# Supporting Material

Hongjing Lu [1,2], Randall R. Rojas [3], Tom Beckers [4,5], and Alan L. Yuille [1,2]

[1]Department of Psychology, University of California, Los Angeles

[2]Department of Statistics, University of California, Los Angeles

[3]Department of Economics, University of California, Los Angeles

[4]Department of Psychology, KU Leuven

[5]Department of Clinical Psychology and Amsterdam Brain and Cognition, University of Amsterdam

## Author Note

Address for proofs: Hongjing Lu

Department of Psychology

University of California, Los Angeles

405 Hilgard Ave.

Los Angeles, CA 90095-1563

Email: hongjing@ucla.edu

Supporting Material

## 1 Summary of the Particle Filters Implementation

We represent the probability distributions $P(\vec{\omega}^t|\mathcal{D}^t, m), P(\vec{\omega}^t|\mathcal{D}^t, m)$ by a set of particles and use particle filters (Liu, 2001) to update the distributions over time. An implementation using particle filters supports the key computations required by our theory – model selection, parameter estimation, and model averaging.

Particle filters approximate distributions like $P(\vec{\omega}_t|\mathcal{D}_t, M)$ by a set of discrete particles $\{\vec{\omega}_t^\mu : \mu \in \Gamma\}$. In simulations reported in the present paper, we used 10000 particles (i.e., $|\Gamma| = 10000$), as we did not obtain significantly different results with a further increase in the number of particles. To test significance, we sampled 10 times (each time with 10000 particles) and computed the variances of the estimates of quantities of interest (e.g., our estimates of the weights).

We use particle filters to perform the sequential Bayesian updates. We initialize by drawing samples $\{\vec{\omega}_1^\mu : \mu \in \Gamma\}$ from the prior distribution $P(\vec{\omega})$, which is a Gaussian distribution with a mean of zero and a small variance. Then we proceed recursively following the prediction and measurement stages. Let $\{\vec{\omega}_t^\mu : \mu \in \Gamma\}$ be the set of particles representing $P(\vec{\omega}_t|D_t)$ at time $t$. We sample from the Gaussian distribution $P(\vec{\omega}_{t+1}|\vec{\omega}_t^\mu)$ for each $\mu$ to give a new set of particles $\{\vec{\bar{\omega}}_t^\mu : \mu \in \Gamma\}$, which represents the prediction $P(\vec{\omega}_{t+1}|D_t)$.

In order to perform the correction step, we compute the importance weights $\lambda^\mu = P(O_{t+1}|\vec{\bar{\omega}}_{t+1}^\mu, \vec{x}_{t+1})$ and normalize them to obtain $\bar{\lambda}^\mu = \lambda^\mu/(\sum_\mu \lambda^\mu)$. Then we re-sample with replacement from the set $\{\bar{\omega}_{t+1}^\mu : \mu \in \Gamma\}$ using probability $\bar{\lambda}^\mu$. This gives a new set $\{\vec{\omega}_{t+1}^\nu : \nu \in \Gamma\}$ of particles, which represent $P(\vec{\omega}_{t+1}|D_{t+1})$.

To implement *parameter estimation* we estimate the causal weights by computing the average with respect to the posterior distribution $P(\vec{\omega}^t|\{O^t\}, \{\vec{x}^t\})$. This can be estimated using the particles:

$$\int d\vec{\omega}(\vec{\omega}_t)P(\vec{\omega}_t|\mathcal{D}_t) = (1/|\Gamma|)\sum_{\mu\in\Gamma}(\vec{\omega}_t^{\mu}). \tag{1}$$

The use of particle filters provides a potential way to study the robustness of the model – i.e., how its performance would be affected by small inaccuracies in the model or degradations due to limited neuronal resources during computation, i.e., reducing the number of particles (Courville & Daw, 2007; Brown, & Steyvers, 2009; Sanborn, Griffiths, & Navarro, 2010). For the simulations reported in the paper, we used a large number of particles (i.e., 10000) to ensure the precision of the inference. Figure 1 illustrates the model estimations using a noisy-or rule for the forward blocking paradigms discussed in section 4. The simulation results indicate that the estimate of mean causal weights are approximately constant,b ut their associated variance is increased when the number of particle filters is reduced. When reducing the number of particles, the estimated causal weights increase their variability .

We can also use particles to compute the *model evidence*. The model evidence is expressed as $P(d_t|\mathcal{D}_{t-1})P(d_{t-1}|\mathcal{D}_{t-2})...P(d_1)$, where $d_t = (O_t, \vec{x}^t)$. We evaluate each term $P(d_{t+1}|\mathcal{D}_t) = \int d\vec{\omega}_{t+1}P(d_{t+1}|\vec{\omega}_{t+1})P(\vec{\omega}_{t+1}|\mathcal{D}_t)$ by $P(d_{t+1}|\mathcal{D}_t) = \frac{1}{|\Gamma|}\sum_{\mu\in\Gamma}P(d_{t+1}|\vec{\omega}_{t+1}^{\mu})$.

To implement *model averaging*, we perform two steps. First, we estimate the expected weights of each model as described above. Next we compute the model evidence, as above, and then compute $P(m|\mathcal{D})$. Finally we average with respect to

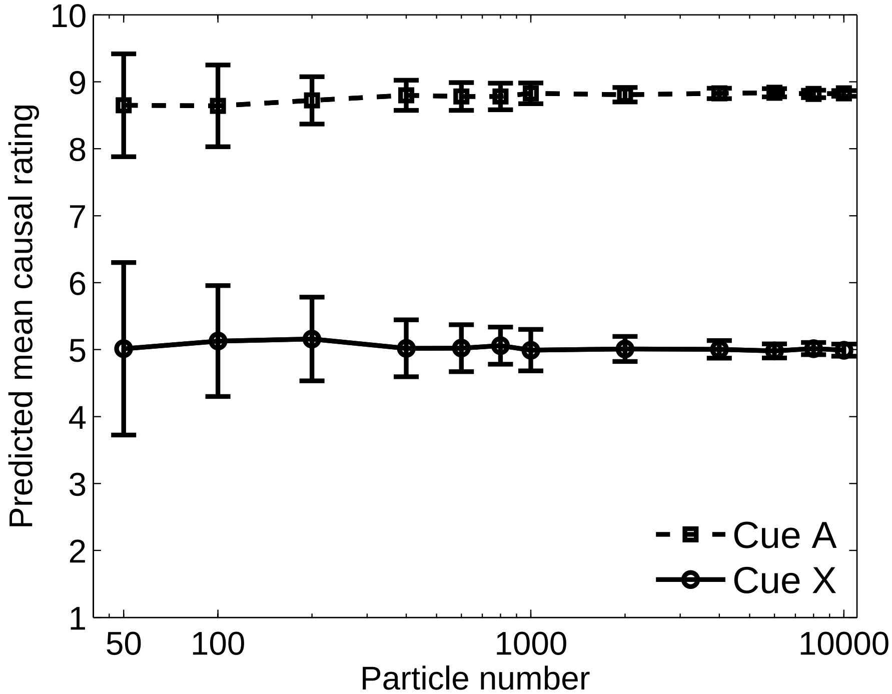


*Figure 1*. Simulation results for forward blocking as a function of particle numbers. The forward blocking paradigm (6A+, 6AX+) is adopted from the study by Vandorpe and De Houwer (2005). The results are based on 100 simulation runs. The error bars indicate standard devisions

these to obtain the result.

From a computational perspective, post-training differs from pre-training in that post-training precludes model selection prior to the blocking session. We assume that the learner therefore proceeds by model averaging: running both the linear-sum and noisy-max sequential models and then combining their estimate weighted by the probability that each model could explain the post-training data.

This procedure yields $\langle\omega\rangle = P(m_1|\mathcal{D}_{post})\bar{\omega}_{m_1} + P(m_2|\mathcal{D}_{post})\bar{\omega}_{m_2}$, where $\mathcal{D}_{post}$ are the data in post-training (i.e., Phase 3), $m_1$ and $m_2$ represent the linear-sum and noisy-max models, respectively, and $P(m_i|\mathcal{D}_{post})$ is the evidence for each model based on observations in the post-training phase. $\bar{\omega}_{m_i}$ is the estimated mean value of causal strength using each model based on observations in the first two training phases.

## 2 Parameters used in the simulations

The parameter $\sigma_T^2$ plays an important role in the temporal prior (see Eq. 12) to control the amount of variation for the weights to change from trial to trial. We set $\sigma_T$ as 0.4 for all simulations to account for the blocking effect presented in section 4 and the abstract transfer effects in section 5, so that the contributions from the dynamic module is equated for all the models using different causal integration rules. For the linear-sum and noisy-max model, two additional parameters $\sigma_h$ and $\sigma_m$ (in Eq. 1 and 3) were set as 0.05. These two parameters were used in the likelihood term to control the uncertainty in associating cues with outcome variables. The parameter $T$ in Eq. 6 for the noisy-max rule is 0.6. The parameter values were selected to provide the best account for the pre-training experiment in Section 5. The same set of parameter values was used to account for the other experimental findings using deterministic causes in Section 4 and 5. In section 6 modeling the primacy effect, the parameter $\sigma_T^2$ relevant to the learning rate was reduced to 0.3 for learning the probabilistic causes.