Revisiting Cost Functions in Sensorimotor Decision-Making

Tobias F. Niehues, Dominik Straub, Constantin A. Rothkopf
Centre for Cognitive Science, Technical University of Darmstadt

Abstract

Human decision-making in sensorimotor tasks is characterized by perceptual uncertainty, motor variability, prior beliefs, and the goal of a task, as well as by other factors like the effort required to act. Distributions and costs in these tasks are usually assumed to be normally distributed or of quadratic shape to maintain analytical tractability for mathematical convenience. However, there are no guarantees whether these assumptions correctly represent the functional relations underlying human behavior. Recent work on inverse decision-making makes it possible to overcome these limitations while still allowing inference of behavioral parameters from data with arbitrary cost functions and sensory encoding. Here, we extended this approach to a hierarchical model, thereby allowing model comparison at the task level instead of a per-subject level. In all data sets, asymmetric cost functions describe human behavior better than quadratic costs, and in four out of five cases, the cost function contains explicit effort costs, contrary to previous investigations.

Keywords: decision-making; perception and action; sensorimotor control; inverse modeling

Background

Bayesian actor models provide a framework for decision-making in a multitude of sensorimotor tasks. They incorporate subjects' perceptual uncertainty about the world (Kersten et al., 2004), motor variability in the subjects' actions (Van Beers et al., 2004; Trommershäuser et al., 2008), and subjective goals and costs. Common approaches do not incorporate these properties but only rely on normal distributions and quadratic costs for mathematical convenience.

Perceptual components of actor models should have signal-dependent properties in encoding sensory stimuli as captured by Weber's Law (Weber, 1831). This can be implemented in several ways (Zhou et al., 2024). Some approaches choose logarithmic mappings (Stocker & Simoncelli, 2006; Petzschner et al., 2015), some use efficient coding (Wei & Stocker, 2015), and others use a generally parameterized mapping (Acerbi et al., 2014). Human movements exhibit similar signal-dependency, as their variability scales with their magnitude (Harris & Wolpert, 1998; Todorov & Jordan, 2002).

Commonly used cost functions are quadratic (Sohn & Jazayeri, 2021), which have been shown to be insufficient to formalize goals in some tasks (Körding & Wolpert, 2004; Sims, 2015), and do not incorporate effort costs. Yet, the assumed cost function should incorporate internal, e.g., computational, costs (Lewis et al., 2014; Lieder & Griffiths, 2020; Gershman et al., 2015) and external costs like effort exerted on the body of the actor (Hoppe & Rothkopf, 2016; Straub & Rothkopf, 2022).

Inverse modeling is a general method to infer the parameters that underlie observed behavior (Aitchison et al., 2015; Rothkopf & Ballard, 2013; Kwon et al., 2020). However, models accounting for the sensorimotor and cost characteristics described above lose their analytical tractability. But the inverse problem may become solvable by amortizing the inference (Radev et al., 2020; Govindarajan et al., 2022) or the decision-making problem, which is amortized with a neural network in the model we use.

Here, we propose a hierarchical model for inverse decision-making, allowing inference and model comparison on the task instead of the subject level.

Method

Our recently proposed method (Straub et al., 2025) overcomes the above limitations and allows for efficient inference. However, it only supports analyses on a per-subject level. Here, we extended it to a hierarchical model, allowing task-level analysis of the inferred cost functions.

Decision-Making Problem

Subjects receive a sensory measurement $m \sim \text{Lognormal}(s, \sigma)$ of a physical stimulus s. They need to find the optimal action a^*, which minimizes the expected cost over a cost function $\ell_\theta(r,s)$ parameterized by cost parameters θ:

$$a^* = \arg\min_a \mathbb{E}_{p(s \mid m)} \left[\mathbb{E}_{p(r \mid a)} \left[\ell_\theta(r,s) \right] \right],$$

given a motor response $r \sim \text{Lognormal}(a^*, \sigma_r)$ and posterior belief over the stimulus $p(s \mid m) \propto p(m \mid s)p(s)$ inferred from a prior belief over $s \sim \text{Lognormal}(\mu_0, \sigma_0)$. See Fig. 1A.

Inverse Decision-Making

The researcher's goal is to infer parameters θ, σ, σ_r, μ_0, σ_0 of each subject given the data $\mathcal{D} = \{r_i, s_i\}$. We extend the previous model by adding hyperpriors Π to all inferred parameters. This allows us to perform model comparison on the task rather than the subject level (Fig. 1A). We sample from the posterior using NUTS (Hoffman et al., 2014) with 5,000 warm-up steps and 15,000 samples. We then perform model comparison using PSIS-LOO (Vehtari et al., 2024).

Data Sets

We use data from published work involving the perception of a stimulus and acting based upon this percept. The used data sets contain behavioral data from bean bag throwing (Willey & Liu, 2018), force production (Onneweer et al., 2016), time interval reproduction (Birkenbusch et al., 2015), puck sliding (Neupärtl et al., 2020) and bicycling distance reproduction tasks (Sun et al., 2004) (Fig. 1B). Because all stimuli are magnitude-like variables, they can be described by a log-normal sensory encoding. Incomplete subject data was removed.

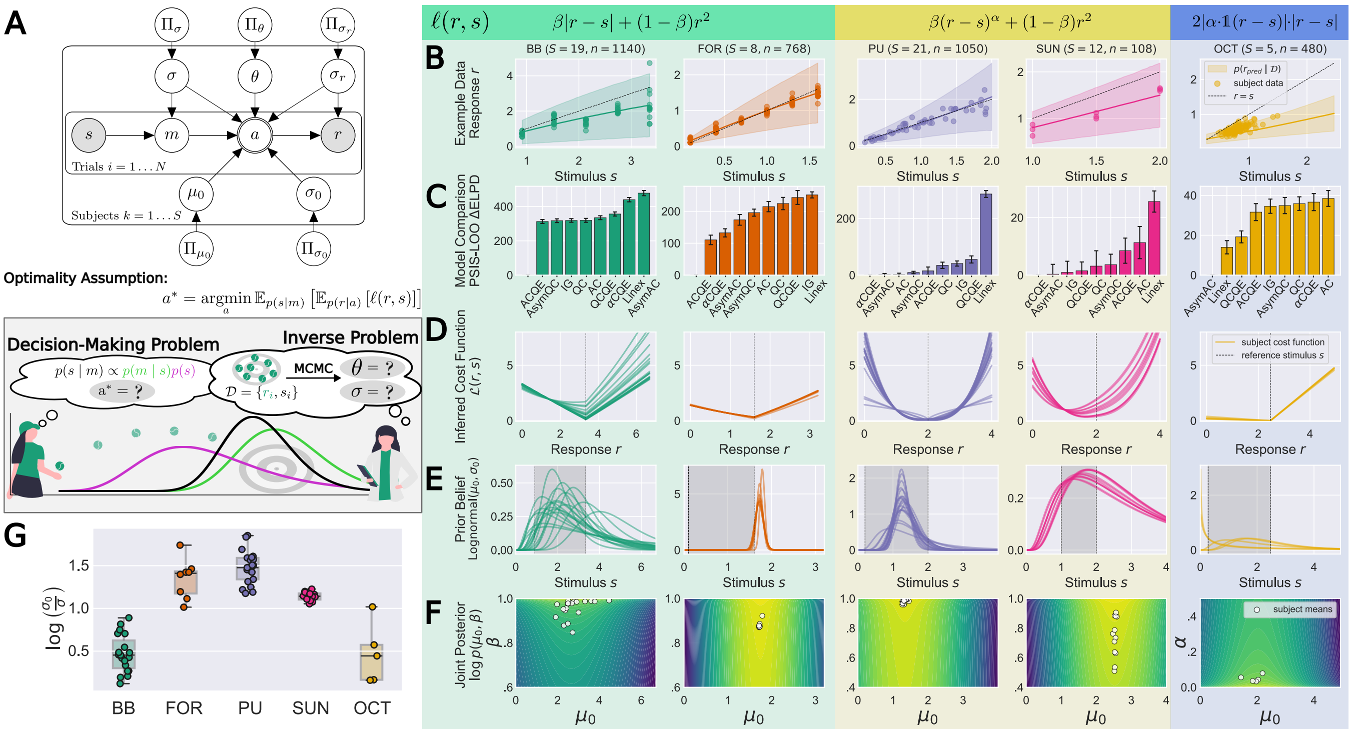

Figure 1: **A** Graphical model from the researcher's perspective and general problem. See Method section for details. **B** Example data for one subject of each task. Dots are observed stimulus-response pairs (r_i, s_i), shaded line and area show mean and 97%-CI of posterior predictive samples. **C** Model comparison over different cost functions. Lower scores indicate better fits. Best fitting cost functions are grouped by background color. **D** Best cost function fit ℓ_θ for each task. Single lines show the cost functions of a single subject. **E** Inferred beliefs $s \sim \text{Lognormal}(\mu_0, \sigma_0)$ per subject. Shaded regions denote the true stimulus ranges. **F** Log probability densities for inferred joint hyperpriors of belief μ_0 and cost parameters β and α, respectively. Yellow denotes high, dark blue denotes low probability. **G** Log ratio of prior uncertainty σ_0 and perceptual uncertainty σ for all subjects per task. Cost Legend: C=Cost, Q=Quadratic, A=Absolute, Asym=Asymmetric, IG=Inverted Gaussian, α=parameterization of the exponent.

Results

Cost Function

We compare several cost functions including different forms of asymmetry and explicit motor effort based on prior studies described above. Model comparison yields each task's best fitting cost functions (Fig. 1C). None of the tasks are best described by quadratic costs. Instead, asymmetric cost functions fit best on all tasks, with most tasks including explicit effort costs (Fig. 1D). In task PU, we find all $\beta \approx 1$ (Fig. 1F), which indicates low effort costs, but the results still highlight the necessity of non-quadratic costs functions with $(r - s)^\alpha$.

Prior Belief

The inferred prior beliefs overlap considerably with the true stimulus ranges in most tasks (Fig. 1E). Subjects systematically overestimate the stimulus distribution in task FOR and have rather uncertain beliefs in task SUN. In general, subjects rely more on sensory information than on prior beliefs, indicated by $\log(\sigma_0/\sigma) > 0$ for all tasks (Fig. 1G). Remarkably, the variability of the log ratio is smallest in task SUN, where people sensed the stimulus via proprioception in contrast to the other tasks, which used visual or auditory stimuli.

Identifiability

Generally, we find no identifiability problems. In tasks FOR and PU, group-level posteriors for β and μ_0 are wide, but low variability on an individual level (Fig. 1F). This can indicate identifiability issues, which could be resolved by adding levels of perceptual uncertainty in the experiments (Wei & Hahn, 2024; Straub et al., 2025).

Discussion

We extended an approach that inferred cost functions from single subjects to a hierarchical model to use data from all subjects in a task. Applied to a collection of sensorimotor tasks, we find that models require cost functions other than quadratic, including effort costs, and need to accommodate idiosyncratic differences to describe human behavioral data. The current model is limited by a logarithmic encoding of stimuli, which should be extended to arbitrary sensory encodings. Future model extensions could include task-level inference with different cost functions per subject. Investigations of intra-individual variability, where the same subjects solve various tasks, can support this. This data should incorporate multiple levels of perceptual uncertainty to disambiguate priors and effort costs.

Acknowledgments

We thank Chéla Willey and Zili Liu for sharing their beanbag throwing data (Willey & Liu, 2018), Bram Onneweer for sharing the force reproduction data (Onneweer et al., 2016), and Fabian Tatai for help with the puck sliding data (Neupärtl et al., 2020). This research was supported by the European Research Council (ERC; Consolidator Award "ACTOR"-project number ERC-CoG-101045783), by the "The Adaptive Mind", funded by the Excellence Program of the Hessian Ministry of Higher Education, Science, Research and Art, and the Hessian research priority program LOEWE within the project "WhiteBox". We acknowledge the computing time provided to us on the high-performance computer Lichtenberg at the NHR Centers NHR4CES at TU Darmstadt.

References

Acerbi, L., Ma, W. J., & Vijayakumar, S. (2014). A framework for testing identifiability of bayesian models of perception. *Advances in neural information processing systems*, *27*.

Aitchison, L., Bang, D., Bahrami, B., & Latham, P. E. (2015). Doubly bayesian analysis of confidence in perceptual decision-making. *PLoS computational biology*, *11*(10), e1004519.

Birkenbusch, J., Ellermeier, W., & Kattner, F. (2015). Octuplicate this interval! Axiomatic examination of the ratio properties of duration perception. *Attention, Perception, & Psychophysics*, *77*(5), 1767–1780.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.

Govindarajan, L. N., Calvert, J. S., Parker, S. R., Jung, M., Darie, R., Miranda, P., . . . Serre, T. (2022). Fast inference of spinal neuromodulation for motor control using amortized neural networks. *Journal of neural engineering*, *19*(5), 056037.

Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*(6695), 780–784.

Hoffman, M. D., Gelman, A., et al. (2014). The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, *15*(1), 1593–1623.

Hoppe, D., & Rothkopf, C. A. (2016). Learning rational temporal eye movement strategies. *Proceedings of the National Academy of Sciences*, *113*(29), 8332–8337.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as bayesian inference. *Annu. Rev. Psychol.*, *55*, 271–304.

Körding, K. P., & Wolpert, D. M. (2004). The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences*, *101*(26), 9839–9842.

Kwon, M., Daptardar, S., Schrater, P. R., & Pitkow, X. (2020). Inverse rational control with partially observable continuous nonlinear dynamics. *Advances in neural information processing systems*, *33*, 7898–7909.

Lewis, R. L., Howes, A., & Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science*, *6*(2), 279–311.

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences*, *43*, e1.

Neupärtl, N., Tatai, F., & 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.

Onneweer, B., Mugge, W., & Schouten, A. C. (2016). Force reproduction error depends on force level, whereas the position reproduction error does not. *IEEE Transactions on Haptics*, *9*(1), 54-61. doi: 10.1109/TOH.2015.2508799

Petzschner, F. H., Glasauer, S., & Stephan, K. E. (2015). A Bayesian perspective on magnitude estimation. *Trends in Cognitive Sciences*, *19*(5), 285–293. doi: 10.1016/j.tics.2015.03.002

Radev, S. T., Mertens, U. K., Voss, A., Ardizzone, L., & Köthe, U. (2020). Bayesflow: Learning complex stochastic models with invertible neural networks. *IEEE transactions on neural networks and learning systems*, *33*(4), 1452–1466.

Rothkopf, C. A., & Ballard, D. H. (2013). Modular inverse reinforcement learning for visuomotor behavior. *Biological cybernetics*, *107*, 477–490.

Sims, C. R. (2015). The cost of misremembering: Inferring the loss function in visual working memory. *Journal of vision*, *15*(3), 2–2.

Sohn, H., & Jazayeri, M. (2021). Validating model-based bayesian integration using prior–cost metamers. *Proceedings of the National Academy of Sciences*, *118*(25), e2021531118.

Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*(4), 578–585. doi: 10.1038/nn1669

Straub, D., Niehues, T. F., Peters, J., & Rothkopf, C. A. (2025). Inverse decision-making using neural amortized bayesian actors. In *The thirteenth international conference on learning representations.* Retrieved from `https://openreview.net/forum?id=zxO4WuVGns`

Straub, D., & Rothkopf, C. A. (2022). Putting perception into action with inverse optimal control for continuous psychophysics. *Elife*, *11*, e76635.

Sun, H.-J., Campos, J. L., & Chan, G. S. W. (2004). Multisensory integration in the estimation of relative path length. *Experimental Brain Research*, *154*(2), 246–254.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, *5*(11), 1226–1235.

Trommershäuser, J., Maloney, L. T., & 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., & Wolpert, D. M. (2004). The Role of Execution Noise in Movement Variability. *Journal of Neurophysiology*, *91*(2), 1050–1063.

Vehtari, A., Simpson, D., Gelman, A., Yao, Y., & Gabry, J. (2024). Pareto smoothed importance sampling. *Journal of Machine Learning Research*, *25*(72), 1–58.

Weber, E. H. (1831). *De pulsu, resorptione, auditu et tactu: Annotationes anatomicae et physiologicae*. Koehler.

Wei, X.-X., & Hahn, M. (2024). The identifiability of bayesian models of perceptual decision. *Journal of Vision*, *24*(10), 572–572.

Wei, X.-X., & Stocker, A. A. (2015). A bayesian observer model constrained by efficient coding can explain'anti-bayesian'percepts. *Nature neuroscience*, *18*(10), 1509–1517.

Willey, C. R., & Liu, Z. (2018). Long-term motor learning: Effects of varied and specific practice. *Vision research*, *152*, 10–16.

Zhou, J., Duong, L. R., & Simoncelli, E. P. (2024). A unified framework for perceived magnitude and discriminability of sensory stimuli. *Proceedings of the National Academy of Sciences*, *121*(25), e2312293121. doi: 10.1073/pnas.2312293121