Learning Online via Prompts to Explain

Joseph Jay Williams
Stanford University
Littlefield 253, 365 Lasuen St
Stanford, CA, 94305 USA
josephjaywilliams@stanford.edu

Geza Kovacs
Stanford University
Stanford, CA, USA
gkovacs@stanford.edu

Caren Walker
University of California Berkeley
3410 Tolman Hall
Berkeley, CA, 94720
Caren.walker@berkeley.edu

Samuel G. Maldonado
University of California Berkeley
3410 Tolman Hall
Berkeley, CA, 94720
Samuel.g.maldonado@gmail.com

Tania Lombrozo
University of California Berkeley
3410 Tolman Hall
Berkeley, CA, 94720
lombrozo@berkeley.edu

Abstract

Prompting learners to explain their beliefs can help them correct misconceptions upon encountering *anomalies* -- facts and observations that conflict with learners' current understanding. We have developed a way to augment online interfaces for learning by adding prompts for users to explain a fact or observation. We conducted two experiments testing the effects of these explanation prompts, finding that they increase learners' self-correction of misconceptions, though these benefits of explaining depend on: (1) How many anomalies the prompts require people to explain, and (2) Whether anomalies are distributed so that individual observations guide learners to correct ideas by conflicting with multiple misconceptions at once.

Author Keywords

Explanation; self-explanation; learning; generalization; statistics; misconceptions; anomalies;

ACM Classification Keywords

H.4 Information Systems Applications; H.5 Information interfaces and presentation; K.3.1 Computer Uses in Education; J.4 Social and Behavioral Sciences

Introduction

People are constantly learning from interactions with software, which can often involve revising their existing beliefs based on anomalies or facts that contradict what they currently believe. For example, in online contexts, users learn right/wrong answers while solving exercises, informally learn surprising facts while browsing like Wikipedia articles, and also learn how to work with interfaces when their observations conflict with their current understanding.

In this work, we design and evaluate interfaces to help learners effectively resolve conflicts between what they understand or predict, and what they observe.

One strategy for helping learners realize misconceptions is to prompt them to reflect and provide explanations for unusual facts. Prompts asking learners for explanations are also a valuable tool for augmenting interfaces, because they can easily be added to existing content, and because prompts can help guide learners' processing of information, while still letting learners take charge in interacting with the information being presented.

Related Work

While it might be more intuitive that people learn by *receiving* explanations, evidence from psychology and education suggests that certain ways of prompting people to explain and answer questions can help learning, even without feedback on correctness of learners' explanations [1, 5]. Studies document the benefits of explaining broadly, in populations, tasks, topic ranging from five year olds learning math to professionals learning to use Excel [1, 3].

Helping learners realize their misconceptions by prompting for explanations is challenging, because people often ignore contradictory information [2]. Or, asking them to elaborate could lead them to entrench beliefs. How can prompts be designed that ensure users interact with and processes facts that contradict existing knowledge in a way that leads to productive learning and behavior? [3]

The *Subsumptive Constraints Account* [5] proposes that prompts to explain "why?" will be particularly effective, because they do not merely boost engagement, but selectively drive learners to *subsume* what is being explained as an instance of an underlying principle. We therefore focus on comparisons of "why?" to other prompts, and explore further questions about how the pattern-seeking constraints of "why?" prompts can be effectively leveraged by presenting anomalous information that guides learners to from misleading to reliable patterns.

Current Research

We conducted two experiments in the context of learning statistical rankings [4]. These experiments looked at how effects of explaining depend on amount of contradictory information, and whether the quantity of contradictory information is the dominant factor, or whether one designs the presentation of information to specifically rule out existing beliefs.

Our findings show that prompting learners for explanations help learners realize and correct their misconceptions, though the benefits of this intervention depend on the number and distribution of anomalies that users are asked to explain.

Experiment 1

The learning task in both experiments was adapted from [4]. Participants had to learn from observations of the relative ranking of five/six pairs of samples from different populations (grades of two students from different classes) what the basis for ranking was and

Tom was ranked more highly by the university than Sarah.

Sarah got 85% in a Sociology class, where the average score was 79%, the average deviation was 8%, the minimum score was 67%, and the maximum score was 90%.

Tom got 69% in a Art History class, where the average score was 65%, the average deviation was 3%, the minimum score was 42%, and the maximum score was 87%.

Explain why this student was ranked higher. [*Explain* condition]
Write out any thoughts you have about this information. [*Write Thoughts*]

Sarah was ranked more highly by the university than ***Tom***.

Sarah got 85% in a Sociology class, where the average score was 79%, the average deviation was ***3%***, the minimum score was 67%, and the maximum score was 90%.

Tom got 69% in a Art History class, where the average score was 65%, the average deviation was ***8%***, the minimum score was 42%, and the maximum score was 87%.

Figure 1: (a) illustrates what participants saw as they observed each ranked pair and answered a question prompt. Only one prompt was shown, either the *explain* **or** the *write thoughts* prompt. Table 1 depicts how the ranked pair in **(a)** would have been ranked by each of the four principles. This reveals that the correct ranking for the pair in **(a)** is only consistent with the "more deviations above the average" principle, because **(a)** is an *anomaly* with respect to the "higher score", "greater distance from average", and "closer to maximum" rules.

As is discussed in Experiment 2, ranked pairs like **(a)** were used in the *overlapping anomalies* condition because the same observation is an anomaly to all three of the non-normative rules.

(b) shows how a near-identical ranked pair could be produced that was instead consistent with all four rules (and so not an anomaly with respect to any rule), simply by switching the class average deviations of the pair in **(a)**. This was how number of anomalies was manipulated.

Type of Information	Sarah	Tom
Personal Score	85%	69%
Class Average	79%	65%
Class Maximum	90%	87%
Class Deviation	8%	3%

Ranking Rule	Use of rule	Higher ranked
Higher score	85 > 69	Sarah
Greater distance from average	(85 – 79) > (69 – 65)	Sarah
Closer to maximum	(90 – 85) < (87 – 69)	Sarah
More deviations above the average	(85-79)/8 < (69-65)/3	Tom

Table 1: How different principles would rank students.

how to similarly rank new observations. This paradigm was useful for investigating people's belief revision because previous research has shown that people often have misconceptions about such statistical rankings [4], and find it challenging to use observations or *anomalies* that contradict these misconceptions in revising beliefs and identify the correct principle, a problem shared with many other important educational tasks [2].

Learning Materials: Ranked Pairs of Students

Participants studied five pairs of students from different classes whose academic performance had been ranked by the university, and were told that their goal was to learn the ranking system employed. Figure 1 shows two ranked example pairs, detailing the information provided. Each pair stated which student was ranked higher by the university, and reported each student's score (e.g., 83%), as well as the class's mean score (e.g., 73%), average deviation (e.g., 8%), and maximum & minimum scores. Participants were given the definition of mean and average deviation in the introduction.

Principles for ranking pairs of students

1 shows four ways in which participants could rank the two students from Figure 1(a). The only principle that was consistent with the observed rankings of all five student pairs was the fourth one. This "more deviations above the average" principle predicted that the higher ranked student would be the one with a score that was more deviations above their class mean. While the other three *non-normative* principles (e.g., "higher score" student is ranked higher) were consistent with some of the observed rankings, they never correctly applied to all. Moreover, these principles are termed *non-normative* because while previous research has found they are commonly used and consistent with the intuitive statistical knowledge many people posses [4], they are less reasonable as a basis for ranking from the perspective of a correct understanding of statistics,

while the “more deviations above the average” corresponds to core concepts like z-scores or standardized normal scores.

Participants

The participants were adults recruited online through the Amazon Mechanical Turk marketplace (659 in Experiment 1 and 261 in Experiment 2). Participants were asked to complete 1 HIT asking them to answer a 20-40 minute survey using an external platform, with compensation around $3.00/hour.

Experiment Design & Procedure

EXPLAIN VERSUS WRITE THOUGHTS

To examine the effect of prompts to explain “why?”, both Experiment 1 and 2 randomly assigned half of the participants to have one of two kinds of question prompts displayed below the ranked pairs. These are shown in Figure 1. Responses were typed into a text box below the prompt.

In the *explain* condition the prompt was to “Explain why this student was ranked higher”. In the *write thoughts* condition the prompt was to “Write out any thoughts you have about this information”, so they could use a range of strategies to engage with the observation, while ensuring they paid attention and engaged in comparable levels of verbalization.

NUMBER OF ANOMALIES

To explore how the effect of explaining “why?” was enhanced by the way in which information contradicting existing beliefs was presented, both Experiment 1 and 2 manipulated whether there were few or many observations that contradicted or were anomalies with respect to common misconceptions.

In Experiment 1, in the *1 out of 5* condition, for each of the three non-normative principles, there was one ranked pair that was an anomaly with respect to the principle (and four ranked pairs consistent with the principle). In the *4 out of 5* condition, there were four ranked pairs that were anomalies with respect to each of the principles, and one that was consistent with it. All five ranked pairs were consistent with the “more deviations above the average” principle. Also, note that the anomalies with respect to the non-normative rules were *overlapping*, meaning that each ranked pair observation that was an anomaly with respect to one non-normative principle was *also* an anomaly with respect to the others, so the correct “more deviations above the average” principle was the only one not ruled out. In Experiment 2 we directly manipulated this factor, using *distributed* anomalies.

Measure of belief revision

Participants had to rank four unranked pairs of students that pitted the “more deviations above average” principle against the three non-normative principles, both *before (Pre-Test)* and *after (Post-Test)* studying the ranked pairs. Scoring “accuracy” as a participant ranking in accordance with the “more deviations above average” principle, accuracy change from before or after study served as the main dependent measure of belief revision and is shown in the Figures as a function of experimental condition.

Explanation’s effects depended on number of anomalies

Figure 2 shows the change in accuracy, as a function of the two independent factors, which was analyzed using a 2 (task: explain vs. free study) x 2 (number of anomalies: single vs. multiple) ANOVA.

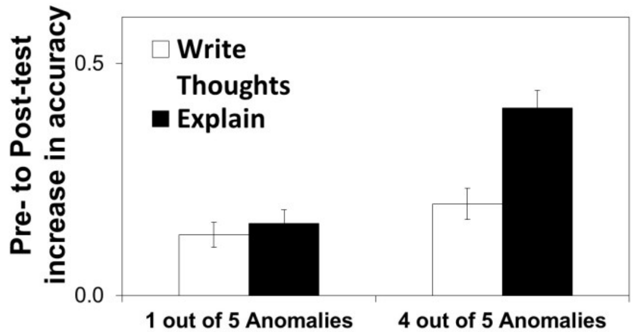

Figure 2: Change from pre-test to post-test in accuracy on anomalous items. Bars represent +/- 1 standard error of the mean.

While both receiving multiple anomalies and engaging in explaining overall promoted learning, these main effects ($ps < 0.01$) were superseded by the interaction between explaining and number of anomalies, $F(1, 659) = 8.20$, $p < 0.01$).

Explaining was beneficial when multiple anomalies were present, so that it promoted revision of beliefs about the non-normative principles, and discovery and use of the correct relative-to-deviation principle.

Experiment 2

Experiment 2 further manipulated whether the participants worked with observations that were designed so that contradictions to all the misconceptions *overlapped* in the same observation, so that information was presented to strongly rule out all but the correct principle, or whether such contradictions were *distributed* across multiple observations.

Experiment 2		2 out of 6 anomalies condition											
		Overlapping condition						*Distributed* condition					
		1	2	3	4	5	6	1	2	3	4	5	6
Ranking Rule	Higher score	✖	✖	✔	✔	✔	✔	✖	✖	✔	✔	✔	✔
	Greater distance from average	✖	✖	✔	✔	✔	✔	✔	✔	✖	✖	✔	✔
	Closer to maximum	✖	✖	✔	✔	✔	✔	✔	✔	✔	✔	✖	✖
	More deviations above average	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔
		4 out of 6 anomalies condition											
		Overlapping condition						*Distributed* condition					
		1	2	3	4	5	6	1	2	3	4	5	6
	Higher score	✖	✖	✖	✖	✔	✔	✖	✖	✖	✖	✔	✔
	Greater distance from average	✖	✖	✖	✖	✔	✔	✔	✔	✖	✖	✖	✖
	Closer to maximum	✖	✖	✖	✖	✔	✔	✖	✖	✖	✖	✔	✔
	More deviations above average	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔

Table 2: The difference between having anomalies be *overlapping* versus *distributed* in Experiment 2.

Table 2 shows how ranked pairs were consistent and anomalous with respect to the principles in both the *overlapping* and *distributed* anomaly conditions. In the ***Overlapping*** condition the same ranked student pair was anomalous with respect to all three non-normative rules, as shown in. In the ***Distributed*** condition a ranked student pair that was anomalous with respect to one non-normative rule was consistent with one or more of the *other* non-normative rules.

Results

Figure 3 shows the change in accuracy, as a function of the three independent factors, which was analyzed using a 2 (task: explain vs. free study) x 2 (number of anomalies: single vs. multiple) x 2 (distribution of anomalies: overlapping vs. distributed) ANOVA.

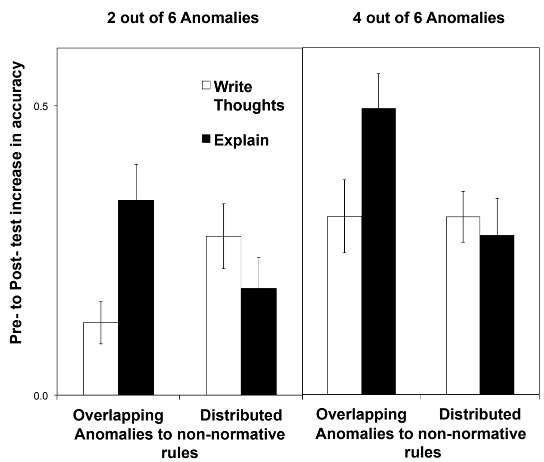

Figure 3: Change in accuracy on anomalous items from pre-test to post-test in Experiment 2. Bars represent +/- 1 standard error of the mean.

As in Experiment 1, receiving multiple anomalies boosted learning, $F(1, 261) = 8.94$, $p < 0.01$. There was also an interaction of task with the distribution of anomalies. Explaining benefited participants when the anomalies overlapped, but explaining was no longer beneficial when the anomalies were distributed, $F(1, 261) = 11.23$, $p < 0.01$.

The quantity of anomalous information was therefore not the key determiner in prompts to explain producing belief revision. Instead, it was critical that observations were designed so that prompts to explain could help people rule out the non-normative principles.

References

[1] Fonseca, B.A. & Chi, M.T.H. (2012). Instruction based on self-explanation. In Mayer, R. & Alexander, P. (Eds.), The Handbook of Research on Learning and Instruction. Routledge Press.

[2] Chinn, C. A., & Brewer, W. F. (1993). The role of anomalous data in knowledge acquisition: A theoretical framework and implications for science education. *Review of Educational Research*, 63, 1-49.

[3] Legare, C.H., Gelman, S.A., & Wellman, H.M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Development*, 81, 92-944.

[4] Schwartz, D.L., & Martin, T. (2004). Inventing to Prepare for Future Learning: The Hidden Efficiency of Encouraging Original Student Production in Statistics Instruction. *Cognition and Instruction, 22*(2), 129-184.

[5] Williams, J.J. & Lombrozo, T. (2010). The role of explanation in discovery and generalization: evidence from category learning. *Cognitive Science*, 34, 776-80.