

Task-set Selection in Probabilistic Environments: a Model of Task-set Inference

Ian W. Eisenberg (ieisenbe@stanford.edu)

Department of Psychology, Stanford University

Russell A. Poldrack (russpold@stanford.edu)

Department of Psychology, Stanford University

Abstract

To act effectively in a complicated, uncertain world, people often rely on task-sets (TSs) that define action policies over a range of stimuli. Effectively selecting amongst TSs requires assessing their individual utility given the current world state. However, the world state is, in general, latent, stochastic, and time-varying, making TS selection a difficult inference for the agent. An open question is how observable environmental factors influence an actor's assessment of the world state and thus the selection of TSs. In this work, we designed a novel task in which probabilistic cues predict one of two TSs on a trial-by-trial basis. With this task, we investigate how people integrate multiple sources of probabilistic information in the service of TS selection. We show that when action feedback is unavailable, TS selection can be modeled as “biased Bayesian inference”, such that individuals participants differentially weight immediate cues over TS priors when inferring the latent world state. Additionally, using the model's trial-by-trial posteriors over TSs, we calculate a measure of decision confidence and show that it inversely relates to reaction times. This work supports the hierarchical organization of decision-making by demonstrating that probabilistic evidence can be integrated in the service of higher-order decisions over TSs, subsequently simplifying lower-order action selection.

Keywords: task-sets; structure learning; Bayesian cognition; model-based; decision making

Introduction

Humans face the daily challenge of making decisions in an uncertain and open-ended world. In such a world, caching independent stimulus-response mappings is impractically slow and fails to capitalize on the structure inherent in natural tasks. Many actions learned at one time may generalize to new environments: our experiences dealing with petulant adults may help us placate ill-tempered children or vice versa, expertise using our own computer seamlessly translates to computer expertise in general, and so on. In general, the organized structure of the world allows agents to fruitfully group learned actions into higher-order action policies which can be retrieved to avoid redundant learning. Such action policies are often referred to as task-sets (TSs), and their use potentially eases the learning problem by providing flexible representations that can be leveraged across environments.

Much of the computational work on TSs relates to model-based decision making (Daw et al., 2012; Solway & Botvinick, 2012). In this framework, a person learns a goal-sensitive action policy based on an internal model of the

world, which consists of a set of states, transition probabilities, and actions. Through exploration and feedback, the agent gradually develops an action policy which determines their behavior. While this general concept has been informative in outlining how TSs may be learned given a model of the world, model-based decision making has largely ignored how the agent's internal models are developed and selected when multiple models may apply in a particular environment (so-called “structure learning”).

Some research has directly addressed the problem of structure learning, proposing that agents simultaneously infer the latent causal structure of the world while identifying the appropriate TS given that inferred structure (Gershman & Niv, 2010; Redish et al, 2007). From this perspective, structure learning is intimately tied with stimulus-response learning, leading to the compelling prediction that people will reuse TSs whenever they infer that the latent structure of the world conforms to the structure in which the TS was first learned.

Inferring the latent world state is closely related to categorization (Gershman et al. 2010), the cognitive process by which people use an organizational framework to assign discrete instances (objects, events, emotions, etc.) to groups that are functionally or perceptually equivalent on some level of abstraction (Anderson, 1991; Shafto et al. 2011). The central idea is that latent categories stochastically generate observable features conforming to some characteristics distribution. If people represent, on some level, a generative model of the environment that constitutes a hypothesis space over possible categories, then they can infer the underlying category given uncertain evidence (Fei-Fei et al., 2007; Tenenbaum et al., 2006). Moreover, they can categorize novel observations by appealing to these generative models. For decision-making, useful categorical boundaries are defined by states which call for different action policies. To capture TS selection in such a scenario, we use a task where the agent knows that multiple task-relevant states exist, such that establishing the latent world state is equivalent to establishing the best TS. If the agent can uncover the structure underlying state transitions, they can greatly simplify the task and improve their performance.

Empirical and computational support for probabilistic inference over TSs comes from work by Collins and colleagues, who have shown that people reuse TSs in an approximately optimal way based on contextual support (Collins & Koechlin, 2012; Collins & Frank, 2013). Collins & Koechlin have put forward a model where a small number of TSs are held in a working memory-like cache

where they are evaluated to assess their individual “reliabilities.” TSs are selected when they prove reliable, and otherwise discarded, replaced by new TS propositions constructed combinatorially from TSs held in long term memory. In a similar vein, Frank & Badre (2012) propose a Bayesian “mixture of experts” model of TS selection (see Doya et al., 2012 for a similar idea). In this framework, multiple competing TS hypotheses govern people’s behavior as they search for higher-order rules in a hierarchically structured decision making task (Badre et al. 2010).

These experiments used deterministic cues to indicate the appropriate TS, a simplification that potentially obscures the factors underlying TS inference in general. In this project, we aimed to resolve these problems and clarify the process underlying TS inference. We introduce a novel task-switching paradigm that required participants to reason over probabilistic environmental cues to select the appropriate TS on a trial-by-trial basis. With this paradigm, we anticipated that participants would use multiple sources of information when selecting TSs such that their decisions related both to contextual cues and TS transition probabilities. As these different sources contribute to behavior in subtle ways, we develop an explicit quantitative model to assess the information participants access to infer TSs. We hypothesize that while there will be substantial individual differences in how people integrate information, TS inference can be characterized by Bayesian inference with minimal free parameters reflecting individual information-processing biases.

Method

Task Description

49 participants completed the Probabilistic Context Task, a task-switching experiment (Figure 1) composed of two phases: training (832 trials: 45 min) and testing (800 trials: 30 min) On each trial, participants were required to select one of four keys in response to two-dimensional stimuli varying in color (red or blue) and shape (circle or square). Each key was mapped to one of these feature (e.g. key 1 for blue stimuli, key 2 for circles), which were randomized across participants. Participants had 1.5 s to respond. The correct response was determined by a latent TS that established the relevant feature, which changed from trial-to-trial. There were two TSs: the shape TS (STS) and the color TS (CTS). Correct responses conformed to both the stimulus and the TS (pressing the red key for a red circle while the CTS was operating). Correct responses earned a point which was presented for .5 s during the training phase followed by a variable intertrial interval. During the testing phase participants received no feedback. Overall, each training trial lasted 3-3.5 s, and each test trial lasted 2-2.5 s.

While there were no deterministic cues indicating the current TS, the task was designed to allow inference of the trial-by-trial TSs using probabilistic information. TSs switched probabilistically on each trial such that $P(TS_{t+1} = TS_t)$, the probability of the TS remaining the same from trial-to trial, was 90%, referred to for the remainder of the

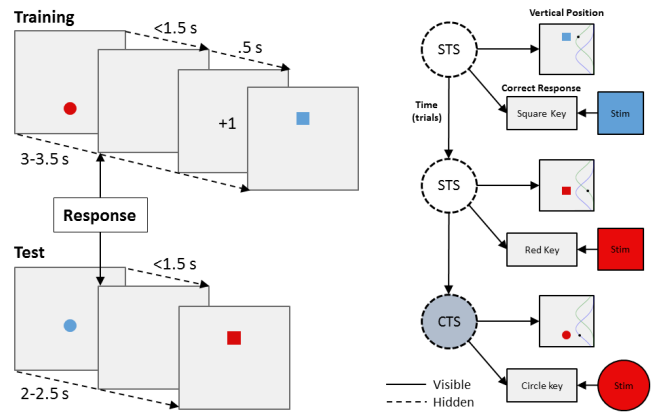


Figure 1: (Left) On each trial, a shape appears at one of 12 vertical positions. The participant responds with one of four keys corresponding to the two features of the stimuli (red, blue, circle, square). During training they get deterministic feedback after they respond. (Right) Schematic of the latent trial structure. Stimulus vertical position is drawn one of two distributions (shown in green and blue) corresponding to the current TS (STS: green, CTS: blue). The current TS has a 10% chance of changing from trial-to-trial, otherwise it remains the same. The stimuli are randomly drawn on each trial and are observable while the current TS is latent. Thus to correctly respond to the stimulus the participant must infer the current TS based on the stimulus position and task history.

paper as the recursion probability. Additionally, on each trial the stimulus’s vertical position on the screen was drawn from a truncated Gaussian distribution (limits 1 and -1, corresponding to the top and bottom of the screen, respectively) parameterized by the current TS. Stimulus position Gaussians had the same standard deviation ($\sigma=.37$), but had different means ($\mu = .3$ or $-.3$) depending on the TS. These Gaussians were discretized into 12 bins spanning the screen. The TS that was primarily associated with the top of the screen was counterbalanced across participants. For simplicity, for the remainder of the paper we will assume the STS was primarily associated with the top of the screen.

Before training participants were explicitly told about both color and shape TSs and were given an opportunity to practice using a separate set of practice stimuli and key mappings. When training started, participants knew that one key would correspond to each stimulus feature (red, blue, circle square), and only one TS operated on each trial, but did not know what the response mappings were or how the TSs switched. They were told that their goal was to learn on which trials they should respond based on color or shape. They were also told that they should use the feedback to learn during training, but not to rely on it, as it would be removed during the test. Participants were encouraged to respond as quickly and accurately as possible. Participants also knew that their performance during training and test determined their bonus payment, which could range from

\$0-\$5. A post-task questionnaire probed participant’s explicit understanding of the task, including their estimates of TS transition probabilities.

In summary, during the training phase, participants received deterministic feedback which they could use to learn the mapping between stimuli features and response keys as well as the determinants underlying TS switches. Because feedback was omitted during the subsequent test phase, participants had to respond based on their understanding of the TSs’ relationship with the probabilistic cues in the environment (stimulus position and the previous operating TS), or based on irrelevant factors fabricated by the participant (e.g. deterministic switches every 5 trials, switch after a red square). Above chance performance during the test phase would indicate that the participant had internalized some true aspects of the task structure.

Behavioral Analysis

Participants responses were assigned to either the CTS or STS based on whether one of the two colors keys or shapes keys was pressed, respectively. Because we were interested in TS selection, it was necessary that participants knew which keys corresponded to which features by the beginning of the test phase. To ensure this, we coded each response as either conforming or not conforming to one of the two dimensions of the stimulus and excluded participants whose average stimulus conformance fell below 75% during the test phase. This exclusion criterion ensured that all analyzed participants knew the two appropriate responses for each stimulus (e.g. either the red or circle keys for a red circle)

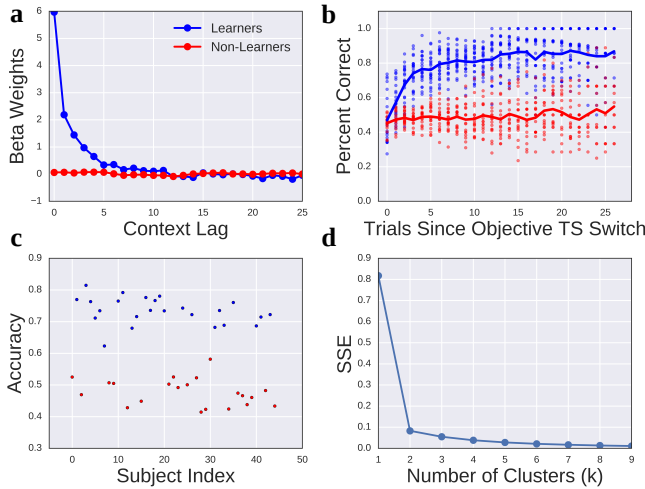


Figure 2: Summary of learners (blue) and non-learners (red). (a) Output of regression predicting participant choice by current context and context history. (b) Participant accuracy as a function of trials since objective (latent) TS switch. Each point is an individual participant’s accuracy at that delay. (c) Clustering of participants using k-means on participants accuracy with two centroids initialized at .49 and .51. (d) Sum of squared errors for different k values with random initialization averaged over 1000 iterations.

and only had to determine the operating TS to successfully respond. Of our 49 participants, 4 were excluded based on this criterion. All remaining participants conformed to the stimulus >90% of the time. For the remaining participants, we collapsed their responses from the four choices to the two TSs resulting in a binary choice vector across trials.

Our analysis was principally concerned with how people weighed the probabilistic information potentially relevant to TS selection. Therefore it was necessary that behavior related in some way to task variables. We used k-means clustering to divide the participants based on overall accuracy, resulting in a clear separation between two groups of participants: 24 “learners” and 21 “non-learners” (Figure 2c). While we presume that there is structure in all participant’s behavior, this paper is solely exploring the correspondence between task structure and behavior, rather than a complete evaluation of participant behavior. Modeling work was therefore restricted to the learners.

Prior to modeling we fit mixed-effects logistic regressions to both groups to assess the impact of context, context history and prior choice on TS selection.

Computational Modeling Optimal TS inference during test can be formalized as Bayesian inference over probabilistic cues and task history. The optimal prior over TSs on trial_t is proportional to the posterior over TSs after trial_{t-1}. Specifically, the prior is the product of the transition probability matrix between TSs and the posterior vector over TSs. In other words, if the person was absolutely confident in the TS on trial_{t-1}, then the prior conforms to the transition probabilities associated with that TS; if the person was completely unsure on trial_{t-1}, the prior over TSs is uniform. This prior information is then combined with the stimulus position’s likelihood under each TS’s positional distribution to arrive at a posterior over TSs on each trial.

Our main hypothesis is that most participants will integrate both transition probabilities and positional distributions to select TSs, but individuals may be biased in their weighting such that their choices unequally favor the probabilistic cue. This is a soft form of base-rate neglect, where the prior is down-weighted in favor of the likelihood. Our model (the *bias-2* observer, below) instantiates this idea by fitting two variables, r_1 and r_2 , which together define a participant’s transition probability matrix.

$$P(TS)_t = \frac{\begin{pmatrix} r_1 & (1-r_2) \\ (1-r_1) & r_2 \end{pmatrix} \cdot \begin{pmatrix} P(TS_1)_{t-1} \\ P(TS_2)_{t-1} \end{pmatrix} * \begin{pmatrix} P(context_t|TS_1) \\ P(context_t|TS_2) \end{pmatrix}}{N}$$

We can define a number of cases corresponding to different inference strategies: if r_1 and r_2 equal .9 (the true recursive probability) the participant is Bayes optimal; if r_1 and r_2 are less than .9 the participant overweights the probabilistic cue (with $r_1 = r_2 = .5$ being the special “base-rate neglect” case); if r_1 and r_2 are greater than .9 the participant overestimates the transition probabilities. In

reality, the transition matrix should be symmetric, but we estimate the transitions associated with each TS separately to allow for participant bias such that they prefer to choose one TS over the other without any evidence.

Simpler models are possible by fixing some of the parameters. We contrast the *bias-2* observer to three related models: one where the transition matrix is forced to be symmetric (*bias-1*), an optimal observer where the transition matrix is fixed based on the training run (*optimal*), and a base-rate neglect model with a fixed transition matrix as defined above (*base-rate neglect*).

The likelihood on each trial was calculated based on TS positional distributed, which were defined by the mean and standard deviation for each TS observed during the training phase. While this assumption is optimistic in regards to the task statistics participants encoded during training, as long as the participant's estimation errors are not systematically biased away from the true statistics, the models should reflect participant performance in the aggregate.

Each model resulted in a vector of trial-by-trial TS posteriors for the testing phase of each participant. While an optimal decision maker would select the most likely TS, we assume some noise in translating posterior estimates into action. Thus we fit an ϵ parameter to each model which reflects the probability of randomly choosing a TS. This led to a final vector of trial-by-trial TS choice likelihoods, which were used to fit each model.

Individual participant parameter estimates were fit using python's `lmfit` module's L-BFGS method, with the cost equal to the $-\log$ likelihood of that participant's TS selections. Model selection was accomplished using Bayesian information criterion (BIC: Schwartz, 1978), and by fitting the models on either the first or last half of the data and testing on the left out half. When discrete TS choices were needed, they were defined by the maximum likelihood on each trial across the posterior.

We were also interested in whether more difficult decisions were related to reaction times (RT), as predicted by a number of studies relating choice confidence and RT (e.g. Henmon, 1911; Roitman & Shadlen, 2002). We defined a trial-by-trial estimate of model choice confidence based on the average distance from .5 across the TS posteriors, ranging from 0 (indifference) to 1 (certainty). Because there were only two possible TSs, this is equivalent to calculating the distance from the choice boundary between the two TSs. We assessed this relationship with a mixed-effects linear regression, using the `lme4` package.

Results

Context and aspects of context history significantly predicted TS choice for learners ($p < .001$), but not for nonlearners (Figure 2a). In addition, prior choice was significantly predictive in both groups. When included in the same model, prior choice abolished the effect of context history on participant choice in the learner group.

Model comparison across the population showed that the *bias-2* observer was a significantly better fit than any of the

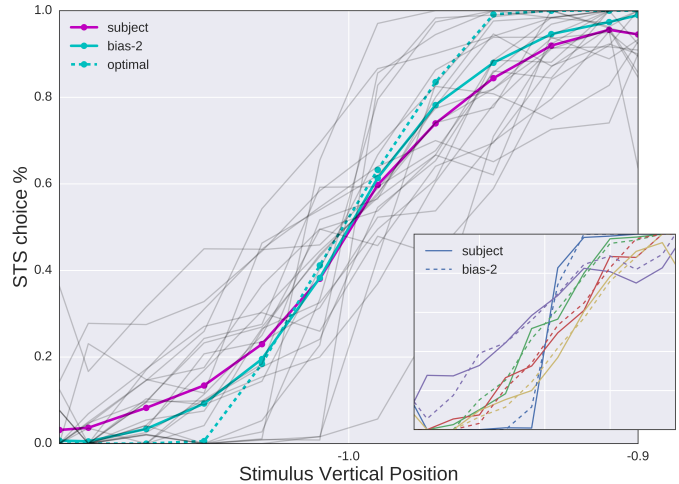


Figure 3: Percentage of trials for each of the 12 vertical position bins where responses reflected STS selection. The stimulus was never shown exactly at the midline. The purple line shows the average percentage chosen across all participants. The teal lines show the *bias-2* and optimal model performance. Though not shown the *bias-1* lies between these two curves. Individual participants curves are shown in light gray. (Inset) 5 example participant fits.

comparison models (Table 1). Moreover, individual analysis showed that both the *bias-2* and *bias-1* observers fit better than the base-rate neglect or optimal models for participants. Converting the *bias-2* posteriors into TS choices, we found that model choices matched individual participant's choices well ($\mu=87.6\%$, $\sigma=4.5$). Competing models were also relatively successful at capturing participant choices (*bias-1*: $\mu=86.1\%$, $\sigma=5.0$; *optimal*: $\mu=82.7\%$, $\sigma=5.3\%$). Each model's likelihood for individual participants is shown in Figure 2.

Parameter estimates showed no systematic preference for one TS over another as measured by the parameter estimates of r_1 ($\mu=88.9\%$, $\sigma=13.8$) and r_2 ($\mu=86.8\%$, $\sigma=16.9$). Overall, the population average transition matrix is quite similar to the true recursion probability of 90%, though there is a slight population-wide bias to overvalue the stimulus's vertical position. However, while this population-wide estimate is close to the true statistics (and therefore close to optimality), there is incredible variability across participants indicating that the population summary may mask consistent inferential biases that are distributed around an optimal strategy. While we interpret these differences as relation to biases in the TS inference process, it is possible that they instead stem from individual differences in encoding the environmental structure, which would affect TS inference. To address this we looked at the participant-reported estimates of the task statistics during a post-task questionnaire. These estimates were less accurate (STS: $\mu=68.2\%$, $\sigma=21.5$; CTS: $\mu=70.6\%$, $\sigma=19.8$) than the parameter estimates and were not significantly related to the *bias-2* observer parameter estimates (CTS: $r = -.13(24)$, $p = .53$; STS: $r = .25(24)$, $p = .23$).

Table 1: Model *BIC* across participants, $n = 18,906$

Model	BIC
Bias-2	11905
Bias-1	12841
Optimal	14470
Base-Rate Neglect	18167

To visualize the model fits, we calculated the proportion of times participants selected the STS at each vertical position (Figure 3). All models predict that the STS should be chosen more frequently for higher contextual values. Also shown are individual traces (in gray) highlighting the large heterogeneity in individual performance as well as example individual bias-2 fits (inset), which demonstrate the flexibility of the model to capture these large individual differences.

Finally, decision confidence as estimated by *bias-2* was inversely related to RT ($\beta = -.35(.01)$, $t = -25.50$). The regression predicts that when choice confidence equaled 1, participants responded 254 ms faster than when it equaled 0. Random effects analysis showed that this trend was true for all but one subject. Five representative participants are shown in Figure 4.



Figure 4: Two sample participant reaction times plotted against the *bias-2* model confidence. 0 indicates that both CTS and STS had a posterior probability of .5, while 1 indicates that either CTS or STS had a posterior probability of 1. Individual regressions are also shown.

Discussion

Using a novel task-switching task, we investigated how people integrate probabilistic evidence in the service of task-set (TS) selection. We found that of the people classified as “learners”, most based their decisions on both the probabilistic cues and transition probabilities, consistent with their internalizing the latent structure of the environment. On the population level, participants appeared to correctly identify the true statistics of the environment, giving the impression that they behaved in accordance with optimal Bayesian inference. However, individual participants differed greatly in their weighting of the two sources of information, such that some overvalued the probabilistic cue when making their choice.

The importance of this distinction is particularly clear when predicting RT from model estimates of trial-by-trial choice confidence. As choice confidence is a continuous

metric, it is particularly sensitive to specific trial sequences, as well as parameter estimates. During test, each trial’s TS is estimated based on the encoded transition probabilities, the posteriors over TSs on the previous trial, and the probabilistic cue. Thus it is imperative to have an individual-specific estimate of the encoded transition probabilities to analyze trial-by-trial performance. In this task, the estimate of model confidence inversely related to RT. This parameter was defined by the absolute distance from the choice boundary between the two TSs, suggesting that this distance may relate to the speed of evidence accumulation in a way analogous to perceptual decision tasks (Roitman & Shadlen, 2012). Evidence accumulation in higher-level decision making has been suggested before by Shadlen & Kiani (2013), where they forward the idea that accumulators may serve as a general algorithmic description of many cognitive computations. The relationship between RT and choice confidence would support this description.

A related idea is that RT relates to decision conflict. Difficult decisions are, by definition, closer to the choice boundary indicating that the evidence does not clearly favor a particular action. On a neural level this conflict may stem from the concurrent representation of multiple TSs which must compete in a winner-take-all fashion before gating lower-level actions (Collins & Frank, 2013). If this competition is probabilistically resolved in proportion to each TSs representational strength, this idea is just a restatement of evidence accumulation for mutually exclusive alternatives.

The best fit model had two free parameters, which together represent a bias towards the STS or CTS (reflected in an asymmetrical transition matrix) and the encoded recursion probability. Differences in the recursion probabilities may either reflect individual differences in encoding of the task statistics or biased weighting during decision-making. For instance, if participants encoded the true transition probabilities, but only attended to the stimulus position when making a choice, the model would estimate an “encoded” recursion probability of .5. While it is impossible to completely disentangle these two alternatives, the lack of correspondence between the parameter estimates and the participant estimates on the post-task questionnaire suggests a decision bias, rather than an encoding bias. However, due to the possibility that encoded task statistics are not directly available to semantic retrieval during the questionnaire, we cannot rule out either possibility.

Regardless of whether variability is linked to encoding or the decision process, an obvious question emerges: what underlies these individual differences? Participants undoubtedly arrived at the experiment with different prior expectation for the kinds of rules that may be operating within a psychology experiment. While we attempted to normalize their expectations by orienting them to the TSs of interest (shape or color), the prior expectations for higher order rules may have prevented some people from appropriately integrating certain information. This may

partially explain why some people were unable to learn any rule at all - their prior beliefs constrained the search space, preventing the encoding of the relevant variables.

Similarly, early identification of a particular pattern may have stifled later learning - a type of confirmation bias that may have attentional roots. Participants who identified the relationship between TS and vertical position may have been less motivated to search for more complicated relationships. While we expect that the relationship between transition probabilities and TS selection relates more to unconscious statistical reasoning than explicit rules, it may be that explicit adherence to a particular rule overwhelmed other potential factors. In addition, lower level processes like perseverance may compete with these cognitive strategies, as suggested by the significant relationship between prior choice and TS choice in the non-learner group. Further work exploring their effects may refine our understanding of TS selection and allow us to account for the behavior of the substantial portion of non-learners.

In this work we compared model choices to participants with the simple maximum likelihood linking function. Our success in fitting participants without relying on a softmax rule indicates that this decision behavior may deviate from the probability matching widely reported in the decision-making literature (Erev & Barron, 2005). From the perspective of hierarchical reinforcement learning, there is no particular reason to believe that a single decision rule underlies decision-making at various levels of abstraction. It is possible that TSs are selected by a qualitatively different process than lower-level action selection, as proposed by Collins & Koehlin (2012). One alternative hypothesis is that higher-order action constructs like TSs are simply less noisy than lower-order decisions. Conflict resolution would consistently favor the stronger (more supported) TS, leading to the appearance of maximization behavior without appealing to fundamentally different computations. Future work could pursue this hypothesis by selectively degrading the observable evidence that contributes to TS selection.

Conclusion

We have shown that people can successfully leverage probabilistic information to infer a decision-relevant world state. While the group results seem to indicate that people act in accordance with Bayesian optimality, individual analysis reveals large heterogeneity in the inference strategy. The bias-2 model was able to capture much of this variation, suggesting that TS inference can be viewed as a biased integration over multiple information sources.

References

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3), 409–429.
Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, 66(2), 315–326.

Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: creating, clustering, and generalizing TS structure. *Psychological Review*, 120(1), 190–229.
Collins, A., & Koehlin, E. (2012). Reasoning, Learning, and Creativity: Frontal Lobe Function and Human Decision-Making, 10(3).
Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Raymond, J. (2012). Model-based influences on humans' choices and striatal prediction errors, 69(6), 1204–1215.
Doya, K., Samejima, K., Katagiri, K., & Kawato, M. (2002). Multiple Model-Based Reinforcement Learning. *Neural Computation*, 14, 1347–1369.
Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112(4), 912–931.
Fei-Fei, L., Fergus, R., & Perona, P. (2007). Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, 106(1), 59–70.
Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: computational analysis. *Cerebral Cortex (New York, N.Y. : 1991)*, 22(3), 509–26.
Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, 117(1), 197–209.
Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*, 20(2), 251–6.
Henmon, V. a. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review*, 18(3), 186–201.
Redish, a D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychological Review*, 114(3), 784–805.
Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 22(21), 9475–9489.
Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, 80(3), 791–806.
Shafto, P., Kemp, C., Mansinghka, V., & Tenenbaum, J. B. (2011). A probabilistic model of cross-categorization. *Cognition*, 120(1), 1–25.
Solway, A., & Botvinick, M. M. (2012). Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. *Psychological Review*, 119(1), 120–54.
Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10(7), 309–318.