

Affective Sensors, Privacy, and Ethical Contracts

Carson Reynolds

MIT Media Lab E15-120f
20 Ames St., Cambridge, MA 02139
carsonr@media.mit.edu

Rosalind W. Picard

MIT Media Lab E15-020g
20 Ames St., Cambridge, MA 02139
picard@media.mit.edu

ABSTRACT

Sensing affect raises critical privacy concerns, which are examined here using ethical theory, and with a study that illuminates the connection between ethical theory and privacy. We take the perspective that affect sensing systems encode a designer's ethical and moral decisions: which emotions will be recognized, who can access recognition results, and what use is made of recognized emotions. Previous work on privacy has argued that users want feedback and control over such ethical choices. In response, we develop ethical contracts from the theory of contractualism, which grounds moral decisions on mutual agreement. Current findings indicate that users report significantly more respect for privacy in systems with an ethical contract when compared to a control.

Author Keywords

Affective Computing, Sensors, Privacy, Ethics, Emotion Recognition, Contractualism.

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation]: User Interfaces — theory and methods; K.4.1 [Public Policy Issues] — ethics, privacy.

ETHICS AND AFFECT

Would it be ethical for a computer to sense a user's emotions? If a perceptual user interface has the capability to detect emotions, would this be an invasion of privacy? Are users comfortable with having their emotions sensed?

Within these weighty and complex questions are critical issues for designers of interfaces that sense affect. Unfortunately, only modest amounts of information exist to help guide designers encountering these thorny issues.

Computers, as they are currently designed, do not have the

capability to be ethical on their own. Lacking free will, machines cannot make moral choices between “good” and “bad.” Instead, they largely carry out their designer's choices. This means that if a designer makes “bad” choices from the user's perspective, the resulting interaction could be viewed as unethical.

Right now affective sensing systems are being designed and used in ways that raise important ethical concerns [1]. For instance, “Integrated System for Emotional State Recognition for the Enhancement of Human Performance and Detection of Criminal Intent,” is the subject of a recent DARPA SBIR [2]. This initiative emphasizes technologies that can be used without the consent or knowledge of users. Such developments introduce ethical concerns for those whose information is being collected.

The examination of ethical concerns of affective sensors is the purpose of this paper. The approach taken is the combination of contractual ethical theory with studies where scenarios with affective sensors are evaluated from the user's standpoint. This approach is significant because our literature review revealed no studies of affective computing systems evaluated using ethical theory.

To provide a context for this work, we begin with a brief discussion of varieties of affect sensors. What follows is a development of a contractualist ethics suitable for human-computer interaction. The results of evaluations we performed investigating the use of affect sensors with and without ethical contracts are then presented.

AFFECT SENSORS

An affect sensor is a device that takes an input signal and processes it for some evidence of emotions. There are many techniques and modalities used to detect affect: physiological signals, facial expression recognition, speech prosody recognition, and pressure sensors have all been developed [3].

Affect sensors have shown the capability to detect emotions (albeit with less-than-perfect accuracy). This new capability however means that we need to perform new assessments about how users feel about these sorts of technologies. Do users find systems that detect their emotional state ethically acceptable?

COMPUTER ETHICS

To date, the emphasis in affective technologies has been on building systems that work – that recognize, express, help communicate, and respond to human emotion. An important growing concern, however, is how users feel about such technology – whether it feels respectful of their privacy and other needs, and on what basis it is acceptable or not. Thus, the emphasis here is not on what *can* be done, but rather on helping illuminate what users think *should* be done.

Morr, in the classic article “What is Computer Ethics?” [4] conceptualizes computer ethics as dealing with the policy vacuums and conceptual muddles raised by information technology. The introduction of affect sensors has created a policy vacuum. Although several descriptions of potentially harmful or unethical uses of affective technology have been broached [1], the ethical consequences of sensing user emotion are unstudied, and methods for dealing with them in a manner users and designers see as ethically acceptable are absent. This absence of policy has led us to the development of a theory of ethics relevant to interaction design.

DESIGN CONTRACTUALISM

We theorize that interaction technologies represent an implicit or provide an explicit contract between the designer and user. The designer makes a variety of moral and ethical decisions in the development of an interaction technology. These decisions are encoded in the technology, and experienced by the user. When users encounter a new technology, they are at the outset in an initial bargaining position. Users can choose to use or not use a technology based on their evaluation of the encoded ethical stance.

When the contract is implicit, users may evaluate an interaction technology for a period of time before accepting or rejecting its use. Their comfort with the technology could be taken as an indicator for how the designer’s ethical decisions match their own expectations of what is acceptable.

When the contract is explicit, in terms of a privacy policy or social contract, users can assess the designer’s ethical decisions more immediately. For instance, a user of an open source technology may accept software adhering to a particular contract, but reject closed-source alternatives, sometimes without using either.

This theory is an extension of contractualism, which has been used in both moral and political philosophy. As a foundation for our ethical analysis, we take the position the computer ethics should be approached in a weakly conservative manner. Namely, we accept that computers could give rise to entirely new varieties of ethics, but for present purposes we will treat computer ethics as a microethics, akin to medical ethics or environmental ethics. We will thus apply contractualism as a macroethics to serve as the foundation of an ethical theory for affect sensing.

Contractualism finds ethical evaluations on a hypothetical or real contract formed between groups or individuals. Cudd describes the contractual macroethical position in the following manner: “Contractualism, which stems from the Kantian line of social contract thought, holds rationality requires that we respect persons, which in turn requires that moral principles be such that they can be justified to each person” [5]. Thus, we should offer our moral decisions in public and seek to justify them to each user.

Encoded Ethical Decisions

Our theory asserts that ethical decisions are encoded by interaction technology. A study by Friedman provides some support for this idea [6]. In interviews with computer-science students, Friedman found that 75% attributed “decision-making” to computers. But only 21% held the computer “morally responsible” for errors. These results indicate that the majority of the interviewees thought a computer could make decisions but a minority blamed the computer for the consequences of bad actions. One participant was quoted as saying “the decisions that the computer makes are decisions that somebody else made before and programmed into the computer . . .” Friedman concludes by noting that “designers should communicate through a system that a (human) who and not a computer (what) – is responsible for the consequences of computer use.” We think designers should communicate and disclose the moral decisions made during the design of interaction technology and embedded into the technology.

When designing interaction with affective sensors, we make a series of decisions that could affect ethical acceptability. These include decisions related to privacy like: who will collect emotional data, what type of emotions are recognized, and for what task the recognized emotional data is used. If we design a system that detects fear by using electrodes and video cameras, and make this data available to anyone who requests it, then we speculate users will feel that their privacy is invaded.

Examples of Design Contractualism

We see the proliferation of privacy policies and technologies related to them as evidence that explicit design contractualism is already occurring. Many websites offer users privacy policies which explain how the private data collected will be used. The Platform for Privacy Preferences (P3P) provides an internet protocol to help standardize and automate the process of accepting or rejecting a particular website’s privacy policy [7]. Some open source communities also offer explicit Social Contracts describing “a set of commitments” to which designers will abide. The Debian Linux social contract describes the ethics and commitments which the designers will commit themselves to in the design of the Linux distribution [8].

EVALUATION

To better understand users' perceptions of affect sensors in light of design contractualism, we conducted an evaluation that examines privacy effects of affect sensors. These studies presented users with a variety of scenarios involving particular emotions, various sensors, and contracts between the designer and user. We chose to use web-based questionnaires as a method of inquiry. As a research method, questionnaires have many shortcomings, but also some strengths. Questionnaires allow for exposure to a variety of hypothetical situations, while a subject is (presumably) comfortably seated in his or her natural computing environment. They also allow us to quickly expose a large number of people to a set of ideas, while avoiding scheduling conflicts, no-shows, biases due to experimenter interaction, and other attitude-altering phenomena associated with in-person interactive experiences. At the same time, we do lose control over whether or not a subject fills out more than one questionnaire (although the task is boring, and there is no compensation, so we do not expect this to happen.) We also recognize that the way people answer how they think they would behave can differ from the way they will actually behave. Nonetheless, because a contract is a consciously engaged in endeavor, and our initial interest is to assess what people think about affect sensing and the effect of contract use, questionnaires are an appropriate tool.

Hypotheses

In the broadest sense, we wanted to understand more about the ethical acceptability of affect sensors. We considered respect for privacy, willingness to use, and comfort as factors that might indicate ethical acceptability. We speculated that subjects will report a greater sense of privacy invasiveness if they think affect sensing is unethical. We hypothesized that the introduction of an ethical contract would improve ethical acceptability.

Pilot Studies

Two pilot studies were performed as we developed our understanding of the relationship between ethical acceptability and contractualism. The first was a paper questionnaire distributed to visitors to our laboratory. Although the results supported our hypotheses, the participants were solicited in a non-random manner and much of the data was incomplete so the experiment was redesigned. The second was an internet questionnaire that corrected these shortcomings, but did not balance across conditions appropriately. We arrived at the experimental design for the survey described below by revising the design of these pilot studies.

Main study

Methodology

Our evaluation took the form of a 2x4 design. Participants were first randomly assigned to treatment (Ethical Contract)

or control (No Contract) conditions. Demographic information was collected initially. Then, each subject was presented with four scenarios. Participants in the treatment condition also received a contract with each scenario. Participants in the control condition were never presented an ethical contract. Each scenario involved one of four emotions and one of two application contexts. The emotions were Joy, Anger, Sadness, and Excitement. The application contexts were music and news recommendation systems. The design was balanced: each subject received the same emotions and contexts, but in a random order.

For example, a scenario for the emotion joy in the context of music recommendations was:

"Imagine yourself using a music listening program at your home. The current song that is playing evokes excitement. The computer detects your excitement and sends information to a recommendation website. This website is a resource used by music listeners to categorize and recommend music. It provides you with music recommendations associated with particular emotions, together with other features based on preferences you set. For example, it might be useful to search this website for music that evoked excitement in listeners. A combination of camera and pressure-sensitive mouse are used to assess your excitement. Your excitement is assessed by recognizing facial expressions from video and muscle tension from pressure sensors on the mouse."

If the participant was assigned to the Contract condition, they additionally saw the following text:

"The recommendation website has a contract that describes who will have access to your emotional data, what this data will be used for, and exactly how emotions like excitement are assessed: The excitement recorded by the system will only be used to recommend content to members of this website. This data is anonymous and can only be accessed by users of the website. Your excitement is assessed by recognizing facial expressions from video and muscle tension from pressure sensors on the mouse. You have the opportunity to interact with the system with the contract for a month and find that it behaves just as the contract dictates."

Participants in both conditions were then asked to report if they would use the system, how comfortable they would feel, if their privacy was affected, and how certain they were of these responses. The questions were presented on a seven-point Likert scale.

Results

There were 125 total responses to the internet questionnaire. Of these, 64 had been filled in completely without error. Participants were solicited with emails to departmental emails lists and postings on community websites. Participants were told the survey would take approximately 15 minutes and that they would be asked "a variety of questions about scenarios motivated by recent

research developments.” There were 30 Female and 34 Male participants. The average age of the participants was 30. The participants listed 12 countries as their nationalities, with the majority (71%) from the United States. In terms of education, 56% reported post-graduate education, 39% undergraduate, and 5% secondary level. The participants were randomly assigned to a condition, with 33 assigned to the control condition and 31 assigned to the effect condition.

ANOVAs were performed to compare the response variables with experimental treatment. These showed a difference in “respect for privacy” reported by participants in the contract condition. ($F=14.57$, $df=1$, $p=0.0002$). Because the data was reported as 7-point Likert scores, the normality assumption required for ANOVA does not hold. Consequently, the Kruskal-Wallis non-parametric test was also applied. This test also reported a difference in “respect for privacy” ($\chi^2=12.37$, $df=1$, $p\text{-value}=0.0004$). With a difference in means 0.73, Cohen’s d was 0.48, leaving an effect size r of 0.23. A trend towards an improvement in comfort for the contract condition was also observed, but the Kruskal-Wallis test was not significant ($\chi^2=2.98$, $df=1$, $p=0.08$). No significant difference was observed in questions dealing with willingness to use and certainty.

Discussion

Although this experiment only considered two contexts: music recommender systems and news recommender systems, the impact of a contract was found to be significant in both. Regardless of which emotion was detected, participants who had a contract felt their privacy more respected. Participants were asked: “Do you think your privacy would be affected by using a system like the one above?” Choosing among 1 “completely invaded” to 7 “completely respected,” participants without a contract reported a mean of 2.6, towards the “invaded” side of the scale, while participants with a contract reported a mean of 3.4, which is more neutral, although still far from “completely respected”. Thus, a designer of a system involving affect sensing may find that subjects are biased toward feeling a violation of privacy when using that system. Our findings indicate that this bias may be alleviated for many subjects if the designer includes a contract. Without ethical contracts, participants report that their privacy is invaded, and with contracts, report an increase in respect.

CONCLUSION

Although this is the first study to look at ethical and privacy implications of affect sensors, the findings were significant. Specifically, our participants displayed a significant preference for ethical contracts. The results recommend that

designers seeking to improve respect for privacy should consider including an ethical contract.

This study represents a tentative first step in a new style of evaluation motivated by ethics. An important next move is to investigate how our results differ when people use real affect sensors, instead of simply being asked about them in the abstract. New evaluations should consider what users describe as indicators of ethical interaction, instead of focusing on privacy invasion. There are likely to be a number of different factors bearing on users’ impressions of what is ethical, and these need more exploration.

We’ve just considered contractualism as a theoretical basis for an ethical theory applicable to affect sensing. This mode of inquiry has given us interesting insights into privacy boundaries. New examinations of different approaches to ethics, such as utilitarianism, consequentialism, or deontologism suggest novel styles of ethical interface evaluation.

ACKNOWLEDGMENTS

This work was supported in part by the MIT Media Lab Things That Think consortium.

REFERENCES

1. Picard, R.W. and Klein, J. (2002). Computers that Recognise and Respond to User Emotion: Theoretical and Practical Implications. *Interacting with Computers*, 14(2) (2002), 141-169.
2. DARPA SB032-038 TITLE: Integrated System for Emotional State Recognition for the Enhancement of Human Performance and Detection of Criminal Intent. <http://www.dodsbir.net/solicitation/darpa032.htm>
3. Reynolds, C. (2001) The Sensing and Measurement of Frustration with Computers. Master’s thesis, MIT.
4. Moor, J.H. (1985). What is computer ethics? *Metaphilosophy*, 28(3) 266-275.
5. Cudd, A. (2000). Contractarianism. <http://plato.stanford.edu/entries/contractarianism>
6. Friedman, B. (1995). It’s the computer’s fault: reasoning about computers as moral agents. *CHI ‘95 Conference Companion 1995*: 226-227.
7. Cranor, L. and Reagle, J. (1998). Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project, in J.K. MacKie-Mason and D. Waterman (eds.) *Telephony, the Internet, and the Media*. Lawrence Erlbaum Associates, Mahwah, NJ, 1998.
8. Debian Social Contract http://www.debian.org/social_contract