Learning Hidden Causal Structure from Event Sequences

Simon Valentin, Neil R. Bramley, Christopher G. Lucas
University of Edinburgh

People have an impressive ability to infer the hidden causal structure behind everyday observations, often basing this on sparse, ambiguous signals that may be far apart in time and space. Much of the empirical research on causal structure learning has focused on static presentations of stimuli and modeling approaches that exploit a statistical-contingency-based causal graphical model framework (Pearl, 1988). However, in their daily lives, people do not encounter contingencies directly, but rather events occurring over time, often without the auxiliary information that would allow one to build a contingency table. There has been a renewed interest in the role of time in causal structure learning (e.g., Bramley et al., 2018), but early work goes as far back as Michotte (1946).

Past research in this area has shown that people readily use temporal information to infer causal relationships among observed events, based on both temporal order (Rottman & Keil, 2012) and delay information (Bramley et al., 2018). However, little is known about how and when people use temporal information to discover hidden causes and causal cycles, despite studies showing that people – including children as young as 10 months – can use covariation information and interventions to learn about hidden causes (Kushnir et al., 2010; Lucas et al., 2014).

We build on prior work (Valentin et al., 2020) to study how people use order and delay information as low-level perceptual cues to infer both the hidden causal structure any hidden causes behind sequences of events. Here, the causal structures considered may include a common hidden cause or cyclic relationships, as presented in Fig. 1.

Methods Our experimental task was to learn the causal structure governing sequences of events of two types, and what forms those relationships take. The sequences were sampled from a generative model based on causal connections with gamma-distributed delays. For each 35s video of events, participants made a forced-choice judgment as to which causal structure they believe generated the data, and an associated confidence. In the first experiment ($N = 50$), all participants were given one cover story (bioluminescent bacteria), which we replicated and extended with three new cover stories in a second experiment ($N = 200$), in order to explore the influence of different context-specific priors.

Our computational models are based on probabilistic finite state machines in the case of the order model, and dynamic Bayesian networks over gamma-distributed event delays in the case of the delay model. Here, the delay model provides an ideal-observer model comparison to human judgments.

Results and discussion Overall, our findings support the idea that people rely on the delays between events rather than order information alone to identify hidden causal structure from low-level temporal information. In particular, participants were able to infer the presence of a common hidden cause from the order and timing of event sequences. This finding extends previous work on fully observed and acyclic structures, which showed that people use order information to rule out incompatible causal structures, and the duration and variability of causal delays to make more fine-grained judgments (Bramley et al., 2018).

Across different cover stories, participants appear to combine their prior beliefs with the observed evidence in a manner that is broadly consistent with Bayesian updating. Meanwhile, the computational problem of inferring causal relationships from temporal data is challenging, and is only exacerbated by the possibility of hidden causes. We argue that exploring process models is an exciting next step in understanding how people solve the inductive problems considered here.

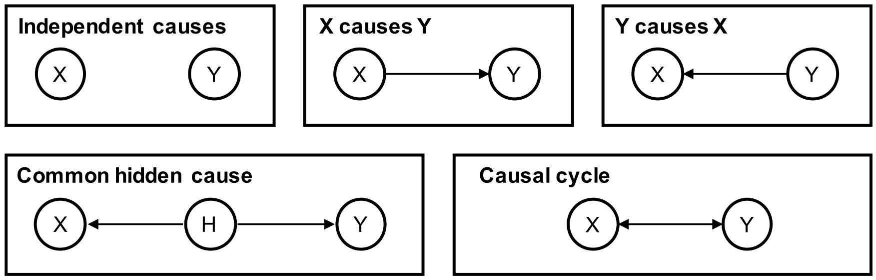

Figure 1: Causal structures considered.

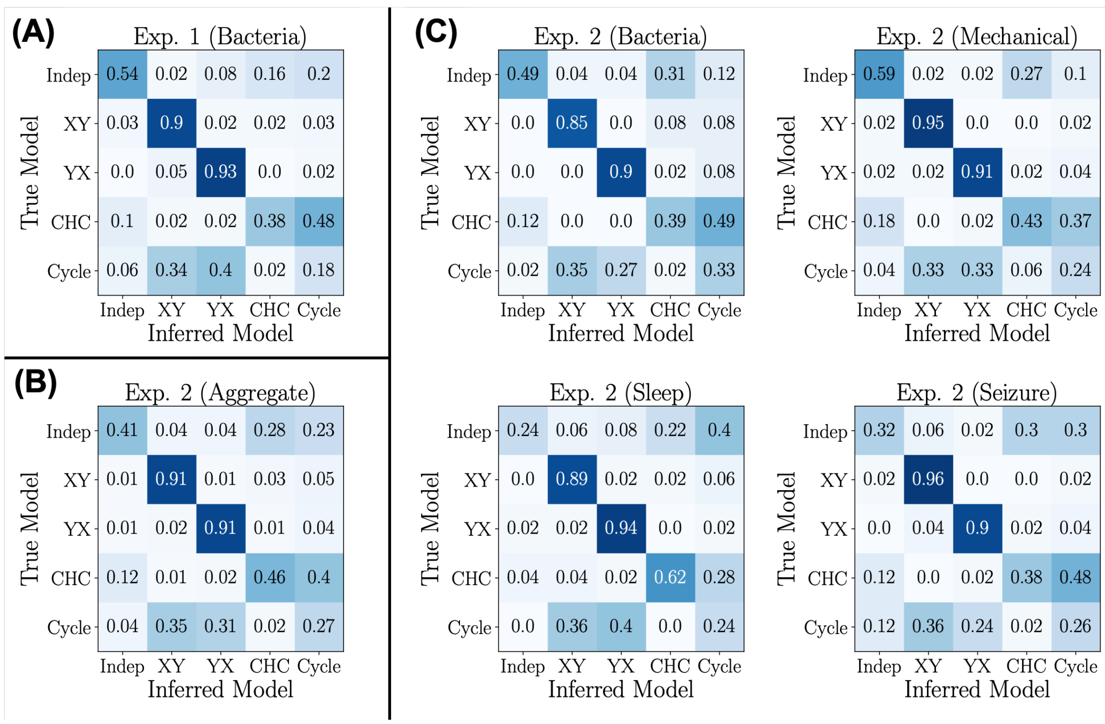

Figure 2: Confusion matrices for people's structure inferences in experiments 1 (A) and 2, in aggregate (B) and per cover-story (C).

References

Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (2018). Time in causal structure learning. *Journal of Experimental Psychology: Learning Memory and Cognition*, *44*(12), 1880–1910.

Kushnir, T., Gopnik, A., Lucas, C., & Schulz, L. (2010). Inferring hidden causal structure. *Cognitive Science*, *34*(1), 148–160.

Lucas, C. G., Holstein, K., & Kemp, C. (2014). Discovering hidden causes using statistical evidence. *Proceedings of the Annual Meeting of the Cognitive Science Society*.

Michotte, A. (1946). *La perception de la causalité*. Louvain: Publications Universitaire.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann Publishers Inc.

Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, *64*(1-2), 93–125.

Valentin, S., Bramley, N., & Lucas, C. (2020). Learning hidden causal structure from temporal data. *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*.