How do the semantic properties of visual explanations guide causal inference?

Holly Huey, Caren M. Walker, and Judith E. Fan

University of California, San Diego, CA 92092, USA
hhuey@ucsd.edu

Abstract. What visualization strategies do people use to communicate abstract knowledge? We developed a drawing paradigm to elicit visual explanations about novel machines and found that these explanatory drawings contained: (1) greater emphasis on causally relevant parts of the machines, (2) less emphasis on structural features, and (3) more symbols than baseline drawings intended to communicate the machines' visual appearance. However, this overall pattern of emphasis did not necessarily improve naive viewers' ability to infer how to operate the machines, nor their ability to identify them. This suggests a potential mismatch between what people believe should be included in an effective visual explanation and what may actually be most useful. Together, our findings advance our understanding of how communicative goals constrain the visual communication of abstract knowledge across behavioral contexts.

Keywords: visual communication · explanation · object identification.

1 Introduction

Our ability to learn the causal structure of our world is a fundamental aspect of human cognition. From infants exploring the objects in their environments to scientists exploring the frontiers of our solar system, we seek causal knowledge to explain our observations and generate desired outcomes [1]. However, acquiring this knowledge firsthand is often costly in time and effort [2], highlighting the importance of visualization tools (e.g., diagrams, graphs) to efficiently transmit knowledge to others. Prior work has largely focused on how visualizations convey perceptual knowledge about objects and spaces, by exploiting visual and spatial cues that enable viewers to intuitive grasp the correspondence between physical features of the world and elements in the representation of it [3]. By contrast, however, relatively little is known about the mechanisms that enable humans to produce visualizations that communicate more abstract knowledge to others. Here we explore the hypothesis that producing effective visual explanations of causal phenomena relies on combining information about structure (i.e., what kinds of entities there are) with information about function (i.e., how these entities interact). Specifically, we predicted that effective visual explanations tend to highlight causally relevant information for the function of objects, while preserving enough structural information to establish a mapping to the target system [see preprint, 4].

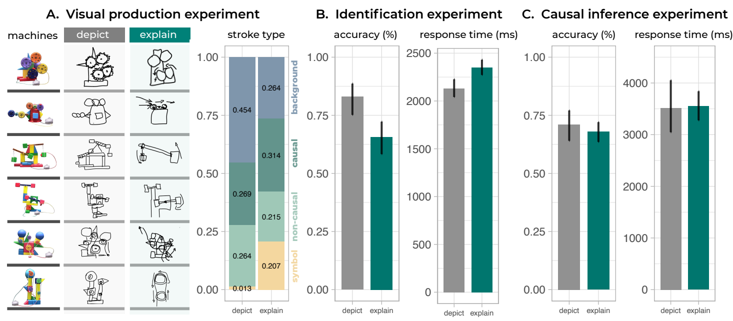

Fig. 1. A: Visual production experiment: *left,* stimuli and example drawings; *right,* stroke allocation. **B:** Identification experiment: accuracy and response time. **C:** Causal inference experiment: accuracy and response time. All error bars represent 95% CIs.

2 What information is prioritized in visual explanations of causal knowledge?

We first investigated the semantic properties that characterize visual explanations of mechanical systems. **Visual production experiment:** 52 participants completed a drawing task in which they saw video demonstrations of 6 novel machines with simple mechanical elements (i.e., gears, levers, pulleys), and generated two kinds of drawings: *visual explanations* to help naive viewers learn how the machines functioned to activate a light; and *depictions* to help naive viewers identify the machines by their appearance (Fig. 1A, left). To identify properties distinctive of visual explanations, depictions served as a baseline measure of drawing without the intent to communicate causal knowledge. **Results:** We found that visual explanations: (1) placed greater emphasis on components that were causally relevant for operating the machines (explain: 59%, depict: 50%, $b = 0.416$, $z = 3.93$, $p = 8.56e-05$), (2) placed less emphasis on structural but visually salient features (explain: 26.4%, depict: 45.4%, $b = -7.58$, $t = -5.49$, $p = 1.16e-07$), and (3) included more symbols (e.g., arrows, motion lines; explain: 20.7%, depict: 1.3%, $b = 9.91$, $t = 4.62$, $p = 1e-05$) than depictions (Fig. 1A, right). These results suggest that the goal of communicating causal knowledge systematically shifts drawings toward more abstract, functional information, even at the expense of fidelity to other visually salient features.

3 How well do visual explanations support downstream behaviors?

To understand how useful these visual explanations were for helping naive viewers to: (1) identify the machines, and (2) understand their function, we conducted two additional experiments. **Visual identification experiment:** 52 participants completed a visual search task with 300 trials, in which they identified machines by matching photos to drawings from the visual production experiment. To the extent that visual explanations place less emphasis on structural features, we predicted that participants would be slower and more error prone when cued with visual explanations than with depictions. **Results:** Consistent with our predictions, participants were less accurate (explain: 66%; depict: 83%; $b = -0.737$, $z = -2.32$, $p = 0.0201$; Fig. 1B, left) when viewing visual explanations and slower to respond (correct trials only: explain: 2351ms; depict: 2132ms; $b = 9.79e{-}02$, $t = 3.091$, $p = 0.002$; Fig. 1B, right) when viewing visual explanations, relative to depictions. These results show that the differences in semantic information contained in these drawings have distinct behavioral consequences: visual explanations, which prioritized functional information at the expense of other visually salient information like structural features, were less informative to naive viewers about the identity of the target machine than depictions were. **Causal inference experiment:** 305 participants completed a causal inference task with 6 trials, in which they inferred how to operate the machines based on their interpretation of the drawings. Insofar as more effective visual explanations placed greater emphasis on causally relevant parts of the machine and included more symbols, we predicted that participants would be faster and more accurate when inferring which interventions would activate the machine when cued with visual explanations, relative to depictions. **Results:** Surprisingly, participants were were actually less accurate (explain: 68%; depict: 72%, $b = -0.792$, $z = -3.34$, $p = 0.832e - 3$; Fig. 1C, left) and slower to respond (correct trials only: explain: 3629ms; depict: 3580ms, $b = -0.096$, $z = -0.948$, $p = 0.343$; Fig. 1C, right) when cued with visual explanations, relative to depictions.

4 Conclusion

Leveraging semantic analyses on the characteristics distinguishing visual explanations from depictions, we found the collected explanatory drawings contained distinct emphasis on abstract, functional information. However, results suggest there may be a potential mismatch between what people believe should be included in an effective visual explanation and what semantic content may *actually* be most useful for conveying causal knowledge. Taken together, insights from these studies may lead to both improved visual communication tools, as well as a deeper understanding of how we encode abstract knowledge in visual form.

References

1. Pearl, J.: Causality. Cambridge University Press, Boston (2009)
2. Boyd, R., Richerson, P.J., Henrich, J.: The cultural niche: Why social learning is essential for human adaption. Proceedings of the National Academy of Sciences **108**(Supplement 2), 10918–10925 (2011)
3. Tversky, B.: The cognitive design of tools of thought. Review of Philosophy and Psychology **6**(1), 99–116 (2015)
4. Huey, H., Walker, C., & Fan, J.: How do the semantic properties of visual explanations guide causal inference? (2021)