
Approximate Bayesian Inference of Parametric Cost Functions in Continuous Decision-Making

L^AT_EX using TU Darmstadt's Corporate Design

Master thesis by Tobias F. Niehues

Date of submission: December 29, 2023

1. Review: Prof. Jan Peters, PhD
2. Review: Prof. Constantin A. Rothkopf, PhD
Darmstadt



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Computer Science
Department

Centre for Cognitive Science
Psychology of Information
Processing

Erklärung zur Abschlussarbeit gemäß § 22 Abs. 7 APB TU Darmstadt

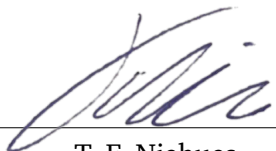
Hiermit erkläre ich, Tobias F. Niehues, dass ich die vorliegende Arbeit gemäß § 22 Abs. 7 APB der TU Darmstadt selbstständig, ohne Hilfe Dritter und nur mit den angegebenen Quellen und Hilfsmitteln angefertigt habe. Ich habe mit Ausnahme der zitierten Literatur und anderer in der Arbeit genannter Quellen keine fremden Hilfsmittel benutzt. Die von mir bei der Anfertigung dieser wissenschaftlichen Arbeit wörtlich oder inhaltlich benutzte Literatur und alle anderen Quellen habe ich im Text deutlich gekennzeichnet und gesondert aufgeführt. Dies gilt auch für Quellen oder Hilfsmittel aus dem Internet.

Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Mir ist bekannt, dass im Falle eines Plagiats (§ 38 Abs. 2 APB) ein Täuschungsversuch vorliegt, der dazu führt, dass die Arbeit mit 5,0 bewertet und damit ein Prüfungsversuch verbraucht wird. Abschlussarbeiten dürfen nur einmal wiederholt werden.

Bei einer Thesis des Fachbereichs Architektur entspricht die eingereichte elektronische Fassung dem vorgestellten Modell und den vorgelegten Plänen.

Darmstadt, 29. Dezember 2023



T. F. Niehues

Acknowledgements

This work marks the end of seven wonderful years that introduced me to the academic world. A lot of people came and went during that journey but some stayed – it is the very least I can do to mention and honor them here.

First, I want to thank Constantin for his constant support and understanding for any problems that I encountered during my work in the department and for helping me find the academic direction I want to pursue. I am immensely grateful to Dominik and Florian, whose unwavering dedication and guidance allowed me to write two theses I am incredibly proud of. Moreover, they showed me how enjoyable academia can be with the right colleagues, adding a valuable dimension to my academic journey.

I also want to thank the department's secretary Inge for settling any formal matters easily and for all the delightful chats before and after exhausting laboratory sessions.

The support from my family and friends was enormous and something I could always depend on. I am incredibly grateful for the steadfast emotional support, both in and out of my educational pursuits, all the amazing people I have welcomed into my heart, and the wonderful experiences we got to share.

Thank you, Jana-Sophie, Janik, Lydia, Thabo, Marius, and Isabelle (special credit to her for reading my thesis and improving it with her excellent English skills) for always being there in any matters. Your presence was and is a constant source of security and warmth. Heartfelt thanks to Leonie, Sebastian, Levin, and Anestine for standing by me since the very beginning of my time as a cognitive science student and for the countless coffees and chats in the bistro and café. To Luca, Rhea, Elisa, Cedric, Lea, and Jakob, I express my gratitude for the fantastic times and support outside the realms of cognitive science. Last but not least, thank you for a great mentoring experience in the last months of my time as a student to my fellow mentors and meanwhile dear friends Lena and Fabienne.

I am deeply grateful for each and every one of you. Without your friendship and love, I would not be where or who I am today.

Abstract

Bayesian observer and actor models have provided normative explanations for behavior in many perception-action tasks including discrimination tasks, cue combination, and sensorimotor control by attributing behavioral variability and biases to factors such as perceptual and motor uncertainty, prior beliefs, and behavioral costs. However, it is unclear how to extend these models to more complex tasks such as continuous production and reproduction tasks, because inferring behavioral parameters is often difficult due to analytical intractability. Here, we overcome this limitation by approximating Bayesian actor models using neural networks.

Because Bayesian actor models are analytically tractable only for a very limited set of probability distributions, e.g. Gaussians, and cost functions, e.g. quadratic, one typically uses numerical methods. This makes inference of their parameters computationally difficult. To address this, we approximate the optimal actor using a neural network trained on a wide range of different parameter settings. The pre-trained neural network is then used to efficiently perform sampling-based inference of the Bayesian actor model's parameters with performance gains of up to three orders of magnitude compared to numerical solution methods.

We validated our proposed method on synthetic data, showing that recovery of sensorimotor parameters is feasible. Importantly, individual behavioral differences can be attributed to differences in perceptual uncertainty, motor variability, and internal costs.

We finally analyzed real data from a task in which participants had to throw beanbags towards targets at different distances and from a task in which subjects needed to propel puck to different target distances. Behaviorally, subjects differed in how strongly they undershot and overshoot different targets and whether they showed a regression to the mean over trials. We could attribute these complex behavioral patterns to changes in priors because of learning and undershoots and overshoots to behavioral costs and motor variability.

Taken together, we present a new analysis method applicable to continuous production and reproduction tasks, which remains computationally feasible even for complex cost functions and probability distributions.

Contents

1. Introduction	7
1.1. Overview	10
2. Background	11
2.1. Perception & Action	12
2.1.1. Psychophysics and the Weber-Fechner Law	12
2.1.2. Perception as Inference	14
2.1.3. Interdependence of Perception and Action	14
2.1.4. Perceptual Uncertainty and Motor Uncertainty	15
2.1.5. Response Patterns in Human Behavior	16
2.2. Bayesian Approaches	18
2.2.1. Overview	18
2.2.2. Ideal Observer Model	20
2.2.3. Bayesian Actor Model	21
2.3. Experimental Paradigms in Decision-Making	23
2.3.1. Yes-No, 2AFC and 2IFC Tasks	23
2.3.2. (Re-)Production Tasks	23
2.4. Inference Methods for Behavioral Parameters	24
2.4.1. Identifiability Problems	25
2.5. Bayesian Approaches in Psychophysics and Control	26
2.5.1. Similar Approaches	26
2.5.2. Applications	27
3. Methods	29
3.1. Observer Model	34
3.2. Cost Function and Optimization	36
3.2.1. Cost Function Parameterization	36
3.2.2. Optimization	37
3.2.3. Analytical Solution	39
3.2.4. Solvability	42
3.3. Neural Network	44
3.3.1. Architecture	44
3.3.2. Data, Training and Evaluation	45
3.4. Parameter Inference	49
3.4.1. Probabilistic Model	49
3.4.2. Prior Predictive Checks	51
3.4.3. Sampling Process	51

4. Results	53
4.1. Synthetic Data	53
4.1.1. Numerical Versus Analytical Results	55
4.1.2. Recovery of Characteristic Parameter Configurations	57
4.1.3. Inference of Sensory Information	61
4.1.4. Recovery of Parameters from Randomly Generated Data Sets	62
4.1.5. Influence of the Number of Trials on the Results	64
4.2. Experimental Data	66
4.2.1. Bean Bag (BB) Data Set	66
4.2.2. Puck (PU) Data Set	71
5. Discussion	76
5.1. Constraints and Limitations	77
5.1.1. Identifiability Problems	78
5.2. Theoretical Grounding of the Framework	80
5.3. Future Work	81
6. Conclusion	82
A. Appendix	93

1. Introduction

As humans navigate through the complexities of everyday life, they heavily rely on their perception and action skills. Even simple tasks, like grabbing an object, require precise capturing of an object's position and a carefully planned and executed trajectory to finally reach for it.

In experimental psychology, a lot of these tasks can be formalized as so-called production or reproduction tasks (Wixted and Thompson-Schill, 2018). That is, tasks where subjects need to produce an action to hit a target or reproduce an action on a continuous scale. For instance, this could be simply navigating to a certain goal (Mittelstaedt and Mittelstaedt, 2001; Harris et al., 2000) or listening to a sound and holding down a button for the same duration in order to reproduce it (Fraisse, 1984; Buhusi and Meck, 2005). These tasks belong to the scope of continuous decision-making, since subjects need to actively decide for a certain action in continuous time and space.

One famous example of a production task is the story of Wilhelm Tell. For his disobedience towards the regional governor, he receives the punishment to shoot an apple off his son's head with a crossbow. In the context of a production task, this can also be divided into a perception and an action stage. When facing this challenge, Wilhelm needs to identify the apple on his son's head as the target he needs to calibrate his action to. Since we are all subject to the limitations of our sensory systems, Wilhelm will only be able to obtain a noisy visual estimate of the true position of the apple (Van Bergen et al., 2015). He then needs to plan his action according to this noisy estimate of the apple's position by also incorporating prior knowledge about the world – for example, he will approximately know the height of his son and can get a good estimate of the apple's position from that without any visual information.

An established approach to this action planning and selection assumes that people employ cost functions that assign costs to different actions and their respective outcomes (Cohen et al., 2010; Körding and Wolpert, 2004b). Depending on the relationship with his son, Wilhelm will most likely assign high costs to actions that result in missing the apple and shooting his offspring instead. Lesser costs will be assigned to missing the apple but also not hitting his son and he might even gain a reward instead of paying a cost for hitting the apple without harming his child. The physical effort or mental stress that certain actions put on his body and mind will most likely also play a role and add to the cost of hitting a certain object (or body). For action selection, Wilhelm will try to minimize the sum of these costs. Even when carefully planning his action, he will most likely not hit the exact point he was aiming at, since his muscles, joints and tendons and the neural activations that actuate them exhibit natural, intrinsic noise (Van Beers et al., 2004; De C. Hamilton et al., 2004) and he only had a noisy estimate of his target to aim for.

People's actions are substantially influenced by the forms said cost functions can possibly take on. In the context of psychological experiments, these cost functions are usually defined by the experimental design the researcher's choose. In the case of Wilhelm, his situation and the resulting cost function will probably introduce a tendency to overshoot the target, since he rather wants to shoot above the apple instead of hitting his son.

In contrast to this, a majority of experimental research results show that animals (Lejeune and Jasselette, 1986; Lejeune and Richelle, 1982; Lowe et al., 1979) and people (Willey and Liu, 2018; Harris, 1995; Bergmann

et al., 2011; Elliott et al., 2010, 2004; Engelbrecht et al., 2003; Sun et al., 2004b,a; Zeiler and Hoyert, 1989; Becker and Fuchs, 1969; Weber and Daroff, 1971) are often likely to undershoot targets in the context of perception-action-tasks.

To explore the origins of these behaviors and ultimately understand human behavior holistically, it is crucial to have a reliable estimation and representation of the costs, prior beliefs, uncertainties and the motor variability that actually shape human behavior. The models that represent human perception, decision-making and action became increasingly complex over the last two centuries.

The modern approach in modeling perception is to pose it as an inference problem over the true state of the world, given only limited sensory information – for instance, when Wilhelm combined knowledge about the height of his son and his visual information to get an estimate of the apple’s true position. The development of the research work within this domain led to Bayesian models being the dominant modeling approach to human perception, known as so-called *ideal observers* (Knill and Richards, 1996; Ernst and Banks, 2002; Kersten et al., 2004; Kersten and Mamassian, 2009).

Extending from the perception model, we need to also consider the decision-making process and the resulting action together with the perception stage. This is an increasingly prominent view in the modeling of human behavior and decision-making (cf. the idea of *affordances* by Gibson (1966)). Perception is pointless without any consequences that can be taken and action without perception does not work. For instance, it has been shown that neural evidence from perception is directly transmitted to the motor system (Hagura et al., 2017) and that perceptual uncertainty (Schlicht and Schrater, 2007; Barthelmé and Mamassian, 2009; Van Bergen et al., 2015; Hagura et al., 2017; Knill and Pouget, 2004) and likewise action uncertainty (Trommershäuser et al., 2003; Harris and Wolpert, 1998) are incorporated into decision-making for movements. This led to the emergence of so-called *Bayesian actor models* that extend the ideal observer to more complex decision-making and acting within the agent’s environment (Körding, 2007; Trommershäuser et al., 2008; Wolpert and Landy, 2012; Acerbi et al., 2014). They also implement the previously introduced cost functions. Moreover, they implement a probabilistic account to perception and behavior by incorporating probability distributions describing how the subject perceives their environment and how they generate their actions.

Yet, most models and experimental paradigms only consider discrete stimulus and action spaces, since understanding and modeling human behavior becomes increasingly hard when moving into continuous space because the number of stimuli a subject can be exposed to and the number of actions they can choose from are not longer finite in continuous space.

Challenged by the complexities of such continuous models, most researchers constrained themselves to simple cost functions (e.g. quadratic) and probability distributions (e.g. Gaussian) used in their models (Körding and Wolpert, 2004b). Thereby, the model remains analytically tractable and can be used in all kinds of computations. Yet, it is not clear whether quadratic costs and Gaussian distributions describe the aspects of human perceptuomotor behavior adequately. Cost functions may be better approximated by functions that are more robust to outliers (Körding and Wolpert, 2004b) or distributions may be grounded in psychophysical (Weber, 1831; Dehaene, 2003; Portugal and Svaiter, 2011) or information-theoretic (Wei and Stocker, 2015, 2017) laws. To incorporate these theories into models, the actor model usually loses its analytical tractability and thus can only be solved via numerical methods. However, numerical methods are usually computationally expensive and hence slow, limiting the models’ use.

When modeling human behavior, the underlying parameters that generate the behavior observed in a subject are usually unknown. They can be recovered using, for instance, sampling-based methods like *Markov Chain Monte Carlo (MCMC)* methods (Hastings, 1970; Robert and Casella, 2004). These methods rely on sampling parameter proposals and solving the model with different proposals thousands of times to guarantee statistical convergence and thereby obtain a probability distribution of the model’s underlying parameters. With arbitrary

cost functions and probability distributions that need to be solved numerically, these methods cannot be used anymore, since solving simply takes too long.

We propose a novel approach for approximate Bayesian inference in actor models by approximating the decision-making process with a pre-trained neural network, inspired by Neupärtl and Rothkopf (2021). This network treats the perception & action problem from the perspective of the subject with no access to the true state of the world. Here, we also set some high-level normative assumptions on how the subject's perception and decision-making work, for instance by introducing a cost function which the subject tries to minimize. To infer the parameters underlying the subject's behavior, we need another probabilistic model where we integrate said network to perform inference. This model is designed from the researchers' perspective who have no access to the measurements of the subject's sensory system but can measure the true state of the world and the responses of the subject. This model solely serves the analysis and inference of behavior, not imposing any normative constraints on how the subject's behavior is implemented in the brain.

The neural network receives the parameters the model should be solved for and returns the parameters of an action distribution that ultimately generates the subject's actions. It is trained on a large data set consisting of numerically solved configurations of the model. This speeds up solving the model massively and keeps it computationally feasible while allowing for arbitrary cost functions and probability distributions. Thereby, researchers are free of any computational constraints in their modeling choices and the model can be integrated into any desired use case.

In our model, we employ a log-normal distribution instead of a Gaussian distribution in the perception stage of our model, grounded in the Weber-Fechner law. The cost function in our model is defined by a set of our parameters that allow for functional forms beyond quadratic costs and are inferred along with other behavioral parameters. The action generation in our model is also implemented with a log-normal distribution as well to incorporate signal-dependent noise (Harris and Wolpert, 1998).

We evaluated our method on synthetically generated data as well as on empirical data from (re-)production tasks (Willey and Liu, 2018; Neupärtl et al., 2020) showing characteristic response patterns of human behavior (Poulton, 1968; Petzschnier et al., 2015), e.g. pervasive undershoots in these kinds of tasks. Our method recovers perceptual uncertainty, motor variability, cost function parameters and prior beliefs reliably on all data sets. We created a small set of analytically tractable data sets to compare the performance of our approximate model with the results of a model using the analytical solution. The inferred posteriors of our method are close to the ones obtained with an analytical solver. In a large-scale evaluation on randomly generated data sets, the posterior 95% HDIs contain the true value 80-90% percent of the time.

In experimental data from a bean bag throwing (Willey and Liu, 2018) and a puck sliding task (Neupärtl and Rothkopf, 2021), our method infers parameters reliably reproducing the observed behavior.

Still, our models exhibit some weaknesses as it shows statistical issues like eventually non-converging MCMC chains and questionable MCMC statistics. We treat the internal representation of each stimulus as latent sensory variables that are also inferred, thereby making the inference problem high-dimensional and thus really hard, potentially being the cause for said issues.

We propose a novel method increasing researcher's freedom in their modeling approaches for human behavior in perception-and-action tasks while keeping these complex actor models solvable in reasonable time with neural networks. Our method infers parameters on synthetic and empirical data that reliably produces the observed behavior. Still, there are some minor issues with the statistical properties of our framework. We plan to resolve these deficiencies in the future and thus provide a tool that aids researchers in testing various theories in the interplay of human perception, decision-making and action.

1.1. Overview

Section 2 introduces all necessary background information related to this work. It starts with the roots and interplay of perception and action (Section 2.1), introduces modern Bayesian approaches (Section 2.2) as well as experimental paradigms in (sensorimotor) decision-making (Section 2.3). The fundamentals of inference of behavioral parameters are explained in Section 2.4 and Bayesian approaches especially in a sensorimotor context and their applications are highlighted in Section 2.5.

Our approach is introduced in Section 3. It gives a detailed overview over the different components of our framework, from the cost function and optimization process (Section 3.2) over architecture, training and evaluation of the neural network (Section 3.3) to the implementation of the Markov Chain Monte Carlo inference (Section 3.4).

Section 4 presents the results of our method, applied to synthetically created data in Section 4.1 and to experimentally collected data in Section 4.2.

The limitations and potential of our framework is discussed in Section 5. First, its constraints are reviewed in section 5.1. An overall discussion of the theoretical grounding and implications of our Framework in the Bayesian context is provided in Section 5.2, before Section 5.3 closes with an overall outlook on eventual future work.

2. Background

Given the fact that our work spans a wide range of different interrelated fields of research, this chapter is intended to provide a short overview of important concepts and relevant studies.

Section 2.1 starts with the history of research regarding the perception of physical stimuli, known as *psychophysics*. We then relate this research to how humans plan their actions based on their sensory perceptual input and how perception and action depend on each other. Since it has become a common view that perception not only depends on sensory input alone, but also incorporates prior beliefs about the environment of the subject, Section 2.2 introduces Bayesian approaches to modeling perception and action. Bayesian models are capable of representing prior knowledge and how it is combined with sensory evidence. This is modeling from the subject's perspective. Furthermore, Bayesian approaches allow for probabilistic modeling involving distributions of various quantities, providing us with measures of uncertainty and not only scalar estimates of different parameters. This is modeling from the researcher's perspective. Bayesian models on the subject level and the researcher level are independent of each other – one can use Bayesian or non-Bayesian methods alike on both levels. We use Bayesian approaches for the model of the subject as well as for the inference of the parameters of the model. Section 2.3 introduces the most popular experimental setups for measuring performance and collecting behavioral data in psychophysical and perception-and-action tasks. How these observed data can be used by inferential methods to recover the underlying parameters of a probabilistic model, especially in modeling human behavior, is explained in Section 2.4. This chapter concludes with a short overview of related work and different use-cases in Section 2.5.

2.1. Perception & Action

When humans act and solve tasks in the world, the resulting action is the product of a pipeline that spans from (1) perception over (2) decision-making in planning to (3) motor execution. Since all of these stages are interconnected, it makes sense to always consider the actions of subjects with regard to the foregoing perception process.

This Section starts with the early history of perception models (Section 2.1.1) and how they were extended to general inference about the world and how they incorporate prior knowledge instead of solely relying on information streams of the human sensory system (Section 2.1.2). The general interdependence of perception and action is addressed in Section 2.1.3 and the role of perceptual uncertainty and motor variability for human behavior and thus also for computational models of behavior is introduced in Section 2.1.4.

2.1.1. Psychophysics and the Weber-Fechner Law

We undergo a plethora of sensory stimuli in our everyday life. How we act in the world crucially depends on our internal model of the outside world, which we can only access via our sensory system. For example, to grab a cup, we first need to determine its position and shape, plan the grabbing movement and execute it accordingly. How we perceive the physical world and represent it in our mind is captured by the field of *psychophysics*. One of the first to investigate the quantitative relation between real stimulus magnitudes and their perceived intensities was Gustav Theodor Fechner, who proposed the so-called *Weber-Fechner Law*. The Weber-Fechner Law is a composite of the Weber law and the Fechner law. We will provide a short introduction on them here, since they are fundamental and widely discussed laws in psychophysics research (Weber, 1831; Fechner, 1860). Weber's law was one of the first quantitative approaches to perceived stimulus magnitude in humans. It states that the *Just-Noticeable Difference (JND)* between two stimuli is proportional to the stimulus magnitude S . The JND describes the change that needs to occur in a stimulus S to be recognizable by a subject. The ratio between the JND dS and the stimulus intensity S is proposed to be constant by Weber's law. This also means that bigger stimuli are much harder to distinguish than stimuli at lower intensities. The JND dS can also be described by the relation shown in equation (2.1), where K is a constant.

$$dS = K \cdot S \quad (2.1)$$

Fechner then extended this law to the so-called Weber-Fechner law in equation (2.2). By introducing the perceived intensity p , we obtain a variable describing the subjective intensity of a physical stimulus S . Fechner also assumes that this perceived quantity p is equal to zero for some threshold S_0 . Again, k is a constant and dependent on the sensory modality – i.e. vision, hearing, tactility, etc. – and the type of stimulus.

$$p = k \cdot \ln \frac{S}{S_0} \quad (2.2)$$

While the Weber-Fechner law seems to account for psychophysical phenomena quite well within the middle ranges of different stimuli (Roberts, 2006; Yi, 2009), it fails at extreme stimulus values, i.e. either very low or very high stimulus intensities (Krueger, 1989).

Some studies even go beyond psychophysics following a logarithmic pattern and claim that the mental number line in the mind is overall organized logarithmically (Dehaene, 2003; Longo and Lourenco, 2007; Scheler, 2017), advocating for evolutionary optimality of a logarithmic mental scale following the Weber-Fechner law (Portugal and Svaiter, 2011). Laski and Siegler (2007) propose that the representation of numbers and magnitudes in the mind is naturally logarithmic but is overwritten when getting used to mathematics and simple algebra, since it is historically established in the Western world to perform mathematical operations on a linear number scale. For example, especially children use rulers a lot in their early development in kindergarten and elementary school. The ruler is one of the tools that shows most prominent how we are getting used to linear number scales in our early development already. Dehaene et al. (2008) show how number perception is different and organized much more linearly in Western than in Indigenous cultures, even though they received some criticism by Cantlon et al. (2009), who stated that their results are not as clear as they claim.

As an alternative to the Weber-Fechner law, Stevens (1961) formulated Stevens' Power Law, which assumes that the relation between perceived stimulus intensity p and the stimulus S is not logarithmic but follows an exponential relation. This is shown in equation (2.3), where the exponent a depends on the stimulus type and the constant k depends on the units in which the stimulus is measured.

$$p = k \cdot S^a \quad (2.3)$$

Yet, Stevens' Power law is generally criticized for being nothing more than merely a good fit of data points averaged over multiple subjects without any underlying assumptions or explanatory power.

Some researchers argue that both laws do not contradict each other but are instead based on different conceptualizations of how sensation can be measured and were applied to different experimental designs (discrimination versus matching tasks) (MacKay, 1963; Shepard, 1981). Meanwhile, other researchers even go so far as to claim that both Weber's law and Steven's law are incorrect and propose a new relationship, also based on a power function but parameterized differently than Stevens' law (Krueger, 1989).

Poulton (1968), one of Stevens' students, criticized Stevens' law for neglecting several other influences in magnitude estimation tasks, since he found, for instance, the characteristics of the stimulus distribution and other parameters of the experimental design to influence subjects' performance.

The modern trend within psychophysics research argues that perception based on the Weber-Fechner law is not sufficient, for example since it fails for extreme stimulus values. Recent research rather relies on Bayesian principles that combine sensory evidence following the Weber-Fechner'ian concept with prior knowledge about the environment in order to produce the perceived stimulus intensity (Pizlo, 2001). These frameworks can even show that the Weber-Fechner law is Bayes-optimal in information-theoretic terms (Sun et al., 2012). The Bayesian approach to psychophysics is also able to account for various phenomena arising in magnitude estimation tasks across different sensory modalities (Poulton, 1968; Petzschner et al., 2015). These phenomena can be classified into different characteristic biases which can be observed in empirical data. Please refer to Section 2.1.5 for a detailed description of these patterns.

2.1.2. Perception as Inference

As introduced in the previous Section, the Weber-Fechner law describes the relationship between the physical and the perceived intensity of stimuli quite well, at least in moderate ranges of stimulus values. Yet, this is not sufficient to capture the perceptual process holistically. Perception is more than just the mapping of a physical quantity onto a subjective, sensory space. Rather, it is complex reasoning about what the world really looks like.

Helmholtz (1867) already described perception as *unconscious inference* about the real world, an idea that still lives on today (Rock, 1983). We need more information than just the sensory information of the world currently surrounding us in order to perform inference, especially since this information does not represent the ground truth. (Knill and Richards, 1996; Clark and Yuille, 2013)

Because of the limitations of the human sensory system, this sensory information is naturally noisy and ambiguous. For example, consider the visual system: The resolution of retina images is limited by the amount of photoreceptors on the retina. In addition to this sensory noise, there is also neural noise arising from the dynamic electrical activity of neurons in the brain. Also, the same two-dimensional retinal image can be produced by different three-dimensional geometries of objects and thus also challenges the human sensory system with ambiguity. To overcome this problem, humans need to incorporate their prior knowledge about what the world usually looks like and update their estimate of the state of the environment by combining this prior knowledge with their current sensory image of the world. Besides, they need to have an additional understanding for different concepts and categories in the world (Ashby and Maddox, 2005). Nevertheless, we still need to know about the ‘usual’ structure of the world to get a sense of scene understanding. Consider the case that we receive a mental image from our sensory system that contains no noise and accurately represents the real world. We can see edges and colors within this picture, but we do not gain any new knowledge about our current situation – without prior knowledge about the world, we cannot derive any semantic meaning from this image. In reality, humans have this prior knowledge and it contains several facts and rules about the real world, for example that natural object surfaces usually tend to have smooth, continuous changes within their brightness. This and other rules about color, reflectivity of surfaces or other natural statistics of the world were most likely acquired by evolution (Simoncelli, 2003). Leveraging this knowledge, edges and other perceived properties in the mental image can be combined to higher-level features that form objects. Consider a visual scene containing a vertically-oriented brown structure, which allows us to possibly infer by the texture and color that it must be made of wood. If we now combine the vertically oriented edges and the smooth flat surface on top of them, as well as our previous experience that furniture is mostly made of wood, we might come to the correct conclusion that it is not just an image of color and edges, but actually a depiction of a wooden chair. Without this semantic knowledge, we would only be able to see the structural information of an image but could not attribute any meaning to it. This concept of combining previous knowledge with sensory information is captured perfectly by Bayesian statistics (cf. Section 2.2).

2.1.3. Interdependence of Perception and Action

The information about the state of the world obtained by a subject’s sensory system propagates through a pipeline from the perceived inputs (perception stage) over the decision-making forming the final response (planning stage) to ultimately execute the response of an agent (action stage). Thus, it makes sense that perception and action always need to be considered together, which is also backed up by recent studies. For instance, Schlicht and Schrater (2007) have shown that visual uncertainty influences the planning of grasping movements. Namely, they measured the grip aperture – i.e. the distance between the index finger and thumb grabbing the target object – during grasping movements and observed a positive correlation between grip

aperture and visual uncertainty. Hagura et al. (2017) found evidence that suggests that action costs even influence our perception, since the respective motor costs for different movements in a reaching task biased the subjects towards the less costly options. This is further backed up by neural evidence that information about uncertainty in our sensory system is directly transmitted to the motor system, where said uncertainty influences the planning in action control (De Lange et al., 2013).

Considering actions mutually with sensory information is also backed up from a technical perspective, since state-of-the-art approaches in *control theory* also continuously integrate sensory inputs into their planning and action execution. *Optimal control* is the domain of control theory that tries to find movement trajectories minimizing some time-continuous cost function (Lewis et al., 2012). For the most part, models for human behavior in perception-and-action tasks grounded in control theory involve a closed feedback-loop, connecting inputs (percepts) and outputs (actions) (Todorov and Jordan, 2002; Wolpert and Ghahramani, 2000). We can also observe the same combination of noisy information from prior beliefs and sensory information in these technical models. For example, the *Kalman Filter* is a method for estimating the true state of the world based on prior knowledge from the internal model of the world (prior beliefs) and sensory evidence (Kalman, 1960; Wolpert and Ghahramani, 2000), which is also often used in Bayesian models of human behavior (Vilares and Kording, 2011).

2.1.4. Perceptual Uncertainty and Motor Uncertainty

As introduced in the previous Sections, there are multiple sources of uncertainty and noise within the process of perceiving an stimulus, planning a response to that stimulus and ultimately executing the action corresponding to that response. In order to produce behavior in a noisy system that succeeds in solving the task at hand, humans need to have an estimate of their perceptual uncertainty and motor variability (Knill and Pouget, 2004). This estimate must not necessarily be consciously accessible, but at least be implicitly involved in the planning process of the movement.

There is a plethora of studies regarding whether and how well humans can access their own perceptual and motor uncertainty, especially in (re-)production tasks. How well humans can incorporate an internal estimate of their own uncertainty into their planning processes can easily be probed by designing experiments with an asymmetric cost function. If the cost for one of two possible responses is much greater than for the other response, the less-costly response should proportionally be chosen more often the more uncertain the subject is (Whiteley and Sahani, 2008). In metacognitive approaches, the subjects are directly asked about their uncertainty (Barthelmé and Mamassian, 2009; De Gardelle and Mamassian, 2015). Even more evidence can be found in other studies, mostly in tasks regarding visual perception (Körding and Wolpert, 2004a; Sims and Gray, 2008). This is further backed up by neuroscience, which shows that there are populations of neurons – called *neural population codes* – in the human brain that encode the uncertainty about a quantity or even its probability distribution as a whole (Van Bergen et al., 2015; Zemel et al., 1998; Beck et al., 2008). Even the confidence about this uncertainty seems to be encoded in the brain, since people with orbitofrontal lesions are not able to assess the risk associated with different responses (Hsu et al., 2005).

The hypothesis that humans have access to their own uncertainties is also evident when we investigate tasks that require multimodal integration – i.e., tasks that need to be solved by relying on cues from different sensory domains, for example of visual and haptic nature. The statistically optimal behavior would attribute less weight to the more uncertain sensory inputs and rely more heavily on cues with the lowest uncertainty – a process that has been shown to be employed by humans in multiple studies (Ernst and Banks, 2002; Ernst

and Bühlhoff, 2004; Bresciani et al., 2006; Hillis et al., 2004; Fetsch et al., 2012). The very same concept also applies to the ratio of variability in the prior and the likelihood. Depending on the variability within both distributions, the subject should depend more on the prior or likelihood, whichever of both has lower variance – at least when assuming that the subject is statistically optimal. That is, the subject needs to rely more on the sensory information arising from their observations if they do not have a strong prior. Vice versa, they should rely more on their prior beliefs if they can only make noisy and possibly ambiguous observations. (Mamassian and Landy, 2001; Tassinari et al., 2006)

Furthermore, variability does not arise from perceptual uncertainty alone. As soon as the subject has a reliable representation of the world around them, they need to act upon it. When humans want to execute a movement, they are subject to motor noise which prevents them from executing the movement exactly as planned. This noise is primarily induced by the physiological constraints of joints, tendons and muscles - Van Beers et al. (2004) have shown that the noise in movement trajectories seems to stem completely from the execution of the movement itself and not from the planning process. The strength of the noise can be shown to be dependent on the joint strength and to correlate with the number of motor units within the muscles, which also correlates with muscle strength (De C. Hamilton et al., 2004). Moreover, motor noise was found to be signal-dependent noise, which means that the amount of noise within a movement scales with the extent of the movement (Harris and Wolpert, 1998; Jones et al., 2002; De C. Hamilton et al., 2004). Larger movements, or movements that require more strength, are subject to bigger amounts of noise, a fact that was important to consider within our model.

2.1.5. Response Patterns in Human Behavior

Research in the area of continuous responses, for example via the (re-)production paradigm that will be introduced in Section 2.3.2, revealed characteristic behavioral patterns of subjects in psychophysical tasks. The following section will provide a brief summary of these tendencies and biases.

Poulton (1968) described six different effects that the experimental design and other external conditions have on the behavior elicited in a subject. Petzschner et al. (2015) narrowed these six effects down to the four most prominent behavioral phenomena, which we will introduce here (cf. Figure 2.1). In the following paragraph we assume that subjects needed to produce a response on a continuous scale and that the target stimulus is drawn from a stimulus distribution of a certain range during the experiment.

(1) The so-called *regression effect* or *regression-to-the-mean effect* describes the tendency of subjects to produce responses that are biased towards the mean of the stimulus distribution. This becomes apparent in overshoots for stimuli below the stimulus mean and undershoots for stimuli above the stimulus mean. (2) The *range effect* is a property of the regression effect, which states that the regression effect becomes stronger for broader stimulus distributions. Therefore, the regression and range effect should always be considered in tandem (Teghtsoonian and Teghtsoonian, 1978).

(3) The responses of the subjects exhibit *scalar variability*, i.e. the variability in the responses scales with the mean of the response (Harris and Wolpert, 1998; Jones et al., 2002; De C. Hamilton et al., 2004). This was already introduced in Section 2.1.4 as *signal-dependent noise* and is in line with the logarithmic relationship between stimulus and sensation formulated by the Weber-Fechner law (cf. Section 2.1.1).

(4) *Order effects* and general effects of subsequent responses can be observed in some kind of *hysteresis effect* (Eisler and Ottander, 1963; Hock et al., 2005). That is, responses after a large stimulus or response tend to be larger than after a small(er) stimulus and vice versa.




Figure 2.1.: **Behavioral Phenomena in Magnitude Estimation and (Re-)Production Tasks** - Behavioral effects can be summarized as the regression and range effect, being subject to scalar variability and showing a hysteresis effect for subsequent tasks. Figure taken from Figure 1B in Petzschner et al. (2015).

Furthermore, there is a substantial amount of literature showing undershoots for a large variety of different tasks where subjects needed to produce an action on a continuous scale. This undershoot behavior could be observed for treadle pressing tasks (Lejeune and Richelle, 1982; Lejeune and Jasselette, 1986) and time reproduction tasks (Zeiler and Hoyert, 1989) in pigeons and for time estimation tasks in pigeons and rats (Lowe et al., 1979). In humans, this behavior was observed in tasks where objects needed to be placed relative to a target stimulus (Elliott et al., 2004), in tasks involving reaching movements (Elliott et al., 2010), in locomotion and driving a bicycle (Bergmann et al., 2011; Sun et al., 2004b,a) or when controlling objects on a computer screen (Engelbrecht et al., 2003). Research in active vision extended this list by identifying the same undershooting pattern in saccades (Becker and Fuchs, 1969; Harris, 1995; Weber and Daroff, 1971).

The phenomena introduced here can be explained quite well by using Bayesian models (Petzschner et al., 2015), for example by showing that the regression and range effect found in empirical data can be reproduced by showing a strong reliance on the prior, assuming that the prior was learnt during the experiment to approximately represent the real stimulus distribution. Therefore, a Bayesian framework was chosen as a baseline for our work (cf. Section 3).

2.2. Bayesian Approaches

As introduced in Section 2.1.2, perception does not solely depend on sensory information but also on prior knowledge and beliefs about the environment the subject is situated in. This can be posed as an inference problem on the true state of the world, which can be formalized using Bayesian concepts. We provide a short overview over Bayesian models in Section 2.2.1 and introduce the *ideal observer* and *Bayesian actor* model in Sections 2.2.2 and 2.2.3, respectively.

2.2.1. Overview

Bayesian approaches are a prominent tool for modeling a variety of different cognitive processes, since they provide a simple method for fusing bottom-up sensory information with top-down prior knowledge. Sufficient inference about the structure of the real world is only possible with both of these information streams. Bayesian models first were first introduced in computer vision (Geman and Geman, 1984; Bolle and Cooper, 1984) and were soon extended to computational models of human perception (Knill and Richards, 1996; Clark and Yuille, 2013), for example in contour detection (Geisler et al., 2001) or the perception of visual motion (Simoncelli, 1993).

Logically, an agent embodied in the real world wants to act within this world. Thus, Bayesian models were soon extended to not only account for perception but also motor control, thereby accounting for the whole area of *sensorimotor control*. Perception and action are connected by the decision-making and planning processes that produce an action based on sensory evidence and prior beliefs (Körding and Wolpert, 2004a; Körding, 2007; Tassinari et al., 2006).

Other approaches were applied to general cognition (Griffiths et al., 2008), e.g. for modeling predictions based on subjects' experiences and their memory of various natural statistics (Griffiths and Tenenbaum, 2006) or how humans attribute causality in relations between different entities (Sloman and Lagnado, 2015). Other probabilistic models implementing Bayesian concepts are applied in the domain of language processing and acquisition (Chater and Manning, 2006).

There is also support for the hypothesis that the human brain itself carries out Bayesian computations. For example, Ma et al. (2006) proposed an explanation on how the brain can neurally implement Bayesian inference processes encoded neural population codes.

Still, Bayesian models received a lot of criticism for various reasons (Bowers and Davis, 2012) and their real contribution to explaining human behavior is part of an ongoing debate (Jones and Love, 2011). A lot of this criticism is directed towards studies that claim to prove that the brain implements Bayesian computations. While it is true that there is a lack of clarity in most literature, we need to distinguish between descriptive and normative roles of Bayesian models. While descriptive models are used for data analysis without making any claims about the algorithmic or implementational details underlying the decision-making process in the human mind (Kruschke, 2015; Tauber et al., 2017), normative models propose how humans *should* behave to be statistically optimal and assume that humans' mental processes actually follow Bayesian principles (Knill and Pouget, 2004; Van Den Berg et al., 2012). Another argument against Bayesian approaches and models is that Bayesian models offer too many degrees of freedom in the selection of the likelihood or the prior distributions and therefore lack explanatory power. This was defended again by Griffiths et al. (2012), for example by arguing that the choice of priors is not arbitrary but actually guided by knowledge about how human beliefs and mental processes in the respective domain are structured (Jones and Love, 2011) or

by the natural, environmental statistics of the quantity in question (Simoncelli, 2003). Another solution to resolve this conflict is to not assume statistical optimality in humans but to accept that they are sub-optimal to a certain degree in their decision-making and planning processes (Rahnev and Denison, 2018; Rosas and Wichmann, 2011) – several models represented empirical data more accurately when flaws were deliberately introduced into the model (Stengård and Van Den Berg, 2019). This could be grounded in an inherently noisy inference process in the human brain, thus introducing another source of variability besides perceptual uncertainty and motor variability (Beck et al., 2012; Drugowitsch et al., 2016; Faisal et al., 2008; Renart and Machens, 2014; Wyart and Koechlin, 2016).

We take an agnostic view to whether normative Bayesian models are an appropriate approach to model human behavior. Certainly, descriptive Bayesian approaches are a great toolbox for analyzing data without making any assumptions on the structure of the data or how it came about. We employ non-normative Bayesian methods for the inference of parameters and only include a normative Bayesian component for our perception model – however, we primarily model perception as inference due to the prominence of this approach and it could be substituted by any other perception model. Our framework is intended to aid resolving this conflict on the distinction between descriptive and normative Bayesian methods.

Some studies suggest that certain behavior cannot be explained by Bayesian principles at all (Brayanov and Smith, 2010) and termed this phenomenon ‘*anti-Bayesian*’. Later on, this was reconciled by new Bayesian models guided by information theory which explained these supposedly anti-Bayesian percepts quite well (Wei and Stocker, 2015) by employing encoder-decoder structures for encoding stimuli with respect to information-theoretic principles in the likelihood (Geisler et al., 2009) and grounding the prior belief of the subjects in the natural statistics of the environment (Simoncelli, 2003).

Despite this ongoing debate about the real explanatory power of Bayesian models, a lot of them offer explanations of how illusions (Weiss et al., 2002; Shams et al., 2005; Notredame et al., 2014) or systematic biases in human perception come about. For example, humans have a bias to assume that light sources are often placed above a scene (Ramachandran, 1988; Adams, 2008) or that objects are more likely of convex shape (Langer and Bühlhoff, 2001) or at least more accurately recognized if they had a convex shape (Burge et al., 2010). Moreover, humans seem to judge orientations of shapes more accurately along cardinal axes (horizontally and vertically), which could be explained by the prominence of cardinal orientations in the natural statistics of the world (Girshick et al., 2011).

Knowledge and Beliefs are not necessarily consciously accessible - The “knowledge” and beliefs we refer to here can easily be interpreted as explicit knowledge which is consciously available during perceptual processes. This is most likely not the case and usually not what is meant when referring to this world knowledge. It is rather part of an automatic process during perception and exists unconsciously. For example, the natural statistics of how edges are distributed in scenes in the real world (Girshick et al., 2011) are a form of implicit knowledge and can probably not be accessed consciously, but were shown to shape our interpretation of our perceptual input stream unconsciously all the time.

2.2.2. Ideal Observer Model

In Bayesian statistics, the previously introduced concepts of incorporating prior knowledge with our subjective sensory sensations are integrated within the so-called *Ideal Observer*. The observer is called ideal in the sense that it Bayes-optimally combines the available information – that is, it produces a prediction of the real world without being subject to any physiological constraints or psychological biases, purely on the computational level (Knill and Richards, 1996; Kersten et al., 2004; Kersten and Mamassian, 2009). This is based on *Bayes' Theorem* (Bayes, 1763), which is shown in equation (2.4). Generally, we want to infer the true state s of the world, given some sensory, latent input quantity z .

$$p(s | z) = \frac{p(z | s) \cdot p(s)}{p(z)} \quad (2.4)$$

The left-hand side of the formula is known as the *posterior*, which is a combination of the *likelihood* $p(z | s)$ describing the probability of observing some sensory information z and the *prior* $p(s)$, which represents a person's beliefs and knowledge about the world. The marginal probability $p(z)$ is called the *evidence* and serves as a normalizing constant to the posterior distribution $p(s | z)$.

An ideal observer chooses a response to a stimulus, in a manner that it is statistically optimal. That is, we have a task within the domain of psychophysics, where humans need to perceive some stimulus and then respond with their belief about the true magnitude of this stimulus. In ideal observer models, this is captured via a rule that chooses an action $a(z)$ as a response to the stimulus based on their posterior belief about the true state s , as shown in Figure 2.2. Formally, we assume a loss function $\mathcal{L}(a, s)$ in ideal observer models, that assigns a value to each action-state-tuple. We refer to this value as *loss* or *cost* interchangeably, because both terms are commonly used within the literature but mean the same: an abstract, numerical representation of how 'bad' a certain action a is in a certain state s . The resulting action is then chosen by minimizing the expected loss over all possible states in the posterior distribution, given the measurement z . The expected loss is computed by a convolution of the posterior distribution $p(s | z)$ with the loss function, as shown in equation (2.5). This formalization is further supported by a variety of studies from neuroscience that show how risk, expected utility and other properties of distributions like modality specific uncertainties can be encoded in the brain (Schultz et al., 1997; Knutson et al., 2001; McCoy and Platt, 2005; Knutson and Peterson, 2005; Kable and Glimcher, 2007; Körding, 2007; Hosokawa et al., 2013).

$$a(z) = \underset{a}{\operatorname{argmin}} \mathbb{E}[\mathcal{L}(a, s)] = \underset{a}{\operatorname{argmin}} \int_{-\infty}^{\infty} \mathcal{L}(a, s) p(s | z) ds \quad (2.5)$$

The loss can also contain additional factors influencing human behavior. For example, when reporting a response with different buttons, of which one frequently gets stuck, the cost for pressing this button will be higher because of the discomfort it exerts on the subject. Equivalent to minimizing the loss, some models are formulated with the intended goal of maximizing the expected gain. We are using a cost-minimization approach here, but it could also be transformed into a gain-maximization by using the negative cost values as gain. (Körding and Wolpert, 2004a; Trommershäuser et al., 2006; Izawa et al., 2008)

The functional form of the loss function then determines how the resulting action a is chosen from the posterior distribution. For example, when assuming a quadratic loss, the mean of the posterior distribution minimizes the expected loss. When assuming a loss that is low when hitting the true target and uniformly high about all other actions, the so-called *maximum a-posteriori (MAP)* value, which is equal to the mode of the posterior distribution, minimizes the expected loss. While a quadratic functional form of the loss is one of the most

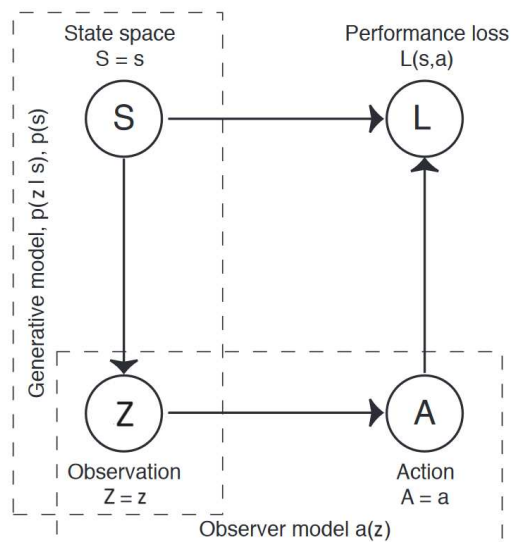


Figure 2.2.: **Ideal Observer Model** - The true state s of the world cannot be observed. The subject only has access to a measurement z which is generated according to a generative model consisting of a prior $p(s)$ about the distribution of possible world states and the likelihood $p(z | s)$ which specifies how the measurement comes about. Based on the measurement z the observer chooses an action $a(z)$ such that the loss $L(s, a)$ becomes minimal. Taken and modified from Figure 1A in Kersten and Mamassian (2009).

prominent loss function types, some research suggests that this does not hold for bigger deviations from the optimal response. Körding and Wolpert (2004b) have shown that in a pointing task the loss is quadratic for small errors but behaves smaller than quadratic for bigger errors. They propose that, for instance, an inverse Gaussian curve could be a better fit than a quadratic curve, motivated by its higher robustness towards outliers.

2.2.3. Bayesian Actor Model

The exact boundaries of the Ideal Observer Model introduced in Section 2.2.2 are not well-defined. Since the ideal observer is usually defined involving a response and an associated loss for each possible response, it is not only an observer but already includes some kind of action in the world. Yet, to draw a clear line between the term of the ideal observer, which we use for perceptual tasks, and models that explicitly include more complex planning regarding an agent's actions, we introduce the term *Bayesian Actor Model*. This is especially important in regard to our own model (cf. Section 3), which consists of an ideal observer for the perception-stage of the task but extends its action selection to a more complex optimization over the parameterization of an action distribution that ultimately yields the subject's responses.

The Bayesian actor also breaks with another assumption that most ideal observer models come with. Usually, stimulus and response of the ideal observer lie in the same domain. For instance, subjects need to perceive a target distance and then somehow reproduce this distance by, e.g., throwing a puck (cf. *reproduction tasks*, Section 2.3.2). Thereby, stimulus and action are both a spatial quantity, namely a distance - but this does not necessarily need to be the case. Consider a model that does not measure the subjects' actions as the produced distance but as the time holding down a button to control the effort applied to slide a puck towards the

target. To transform the duration of the button press from the temporal to the spatial domain of the distance travelled by the puck, we need to incorporate an appropriate transformation. This could also be extended with a physics-based Newtonian model that integrates the relation between friction and mass of the puck, as done by Neupärtl et al. (2020). Because of the inherent complexity of the action generation conditional on the sensory information in the measurement z , we reckon the term *ideal observer* as insufficient and thus would name a model like this a *Bayesian actor model*.

Furthermore, while the ideal observer (merely) considers the perceptual uncertainty of the agent's sensory system, a potential feature of the Bayesian actor is to also consider the variability inherent to its motor system. Bayesian decision-making for planning an action under risk with respect to this motor variability was already introduced and applied by Trommershäuser et al. (2008) and Wolpert and Landy (2012). Yet, in most approaches motor variability is often only considered as additive noise to the planned action (e.g. Acerbi et al. (2014)) and not as an intrinsic variability of a distribution that is also considered in the planning process. Furthermore, if the uncertainty is modeled non-additively, usually only Gaussian distributions are used, which does not leave any room for using more complex or non-linear distributions. Our framework overcomes this limitation and allows for arbitrary action distributions while considering their motor variability in the planning process (cf. Section 3). Either way, a model that also considers motor variability transcends the capabilities of an ideal observer in our understanding.

Moreover, there is consistent evidence that the cost function does not only consider the accuracy in the task but also the effort of executing the chosen action, for example the metabolic cost and thus the effort required for executing farther-reaching movements, which have also been found to be encoded in the brain (Rudebeck et al., 2006; Prévost et al., 2010; Walton et al., 2006, 2003; Huang et al., 2012; Cohen et al., 2010; Izawa et al., 2008).

While we introduced Bayesian approaches here as the successor of the first psychophysical research results, like the Weber-Fechner law, they do not stand in contrast to each other. Rather, we believe that they can complement each other by guiding design choices within Bayesian models. The logarithmic relation between stimulus and sensation, as proposed by Fechner (1860), can be implicitly integrated by assuming a log-normal distribution for the likelihood (e.g. Battaglia et al. (2011)).

2.3. Experimental Paradigms in Decision-Making

Studying human perception naturally poses the question of how reasonable experimental paradigms should look like to allow for quantification of the human perceptual system. This Section provides a brief insight into which experimental paradigms were introduced in the history of perception studies (Section 2.3.1) and how common experimental setups look like today (Section 2.3.2).

2.3.1. Yes-No, 2AFC and 2IFC Tasks

Traditionally, there were three experimental paradigms most commonly used in psychophysics research: (1) *yes-no* tasks, (2) *two-alternative-forced-choice* tasks (2AFC) and (3) *two-interval-forced-choice* tasks (2IFC). All of these paradigms have in common that they require a response by the subject, regardless of their level of uncertainty about the correct answer. In a *yes-no* task, subjects are presented with stimuli of different classes and they must respond if a single stimulus belongs to a certain class ("yes") or not ("no"). For example, they are presented with a picture of an animal and must answer if the animal is a cat or not. In a 2AFC task, subjects are presented with spatially separated pairs of stimuli and they must determine which of the two stimuli fulfills a certain condition. Staying with our animal example, they could be shown two images containing animals and they need to decide which of the two images contain a cat. The 2IFC paradigm can be seen as a variant of 2AFC tasks, where the two stimuli are not separated spatially but temporally by an interval of a certain duration with the stimuli being presented to the subject successively. Both the 2AFC and the 2IFC paradigms can be extended to trials with an arbitrary number of n stimuli per trial between which the subject must decide. In this case, the paradigms are referred to as n AFC or n IFC tasks, respectively. With data gathered from a sufficiently large amount of trials, psychometric functions based on the probabilities for different responses with respect to different stimuli can be estimated. (Prins and others, 2016; Wixted and Thompson-Schill, 2018)

One limit of these paradigms is that they are constrained to discrete responses. While psychometric functions can be used to interpolate between different stimuli, it stands to reason that conducting experiments on a continuous response scale can offer even more insights on human behavior.

2.3.2. (Re-)Production Tasks

Extending the discrete response space from the *yes-no*, 2AFC and 2IFC tasks to tasks involving a continuous response space (Yoo et al., 2021), we can also test subjects' performance and behavior in so-called *production* and *reproduction tasks*. Both task types have in common that they require the subject to produce an action in order to hit a target. Hitting a target can be understood literally, as the subjects, for instance, are asked to throw a puck at a certain target distance (Neupärtl et al., 2020), but can also consist of any other physical quantity that should be produced in any way, for example by holding down a button for a certain duration. While production tasks require the subjects to produce a certain magnitude of some physical quantity (e.g. distance, time), reproduction tasks require the subjects to first perceive and then reproduce a reference stimulus of a certain magnitude, for example by first listening to a sound of a given duration and subsequently press a button for the same duration (Birkenbusch et al., 2015). This enables us to measure human behavior in a more natural setup, since all types of movements the human body can naturally execute are of continuous and smooth nature and not discrete. However, this also introduces more complexity in the analysis and modeling of the planning and action process, for we need to consider each subjects' action variability and cost functions on a continuous scale as well, which is typically more complex than in discrete space.

2.4. Inference Methods for Behavioral Parameters

This section shortly introduces the methods that can be used for retrieving parameters from behavioral data. In addition to this, we shortly sketch the problem of identifiability in the context of Bayesian actor models in Section 2.4.1.

The previously introduced Bayesian modeling approaches allow for complex models incorporating sensory information and prior beliefs that can be used to accurately represent the relation between perception and action in human behavior. Yet, we often do not know how certain behavioral parameters in the model (like perceptual uncertainty, motor variability or the parameterization of the subject's prior beliefs) need to be tuned such that the model accurately fits the data and can reliably reproduce it. To this end, we need methods that allow us to approximate the true values of these parameters, based on the quantities we can observe as researchers, which are usually only the stimuli of the experiment and the responses they elicit in the subject.

We are already using Bayesian concepts for modeling subject behavior (cf. perceptual posterior in ideal observers, Section 2.2.2), but can also make use of Bayes' theorem for inferring the parameters underlying our model. Assuming that we have a set of observed data \mathcal{D} and we want to infer the parameters θ that define our model, we can obtain the probabilities of different parameters conditional on our data analogously to equation (2.4) as shown in equation (2.6).

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta) \cdot p(\theta)}{p(\mathcal{D})} \quad (2.6)$$

Based on equation (2.6) we could only use the likelihood $p(\mathcal{D} | \theta)$ and try to find parameters θ such that they maximize the likelihood when we plug in our observed data \mathcal{D} . This is known as *maximum-likelihood estimation (MLE)*. If we want to incorporate our prior knowledge as researchers about which parameter ranges are realistic or at least most likely, we can employ *maximum a-posterior (MAP)* methods, which were previously introduced as a way of obtaining the optimal action in an ideal observer model (cf. Section 2.2.2). Now, in this case, one would combine the prior and likelihood distributions according to Bayes to obtain a posterior distribution over the possible parameter space conditional on the observed data. We can now also derive the values for the parameters θ that maximize this posterior probability. There are two limitations with this approach: (1) MLE and MAP estimates only provide us with scalar estimates for the parameter values and we do not have any measure of confidence about our estimate. (2) The MAP estimate can only be solved analytically if likelihood and prior distribution are conjugate to each other – i.e. if there exists a closed-form solution describing the posterior resulting from them.

Another approach to recover the underlying parameters of a model is *Markov Chain Monte Carlo (MCMC) sampling* (Hastings, 1970; Robert and Casella, 2004). It provides us not only with a scalar estimate for the underlying parameter values but with a full distribution over the parameter we want to recover, thereby resolving problem (1). Furthermore, it is a numerical method that does not require analytical tractability of the posterior, so that problem (2) is also of no further concern.

Markov Chain Monte Carlo (MCMC) sampling is – as the name suggests – a sampling-based inference method, again based on Bayesian concepts (Hastings, 1970; Robert and Casella, 2004). A welcome feature of MCMC sampling is its ability to not only infer scalar estimates for the true parameters, but to approximate whole posterior distributions over the model parameters. This allows us to quantify the uncertainty of our parameters and to identify ambiguities. For instance, if the observed data can be explained by two different sets of

parameters equally well, MLE or MAP methods would only provide one of the two possible solutions, while MCMC can provide us with a multimodal distribution showing both possible parameter configurations. The basic idea behind MCMC sampling is computationally expensive because it relies on drawing thousands of samples to approximate the posterior distribution, but yields impressive results given its simple concept. In MCMC sampling, we draw proposals for possible values of the true parameter values and compute their posterior probability based on the observed data and our prior beliefs. Based on the posterior probability, the proposed samples are either discarded or accepted as a ‘true’ posterior sample. By that, a Markov Chain can be constructed to be stationary and represent a distribution as estimate for the true posterior distribution. The exact computational principles of MCMC sampling are beyond the scope of this work and will not be explained further here.

2.4.1. Identifiability Problems

If we want to retrieve the parameters of a (Bayesian) model we can encounter identifiability problems for the parameters in question. Depending on the general structure of the model, the complex and non-trivial interactions between different cost function parameters and the overall shape of likelihood and prior form an ill-defined problem since multiple parameter configurations can produce the same behavioral data (Acerbi et al., 2014; Sohn and Jazayeri, 2021).

Sohn and Jazayeri (2021) introduce of so-called *prior-cost metamers*, as illustrated in Figure 2.3. Given the original prior probability density $p_o(s)$ and an original, quadratic cost function $c_o(s, a)$, we can construct a metamer consisting of a metamer prior $p_m(s)$ and a metamer cost $c_m(s, a)$ by shifting prior and cost function inversely proportionally, which results in the same behavior, since the expected cost determines the resulting action and is formed by the convolution of prior and cost.




Figure 2.3.: **Prior-Cost Metamers** – In the common formulation of Bayesian actor models with quadratic cost and Gaussian prior, one can encounter identifiability problems in pairs of prior and cost. An original behavior (top row) can be reproduced with a metamer (bottom row) with different parameters by inversely proportionally shifting cost and prior. This makes the inference problem harder because it does not have a unique, well-defined solution. Inspired by Figure 1B from Sohn and Jazayeri (2021).

Acerbi et al. (2014) resolve this problem with several proposals for the experimental design and the formulation of the computational model, in order to avoid degeneracies in the identification of behavioral parameters. Sohn and Jazayeri (2021) created an experiment with transitions in the prior or cost function between different blocks of trials that do not require a change of the agent’s actions to still perform well in the new block. In some experiments, they found a change in subject behavior at these transitions, in some they did not - this suggests that different experiments trigger different processing pathways in the human mind, either explicitly or implicitly learning prior distributions and cost functions, respectively. Which one of these mechanism will be used, might be dependent on various external, experimental conditions.

2.5. Bayesian Approaches in Psychophysics and Control

There is a multitude of work in modeling perception and action and inference of behavioral parameters that employs methods or concepts similar to the ones used in this work. Section 2.5.1 will give a short overview to related work and Section 2.5.2 introduces some potential applications in, for instance, psychiatric diagnostics and meta-cognition research.

2.5.1. Similar Approaches

This work is mainly inspired and heavily influenced by Neupärthl and Rothkopf (2021). They pursued the same goal and employed the same methods that we use now, but with a few limitations. They also consider motor variability in an optimization problem to find the optimal action in a Bayesian actor model and perform approximate inference via neural networks over the underlying behavioral parameters. A limitation that arises in this approach lies in the perception stage of the model, which does not allow for prior beliefs to be considered in the subject's inference of the true target. We integrate a prior distribution in our model and introduce latent sensory variables that represent the measurements of each stimulus which the subjects combine to obtain a noisy *percept* of the real stimulus. For more details on this, consider Section 3.1.

Further work by Neupärthl et al. (2020) considers perceptual uncertainty and motor variability in a Bayesian model. They apply their model to a puck sliding task which we also use for the evaluation of our method on empirical data (Section 4.2.2). However, their approach is different to ours, since they focused on testing whether and how well subjects use internal physical models in the calculation of their action in (re-)production tasks.

Houlsby et al. (2013) introduced their so-called *Cognitive Tomography* as a method of inferring the priors of subjects in a face perception task. They formulate their problem as an inverse ideal observer. They do not infer any behavioral parameters other than the priors, but provide an in-depth analysis of whether priors are task-dependent or invariant to the task-specific context. Furthermore, they consider actions as simple, discrete choices instead of on a continuous scale.

The use of log-normal distributions as likelihood for sensory evidence inspired by the Weber-Fechner law (cf. Section 2.1.1) was also part of the work of Battaglia et al. (2011). Furthermore, they estimated perceptual uncertainty and motor variability in a distance estimation task involving cue combination of haptic and visual sensory information.

Petzschner and Glasauer (2011), again, use a log-normal relation between stimulus and subjective sensation and model data specifically showing the typical regression and range effects (cf. Section 2.1.5) in a human path integration task. While they only use a perceptual task – distance estimation during locomotion –, not modeling the action and planning process, they use a Kalman-Filter to iteratively update the posterior. This makes the posterior adaptive to each single trial and thus allows for modeling the learning of the prior distribution, which our model does not yet consider. However, they do not perform parameter inference about their models.

Tractability of Bayesian models is a common problem, which can be overcome by using numerical methods. The branch concerned with this is called *Approximate Bayesian Computation (ABC)* and the majority of these works addresses non-tractable likelihoods. Recently, because of the general advance of neural networks, a subdomain of ABC evolved, involving neural approaches (Sisson et al., 2018).




Figure 2.4.: **Müller-Lyer illusion** - Most people experience the upper line to be shorter than the bottom line, despite their actually equal length. The fins at the ends of the two parallel lines modulate the prior of the subject to favor either longer or smaller lines. This probably arises from experience on the spatial organization of the world. Taken from Barrett (2017).

Mostly, these methods focus on estimating unavailable likelihoods, usually by simulating data from the model to learn synthetic likelihoods (Papamakarios et al., 2019; Lueckmann et al., 2019; Fengler et al., 2021). Moreover, there are methods for estimating posterior densities (Blum and François, 2010; Papamakarios and Murray, 2016). While we also use a neural network in our model, our approach is much less sophisticated since we solve an optimization problem to estimate parameters of the subject’s action distribution that produces their responses to different stimuli, instead of approximating the whole distribution. In our model, we still assume a certain functional form of the resulting distribution.

2.5.2. Applications

One of the main features of our framework is the disentanglement of different sources of variability and prior beliefs. While solving (re-)production tasks, subjects go through a complex process involving perception, movement planning and execution of the planned response. Each of these stages adds variability to the resulting action, which is the only observable quantity for us as researchers. At first glance, it seems impossible to retrace this accumulated noise to its different sources. Yet, our approach resolves this problem by confidently estimating each stage’s share on the total resulting noise. Thereby, we can attribute failures in performance to whether there was a problem with the execution of the movement, the perception of the environment or if the person has a mismatch between their prior beliefs and the natural statistics of the world. This is especially useful in medical applications, for example in the early detection of mental or sensorimotor disorders. While it is straight-forward that deviations in the perceived stimuli or the executed actions are reliable indicators of perceptual or motor disorders and probably do not need a thorough analysis through frameworks like ours, the prior beliefs within a Bayesian framework can also be good indicators for mental diagnostics or explain the dynamics of mental disorders (Chambon et al., 2011; Notredame et al., 2014). Since visual illusions are good examples of where our perceptual system fails, illusory paradigms are a good methodological approach to test and compare healthy subjects to subjects with a mental disorder in terms of their performance and behavior (Pessoa et al., 2008). Usually, visual illusions arise because of a mismatch in the weighting of sensory evidence and prior beliefs, for example the Müller-Lyer illusion shown in Figure 2.4 is, according to Gregory (1997), induced by the context of the fins of the parallel lines. Experience in how the world is spatially structured makes the lines appear as if they were of different lengths - this is, in the Bayesian context, an overweighting of the prior belief that is modulated by the spatial context of the lines. This makes the Bayesian framework a well-suited tool for studying visual illusions, especially in the context of mental disorders. While still debated, there is increasing evidence for lesser susceptibility to visual illusions in patients suffering from schizophrenia (Pessoa et al., 2008; Brown et al., 2013; Notredame et al., 2014). This is especially the case in illusions that arise because of heavy prior expectations. In patients with schizophrenia, this can for instance

be explained by too heavy or too weak influence of the prior in the inference of the stimulus, depending on the exact task (Chambon et al., 2011; Notredame et al., 2014). As another example, in autistic subjects an exaggerated reliance on sensory evidence was found (Pellicano and Burr, 2012). This paradigm can be extended from the sole perception aspect towards tasks involving perception and action, since there may be not only perceptual but also sensorimotor impairments (Shergill et al., 2005; Chen et al., 2011; Brown et al., 2013), which perfectly maps to the use-case of our framework.

The domain of meta-cognition is concerned with how accessible cognitive processes are, essentially it addresses '*cognition about cognition*'. Since our work puts different sources of uncertainty in planning processes and how they can be disentangled in its focus, it is worth taking a look at metacognition in perceptual tasks. De Gardelle and Mamassian (2015) investigated whether the reported uncertainty of subjects matched the inherent variability of a task's stimuli. They further examined whether confidence judgments were similar across different sensory modalities (De Gardelle and Mamassian, 2014). Our work could contribute to extending their findings from perceptual tasks to tasks involving sensorimotor control. Furthermore, our model could be modified to account for meta-cognitive uncertainty about the subjects' confidence in judging their own uncertainty. Additionally, it might help to explain to which extent the meta-cognitive uncertainty is composed of perceptual uncertainty or motor variability and how much noise is added only in the retrieval and judgment of a subject's own uncertainty.

3. Methods

Bayesian actor models (cf. Section 2.2) are a popular approach for modeling human behavior in tasks involving perception and action (cf. Section 2.1). Most approaches were designed for tasks that only allow for simple, discrete responses (cf. Section 2.3.1) and reach their limits when we want to extend them to more complex, continuous production and reproduction tasks (cf. Section 2.3.2). Especially because inferring behavioral parameters is analytically intractable in most cases (cf. Section 2.4). We overcome this limitation by approximating the decision-making within Bayesian actor models with neural networks.

In reproduction tasks, from the researchers' perspective, only the target stimulus t and the subject's response action a are observable. All other variables, like the subject's sensory representation of the stimulus – which we call *measurement* or *sensory information* z – or behavioral parameters that determine the subject's behavior, are latent, internal variables and thus can only be approximated by inference. Inferring these parameters is crucial to identify and quantify differences between subjects and understanding human behavior. Our model consists of different components, which are shown in Figure 3.1 and will be introduced here briefly to give a general understanding of how all of these components are interconnected. First, we model the perception of the subject, which can be posed as an inference problem over the true state of the world. The subjects only have access to the measurement z , which represent their sensory information about the stimulus. Formalized in Bayesian statistics, this measurement is combined with their prior beliefs about the world to form a posterior distribution that represents their belief about the real world and thereby forming a so-called *percept*. Our perceptual modeling is also psychophysically grounded in the Weber-Fechner law. How this process is implemented exactly is addressed in Section 3.1.

To quantify human behavior and attribute its origins to different aspects in a subject's performance, we assume that human behavior can be described by a *loss function* $\mathcal{L}(a, t)$ involving both a loss for missing the target t as well as a loss for the effort of producing the respective action a . We will also refer to loss functions as *cost functions*; the terms will be used interchangeably in this work. The weighting of this trade-off between task accuracy and required effort is a highly subjective measure and modeled by the free parameters of the employed cost function. Since there is plenty of evidence that humans can (at least unconsciously) access their own perceptual uncertainty and motor variability, these also need to be considered in modeling the decision-making for the final action a . To incorporate sensory variability into a model, as it is done, e.g., in *ideal observer models* (cf. Section 2.2.2), usually the expected loss over the posterior from the perception state is computed and minimized. Our model is more sophisticated in that it also considers motor variability and not only sensory variability, therefore we consider our model to be a *Bayesian actor model* (cf. Section 2.2.3). The final action is generated from a distribution parameterized by a set of parameters θ_a .

Usually, the distributions used in these models are Gaussian and the cost functions are assumed to be quadratic. This is often due to computational convenience since it keeps the model analytically tractable. We employ log-normal distributions since these are in line with Weber-Fechner'ian principles of perception and we use a more complex cost function to factor in more aspects of behavior. This does away with said analytical

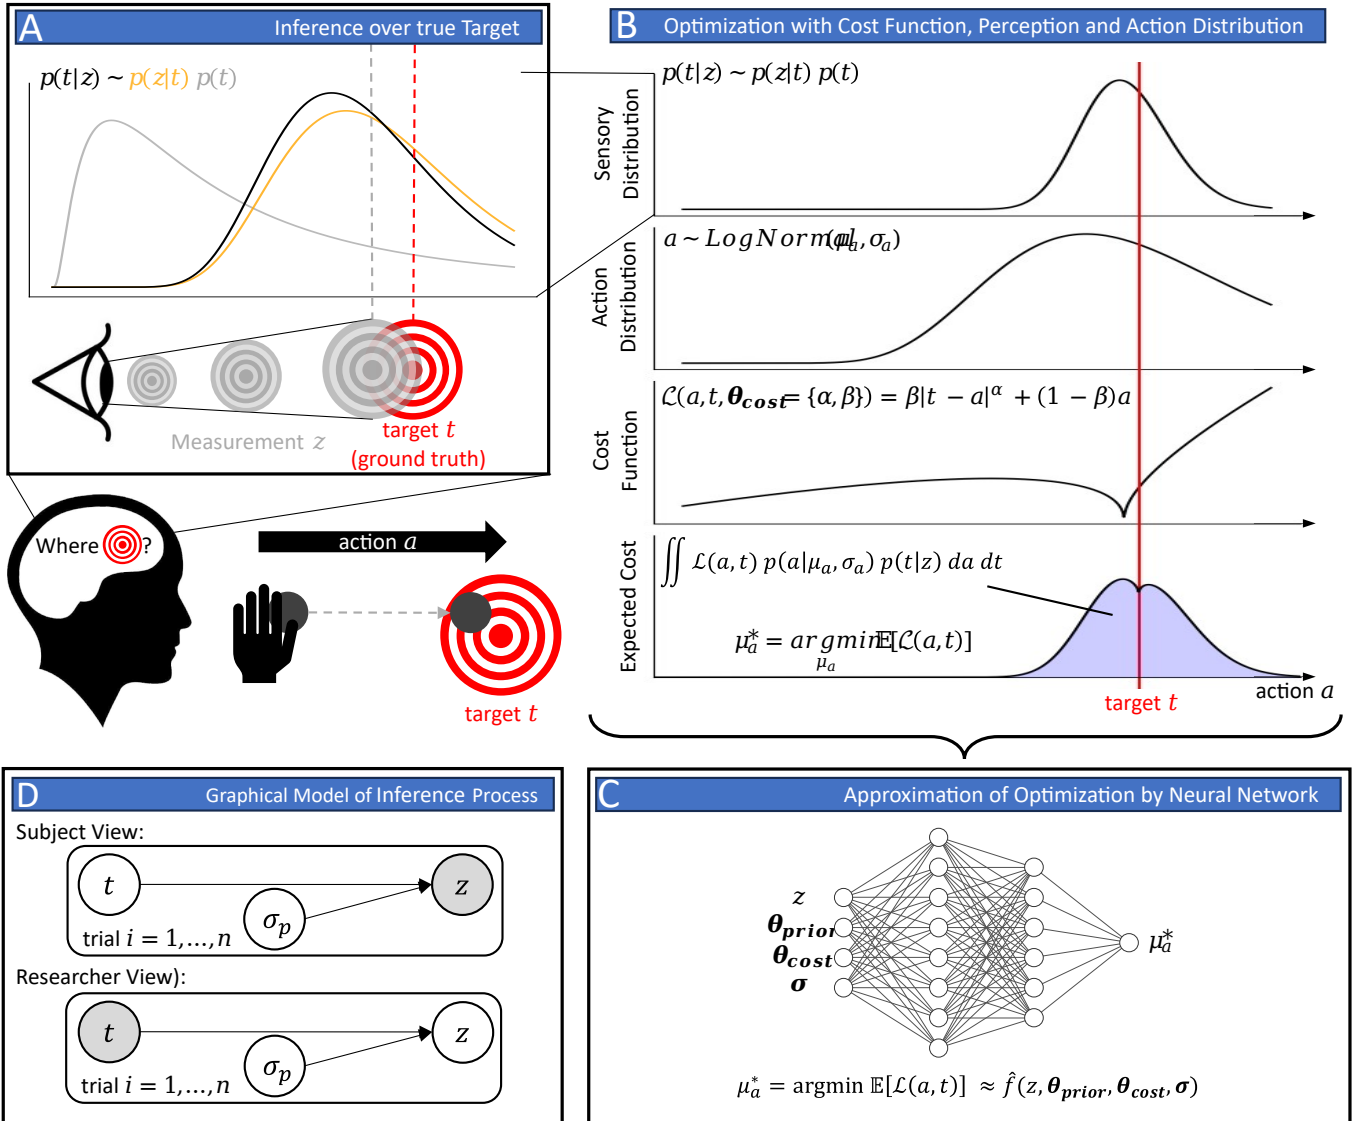


Figure 3.1.: **Bayesian Actor Model of our Framework** – (A) The subject only has access to a sensory estimate z of the true target stimulus t . In order to solve the task successfully, they need to perform inference to obtain a posterior over the true value of the target t . (B) The cost function is convoluted with the respective sensory and action distributions to form the expected cost over both distributions. The optimal action distribution parameter μ_a^* is given by the value of μ_a for which the expected cost $\mathbb{E}[\mathcal{L}(a, t)]$ becomes minimal. (C) To speed up computation and make MCMC feasible, the convolution and optimization process from (B) is approximated by a neural network that takes the sensory information z and other behavioral parameters θ_{cost} , θ_{prior} and σ as input to approximate the optimal action distribution parameter μ_a^* . In our case, the action distribution only has one free parameter, but this could be easily extended to probability distributions with higher dimensional parameterizations. (D) While the subject needs to infer the true target t based on its measurement z , the researcher can only access the true target t in each experimental trial and needs to infer the subject’s perceptual uncertainty σ_p and their sensory information z from that.

tractability. Computing the optimal parameters θ_a^* of the action distribution, such that the expected loss is minimized, is based on a double convolution which can only be solved numerically in most cases. The exact design of the cost function and how the numerical solution to the optimization problem is found is introduced in Section 3.2.

Numerical solutions like these are usually too computationally expensive to perform inference over the models that implement them. We overcome this limitation by approximating the numerical solution with neural networks, trained on a wide range of possible parameters for our actor model. Thereby, we achieve a speedup of about three orders of magnitude and thus make it feasible to employ sampling-based inference methods in order to retrieve the true values of the underlying parameters describing a subject's behavior. The architecture, training process and evaluation of the neural network are presented in Section 3.3.

Since our ultimate goal is to not only appropriately model human behavior but also to infer the parameters shaping this behavior from experimental data, we use *Markov Chain Monte Carlo (MCMC) Sampling* in our framework. Thereby, we can infer whole posterior distributions over our parameters given the observed actions of the subject. The exact probabilistic model we assume, as well as the technical details for our inference method, are shown in Section 3.4. To avoid confusion in the inference process about which variables are observed and which are not, we need to distinguish between the perspective of the subject and of the researcher:

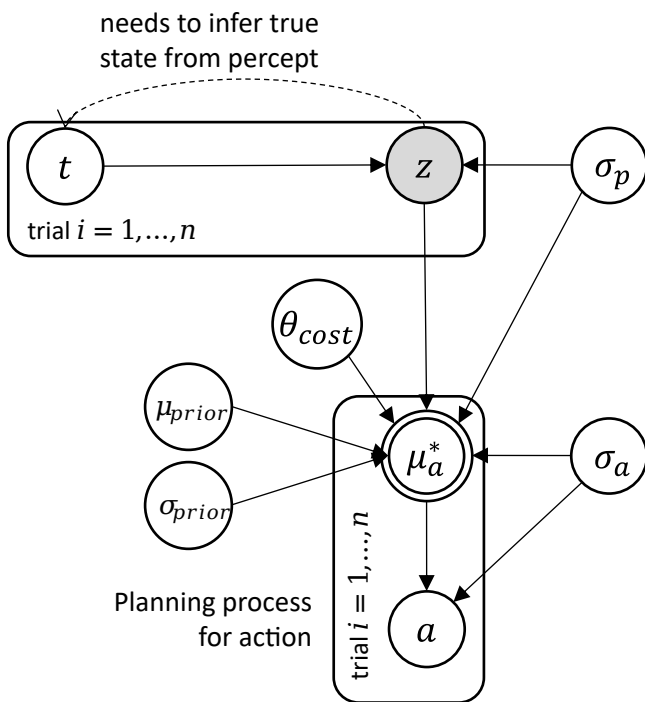
- **Subject Perspective** - From the view of the subject, the target t is unknown, since they only observe their internal measurement z , which they need to use to perform inference about the true stimulus t . All other parameters are assumed to be known to the subject, since there is a large amount of empirical evidence showing that subjects are (at least implicitly) aware of their perceptual and action uncertainties and incorporate them into their movement planning (cf. Section 2.1.4). This also applies to the cost function, prior and action distribution parameters.
- **Researcher Perspective** - The researcher strives to obtain estimates of the behavioral parameters of the subject and is thus also performing inference, but not in the same way the subject does. While the subject performs inference only over the target stimulus t depending on their observation z of the stimulus, the researcher uses the targets t known from the experimental setup and the observed actions a of the subject in order to perform inference over the whole process from perception to action. They want to infer all behavioral parameters $\theta = \{\theta_{\text{cost}}, \theta_{\text{prior}}, \sigma\}$ and thereby obtain quantitative measures for the subject's behavior. The latent sensory variables z and action distribution's parameters θ_a are also latent and need to be inferred from the researcher, but the behavioral parameters θ shaping the overall behavior are of primary interest to them.

Model Structure

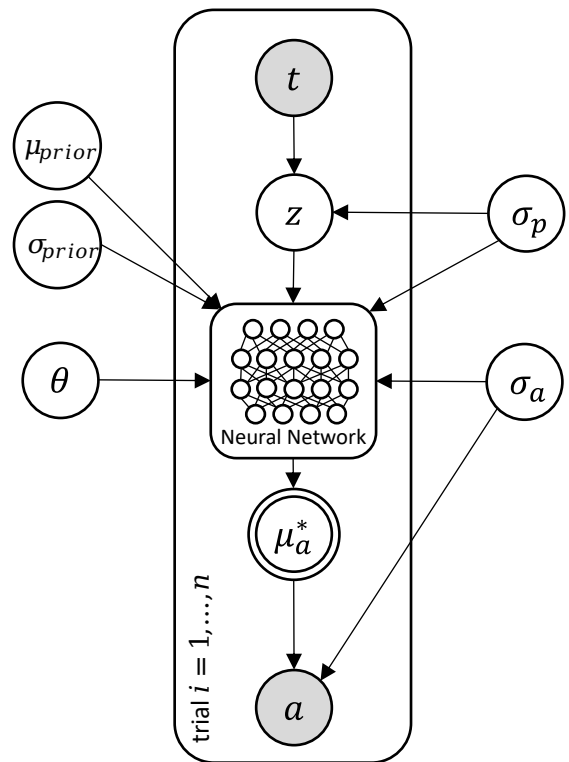
The whole model consists of four main components, the first three of which are shown in Figure 3.1:

1. **Bayes-Optimal perception** of the target, representing subject-side inference over the true value of the target t given a noisy measurement z which is latent and cannot be accessed from the researchers' perspective. We assume behavior of an ideal observer and thus that people have a-priori knowledge about how the targets are distributed and combine it with their sensory information received during the task to obtain an estimate over the true value of the target (Section 3.1).
2. **Numerical minimization** of the expected value of the cost function $\mathbb{E}[\mathcal{L}(a, t)]$ over the values of the action distribution's parameters θ_a which determine the actions taken by the subject (Section 3.2).
3. Training a **Neural Network** leveraging the results of the minimization in order to speed up the optimization process. This step has no purpose besides speeding up computation, since the numerical minimization is computationally expensive (Section 3.3). By approximating the minimization with a neural network, it becomes feasible to use Markov Chain Monte Carlo Sampling, which would not be possible otherwise.
4. **Markov Chain Monte Carlo Sampling** for inference of parameters, making use of said neural network. With this method we can reliably extract the values of the parameters of the cost function θ_{cost} as well as behavioral parameters like the uncertainties of the perceptual and action distribution $\sigma = \{\sigma_p, \sigma_a\}$ and the parameters of the prior distribution θ_{prior} . The totality of all inferred parameters is given by $\theta = \{\theta_{\text{cost}}, \theta_{\text{prior}}, \sigma\}$ (Section 3.4).

By assembling all of these concepts in our framework, we are providing a toolbox for any researcher concerned with continuous decision-making to obtain estimates of the parameters shaping the subject's behavior and thus making inter-subject comparisons possible on a quantitative level. Furthermore, careful shaping of the cost function and the perception stage as Bayesian inference allows for the interpretation of cost function parameters and prior parameters. For example, how strongly a subject weights effort cost versus missing the target or how heavily they rely on a prior versus sensory evidence to plan their action and what this prior looks like. By substituting different components based on different assumptions about the underlying perceptual / action distributions and parameterizations of the cost function, our framework encourages researchers to find a model best describing a subject's behavior and potentially gain insights and ideas about the inner workings of the subject's mind on an implementational or algorithmic level. Thereby, our tool enables researchers to identify behavioral differences between subjects.



Subject View



Researcher View

Figure 3.2.: **Bayesian Models from Different Perspectives** – The subject only has (explicit) access to their internal measurement z and the action a they execute. The subjects need to infer the ground truth of the stimulus t given the sensory information z . They then need to generate an action a from that. The researcher needs to infer the latent measurement of the subject given the stimulus t they used in each trial. For that, we employ a neural network to mimic the inference of the true stimulus and the optimization process happening in the subject’s mind to choose the action distribution’s parameter μ_a^* which yields the action a in the end. While the subject is oblivious of their internal parameters θ and σ , these are the exact parameters the researcher tries to infer to describe and analyze their behavior.

3.1. Observer Model

Our model is based on an ideal observer model as introduced in Section 2.2.2. Thus, we are modeling the perception stage also in a Bayesian framework by computing the posterior distribution over the true target t based on the measurement z of the target t the subject received from their sensory system. In the introduction of Bayesian approaches in Section 2.2 we used the variable s for the overall state of the environment which the ideal observer tries to infer and the Bayesian actor acts upon. In the following, we will use the variable t for the target, which is basically equivalent to the state s , but we only consider (re-)production tasks for our framework, thus we are always only inferring the value of a specific target t and not an extensive state s . It is important to note here that we cannot use the true value of the target t in our model directly, since we assume that humans do not have direct access to the ground truth of the world structure surrounding them but only to a noisy measurement z . They need to use this information from their sensory system to infer the true state of the world and we try to mimic this process with our ideal observer (see Figure 3.2).

To do so, we assume a log-normal likelihood $\log \mathcal{N}(z | \log t, \sigma_p)$ for the measurement z , using the target t as median of the resulting log-normal distribution and the subject's individual perceptual uncertainty σ_p . The main reason why we chose a log-normal likelihood is that it is an elegant way to integrate the Weber-Fechner law into a Bayesian model. Since Weber-Fechner proposes a logarithmic relation between stimulus and sensation (cf. Section 2.1.1), we can implicitly model this relation in a log-normal distribution for the sensory likelihood. Furthermore, log-normal distributions are already naturally constrained to positive values, which only make sense in perception-and-action tasks, for negative measurements of the stimuli or negative actions are usually not possible in common formalizations of these problems. Last but not least, a core property of the log-normal distribution is that it mimics signal-dependent noise since its variability scales with the distribution's mean. Thereby, we can implicitly incorporate insights from human motor control research that shows that human movements are usually subject to signal-dependent noise (cf. Section 2.1.4).

For computational convenience, we also chose the subject's prior as a log-normal distribution with two free parameters $\log \mathcal{N}(\mu_{\text{prior}}, \sigma_{\text{prior}})$. A log-normal distribution is the conjugate prior to a log-normal likelihood, which results in an analytically tractable posterior distribution that also takes on the form of a log-normal distribution. Bayes' rule describes this relation, as introduced in Section 2.2:

$$P(t | z) = \frac{P(z | t) \cdot P(t)}{P(z)}$$

One should keep in mind that a subject's prior beliefs are not necessarily extensive knowledge about the true characteristics of the world but can be as simple as their own idea about the distribution of the experiment's stimulus.

The parameters of the posterior distribution are then given by equation (3.1)

$$\mu_{\text{post}} = \sigma_{\text{post}}^2 \left(\frac{\log \mu_{\text{prior}}}{\sigma_{\text{prior}}^2} + \frac{\log z}{\sigma_p^2} \right) \quad \sigma_{\text{post}}^2 = \frac{1}{\frac{1}{\sigma_{\text{prior}}^2} + \frac{1}{\sigma_p^2}} \quad (3.1)$$

The resulting posterior is now a product of the likelihood and prior, where the respective uncertainties σ_p and σ_{prior} define which of the two has more influence on the posterior. By definition, the posterior is always more skewed towards the one of the two distributions with lower variability. Because of the lower variability, the respective distribution can be interpreted as the 'more reliable' source of information and thus is given more

weight. This concept is also present in Bayesian scenarios of cue integration or multimodal integration, where the cue with the smallest variability is usually attributed the strongest weight in the formation of the posterior and there is strong evidence that the human brain works in a similar way (cf. Section 2.4). The resulting posterior is shown exemplarily in Figure 3.3.




Figure 3.3.: **Inference of the Posterior Distribution** - Example of the resulting posterior of the ideal observer when given a sensory measurement z . The likelihood is constructed around this measurement. In this example, the variability of the likelihood is smaller than the variability of the prior and thus is skewed towards the likelihood. Here, the prior draws the posterior away from the true target value t .

At this point, we extend the classical ideal observer model to a Bayesian actor by introducing an action distribution that also includes the motor variability σ_a . Since we now have two sources of uncertainty – the sensory as well as the action uncertainty –, we need to incorporate both distributions in the calculation of the expected loss. We model the planning of the movement as a minimization over the expected value of the loss and thereby determine the parameters of the action distribution. This process will be the focus of the next Section.

3.2. Cost Function and Optimization

The first component of our model consists of finding parameters for the distribution generating the subjects' actions, such that the expected loss $\mathbb{E}[\mathcal{L}(a, t)]$ over our loss function $\mathcal{L}(a, t)$ becomes minimal. The choice and parameterization of the cost function are discussed in Section 3.2.1, while the optimization process is described in Section 3.2.2. Usually, the expected cost $\mathbb{E}[\mathcal{L}(a, t)]$ is not analytically tractable, but for some special cases we can derive an analytical solution. This is shown in Section 3.2.3.

3.2.1. Cost Function Parameterization

The choice of the cost function depends on a lot of parameters. It should contain a cost for missing the target as well as a regularizing term penalizing increased effort. This effort represents physical as well as mental stress exerted on the subject by taking a certain action. For example, throwing a ball at a target farther away requires more physical strength or holding down a button for a certain period of time tests the patience of the subjects and thus creates mental stress.

In order to portray these costs accurately in our loss function $\mathcal{L}(a, t)$, we need to choose from a range of possible functional forms. Besides that, we also need to account for solvability of the integral, since the functional form and choice of parameters also influences whether a solution exists at all. The main focus of the loss function should be to represent the goal of the task well - that is, how well the action hits the target in terms of accuracy. This is usually implemented using quadratic loss. Yet, how accurate the target is hit is not the only factor guiding human behavior. For example, the mental and physical effort for performing a chosen action should be represented by the cost function as well. The cost function we ultimately decided on is given by equation (3.2) and visualized in Figure 3.4a. It uses a weighted combination of the effort a and a penalty for missing the target.

$$\mathcal{L}(a, t) = \beta|a - t|^\alpha + (1 - \beta)a \quad (3.2)$$

Since we are considering perceptual as well as motor uncertainty, we need to account for variability in the perception of the target t and in the execution of the action a . Therefore, we have a perceptual distribution representing the sensory uncertainty and an action distribution representing the motor variability as given in equation (3.3). We assume by definition that each subject's actions distribution is parameterized by the optimal value μ_a^* , which minimizes the expected value of the subject's cost function with respect to its parameters $\theta_{\text{cost}} = \{\alpha, \beta\}$, uncertainties $\sigma = \{\sigma_p, \sigma_a\}$ and prior beliefs $\theta_{\text{prior}} = \{\mu_{\text{prior}}, \sigma_{\text{prior}}\}$.

$$\begin{aligned} t &\sim \log \mathcal{N}(\mu_{\text{post}}, \sigma_{\text{post}}) \\ a &\sim \log \mathcal{N}(\mu_a^*, \sigma_a) \quad \text{with} \quad \mu_a^* = \underset{\mu_a}{\text{argmin}} \mathbb{E}[\mathcal{L}(a, t)] \end{aligned} \quad (3.3)$$

The expected cost is then given by equation (3.4).

$$\mathbb{E}[\mathcal{L}(a, t)] = \int_0^\infty \int_0^\infty \mathcal{L}(a, t) p(a | \mu_a, \sigma_a) p(t | \mu_p, \sigma_p) p(t) dt da \quad (3.4)$$

We can now use this equation to find the optimal value μ_a^* for μ_a which minimizes the expected cost. This is further described in Sections 3.2.2 and 3.2.3.

Alternative Parameterizations

Bear in mind that our choice of the cost function is motivated by implementations of cost and effort commonly used in other studies (cf. Section 2.2.3), but the plenty possible parameterizations also allow for a lot of other cost functions that might give even better results. We also see our framework as a possibility of exploring different cost functions in order to find which one resembles empirical data best. A small overview over possible cost function parameterizations is given in Table 3.1. For example, while quadratic loss functions are most common in research for computational modeling of human behavior or optimal control, Körding and Wolpert (2004b) propose an inverse Gaussian curve as a possible loss function. Since they observed that the cost function is less sensible to big errors than prescribed by a quadratic cost, the inverse Gaussian may be a better fit because it is more robust towards big errors. This cost function is visualized in Figure 3.4b.

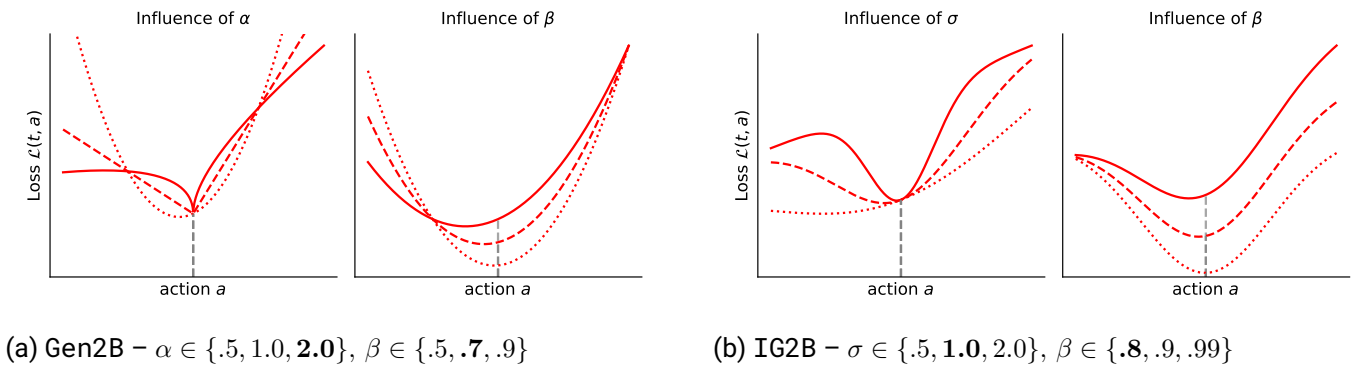


Figure 3.4.: **Cost Functions** – Exemplary plots for different parameterizations of the Gen2B and IG2B cost functions. The dashed grey line represents the target t . In each pair plot β is constant on the left and α is constant on the right (respective bold value in subfigure caption). Parameter values increase from dotted to solid line styles.

It is important to note that each parameterization can change the range of values that are valid for each parameter in the sense that an invalid combination of parameter values can lead to the non-existence of an optimal solution for the problem.

3.2.2. Optimization

Since we train the neural network for our model in a supervised fashion, we need a large data set with solutions to the optimization problem to use in the training. However, an analytical solution does not exist in the general case. Therefore, we need to rely on numerical methods to find the optimal μ_a^* minimizing the expected cost $\mathbb{E}[\mathcal{L}(a, t)]$. We need a numerical tool (1) for the integration over the probability distribution and the cost function and (2) for finding the optimal μ_a^* , as explained in the following Sections. For further details on a special case yielding an analytical solution, see Section 3.2.3.

ID	Cost Function $\mathcal{L}(a, t)$	Parameters θ_{cost}	Components & Properties
Gen2B	$\beta a - t ^\alpha + (1 - \beta)a$	(α, β)	General form that we use here for all examples and which gave us the best results. Trade-off between accuracy (1st term) and effort (2nd term) is weighted by $\beta \in [0, 1]$ in both two terms.
Gen2BExt	$\beta a - t ^\alpha + (1 - \beta)a^\gamma$	(α, β, γ)	Extended version of Gen2B. Introduces an additional parameter γ which further transforms the taken action a into a value representing subjective effort.
Gen1B	$ a - t ^\alpha + \beta a$	(α, β)	Modified version of Gen2B using β not for a weighted average of accuracy and effort, but only as weight for one term - the effort. This changes the value range of β , since it is not longer constrained to $[0, 1]$ but can take any positive real value $\beta \in \mathbb{R}^+$. E.g., if the first term is big, we also need to increase the value of β , since the effort would have no influence on the overall cost otherwise.
IG2B	$\beta e^{-\frac{(a-t)^2}{2\sigma^2}} + (1 - \beta)a$	(σ, β)	Functional form of the error given by an Inverted Gaussian instead of the general form with an exponentiated absolute difference.

Table 3.1.: Overview of different cost functions with their respective IDs by which they are referenced throughout this work. We mostly use version Gen2B for our examples cases. Each cost function has their own set of parameters that all need to be inferred in the inference step (cf. Section 3.4). Regarding the value ranges of the parameters and the solvability of the problem w.r.t. the chosen parameter ranges, consult Section 3.2.4. Plots for the other versions of cost functions are shown in Appendix A.1 and A.2.

Numerical Integration

For solving the double integral in the computation of the expected cost given by equation (3.4), we use the `dblquad` function as provided by the `SciPy`¹-package. Thus, the integral is solved with an adaptive 21-point Gauss-Kronrod-Quadrature in combination with the epsilon algorithm by Wynn (1956).

The integral from equation (3.4) that we are solving has indefinite upper bounds, because we are using log-normal distributions which have infinite positive support. To ensure convergence of the integral, we implement definite bounds. We assume that we can integrate over the $[0.001, 0.999]$ -quantiles of the respective random variable's dimension without considerable loss of accuracy. Thus, we are actually numerically solving the integrals in (3.5), where q_x^a and q_x^t describe the distributions' x -quantiles, respectively. Moreover, computation of definite integrals is usually faster than computation of indefinite integrals, regardless of the method used.

$$\mathbb{E}[\mathcal{L}(a, t)] = \int_{q_{0.001}^a}^{q_{0.999}^a} \int_{q_{0.001}^t}^{q_{0.999}^t} \mathcal{L}(a, t) p(a \mid \mu_a, \sigma_a) p(t \mid \mu_{\text{post}}, \sigma_{\text{post}}) p(t) dt da \quad (3.5)$$

Numerical Optimization

Our framework assumes that μ_a of the subject's action distribution is chosen such that the expected loss $\mathbb{E}[\mathcal{L}(a, t)]$ is minimized, as given in equation (3.3). Again, since the integrals are not analytically tractable in the general case, we employ numerical methods to minimize the expected loss.

We are using a bounded version of *Brent's Method*² (Brent, 2013) to find the optimal μ_a^* which minimizes the expected loss $\mathbb{E}[\mathcal{L}(a, t)]$. The boundaries allow us to constrain the solution space to positive values for μ_a^* , since we parameterized our log-normal distributions $\log \mathcal{N}(\mu, \sigma)$ such that $\mu \in (0, \infty)$. Even though the values of μ_a^* should not be much bigger than the biggest values for the stimuli t , we do not employ any such constraints into the optimizer's bounds. Therefore we choose the upper bounds generously. This leads to the bounds $[\epsilon, 100]$, where we choose $\epsilon > 0$ such that it is as small as our implementation allows without causing numerical problems³.

3.2.3. Analytical Solution

Even though the problem is not solvable analytically for most cases, there are some special cases for which an analytical solution exists. While this analytical solution is independent of the probability distributions used, the solution (and with it also the solvability) of the problem changes with a change of the cost function. We will demonstrate the derivation of the analytical solution for cost function `Gen2B` given in equation (3.2). We also assume the log-normal distributions in equation (3.3). For convenience, we show them here again:

$$\mathcal{L}(a, t) = \beta |a - t|^\alpha + (1 - \beta)a$$

¹<https://scipy.org/>

²We use the method `minimize_scalar` of the package `SciPy` with parameter `method='bounded'`.

³We choose ϵ equal to ten times the machine epsilon for `float64` in Python, as provided by the package `NumPy`'s method `np.finfo(float).eps`.

With the following distribution of random variables:

$$t \sim \log \mathcal{N}(\mu_{\text{post}}, \sigma_{\text{post}})$$

$$a \sim \log \mathcal{N}(\mu_a^*, \sigma_a) \quad \text{with} \quad \mu_a^* = \underset{\mu_a}{\text{argmin}} \mathbb{E}[\mathcal{L}(a, t)]$$

We want to solve for the expected loss given by

$$\mathbb{E}[\mathcal{L}(a, t)] = \int_0^\infty \int_0^\infty \mathcal{L}(a, t) p(a \mid \mu_a, \sigma_a) p(t \mid \mu_{\text{post}}, \sigma_{\text{post}}) p(t) dt da$$

While solving the integral directly is quite laborious, the solution is straightforward if we express it as the expected value and make use of its linearity:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(a, t)] &= \mathbb{E}[\beta |a - t|^\alpha + (1 - \beta)a] \\ &= \beta \mathbb{E}[|a - t|^\alpha] + (1 - \beta) \mathbb{E}[a] \end{aligned}$$

Since the expected value of the absolute difference in the first term cannot be solved for directly without solving the integral (which is also quite demanding and does not generalize well to other cost functions), we are now constraining ourselves to the cases where we do not need to resolve the absolute difference. By constraining ourselves to $\alpha = 2n, \forall n \in \mathbb{N}^+$, the absolute difference can be ignored because of the even exponent:

$$\mathbb{E}[\mathcal{L}(a, t)] = \beta \mathbb{E}[(a - t)^\alpha] + (1 - \beta) \mathbb{E}[a]$$

Now, we still need to solve for the exponentiated difference. For natural values of α , we can do so by making use of the binomial coefficient given by

$$(a - t)^\alpha = \sum_{k=0}^{\alpha} \binom{\alpha}{k} a^k (-t)^{\alpha-k}$$

Of course, the parameter α can generally take on any real value in reality, but this approach only works for natural values.

For the sake of simplicity, we are constraining ourselves to $\alpha = 2$ for the remaining steps. This leads us to the following expression:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(a, t)] &= \beta \mathbb{E}[(a - t)^2] + (1 - \beta) \mathbb{E}[a] \\ &= \beta \mathbb{E}[a^2 - 2at + t^2] + (1 - \beta) \mathbb{E}[a] \\ &= \beta(\mathbb{E}[a^2] - \mathbb{E}[2at] + \mathbb{E}[t^2]) + (1 - \beta) \mathbb{E}[a] \end{aligned}$$

If we have exponents $\alpha > 2, \alpha \in \mathbb{N}^+$, the binomial coefficient becomes much more complicated and its solution contains exponents up to the value of α . Still, we can further resolve this equation for any arbitrary probability distribution for a and t , as long as we know the closed form expressions for sufficiently high moments (up to $\mathbb{E}[x^\alpha]$) of the distributions.

Besides, we can assume independence between any occurrences of random variables. This may seem naive, since the action a clearly depends on the target t by definition. However, because the objective function is to be optimized with respect to a parameter of the action distribution, we are (at least in the context of finding the optimum) free to choose this parameter, which eliminates this dependence. Exploiting this property, we can further simplify the expression:

$$\mathbb{E}[\mathcal{L}(a, t)] = \beta(\mathbb{E}[a^2] - 2\mathbb{E}[a]\mathbb{E}[t] + \mathbb{E}[t^2]) + (1 - \beta)\mathbb{E}[a]$$

The first two moments of any distribution are given by the following expressions:

$$\int_{-\infty}^{\infty} x^2 p(x) dx = \mathbb{E}[X^2] = \text{Var}[X] + \mathbb{E}[X]^2 \quad \int_{-\infty}^{\infty} x p(x) dx = \mathbb{E}[X]$$

For the case considered here, we can just plug in the respective values for the log-normal distribution:

$$\mathbb{E}[X^2] = \int_{-\infty}^{\infty} x^2 p(x) dx = (e^{\sigma^2} - 1)e^{2\mu + \sigma^2} + e^{2\mu + \sigma^2} = e^{2(\mu + \sigma^2)} \quad \mathbb{E}[X] = \int_{-\infty}^{\infty} x p(x) dx = e^{\mu + \frac{\sigma^2}{2}}$$

Plugged into our formula, this yields:

$$\mathbb{E}[\mathcal{L}(a, t)] = \beta(e^{2(\mu_a + \sigma_a^2)} - 2e^{\mu_a + \mu_{\text{post}} + \frac{\sigma_a^2 + \sigma_{\text{post}}^2}{2}} + e^{2(\mu_{\text{post}} + \sigma_{\text{post}}^2)}) + (1 - \beta)e^{\mu_a + \frac{\sigma_a^2}{2}}$$

We can now differentiate this and solve for its roots in order to find the optimum μ_a^* minimizing the expected loss:

$$\frac{d}{d\mu_a} \mathbb{E}[\mathcal{L}(a, t)] = 2\beta(e^{2(\mu_a + \sigma_a^2)} - 2e^{\mu_a + \mu_{\text{post}} + \frac{\sigma_a^2 + \sigma_{\text{post}}^2}{2}}) + (1 - \beta)e^{\mu_a + \frac{\sigma_a^2}{2}} \stackrel{!}{=} 0$$

Solving this for μ_a leads to

$$\mu_a^* = \ln\left(e^{\mu_{\text{post}} + \frac{\sigma_{\text{post}}^2}{2}} + \frac{\beta - 1}{2\beta}\right) - \frac{3}{2}\sigma_a^2 = \ln\left(ze^{\frac{\sigma_{\text{post}}^2}{2}} + \frac{\beta - 1}{2\beta}\right) - \frac{3}{2}\sigma_a^2 \quad (3.6)$$

This closed-form solution already shows that the problem is not solvable for every arbitrary configuration of parameters, since this solution is only defined if the logarithmized term is positive:

$$ze^{\frac{\sigma_{\text{post}}^2}{2}} + \frac{\beta - 1}{2\beta} > 0$$

Which is equivalent to

$$ze^{\frac{\sigma_{\text{post}}^2}{2}} > \frac{1 - \beta}{2\beta}$$

This solution implies that the smaller β is, the higher the posterior uncertainty σ_{post} , the target t or ultimately its measurement z need to be in order for a solution to exist. Furthermore, the result of the logarithm in equation (3.6) also needs to be $\ln(\dots) > \frac{3}{2}\sigma_a^2$, since $\mu_a \leq 0$ is not allowed by definition - note that we parameterize our log-normal distributions in positive space before they are transformed to log-space $(-\infty, \infty)$ and not directly in log-space, as it is often done. A more thorough investigation of the solvability of the problem for different parameter configurations is presented in the next Section.

3.2.4. Solvability

The solvability of the problem is heavily influenced by the complex dynamics and interactions of the prior, cost and uncertainty parameters. Since an analytical solution only exists in special cases (cf. Section 3.2.3), we need other methods to explore the constraints for when a well-defined solution exists and when it does not. To do so, we used the data set generated for training and testing the neural network (cf. Section 3.3) and especially compared the parameter configurations for which the numerical solver could not find a solution with the configurations for which it did find one. Since the numerical solver we used always outputs a proposed solution regardless of whether a solution really exists, we discarded all solution proposals that lie exactly on the bounds of the solution space that we constrained the numerical solver to or that heavily exceeded a realistic range. For instance, we discarded all solutions with values $\mu_a^* > 40$ since our stimuli all lie in the interval $[1.2, 4.8]$ and the value for μ_a should not really exceed this upper bound in the most cases. The parameter configuration for which a solution exists is shown for cost function Gen2B in Figure 3.5. For the solvability of the model with other cost functions, please consult Appendix A.3 and A.4.

In case of the Gen2B scenario, we can identify an interaction between α and β . There is a linear threshold between the values of the two parameters, such that smaller values of β require bigger values of α in order for a solution to exist. In general, both cost function parameters $\theta_{\text{cost}} = \{\alpha, \beta\}$ are not allowed to become too small or else the problem is not solvable. The value of β seems to almost categorically need to surpass the value of 0.5 to guarantee existence of a solution. This is not that clear in the plot in Figure 3.5 since we already incorporated these insights into the generation of the data set, which hence does not contain a lot of samples with low β .

It is also noteworthy that a solution cannot be found for small targets t if the value of α gets too big. This is most likely not a problem of the model specification itself but rather a numerical issue, since α defines the exponent of the cost function's accuracy-term which can lead to extremely high or extremely low values, depending on the exponentiated value.

The solvability always depends on more than the relationship of a single pair of parameters but also other parameters. However, since the parameter space of our model is quite high-dimensional, it is difficult to identify all of these interdependencies between the different parameters.

Please keep in mind that these results and insights only apply to the Gen2B scenario. If one wishes to use another cost function or exchange other components of the model, this will most likely influence the subsets of the parameter space for which a solution exists. Therefore, one needs to examine the dynamics of their model's parameters carefully for each new model specification and adjust the inputs on which the neural network is trained accordingly, in order for our framework to produce good results. Our method can only work reliably if the neural network was trained on a set of parameters for which solutions exist and that are realistic in the sense that the combination of model and parameters can re-produce empirical data.

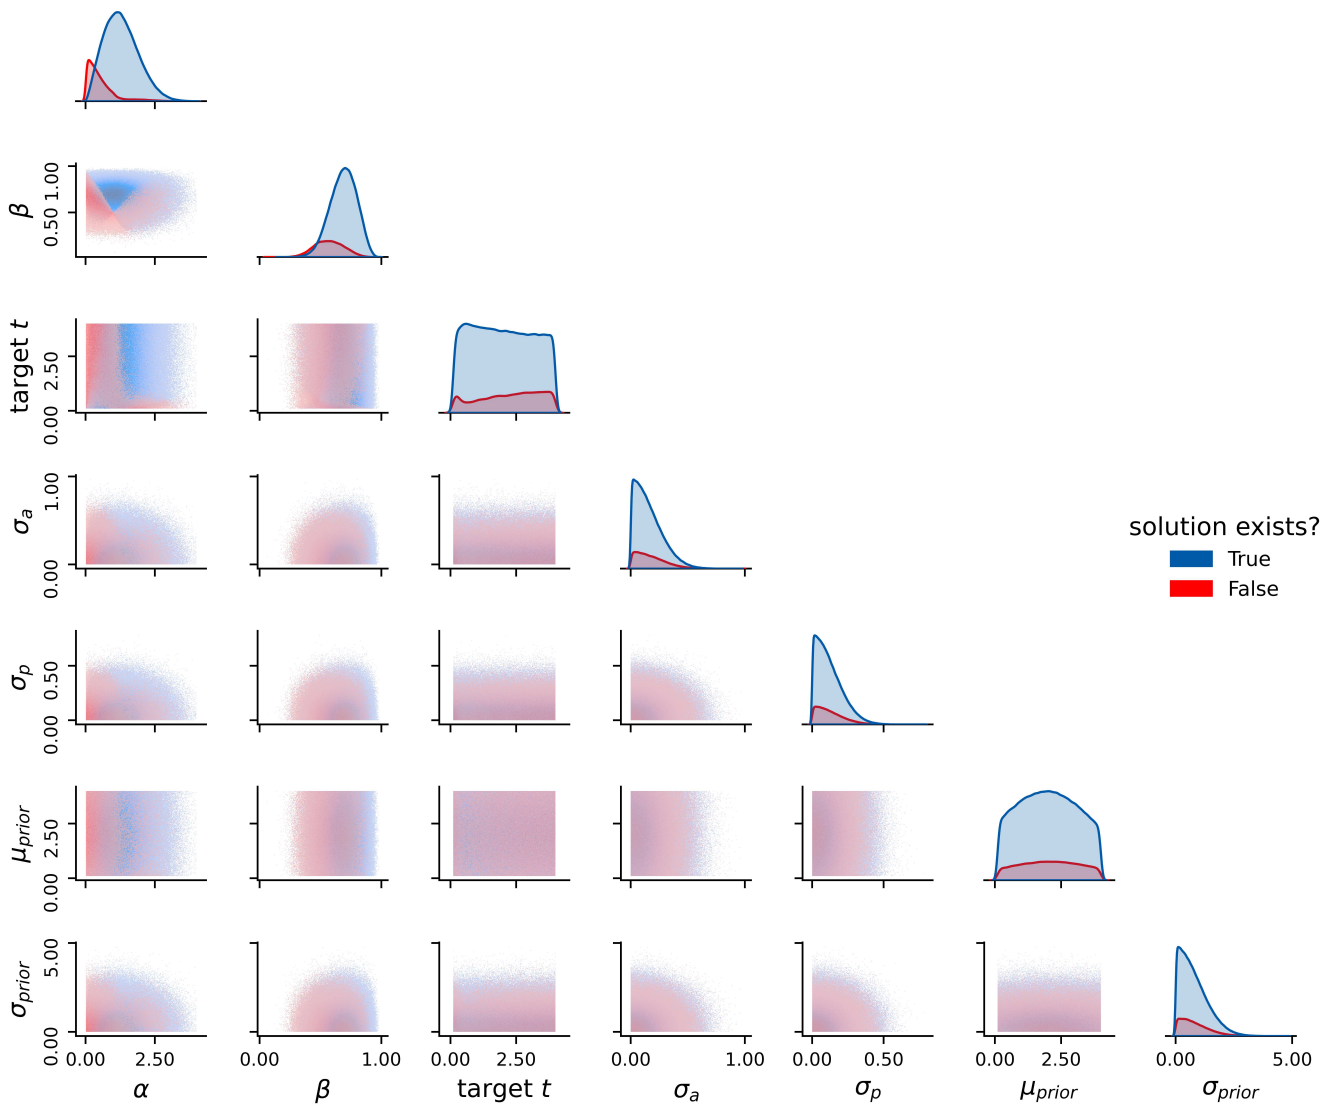


Figure 3.5.: **Solvability of the Problem for different Parameter Configurations** - Histograms of samples with (blue) and without (red) existing solution. The histogram of samples without existing solutions is plotted above the histogram of samples with existing solutions, since we focus on the cases where a solution does not exist here. There can be complex and unforeseen interactions between different parameters that determine whether a solution exists and can be found by the numerical solver.

3.3. Neural Network

The numerical optimization introduced in Section 3.2.2 is very costly in terms of computation time. Since we need the optimization within our Markov Chain Monte Carlo Sampling (cf. Section 3.4) to calculate μ_a for each sample proposal, we need a means to speed up the computation, to make MCMC Sampling feasible. For that purpose, we employ a neural network $f(z, \theta_{\text{cost}}, \theta_{\text{prior}}, \sigma)$ as an estimator for the optimal μ_a^* given the cost function parameters $\theta_{\text{cost}} = \{\alpha, \beta\}$, the latent measurement z , the respective perceptual uncertainty and motor variability $\sigma = \{\sigma_p, \sigma_a\}$ and the parameters defining the prior distribution $\theta_{\text{prior}} = \{\mu_{\text{prior}}, \sigma_{\text{prior}}\}$. The parameters of the cost function and prior distribution given here can generally be substituted with any other arbitrary cost function or distribution, we would only need to train a new network in order to do so. Section 3.3.1 introduces the architecture and details of our implementation, followed by training and evaluation procedures in Section 3.3.2.


3.3.1. Architecture

The architecture of our network consists of one input, four hidden and one output layer. The number of inputs (for the Gen2B cost function) is given by all cost parameters $\theta = \{\alpha, \beta\}$, the value of the sensory variable z , the uncertainties of the perceptual and action stage $\{\sigma_p, \sigma_a\}$ and the parameters $\theta_{\text{prior}} = \{\mu_{\text{prior}}, \sigma_{\text{prior}}\}$ defining the subject's prior belief about the stimulus distribution, thus $n_{\text{input}} = 7$. The architecture is shown in Figure 3.6.

Intuitively, one would use the target t as input for the neural network. Yet, this is not what we need here. The neural network ought not to approximate a general optimization for the optimal action to minimize the expected cost with respect to the target t but with respect to the subject's sensory estimate z of the true target. Thus, the network approximates not only the optimization process but also the inference process of the subject over the true target t , given only the noisy measurement z provided by their sensory system. It is true that we do not have access to the true value of the measurement z , so how can we use it as input to our neural network? By integrating the neural network as a component in our probabilistic model, we can treat the measurement z as latent variable in the MCMC sampling process and perform inference over its value as well. Thereby, we also get estimates of how the target is perceived through the subject's sensory system. Still, this measurement is only an auxiliary mathematical variable to represent the subject's sensory information in their inference over the true state of the world. We do not make any claims about whether these measurements are or can be explicitly represented in the brain.

The four hidden layers consist of [16, 64, 16, 8] nodes, respectively. The last hidden layer has a sigmoid activation function, while the first three hidden layers have tanh activation function.

The output of the network is given by a single node, which approximates the value of μ for which the expected loss becomes minimal – to be exact, it approximates the median \hat{a} of the action distribution that minimizes the expected loss $\mathbb{E}[\mathcal{L}(a, t)]$, which can then be used to calculate the value of $\mu = \log \hat{a}$ which produces a log-normal distribution with said median. We chose the neural network to output the median of the resulting distribution instead of the distribution parameter μ directly, since the parameter μ of a log-normal distribution is hard to interpret in a straight-forward fashion, while the median allows to do so more intuitively.



$$\mu_a^* = \operatorname{argmin} E[\mathcal{L}(a, t)] \approx \hat{f}(z, \theta_{\text{prior}}, \theta_{\text{cost}}, \sigma)$$

Figure 3.6.: **Network Architecture** – The network receives the cost function parameters θ_{cost} , prior parameters θ_{prior} as well as the sensory information z and the uncertainties $\sigma = \{\sigma_p, \sigma_a\}$ as input (thus $n_{\text{input}} = 7$) and outputs the value of μ_a^* that minimizes the expected loss, trained on a data set of samples which were numerically solved for μ_i , according to Section 3.2.2. There are four hidden layers with [32, 64, 16, 8] nodes each. Hidden layers all have tanh-activation functions except the last layer which has ReLU activation. The output layer also has ReLU activation to constrain the output to the positive domain. The number of nodes in the image does not correspond to the actual number of nodes for all layers for visual reasons.

3.3.2. Data, Training and Evaluation

Especially in scenarios like ours, where the existence of a solution strongly depends on the input parameters (cf. Section 3.2.4) and we can easily leverage knowledge about realistic parameter values (e.g. the perceptual uncertainty will be rather small in a healthy subject), the training and test data needs to be chosen carefully. Thereby, we can achieve faster and more stable training than if we were to just sample the data in the complete domain of real numbers. We evaluate the performance of the network in terms of accuracy as well as computation time in contrast to the numerical solver.

Training and Test Data

The training and test data for our network was generated by randomly sampling the cost function parameters θ_{cost} , the uncertainties $\{\sigma_a, \sigma_p\}$, the prior parameters θ_{prior} and the target t as well as its latent measurement z for each synthetic trial. The distribution we use for generating these samples are constrained to produce realistic values that are most likely to occur in the sampling process when using real empirical data. For instance, our network does not need to be trained for high perceptual uncertainties σ_p , since we can assume that the sensory system of a healthy subject can produce a somewhat reliable measurement of the real stimulus. The exact priors for cost function Gen2B are shown in table 3.2.

The chosen support of each distribution is based on our experiences, regarding which values for different cost function parameters are most likely to result in realistic cost functions or are possible with the chosen parameterization at all. We also confirmed our assumptions with prior predictive checks, which are introduced along the probabilistic model in Section 3.4.1.

The support of the target distribution would either need to be adjusted for experiments with targets in a

completely different range or the targets of such experiments would need to be transformed into the scale supported by our model. In this case, all parameters must be interpreted on this transformed scale of target values or an inverse transformation for the estimated parameters needs to be derived.

Parameter	Distribution	Distribution Parameters
α	Truncated Normal $\mathcal{N}_{\text{trunc}}(\mu, \sigma)$	$(\mu, \sigma) = (1.0, 0.8)$ in $[0.01, 4]$
β	$x \sim \text{Beta}(\alpha, \beta)$ with $\beta = 0.4 + 0.6x$	$(\alpha, \beta) = (3, 2)$
σ_a	Half-Normal $\mathcal{H}(\sigma)$	s.t. 95%-quantile $q_{.95} = .2$
σ_p	Half-Normal $\mathcal{H}(\sigma)$	s.t. 95%-quantile $q_{.95} = .2$
t	Uniform $\mathcal{U}(a_t, b_t)$	$(a_t, b_t) = (.1, 4.8)$
z	$\log \mathcal{N}(\mu, \sigma)$	$(\mu, \sigma) = (t, \sigma_p)$
μ_{prior}	Truncated Normal $\mathcal{N}_{\text{trunc}}(\mu, \sigma)$	$(\mu, \sigma) = (\frac{b_t - a_t}{2}, 2.0)$ in $[.1, 4.8]$
σ_{prior}	Half-Normal $\mathcal{H}(\sigma)$	s.t. 95%-quantile $q_{.95} = .2$

Table 3.2.: **Priors for Data Generation** - The different distributions for generating synthetic data for training the neural network. Note that the distribution-specific parameters in the right columns do not correspond to the model parameters in the left-most column in any way. This table only displays the priors for the cost function Gen2B.

We then solve each sampled configuration of parameters $\{\theta_{\text{cost}}, z, \theta_{\text{prior}}, \sigma\}$ and pass it on to our numerical solver, which is described in Section 3.2.2. Thereby, we obtain the numerical solution for the optimal μ_a^* minimizing the expected cost for each particular parameter configuration. We use these as the desired output in our training and test data.

The complete data set consists of $n_{\text{data}} = 1,000,000$ data points, of which $n_{\text{failed}} = 165,145$ samples were discarded because the optimization failed. This can happen because some parameter configurations might not have a solution due to the natural properties of the model they define. We tried to avoid this by carefully analyzing the solvability of the problem in different subspaces of the complete parameter space and adjusted our data generation priors accordingly. Still, this could not be avoided completely. For further details on the numerical solving and optimization process, please consult Section 3.2.

From these remaining $n_{\text{solvable}} = 834,855$ samples, 80% were taken to form the training set and the other 20% to comprise the test set, resulting in $n_{\text{train}} = 661,575$ training samples and $n_{\text{test}} = 173,280$ test samples. Again, these numbers only apply to cost function Gen2B.

Training Procedure

For the training, we provided all of the training data set with a batch size of 32 for 600 epochs and used a validation split of 0.2. The loss for the neural network \mathcal{L}_{NN} during training is defined by the Mean Squared Error (MSE) between the numerical solution for the optimal μ_a^* and the network’s prediction $\mu_a^{\text{predicted}}$, thus

$$\mathcal{L}_{\text{NN}}(\mu_a^*, \mu_a^{\text{predicted}}) = \|\mu_a^* - \mu_a^{\text{predicted}}\|_2$$

We also use early stopping on the MSE loss with the patience set to 60, usually leading to early convergence before the 600 epochs are fully rolled out. For the tuning of weights, we use the *RMSprop*-optimizer (Tieleman et al., 2012) with a learning rate of 0.001 and $\rho = 0.9$.

Evaluation Procedure

In order to validate that our networks approximate the correct solution, we use the numerically solved samples in our test set, consisting of $n_{\text{test}} = 173,280$ samples for cost function Gen2B and evaluated the difference to the numerical solutions. Since we use numerical methods to solve for each set of parameters (see Section 3.2.2 for more details), the solutions in our test set do not contain the exact solution, but only a numerical approximation. While it is crucial to design an accurate neural network in order to prevent strong error propagation, we still only need a solution that is sufficiently good in order for the MCMC Sampling to work reliably. The results of the evaluation of the neural network are shown in Figure 3.7. On the training data, our network achieved a mean absolute error (MAE) of 0.016, a root mean squared error (RMSE) of 0.028 and a mean absolute relative error (MARE) of 0.045. On the test set, it achieved a MAE of 0.016 as well and a slightly better RMSE of 0.028 and MARE of 0.041. Since the targets in most of our experimental and generated synthetic data (cf. Section 4) lie within the range of 1 to 5 (arbitrary unit), we judge an accuracy of about one to two hundredths to be sufficiently accurate.

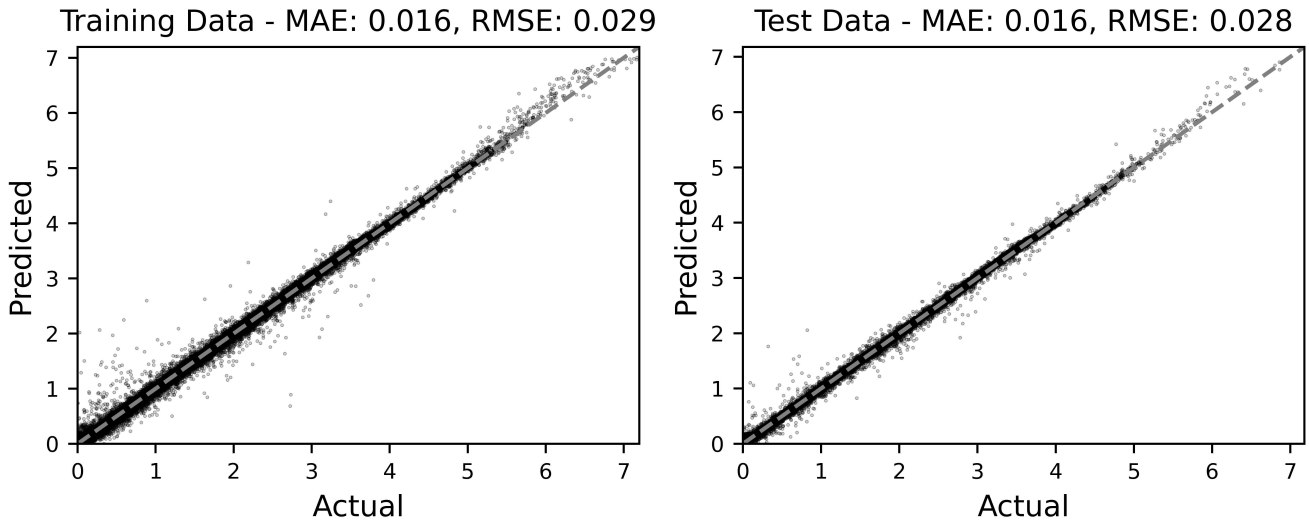


Figure 3.7.: **Performance of the Neural Network** - On the training set consisting of $n_{\text{train}} = 661,575$ samples our network achieved a MAE of 0.016, a RMSE of 0.029 and a MARE of 0.045 when predicting the numerically found optimal μ_a^* . On the test set consisting of $n_{\text{test}} = 173,280$ samples, our network achieved a MAE of 0.016, a RMSE of 0.028 and a MARE of 0.041. The grey dashed line shows the optimal performance, on which all points would lie if each prediction was exactly equal to the ground truth.

Speed-Up in Computation Time

By using a neural network as an approximator for the solution of the optimization problem, we achieve a strong boost in the necessary computation time. Figure 3.8 shows the necessary computation time (y-axis, please mind the log-scaling) for solving a set of n samples (x-axis) with the numerical solver or the neural network either as a batch of all samples or by iteratively passing each parameter configuration as a batch with only one data point. By using a neural network, we are about 1,000 to 10,000 times faster than the numerical

solver. This makes it feasible to use our problem within a Markov Chain Monte Carlo setting, since MCMC sampling needs thousands of samples to ensure convergence to the stationary posterior distribution. If we were not using a neural network, inferring the model parameters would take an unreasonable amount of time. We still need to run the numerical solver on a huge amount of data points in order to be able to train the network, which will take a lot of time, but the respective network can be used as many times as needed after being trained once. The speed up in the MCMC sampling is twofold. Obviously, a forward pass in a neural




Figure 3.8.: **Computation Time with and without Neural Network** - Computation time of the numerical solver and of the neural network with data batches of different sizes. We passed the data to the neural network as one batch with all data points or iteratively as single-datum-batches. Either way, we observe a speed-up factor of about 1,000-10,000.

network needs less time to compute a solution than using a numerical solver which potentially uses iterative methods. Furthermore, by using a neural network in our probabilistic model, we preserve differentiability of the model which enables us to use gradient-based kernels in the MCMC inference. Since gradient-based approaches to MCMC methods can employ additional knowledge about the parameter space it draws the posterior samples from, we also gain increased sample efficiency which reduces the number of necessary samples to obtain good results.

3.4. Parameter Inference

For inference of the cost function parameters θ as well as the perceptual and motor uncertainties $\sigma = \{\sigma_a, \sigma_p\}$, we use Markov Chain Monte Carlo (MCMC) sampling. Since MCMC sampling requires drawing a lot of proposals for the parameters we want to infer and solving the optimization each time numerically is computationally very expensive, we trained a neural network which approximates the numerical solutions as described in Section 3.3. As shown in Section 3.3.2, this approach speeds up computation by about three orders of magnitude.

By employing MCMC, we obtain estimations of the true cost function parameters, the uncertainties and the priors beliefs of the subject. Further, since MCMC is a Bayesian method, we also have a measure of confidence for our estimations.

In Section 3.4.1 the probabilistic model including our choice of the prior and likelihood distributions is introduced, which is then validated by prior predictive checks in Section 3.4.2. Section 3.4.3 describes the sampling process including our hyperparameter and kernel choices. The whole MCMC process was implemented in Python using *TensorFlow Probability* (Dillon et al., 2017).

3.4.1. Probabilistic Model

The probabilistic model used for the inference by Markov Chain Monte Carlo sampling is shown in Figure 3.9. The prior distributions underlying each variable in the model are roughly the same as those used in the data generation for the neural network, except for the targets t_i , which – together with a set of action observations a_i for each trial i – are part of the observed data and thus do not need a prior for they are not inferred. The prior distributions for each latent variable were first chosen as educated guesses based on solvability of the problem for different ranges of the parameters as shown in Section 3.2.4. For example, the distribution for β is transformed by scaling and shifting such that the overall range of the samples is not $[0, 1]$, as it is the case for the original beta distribution, but $[0.4, 1.0]$. This stems from our solvability analysis which has shown that most parameter configurations with $\beta < 0.4$ do not have an optimal solution. The uncertainties σ are chosen to be very small, with the 95%-quantile $q_{0.95}$ of Half-Normal distributions $\mathcal{H}(\sigma)$ chosen such that $q_{0.95} = 0.2$, since empirical data shows only weak variability within the perceptual and motor dimensions. It is essential here, that σ_p and σ_{prior} are sampled from the same range in order to allow for similar magnitudes and thus reasonable values of ratios $\sigma_p/\sigma_{\text{prior}}$, since this crucially determines how strongly the subject relies on their prior knowledge in contrast to their sensory information from the measurements z_i .

The range for μ_{prior} is chosen such that it cannot exceed the biggest values of the observed targets t_i , since we judge prior beliefs outside of the stimulus distribution to be nonsensical. Our choice of priors is also further validated with prior predictive checks in Section 3.4.2. The likelihoods of the measurements z_i and actions a_i are chosen as log-normal distributions. The likelihood distribution for the measurements can easily be justified to be log-normal because log-normal distributions implicitly implement the Weber-Fechner-Law (cf. Section 2.1.1). Moreover, the log-normal distribution's probability mass density is naturally skewed towards smaller values which can also help to explain undershoots frequently found in (re-)production tasks and perception-action-tasks (cf. Section 2.1.5). The likelihood of the action distribution is also chosen as a log-normal distribution since we can observe signal-dependent noise in most behavioral data. That is, the noise of a distribution increases with increased signal strengths. In our context, this means that we observe an increase in motor variability for larger stimuli – a property that log-normal distributions also naturally possess.

$$z_i \sim \log \mathcal{N}(t_i, \sigma_p)$$

$$\theta_{\text{cost}} = \{\alpha, \beta\}$$

$$\alpha \sim \log \mathcal{N}(1.0, 0.8) \text{ in } [0.01, 4]$$

$$\beta \sim 0.4 + 0.6x \text{ with } x \sim \text{Beta}(3, 2)$$

$$\sigma_a \sim \mathcal{H}(0.1)$$

$$\sigma_p \sim \mathcal{H}(0.1)$$

$$\theta_{\text{prior}} = \{\mu_{\text{prior}}, \sigma_{\text{prior}}\}$$

$$\mu_{\text{prior}} \sim \log \mathcal{N}\left(\frac{t_{\text{max}} - t_{\text{min}}}{2}, 2.0\right) \text{ in } [0.1, 4.8]$$

$$\sigma_{\text{prior}} \sim \mathcal{H}(0.1)$$

$$\mu_{a,i}^* = \underset{\mu_{a,i}}{\text{argmin}} \mathbb{E}[\mathcal{L}(a_i, t_i)]$$

$$a_i \sim \log \mathcal{N}(\mu_{a,i}^*, \sigma_a)$$

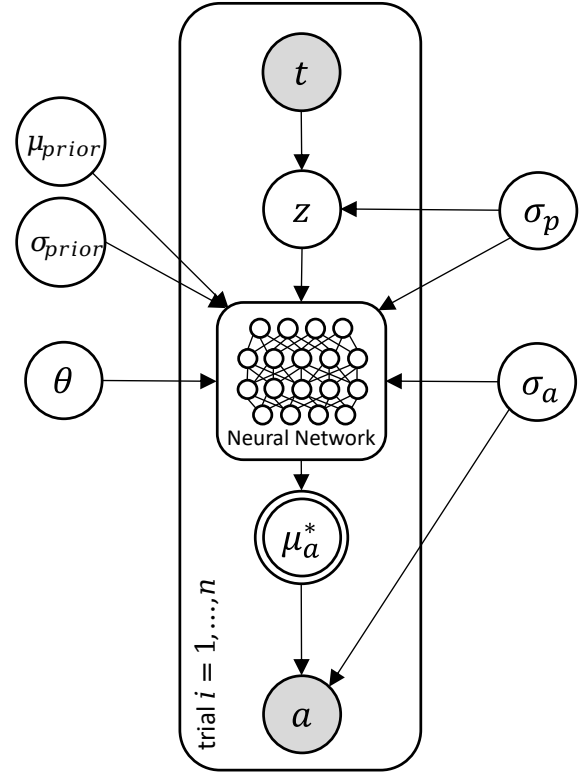


Figure 3.9.: **Probabilistic Model** - Model in our framework using n trials of an arbitrary experiment and as implemented in the MCMC sampling. The observed values are sets of targets t_i and actions a_i . Treating the measurements z_i as latent variables makes the problem strongly high-dimensional, since one needs to be inferred for each target observation. The neural network receives the perceptual uncertainty σ_p , motor variability σ_a , cost function parameters θ_{cost} and prior parameters ($\mu_{\text{prior}}, \sigma_{\text{prior}}$ and returns the approximately optimal action distribution parameter μ_a^* . Distributions \mathcal{H} denote Half-Normal distributions and some log-normal distributions are truncated to intervals $[a, b]$.

The set of latent variables that need to be inferred is six-dimensional already with the free parameters of the model – the prior parameters θ_{prior} , uncertainties $\sigma = \{\sigma_p, \sigma_a\}$ and cost parameters $\theta_{\text{cost}} = \{\alpha, \beta\}$. But since we treat the subject's measurements z_i as latent variables as well, the dimensionality of the model increases rapidly. A measurement z_i is sampled in the sensory space given by a stochastic log-normal transformation for each target t_i . We do not take any position regarding the existence of explicit scalar representations of this sensory information in the brain, but sampling the measurements explicitly has the advantage of easily interpretable subjective values for each target, which may help explain why certain trials resulted in certain actions. This also allows for noisy inputs to the optimization problem which aligns very well with our assumption that subjects do not have access to the ground truths of the target values.

Yet, consider a data set consisting of 200 trials, which we usually used in our synthetic data sets for evaluation of our methods (cf. Section 4.1). Since each trial has its own target t_i and action a_i , it also has a latent measurement z_i . This makes the parameter space in which we need to perform the inference 206-dimensional and thus a really hard inference problem.

3.4.2. Prior Predictive Checks

In order to validate our choice of priors and not use handwavy educated guesses, we performed prior predictive checks with the priors we ultimately decided for. We sampled 1,000 configurations of parameters as prior predictive samples from our prior distributions and solved the model for an experimental setup of eight targets with 25 trials each. Therefore, we generated 200 actions in total for each prior predictive sample. The results are shown in Figure 3.10.

The prior predictive checks show that our choice of priors allows for a wide range of possible actions. The mean of the probability mass of the joint prior distribution is distributed in a way that favors undershoots a little, which is exactly the behavior we want to model. Still, the priors allow for freedom in the way the actions are distributed around the target. Also, the priors should represent prior knowledge we have as researchers. This includes the human tendency to show certain response patterns like regression-to-the-mean behavior or undershoots (cf. Section 2.1.5). In summary, the prior predictive checks validate our choice of priors.

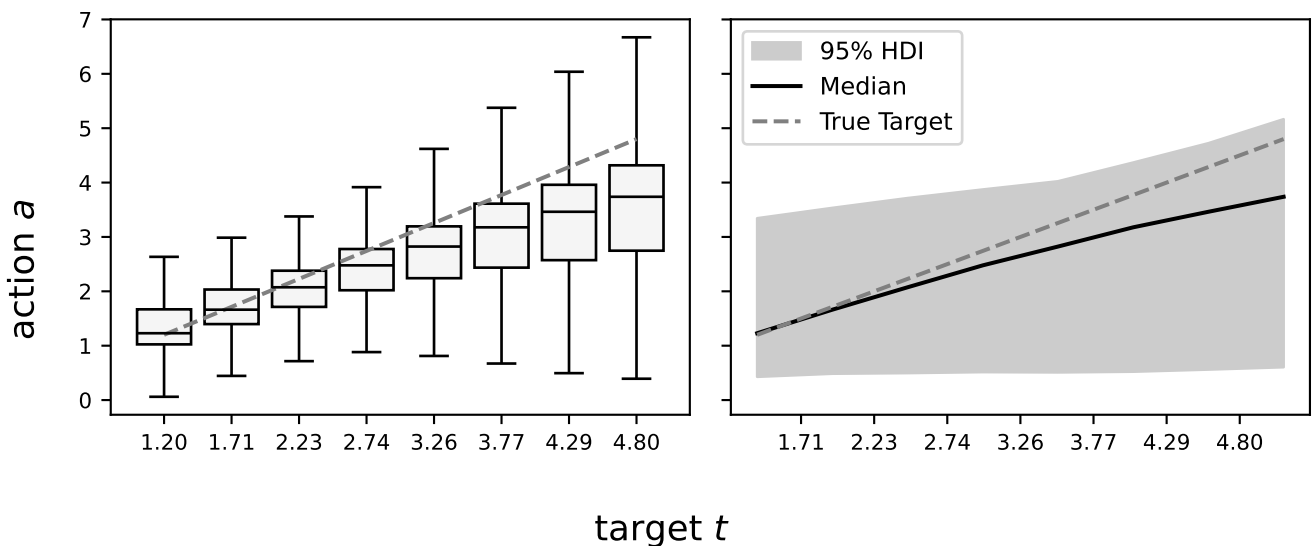


Figure 3.10.: **Prior Predictive Checks** – We have drawn 1,000 prior predictive samples from our prior distributions. Each sampled parameter configuration was then fed into our model and we generated actions for 8 targets with 25 trials each per parameter configuration. The left plot shows box plots for each target, while the right plot shows the median action and the 95% interval of all 200,000 generated actions. The grey dashed line denotes equality of target t and action a .

3.4.3. Sampling Process

The way in which the posterior distribution is inferred in MCMC sampling is strongly dependent on the used kernel. For example, we could use random-walk methods for exploring the posterior distribution, like the classic *Metropolis-Hastings-Algorithm (MH)* (Metropolis et al., 1953; Hastings, 1970). In most cases, those MH-methods are outperformed by *Hamiltonian Monte Carlo (HMC)* approaches which make use of the geometrical properties of the target distributions spaces they try to explore and thus usually converge to the target distributions much faster (Betancourt, 2018). Ultimately, we decided for the gradient-based *No-U-Turn-Sampling (NUTS)*-algorithm (Hoffman et al., 2014), which is a state-of-the-art algorithm in the domain of

Hamiltonian Monte Carlo, applying a more sophisticated method for determining the next sampling steps and needing much less manual tuning. In order to exploit the geometrical properties of the target distribution, our model must be fully differentiable for NUTS to work reliably. While most numerical solutions work iteratively and are thus usually not differentiable, simply using a numerical solver within our probabilistic model will not allow us to make use of the advantages of HMC-methods. By approximating the numerical solution with a neural network, we not only speed up the computation considerably, but also maintain differentiability, since neural networks are differentiable by definition in order to be trainable via backpropagation at all. For further details on our model, consult the previous Section 3.4.1.

For each MCMC run, we ran eight chains with 20,000 steps after warming up our model with 5,000 burn-in steps to ensure that the chains are actually sampling from a stationary posterior distribution. The initial states for each chain were drawn from the prior distributions of each parameter and the sensory variables z_i were initialized on noisy versions of the observed targets, given by $\mathcal{N}(t_i, 0.1)$.

As introduced already, we are using a *NUTS*-kernel. This is supported by a Dual-Averaging-Step-Size-Adaptation (Section 3.2 of Hoffman et al. (2014)) with 4,000 adaptation steps and a target acceptance probability of 0.6. Additionally, we use bijectors on the parameters to map them from the unconstrained parameter samples from which MCMC is sampling into their respective supported domains. Parameters with only positive support (i.e. α , uncertainties σ_a and σ_p , prior distribution parameters μ_{prior} and σ_{prior} , measurements z_i), are constrained to the real, positive domain using a softplus bijector, which is a smoothed version of a ReLU-function. For a random variable X transformed via the softplus bijector $g(\cdot)$ to a positive real variable Y , the forward transformation is given by $Y = g(X) = \log(1 + e^X)$.

Since the parameter β is responsible for the weighted combination of the cost for deviating from the target t and the effort cost of performing the action a itself, it is already constrained to the range $[0, 1]$. Further investigation of the solvability for different parameter configurations (cf. Section 3.2.4) constrained the range of β to $[0.4, 1.0]$ in order for a solution to the optimization problem to exist at all. Thus, we use a sigmoid bijector which applies the logistic sigmoid function to the non-transformed parameter proposals from $(-\infty, \infty)$ to obtain a transformed version of them constrained to $[0, 1]$.

Since the inference problem becomes highly multidimensional by treating the measurement z_i for each single target t_i as a latent variable that also needs to be inferred, the inference becomes extremely difficult. In some cases, the sampled Markov chains were stuck in their initial state for the whole sampling process. To avoid distorting the inferred posterior distributions from the converged stationary chains with the stuck chains, we removed these after each run. To achieve full chain-efficiency for inferring the posterior distributions, further fine tuning of our model might be necessary.

4. Results

In this chapter, we benchmark our method over different synthetically generated data sets. We generated data sets representing characteristic response patterns as described in Section 2.1.5 and randomly generated parameter configurations over wide ranges of the respective parameters. The complete evaluation on synthetic data is shown in Section 4.1. Furthermore, we retrieved two data sets with empiric data of a bean bag throwing task and a puck sliding task. The results of our method on these data is presented in Section 4.2.

4.1. Synthetic Data

Before evaluating our method on empirical data, we generated synthetic sets of data for which we know the ground truth values that should be recovered by our approach. Thereby, we can explore the limits of our implementation within a controlled environment.

The performance of our method on manually tuned regimes representing different characteristic behavior like the ones introduced in Section 2.1.5 is evaluated in Section 4.1.2. Results for a large-scale evaluation of our methods on randomly generated synthetic data is then presented in Section 4.1.4. Furthermore, we also examined the influence of the amount of trials in the experimental setup on the performance of our method in Section 4.1.5.

The synthetic data sets can be divided into a first, manually tuned data set for which we identified different parameter configurations that result in qualitatively different behavior. We will refer to these different data sets as different ‘regimes’ in the following. The regimes can be divided into the categories ‘*Regression to the Mean*’, ‘*Undershoot via Prior*’, ‘*Undershoot via Cost*’ and ‘*Hitting Target*’. For designing these regimes, we identified different patterns from empirical behavioral data (Petzschnr and Glasauer, 2011; Jazayeri and Shadlen, 2010; Neupärtl et al., 2020; Willey and Liu, 2018) and decided on how they can come about with different parameter configurations of our model by simulating the results with predictive forward passes through our model. During these simulations, we learned that the cost function parameter α has only minor influence on the resulting behavioral data. Because of that, we decided to fix $\alpha = 2.0$ and by that preserve the possibility of comparing our model using an analytically tractable closed-form solution (cf. 3.2.3) and the neural network as approximator for the parameter μ_a^* of the Bayesian actor’s action distribution. The identified regimes and their parameter configurations are shown in Table 4.1.

The *Hitting Target* regime represents a subject which is capable of hitting the targets accurately. They completely rely on the sensory evidence by having low perceptual uncertainty with respect to their prior uncertainty. They also have a rather high value of β and thus put a lot of emphasis on the accuracy of hitting the target without applying much cost for the necessary effort for performing the action.

Regime	Parameters	Generated Data
Hitting Target	$\alpha = 2.00$ $\beta = 0.90$ $\mu_{\text{prior}} = 3.0$ $\sigma_{\text{prior}} = 2.0$ $\sigma_a = 0.10$ $\sigma_p = 0.20$	
Undershoot via Prior	$\alpha = 2.00$ $\beta = .95$ $\mu_{\text{prior}} = 1.50$ $\sigma_{\text{prior}} = 1.0$ $\sigma_a = 0.10$ $\sigma_p = 0.10$	
Undershoot via Cost	$\alpha = 2.00$ $\beta = 0.60$ $\mu_{\text{prior}} = 2.50$ $\sigma_{\text{prior}} = 0.30$ $\sigma_a = 0.10$ $\sigma_p = 0.10$	
Regression to the Mean	$\alpha = 2.0$ $\beta = 0.80$ $\mu_{\text{prior}} = 3.40$ $\sigma_{\text{prior}} = 0.05$ $\sigma_a = 0.10$ $\sigma_p = 0.10$	

Table 4.1.: **Synthetic Regimes** – Manually defined regimes that resemble characteristic response patterns of human behavior in (re-)production tasks. The plot on the right shows the median actions with their 95%-HDI over the targets. Orange dots are the generated actions, red dashed line denotes equality between action and target and the solid grey line shows the value of μ_{prior} .

The *Undershoot via Prior* regime shows undershoots as they are observed quite often in empirical data (cf. Section 2.1.5). These undershoots are generated by having a high reliance on the prior in contrast to the sensory evidence by having a low prior uncertainty. We also use a high value for β , such that the observed effect is really only caused by the prior and not by effort costs.

In contrast to this, the *Undershoot via Cost* regime also generated undershoots, but caused by the parameter β instead of the prior. The main difference herein is that the generated actions do not tend towards the prior but are rather consistently set below the real target, since the low value of β puts the main focus on the effort cost instead of the cost for missing the target.

The *Regression to the Mean* regime generates a regression to the mean behavior by really strong strong reliance on the prior. Also, the mean of the prior is much higher than in the *Undershoot via Cost* regime, thus not only creating undershoots but also overshoots for lower stimulus intensities.

A second data set consists of 100 different parameter configurations randomly drawn from the priors also used in our MCMC process and the respective simulated actions (8 targets á 25 trials each). Thereby, we can acquire statistics on how well our method performs on a larger scale, e.g. by determining how often the true value lies within the 94% high density interval (HDI) of the inferred posterior distribution or by calculating the mean absolute error (MAE) or root mean squared error (RMSE) to the true value when using the inferred posterior's mean as a predictor.

4.1.1. Numerical Versus Analytical Results

As a proof of concept, we tried to tune most of the synthetic regimes such that they remain analytically tractable and compare the inference results incorporating the analytical solution or the neural network in our probabilistic model (cf. Section 3.4.1). Since we have shown in Section 3.2.3 that there is a closed form solution to the problem if we choose $\alpha = 2$ and because our experience with the data has shown that the influence of α is not that big in comparison to the other parameters, we also tried to choose $\alpha = 2$ in all manually tuned regimes. Thereby, we can define our usual probabilistic model using the pre-trained neural network (cf. Section 3.3) and another model using the closed-form analytical solution to perform inference over the data. Figure 4.1 shows the results of the model with the analytical solution, which we assume to be the ground truth we want to estimate with our approximate inference, and with the network solution for the *Regression to the Mean* data set. The results show that the posteriors inferred by our model using the neural network are the same or at least similar to the posteriors inferred by a model using the analytical solution. The posterior for the parameters β and σ_a almost completely cover each other with no remarkable differences. For parameters μ_{prior} , σ_{prior} and σ_p , the posterior distributions are a little more skewed or shifted from the results with the analytical model but still capture approximately the same shape and range of values. We cannot compare the parameter α , since it is held constant at its true value of 2.0 for the analytical model to enable us to use the analytical solution, since we only know the closed-form solution for $\alpha = 2.0$ (cf. Section 3.2.3).

Both models reproduce the observed actions quite well. The mean and 95%-HDIs of the posterior predictive samples are almost exactly the same for the analytical model and the model using our neural network.

The same results for the other regimes can be found in Appendix A.5 to A.7. In all other regimes, the analytical model and our model both produce the same means and HDIs for the posterior predictive samples, which almost completely capture the set of observed actions. However, the inferred posterior distributions differ. For the *Undershoot via Prior* regime, both models produce very similar posterior distributions. The differences increase for the *Undershoot via Cost* regime, but still the means of the inferred posterior distributions behave

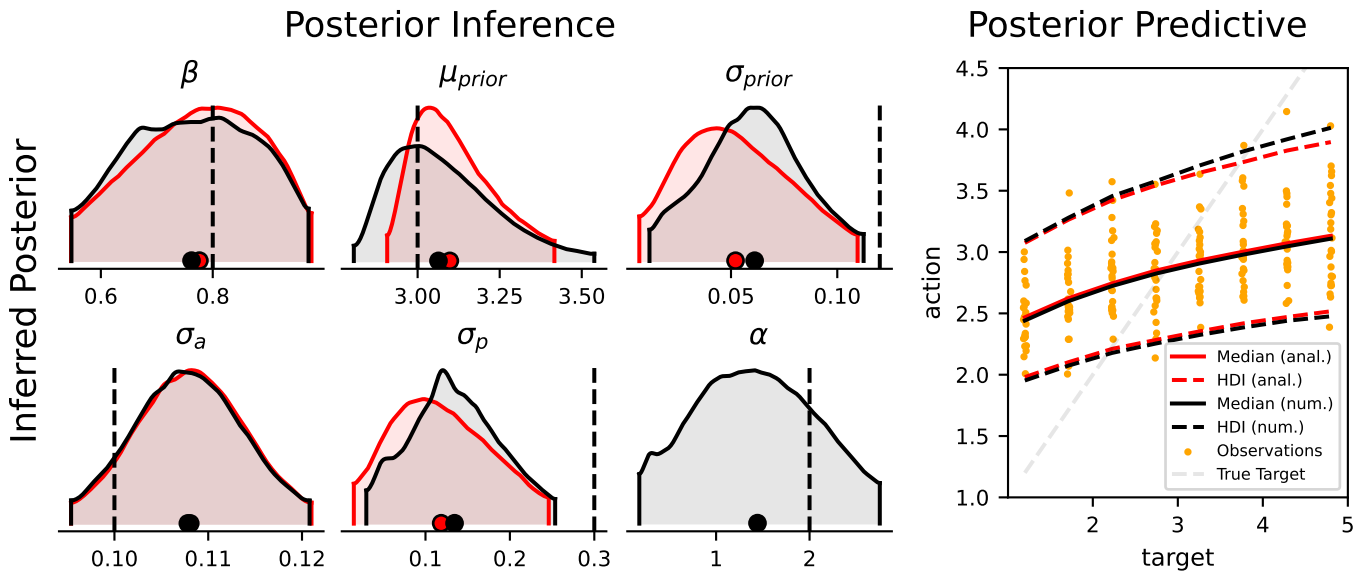


Figure 4.1.: **MCMC Results with Analytical vs. Neural Network in the Bayesian Actor Model** - Inferred posterior distributions on the *Regression to the Mean* data set over the model parameters θ using an analytical solver (red) and a neural network (black) within the model to calculate the action distribution's optimal parameter μ_a^* . Markers on the axis denote the bounds of the 95% high density interval for the respective parameters. Dashed black lines show the ground truth within the used data set. Black and red dots show the mean of the inferred posterior distribution with the analytical solver and the neural network, respectively. The right-most plot shows the posterior predictive samples for each inference method in the respective colors as well as the true observations in orange. HDIs shown are the 95%-HDIs. Results for other regimes are shown in Appendix A.5 to A.7.

about the same.

For the *Hitting Target* regime, the results seem to differ more strongly. For the prior parameters, the posterior are about the same. For σ_a and σ_p the differences seem even more extreme but are probably due to strong peaks for very low values, which we did not find a detailed explanation for yet. They probably come about because the results of the *Hitting Target* regime are the most unstable and overall worst of all regimes – as we will present in the next section –, which is probably due to our model being designed for regression to the mean and undershoot behavior and not for subjects hitting the target on average over the whole stimulus range.

Furthermore, these results should be examined again without inferring α in the model with our neural network, since it is also not inferred in the model using the analytical solution and thus our model needs to solve a harder inference problem, which reduces comparability between the two models.

4.1.2. Recovery of Characteristic Parameter Configurations

The results from the four regimes introduced in Table 4.1 are shown in Figure 4.2 and the exact values are listed in Table 4.2. Overall, the inference works quite well but shows some weaknesses in different regimes. Moreover, not all regimes exhibit favorable MCMC diagnostics.

For the *Hitting Target* regime, the cost function is recovered really accurately. However, the prior is completely off, which can be explained by its weak influence on the resulting actions, since this regime is constructed with a strong focus of the subject on its sensory information. Also, the perceptual uncertainty and motor variability show rather high inference errors in comparison to the error in these parameters for other regimes. When we take a look at the MCMC diagnostics, we find really concerning effective samples sizes (ESS) and values of \hat{R} . Remembering that we draw 20,000 samples per chain, effective samples sizes of 10 to 30 are astonishingly bad. The convergence diagnostic \hat{R} should be as close to 1.0 as possible, in the best case not bigger than 1.02. This is violated for all parameters in this regime. The chains of this inference run also converged to two different points. We see that the sampling process is overall really unstable here and it is somehow remarkable that the inference yields results anywhere near the true values for some parameters. The bad results in this regime probably become about because it represents a case our model was not designed for, since it the regime does not correspond to any of the characteristic response patterns humans show in reproduction tasks, e.g. regression to the mean behavior or undershoots. Especially the high value of σ_{prior} is something our priors are not suited for. While our model allows for high ratio between σ_p and σ_{prior} to allow for low reliance of the subject on either the prior or the sensory evidence, our priors do not allow for high absolute value like $\sigma_{\text{prior}} = 2.0$ (cf. Table 3.2). If one can find another parameter configuration giving rise to the same behavior, this experiment should be repeated.

In the *Undershoot via Prior* regime, the cost function is not recovered very well. The 95%-HDI of the inferred cost function barely contains the true cost function, which probably arises because the inference fails to capture the increased value of β with an error of 0.2. This also makes sense in a regime where the behavior is primarily generated through influence of the prior. Since the prior has the most influence on the resulting behavior, it is easily inferred and vice versa it is much harder to infer the true values of the cost function.

The inference of α fails with an error of 0.48. Inaccurate inference of α is a pattern that we will frequently face in the results. This speaks for a small influence of alpha on the resulting cost function overall and could be improved by trying an alternative parameterization or functional form of the cost function used in our model. However, the inference of the variabilities σ_a , σ_p and σ_{prior} is accurate with errors not exceeding 0.02 and the inferred posterior mean of σ_a being equal to the ground truth value. The value for μ_{prior} is also good with an error of only 0.13.

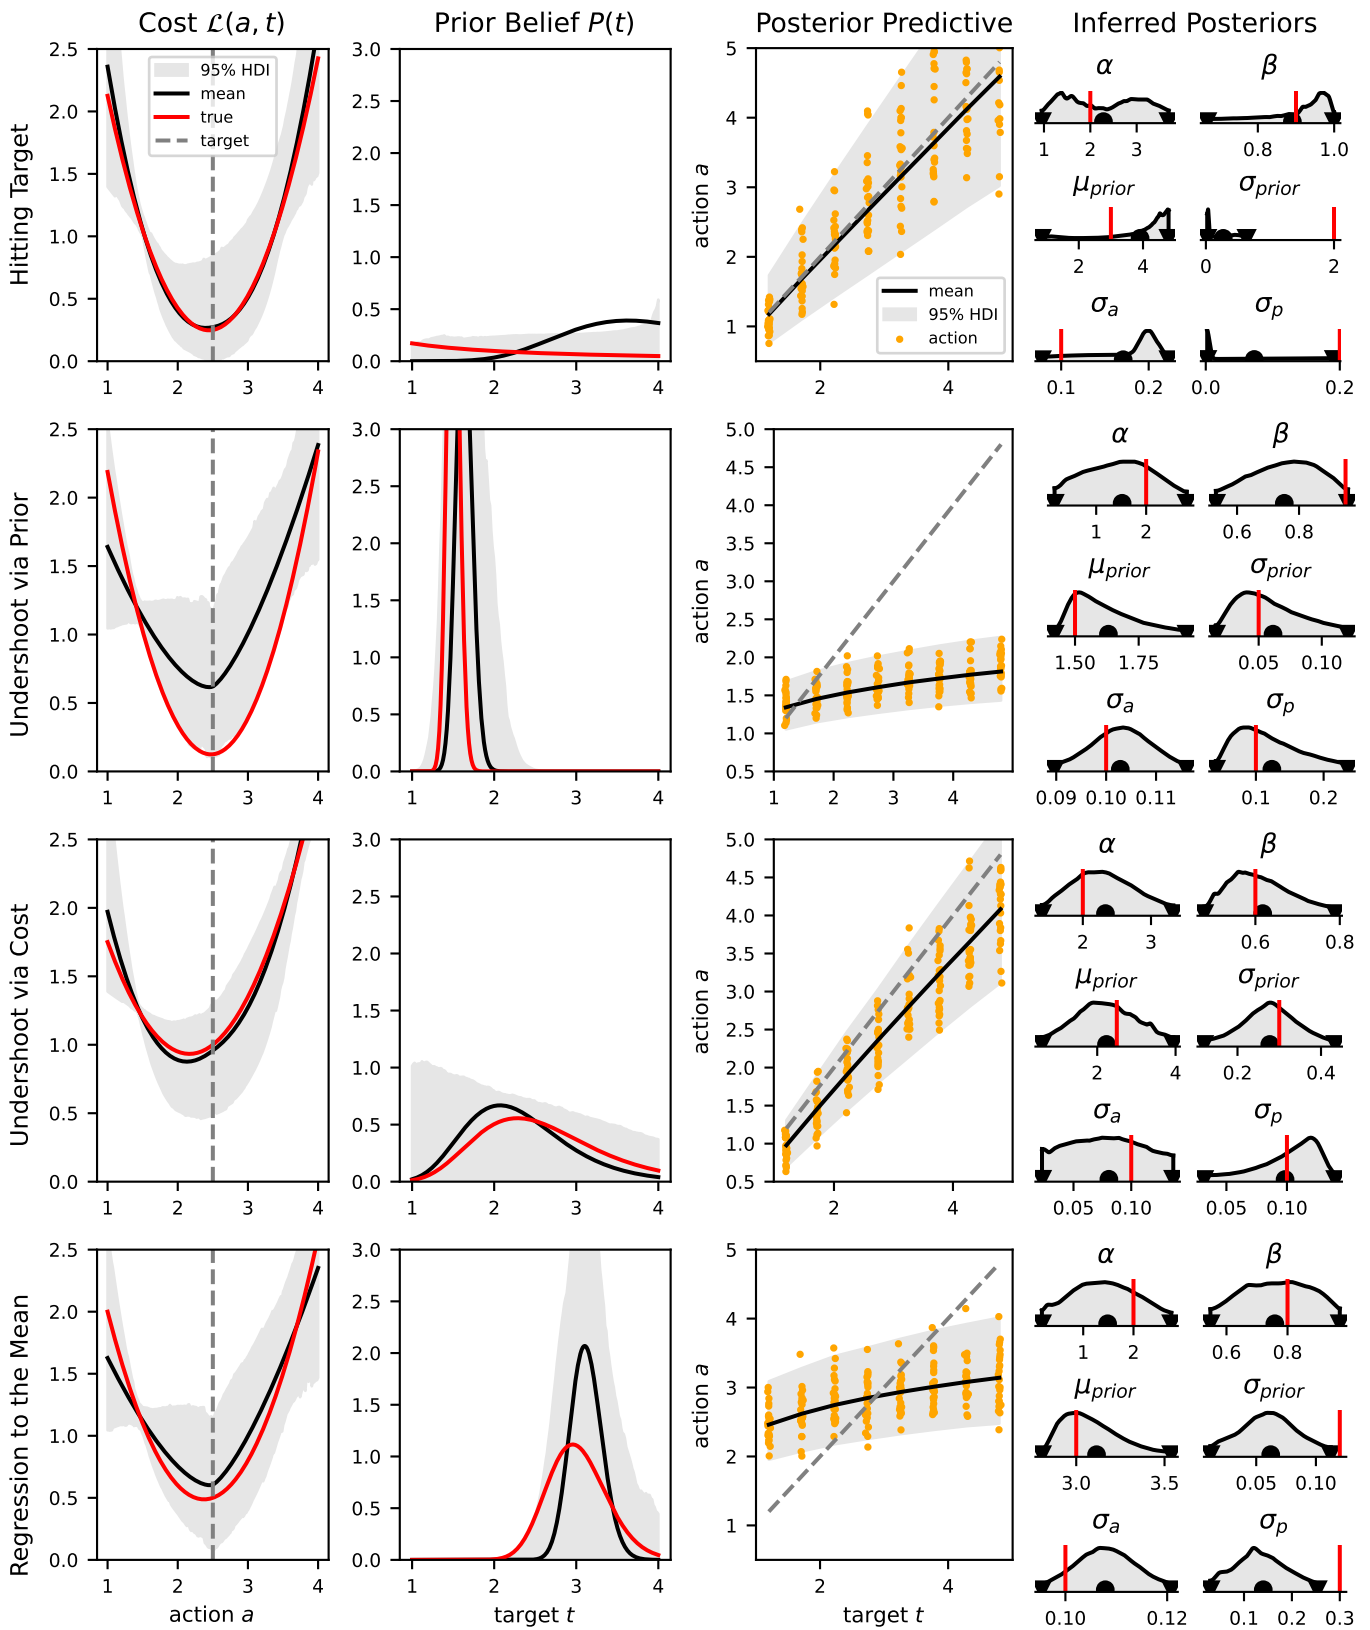


Figure 4.2.: **Results from Inference over Regimes** – Inferred cost functions and priors with posterior predictive checks and inferred posterior parameters.
 Color legend – True values: red; Inferred means: black; 95%-HDI: grey; targets: dashed grey lines; actions: orange. Black dots at the x-axes denote posterior means.

The good quality of this results is also evident in the promising MCMC diagnostics. The values of \hat{R} are all equal to 1.00 and the ESS are as high as for no other regime. The parameters α , β and μ_{prior} even exceed 10,000 effective samples. Only σ_p and σ_{prior} cannot accumulate a thousand effective samples, while still being inferred reliably. Moreover, the ESS for σ_p is much higher than for σ_p and σ_{prior} with a value of about 4,500. This, again, is a common pattern in our results. Usually, σ_a is easier to infer than the other variabilities because it can be inferred from the observed actions directly and is not hidden deeper in the model like it is the case for σ_p and σ_{prior} . This also shows in the MCMC diagnostics.

The results of the *Undershoot via Cost* regime are also very good. Both, the cost function as well as the prior distribution are inferred quite well. The same reasoning why the cost function was hard to infer in the *Undershoot via Prior* regime applies here in the positive direction; because the cost function primarily shapes the resulting behavior here, it is much easier to infer the respective cost function parameters.

In line with that, the inferred mean of the perceptual uncertainty σ_p is exactly at the true value and the inferred means for the parameters β , σ_a and σ_{prior} do not deviate more than 0.02 from the respective true values. Only α and μ_{prior} are not inferred that well.

The inference of α is off with an error of 0.33 but, again, does not seem to have too much influence on the resulting behavior and cost function.

The inference of μ_{prior} is still good with an error of 0.26. Potentially even the small inference errors in the other parameters can lead to a stronger deviation in another parameter – in this case μ_{prior} – in order to still capture the observed behavior well.

The MCMC diagnostics are also good. The values of \hat{R} are either close to one or near enough to be acceptable. The effective samples sizes are still not very large, but are also acceptable for the parameters α , β and μ_{prior} . They are only in the range of a few hundreds for parameters σ_a , σ_p and σ_{prior} , which usually is not enough since the ESS should be at least at 2,000 and which also shows in increased values of \hat{R} . However, this does not seem to disturb the inference of the concerned parameters. Again, this is a common pattern in our results: Some parameters, mostly the variabilities σ exhibit very low ESS values.

For the *Regression to the Mean* regime, the inferred cost function is estimated well and while the inferred prior distribution is a little narrower than the original distribution, the overall behavior is also captured quite well here.

We see, that the inferred cost function captures the behavior of the original cost function well without inferring the parameter α very well, since its inferred value shows an error of 0.51. Again, this speaks for an overall low influence of α in the current parameterization of the model. Parameters β and σ_a are inferred with errors not bigger than 0.04. The inference of μ_{prior} is also quite accurate with an error of only 0.11, however σ_{prior} has an error of 0.06. The inferred value for σ_p is also far off with an error of 0.16, since the variabilities usually always take on values $\sigma \leq 0.3$.

The MCMC diagnostics also look good with $\hat{R} \leq 1.01$ for all parameters and the effective samples sizes are again really low for the variabilities σ_p and σ_{prior} . While the results of the inference in this regime are not the prime example of good inference, the results are still not bad. However, it is a little bit concerning that the effective sample size for μ_{prior} is only at about 760 and that the error for σ_{prior} is rather high. Especially in a regime that relies a lot on the prior for generating the observed behavior, the prior parameters should be easy to recover, which does not seem to be the case here.

It is noteworthy, that the 95%-HDI of the inferred cost function in Figure 4.2 becomes almost infinitesimal at the point where $a = t - 1$. This is most likely due to the nature of the cost function. As one can see in equation (3.2), the exponentiated term in the first term of the cost function becomes $a - t = 1$. Since 1 does not change to the power of anything, there is no variability there except from the effort term $(1 - \beta)a$ which

Regime Name	Parameter	True Value	95% HDI	Accuracy		MCMC Statistics	
				Mean	Error	\hat{R}	ESS
Hitting Target	α	2.00	[0.95, 3.69]	2.28	0.28	1.16	24.84
	β	0.90	[0.67, 1.00]	0.89	0.01	1.22	19.48
	σ_a	0.10	[0.84, 4.80]	0.17	0.07	1.28	16.49
	σ_p	0.20	[0.02, 0.64]	0.07	0.13	1.65	9.74
	μ_{prior}	3.00	[0.08, 0.22]	3.91	0.91	1.14	27.94
	σ_{prior}	2.00	[0.00, 0.19]	0.27	1.73	1.63	9.77
Undershoot via Prior	α	2.00	[0.16, 2.81]	1.52	0.48	1.00	15,227.92
	β	0.95	[0.53, 0.96]	0.75	0.20	1.00	14,287.70
	σ_a	0.10	[0.09, 0.12]	0.10	0.00	1.00	4,458.11
	σ_p	0.10	[0.04, 0.24]	0.12	0.02	1.00	949.23
	μ_{prior}	1.50	[1.42, 1.94]	1.63	0.13	1.00	11,262.45
	σ_{prior}	0.05	[0.02, 0.12]	0.06	0.01	1.00	970.19
Undershoot via Cost	α	2.00	[1.40, 3.33]	2.33	0.33	1.00	1,984.55
	β	0.60	[0.48, 0.79]	0.62	0.02	1.00	2,190.42
	σ_a	0.10	[0.02, 0.14]	0.08	0.02	1.03	370.97
	σ_p	0.10	[0.03, 0.14]	0.10	0.00	1.02	432.58
	μ_{prior}	2.50	[0.59, 3.94]	2.24	0.26	1.00	3,652.74
	σ_{prior}	0.30	[0.12, 0.43]	0.28	0.02	1.01	950.10
Regression to the Mean	α	2.00	[0.18, 2.75]	1.49	0.51	1.01	1,198.10
	β	0.80	[0.55, 0.97]	0.76	0.04	1.00	7,720.34
	σ_a	0.10	[0.10, 0.12]	0.11	0.01	1.00	3,997.34
	σ_p	0.30	[0.03, 0.25]	0.14	0.16	1.01	779.48
	μ_{prior}	3.00	[2.80, 3.54]	3.11	0.11	1.01	760.48
	σ_{prior}	0.12	[0.01, 0.11]	0.06	0.06	1.01	800.52

Table 4.2.: **MCMC Results and Diagnostics for Synthetic Regimes** – Results and MCMC diagnostics for each parameter in every regime. The error shown is the absolute error.

only has little influence there. Since there is no variability in the term, there is also no uncertainty in the estimation of the cost function there, since it will always be 1 no matter the value of α .

Even if some parameters cannot be reliably estimated, this evaluation showed us that the concept of our method works and the posterior predictive checks show us that the observed data can be reliably captured by the estimated parameters. Of course, it is always the goal to recover the exact parameters to be able to attribute every component its respective role, but it shows us that there might be identifiability problems in the estimation, since different sets of parameters can produce the same behavior. For more on that, please consult Sections 2.4.1 and 5.1.1.

4.1.3. Inference of Sensory Information

Until now and also for the most part of the following results we focus on the inference of the parameters. However, we also infer the values of the subject’s measurements z_i of the true targets t_i . This makes the inference a really hard problem, since the parameter space the inference must be performed over gets very high-dimensional through that. We will take a closer look on this part of the inference in this section.

Again, we are using the four regimes introduced in the beginning of Section 4.1. Each of these data sets consists of 8 targets with 25 trials each, thus the subject creates 200 measurements per data set in total. The results of the inference are shown in Figure 4.3. Despite its catastrophic MCMC diagnostics and its improvable

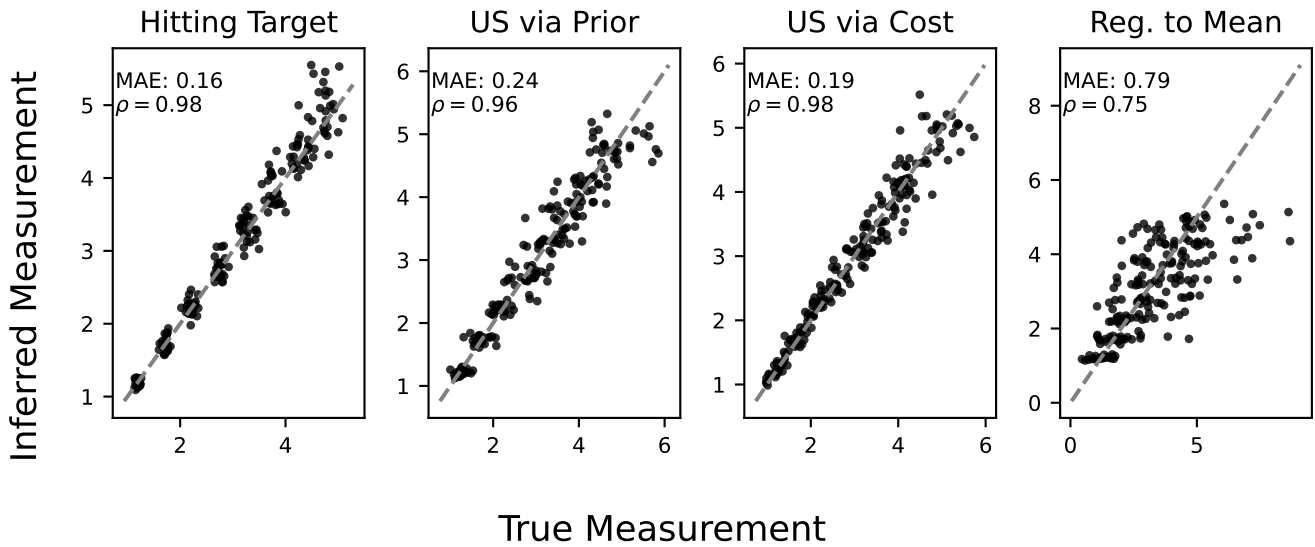


Figure 4.3.: **Inference of Subject’s Measurements z** – Means of the inferred posterior distributions of the subject’s measurements z_i over the ground truth for the following data sets: *Hitting Target*, *Undershoot via Cost*, *Undershoot via Prior* and *Regression to the Mean*. The grey dashed lines indicate optimal inference of each measurement.

inference results, the *Hitting Target* regime has the lowest MAE of 0.16. However, the MAE neglects the spread of the data points around the optimum (dashed grey line in Figure 4.3). A much better measure for evaluating this is the *Pearson correlation coefficient* ρ (Cohen et al., 2009) which measures the correlation between two variables – here, the ground truth and the inferred values – and does not only average over the inference errors. Whenever we refer to the *correlation* of two variables in the following, we mean the Pearson correlation coefficient.

Investigating the correlation of the inferred values of each regime with the ground truth, the *Hitting Target* regime and the *Undershoot via Cost* regime perform best with correlations of $\rho = 0.98$ each. Close in performance is also the *Undershoot via Prior* regime with a correlation of $\rho = 0.96$. The *Regression to the Mean* regime stands out here with a visibly bad correlation of only $\rho = 0.75$ between the inferred values of the latent measurements z_i and their ground truths.

We explain this difference primarily with the difference influence of the prior on the resulting behavior. Since the *Regression to the Mean* regime is the regime where the prior has the most influence on the resulting behavior and sensory information is almost neglected, the measurements are also not really considered in the decision-making and are thus also difficult to infer. This trend is also observable in the other regimes, only on a smaller scale: the higher the ratio between $\sigma_p/\sigma_{\text{prior}}$ (in other words, the smaller the influence of the sensory information), the higher the MAE of the measurement inference.

4.1.4. Recovery of Parameters from Randomly Generated Data Sets

To get a measure of how well our methods performs on different sets, that is, different configurations of parameters, we generated 100 data sets by randomly sampling parameter values and generating actions using the numerical solver introduced in Section 3.2.2. Since we want to evaluate our approximate inference method with the neural network here, we needed to use the numerical solver and not the neural network for the data generation to obtain a more accurate solution. The data set simulates an experimental setup with eight discrete target distances equidistantly in $[1.2, 4.8]$ with 25 trials each, since we judge this to be a realistic experimental setup for any experiment conducted in reality involving a perception-and-action task. The parameter values are drawn randomly from the probability distributions shown in Table 3.2. We removed three inference runs where the parameter estimates were more than ten times bigger than the real value, since we judge these runs to have run into numerical problems during the MCMC sampling.

The posterior means of our inference method and the ground truth values are shown in Figure 4.4. The parameter with the visibly best correlation between inferred value and true values is σ_a ($\rho = 0.942$). This is followed by μ_{prior} ($\rho = 0.747$), σ_p ($\rho = 0.635$) and σ_{prior} ($\rho = 0.566$). The correlation gets worse for the parameter β , which only has a correlation of $\rho = 0.403$. The parameter α shows almost no correlation with the ground truth ($\rho = 0.145$), which is line with our previous findings that α is generally hard to infer, most likely because it has not that much impact on the resulting behavior – which also makes it less severe if we cannot infer it reliably.

However, this ranking in performance changes a little when we consider the MAE instead of the correlation. Please keep in mind that the correlation also accounts for the variability while the MAE can average it out – we also provide the MAE for completeness here but do not rely on it too much out of said reasons. We also computed the relative errors of the inferred posterior means to the true values per parameter and run. The lowest inference errors arise in the inference of the cost parameter β (Mean = 0.13, SD = 0.10), the prior parameter μ_{prior} (Mean = 0.33, SD = 0.57) and the motor variability σ_a (Mean = 0.41, SD = 0.70). For the motor variability and μ_{prior} , a lot of probability mass lies between the median and the 75%-quartile though, thus these are more likely to be overestimated. The relative inference error is higher for α (Mean = 0.43, SD = 0.53), the perceptual uncertainty σ_p (Mean = 0.68, SD = 1.02) and the prior uncertainty σ_{prior} (Mean = 0.74, SD = 1.03). The upper bound of the observed errors is almost always farther away from the median than the lower bound because of the nature of the relative error. Especially when considering the variabilities, we observe a lot of very small values which can easily lead to high relative errors if an inferred estimate is a little off.

As another quantitative evaluation, we calculated the 95%-HDI for each inferred posterior and counted in how many cases the 95%-HDI contained the true value of the respective parameter on the right of Figure 4.4. We also calculated the standard error of the mean (SE) for the number of times a HDI covered the true value of a parameter. Based on this evaluation, the inference of the cost function parameters α (89.88% in HDI, SE = 0.0321) and β (89.88% in HDI, SE = 0.0321) works quite well with inferred values within the HDI in about 90% of the cases. For μ_{prior} (86.52% in HDI, SE = 0.0364), this is a little less, and the variabilities σ_{prior} (80.90% in HDI, SE = 0.0419), σ_a (83.15% in HDI, SE = 0.0399) and σ_p (82.02% in HDI, SE = 0.0409) are inferred correctly in about 80% of the cases.

Please note that it is very difficult for the MCMC sampler to infer values that have only little influence on the observed actions. For example, in a parameter configuration that has a high ratio between σ_p and σ_{prior} and thus leads to the prior having no influence on the subject's behavior, it is often hard to infer the prior's parameters. This also makes inferring σ_{prior} and σ_p really difficult, since there potentially are identifiability problems between the prior parameters μ_{prior} and σ_{prior} (cf. Section 5.1.1). Moreover, these synthetic data sets were sampled randomly and besides the tuning of the sampling distributions do not have constraints to really

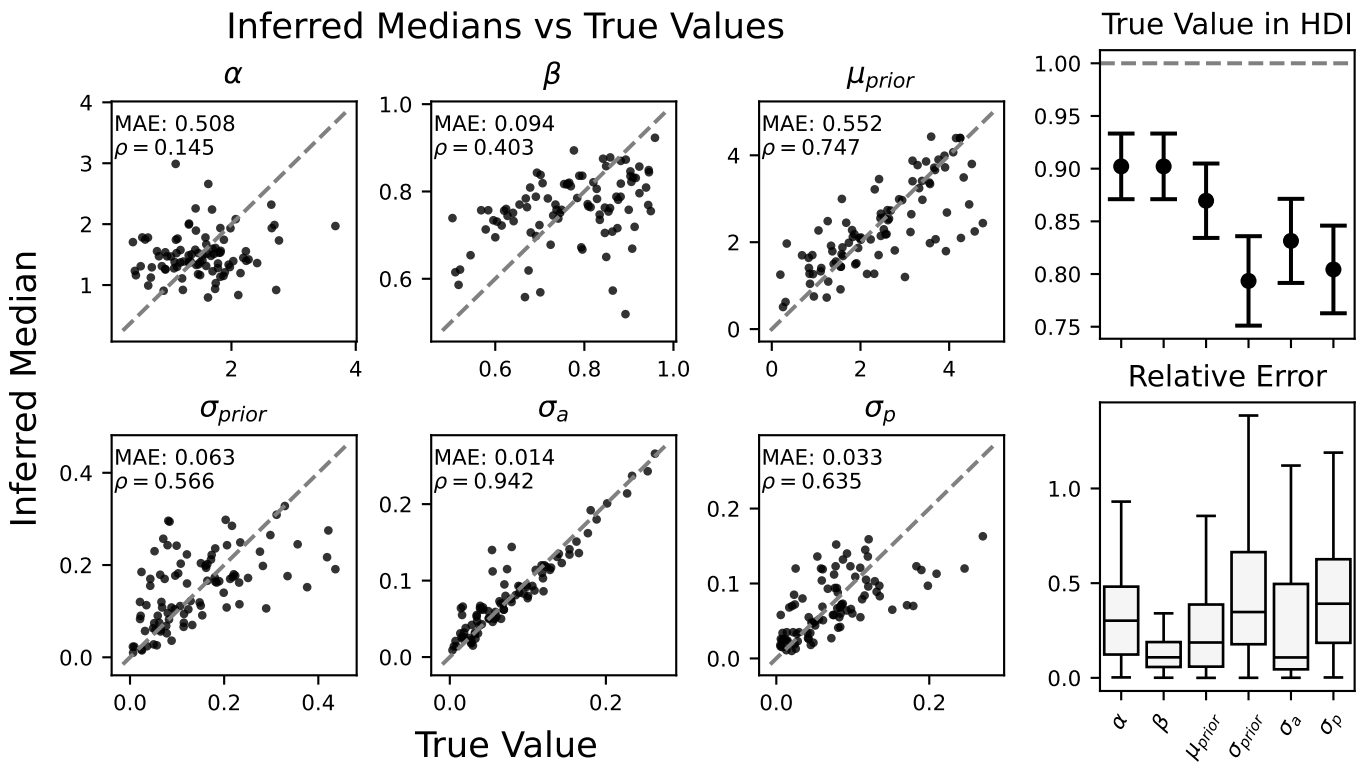


Figure 4.4.: **Inferred Posterior Means vs. Ground Truth in Random Synthetic Data** – Means of the inferred posteriors versus the ground truth value per model parameter. With optimal and exact inference, all samples would fall on the grey dashed diagonal line denoting equality between the ground truth and the inferred value. On the top right, the amount of samples in the 95%-HDI is shown per parameter with the error bars denoting the standard error of the mean. The grey dashed line denotes the best case scenario where the HDI always contains the true value. On the bottom right, the relative absolute error as difference between ground truth and inferred mean is shown for each parameter. Boxes show interquartile ranges with the medians denoted as a solid black line. The whiskers of each box extend to the full range of each error distribution.

represent realistic data. Most likely, some parameter configurations do not generate realistic behavior, but we still relied on this evaluation method in order to cover the majority of possible parameter configurations, regardless of how likely they are to describe real behavioral data.

MCMC Statistics

For the most part, the diagnostics of our MCMC runs were good. That is, we have a good value of \hat{R} (a measure for the convergence of chains) which should be close to 1.0 and sufficiently large effective sample sizes (ESS). The statistics of the runs on the synthetically generated data set used in Section 4.1.4 are shown in Figure 4.5. The median of \hat{R} for all parameters is at about 1.01 or 1.02, which indicates that the values of \hat{R} are mostly close to 1.0. If one observes $\hat{R} > 1.05$, this is a sign that the chains of the respective run did not converge and the posterior estimates are more likely to not be close to the ground truth. Since we have maximum \hat{R} values that are even bigger than 2.0, not all of our MCMC runs converged. This makes sense regarding the results from Section 4.1.4, where the HDI did not always cover the true values – most likely, the runs where the HDI fails to contain the ground truth are also the ones where the chains did not converge.




Figure 4.5.: **MCMC Statistics** – Statistics of the MCMC runs on the synthetic data set from Section 4.1.4. The table on the left shows the statistics of \hat{R} over all 100 data sets per parameter, also depicted in the center figure. The figure on the right shows the median values of the effective sample sizes (ESS) over all runs per parameter.

Moreover, the effective samples sizes differ a lot for some parameters. The ESS is a measure for the independent information gain per MCMC sample (Kruschke, 2015) and should be as high as possible. We only have a median ESS of 1,000 to 1,500 in the bulk of the distribution and of about 1,500 to 2,500 in the tail of the distribution for parameters α , β , μ_{prior} . This is not a lot, especially considering that our MCMC sampling ran for 20,000 sampling steps. Still, it is sufficient to reliably obtain estimates close to the true values (Bürkner, 2017). In line with the the slightly worse results for the variabilities σ_{prior} , σ_a and σ_p in Section 4.1.4, the effective samples sizes are especially low for these parameters. This is really concerning and we should aim to improve these statistics to improve the overall performance and ultimately better the accuracy of our method. Anyhow, we did this evaluation only on the randomly generated data set and the results may be way better on a data set that is more carefully tuned towards representing realistic parameter values instead of random parameter configurations.

4.1.5. Influence of the Number of Trials on the Results

The experimental setup substantially influences how well our inference method performs, since the setup determines how many observations of real target-action-pairs and how they are spatially distributed. If our setup only provides us with few observations, our method will perform worse or at least provide broader posterior distributions and thus less confident estimates, and the estimation accuracy will increase with a higher number of samples. On the other hand, since we introduce a latent variable representing a percept for each target, the dimensionality of the inference problem increases linearly with the number of observations in our data set. Thus, the inference problem gets harder with larger observation data sets which might impact performance negatively. The inference results for data sets with eight equidistant targets in $[1.2, 4.8]$ with $[10, 25, 50, 75, 100]$ trials each are shown in figure 4.6. For each number of trials we created 30 data sets. We expected an increase in accuracy for an increasing amount of trials per target, since this also increases the number of observations our inference method can use for its estimations. Usually, more data gives rise

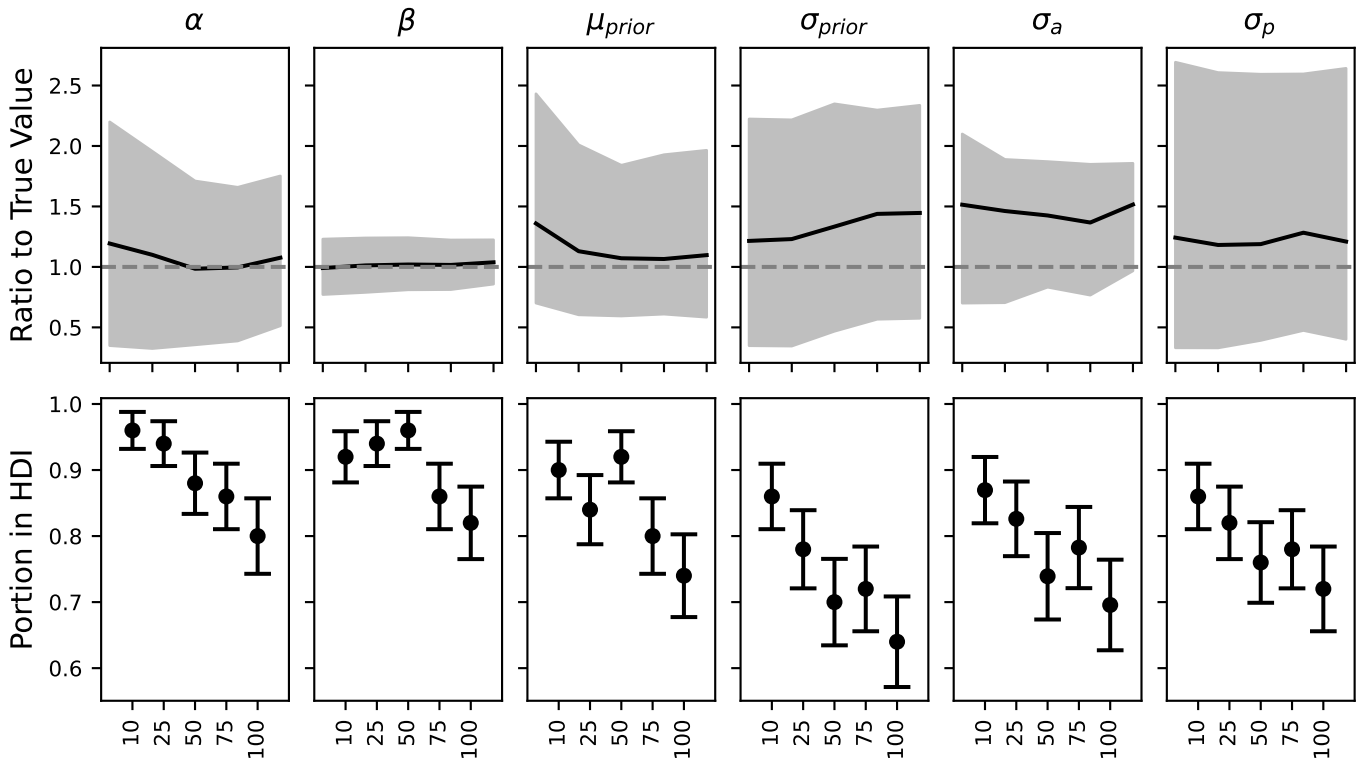


Figure 4.6.: **Influence of Number of Trials on Accuracy** – In the top row, the ratios of posterior mean to true value are shown for each parameter over a varying amount of trials per target in the observed data (solid black line). The grey shading denotes the 95%-HDI for the estimation. With optimal and exact inference, all values would be equal to 1.0 (grey dashed line). In the bottom row, the portion of runs where the 95%-HDI covered the ground truth value is shown. Error bars indicate the standard error. All inferences are based on 30 data sets per number of trials.

to higher accuracy. This is not the case here. The posterior HDIs do not shrink with an increasing amount of trials and the variabilities σ are constantly overestimated, regardless of the number of trials. Only the parameter β is reliably estimated. This trend is even more evident in the portion of inferred HDIs that cover the true value of the respective parameter. With an increasing number of trials, less HDIs include the ground truth parameter value.

We have no explanation for the systematic overestimation of the variabilities σ , but it could arise due to numerical issues in the sampling process, since the values of the uncertainties are naturally rather small in the modeling of human behavior. We explain the decrease in performance with the increasing hardness of the inference problem with more observations. Although the accuracy of most inference methods increases with more observations, the inference problem in our framework becomes higher-dimensional with more observations. Since we treat the measurement z_i of each target t_i as own latent variable that needs to be inferred and is not observable, each observation adds a dimension to the parameter space we need to perform inference on. With the six cost function parameters (α, β), prior parameters ($\mu_{\text{prior}}, \sigma_{\text{prior}}$) and uncertainties (σ_a, σ_p) and an experimental setup consisting of 8 targets with 10 trials each, the inference space is 86-dimensional. For 100 trials per target, this increases to a 806-dimensional inference space. This is an issue that we need to address in future work to improve our method such that it scales well with observed data.

4.2. Experimental Data

Since a benchmark on synthetic data can only serve as a proof of concept, we still need to validate whether our method can successfully recover parameters from empirically gathered data and if it enables us to qualitatively interpret and identify different patterns. To that end, we chose two data sets where the characteristic patterns described in Section 2.1.5 can be observed. That is, a data set of a bean bag throwing task (Willey and Liu, 2018) and a data set of a puck sliding task (Neupärtl et al., 2020).

4.2.1. Bean Bag (BB) Data Set

In this section, we apply our method to a data set from a bean bag throwing task by Willey and Liu (2018). 20 participants were asked to throw a bean bag at five different target distance from 3 to 11 feet with 2 feet increments (.9144 to 3.3528 meters). The participants solved the task in blocks of twelve trials each with random ordering of the five blocks. The subjects did not receive any feedback, neither visually nor verbally.

The inferred parameters are shown here for four selected subjects in Table 4.3 and Figure 4.7 and in the appendix in Figures A.8, A.9 for all subjects.

We selected the subjects which we will review here in more detail by qualitative differences in exhibited behavior and the inferred parameters. Subjects 12 and 02 show undershoot behavior but with different cost functions, subject 06 shows a regression-to-the-mean pattern and subject 17 hits the targets quite well over the full stimulus range.

Subject 12 shows systematic undershoots with a small increase in the difference between the inferred median of their actions and the real target values, which could be explained by signal-dependent noise – because the variability of their actions increases with increasing stimulus intensities, they are more safe in avoiding high costs for overshoots by rather aiming further below the real stimulus value.

The cost function inferred for subject 12 is clearly asymmetric, which comes about by the low value of β (Mean = 0.59, $SD = 0.10$) which puts a lot of weight on penalizing the effort necessary for executing the action. The true value of β might be even smaller than the inferred mean, if it coincides with the peak in the inferred posterior for β . The overall form of the cost function is slightly super-quadratic given the value of α (Mean = 2.40, $SD = 0.62$).

The prior distribution further amplifies the undershoots by putting a lot of probability mass on low stimulus values with $\mu_{\text{prior}} = 1.47$ ($SD = 0.72$) and $\sigma_{\text{prior}} = 0.35$ ($SD = 0.10$). However, the prior should not have too much influence since the uncertainty in the prior is about double the value of the perceptual uncertainty σ_p (Mean = 0.17, $SD = 0.05$). The spread of the subject's actions per target is rather small in comparison to the other subjects and this also shows in a small motor variability σ_a (Mean = 0.12, $SD = 0.06$).

While subject 02 also exhibits undershoot behavior, their cost function looks a bit different. The observed behavior is about the same as for subject 12 but with a little stronger undershoots and increased spread of actions, which directly shows in an increased value for the motor variability σ_a (Mean = 0.22, $SD = 0.04$).

In line with the undershoots, the inferred cost function is also asymmetric, but it is far less quadratic. The value of α is nearly 1.0 (Mean = 1.16, $SD = 0.67$), thereby the cost almost scales with the absolute error directly. The inferred posterior even shows a skew towards small values $\alpha < 1$, thereby indicating that it might also show behavior of a root function in reality.

The undershoots are again further supported by a prior with a lot of probability mass at low values. While the

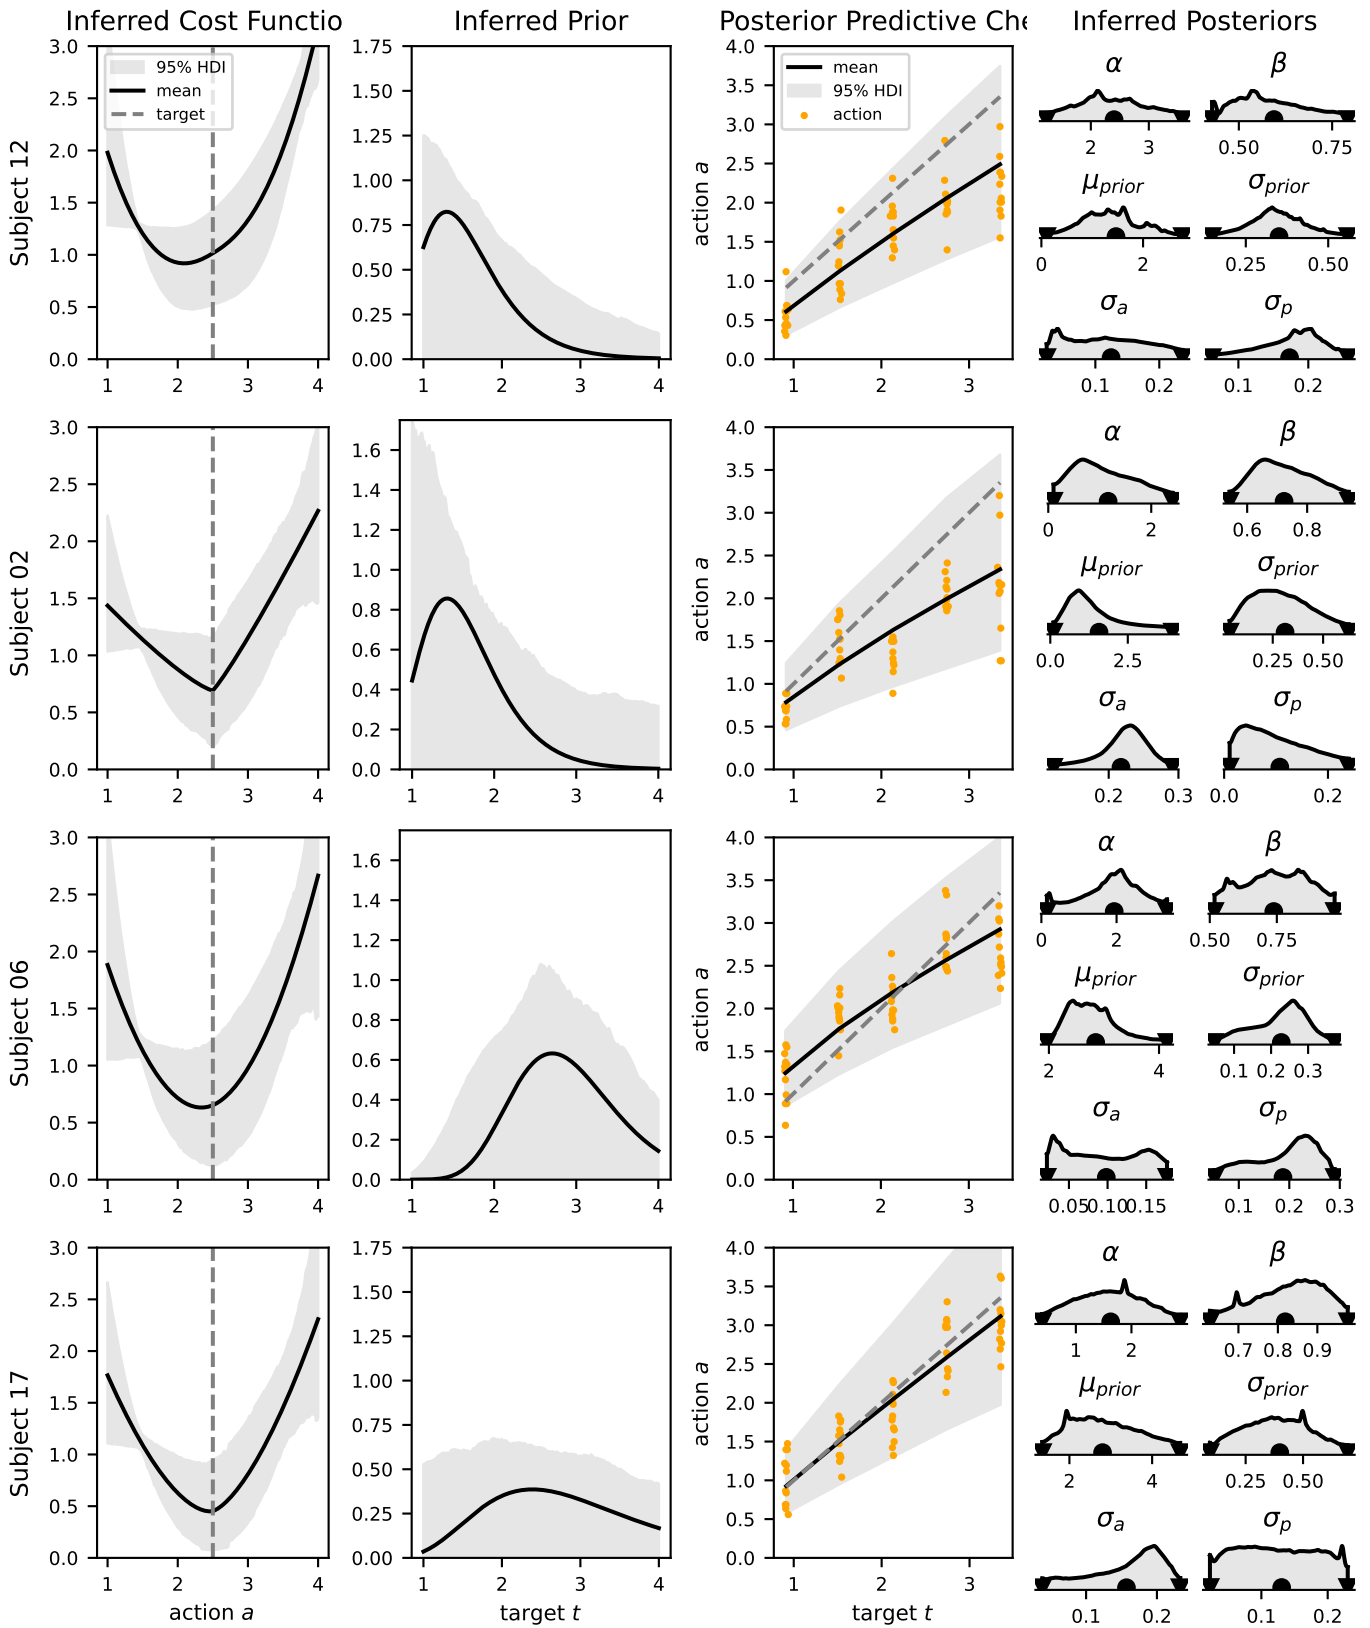


Figure 4.7.: **Exemplary Results from BB Data Set** – Inferred cost functions and priors with posterior predictive checks and inferred posterior parameters on data set BB. Color legend – True values: **red**; Inferred means: **black**; 95%-HDI: **grey**; targets: **dashed grey lines**; actions: **orange**. **Black dots** at the x-axes denote posterior means. Inferred values listed in Table 4.3

log-normal distribution is shifted a little towards the right this time with $\mu_{\text{prior}} = 1.57$ ($SD = 1.08$), this is cancelled out due to a bigger ratio between the perceptual uncertainty σ_p (Mean = 0.11, $SD = 0.07$) and the prior uncertainty σ_{prior} (Mean = 0.31, $SD = 0.17$).

Subject 06 qualitatively differs in its behavior by showing regression-to-the-mean patterns, which usually comes about by stronger reliance on the prior and less reliance on the sensory information. This also shows in the inferred data, since the perceptual uncertainty σ_p (Mean = 0.19, $SD = 0.07$) is highest for subject 06 while their prior uncertainty σ_{prior} (Mean = 0.23, $SD = 0.08$) is the lowest of all subjects.

The inferred value of μ_{prior} (Mean = 2.85, $SD = 0.54$) puts most of its probability mass into $[2.0, 4.0]$ and makes the prior belief not match the stimulus distribution, which covers about the range $[0.9, 3.4]$. This might be due to different reason. First of all, the subject could have really learnt a mismatched representation of the stimulus distribution, but we think that this is most unlikely. It is rather a common phenomenon in our method, that a strong effect of the prior is rather brought about by high values of μ_{prior} instead of a higher ratio $\sigma_{\text{prior}}/\sigma_p$, which would capture the stronger reliance on the prior more appropriately. Moreover, it could be due to technical reasons, since we did not adapt the stimulus range for our MCMC priors (cf. Table 3.2) and thus the inference ran with the prior distributions for μ_{prior} conditioned on the stimulus range of $[1.2, 4.8]$ from our synthetic data. It is necessary to repeat this experiment with an adapted prior for μ_{prior} and see how this changes the inferred value for μ_{prior} .

The cost function inferred for subject 06 shows no specialties, it is slightly asymmetrical because of the value of β (Mean = 0.74, $SD = 0.13$) and its overall form is quasi-quadratic with $\alpha = 1.94$ ($SD = 0.82$).

The variability in actions is lowest for this subject and it also shows in the motor variability σ_a (Mean = 0.10, $SD = 0.05$).

The last subject with ID 17 was able to hit the targets quite well on average. Usually, this comes about with low reliance on the prior, small perceptual uncertainty and low emphasis on effort costs.

We can observe exactly this pattern in the values our method inferred for this subject. The prior is the widest of all inferred priors with a variability of $\sigma_{\text{prior}} = 0.40$ ($SD = 0.16$). The value of μ_{prior} (Mean = 2.80, $SD = 0.90$) is rather high but of no further importance since it has no real influence on the resulting behavior anyway. Because of the small value of σ_p (Mean = 0.13, $SD = 0.06$) in contrast to the high value of σ_{prior} , this subject relies a lot on its sensory information instead of the prior.

The cost function also resembles the low effort cost that is necessary to obtain good accuracy in the task with $\beta = 0.82$ ($SD = 0.10$). This puts a lot of weight on the accuracy term of the cost function and less on the effort term. The overall form of the cost function is between linear and quadratic due to the value of parameter α (Mean = 1.63, $SD = 0.66$).

The motor variability for this subject lies in the middle of the field (Mean = 0.16, $SD = 0.05$).

In all inferred data sets showing undershoots (subjects 02 and 12) or a well-tuned response to the stimuli (subject 17), the resulting behavior was explained with the cost function parameters, for example the undershoots of subject 12 with $\beta = 0.59$ or the hits of subject 17 with $\beta = 0.82$. Strong reliance on the prior led to regression to the mean behavior in subject 06 with the lowest uncertainty $\sigma_{\text{prior}} = 0.23$ in the prior of all selected subjects.

Also, the prior was much lower for subjects 12 and 02 ($\mu_{\text{prior}} \in \{1.47, 1.57\}$) in contrast to subjects 06 and 17 ($\mu_{\text{prior}} \in \{2.85, 2.80\}$). This could be a problem in some cases though because it might not always clearly be identifiable whether undershoots come about by low effort costs or high reliance on the prior. For example, the value of β of subject 02 is not really low (as it usually is the case for undershoots via effort cost) and the prior also has a lot of probability mass at low values. While it can surely be that the observed behavior comes about by a mixture of effort cost and reliance on prior beliefs, we need to make sure that more extreme values can also be reliably extracted, which would make the interpretation of the results also more easy and clear.

Subject ID	α	β	σ_a	σ_p	μ_{prior}	σ_{prior}
12 (Undershoot)	2.40(± 0.62)	0.59(± 0.10)	0.12(± 0.06)	0.17(± 0.05)	1.47(± 0.72)	0.35(± 0.10)
02 (Undershoot)	1.16(± 0.67)	0.72(± 0.11)	0.22(± 0.04)	0.11(± 0.07)	1.57(± 1.08)	0.31(± 0.17)
06 (Reg. to Mean)	1.94(± 0.82)	0.74(± 0.13)	0.10(± 0.05)	0.19(± 0.07)	2.85(± 0.54)	0.23(± 0.08)
17 (Hit)	1.63(± 0.66)	0.82(± 0.10)	0.16(± 0.05)	0.13(± 0.06)	2.80(± 0.90)	0.40(± 0.16)

Table 4.3.: **Exemplary Results from BB Data Set** – Results of our methods applied to the BB data for selected subjects that show different results and behaviors. Each estimated parameter’s posterior mean is shown with the respective standard deviation of the inferred posterior in brackets. Subject 12 and 02 exhibit undershoots while subjects 06 and 17 show regression to the mean behavior and hit the targets quite well, respectively. Results are visualized in Figure 4.7.

The ranges of inferred values over all 20 subjects are shown in Table 4.4. We inferred values of α in $[1.04, 2.70]$, thereby not finding any root functions as cost function which all behave at least linearly. The inferred value range for β is $[0.51, 0.86]$, thereby spanning a big range of how strongly effort costs are considered and avoiding value $\beta < 0.5$ for which the model usually does not have a solution.

For the prior distributions, we inferred values of σ_{prior} in $[0.15, 0.59]$, which seems reasonable. The values of μ_{prior} lie within $[1.02, 4.33]$, where the upper bound seems a little bit high considering that the stimuli range is about $[0.91, 3.35]$. This could be due to the previously mentioned un-adjusted prior distribution for μ_{prior} , which is still set on stimuli in the range of $[1.2, 4.8]$ as it was the case in the synthetic data.

The inferred values for σ_a are in $[0.05, 0.24]$ and for σ_p are in $[0.07, 0.24]$, which also seem reasonable.

Parameter	Inferred Value		
	Min	Median	Max
α	1.038	1.631	2.669
β	0.505	0.741	0.855
μ_{prior}	1.018	2.440	4.334
σ_{prior}	0.153	0.326	0.594
σ_a	0.051	0.134	0.236
σ_p	0.070	0.135	0.237

Table 4.4.: **Range of Inferred Values for Data Set BB** – The minimum and maximum value along with the median value over the range of all inferred values per parameter.

MCMC Diagnostics

It is also worth taking a look at the MCMC diagnostics of the inferences performed on the empirical data of all 20 subjects. The values of \hat{R} and the effective sample sizes are shown in Figure 4.8. It is remarkable that the MCMC diagnostics look much better on real data than on the synthetic data set in Section 4.1.4. Here, we have $\hat{R}_{\text{median}} < 1.025$ in the median for all parameters over all runs. The maximum values of \hat{R} also do not exceed 1.18 at most for σ_a . Given that, we can assume that the chains for all inference processes on the empirical data converged.

However, the effective sample sizes still are really low. The effective sample sizes are highest for parameters α and μ_{prior} with median values of 2,096.0 and 2,660.5 in the bulk, respectively. We are primarily looking

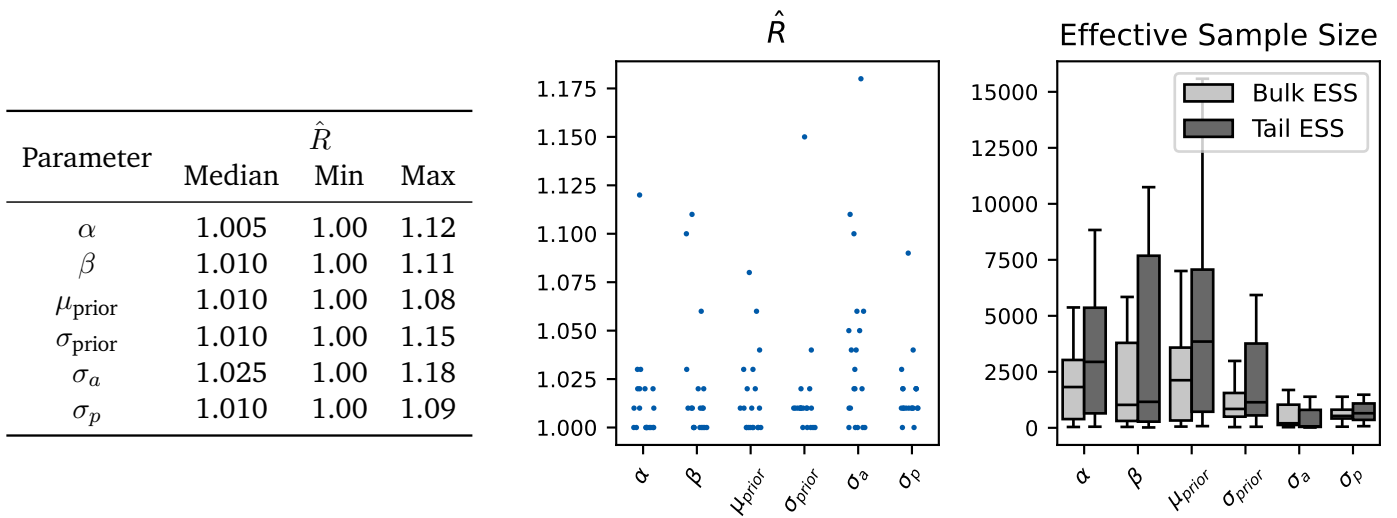


Figure 4.8.: **MCMC Statistics for the BB Data Set** – Statistics of the MCMC runs on the synthetic data set from Section 4.1.4. The table on the left shows the statistics of \hat{R} over all 20 participants in the BB data set, also depicted in the center figure. The figure on the right shows the median values of the effective sample sizes (ESS) over all runs per parameter.

at the bulk ESS in the following since this value is of most importance to us and it is commonly used when investigating the ESS. The values for α and μ_{prior} are still acceptable.

We then observe a much smaller ESS for β and σ_{prior} with values of 1,075.5 and 1,065.0, respectively. This is a little low, but can still be accepted, since at least 1,000 are necessary to obtain stable estimates (Bürkner, 2017).

However, for the perceptual uncertainty and motor variability we drop below 1,000 effective samples, which is definitely too low. The motor variability σ_a has a median ESS of 151.5 and the perceptual uncertainty σ_p shows a median ESS of 564.0. According to Bürkner (2017), these values are too low to perform stable inference over these parameters. Yet, we obtain reliable parameter estimations for most runs. We definitely need to look into this by examining to which extent the ESS is limited by the inference process and by our model itself, in order to increase the effective sample sizes, especially for σ_a and σ_p , in future versions.

4.2.2. Puck (PU) Data Set

In this continuous visuomotor task by Neupärtl et al. (2020), 16 participants were asked to propel a puck towards a target in a virtual environment with simulated friction on the surface the pucks were slid on. The initial force and thus also initial acceleration applied to the puck was determined by how long the subjects held down a button in each trial. Initial positions of the puck and positions of the targets were randomly drawn to virtually cover a distance between 1.0 and 5.0 meters. Here, we only analyze the data of the second of four different phases. The subjects were familiar with the task after finishing the first phase without any visual feedback. In the second phase, they then received visual feedback on where the thrown pucks landed. All subjects solved the task with two different conditions ('D' and 'F') in which they were given pucks with different masses, thus influencing the friction with which the pucks slid over the surface. Each subject solved 200 trials.

The inferred parameters are shown here for selected subjects in Table 4.5 and Figure 4.9 and in the appendix in Figures A.10, A.11 for all subjects.

Again, as for the results for data set BB, we selected four subjects which either differ in their qualitative behavior or the inferred parameter values, to showcase the overall behavior of our method on data set PU. Subjects *KA01NS* and *OK18ER* show undershoot behavior, while subject *KA06IN* shows a regression-to-the-mean pattern in their actions. Subject *KA09NG* also exhibits undershoots, but this only applies to the upper half of the stimulus range. They hit the targets well on average for the lower half of the stimulus range.

Subject *KA01NS* with the 'D'-condition puck exhibits undershoots which most likely come about through the asymmetric cost function given by the inferred values of the parameters α (Mean = 2.33, $SD = 0.82$) and β (Mean = 0.74, $SD = 0.15$), the latter of which is primarily responsible for the undershoots.

The inferred prior distribution lies approximately in the middle of the stimulus range. The inferred parameters are $\mu_{\text{prior}} = 3.43$ ($SD = 1.21$) and $\sigma_{\text{prior}} = 0.18$ ($SD = 0.09$) but the inferred shape of the prior distribution is not really certain, as the 95%-HDI of the inferred prior covers the whole stimulus range approximately uniformly.

The inferred motor variability σ_a (Mean = 0.42, $SD = 0.03$) is rather high in contrast to the values inferred in the BB data set. The perceptual uncertainty σ_p (Mean = 0.04, $SD = 0.05$) is much lower than the prior uncertainty σ_{prior} , which reduces the influence of the prior on the produced behavior.

The subject *OK18ER* with the 'D'-condition puck also produces undershoots but with less variability in its actions and the undershoots are a little stronger than for subject *KA01NS*.

The cost function inferred for subject *OK18ER* is strongly asymmetric and puts a lot of weight on the effort cost via parameter β (Mean = 0.50, $SD = 0.11$). The overall shape of the form is super-quadratic with $\alpha = 2.61$ ($SD = 0.57$).

Besides the high influence of the effort cost, the undershoots are further supported by a prior with a lot of probability mass at the lower end of the stimulus range. The prior of the subject is parameterized by $\mu_{\text{prior}} = 2.10$ ($SD = 0.65$) and $\sigma_{\text{prior}} = 0.44$ ($SD = 0.11$). The influence of the prior is of moderate strength because of its high uncertainty. Although the perceptual uncertainty σ_p is also rather high (Mean = 0.30, $SD = 0.09$), the ratio between σ_{prior} and σ_p is at about 1.5, which usually does not lead to strong influences of the prior. The motor variability σ_a of subject *OK18ER* is the lowest of the four subjects featured here (Mean = 0.22, $SD = 0.09$).

We can also observe regression-to-the-mean behavior in the PU data set, for instance in subject *KA06IN* with the 'F'-condition puck. Regression to the mean usually comes about through strong reliance on the prior, which can also be observed here.

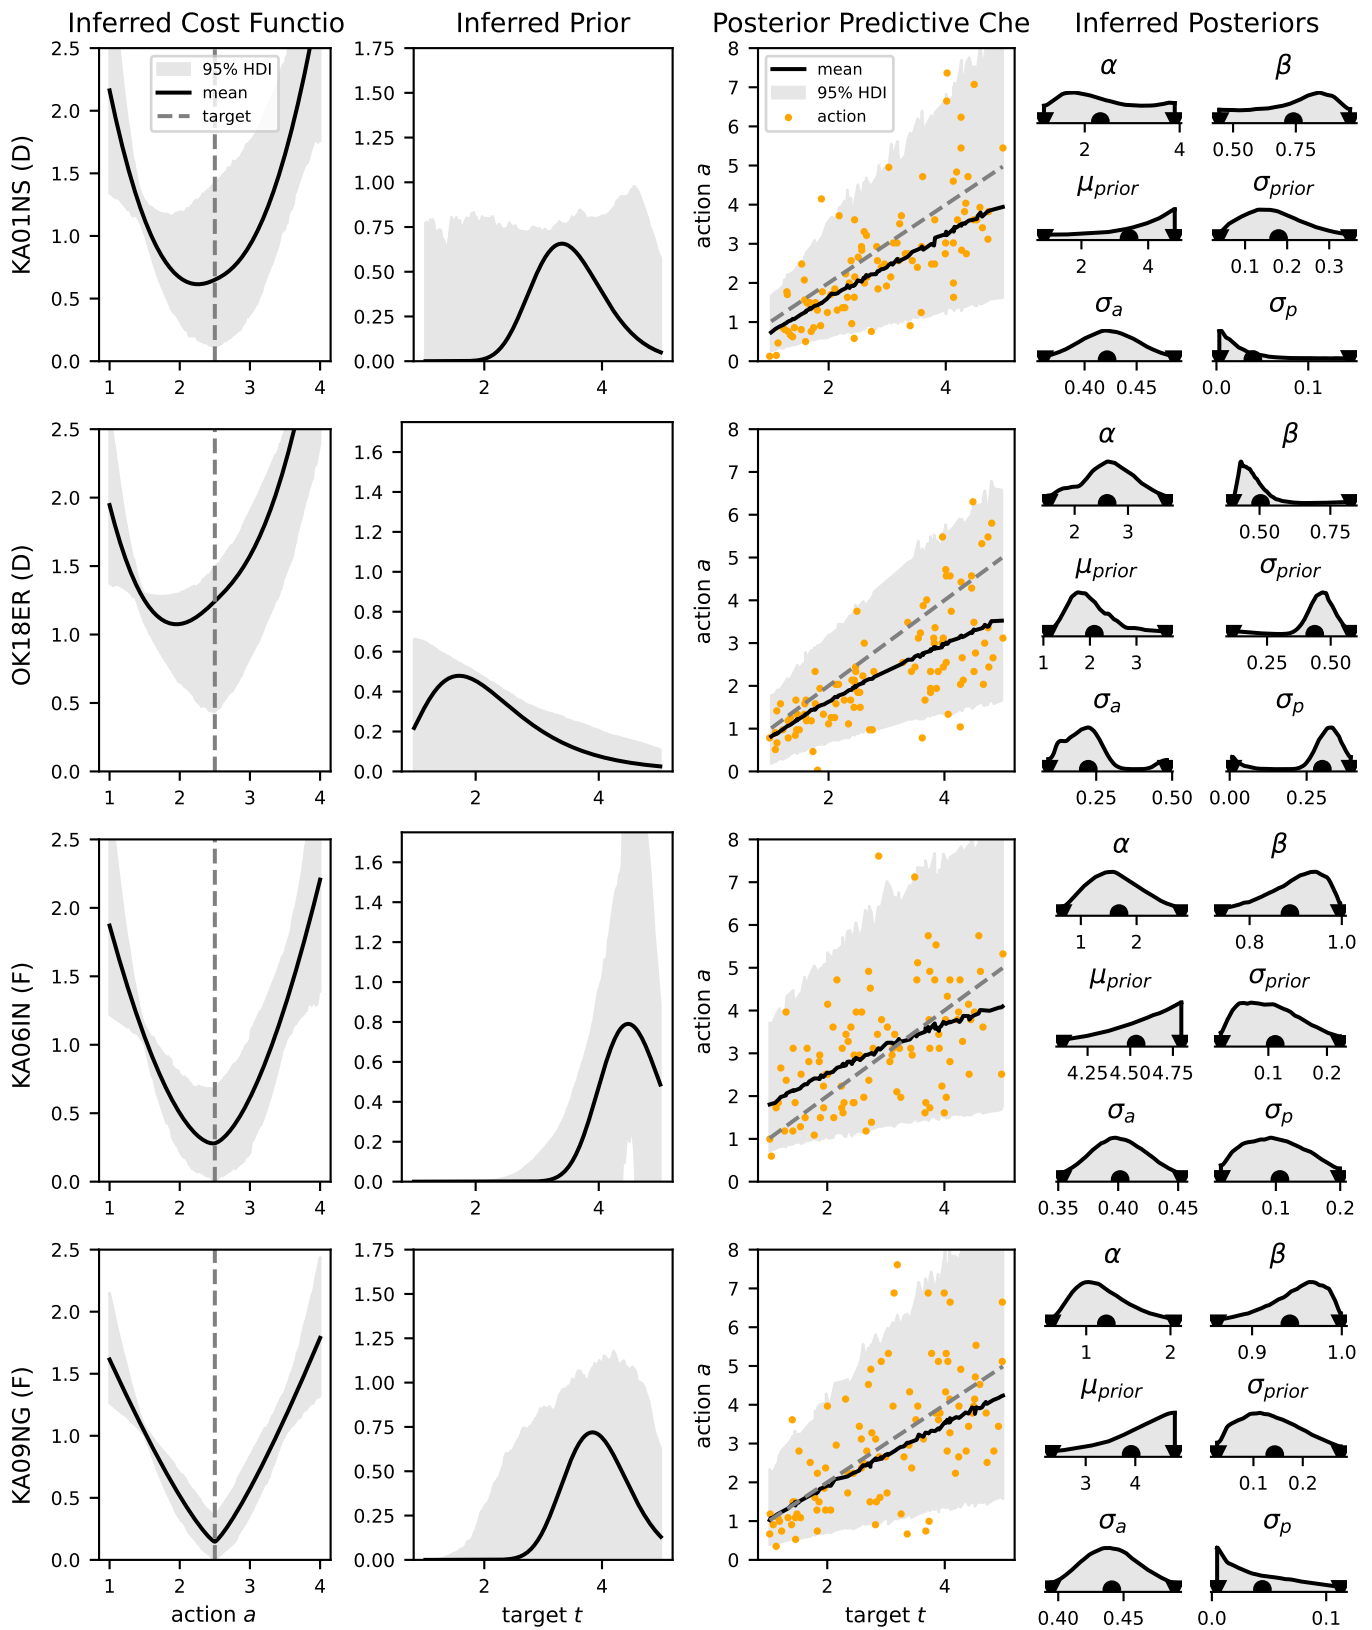


Figure 4.9.: **Exemplary Results from PU Data Set** – Inferred cost functions and priors with posterior predictive checks and inferred posterior parameters on data set PU. Color legend – True values: **red**; Inferred means: **black**; 95%-HDI: **grey**; targets: **dashed grey lines**; actions: **orange**. **Black dots** at the x-axes denote posterior means. Inferred values listed in Table 4.5.

Subject ID	α	β	σ_a	σ_p	μ_{prior}	σ_{prior}
KA01NS (D)	2.33(± 0.82)	0.74(± 0.15)	0.42(± 0.03)	0.04(± 0.05)	3.43(± 1.21)	0.18(± 0.09)
OK18ER (D)	2.61(± 0.57)	0.50(± 0.11)	0.22(± 0.09)	0.30(± 0.09)	2.10(± 0.65)	0.44(± 0.11)
KA06IN (F)	1.68(± 0.57)	0.89(± 0.08)	0.40(± 0.03)	0.11(± 0.05)	4.53(± 0.22)	0.11(± 0.06)
KA09NG (F)	1.24(± 0.41)	0.94(± 0.04)	0.44(± 0.02)	0.04(± 0.04)	3.92(± 0.78)	0.14(± 0.07)

Table 4.5.: **Exemplary Results from PU Data Set** – Results of our methods applied to the PU data for selected subjects that show different results and behaviors. Each estimated parameter’s posterior mean is shown with the respective standard deviation of the inferred posterior in brackets. Subjects *KA01NS* and *OK18ER* show undershoot behavior, while subject *KA06IN* shows a regression-to-the-mean pattern. Subject *KA09NG* also shows undershoots but hits the targets well on average for the lower half of the stimulus range. Results are visualized in Figure 4.9.

The prior distribution of subject *KA06IN* has a really high value for μ_{prior} (Mean = 4.53, $SD = 0.22$) which can strongly increase the value of the produced action if the prior has enough influence. The variability σ_{prior} of the prior is also rather small (Mean = 0.11, $SD = 0.06$) and equal to the perceptual uncertainty σ_p (Mean = 0.11, $SD = 0.05$), thereby exhibiting a strong influence on the final response which leads to the observed regression to the mean.

The cost function is parameterized by $\alpha = 1.68$ ($SD = 0.57$) and $\beta = 0.89$ ($SD = 0.08$), thereby putting strong emphasis on the accuracy term of the cost function. However, this is not of so much importance here since the behavior is primarily shaped by the prior anyway.

The motor variability of the action is, again, rather high with a value of $\sigma_a = 0.40$ ($SD = 0.03$).

Subject *KA09NG* with the ‘F’-condition puck also shows undershoot behavior but with a little different values for the cost function parameters.

The inferred value for $\alpha = 1.24$ ($SD = 0.41$) is one of the lowest observed in the whole data set and the inferred value for $\beta = 0.94$ ($SD = 0.04$) is one of the highest values in the whole data set for the respective parameter. This leads to a cost function which behaves close to linear to the error and is very symmetric since it does not put a lot of weight on the effort term.

The inferred prior is primarily located towards the upper end of the stimulus range because of μ_{prior} (Mean = 3.92, $SD = 0.78$) but only has minor influence on the resulting behavior because of the ratio between σ_{prior} (Mean = 0.14, $SD = 0.07$) and σ_p (Mean = 0.04, $SD = 0.04$).

The inferred motor variability is typically high for this data set with $\sigma_a = 0.44$ ($SD = 0.02$).

Subject *KA09NG* probably hits the targets well at the lower end of the stimulus range because the prior increases the intensity of weak actions more strongly. Yet, since the parameters are not that clear here, it could also come about by the random distribution of actions for smaller stimuli, which could easily be the case because of the higher motor variability in this task overall.

In this data set, the undershoots for subjects *KA01NS* and *OK18ER* could be explained by low values for the parameter $\beta \in \{0.74, 0.50\}$. The regression-to-the-mean behavior by subject *KA06IN* comes about by strong reliance on a high prior with $\mu_{\text{prior}} = 4.53$ and $\sigma_{\text{prior}} = \sigma_p = 0.11$. The accurate performance in low stimulus ranges and undershoots for higher stimuli shown by subject *KA09NG* cannot be explained clearly, but the good performance corresponds to high weight on the parameter $\beta = 0.94$. The overall behavior could also be a little unclear because of the generally high motor variability in this task, with $\sigma_a = 0.44$ for subject *KA09NG*.

The ranges of inferred values over all 16 subjects with both puck conditions are shown in Table 4.6 and we will compare the inferred values here with the bean bag task (BB).

The parameter α was inferred in the range [1.19, 3.51] and thereby has a much higher maximum value than in the BB task while the median ($\alpha^{\text{median}} = 0.861$) is about the same. For parameter $\beta \in [0.50, 0.96]$, there is an extended upper bound as well. However, the median ($\beta^{\text{median}} = 0.861$) of the values of beta is also increased in contrast to the values inferred in the BB task.

The inferred values for the prior distribution's parameter μ_{prior} are much higher than in the BB task with values in [1.54, 4.66]. This is only a logical sequence of the increased stimulus range with higher values in the PU task. The variability of the prior lies in [0.07, 0.52], where the minimum value and the median ($\sigma_{\text{prior}}^{\text{median}} = 0.152$) are only half of the values in the BB task. This would usually suggest that there is a stronger reliance on the prior in this task overall, yet the values for σ_p behave similarly.

The values inferred for σ_p lie in [0.01, 0.30] with the median ($\sigma_p^{\text{median}} = 0.053$) being less than half the value of the median in the BB task. The motor variability is generally higher in the PU task with $\sigma_a \in [0.22, 0.67]$, which can make the inference problem a little harder because it may require more observation to make reliable inferences about the underlying behavior.

Parameter	Inferred Value		
	Min	Median	Max
α	1.194	1.682	3.514
β	0.495	0.861	0.955
μ_{prior}	1.544	3.466	4.656
σ_{prior}	0.067	0.152	0.517
σ_a	0.223	0.358	0.665
σ_p	0.010	0.053	0.301

Table 4.6.: **Range of Inferred Values for Data Set PU** – The minimum and maximum value along with the median value over the range of all inferred values per parameter.

Parameter	Inferred Means	
	<i>Diamond</i> (D)	<i>Five Dots</i> (F)
α	1.884	1.850
β	0.798	0.845
μ_{prior}	2.974	3.867
σ_{prior}	0.197	0.190
σ_a	0.357	0.394
σ_p	0.078	0.086

Table 4.7.: **Inferred Means for different Pucks** – Differences in inferred parameters for the two different puck conditions 'D'=Diamond and 'F'=Five Dots.

As a small example how our method can also be used, we did a brief comparison of the inferred posterior means per parameter for the different puck conditions D and F . The pucks differed in their weight, with the D -pucks being 25% heavier than the F -pucks. We only compare the mean of the inferred means over all subjects in each condition in Table 4.7. While this is neglectful of the variance of each parameter and does not take into consideration, for instance, whether one puck produced data for which the parameter set was generally more difficult to infer, it still shows some differences between the two conditions.

The values for parameter α are about the same in both conditions. For parameter β , the values is slightly higher in the F -condition, which puts less emphasis on the effort in the F -condition. This makes sense, since the F -pucks are lighter than the D -pucks and thereby the subjects need to put less effort into them to hit the same targets. There are no other remarkable differences between the parameters across the two conditions, except for the location of the prior distribution μ_{prior} , which is much higher for the F -pucks. There could either really be a shift in the prior distribution from one condition to the other, or farther sliding pucks because of the lower weight are wrongfully attributed to the subject's prior beliefs by our method. This could have a lot of reasons and would need more thorough investigation for which we do not have the room here, since we only want to showcase an example for analysis with our method here.

MCMC Diagnostics

It is also worth taking a look at the MCMC diagnostics of the inferences performed on the empirical data of all 32 inference runs. The values of \hat{R} and the effective sample sizes are shown in Figure 4.10.

The MCMC diagnostics in the PU data set are worse than in the BB data set. While the values of \hat{R} were almost perfect in the BB data set, we have a lot more outliers in the PU data set. While the median of all values of \hat{R} is below 1.035 for all parameters, there are also some runs with values $\hat{R} > 1.1$, which is a sign of bad convergence. This impression is also confirmed when looking at the traces of the runs. For example Figure A.12 in the appendix shows a bimodal distribution for subject OK18ER in condition F . In this case, some chains converged to one parameter configuration and some chains converged to another. This is one of the identifiability problems introduced in Section 2.4.1 and will be discussed in more detail in Section 5.1.1. Besides bimodal posterior distributions, the chains generally did not converge to exactly the same shapes for various runs in the inference on the PU data set, which probably causes the bad \hat{R} values.




Figure 4.10.: **MCMC Statistics for the PU Data Set** – Statistics of the MCMC runs on the synthetic data set from Section 4.1.4. The table on the left shows the statistics of \hat{R} over all 16 participants and 2 conditions in the PU data set, also depicted in the center figure. The figure on the right shows the median values of the effective sample sizes (ESS) over all runs per parameter.

The effective samples size behaves the same way as in the BB data set and in our synthetically generated data, except for the ESS of σ_a . The effective sample size for σ_a is huge for some runs. We needed to clip the value to the number of maximum samples per inference run, which is $n_{\text{samples}} = 20,000$. The effective sample size can exceed the number of actual samples if the samples drawn during the MCMC inference are anti-correlated. For visual reasons and for comparison with the other ESS values, we clipped the value at 20,000. The only explanation we have for the much better ESS for σ_a is the increased action variability in the PU task, through which the more widespread action observations could provide the sampling process with more useful information.

For the other parameters, we get the same results as every time before when examining the effective sample size. For parameters α , β and μ_{prior} , the ESS is sufficient. However for the variabilities σ_{prior} , σ_a and σ_p , the ESS decreases below 1,000 samples and thus indicates unstable estimates. Since this happened on all data sets, the problem seems to be model specific and invariant to the data we are performing inference on.

5. Discussion

We developed a new framework extending approximate inference to (re-)production tasks with continuous action spaces in complex Bayesian actor models. Inference over existing models was usually difficult because of analytical intractability of the models. This limited existing models to use quadratic cost functions and Gaussian probability distributions. By using a pre-trained neural network, we provide a method of obtaining numerical solutions to these intractable models, which is fast enough to be used in sampling-based inference methods.

With our framework, researchers from the domain of perception and action can explore the degrees of freedom in modeling human behavior with arbitrary cost function and probability distributions. The general perception and action generation stage can be exchanged completely and the decision-making stage can be adjusted to one's liking as long as the key components like the cost function and decision-making through optimization remain in place.

Using our model, researchers have a new tool to investigate learning. By applying our method to pre-test and post-test data, gathered before and after a training session, one can find the influence of certain kinds of training on the prior beliefs, cost function, perceptual uncertainty and motor variability of the subjects. In addition to these intra-individual differences, we can also examine inter-individual differences by comparing the inferred parameters for different subjects. Or, for instance, the influence of external circumstances on the subjects' parameters can be measured, like we superficially demonstrated with the different puck conditions in the PU data set (cf. Section 4.2.2).

Still, there are some limitations which we discuss in Section 5.1. A common problem in inference over Bayesian actor models incorporating prior beliefs and cost functions is the identifiability problem which arises because different parameter configurations can result in the same behavior, making it hard to discern which of the estimated parameter configuration corresponds to the real parameter values. This problem was already introduced in Section 2.4.1 and is discussed with regard to our method in Section 5.1.1. Section 5.2 goes a little beyond the scope of this work by elaborating on the theoretical assumptions and implications our framework entails. This chapter closes with a short outlook on future work in Section 5.3.

5.1. Constraints and Limitations

One of the most concerning problems of our method is the fact that MCMC chains are regularly getting stuck in their initial states. Usually, we ran eight chains per data set and one to two chains were stuck at most. As described in Section 3.4.3, we filtered all chains that were stuck throughout the whole inference process and only used unstuck chains for the inference of the posterior distributions. Still, it is highly desirable that this does not happen at all. We also tried different initialization methods, e.g. not initializing chains by drawing the initial states from the prior (as we did in all results presented in Section 4) but also initializing with the median of a small number of samples from the prior or initializing uniformly in a symmetric interval around zero in the unconstrained parameter space as recommended by the inventors of NUTS themselves (Hoffman et al., 2014). Usually, chains getting stuck can be a sign of a ill-defined probabilistic model. Despite several checks, we could not find any flaws of that kind in our implementation though. Still, further work should be put into it to ensure a correct formalization and implementation of the probabilistic model.

Furthermore, we observed insufficiently small effective small sizes for the variabilities σ_{prior} , σ_a and σ_p . This is most likely a problem with the model or the inference process, since it happened for all data sets – synthetic and empirical data alike. We need to find how this comes about and how to mitigate this.

Moreover, we observed a decrease in performance for bigger sets of observations in our synthetic data sets. This is not too problematic, since researchers need to collect the data experimentally and thus usually do not have too many data points at hand for each subject anyway.

Nevertheless, more data usually can be used for more accuracy. We should strive for the same behavior of our method and make sure it scales well with more data. The decrease in performance probably comes about by a bigger and harder inference problem when more observations of targets t_i and actions a_i are added, since we need to infer one more latent sensory variable z_i for each new target-action-pair. Perhaps this could be done by hierarchically modeling or inferring the measurements z_i in some way. However, we have no vision of how this could exactly look like, since we also need to preserve single measurements z_i per target-action-pair (t_i, a_i) and pass them into our neural network and thus through the whole model independent from each other.

While our implementation allows for arbitrary cost functions and distributions other than normal distributions, there are still some limitations on the distributions. Since no parameters in the model depend on the action distributions but its only purpose is the generation of action according to some arbitrary distribution form, the form of the action distribution can be chosen freely. We chose a log-normal for the action distribution since it naturally implements signal-dependent noise but also any other distribution could be used. Yet, there is a constraint on the form of the prior distribution in the perception stage of the model where the subject performs inference over the true state of their environment. The likelihood used in the subject’s inference problem must be log-normal based on one of our core assumptions that perception obeys the Weber-Fechner law (cf. Section 2.1.1). For computational convenience, we chose the prior distribution representing the subject’s belief about the experimental statistics (i.e. the distribution of the stimuli and targets t_i) to also be a log-normal distribution since it is a conjugate prior to the log-normal likelihood, also resulting in a log-normal posterior distribution with closed-form solutions for the posterior parameters. A log-normal distribution may not be an appropriate choice for the subject’s prior belief though. For now, our method is yet limited to priors that are conjugate to a log-normal likelihood. This limitation could be overcome by using *Variational Inference* (Blei et al., 2017) for approximating the perceptual posterior in the subject to allow for arbitrary prior distributions in combination with the log-normal likelihood. This was already used in Bayesian Decision-Making by Kuśmierczyk et al. (2019).

In Section 3.2.1 we introduced different functional forms of cost functions that could represent the costs in human decision-making quite well. Yet, we only used version Gen2B in the evaluation of our method. We

want to emphasize here that our framework is not limited to this kind of cost function. Still, since we did not evaluate our method with different cost functions against each other, there might be better choices in the design of the cost function. This should definitely be done in the future to confirm or falsify theories about which functions describe different human costs best.

Moreover, despite its prevalence in control theory and optimal control, our model does not contain any feedback loops but is feed-forward only. A core property of feedback loops is that the difference between target action and actual action is fed back into the system to allow for more accurate regulation of the action. It can be fundamentally debated whether models motivated from control theory, Bayesian models with only feed-forward action generation, a mixture of both or a completely different concept describes human behavior best. Still, it is obvious that people probably incorporate some kind of feedback into their decision-making which we do not model at all. When people utilize feedback on their actions – be it extrinsic from a researcher supervising an experiment or intrinsic because they see for themselves how they missed the target –, this probably results in a per-trial adaptation of the internal model and its parameters. For example, Izawa et al. (2008) propose that humans learn a forward-model describing the sensorimotor consequences of their actions and thus also try to learn an optimal controller. Cohen et al. (2010) also showed that with more practice, the influence of the visual uncertainty got less. Since the visual uncertainty is determined by the properties of the visual system which does not change short-term, this could be modeled with higher reliance on the prior beliefs instead of the sensory uncertainty - that is, the prior uncertainty becomes much less than the sensory uncertainty over time. This is further supported by Petzschner and Glasauer (2011), who obtained accurate results with a model incorporating adaptive priors using a Kalman-Filter. By learning good prior beliefs about the experimental statistics during the course of the experiment, the prior uncertainty gets less and thus the sensory information loses value.

Our method assumes that the behavioral parameters we want to infer are constant over all trials and the respective action observations. This can only hold true if subjects were sufficiently trained on the task and we can assume that any adaptations to the task and their overall behavior converged to an equilibrium state where they perform qualitatively equal in every trial. We could only investigate the development of the behavioral parameters by applying our method block-wise (e.g. the blocks in which the trials were presented to the subject). Still, this probably reduces the accuracy of our method because we have less observations to use in the inference (cf. Section 4.6) and still only models these changes on block-level instead of trial-level, which is probably more realistic.

5.1.1. Identifiability Problems

We already introduced identifiability problems in Section 2.4.1. Usually, these problems arise in models with quadratic cost functions and Gaussian distributions, since the linearity of the Gaussians interacts with the quadratic cost function in a way that multiple prior-cost-metamers emerge that can all describe the same behavior equally well (cf. Figure 2.3, Acerbi et al. (2014)). This usually shows in the form of multimodal posterior distributions over the model parameters or through sampling difficulties due to perfect correlations between two model parameters. We already encountered bimodal posterior distributions when we applied our method to the puck data set (cf. Figure A.12).

The results of our work indicate that there could be some form of identifiability problem in our model but not necessarily concerning the prior-cost-metamers introduced by Acerbi et al. (2014). For example, consider the results from the *Undershoot via Beta* regime in Section 4.1.2 or, the *Regression to the Mean* regime (cf. Section 2.1.5) in general, which can be described by a too strong reliance on the prior representing the distribution of the target stimuli. The mean of the distribution should correspond to the mean of the stimulus distribution and the variability of the prior should become less the more the subject relies on it. Yet, how strongly the

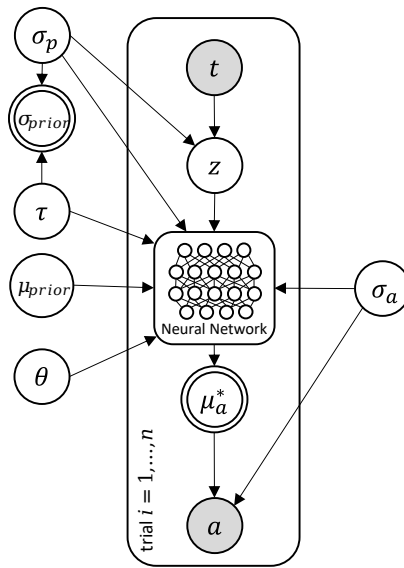


Figure 5.1.: **Modified Probabilistic Model** – Instead of inferring the prior uncertainty σ_{prior} directly, we infer σ_p and the new parameter τ .

resulting action is shifted from the sensory information towards the prior can be influenced by two factors: (1) the variability σ_{prior} in the prior decreases with respect to the sensory uncertainty σ_p and thus the subject relies on the prior more heavily. Or (2) the variability in the prior remains constant but the parameter μ of the log-normal $\log \mathcal{N}(\mu, \sigma)$ increases. Thereby, the resulting action is shifted more strongly towards the priors even though the ratio of uncertainty in the prior and in the likelihood stays the same. Overall, since all the parameters in our model seem to interact in complex ways, any change in the resulting behavior can probably be brought about by many different changes in the parameter configuration. We cannot completely rule out, that we occasionally infer parameter values that describe the observed data well but do not correspond to the ground truth, but the evaluation on synthetic data suggests that our model works reliably well in that matter.

Also, some of these identification problems could eventually be solved by reparameterizing the probabilistic model and changing the variables we infer. For example, we assume that there is a potential identifiability problem between the parameters σ_{prior} and σ_p . Since the ratio of these two variabilities defines the influence of the prior versus the influence of the sensory information and the same ratio can be created with arbitrary values for σ_p and σ_{prior} , the pair of these parameters is not uniquely determined.

We could potentially resolve this issue by introducing a new variable $\tau = \sigma_{\text{prior}}/\sigma_p$ and modifying our probabilistic model as shown in figure 5.1. We do not infer σ_p and σ_{prior} directly but instead try to infer the correct value for σ_p and the new parameter τ , from which the value of σ_{prior} can then be calculated. However, it needs to be tested extensively whether this really improves performance and in which cases it makes the inference more or less accurate.

5.2. Theoretical Grounding of the Framework

Besides providing another option for modeling human behavior in continuous perception-and-action tasks, our framework could contribute to the debate how well suited Bayesian approaches are for describing human behavior overall (cf. Griffiths et al. (2008); Bowers and Davis (2012); Griffiths et al. (2012); Tauber et al. (2017)). Please note that all implications of our work are only valid on the functional / computational level, as in *why* certain behavior emerges (cf. Marr’s three levels of analysis, Marr (2010)). Our model does not imply anything for the algorithmic and implementational levels of human behavior, e.g. *how* certain behavior emerges in the brain. If our work continues to deliver good results, also when applied to many other sets of empirical data, and allows for good interpretation of the inferred parameters, our framework delivers a strong argument that at least human perception indeed is statistically optimal, as for example assumed and investigated by Körding and Wolpert (2004a), Knill and Pouget (2004) and Drugowitsch et al. (2016). We can falsify or support normative Bayesian models only for the perception stage since it is the only component on which we impose Bayesian assumptions on.

Generally, we take on an agnostic position on normative claims whether Bayesian methods really are implemented in the brain. Nevertheless, we see strong value in Bayesian methods as descriptive means from the researcher’s perspective (cf. Section 2.2.1). While MCMC Sampling methods are also a Bayesian method, they are a purely descriptive tool for obtaining estimates of the parameters in our model and do not impose any normative assumptions on how human perception and action works.

However, we also use Bayesian modeling from the subject’s view in the perceptual component, which we modeled as Bayesian inference over the true state of the world, since it is a common view in modern perception research (cf. Section 2.1.2). Nonetheless, this component could also be substituted with any other mechanism that yields a percept of the true world state for the subject – be it Bayesian or not – without interfering with our framework’s assumptions or concepts.

While there is empirical evidence from neuroscience that brains have the ability to encode probability distributions (Hsu et al., 2005; Beck et al., 2008; Fiser et al., 2010) and utilities or costs (Schultz et al., 1997; Knutson and Peterson, 2005; Rudebeck et al., 2006; Walton et al., 2006; Prévost et al., 2010; Hosokawa et al., 2013), we do not take any position regarding the real mechanism in the human brain. There may or may not be an explicit neural representation of cost functions. For us, these cost functions are solely a tool for interpreting and explaining human behavior quantitatively in the form of abstract parameters. The same applies to the measurements z_i of the true stimuli t_i , the subject is provided with by its sensory system. It is possible that there are scalar values encoded in the human brain somehow or that these measurements do not take on a scalar form but are encoded as a probability distribution in the brain. These are important questions but we leave it to neuroscience to answer them.

5.3. Future Work

First of all, future work regarding our framework should focus on finding a way to avoid the MCMC chains getting stuck in their initial states to have posterior estimates build from all run chains and not only a subset of them. Furthermore, it shows that our probabilistic model is correctly specified.

Also, our method should be applied to more empirical data sets to identify strengths and especially weaknesses in our method. The two experimentally collected data sets we used for our evaluation were quite similar because they both involved subjects moving an object (bean bag or puck) to a certain distance that is visually perceived. It would be interesting to apply our method also to different kinds of data, where the stimulus are perceived through a different modality, for instance the auditory system, and action are of a different kind, for example by holding down a button for a certain duration instead of accelerating one object with a quick movement. We need to know under which circumstances it fails to perform well or how experiments should be designed to guarantee that our method can be applied to the observed data successfully.

Additionally, our method should again be evaluated with different cost functions to find a functional form that describes human behavior best, for example with the IG2B-form introduced in Section 3.2.1. Besides, while the log-normal distribution for the likelihood and action distribution are well grounded in the Weber-Fechner law and signal-dependent noise, respectively, it makes sense to enable an extension of the prior distribution to other distributions than a log-normal. We have no reason to assume a log-normal form of the prior distribution other than that it is conjugate to the log-normal likelihood. In order to do so, one could use variational inference methods to approximate the posterior for non-conjugate priors. Also, we do not incorporate any learning processes in our model yet. This could be extended with sequential Monte Carlo methods (Cappe et al., 2007) over trials or with approaches similar to estimating the prior parameters in each trial with a Kalman-Filter as done by Petzschner and Glasauer (2011).

Another possible extension of our model towards more general perception-action-models is the integration of multimodal perception. While we currently only rely on one modality – mostly vision because it is one of the most dominant senses –, it makes sense to extend the perception stage to receive sensory input in the form of, e.g., auditory or haptic information. This can be easily done in our framework, for example the resulting inference process of the subject for the true state of the world could work with a percept from vision and audio. The resulting posterior belief about the true state of the world can then be based on a cue-combination model (Ernst and Bühlhoff, 2004; Fetsch et al., 2012). Just as the neural network is now trained on the joint process of world state inference and optimization over the action distribution's parameters, it would then be trained with the cue-combination model for inference of the world state.

6. Conclusion

We introduce a new model for human behavior in (re-)production tasks. Usually, Bayesian actor models in continuous action and stimulus space are constrained to loss functions of certain forms (e.g. quadratic) and distributions of certain forms (e.g. Gaussian) to preserve analytical tractability. We free researchers from these constraints by making previously intractable Bayesian actor models in continuous space solvable with little computational demand by approximating optimization and inference processes within the subject using a neural network. This makes sampling-based inference methods feasible to be used on our model, which was previously not possible due to the computational demand of numerically solving a model that has no analytical solution. Thus, we can infer parameters describing behavioral costs, perceptual uncertainty, motor variability and prior beliefs of the subject solving the task. We show that our method can reliably extract parameters on synthetic data sets as well as on empirically collected data sets. Still, our method should be tested and evaluated on various other data sets of empirical data to identify the conditions under which it performs better or worse.

Moreover, there is potential in optimizing the probabilistic model and the inference process, e.g. to avoid MCMC chains getting stuck and improve MCMC diagnostics. Because of the generality of our approach, it can be easily extended to, for instance, account for learning effects or incorporate multimodal perception. We allow researchers to model and analyze behavioral data from continuous decision-making tasks with more degrees of freedom and enable them to test more complex theories in the perception, planning and action stages that constitute human behavior.

References

- Acerbi, L., Ma, W. J., and Vijayakumar, S. (2014). A Framework for Testing Identifiability of Bayesian Models of Perception. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Adams, W. J. (2008). Frames of reference for the light-from-above prior in visual search and shape judgements. *Cognition*, 107(1):137–150.
- Ashby, F. G. and Maddox, W. T. (2005). Human Category Learning. *Annual Review of Psychology*, 56(1):149–178.
- Barrett, L. F. (2017). *How emotions are made: The secret life of the brain*. Pan Macmillan.
- Barthelmé, S. and Mamassian, P. (2009). Evaluation of Objective Uncertainty in the Visual System. *PLoS Computational Biology*, 5(9):e1000504.
- Battaglia, P. W., Kersten, D., and Schrater, P. R. (2011). How Haptic Size Sensations Improve Distance Perception. *PLoS Computational Biology*, 7(6):e1002080.
- Bayes, T. (1763). LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, F. R. S. communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S. *Philosophical Transactions of the Royal Society of London*, 53:370–418.
- Beck, J., Ma, W., Pitkow, X., Latham, P., and Pouget, A. (2012). Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron*, 74(1):30–39.
- Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., Shadlen, M. N., Latham, P. E., and Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, 60(6):1142–1152. Publisher: Elsevier.
- Becker, W. and Fuchs, A. (1969). Further properties of the human saccadic system: Eye movements and correction saccades with and without visual fixation points. *Vision Research*, 9(10):1247–1258.
- Bergmann, J., Krauß, E., Münch, A., Jungmann, R., Oberfeld, D., and Hecht, H. (2011). Locomotor and verbal distance judgments in action and vista space. *Experimental Brain Research*, 210(1):13–23.
- Betancourt, M. (2018). A Conceptual Introduction to Hamiltonian Monte Carlo. arXiv:1701.02434 [stat].
- Birkenbusch, J., Ellermeier, W., and Kattner, F. (2015). Octuplicate this interval! Axiomatic examination of the ratio properties of duration perception. *Attention, Perception, & Psychophysics*, 77(5):1767–1780.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877.
- Blum, M. G. B. and François, O. (2010). Non-linear regression models for Approximate Bayesian Computation. *Statistics and Computing*, 20(1):63–73.

-
- Bolle, R. M. and Cooper, D. B. (1984). Bayesian Recognition of Local 3-D Shape by Approximating Image Intensity Functions with Quadric Polynomials. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(4):418–429.
- Bowers, J. S. and Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3):389–414.
- Brayanov, J. B. and Smith, M. A. (2010). Bayesian and “Anti-Bayesian” Biases in Sensory Integration for Action and Perception in the Size–Weight Illusion. *Journal of Neurophysiology*, 103(3):1518–1531.
- Brent, R. P. (2013). *Algorithms for minimization without derivatives*. Courier Corporation.
- Bresciani, J.-P., Dammeier, F., and Ernst, M. O. (2006). Vision and touch are automatically integrated for the perception of sequences of events. *Journal of Vision*, 6(5):2.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14(4):411–427.
- Buhusi, C. V. and Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience*, 6(10):755–765.
- Burge, J., Fowlkes, C. C., and Banks, M. S. (2010). Natural-Scene Statistics Predict How the Figure–Ground Cue of Convexity Affects Human Depth Perception. *The Journal of Neuroscience*, 30(21):7269–7280.
- Bürkner, P.-C. (2017). **brms** : An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1).
- Cantlon, J. F., Cordes, S., Libertus, M. E., and Brannon, E. M. (2009). Comment on “Log or Linear? Distinct Intuitions of the Number Scale in Western and Amazonian Indigene Cultures”. *Science*, 323(5910):38–38.
- Cappe, O., Godsill, S. J., and Moulines, E. (2007). An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924.
- Chambon, V., Pacherie, E., Barbalat, G., Jacquet, P., Franck, N., and Farrer, C. (2011). Mentalizing under influence: abnormal dependence on prior expectations in patients with schizophrenia. *Brain*, 134(12):3728–3741.
- Chater, N. and Manning, C. D. (2006). Probabilistic models of language processing and acquisition. *Trends in cognitive sciences*, 10(7):335–344. Publisher: Elsevier.
- Chen, Y., McBain, R., Norton, D., and Ongur, D. (2011). Schizophrenia patients show augmented spatial frame illusion for visual and visuomotor tasks. *Neuroscience*, 172:419–426.
- Clark, J. J. and Yuille, A. L. (2013). *Data fusion for sensory information processing systems*, volume 105. Springer Science & Business Media.
- Cohen, I., Huang, Y., Chen, J., Benesty, J., Benesty, J., Chen, J., Huang, Y., and Cohen, I. (2009). Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4.
- Cohen, R. G., Biddle, J. C., and Rosenbaum, D. A. (2010). Manual obstacle avoidance takes into account visual uncertainty, motor noise, and biomechanical costs. *Experimental Brain Research*, 201(3):587–592.
- De C. Hamilton, A. F., Jones, K. E., and Wolpert, D. M. (2004). The scaling of motor noise with muscle strength and motor unit number in humans. *Experimental Brain Research*, 157(4):417–430.

-
- De Gardelle, V. and Mamassian, P. (2014). Does Confidence Use a Common Currency Across Two Visual Tasks? *Psychological Science*, 25(6):1286–1288.
- De Gardelle, V. and Mamassian, P. (2015). Weighting Mean and Variability during Confidence Judgments. *PLOS ONE*, 10(3):e0120870.
- De Lange, F. P., Rahnev, D. A., Donner, T. H., and Lau, H. (2013). Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *The Journal of Neuroscience*, 33(4):1400–1410.
- Dehaene, S. (2003). The neural basis of the Weber–Fechner law: a logarithmic mental number line. *TRENDS in Cognitive Sciences*, 7(4):145–147.
- Dehaene, S., Izard, V., Spelke, E., and Pica, P. (2008). Log or Linear? Distinct Intuitions of the Number Scale in Western and Amazonian Indigene Cultures. *Science*, 320(5880):1217–1220.
- Dillon, J. V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., Patton, B., Alemi, A., Hoffman, M., and Saurous, R. A. (2017). TensorFlow Distributions. arXiv:1711.10604 [cs, stat].
- Drugowitsch, J., Wyart, V., Devauchelle, A.-D., and Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron*, 92(6):1398–1411.
- Eisler, H. and Ottander, C. (1963). On the problem of hysteresis in psychophysics. *Journal of Experimental Psychology*, 65(6):530–536.
- Elliott, D., Hansen, S., Grierson, L. E. M., Lyons, J., Bennett, S. J., and Hayes, S. J. (2010). Goal-directed aiming: Two components but multiple processes. *Psychological Bulletin*, 136(6):1023–1044.
- Elliott, D., Hansen, S., Mendoza, J., and Tremblay, L. (2004). Learning to Optimize Speed, Accuracy, and Energy Expenditure: A Framework for Understanding Speed-Accuracy Relations in Goal-Directed Aiming. *Journal of Motor Behavior*, 36(3):339–351.
- Engelbrecht, S. E., Berthier, N. E., and O’Sullivan, L. P. (2003). The Undershoot Bias: Learning to Act Optimally Under Uncertainty. *Psychological Science*, 14(3):257–261.
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433.
- Ernst, M. O. and Bühlhoff, H. H. (2004). Merging the senses into a robust percept. *Trends in Cognitive Sciences*, 8(4):162–169.
- Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303.
- Fechner, G. T. (1860). *Elements of psychophysics*, 1860. Publisher: Appleton-Century-Crofts.
- Fengler, A., Govindarajan, L. N., Chen, T., and Frank, M. J. (2021). Likelihood approximation networks (LANs) for fast inference of simulation models in cognitive neuroscience. *eLife*, 10:e65074.
- Fetsch, C. R., Pouget, A., DeAngelis, G. C., and Angelaki, D. E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. *Nature Neuroscience*, 15(1):146–154.
- Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130.
- Fraisse, P. (1984). Perception and Estimation of Time. *Ann.Rev. Psychol.*, 35:1–36.

-
- Geisler, W. S., Najemnik, J., and Ing, A. D. (2009). Optimal stimulus encoders for natural tasks. *Journal of Vision*, 9(13):17–17.
- Geisler, W. S., Perry, J. S., Super, B., and Gallogly, D. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision research*, 41(6):711–724. Publisher: Elsevier.
- Geman, S. and Geman, D. (1984). Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741.
- Gibson, J. J. (1966). The senses considered as perceptual systems.
- Girshick, A. R., Landy, M. S., and Simoncelli, E. P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14(7):926–932.
- Gregory, R. L. (1997). Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 352(1358):1121–1127.
- Griffiths, T., Kemp, C., and Tenenbaum, J. (2008). Bayesian models of cognition. Publisher: Carnegie Mellon University.
- Griffiths, T. L., Chater, N., Norris, D., and Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012). *Psychological Bulletin*, 138(3):415–422.
- Griffiths, T. L. and Tenenbaum, J. B. (2006). Optimal Predictions in Everyday Cognition. *Psychological Science*, 17(9):767–773.
- Hagura, N., Haggard, P., and Diedrichsen, J. (2017). Perceptual decisions are biased by the cost to act. *eLife*, 6:e18422.
- Harris, C. M. (1995). Does saccadic undershoot minimize saccadic flight-time? A Monte-Carlo study. *Vision Research*, 35(5):691–701.
- Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–784.
- Harris, L. R., Jenkin, M., and Zikovitz, D. C. (2000). Visual and non-visual cues in the perception of linear self motion. *Experimental Brain Research*, 135(1):12–21.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. Publisher: Oxford University Press.
- Helmholtz, H. v. (1867). Concerning the perceptions in general. Publisher: Appleton-Century-Crofts.
- Hillis, J. M., Watt, S. J., Landy, M. S., and Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, 4(12):1.
- Hock, H., Nichols, D., Huisman, A., Rivera, M., and Bukowski, L. (2005). Dynamical vs. judgmental comparison: hysteresis effects in motion perception. *Spatial Vision*, 18(3):317–335.
- Hoffman, M. D., Gelman, A., and others (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623.
- Hosokawa, T., Kennerley, S. W., Sloan, J., and Wallis, J. D. (2013). Single-Neuron Mechanisms Underlying Cost-Benefit Analysis in Frontal Cortex. *The Journal of Neuroscience*, 33(44):17385–17397.
- Houlsby, N., Huszár, F., Ghassemi, M., Orbán, G., Wolpert, D., and Lengyel, M. (2013). Cognitive Tomography Reveals Complex, Task-Independent Mental Representations. *Current Biology*, 23(21):2169–2175.

-
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C. F. (2005). Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making. *Science*, 310(5754):1680–1683.
- Huang, H. J., Kram, R., and Ahmed, A. A. (2012). Reduction of Metabolic Cost during Motor Learning of Arm Reaching Dynamics. *The Journal of Neuroscience*, 32(6):2182–2190.
- Izawa, J., Rane, T., Donchin, O., and Shadmehr, R. (2008). Motor Adaptation as a Process of Reoptimization. *The Journal of Neuroscience*, 28(11):2883–2891.
- Jazayeri, M. and Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, 13(8):1020–1026.
- Jones, K. E., Hamilton, A. F. D. C., and Wolpert, D. M. (2002). Sources of Signal-Dependent Noise During Isometric Force Production. *Journal of Neurophysiology*, 88(3):1533–1544.
- Jones, M. and Love, B. C. (2011). Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(4):169–188.
- Kable, J. W. and Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10(12):1625–1633.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems.
- Kersten, D. and Mamassian, P. (2009). Ideal Observer Theory. *Encyclopedia of neuroscience*, 5:89–95. Publisher: Oxford University Press Oxford, UK.
- Kersten, D., Mamassian, P., and Yuille, A. (2004). Object Perception as Bayesian Inference. *Annual Review of Psychology*, 55(1):271–304.
- Knill, D. C. and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719.
- Knill, D. C. and Richards, W. (1996). *Perception as Bayesian inference*. Cambridge University Press.
- Knutson, B., Adams, C. M., Fong, G. W., and Hommer, D. (2001). Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *The Journal of Neuroscience*, 21(16):RC159–RC159.
- Knutson, B. and Peterson, R. (2005). Neurally reconstructing expected utility. *Games and Economic Behavior*, 52(2):305–315.
- Krueger, L. E. (1989). Reconciling Fechner and Stevens: Toward a unified psychophysical law. *Behavioral and Brain Sciences*, 12(2):251–267.
- Kruschke, J. K. (2015). *Doing Bayesian Data Analysis*. Elsevier, 2 edition.
- Kuśmierczyk, T., Sakaya, J., and Klami, A. (2019). Variational Bayesian Decision-making for Continuous Utilities. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F. d., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Körding, K. (2007). Decision Theory: What "Should" the Nervous System Do? *Science*, 318(5850):606–610.
- Körding, K. P. and Wolpert, D. M. (2004a). Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247.
- Körding, K. P. and Wolpert, D. M. (2004b). The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences*, 101(26):9839–9842.

-
- Langer, M. S. and Bühlhoff, H. H. (2001). A Prior for Global Convexity in Local Shape-from-Shading. *Perception*, 30(4):403–410.
- Laski, E. V. and Siegler, R. S. (2007). Is 27 a Big Number? Correlational and Causal Connections Among Numerical Categorization, Number Line Estimation, and Numerical Magnitude Comparison. *Child Development*, 78(6):1723–1743.
- Lejeune, H. and Jasselette, P. (1986). Accurate DRL performance in the pigeon: Comparison between perching and treadle pressing. *Animal Learning & Behavior*, 14(2):205–211.
- Lejeune, H. and Richelle, M. (1982). Differential reinforcement of perching duration in the pigeon : a comparison with differential-reinforcement-of-low-rate key-pecking. *Behaviour Analysis Letters*, 2. Publisher: Elsevier.
- Lewis, F. L., Vrabie, D., and Syrmos, V. L. (2012). *Optimal control*. John Wiley & Sons.
- Longo, M. R. and Lourenco, S. F. (2007). Spatial attention and the mental number line: Evidence for characteristic biases and compression. *Neuropsychologia*, 45(7):1400–1407.
- Lowe, C. F., Harzem, P., and Spencer, P. T. (1979). Temporal Control of Behavior and the Power Law. *Journal of the Experimental Analysis of Behavior*, 31(3):333–343.
- Lueckmann, J.-M., Bassetto, G., Karaletsos, T., and Macke, J. H. (2019). Likelihood-free inference with emulator networks. In Ruiz, F., Zhang, C., Liang, D., and Bui, T., editors, *Proceedings of The 1st Symposium on Advances in Approximate Bayesian Inference*, volume 96 of *Proceedings of Machine Learning Research*, pages 32–53. PMLR.
- Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11):1432–1438.
- MacKay, D. M. (1963). Psychophysics of Perceived Intensity: A Theoretical Basis for Fechner’s and Stevens’ Laws. *Science*, 139(3560):1213–1216.
- Mamassian, P. and Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research*, 41(20):2653–2668.
- Marr, D. (2010). *Vision: a computational investigation into the human representation and processing of visual information*. MIT Press, Cambridge, Mass. OCLC: ocn472791457.
- McCoy, A. N. and Platt, M. L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience*, 8(9):1220–1227.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092. Publisher: American Institute of Physics.
- Mittelstaedt, M.-L. and Mittelstaedt, H. (2001). Idiopathic navigation in humans: estimation of path length. *Experimental Brain Research*, 139(3):318–332.
- Neupärtl, N. and Rothkopf, C. A. (2021). Inferring perceptual decision making parameters from behavior in production and reproduction tasks. Publisher: arXiv Version Number: 1.
- Neupärtl, N., Tatai, F., and Rothkopf, C. A. (2020). Intuitive physical reasoning about objects’ masses transfers to a visuomotor decision task consistent with Newtonian physics. *PLOS Computational Biology*, 16(10):e1007730.

-
- Notredame, C.-E., Pins, D., Deneve, S., and Jardri, R. (2014). What visual illusions teach us about schizophrenia. *Frontiers in Integrative Neuroscience*, 8.
- Papamakarios, G. and Murray, I. (2016). Fast ϵ -free Inference of Simulation Models with Bayesian Conditional Density Estimation. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Papamakarios, G., Sterratt, D., and Murray, I. (2019). Sequential Neural Likelihood: Fast Likelihood-free Inference with Autoregressive Flows. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 837–848. PMLR.
- Pellicano, E. and Burr, D. (2012). When the world becomes ‘too real’: a Bayesian explanation of autistic perception. *Trends in Cognitive Sciences*, 16(10):504–510.
- Pessoa, V., Monge-Fuentes, V., Simon, C., Suganuma, E., and Tavares, M. (2008). The Müller-Lyer Illusion as a Tool for Schizophrenia Screening. *Reviews in the Neurosciences*, 19(2-3).
- Petzschnner, F. H. and Glasauer, S. (2011). Iterative Bayesian Estimation as an Explanation for Range and Regression Effects: A Study on Human Path Integration. *The Journal of Neuroscience*, 31(47):17220–17229.
- Petzschnner, F. H., Glasauer, S., and Stephan, K. E. (2015). A Bayesian perspective on magnitude estimation. *Trends in Cognitive Sciences*, 19(5):285–293.
- Pizlo, Z. (2001). Perception viewed as an inverse problem. *Vision Research*, 41(24):3145–3161.
- Portugal, R. D. and Svaiter, B. F. (2011). Weber-Fechner Law and the Optimality of the Logarithmic Scale. *Minds and Machines*, 21(1):73–81.
- Poulton, E. C. (1968). The new psychophysics: Six models for magnitude estimation. *Psychological Bulletin*, 69(1):1–19.
- Prins, N. and others (2016). *Psychophysics: a practical introduction*. Academic Press.
- Prévost, C., Pessiglione, M., Météreau, E., Cléry-Melin, M.-L., and Dreher, J.-C. (2010). Separate Valuation Subsystems for Delay and Effort Decision Costs. *The Journal of Neuroscience*, 30(42):14080–14090.
- Rahnev, D. and Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41:e223.
- Ramachandran, V. S. (1988). Perception of shape from shading. *Nature*, 331(6152):163–166.
- Renart, A. and Machens, C. K. (2014). Variability in neural activity and behavior. *Current Opinion in Neurobiology*, 25:211–220.
- Robert, C. P. and Casella, G. (2004). *Monte Carlo statistical methods*. Springer Texts in Statistics. Springer-Verlag, New York, second edition.
- Roberts, W. A. (2006). Evidence that pigeons represent both time and number on a logarithmic scale. *Behavioural Processes*, 72(3):207–214.
- Rock, I. (1983). The logic of perception.
- Rosas, P. and Wichmann, F. A. (2011). Cue combination: Beyond optimality. *Sensory cue integration*, pages 144–52. Publisher: Oxford University Press.

-
- Rudebeck, P. H., Walton, M. E., Smyth, A. N., Bannerman, D. M., and Rushworth, M. F. S. (2006). Separate neural pathways process different decision costs. *Nature Neuroscience*, 9(9):1161–1168.
- Scheler, G. (2017). Logarithmic distributions prove that intrinsic learning is Hebbian. *F1000Research*, 6:1222.
- Schlicht, E. J. and Schrater, P. R. (2007). Effects of visual uncertainty on grasping movements. *Experimental Brain Research*, 182(1):47–57.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, 275(5306):1593–1599.
- Shams, L., Ma, W. J., and Beierholm, U. (2005). Sound-induced flash illusion as an optimal percept.
- Shepard, R. N. (1981). Psychological relations and psychophysical scales: On the status of “direct” psychophysical measurement. *Journal of Mathematical Psychology*, 24(1):21–57.
- Shergill, S. S., Samson, G., Bays, P. M., Frith, C. D., and Wolpert, D. M. (2005). Evidence for Sensory Prediction Deficits in Schizophrenia. *American Journal of Psychiatry*, 162(12):2384–2386.
- Simoncelli, E. P. (1993). Distributed analysis and representation of visual motion. *Unpublished doctoral dissertation, MIT*.
- Simoncelli, E. P. (2003). Vision and the statistics of the visual environment. *Current Opinion in Neurobiology*, 13(2):144–149.
- Sims, C. R. and Gray, W. D. (2008). Adaptation to embodied dynamics: Evidence from Bayes’ ball. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 30. Issue: 30.
- Sisson, S. A., Fan, Y., and Beaumont, M. (2018). *Handbook of approximate Bayesian computation*. CRC Press.
- Sloman, S. A. and Lagnado, D. (2015). Causality in Thought. *Annual Review of Psychology*, 66(1):223–247.
- Sohn, H. and Jazayeri, M. (2021). Validating model-based Bayesian integration using prior–cost metamers. *Proceedings of the National Academy of Sciences*, 118(25):e2021531118.
- Stengård, E. and Van Den Berg, R. (2019). Imperfect Bayesian inference in visual perception. *PLOS Computational Biology*, 15(4):e1006465.
- Stevens, S. S. (1961). To Honor Fechner and Repeal His Law: A power function, not a log function, describes the operating characteristic of a sensory system. *Science*, 133(3446):80–86. Publisher: American Association for the Advancement of Science.
- Sun, H.-J., Campos, J. L., and Chan, G. S. W. (2004a). Multisensory integration in the estimation of relative path length. *Experimental Brain Research*, 154(2):246–254.
- Sun, H.-J., Campos, J. L., Young, M., Chan, G. S. W., and Ellard, C. G. (2004b). The Contributions of Static Visual Cues, Nonvisual Cues, and Optic Flow in Distance Estimation. *Perception*, 33(1):49–65.
- Sun, J. Z., Wang, G. I., Goyal, V. K., and Varshney, L. R. (2012). A framework for Bayesian optimality of psychophysical laws. *Journal of Mathematical Psychology*, 56(6):495–501.
- Tassinari, H., Hudson, T. E., and Landy, M. S. (2006). Combining Priors and Noisy Visual Cues in a Rapid Pointing Task. *The Journal of Neuroscience*, 26(40):10154–10163.
- Tauber, S., Navarro, D. J., Perfors, A., and Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review*, 124(4):410–441.

-
- Teghtsoonian, R. and Teghtsoonian, M. (1978). Range and regression effects in magnitude scaling. *Perception & Psychophysics*, 24(4):305–314.
- Tieleman, T., Hinton, G., and others (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31.
- Todorov, E. and Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235.
- Trommershäuser, J., Landy, M. S., and Maloney, L. T. (2006). Humans Rapidly Estimate Expected Gain in Movement Planning. *Psychological Science*, 17(11):981–988.
- Trommershäuser, J., Maloney, L. T., and Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *J. Opt. Soc. Am. A*, 20(7):1419–1433. Publisher: Optica Publishing Group.
- Trommershäuser, J., Maloney, L. T., and Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, 12(8):291–297.
- Van Beers, R. J., Haggard, P., and Wolpert, D. M. (2004). The Role of Execution Noise in Movement Variability. *Journal of Neurophysiology*, 91(2):1050–1063.
- Van Bergen, R. S., Ji Ma, W., Pratte, M. S., and Jehee, J. F. M. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*, 18(12):1728–1730.
- Van Den Berg, R., Vogel, M., Josić, K., and Ma, W. J. (2012). Optimal inference of sameness. *Proceedings of the National Academy of Sciences*, 109(8):3178–3183.
- Vilares, I. and Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain: Bayesian models and the world. *Annals of the New York Academy of Sciences*, 1224(1):22–39.
- Walton, M., Kennerley, S., Bannerman, D., Phillips, P., and Rushworth, M. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, 19(8):1302–1314.
- Walton, M. E., Bannerman, D. M., Alterescu, K., and Rushworth, M. F. S. (2003). Functional Specialization within Medial Frontal Cortex of the Anterior Cingulate for Evaluating Effort-Related Decisions. *The Journal of Neuroscience*, 23(16):6475–6479.
- Weber, E. H. (1831). *De Pulsu, resorptione, auditu et tactu: Annotationes anatomicae et physiologicae*. Koehler.
- Weber, R. B. and Daroff, R. B. (1971). The metrics of horizontal saccadic eye movements in normal humans. *Vision Research*, 11(9):921–IN2.
- Wei, X.-X. and Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nature Neuroscience*, 18(10):1509–1517.
- Wei, X.-X. and Stocker, A. A. (2017). Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences*, 114(38):10244–10249.
- Weiss, Y., Simoncelli, E. P., and Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6):598–604.
- Whiteley, L. and Sahani, M. (2008). Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of Vision*, 8(3):2.
- Wiley, C. R. and Liu, Z. (2018). Long-term motor learning: Effects of varied and specific practice. *Vision Research*, 152:10–16.

-
- Wixted, J. T. and Thompson-Schill, S. L. (2018). *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, Language and Thought*, volume 3. John Wiley & Sons.
- Wolpert, D. M. and Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, 3(S11):1212–1217.
- Wolpert, D. M. and Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, 22(6):996–1003.
- Wyart, V. and Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, 11:109–115.
- Wynn, P. (1956). On a device for computing the e m (S n) transformation. *Mathematical Tables and Other Aids to Computation*, pages 91–96. Publisher: JSTOR.
- Yi, L. (2009). Do rats represent time logarithmically or linearly? *Behavioural Processes*, 81(2):274–279.
- Yoo, S. B. M., Hayden, B. Y., and Pearson, J. M. (2021). Continuous decisions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1819):20190664.
- Zeiler, M. D. and Hoyert, M. S. (1989). TEMPORAL REPRODUCTION. *Journal of the Experimental Analysis of Behavior*, 52(2):81–95.
- Zemel, R. S., Dayan, P., and Pouget, A. (1998). Probabilistic Interpretation of Population Codes. *Neural Computation*, 10(2):403–430.

A. Appendix




Figure A.1.: **Cost Function Gen1B** – $\alpha \in \{.5, 1.0, \mathbf{2.0}\}$, $\beta \in \{.5, .7, .9\}$

The dashed grey line represents the target t . Parameter β is constant on the left and α is constant on the right (respective bold value). Parameter values increase from dotted to solid line styles.




Figure A.2.: **Cost Function IG1B** – $\sigma_{\text{cost}} \in \{.5, \mathbf{1.0}, 2.0\}$, $\beta \in \{.2, \mathbf{.4}, .6\}$
 The dashed grey line represents the target t . Parameter β is constant on the left and σ_{cost} is constant on the right (respective bold value). Parameter values increase from dotted to solid line styles.

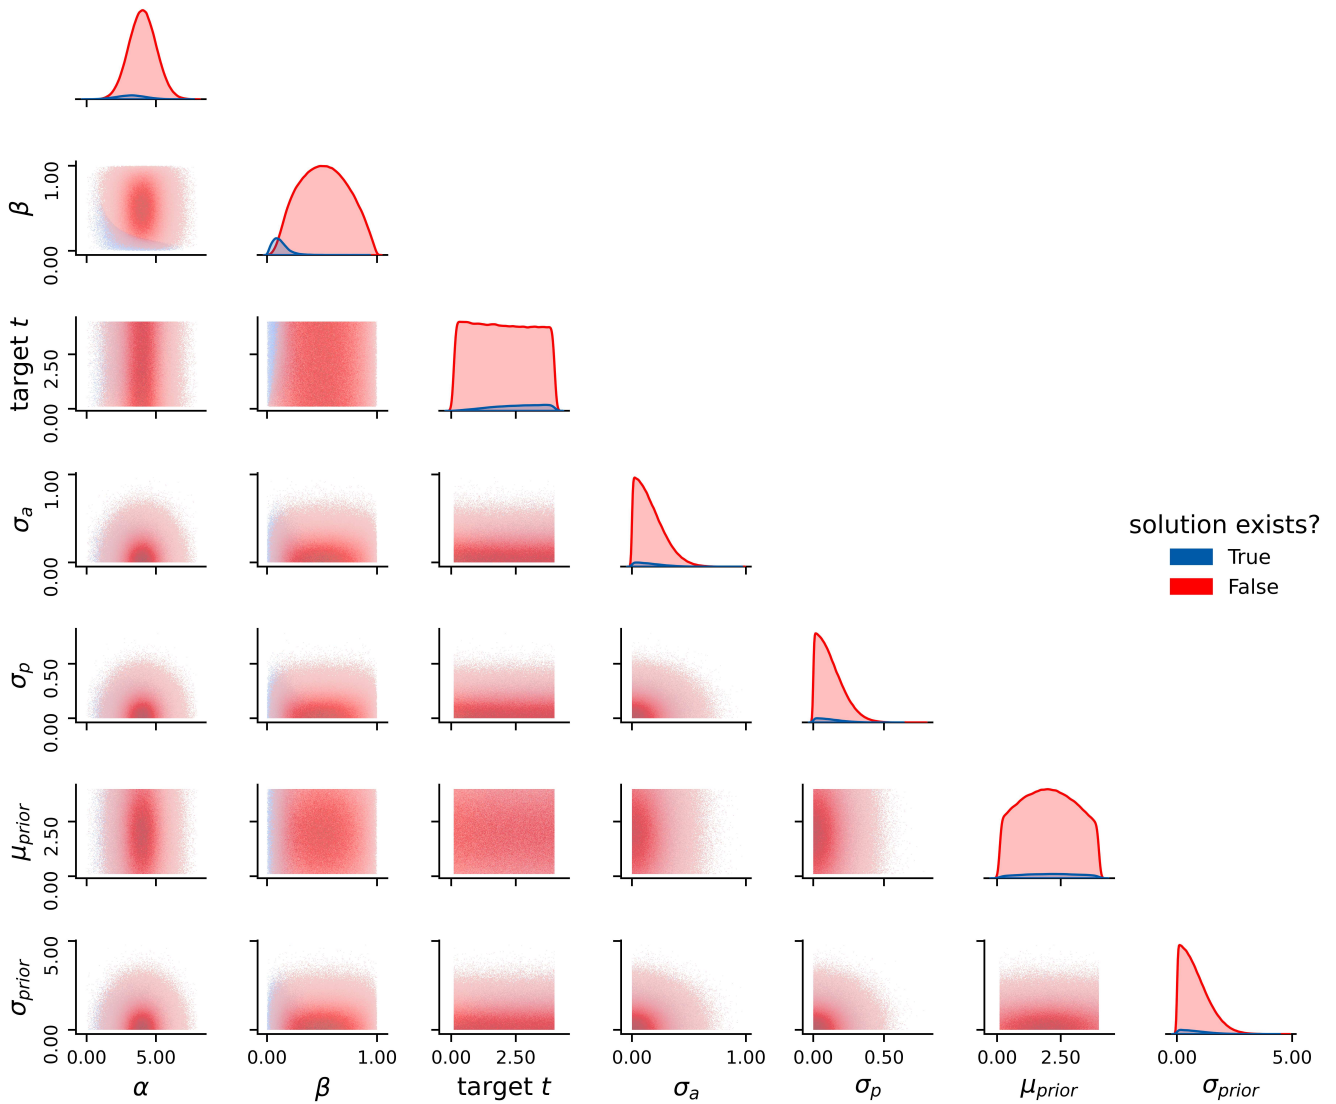


Figure A.3.: **Solvability Plot for Gen1B** – Histograms of samples with (blue) and without (red) existing solution. The histogram without existing solutions is plotted above the histogram with existing solutions.

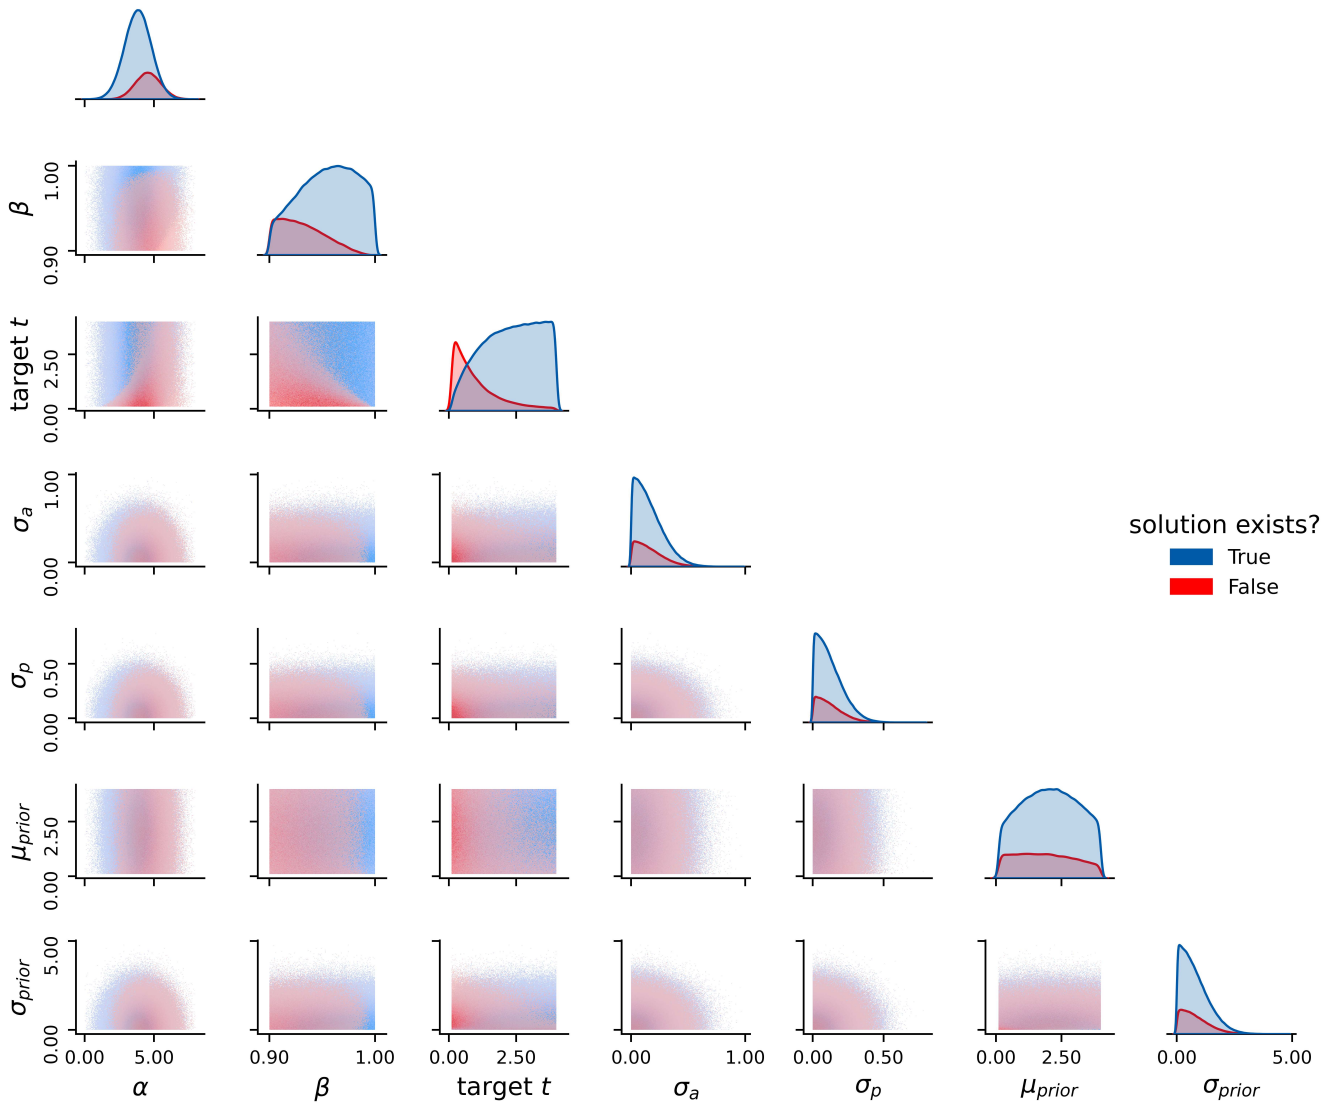


Figure A.4: **Solvability Plot for IG1B** – Histograms of samples with (blue) and without (red) existing solution. The histogram without existing solutions is plotted above the histogram with existing solutions. Any values of β lower than 0.9 did not lead to solvable problems, therefore we only sampled these values in $[0.9, 1.0]$. This is a perfect example of how the form and parameterization of the cost function changes valid parameter ranges.

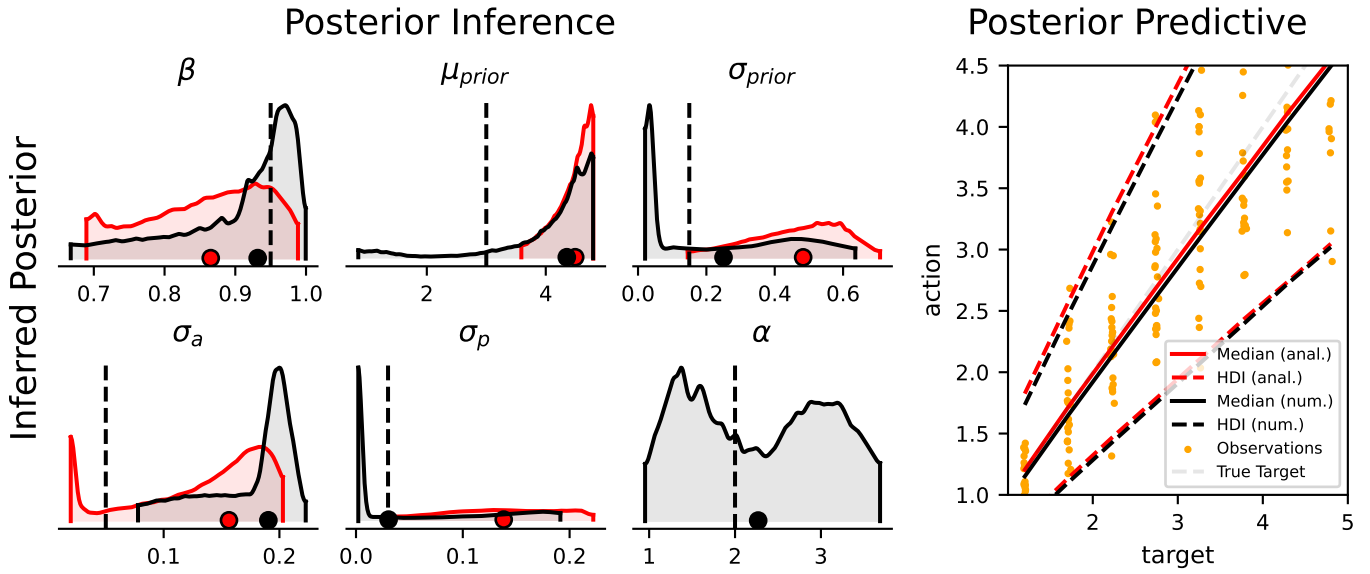


Figure A.5.: MCMC Numerical versus Analytical Solver: *Hitting Target Regime*

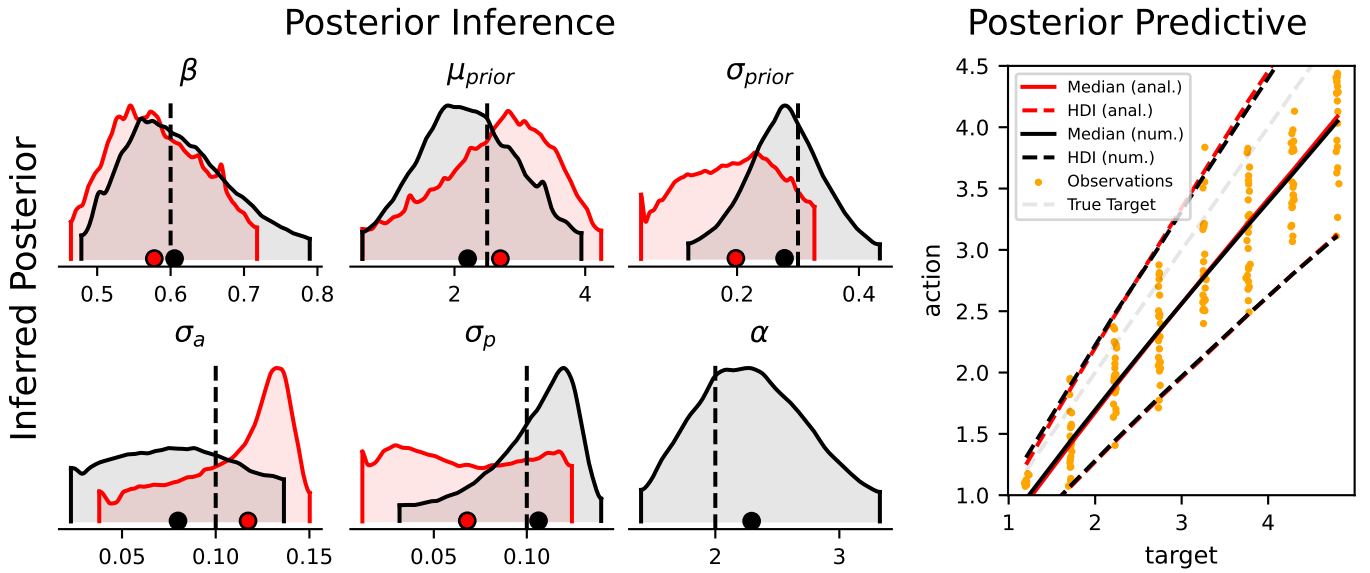


Figure A.6.: MCMC Numerical versus Analytical Solver: *Undershoot via Cost Regime*

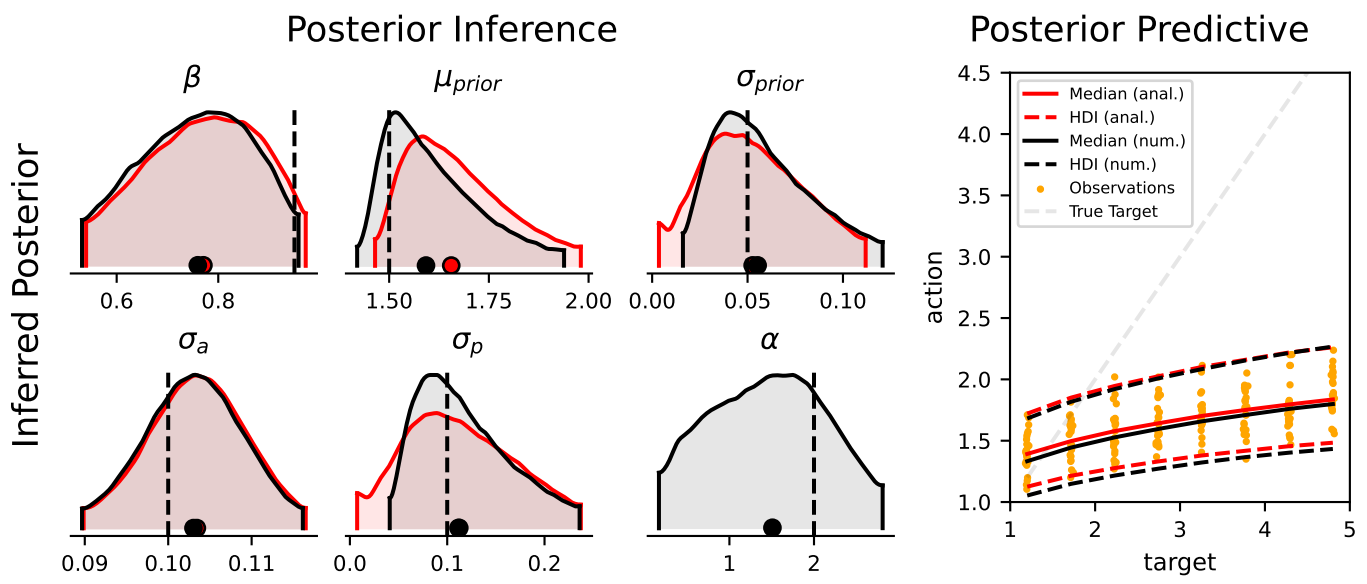


Figure A.7.: **MCMC Numerical versus Analytical Solver: Undershoot via Prior Regime**




Figure A.8.: **Inference Results for Bean Bag (BB) Data** – 95% HDI for each parameter per subject. Dots denote the inferred posterior means, and thick lines denote the area between the central quartiles.




Figure A.9.: **Inference Results for Bean Bag (BB) Data** – 95% HDI for each parameter per subject. Dots denote the inferred posterior means, and thick lines denote the area between the central quartiles.




Figure A.10.: **Inference Results for Puck (PU) Data** – 95% HDI for each parameter per subject. Dots denote the inferred posterior means, and thick lines denote the area between the central quartiles.

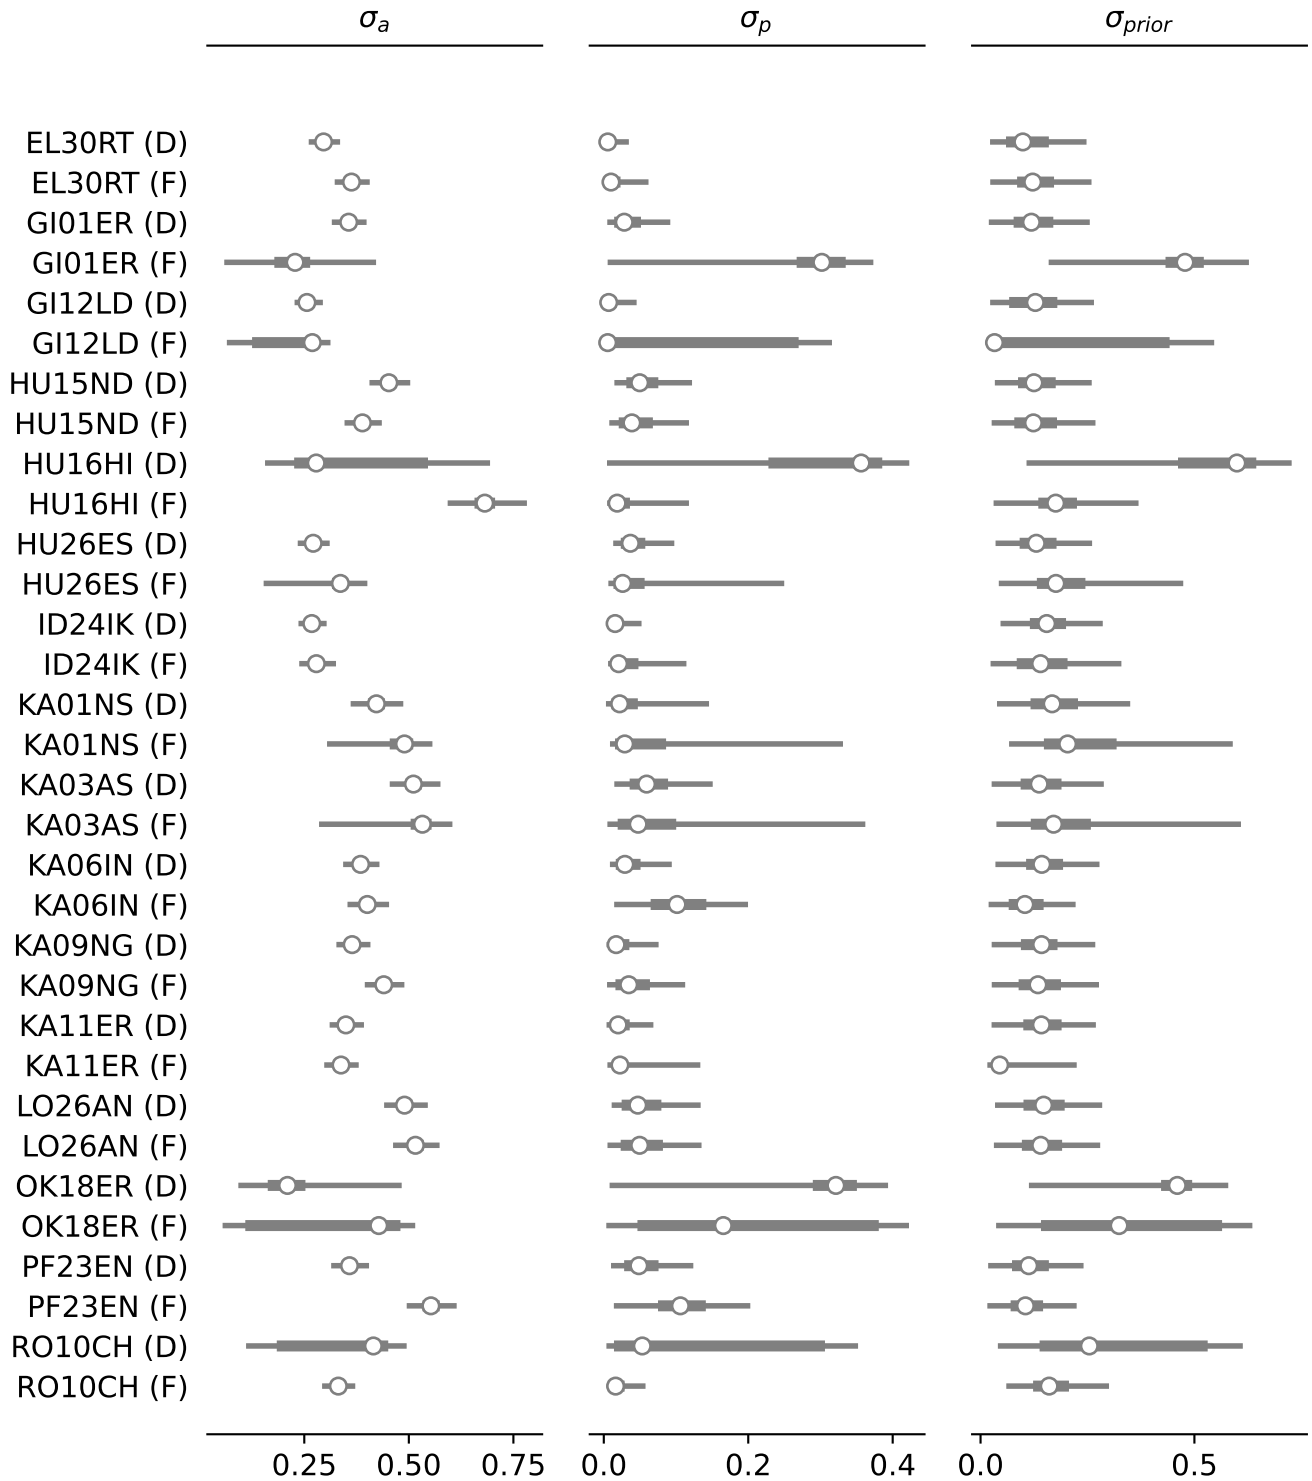


Figure A.11.: **Inference Results for Puck (PU) Data** – 95% HDI for each parameter per subject. Dots denote the inferred posterior means, and thick lines denote the area between the central quartiles.

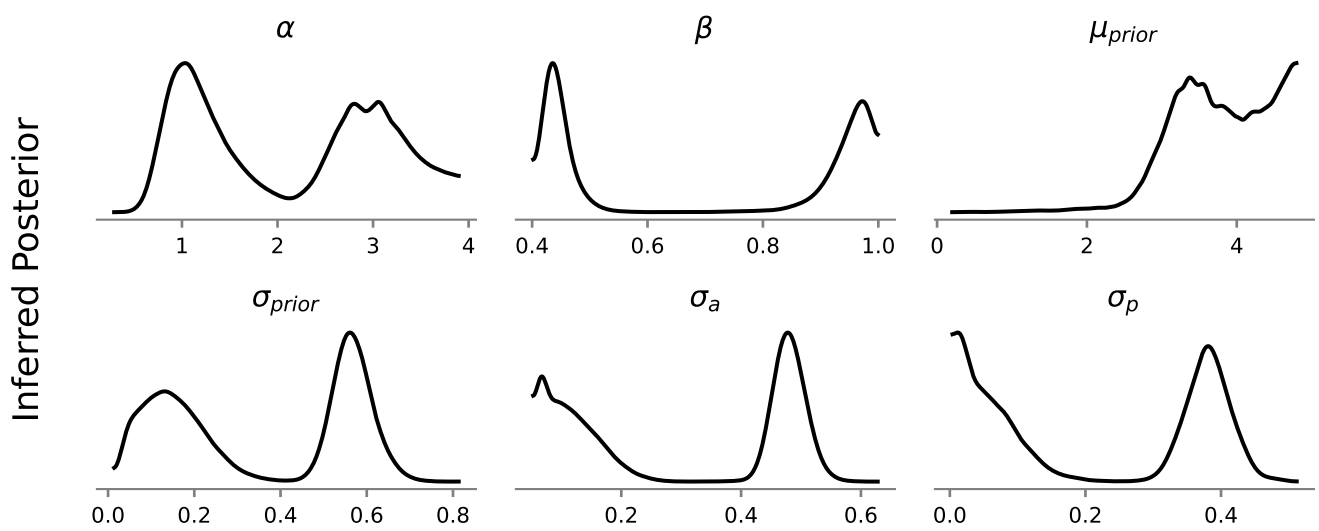


Figure A.12.: **Example of a Bimodal Posterior** – Bimodal posterior observed for subject *OK18ER* with puck *F*.