Guiding Visualization Users Towards Improved Analytic Strategies Using Small Interface Changes

Radu Jianu* David H. Laidlaw†

Brown University

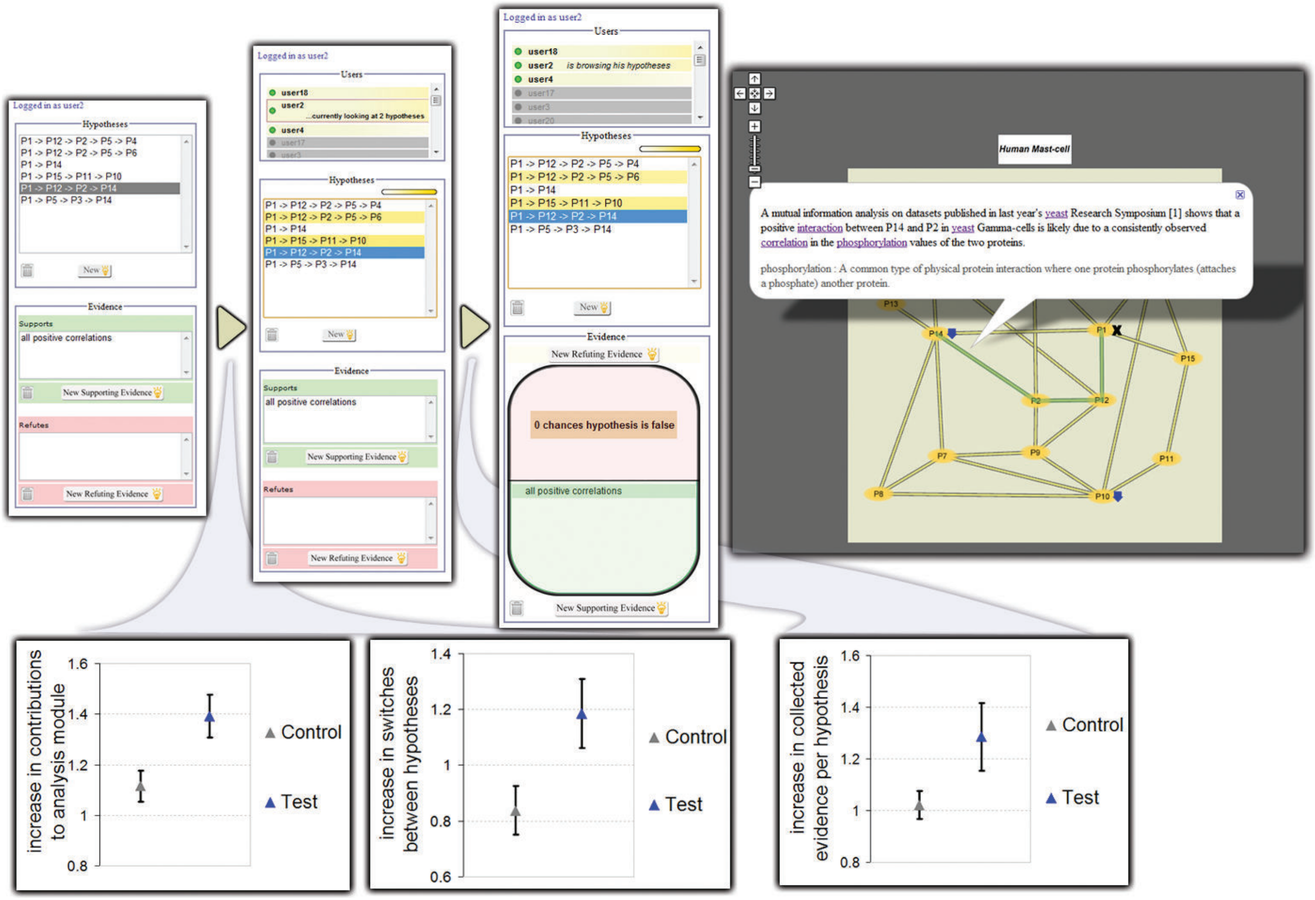

Figure 1: Subtle, non-functional interface changes in an analysis support module (top, three left panels) generated significant changes in users' analysis of a visual problem solving task (top right). A first set of changes nudged subjects to increase their use of the analysis module by 27% (bottom left, p=0.02) in an attempt to expand users' working memory. It also caused them to switch between hypotheses 27% more often (bottom center, p=0.03), indicating more consideration of alternative hypotheses. A second set of changes then lead subjects to gather 30% more evidence per hypothesis (bottom right, p=0.02).

Abstract

We provide quantitative evidence that subtle changes in a visualization system's interface can be used to alter users' analytic behaviors in targeted ways. In a controlled study subjects completed three analyses, at one week intervals, using a system consisting of a visualization and an analysis support module. A control group used one interface for all three analyses. A test group started with the same interface but then used modified versions in the following two sessions. A first set of changes, included before the second session, aimed to increase subjects' use of the system and increase their consideration of alternative hypotheses. The second set of changes, added before the last session, aimed to increase the amount of evidence collected. After the first set of changes, test subjects used the interface 27% more and switched between hypotheses 35% more than a control group. After the second set of changes test subjects collected 27% more evidence than control subjects. All observed increases are significant (p1=0.02, p2=0.03, p3=0.02). We hypothesize this approach can be used to guide visualization users unobtrusively towards improved analytic strategies.

Index Terms: Visual Analytics, Analytic Biases

*e-mail: jr@cs.brown.edu
†e-mail:dhl@cs.brown.edu

1 INTRODUCTION

Cognitive science studies have shown that human thinking is subject to heuristics and biases that may lead to suboptimal decision making [5]. Such effects have also been documented in the context of hypothesis driven analysis, an area of interest in visualization research. For example, *satisficing* [7] limits analysis to a hypothesis that is good enough, while *confirmation bias* conditions us to confirm hypotheses rather than disconfirm them [10].

Motivated by these results, we posit the following hypothesis: subtle, targeted changes in interfaces of visualization systems can unobtrusively guide users towards better analytic strategies. In support of this hypothesis we report results from a study in which subjects used a visualization system to solve three analyses at one week intervals. A control group used a single interface for all sessions while a test group was given slightly altered interfaces in sessions two and three. As hypothesized, the evolution of performance measures over the three sessions differed between the two groups. Test subjects used the system more, they considered hypotheses in parallel, and searched for more evidence. These results suggest that we can leverage interfaces to overcome analytic shortcomings.

2 RELATED WORK

Our approach is based on previous work in behavioral economics and human-computer interaction (HCI). Specifically, Thaler and Sunstein [9] (behavioral economics) popularized the terms *choice architecture* — how choices are presented to consumers —, and *libertarian-paternalism* — designing choice architectures that "nudge" consumers towards decisions in their own interest. Fogg [4] (HCI) defines *persuasive technology* as "interactive information technology designed for changing users' attitudes or behavior". Sunstein, Thaler and Fogg motivate this approach with two arguments. First, any choice architecture or computer interface necessarily influences decision-making behavior, whether intentionally or not. Second, people's choices and behaviors are not necessarily aligned with their goals. Both arguments are supported by empirical cognitive science evidence. An array of scientific results have validated the feasibility of the "nudge" approach [8, 6]. We introduce the nudge paradigm in the visualization domain and demonstrate empirically that it can further the visual analytics agenda.

3 METHODS

We asked 32 undergraduate and graduate students to solve three analyses of a similar form at one week intervals. The tasks where inspired by the proteomic domain: using evidence linked to edges and nodes of a protein interaction network to explain interdependencies of pairs of not directly connected proteins (Fig. 1, top right). Our networks borrowed proteomic terminology but the likelihood of an interaction path depended on a reduced set of rules which were explained at the beginning of the study.

We separated our subjects into control(18) and test(14) groups. The control group solved all three tasks using an analysis support module with three lists: one for storing hypotheses, and two for recording confirming and disconfirming evidence for each hypothesis (Fig. 1, top left). For test subjects this base interface was altered before the second and last sessions. We hypothesized that changes between sessions would be observed in both groups due to task-learning but that test subjects would exhibit additional artifacts which could be attributed to the interface alterations.

We designed three nudges to target three hypothesized analytic deficiencies. Due to time constraints the first two were implemented in the second session (Fig. 1 top, 2nd from left). The last was added in the third (Fig.1 top, 3rd from left). The **first nudge** aimed to relieve users' memory by increasing their reliance on the system. Our design leveraged conformity effects and motivational factors for online contributions [2, 1]: if subjects saw others actively using the interface they would do so as well. To this end online users were listed at the top of the interface and their interactions were reflected in a public status message (e.g., has entered new evidence). The **second nudge** encouraged users to consider hypotheses in parallel. We assigned each hypothesis a recency score that decayed over time and increased when users interacted with the hypothesis. Active hypotheses were then highlighted thus offering a visual reward. Finally, the **third nudge** encouraged subjects to gather more evidence. First, the evidence lists were made visually distinct. Second, if no evidence had been entered for a hypothesis, the two lists would read "0 chances that hypothesis is false" or "hypothesis is unlikely": committing to extreme cases is avoided by humans [3]. This nudge could be restricted to disconfirming evidence only, in which case it may alleviate confirmation biases [10].

4 RESULTS

We measured three indicators to test our hypothesis. The number of entered hypotheses and evidence, normalized by analysis time, served as a proxy for subjects' reliance on the interface. The number of times a subject switched between hypotheses, normalized by hypotheses count, indicated the degree to which hypotheses were considered in parallel. Finally, we recorded the number of evidence-items and divided it by number of hypotheses.

The findings are summarized in Figure 1. Test subjects contributed 35% more hypotheses and evidence to the analysis module in the second session than in the first. This compares to an increase of only 8% in the control group. The difference in switches between hypotheses was an increase of 18% in test subjects versus a decline of 17% in control subjects. The amount of evidence collected per hypothesis remained fairly constant between sessions in the control group (+2%). Test subjects however, gathered on average 29% more evidence per hypothesis in the third condition than the second. A t-test captures the significance of the observed differences: $p1=0.02$, $p2=0.03$, $p3=0.02$.

5 CONTRIBUTIONS

We introduce the *nudge* concept in the visual analytics domain: subtle changes in visualization interfaces can be used in controlled ways to guide users towards better analysis. We present results from a quantitative user study demonstrating that this approach is viable.

REFERENCES

[1] M. Ames and M. Naaman. Why we tag: motivations for annotation in mobile and online media. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, page 980. ACM, 2007.

[2] R. Cialdini and N. Goldstein. Social influence: Compliance and conformity. 2004.

[3] W. DuCharme. Response bias explanation of conservative human inference. *Journal of Experimental Psychology*, 85(1):66–74, 1970.

[4] B. Fogg. Persuasive computers: perspectives and research directions. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 225–232. ACM Press/Addison-Wesley Publishing Co., 1998.

[5] R. Hastie and R. Dawes. Rational choice in an uncertain world. *Journal of the Indian Academy of Applied Psychology*, page 107, 2003.

[6] M. Kumar and T. Kim. Dynamic speedometer: dashboard redesign to discourage drivers from speeding. In *CHI'05 extended abstracts on Human factors in computing systems*, page 1576. ACM, 2005.

[7] H. Simon. Rationality as Process and as Product of Thought. *The American Economic Review*, 68(2):1–16, 1978.

[8] R. Thaler and S. Benartzi. Save More Tomorrow: using behavioral economics to increase employee saving. *Journal of political Economy*, pages 164–187, 2004.

[9] R. Thaler and C. Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Yale Univ Pr, 2008.

[10] P. Wason. Reasoning about a rule. *The Quarterly Journal of Experimental Psychology*, 20(3):273–281, 1968.