

Causal Beliefs Influence the Perception of Temporal Order

Philip M. Fernbach (philip_fernbach@brown.edu)
Preston Linson-Gentry (preston_linson-gentry@brown.edu)
Steven A. Sloman (steven_sloman@brown.edu)

Brown University, Department of Cognitive and Linguistic Sciences, Box 1978
Providence, RI 02912 USA

Abstract

Philosophers and scientists have long recognized how useful people find the temporal order of events as a guide to uncovering the causal relations among those events (e.g., Hume, 1739; Lagnado & Sloman 2004, 2006; White, 2006). In this paper we provide evidence for the converse, that beliefs about causation influence the perception of temporal order. Participants that learned the structure of causal relations among a set of slider bars on a computer screen were biased to report the movements of the sliders as conforming to a temporal order consistent with their causal beliefs despite the absence of a correlation between temporal order and causal structure in the stimuli.

Keywords: Temporal Perception; Causal Learning;; Intervention; Causal Models

Introduction

Consider a war-zone reporter that has witnessed a confrontation between two enemies breaking a cease-fire agreement. There is disagreement about who fired the first shot, and this issue has diplomatic implications. Suppose that the reporter considers one group to be the aggressor in the conflict and the other to be defending itself. The reporter is ethical and always reports precisely what he observes, and in this case he claims that the aggressive group fired first. Is it possible that his perception of the events was biased and that the defensive group actually fired the first shot? Can beliefs about the causal relations responsible for generating events change the perception of those events?

It is well known that beliefs about the lexical content of an utterance can change how the sound stream is perceived (Warren, 1970), and the predictability of an outcome can affect the sensory experience of that outcome (e.g., Blakemore et al., 2002; Witney et al., 1999). Cognitive psychologists have long recognized that perception is not determined wholly by the stimulus but also by prior knowledge, expectations, and other contextual factors (e.g., Bruner, 1974). But is this true when understanding is abstract and the relevant events involve a salient sensory experience?

Recent work (Haggard & Clark 2003; Haggard, Clark, & Kalogeras, 2002) has shown that auditory perception can be modulated by the intention to act. Someone who believes that he or she has caused an auditory event reports that event to have occurred earlier than it actually did, and earlier than someone who does not believe he or she was the cause. Moreover, the putative cause is reported to have occurred

later than it did. The perceived time of two events converges when the first is known to be the cause of the second.

In this paper we consider whether causal beliefs affect time perception in the absence of actions that control the events. There's good reason to believe the converse. Recent studies have shown that people put substantial weight on temporal cues when inferring causal structure from observations of events (e.g., Lagnado & Sloman, 2006), but little work addresses whether beliefs about causal structure mediate temporal perception.

We taught people that events conformed to a specified causal structure and then showed the events in a temporal order that did or did not correspond to that causal structure. We examined whether the temporal order that people reported was biased by the causal beliefs that were induced. Unlike much previous work, the events were visual (movements of sliders on a computer screen), not auditory. Nevertheless, they occurred quickly enough to leave participants with some uncertainty about temporal order. Faced with uncertainty, we expected responses to be informed by prior beliefs about what to expect, beliefs in the form of causal models. In other words, we predicted that beliefs about the causal structure relating visual events would affect the order in which those events were perceived.

Experiment

To assess the question of whether beliefs about causal structure influence the perception of temporal order, we conducted an experiment in which we induced beliefs about a causal system, a set of three causally dependent slider bars. Participants were then presented with the sliders moving in temporal orders that were uncorrelated with the induced causal beliefs and reported the temporal orders that they observed. The temporal delays were short enough that the task was non-trivial and we hypothesized that causal beliefs would influence judgments given uncertain observations.

To induce causal beliefs we used an interventional learning task that forced participants to interact with the sliders to infer the true causal structure. Because of the ease of learning with intervention in this paradigm (Lagnado & Sloman 2004, 2006), we expected this task to produce salient causal beliefs.

Method

Participants. Participants were 17 Brown University students recruited by an internet advertisement and paid 8 dollars for participation. One participant was excluded from analysis due to computer failure.

Stimuli. Stimuli were slider bars and arrows presented on a standard computer display. During the intervention and observation portions of the experiment the slider bars were arrayed in an equilateral triangle and labeled A, B and C with A at bottom left, B at the top and C at bottom right. Each slider bar was 2.5 cm in width and 7 cm in height and had two positions, up and down. During the model selection portions of the experiment, 6 grey unidirectional arrows connecting each of the sliders to the others were present. These arrows represented each of the possible causal relations in the system and were used by participants to report their inferences after making interventions. The arrows were each 6 cm long, and turned red when clicked by the participant.

Design All participants viewed the same causal models and the same temporal orders for each. Models were presented to each participant in a random order. Each of the possible three variable acyclic two-link causal models was tested, for a total of 12 models: 6 chain models, 3 common cause models and 3 common effect models (Figure 1).




Figure 1: The 12 causal models tested in the experiment.

In the observation portion of the experiment, ten temporal orders of slider movement were presented for each model (Table 1), also in random order.

Table 1: The ten temporal orders presented for each model

Temporal Order	Moving First	Moving Second	Moving Third
111	A B C	-	-
112	A B	C	-
121	A C	B	-
211	C B	A	-
123	A	B	C
132	A	C	B
213	B	A	C
231	C	A	B
312	B	C	A
321	C	B	A

*The numbers associated with the temporal order denote the temporal position of A, B and C respectively. For example, 123 means that A moves first, B second and C third.

Procedure Participants were told that they would be participating in a visual task involving the computer and were given a chance to put on any corrective lenses. A maximum of two participants were run at a time on different computers. Introductory instructions were presented on the screen:

Thank you for participating in this experiment. Please read these instructions carefully. In a moment you will be presented with a set of sliders. These sliders move up and down and have hidden connections between them. This means that moving one slider might or might not cause one or both of the other sliders to move. Your job is to figure out how the sliders are connected.

Participants were then presented with three sliders arrayed in an equilateral triangle and all in the down position (Figure 2). For each causal model, participants were given the opportunity to make as many interventions as they wanted until they believed they understood the underlying causal relations. Causal relations were deterministic, so when a slider was intervened on all of its effects were necessarily activated. Participants intervened on a slider by clicking the top portion of it with the mouse. When an intervention occurred, all activated sliders moved simultaneously, remained in the up position for 1 second and then returned to the down position. When ready to report the causal model, the participant clicked a button marked 'ready'.




Figure 2: The screen as it appeared during the intervention task.

When the 'ready' button was clicked, a new screen appeared with the same three sliders and arrows from each slider pointing to the other two (for example the A slider had an arrow pointing from it toward the B slider and another arrow pointing from it toward the C slider) for a total of six arrows, each representing a potential relation from cause to effect (Figure 3).




Figure 3: The GUI for inputting intervention inferences.

Participants were instructed to highlight the arrows by clicking on them to indicate how they thought the sliders were connected. When clicked, the arrows changed from grey to red. When the participant had chosen all of the arrows he or she wished and clicked the 'finished' button, text appeared on the screen indicating whether the model was right or wrong. If the participant's choice was incorrect, he or she was returned to the intervention task to try again. This process iterated until the participant chose the correct model. Once this occurred another set of directions appeared:

That is Correct!!! Next you will be presented with the same three sliders. This time they will move on their own. These sliders have the same connections as the ones you just figured out, however the order in which they move may not reflect those connections. This time your job will be to determine the order in which the sliders moved. You will see the slider move 10 times. After each time you will indicate the order in which they moved. Sometimes they will move right away, other times there will be a delay, so pay close attention.

In the temporal order task, the same sliders were presented along with the correct directional arrows in red indicating the true causal model. These arrows remained on the screen while participants completed the temporal order task to ensure that the most current causal model was salient (Figure 4).




Figure 4: The screen as it appeared during temporal order presentations.

Upon initiation of a temporal order, sliders moved one, two, or three at a time. As in the intervention task, these moved from the down position to the up position. The initiation of movement occurred at randomized intervals by sampling the onset time from a normal distribution. When sliders did not move simultaneously there was a delay of 100 milliseconds between movements. This delay was chosen to make the task difficult enough to avoid ceiling effects but not so difficult that participants were responding at chance.

Once all the sliders had moved a new screen appeared for the participant to input his or her response. This screen had three horizontal sliders marked A, B, and C (Fig 5).




Figure 5: The GUI for reporting temporal orders.

Participants placed each slider in one of three spatial positions to indicate the order in which they saw the sliders on the previous screen move. The position furthest to the left represented the first movement, the middle position the second movement, and the furthest right represented the third movement. If two sliders moved at once participants placed the corresponding sliders in the same position. If they thought all sliders moved simultaneously they placed all the sliders at the furthest left position. After indicating their choices they clicked a button that took them to the next temporal order. Once all 10 temporal orders were completed for a causal model, participants repeated the intervention task for the next causal model.

Upon completion of all causal models, participants were given a questionnaire sheet with the following questions: "While completing the experiment, did you rely more upon temporal cues or the connectivity of the sliders to guide you in reconstructing the order? While completing the experiment, did you think the connectivity of the sliders that you found was affecting the order in which they were moving?" Further clarification was given upon request.

Analysis

Analysis of participant responses was conducted by assessing the extent to which incorrect responses were biased in the direction of the temporal order implied by the relevant causal model. For each response there were three variables: the temporal order presented, the temporal order reported by the participant, and the causal model learned prior to the set of observations. For analysis, each of these variables is

represented as a three-element vector denoting the temporal relations. The first element represents the temporal relation between A and B, with a 1 meaning that A occurs before B, a 0 meaning that A and B occur simultaneously, and a -1 meaning that B occurs before A. The second element represents the A-C relation, and the third represents the B-C relation. We represented causal models via temporal relations derived from them. To derive the temporal order implied by a causal model, we assumed that causes precede effects. For instance, if in causal model X , A is a cause of B, then A should precede B and the value of the first element of the vector representing X is 1. A 0 in a causal model representation means that the model does not specify the variables' temporal relation as in the case of causes of a common effect or effects of a common cause. Figure 6 shows the vector representations of all of the causal models and temporal orders tested in the experiment.

Table 2: Vector Representations

Temporal Orders	A-B Link	A-C Link	B-C Link	Causal Models	A-B Link	A-C Link	B-C Link
111	0	0	0	ABC Chain	1	1	1
112	0	1	1	ACB Chain	1	1	-1
121	1	0	-1	BAC Chain	-1	1	1
211	-1	-1	0	BCA Chain	-1	-1	1
123	1	1	1	CAB Chain	1	-1	-1
132	1	1	-1	CBA Chain	-1	-1	-1
213	-1	1	1	ABC Common Cause	1	1	0
231	1	-1	-1	BAC Common Cause	-1	0	1
312	-1	-1	1	CAB Common Cause	0	-1	-1
321	-1	-1	-1	BCA Common Effect	-1	-1	0
				ACB Common Effect	1	0	-1
				ABC Common Effect	0	1	1

Given these vector values analysis proceeded in two steps. First the difference between the reported and presented temporal orders was calculated by subtracting the vectors

$$\vec{D} = \vec{R} - \vec{P} \quad (1)$$

where \vec{D} is the difference vector, \vec{R} is the order reported by the participant and \vec{P} is the order presented.

\vec{D} represents the departure of the participant's response from the true temporal order. Next the cosine of the angle between the difference vector and the vector representing the prescribed temporal relations of the causal model that generated the data is calculated.

$$\cos \theta = \frac{\vec{D} \cdot \vec{C}}{\|\vec{D}\| \|\vec{C}\|} \quad (2)$$

where θ is the angle between the vectors, \vec{C} is the vector representing the temporal relations of the causal model, and $\|\vec{D}\|$ and $\|\vec{C}\|$ are the magnitudes of the difference vector and causal model vector respectively in the Euclidean norm. The cosine of the angle between the difference vector and the causal model provides a measure of the extent to which an incorrect response is in the direction of the temporal relations implied by the causal model. It varies from 1 to -1 with 1 implying that the response is highly consistent with the causal model, -1 implying that the response is highly inconsistent with the causal model and 0 implying that the response is neither consistent nor inconsistent with the model.

Correct responses result in a difference vector of magnitude zero as the presented and reported orders are the same. Since the direction of such a vector is not defined, a cosine cannot be calculated and the response is not included in the analysis. Thus the average cosine between difference vectors and causal models for incorrect responses serves as the dependent measure. If there is no bias the average cosine should not be significantly different from zero as consistent and inconsistent responses should be equally likely. More formally, if all participants sample responses from the set of all possible temporal orders independently of causal model, then the average expected value of the cosine between \vec{C} and \vec{D} across participants, causal models, and temporal orders can all be shown to equal zero.

Results

Figure 6 depicts the average cosine of the angle between the difference vector and the causal model for the results for each participant, temporal order presented, and causal model learned.




Figure 6: Average cosine of the angle between the difference vector and the causal model by participant (a), temporal order presented (b) and causal model learned (c).

Single sample t-tests revealed average cosine values significantly greater than zero, implying a bias in favor of the temporal relations implied by the model, across participants (mean = .059; SE=.024, $t_{15}=2.49$, $p=.025$) and presented orders (mean=.083; SE=.027, $t_9=3.08$, $p=.013$). The effect across causal models was in the right direction but not significant (mean = .057; SE=.033, $t_{11}=1.75$, $p=.112$).

Correct responses were 52.6% of total responses with a range of 83.3% to 12.5% over participants, and a standard deviation of 17.3%. Linear regression revealed that percent correct was uncorrelated to degree of bias ($r^2=.09$, $F_{1,14}=1.40$, $p=.256$). Causal model learned had no significant effect on percent correct. Temporal order was generally unrelated to percent correct though there were two exceptions: The temporal order in which all three sliders moved simultaneously was substantially easier than the others, with 94% correct responses. The order in which B and C moved simultaneously and then A moved last was substantially harder than the rest with only 22% correct.

Beyond the bias for reported orders to be consistent with causal models, there was a general bias to treat A as occurring before B and C and to treat B as occurring before C. This can be illustrated by looking at the average values of the vectors reported by participants ($\bar{\mathbf{R}}$). Since an equal number of temporal orders were presented in which any variable preceded any other variable, the average value of the presented vectors ($\bar{\mathbf{P}}$) was 0. However, all three elements of participants' average reported vectors were greater than 0. The average values of the elements representing the A-B, A-C and B-C temporal relations were .07, .12 and .03 respectively implying a bias toward A in the A-B relation which was highly significant ($t_{1919}=3.36$, $p<.001$), toward A in the A-C relation which was highly significant ($t_{1919}=5.79$, $p<.001$), and toward B in the B-C relation which was marginally significant ($t_{1919}=1.70$, $p=.089$).

All 16 participants indicated that their responses were based on the perceived temporal order and not on the causal structure.

Dicussion

In summary, this experiment suggests that the temporal order of events that people report is affected by beliefs about the causal model generating those events. When in a state of uncertainty about what happened, people take advantage of abstract beliefs about what should have happened to make a report. This is analogous to using prior distributional knowledge as information about when or how an event will occur (Sobel et al., 2004; Tenenbaum & Griffiths, 2001).

The critical open question is whether the bias we have reported is perceptual or operative at the response stage. Are causal models inducing people to see events in a different order or affecting how they decide to respond given that they are not sure of their answer? The fact that every one of the participants claimed that they were not using the causal model to inform them suggests that, if the effect is due to a process occurring during response, that process is not explicitly

available to participants. Participants think that they are reporting what they see whether or not they are.

One possible mechanism for how causal beliefs affected the responses of out participants is attentional. On this view participants used causal information to make predictions about the time course of events. If A was a root cause then they expected it earlier and attended to it first. This suggests that errors may have been driven by failures of attention. Such attentional failures are consistent with both the perceptual and response bias interpretations. On the perceptual view, events that are not attended directly are filled-in by prior beliefs. On the response-bias view, events that are not attended directly are more likely to be guessed using abstract knowledge at response time. Previous work has shown that attention can modulate temporal order perception (Stelmach & Herdman, 1991; Haddad, Carreiro & Baldo, 2002), a finding that supports the perceptual interpretation. Nonetheless a definitive answer requires further work.

The presence of a general bias to treat A as occurring before B and C and B as occurring before C suggests that causal beliefs are not the only source of information drawn on to resolve uncertainty. Other sources might include linguistic considerations, the spatial arrangement of the sliders, and the spatial arrangement of the user interface for responding.

Our study suggests that not only do temporal cues influence the learning of causal structure, but causal knowledge influences the perception of temporal cues. This could be taken as evidence that learning theories that rely on temporal cue information are circular. We take it instead as evidence that, like so many learning processes, causal learning involves bootstrapping between uncertain abstract knowledge gleaned from the past with uncertain perceptions about the present.

Acknowledgments

This work was supported by NSF Award 0518147. We thank David Buchanan, Adam Darlow, Jessica Greenbaum, Jared Hotaling, Dave Lagnado, Ju-Hwa Park, and John Santini for helpful discussion and input.

References

- Frith, C.D., Blakemore, S.J. & Wolpert, D.M. (2002). Abnormalities in the control and awareness of action. *Trends in Cognitive Science*, 6, 237–242
- Bruner, J. (1974). *Beyond the Information Given: Studies in the Psychology of Knowing*. J. Anglin (ed.), London: Allen and Unwin.
- Haddad Jr., H., Carreiro L. R. R. & Baldo, M. V. C. (2002). Modulation of the perception of temporal order by attentional and pre-attentional factors. *Brazilian Journal of Medical and Biological Research* 35: 979-983.
- Haggard, P. & Clark, S (2003). Intentional action: Conscious experience and neural prediction. *Consciousness and Cognition*, 12, 695-707.
- Haggard, P, Clark, S & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, 5 (4), 382-385.

- Hume, D. (1739), *A Treatise on Human Nature*. New York: Oxford University Press.
- Lagnado, D. & Sloman, S.A. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 856-876.
- Lagnado, D.A. & Sloman S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory and Cognition* 32, 451-460.
- Sobel, D. M., Tenenbaum J.B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, 28, 303-333.
- Stelmach, L. B. & Herdman, C.M. (1991). Directed attention and perception of temporal order. *Journal of Experimental Psychology: Human Perception and Performance*, 17(2), 539-550.
- Tenenbaum, J., & Griffiths, T. L. (2001). Structure learning in human causal induction. *Advances in Neural Information Processing Systems*, 13, Cambridge, MA.
- Warren, R. M. (1970). Perceptual restorations of missing speech sounds. *Science*, 167, 392-393.
- Witney, A.G., Goodbody, S.J. & Wolpert, D.M. (1999) Predictive motor learning of temporal delays. *Journal of Neurophysiology*. 82, 2039-2048.
- White, P. A. (2006). How well is causal structure inferred from cooccurrence information? *European Journal of Cognitive Psychology*, 18 (3), 454-480.