Title
Searching for the Cause: Search Behavior in Explanation of Causal Chains

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN
1069-7977

Authors
Czarnowski, Daniel W
Marsh, Jessecae

Publication Date
2021

Peer reviewed

Searching for the Cause: Search Behavior in Explanation of Causal Chains

Daniel W. Czarnowski (dwc319@lehigh.edu)
Department of Psychology, 17 Memorial Drive East
Bethlehem, PA 18015

Jessecae K. Marsh (jessecae.marsh@lehigh.edu)
Department of Psychology, 17 Memorial Drive East
Bethlehem, PA 18015

Abstract

Understanding cause and effect relationships gives power to produce desired effects and avoid negative outcomes. Despite the power of causal explanations, people often lack full understanding of how causes relate to or produce their effects. In two experiments, we explored how people search for information to enrich their causal explanations of real-world phenomena when given the chance. Participants completed an information search task that provided a causal relationship where they could seek out mechanistic information at different steps between the cause and the effect. We measured where people searched in the causal chain of events that made the explanation. We found that when allowed to search freely (Experiment 1) or when instructed that they must search for information (Experiment 2) participants consistently sought out information closest to the root cause in the explanation. We discuss implications for how to improve the teaching of new explanations to maximize the informational desires of the learner.

Keywords: causal reasoning; information search; explanation

Introduction

Knowing how causes bring about their effects allows us to make predictions, perform interventions, and formulate explanations for the phenomena we experience in the world (Shanks, 2004; Hagmayer et al., 2007; Hagmayer & Sloman, 2009; Fernbach et al., 2011). However, it is not always the case that our understanding of cause and effect is correct (e.g., based on fabricated information; Lazer et al., 2018) or complete (e.g., Illusion of Explanatory Depth; Rozenblit & Keil, 2002). These gaps in our understanding can be problematic, as causal knowledge, accuracy notwithstanding, often informs complex decisions (e.g., vaccination preferences; Cooper et al., 2008; Dube et al., 2015). When people sense gaps in their knowledge, they may search for more information to expand their causal understanding. Multiple theories of why people engage in information-seeking behavior have been proposed (e.g., functional explanation, Liquin & Lombrozo, 2020; information gap theory, Lowenstein, 1994; map-mismatch, Wong & Yuddel, 2015). When people have found motivation to search, an open question is what information they preferentially search for to support their causal explanations. In this paper, we explore where in a causal sequence of an explanation people choose to learn more information to support their existing causal explanations.

Causal explanations serve to answer "how" and "why" questions regarding the occurrence of a phenomenon by describing the process through which it is produced (Lombrozo & Vasilyeva, 2017; Wellman, 2011). Causes give rise to effects through intermediary processes that link the cause and the effect (e.g., mechanism; Ahn & Kalish, 2000). These intermediary processes can range from simple (e.g., a collision of billiard balls is mediated by the transfer of momentum from the first ball to the second) to more complex chains of events. Imagine the example of generating an explanation for how prolonged exposure to sunlight causes sunburn. Simply identifying UV radiation in the sunlight as a causal factor is not a sufficient explanation as to why it causes sunburn. Instead, a good explanation traces the sequence of events that occurs between the initial exposure to sunlight (cause) and the onset of sunburn (effect). Even the simplest explanation involves a series of intermediary mechanistic events: prolonged UV radiation exposure damages DNA in skin cells, specialized cells release melanin to protect damaged cells, unprotected cells undergo programmed cell death, immune system senses cell death and triggers inflammatory response to heal (sunburn). For simplicity, we will use the term "intermediary steps" to refer to those mechanistic links that connect a cause and effect in a causal explanation.

To create a full explanation of a phenomenon like UV radiation causing sunburn, people should have knowledge of the intermediary steps that link the cause and effect. In searching for information to enrich their causal explanation, participants could search for information that relates to any intermediary steps. However, the sequential nature of the steps that link a cause and an effect in an explanation means that there are inherent differences regarding the role each step plays in the overall causal chain. The intermediary step most proximal to the root cause is "responsible" for both the greatest number of subsequent steps in the causal chain sequence, as well as the ultimate effect indirectly through those steps. The step most proximal to the effect, on the other hand, is directly responsible for producing the effect, but has no bearing on other steps in the explanation. In short, where people learn information in a causal chain sequence can provide differing amounts of knowledge about the other steps in a causal explanation.

Given the different informational value of intermediary steps in an explanation, where would people seek out information to enrich their explanations? One possibility is that people seek out steps that are closest to the root cause of an explanation. Categorization research has found that people place strong weight on root causes in causal chains (e.g., Ahn & Kim, 2000; Ahn et al., 2000; Marsh & Ahn, 2006). Likewise, people have shown a preference to intervene on root causes to bring about change in a causal system (Bohm & Pfister, 2000; Edwards et al., 2008; 2015; Green & McManus, 1995; Hagmayer & Sloman, 2005; 2009; Yopchick & Kim, 2009). This weighting of root causes may result in a general preference to search for information that is closest to the root cause of an explanation, or the most root of the intermediary steps.

Not all research has found that people have this preference for root causes, however. Under specific conditions, people show preferences for immediate causes. Previous work has shown a flip to a preference for intervening on immediate causes when implementing short-term solutions (Edwards et al., 2008; 2015). More generally, in complex causal networks, such as ecological food webs, people prefer to intervene on immediate causes, believing that the power of any interventions on root causes will attenuate as the effects spread throughout the complex network (White, 1997). Combined, these findings suggest limits to the preference for root causes. If people see a causal explanation that contains many intermediate steps as complex, they may focus on searching for information about intermediary steps that most directly produce the effect. In this way, people may prefer to search for information closest to the effect they are trying to explain.

In this set of experiments, we test where people search for information when learning more about causal explanations. We developed an information search task that presented participants with a cause (e.g., sunlight exposure) and an effect (e.g., sunburn) that were linked by a varying number of causal intermediary steps. We allowed people to select steps in that causal sequence to learn more about the mechanism that linked the cause and the effect. Through this paradigm we can test what type of searching people do in enriching their explanations for different causal relationships. Given the large amount of work that suggests a preference for root causes, we predict that participants will demonstrate a preference for learning information closest to the root cause, with subsequent searches following the steps along the sequence from cause to effect. We test that prediction through the following two experiments.

Experiment 1

In Experiment 1 we present a first test of where people search for information in a causal explanation sequence. We present people with varying amounts of intermediary steps between a root cause and an effect to determine if the total amount of information represented in a causal explanation changes how people search for information (e.g., because the relationship is seen as more complex). We also test search across domains to determine if search varies as a factor of the type of information being searched for.

Method

Participants We recruited 49 (30 female, 17 male, and 2 non-binary) undergraduate students who were enrolled in an introductory psychology course. Participants received course credit as compensation. All participants completed the study remotely. Across both experiments, participants were screened to ensure that they had normal or corrected-to-normal vision and were fluent in English. Approval for both studies was obtained from the authors' university IRB.

Materials We created a set of causal chains that each depicted a cause and an effect linked by a series of intermediary causal steps. A causal chain was presented to a participant as a series of boxes that represented the various pieces of information in a causal relationship. The boxes on the left and right sides of the screen contained text describing either the root cause or the final effect of the causal relationship (with positions counterbalanced). The boxes containing the root cause and the effect were separated by a sequence of either three, four, or five intermediary boxes, each labeled with the word "Step" and a number that indicated the box's position along the causal pathway from the root cause to the effect (e.g., Step 2; see Figure 1).[1] Each intermediary box occluded a piece of information that explained what occurs at that particular step in the causal pathway from the root cause to the effect.

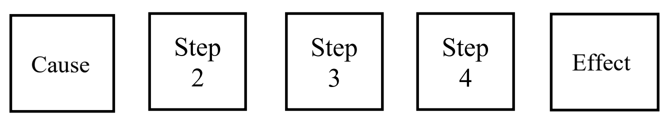

Figure 1: Example information search presentation.
Note. For the different causal chains used in the experiment, the labels "cause" and "effect" would be replaced by appropriate labels for that causal chain (e.g., "UV Radiation" and "Sunburn", respectively).

We developed causal chains representing three different domains: natural phenomena, mechanics, and personal health. We selected three scientifically-based causal chains in each domain (e.g., health: "How does UV radiation make you sunburned?"; mechanics: "How does a toilet flush?"; natural: "How is lightning generated in clouds?"). We additionally developed a set of causal chains that depicted an overall true causal relationship (artificial domain: health: "How do diets make someone lose weight?"; mechanics: "How do digital cameras take pictures?"; natural: "How

[1] Labels for intermediary boxes began with "Step 2" rather than "Step 1" to avoid confusing participants who may have intuited the root cause to be "Step 1" by default.

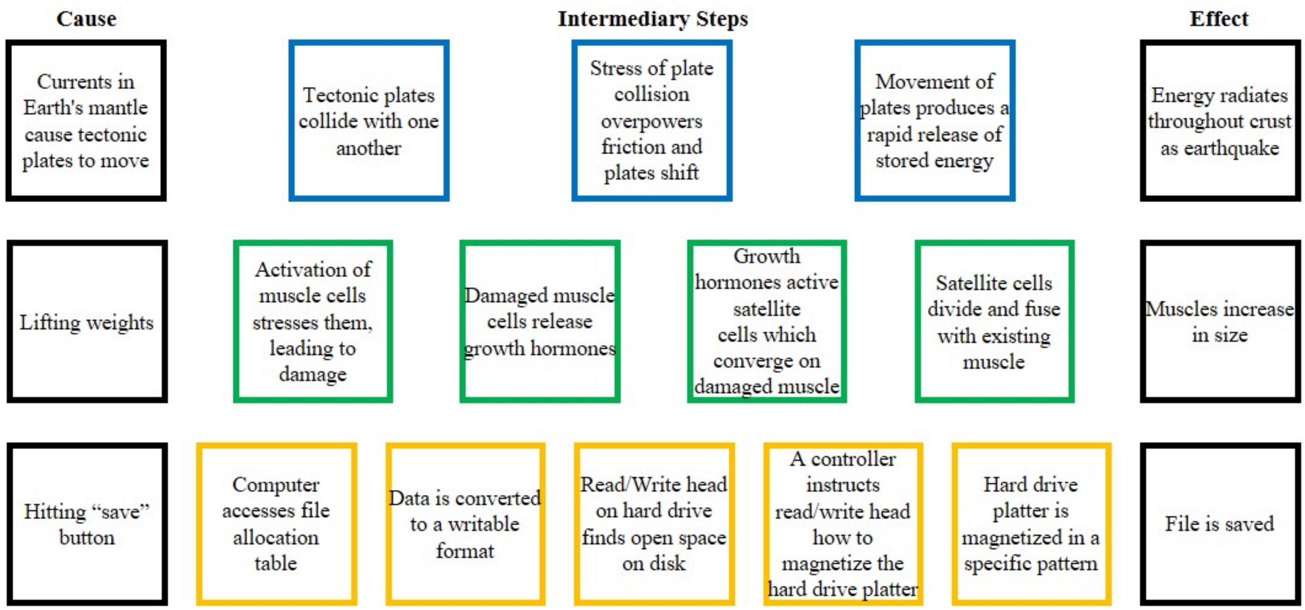

Figure 2: Example stimuli from natural (3-step), health (4-step), and mechanical (5-step) domains.
Note. Colors highlight the differing number of intermediatory steps, but were not used in the actual experiment.

does the sun influence ocean currents?"), but the information in the steps did not represent true scientific evidence of how that cause produced the given effect. We included artificial chains to test if subsequent information search differed when participants received a surprising piece of information on their first search. Because of the different research question relevant to the artificial chains, we do not include data for those chains in further analyses.

Within each domain, we developed a chain that had 3 steps between the root cause and the effect, 4 steps between, and 5 steps between. Overall, we developed a total of 12 causal chains, with participants seeing 3 chains, each of different lengths, in each of the four domains. See Figure 2 for example chains of each length and in each domain.

Design and Procedure The study was programmed using PsychoPy v2020.2.10 and was distributed through the Pavlovia online experiment repository. Participants began the experiment by reading a set of instructions that informed them they would be presented with a series of cause and effect relationships that they would be able to learn more about. The pieces of information they would be able to learn were the steps in the causal chain between the root cause and the effect. Participants were instructed to imagine that they had to explain the relationship to an expert in the field when considering whether or not to search for more information about the relationships.[2] They were then given an example search task. After becoming familiarized with the task, participants were told that they might be prompted to produce a written explanation for one of the relationships they have seen. This instruction was included to encourage participants to engage with the task; no such written explanations were requested.

Participants then started the information search tasks. Prior to seeing the relationship in each search task, a prompt reminded participants to imagine that they had to explain the relationship to an expert in the field when considering whether or not they would like to search for additional information. In a search task, participants were free to search for as much information about the causal relationship as they wanted. To search for information at a particular step, participants pressed the appropriate number on their keyboard that corresponded with the step they would like more information about, which would reveal the information occluded by the box. After selecting a box, participants could select to search for more information and choose another box. When they searched for the next piece of information, the previously searched information pieces went back to being occluded. In this way, participants could only see one intermediary step at a time on the screen with the cause and effect. Participants could request to see a step multiple times. When they deemed themselves ready, participants pressed a key that indicated they were done with searching and ready to move on. Once participants finished information searching for a given chain, they advanced to a screen and rated how confident they were in their knowledge of the relationship using a 1 (not at all confident) to 10 (extremely) scale. After making their confidence judgment, participants moved on to searching

[2] We used this instruction because in pilot work participants said they would be most likely to search for information if they had to make an explanation to an expert. We therefore used this goal to try to maximize searching for any information.

for information in the next chain. Once all 12 chains were completed, participants were asked to report how familiar they were with each of the 12 causal relationships prior to their participation in the experiment using a 1 (not at all familiar) to 10 (extremely familiar) scale.

The presentation order of all 12 chains was randomized for each participant. The left-right screen positions of the root cause and the effect were counterbalanced within participants such that each chain had an equal likelihood of being presented with the root cause on the left (arrows flowing from left to right) or on the right (arrows flowing from right to left). Collapsing across this position manipulation should account for any preference of searching left to right in the order of English reading.

Results

Amount of information searched We first explored how much information search participants were doing overall. Most participants (87.8%) searched for data in at least one of the 9 causal chains of interest in these analyses (number of chains searched: M = 5.31). A small number of participants (12%) did not search for information in any chain. We next analyzed whether for participants who did search for information, if the length of the causal chain (3, 4, or 5 intermediate steps) influenced the amount of information search participants completed. Because of the missing nature of the data (participants could search in 1 to 9 total chains, with data missing for any chains they did not search), we used mixed linear modeling (MLM) to analyze the data. MLM models are robust to missing data and can be used to analyze repeated measures data in an ANOVA-style format. We entered domain (natural, mechanical, health) and intermediate steps (3, 4, or 5) as fixed effects into the MLM model, using a diagonal covariance matrix. Again, participants who did not search for any information in any chain were not included in this analysis. We found a significant main effect of intermediate steps, $F(2, 136.3) = 11.6$, $p < .001$. These findings reflect that participants searched for more information when there was more information to search through across domains (3-step: M = 2.97, SE = 0.18; 4-step: M = 3.44, SE = 0.21; 5-step: M = 4.48, SE = 0.26). We did not find a significant main effect of domain or a significant interaction, ps > .27.

We additionally examined whether our measures of familiarity and confidence correlated with the number of steps searched when participants chose to search for information. We calculated mean familiarity ratings, confidence, and number of steps searched across the three domains at each step level for each participant. While all correlations were negative, we did not find significant correlations between familiarity and number of steps searched for any chain lengths, ps > .14. We likewise did not find significant correlations between confidence and steps searched for any chain length, ps > .21.

Location of search We next turned to our main question of where in a causal chain people were searching for information. We tallied the number of times people chose each step in the causal sequence for a given chain on their first choice of information search (first selections), second choice (second selections), third (third selections), and so on. As a reminder, participants could search for as much or as little information as they liked. Therefore, some participants did not make first choices in at least one chain (n = 6) and many participants did not make second choices in at least one chain (n = 11). Table 1 shows the proportion of choices made for each step in a causal chain. We do not present data beyond the third selection. A fourth selection for a 4- or 5-step chain could provide novel information. However, in a 3-step chain if participants have already looked at the three intermediate steps, a fourth selection could not provide novel information. Thus, comparing searches beyond three selections is not equated in new information value between different step-lengths so we do not analyze those data here.

Table 1: Proportion of each step selected in Experiment 1

	Closest to Cause	Intermediate Steps			Closest to Effect
1st Selection					
3-step (n=85)	**.72**		.13		.15
4-step (n=84)	**.76**	.11		.06	.07
5-step (n=91)	**.73**	.08	.10	.04	.05
2nd Selection					
3-step (n=65)	.05		**.88**		.08
4-step (n=60)	.05	**.85**		.08	.02
5-step (n=70)	.03	**.84**	.03	.07	.03
3rd Selection					
3-step (n=58)	.08		.06		**.84**
4-step (n=53)	.01	.07		**.86**	.03
5-step (n=67)	.02	.05	**.86**	.01	.02

Note. The total number of information selections across participants is listed in parentheses next to each step sequence. Bolding indicates the most commonly selected step for each sequence length.

To examine search patterns, we compared the proportion of participants who chose the intermediary step closest to the root cause to the proportion who chose the step closest to the target effect through Fishers' Exact test. We present these analyses collapsed across domain because the same pattern of results was found within each domain. Participants preferentially chose to search for the step closest to the root cause for all three step lengths in their first selections, ps < .001.

After having learned about the intermediary step closest to the root cause in their first selections, would participants then choose to "skip ahead" to the intermediary step closest to the effect or search sequentially along the causal chain? For second selections, participants were significantly more likely to select the subsequent step along the causal chain (next-closest to the cause) than the step closest to the effect for all step lengths, *p*s < .001. For third selections in 4- and 5-step chains, participants again chose the step next in the sequence instead of closest to the effect, *p*s < .001.[3]

While the majority of participants began searching for information at the intermediary step closest to the root cause and then selected consecutive steps when subsequent selections were made, a subset of participants made their first selection somewhere else in the chain. A small percentage of participants started at the effect and searched back toward the cause (3-step: 3.5%; 4-step: 1.2%; 5-step: 3.3%). Another small group of participants selected a step in the middle and searched sequentially towards the effect from there (3-step: 1.2%; 4-step: 4.8%; 5-step: 1.1%). The remaining participants who did *not* select the step closest to the cause during their first search exhibited random search behavior across all three chain lengths where they did not go in one specific direction in their selections (3-step, 23.5%; 4-step, 17.9%; 5-step, 23.1%).[4]

Discussion

We demonstrated a search preference for information closest to a chain's root cause and subsequently searching for information that flows next in the chain. One limitation of the previous experiment is that many participants did not search for information in a sizeable number of chains. Would the patterns of search differ if people felt obligated to search? It is possible that people may engage in more exploratory behavior, searching for information closer to the terminal effect or choosing information more at random. We explore this possibility in Experiment 2.

Experiment 2

In Experiment 2 we tested whether requiring people to search for information would change search patterns. It is possible that any effects of searching closer to terminal effects were lost because our participants were not searching as often as they could in Experiment 1.

Method

Participants We recruited 48 (28 female, 19 male, 1 preference to self-describe) undergraduate students who participated in exchange for course credit.

Materials, Design, and Procedure The method for Experiment 2 was identical to Experiment 1 with the exception that the instructions told participants they had to select at least one piece of information on each trial.

Results and Discussion

Amount of information searched Our instructions told participants they had to search for information and we did see a corresponding increase in search from Experiment 1, with all participants searching for data in at least one of the 9 causal chains. Because an individual trial allowed participants to move on without searching, not all participants searched in every single chain (number of chains searched M = 7.14). We analyzed whether the length of the causal chain (3, 4, or 5 intermediate steps) influenced the amount of information search participants completed. We again used mixed linear modeling (MLM) to account for missing data. Participants who did not search for any information were excluded. We entered domain (natural, mechanical, health) and intermediate steps (3, 4, or 5) as fixed effects into the MLM model, using a diagonal covariance matrix. As in Experiment 1, we found a significant main effect of intermediate steps, $F(2, 193.1) = 17.6$, $p < .001$. Again, participants are searching for more information when there is more information to search through across domains (3- step: M = 2.56, SE = 0.15; 4-step: M = 3.28, SE = 0.18; 5- step: M = 4.17, SE = 0.23). We did not find a significant main effect of domain or a significant interaction, *p*s > .09.

We again examined whether our measures of familiarity or confidence correlated with the number of steps searched for by participants who chose to search. We calculated mean familiarity ratings, confidence, and number of steps searched across the three domains at each step level for each participant. We did not find significant correlations between familiarity and number of steps searched for 3- or 4-step chain lengths, *p*s > .07. There was a significant correlation between familiarity and number of steps searched for 5-step chains, $r(46) = -.329$, $p = .022$. We did not find significant correlations between confidence and steps searched for any chain length, *p*s > .16.

Location of search To examine participants' patterns of search, we once again ran a series of Fisher's Exact tests to compare the proportion of participants who chose the intermediary step closest to the root cause to the proportion who chose the step closest to the target effect. Results across domain mirrored those found in Experiment 1, so we once again present these analyses collapsed across domain.[5] Table 2 shows the proportion of choices made for each step.

[3] The third step is the step closest to the effect in 3-step chains.

[4] These percentages reflect the percent of chains searched using the corresponding search pattern.

[5] We see the same pattern of means in all three domains. Additionally, the significance holds in all chains, except for the 3-step chain in the health domain. The difference here was not significant while the pattern of means was in the same direction as all other domains.

We replicate the basic pattern of Experiment 1. In their first selections, participants chose the intermediary step closest to the root cause over the step closest to the effect for all three chain lengths, $ps < .001$. Comparisons of participants' second and third selections reveal the same pattern; participants are once again searching for the next closest step to the root cause, $ps < .001$.

We also looked at the subset of participants who did not select the step closest to the cause when making their first selections. As was the case in Experiment 1, the majority of these participants exhibited random search behavior (3-step, 36.8%; 4-step, 37.7%; 5-step, 35%). The remaining few participants either searched sequentially from effect to cause (3-step, 4.3%; 4-step, 5.3%; 5-step, 5.8%), or selected a step in the middle and searched sequentially towards the effect from there (3-step, 3.4%; 4-step, 7%; 5-step, 5%).

Table 2: Proportion of each step selected in Experiment 2

	Closest to Cause	Intermediate Steps			Closest to Effect
1st Selection					
3-step (n=117)	**.56**		.28		.17
4-step (n=113)	**.50**	.25		.11	.15
5-step (n=120)	**.54**	.14	.13	.08	.08
2nd Selection					
3-step (n=69)	.07		**.78**		.14
4-step (n=82)	.05	**.66**		.22	.08
5-step (n=89)	.08	**.67**	.07	.16	.02
3rd Selection					
3-step (n=60)	.13		.1		**.76**
4-step (n=71)	.04	.16		**.66**	.12
5-step (n=81)	.07	.04	**.72**	.04	.08

Note. The total number of information selections across participants is listed in parentheses next to each step sequence. Bolding indicates the most commonly selected step for each sequence length.

General Discussion

Across two experiments, we demonstrated a desire to search for information close to root causes in causal chains. This finding held when participants could search for as much information as they liked (Experiment 1) or were requested to search in every instance (Experiment 2). This finding held for the first piece of information searched for, as well as the second and third. Overall, this suggests that when seeking out information about how causes produce effects, participants seek out information proximal to the cause and work toward the terminal effect.

In comparing the findings of Experiment 1 and 2, it seems that participants had less of a bias toward the closest-to-cause piece of information when they were requested to search as in Experiment 2. This pattern may have arisen from a subset of participants in Experiment 2 simply selecting a step at random during their first "forced" selection. Looking at the differences between the total number of participants who made first and second selections across Experiments 1 and 2 suggests a greater number of participants in Experiment 2 made first selections. However, the rates of second selections are similar across experiments. While it is possible that participants only wanted to search for a single step of information, it might also be the case that this pattern arose from a subset of participants not engaging with the task and merely selecting a random step out of obligation and then continuing on. It is a question for future research how forcing people to search out more information for an explanation (e.g., not allowing participants to move on until they have selected at least three pieces of information) may alter search patterns.

The desire to work from cause to effect in searching for information provides additional support for the importance of root causes in reasoning about causality (Ahn & Kim, 2000; Hagmayer & Sloman, 2009; Marsh & Ahn, 2006). While support for preferring immediate causes has been shown (e.g., White, 1997; Edwards et al., 2008; 2015), the nature of our task did not invoke the conditions under which these effects arise. Future research can explore if primed to think an explanation will be very complex, will participants then shift away from this preference for root causes.

While not our main measure of interest, we do see evidence across experiments that participants are sensitive to the number of possible steps and seek more information when there are more possible steps. Interestingly, familiarity or confidence did not correlate with the amount of information searched. It is possible that people searched for different reasons: some may have searched because they did not have knowledge of that relationship and some may have searched because they were looking to validate knowledge, they believed they held. Future work can explore what exactly motivates a person to look for a given step in filling out a causal explanation.

An open question from our findings is whether participants are conducting efficient search in moving from the root cause toward the effect. For example, is it best to include all steps in an explanation? It is possible that a strategy of searching out the steps closest to the cause and closest to the effect first could have allowed participants to guess or fill in the remaining intermediate steps. Likewise, selecting the most middle step in the sequence may have allowed for the easiest guessing to fill out the rest of the chain. Whether either of these other strategies would be more efficient to completely understand the causal relationship is a question for future research.

Future research could also examine how different motivations for search change search behavior. We instructed participants to generate an explanation suitable to give to an expert in the field. Giving participants a different goal for search (e.g., to search until they were able to generate an explanation for a peer or an explanation they personally found suitable), could result in different search patterns. Beyond who the explanation is for, future work could explore how the use of the explanation could influence search. For example, participants could be given the goal of explaining an intervention to prevent the effect from occurring. People's preferences for where to intervene in causal systems can vary by whether the goal is a short- or long-term intervention (Edwards et al., 2015). These findings suggest that instructing participants of a specific intervention goal has potential to change search. More generally, future work can determine how people go about finding information as their goals change in relation to explaining a causal relationship.

In sum, our findings suggest that when people search for information to support existing explanations of phenomena, they exhibit an initial preference for steps closest to the root cause and work towards the effect from there. In practice, these findings have implications in learning environments. If people are preferring to learn information closest to the root cause, then teaching explanations of novel phenomena may benefit from providing mechanistic information from cause-to-effect in a step-by-step manner when describing the process by which the phenomenon is brought about.

References

Ahn, W., & Kim, N. S. (2000). The causal status effect in categorization: An overview. In D. L. Medin (Ed.), *Psychology of Learning and Motivation, 40*, (p. 23-65) New York: Academic Press.

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality, *Cognitive Psychology, 41*, 361-416

Böhm, G., & Pfister, H. R. (2000). Action tendencies and characteristics of environmental risks. *Acta Psychologica*, *104*(3), 317-337.

Cooper, L. Z., Larson, H. J., & Katz, S. L. (2008). Protecting public trust in immunization. *Pediatrics*, *122*(1), 149-153.

Dubé, E., Vivion, M., & MacDonald, N. E. (2015). Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: influence, impact and implications. *Expert review of vaccines*, *14*(1), 99-117.

Edwards, B., Burnett, R., & Keil, F. (2008). Structural determinants of interventions on causal systems. In *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (pp. 1138-1143).

Edwards, B. J., Burnett, R. C., & Keil, F. C. (2015). Effects of causal structure on decisions about where to intervene on causal systems. *Cognitive science*, *39*(8), 1912-1924.

Green, D. W., & McManus, I. C. (1995). Cognitive structural models: The perception of risk and prevention in coronary heart disease. *British Journal of Psychology*, *86*(3), 321-336.

Hagmayer, Y., & Sloman, S. A. (2005). Causal models of decision making: choice as intervention. In B. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 881-886). Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc.

Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General*, 138, 22-38.

Hagmayer, Y., Sloman, S. A., Lagnado, D. A., & Waldmann, M. R. (2007). Causal reasoning through intervention. *Causal learning: Psychology, philosophy, and computation*, 86-100.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., ... & Zittrain, J. L. (2018). The science of fake news. *Science*, *359*(6380), 1094-1096.

Liquin, E. G., & Lombrozo, T. (2020). A functional approach to explanation-seeking curiosity. *Cognitive Psychology*, *119*, 101276.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin, 116*(1), 75–98.

Lombrozo, T., & Vasilyeva, N. (2017). Causal explanation. In M. Waldmann (Ed.), *Oxford Handbook of Causal Reasoning*. Oxford: Oxford University Press.

Marsh, J. K., & Ahn, W. (2006). The role of causal status versus inter-feature links in feature weighting. *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 561-566). Mahwah, New Jersey: Lawrence Erlbaum Associates.

Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science, 26*(5), 521–562.

Shanks, D. R. (2004). Judging covariation and causation. *Blackwell handbook of judgment and decision making*, 220-239.

Wellman, H. M. (2011). Reinvigorating explanations for the study of early cognitive development. *Child Development Perspectives*, *5*(1), 33-38.

White, P. A. (1997). Naive ecology: Causal judgments about a simple ecosystem. *British Journal of Psychology*, *88*(2), 219-233.

Wong, W., & Yudell, Z. (2015). A normative account of the need for explanation. *Synthese, 192*(9), 2863–2885.

Yopchick, J. E., & Kim, N. S. (2009). The influence of causal information on judgments of treatment efficacy. *Memory & Cognition, 37*, 29-41.