

Crowdsourcing Workshop: The Emergence of Affective Crowdsourcing

Robert Morris
MIT Media Lab
75 Amherst St
Cambridge, MA 02139
rmorris@media.mit.edu

ABSTRACT

Affective computing is a multidisciplinary field that integrates theories, methods, and technologies from a variety of different disciplines, including affective science, machine learning, signal processing, and philosophy. This paper argues that affective computing would be well-served to embrace yet another discipline – crowdsourcing. Likewise, this paper argues that crowdsourcing, itself a multidisciplinary field, can be well-served by incorporating theories and methods from affective computing. The commingling of affective computing and crowdsourcing creates exciting new research opportunities, several examples of which are outlined in this paper.

Author Keywords

Affective computing, crowdsourcing, human computation.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Affective computing is a form of computing that “relates to, arises from, or deliberately influences human emotion or other affective phenomena.[1]” Put another way, affective computing explores emotionally intelligent technologies – namely, technologies that recognize emotions, express emotions, influence emotions, and respond intelligently to the emotions of others. The goals of affective computing research, by virtue of their scope, necessitate a multidisciplinary approach and indeed affective computing incorporates theories from many different disciplines, including signal processing, computer vision, and psychophysiology, to name just a few.

Crowdsourcing [2], a field that aims to cultivate and coordinate the collective efforts and wisdom of large groups of people, is also vastly multidisciplinary. The taxonomy of

crowdsourcing is still being debated [3], and researchers often use some terms interchangeably (human computation is often used in place of crowdsourcing, for example). For the purposes of this paper, human computation work that involves large groups of people will be considered a form of crowdsourcing. Research in crowdsourcing, defined thusly, draws insights from fields such as machine learning, game studies, and management science, among others [4-6]. Crowdsourcing is partly motivated by the assumption that, under the right conditions, a diverse group of individuals can wield a collective intelligence that is greater than the intelligence of any one of its individual constituents [7]. As such, the field is, in a sense, theoretically poised to embrace other disciplines, so that it too can benefit from a crowd of diverse perspectives.

Unfortunately, despite their multidisciplinary profiles, the fields of affective computing and crowdsourcing have yet to fully partner with each other. To date, there have only been a handful of projects that lie at the intersection of affective computing and crowdsourcing (for recent examples, see [8], [9], and [10]). This paper argues that these projects, though few in number, portend a powerful confluence between affective computing and crowdsourcing. The commingling of these disciplines – something I will call *affective crowdsourcing*, for lack of a more inventive name - will create many exciting new areas of research. Ideally, this interdisciplinary approach will benefit both affective computing and crowdsourcing; ideas and methods from one discipline will enrich the other, and vice versa. This paper describes the benefits of this interdisciplinary approach and outlines future research avenues that combine affective computing with crowdsourcing.

AFFECTIVE COMPUTING APPLICATIONS

Many facets of crowdsourcing are applicable to affective computing research. By leveraging tools and techniques from the field of crowdsourcing, affective computing researchers can explore many new avenues of research. While it is beyond to the scope of this article to outline all these new research possibilities, this paper will describe several specific projects that can emerge when affective computing gets crowdsourced.

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

Emotion-elicitation Corpus

To evaluate affective computing technologies, researchers often need to artificially evoke different emotions in research study participants. Many researchers use film clips to elicit emotions in the lab, and there are several well-known clip databases that are used in affective computing research, and indeed throughout all of affective science [11]. Unfortunately, many of these film databases are outdated, and some of the most widely used sources are over ten years old. Sensitivity to affectively charged media has no doubt changed over the past decade, and film clips that were once highly emotional in the mid-90s may now seem sterile.

A new corpus of reliable, emotion-eliciting films is sorely needed for investigators in affective computing, and this corpus may need to be updated every few years. New crowdsourcing tools such as Amazon's Mechanical Turk (MTurk) offer new ways to compile and evaluate a large corpus of emotion-eliciting films. While the demographics of MTurk are fluid and ever-changing [12], there is arguably far more variety in MTurk workers than there is in the typical undergraduate research cohort that usually evaluates emotional stimuli. Annotations from a diverse group of MTurk workers could begin to assess how different videos affect different groups (for a longer discussion on the advantages of crowdsourcing diversity in this context, see [8]). Further, in addition to simply labeling affective content in multimedia, crowds could also be incentivized to collect this content. MTurk workers, for example, could be tasked to not only rate affective content, but to also supply it.

New Data Sets for Affect Recognition Technologies

To detect human emotion in real-time, computers need sophisticated sensing mechanisms and advanced pattern recognition algorithms. To be robust, the algorithms that drive these systems need to be trained on immense sets of affective data. Traditionally, these data sets have come from actors or research subjects, and have been collected in controlled, artificial settings. Unfortunately, it is hard to collect vast amounts of data in this manner, and the data collected is not always ecologically valid.

Crowdsourced communities and workforces, when harnessed appropriately, can generate massive amounts of affective data, and these data can be collected from real-world settings (for example, while users are sitting in front of a computer, at home, or at work, where there might be extraneous noise or unpredictable lighting conditions). New advances in computer vision have paved the way for massive data collection from distributed crowds online, even under difficult settings. Several behavioral and physiological correlates of emotion (e.g., facial expression [13], and even heart-rate [14]) can now be monitored unobtrusively through an ordinary computer webcam. To collect this data, crowds could be incentivized to view emotion-eliciting media while a webcam records video and

audio streams. This data, in turn, could be used as a help train the next generation of real-time affect sensing technologies. It is important to note that this data should never be collected without a person's awareness and an opt-in system should be applied for these types of data collection projects.

Crowdsourcing Subjective Well-Being

The ability to regulate emotions (and, in particular, to up-regulate positive emotions and down-regulate negative ones) is a crucial component of subjective well-being and overall physical and mental health [15]. And while philosophy, religion, and science have long been concerned with this topic, there is still much to learn about how humans regulate emotions and how technologies might provide a supporting role.

There are many different ways to regulate emotions, but considerable research attention has been given to cognitive reappraisal – a technique that involves reframing thoughts or situations so as to decrease their emotional impact [16]. Habitual use of reappraisal is linked to high levels of positive affect, interpersonal functioning, self-esteem, life satisfaction, and well-being [17]. The basic concept of cognitive reappraisal also forms the backbone for many well-validated forms of psychotherapy, most notably cognitive therapy [18]. It is important to avoid confusing cognitive reappraisal with Pollyannaism, or naïve optimism. Instead, adaptive reappraisal techniques, as defined in the psychological literature, strive to put a *rational*, yet positive, frame around situations that might otherwise be viewed as emotionally unpleasant.

Unfortunately, not everyone is blessed with effective reappraisal skills, and there is a great need for technologies that help people practice reframing techniques. For instance, an ideal new technology might be a mobile device that helps people reframe emotion-eliciting thoughts and situations. Such a tool could not only be a useful training aid, but could also provide near real-time therapeutic support. For individuals with clinical features, this tool could be a powerful adjunct to cognitive-behavioral therapy and/or dialectical-behavior therapy.

Advances in human-based computation, combined with the rise of online, crowdsourced workforces (e.g., Amazon's Mechanical Turk), suggest that this type of tool is now within our grasp. Mechanical Turk workers, for example, could be recruited to provide feedback for people in need of emotion regulatory assistance. Human-based computation algorithms could distribute these tasks amongst many different workers, and across many different parallel and iterative stages. New algorithms and verification strategies could be used to ensure that feedback is impartial, rational, and tactful. Research in this direction builds on other near real-time crowdsourced interface technologies, such as *VizWiz* [19], but this work extends

these methodologies to the domain of emotion regulation and computer-mediated psychotherapy.

CROWDSOURCING APPLICATIONS

The commingling of affective computing and crowdsourcing will convey benefits for both disciplines, not just affective computing. This section describes several ways to use affective computing methods to support the field of crowdsourcing.

Mood Induction

Under the right conditions, crowdsourcing can generate immensely creative output. Threadless, Linux, and IdeaStorm are just a few examples of the creative force that can come from massively distributed crowds. However, cultivating crowdsourced creativity is a delicate art and it largely depends on creating the right incentives. Reputation, fun, commitment to a community, and allegiance to a brand are all forms of incentives that can spur creativity in the crowd. Unfortunately, incentives like these are virtually nonexistent in anonymous, online labor markets like MTurk. Yet, there are still several ways to enhance creative output amongst MTurk users. For instance, research in psychology and creativity has long shown that positive emotions spur creative thinking [20]. Affective computing technologies can be used to elicit and monitor positive emotions in workers before and during MTurk tasks. Workers in positive moods should be better positioned to think divergently and may contribute more creative work on crowdsourced tasks.

Mood-Based Task Allocation

One implication of the broaden-and-build hypothesis is that individuals in a positive mood may express greater divergent thinking whereas individuals in negative moods may show greater abilities at narrow, circumscribed tasks such as rote search and retrieval [21]. Affective computing technologies can assess an online worker's mood and allocate him/her to tasks that are appropriate to his/her mood. In other words, instead of directly eliciting different mood states, affective technologies could assess a worker's baseline mood and allocate tasks accordingly. In addition, affective computing technologies could also monitor arousal levels in crowds, and could, for example, reallocate work if a worker appears too bored or tired.

Affective Evaluation of Crowdsourced Tasks and Tools

Crowdsourced tasks that are too stressful or too boring can produce dismal results. In many cases, researchers or employers might want to know whether certain tasks are creating undue levels of boredom or stress for crowdsourced workers. Affective computing technologies can be used to measure these affective states in real-time and provide feedback to those coordinating the crowd.

Cognitive-Affective Adjudication

Many projects on MTurk and other sites use voting systems to sift through the vast contributions generated by the crowd. To date, these methods of adjudication have been largely based on cognitive measures (such as self-report). But, cognitive measures are often prone to bias and may not accurately reflect the true feelings of the respondent. In the future, these cognitive ratings could be combined with data from multimodal affective sensing technologies. Affective cues from the face or from sympathetic nervous system activity can provide a layer of data that is largely impervious to the biases that plague purely cognitive-based self-report measures. By combining cognitive and affective measures, crowdsourced methods of adjudication can be made more reliable.

BIOGRAPHY

I am a PhD student at the MIT Media Lab. I work in the Affective Computing lab, under the direction of Prof. Rosalind Picard, and I am currently pursuing new research at the intersection of affective computing and crowdsourcing. I am also actively researching topics related to affective computing, computer-mediated psychotherapy, and electronic musical interfaces. Before coming to MIT, I worked at Harvard and MGH, where I conducted fMRI research on deception and addiction disorders, respectively. I have an M.A. in Media Arts and Sciences from MIT and a B.A. in Psychology from Princeton University. I am also a performing and recording musician and I am currently collaborating with Gibson and Avid on the development of new gestural and haptic interfaces for the electric guitar.

EXPERIENCE WITH CROWDSOURCING

I have numerous crowdsourcing projects that are in progress. As I type this, MTurk workers are completing several tasks that combine affective computing techniques with crowdsourcing methods. For instance, I am monitoring the emotional state of MTurk workers and am looking for correlations between their self-reported mood and performance on creative tasks. I am also actively administering reframing tasks to MTurk workers and I'm testing various ways to coordinate this form of human computation. I am also testing the ability of MTurk workers to identify cognitive distortions, such as overgeneralization, catastrophization, and all-or-nothing thinking. In addition, I am collecting data from an MTurk task that requires workers to record brief emotional statements into a computer microphone. For example, some workers are asked to "describe a time when [they] felt really connected with other people." This task may have implications for future emotion-elicitation experiments.

Finally, I am also using MTurk to create crowdsourced audio files for musical compositions. This work is similar to Aaron Koblin's "Bicycle Built For Two-Thousand" project, except that it uses crowdsourced audio to augment existing compositions, rather than to create them from the bottom-

up. This work will be used in a musical composition that will be released in the summer of 2011.

REFERENCES

- [1] R. W. Picard, *Affective Computing*, 1st ed. The MIT Press, 2000.
- [2] J. Howe, *Crowdsourcing: Why the Power of the Crowd Is Driving the Future of Business*, 1st ed. Crown Business, 2008.
- [3] A. Rouse, "A Preliminary Taxonomy of Crowdsourcing," *ACIS 2010 Proceedings*, Jan. 2010.
- [4] V. C. Raykar et al., "Supervised learning from multiple experts: whom to trust when everyone lies a bit," in *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 889–896, 2009.
- [5] L. von Ahn and L. Dabbish, "Designing games with a purpose," *Communications of the ACM*, vol. 51, no. 8, p. 57, 2008.
- [6] T. W. Malone, R. Laubacher, and C. Dellarocas, "Harnessing Crowds: Mapping the Genome of Collective Intelligence," *SSRN eLibrary*, Feb. 2009.
- [7] J. Surowiecki, *The Wisdom of Crowds*. Anchor, 2005.
- [8] M. Soleymani and M. Larson, "Crowdsourcing for Affective Annotation of Video: Development of a View-reported Boredom Corpus," in *SIGIR 2010 Workshop on Crowdsourcing for Search Evaluation*, 2010.
- [9] S. D. Kamvar and J. Harris, "We feel fine and searching the emotional web," in *Proceedings of the fourth ACM international conference on Web search and data mining*, pp. 117–126, 2011.
- [10] B. G. Morton, J. A. Speck, E. M. Schmidt, and Y. E. Kim, "Improving music emotion labeling using human computation," in *Proceedings of the ACM SIGKDD Workshop on Human Computation*, pp. 45–48, 2010.
- [11] J. Rottenberg, R. Ray, and J. Gross, "Emotion Elicitation Using Films," in *Handbook of emotion elicitation and assessment*, Oxford University Press US, 2007.
- [12] J. Ross, L. Irani, M. S. Silberman, A. Zaldivar, and B. Tomlinson, "Who are the crowdworkers?: shifting demographics in mechanical turk," in *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, pp. 2863–2872, 2010.
- [13] P. Michel and R. El Kaliouby, "Real time facial expression recognition in video using support vector machines," in *Proceedings of the 5th international conference on Multimodal interfaces*, pp. 258–264, 2003.
- [14] M. Poh, D. J. McDuff, and R. W. Picard, "Non-contact, automated cardiac pulse measurements using video imaging and blind source separation," *Optics Express*, vol. 18, no. 10, pp. 10762–10774, May. 2010.
- [15] J. J. Gross, "The emerging field of emotion regulation: An integrative review," *Review of General Psychology*, vol. 2, no. 3, pp. 271–299, 1998.
- [16] K. N. Ochsner and J. J. Gross, "The cognitive control of emotion," *Trends in Cognitive Sciences*, vol. 9, no. 5, pp. 242–249, May. 2005.
- [17] J. J. Gross and O. P. John, "Individual differences in two emotion regulation processes: implications for affect, relationships, and well-being," *Journal of Personality and Social Psychology*, vol. 85, no. 2, pp. 348–362, Aug. 2003.
- [18] Campbell-Sills, L. & Barlow, D.H., "Incorporating emotion regulation into conceptualizations and treatments of anxiety and mood disorders." In J.J. Gross Ed., *Handbook of Emotion Regulation*, pp. 542–560, 2007.
- [19] J. P. Bigham et al., "VizWiz," in *Proceedings of the 23rd annual ACM symposium on User interface software and technology - UIST '10*, p. 333, 2010.
- [20] A. M. Isen, M. M. Johnson, E. Mertz, and G. F. Robinson, "The influence of positive affect on the unusualness of word associations.," *Journal of Personality and Social Psychology*, vol. 48, no. 6, pp. 1413–1426, 1985.
- [21] B. L. Fredrickson, "The broaden-and-build theory of positive emotions.," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 359, no. 1449, pp. 1367–1378, Sep. 2004.