An Instructional Factors Analysis of an Online Logical Fallacy Tutoring System

Nicholas Diana, John Stamper, and Ken Koedinger

Carnegie Mellon University,
5000 Forbes Ave, Pittsburgh, PA 15213, United States
ndiana@cmu.edu,john@stamper.org,koedinger@cmu.edu
http://www.springer.com/lncs

Abstract. The proliferation of fake news has underscored the importance of critical thinking in the civic education curriculum. Despite this recognized importance, systems designed to foster these kinds of critical thinking skills are largely absent from the educational technology space. In this work, we utilize an instructional factors analysis in conjunction with an online tutoring system to determine if logical fallacies are best learned through deduction, induction, or some combination of both. We found that while participants were able to learn the informal fallacies using inductive practice alone, deductive explanations were more beneficial for learning.

Keywords: Informal Logic, Instructional Factors, Analysis, Online Tutoring Systems, Argumentation, Ill-defined Domains

1 Introduction

In late November of 2016, Ipsos Public Affairs surveyed Americans about the accuracy of various real and fake news headlines. They found that respondents rated fake news headlines as "somewhat" or "very" accurate 75% of the time [19]. Given the fact that most (62%) of adults get their news from social media outlets where fake news is most rampant [9], the need for a citizenry capable of evaluating evidence and arguments is more crucial than ever. We propose that educational technology provides an opportunity for accessible, evidence-based instruction on these essential critical thinking skills. To test this claim, we built an online tutoring system designed to teach people how to identify informal logical fallacies.

The recognized importance of critical thinking skills is not new. In 1972, a study conducted by the American Council on Education found that 97% of the 40,000 faculty members interviewed considered fostering critical thinking skills to be the most important goal of undergraduate education [16]. Over two decades later, a similarly large study by Paul et al. [17] of 66 colleges and universities found that 89% of faculty saw critical thinking as a primary objective of their instruction. Note that these faculty members are reflecting on a world where "fake news" was an article about lizard people in the National Enquirer. Citizens can no longer simply consume information, assuming that a wide distribution or high production value implies a certain level of legitimacy. Being an informed citizen, the foundation of civic engagement, requires evaluating sources of information and recognizing poorly constructed arguments.

The ability to recognize when an argument is built upon a faulty premise is a key facet of critical thinking. In the Common Core Standards for English Language Arts & Literacy, the ability to identify fallacious reasoning or distorted evidence is listed alongside basic communication skills like "evaluating a speaker's point of view" and "[their] use of evidence and rhetoric" as key measures of a student's career or college readiness. The same standards suggest that the cost of failing to adequately teach these kinds of reasoning skills is high. In the introduction to the standards, the authors stress that the importance of these skills extends well beyond the students' academic lives, arguing that students must "reflexively demonstrate the cogent reasoning and use of evidence that is essential to both private deliberation and responsible citizenship in a democratic republic" [12].

There is, unfortunately, little evidence to suggest these aspirational goals are met in practice. In the same study of 66 colleges and universities, Paul et al. [17] found that only a small percentage of faculty members (9%) were teaching for critical thinking on a daily basis. Even then, these generally underwhelming efforts to teach critical thinking skills are only available to students attending colleges and universities. Few opportunities for learning these skills exist for citizens not receiving a post-secondary education. Citizens lack accessible, evidence-based ways to learn critical thinking skills. We propose that educational technology (e.g., educational games, intelligent tutoring systems, etc.) may play a role in filling that vacuum.

Unfortunately, most research and interventions that utilize intelligent tutoring systems focus on well-defined domains such as math and science [7]. This bias towards well-defined domains may be due to an increased cultural focus on STEM education [10], or simply due to the fact that problems in well-defined domains tend to have solutions that are (generally) clear-cut and therefore more amenable to interpretation by a computer system. That said, there has been some work demonstrating that intelligent tutoring systems can be effective learning tools in ill-defined domains. For example, Ashley and Aleven [3], have demonstrated that intelligent tutoring systems can be used to teach law students to argue with cases. Similarly, Easterday et al. [8] has shown that educational games can be used to teach skills such as policy argumentation. With respect to argumentation specifically, research has shown the effectiveness of digital argument diagramming tools on teaching argumentation [18], and fostering critical thinking skills [11].

That is not to say that building educational technology for ill-defined domains does not present unique challenges. In their review of research on intelligent tutoring systems for ill-defined domains, Lynch et al. [15] describe how characteristics such as a lack of formal theories or the inability to verify "correct" solutions make designing systems to teach these domains challenging. In our domain of informal logic, for instance, it would be problematic to simply ask participants if an argument is fallacious because an infinite number of factors could contribute to whether or not a participant considers an argument is valid. Take the following argument for example:

I was just outside; it's raining.

This incredibly innocuous statement would rarely elicit a critical thought in normal conversation. However, in the context of a tutoring system designed to test critical thought, even this mild statement might be met with critiques like: *How long have you been inside? Maybe it stopped raining. It may be raining there, but it's not raining everywhere. How do you define raining? Maybe it's just misting.*

To overcome some of these challenges and make teaching informal logic more tractable, we narrow our focus to the relatively more structured but under-investigated area of informal logical fallacies. By focusing on fallacies, we can avoid the problems associated with focusing on how valid an argument is, and instead focus on teaching learners to identify specific patterns of faulty argumentation. Instead of asking if an argument is fallacious, we can ask if the argument contains a specific fallacy. This unfortunately does not solve the problem of ambiguity completely. Even the presence or absence of a specific fallacy in an argument can be debatable if the argument is sufficiently nuanced. We mitigate this concern by making the arguments we present as unambiguous as possible, albeit at the potential expense of authenticity (see [4]).

In addition to examining the feasibility of teaching informal logical fallacies using an online tutoring system, we also demonstrate the utility of tutoring systems as a platform for researching how students learn to identify patterns of faulty reasoning. Most textbooks teach informal fallacies with a combination of general definitions and specific examples. However, the relative effectiveness of these two different kinds of instruction is unclear. The inclusion of a definition for the fallacy seems intuitive, but it may be the case that students can learn to identify fallacies simply by seeing many different examples (i.e., through induction alone). Note that the Common Core Standards called for these reasoning skills to be *reflexive*, suggesting an automaticity that corresponds to inductive skills. We frame our investigation using the Knowledge-Learning-Instruction (KLI) framework [13], which suggests that the best instruction for teaching a specific skill depends on the type of process used to learn that skill. With respect to the current study, we ask whether identifying informal fallacies is primarily an inductive process or a deductive, sense-making process.

In this work we utilize the agile nature of online tutoring systems to explore how people learn to identify logical fallacies. We tested five different instructional designs. Each design shares some instructional features with one or more of the others. Rather than compare these designs directly, we can leverage the degree to which the different designs overlap by using an instructional factors analysis. The instructional factors analysis determined the relative effectiveness of each of the three main instructional components (inductive practice, expert-explanations, and self-explanations) present in different combinations across the five designs. We found that:

1. We can successfully teach the Appeal to Ignorance fallacy using an online tutor.
2. Participants could learn the fallacy through inductive practice alone.
3. However, deductive explanations (via Expert-Explanations) were more effective than inductive practice.

The main contribution of this work is the use of an instructional factors analysis to determine the relative impact of two traditional types of instruction on teaching logical fallacies. Our results will inform the design of future informal fallacy tutoring systems, and demonstrate the usefulness of intelligent tutoring systems for teaching and researching informal logical fallacies.

2 Methods

A total of 86 participants were recruited using Amazon Mechanical Turk [5]. Participants were required to be located in the United States and were compensated at a rate of $10 USD/hour to participate in the experiment. Demographic information was collected during a post-test questionnaire. Of all participants, 45% were female, 46% were college-educated, and the average age was 31 years old. 77% of participants identified as Caucasian, 8% as Black or African American, 8% as Asian, 3% as Hispanic, and 2% did not identify with the listed options or identified with more than one. None of these demographic factors were significant predictors of performance.

2.1 Informal Logical Fallacies

Informal logical fallacies are patterns of bad argumentation, where the premises fail to support the conclusion. Informal fallacies are distinct from formal fallacies, which are errors in the *form* of an argument (e.g., If P then Q; Q; Therefore P). In contrast, informal fallacies more often contain errors in the content of the argument (e.g., mischaracterizing an opponent's argument). While there are many types of informal fallacies, some are more common than others. Ad Hominem (attacking the person rather than their argument), for example, has become mainstream enough to be mentioned by name during U.S. Presidential Debates. Because prior knowledge and conceptions of well-known fallacies might impact our results, for this work we chose to focus on a lesser known informal fallacy: Appeal to Ignorance.

The Appeal to Ignorance Fallacy Appeal to Ignorance is an informal logical fallacy that involves using the absence of evidence as evidence itself. For example, if I were to argue, "Bigfoot exists because nobody has proven he doesn't exist" I would be employing the Appeal to Ignorance fallacy. While this simple example is illustrative, in reality use of the Appeal to Ignorance is often more subtle. During one of his witch hunts in the 1950s, Joseph McCarthy produced a list of 81 names of people he claimed to be Communists working inside the State Department. When asked about one of the names on the list, McCarthy infamously said:

> *I do not have much information on this except the general statement of the agency that there is nothing in the files to disprove his Communist connections.*

As with most informal logical fallacies, the boundary of what is and isn't fallacious is also often less clear in the real world. For example, the justice system in the United States operates under the assumption of innocence until proven guilty. While this assumption appears to be directly at odds with evidence-based logic, the distinguishing feature here is the thorough, methodical investigation that (at least theoretically) is present in every case. Tindale [20] suggests that we can distinguish an Appeal to Ignorance by asking if there has "been a reasonable effort to search for evidence, or is the absence of evidence for or against something really negative evidence arising from the attempts to show otherwise?" As mentioned previously, we deliberately avoided these kinds of subtleties when designing the problems used in the tutoring system to make the correct answer as clear and unambiguous as possible. Examples of the kinds of arguments implemented in the tutoring systems can be seen in Figures 1, 2, and 3.

2.2 The Fallacy Tutor

In order to test the relative effect of different kinds of instruction on teaching logical fallacies, we built a simple online tutoring system for teaching one kind of fallacy (Appeal to Ignorance). The online tutoring system was built using the Cognitive Tutor Authoring Tools (CTAT) [1], and hosted on TutorShop, a web-based learning management system. Log data was sent from TutorShop to DataShop [14] for storage and analysis.

Inside the tutor, participants could encounter three types of problems: Fallacy/No Fallacy problems, Expert-Explanation problems, or Self-Explanation problems. The number of each type of problem the participant encountered was determined by the experimental condition the participant was assigned to. We tested five different instructional designs, each with a different number of each problem type (see Table 1).

Fallacy/No Fallacy Problems Fallacy/No Fallacy (FNF) problems involved presenting the participant with an argument, and asking whether the argument contains an Appeal to Ignorance or not. After selecting an answer, participants received correctness feedback (i.e., correct or incorrect). Unlike the other kinds of problems (Expert-Explanation and Self-Explanation), FNF problems did not provide an explanation of why the argument does or does not contain an Appeal to Ignorance. Explanations and definitions were intentionally omitted from FNF problems, as they were designed to promote inductive rather than deductive reasoning.

Table 1. Number of Problem Types for Each Condition

Condition	Instruction	Practice
Baseline	6 Fallacy/No Fallacy	6 Fallacy/No Fallacy
4EE 2SE	4 Expert-Explanation, 2 Self-Explanation	6 Fallacy/No Fallacy
2EE 2SE	2 Expert-Explanation, 2 Self-Explanation	6 Fallacy/No Fallacy
4EE	4 Expert-Explanation	6 Fallacy/No Fallacy
2EE	2 Expert-Explanation	6 Fallacy/No Fallacy

Expert-Explanation Problems Some participants received either two or four Expert-Explanation problems (depending on condition). Expert-Explanation problems involved presenting the participant with an argument, indicating that it does or does not contain an Appeal to Ignorance, and then providing an explanation as to why it does or does not. In the context of our tutor, these Expert-Explanations provided direct instruction and were designed to promote deductive reasoning.

Self-Explanation Problems In addition to Expert-Explanation problems, some participants received two Self-Explanation problems (depending on condition). Requiring students to check their understanding by providing an explanation in their own words has been shown to be an effective instructional practice [2]. In our tutor, Self-Explanation problems involved presenting the participant with an argument, indicating that it does or does not contain an Appeal to Ignorance (as additional scaffolding), and then asking them to explain why it does or does not contain an Appeal to Ignorance. After providing their explanation, they were given an expert explanation that they could compare their explanation to. Participants received no correctness feedback from the system about their explanation.

2.3 Instructional Factors Analysis Model

To determine the relative effectiveness of these different types of problems, we generated an Instructional Factors Analysis Model (IFM). IFM is a cognitive modeling approach that is useful for modeling student performance when more than one instructional intervention is used. IFM has been shown to outperform other cognitive modeling approaches such as Additive Factor Models (AFM) and Performance Factor Models (PFM) when multiple instructional interventions were involved [6].

In our case, the instructional factors of interest are the three different types of problems participants may see in the tutoring system. The general goal of

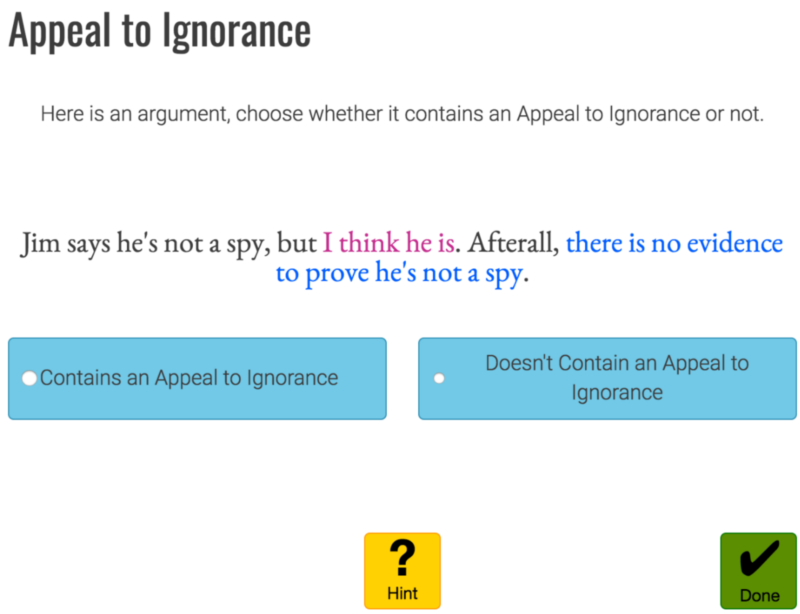

Fig. 1. Screenshot of the tutor interface during a Fallacy/No Fallacy problem. After selecting an answer, participants will be given correctness feedback only.

this model is to discover which types of problems are the most beneficial for learning. More specifically, we were interested in whether problems that promote deductive reasoning (Expert-Explanation and Self-Explanation problems) are more effective than inductive practice (Fallacy/No Fallacy problems). This approach has two key advantages. First, if we compare the conditions to one another directly, we fail to account for any instructional overlap across conditions. Instead, an IFM model deconstructs each condition into the relevant features, giving us more detailed insights into which instructional factors are effective, regardless of condition. Second, IFM does not require that a direct observation of student performance is generated from each instructional intervention. This is crucial because both Expert-Explanation and Self-Explanation problems (as they are presently implemented) do not generate direct observations of student performance.

To implement an IFM, we first generated a table where each row corresponded to a student's attempt at a problem (see Table 2 for an example). The columns of the table corresponded to the factors of a mixed-effect model. Our fixed effects were the number of prior opportunities of each of the three kinds of problems (Fallacy/No Fallacy, Expert-Explanation, and Self-Explanation). We used *student* as a random effect. To calculate our outcome variable (Error Rate) we first calculated the Assistance Score, which is equal to the number of incorrect attempts and hint requests for a particular FNF problem. The Assistance Score

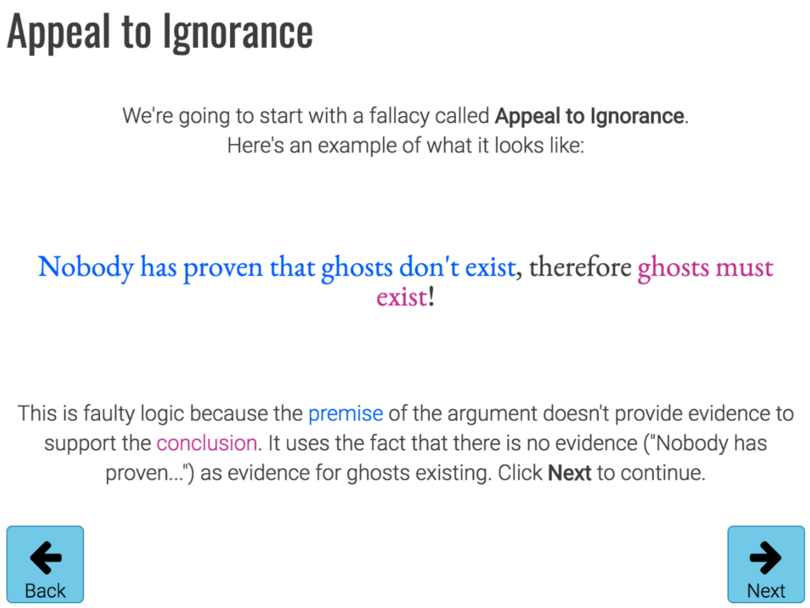

Fig. 2. Screenshot of the tutor interface during an Expert-Explanation problem which indicates whether an Appeal to Ignorance is present and explains why.

is then divided by the total number of attempts and hint requests to produce the Error Rate.

We then implemented the model using the Python library *StatsModels* using the following formulation:

$$mixedlm(ErrorRate \sim \ FNF + EE + SE + Error(student)) \tag{1}$$

Where FNF, EE, and SE represent the number prior FNF, EE, and SE problems. The term $Error(student)$ represents our inclusion of the variable *student* as a random effect. Note that our IFM implementation is slightly different from the implementation reported in [6] in that we use linear regression (rather than logistic regression) to accommodate our continuous outcome variable (Error Rate).

3 Results & Discussion

In order to determine which problem type was most effective for learning, we generated an instructional factors analysis model (IFM). Controlling for the time spent in the tutor, we found that the number of prior Fallacy/No Fallacy (FNF) problems and the number of prior Expert-Explanation (EE) problems were significant predictors of performance ($p < .001$), while the number of prior Self-Explanation (SE) problems was not. Though both FNF and EE problems seem

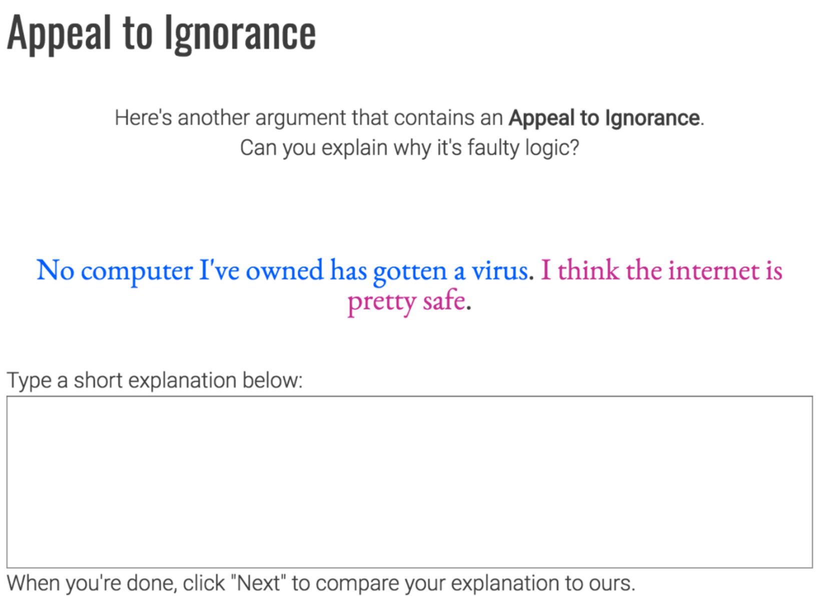

Fig. 3. Screenshot of the tutor interface during a Self-Explanation Problem.

Table 2. Excerpt of Data Inputed into the IFM. Here students 1, 2, and 3 represent students in the Baseline, 4EE 2SE, and 2EE 2SE conditions, respectively.

Student	Problem	Prior FNF	Prior EE	Prior SE	Error Rate
Stu_1	1	0	0	0	.5
Stu_1	2	1	0	0	.75
Stu_2	1	0	4	2	0
Stu_3	1	0	2	2	.6

to be instructional, EE problems had more than twice the impact ($\beta = -0.034$) on reducing Error Rate than FNF problems ($\beta = -0.015$). These results seem to suggest that instruction aimed at promoting deductive reasoning is more effective than inductive practice. While true in this case, the relationship between deductive and inductive instruction is likely different for different fallacies. Fallacies that are difficult to articulate may be more easily taught through inductive examples. In our future work, we plan to expand the tutoring system to include many different kinds of informal fallacies. This serves the dual purpose of discovering the best kind of instruction for each fallacy, while also potentially revealing the features of a fallacy that inform the kinds of instruction that should be prioritized when teaching it.

It is possible that we did not see an effect for Self-Explanation problems because of the constraints of the experimental design. Because Self-Explanation problems require participants to have a working definition of the fallacy, they must come after Expert-Explanation problems. In each of the conditions with Self-Explanation problems, participants will have seen at least two Expert-Explanation problems before they are required to explain the faulty logic themselves. It may be the case that the novelty of testing one's own understanding does not outweigh the diminishing returns of seeing another explanation of the fallacy. One can imagine a (frustratingly difficult) design that begins with Self-Explanation problems, asking participants to explain why an argument contains an Appeal to Ignorance without any explanation of what an Appeal to Ignorance is. In this hypothetical case, we may see Self-Explanation problems having a measurable, positive effect similar to or greater than that of Expert-Explanation problems. This is an avenue of research for future work.

We have demonstrated that both the deductive EE problems and the inductive FNF problems are effective instructional interventions. What remains to be seen is if participants can learn Appeal to Ignorance in the absence of any general definitions or explanations (i.e., through inductive practice alone). Our Baseline condition was specifically designed to answer this question. Recall that the Baseline condition contains only FNF problems. Participants in this condition never received a definition of Appeal to Ignorance or any explanations of why it was or was not present in an argument. The only feedback they received was whether or not they answered the problem correctly. If it is possible to learn Appeal to Ignorance through inductive practice alone, we would expect the number of prior practice opportunities to predict performance in the Baseline condition, and we found that this is indeed the case. If we consider only participants in the Baseline condition, the number of prior FNF problems is a significant predictor of performance ($p < .001$). This suggests that while deductive instruction may be more beneficial for learning, participants were still able to learn Appeal to Ignorance through inductive practice alone.

3.1 Limitations and Future Work

The ultimate goal of this work is to develop a tutor that could be deployed to late high school classrooms and freely accessible via the web for older adults. While Amazon's Mechanical Turk is a great resource for testing various implementations of the tutor, there are quality limitations that make collecting data from a classroom preferable. In our future work, we plan to expand the tutor to include many more types of informal fallacies. From this variety we hope to uncover the hidden features that make a fallacy more easily learned through either induction or deduction. Another simple, but important addition to our experiment is a measure of enjoyment. It may be the case that you can learn a fallacy through induction alone, but that the act of blindly searching for a pattern is frustrating.

4 Conclusion

The online tutoring system presented here is an initial foray into the vast and complex domain of informal logic. Nevertheless, this relatively simple system has allowed us to gain several key insights: First, it is possible to teach at least one kind of informal fallacy in an online tutoring system. Second, it is possible to learn Appeal to Ignorance using only inductive practice. However, the results from our instructional factors analysis suggest that instruction aimed at promoting deduction is more valuable than inductive practice. These insights are not only useful for the development of future tutoring systems, but offer a promising glimpse into the role that educational technology can play in creating accessible, evidence-based critical thinking instruction.

References

1. V. Aleven, B. M. McLaren, J. Sewall, and K. R. Koedinger. The cognitive tutor authoring tools (ctat): Preliminary evaluation of efficiency gains. In *Proceedings of the 8th International Conference on Intelligent Tutoring Systems*, ITS'06, pages 61–70, Berlin, Heidelberg, 2006. Springer-Verlag.
2. V. A. Aleven and K. R. Koedinger. An effective metacognitive strategy: Learning by doing and explaining with a computer-based cognitive tutor. *Cognitive science*, 26(2):147–179, 2002.
3. K. D. Ashley and V. Aleven. Toward an intelligent tutoring system for teaching law students to argue with cases. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 42–52. ACM, 1991.
4. M. Boudry, F. Paglieri, and M. Pigliucci. The fake, the flimsy, and the fallacious: demarcating arguments in real life. *Argumentation*, 29(4):431–456, 2015.
5. M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon's mechanical turk. *Perspectives on Psychological Science*, 6(1):3–5, 2011. PMID: 26162106.
6. M. Chi, K. R. Koedinger, G. J. Gordon, P. Jordon, and K. VanLahn. Instructional factors analysis: A cognitive model for multiple instructional interventions. 2011.
7. A. T. Corbett, K. R. Koedinger, and W. Hadley. Cognitive tutors: From the research classroom to all classrooms. *Technology enhanced learning: Opportunities for change*, pages 235–263, 2001.
8. M. W. Easterday, V. Aleven, R. Scheines, and S. M. Carver. Using tutors to improve educational games. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6738 LNAI:63–71, 2011.
9. J. Gottfried and E. Shearer. News use across social media platforms 2016. *Pew Research Center*, 26, 2016.
10. J. Handelsman and M. Smith. Stem for all — whitehouse.gov. `https://obamawhitehouse.archives.gov/blog/2016/02/11/stem-all`, February 2016. (Accessed on 09/28/2017).
11. M. Harrell. Assessing the efficacy of argument diagramming to teach critical thinking skills in introduction to philosophy. *Inquiry: Critical Thinking Across the Disciplines*, 27(2):31–39, 2012.
12. C. C. S. S. Initiative et al. *Common Core State Standards for English Language Arts & Literacy InHistory/Social Studies, Science, and Technical Subjects.* 2010.

13. K. R. Koedinger, A. T. Corbett, and C. Perfetti. The Knowledge-Learning-Instruction Framework: Bridging the Science-Practice Chasm to Enhance Robust Student Learning. *Cognitive Science*, 36(5):757–798, 2012.
14. K. R. Koedinger, J. C. Stamper, B. Leber, and A. Skogsholm. Learnlab's datashop: A data repository and analytics tool set for cognitive science. *Topics in cognitive science*, 5(3):668–669, 2013.
15. C. Lynch, K. Ashley, V. Aleven, and N. Pinkwart. Defining ill-defined domains; a literature survey. In *Proceedings of the workshop on intelligent tutoring systems for ill-defined domains at the 8th international conference on intelligent tutoring systems*, pages 1–10, 2006.
16. O. Milton. Will that be on the final?. 1982.
17. R. W. Paul, L. Elder, and T. Bartell. California teacher preparation for instruction in critical thinking: Research findings and policy recommendations. 1997.
18. C. Reed, D. Walton, and F. Macagno. Argument diagramming in logic, law and artificial intelligence. *The Knowledge Engineering Review*, 22(01):87, 2007.
19. C. Silverman and J. Singer-Vine. Most americans who see fake news believe it, new survey says. *BuzzFeed News (www. buzzfeed. com). URL: https://www. buzzfeed. com/craigsilverman/fake-newssurvey*, 2016.
20. C. W. Tindale. *Fallacies and Argument Appraisal*, volume 23. Cambridge University Press, Cambridge, mar 2007.