Progressive Disclosure

Empirically Motivated Approaches to Designing Effective Transparency

Aaron Springer
Computer Science
University of California Santa Cruz
Santa Cruz, CA, USA
alspring@ucsc.edu

Steve Whittaker
Psychology
University of California Santa Cruz
Santa Cruz, CA, USA
swhittak@ucsc.edu

ABSTRACT

As we increasingly delegate important decisions to intelligent systems, it is essential that users understand how algorithmic decisions are made. Prior work has often taken a technocentric approach to transparency. In contrast, we explore empirical user-centric methods to better understand user reactions to transparent systems. We assess user reactions to transparency in two studies. In Study 1, users anticipated that a more transparent system would perform better, but retracted this evaluation after experience with the system. Qualitative data suggest this arose because transparency is distracting and undermines simple heuristics users form about system operation. Study 2 explored these effects in depth, suggesting that users may benefit from initially simplified feedback that hides potential system errors and assists users in building working heuristics about system operation. We use these findings to motivate new progressive disclosure principles for transparency in intelligent systems.

CCS CONCEPTS

- Human-centered computing-Human computer interaction (HCI)

KEYWORDS

Transparency, Intelligibility, Intelligent Systems, Machine Learning, Mood, Explanation, Error, Progressive Disclosure.

ACM Reference format:

Aaron Springer and Steve Whittaker. 2019. Progressive Disclosure: Empirically Motivated Approaches to Designing Effective Transparency. In *24th International Conference on Intelligent User Interfaces (IUI'19)*, March 17–20, 2019, Marina del Rey, CA, USA. ACM, NY, USA, 14 pages.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

IUI '19, March 17–20, 2019, Marina del Rey, CA, USA

© 2019 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-6272-6/19/03 \$15.00

<https://doi.org/10.1145/3301275.3302322>

<https://doi.org/10.1145/3301275.3302322>

1 Introduction

Machine learning algorithms underlie the many intelligent systems we routinely use. These systems provide information ranging from routes to work to recommendations about criminal parole [3,8]. As humans with limited time and attention, we increasingly defer responsibility to these systems with little reflection or oversight. Nevertheless, intelligent systems face mounting criticisms about how they make decisions. These criticisms are exacerbated by recent machine learning advances like deep learning that are difficult to explain in human-comprehensible terms. Major public concerns have arisen following demonstrations of bias in algorithmic systems with regards to gender, race, and other characteristics [9,75,79]. These issues have led to calls for transparency as a solution to the unintelligible algorithms that impair the adoption of intelligent systems [32,63,83].

Algorithmic transparency is needed for many reasons. Greater transparency potentially increases end user control and improves acceptance of complex algorithmic systems [43]. It can also promote user learning and insight from complex data, which is important as humans increasingly work with complex inferential systems for analytic purposes [43,71]. Transparency can also enable oversight by system designers. Without such transparency it may be unclear whether an algorithm is optimizing the intended behavior [33,51], or whether an algorithm accidentally promotes negative, unintended consequences (e.g. filter bubbles in social media; [10,64]). Given these issues, it is increasingly possible that transparency, i.e. “a right to explanation”, may become a legal requirement in some contexts [29]. These issues have led researchers to argue that machine learning must be ‘interpretable by design’ [1] and that transparency is essential for the adoption of many intelligent systems, e.g. for medical diagnoses [32,83].

While such calls for transparency are well-motivated, it remains unclear exactly how to enact them in practice. Extensive research about operationalizing transparency has emerged in the machine learning community but no clear consensus has resulted [1,21,82]. Deciding exactly how to implement transparency is difficult—there are numerous implementation trade-offs involving accuracy and fidelity. Making a complex

algorithm understandable to end users might require simplification, which often comes at the cost of reduced accuracy of explanation [46,71]. For example, methods have been proposed to explain deep neural network algorithms in terms of more traditional machine learning approaches, but these explanations necessarily present approximations of the original algorithms deployed [53]. These studies often approach transparency from a technical perspective, asking: “what is possible from an algorithmic standpoint?” rather than “what does the user need?”

However, some recent empirical studies attempt to examine the effects of transparency on users. But these studies reveal puzzling and sometimes contradictory effects. In some settings there are expected benefits: transparency improves algorithmic perceptions because users better understand system behavior [42,43,49]. But in other circumstances, transparency has other quite paradoxical effects. Transparency may erode confidence in a system, with users trusting it less because transparency led them to question the system even when it was correct [49]. Providing system explanations may also undermine user perceptions when users lack the attentional capacity to process complex explanations, for example while they are executing a demanding task [12,84]. Overall, these results indicate mixed evidence for the benefits of transparent systems.

The above research suggests that we have yet to identify the appropriate interaction paradigms to present transparency. Machine learning research communities are forging ahead with foundational research on how to generate transparent systems [82] but studies often stop short of actually testing these systems with users [46,48]. User evaluation is critical because, as we have seen, user reactions to transparency show quite contradictory results [12,41,47]. Some research suggests that the way we present transparency may account for these contradictory results [26,42]. Our research seeks to bridge this gap between generating explanations and user-centric presentation. In two studies, we explore users’ direct reactions to a transparent personal informatics system that interprets emotions that users express in written text. The domain of emotional personal informatics is fruitful for transparency research for several reasons. Emotional analytics is perceived as a key application by many users [38]. And users are the ultimate experts on their own emotions, providing them with ground truth assessments about whether the system is correct. Further, making accurate predictions about emotions is also difficult, allowing us to examine a key challenge for intelligent systems: how to deal with system error.

We examine user preferences and reactions to transparency by comparing two versions of a working analytics system that predicts users’ emotions from written text. Each version uses the same underlying algorithm. The first, transparent, version makes overall predictions about the user’s affective state and supplements these overall predictions with incremental real-time feedback showing how the algorithm made these judgments. The second non-transparent version simply predicts the user’s overall affective state, without underlying transparency feedback. We also examine the role of cognitive load in

explaining these preferences and explore problems with current presentations of transparency. We address the following research questions:

- RQ1: Do users prefer to use transparent systems? Do they prefer systems providing transparent feedback or those that simply present overall predictions? (Study 1)
- RQ2: Do cognitive load and distraction affect preferences for transparency? What other factors influence reactions to transparency? (Studies 1 and 2)
- RQ3: How might we support effective transparency? (Studies 1 and 2)

1.1 Contribution

Much recent work on transparency has focused on technical explorations of self-explanatory systems. In contrast, we take an empirical user-centric approach to better understand how to design transparent systems. Two studies provide novel data concerning user reactions to systems offering transparent feedback vs overall prediction information. In Study 1, users anticipated that a transparent system would perform better, but retracted this evaluation after experience with the system. Qualitative data suggest this may be because incremental transparency feedback is distracting, potentially undermining simple heuristics users form of system operation. Study 2 explored these effects in more detail suggesting that users may benefit from simplified feedback that hides potential system errors and assists users in building working heuristics about system operation. We use these data to motivate new progressive disclosure principles for presenting transparency in intelligent systems.

2 Related Work

2.1 Folk Theories of Algorithms

A wealth of prior work has explored issues surrounding algorithm transparency in the commercial deployments of systems for social media and news curation. Social media feeds are often curated by algorithms that may be invisible to users (e.g., Facebook, Twitter, LinkedIn). At one point, many users were unaware that Facebook newsfeeds filtered the posts that their friends made [24]. These users reacted in surprise and sometimes anger when they were shown the filtered posts that were ‘missing’ from their newsfeed. Later research showed that many Facebook users develop ‘folk theories’ of their social feed [22], which are imprecise heuristics about how the system works, even going so far as to make concrete plans based upon their folk theories. This work also showed that making the design more transparent or seamful, allowed users to generate multiple folk theories and more readily compare and contrast between them [22].

Other work has illustrated issues regarding incorrect folk theories in the domain of intelligent personal informatics systems, showing specific challenges in how users understand these systems. Users are prone to blindly believing outputs from

algorithmic systems, a phenomena referred to as algorithmic omniscience [23,38,76] and automation bias [17,58]. For example, KnowMe [81] is a program that infers personality traits from a user's posts on social media based on Big Five personality theory. KnowMe users were quick to defer to algorithmic judgment about their own personalities, stating that the algorithm is likely to have greater credibility than their own personal statements (e.g., "...At the end of the day, that's who the system says I am..."). Similar results were shown in [38]; participants expected intelligent personal informatics systems to serve as ground truth for their experiences and even attributed superhuman qualities to these devices, e.g., "...[it] could tell me about an emotion I don't know that I am feeling...". Other experiments indicate the risk of such trust, showing that users may believe even entirely random system outputs are moderately accurate [76]. Similarly, giving users placebo controls over an algorithmic interface shows corroborating results [80]; users with placebo controls felt more satisfied with their newsfeed. Without a standard of transparency in intelligent systems, it may be easy to deceive end-users into believing they are using a real system; this is a dangerous proposition when apps can be so easily distributed.

2.2 Transparency

There is a long history of studying transparency and intelligibility in automated systems [6]. However, the results often indicate contradictory effects on user perceptions. Many experiments have indicated that transparency improves user perceptions of the system [20,48]. Others have shown that interventions that simply show prediction confidence improve users system perceptions [4]. In extreme cases, animations that simulate transparency can cause users to be overconfident about systems even when they err [26].

Other studies show less positive effects for user perceptions of transparent systems. Exposing users to a hypothetical transparent system can lead them to question the system, resulting in reduced agreement with the system [49]. However, the effect may be opposite for high certainty systems—transparency only increases user agreement. Other work concludes that any hint of error in an automated system will decrease trust [59]. More recent work indicates other effects; explanations of how a system is working may increase trust [42] but explanations that include too much information may be harmful to user trust. However, these effects are dependent upon the amount of expectation violation that a user experiences. User expectation violation occurs when the system makes a prediction that a user did not expect [42,76,80]. Transparency increases positive perceptions of a system when the user experiences expectation violation; however, transparency decreases positive system perceptions when expectations are not violated [78]. Ideally transparency should build user confidence in a system, whether or not the user is experiencing expectation violation. However, it seems that research communities have not found the correct interaction paradigms to achieve this.

Recently, the machine learning community has begun grappling with issues of transparency. This may be due to the

rise of more inscrutable methods like deep learning or from legal requirements arising from the European Union's GDPR. Some machine learning models are "inherently understandable" such as linear models and Generalized Additive Models [52,82]. These understandable models can be "explained" to users simply through the linear contributions of their features. Other algorithms such as deep neural nets and random forests are inscrutable, making it difficult to explain how input features match to output predictions [82]. For example, seemingly random noise that is unnoticeable to humans can cause deep neural network image classifiers to fail catastrophically [28]. Various approaches aim to make these inscrutable algorithms understandable. These approaches often rely on approximating the inscrutable algorithm through a simple local or linear model than can be explained to the end user [53,71]. However, even with these methods that result in "inherently understandable" models, there is no clear consensus how to convey these models to users in an understandable way. Many such attempts at transparency are not tested with users or simulate user studies [5,57,60]. The absence of real user feedback makes it challenging to operationalize transparency in ways that positively impact users.

2.3 Explanation and Persuasion Theory

People interact with computers and intelligent systems in ways that mirror how they interact with other people [62,70]. Given that transparency is essentially an explanation of why a model made a given prediction, we can turn to fields such as psychology and sociology for guidance about operationalizing explanations. These fields have a long history of studying explanation. One approach is to model causal explanation as a form of conversation which is governed by common-sense conversational rules [36] such as Grice's maxims [30]. In addition, when explanation is needed and a communication breakdown occurs this is remedied by a phenomenon known as *conversational repair*. Conversational repair is interactional, participants in the conversation collaborate to achieve mutual understanding; this often happens in a turn-by-turn structure with repeated questions and clarifications [73]. These theories would indicate that we should operationalize transparency in ways that fit human communication and repair strategies.

Additionally, we see parallels between how people interact with intelligent systems and theories of persuasion. The Elaboration Likelihood Model (ELM) is a dual process model of persuasion [66]. The ELM posits that two parallel processes are engaged when a person evaluates an argument, similar to Kahneman's conception of System 1 and 2 thinking [40]. The central processing route involves careful consideration of the argument and complex integration into a person's beliefs. The central route is often engaged in high stakes decisions. The peripheral route in contrast focuses on heuristic cues such as the attractiveness of the speaker, the listener's current affect, the number and length of the arguments, and other cues not directly related to the content of the argument. Prior work on intelligent systems seems to align with this dual process model [42], people understand systems through peripheral routes if their

expectations are met, only engaging in central processing when their expectations are violated. This is also demonstrated in the context of Google search suggestions; where users can feel the cost of processing explanations outweighs their benefits [12]. Transparency needs to be operationalized in ways that allow users to understand transparency through both cursory heuristic routes and also through focused effort.

2.4 Emotional Analytics

Our focus is on how users interact with intelligent personal informatics (PI) systems which are being increasingly deployed within commercial [86–88] and research domains [7,25,37,54,85]. PI systems track how a person behaves on some dimension, whether physical, emotional, or mental, and may suggest improvements to this behavior through customized feedback and recommendations [37,68]. Such personally relevant data potentially allows users to analyze and modify their behaviors to promote well-being [15,16,39]. In the current study, we explored user reactions to textual emotional analytics, where an algorithm interprets the emotional characteristics of a reported personal experience. Other work has shown that emotional analytics is evaluated by users as one of the key applications of PI technology [38].

Figure 1: The non-transparent E-meter shows only the overall emotion ratings from the written text.

Emotional analytics is a fruitful domain for transparency research for other reasons too. First it allows users to directly engage with personally relevant data. Other transparency work has often asked users to evaluate algorithms in hypothetical scenarios where participants read about or watch algorithmic deployments and decisions [48–50]. Instead our aim was to have users directly experience the algorithm, evaluating it in situ as it made decisions about personally generated data [41]. A further critical aspect of emotional interpretation is that users are knowledgeable about their own feelings and experiences, allowing them to directly compare algorithmic interpretations with their own personal evaluations of those emotional experiences. This contrasts with other applications of smart algorithms, such as medical diagnostics, where the end user may

not be a domain expert. In these contexts, users might be less able to evaluate the results of algorithmic interpretations. In addition, emotion is highly variable between individuals and previous research demonstrates difficulty in accurately predicting emotion from text [47,69]. One of the major challenges with intelligent systems is handling errors; explaining algorithmic prediction in this difficult domain of emotional analytics allows us to better understand how users make sense of output that contains errors.

3 Research System: E-meter

We designed a system called the E-meter that tests users’ reactions to a transparent system that actively interprets their own data. The E-meter system and protocol have been successfully deployed in multiple previous experiments [76,78]. The E-meter predicts the emotional valence of a user’s written personal experience. The E-meter is shown in Figures 1 and 2, showing non-transparent and transparent versions of the same algorithm. The system was described to users as an “algorithm that assesses the positivity/negativity of [their] writing”. It was built as a client-side web application using HTML5 and JavaScript, technologies that give us flexibility in deploying to various populations including in-lab studies and Amazon Mechanical Turk experiments.

There are a number of tradeoffs that need to be made when testing user reactions to transparency. In contrast to many previous hypothetical scenario-based studies, where users are presented with descriptions of a system along with its outputs [48–50] we instead chose to deploy a working system to better engage and elicit direct user reactions to personally relevant data. The system also provided real-time feedback which indicated to users that the system was authentic and not a deception [76]. But implementing a real-time transparent system imposes trade-offs regarding performance of the underlying model. In particular, we needed to sustain performance across multiple hardware platforms. We therefore limited the complexity of our machine learning model and constrained our underlying feature set so that it could run in real-time within a browser.

3.1 Machine Learning Model

Little empirical work has examined how users respond to even simple operating models of transparency. This motivates our current approach, which we have successfully deployed in several prior studies [76,78]. We chose a modeling approach that is accurate enough to be convincing, but simple enough to explain to end-users. We therefore implemented a unigram-based regression model. While such a model may be less accurate than state of the art methods using deep neural networks [19], it provides a straightforward operationalization of transparency that is intuitively understandable while still appearing accurate and convincing to users [78]. And although the algorithm makes some errors, evaluating user reactions to these errors is an important element of our approach.

The emotion detection algorithm works as follows: each word written by the user is evaluated for its positive/negative emotion association in our model. If a word is found in the model, the overall mood rating in the system is updated. This constitutes an incremental linear regression that recalculates each time a word is written. In prior experiments, users judged the same algorithm to be accurate but not perfect; the median user response for accuracy was ‘Accurate’ (6 on a 7-point Likert scale with the highest rating (‘7’) being ‘Very Accurate’) [78].

Models were trained on text gathered the EmotiCal deployment [37,77]. In that deployment, users reported and analyzed activities and emotions over several weeks by writing short textual entries about daily experiences and directly evaluating their mood in relation to those experiences. This data gave us a gold-standard supervised training set where users labeled their own descriptions for underlying affect. We trained the linear regression on 6249 textual entries and mood scores from 164 EmotiCal users. Text features were stemmed using the Porter stemming algorithm [67] and then the top 600 unigrams were selected by F-score, i.e. we selected the 600 words that were most strongly predictive of user emotion ratings. Using a train/test split of 85/15 the linear regression tested at $R^2 = 0.25$; mean absolute error was .95 on the target variable (mood) scale of (-3,3). As we noted above, previous users found the system to be accurate to their experience, judging median accuracy to be 6 out of 7 on a scale where 7 is ‘Very Accurate’.

Figure 2: The transparent E-meter shows how the overall emotion rating was calculated by revealing which words were weighted positively or negatively.

3.2 System Versions

The underlying machine-learning model was evaluated by users for two different system versions: non-transparent and transparent.

Non-transparent version: As users wrote, the E-meter showed the system’s overall, document-level, interpretation of the emotional valence of their writing. If the overall text was interpreted negatively, the gauge emptied to the left and turned more red (see Figure 1). If the text was judged to be positive, the

meter filled the gauge to the right and turned more green. This continuous scale feedback represents the coarse global information that many machine learning systems currently display. Such systems give an overall rating but do not offer the user an insight into the detailed workings of the underlying algorithm.

Transparent version: In contrast, the transparent version provided fine-grained feedback about how the algorithm was making its prediction. This transparency was operationalized as word-level feedback; in this system version the E-meter signaled the affective association of each word the user typed. Individual words are highlighted and color coded according to how the underlying algorithm interpreted that word’s affect (Figure 2). If a word is highly associated with positive mood then it will be highlighted green, whereas a word associated with negative mood will be highlighted red. This incremental feedback allows users to see how each individual word they write contributes to the overall E-meter rating. Furthermore, words remain highlighted as users continue to type allowing them to assess each word’s relative contribution to the overall score.

This form of transparency offers users insight into the underlying word-based regression model driving the E-meter evaluation; it depicts how the regression model correlates each word with positive or negative emotion to arrive at an overall weighting for the entire text that the user has entered. The fact that the visualization is persistent also allows users to reexamine what they have written, reconciling the overall E-meter rating with the fine-grained word-level transparency.

Of course, we could have operationalized transparency in other ways. Other researchers have implemented transparency through natural language explanations [42] and diagrams [49]. However, in our case we can convey key aspects of the underlying system through word highlighting. In addition, our operationalization allows the answering of counterfactual questions, an important part of explanation [56,82]. Highlighting the text helps directly convey to the user what drives the algorithm and gives clear clues about the underlying linear model. In addition, by varying the colors of the highlighting we also show how the model is interpreting the specific words.

4 Study 1

Study 1 was intended to explore user reactions to transparent and non-transparent versions of the algorithm. For each version we assessed user evaluations both before and after using the system as well as their perceptions of cognitive load.

4.1 Method

Participants experienced both versions of the E-meter system. After using each version, they were asked a series of follow-up questions. The study design was counterbalanced; half the users experienced the non-transparent system version first. Questions and procedure were carefully piloted and had been used before in multiple prior studies.

4.1.1 Users. We recruited 100 users who had previously completed a subset of the Psychological General Well-being Index (PGWBI) [31]. Following recommendations from an IRB panel, the PGWBI screening was intended to exclude users who might find the emotional reflection distressing. Users were recruited from Amazon Mechanical Turk and paid \$3.33. The evaluation took 14.68 minutes on average. Following prior methodological recommendations in [18], we eliminated 26 respondents based on their responses to open ended questions, leaving us with a sample of 74 users. This study was approved by an Institutional Review Board.

4.1.2 Measures. Before actually using the system, participants saw simulations of both transparent and non-transparent versions of the system and were asked to *predict accuracy* for each. We simulated both versions of the system showing what happened as filler Latin text was typed. We used Latin text because we wanted high-level system comparisons that were not based on users' reactions to specific English words. Following the animations, we asked the *predicted accuracy* question "This program evaluates the positivity/negativity of emotional experiences that users write about. How accurate or inaccurate do you think this program would be for you? The program works with English also." Users then provided qualitative explanations for their ratings—"Please give 2 or more reasons for the accuracy ratings you made on the previous page."

Participants then began their writing activity. Following a protocol deployed in prior work, users were presented with one of the two system versions with the instructions "Please write at least 100 words about an emotional experience that affected you in the last week." After this system experience, users completed the TLX *workload* assessment [34]. Users then answered the questions: "How positive or negative did you feel your writing was?" (*subjective affect*), "How positive or negative did the E-meter assess your writing to be?" (*system affect*), "How accurate or inaccurate was the E-meter in its assessment of your writing?" (*retrospective accuracy*), "How trustworthy or untrustworthy did you find the E-meter system?" (*subjective trust*). These questions were all 7-point Likert items. For example, the subjective trust questions items were from "Very Untrustworthy" to "Very Trustworthy".

Users then repeated this process for the other system version. After using both versions, users answered a final *experience-based system preference* question "If you were to use the E-meter again, which system would you prefer?". They then supplied reasons for this: "Please give 2 or more reasons for the choice you made above".

4.2 Results

System Perceived as Moderately Accurate and Trustworthy: Overall, the median user found the E-meter to be 'Slightly Accurate' and 'Slightly Trustworthy'. Both retrospective accuracy and subjective trust distributions were bimodal (due to the 7-point Likert item containing an infrequently used neutral item) and there was no difference between conditions ($p = .24$, $p = .41$). This replicates our previous work [78] where the median user found the E-meter to be 'Accurate'.

Transparent version is predicted to be more accurate before usage: Before any hands-on experience with the system, participants generated *predicted accuracy* judgments for both the transparent and non-transparent system versions. We statistically compared the difference in these system evaluations using a paired t-test. Participants anticipated greater accuracy for the transparent system, $t(73) = 5.452$, $p = 0.022$, although the effect was small—means were 4.24 and 4.57 respectively. Qualitative user comments supported these anticipated benefits. On being asked to explain why they expected the transparent system to be more accurate, users drew attention to the benefits of transparent color-coded feedback giving a clearer sense of how the system was operating and boosting their confidence that the system was operating appropriately. Participant 73 wrote "The second one had a legend with it and actually changed the color of the words I would have written. ... It was also more catchy and the colors stood out to me." In the same vein, participant 50 wrote: "One meter is more transparent than the other. I can see how it works. I feel more confident in knowing exactly how it comes up with its answers. I tend to think it is more reliable." Overall then, before actually using either E-meter version, users anticipated that a system offering word-level transparency would be more accurate.

Recall that after experiencing each system version, we asked users for *retrospective accuracy*, and a final *experience-based preference* about which system version they would choose for future usage. User perceptions of *retrospective accuracy* with both versions of the system highly correlate with their final *experience-based system preference* in a logistic regression model (p 's = [0.019, 0.0001]). Given only retrospective accuracy scores from both versions, we can predict the version choice with 69.5% accuracy in a 5-fold cross-validated test. Therefore, knowing users' *predictive accuracy* was higher for the transparent version, we would expect that this would lead to users ultimately preferring the more transparent version of the system.

After experiencing the system there was no preference for either system version. However, this positive predictive evaluation of transparent system accuracy did not persist after actual experience with the system when we analyzed final *experience-based preferences*. As we expected, many of those who preferred the transparent system after usage did so because it illustrated the inner workings of the system. Participant 72 said "I know what the first [transparent] version is doing. I cannot tell what the second version [non-transparent] is doing. Because the second version does not give real feedback, I cannot make an informed decision when writing if I should be using it or not." Participant 28 concurred, saying "I think it [the transparent version] provides more engaging feedback and helps me better understand the reasons it gives for the amount in the meter."

However, to our surprise, the benefits of transparency did not generalize to all users. After experience with both systems overall users were evenly split in which version of the system they preferred to use in the future: fifty percent of participants (37) said that they would prefer the non-transparent version if they were to use the system again. The other 50% (37) chose the transparent condition. Consistent with these preference

judgments, participants also showed no overall differences in trust after using the two different system versions ($t(73)=.910$, $p=.343$). How can we interpret these findings? While users had better impressions of the transparent condition initially, those preferences disappear after using both systems. It is important to note here that although the only objective differences between the systems lay in their transparency, users seemed to treat them as operating quite differently. As we will see later, even though participants knew both versions of the system existed at the start of the experiment, they seemed to attribute different qualities to each version after experiencing them.

Overall, both the final system choice and trust showed no differences between system versions despite people being confident initially that the transparent version would be more accurate. What might explain these changed perceptions after usage?

Role of Cognitive Load: One possible explanation for this changed perception is cognitive load. Transparency may demand attention and distract participants, in contrast to the non-transparent system version which presents less information [12,82]. Some participants' explanations for their *experience-based preference* seem to support this. These participants ($n=13$) cited the distracting nature of the word highlighting. Users called word-based transparent feedback “*annoying*” (P3) and “*obtrusive*” (P7). Participant 20 said that the non-transparent version was “*a lot less distracting*”. P57 said “*The individual highlighting of the words was distracting during writing; I wouldn't have minded it as much if I could turn it on and off.*” However, these subjective reports were not borne out by our quantitative analysis of cognitive load as assessed by the TLX survey. A paired t-test comparing the overall TLX measures for both versions of the system indicated no difference in workload: $t(73) = -.05$, $p=.95$.

Reduced transparency may lead users to overestimate system capabilities: Another potential reason why users may prefer the non-transparent system relates to user inferences about algorithmic capability. Our qualitative analysis of *experience-based preference* suggests that users who know less about the working of the system, seem to ascribe more advanced abilities to it. Nearly a quarter (24%) of users who chose the non-transparent E-meter as their preferred version stated that they preferred it because it took their overall writing context into account, incorporating information beyond simple lexical weightings. Participant 66 said, “*I think the [non-transparent version] takes into account everything you are writing and makes a decision better than just by focusing on word choice.*” Participant 17 concurred, saying “*I like the second [non-transparent version] as it seems to focus on the whole and not each word.*” While these non-transparent inferences are positive, they are also inaccurate. Recall that both systems use the same underlying machine learning model which uses solely individual word features.

What might explain this overestimation of system capabilities? One possibility is that non-transparent feedback can hide low-level errors from users. In contrast, many transparent condition users identified highlighted words they felt were misclassified, leading them to downgrade their system

evaluation. For example, participant 40 chose the non-transparent version, justifying it by saying: “*...the biggest reason is that the most negative thought I had was expressed by the word “isolated” in the text I wrote and the e-meter marked that one word as “Unimportant” I couldn't get past that.*” Participant 70 said: “*Some associations don't make any sense, while others do.*” In contrast, non-transparent feedback did not expose these errors. If the algorithm was behaving consistently with their overall expectations, users in the non-transparent condition judged it very positively. This did not seem to be the case in the transparent condition, users judged the system more negatively even if the overall prediction was correct.

Together these observations suggest that error hiding and the absence of detailed information in non-transparent feedback leads some participants to form approximate but positive working heuristics about how the system operates.

4.3 Discussion

Our initial hypothesis was that providing detailed, transparent word-level feedback would be more helpful to users than non-transparent overall predictions. However, our first study unearthed some unexpected findings, showing that user interpretations of transparency feedback are far from straightforward. Consistent with our initial expectations and the prior literature on transparency [26,48], users anticipated a preference for transparency before using the system. Users justified this preference by making arguments that such feedback would provide detailed incremental information about algorithmic decisions. But to our surprise, many participants did not retain this preference after using the system, at which point participants were evenly split between the systems in their trust and transparency preferences. As we had originally anticipated, some users continued to prefer the transparent version and they were likely to cite the increased insight that it facilitated into the algorithm's underlying operation. In contrast, others preferred the non-transparent version, but offered very different reasons for their preferences. Some of these users seemed to find highlighting to be distracting, although our cognitive load results do not support this. Others may have preferred non-transparent feedback because it did not expose word-level errors, potentially leading users to overestimate the competence of the underlying algorithm with the consequence that they believed it to be more advanced than it was. For these users who preferred the non-transparent version, it seemed that incremental feedback was providing more information than they required [12,42].

5 Study 2

Ideally a system should mitigate the negative distracting elements of incremental transparency while providing improved understanding to users. However, it is difficult from our initial study to know how to operationalize transparency in a way that achieves this. In order to better understand how to convey

transparency to users in effective ways we employed a semi-structured interviewing process in our second study.

We used think-aloud interviewing methods to examine in depth how the type and timing of algorithmic transparency can inform decisions about how to design effective transparency. In particular, given that Study 1 indicated that some users felt incremental feedback provided too much information, Study 2 examines letting users view increased transparency only when they explicitly request it, after they have finished writing. Overall Study 2 gathered richer contextual qualitative data to illuminate what factors influence the interpretation and uptake of transparency information. In particular we wanted to better understand why ostensibly richer feedback was not providing anticipated benefits to some participants.

5.1 Method

5.1.1 Users. Twelve users were recruited from an internal participant pool at a large United States west-coast university. They received course credit for participation. Participants average age was 19.54 years (s.d.=1.52) and 7/12 identified as female. This study was approved by an Institutional Review Board.

5.1.2 Measures. Users again completed a shortened version of the PGWBI to screen for mental health before the study [31]. Users answered a similar set of survey questions to Study 1 in a think-aloud style; however, these were primarily used to prompt explanation and structure the interview. As such, we do not present the results in this paper.

5.1.3 Procedure. The participants were randomly divided into one of two conditions. Both groups were given overall prediction feedback from the E-meter (Figure 1).

Passive Transparency: Six participants received real-time transparent word-level feedback about the algorithm's interpretation of their affect as they typed each word.

Requested Transparency: The other six only obtained transparent word-level feedback after they had finished the writing task; these users explicitly requested transparency by clicking a button labeled "How was this rating calculated?".

The researcher explained the experiment and think-aloud procedure, demonstrating a think-aloud on an email client. As in Study 1, the researcher asked participants to "Please write at least 100 words about an emotional experience that affected you in the last week." After the think-aloud writing exercise, the experimenter conducted a semi-structured interview that included an on-screen survey. After the survey, participants in the initial non-transparent condition 1 saw exactly the same screen with an added button labeled "How was this rating calculated" which they pressed to reveal the transparent word-level highlighting. The entire process took around 50 minutes.

5.1.6 Analysis. Interviews were recorded using both audio and screen-recording. Two interviews (one from each condition) were not audio recorded, thus only the remaining 10 are used for the analysis. For this qualitative analysis, responses were coded using thematic analysis [11] specifically targeting RQ3: What problems must be mitigated to support effective transparency?

5.2 Results

More Transparency is Not Always Better: The transparent version of the E-meter again operationalized transparency using color highlighting to show how each word contributes to the calculation of the overall mood score. However, several users took issue with this level of detail. P10 felt that only the "*big emotionally heavy words*" should be highlighted. Other users felt similarly, P4 talked about select "*trigger*" words that that "*trigger the foundation of the issue*" and were essential to understanding the text. P8 felt that there were a few important words and the rest just added noise: "*it's taking into account words like 'stressful' and 'regretful' and stuff but then like everything else in between adds like an extra layer that complicates it*". While the machine learning model was limited to the 600 most predictive words of mood, that model still seemed to present too many extraneous words that users felt were unimportant. From a design perspective, it may be that users need to identify a small number of clear examples of words showing strongly positive and negative affect, in order to form a working model of system operation.

Transparency May Violate User Expectations Even When the System is Correct: Similar to how participants felt there were many extraneous highlighted words, participants also often focused on specific words that they judged had been misinterpreted by the system. Many users wrote about their experiences in the first person, often using the word "I". In our machine learning model, "I" has a slightly negative connotation which confused our users, because many of them thought "I" should be neutral. P8 said "*So I was gonna say that yellow words would be neutral because it has highlighted 'I'...*". Along the same lines, when analyzing the highlighting of different words, P6 said "*I? Mmm, I don't understand that either*". Participant 5 even started conjecturing about the actual system model saying "*I doesn't seem like it would have...[participant trails off] unless people speak in objective terms when they're talking about more positive experiences*." Clearly these users don't feel that the word "I" has negative connotations. However, extant literature confirms that system feedback is correct as high usage of first person singular pronouns is correlated with depression and negative mood [65,72]. These examples indicate a problem when system feedback reveals information that contradicts the users' expectations. Even if the system is objectively correct when displaying transparency information, it can still cause users to take issue and result in poorer perceptions of that system.

User Heuristics May Contradict Transparency: Other problems arose because users formed working heuristics of how the system operated, which were sometimes contradicted by word-level feedback. Four users felt that there were discrepancies between the overall rating and the transparent word-level feedback within the system. When Participant 8 viewed the word-level transparency they started off by saying "*I'm very confused*" and then explained how they felt the overall rating should just be calculated as a ratio of positive/negative words—"*Well I just assumed that some words would be coded as positive or negative and then it would just like do a ratio of those two*."

Participant 6 explained it similarly, the overall rating showed a slightly negative emotion rating for their written passage but, in participant 6's words: *"So when you look at the comparison with the meter, and you look at my paragraph itself, right? There's more words that are highlighted in green."* These users are using simple heuristics to relate the overall document rating and word-level transparency. They feel that the overall rating should reflect a simple ratio of the word-level transparency. For example, if the highlighted words are primarily green, then the overall rating should be very positive. However, in our machine learning model, a single word such as "angry" could be rated negatively enough that it would cancel out multiple mildly positive words. Some users arrive at this correct model after consciously engaging with the system. For example, recall how participant 6 talked about how they felt the word-level highlighting and the overall rating were incongruent; however, after thinking more deeply, this participant later said: *"... it's weighing certain words, right? Because obviously these two words, right? Like 'upset' and like—er yeah, 'upset' really polarize the meter."* This quote demonstrates that this user has moved beyond their initial heuristic that all words are weighted equally; they are now noting how one word seemed to have a larger effect in the system.

Together these observations suggest that users initially form simple working hypotheses about system operation. Users seem to engage with transparency first by operating with these simple hypotheses and only scrutinizing these when their expectations aren't met. As with other areas of reasoning, it may be that in interacting with a system users first engage in rapid, approximate, System 1 thinking and only engage in deeper, more analytic, System 2 thinking when directly prompted or confronted with clearly anomalous information [40,66].

5.3 Discussion

The second study again revealed the complexity involved in presenting transparency information. Confirming our first study, we again showed that providing more detailed transparency information isn't always better. While some users saw potential benefits to word level transparency information, others argued against "complete" transparency, preferring to see only a subset of 'important' words. This observation suggests a principle for transparency presentation that limits the amount of information presented. This principle might involve explaining as much variation as possible using the smallest number of explanatory features. We might therefore aim to weigh the overall number of features we present against the information they provide. This is consistent with machine learning approaches to developing models with high dimensional feature sets that aim to identify features with the greatest explanatory power. While we know of technical methods that support this [55], we have not seen such transparency actually tested with users.

A second observation is that users may take issue with detailed transparency and predictions, even when underlying system models are objectively correct. This creates quite a difficult problem for system designers. If it were possible to know which features prompt mistaken beliefs, then these

features could be excluded from transparency. Unfortunately, short of testing user beliefs about all features, this may be very hard to do. We also saw that users don't actively interrogate transparency information to deeply analyze all its implications. Instead, users often look for quick heuristic routes to confirm or discredit simple working theories. We should therefore design intelligent systems in ways that allow users to develop simple working heuristics but also invite them to evolve more accurate mental models when they are motivated to do so.

6 Discussion and Conclusions

Unlike much technically oriented work that aims to develop new transparency algorithms, we explored user perceptions of, and reactions to, transparency. Specifically, we examined users engaging with a real system that interprets self-generated personally relevant emotional data. We deployed a necessarily simple algorithm that users find accurate [78] in order to create simple transparent feedback that conveys the underlying operation of the algorithm. Both exploratory studies indicate that developing and deploying transparent smart systems is complex in practice. In both cases, we observed unexpected user reactions to our attempts to provide detailed information about algorithmic operation. In Study 1 we found that before actual usage, participants initially anticipate greater accuracy for the transparent version of the E-meter but this is altered by their system experiences. After *using* both system versions, users are split 50-50 in their preference, and trust data showed similar ambivalence, even though both versions of the algorithm were perceived as accurate. We identified several possible reasons for this shift. As anticipated, participants who preferred the transparent version valued the increased system understanding that transparency afforded. But two different reasons led other users to prefer the non-transparent version: some attributed more sophisticated abilities to the less transparent system while others felt distracted by transparency. Although many user comments mentioned how incremental word level affective feedback was highly distracting, these comments were not reflected by a reliable overall difference in cognitive load between transparency conditions as assessed by NASA TLX. However more sensitive real-time measures of cognitive load might yield different results.

These initial results raise the question of how we can operationalize transparency in ways that don't distract, while simultaneously allowing users to engage with the system using simple heuristics and facilitating advanced understanding. Study 2 gathered richer contextual data to illuminate exactly how to operationalize transparency to fit the goals from Study 1. To address distraction and improve clarity, we discovered that some users prefer to see only major contributing features to the overall algorithm rating. More detailed transparency information can lead some users to falsely believe the system is operating incorrectly. Furthermore, users often evaluate transparency using simple heuristics rather than deep reflection. Together these results suggest that supporting transparency is complex and there are myriad decisions that affect the user experience

when deciding how to operationalize it. We now discuss future design approaches that build on these observations.

6.1 Meeting the Competing Needs of Transparency Through Progressive Disclosure

Studies 1 and 2 revealed requirements that transparency must meet to be effective for users. This is a challenge because some of these requirements seem internally inconsistent. How can we allow users to develop preliminary working heuristics, while at the same time facilitating detailed understanding for those users who value it? One design solution is suggested by an interaction that took place in Study 2. Recall that in Study 2, some users described their emotional experience using the non-transparent version and only later saw the word-level transparency after clicking a button labeled “How was this rating generated?” After clicking the button to reveal transparency, Participant 11 had further questions. Quite naturally, the participant pointed at the button again and had this exchange with the researcher (R):

P11: *Can I click this? Does this...?*

R: *I don't think it shows any more than that.*

P11: *Damn it.*

R: *Yeah. If you were to click it, what would you expect to see more about?*

P11: *I want bullet points to tell me why it works the way it does.*

Even after seeing the exhaustive transparent feature contributions to the overall rating, this user still had more questions about the system's inner workings. This data point suggests an interaction paradigm that meets the competing needs these two studies have generated: progressive disclosure.

Progressive disclosure has a long history in UI research, dating back to the Xerox Star and early word processing systems [14,61,74]. The original concept involved hiding advanced interface controls; allowing users to make fewer initial errors and learn the system more effectively [14]. In other words, advanced information and explanation is provided on an ‘as needed’ basis, only when the user requests it. Progressive disclosure is also consistent with literature on explanation from the social sciences, which argues that in human-human interaction, explanations are ‘occasioned’, being provided only when the situation demands it [27,30,36,73]. We can apply the principles of progressive disclosure directly to transparency in intelligent systems. For example, similar to Study 2, the E-meter could show a “How was this rating calculated?” button. In this setting, the E-meter might start with only a document-level rating, which reduces distraction and avoids unnecessary complexity. Upon first press of the “How was this rating calculated?” button, the E-meter shows a brief natural language explanation e.g. “This rating was calculated using the positive and negative weighting of the words you have written.” On the next press, the E-meter might show a subset of the high confidence and high impact words. This second press satisfies the users in Study 2 who only wanted to see the major factors influencing the algorithm. However, some users (like those in Study 1) may want yet more transparency. Another press of the

button could reveal further word features that contributed to the overall score. More presses could reveal details about how the training data was collected, or a textual summary of the machine learning model. In this progressive disclosure approach, following UX and human conversational principles, explanation is presented as two-way communication with the user driving exactly when and how explanations are provided. Note too that because transparency is provided ‘on demand’ this removes confusions and inefficiencies arising from spurious, unwanted explanations, and adjusts explanations to the users’ requirements. These design suggestions are consistent with recent work on folk theories in algorithmic systems [18,22] as well as a large body of social science theory; these theories show that people are content to operate with simple (often inaccurate) situational heuristics unless they are deeply invested in a decision or the situation is strikingly anomalous [40,66].

Progressive disclosure of transparency is not limited to predictions from text. For example, other researchers have examined transparency in the context of deep learning models predicting patient outcomes in the medical realm [45]. In this context, patients have a predicted diagnosis risk of a disease and a high number of features (e.g. previous medical diagnoses, number of doctor visits, type of care). Each feature can be visualized and ranked for its contribution to the overall predicted diagnosis risk score. In addition, users can compare outcomes between patients and conduct what-if analyses by modifying patient attributes and seeing how predicted risk changes. Our results suggest that exposing such complexity at once could overwhelm certain users, leading them to reject the tool. To operationalize progressive disclosure in this setting, we might present the predicted risk score by default with a short natural language explanation following natural language explanations in [35]. Should the user request more information, we can incorporate visualizations of features that most contribute to the predicted risk score (e.g. [71] and our transparent version of the E-meter). A request for further information could show similar patients and previous outcomes that have informed the current prediction. Finally, if the user continues requesting increasing amounts of disclosure about a prediction, we can infer that user is invested enough to truly engage with the system and we can present an exploratory explanation (e.g. the explorable predictions in [37]). Such a tool might allow users to modify a patient's features and see how this changes the risk prediction, helping the user to consciously build an accurate mental model of the machine learning predictions. Again, the major contribution of progressive disclosure is avoiding overwhelming users with information, but instead slowly increasing transparency as users indicate a willingness to engage meaningfully with it. We believe this can improve the acceptance of intelligent systems in many realms.

6.2 Impact for Future Transparency Research

Another issue arising from our research is the role of individual differences. It was apparent that different users have varying reactions to, and expectations about intelligent systems. One possibility is that these differences arise from individual

user traits, such as need for control [2] or need for cognition [13]. Future work might examine the relation between such individual traits and reactions to transparency, allowing designers to profile users and deploying personalized system versions.

Our results also have important methodological implications. One method used in prior studies is to provide potential system users with hypothetical scenarios describing system operation and eliciting reactions to those systems. These methods offer ways to collect controlled user data at scale [26,49]. However, results from Study 1 indicate the importance of direct user experience when making system evaluations; users' perceptions of the system were very different following actual usage compared with their projected reactions prior to usage. Care needs to be taken with usage of scenario-based methods.

6.3 Limitations

The current study examines the important algorithmic domain of emotion analytics, but clearly other contexts need to be explored. Furthermore, our deployment of a working algorithm meant that results were obtained for situations where our algorithm generated moderate numbers of errors—future research should compare contexts where there are different levels of errors [46,47]. Additionally, while users generated their own data in our system, results were not directly used to inform other aspects of the user's personal behavior so the costs of system errors were low. While this is appropriate for exploring understanding of initial algorithms with moderate error rates, future work might explore user reactions to transparency in higher risk contexts. While we believe our implications regarding progressive disclosure are generalizable, our work derived these insights from one operationalization of transparency, namely a dynamic visualization of an algorithm. There are many other ways to depict how an algorithm operates including verbal explanations, concrete user exploration, and so forth [42,44,50,71].

6.4 Conclusion

Overall, our data suggest empirically motivated challenges in designing effective UX methods to support transparency for complex algorithms. Our results also reveal potential new research questions regarding user traits, system heuristics and workload. The current study indicates a promising design approach involving progressive disclosure which we intend to explore in future work. It is critical to answer these questions as we continue to deploy intelligent systems with increasing ubiquity and impact.

ACKNOWLEDGMENTS

We would like to thank Victoria Hollis, Ryan Compton, and Lee Taber for their feedback on this project. We would also like to thank the anonymous reviewers for their insightful comments that helped refine this work.

REFERENCES

- [1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, 1–18. <https://doi.org/10.1145/3173574.3174156>
- [2] Icek Ajzen. 1991. The theory of planned behavior. *Organizational Behavior and Human Decision Processes* 50, 2: 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- [3] Julia Angwin and Jeff Larson. 2016. Machine Bias. *ProPublica*. Retrieved October 27, 2017 from <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [4] Stavros Antifakos, Nicky Kern, Bernt Schiele, and Adrian Schwaninger. 2005. Towards improving trust in context-aware systems by displaying system confidence. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services - MobileHCI '05*, 9. <https://doi.org/10.1145/1085777.1085780>
- [5] Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2017. “What is relevant in a text document?”: An interpretable machine learning approach. *PLOS ONE* 12, 8: e0181142. <https://doi.org/10.1371/journal.pone.0181142>
- [6] Victoria Bellotti and Keith Edwards. 2001. Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Human-Computer Interaction* 16, 2: 193–212. https://doi.org/10.1207/S15327051HCI16234_05
- [7] Frank Bentley, Konrad Tollmar, Peter Stephenson, Laura Levy, Brian Jones, Scott Robertson, Ed Price, Richard Catrambone, and Jeff Wilson. 2013. Health Mashups: Presenting Statistical Patterns between Wellbeing Data and Context in Natural Language to Promote Behavior Change. *ACM Transactions on Computer-Human Interaction* 20, 5: 1–27. <https://doi.org/10.1145/2503823>
- [8] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. “It’s Reducing a Human Being to a Percentage”; Perceptions of Justice in Algorithmic Decisions. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*: 1–14. <https://doi.org/10.1145/3173574.3173951>
- [9] Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam T. Kalai. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems*, 4349–4357.
- [10] Engin Bozdag. 2013. Bias in algorithmic filtering and personalization. *Ethics and Information Technology* 15, 3: 209–227. <https://doi.org/10.1007/s10676-013-9321-6>
- [11] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2: 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- [12] Andrea Bunt, Matthew Lount, and Catherine Lauzon. 2012. Are explanations always important?: a study of deployed, low-cost intelligent interactive systems. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces*, 169–178. Retrieved April 25, 2017 from <http://dl.acm.org/citation.cfm?id=2166996>
- [13] John T. Cacioppo, Richard E. Petty, and Chuan Feng Kao. 1984. The Efficient Assessment of Need for Cognition.

- Journal of Personality Assessment* 48, 3: 306–307. https://doi.org/10.1207/s15327752jpa4803_13
- [14] John M. Carroll and Caroline Carrithers. 1984. Training Wheels in a User Interface. *Commun. ACM* 27, 8: 800–806. <https://doi.org/10.1145/358198.358218>
- [15] Eun Kyoung Choe, Bongshin Lee, Haining Zhu, Nathalie Henry Riche, and Dominikus Baur. 2017. Understanding Self-reflection: How People Reflect on Personal Data Through Visual Data Exploration. In *Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '17)*, 173–182. <https://doi.org/10.1145/3154862.3154881>
- [16] Eun Kyoung Choe, Nicole B. Lee, Bongshin Lee, Wanda Pratt, and Julie A. Kientz. 2014. Understanding quantified-selfers’ practices in collecting and exploring personal data. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, 1143–1152. <https://doi.org/10.1145/2556288.2557372>
- [17] Mary L. Cummings. 2004. Automation bias in intelligent time critical decision support systems. In *AIAA 1st Intelligent Systems Technical Conference*, 557–562.
- [18] Michael A. DeVito, Jeremy Birnholtz, Jeffery T. Hancock, Megan French, and Sunny Liu. 2018. How People Form Folk Theories of Social Media Feeds and What it Means for How We Study Self-Presentation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, 1–12. <https://doi.org/10.1145/3173574.3173694>
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]*. Retrieved December 27, 2018 from <http://arxiv.org/abs/1810.04805>
- [20] Mary T. Dzindolet, Scott A. Peterson, Regina A. Pomranky, Linda G. Pierce, and Hall P. Beck. 2003. The Role of Trust in Automation Reliance. *Int. J. Hum.-Comput. Stud.* 58, 6: 697–718. [https://doi.org/10.1016/S1071-5819\(03\)00038-7](https://doi.org/10.1016/S1071-5819(03)00038-7)
- [21] Malin Eiband, Hanna Schneider, Mark Bilandzic, Julian Fazekas-Con, Mareike Haug, and Heinrich Hussmann. 2018. Bringing Transparency Design into Practice. In *23rd International Conference on Intelligent User Interfaces (IUI '18)*, 211–223. <https://doi.org/10.1145/3172944.3172961>
- [22] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I like it, then I hide it: Folk Theories of Social Feeds. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, 2371–2382. Retrieved April 25, 2017 from <http://dl.acm.org/citation.cfm?id=2858494>
- [23] Motahhare Eslami, Sneha R. Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating Algorithmic Process in Online Behavioral Advertising. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, 432:1–432:13. <https://doi.org/10.1145/3173574.3174006>
- [24] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. “I Always Assumed That I Wasn’t Really That Close to [Her]”: Reasoning About Invisible Algorithms in News Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, 153–162. <https://doi.org/10.1145/2702123.2702556>
- [25] Mads Frost, Afsaneh Doryab, Maria Faurholt-Jepsen, Lars Vedel Kessing, and Jakob E. Bardram. 2013. Supporting disease insight through data analysis: refinements of the monarca self-assessment system. 133. <https://doi.org/10.1145/2493432.2493507>
- [26] Pedro García García, Enrico Costanza, Jhim Verame, Diana Nowacka, and Sarvapali D. Ramchurn. 2018. Seeing (Movement) is Believing: The Effect of Motion on Perception of Automatic Systems Performance. *Human-Computer Interaction* 0, 0: 1–51. <https://doi.org/10.1080/07370024.2018.1453815>
- [27] Harold Garfinkel. 1991. *Studies in Ethnomethodology*. Wiley.
- [28] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and Harnessing Adversarial Examples. *arXiv:1412.6572 [cs, stat]*. Retrieved January 7, 2019 from <http://arxiv.org/abs/1412.6572>
- [29] Bryce Goodman and Seth Flaxman. 2017. European Union Regulations on Algorithmic Decision-Making and a “Right to Explanation.” *AI Magazine* 38, 3: 50–57. <https://doi.org/10.1609/aimag.v38i3.2741>
- [30] H. P. Grice. 1975. *Logic and conversation*.
- [31] Enzo Grossi, Nicola Groth, Paola Mosconi, Renata Cerutti, Fabio Pace, Angelo Compare, and Giovanni Apolone. 2006. Development and validation of the short version of the Psychological General Well-Being Index (PGWB-S). *Health and Quality of Life Outcomes*: 8.
- [32] Chloe Gui and Victoria Chan. 2017. Machine learning in medicine. *University of Western Ontario Medical Journal* 86, 2: 76–78. <https://doi.org/10.5206/uwomj.v86i2.2060>
- [33] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of Opportunity in Supervised Learning. *arXiv:1610.02413 [cs]*. Retrieved from <http://arxiv.org/abs/1610.02413>
- [34] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*, Peter A. Hancock and Najmedin Meshkati (eds.), North-Holland, 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- [35] Lisa Anne Hendricks, Ronghang Hu, Trevor Darrell, and Zeynep Akata. Grounding Visual Explanations. 16.
- [36] Denis J. Hilton. 1990. Conversational processes and causal explanation. *Psychological Bulletin* 107, 1: 65–81. <https://doi.org/10.1037/0033-2909.107.1.65>
- [37] Victoria Hollis, Artie Konrad, Aaron Springer, Chris Antoun, Matthew Antoun, Rob Martin, and Steve Whittaker. 2017. What Does All This Data Mean for My Future Mood? Actionable Analytics and Targeted Reflection for Emotional Well-Being. *Human-Computer Interaction*. <https://doi.org/10.1080/07370024.2016.1277724>
- [38] Victoria Hollis, Alon Pekurovsky, Eunika Wu, and Steve Whittaker. 2018. On Being Told How We Feel: How Algorithmic Sensor Feedback Influences Emotion Perception. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 3: 114:1–114:31. <https://doi.org/10.1145/3264924>
- [39] Ellen Isaacs, Artie Konrad, Alan Walendowski, Thomas Lennig, Victoria Hollis, and Steve Whittaker. 2013. Echoes

- from the past: how technology mediated reflection improves well-being. 1071. <https://doi.org/10.1145/2470654.2466137>
- [40] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Macmillan.
- [41] SeungJun Kim, Jaemin Chun, and Anind K. Dey. 2015. Sensors Know When to Interrupt You in the Car: Detecting Driver Interruptibility Through Monitoring of Peripheral Interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, 487–496. <https://doi.org/10.1145/2702123.2702409>
- [42] René F. Kizilcec. 2016. How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface. 2390–2395. <https://doi.org/10.1145/2858036.2858402>
- [43] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces - IUI '15*, 126–137. <https://doi.org/10.1145/2678025.2701399>
- [44] Todd Kulesza, Weng-Keen Wong, Simone Stumpf, Stephen Perona, Rachel White, Margaret M. Burnett, Ian Oberst, and Andrew J. Ko. 2008. Fixing the program my computer learned: barriers for end users, challenges for the machine. In *Proceedings of the 13th international conference on Intelligent user interfaces - IUI '09*, 187. <https://doi.org/10.1145/1502650.1502678>
- [45] B. C. Kwon, M. Choi, J. T. Kim, E. Choi, Y. B. Kim, S. Kwon, J. Sun, and J. Choo. 2018. RetainVis: Visual Analytics with Interpretable and Interactive Recurrent Neural Networks on Electronic Medical Records. *IEEE Transactions on Visualization and Computer Graphics*: 1–1. <https://doi.org/10.1109/TVCG.2018.2865027>
- [46] Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. 2017. Interpretable & Explorable Approximations of Black Box Models. *arXiv:1707.01154 [cs]*. Retrieved September 18, 2018 from <http://arxiv.org/abs/1707.01154>
- [47] Shoushan Li, Lei Huang, Rong Wang, and Guodong Zhou. 2015. Sentence-level Emotion Classification with Label and Context Dependence. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1045–1053. Retrieved September 21, 2018 from <http://www.aclweb.org/anthology/P15-1101>
- [48] Brian Y. Lim and Anind K. Dey. 2009. Assessing demand for intelligibility in context-aware applications. In *Proceedings of the 11th international conference on Ubiquitous computing*, 195–204. Retrieved May 8, 2017 from <http://dl.acm.org/citation.cfm?id=1620576>
- [49] Brian Y. Lim and Anind K. Dey. 2011. Investigating intelligibility for uncertain context-aware applications. In *Proceedings of the 13th international conference on Ubiquitous computing*, 415–424. Retrieved April 25, 2017 from <http://dl.acm.org/citation.cfm?id=2030168>
- [50] Brian Y. Lim, Anind K. Dey, and Daniel Avrahami. 2009. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI '09*, 2119. <https://doi.org/10.1145/1518701.1519023>
- [51] Zachary C. Lipton. 2016. The Mythos of Model Interpretability. *arXiv:1606.03490 [cs, stat]*. Retrieved September 21, 2018 from <http://arxiv.org/abs/1606.03490>
- [52] Yin Lou, Rich Caruana, and Johannes Gehrke. *Intelligible Models for Classification and Regression*.
- [53] Scott M Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. 10.
- [54] Daniel McDuff, Amy Karlson, Ashish Kapoor, Asta Roseway, and Mary Czerwinski. 2012. AffectAura: an intelligent system for emotional memory. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 849–858. Retrieved March 16, 2016 from <http://dl.acm.org/citation.cfm?id=2208525>
- [55] B. Micenkova, R. T. Ng, X. Dang, and I. Assent. 2013. Explaining Outliers by Subspace Separability. In *2013 IEEE 13th International Conference on Data Mining*, 518–527. <https://doi.org/10.1109/ICDM.2013.132>
- [56] Tim Miller. 2017. Explanation in Artificial Intelligence: Insights from the Social Sciences. *arXiv:1706.07269 [cs]*. Retrieved September 17, 2018 from <http://arxiv.org/abs/1706.07269>
- [57] Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. 2018. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing* 73: 1–15. <https://doi.org/10.1016/j.dsp.2017.10.011>
- [58] Kathleen L. Mosier, Linda J. Skitka, Susan Heers, and Mark Burdick. 1998. Automation Bias: Decision Making and Performance in High-Tech Cockpits. *The International Journal of Aviation Psychology* 8, 1: 47–63. https://doi.org/10.1207/s15327108ijap0801_3
- [59] Bonnie M. Muir and Neville Moray. 1996. Trust in automation. Part II. Experimental studies of trust and human intervention in a process control simulation. *Ergonomics* 39, 3: 429–460. <https://doi.org/10.1080/00140139608964474>
- [60] Saurabh Nagrecha, John Z. Dillon, and Nitesh V. Chawla. 2017. MOOC Dropout Prediction: Lessons Learned from Making Pipelines Interpretable. In *Proceedings of the 26th International Conference on World Wide Web Companion - WWW '17 Companion*, 351–359. <https://doi.org/10.1145/3041021.3054162>
- [61] Lloyd H. Nakatani and John A. Rohrlich. 1983. Soft machines: A philosophy of user-computer interface design. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems - CHI '83*, 19–23. <https://doi.org/10.1145/800045.801573>
- [62] Clifford Nass and Youngme Moon. 2000. Machines and Mindlessness: Social Responses to Computers. *Journal of Social Issues* 56, 1: 81–103. <https://doi.org/10.1111/0022-4537.00153>
- [63] Cathy O’Neil. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown.
- [64] Eli Pariser. 2011. *The Filter Bubble: What The Internet Is Hiding From You*. Penguin Books Limited.
- [65] James W. Pennebaker. 2011. *The Secret Life of Pronouns: What Our Words Say About Us*. Bloomsbury Press.
- [66] Richard E. Petty and John T. Cacioppo. 1986. The Elaboration Likelihood Model of Persuasion. *Advances in*

- Experimental Social Psychology* 19: 123–205. [https://doi.org/10.1016/S0065-2601\(08\)60214-2](https://doi.org/10.1016/S0065-2601(08)60214-2)
- [67] M. F. Porter. 1980. An algorithm for suffix stripping. *Program* 14, 3: 130–137. <https://doi.org/10.1108/eb046814>
- [68] Mashfiqui Rabbi, Min Hane Aung, Mi Zhang, and Tanzeem Choudhury. 2015. MyBehavior: automatic personalized health feedback from user behaviors and preferences using smartphones. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 707–718. Retrieved May 4, 2017 from <http://dl.acm.org/citation.cfm?id=2805840>
- [69] Lena Reed, Jiaqi Wu, Shereen Oraby, Pranav Anand, and Marilyn Walker. 2017. Learning Lexico-Functional Patterns for First-Person Affect. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 141–147. Retrieved September 18, 2018 from <http://aclweb.org/anthology/P17-2022>
- [70] Byron Reeves and Clifford Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places*. Cambridge University Press.
- [71] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [72] Stephanie Rude, Eva-Maria Gortner, and James Pennebaker. 2004. Language use of depressed and depression-vulnerable college students. *Cognition and Emotion* 18, 8: 1121–1133. <https://doi.org/10.1080/02699930441000030>
- [73] Emanuel A. Schegloff. 1992. Repair After Next Turn: The Last Structurally Provided Defense of Intersubjectivity in Conversation. *American Journal of Sociology* 97, 5: 1295–1345.
- [74] David Canfield Smith. Designing the Star User Interface. 21.
- [75] Aaron Springer and Henriette Cramer. 2018. “Play PRBLMS”: Identifying and Correcting Less Accessible Content in Voice Interfaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, 296:1–296:13. <https://doi.org/10.1145/3173574.3173870>
- [76] Aaron Springer, Victoria Hollis, and Steve Whittaker. 2017. Dice in the Black Box: User Experiences with an Inscrutable Algorithm. Retrieved April 24, 2017 from <https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15372>
- [77] Aaron Springer, Victoria Hollis, and Steve Whittaker. 2018. Mood modeling: accuracy depends on active logging and reflection. *Personal and Ubiquitous Computing*: 1–15. <https://doi.org/10.1007/s00779-018-1123-8>
- [78] Aaron Springer and Steve Whittaker. 2018. What are You Hiding? Algorithmic Transparency and User Perceptions. In *2018 AAAI Spring Symposium Series*.
- [79] Rachael Tatman. 2017. Gender and Dialect Bias in YouTube’s Automatic Captions. *EACL 2017*: 53.
- [80] Kristen Vaccaro, Dylan Huang, Motahhare Eslami, Christian Sandvig, Kevin Hamilton, and Karrie Karahalios. 2018. The Illusion of Control: Placebo Effects of Control Settings. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, 16:1–16:13. <https://doi.org/10.1145/3173574.3173590>
- [81] Jeffrey Warshaw, Tara Matthews, Steve Whittaker, Chris Kau, Mateo Bengualid, and Barton A. Smith. 2015. Can an Algorithm Know the “Real You”? Understanding People’s Reactions to Hyper-personal Analytics Systems. 797–806. <https://doi.org/10.1145/2702123.2702274>
- [82] Daniel S. Weld and Gagan Bansal. 2018. The Challenge of Crafting Intelligible Intelligence. *arXiv:1803.04263 [cs]*. Retrieved September 20, 2018 from <http://arxiv.org/abs/1803.04263>
- [83] Jenna Wiens and Erica S. Shenoy. 2018. Machine Learning for Healthcare: On the Verge of a Major Shift in Healthcare Epidemiology. *Clinical Infectious Diseases* 66, 1: 149–153. <https://doi.org/10.1093/cid/cix731>
- [84] Rayoung Yang and Mark W. Newman. 2013. Learning from a learning thermostat: lessons for intelligent systems for the home. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing - UbiComp '13*, 93. <https://doi.org/10.1145/2493432.2493489>
- [85] Miriam Zisook, Sara Taylor, Akane Sano, and Rosalind Picard. 2016. SNAPSHOT Expose: Stage Based and Social Theory Based Applications to Reduce Stress and Improve Wellbeing. Retrieved March 10, 2017 from <https://pdfs.semanticscholar.org/b283/53899d0c5059a31c9bc69c364e62bc6c7ff5.pdf>
- [86] Home | Pip. Retrieved September 21, 2018 from <https://thepip.com/en-us/>
- [87] Real-time physiological signals | E4 EDA/GSR sensor. Empatica. Retrieved September 21, 2018 from <https://www.empatica.com/research/e4>
- [88] Woebot - Your charming robot friend who is here for you, 24/7. Retrieved September 21, 2018 from <https://woebot.io>