SUPPLEMENTARY MATERIALS: Toward a Common Representational Framework for Adaptation

Brandon M. Turner[a,*]

[a] *The Ohio State University*

Abstract

We develop a computational model – the adaptive representation model (ARM) – for relating two classic theories of learning dynamics: instance and strength theory. Within the model, we show how the principles of instance and strength theories can be instantiated, so that the validity of their assumptions can be tested against experimental data. We show how under some conditions, models embodying instance representations can be considered a special case of a strength-based representation. We discuss a number of mechanisms for producing adaptive behaviors in dynamic environments, and detail how they may be instantiated within ARM. To evaluate the relative strengths of the proposed mechanisms, we construct a suite of 10 model variants, and fit them to single-trial choice response time data from three experiments. The first experiment involves dynamic shifts in the frequency of category exposure, the second experiment involves shifts in the means of the category distributions, and the third experiment involves shifts in both the mean and variance of the category distributions. We evaluate model performance by assessing model fit, penalized for complexity, at both the individual and aggregate levels. Our general conclusion is that the mechanisms of prediction error and lateral inhibition are strong contributors to

*Corresponding Author

Email address: turner.826@gmail.com (Brandon M. Turner)

This research was supported by National Science Foundation grant SMA-1533500. We would like to thank Per Sederberg and Trisha Van Zandt for generously sharing their lab resources, and Giwon Bahg, Scott Brown, Marc Howard, Ashley Liu, Tony Marley, Jay Myung, James Palestro, Vladimir Sloutsky, and Trisha Van Zandt for insightful comments that improved an earlier version of this manuscript.

the successes of the model variants considered here. Our results suggest that the joint distribution of choice and response time can be thought of as an emergent property of an evolving representation mapping stimulus features to their appropriate response assignment.

Keywords: adaptation, learning, categorization, dynamic stimuli, cognitive modeling

---

In these Supplementary Materials, we report four additional summaries from the main text. First, we report a brief version of a parameter recovery study we performed to verify whether or not all model parameters were identifiable. For this result, we only report the recovery of the SGP model, but other model recoveries were similarly successful. Second, we report the qualitative model fits obtained from Experiment 1 for all models in the suite of ARM variants investigated in the main text. Third, we report similar qualitative model fits for Experiment 2. Fourth, we report qualitative model fits for Experiment 3.

1. Recovery of Model Parameters

To first test the adequacy of the model-fitting procedure and the identifiability of the model, we assessed whether or not we could simultaneously estimate all of the model parameters accurately. Although we tested all eight model variants reported in the main text, for sake of brevity, we report only the results from the SGL model here, as this model provided a good fit to the data across all three experiments, and was the most flexible in terms of the number of parameters.

We began by generating hypothetical data from the SGL model in the context of the "mean shift" environment of Experiment 2. Recall that this experiment involved shifting the means of the two stimulus generating distributions for each category. The shift occurred in blocks occurring at every 100 trials, where Blocks 1, 3, and 5 had means of 60 and 80, Block 2 had means of 50 and 70, and Block 4 had means of 70 and 90. Each category was equally likely to be sampled, and shared a standard deviation setting of 10.

To generate the data, we simulated the model using the following parameter setting for one random experience of the stimulus environment: $\alpha = 0.8$, $\lambda = 0.99$, $\delta = 0.35$, $\beta = 0.1$, $\theta = 3$, $I_0 = 0.5$, and $\tau = 0.05$. Within the racing diffusion process, we set $\sigma = 1$ for both accumulators to satisfy mathematical scaling properties of the model. We chose these values because they typically produced data that roughly corresponded to data from Experiment 2. It is worth noting that these values constituted global settings of the model that were used as initial values in all fitting routines when these parameters were freely estimated. Otherwise, they were set to specific values depending on the specific model variant (e.g., $\beta = 0$ under the SG model).

Figure 1 shows the simulated data from the SGL model for response time (top) and accuracy (bottom). Within each panel, the time series data has

been smoothed by applying a simple moving average with a window width of 10 trials. Each block is represented as a different color code to reflect that the distributions of features are either unique (i.e., Blocks 1, 2, and 4) or equivalent (i.e., Blocks 3 and 5 are the same as Block 1). The figure shows that these data are approximately consistent with the aggregated data from Experiment 2.

To fit the SGL model to the simulated data, we used the same routine reported in the main text. Specifically, we use a Differential Evolution Markov chain Monte Carlo (DE-MCMC; Turner et al., 2013) sampler with 32 chains, and ran the sampler for 4,000 iterations after a burnin period of 1,000 samples. During the first 1,000 samples (i.e., the burnin period), we used a migration probability of 0.1 to reduce the possibility of outlier chains. The scaling parameter γ was uniformly sampled from the interval $[0.5, 1.0]$ for each posterior sample, and the random noise term b was set to 0.001.

Figure 2 shows trace plots of the sampling algorithm for each of the 32 chains on the last 4,000 iterations. Each panel shows the sampler for a given model parameter, but the bottom right panel shows the the posterior density on the logarithmic scale. All trace plots are consistent in that they appear to have converged to a stationary target distribution, an impression that is corroborated by the posterior density trace plot (i.e., the bottom right panel).

Although Figure 2 shows that the sampling algorithm has converged to a target distribution, it does not suggest that the target distribution is accurate. Figure 3 shows the estimated marginal posterior distributions on the diagonal entries, and the estimated pairwise joint posterior distributions in the lower triangle. The upper triangle reports the estimated correlation for each pairwise joint posterior distribution. The true values of the model parameters used to generate the data are shown in the diagonal entries as the vertical red line, and in the lower triangle as the black "X" symbol. Comparing both the marginal and joint posterior distributions, we see that there is good correspondence between the central tendency of the estimated posteriors and the true parameters values used to generate the data. While it is impossible to say that the correct posterior has been obtained, the corresponds between these metrics suggests that the parameters can be accurately recovered when the true model is known.

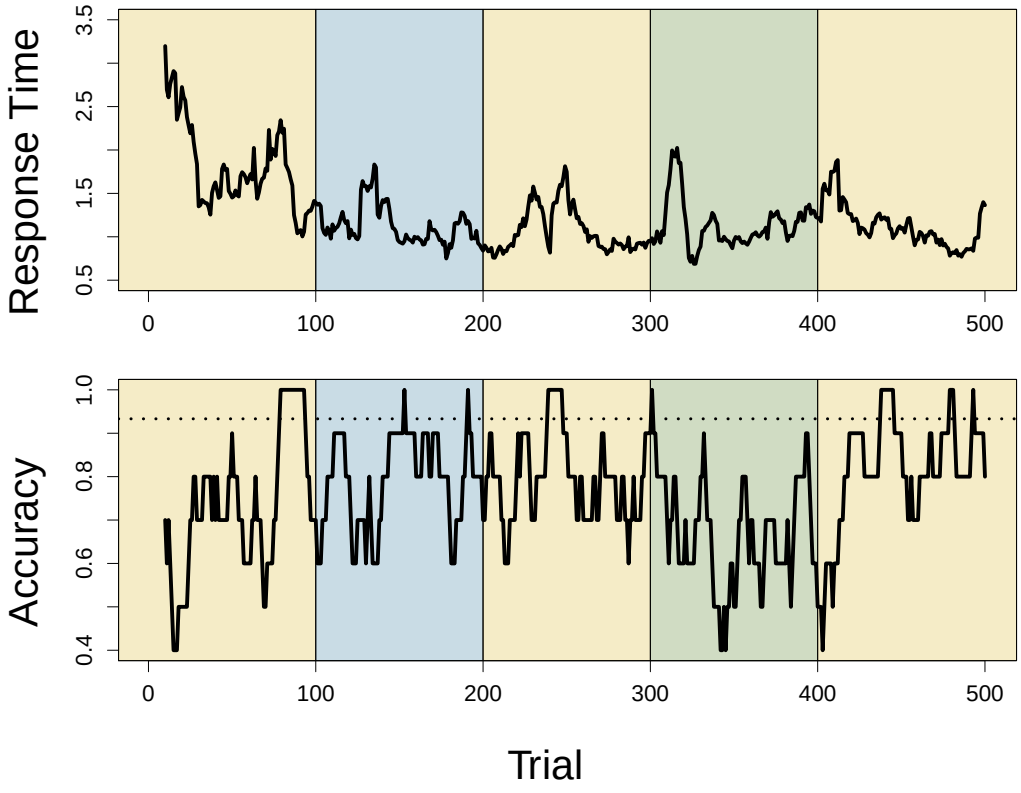

Figure 1: **Hypothetical Data Generated by the SGL Model**. Each panel shows simulated data from the SGL model variant for response times (top panel) and accuracy (bottom panel). Each block consisted of 100 trials and a different context for the means of the stimulus generating distributions, illustrated with different color codes for each new context. All data have been smoothed using a moving average of width 10.

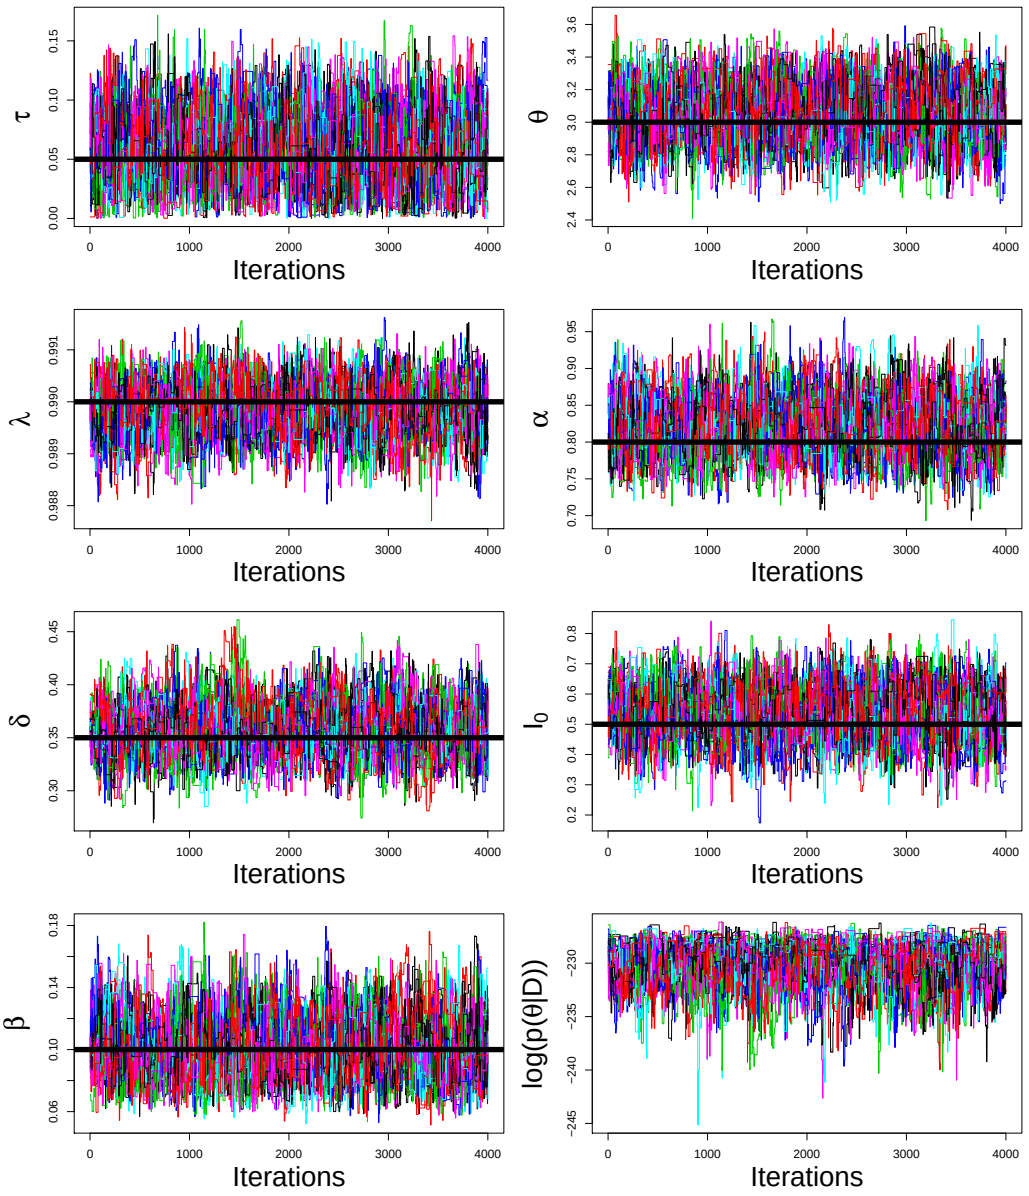

Figure 2: **Trace Plots of the Sampling Algorithm**. Each panel shows trace plots of 32 chains run for 4,000 iterations for each of the 7 parameters in the SGL model. The bottom right panel shows the log of the posterior density corresponding to each parameter proposal. The black lines represent the true values of the parameter value used to generate the data.

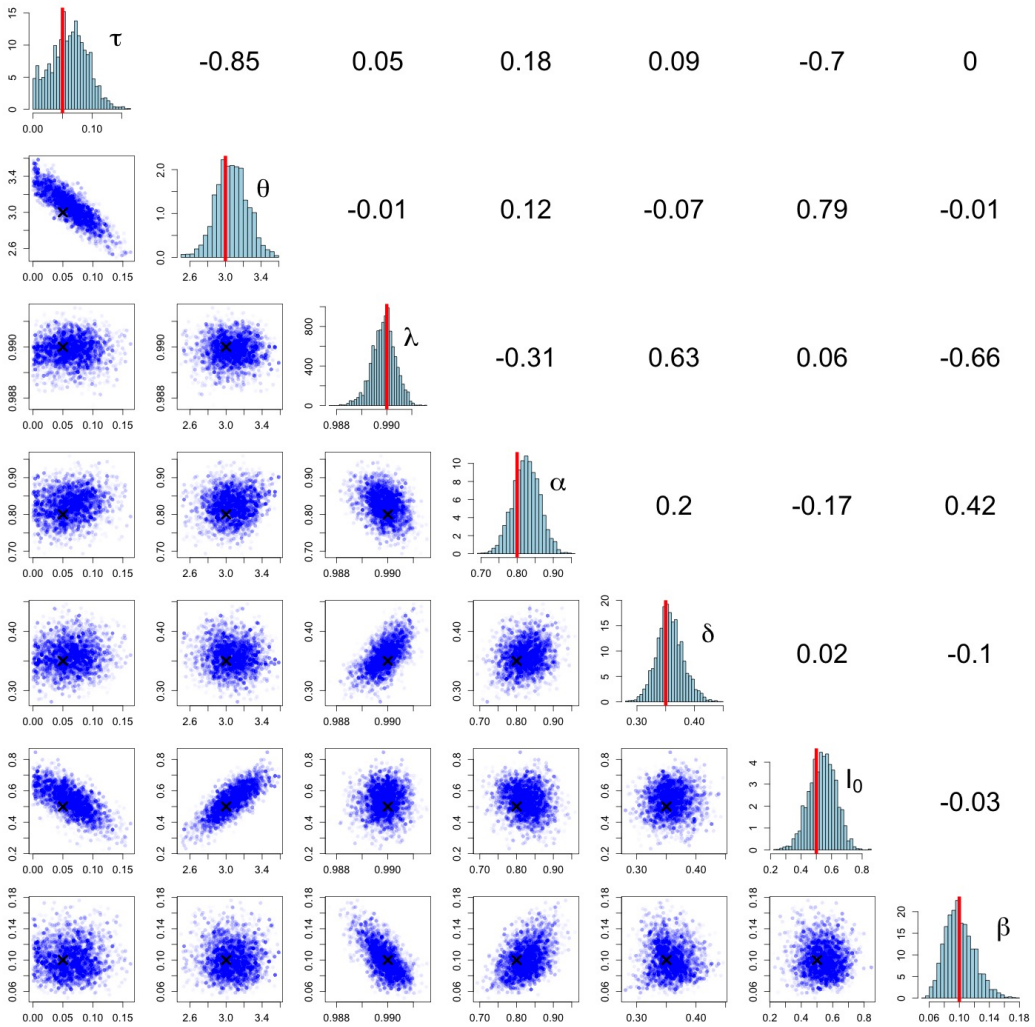

Figure 3: **Posterior Estimates for Each Model Parameter**. The plots on the diagonal show histograms of the marginal posterior distribution for each model parameter. The plots on the lower diagonal show each pairwise joint posterior distribution, whereas the text on the upper diagonal reports each pairwise correlation among the model parameters. The true values used to generate the data are represented as the vertical red line in the diagonal plots, and as black "X" symbols in the lower diagonal.

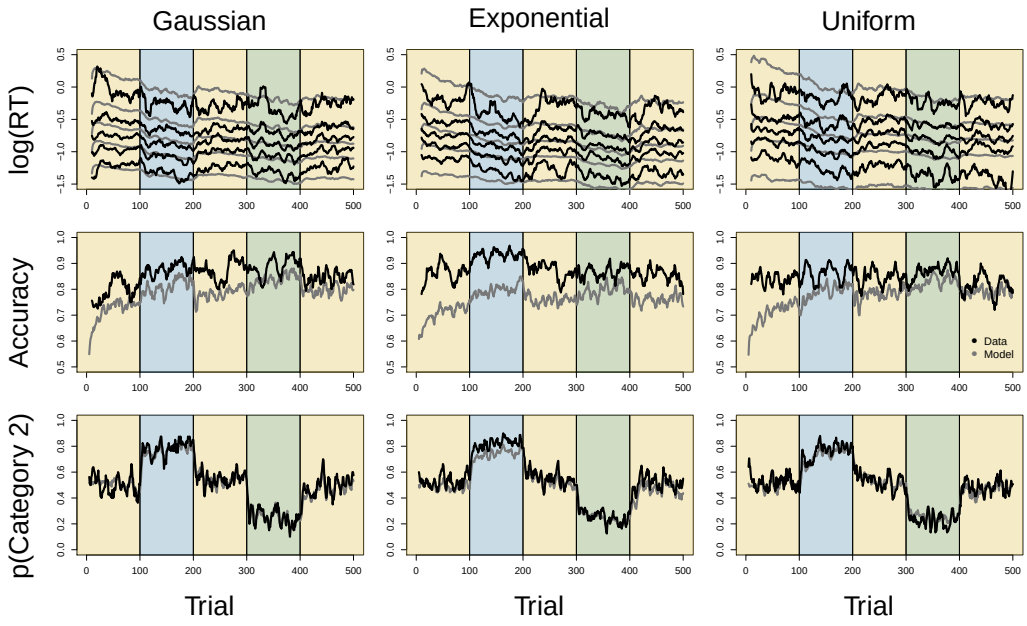

Figure 4: **Aggregated IBM1 Model Predictions and Data from Experiment 1**. Aggregated predictions from the IBM1 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

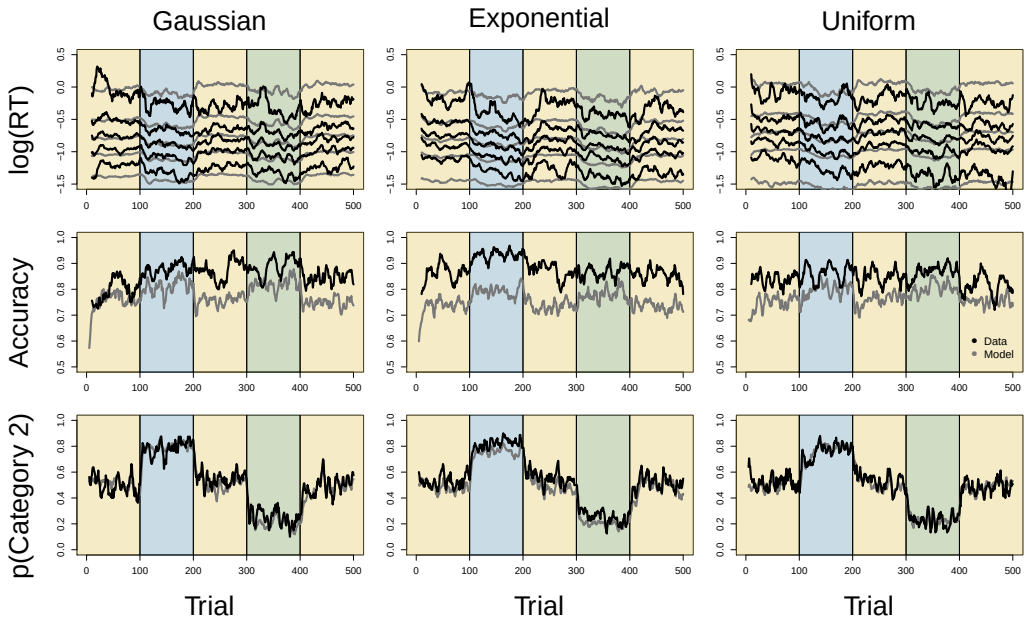

Figure 5: **Aggregated IBM2 Model Predictions and Data from Experiment 1**. Aggregated predictions from the IBM2 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

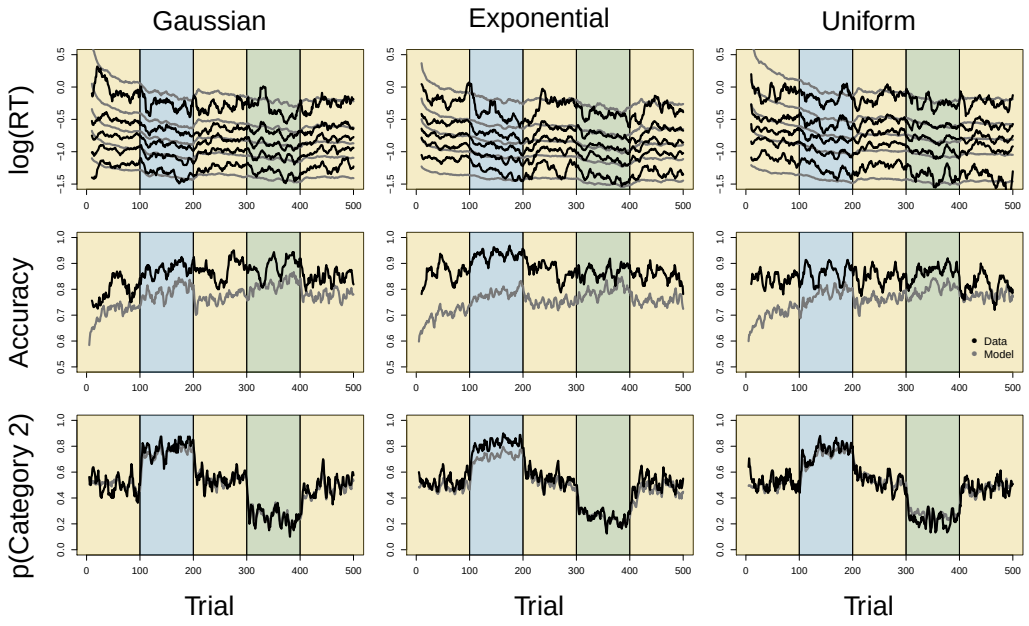

Figure 6: **Aggregated ICM1 Model Predictions and Data from Experiment 1**. Aggregated predictions from the ICM1 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

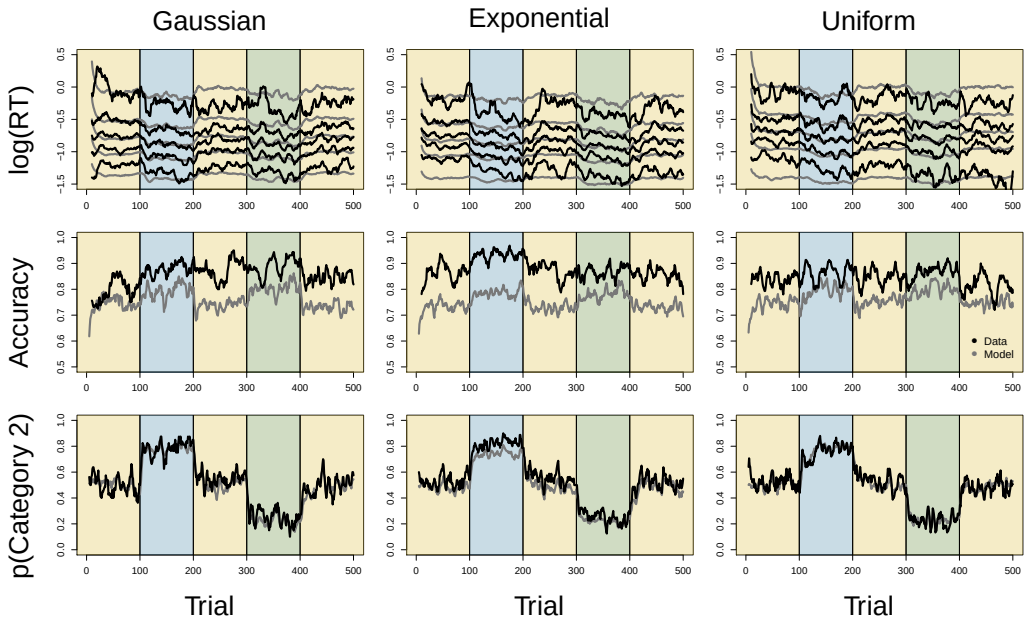

Figure 7: **Aggregated ICM2 Model Predictions and Data from Experiment 1**. Aggregated predictions from the ICM2 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

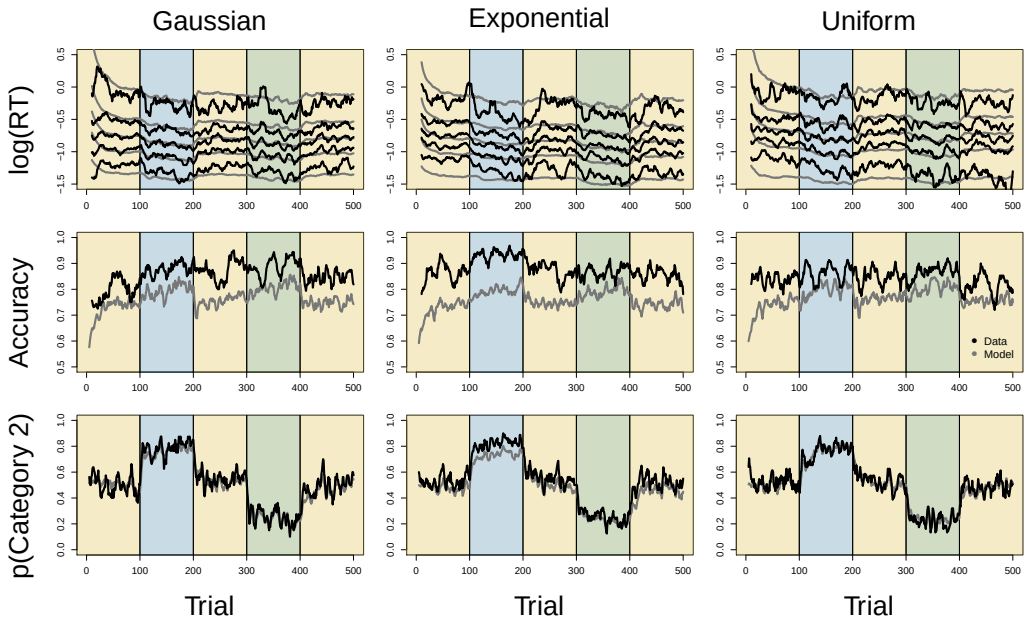

Figure 8: **Aggregated SE Model Predictions and Data from Experiment 1**. Aggregated predictions from the SE model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

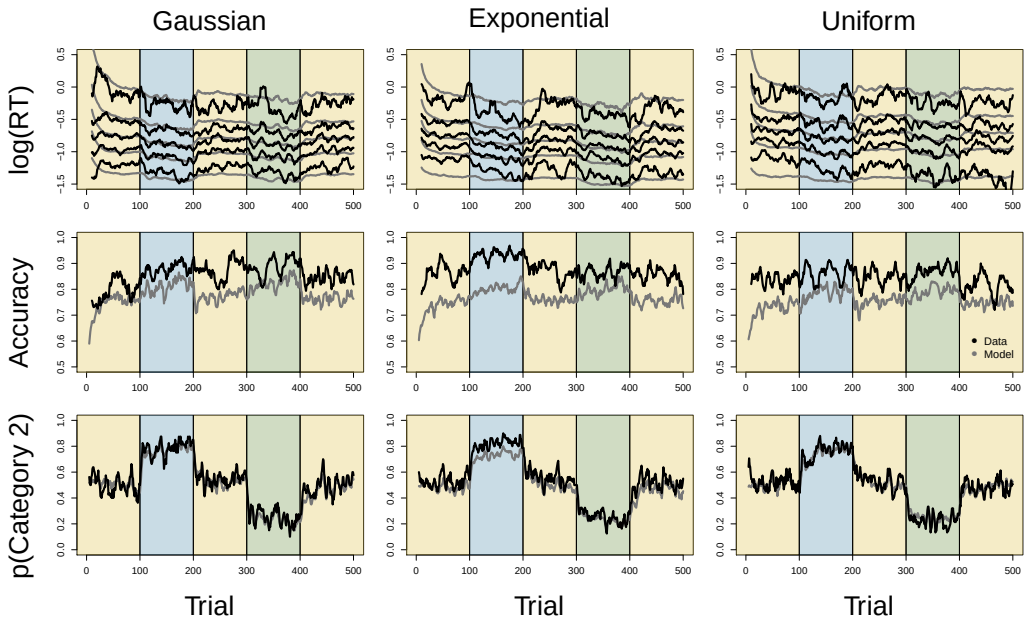

Figure 9: **Aggregated SG Model Predictions and Data from Experiment 1**. Aggregated predictions from the SG model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

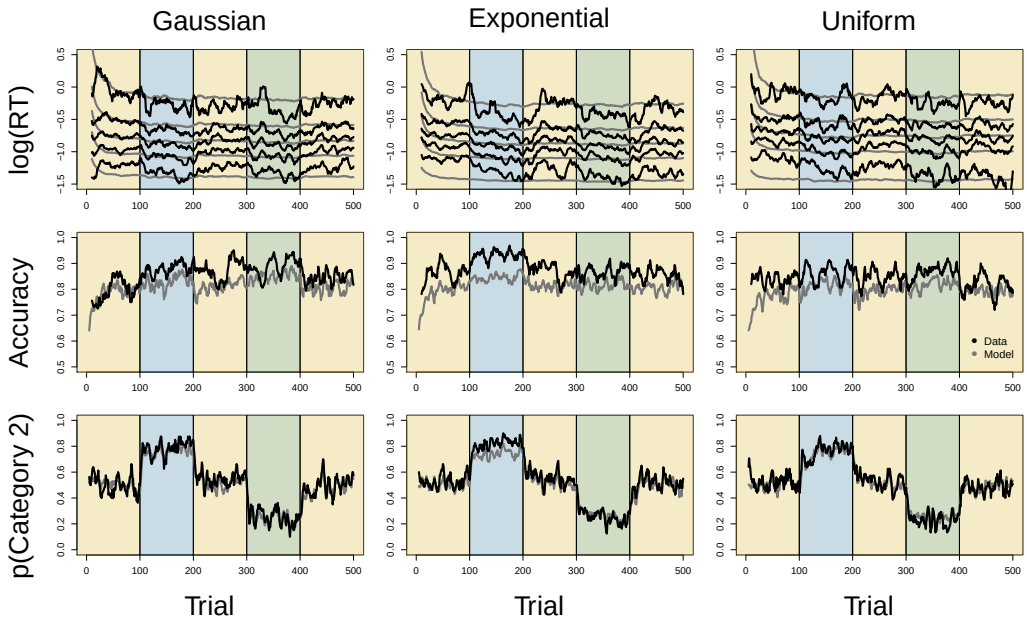

Figure 10: **Aggregated SEP Model Predictions and Data from Experiment 1**. Aggregated predictions from the SEP model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

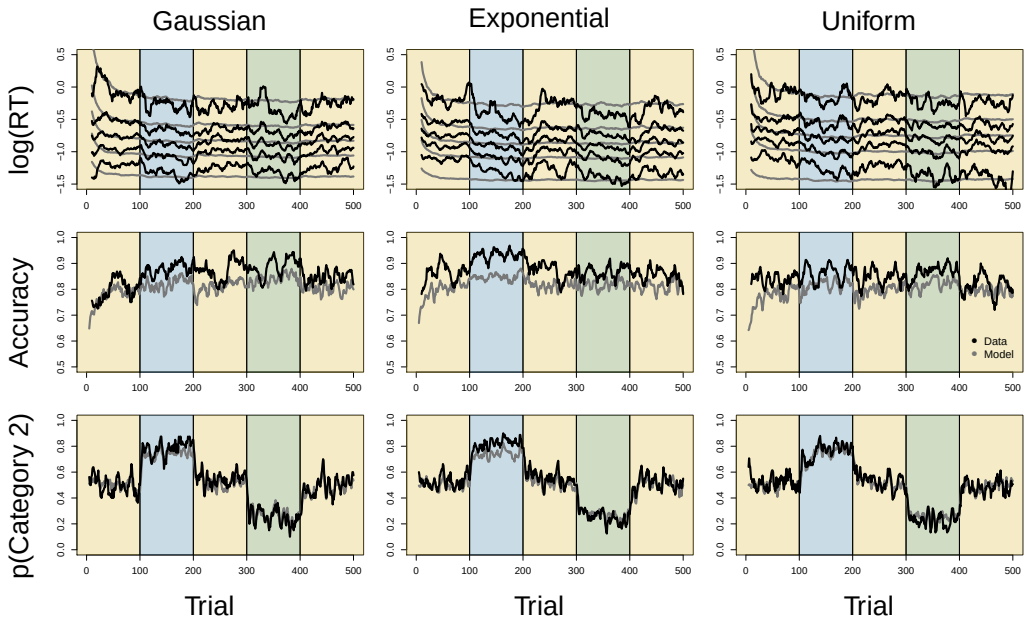

Figure 11: **Aggregated SGP Model Predictions and Data from Experiment 1**. Aggregated predictions from the SGP model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

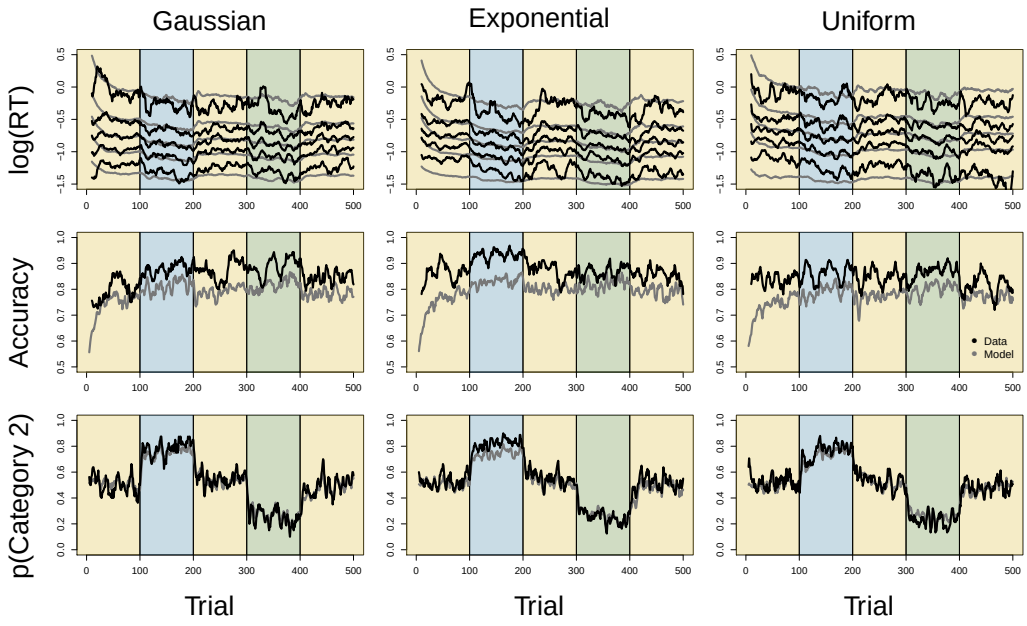

Figure 12: **Aggregated SEL Model Predictions and Data from Experiment 1**. Aggregated predictions from the SEL model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

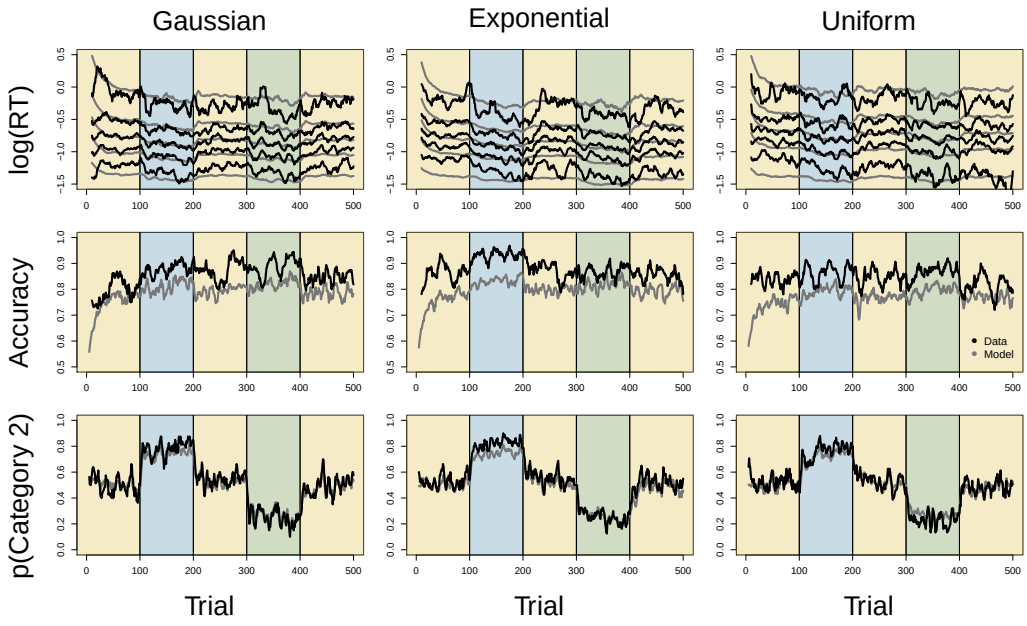

Figure 13: **Aggregated SGL Model Predictions and Data from Experiment 1**. Aggregated predictions from the SGL model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for three behavioral metrics: response time (top row), accuracy (middle row), and response frequency (bottom row). The behavioral metrics are separated by the type of distribution (columns) that generated the category features: Gaussian (left), exponential (middle), and uniform (right). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

2. Experiment 1: Model Fits

3. Experiment 2: Model Fits

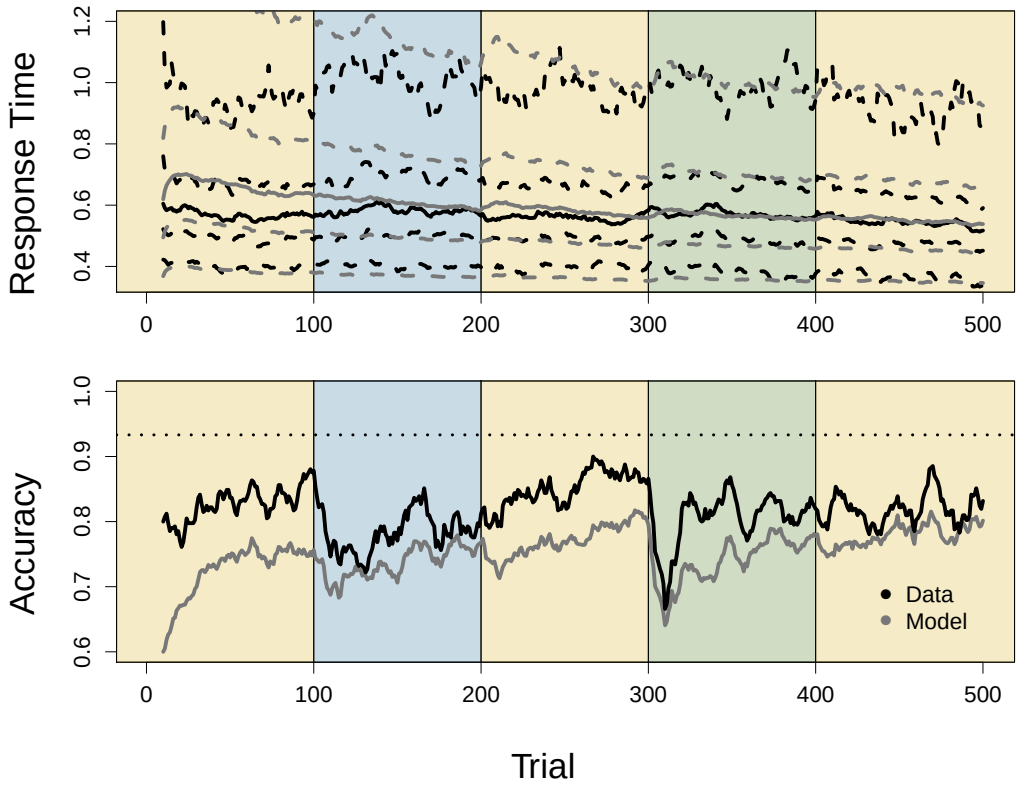

Figure 14: **Aggregated IBM1 Model Predictions and Data from Experiment 2**. Aggregated predictions from the IBM1 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

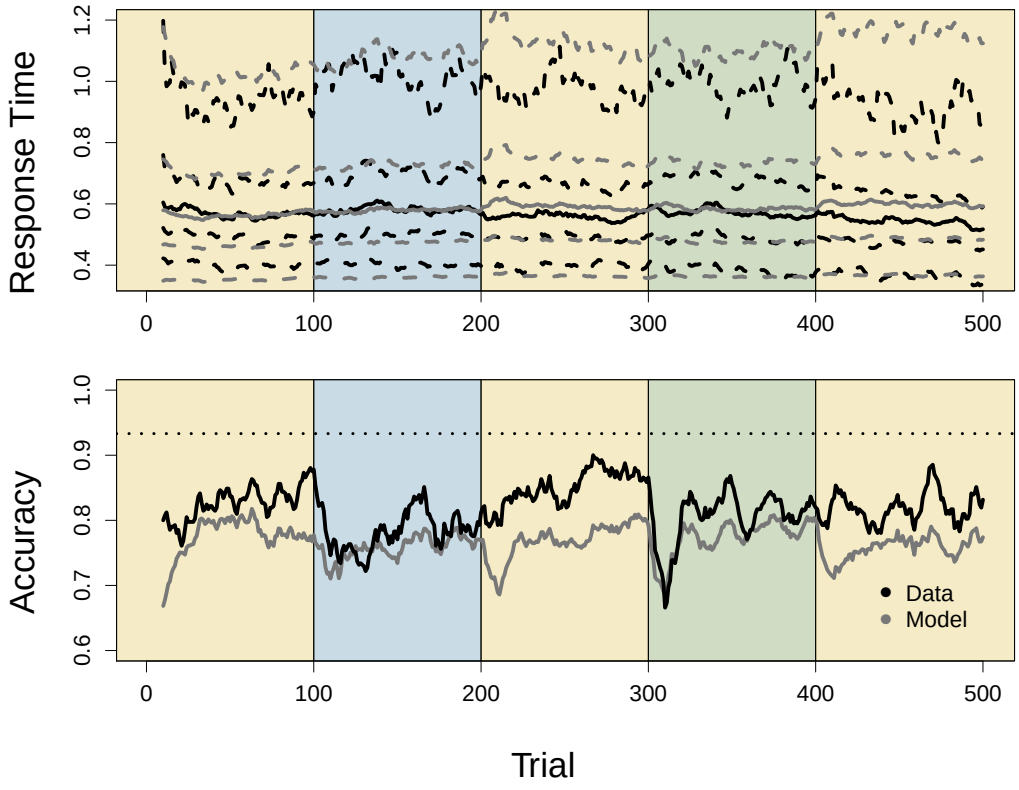

Figure 15: **Aggregated IBM2 Model Predictions and Data from Experiment 2**. Aggregated predictions from the IBM2 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

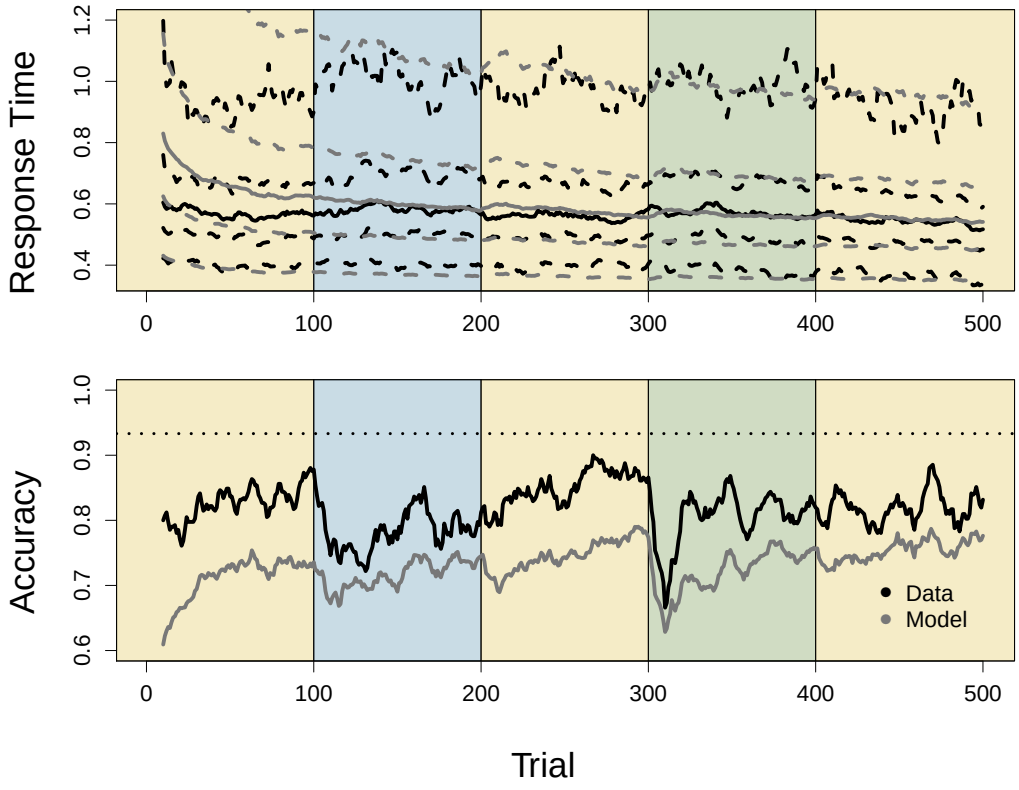

Figure 16: **Aggregated ICM1 Model Predictions and Data from Experiment 2**. Aggregated predictions from the ICM1 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

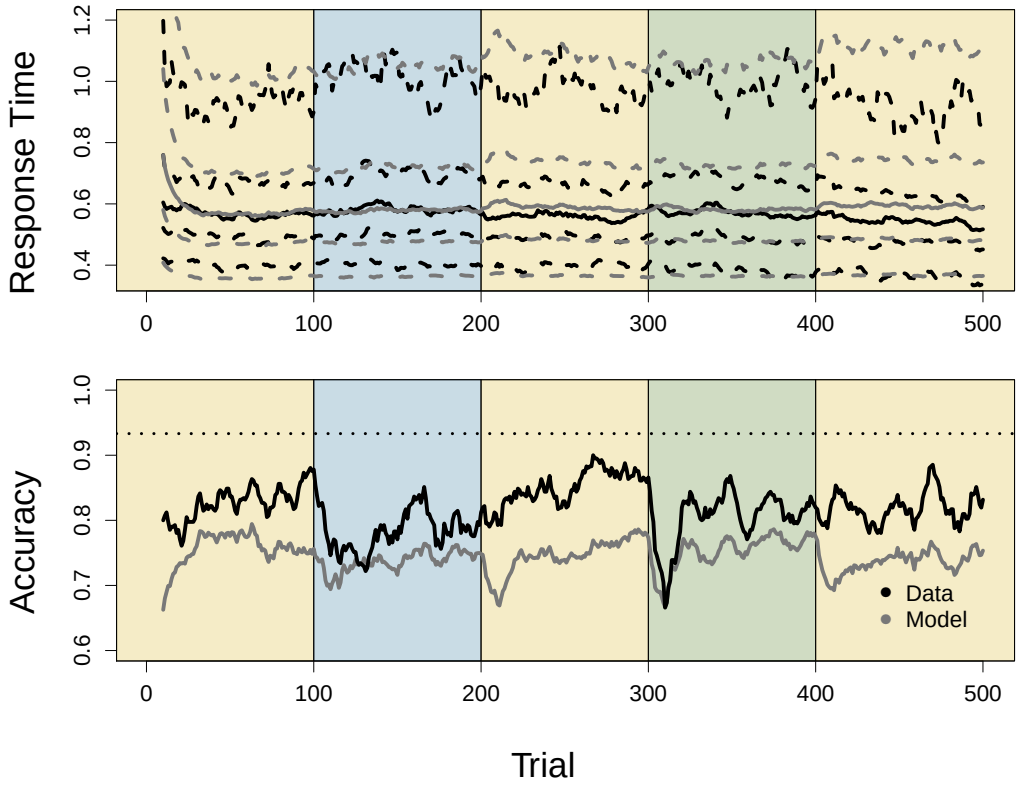

Figure 17: **Aggregated ICM2 Model Predictions and Data from Experiment 2**. Aggregated predictions from the ICM2 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

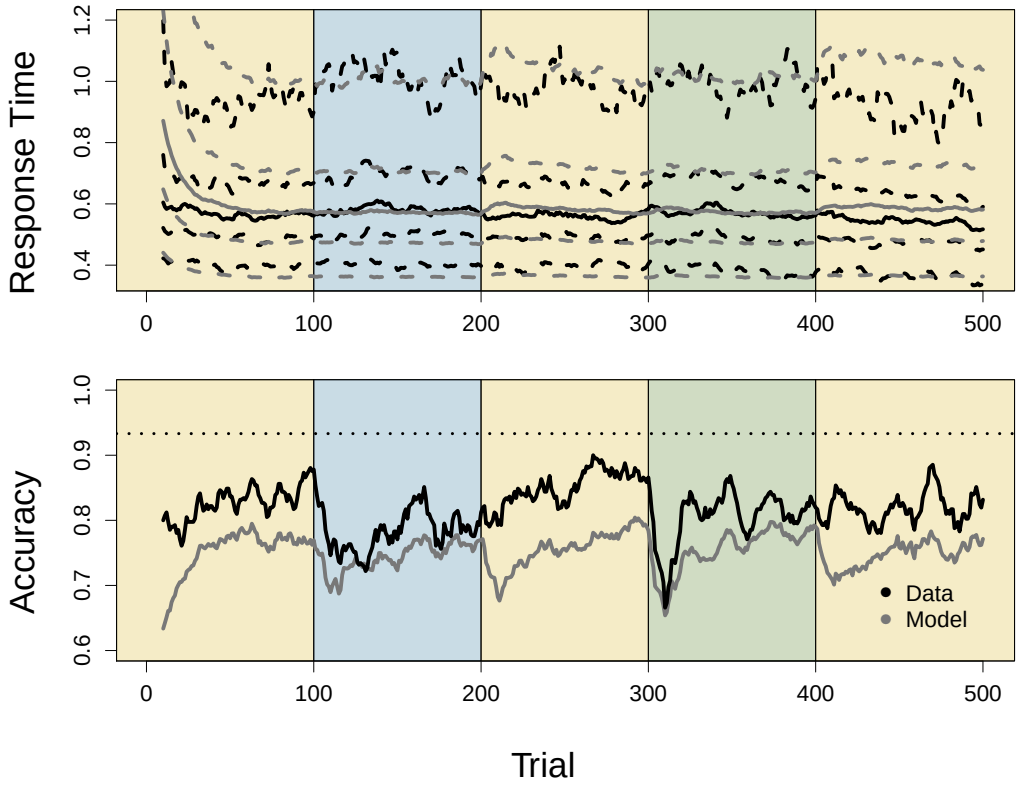

Figure 18: **Aggregated SE Model Predictions and Data from Experiment 2**. Aggregated predictions from the SE model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

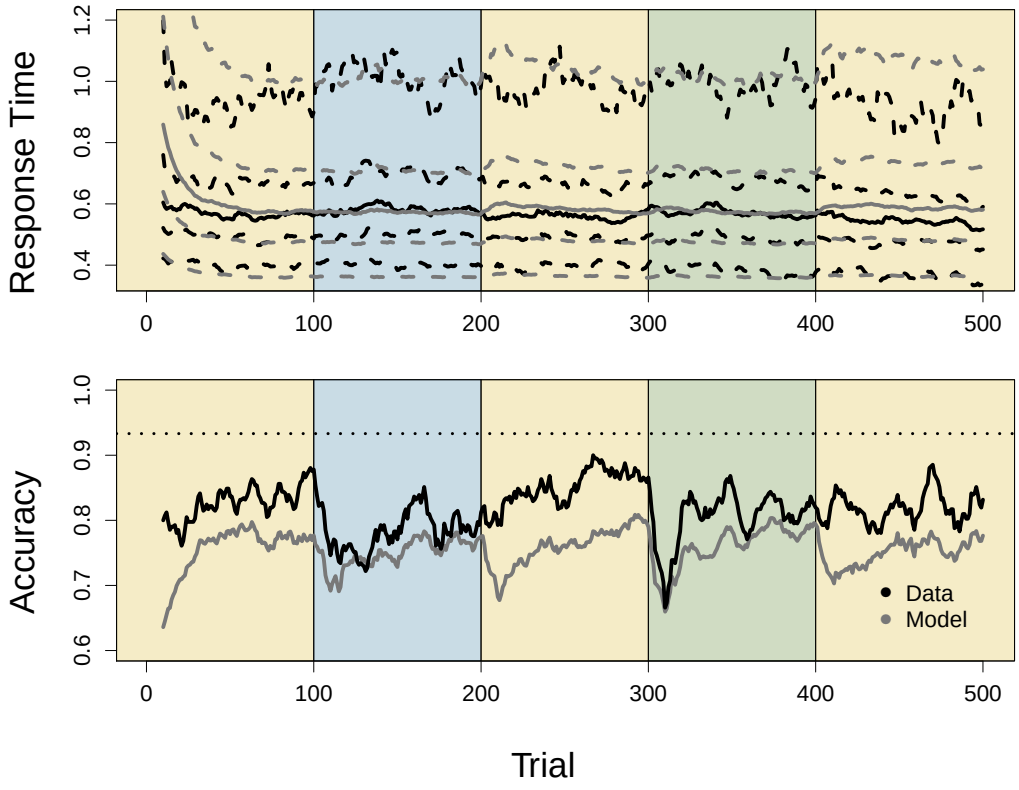

Figure 19: **Aggregated SG Model Predictions and Data from Experiment 2**. Aggregated predictions from the SG model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

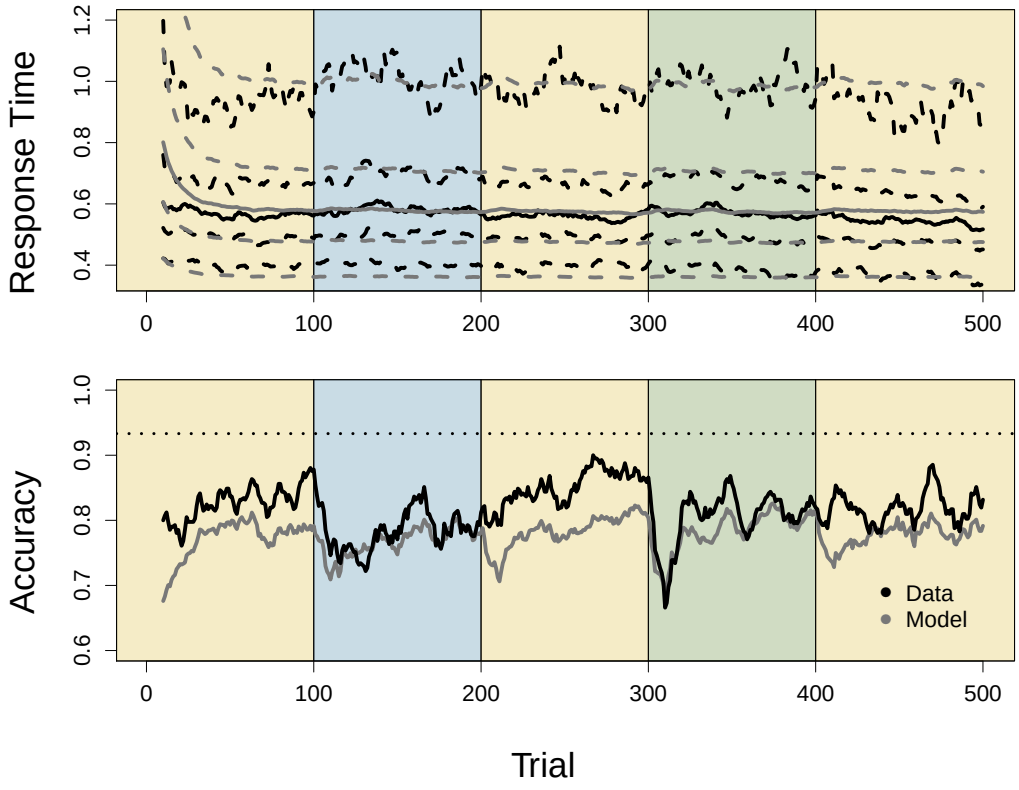

Figure 20: **Aggregated SEP Model Predictions and Data from Experiment 2**. Aggregated predictions from the SEP model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

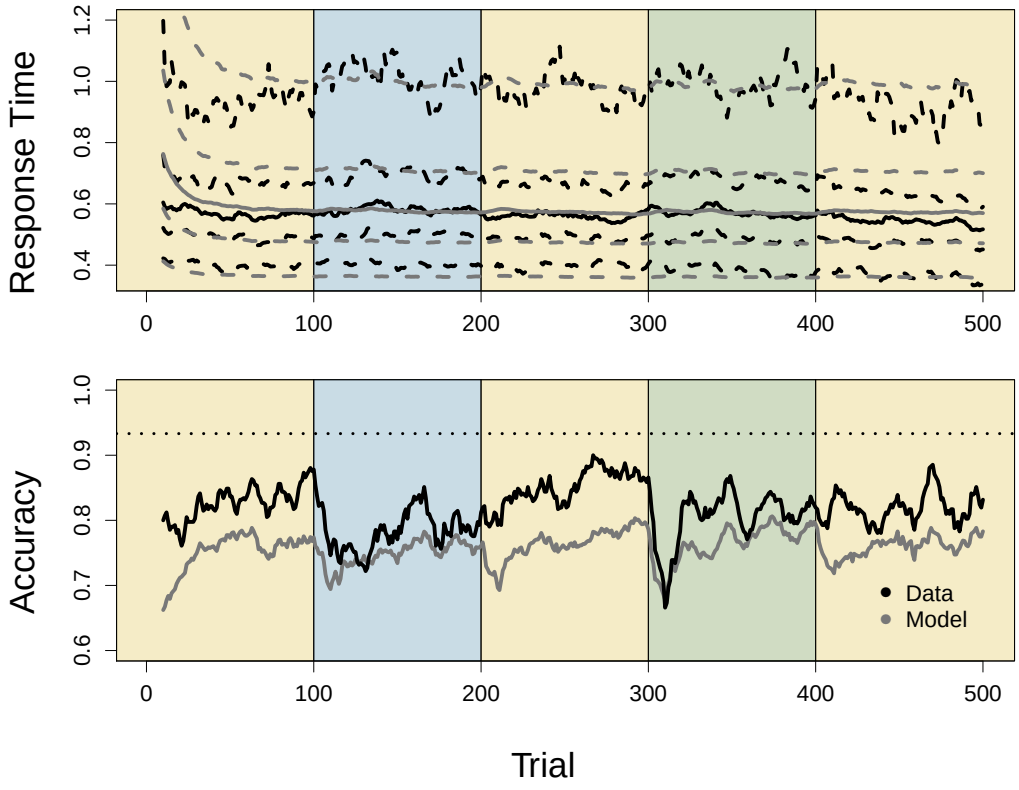

Figure 21: **Aggregated SGP Model Predictions and Data from Experiment 2**. Aggregated predictions from the SGP model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

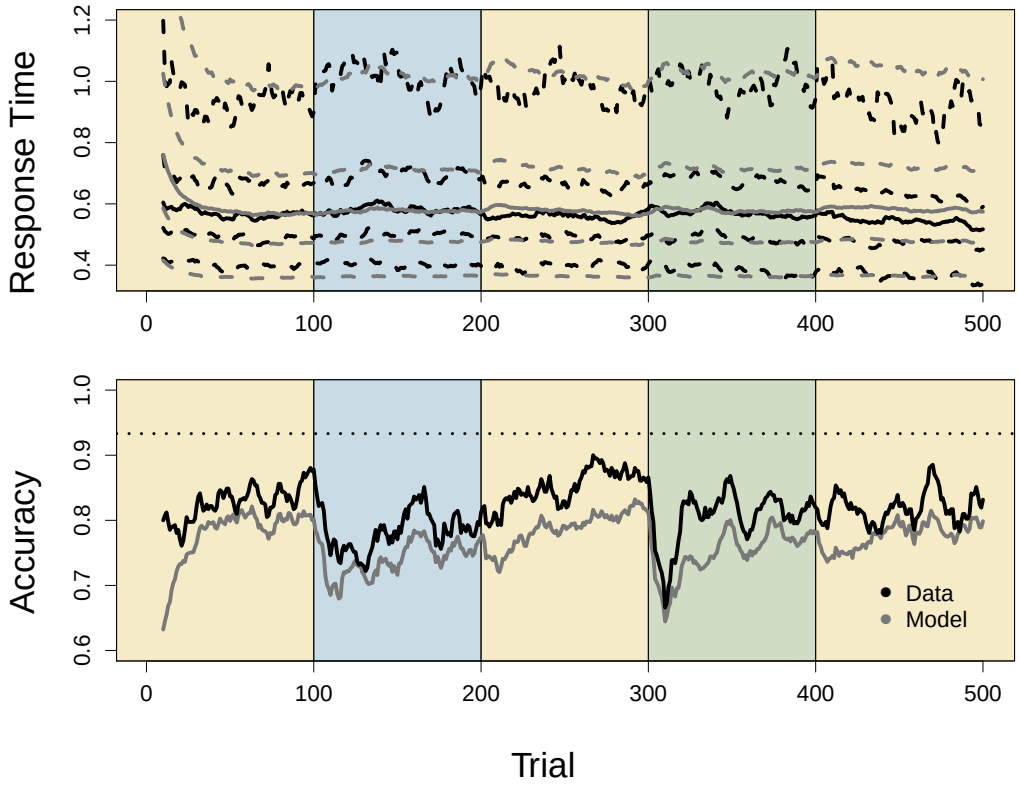

Figure 22: **Aggregated SEL Model Predictions and Data from Experiment 2**. Aggregated predictions from the SEL model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

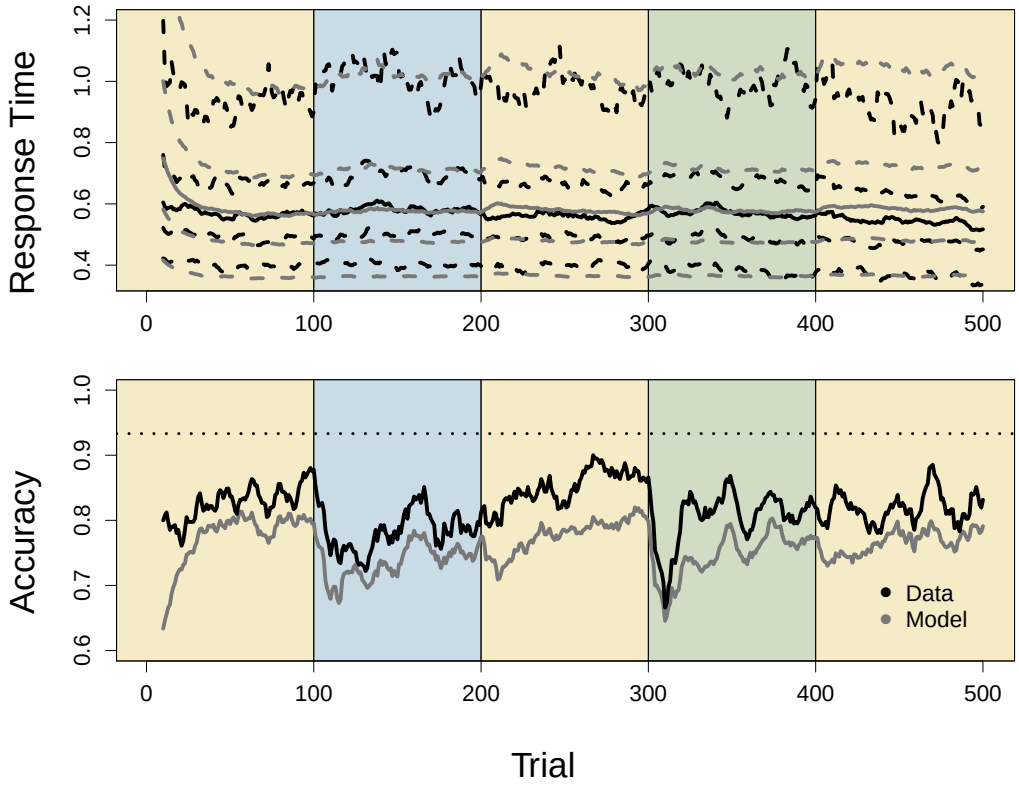

Figure 23: **Aggregated SGL Model Predictions and Data from Experiment 2**. Aggregated predictions from the SGL model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

4. Experiment 3: Model Fits

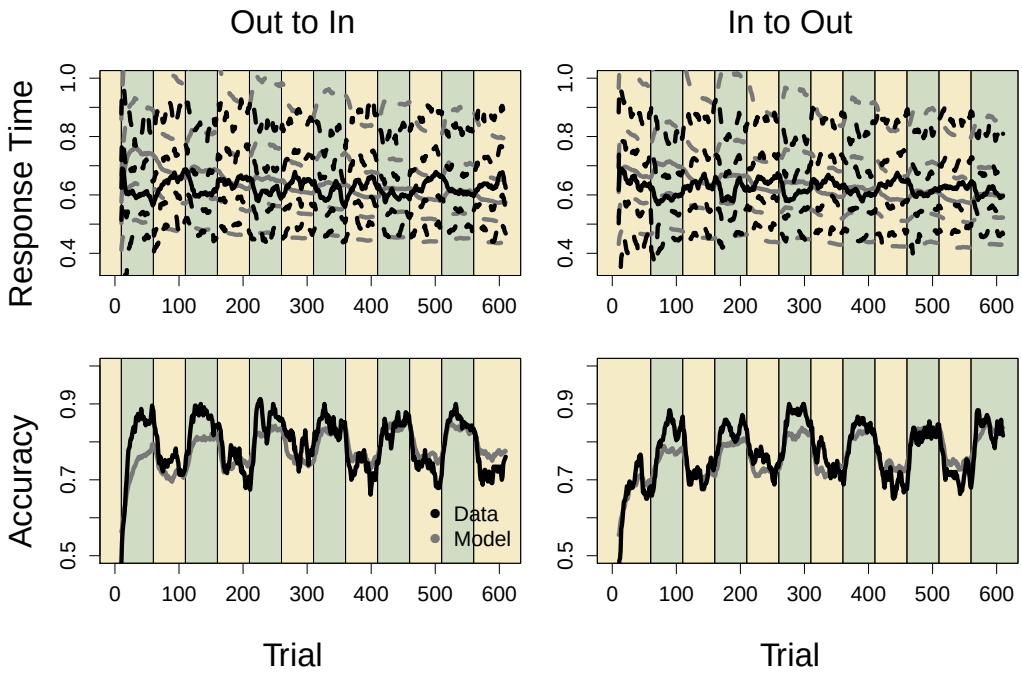

Figure 24: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the IBM1 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

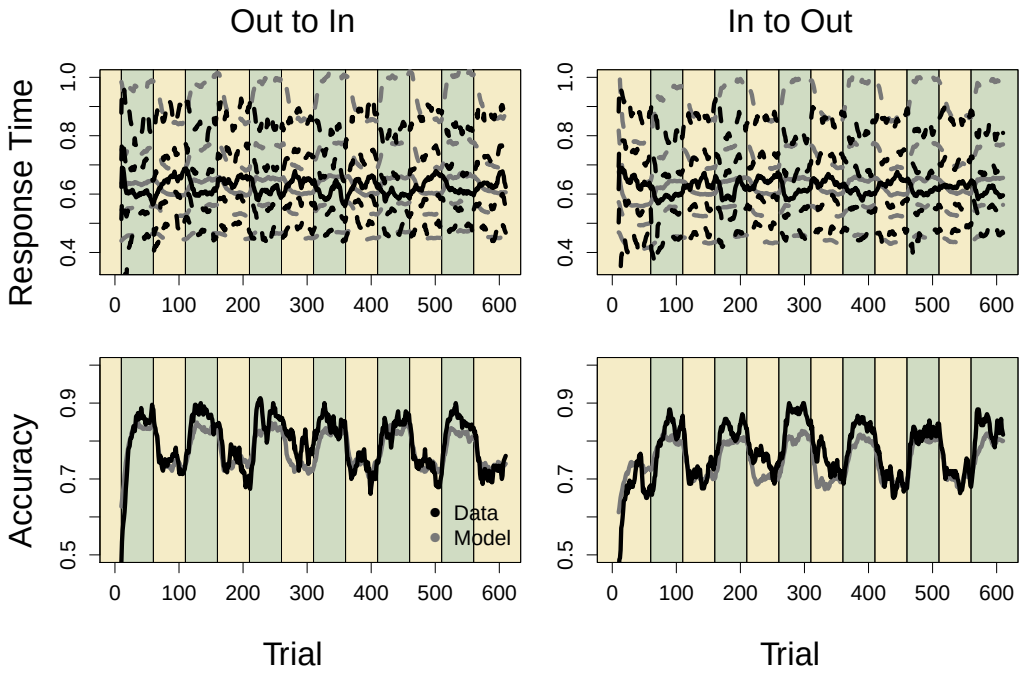

Figure 25: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the IBM2 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

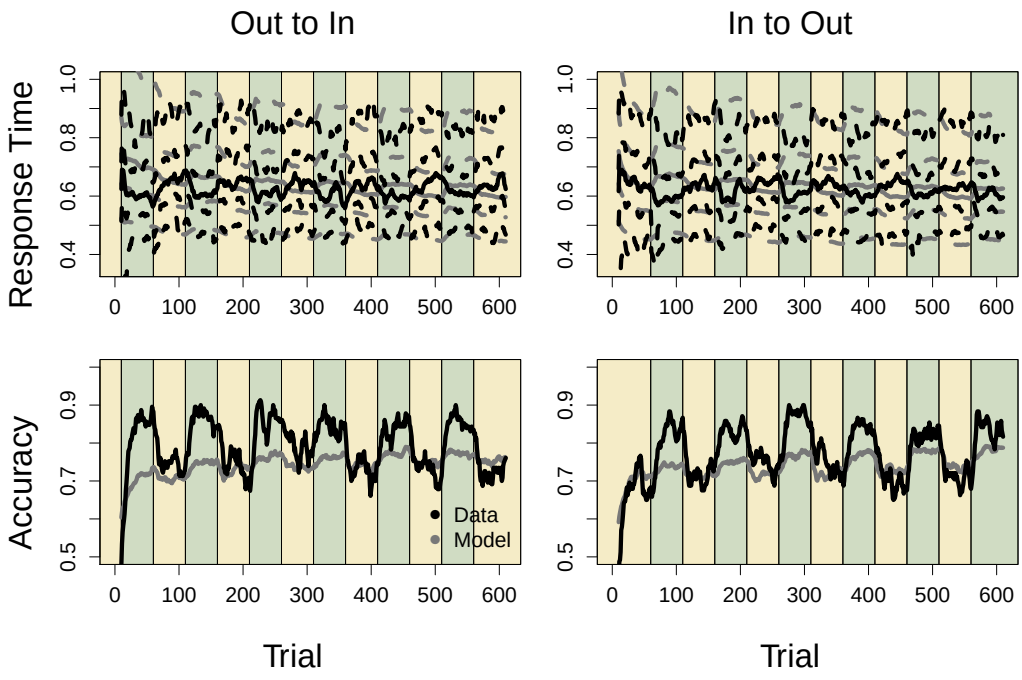

Figure 26: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the ICM1 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

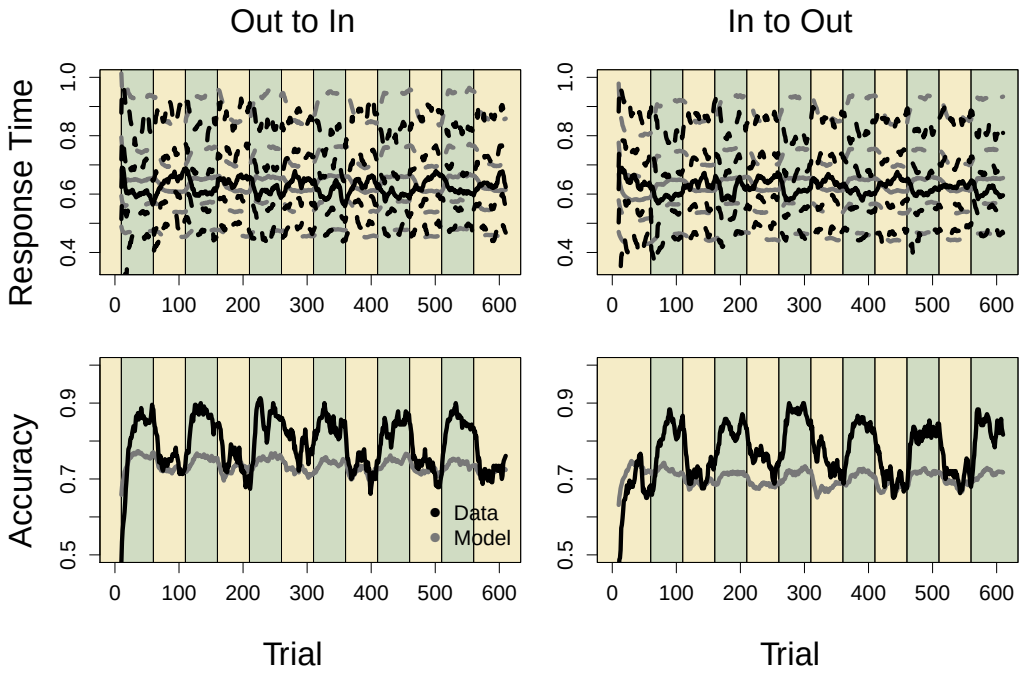

Figure 27: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the ICM2 model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

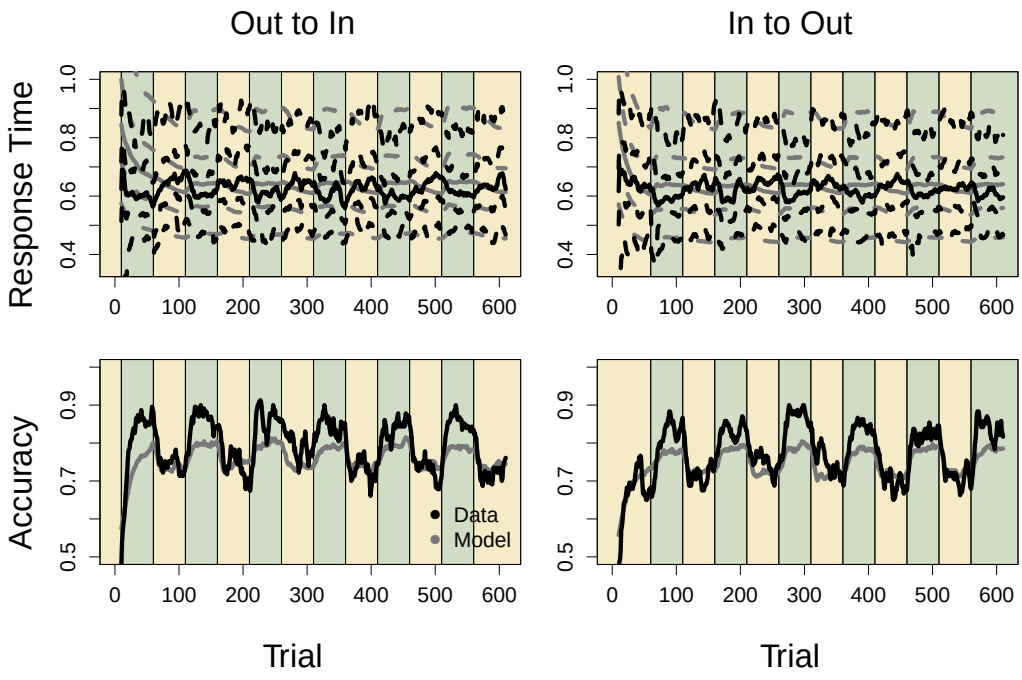

Figure 28: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the SE model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

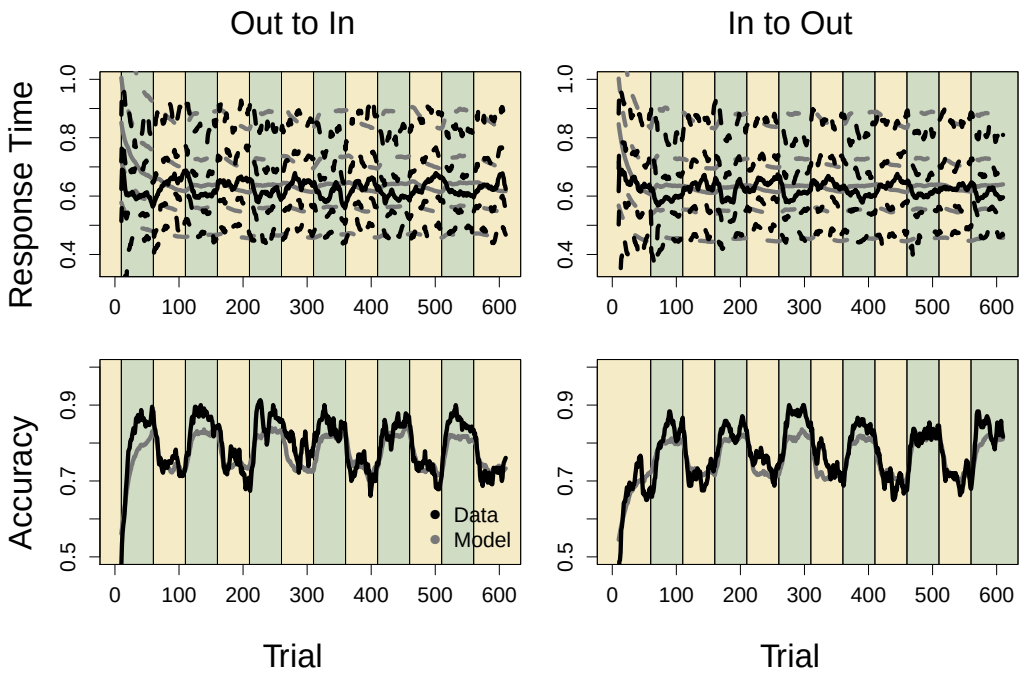

Figure 29: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the SG model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

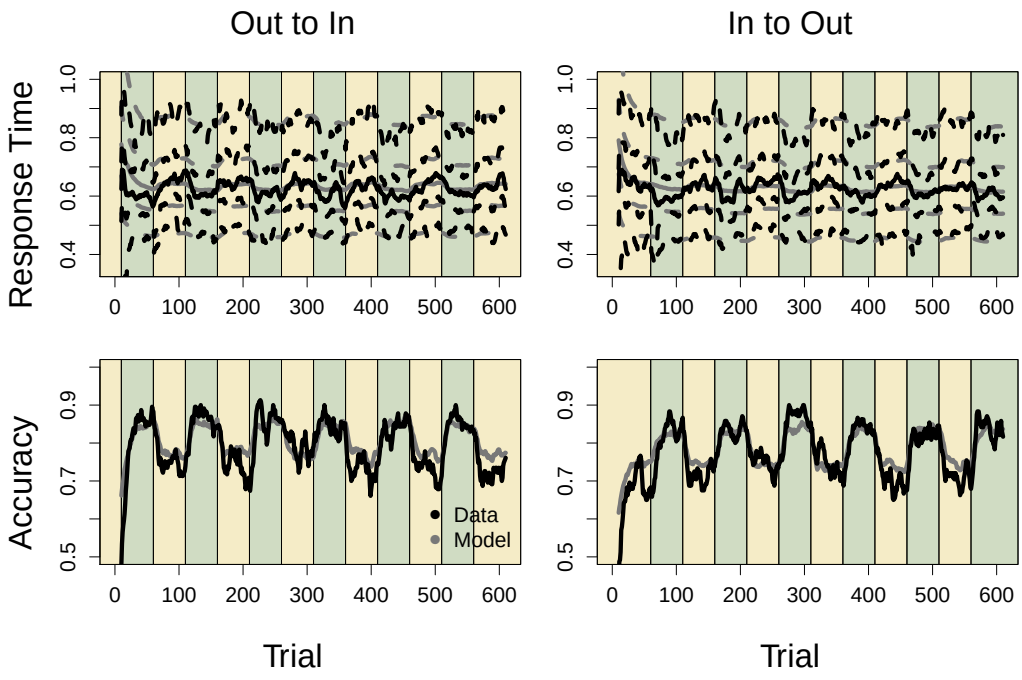

Figure 30: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the SEP model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

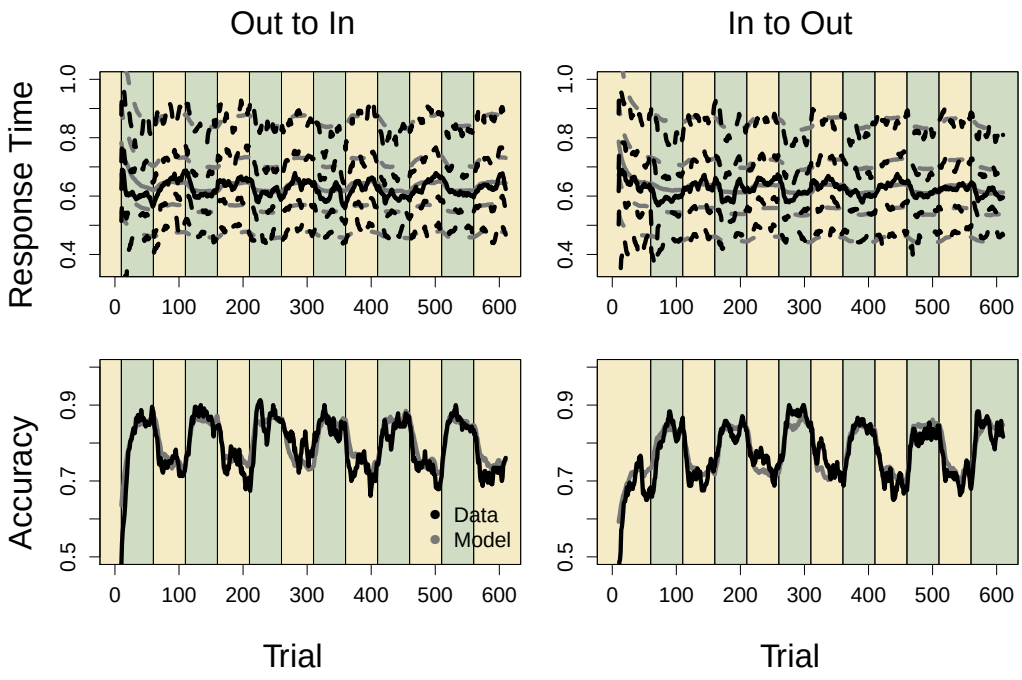

Figure 31: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the SGP model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

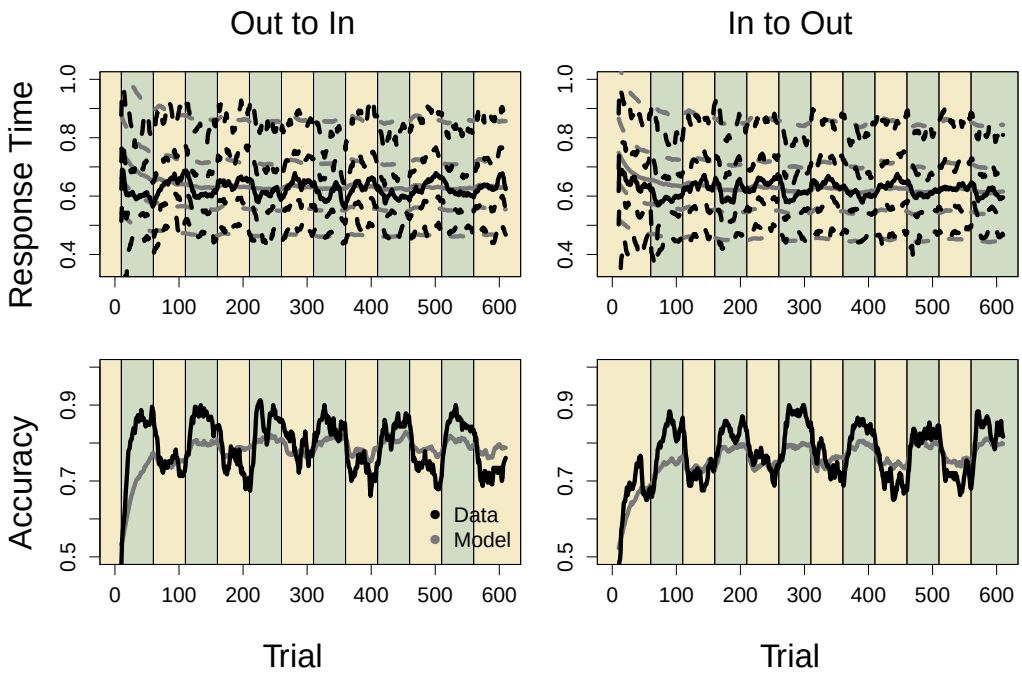

Figure 32: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the SEL model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

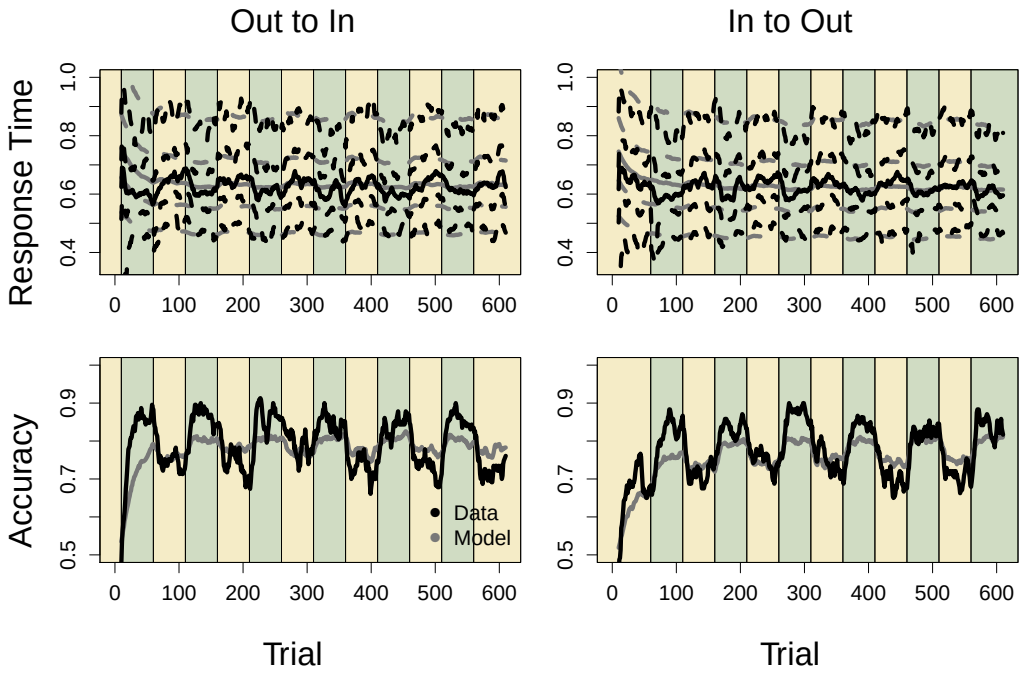

Figure 33: **Aggregated Model Predictions and Data from Experiment 3**. Aggregated predictions from the SGL model variant (gray lines) are shown against the aggregated data from the experiment (black lines) for two behavioral metrics: response time (top row), and accuracy (bottom row). The columns represent the two different conditions. Within each panel, the blocks of the experiment are color coded to represent the contexts of each stimulus environment.

5. References

Turner, B. M., Sederberg, P. B., Brown, S. D., Steyvers, M., 2013. A method for efficiently sampling from distributions with correlated dimensions. Psychological Methods 18, 368–384.