dynamically changing environments? Across four empirical studies combining eye-tracking experiments, Bayesian modeling, and cognitive architecture simulations, This thesis identified mechanisms operating at three interconnected levels that together enable adaptive action control.

At the behavioral level, effective control requires implementing action goals that are grounded in the direct environment to varying degrees. The discovery of two distinct fixation types (Close fixations that maintain a safety envelope around immediate action and Distant fixations that anchor attention to specific future task-relevant locations) reveals how the visual system manages temporal demands inherent to uncertain and dynamically changing environments. These fixation types do not represent optional strategies but are functionally necessary components of effective control. Close fixations provide intervals for formulating concrete action plans while maintaining diffuse situational monitoring, whereas Distant fixations embody those plans through sustained attention to upcoming environmental changes. Their coordinated use demonstrates that hierarchical intentions emerge directly from the temporal and spatial organization of perception and action.

At the cognitive level, adaptive control depends on the quality and use of internal forward models that support both reactive and proactive modes of action selection. The critical insight concerns not whether agents possess forward models, but how those models maintain both precision and coherence under uncertainty. Precision reflects the reliability of predictions given correct model structure, while coherence reflects whether the model correctly represents the causal structure of the environment. This distinction explains why moderate environmental noise preserves SoC even as performance declines: agents retain valid causal understanding despite execution uncertainty. In contrast, invisible forces that violate model structure produce immediate SoC loss. Crucially, SoC functions as a metacognitive signal that tracks model quality and gates the shift from reactive responses to anticipatory control strategies. When SoC remains intact, agents can leverage their internal forward models proactively through two distinct pathways: externally-cued proactivity that uses environmental signals to scaffold anticipation even with uncertain internal models, and internally-enabled proactivity that relies on sufficiently precise learned representations to support forward simulation. The interplay between these pathways determines whether agents can anticipate future states or must fall back on reactive control.

At the computational level, the difference between how agents respond to moderate noise versus how they respond to perturbations that violate forward model predictions implies hierarchical learning mechanisms that flexibly shift between parameter refinement and structural revision. When environmental dynamics fall within an internal model's representational bounds, prediction errors drive parameter updates that improve precision while preserving model structure. This parametric flexibility explains representational inertia, the tendency for internal models and SoC to remain stable across a range of perturbations. Only when environmental dynamics exceed this flexibility must agents engage in structural revision and fundamentally rethink the causal variables and relationships represented in their internal forward models. The ACT-R model's error-driven parameter learning highlights the first stage of this hierarchy, in that it shows how uncertainty becomes productive when it generates prediction errors rather than being circumvented by falling back on prediction-independent heuristics. The two-layer hierarchy of action control proposed by Kahl et al. (2022) could be used to implement this computationally. The sensorimotor control layer continuously updates parameters of the internal forward models and an associated low-level SoC when it encounters prediction errors. When prediction errors accumulate the cognitive control layer activates, which then revises the structure of the forward models themselves.

These three levels are not independent but form an integrated system. Behavioral mechanisms implement the hierarchical structure of intentions through coordinated gaze control. Cognitive mechanisms maintain and update internal forward models that guide both behavioral strategies and subjective experiences of control. Computational mechanisms determine when and how those models are revised, providing both the stability needed for coherent action and the flexibility required for adaptive control. Effective action control in uncertain environments is therefore neither purely reactive nor rigidly predictive. Instead, it emerges from the dynamic interplay between multiple mechanisms operating at different timescales and levels of abstraction.

The implications of this dissertation extend beyond the specific task environment studied. They offer a framework for understanding how adaptive systems (whether biological or artificial) navigate the fundamental tension between exploiting current knowledge and remaining open to environmental changes that require structural learning. It suggests that shielding agents from uncertainty, while reducing immediate errors, paradoxically limits the very experiences needed to build robust, flexible internal models. And it provides concrete targets for intervention when action control breaks down, whether in clinical populations experiencing diminished agency or in artificial systems that fail to generalize beyond training conditions.

Understanding how humans maintain control and agency in uncertain, changing environments remains a challenge that spans behavioral, cognitive, and computational levels of analysis. This work establishes the groundwork for advancing theories of human action control by revealing how these levels interact to enable seamless, ongoing adjustments. I conclude with the hope that these insights will inform future efforts to apply human adaptability to the design of effective systems.

List of Figures

Bibliography

Abalakin, Nikita, Nils Wendel Heinrich, Annika Österdiekhoff, Stefan Kopp, and Nele Russwinkel (2024). “Allocation of Fixational Eye Movements in Response to Uncertainty in Dynamic Environments”. *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 46.

Abramson, Lyn Y., Martin E. P. Seligman, and John D. Teasdale (1978). “Learned helplessness in humans: Critique and reformulation”. *Journal of Abnormal Psychology* 87.1, pp. 49–74.

Anderson, John R., Daniel Bothell, Michael D. Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin (2004). “An integrated theory of the mind”. *Psychological Review* 111.4, pp. 1036–1060.

Anderson, John R. and Christian Lebiere (1998). *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.

Botvinick, Matthew M., Todd S. Braver, Deanna M. Barch, Cameron S. Carter, and Jonathan D. Cohen (2001). “Conflict monitoring and cognitive control”. *Psychological Review* 108.3, pp. 624–652.

Bratman, Michael E. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.

Braver, Todd S. (2012). “The variable nature of cognitive control: A dual mechanisms framework”. *Trends in Cognitive Sciences* 16.2, pp. 106–113.

Camacho, Eduardo F. and Carlos Bordons (2007). “Constrained model predictive control”. *Model Predictive Control*. Springer, pp. 177–216.

Carver, Charles S. and Michael F. Scheier (2013). “Goals and emotion”. *Handbook of Cognition and Emotion*. New York: Guilford Press, pp. 176–194.

Clark, Andy (2013). “Whatever next? Predictive brains, situated agents, and the future of cognitive science”. *Behavioral and Brain Sciences* 36.3, pp. 181–204.

Cooper, Richard and Tim Shallice (2000). “Contention scheduling and the control of routine activities”. *Cognitive Neuropsychology* 17.4, pp. 297–338.

Dale, Rick, Caitlin Kehoe, and Michael J. Spivey (2007). “Graded motor responses in the time course of categorizing atypical exemplars”. *Memory & Cognition* 35.1, pp. 15–28.

Dickinson, Anthony and Bernard W. Balleine (1994). "Motivational control of goal-directed action". *Animal Learning & Behavior* 22.1, pp. 1–18. DOI: `10.3758/BF03199951`.

Elsner, Birgit and Bernhard Hommel (2001). "Effect anticipation and action control". *Journal of Experimental Psychology: Human Perception and Performance* 27.1, pp. 229–240.

Engel, Andreas K., Alexander Maye, Martin Kurthen, and Peter König (2013). "Where's the action? The pragmatic turn in cognitive science". *Trends in Cognitive Sciences* 17.5, pp. 202–209. DOI: `10.1016/j.tics.2013.03.006`.

Fischer, John Martin and Mark Ravizza (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.

Friston, Karl (2005). "A theory of cortical responses". *Philosophical Transactions of the Royal Society B: Biological sciences* 360.1456, pp. 815–836.

Friston, Karl (2010). "The free-energy principle: A unified brain theory?" *Nature Reviews Neuroscience* 11.2, pp. 127–138.

Gentry, Hunter R. (2021). "Extended control systems: A theory and its implications". *Philosophical Psychology* 34.3, pp. 345–373.

Gershman, Samuel J. (2015). "Do learning rates adapt to the distribution of rewards?" *Psychonomic Bulletin & Review* 22.5, pp. 1320–1327. DOI: `10.3758/s13423-014-0790-3`.

Giersiepen, Maren, Nils Wendel Heinrich, Annika Österdiekhoff, Stefan Kopp, Nele Russwinkel, Simone Schütz-Bosbach, and Jakob Kaiser (2025). "Am I in control? The dynamics of sensory information, performance feedback, and personality in shaping the sense of control". *Consciousness and Cognition* 135, p. 103938.

Gilbert, Jessica R., Christina Wusinich, and Jr. Zarate Carlos A. (2022). "A predictive coding framework for understanding major depression". *Frontiers in Human Neuroscience* 16, p. 787495.

Gratton, Gabriele, Michael G. H. Coles, and Emanuel Donchin (1992). "Optimizing the use of information: Strategic control of activation of responses". *Journal of Experimental Psychology: General* 121.4, pp. 480–506.

Grünbaum, Thor and Mark Schram Christensen (2020). "Measures of agency". *Neuroscience of consciousness* 2020.1, niaa019.

Haaren-Mack, Birte von, Carolin Bastemeyer, and Nils Wendel Heinrich (2022). "Affective reactivity to real-life stressors in sports students—the role of context and physical activity". *Poly-Pill-Bewegung—Chancen, Risiken und Nebenwirkungen: Jahrestagung dvs-Kommission Gesundheit*, p. 12.

Heinrich, Nils Wendel, Rebecca von Engelhardt, Annika Österdiekhoff, Stefan Kopp, and Nele Russwinkel (2025). "Towards a Bayesian Cognitive Model of Self-Belief Updating". *Proceedings of International Conference on Cognitive Modeling (ICCM 2025).*

Heinrich, Nils Wendel, Amir Jahanian Najafabadi, and Jairo Pérez-Osorio (2025). “Unsupervised Clustering Uncovers Two Distinct Types of Fixational Eye-Movements in Dynamic Environments”. *2025 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*. IEEE.

Heinrich, Nils Wendel, Annika Österdiekhoff, Stefan Kopp, and Nele Rußwinkel (2023). “A Straightforward Implementation of Sensorimotor Abstraction in a Two-Layer Architecture for Dynamic Decision-Making”. *Proceedings of International Conference on Cognitive Modeling (ICCM 2023)*.

Heinrich, Nils Wendel, Annika Österdiekhoff, Stefan Kopp, and Nele Rußwinkel (2024). “Developing and Evaluating a Computational Cognitive Model of Sensorimotor Grounded Action Selection Based on Eye-movement Behavior”. *Proceedings of International Conference on Cognitive Modeling (ICCM 2024)*.

Heinrich, Nils Wendel, Annika Österdiekhoff, Stefan Kopp, and Nele Rußwinkel (2025). “Using Eye Movements to Understand Sense of Control in Situated Action”. *Cognitive Science* 49.12, e70154. DOI: `10.1111/cogs.70154`.

Heinrich, Nils Wendel, Annika Österdiekhoff, Stefan Kopp, and Nele Russwinkel (2024). “Goal-directed Allocation of Gaze Reflects Situated Action Control in Dynamic Tasks”. *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 46.

Heinrich, Nils Wendel, Annika Österdiekhoff, Martin Stuwe, Stefan Kopp, and Nele Rußwinkel (2026). “Error-Driven Learning in Human Action Control”. Manuscript submitted for publication.

Hommel, Bernhard, Jochen Müsseler, Gisa Aschersleben, and Wolfgang Prinz (2001). “The theory of event coding (TEC): A framework for perception and action planning”. *Behavioral and Brain Sciences* 24.5, pp. 849–878.

Hyland, Michael E. (2020). “A reformulated contextual model of psychotherapy for treating anxiety and depression”. *Clinical Psychology Review* 80, p. 101890.

James, William (1890). *The Principles of Psychology*. Vol. 2. New York: Henry Holt and Company.

Kaelbling, Leslie Pack and Tomás Lozano-Pérez (2013). “Integrated task and motion planning in belief space”. *The International Journal of Robotics Research* 32.9-10, pp. 1194–1227.

Kahl, Sebastian, Sebastian Wiese, Nele Russwinkel, and Stefan Kopp (2022). “Towards autonomous artificial agents with an active self: Modeling sense of control in situated action”. *Cognitive Systems Research* 72, pp. 50–62.

Körding, Konrad P. and Daniel M. Wolpert (2004). “Bayesian integration in sensorimotor learning”. *Nature* 427.6971, pp. 244–247. DOI: `10.1038/nature02169`.

Krumm, Carolin, Nils Wendel Heinrich, and Birte von Haaren-Mack (2023). “Affective reactions to real-life stressors and the role of physical activity in sports students—An Ambulatory Assessment study”. *Psychology of Sport and Exercise* 69, p. 102503.

Land, Michael and Benjamin Tatler (2009). *Looking and Acting: Vision and Eye Movements in Natural Behaviour*. Oxford: Oxford University Press.

Lashley, Karl Spencer (1951). *The Problem of Serial Order in Behavior*. Indianapolis: Bobbs-Merrill.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (2015). “Deep learning”. *Nature* 521.7553, pp. 436–444.

Lee, John D and Katrina A See (2004). “Trust in automation: Designing for appropriate reliance”. *Human factors* 46.1, pp. 50–80.

Levine, Sergey, Chelsea Finn, Trevor Darrell, and Pieter Abbeel (2016). “End-to-end training of deep visuomotor policies”. *Journal of Machine Learning Research* 17.39, pp. 1–40.

Logan, Gordon D. and William B. Cowan (1984). “On the ability to inhibit thought and action: A theory of an act of control”. *Psychological Review* 91.3, pp. 295–327.

Luria, Alexander Romanovich (1966). *Higher Cortical Functions in Man*. New York: Basic Books.

Manchingal, Shireen Kudukkil, Andrew Bradley, Julian FP Kooij, Keivan Shariatmadar, Neil Yorke-Smith, and Fabio Cuzzolin (2025). “Epistemic Artificial Intelligence is Essential for Machine Learning Models to TrulyKnow When They Do Not Know’”. *arXiv preprint arXiv:2505.04950*.

Marian, Mihai and Letitia Filimon (2010). “Cognitive restructuring and improvement of symptoms with cognitive-behavioural therapy and pharmacotherapy in patients with depression”. *Annals of General Psychiatry* 9.Suppl 1, S173.

Moore, James W. and Patrick Haggard (2007). “Awareness of action: Inference and prediction”. *Consciousness and Cognition* 16.2, pp. 391–399. DOI: `10.1016/j.concog.2006.09.004`.

Newell, Allen and Herbert A. Simon (1976). “Computer science as empirical inquiry: Symbols and search”. *Communications of the ACM*. Vol. 19. 3, pp. 113–126.

Norman, Donald A. and Tim Shallice (1986). “Attention to action: Willed and automatic control of behavior”. *Consciousness and Self-Regulation: Advances in Research and Theory*. Ed. by Richard J. Davidson, Gary E. Schwartz, and David Shapiro. Vol. 4. Springer, pp. 1–18.

Österdiekhoff, Annika, Nils Wendel Heinrich, Nele Rußwinkel, and Stefan Kopp (2024). “Model-Based Reinforcement Learning with Hierarchical Control for Dynamic Uncertain Environments”. *Intelligent Systems Conference*. Springer, pp. 626–642.

Österdiekhoff, Annika, Nils Wendel Heinrich, Nele Russwinkel, and Stefan Kopp (2024). "Sense of Control in Dynamic Multitasking and its Impact on Voluntary Task-Switching Behavior". *Proceedings of the Annual Meeting of the Cognitive Science Society.* Vol. 46.

Österdiekhoff, Annika, Nils Wendel Heinrich, Nele Russwinkel, and Stefan Kopp (2026). "Giving AI agents a sense of control facilitates reinforcement learning in multitasking scenarios". *PloS one* 21.2, e0342305.

Pacherie, Elisabeth (2006). "Towards a dynamic theory of intentions". *Does Consciousness Cause Behavior?* Ed. by Susan Pockett, William P. Banks, and Shaun Gallagher. MIT Press, pp. 145–167.

Pacherie, Elisabeth (2008). "The phenomenology of action: A conceptual framework". *Cognition* 107.1, pp. 179–217.

Pezzulo, Giovanni, Francesco Rigoli, and Karl J. Friston (2018). "Hierarchical active inference: A theory of motivated control". *Trends in Cognitive Sciences* 22.4, pp. 294–306.

Prinz, Wolfgang (1997). "Perception and action planning". *European Journal of Cognitive Psychology* 9.2, pp. 129–154.

Reason, James (1984). "Little slips and big disasters". *Interdisciplinary Science Reviews* 9.2, pp. 179–189.

Rosenbaum, David A. (1980). "Human movement initiation: Specification of arm, direction, and extent". *Journal of Experimental Psychology: General* 109.4, pp. 444–474.

Rotter, Julian B. (1966). "Generalized expectancies for internal versus external control of reinforcement". *Psychological Monographs: General and Applied* 80.1, pp. 1–28.

Russell, Stuart J. and Peter Norvig (2020). *Artificial Intelligence: A Modern Approach.* 4th ed. Hoboken, NJ: Pearson.

Schneider, Brooke C., Ruth Veckenstedt, Evangelos Karamatskos, Lara Guedes Pinho, Bruno Morgado, César Fonseca, Steffen Moritz, and Lena Jelinek (2023). "Negative cognitive beliefs, positive metacognitive beliefs, and rumination as mediators of metacognitive training for depression in older adults (MCT-Silver)". *Frontiers in Psychology* 14, p. 1153377.

Schulz, Eric and Samuel J. Gershman (2019). "The algorithmic architecture of exploration in the human brain". *Current Opinion in Neurobiology* 55, pp. 7–14.

Shallice, Tim and Paul W. Burgess (1991). "Deficits in strategy application following frontal lobe damage in man". *Brain* 114.2, pp. 727–741.

Shapiro, Larry (2007). "The embodied cognition research programme". *Philosophy Compass* 2.2, pp. 338–346.

Shepherd, Joshua (2014). "The contours of control". *Philosophical Studies* 170.3, pp. 395–411.

Siciliano, Bruno and Oussama Khatib (2016). *Springer Handbook of Robotics.* 2nd ed. Berlin: Springer.

Spivey, Michael J., Marc Grosjean, and Günther Knoblich (2005). "Continuous attraction toward phonological competitors". *Proceedings of the National Academy of Sciences* 102.29, pp. 10393–10398.

Sternberg, Saul (1969). "The discovery of processing stages: Extensions of Donders' method". *Acta Psychologica* 30, pp. 276–315.

Sutton, Richard S. and Andrew G. Barto (2018). *Reinforcement Learning: An Introduction.* 2nd ed. Cambridge, MA: MIT Press.

Synofzik, Matthis, Gottfried Vosgerau, and Albert Newen (2008). "Beyond the comparator model: a multifactorial two-step account of agency". *Consciousness and cognition* 17.1, pp. 219–239.

Tatler, Benjamin W., Mary M. Hayhoe, Michael F. Land, and Dana H. Ballard (2011). "Eye guidance in natural vision: Reinterpreting salience". *Journal of Vision* 11.5, p. 5.

Tenenbaum, Joshua B., Thomas L. Griffiths, and Charles Kemp (2006). "Theory-based Bayesian models of inductive learning and reasoning". *Trends in Cognitive Sciences* 10.7, pp. 309–318. DOI: `10.1016/j.tics.2006.05.009`.

Todorov, Emanuel and Michael I. Jordan (2002). "Optimal feedback control as a theory of motor coordination". *Nature Neuroscience* 5.11, pp. 1226–1235.

Uesato, Jonathan, Ananya Kumar, Csaba Szepesvari, Tom Erez, Avraham Ruderman, Keith Anderson, Nicolas Heess, Pushmeet Kohli, et al. (2018). "Rigorous agent evaluation: An adversarial approach to uncover catastrophic failures". *arXiv preprint arXiv:1812.01647.*

Verbruggen, Frederick and Gordon D. Logan (2008). "Response inhibition in the stop-signal paradigm". *Trends in Cognitive Sciences* 12.11, pp. 418–424.

Voigt, Laura, Jannis Friedrich, Patricia Grove, Nils Wendel Heinrich, Sandra Ittlinger, Maša Iskra, Lisa Koop, Alexej Michirev, Simone Sparascio, and Markus Raab (2023). "Advancing judgment and decision-making research in sport psychology by using the body as an informant in embodied choices". *Asian Journal of Sport and Exercise Psychology* 3.1, pp. 47–56.

Wennerhold, Lasse and Malte Friese (2023). "Challenges in the conceptualization of trait self-control as a psychological construct". *Social and Personality Psychology Compass* 17.3, e12726.

Wit, Sanne de and Anthony Dickinson (2009). "Associative theories of goal-directed behaviour: A case for animal–human translational models". *Psychological Research* 73.4, pp. 463–476.

Wolpert, Daniel M., Zoubin Ghahramani, and Michael I. Jordan (1995). “An internal model for sensorimotor integration”. *Science* 269.5232, pp. 1880–1882. DOI: `10.1126/science.7569931`.

Wolpert, Daniel M. and Michael S. Landy (2012). “Motor control is decision-making”. *Current Opinion in Neurobiology* 22.6, pp. 996–1003.

Würtz, Felix, Tobias Kube, Marcella L. Woud, Jürgen Margraf, and Simon E. Blackwell (2024). “Reduced belief updating in the context of depressive symptoms: An investigation of the associations with interpretation biases and self-evaluation”. *Cognitive Therapy and Research* 48.2, pp. 225–241.