Journal of
Social
Media for
Organizations

A Roadmap for Open Innovation Systems

Mark Klein, Gregorio Convertino

Volume 2, Number 1

Published by the MITRE Corporation

A Roadmap for Open Innovation Systems

Mark Klein, *m_klein@mit.edu*
Gregorio Convertino, *gconvertino@informatica.com*

ABSTRACT

Open innovation systems have provided organizations with unprecedented access to the "wisdom of the crowd," allowing them to collect candidate solutions for problems they care about, from potentially thousands of individuals, at very low cost. These systems, however, face important challenges deriving, ironically, from their very success: they can elicit such high levels of participation that it becomes very challenging to guide the crowd in productive ways, and pick out the best of what they have created. This article reviews the key challenges facing open innovation systems and proposes some ways the research community can move forward on this important topic.

KEYWORDS

Open Innovation Systems, Idea Management Systems, Large-Scale Deliberation

WHAT ARE OPEN INNOVATION SYSTEMS?

"No matter who you are, most of the smartest people work for someone else"
Sun Microsystems co-founder Bill Joy (Lakhani & Panetta, 2007)

Since at least World War II, the dominant model for innovation in many large organizations, especially in the private sector, has been closed in at least two ways (Chesbrough, 2003). First, it excludes sources external to the organizations; a "do it all yourself" model (Chesbrough, Vanhaverbeke, & West, 2008). Second, even innovation *within* the organization has been relegated to a limited number of roles such as R&D lab staff members, CTOs, or managers in charge of innovation and technology transfer. This model thus excludes most employees, as well as customers, partners and other stakeholders in the organization, who remain an untapped resource. In 2004, for example, R&D employees represented 3.8% of the total employees in US multinational corporations, implying that 96% of the employees were not engaged in innovation for the organization (Moris, 2004; Yorgason, 2007).

Open innovation systems[1] represent a promising emerging approach to addressing this important limitation. In such systems, a customer (e.g., a manager in a company or a leader in a public organization) describes a problem they want to solve (e.g., "we want ideas for new beverage products") and provides an online tool that allows potentially thousands of individuals to submit proposed solutions, as well as rate (and sometimes critique) other people's proposed solutions. In some cases, these systems frame the engagement as a competition, where the authors of winning ideas receive an award, financial or otherwise (Morgan & Wang, 2010). They thus allow an organization, at very low cost, to extend the sources of innovation to include a much broader selection of its own employees, in addition to customers, suppliers, stakeholders, and other interested parties (Klein & Convertino, 2014; Chesbrough, 2003; Von Hippel, 2009; West & Lakhani, 2008; M. l. Tierney & Drury, 2015; B.P. Bailey & Horvitz, 2010; Bjelland & Wood, 2008; Chesbrough et al., 2008; Morgan & Wang, 2010; Westerski, Iglesias, & Nagle, 2011; Hrastinski, Kviselius, Ozan, & Edenius, 2010; Piller & Walcher, 2006; Enkel, Perez-Freije, & Gassmann, 2005; Gassmann, 2006; Lakhani & Panetta, 2007; Ogawa & Piller, 2006). Open innovation systems are distinct from related groupware technologies such as brainstorming (Paulus & Nijstad, 2003), argumentation (Klein, 2007; Moor & Aakhus, 2006) and group decisions support systems (Vetschera, 1990) along the critical dimensions of *scale* of users and *goal* of system. Such groupware systems are aimed at enabling *collaboration and problem solving* with small and medium size *groups* (i.e., where the number of members ranges typically from 3 to 30 people), while open innovation systems are aimed at enabling *innovation* by *crowds* (i.e., where there are typically on the order of thousands of participants). That is, open innovation systems assume, by design, a larger scale of users and a narrower goal (i.e., supporting innovation for an organization). While a closer analogy could be drawn with decision-

[1] Sometimes also known as "idea management," "social ideation," "idea contest," or "idea competition" systems.

support systems for idea-generation tasks (Dennis & Valacich, 1993), as we shall see below, supporting innovation at crowd scales introduces daunting, qualitatively new challenges as compared to supporting innovation at the group level.

Open innovation systems have proved to be an emerging collective intelligence success story. Many open innovation platforms (such as Spigit, Imaginitik, Nosco, BrightIdea, Salesforce, and Ideascale) have emerged and have been used widely in domains ranging from government (the US White House, the UK National Health Service, the Danish central government) to industry (Intel, IBM, Dell, Xerox, Cisco, and P&G) and beyond. These organizations have found that crowds can contribute large volumes of ideas for problems they care about, often uncovering solutions superior to those developed in-house (Lakhani & Jeppesen, 2007). A recent survey found that one in four companies plan to utilize open innovation systems within the next twelve months, and this figure is growing (Thompson, 2013). The six-day IBM "Innovation Jam" in 2006, for example, involved over 150,000 participants from 104 countries in identifying 46,000 product ideas for the company (Bjelland & Wood, 2008). Dell's ongoing Ideastorm website (Di Gangi & Wasko, 2009) has received, to date, over 20,000 suggestions for improved Dell products and services. In the early weeks of his first term, President Obama asked citizens to submit questions on his web site change.gov, and promised to answer the top five questions in each category in a major press conference (Phillips, undated). Over 70,000 questions were submitted. Google's 10 to the 100th project received over 150,000 suggestions on how to channel Google's charitable contributions (Buskirk, 2010), while the 8,000 participants in the 2009 Singapore Thinkathon generated 454,000 ideas (Butterworth, 2005).

Such large-scale participation enables, in turn, such powerful emergent phenomena as:

- *The long tail*: crowds can generate a much *greater diversity* of ideas, including potentially groundbreaking "out of the box" contributions, than we could easily access otherwise. One study found, for example, that 30 percent of the challenges that confounded experienced corporate researchers were solved by participants in an open innovation system that came from outside of the customer organization and generally did not have "expert" credentials in the problem domain (Lakhani & Jeppesen, 2007). In general, the larger the crowd, the greater the probability of encountering truly superior out of the box solutions (Lakhani & Jeppesen, 2007).
- *Many eyes*: crowd participants can check and correct each other's contributions, enabling remarkably high quality results very inexpensively (Raymond, 1999).
- *Wisdom of the crowd*: crowds can collectively make better judgments than the individuals that make them up, often exceeding the performance of experts (Surowiecki, 2005). The reason for this is that the *errors* made by different crowd members tend, under the correct conditions, to be independent and therefore to cancel each other out, while the shared "signal" is strengthened by aggregating over the many contributions.

All is not rosy, however. Open innovation systems face important challenges deriving, ironically, from their very success: they can elicit such high levels of participation that it becomes very difficult to guide the crowd in productive ways, and pick out the best of what they have created. This results in such serious problems as low signal-to-noise ratios, insular ideation, non-comprehensive coverage, poor evaluation, and poor idea filtering. In the following sections this article reviews these key challenges in detail, and then proposes a few promising strategies (including semi-formalized ideation systems, careful micro-task design, and attention mediation) that the research community can use to address these challenges and inform the design of more effective open innovation systems.

KEY CHALLENGES

We can organize the key challenges facing open innovation systems according to which part of the innovation process they apply to (Figure 1).

Solution Generation

The solution generation phase of open innovation is prone to critical challenges that include:

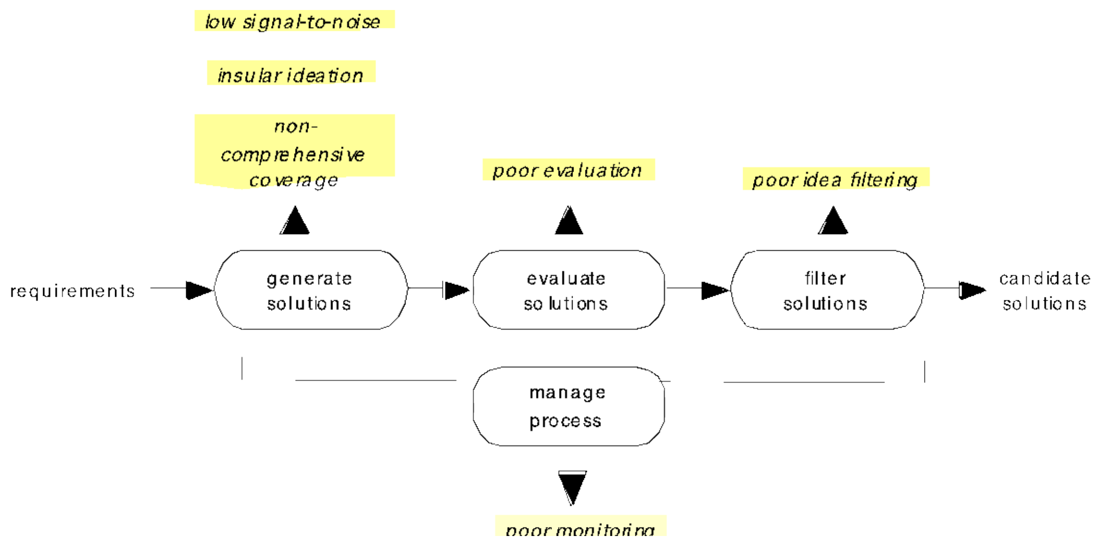

Figurehttp://conference.uxpaboston.org/ 1. Challenges (highlighted) for the different parts of the open innovation process.

- *Low signal-to-noise ratios:* Open innovation engagements tend to generate idea corpuses that are highly redundant, and of highly variable quality (Riedl, Blohm, Leimeister, & Krcmar, 2010; Schulze, Indulska, Geiger, & Korthaus, 2012; Westerski, Dalamagas, & Iglesias, 2013; Blohm, Bretschneider, Leimeister, & Krcmar, 2011a; Bjelland & Wood, 2008; Di Gangi & Wasko, 2009; B.P. Bailey & Horvitz, 2010). Previous research suggests that only about 10-30% of the ideas from open innovation engagements are considered, by the customers, as being of high quality (Blohm et al., 2011a). Convening a group of experts to identify the best ideas from these corpuses can be prohibitively expensive and time-consuming. Nearly 100 IBM senior executives, for example, spent weeks sifting through the tens of thousands of postings generated by their Web Jam (Bjelland & Wood, 2008). Google had to recruit 3000 Google employees to filter the unexpected deluge of ideas for the 10 to the 100th project, a process that put them nine months behind schedule. The change.gov website, finally, had to be shut down prematurely because the huge volume of contributions overwhelmed the staff's ability to meaningfully process it.
- *Insular ideation:* Ideas in open innovation engagements are typically generated quickly by single individuals without reference to other submitted ideas (Bjelland & Wood, 2008). It seems clear, however, that the customers of open innovation engagements would be better served by a smaller number of more deeply considered ideas that have benefited from critiques and refinements by multiple contributors (Blohm et al., 2011a). While some existing innovation systems (Shum, 2008) allow users to easily cite existing ideas when describing new ones, this system feature is not sufficient by itself to address the broader challenge of insular ideation. When there are thousands (or tens of thousands) of ideas in a corpus, it can be very difficult for a contributor to find the existing ideas that best combine with, extend, or inspire her/his own, even with recommender systems.

- *Non-comprehensive coverage:* Open innovation systems have no inherent mechanism for ensuring that the ideas submitted comprehensively cover the most critical facets of the problem at hand, so the coverage is hit-or-miss and may not align with the customer's needs. The space of possible solutions is generally not specified up front, and there is no easy way for potential contributors to see which problem facets remain under-covered. As a result, open innovation systems often produce very spotty coverage. If we ask a crowd, for example, to propose ideas for new laptop designs, they may offer many ideas on displays and processor features, but pay little attention to battery and power management issues, despite the fact that the latter are also critical in laptop design. Open innovation systems based on competitions (Morgan & Wang, 2010) are more likely to include comprehensive entries because they typically include very specific problem statements and provide an incentive for superior entries. This approach is not a panacea, however, because an opportunity is lost: the wisdom of the crowd is not exploited to help ensure that the problem statements are comprehensive.

Solution Evaluation

Solution evaluation in current open innovation systems is also prone to several important weaknesses. Participants often rate ideas with respect to their own, rather than the customer's criteria (Jouret, 2009), little support is provided to help raters correct and build upon each other's facts and reasoning, and analytic tools are rarely provided to bolster the evaluation process with objective assessments of the proposed ideas. In cases where participants have a stake in the outcome, it is often possible for them to "game" the system. A single idea post with coordinated voting, for example, may beat out a better idea that had its votes spread over many redundant instantiations.

Idea Filtering

Many techniques have been proposed to help filter out the most promising ideas from the large corpuses often generated by open innovation systems, but these approaches all currently face important limitations. To understand this better, we will systematically review the space of existing idea filtering approaches. These approaches include (Figure 2):

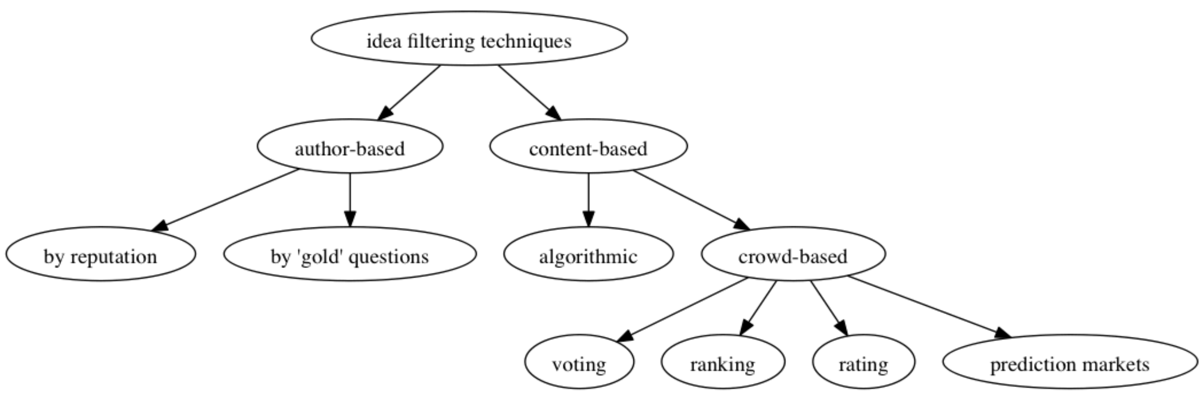

Figure 2. A taxonomy of idea filtering techniques.

- *Author-based* filtering filters ideas based on *who* contributed them. Authors can be excluded, for example, based on their reputation (Kittur et al., 2013) or their responses to test questions (Oleson et al., 2011). Gathering a robust contributor profile can be difficult, however, and often the best ideas can come from unexpected quarters, so filtering by author can have a substantial error rate.
- *Content-based* filtering filters ideas based on their content, rather than their author. Software *algorithms* can be used, for example, to estimate the creativity of a contribution based on the presence of rarely-used words in its description (Walter & Back, 2013). (Westerski et al., 2013) derive idea quality metrics based on manually- as well as machine-generated idea annotations (e.g., concerning what triggered the idea). Such techniques are fundamentally limited, however, by the fact that current natural language processing algorithms have only a shallow understanding of natural language, and thus can be easily fooled. In the Westerski work, for example, the automatically-generated idea quality metrics only achieved a correlation of 0.1 with the quality scores given by human experts.
- *Crowd-based* filtering attempts to transcend these difficulties by asking members of a crowd to identify superior ideas, since humans can potentially understand the content much more deeply than software algorithms. It has indeed been shown that crowds, under the right circumstances, can solve classification problems like that with accuracy equal to or even better than that of experts (Surowiecki, 2005). This can be done in many ways, including voting, rating (Likert, 1932; Salminen & Harmaakorpi, 2012), ranking (which ask participants to provide *relative* rankings of idea pairs) ; Miller, Hemmer, Steyvers, & Lee, 2009; Baez & Convertino, 2012; Saaty, 1990), and prediction markets

(where users buy and sell stocks representing predicted winners, knowing they will receive a payoff if they own stocks that are eventually selected as winners) (Bothos, Apostolou, & Mentzas, 2009; Slamka, Jank, & Skiera, 2012; Blohm, Riedl, Leimeister, & Krcmar, 2011b; Berg & Rietz, 2003; Wolfers & Zitzewitz, 2004; Soukhoroukova, Spann, & Skiera, 2012; Dahan, Kim, Lo, Poggio, & Chan, 2011).

All these approaches, however, face serious difficulties in terms of accuracy and the time demands they place on crowd members, especially when applied to large sets of alternatives. Voting systems, for example, are provably unable to satisfy reasonable outcome criteria (Arrow, 1963) (Kostakos, 2009). Rating systems tend to do a poor job of distinguishing between good and excellent ideas (Bao, Sakamoto, & Nickerson, 2011) and can lock into arbitrary rankings (Salganik, Dodds, & Watts, 2006; Riedl, Blohm, Leimeister, & Krcmar, 2013; Bjelland & Wood, 2008; Yang, 2012). Ranking systems require a number of comparisons that grows exponentially with the size of the idea set. Idea filtering accuracy can be improved by asking participants to evaluate ideas with respect to multiple customer-defined criteria (Dean, Hender, Rodgers, & Santanen, 2006; Riedl et al., 2010; Riedl et al., 2013) but at the cost of increasing the time and cognitive complexity demands for the raters and thus potentially reducing participation. These approaches also face a *criteria alignment* problem: raters often assess ideas based on their personal criteria, or even self-interest, rather than the interests of the open innovation customer (Spears, LaComb, Interrante, Barnett, & Senturk-Dogonaksoy, 2009; Newell, Rakow, Weston, & Shanks, 2004; Forsythe, Rietz, & Ross, 1999). Prediction markets attempt to address this issue by providing incentives for raters to use the same criteria as decision-makers, but require traders to participate in cognitively complex and time-consuming tasks (Blohm et al., 2011b; Bothos, Apostolou, & Mentzas, 2012), are prone to manipulation (Forsythe et al., 1999; Hanson, 2004; Wolfers & Zitzewitz, 2004; Wolfers & Leigh, 2002), face substantial scalability challenges, and often generate too little trading activity to get meaningful stock prices, especially since the benefits to the traders of getting the correct portfolio are often too nominal to merit a substantial ongoing time investment (Hanson, 2003; Healy, Linardi, Lowery, & Ledyard, 2010). Studies have in fact reported low correlations between idea market share prices and rankings by human experts, e.g. 0.33 in (Blohm et al., 2011b), 0.43 in (LaComb, Barnett, & Pan, 2007) and 0.10, 0.39 and 0.46 in (Soukhoroukova et al., 2012).

Process Management

Existing open innovation systems solicit ideas in natural language form that is difficult for computer algorithms to summarize accurately. It can, as a result, be very difficult for the customers of these systems to assess how well their innovation engagement is going (e.g., in terms of how complete the idea coverage is for different challenges) and whether or not some areas merit additional attention from the crowd. There are typically simply too many ideas for it to be practical for the customer who initiated the innovation process to read and summarize on an ongoing basis.

PROMISING DIRECTIONS

Open innovation systems, as we can see, face a critical dilemma. Finding truly superior out-of-the-box ideas requires eliciting contributions from large crowds, the larger the better; but the process of getting the best ideas from such crowds is, at present, expensive and error-prone. Addressing this problem will, we believe, require shifting much of the burden of organizing and filtering the crowd's open innovation contributions from the customer to some combination of software and the crowd itself. We propose three possible strategies for achieving this goal:

- *Partial formalization* of the crowd's contributions, supported by appropriate *incentives*
- Careful *micro-task* design
- *Attention mediation* to help optimize the crowd's contributions

These strategies, which can be adopted independently or in combination, are described in the sections below.

Semi-Formalization

One promising strategy is for the crowd to make their contributions in a semi-formalized form, i.e. one that is structured enough to substantively increase the level of algorithmic support that can be provided for understanding and filtering the idea corpus, without making unreasonable demands on the crowd who, after all, generally will not be experts at formal knowledge representations. We believe that deliberation maps, also known as argument maps (Conklin, 2005; Klein &

Iandoli, 2008; Kirschner, Shum, & Carr, 2003; De Liddo, Sándor, & Shum, 2012) represent a promising approach. The concept is simple. Participants are asked to organize their contributions as tree structures made up of issues (challenges to be addressed), ideas (possible solutions for a challenge), and arguments (statements that support or contradict an idea or other argument). Issues can have criteria attached to them that describe the attributes of a desirable solution, and arguments can reference these criteria. The issue "what design for a water pump?,"[2] for example, can specify that a good design will satisfy the criterion "safe to operate," and the con "exhaust is dangerous ..." can note that a gasoline-powered pump violates that criterion (see Figure 3, below).

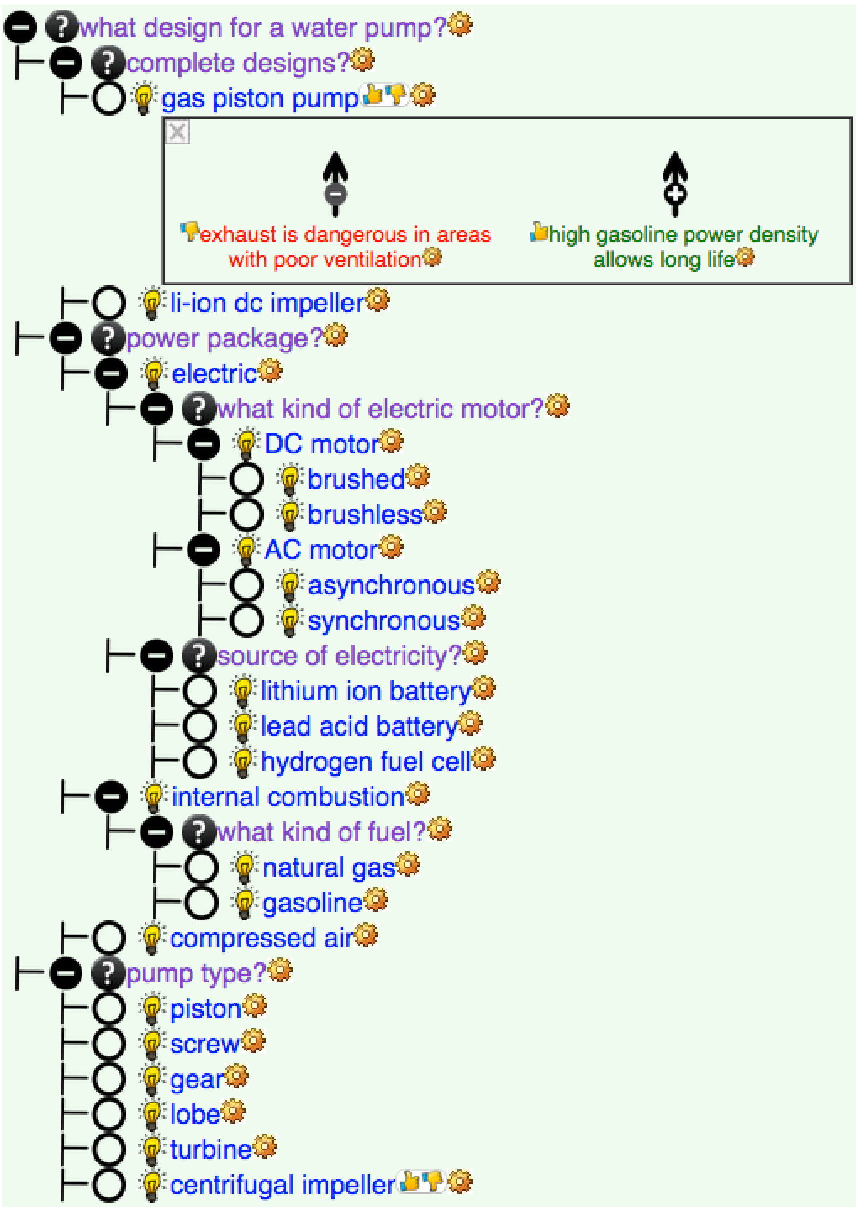

[2] NB: We selected this example for simplicity, but crowd-scale innovation will typically address more open-ended and multi-disciplinary topics such as, for example, designing public policies to respond to the problem of climate change.

Figure 3. A deliberation map applied to an engagement on designing a water pump.

In this kind of structure, all content appears under the posts it logically refers to. This makes it easy to find what has and has not been said on any topic, without having to first perform a time-consuming analysis of the idea corpus. This can be especially important in new or rapidly changing communities that do not have substantive pre-existing “communal common ground” (Clark, 1996). The compact, well-organized map structure also makes it much easier to avoid adding redundant ideas (thereby radically improving the signal-to-noise ratio), as well as to ensure more systematic coverage (because gaps, such as issues without ideas, or ideas without arguments, are easy to spot). Careful critical thinking is encouraged, because users are implicitly encouraged to express the evidence and logic in favor of the options they prefer (Carr, 2003), and the community can rate each element of their arguments piece-by-piece. The customer who initiates the innovation process can pre-populate the argument map with issues ideas and arguments to serve as a guide for the crowd’s activity. Participants, finally, can collaboratively refine proposed solutions. One user can, for example, propose an idea, a second can raise an issue concerning how some aspect of that idea can be implemented, and a third propose possible solutions for that issue. The result is that, rather than getting a long list of relatively shallow ideas, the ideation customer is more likely to receive a shorter list of more deeply considered, collaboratively refined ideas.

Another key advantage of deliberation maps is that they allow crowds to systematically create new ideas by recombining and refining existing ideas (Bernstein, Klein, & Malone, 2003). In Figure 3 above, for example, there are two top-level sub-issues: (1) how to power the pump, and (2) what type of pump will be used. Since these issues represent largely orthogonal design dimensions, that means that you can posit a unique pump design for every combination of power packages and pump types. Each pump design will therefore be made up of a “bundle” of other ideas representing different pump components. Since some bundles will suit the customer’s needs better than others, and the open innovation participants can accordingly attach arguments to the idea bundles concerning which combinations they think are best. This ability to frame design as the recombination of existing ideas is powerful in several ways. For one, it is widely recognized that innovations are often the result of combining previously disparate ideas, and many “creativity enhancement” tools have been created to systematize this process (Yu & Nickerson, 2013). A deliberation map extends this approach into crowd-scale open innovation engagements. Credit assignment remains straightforward, because each proposal is built from components with clear authorship. Another key point is that the value of a deliberation map can extend far beyond the innovation engagement it was initially generated for, because it represents an entire *space* of possible solutions that can be readily harvested, refined and re-combined by other communities facing similar problems.

We can take this idea a step farther by extending open innovation platforms with easy-to-use design tools that allow contributors to express their ideas as simple models, rather than just as natural language text. Imagine, for example, if an open innovation engagement about the design of a city park allowed people to create candidate park designs using a simple drag-and-drop interface. Computer-Aided Design (CAD) systems intended for widespread use by non-experts are becoming increasingly available, ranging for example from Google SketchUp (http://www.sketchup.com/) to the builder tools in video games, to the protein chain design capabilities in Foldit (http://fold.it/). This envisioned capability of expressing ideas in terms of simple models can makes it easier for software tools to support the users via algorithms in finding similar ideas (thus enabling better map organization and avoiding redundancy) and even analyzing these ideas to evaluate how well they meet the customer's criteria. Such design tools can also implicitly orient user contributions towards customer objectives; that is the crowd would be more likely to contribute the type of ideas that the design tools make easier to express.

An important consideration for this kind of approach concerns *incentives*. A semi-formal open innovation system asks more of the crowd than existing approaches: they are also required to adhere a systematic structure and assess how their ideas relate to already-existing content. Will the crowd participate as actively with semi-formal systems as they do with existing approaches? We argue that the answer is "yes,” when the size of the crowd grows large enough. This claim is motivated by the empirical and theoretical reasons that we describe below.

Our empirical evidence stems from our experiences, since 2007, with applying an implemented semi-formal open innovation system known as the Deliberatorium (Klein & Iandoli, 2008) to real-world problems and user communities.

The first evaluation was performed at the University of Naples, where 220 masters students in the information engineering program posted over 3000 issues, ideas and arguments on the topic of bio-fuel use in Italy (Klein & Iandoli, 2008). This pattern of high participation held true for later evaluations including, most recently, one where 350 randomly-selected members of the Italian Democratic Party posted hundreds of issues ideas and arguments concerning electoral law reform, contributing *just as actively* as a demographically-matched group of 350 users who used a conventional web-forum type approach.

We have since developed a better understanding of why we were able to achieve such high participation with a semi-formal system. It has been found (Hars & Ou, 2002; Lakhani & Wolf, 2005; Roberts, Hann, & Slaughter, 2006; Bonaccorsi & Rossi, 2004) that contributors to collective intelligence systems are incentivized by factors that range from self-interest (e.g. making money, getting useful information, making valuable connections) to selflessness (contributing to a problem or community they care about). Active participation tends to occur when the benefits, for individual contributors, substantially exceed their costs (Benkler, 2006). With small numbers of participants, an informal approach clearly minimizes participation costs and thereby allows high participation levels. But this picture changes *as the scale of the crowd grows*. Our previous work, based on simulations of the open innovation process, suggests that the *benefits* of argument mapping (in terms of improved signal-to-noise ratio and co-location of related content) grows linearly with the size of the community, while the *cost* to participants of argument mapping grows only logarithmically with community size. We can thus expect that, as the scale of the discussion grows, that the cost/benefit ratio for participants will increase as well, so users will increasingly recognize the opportunity to “be a hero” to the community by creating an easy-to-harvest deliberation map (Klein, 2012). While we have not quantified the precise value for this cross-over point, it seems clear that many open innovation engagements have already crossed into the realm where a semi-formal approach would be perceived, by the participants, as representing an attractive tradeoff of cost and benefits compared to unstructured systems.

Micro-Tasks

The success of current collective intelligence systems has relied upon creating contexts where large-scale goals can be achieved by asking the crowd to perform many, individually small, (“micro”) tasks (Howe, 2006; Shirky, 2009). Wikipedia articles, for example, are typically built from many small contributions where each user may individually just add a few sentences, correct spelling, check a fact, or add a reference (Viegas, Wattenberg, & Dave, 2004). When tasks are small, the “entry barrier” is low and as a result more people are likely to participate. It also helps to have a diverse suite of micro-tasks, thereby broadening the portion of the crowd that is willing and able to contribute.

Micro-tasks can play an important role in open innovation systems in many ways. They can, for example, enable scalable *moderation* in open innovation engagements. While semi-formal structures offer clear advantages over conventional open innovation systems, they do raise an important challenge: how can we ensure that “messy” crowds, made up largely of individuals who are not experts in knowledge formalization, will create well-structured deliberation maps? One alternative is to engage experts in deliberation mapping to ensure that author contributions to the map are correctly structured and non-redundant. This approach has been applied successfully to moderate-size (~200 author) open innovation engagements at relatively low cost (Klein, Spada, & Calabretta, 2012) but does, however, require moderators with a relatively specialized skillset, limiting the scalability of this approach. The question then becomes: can we replace an expert moderator with non-expert crowd members that each perform relatively simple moderation micro-tasks? It has been demonstrated many times that crowds with even modest skills can, in the aggregate, perform challenging judgment tasks as well as or even better than the experts that have traditionally done them (Surowiecki, 2005) (Smith, Lynn, & Lintott, 2013). We can, accordingly, imagine implementing the moderation process that way as well. The moderation micro-tasks could include, for example, *unbundling* contributions (i.e. breaking them down into individual issues, ideas, and arguments), *tagging* contributions with keywords, proposing *revisions* to clarify contributions (i.e., as experienced contributors do in Wikipedia or StackOverflow), *finding* the right place to place each element, and checking whether that element is *redundant* or not. Crowd programming techniques such as majority-vote and find-fix-verify can be used to ensure high-quality moderation even when crowd members have highly variable skill levels (Bernstein et al., 2010).

Better micro-tasking can also play an important role in improved crowd-based *idea filtering*. As noted above, current idea filtering techniques can suffer from poor filtering accuracy, make unrealistic demands on crowd members, or both. One reason for this, we argue, is that existing mechanisms do not fully break idea filtering into micro-tasks well suited to diverse crowds. Idea filtering generally requires evaluating ideas with respect to *multiple* customer criteria. Requiring raters to evaluate each idea with respect to every criterion, however, increases the cognitive and time demands on each rater, thereby potentially reducing participation, and may reduce accuracy when raters are forced to evaluate ideas with respect to criteria they are not expert on. One example of a more finely micro-tasked approach is discussed in (Klein & Garcia, 2014), where raters were incentivized financially to accurately allocate a limited budget of tokens ("lemons") to inferior ideas. This approach achieved substantially better speed and filtering accuracy than conventional rating, because users could limit themselves to the micro-tasks they knew well: they only had to find the ideas that they were confident failed with respect to at least one criterion.

Attention Mediation

A final important capability for effective large-scale open innovation, we believe, is attention mediation. In engagements with hundreds or thousands of participants and contributions, it becomes very difficult for users to get a "big picture" of what is going on, e.g., to find the issues, ideas, and arguments that they can best contribute to, as well as to assess whether an engagement is progressing well, where problems are occurring and should be addressed, and when the results are mature and ready to "harvest." Without this big picture, however, we run the risk of severely under-utilizing the collective intelligence of the crowd.

Addressing this need will require, we believe, developing analytics that operate as diagramed below in Figure 4:

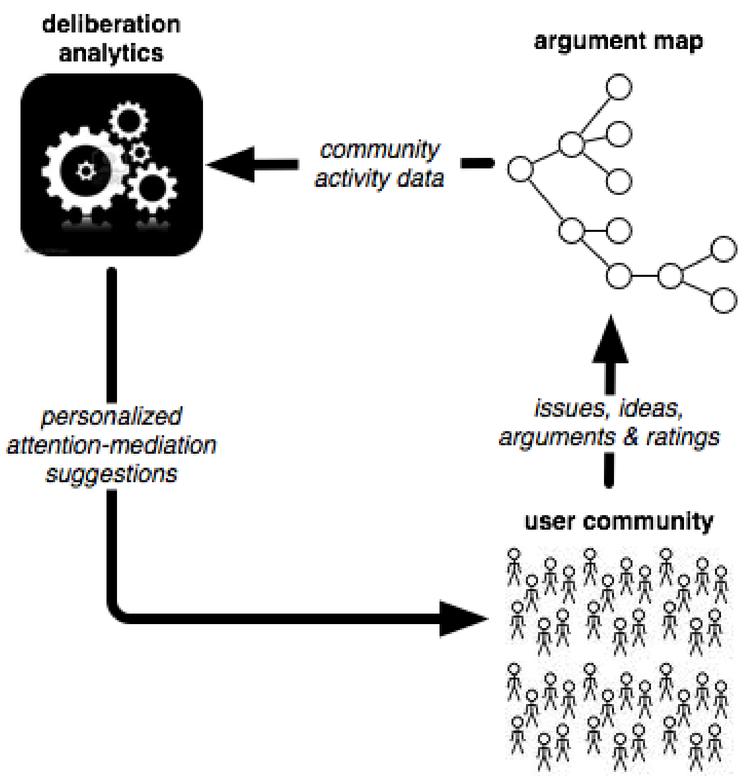

Figure 4. Using analytics to mediate attention in large-scale open innovation.

In this approach, algorithms data-mine the crowd's activity in an open innovation engagement to develop a model of how well the process is going, as well as to offer participants personalized suggestions concerning where, in the deliberation map, they can make the most helpful contributions. In this way, the crowd can enable the creation of a complete, high-quality, set of ideas and associated evaluations. Such personalized attention-mediation suggestions can work in many ways. We list just a few examples below:

- *Guidance for idea and argument generation*: analytics can be used to guide participants to the parts of the deliberation map that, based on their previous activity, they can best contribute to. Users, for example, will probably tend to have interest and expertise in new contributions that are close, in the deliberation map, to posts that they (or people with similar interests) have already attended to in the past. I might be interested, for example, in knowing when someone has attached a "con" to one of my posts, or when a competitor to my idea is garnering positive ratings and supporting arguments. Attention mediation algorithms could also be used to suggest fruitful directions for idea re-combination, for example by pointing authors to idea components that have been part of successful hybrids in the past. As previous work on crowdsourcing has shown, the attention mediation algorithms, like any other optimization algorithm, will need to make careful tradeoffs between exploration (i.e. encouraging the crowd to generate new out-of-the-box ideas in parallel) and exploitation (encouraging the crowd to iteratively refine existing ideas to maximize their value) (Little, 2011; Dow et al., 2011). This approach can be used to create a kind of evolutionary optimization process, where the crowd's attention is drawn to refine and recombine the currently most successful ideas into new ones (Yu & Nickerson, 2013). Given that the number of possible idea hybrids grows exponentially with the number of available building blocks, computational support for optimizing search in such large spaces can be critical.
- *Guidance for idea filtering*: Most open innovation systems allow users to select which ideas they do or do not rate, which can result in gaps and inefficiencies (e.g. if users tend to rate ideas that already have many ratings) as well as rating dysfunctions (if users tend only to rate ideas that already have high average ratings) (Salganik et al., 2006). A promising direction is to develop algorithms, based on information-theoretic concepts, that guide participants towards doing the ratings that help identify the best ideas as quickly as possible (Salganik & Levy., 2012; Toubia & Florès, 2007).
- *Guidance for management*: analytics can help reveal emergent patterns in open innovation engagements that managers may want to intervene in to maximize the value of the engagement to the customer. Analytics can, for example, reveal when promising ideas are being neglected by the crowd, when the crowd tends to ignore or reflexively down-rate competing ideas, when an issue or idea is becoming increasingly controversial and polarized, when users are rating posts without reading, or accounting for, relevant arguments, and so on (Klein, 2012). When these situations are detected, open innovation managers can intervene; e.g., by closing certain topics, implementing incentives to focus on certain branches of the deliberation map, encouraging "new faces" to join the discussion on a given issue, and so on.

The authors of this article have begun exploring how to create effective attention mediation algorithms for large-scale open innovation (Klein, 2012). We have been using a process called process-commitment-violation analysis (Klein & Dellarocas, 2003) to systematically identify what properties characterize an ideal open innovation engagement, have identified metrics for assessing how well an open innovation engagement approaches this ideal, and have identified techniques that can help address the shortcomings these metrics reveal. Our experience to date suggests that deliberation maps enable a whole spectrum of analytics in support of attention mediation that would not be possible with conventional, unstructured innovation tools.

CONCLUSIONS AND FUTURE WORK

This article has argued that current open innovation technologies face a range of shortcomings that seriously limit their effectiveness for engagements with complex issues and large numbers of participants. We have also proposed three potentially powerful strategies for avoiding these shortcomings, including:

- *Semi-formalized* structures (such as deliberation maps) to enable more powerful computational services
- *Micro-tasks* that match the crowd's motivations and capabilities
- *Attention mediation* to guide participants to the tasks where they can do the most good

While elements of these strategies have been broached in other settings, this article is the first, to the best of our knowledge, to explore in detail how they can be applied to the important challenge of enabling effective innovation at *crowd* scales. Semi-formalized structures such as argument maps have been used, in the past, to support only individuals or small teams (Klein, 2007), but not crowds. Micro-task design and attention mediation have been important in creating crowd-scale social-computing applications, but have not been applied to the kinds of *semi-formalized* interaction structures that we consider crucial for transcending the serious limitations in current open innovation technology (Klein, 2012).

All of these strategies, however, raise important, interesting and tightly intertwined research questions. For example:

- What kinds of semi-formalized structures represent the best tradeoff between ease-of-use (for humans) and formal expressiveness (for computers)?
- What is the impact of the properties of the problem (e.g., complexity, amount or prior knowledge) or crowd (e.g., size, diversity)?
- To what extent can a crowd use micro-tasks to formalize knowledge, which is typically accomplished by a trained expert?
- What combination of interaction structures, micro-tasks and incentives can best harness the crowd's collective intelligence to develop deeply considered and collaboratively developed solutions, as opposed to a long list of relatively shallow and redundant single-user ideas?
- What kinds of attention mediation work best if we wish to maximize the crowd's creativity (exploring new out-of-the-box solutions) versus the quality of the solutions (verifying and refining the solutions that have already been identified)?
- Can insights derived from developing open innovation systems be used for other large-scale formalization tasks, such as populating the semantic web (Berners-Lee, Hendler, & Lassila, 2001) or developing knowledge bases to support artificial intelligence systems?
- Which attributes of the open innovation system and process impact speed, cost, and quality of the outcome, and how can they best be detected and purposefully engineered?

The authors invite the research community to explore these and many other compelling questions with us as we engage in developing tool and approaches for the next generation of open innovation systems.

ACKNOWLEDGEMENTS

This work was supported, in part, by funding from the European Union's Seventh Framework Program under grant agreement no. 6611188 - the CATALYST project.

REFERENCES

1. Arrow, K. J. (1963). *Social choice and individual values.* New York, NY, USA: Wiley.
2. B.P. Bailey, Horvitz, E. (2010). What's Your Idea? A Case Study of a Grassroots Innovation Pipeline within a Large Software Company. *Proceedings of Computer-Human Interaction.*
3. Baez, M., & Convertino, G. (2012). Designing a facilitator's cockpit for an idea management system. *Proceedings of ACM 2012 conference on Computer Supported Cooperative Work Companion*.
4. Bao, J., Sakamoto, Y., & Nickerson, J. V. (2011). Evaluating design solutions using crowds. *Proceedings of Seventeenth Americas Conference on Information Systems*. Detroit, Michigan, USA.
5. Benkler, Y. (2006). *The Wealth of Networks: How Social Production Transforms Markets and Freedom.* New Haven, CT USA: Yale University Press.

6. Berg, J. E., & Rietz, T. A. (2003). Prediction markets as decision support systems. *Information Systems Frontiers, 5*(1), 79-93.

7. Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. *Scientific American*, May 17, 2001.

8. Bernstein, A., Klein, M., & Malone, T. W. (2003). The Process Recombinator: A Tool for Generating New Business Process Ideas. In T. W. Malone, K. Crowston, & G. A. Herman (Eds.), *Organizing Business Knowledge: The MIT Process Handbook.* Cambridge MA USA: MIT Press.

9. Bernstein, M., Little, G., Miller, R. C., Hartmann, B., Ackerman, M., Karger, D. R. et al. (2010). Soylent: A Word Processor with a Crowd Inside. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2010).*

10. Bjelland, O. M., & Wood, R. C. (2008). An Inside View of IBM's 'Innovation Jam'. *Sloan Management Review*, *50*(1).

11. Blohm, I., Bretschneider, U., Leimeister, J. M., & Krcmar, H. (2011a). Does collaboration among participants lead to better ideas in IT-based idea competitions? An empirical investigation. *International Journal of Networking and Virtual Organisations, 9*(2), 106-122.

12. Blohm, I., Riedl, C., Leimeister, J. M., & Krcmar, H. (2011b). Idea Evaluation Mechanisms for Collective Intelligence in Open Innovation Communities: Do Traders outperform Raters? *Proceedings from ICIS (International Conference on Information Systems).*

13. Bonaccorsi, A., & Rossi, C. (2004). Altruistic individuals, selfish firms? The structure of motivation in Open Source software. *First Monday*, *9*(1).

14. Bothos, E., Apostolou, D., & Mentzas, G. (2009). IDEM: A Prediction Market for Idea Management. In *Designing E-Business Systems. Markets, Services, and Networks*, 1-13, Springer.

15. Bothos, E., Apostolou, D., & Mentzas, G. (2012). Collective intelligence with web-based information aggregation markets: The role of market facilitation in idea management. *Expert Systems with Applications, 39*(1), 1333-1345.

16. Buskirk, E. V. (2010). Google Struggles to Give Away $10 Million. *Wired Magazine.* June 28, 2010, http://www.wired.com/business/2010/06/google-struggles-to-give-away-10-million/all/.

17. Butterworth, M. S. J. (2005). *Using mass idea brainstorming as an organizational approach to jumpstarting innovation initiative.* Unpublished PhD thesis, University of South Australia.

18. Carr, C. S. (2003). Using computer supported argument visualization to teach legal argumentation. In P. A. Kirschner, S. J. B. Shum, & C. S. Carr (Eds.), *Visualizing argumentation: software tools for collaborative and educational sense-making*, 75-96. Springer-Verlag.

19. Chesbrough, H., Vanhaverbeke, W., & West, J. (2008). *Open Innovation: Researching a New Paradigm.* Oxford: Oxford University Press.

20. Chesbrough, H. W. (2003). *Open innovation: The new imperative for creating and profiting from technology.* Cambridge, MA, USA: Harvard Business Press.

21. Clark, H. H. (1996). *Using language.* Cambridge, MA, USA: Cambridge University Press.

22. Conklin, J. (2005). *Dialogue Mapping: Building Shared Understanding of Wicked Problems.* John Wiley and Sons, Ltd.

23. Dahan, E., Kim, A. J., Lo, A. W., Poggio, T., & Chan, N. (2011). Securities trading of concepts (STOC). *Journal of Marketing Research*, *48*(3), 497-517.

24. De Liddo, A., Sándor, Á., & Shum, S. B. (2012). Contested collective intelligence: Rationale, technologies, and a human-machine annotation study. *Computer Supported Cooperative Work (CSCW), 21*(4-5), 417-448.

25. Dean, D. L., Hender, J. M., Rodgers, T. L., & Santanen, E. L. (2006). Identifying quality, novel, and creative ideas: Constructs and scales for idea evaluation. *Journal of the Association for Information Systems*, *7*(1), 30.

26. Dennis, A. R., & Valacich, J. S. (1993). Computer brainstorms: More heads are better than one. *Journal of Applied Psychology*, *78*(4), 531.

27. Di Gangi, P. M., & Wasko, M. (2009). Steal my idea! Organizational adoption of user innovations from a user innovation community: A case study of Dell IdeaStorm. *Decision Support Systems*, *48*(1), 303-312.

28. Dow, S., Kulkarni, A., Bunge, B., Nguyen, T., Klemmer, S., & Hartmann, B. (2011). Shepherding the crowd: managing and providing feedback to crowd workers. *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, 1669-1674.

29. Enkel, E., Perez-Freije, J., & Gassmann, O. (2005). Minimizing market risks through customer integration in new product development: Learning from bad practice. *Creativity and Innovation Management*, *14*(4), 425-437.

30. Forsythe, R., Rietz, T. A., & Ross, T. W. (1999). Wishes, expectations and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior & Organization*, *39*(1), 83-110.

31. Gassmann, O. (2006). Opening up the innovation process: Towards an agenda. *R&D Management*, *36*(3), 223-228.

32. Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers*, *5*(1), 107-119.

33. Hanson, R. (2004). Foul play in information markets. In R.W. Hahn & P. C. Tetlock, *Information Markets: A New Way of Making Decisions*. Washington DC, USA: AEI Press, 126-141.

34. Hars, A., & Ou, S. (2002). Working for Free? Motivations for Participating in Open-Source Projects. *International Journal of Electronic Commerce*, *6*(3), 25 - 39.

35. Healy, P. J., Linardi, S., Lowery, J. R., & Ledyard, J. O. (2010). Prediction markets: Alternative mechanisms for complex environments with few traders. *Management Science*, *56*(11), 1977-1996.

36. Howe, J. (2006). The rise of crowdsourcing. *Wired magazine*, *14*(6), 1-4.

37. Hrastinski, S., Kviselius, N. Z., Ozan, H., & Edenius, M. (2010). A review of technologies for open innovation: Characteristics and future trends. *Proceedings of 2010 43rd Hawaii International Conference on System Sciences (HICSS)*, 1-10.

38. Jouret, G. (2009). Inside Cisco's search for the Next Big Idea. *Harvard Business Review*, *87*(9), 43-45.

39. Kirschner, P. A., Shum, S. J. B., & Carr, C. S. (2003). *Visualizing Argumentation: Software tools for collaborative and educational sense-making.* Springer.

40. Kittur, A., Nickerson, J. V., Bernstein, M., Gerber, E., Shaw, A., Zimmerman, J. et al. (2013). The future of crowd work. *Proceedings from Computer Supported Cooperative Work.*

41. Klein, M. (2007). Achieving Collective Intelligence via Large-Scale On-line Argumentation. *Second International Conference on Internet and Web Applications and Services (ICIW'07).*

42. Klein, M., & Dellarocas, C. (2003). Designing Robust Business Processes. In T. W. Malone, K. Crowston, & G. A. Herman (Eds.), *Organizing Business Knowledge: The MIT Process Handbook.* MIT Press. Cambridge, MA, USA.

43. Klein, M., & Iandoli, L. (2008). Supporting Collaborative Deliberation Using a Large-Scale Argumentation System: The MIT Collaboratorium. *Directions and Implications of Advanced Computing; Conference on Online Deliberation (DIAC-2008/OD2008)*.

44. Kostakos, V. (2009). Is the crowd's wisdom biased? A quantitative analysis of three online communities. *Proceedings from Computational Science and Engineering 2009.*

45. LaComb, C. A., Barnett, J. A., & Pan, Q. (2007). The imagination market. *Information Systems Frontiers*, *9*(2-3), 245-256.

46. Lakhani, K. R., & Jeppesen, L. B. (2007). R&D: Getting unusual suspects to solve R&D puzzles. *Harvard Business Review*, *85*(5), 30-32.

47. Lakhani, K. R., & Wolf, R. G. (2005). Why hackers do what they do: Understanding motivation and effort in free/open source software projects. In J. Feller, B. Fitzgerald, S. Hissam, & K. R. Lakhani (Eds.), *Perspectives on Free and Open Source Software.* Cambridge, MA, USA: MIT Press.

48. Lakhani, K. R., & Panetta, J. A. (2007). The principles of distributed innovation. *innovations*.

49. Likert, R. (1932). A technique for the measurement of attitudes. *Archives of Psychology*, *140*, 1-55.

50. Little, G. D. G. (2011). *Programming with human computation.* Cambridge, MA, USA: Massachusetts Institute of Technology.

51. Klein, M. (2012). Enabling large-scale deliberation using attention-mediation metrics. *Computer-Supported Collaborative Work*, *21*(4), 449-473.

52. Klein, M., & Garcia, A. C. B. (2014). The bag of stars: High-speed idea filtering for open innovation. http://ssrn.com/abstract=2387180.

53. Klein, M., & Convertino, G. (2014). An embarrassment of riches: A critical review of open innovation systems. *Communications of the ACM*, *57*(11), 40-42.

54. Klein, M., Spada, P., & Calabretta, R. (2012). Enabling deliberations in a political party using large-scale argumentation: A preliminary report. *Proceedings from 10th International Conference on the Design of Cooperative Systems (COOP 2010)*. Marseille, France.

55. Miller, B., Hemmer, P., Steyvers, M., & Lee, M. D. (2009). The wisdom of crowds in rank ordering problems. *Proceedings from Cognitive Modeling 2009.*

56. Moor, A. D., & Aakhus, M. (2006). Argumentation support: From technologies to Ttols. *Communications of the ACM*, *49*(3), 93.

57. Morgan, J., & Wang, R. (2010). Tournaments for ideas. *California Management Review*, *52*(2), 77.

58. Moris**,** F. (2004). Industrial R&D employment in the United States and in US multinational corporations. http://www.nsf.gov/statistics/infbrief/nsf05302/.

59. Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of “success”. *Journal of Behavioral Decision Making*, *17*(2), 117-137.

60. Ogawa, S., & Piller, F. T. (2006). Reducing the risks of new product development. *MIT Sloan Management Review*, *47*(2), 65.

61. Oleson, D., Sorokin, A., Laughlin, G. P., Hester, V., Le, J., & Biewald, L. (2011). Programmatic gold: Targeted and scalable quality assurance in crowdsourcing. *Human Computation*, 11, 11.

62. Paulus, P. B., & Nijstad, B. A. (2003). *Group creativity: Innovation through collaboration.* Oxford, UK: Oxford University Press.

63. Phillips, M. (undated). Open for questions: President Obama to answer your questions on Thursday. http://www.whitehouse.gov/blog/09/03/24/Open-for-Questions-President-Obama-to-Answer-Your-Questions-on-Thursday.

64. Piller, F. T., & Walcher, D. (2006). Toolkits for idea competitions: A novel method to integrate users in new product development. *R&D Management, 36*(3), 307-318.

65. Raymond, E. S. (1999). *The Cathedral and the Bazaar.* O'Reilly Media. Sebastopol, CA.

66. Riedl, C., Blohm, I., Leimeister, J. M., & Krcmar, H. (2010). Rating scales for collective intelligence in innovation communities: Why quick and easy decision making does not get it right. *Proceedings of the International Conference on Information Systems 2010*.

67. Riedl, C., Blohm, I., Leimeister, J. M., & Krcmar, H. (2013). The effect of rating scales on decision Quality and user attitudes in online innovation communities. *International Journal of Electronic Commerce*, *17*(3), 7-36.

68. Roberts, J., Hann, I.-H., & Slaughter, S. (2006). Understanding the motivations, participation and -performance of open source software developers: A longitudinal study of the Apache Projects. *Management Science*, *52*(7), 984-999.

69. Saaty, T. (1990). *The Analytic Hierarchy Process.* New York, NY, USA: McGraw-Hill.

70. Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*(5762), 854-856.

71. Salganik, M. J., & Levy., K. E. C. (2012). Wiki surveys: Open and quantifiable social data collection. *arXiv*. http://arxiv.org/abs/1202.0500.

72. Salminen, J., & Harmaakorpi, V. (2012). Collective intelligence and practice-based innovation: An idea evaluation method based on collective intelligence. In *Practice-Based Innovation: Insights, Applications and Policy Implications*, 213-232. Springer.

73. Schulze, T., Indulska, M., Geiger, D., & Korthaus, A. (2012). Idea assessment in open innovation: A state of practice. *Proceedings of the 20th European Conference on Information Systems.*

74. Shirky, C. (2009). *Here Comes Everybody: The Power of Organizing Without Organizations.* Penguin.

75. Shum, S. B. (2008). Cohere: Towards web 2.0 argumentation. *Proceedings of COMMA 2008,* 97-108.

76. Slamka, C., Jank, W., & Skiera, B. (2012). Second-generation prediction markets for information aggregation: A comparison of payoff mechanisms. *Journal of Forecasting*, *31*(6), 469-489.

77. Smith, A. M., Lynn, S., & Lintott, C. J. (2013). An introduction to the Zooniverse. *Proceedings from AAAI Conference on Human Computation and Crowdsourcing.*

78. Soukhoroukova, A., Spann, M., & Skiera, B. (2012). Sourcing, filtering, and evaluating new product ideas: an empirical exploration of the performance of idea markets. *Journal of Product Innovation Management*, *29*(1), 100-112.

79. Spears, B., LaComb, C., Interrante, J., Barnett, J., & Senturk-Dogonaksoy, D. (2009). Examining trader behavior in idea markets: An implementation of GE's imagination markets. *The Journal of Prediction Markets, 3*(1), 17.

80. Surowiecki, J. (2005). *The Wisdom of Crowds.* New York, NY: Anchor.

81. Thompson, V. (2013). IDC MarketScape: Worldwide innovation management solutions 2013 vendor analysis. http://idcdocserv.com/240823_spigit.

82. Tierney, M. L., & Drury, J. L. (2015). Continuously improving innovation management through enterprise social media. *Journal of Social Media for Organizations*, *1*(1).

83. Toubia, O., & Florès, L. (2007). Adaptive idea screening using consumers. *Marketing Science*, *26*(3), 342-360.

84. Vetschera, R. (1990). Group decision and negotiation support—A methodological survey. *Operations-Research-Spektrum*, *12*(2), 67-77.

85. Viegas, F. B., Wattenberg, M., & Dave, K. (2004). Studying cooperation and conflict between authors with history flow visualizations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems 2004*.

86. Von Hippel, E. (2009). *Democratizing Innovation*. Cambridge, MA, USA: MIT Press.

87. Walter, T. P., & Back, A. (2013). A text mining approach to evaluate submissions to crowdsourcing contests. *HICSS '13 Proceedings of the 2013 46th Hawaii International Conference on System Sciences.*

88. West, J., & Lakhani, K. R. (2008). Getting clear about communities in open innovation. *Industry and Innovation*, *15*(2), 223-231.

89. Westerski, A., Dalamagas, T., & Iglesias, C. A. (2013). Classifying and comparing community innovation in Idea Management Systems. *Decision Support Systems, 54*(3), 1316-1326.

90. Westerski, A., Iglesias, C. A., & Nagle, T. (2011). The road from community ideas to organisational innovation: A life cycle survey of idea management systems. *International Journal of Web Based Communities*, *7*(4), 493-506.

91. Wolfers, J., & Leigh, A. (2002). Three tools for forecasting federal elections: Lessons from 2001. *Australian Journal of Political Science*, *37*(2), 223-240.

92. Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives*, *18*(2), 107-126.

93. Yang, Y. (2012). *Open innovation contests in online markets: Idea generation and idea evaluation with collective intelligence.* Temple University.

94. Yorgason, D. R. (2007). Research and development activities of U.S. multinational companies: Preliminary results from the 2004 Benchmark Survey. https://bea.gov/scb/pdf/2007/03March/0307RDofMNCs.pdf.

95. Yu, L., & Nickerson, J. V. (2013). An internet-scale idea generation system. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, *3*(1), 2.

AUTHOR BIOGRAPHIES

Mark Klein (m_klein@mit.edu) is a principal research scientist at the MIT Center for Collective Intelligence, as well as a visiting researcher at the Dynamic and Distributed Information Systems Group at the University of Zurich in Switzerland. His research draws from such fields as computer science, economics, operations research, and complexity science in order to develop and evaluate computer technologies that enable greater collective intelligence in large groups faced with complex decisions, especially those in the sustainability realm.

Gregorio Convertino (gconvertino@informatica.com) is a Senior User Researcher at Informatica Corporation (CA, USA). His interests are in Human-Computer Interaction (HCI), Collaborative and Social Computing, and Information Visualization. Since 2008, he has worked as user researcher in industry (Xerox PARC, Xerox Europe, Informatica). His work focuses on designing and innovating enterprise software for data management, analysis and visualization, sharing, and deliberation. He chaired four workshops and two special issues on collective intelligence tools for organizations. More about his work can be found at www.gconvertino.com.