Beyond Detection: Investing in Practical and Theoretical Applications of Emotion + Visualization

R. Jordan Crouser
Smith College
Northampton, MA
jcrouser@smith.edu

Lane Harrison
Worcester Polytechnic Institute
Worcester, MA
lane@cs.wpi.edu

Daniel Afergan
Google Inc.
Mountain View, CA
afergan@google.com

Evan M. Peck
Bucknell University
Lewisburg, PA
evan.peck@bucknell.edu

ABSTRACT
Emotion is a dynamic variable that modulates how we perceive, reason about, and interact with our environment. Recent studies have established that emotion's influence carries to data analysis and visualization, impacting performance in ways both positive and negative. While we are still in the infancy of understanding the role emotion plays in analytical contexts, advances in physiological sensing and emotion research have raised the possibility of creating emotion-aware systems. In this position paper, we argue that it is critical to consider the potential advances that can be made even in the face of imperfect sensing, while we continue to address the practical challenges of monitoring emotion in the wild. To underscore the importance of this line of inquiry, we highlight several key challenges related to detection, adaptation, and impact of emotional states for users of data visualization systems, and motivate promising avenues for future research in these areas.

Author Keywords
Emotion, affect, visualization, adaptation, theory.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
Emotion is a central facet of human experience, coloring our perception and interactions with our environment. Often considered an impediment in analytical domains, emotion may steer us towards biased positions or flawed reasoning, even when the source of emotion is irrelevant to the task at hand [4]. Yet emotion is also known to play a key role in rational reasoning, allowing us to quickly recognize and act on preferences, for example, or to choose favorably among uncertain outcomes in complex planning situations. As neuroscientist Antonio Damasio summarized: "[Emotion] allows the possibility of making living beings act smartly without having to think smartly" [4].

EmoVis 2016 Workshop, March 10, 2016, Sonoma, CA, USA.
Workshop proceedings published by Linköping University Electronic Press.

Given the pervasive role emotion plays in human behavior, it is plausible that emotion would also impact how we analyze data and interact with data visualization systems. Several studies have already begun to substantiate this connection. Recent experiments from Harrison *et al.*, for example, combined affective priming with a classic visualization perception experiment and found that emotions with positive valence (*i.e.* happiness) led to better performance when analyzing charts [7]. Other studies in data visualization and human-computer interaction have pointed to emotional states as a possible explanation for their results [1, 2]. Researchers have specifically suggested connections between emotion and visualization performance in the context of task engagement, memorability, as well as higher level facets such as creativity and decision-making.

These findings raise the question of whether it is possible to design systems that reliably detect emotion and respond to it in an intelligent way. Detection remains challenging, however, even after decades of research investigating diverse information sources and analysis techniques. Common approaches include analyzing facial features, posture, and various physiological signals. Yet none of these approaches work perfectly, nor are they available in all contexts in which data visualization systems are used.

We argue that even moderately accurate sensing techniques, such as today's consumer-grade physiological sensors, can cater to the user in beneficial ways through proper adaptation techniques. Over the past decade, studies using brain sensing have yielded many compelling examples of beneficial adaptation with imperfect sensors (for a review, see [11]). This success is largely due to the development of better models of adaptation alongside better consumer-level hardware and detection algorithms. Taken together, these elements allow system designers to readily connect user tasks and goals with adaptation strategies in creating emotion-aware systems.

Breakthroughs in sensors that detect the body's natural signals allow the user to wear lightweight sensors while having normal interactions with a computer and make physiological sensing possible at consumer-level scale. *Physiological computing* "has the potential to extend the communication bandwidth of

HCI and enable a dynamic, individualised dialogue between user and system," [6] without any effort on the part of the user. By monitoring user biological signals, we can extract information about the user's cognitive state and use this as a system input to control the visual elements onscreen.

In order to create systems that successfully interpret and respond to human emotion, the research community must develop robust models to describe the dynamic influence of emotional state on visualization performance. At the same time, it is critical to consider the potential advances that can be made even in the face of imperfect models and sensing modalities. In this position paper, we highlight several key challenges related to detecting and adaptating to human emotion. We also discuss the impact of emotional states for users of data visualization systems, and motivate several promising avenues for future research in these areas.

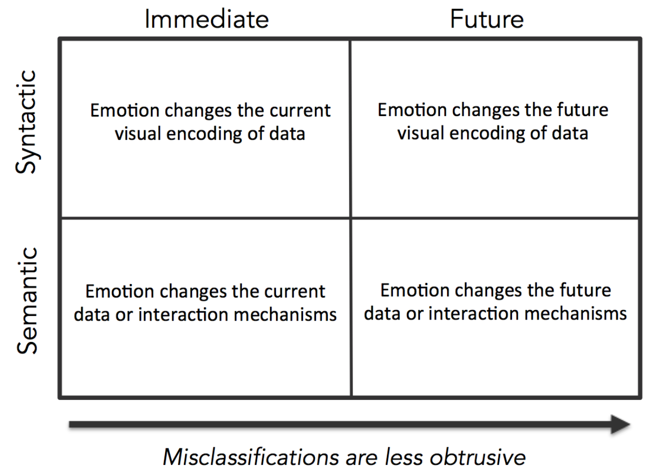

Figure 1. Framework of adaptive systems.

DESIGNING WITH IMPERFECT INPUT

One of the challenges of defining the design space of emotion and visualization is that we often turn to interaction mechanisms that necessitate high levels of accuracy. For example, biofeedback applications are valuable only to the extent that they are accurate. Providing users with visual feedback that actively works against their intuition may cause users to lose a sense of control with the interface, reducing their trust levels. This places innovation in a difficult place - should we wait for detection mechanisms to reach a certain threshold before focusing on certain applications?

Instead, we draw on our experience designing passive brain-computer interfaces (BCIs) with physiological monitoring. Looking more broadly at user state, we built several adaptive applications that relied on unreliable classifications of user workload [9]. The accuracy of our models varied widely between individuals, and there was no way to get ground truth and figure out when the model was misclassifying the user. These systems focus on *implicit input*, user contexts that the system knows is input but that the user does not actively choose to share with the system [10]. However, despite these challenges, our systems quantifiably improved user interaction. This enables us to explore the real-world challenges of passive BCI despite being years away from robust sensing of cognitive state.

We believe that it is critical for researchers to explore the interaction space of emotion + vis *now* instead of waiting until models improve beyond an arbitrary threshold. Because laboratory or consumer-grade physiological monitors are prone to noise and artifacts, and emotion is a complex construct with multiple dimensions that are difficult to capture, we must consider and anticipate the inevitability of these misclassifications, and continue to advance the field using strategies to minimize their impact.

One promising direction for dealing with imperfect input is sensor fusion. Sensor fusion systems integrate multiple physiological sensors into one classification model, with the hopes of improving accuracy or specific inferences that could not be determined with only one sensor [8]. By using techniques such as moving averages, confidence value of predictions (label smoothing), or hybrid sensor fusion, we can increase the reliability of classifications and limit the harm of a small number of misclassifications. As long as we are fairly confident of user state before triggering an adaptation, we can limit the rate of mistakes and provide an overall gain to the user.

WHEN CAN WE RESPOND TO EMOTIONS?

The challenge of identifying emotion in the wild is not trivial. Because of this, there is a danger that the focus of the community overemphasizes the improved detection of affect, with little to show for itself on the other end. To counteract this, we also need to contextualize those improvements within the goals of the visualization or interface: how will better sensing lead to *meaningful improvement* in people's analytical tasks and workflows?

As we move forward applying emotions to visualizations, we can look to the physiological computing community for lessons. Stephen Fairclough has this to say about the state of the art: "Constructing a system that can detect a range of psychological states is pointless if [the] adaptive repertoire of the machine is unable to respond to those psychological states in an intelligible fashion" [5].

We want to make sure that the user is not perturbed by the adaptations a system makes, and that the adaptations be resilient to misclassifications. Because the prediction might not always be correct, a visualization designer should view the physiological input as an augmentation to traditional input devices (keyboard, mouse, touchscreen), and not as the main source of input [11]. The designer should avoid irreversible or mission-critical adaptations, and instead make subtle, helpful changes, that the user might not even recognize or attribute to the system adaptations. Zander proposes that these passive systems can be evaluated along three key dimensions: *complementarity*, or lack of interference with other input mechanisms; *composability*, or potential to stack with other monitors; and *controlled cost*, or the effort of calibration and price of mispredictions [12].

Categorizing Implicit Interactions with Emotion

Building upon the framework for implicit interaction in passive brain-computer interfaces by Solovey et al. [11], we can begin to consider how emotion can be used to potentially drive mixed-initiative systems for visualization. We categorize adaptive mechanisms into four categories - immediate syntactic, immediate semantic, future syntactic, and future semantic - as shown in Figure 1 and discuss applications in visualization.

- **Immediate syntactic:** Emotion is used to change or refine the visual form of the data. This can be accomplished by zooming, filtering or other visualization tasks. For example, a system might detect that a user is frustrated trying to understand data of a particular form, and dynamically alter its representation to aid understanding. Or a visualization could "tag" data with emotional responses as a form of meta-data to be used in later exploration or analysis.

- **Immediate semantic:** Emotion is used to trigger changes in the interaction mechanisms in a visualization. If the user seems frustrated, we can zoom in on a particular subset of data to allow her to focus on a small subset of high-priority information.

- **Future semantic:** A person's emotion is used to inform the system as to what data should be shown next. For example, based on a person's emotional response while engaging with a subset of datapoints, a system might suggest visualizing other data that is related or contains similar properties.

- **Future syntactic:** A person's emotion is used to inform the visual encodings of data seen in the future. For example, an intelligent system might learn that particular affective states modify interaction with a visualization. During a later interaction, the system could recommend a more optimized representation of the same data.

By expressing this design space more explicitly, we may be able to identify opportunities for innovation even as we continue to develop our models for detecting emotion. Looking at the classification above, we identify opportunities to provide for the value for the user, but only if we do it in unobtrusive ways.

For example, constant updates across the entire interface or changes in the display format may be jarring and unsettling for users and disrupt their ability to form cohesive mental models of the system. While immediate adaptations have potential, they requires a high level of accuracy since the adaptations occur directly within a person's focus. This is particularly dangerous in the context of visualizations since our perceptual system is sensitive to changes in certain visual features such as movement or color changes.

Future changes can be effective because we can completely change the system without alerting or disrupting the user. However, it may be more difficult to predict what state the user will be in when the change occurs, making it challenging to evaluate the efficacy of these changes.

MODELING EMOTION'S ROLE AND IMPACT

In addition to the challenges raised in determining when emotion plays a role in data visualization, we presently lack appropriate mechanisms for evaluating the systems we build. In-house experimentation and in situ studies can help us determine whether or not our systems are useful, but they fall short of explaining *why* we see the results we do. This often results in the recycling of known techniques; generalizing our results to new domains is challenging, and we are left to speculate about the role of emotion in producing observed effects on user behavior. To overcome this, it is important that we develop theoretical models and corresponding language of the role of emotion in visualization systems to improve our ability to reason about their performance and design.

In visualization, and particularly in the arena of incorporating and designing for human emotion, interest in the development of real-world implementations [3] has far outpaced the development of theoretical measures. The current trajectory for research in emotion+visualization is largely focused figuring out how to leverage emotion in visualization in ways that measurably impact performance. However, in the absence of a rigorous theoretical framework in which to ground the development of new algorithms, researchers must rely on intuition and some deeply-rooted assumptions about the role of human emotion in order to design new systems. Using tacit knowledge regarding emotional responses to which we believe humans are predisposed, we build systems that capitalize on these responses. These systems will then be used as evidence that the chosen method for adapting to (or exploiting) human emotion works; a sort of "proof-by-example".

We argue that developing a theoretical language for describing the role of emotion is of critical importance to the study and design of visualization systems. Mechanisms for drawing parallels at the algorithmic level and identifying areas where existing approaches are redundant or inefficient will enable us to design more effective systems in the future. In addition, reporting theoretical arguments along with the observed performance of the system would greatly improve study reproducibility, as well as help isolate the effects of interface design and other implementation details.

The importance of understanding human emotion as part of a larger computational framework is not limited to improving visualization design. Augmenting our existing models to incorporate human emotion can expand our understanding of what can be computed, as did the development of probabilistic and parallel computation. The development of theoretical measures for human emotion may play a significant role in the broader acceptance of emotion as an important component of visualization systems.

CALL TO ACTION

Over-emphasizing the detection of emotion is tempting. Models of emotion can be clearly evaluated, and progress can be quantifiably defined by improvements in classification accuracy. However, we should avoid the trap of improving emotion models just for the sake of improving the model. We propose that by bolstering our efforts on the theoretical foundations and practical applications, we can identify high-impact applications of emotion, and begin to answer the question of "how good is good enough?" with respect to our models. We may discover that the most compelling intersections of emotion and visualization do not in fact require high accuracy to achieve the desired results.

From an applied perspective, we should be careful not to view emotion exclusively as a tool for evaluating visualization. While this application is compelling, it requires high levels of trust and accuracy in our interpretation of signals - something that may not be reliable for years. Instead, we believe that focusing on the design space of *emotion-driven* applications has the potential for significant impact even using models with modest accuracy. In the near term, we can consider the applications of implicit interaction on future modifications of a visualization, both at the semantic and syntactic levels. By using emotion to inform the design or presentation of future data presented to the user, we can build personally attentive portraits of data that do not necessarily rely on robust models. Finally, investing in practical applications has the added potential of capturing the imagination of other researchers or developers, propelling this young field forward.

From a theoretical perspective, we need to internalize that emotion exists as part of a larger human-machine interactive system, in the context of a specific task and environment. We cannot consider human performance and cognitive workload as a static measure that is unaffected by the human's current state. The role of emotion cannot be ignored or be considered in a vacuum, nor is it only one-dimensional. Instead, we must strive to consider all aspects of the closed-loop system, and strive to make modifications in real-time as a user's emotional state changes and as the task requirements change.

In assessing the merits of various tools for developing emotion-aware systems, we should not allow the perfect to become the enemy of the good. At the same time, we challenge this interdisciplinary community use these systems as a lens to better understand *how humans fit* into the grand scheme of computational tools, and to develop models that incorporate not only their power but their inherent messiness as well.

REFERENCES

1. Scott Bateman, Regan L Mandryk, Carl Gutwin, Aaron Genest, David McDine, and Christopher Brooks. 2010. Useful junk?: the effects of visual embellishment on comprehension and memorability of charts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2573–2582.
2. Nick Cawthon and Andrew Vande Moere. 2007. The effect of aesthetic on the usability of data visualization. In *Information Visualization, 2007. IV'07. 11th International Conference*. IEEE, 637–648.
3. Daniel Cernea, Andreas Kerren, and Achim Ebert. 2011. Detecting insight and emotion in visualization applications with a commercial EEG headset. In *SIGRAD 2011 Conference on Evaluations of Graphics and Visualization-Efficiency, Usefulness, Accessibility, Usability*. 53–60.
4. Antonio Damasio. 2008. *Descartes' error: Emotion, reason and the human brain*. Random House.
5. Stephen Fairclough. 2015. We Need To Talk About Clippy. (2015). http://physiologicalcomputing.org/2015/03/we-need-to-talk-about-clippy/
6. Stephen H Fairclough. 2009. Fundamentals of physiological computing. *Interacting with computers* 21, 1 (2009), 133–145.
7. Lane Harrison, Drew Skau, Steven Franconeri, Aidong Lu, and Remco Chang. 2013. Influencing visual judgment through affective priming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2949–2958.
8. Danilo P Mandic, Dragan Obradovic, Anthony Kuh, Tülay Adali, Udo Trutschell, Martin Golz, Philippe De Wilde, Javier Barria, Anthony Constantinides, and Jonathon Chambers. 2005. Data fusion for modern engineering applications: An overview. In *Artificial Neural Networks: Formal Models and Their Applications–ICANN 2005*. Springer, 715–721.
9. Evan M Peck, Daniel Afergan, and Robert JK Jacob. 2013. Investigation of fNIRS brain sensing as input to information filtering systems. In *Proceedings of the 4th Augmented Human International Conference*. ACM, 142–149.
10. Albrecht Schmidt. 2000. Implicit human computer interaction through context. *Personal Technologies* 4, 2-3 (2000), 191–199.
11. Erin Treacy Solovey, Daniel Afergan, Evan M Peck, Samuel W Hincks, and Robert JK Jacob. 2015. Designing implicit interfaces for physiological computing: Guidelines and lessons learned using fNIRS. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 6 (2015), 35.
12. Thorsten O Zander, Christian Kothe, Sabine Jatzev, and Matti Gaertner. Brain-Computer Interfaces. In *Brain-Computer Interfaces (Human-Computer Interaction Series)*. Springer London, 181–199.