Exploring the Influence of Particle Filter Parameters on Order Effects in Causal Learning

Joshua T. Abbott
joshua.abbott@berkeley.edu

Thomas L. Griffiths
tom_griffiths@berkeley.edu

University of California, Berkeley

Overview

Question: The order in which people observe data has an effect on their subsequent judgments and inferences. How do we model this phenomenon?

What we know: Most Bayesian models of human behavior do not produce these effects. However, approximation methods for Bayesian inference have been shown to predict certain order effects.

What we don't know: How do the parameters of these approximation methods influence predictions of order effects?

Our contribution: We investigate the role of certain parameters in a sequential Monte Carlo method known as a *particle filter*. In a simple causal learning task, we find a particular parameter setting is responsible for producing different order effects.

Order Effects in Causal Learning

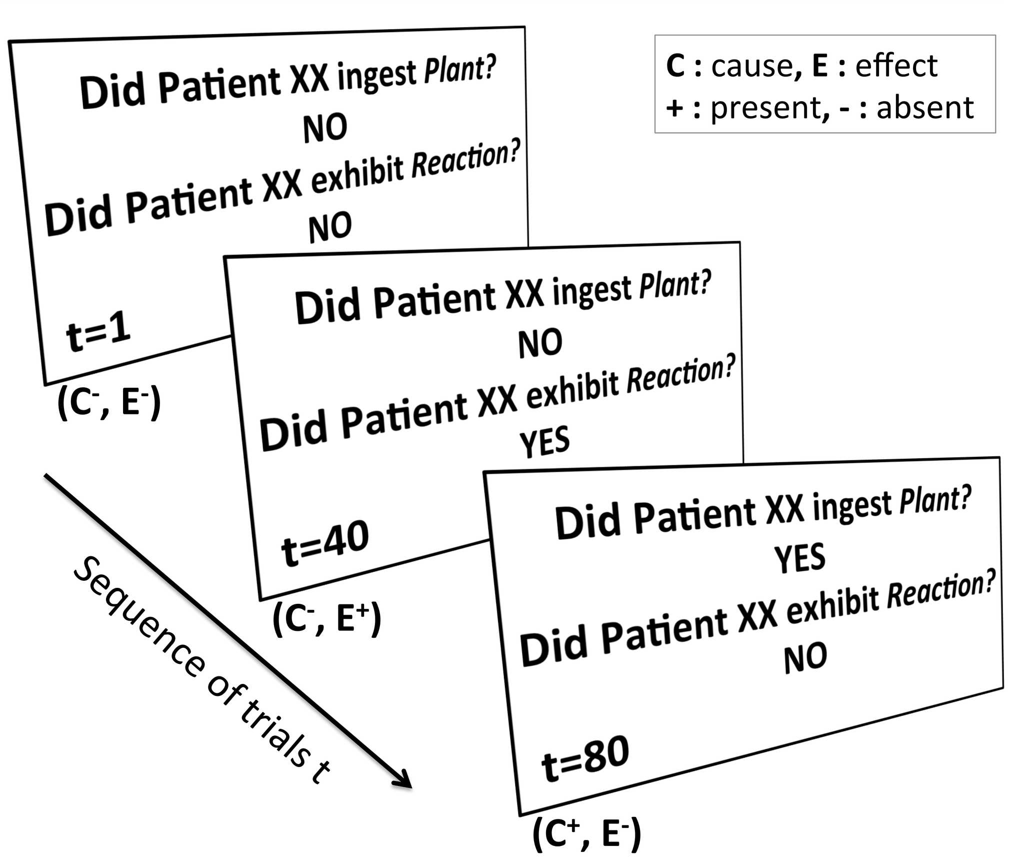

To what extent does *Plant* cause *Reaction*?

Stimuli distribution:

Block 1 (Generative)			Block 2 (Preventative)		
	E^+	E^-		E^+	E^-
C^+	18	2	C^+	2	18
C^-	2	18	C^-	18	2

Primacy effects: initial information has greatest impact on later judgments. Produced when judgment question asked only at the end of the trial sequence. (Dennis and Ahn, 2001)

Recency effects: most recent information has greatest impact on later judgments. Produced when judgment question asked after every 10 trials. (Collins and Shanks, 2002)

Bayesian model of Causal Learning

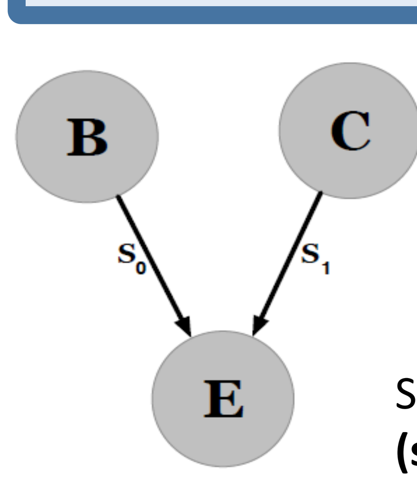

3 binary variables:
(E) Effect
(C) potential cause of interest
(B) a background cause that captures all other causes of **E**

Strength weights on the edges: **(s_0,s_1)** indicating how strongly **B** and **C** influence **E**

What is the probability of the observed data given the strength weights?

C	E	$s_1 \geq 0$	$s_1 < 0$
1	1	$s_0 + s_1 - s_0 s_1$	$s_0(1+s_1)$
1	0	$1-(s_0+s_1-s_0 s_1)$	$1-[s_0(1+s_1)]$
0	1	s_0	s_0
0	0	$1-s_0$	$1-s_0$

Particle Filters

Assume we have a sequence of unobserved latent variables $z_1,\ldots, z_t$ where $\mathbf{z}_{0:t}$ is modeled as a Markov process and each z holds a pair of strength estimates s_0 and s_1. Additionally, we have a sequence of observations $y_1,\ldots, y_t$ representing the covarying events. The posterior distribution $P(\mathbf{z}_{0:t}|\mathbf{y}_{1:t})$ can be obtained recursively as:

$$P(\mathbf{z}_{0:t+1}|\mathbf{y}_{1:t+1}) \propto P(\mathbf{z}_{0:t}|\mathbf{y}_{1:t})P(y_{t+1}|z_{t+1})P(z_{t+1}|z_t)$$

Importance sampling can be used recursively to approximate this distribution by sampling from $P(z_{t+1}|z_t)$ for each value of z_t, weighting each value of z_{t+1} by $P(y_{t+1}|z_{t+1})$, and then resampling from this weighted distribution. This algorithm, in which a set of samples is constantly updated to reflect the information provided by each observation, is known as a particle filter. The samples are referred to as particles.

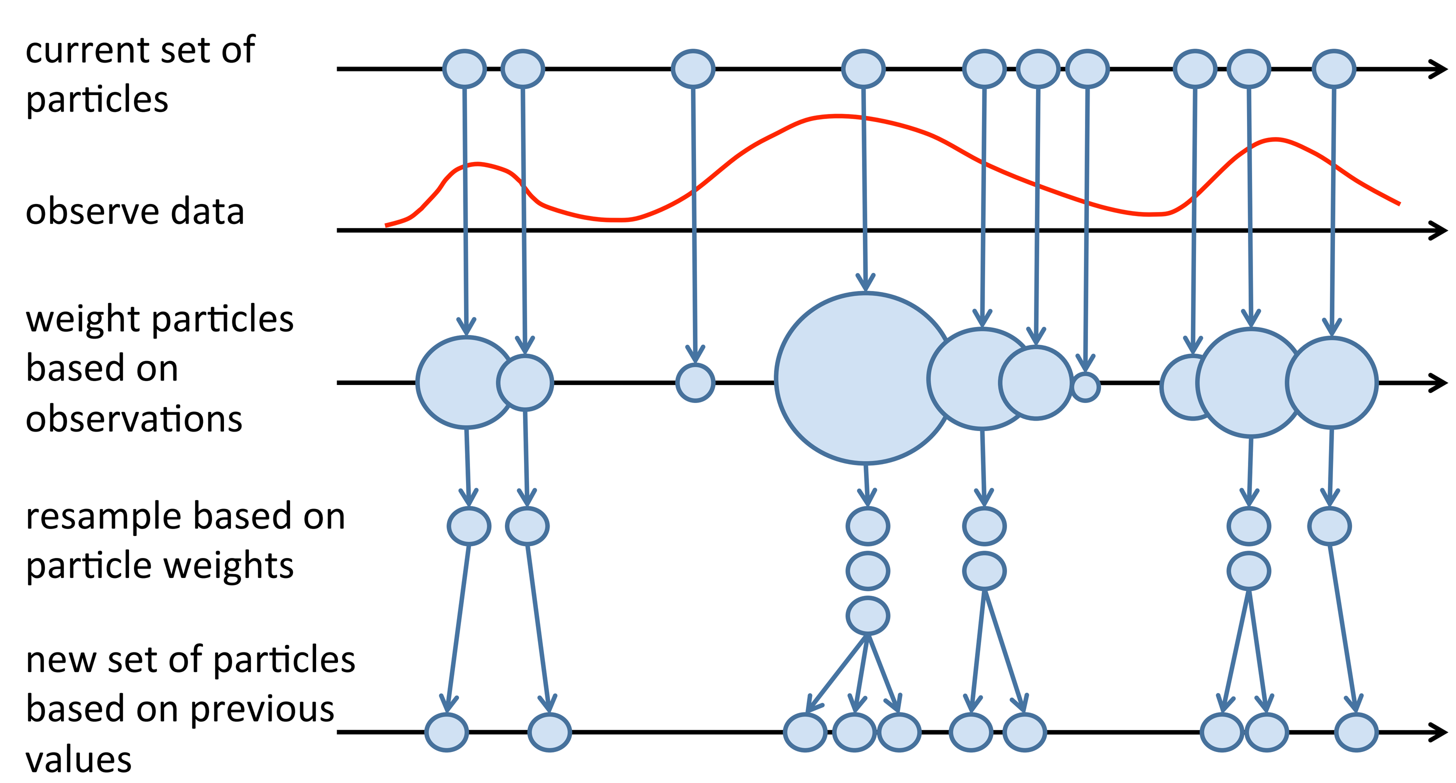

Order Effects and Particle Filter Parameters

How do different methods of resampling influence order effects?

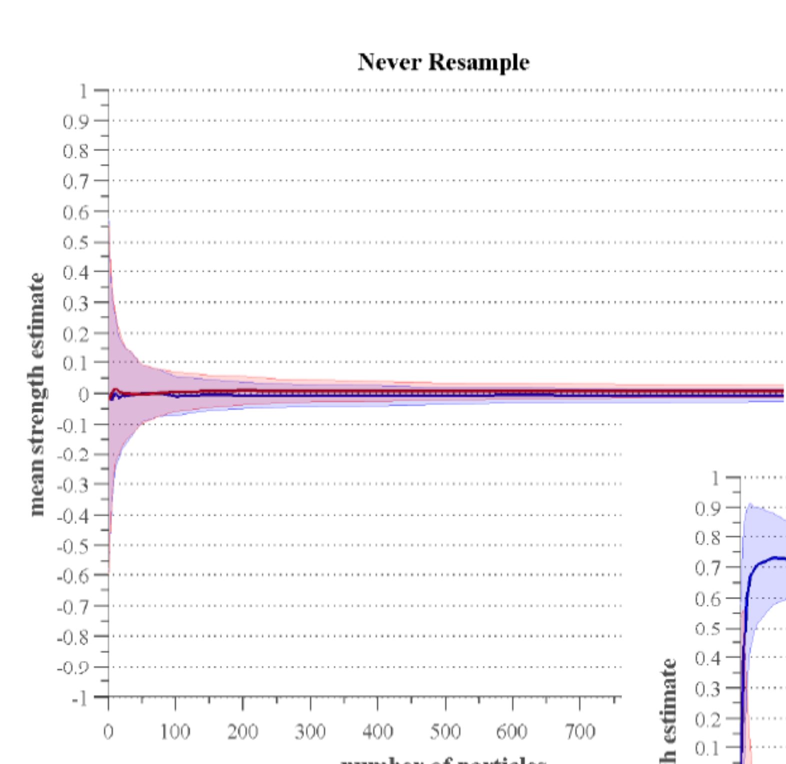

Model 1: Never resample – let the particle weights determine the final strength estimates

generative-preventative stimuli (blue curves)
preventative-generative stimuli (red curves)

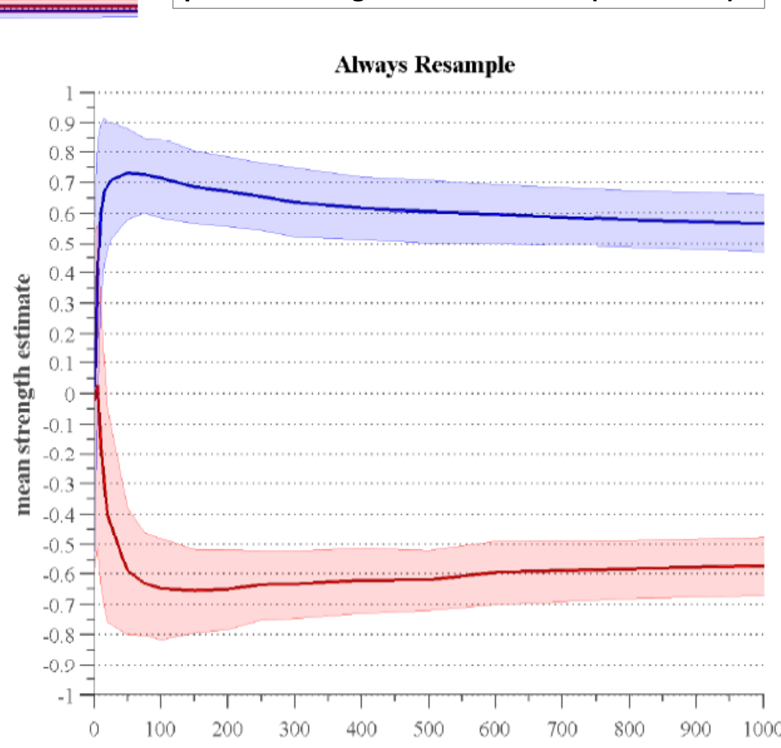

Model 2: Always resample – based on a multinomial distribution defined on the particle weights and at every trial t

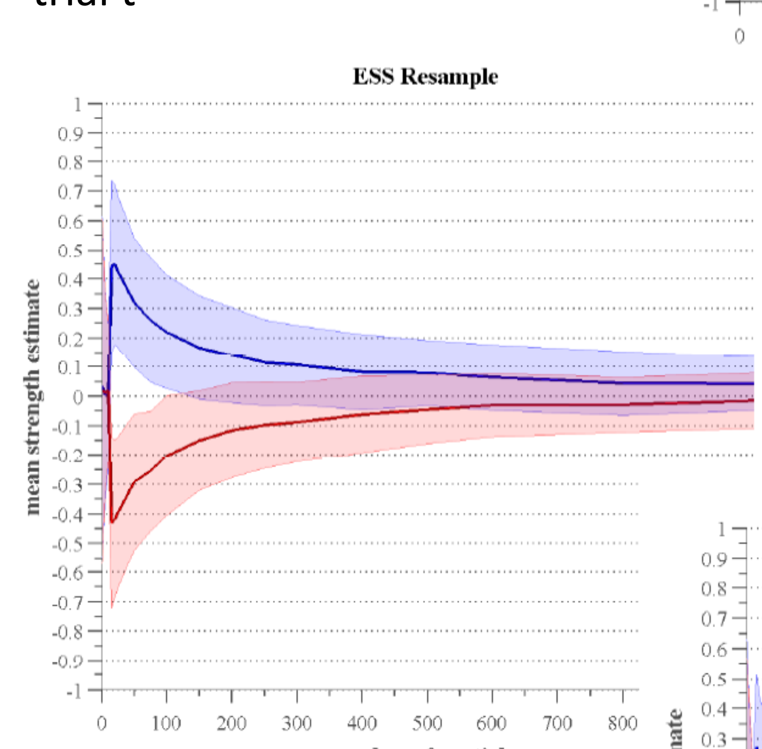

Model 3: ESS resample – based on a multinomial distribution defined on the particle weights and only if the variance of the weights is too large as defined by the Effective Sample Size: ESS = $||w_t||^{-2}$

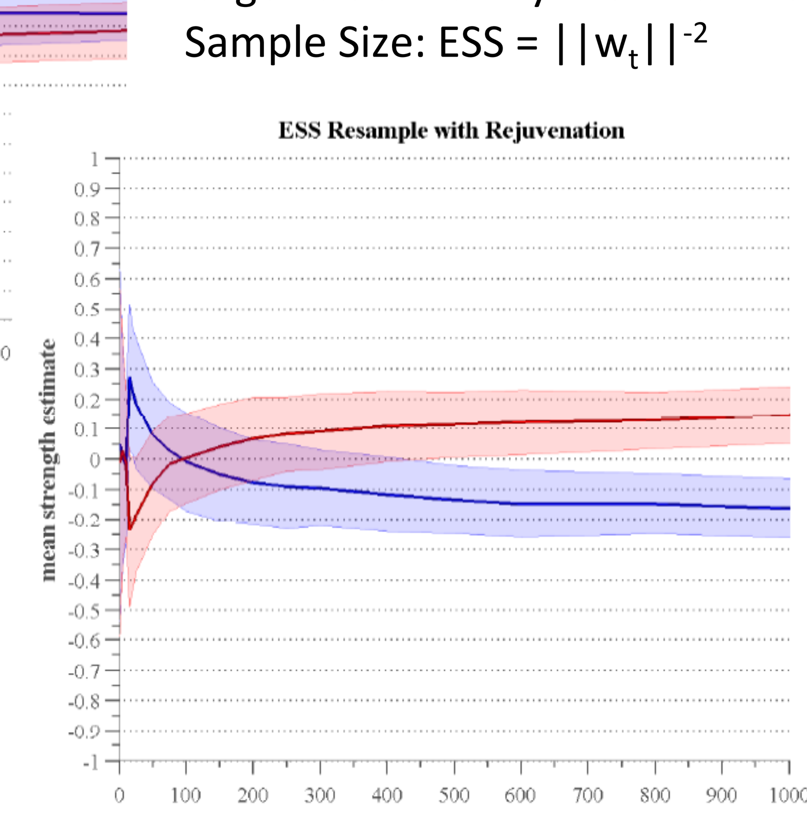

Model 4: ESS resample with rejuvenation – same as Model 3 but with ten iterations of Metropolis-Hastings immediately after resampling.

A Closer Look at Rejuvenation

How are we getting a primacy effect when we don't rejuvenate and a recency effect when we do?

We get a better understanding of the predictions of Models 3 and 4 by focusing on the predictions of 50 particles at each trial *t*.

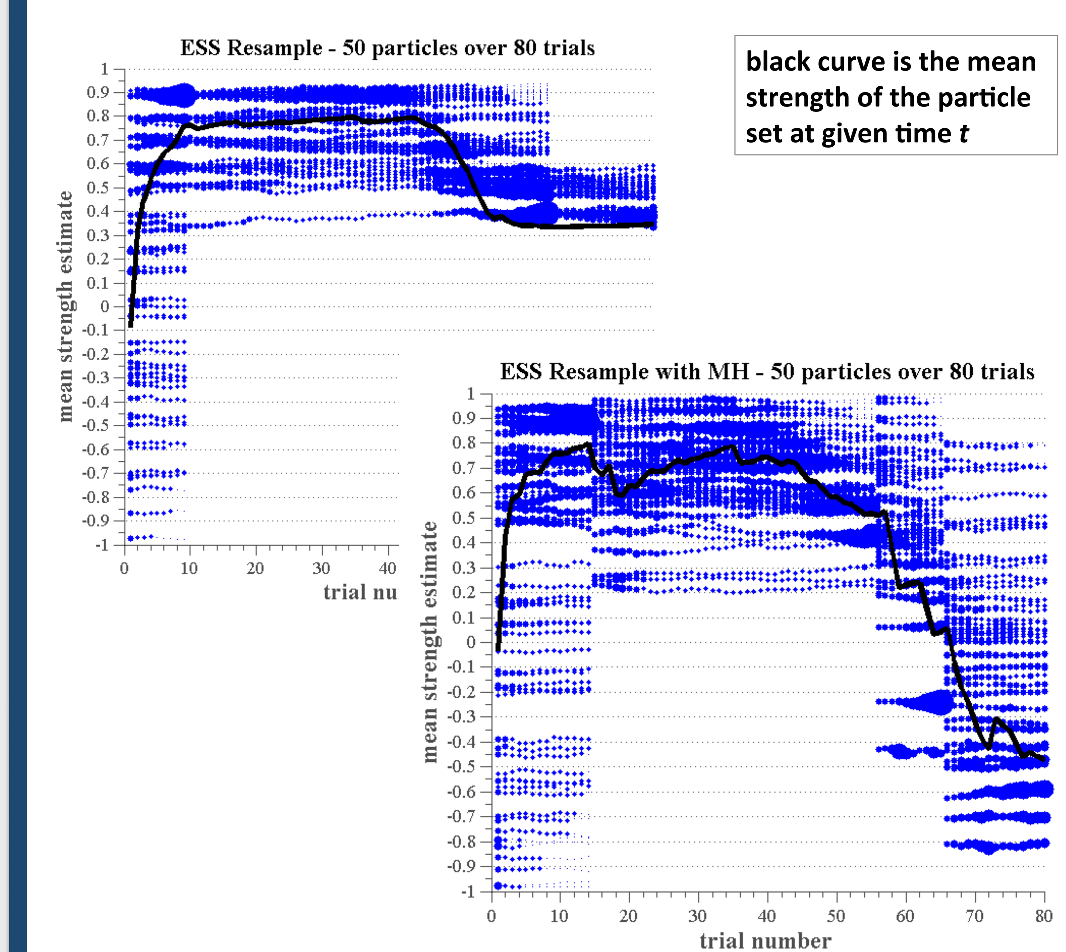

In Model 3, the diversity of the particle set narrows after resampling, resulting in a primacy effect.

In Model 4, the diversity is much broader after the MH rejuvenation step, producing a recency effect.

Modeling Human Data

Using Model 3 with a stronger prior for generative causes, we obtain results similar to reported primacy effects with a small number of particles.

If we add the MH rejuvenation step after every 10 trials, we additionally obtain results similar to reported recency effects with a small number of particles.

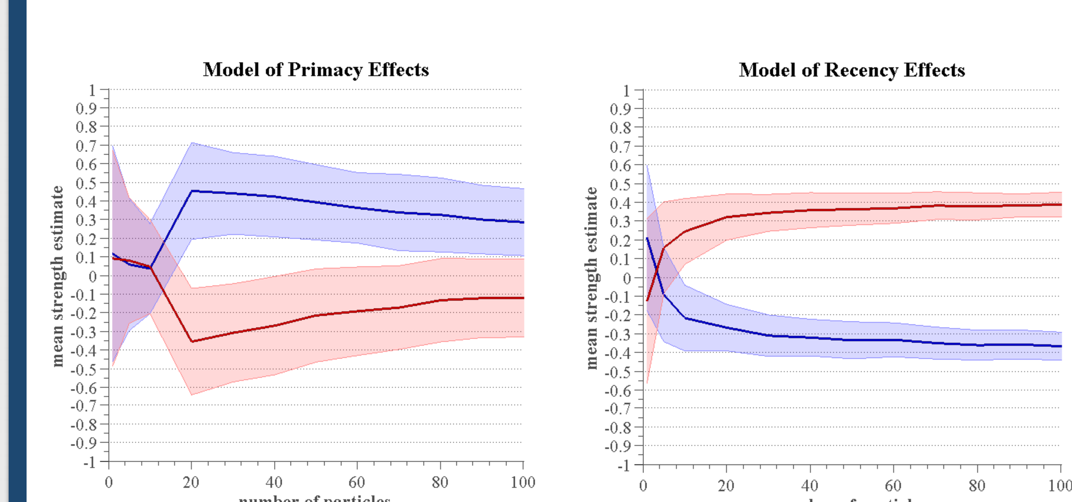

Conclusions

- Different resampling methods in a particle filter can produce different order effects in a causal learning task and provide a more consistent explanation of observed order effects in behavioral data.
- Two key elements interacting: *filtering*, in which we observe one data point at a time, and *rejuvenation*, in which we consider all previously observed data. This interaction may explain why people produce order effects.

References

- Collins, D., & Shanks, D. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition*, 30(7), 1138.
- Dennis, M., & Ahn, W. (2001). Primacy in causal strength judgments:The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*, 29(1), 152.
- Griffiths, T., & Tenenbaum, J. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51(4), 334–384.