Time Reordered: Causal perception Guides the Interpretation of Temporal Order

Christos Bechlivanidis & David A. Lagnado

University College London

Abstract

We present a novel temporal illusion in which the perceived order of events is dictated by their perceived causal relationship. Participants view a simple Michotte-style launching sequence featuring 3 objects, in which one object starts moving before its presumed cause. Not only did participants re-order the events in a causally consistent way, thus violating the objective temporal order, but they also failed to recognise the clip they had seen, preferring a clip in which temporal and causal order matched. Moreover, we show that the effect is not due to lack of attention to the presented events. In explaining the effect, we argue that low-level cues to causality detected in a dynamic sequence as a whole trigger the search for causal hypotheses, selecting the simplest ones even at the expense of ignoring or altering sensory input.

Introduction

Imagine watching a long queue of dominos falling one after another. Apparently, each domino's fall is causing the fall of the next one in the queue. But suddenly one domino falls early, without being touched by the previous domino (Fig. 1). Would you see the domino's premature fall?

Figure 1: The temporal order (the 3/3 domino falls before the 2/2) does not match the assumed causal order (the 2/2 domino causes the fall of the 3/3). How do we resolve this incongruence?

Whether or not causal impressions can influence the experienced temporal order depends on two questions: (1) Can the perception of temporal order be influenced by information other than the order of the percepts themselves, and (2) do causal impressions possess those features necessary to influence presumably lower level percepts?

For some philosophers (Hoerl, 2013; Phillips, 2014; Soteriou, 2010), the answer to the first question is negative: the order of our experiences mirrors or inherits the temporal structure of the environment. Thus, to experience event A happening before B, we must perceive that particular temporal order, even if it is illusory, such as when a thunder is seen before it is heard. Others, however, have argued for a more constructed view of temporal order (Dainton, 2010; Grush, 2007), i.e. for temporal order as a second-order judgement.

Despite this ongoing debate, experimental evidence from multisensory integration suggests that experienced temporal order is in fact malleable: when two successive bimodal stimuli are assumed to originate from the same source, the perceived timing of each stimulus is shifted so that the two events appear simultaneous (King, 2005; Spence & Squire, 2003). However, unlike the domino example above, the order here is collapsed rather than reversed, and, most importantly, the bimodal nature of the stimuli and the requirement for sensory integration introduce additional complications and allow for multiple interpretations.

With regard to the second question, the role of causality in perception, several recent studies show that judgments of spatial relations (Scholl & Nakayama, 2004), size (Buehner & Humphreys, 2010) trajectory (Kim, Feldman, & Singh, 2013) and, more relevantly, temporal duration (Buehner, 2012; Humphreys & Buehner, 2010) are sensitive to impressions of causality. In most cases, however, causal beliefs led to quantitative shifts, whereas reversing the order of events requires a stronger qualitative change.

Nevertheless, Bechlivanidis and Lagnado (2013) show that recently acquired causal knowledge can switch temporal order judgements. In their study participants played a computer-based puzzle game which required learning a novel causal relation between two events. After training, participants observed events happening in an order that violated the causal order of the learned relationship. When asked to report the order they saw, the majority preferred the order that matched their acquired causal beliefs, thus distorting the objective order of presentation.

Does this reordering effect depend on recently learned causal rules or does it generalize to any case where strong causal beliefs are present? Similarly, does it depend on inference or can instances of causal perception (Michotte, 1963) also result in the reordering of events? Finally, did participants simply report the most plausible order of events guided by learned causal rules or did they actually *perceive* a different order to that presented to them?

Experiment 1

To evaluate the sensitivity of temporal order judgments to contradictory causal impressions, we modified the classic Michottean launching sequence (Michotte, 1963) by adding a third object. Participants observed a three-object pseudo-collision where, critically, the third object starts moving before the second object, i.e. the effect takes place before its presumed cause

(Fig 2a). Following presentation, participants were asked to report the temporal order of the events they had just witnessed.

Furthermore, to ensure that the order of events is perceptually distinguishable, as well as using relatively low speeds and long delays between critical events, we included a control condition in which object A is not present (Fig. 2b). In this case, we expected participants to report the veridical order of events since in the absence of A the causal direction is ambiguous.
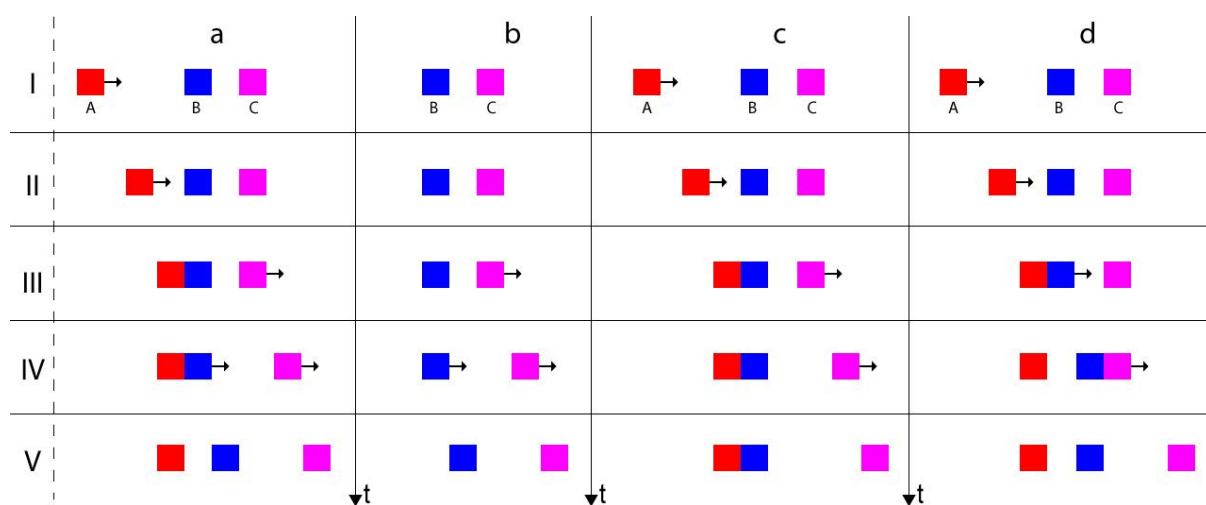

Figure 2: The sequences used in the two experiments: -a. Object A approaches B (I-II) and stops next to it at which point object C starts moving (III). After 350ms Object B starts moving (IV) and stops to the left of object C's original position (V), -b. Identical to sequence (a) without object A -c. Identical to (a) but object B remains stationary throughout –d. Realistic 3-object collision shown during the review question in Experiment 2 (the arrows show the direction of movement and were not visible in the experiments)

Participants and Materials

The experiment was programmed in Adobe Flex 4.6 and conducted over the Internet using Amazon Mechanical Turk (both experiments presented here can be seen at http://goo.gl/4noAmR).

Based on past experiments (Bechlivanidis & Lagnado, 2013) and piloting data (using different parameters, however) we expected about 40% of participants in the experimental group to report the temporal order of events. Although we had no data regarding the behaviour of the control group, given the low velocity of objects and the great temporal gap between the critical events we hypothesized that the majority (~80%) would report the

correct temporal order. Thus, with a two sided-significance between 0.05-0.01 and a power of 0.8, we calculated that about 27-38 participants would be required for each group.

On that basis, we recruited 60 participants in total. A single participant was excluded from the analysis for providing a nonsensical answer to the order question (see Design and Procedure), i.e. not identifying correctly the first object that started moving. Of the remaining 59 participants, 39 were male and 20 were female. The mean age was 32.39 (SD=9.96). Each participant was paid $0.50.

Although it is relatively uncommon to conduct perceptual experiments online due to variability in hardware setups, we took special precautions in order to achieve uniform presentation of stimuli. To prevent object size discrepancies due to variable monitor sizes and resolutions each participant went through a calibration session (programmed in DHTML+Ajax). The procedure required participants to place a physical optical disc (CD, DVD etc), credit card or dollar note (all participants were from the US) on their monitor in order to match the size of the respective virtual object that appeared on screen. Using the provided controls they could increase or decrease the size of the virtual object to match it against the physical object. Since the size of these physical objects is standard, we compared the known dimensions of the selected physical object against that of the virtual object as defined by the participant. This gave us the effective ppi (pixels per inch) that we used to adapt the sizes of the objects that appeared in the main experiment and thus achieve consistent object dimensions for all participants. To verify correct calibration participants had to use a physical ruler to measure the size of an horizontal line that was displayed subsequently. Entering an incorrect value forced participants to repeat the calibration session.

After successful calibration participants were asked to report their approximate physical distance from their monitor. The two options were: "I can more or less touch the screen if I extend my arm" and "I am further away from the screen". Only one participant picked the second option.

Regarding the potential temporal discrepancies, we recorded the frame rate while the critical sequence was presented as well as the actual delay between the onset of motion of objects C and B. There were minimal deviations: the average minimum frame rate was 29.70fps (target and maximum was 30fps) and the standard deviation was 0.36. Similarly the average actual delay was 352ms (target was 350ms) and the standard deviation 5.70. We nevertheless repeated the analysis that is reported in the results section after removing the single

participant who stated that she was seated far from the monitor and the 10 participants for whom the minimum frame rate dropped below 29fps at any point or for whom the recorded delay deviated by more than 10% (i.e. 35ms) from the 350ms target. The results of this stricter analysis were the same with those in the reported analysis.

Design and Procedure

The 59 participants were randomly assigned to one the two conditions that differed only in the displayed clip, resulting in 29 participants in condition 1 and 30 in condition 2.

After completing the calibration section (see Participants and Methods) participants were welcomed to the experiment and were asked for some simple demographic data. They were then informed that they would watch a short movie clip and answer some questions about it. They were also asked to be as focused as possible since the clip would be displayed only once.

Participants saw the clip a single time. In condition 1 ("A present") that is shown in figure 2a, three 8x8mm squares fade in slowly (2sec). Object A is located 35mm to the left of object B and object C is located 16mm to the right of object B. The squares remained static for another 2 seconds. Then object A starts moving to the right towards B at a speed of 30mm/sec, i.e. relatively slowly compared, for example, to the 300-400mm/sec in Michotte's original studies. Object A stops adjacent to B and, immediately after, object C starts moving also at 30mm/sec. After 350ms object B starts moving to the right at the same speed and stops to the left of object C's original position. After C travels 35mm it comes to a halt and the clip ends. The clip was designed to be as similar as possible to a normal 3-object collision with the exception of the order of events between B and C.

In condition 2 ("A absent") the clip that was shown was exactly the same but object A was not present (Fig. 2b). Specifically, in this condition, object C starts moving to the right at 30mm/sec and 350ms later B moves also to the right and stops next to C's original position. C travels for 35mm and the clip ends.

For each condition there were two similar versions of the target clip that differed only in the objects' colours and the direction of movement. In the "normal" version the colours were as shown in Figure 2a (red-blue-purple) and the direction of movement was left-to-right as described above. In the "mirrored" version the colours were A=purple, B=red, C=blue and the direction of movement was right-to-left, meaning that the initial position of objects was

mirrored compared to the "normal", i.e. A starts to the right of B and C to the left of object B. Participants in each condition were randomly assigned to one of the two clip versions.

After watching the clip a single time, participants were shown the initial configuration of the objects (i.e. Fig. 2a-I or Fig. 2b-I) and were asked to place the events in the order that they saw them. To do this they had to drag-and-drop the event sentences "The red square started moving" (only in condition 1) , "The blue square started moving" and "The purple square started moving" from their initial container to another box. The order of appearance of the sentences was randomised for each participant. Then participants were asked to indicate their confidence to the selected order by dragging a slider on a scale that was labelled "Not at all confident" to the left and "Very confident" to the rightmost position.

In the next screen, the initial object configuration was shown again and participants were asked for their causal impressions for all possible object pairs (six in condition 1 and two in condition 2). These were expressed by dragging on a slider labelled "Completely Disagree", "Neutral" and "Completely Agree" next to statements of the form "The X square made the Y square move", were X and Y were colour pairs (e.g. "The red square made the blue square move"). Finally, participants were asked for any comment they had regarding the experiment and they were thanked for participating.

Results

Figure 3, shows the proportion of participants that reported the objective temporal (A-C-B) or the causal order (A-B-C) of events (we collapsed the normal and mirrored versions of the clips since no difference was observed). The overwhelming majority (82.76%) preferred the causal order when A was visible while a similar majority (83.33%) preferred the objective temporal order when A was absent, despite the fact that in both conditions the behaviour of objects B and C was identical ($\chi^2(1, N = 59) = 25.77, p < .001$). Furthermore, participants in both conditions were very confident in the order they have reported, with mean confidence ratings 78.76/100 (SD=24.62) for condition 1 and 73.63/100 (SD=21.05) for condition 2.

Figure 3: Proportion of participants that reported the correct (temporal) or the causal order of events in each condition.

The direct causal judgments (Fig. 4) show that when A was present participants thought that it caused B to move (88.45%) and also that B caused C to move (77.28%). Participants were relatively indecisive about the A-C relationship (51.76%) but given the strong endorsement of the A-B and B-C relationships, one can assume that those endorsing it probably referred to the indirect A-C relationship, through B. The judgments for the inverse relationships were, as expected, very low. The C-B relationship is significantly higher than the C-A relationship (t(28)=2.305, p<0.05) and approaches significance compared with B-A (t(28)=2.007, p=0.054) but this is driven by those few participants who reported the correct temporal order of the sequence and any significant difference goes away if these participants are excluded.

Similarly, when A was not present (condition 2) and, thus, when the majority of participants reported the objective temporal order of events, the causal judgments were far weaker. The strongest causal belief is in C making B move. In fact significantly more participants endorsed the C-B causal relationship in condition 2 compared to condition 1 (t(57)=4.837, p=0.000).

Figure 4: Mean causal judgments for each object pair per condition (in Condition 2, object A was not visible, so there are no ratings involving it). A value of 50 corresponds to neutrality – lower ratings indicate disagreement and higher rating indicate agreement with the causal statement (Error bars represent 95% confidence intervals)

Discussion

The results indicate that spontaneous causal attribution can modify the perceived temporal order of events. The overwhelming majority of participants in the first condition reported the order that matched their causal impression, despite the fact that the actual temporal order was clearly perceivable, as both the long delay (350ms) between the events and the veridical ordering in condition 2 indicate. The causal basis of the reordering effect is further demonstrated by the strong endorsement of the statement according to which B made C move.

Before attempting to explain and evaluate this finding, we need to address some potential confounds, especially related to attentional issues. First, the sequence becomes visually simpler in the absence of object A, therefore the erroneous temporal order reported in condition 1 could be attributed to the relative increase of perceptual load. Similarly, since motion and especially the onset of motion are known to attract attention (Abrams & Christ, 2003; Hillstrom & Yantis, 1994), perhaps participants' attention is drawn to object C when it

starts moving, thus completely missing B's behaviour. According to either explanation, participants report not the order that they actually see but rather the most plausible order given the lack of information due to perceptual overload or split attention. In short, if participants miss part of the action, it makes sense to assume a causal relationship between events, given the starting/ending configuration, and thus report the order that matches this assumed relationship. Experiment 2 investigates these possibilities and applies a stricter test to the reordering effect.

Experiment 2

We again presented the 3-object sequence of Experiment 1 (Fig. 2a) but instead of asking participants for an explicit ordering of the events, we presented the same sequence again side-by-side with a realistic collision sequence, i.e. a sequence in which the order of events is congruent with their causal relationships (Fig. 2d). After watching each of these sequences participants were asked to identify which of the two they saw earlier.

In the second condition of this between-group experiment, we presented participants with a very similar sequence that differed only in that object B remains stationary throughout (Fig. 2c). We hypothesised that the lack of motion would diminish the causal link between A and B as well as between B and C. In the absence of a causal interpretation, participants would be better at identifying the sequence they saw when asked to choose between that and a realistic collision. If this is the case, we will have evidence that the reordering effect observed in Experiment 1 and in the first condition of this experiment cannot be explained by lack of attention to B's behaviour.

Participants and Materials

The experiment was programmed in Adobe Flex 4.6 and conducted over the Internet using Amazon Mechanical Turk.

Based on results from Experiment 1, as few as 15 participants per group would suffice for a significant result at 0.01 and a power of 0.8. Nevertheless, given the radically different measure employed, we expected a far less pronounced effect. This prediction was confirmed in 3 pilot studies using speedier clips than before (65-100mm/sec). However, we also observed that while for the control group the proportion of correct answers increased as the speed of objects decreased (56%-64%), no such correlation was observed in the experimental group. Given that our aim was to make the presented clips as perceptually clear as possible,

we decided to keep the parameters the same as in Experiment 1 and also keep the number of participants constant based on our earlier calculations.

Thus, we recruited 60 participants in total. Two were excluded from the analysis because in the critical question they did not watch one of the two sequences that they were asked to choose from so their answers were in fact random (see Design and Procedure). Of the remaining 58 participants, 31 were male and 27 were female. The mean age was 34.57 (SD=12.12). Each participant was paid $0.50.

The same precautionary measures as in Experiment 1 were taken and the same values were recorded during the experiment. The mean minimum frame rate was 29.63fps (SD=0.39) and the mean delay between the movements of objects C and B was 351.03ms (SD=5.22). Lastly, no participant reported that their distance from the monitor was greater than an arm's reach.

As in Experiment 1 we repeated the analysis that is reported in the results section after removing the 6 participants for whom the minimum frame rate dropped below 29fps at any point or for whom the actual delay deviated by more than 10% (i.e. 35ms) from the 350ms target. The results of this stricter analysis remained the same.

Design and Procedure

The 58 participants were randomly assigned to one the two conditions that differed only in the critical clip, resulting in 29 participants in each condition. The introductory screens were the same as in Experiment 1 and participants were asked to pay attention to the clip that would be shown a single time.

In condition 1 ("B moving") the clip was identical to the clip shown in the first condition of Experiment 1 (Fig. 2a). In condition 2 ("B static") the clip was similar with the exception that B remained static throughout the sequence (Fig. 2c). So, object A approaches from the left at 30mm/sec and stops next to B at which point C starts moving at the same speed. The clip ends when object C stops after 35mm. As in Experiment 1 there were two versions of each clip, one with the colours being red, blue and purple and direction left-to-right as in Figure 2 and another version were the colours were shuffled (A=purple, B=red, C=blue) and the direction of movement was right-to-left.

After watching the clip, participants proceeded to the "review" screen in which two clips were displayed side-by-side. One of the clips was the critical clip that they had just seen and

the other was a clip featuring a realistic three-object collision: Object A approaches from the left at 30mm/sec and stops next to B at which point B starts moving to the right and stops next to C, following which C starts moving to the right (Fig. 2d). So, in condition 1 the participants had to choose between two clips that differed only in the order in which B and C started moving (Fig. 2a vs Fig. 2d), while the difference in condition 2 was mainly whether object B moved or not (Fig. 2c vs Fig. 2d). Below each clip there was a "play" button and participants were allowed to watch each clip as many times as they wanted before reporting which of the two they had seen in the previous screen. Then participants were asked to indicate their confidence by dragging a slider on a scale that was labelled "Not at all confident" to the left and "Very confident" to the rightmost position. Finally, participants were asked for direct causal judgments for each pair of objects in the clip, as in Experiment 1.

Results

The proportion of participants that correctly identified the clip they saw was 37.93% for condition 1 and 72.41% for condition 2, as shown in figure 5[1]. These two conditions were significantly different: $\chi^2(1, N = 58) = 6.97$, $p = .008$. Again participants were confident in their choice: the mean confidence rating was 74.10/100 (SD=23.15) for condition 1 and 81.21/100 (SD=22.40) for condition 2.
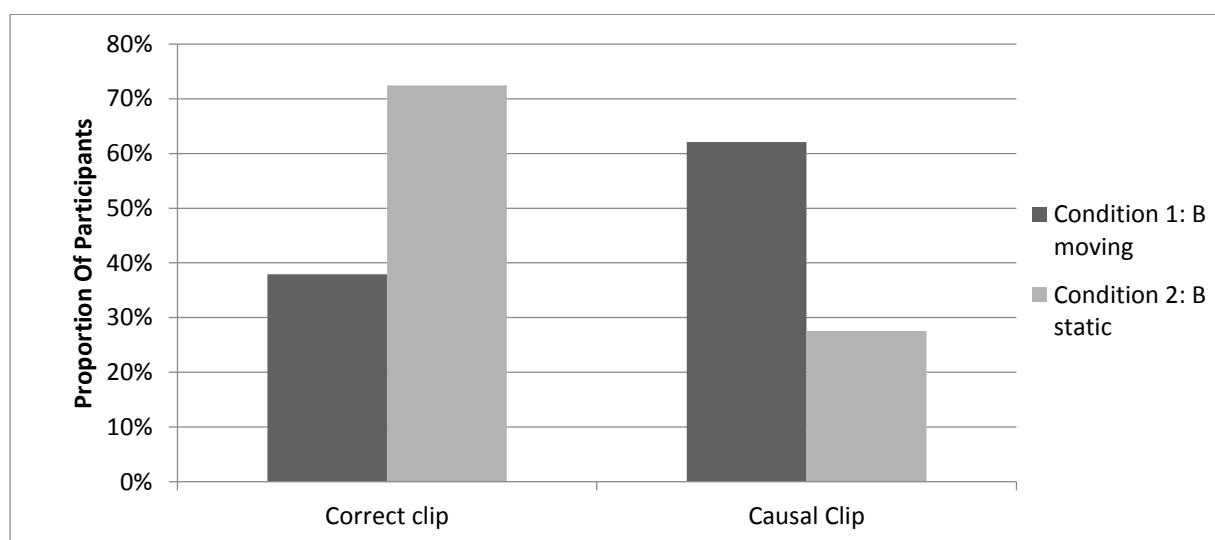

Figure 5: Proportion of participants that selected the correct clip (the one they saw) or the "realistic" collision clip in each condition.

---

[1] Again, we collapsed the responses to the normal and the mirrored versions of the clips since no difference was observed.

Regarding the causal ratings, for condition 1 they are almost identical to the respective ratings in Experiment 1: Participants agree strongly that A caused B to move and that B caused C to move, while being relatively neutral in the indirect A-C relationship and giving very low ratings to the inverse relationships. For condition 2, since B did not move at all, these causal questions are rather ambiguous and the answers participants gave to some extent reflect this ambiguity by being around the midpoint mark for all relationships with compatible temporal order (A-B, B-C and A-C). In any case there does not seem to be a prevalent causal perception in condition 2.


Figure 6: Mean causal judgments for each object pair per condition. A value of 50 corresponds to neutrality – lower ratings indicate disagreement and higher rating indicate agreement with the causal statement (Error bars represent 95% confidence intervals)

Discussion

Compared to Experiment 1 the reordering effect this time was less pronounced but perhaps more impressive, given the different measure that we used. Participants in condition 1 saw a clip featuring relatively slow moving objects (30mm/sec) in which object C moves 350ms before object B, but failed to identify the clip they saw, choosing instead with high confidence a clip in which B moves earlier, and, most critically, appears to be launching C. The fact that when asked to report the order of events (Exp. 1) rather than identify the clip they saw (Exp. 2), participants show an even stronger preference for the causal order can be

explained, in our view, by the nature of the measure. It is likely that some people do detect a 'glitch' in the clip when they first experience it, some deviation from an ideal collision clip, but they still don't identify the deviation to be the order of the events. Thus, when asked to choose between that clip and a realistic collision, these participants might prefer the deviant one, the one they actually saw, not because of the order, as the results of Experiment 1 show, but because it is the one that does not "look right". This hypothesis is supported by the direct causal judgments that remained roughly the same between the two experiments.

The most important finding from this experiment is that the reordering effect cannot be explained by split attention. When B remains stationary in condition 2 the majority of participants detect it and thus are able to correctly identify the clip they saw. This means that in condition 1, where B does move towards C, albeit late, its motion is in fact noticed and the subsequent reordering does depend on that detection. All the events that take place in the sequence are actually registered by the perceptual system and all are necessary for the causal impression to be formed; events need to be seen in order to be reordered.

General Discussion

In two experiments we have provided evidence that spontaneous causal attribution can influence the perception of temporal order. The causal reordering effect (Bechlivanidis & Lagnado, 2013) seems to generalize to situations where strong causal beliefs are present, irrespective of whether the causal links are recently learned or directly perceived.

The effect is not due to limitations of the perceptual system, since the relative onset of the events that are reordered is clearly perceivable in the absence of causal incongruences. Moreover, it is not the case that causation is driving attention away from the critical events: participants are in fact reporting their experience rather than an educated guess based on fragmented perceptual input. All the events are, in fact, perceived and are subsequently reordered to fit a causal interpretation. Finally, we have shown that the effect is strong enough to result in the formation of a mental representation of the reordered sequence: the majority of participants failed to recognise the sequence they experienced seconds earlier and were very confident that they had seen the sequence in which the temporal order matches the causal order.

These results point to an interpreted view of temporal order perception. If, as hypothesized, there exists a distinct mechanism for order discrimination (Mitrani, Shekerdjiiski, &

Yakimoff, 1986), it seems plausible that inputs to that mechanism include richer sources of information besides direct sensory input. Alternatively, rather than causal representations influencing the generation of temporal order judgements, there might be no spontaneous representation of temporal order at all. Since the temporal priority principle (causes precede their effects) allows for order judgements to be inferred retrospectively, then perhaps it is only causality that is represented.

Why are causal impressions formed in the first place though? Even if we ignore the order of events, the target sequence features other extreme spatiotemporal deviations from ideal collisions. According to current models of causal perception (see Rips, 2011 for a review) such deviations should neither be tolerated by a modular input analyser (Michotte, 1963; Scholl & Tremoulet, 2000) nor allow for a match against stored causal schemata (Sanborn, Mansinghka, & Griffiths, 2013; Weir, 1978; White, 2006).

White's (2014) more recent approach may provide a solution to this conundrum by specifying 14 low level cues to causality. Many of those cues are present in our target sequence even if they are not instantiated by the appropriate objects. So, for example, there is "contact between actor and object" (cue 5) and "property transmission" (cue 8) despite the fact that upon contact between objects A and B, the momentum is phenomenally transmitted to a third object C.

Nevertheless, if the presence of those cues triggers the search for a causal interpretation then ignoring spatiotemporal deviations and more interestingly reordering the events will result in the simplest representation of the observed sequence (Chater & Vitányi, 2003; Lombrozo, 2007). Rather than object A launching C from distance and object B moving spontaneously or being pulled by C, A will be represented as launching B followed by B launching C. The latter interpretation is clearly simpler by involving two instances of a single type of causal relationship and furthermore by matching a causal schema, similar to a queue of dominos falling.

In sum, we propose that the detection of abstract low level cues to causality (White, 2014) in a sequence as a whole, triggers the search for familiar causal representations (schemata). If compatible schemata matching the sensory data are unavailable, the simplest representations will be selected at the expense of overlooking deviations, even if it results in the modification of the objective temporal order.

References

Abrams, R. A., & Christ, S. E. (2003). Motion Onset Captures Attention. *Psychological science, 14*(5), 427-432. doi: 10.1111/1467-9280.01458

Bechlivanidis, C., & Lagnado, D. A. (2013). Does the "why" tell us the "when"? *Psychological science, 24*(8), 1563-1572. doi: 10.1177/0956797613476046

Buehner, M. J. (2012). Understanding the Past, Predicting the Future: Causation, not Intentional Action, is the Root of Temporal Binding. *Psychological science, 23*(12), 1490-1497. doi: 10.1177/0956797612444612

Buehner, M. J., & Humphreys, G. R. (2010). Causal contraction: spatial binding in the perception of collision events. *Psychological science, 21*(1), 44-48. doi: 10.1177/0956797609354735

Chater, N., & Vitányi, P. (2003). Simplicity: a unifying principle in cognitive science? *Trends in cognitive sciences, 7*(1), 19-22.

Dainton, B. (2010). Temporal consciousness. In E. N. Zalta (Ed.).

Grush, R. (2007). Time and experience. In T. Müller (Ed.). Frankfurt: Klosterman.

Hillstrom, A. P., & Yantis, S. (1994). Visual motion and attentional capture. *Perception & Psychophysics, 55*(4), 399-411.

Hoerl, C. (2013). 'A Succession of Feelings, in and of Itself, is Not a Feeling of Succession'. *Mind, 122*(486), 373-417. doi: 10.1093/mind/fzt070

Humphreys, G. R., & Buehner, M. J. (2010). Temporal binding of action and effect in interval reproduction. *Experimental Brain Research, 203*(2), 465-470. doi: 10.1007/s00221-010-2199-1

Kim, S.-H., Feldman, J., & Singh, M. (2013). Perceived causality can alter the perceived trajectory of apparent motion. *Psychological science, 24*(4), 575-582. doi: 10.1177/0956797612458529

King, A. J. (2005). Multisensory Integration: Strategies for Synchronization. *Current Biology, 15*(9), R339-341. doi: 10.1016/j.cub.2005.04.022

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive psychology, 55*(3), 232-257. doi: 10.1016/j.cogpsych.2006.09.006

Michotte, A. (1963). *The Perception of Causality*. London: Methuen & Co Ltd.

Mitrani, L., Shekerdjiiski, S., & Yakimoff, N. (1986). Mechanisms and Asymmetries in Visual Perception of Simultaneity and Temporal Order. *Biological Cybernetics, 54*(3), 159-165.

Phillips, I. B. (2014). Experience of and in Time. *Philosophy Compass, 9*(2), 131-144. doi: 10.1111/phc3.12107

Rips, L. J. (2011). Causation From Perception. *Perspectives on Psychological Science, 6*(1), 77-97. doi: 10.1177/1745691610393525

Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2013). Reconciling intuitive physics and Newtonian mechanics for colliding objects. *Psychological Review, 120*(2), 411-437.

Scholl, B. J., & Nakayama, K. (2004). Illusory Causal Crescents: Misperceived spatial relations due to perceived causality. *Perception, 33*(4), 455-469. doi: 10.1068/p5172

Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in cognitive sciences, 4*(8), 299-309.

Soteriou, M. (2010). Perceiving events. *Philosophical Explorations, 13*(3), 223-241. doi: 10.1080/13869795.2010.501904

Spence, C., & Squire, S. (2003). Multisensory Integration: Maintaining the Perception of Synchrony. *Current Biology, 13*(13), R519-R521. doi: 10.1016/S0960-9822(03)00445-7

Weir, S. (1978). The perception of motion: Michotte revisited. *Perception, 7*(3), 247-260. doi: 10.1068/p070247

White, P. A. (2006). The role of activity in visual impressions of causality. *Acta psychologica, 123*(1-2), 166-185. doi: 10.1016/j.actpsy.2006.05.002