Affect and Creative Performance on Crowdsourcing Platforms

Robert R. Morris
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, USA
rmorris@media.mit.edu

Mira Dontcheva
Adobe Research
Adobe
San Francisco, USA
mirad@adobe.com

Adam Finkelstein
Department of Computer Science
Princeton University
Princeton, USA
af@cs.princeton.edu

Elizabeth Gerber
Segal Design Institute
Northwestern University
Evanston, USA
egerber@northwestern.edu

Abstract— **Performance on crowdsourcing platforms varies greatly, especially for tasks requiring significant cognitive effort or creative insight. Researchers have proposed several techniques to address these problems, yet few have considered the role of affect, despite the well-established link between positive affect and creative performance. In this paper, we examine two affective techniques to boost creativity on crowdsourcing platforms – affective priming and affective pre-screening. Across three experiments, we find divergent results, depending on which technique is used. We find that not all happy crowdworkers are alike. Those that are primed to feel happy exhibit enhanced creative performance, whereas those that merely report feeling happy exhibit impaired creative performance. We examine these findings in light of preexisting research on creativity, affect, and mood saliency. Lastly, we show how our findings have implications not only for crowdsourcing platforms, but also for other human-computer interaction scenarios that involve affect and creative performance.**

Keywords— affective computing; crowdsourcing; affective priming; affective self-report; creativity

I. INTRODUCTION

Crowdsourcing offers new ways to recruit and organize labor, both online (e.g., Amazon's Mechanical Turk, Crowdflower) and off (e.g., taskrabbit.com, "communitysourcing" [1]). While crowdsourcing can take many forms, considerable research attention has been devoted to microtask crowdsourcing - an approach characterized by: (1) decomposing large, complex tasks into many small, simple subtasks and (2) routing labor to many workers, either in parallel or across a series of iterative steps. This is an especially algorithmic take on labor production, and, indeed, one of the most popular microtask platforms (Amazon's Mechanical Turk) calls it "artificial, artificial intelligence."

The problem with this conceptualization is that it can overlook the humanity of the workers involved. Crowdworkers aren't merely computational units; they are real people with complex emotions, capable of incredibly creative, higher-order thinking (see [2]). Yet, as of this writing, the most common tasks on Amazon's Mechanical Turk are not particularly creative. They include image labeling, audio transcription, spam filtering, and other, similarly routinized, repetitive tasks.

While these tasks are currently the norm, it is not clear that this will always be the case. As computation becomes more advanced, rote jobs such as these may no longer require human labor. Also, there is already a growing number of crowdsourcing sites that command higher-order thinking and creative insight, such as graphic design (e.g., 99designs.com) and copy-writing (e.g., textbroker.com). Creative crowdsourcing jobs may become more prevalent over time, especially if new techniques emerge to improve the quality of this kind of work.

In this paper, we explore two ways to enhance creativity on crowdsourcing platforms. We focus on affective interventions, because there is a well-established link between positive affect and creative performance. Across three experiments, we examine the effects of both induced and self-reported affect on creative performance. Our results show that affect can significantly influence crowdsourced creativity, but there is a big difference between whether it is induced through priming or merely self-reported through affective questionnaires.

II. BACKGROUND AND RELATED WORK

A. Amazon's Mechanical Turk

We conducted our research on Amazon's Mechanical turk Service (MTurk). On MTurk, employers crowdsource short units of work called "Human Intelligence Tasks" (or HITs). HITs typically take a few seconds to a few minutes to complete. Workers login to the online marketplace, preview the available tasks, and chose a HIT they wish to perform. MTurk is a particularly popular crowdsourcing platform, reaching 500,000 workers and typically offering over 180,000 HITs at any given time [3].

B. Crowdsourcing Creative Work

In recent years, there has been growing interest in creative work on MTurk. For instance, researchers have used MTurk to crowdsource encyclopedia articles [4], poetry translations [5], vacation itineraries [6], fashion tips [7], and even psychotherapeutic insights [8].

Unfortunately, crowdsourcing platforms are often plagued by poor quality [9]. Typical quality control techniques include

administering qualification tests to pre-screen workers, creating gold-standard questions to detect cheating, and filtering results *post hoc*, either through statistical techniques or through crowd-based adjudication. However, subjective tasks, such as those that require creative insight or creative problem solving, are not always amenable to these quality control measures To date, few researchers have considered how a crowdworker's affective state might influence performance, despite considerable evidence that mood influences creative output.

C. Affect and Creative Performance

For decades, researchers have examined links between positive affect and creativity [10]–[12]. Scholars theorize that positive affect supports the broadening of attention to remotely connected ideas and facilitates the recombination of these ideas to solve problems requiring insight [13]. Isen proposed that positive affect influences cognitive activity in three ways [10]. First, positive affect increases the number of cognitive elements available for processing, thereby increasing the number of elements available for association. Second, positive affect supports defocused attention, thereby broadening attention and increasing the number of cognitive elements considered relevant to the problem. Third, positive affect increases cognitive flexibility, thereby increasing the probability of association between diverse cognitive elements [10], [14].

In laboratory studies, Isen and colleagues have established an extensive repertoire (see [10] for a review) of ways to experimentally induce short-term positive affect in non-computing environments. In Isen's studies, participants are given gifts and treats, shown a clip of a comedy film, or played an excerpt from affect-laden music. For example, in two separate experiments with college students, participants were shown a clip of the movie *Gag Reel* or given a small bag of candy to induce affect [14]. In both cases, positive affect improved performance on insight problem-solving tasks, such as Mednick's "Remote Associates Task" (RAT) [15].

D. Music and Affect

While affect can be elicited by many different kinds of stimuli, music is especially evocative. Music is often regarded as the "language of the emotions" [16] and studies report that the most common reason people listen to music is to influence their affective state [17].

In the laboratory, music-evoked affect has been used to influence behavior [18], perception, [19], and attention [19]. Rowe and colleagues found that when participants were exposed to happy music in the laboratory, they solved significantly more insight problems than when exposed to sad music [19]. In our work, we seek to understand whether music can induce affect in a real-world context – in this case crowdsourcing – and if this influences performance.

E. Affective Priming

Affective priming techniques use implicit mechanisms to influence affect and behavior [20]. Primes can be presented at or beneath the level of conscious perception. Whether they can be perceived or not is less important than whether their intended mechanism is detected. Ideally, the intended function of the prime should be cloaked from conscious awareness.

Priming techniques have been explored in psychological research for decades, but they are rarely utilized in real-world human-computer interaction contexts. In part, this is because priming effects are thought to be short-lived. They are also hard to deploy at the right place and the right time. However, in crowdsourcing platforms, where tasks are often short and circumscribed, short-term priming effects, extended across many hundreds of individuals, can reap measurable benefits.

Recently, Lewis et al. [21] showed that affective images can enhance creative performance on MTurk. When happy, positive images were placed within a crowdsourcing task, workers were more creative and generated more novel ideas. In this paper, we build on this research and we examine two affective interventions for crowdsourced creative tasks.

III. Hypotheses

In our paper, we investigate the following hypotheses:

Hypothesis 1: *positive music primes will significantly enhance creative performance on MTurk.*

Hypothesis 2: *individuals on MTurk who report being happy will significantly outperform those who report being unhappy.*

In our first experiment, we address hypothesis 1. We build on previous research on affective priming and crowdsourcing, and we examine whether affective music primes can enhance creative performance on MTurk. We also illustrate ways to streamline the priming process, making it easier to embed primes in online work environments.

In our second and third experiments, we address hypothesis 2. We examine ways to pre-screen workers based on self-reported affective state. If positive affect enhances creative performance, irrespective of how it is induced, it may be possible to forgo priming procedures and simply allocate creative tasks to individuals who already report feeling happy. If a task requires significant creativity, employers could pre-screen workers based on self-reported affect. In this scenario, workers who report feeling happy would be assigned the creative task, while others would be diverted to different work. We examine this technique using two common self-report measures of affective state – the Self Assessment Manikin (SAM) and the Positive and Negative Affect Schedule (PANAS).

IV. Experiment 1: Affective Priming with Music

A. Participants

175 participants from MTurk completed our first experiment. Because our task required mastery of the English language, we restricted enrollment to workers from the United States. While we could not ensure that all our workers were indeed from the United States, we did our best to filter results by eliminating participants whose IP address did not originate in the U.S. This was done to reduce any within-groups variance due to poor English fluency.

B. Materials

To induce affect in this experiment, we used short segments of emotionally evocative music. Our music primes were similar to those used previously in mood induction research [19]. The positive prime was a 30-second excerpt from Bach's *Brandenburg Concerto No. 3* (played by jazz flutist Hubert Laws) . The negative musical prime was a 30-second excerpt from Prokofiev's *Alexander Nevsky: Russian Under the Mongolian Yoke*. Following the procedure described by Rowe and colleagues [9], we slowed the Prokofiev piece to half-speed, while keeping the pitch unchanged. The slow tempo helps emphasize the sad tones of the piece. We did not use a neutral piece as an audio control, because we agree with other music scholars who suggest that it is impossible to find a truly "neutral" musical stimulus [22].

To validate our music primes, we collected 15 affect ratings for each excerpt using workers from MTurk. We conducted this validation with a separate cohort, apart from the group enrolled in our actual experiment. The cohorts were separated because we did not want the rating task to interfere with the priming manipulations used in the actual experiment. Using a separate sample to assess the affective qualities of a stimulus is a common approach used in mood-induction research [23].

To rate the musical excerpts, we used the 9-point Self-Assessment Manikin (SAM) Scale (Figure 1), which uses graphical figures to depict human emotions ranging from smiling (i.e. happy) to frowning (i.e. unhappy) [24]. After listening to the 30-second audio clip, participants were shown the SAM for emotion and arousal and were asked to "*Select the figure that most closely corresponds with how you felt while listening to the previous audio clip."* The results confirmed our expectations and showed that participants felt significantly happier while listening to the positive music (M = 7.8, SD = .86) as compared to the negative music (M = 4.47, SD = 1.19);

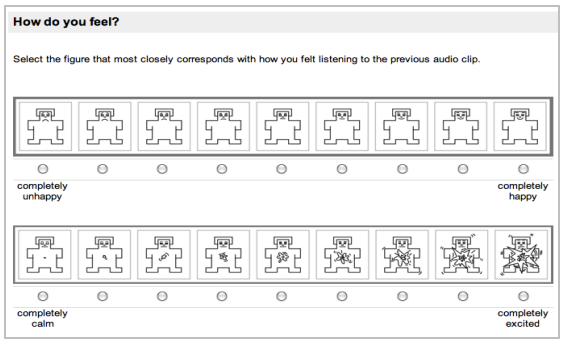

Figure 1. Screenshot of the Self-Assessment Manikin task on MTurk.

$t(28) = 8.80$, $p < .01$. There was no significant difference between the selections in the arousal dimension.

To measure creative performance, we used the Remote Associates Task (RAT) [25]. In the RAT, participants are shown three words and are asked to generate a fourth that forms a two-word phrase or compound with the other three. For example, if given the words "aid," "rubber," and "wagon," the correct answer would be "band", as in "band-aid," "rubber band," or "bandwagon." While hundreds of RAT items have been published, we selected twenty of moderate difficulty, based on norms collected by previous researchers [25].

We chose the RAT task based on several factors. First, the RAT is useful for studying insight and creative performance because it has high predictive validity - that is, achievement on the RAT correlates well with achievement on other creative problem solving tasks [26]. Thus, the results we achieve with the RAT should extend to other crowdsourcing tasks that require creative- and insight-based problem solving. Second, scores on the RAT are unambiguous and easy to interpret. Finally, we chose the RAT because it is considered a test of *convergent* creative thinking, a process where disparate, unrelated ideas must be brought together to solve a problem. Previous work on affective priming on MTurk used a *divergent* creative task, which asked people to generate as many unusual uses of an object as possible [21]. Therefore, we also chose the RAT because we wanted to determine whether affective primes could also enhance convergent creativity on MTurk.

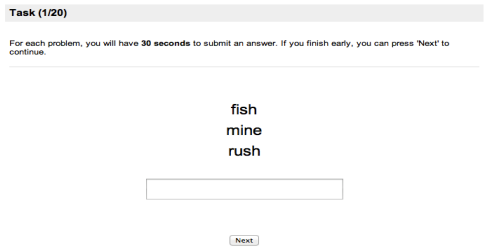

Figure 2. Screenshot of RAT task on MTurk. (Correct answer is **gold**.)

C. Procedure

We paid participants $.35 to complete the task and our HIT description told potential workers that they would "*complete several short word puzzles."* After accepting our HIT, participants read and agreed to an online consent form. Participants were then shown how to complete the RAT task, and were given two examples with solutions provided.

After the instructions phase, participants were randomly assigned to one of three experimental conditions, each disguised as a verification task. Verification tasks are often used on platforms like MTurk to ensure that workers are paying attention. Before a task begins, workers are asked some basic verification questions. If they fail to answer these questions correctly, they can be blocked from completing the work, or the remainder of their answers can be ignored.

In our neutral control condition, the verification task was written as follows: "*To show us that you are paying attention, please write down today's date."* In the positive and negative priming conditions, participants were asked to listen to a 30-second audio clip and describe what they heard in the clip. They were told: "*To show us that you are paying attention, please turn up your computer volume and listen to the following 30 second clip. Please tell us: What do you hear in this audio clip?"* Participants in the positive condition heard the Bach piece, while participants in the negative condition heard the Prokofiev piece.

Immediately following the verification task, participants were given an online version of the RAT task (Figure 2). The presentation order of the 20 RAT items was randomized for each subject. Participants were given up to 30 seconds to complete each question.

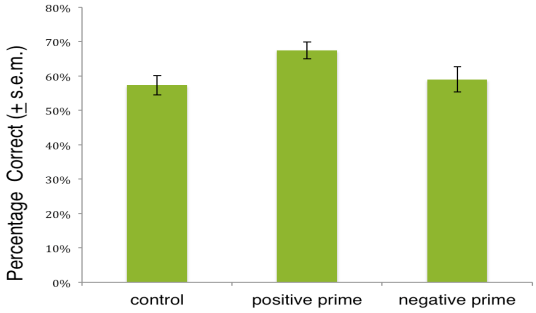

Figure 3. Percentage of RAT items solved correctly, per condition. Error bars reflect standard errors of the mean.

D. Results

A one-way between subjects ANOVA was conducted to examine the effect of music-based affective primes on RAT performance. The means reflect the percentage of test items successfully completed out of 20 (Figure 3). We found a significant effect of priming condition on RAT performance [$F(2,173) = 3.95$, $p = .02$]. Crowdworkers primed with positive music ($M = 67.54$, $SD = 19.44$) significantly outperformed those who were simply asked to write down the date ($M = 57.35$, $SD = 22.84$), $p = .02$, $d = .48$. There was not a significant difference between the negative ($M = 59.04$, $SD = 25.29$) and positive ($M = 67.54$, $SD = 19.44$) conditions. By asking workers to listen and attend to a positive 30-sec music clip, we increased creative performance by over 10%.

E. Discussion

These results extend research findings on affective priming in crowdsourcing environments, showing significant effects of music primes on convergent creative tasks. We also illustrate a new way to deploy primes on these platforms. Previously, researchers tried integrating primes directly within the crowdsourcing tasks. In our design, primes are embedded within an unrelated verification task and therefore decoupled from the target, crowdsourcing task. This design makes it easy to add primes to almost any online crowdsourcing task where verification tasks are appropriate. This approach could also be extended for use in more traditional online work environments. For example, primes could be deployed as part of login processes whenever online workers sign on to complete creative tasks. Future work should examine this possibility.

Our results also show that it is possible to dramatically streamline affective priming techniques to better suit microtask crowdsourcing environments. On crowdsourcing platforms like MTurk, tasks tend to be short and any priming manipulation needs to be done quickly. In our experiment, we improved performance using 30-second music excerpts, whereas previous research has relied upon much longer mood induction procedures. In an experiment conducted by Rowe et al. [19], participants listened to positive and negative musical pieces for a full ten minutes, before performing creative tasks, and they were given booster sessions throughout the task. Participants were also instructed to generate affect-congruent thoughts while listening to the pieces. Our experiment, by contrast, shows that even very short music pieces can act as affective primes and can significantly improve workers' performance.

Yet, it might be possible to further streamline our approach. Instead of taking the trouble to prime individuals, perhaps we can simply allocate work based on quick self-report measures of affect. If positive affect enhances creativity, happy individuals should outperform those who report feeling sad.

V. Experiment 2: Affective Pre-Screening (SAM)

Crowdsourcing platforms make it extremely easy to allocate jobs to specific sets of workers. For instance, MTurk provides ways to recruit workers based on country, past performance, and experience (in terms of the number of jobs successfully completed). For tasks that require creativity, baseline affective state may be just as important as some of the other qualifications one could consider. To maximize the creative output of the workforce, workers could be pre-screened based on affective state.

In our second experiment, we test these ideas by examining whether self-reported affective state might predict creative performance. Specifically, we hypothesized that workers who report feeling happy will outperform other workers.

A. Participants & Procedure

We recruited 367 MTurk workers for this experiment, excluding individuals who had already completed the previous experiment with the RAT and the affective primes. Participants were first asked to fill out the SAM for valence and arousal, indicating how they felt at the current moment. After completing the SAM, the participants were given the same 20-item RAT task that was used in the affective priming study. No primes or verification tasks were employed.

B. Results

To analyze the SAM, we first conducted a linear regression, using SAM scores as the predictor variable. We found that SAM scores for valence significantly predicted performance on the RAT, $b = -.157$, $t(364) = -3.04$, $p<.003$. We also did a quartile analysis to compare performance between individuals who scored in the top and bottom 25% of the SAM (and who thus reported being 'happy' or 'sad', respectively). Contrary to our expectations, those reporting *negative affect* scored significantly better on the RAT ($M = 61.87$ $SD = 18.46$) than those reporting positive affect ($M=53.48$, $SD=21.92$), $t(181) = 2.80$, $p <.01$, $d = .41$. Analyses of SAM scores for arousal showed no significant relationship between self-reported arousal and creative performance.

C. Discussion

Given these unexpected results, we considered alternative hypotheses for why individuals reporting negative moods might outperform those reporting positive moods. The SAM is a simple questionnaire and it is possible that a substantial set of workers answered the questions as quickly as possible. There

might have been some bias for selecting the happy faces (e.g., these were the easiest ones to click on). Moreover, it is possible that lazy or inattentive works made their selections without actually reflecting upon their mood state at all, and these workers may have weakened the overall performance of the happy group. Another hypothesis is that self-reflection of affect disrupts the commonly observed links between positive mood and high creative performance. Perhaps the mere act of filling out the form altered behavior in a significant way. To explore these issues further, we repeated the experiment using a more extensive affective self-report measure.

VI. Experiment 3: Affective Pre-Screening (PANAS)

A. Participants & Procedure

In our third experiment, 237 MTurk workers filled out the brief Positive and Negative Affect Scale (PANAS) – a ten-item questionnaire that yields scores for both positive (PA) and negative affect (NA) [27]. To complete the PANAS, individuals indicate the extent to which they feel certain positive or negative affective states (e.g., "interested," "excited," "hostile," "afraid") at the current moment. These positive and negative state items are interleaved throughout the questionnaire. While workers can still fill out the form without thinking carefully, there is less of a chance that quick, random selections would be biased in the positive or negative direction (as we may have observed with the SAM).

In our experiment, we hypothesized that: (1) individuals reporting high PA would outperform those reporting low PA and (2) individuals reporting low NA would outperform those reporting high NA. As in the SAM experiment, we administered the PANAS immediately prior to the RAT, and we did not introduce any verification tasks in between.

B. Results

A linear regression analysis, with PA and NA as predictor variables, revealed a significant effect of PA on creative performance, b = -.169, $t(235)$ = -2.62, p < .01. However, as with the SAM experiment, the results were in the opposite direction from what we expected. A quartile analysis showed that individuals reporting low PA scored significantly higher on the RAT (M = 56.61, SD = 19.86) than those reporting high PA (M = 49.08, SD = 20.33), $t(117)$ = 2.04, p <.05, d =. 37.

There was no significant effect of NA scores on creative performance. Most individuals reported low NA scores on the PANAS, making it hard to use this as a measure to differentiate between groups. A histogram of the NA scores revealed a non-normal, positively-skewed distribution, which might account for the null results we observed. We found the PA score to be most predictive of creative performance.

C. Discussion

The PANAS results show a similar pattern as the SAM results, suggesting consistency across both self-report measures. The PANAS results also suggest that our findings are probably not an artifact of lazy workers who tried to fill out the form as quickly as possible. While it is possible that some workers filled out the form without thinking, it is unlikely that they were significantly well represented in either the high or low PA group.

Given our findings, we argue that happiness, when assessed by self-report questionnaires, does not correlate with high creative performance on MTurk. In fact, we find a significant *negative relationship* between self-reported happiness and creative performance. On the surface, these results seem to contradict what we observed in the priming study. To explain these divergent findings, it is useful to distinguish how a priming procedure might differ from a self-report procedure.

One of the main differences between the priming and self-report procedures relates to affective self-awareness. Mood clarity should be higher in the latter than the former. Primes are thought to influence mood without significant conscious awareness. Affective self-report measures, by contrast, are designed to call attention to one's mood state. Thus, differences in affective self-awareness might account for our divergent findings. Further, we argue that the mood-as-information model, described below, provides a compelling theoretical basis for this interpretation.

D. Mood-As-Information Model

In their mood-as-information model, Schwarz and Clore [28] argue that moods frequently serve as heuristics for judgment and decision-making. Baseline mood states can sometimes color one's perception of unrelated situations or stimuli in ways we're not aware of. For instance, on a sunny day, our lunch may taste better than usual, not because it tastes any better objectively, but rather because our appraisal is shaped by our pre-existing sunny disposition. Research shows that this effect can also shape our perception of our own performance. Studies show that people in positive moods may be more likely to judge their performance as acceptable than those in negative moods [29].

George and Zhou [30] expand on this model further and show how it relates to creative performance. In their research, they found that creative performance was negatively related to positive mood, but only when clarity of mood and perceived recognition for creativity were high. Their study suggests that, when positive affect is salient, workers put forth less effort on creative tasks. When evaluating how much effort to put forth, workers might confuse their feelings of positive mood with feelings of satisfaction about their performance. This could explain the results we observed in our second and third experiments. Filling out the SAM and the PANAS may have prompted our participants to reflect on their mood state. Participants who observed themselves to be happy may have had rosier perceptions of their own performance and may have put forth less effort as a result.

Future work needs to be done to test this proposition directly. For instance, mood saliency should be assessed before and after affective self-report questionnaires to see what effects, if any, these questionnaires have on affective self-awareness. Also, it might be beneficial to run a multi-factorial study that combines priming techniques with affective self-report measures to see how the two interact with each other. In general, links between affect and creative performance should not be assessed merely in terms of information processing.

While affective states may engage creative cognition, these faculties may not be all that useful if motivation is low.

VII. Conclusion & Design Implications

In this paper, we show how affect can influence creative performance in mircrotask crowdsourcing environments. In our first experiment, we show that positively valenced music can significantly enhance creative performance on MTurk. We also outline ways to embed primes in verification tasks so they can be reused easily. This technique need not be limited to crowdsourcing applications. Verification tasks can be embedded in almost any online work environment, as a way to check if a worker is paying attention. To enhance creative performance, verification tasks should include positive affective stimuli, such as positive music, images, or other media.

In our second and third experiments, we observe a negative relationship between creativity and positive affect, when assessed through self-report measures. Our results provide further support for the mood-as-information model and suggest that affective self-reports cannot be regarded as an unbiased diagnostic tool. Simply reflecting on one's mood state can significantly alter creative performance. To our knowledge, no prior research has observed such a dramatic effect of affective self-report questionnaires on subsequent behavior. These results have implications for researchers in affective computing and other HCI fields that routinely use affective measures to assess interventions and new technologies. Care should be taken when deploying these measures, as they may have substantial consequences on behavior.

As a whole, our findings suggest that affect has important implications for performance on crowdsourcing environments. The rise of crowd-computing might lead some to consider crowdworkers as menial, interchangeable computational units. However, researchers need to remember that crowdworkers are living, breathing people, with complex affective states. For creative crowdsourcing to evolve, researchers should take care to acknowledge the emotions of the crowd.

Acknowledgments

We would like to thank Rosalind Picard, Laura Ramsey and Javier Hernandez for valuable feedback on this work.

References

[1] K. Heimerl, B. Gawalt, K. Chen, T. Parikh, and B. Hartmann, "CommunitySourcing: engaging local crowds to perform expert work via physical kiosks," in *Proceedings of the SIGCHI*, New York, NY, USA, 2012, pp. 1539–1548.

[2] A. Kittur, J. V. Nickerson, M. Bernstein, E. Gerber, A. Shaw, J. Zimmerman, M. Lease, and J. Horton, "The future of crowd work," in *Proceedings of the 2013 conference on Computer supported cooperative work*, New York, NY, USA, 2013, pp. 1301–1318.

[3] "mturk.com." .

[4] A. Kittur, B. Smus, S. Khamkar, and R. E. Kraut, "CrowdForge: crowdsourcing complex work," in *Proceedings of UIST*, New York, NY, USA, 2011, pp. 43–52.

[5] A. Kittur, "Crowdsourcing, Collaboration, and Creativity," *XRDS*, vol. 17, no. 2, pp. 22–26, 2010.

[6] H. Zhang, E. Law, R. Miller, K. Gajos, D. Parkes, and E. Horvitz, "Human computation tasks with global constraints," in *Proceedings of the Proceedings of the SIGCHI*, New York, NY, USA, 2012, pp. 217–226.

[7] M. A. Burton, E. Brady, R. Brewer, C. Neylan, J. P. Bigham, and A. Hurst, "Crowdsourcing subjective fashion advice using VizWiz: challenges and opportunities," in *Proceedings of the ACM SIGACCESS*, New York, NY, USA, 2012, pp. 135–142.

[8] R. R. Morris and R. Picard, "Crowdsourcing Collective Emotional Intelligence," *Collective Intelligence*, Apr. 2012.

[9] A. Kittur, E. H. Chi, and B. Suh, "Crowdsourcing user studies with Mechanical Turk," in *Proceeding of the SIGCHI*, New York, NY, USA, 2008, pp. 453–456.

[10] A. Isen, "Positive Affect," in *The Handbook of Cognition and Emotion*, New York: Wiley, 1999, pp. 521–539.

[11] T. Amabile, S. Barsade, J. Mueller, and B. Staw, "Affect and Creativity at Work," *Administrative Science Quarterly*, vol. 50, pp. 367–403, 2005.

[12] B. Frederickson, "What good are postive emotions?" *Review of General Psychology*, vol. 2, pp. 300–319, 1998.

[13] B. Frederickson, "The broaden-and-build theory of positive emotions," *Philisophical Transactions of the Royal Society*, vol. 359, pp. 1367–1377, 2004.

[14] A. Isen, K. Daubman, and G. P. Nowicki, "Positive affect facilitates creative problem solving," *Journal of Personality and Social Psychology*, vol. 52, pp. 1122–1131, 1987.

[15] M. T. Mednick, S. A. Mednick, and E. V. Mednick, "Incubation of creative performance and specific associative priming," *Journal of Abnormal and Social Psychology*, vol. 69, 1964.

[16] D. Cooke, *The Language of Music*. Oxford University Press, USA, 1959.

[17] P. N. Juslin and D. Västfjäll, "Emotional responses to music: the need to consider underlying mechanisms," *Behav Brain Sci*, vol. 31, no. 5, pp. 559–575; discussion 575–621, 2008.

[18] P. N. Juslin and J. Sloboda, *Handbook of Music and Emotion: Theory, Research, Applications*. Oxford University Press, USA, 2011.

[19] G. Rowe, J. B. Hirsh, and K. A. Anderson, "Positive affect increases the breadth of attentional selection," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 104, pp. 383–388, 2007.

[20] K. C. Klauer and M. Jochen, "Affective priming: Findings and theories," in *The psychology of evaluation: affective processes in cognition and emotion*, 2003, pp. 7–49.

[21] S. Lewis, M. Dontcheva, and E. Gerber, "Affective computational priming and creativity," in *Proceedings of the SIGCHI*, Vancouver, BC, Canada, 2011, p. 735.

[22] M. Roy, J.-P. Mailhot, N. Gosselin, S. Paquette, and I. Peretz, "Modulation of the startle reflex by pleasant and unpleasant music," *International Journal of Psychophysiology*, vol. 71, pp. 37–42, 2009.

[23] A. Isen and A. Erez, "Some Measurement Issues in the Study of Affect," in *Oxford Handbook of Methods in Positive Psychology*, Oxford: Oxford University Press, 2007, pp. 250–265.

[24] M. Bradley and P. Lang, "Measuring Emotion: The Self-Assessment Manikin and the Semantic Differential," *Journal of Behavioral Therapy and Experimental Psychiatry*, vol. 25, pp. 49–59, 1994.

[25] E. M. Bowden and M. Jung-Beeman, "Normative data for 144 compound remote associate problems," *Behavior Research Methods, Instruments, & Computers*, vol. 35, pp. 634–639, 2003.

[26] P. Ansburg, "Individual differences in problem solving via insight," *Current Psychology*, vol. 19, no. 2, pp. 143–146, 2000.

[27] D. Watson, L. Clark, and A. Tellegen, "Development and validation of brief measures of positive and negative affect: The PANAS scale." *Journal of Personality and Social Psychology*, vol. 54, pp. 1063–1070, 1988.

[28] N. Schwarz and G. L. Clore, "How do I feel about it? The information function of affective states." in *Affect, cognition, and social behavior*, Lewinston, NY: 1988, pp. 44–62.

[29] E. J. Johnson and A. Tversky, "Affect, generalization, and the perception of risk," *Journal of Personality and Social Psychology*, vol. 45, no. 1, pp. 20–31, 1983.

[30] J. M. George and J. Zhou, "Understanding when bad moods foster creativity and good ones don't: The role of context and clarity of feelings," *Journal of Applied Psychology*, vol. 87, no. 4, pp. 687–697, 2002.