

Accepted to appear in Journal of Positive Psychology, 2014

Crowd-Powered Positive Psychological Interventions

Robert R. Morris & Rosalind Picard

MIT Media Lab, Massachusetts Institute of Technology, Cambridge, USA

rmorris@media.mit.edu

Crowd-Powered Positive Psychological Interventions

Recent advances in crowdsourcing have led to new forms of assistive technologies, commonly referred to as crowd-powered devices. To best serve the user, these technologies crowdsource human intelligence as needed, when automated methods alone are insufficient. In this paper, we provide an overview of how these systems work and how they can be used to enhance technological interventions for positive psychology. As a specific example, we describe previous work that crowdsources positive reappraisals, providing users timely and personalized suggestions for ways to reconstrue stressful thoughts and situations. We then describe how this approach could be extended for use with other positive psychological interventions. Finally, we outline future directions for crowd-powered positive psychological interventions.

Keywords: crowdsourcing; affective computing; reappraisal; social support

Introduction

Behavioral intervention technologies (BITs) offer exciting new ways to disseminate positive psychological interventions. Self-guided BITs, in particular, have the potential to proliferate widely, reaching underserved populations that might not otherwise have the resources to pursue traditional therapist-led interventions (Schueller, Muñoz, & Mohr, 2013). Unfortunately, when computer-based interventions are purely self-administered, without any oversight from clinicians or experimenters, adherence and motivation drop precipitously (Christensen, Griffiths, & Farrer, 2009; Eysenbach, 2004)

When evaluating the success of self-guided BITs, therefore, researchers must consider both efficacy and adherence. Effective widespread dissemination of self-guided BITs requires designing them to prevent attrition. BITs that are poorly designed and unengaging are of little benefit for the wider population, even if they yield successful outcomes in supervised experiments.

This has important ramifications for self-guided positive psychological activities. According to a recent meta-analysis, self-guided positive psychological interventions are less effective than therapist- or group-led interventions (Sin & Lyubomirsky, 2009). Part of the problem may relate to reduced engagement and motivation. Indeed, research suggests that the success of positive psychological interventions depends, at least in part, on whether one puts forth sufficient effort when doing the activities (Lyubomirsky, Dickerhoof, Boehm, & Sheldon, 2011).

To boost adherence and engagement, BITs should incorporate design strategies from fields such as user interface design (Shneiderman & Plaisant, 2009), user experience design (Garrett, 2010), and persuasive systems design (Törning & Oinas-Kukkonen, 2009), among others. New technologies might also take design inspiration from therapist-led interventions. BITs for positive psychology could be crafted to emulate interactions that support successful engagement in the clinic, such as social

support, accountability and individualized instruction (Mohr, Cuijpers, & Lehman, 2011; Sin, Della Porta, & Lyubomirsky, 2011). For example, an ideal technology might be one that engages conversationally with the user, offering contextually relevant, individualized training combined with on-demand social support.

An interesting challenge for future research, then, is to design BITs that offer some, or perhaps all, of the interactions describe above, without sacrificing the benefits that come from self-administered approaches – namely, accessibility, anonymity and scalability. Incorporating these design features could significantly enhance self-guided interventions for positive psychology and beyond.

In the future, artificial intelligence (AI) may help BITs achieve these goals. Personal assistants could be engineered to not only disseminate intervention concepts and instructions, but to also provide timely support and social interactions that keep the user motivated and entertained. Unfortunately, artificial agents are not yet sophisticated enough to support many of these interactions. They currently cannot construct nuanced, extemporaneous responses that reliably fit the particular context or personal history of the user. Doing so requires tremendous advances in AI, especially in the domains of natural language processing and common-sense reasoning. This is unfortunate because systems with greater human-like intelligence could help fill a huge unmet need for accessible yet powerfully engaging positive psychological interventions.

Because AI cannot yet provide this level of interactivity, some researchers have begun to explore various technological work-arounds, the most common being the use of online crowdsourcing workforces - an approach that is sometimes called “human computation” (von Ahn, 2005). In recent years, these techniques have helped solve challenging problems that cannot be addressed by AI alone. Applying these design paradigms to positive psychology could be an interesting new avenue of research.

In this paper, we review a new breed of assistive crowd-powered interfaces – that is, technologies that leverage human intelligence on demand while still preserving the anonymity and scalability of purely automated tools. We illustrate how these tools might enhance existing positive psychological interventions and foster new approaches that were not feasible before. We believe crowd-powered approaches have the potential to improve the adherence, engagement, and efficacy of self-guided BITs, by offering both social interaction and creative new forms of on demand intervention assistance.

To introduce this approach, we first review existing work in this space, describing technology that crowdsources positive reappraisals in near real-time. We then extrapolate from this example and consider how other positive psychological interventions might similarly benefit from crowd-powered approaches. Finally, we explore future directions for this field, discussing open research questions and new technological opportunities.

Crowdsourcing

Crowdsourcing refers to recruiting and organizing ad hoc labour, usually through an open call for participation (Howe, 2006). Crowds are not bound by contractual obligations and, for the most part, they can come and go as they please. The fluidity of this workforce has its drawbacks, but, when managed carefully, crowdsourcing can be an extremely powerful tool. For example, crowdsourcing systems have helped digitize millions of documents (Ahn, Maurer, McMillen, Abraham, & Blum, 2008), label millions of images (von Ahn & Dabbish, 2004), and classify immense planetary data (Kanefsky, Barlow, Gulick, & Norvig, 2001). In just three weeks, crowdworkers discovered the protein structure of M-PMV – a challenge that had previously eluded scientists for over ten years (Khatib et al., 2011).

In recent years, new platforms have evolved to support the practice of crowdsourcing. One such platform is Amazon's Mechanical Turk service (MTurk), a

crowdsourcing marketplace that employs hundreds of thousands of crowdworkers (<http://www.mturk.com>). Jobs are done online and follow a so-called ‘microtask’ structure; they are short and circumscribed and are often completed for less than a dollar (Kittur, Chi, & Suh, 2008).

Crowd-Powered Assistive Devices

One notable feature of MTurk is its application programming interface (API), which allows anyone to programmatically post jobs, hire workers, accept work, reject work, and award bonuses. Recently, human-computer interaction researchers have begun to exploit these features, creating sophisticated crowd-powered systems that recruit and coordinate human labor through algorithmic processes. *Soylent*, for example, is a word processor built upon these principles (Bernstein et al., 2010). One of *Soylent*’s features is ‘shorten’, which reduces a document to 85% of its original length without compromising the meaning of the text. It accomplishes this by recruiting multiple crowdworkers in succession to remove clutter in the document. One team of workers scans each paragraph and identifies words and sentences that can be deleted or shortened. Another team makes the changes and, finally, yet another group of workers reviews and approves the work. Similar algorithmic techniques have been used to crowdsource encyclopedia articles (Kittur, Smus, Khamkar, & Kraut, 2011), poetry translations (Kittur, 2010), travel itineraries (Zhang et al., 2012), fashion tips (Burton et al., 2012), and visual assistance for the visually impaired (Bigham et al., 2010). In all cases, worker recruitment, worker training, worker coordination, and quality control are coordinated entirely through code. In some instances, the crowd is hidden and the end-user may not even be aware that human workers are powering the system from within. In others, the crowd’s efforts are exposed to the user. This latter approach can be especially useful for assistive devices that might benefit from the engaging properties of direct human interaction and social support.

Quality Control

All of the aforementioned examples contained techniques to ensure the quality of the crowdsourced contributions. These techniques include administering qualification tests to pre-screen workers and employing ‘gold-standard’ questions to detect cheating (Kittur, Chi, & Suh, 2008). Crowdworkers can also be hired to curate the work of their peers, returning only the best responses from the crowd (Bernstein et al., 2010). Other approaches involve priming workers with positive stimuli (e.g., positive music or images) to boost creativity and motivation (Morris, Dontcheva, Finkelstein, & Gerber, 2013), or filtering out workers whose behavioral signatures (e.g., scrolling through instructions too quickly to be actually reading them) predict poor task performance (Rzeszotarski & Kittur, 2011).

Crowd-Powered Positive Psychological Interventions

In many ways, crowd-powered systems are extremely well-suited for positive psychological BITs. Crowd-powered tools could offer on-demand assistance, helping reinforce important positive psychology concepts as needed, while also responding intelligently to the user’s unique situation. When used appropriately, these tools could also help boost adherence and sustain engagement. In the next sections, we contextualize this approach for the positive psychology community, by reviewing a previously published system designed to crowdsource positive reappraisals. We then describe how a similar approach could be used to enhance components of other positive psychological interventions, such as resilience training, goals and planning, using signature strengths, and performing acts of kindness.

Positive Reappraisal

Reappraisal is an adaptive emotion regulatory skill that involves reinterpreting the meaning of a thought or situation to alter its emotional impact (Gross, 1998). Its use is associated with high levels of well-being (Gross & John, 2003), lowered incidence of

depression (Garnefski & Kraaij, 2006), and strong interpersonal skills (Gross & John, 2003). It has also been used within interventions designed to produce positive affect (Moskowitz et al., 2011). Unfortunately, reappraisal strategies do not come naturally to everyone and they can be hard to utilize during high intensity emotional situations (Sheppes & Gross, 2011). While peers can be a great source of reappraisal assistance (Marroquín, 2011; Panzarella, Alloy, & Whitehouse, 2006), social support is not always available and some situations may necessitate anonymity. New technologies could help serve these needs by providing reappraisal support that is at once anonymous, real-time and contextualized to the unique circumstances of the user.

To explore this space, we recently built a crowdsourcing system that helps users generate healthy, positive reappraisals of stressful situations (for a detailed description of the system design and experimental findings, see Morris & Picard, 2012). To illustrate how it works, consider the following user scenario: A woman named Ellen is at work and her boss berates her in front of all her colleagues. While she is familiar with reappraisal skills, she still needs practice and could use some assistance. She opens an application and writes, “My boss yelled at me today and humiliated me in front of my colleagues. I’m such a loser. I feel ashamed.”

Her text is immediately sent to several sets of crowdworkers on MTurk trained to help her respond more positively to the situation. The first set of workers is trained to offer empathetic support and reassurance and is asked to let Ellen know her feelings make sense, given the situation. Another set of workers then attempts to reappraise the situation. For instance, these workers might help Ellen interpret the event in less personal terms, shifting the focus onto her boss’s behavior. One worker might say: “If you made a mistake, it’s your boss’s job to help you, not to make things worse by yelling like an idiot.” Yet another worker might say, “Maybe your boss has some serious anger

issues? I could even feel sorry for this person - not being able to control your temper can be a serious health problem. It also makes you pretty unlikeable!” Because it is assumed that these workers lack training in reappraisal strategies, they are first given a short three to five minute online tutorial about the concept. They are also given a short quiz along with feedback that differentiates between good and bad reappraisals (e.g., ones that help a user think differently vs. ones that simply advocate problem-solving advice). The workers are not asked to administer a particular type of reappraisal, (e.g., 'detaching oneself from the situation' or 'finding the silver lining'; McRae, Ciesielski, & Gross, 2012), but are instead told in general terms to help the user find creative and realistic ways to reinterpret the situation in a more positive light. Future research could examine whether specific reappraisal tactics might be solicited at different times, perhaps depending upon the user's preference or the nature of the user's situation.

Within a few minutes after submitting the text, the user starts receiving responses from the crowd. The latency of the responses depends on various factors, including (a) how much workers get paid, (b) whether workers have previous training and are able to advance beyond the usual, preliminary instructions and (c) whether workers have agreed to wait on retainer and be ready the second a new task arrives in the queue. When designed appropriately, crowd-powered systems can return results extremely quickly. For simple tasks, with workers placed on retainer, researchers have achieved results within two seconds (Bernstein, Brandt, Miller, & Karger, 2011).

Once the crowd has completed the tasks, the user rates the responses, which helps refine the system further. Workers who submit low-ranking responses can be excluded from future tasks, while high-performing workers can be sought out intentionally (e.g., by sending them an email notification whenever a new job is available and by paying them higher wages).

MTurk has an international, multilingual workforce (Ipeirotis, 2010), and many workers might lack the English skills required to compose well-written reappraisals. To address this concern, our system uses several techniques to increase the odds that workers are fluent in English. Workers must be registered as a U.S. worker in the MTurk database, have an IP address from the U.S. (as suggested by Shapiro, Chandler, & Mueller, 2013), and pass a short, two minute SAT verbal exam.

Responses are also filtered by the crowd, and are only returned to the user if they are approved by multiple crowdworkers who are trained to reject malicious or poorly written responses. If a response is deemed inappropriate, it is discarded and a new worker is recruited to try again. The algorithm that guides this entire workflow is completely automated using the MTurk API. The user needs only to submit the original text – there is no need for the user or any other human to coordinate or oversee the subsequent crowdsourcing workflows (see Figure 1).

A prior study of this system showed that crowdsourced responses were rated significantly higher than responses generated from an open-response structure, wherein workers were simply asked to help the user feel better (Morris & Picard, 2012). However, a longitudinal study examining repeated use of the entire system has yet to be conducted. More work is needed to see whether repeated exposure to this kind of intervention might improve reappraisal ability and whether this, in turn, might promote greater subjective well-being for the user. In the following sections, we extrapolate from this example and consider how similar crowd-powered workflows might be applied to other positive psychological interventions.

Putting it in Perspective

One aim from the Penn Resilience program (Reivich, Gillham, Chaplin, & Seligman, 2013) is to help individuals apply ‘realistic optimism’ when considering the implications of stressful situations. Positive reappraisal, as described above, could

certainly help achieve this end. Another technique described by Reivich and colleagues is ‘putting it in perspective’ – an exercise that involves generating best- and worst-case scenarios to adverse events. This technique balances dichotomous or catastrophic interpretations of stressful events with equally extreme best-case scenarios to help individuals recognize the unrealistic nature of their initial responses.

Returning to the example with Ellen, we posted the situation with her boss on MTurk and asked workers to generate extreme best-case and worst-case scenarios, as a simple proof-of-concept pilot test. As a worst-case scenario, one worker wrote, *“One of your coworkers will capture the incident on her cellphone and post it on YouTube with a soundtrack featuring "Loser" lyrics. The video is played at the annual department meeting.”* As a best-case scenario, one worker wrote: *“Your boss apologizes and explains that he hasn’t been sleeping well. You correctly diagnose him as having sleep apnea, he gets better, and as thanks he gives you a bonus in the form of a new Porsche [sic].”* As described previously, proper quality control techniques would likely improve these responses further.

When individuals are first exposed to ‘putting it in perspective,’ it might be useful to have crowdworkers help generate some initial, extreme ideas as shown above. Because the workers are distanced from the situation, they may be better positioned to generate quick and creative ideas. Further, the user might find the content and sheer variety of the crowd responses entertaining, which could lead him/her to engage with the technique more often. Monotony and routine can be a serious problem for some positive psychological activities (Layous & Lyubomirsky, 2012). Crowdsourced interactions, such as the ones shown thus far, could add spontaneity to intervention activities that are at risk of growing stale.

Goal Setting and Planning

Goal setting and planning interventions can significantly enhance well-being (Cheavens, Feldman, Gum, Michael, & Snyder, 2006; MacLeod, Coates, & Hetherington, 2008). Crowd-powered technologies could intervene in at least two cognitive components believed to be important for healthy goal pursuit – “pathways thinking” and “agency thinking” (Snyder, 2002).

Pathways thinking, as described by Snyder (2002), is the process of identifying the specific routes and steps needed to attain a goal. Crowds could be recruited to help users define clear blueprints for goal attainment. They could also be recruited to intervene whenever an individual reaches an impasse and needs to brainstorm alternate pathways to achieve a goal. Research by Zhang et al. (2012) shows how crowds can be recruited to help plan vacation itineraries, making sure each activity is a reasonable precursor to the next (both in terms of the time required to complete the task and the distance needed to travel to the next stop). Similar design approaches could be used to help individuals identify specific, manageable goals and the various routes for achieving those goals. In practice, such an approach would be most appropriate for goals that do not require considerable domain expertise (e.g., finding ways to read more vs. perfecting a machine learning algorithm). However, crowdsourcing platforms that incorporate social network information can match users with respondents who have similar interests and skills (Horowitz & Kamvar, 2010). Systems such as these might be useful for individuals that need help planning goals within specialized domains.

Agency thinking is defined as the perceived ability to achieve one’s goals (Snyder, 2002). This is another component thought to be important in successful goal pursuit. To bolster this faculty, crowds could be hired to congratulate users when goals are achieved or boost their confidence when goals seem unduly challenging. Simple

messages of encouragement, crowdsourced at opportune times, might help individuals feel more confident in their ability to attain their goals.

Using Personal Strengths

Utilizing personal strengths is associated with a host of positive outcomes, such as enhanced well-being (Proctor, Maltby, & Linley, 2011), less stress (Wood, Linley, Maltby, Kashdan, & Hurling, 2011), and greater work satisfaction (Peterson, Stephens, Park, Lee, & Seligman, 2009). Crowd-powered technologies could be used to augment strengths interventions in various ways. For instance, crowds could be trained to help people identify optimal ways to utilize signature strengths. As crowds help people identify goals and planning strategies, as discussed earlier, they could tailor their recommendations to align with the user's signature strengths. Indeed, evidence suggests that strengths use mediates the relationship between goal pursuit and well-being (Linley, Nielsen, Wood, Gillett, & Biswas-Diener, 2010).

Crowds could also help a user brainstorm new ways to use certain types of signature strengths throughout daily life, outside the context of goal interventions. For those whose signature strengths include 'Love of Learning,' for example, crowds could help users find ways to pursue new pathways to knowledge (such as providing links to interesting online courses, blogs, or adult education classes). Previous work in human-computer interaction has outlined ways to successfully crowdsource idea generation in domains as diverse as product ideation (Poetz & Schreier, 2012), poetry translation (Kittur, 2010), and furniture design (Yu & Nickerson, 2011). Crowds could similarly produce useful ideas for strengths use that may not have occurred to the user. While automatic recommendation systems for movies and music are now commonplace and may eventually suggest interesting and apt ways to use strengths, these algorithms require massive amounts of user data to work effectively (Resnick & Varian, 1997). To predict what movies you might like, for instance, recommendation algorithms need to

know a lot about your own movie tastes and how they compare with similar others.

Until individuals are commonly reporting ways in which their strengths align with the activities they pursue, subtle and personalized recommendations for strengths use will likely remain the purview of human intelligence.

Crowds could also be tapped to provide encouragement and positive feedback as individuals pursue strengths use interventions. Buckingham (2010) suggests that only 17% of people report using their strengths ‘most of the time’ each day. Offering individuals motivation, as well as more interesting and varied opportunities to exercise strengths, might help encourage greater use of this technique.

Acts of Kindness

To date, most crowd-powered systems have relied on crowdsourcing markets like MTurk. However, these systems can be more self-sustaining if the end-users are incentivized to contribute themselves, in lieu of for-pay workforces. This approach not only solves throughput issues, it may also have therapeutic benefits for the user as well; applying intervention techniques to others can be rewarding in itself and can be a great way to learn the concepts in a deep and meaningful way. Also, helping others through acts of kindness has been explored as a positive psychological intervention, and research suggests that it can boost mood and well-being, provided the tasks are not too repetitive (Lyubomirsky, Sheldon, & Schkade, 2005). Crowd-powered assistive platforms, such as those described in this paper, offer many avenues to help other people and its quite possible that simply being a respondent on these systems will have significant benefits for well-being.

Future Directions

The crowd-powered approaches described in this paper are still developing and future research is needed to understand how they can be used effectively, especially in the context of positive psychological interventions. Future work should address what types

of people might benefit most from crowdsourced feedback and which interventions are most suitable for this type of approach. The technologies and design choices that support crowd-powered tools are also ripe for additional research and innovation. Advances in natural language processing, for example, can help groom responses from the crowd, ensuring that malicious or inappropriate responses are not returned to the user. Algorithms can also be trained to automatically detect text features that correlate with good responses (e.g., sentiment, grammar), obviating the need for crowdsourced voting. In the future, with enough training data, AI algorithms might even be able to compose intelligent responses as well as vet them. Indeed, one of the benefits of crowd-powered systems is their ability to generate a massive corpus of data. New algorithms can learn to mimic the responses that are most well-received in human-mediated crowdsourcing systems, so that future systems are not wholly dependent on real-time, human contributions.

Another open question relates to the amount of support crowd-powered systems should provide. At early stages in an intervention, crowdsourced suggestions and hints can be valuable and engaging. However, for some interventions, it might be wise to slowly taper-off this support and let the user do all the legwork him or herself. More work is needed to examine how and when crowdsourced assistance should be solicited.

Crowd-powered systems should also move beyond paid labor markets and make efforts to attract unpaid volunteers who are interested in contributing. Individuals who self-select as volunteers should be more intrinsically motivated than those working primarily for money or other extrinsic rewards (Deci et al., 1999). One way to increase this pool of unpaid volunteers, as described previously, is to encourage the users to become contributors themselves.

Finally, more work is needed to determine how these systems can improve intervention engagement. Done poorly, feedback from crowdworkers could have the opposite of its intended effect. As noted by Layous and Lyubomirsky (2012), social support that merely cheerleads a person with generic messages of encouragement could come across as patronizing or nagging, thereby reducing engagement. Crowdsourced messages, however, are highly personalized and nuanced and thus could be extremely motivating, helping users adhere to interventions they might otherwise find challenging or tiresome. Receiving responses back as close to the moment as possible is also important (Mamykina et al., 2011). Fortunately, modern crowdsourcing techniques can produce results within a few minutes or even seconds (Bernstein, Brandt, Miller, & Karger, 2011). Lastly, sharing information about the self is intrinsically rewarding (Tamir & Mitchell, 2012) and a crowd-powered system that offers social sharing opportunities could be especially alluring.

The use of crowd-powered techniques to support positive psychological interventions is a largely unexplored space. The ideas presented in this paper are by no means exhaustive. To realize the potential of this approach, new collaborations are needed between the fields of positive psychology and human-computer interaction (particularly crowdsourcing and social computing). A confluence of approaches from these disciplines could help forge exciting new tools for the advancement of human well-being.

References

- Ahn, L. von, Maurer, B., McMillen, C., Abraham, D., & Blum, M. (2008). reCAPTCHA: Human-Based Character Recognition via Web Security Measures. *Science*, *321*(5895), 1465–1468.
- Bernstein, M., Little, G., Miller, R., Hartmann, B., Ackerman, M., Karger, D., ... Panovich, K. (2010). Soylent: A Word Processor with a Crowd Inside. In Perlin, K., Czerwinski, M., & Miller, R. (Eds.), *Proceedings of the 23rd Symposium on User Interface Software and Technology* (pp. 313-322). New York, NY: ACM Press.
- Bernstein, M. S., Brandt, J., Miller, R., & Karger, D. R. (2011). Crowds in Two Seconds: Enabling Realtime Crowd-Powered Interfaces. In Pierce, J., Agrawala, M., & Klemmer, S. (Eds.), *Proceedings of the 24th Symposium on User Interface Software and Technology* (pp. 33-42). New York, NY: ACM Press.
- Bigham, J. P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R. C., ... Yeh, T. (2010). VizWiz: nearly real-time answers to visual questions. In Perlin, K., Czerwinski, M., & Miller, R. (Eds.), *Proceedings of the 23rd Symposium on User Interface Software and Technology* (pp. 333-342). New York, NY, USA: ACM.
- Buckingham, M. (2010). *Go put your strengths to work: 6 powerful steps to achieve outstanding performance*. New York: Free Press.
- Burton, M. A., Brady, E., Brewer, R., Neylan, C., Bigham, J. P., & Hurst, A. (2012). Crowdsourcing subjective fashion advice using VizWiz: challenges and opportunities. In Huenerfauth, M. & Kurniawan, S. (Eds.), *Proceedings of the 14th international ACM SIGACCESS conference on Computers and accessibility* (pp. 135-142). New York, NY, USA: ACM.
- Cheavens, J. S., Feldman, D. B., Gum, A., Michael, S. T., & Snyder, C. R. (2006). Hope Therapy in a Community Sample: A Pilot Investigation. *Social Indicators Research*, *77*(1), 61–78.
- Christensen, H., Griffiths, K. M., & Farrer, L. (2009). Adherence in internet interventions for anxiety and depression. *Journal of Medical Internet Research*, *11*(2), e13.
- Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, *125*(6), 627–668.

- Eysenbach, G. (2005). The law of attrition. *Journal of Medical Internet Research*, 7(1), e11.
- Garnefski, N., & Kraaij, V. (2006). Relationships between cognitive emotion regulation strategies and depressive symptoms: A comparative study of five specific samples. *Personality and Individual Differences*, 40(8), 1659–1669.
- Garrett, J. J. (2010). *The Elements of User Experience: User-Centered Design for the Web and Beyond* (2nd ed.). Thousand Oaks, CA, USA: New Riders Publishing.
- Gross, J. J. (1998). Antecedent- and Response-Focused Emotion Regulation: Divergent Consequences for Experience, Expression, and Physiology. *Journal of Personality and Social Psychology*, 74(1), 224–237.
- Gross, J. J., & John, O. P. (2003). Individual differences in two emotion regulation processes: implications for affect, relationships, and well-being. *Journal of Personality and Social Psychology*, 85(2), 348–362.
- Horowitz, D., & Kamvar, S. D. (2010). The anatomy of a large-scale social search engine. *WWW*, 431-440.
- Howe, J. (2006, June). The Rise of Crowdsourcing. *Wired*, 14(6), 1-4.
- Ipeirotis, P. (2010). Demographics of Mechanical Turk. *CeDER-10-01 working paper*, New York University.
- Kanefsky, B., Barlow, N. G., Gulick, V. C., & Norvig, P. (2001). Can Distributed Volunteers Accomplish Massive Data Analysis Tasks? In *Proceedings of Lunar and Planetary Science*, Abstract 1272.
- Khatib, F., DiMaio, F., Group, F. C., Group, F. V. C., Cooper, S., Kazmierczyk, M., ... Baker, D. (2011). Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature Structural & Molecular Biology*, 18(10), 1175–1177.
- Kittur, A. (2010). Crowdsourcing, collaboration and creativity. *XRDS*, 17(2), 22–26.
- Kittur, A., Chi, E. H., & Suh, B. (2008). Crowdsourcing user studies with Mechanical Turk. In Czerwinski, M., Lund, A., & Tan, D. (Eds.), *Proceedings of the 2008 ACM annual conference on Human Factors in Computing Systems* (pp. 453–456). New York, NY, USA: ACM.
- Kittur, A., Smus, B., Khamkar, S., & Kraut, R. E. (2011). CrowdForge: crowdsourcing complex work. In Pierce, J., Agrawala, M., & Klemmer, S. (Eds.), *Proceedings of the 24th Symposium on User Interface Software and Technology* (pp. 43–52). New York, NY, USA: ACM.

- Layous, K., & Lyubomirsky, S. (in press). The how, who, what, when, and why of happiness: Mechanisms underlying the success of positive interventions. In J. Gruber & J. Moskowitz (Eds.), *The light and dark side of positive emotions*. New York: Oxford University Press.
- Linley, P. A., Nielsen, K. M., Wood, A. M., Gillett, R., & Biswas-Diener, R. (2010). Using signature strengths in pursuit of goals: Effects on goal progress, need satisfaction, and well-being, and implications for coaching psychologists. *International Coaching Psychology Review*, (5), 6–15.
- Lyubomirsky, S., Sheldon, K.M., & Schkade, D. (2005). Pursuing happiness: The architecture of sustainable change. *Review of General Psychology*, 9, 111-131.
- Lyubomirsky, S., Dickerhoof, R., Boehm, J. K., & Sheldon, K. M. (2011). Becoming happier takes both a will and a proper way: An experimental longitudinal intervention to boost well-being. *Emotion*, 11(2), 391–402.
- MacLeod, A. K., Coates, E., & Hetherington, J. (2008). Increasing well-being through teaching goal-setting and planning skills: results of a brief intervention. *Journal of Happiness Studies*, 9(2), 185–196.
- Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., & Hartmann, B. (2011). Design lessons from the fastest Q&A site in the west. In Tan, D., Fitzpatrick, G., Gutwin, C., Begole, B., & Kellogg, W. (Eds.), *Proceedings of the 2011 ACM annual conference on Human Factors in Computing Systems* (pp. 2857–2866). New York, NY, USA: ACM.
- Marroquín, B. (2011). Interpersonal emotion regulation as a mechanism of social support in depression. *Clinical Psychology Review*, 31(8), 1276–1290.
- McRae, K., Ciesielski, B., & Gross, J. J. (2012). Unpacking cognitive reappraisal: Goals, tactics, and outcomes. *Emotion*, 12(2), 250–255.
- Mohr, D. C., Cuijpers, P., & Lehman, K. (2011). Supportive accountability: a model for providing human support to enhance adherence to eHealth interventions. *Journal of Medical Internet Research*, 13(1), e30.
- Morris, R. R., & Picard, R. (2012, April). *Crowdsourcing Collective Emotional Intelligence*. Paper presented at Collective Intelligence, Cambridge, MA.
- Morris, R.R., Dontcheva, M., Finkelstein, A., & Gerber, E. (2013). Affect and creative performance on crowdsourcing platforms. In *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*, (pp. 250-255). Geneva, Switzerland.

- Moskowitz, J. T., Hult, J. R., Duncan, L. G., Cohn, M. A., Maurer, S. A., Bussolari, C., & Acree, M. (2011). A Positive Affect Intervention for People Experiencing Health-Related Stress: Development and non-Randomized Pilot Test. *Journal of Health Psychology, 17*(5), 676-692.
- Panzarella, C., Alloy, L. B., & Whitehouse, W. G. (2006). Expanded Hopelessness Theory of Depression: On the Mechanisms by which Social Support Protects Against Depression. *Cognitive Therapy and Research, 30*(3), 307–333.
- Peterson, C., Stephens, J. P., Park, N., Lee, F., & Seligman, M. E. P. (2009). Strengths of character and work. In Linley, P.A., Harrington, S., & Garcea, N. (Eds.), *Oxford Handbook of Positive Psychology and Work* (pp. 221-235). Oxford, UK: Oxford University Press.
- Poetz, M. K., & Schreier, M. (2012). The Value of Crowdsourcing: Can Users Really Compete with Professionals in Generating New Product Ideas? *Journal of Product Innovation Management, 29*(2), 245–256.
- Proctor, C., Maltby, J., & Linley, P. A. (2011). Strengths Use as a Predictor of Well-Being and Health-Related Quality of Life. *Journal of Happiness Studies, 12*(1), 153–169.
- Reivich, K., Gillham, J. E., Chaplin, T. M., & Seligman, M. E. P. (2013). From Helplessness to Optimism: The Role of Resilience in Treating and Preventing Depression in Youth. In G. Sam & B. R. B (Eds.), *Handbook of Resilience in Children* (pp. 201–214). New York, New York: Spring Science+Business Media.
- Resnick, P., & Varian, H. R. (1997). Recommender Systems. *Commun. ACM, 40*(3), 56–58.
- Rzeszotarski, J. M., & Kittur, A. (2011). Instrumenting the crowd: using implicit behavioral measures to predict task performance. In Pierce, J., Agrawala, M., & Klemmer, S. (Eds.), *Proceedings of the 24th Symposium on User Interface Software and Technology* (pp. 13–22). New York, NY, USA: ACM.
- Schueller, S. M., Muñoz, R. F., & Mohr, D. C. (2013). Realizing the Potential of Behavioral Intervention Technologies. *Current Directions in Psychological Science, 22*(6), 478–483.
- Shapiro, D. N., Chandler, J., & Mueller, P. A. (2013). Using Mechanical Turk to Study Clinical Populations. *Clinical Psychological Science, 1*(2), 213–220.

- Sheppes, G., & Gross, J. J. (2011). Is Timing Everything? Temporal Considerations in Emotion Regulation. *Personality and Social Psychology Review*, *15*(4), 319–331.
- Shneiderman, B., & Plaisant, C. (2009). *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. London: Pearson PLC.
- Sin, N. L., & Lyubomirsky, S. (2009). Enhancing well-being and alleviating depressive symptoms with positive psychology interventions: a practice-friendly meta-analysis. *Journal of Clinical Psychology*, *65*(5), 467–487.
- Sin, N. L., Della Porta, M. D., & Lyubomirsky, S. (2011). Tailoring Positive Psychology Interventions to Treat Depressed Individuals. In Donaldson, S.I., Csikszentmihalyi, M., & Nakamura, J. (Eds.), *Applied Positive Psychology: Improving Everyday Life, Health, Schools, Work, and Society*. (pp. 79–96). New York, NY: Routledge.
- Snyder, C. R. (2002). Hope theory: Rainbows in the mind. *Psychological Inquiry*, *13*(4), 249–275.
- Tamir, D.I., & Mitchell, J.P. (2012). Disclosing information about the self is intrinsically rewarding. *Proceedings of the National Academy of Sciences*, *109*(21), 8038-8043.
- Torning, K., & Oinas-Kukkonen, H. (2009). Persuasive System Design: State of the Art and Future Directions. In Chatterjee, S. & Dev, P. (Eds.), *Proceedings of the 4th International Conference on Persuasive Technology* (pp. 30:1–30:8). New York, NY, USA: ACM.
- Von Ahn, L. (2005). Human Computation.
- Von Ahn, L., & Dabbish, L. (2004). Labeling images with a computer game. In Dykstra-Erickson, E. & Tscheligi, M. (Eds.), *Proceedings of the 2004 ACM annual conference on Human Factors in Computing Systems* (pp. 319–326). New York, NY, USA: ACM.
- Wood, A. M., Linley, P. A., Maltby, J., Kashdan, T. B., & Hurling, R. (2011). Using personal and psychological strengths leads to increases in well-being over time: A longitudinal study and the development of the strengths use questionnaire. *Personality and Individual Differences*, *50*(1), 15–19.
- Yu, L., & Nickerson, J. (2011). Cooks or cobblers? crowd creativity through combination. In Tan, D., Fitzpatrick, G., Gutwin, C., Begole, B., & Kellogg, W.

(Eds.), *Proceedings of the 2011 ACM annual conference on Human Factors in Computing Systems*. New York, NY, USA: ACM.

Zhang, H., Law, E., Miller, R., Gajos, K., Parkes, D., & Horvitz, E. (2012). Human computation tasks with global constraints. In Konstan, J.A., Chi, E.H., & Hook, K. (Eds.), *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems* (pp. 217–226). New York, NY: ACM.