

Games for Crowds: A Crowdsourcing Game Platform for the Enterprise

Ido Guy*
Yahoo Labs, Israel
idoguy@acm.org

Anat Hashavit
IBM Research – Haifa, Israel
anath@il.ibm.com

Yaniv Corem*
MIT Media Lab, USA
yaniv.corem@media.mit.edu

ABSTRACT

In this paper, we present a crowdsourcing game platform that allows users to play, create, and share simple games that harness the collective intelligence of employees within the enterprise. The platform uses the wizard design pattern to guide users through the process of creating a game. We describe the platform in detail and report our findings from deploying it within a large global organization for a period of three months, in which 34 games were created by 25 employees and played by 339. We combine qualitative and quantitative analysis to understand the characteristics of the different games and their impact on popularity and engagement, to validate our design goals, and to suggest potential enhancements.

Author Keywords

Crowdsourcing; Games; Human Computation

ACM Classification Keywords

H.5.3 Group and Organizational Interfaces: *Computer-supported cooperative work*

INTRODUCTION

Crowdsourcing has come a long way since the days of Sir Francis Galton, who tried to estimate the weight of an ox at a county fair by averaging the individual estimates of people in the crowd [26]. In fact, various applications of crowdsourcing in recent years have shown, not only that crowdsourcing is a viable problem-solving technique, but that for some tasks under certain conditions, it can outperform an individual expert or computer algorithm [31,32,34]. Crowdsourcing is the act of taking a job traditionally performed by a designated agent and outsourcing it to an undefined, generally large group of people, in the form of an open call [16,24]. Among some of the more prominent applications of crowdsourcing are image tagging [31], text digitization [34], and translation [4]. Despite its conceptual simplicity, crowdsourcing

presents many intrinsic challenges that ultimately determine its success. These challenges include recruiting and retaining participants, quality assurance, task assignment, aggregation of results, and more.

The benefits of crowdsourcing have not escaped the attention of large companies as they search for ways to harness the collective intelligence of employees to drive business. Several attempts have been made to apply crowdsourcing techniques within the enterprise to enrich the social graph [15], locate assets [35], or recommend bookmarks [11]. Despite the relative success of these initial attempts, enterprise crowdsourcing is still in its infancy and often remains an untapped resource due to challenges such as the entry barrier and employee engagement in a workplace setting.

In this work, we introduce a crowdsourcing game platform called “Games for Crowds” (G4C) that allows users to play, create, and share simple games that harness the collective intelligence of employees within the enterprise. G4C addresses the challenges of enterprise crowdsourcing in several ways: (1) it packages the crowdsourcing task as a game, making it more fun and engaging for users; (2) it not only allows users to play, but also enables them to create simple crowdsourcing games; and (3) it builds a community around players and game creators, making users feel more involved in the crowdsourcing tasks.

Creating crowdsourcing games in a simple and intuitive manner is a key aspect of G4C. To achieve this, we extract the game’s DNA – goals, logic, rules, feedback system, and entities – and then use a wizard design pattern to guide users through the process of creating it. The wizard helps the user provide a set of parameters that are essential for creating the game. It includes three parts: general information, such as the name and description of the game; the topics and questions, which form the game’s content; and advanced settings, such as the order of questions and topics and the reward mechanism. These parameters are used by the wizard to initialize a game that is similar in its look-and-feel to the GuessWho game [15], which the user can access and share with others through the platform.

G4C was deployed in a large global IT organization for a period of three months. During this period, 34 different games were created by 25 different employees and played by a total of 339 employees, who contributed a total of 28,374 answers. Our evaluation is based on data collected from the G4C database logs, as well as on interviews of 12

*Part of the research was conducted while working at IBM Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSCW 2015, March 14–18, 2015, Vancouver, BC, Canada.

Copyright © 2015 ACM 978-1-4503-2922-4/15/03...\$15.00.

<http://dx.doi.org/10.1145/2675133.2675189>

creators and 12 players who played at least 3 games. We review the different games and their characteristics to give a sense of what kind of games employees are creating. We also try to understand what factors contribute to the popularity of a game and its players' contribution. Finally, our analysis aims at validating the G4C design goals and understanding what has worked well and what can be improved.

RELATED WORK

Crowdsourcing Games

One of the major challenges in crowdsourcing is engagement, i.e., motivating individuals within the crowd to complete a task [5]. Games have already been used to make crowdsourcing tasks more fun and engaging. Luis von Ahn coined the term *Games With A Purpose* (GWAP) [30] to define a type of crowdsourcing game that is (a) fun and engaging, (b) includes a task that can only be completed by humans, and (c) has a goal that is hidden from the player. The first GWAP, the ESP game [31], matched two players against each other and presented them with a random image they were asked to label with as many tags as possible. Whenever players' tags corresponded, both players would receive points. Following the success of the ESP game, other GWAP were created such as Peekaboom [33], a game for locating objects in an image and Verbosity [32], a game used to collect commonsense knowledge. These GWAP and others were later collected to form a game portal where users could play, and in the process, complete tasks. Similarly, G4C includes a portal of crowdsourcing games. However, G4C games are created by users, often without past experience in game creation or crowdsourcing, making its portal a community hub for game creators and players.

Crowdsourcing Platforms

With the rise of crowdsourcing, corresponding platforms and markets have become popular [6,10]. Amazon's Mechanical Turk (MTurk) is an increasingly popular web service for paying people (Turkers), typically a few cents, to carry out simple tasks, known as Human Intelligence Tasks (HITs), that would otherwise be extremely difficult, if not impossible, for computers to perform [17,19]. MTurk has been largely successful for conducting user studies [17], speech-to-text transcription [13], clustering [23], and more. However, the use of monetization as an incentive has sometimes been shown to have a negative effect on the completion time and quality of work produced by Turkers [5,7]. Quadrant of Euphoria is a crowdsourcing platform meant to address the problems of inaccuracy and cost when trying to assess the experience of multimedia end-users. Workers can be recruited either by payment or by publishing to online communities [8]. InnoCentive is an open innovation platform that broadcasts research and development problems for anyone to solve, granting cash awards for the best solutions that meet the criteria [3,18]. In contrast, G4C uses game mechanics, such as leader-boards, to incentivize users to play, and

awards points as compensation for completing micro-tasks.

Crowdsourcing and Games in the Enterprise

The most comprehensive work on crowdsourcing within the enterprise is by Vukovic [35], who presented a sample scenario for enterprise crowdsourcing in the software development domain, in order to derive requirements and a research agenda for a general-purpose crowdsourcing service in the cloud. Among the identified challenges were the task request model, the matchmaking process, the bid evaluation, and the quality of experience measurement. The paper did not focus on games as the incentive mechanism. In his work about video games, Gee [14] argues that the use of game technologies for learning skills in workplaces will become pervasive. Mollick and Rothbard [21] examined whether managerially-initiated games at work can provide the desired benefits for affect and performance. They found that the answer depends on the notion of "consent" – the degree to which employees actively cooperate with management initiatives.

Another body of research focuses on brainstorming at large scale through crowds in order to collect and evaluate ideas. This is often referred to as "crowdstorming" or "ideastorming". The Dell IdeaStorm website [9] was launched in 2007 to help Dell promote innovation by collecting and measuring ideas. Crowdsourcing in this case went beyond the scope of the enterprise, reaching out to the company's end user community. In their recently published book, Abrahamson et al. [2] argue that enterprise crowdstorming tools focused on gathering, assessing, and understanding ideas, may be more beneficial for organizations than "traditional" enterprise social media tools. While some of the games created using G4C focused on idea gathering, others had different goals, such as building knowledge bases, collecting resources, or educating.

A few studies focused on specific enterprise crowdsourcing games. The Dogear game [11] drove employees to create and share bookmarks within the enterprise, while the GuessWho game [15] collected social network information from employees. The G4C system builds on some of the elements included in GuessWho, as described later in this work.

Human-Centered Game Creation Tools

The culture of Web 2.0 has made a significant impact on the video game industry [2,14]. The e-Adventure3D platform [28] was among the first to facilitate the rapid creation of 3D adventure games without any programming background. The popular Minecraft game [12] introduced a block-based editing paradigm where players can create their own games using a collection of meter-square blocks, from tree leaves to floating clouds. In Sony's Little Big Planet [36], players create the levels of the game by using an extensive inventory of game objects, from building materials to physical behavior. Additionally, algorithms

The screenshot shows the 'Games for Crowds' home screen for user Boyd Gossens. The top navigation bar includes 'Home', 'Studio', 'About', and 'Log Out'. The left sidebar features a user profile picture, a welcome message, and statistics: 'Created: 6 games' and 'Played: 3 games'. A vertical 'feedback' button is on the left edge. The main area is titled 'My Games' and is split into 'Recently Played' and 'Recently Created' sections. 'Recently Played' lists 'Scattegories' (Rank 150, Score 8, Answers 4), 'Let's Map!' (Rank 9, Score 3, Answers 2), and 'Linux' (Rank 51, Score 268, Answers 11). 'Recently Created' lists 'My Dear Watson' (10 Players, 201 Answers), 'Book Me!' (175 Players, 13762 Answers), and 'Doctor in the House' (19 Players, 167 Answers). Below this is the 'All Games' section with a search bar and a scrollable library of games: 'GuessWho' by Boyd Gossens, 'Healthy Forms...' by Al Moreno, 'Get Social' by Allie Singh, 'Boom' by Anna Bauer, 'Smartphone Apps' by Bill Jordan, and 'Retail IQ' by Minh Li. The right sidebar contains 'Recent Activity' (Network and All tabs), 'Most Popular' games (Scattegories: 307, Presents for holidays: 291, Busy Bee: 68), 'Recently Created' games (Hobbies: Mar 18, 2012, What would...: Mar 18, 2012, Ender's Game: Mar 18, 2012), and 'Top Scoring Games' (Busy Bee: 6968, Scattegories: 3074, My Dear Watson: 2915).

Figure 1. Personalized home screen.

have been widely used to automatically create levels in games [25]. Mehm et al. [20] provide an overview of state-of-the-art game authoring tools and suggest future directions, such as multiplayer authoring and automated model checking. These game authoring tools substantially lower the entry barrier, allowing novice authors without technical expertise to contribute. Analogously, G4C aims at lowering the entry barrier for creating human computation games, which unlike video games have the hidden agenda of collecting data from the crowd. Hence, more emphasis is put on the game's content and logic than on its UI. Game creators are provided with a predefined UI, in which they can control a limited set of parameters.

To the best of our knowledge, no other work provides a platform that brings together game authoring and crowdsourcing capabilities in the enterprise. We take a user-centered approach to crowdsourcing and extend the role of every individual from worker to “crowdsourcer”, with the power to create crowdsourcing games in the workplace.

GAMES FOR CROWDS

Design Goals

G4C was designed to help users create games that leverage human computation, with the following goals in mind:

- Game lifecycle – provide users with a simple, intuitive, and fast way of creating and managing games
- Data and analytics – enable creators to collect game-generated data and analyze it to gain initial insight
- Engagement – create an entertaining and engaging user experience for players, to provide a crowd for creators
- Community – blur the lines between players and creators and allow them to work (and play) together.

The rest of this section describes G4C and its UI.

Home Screen

After logging in with their intranet ID and password, users are navigated to a personalized home screen (Figure 1), which includes three sections. The left sidebar displays the user's profile including a picture and the number of games played and created. The right sidebar displays four different widgets, which together provide a snapshot of game activity. The Recent Activity widget is a news feed that includes information organized into two tabs – Network and All – about users who recently created or played a game. The All feed includes any public activity from any G4C user, whereas the Network feed is limited to a user's social network (taken from the enterprise social network service). Under the Most Popular widget, users find the three most popular games by the total number of players in the last month. Similarly, the Recently Created widget provides the three games most recently created. Finally, the Top Scoring Games widget displays the three games with the overall top scoring player per game.

The home screen's central section shows the user's three most recently played games and three most recently created games. If no games have been played or created, a message is displayed to the user, prompting him to play or create games. The bottom part contains a library of all games created by the different users. Users can manually browse the library by horizontally scrolling through it; sort the library's games by title, date, creator, or popularity; and search for games using a free-text string.

Studio

G4C provides a studio for users who want to create a game. The left and right sidebars in the studio remain the same as in the home screen, displaying the profile information and gaming activity, respectively. The studio's middle section is shown in Figure 2. It allows the user to type a name and create a new game using the New Game wizard. It also displays all games created by the user and allows to toggle a game online or offline, delete a

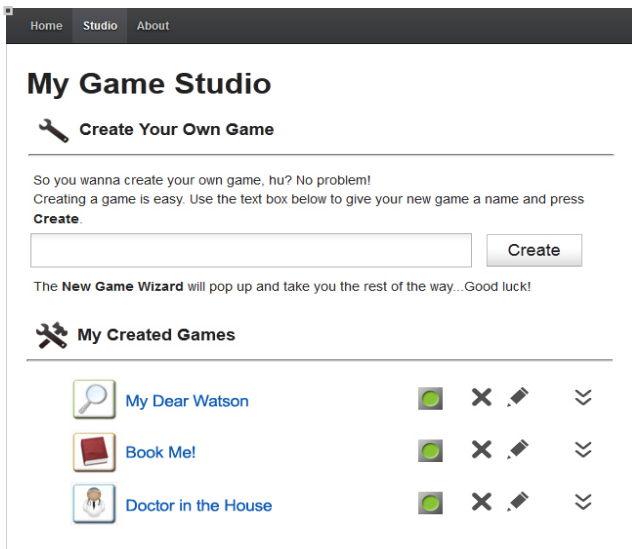


Figure 2. Studio screen.

game, edit game details, and view the game’s statistics by clicking on the arrow at the end of each row.

New Game Wizard

The following section describes the process of creating a new game using the New Game wizard. One of our design goals was to provide users with a simple, intuitive, and fast way of creating games. Therefore, we decided to follow the wizard design pattern [27], which guides the user step-by-step through the process. The wizard, illustrated in Figure 3, includes three parts: (1) General Information (2) Topics and Questions, and (3) Advanced Settings.

General Information

The first step in the game creation process requires general information about the game such as name, description, playing instructions, and a representative thumbnail. If no thumbnail is provided, a default one is used. Some fields, such as Name and Description are marked with a red asterisk to indicate that they are mandatory.

Topics and Questions

The second step in the game creation process is to insert the content of the game, i.e., topics and questions. A topic

and a question are presented to the player in each round of the game, which lasts one minute. Typically the question refers to the topic, for example, if the topic is *Research Division*, then the question can be *Name as many departments as you can under the Research Division*. The user needs to provide as many *good* answers as possible during the time allotted, where *good* is defined by the scoring function, which considers answers provided by other players to the same topic and question. Topics can be embedded within questions using the [topic] placeholder, which is replaced during game runtime by the current topic. In a simple embodiment, there is either one question that refers to many topics or one topic to which many questions refer. In a more advanced embodiment, there are multiple questions and topics. In this case, each question is paired with all topics. For example, the GuessWho game had two questions – *People related to [topic]* and *Tags related to [topic]* – and multiple topics representing different people’s names. During the game, the [topic] placeholder was replaced with a person’s name and the player’s goal was to list as many related people or tags as possible [15].

The user must enter at least one topic and one question, while there is no upper limit on both. Topics and questions can be entered either manually or through an uploaded comma-separated value (CSV) file. Overall, in order to create a game, the user must enter a name, description, at least one topic, and at least one question.

Advanced Settings

The third and final step in the game creation process allows users to set several parameters that control the logic and behavior of the game. *Order* determines the order in which topics and questions are displayed, with three options to choose from: (1) Random – displays the topics and questions in random order (the way it was implemented for GuessWho [15]); (2) Fixed by question – displays a question and then loops through all topics (in a randomized order); (3) Fixed by topic – displays a topic and then loops through all questions (randomized order).

Access determines whether or not a game is available for everyone to play. If Public is selected, the game will

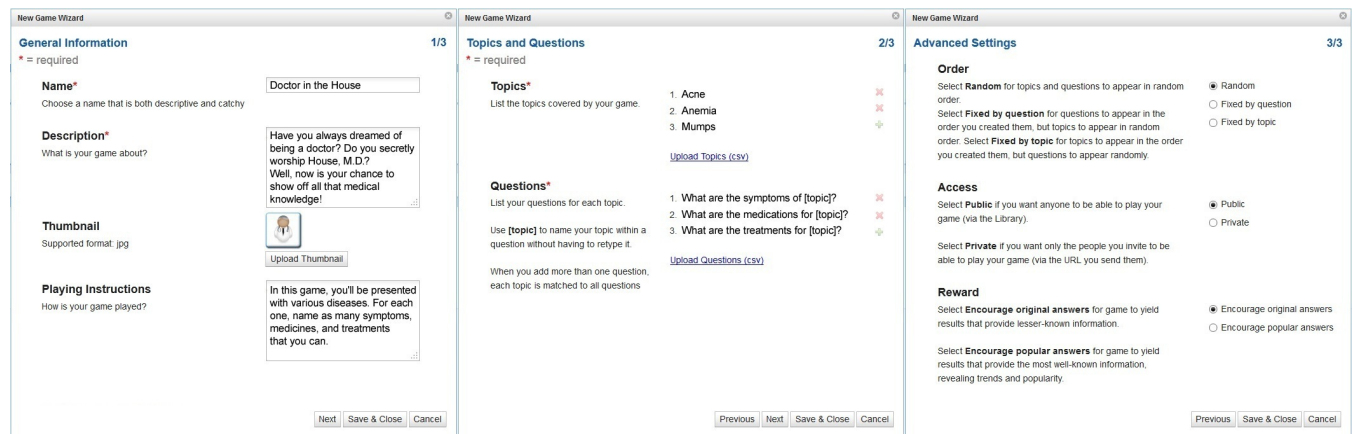


Figure 3. New Game wizard.

appear in the All Games library and various sidebar lists when relevant. If Private is selected, the game will only be accessible through a unique URL. The *Reward* section controls when and how points are awarded to a player for completing a task. If ‘encourage popular answers’ is selected, players will receive points based on the number of times an answer was already mentioned by other players. If ‘encourage original answers’ is selected, the score per player will increase for the first few players, up to a given cutoff threshold, and then decrease towards zero for answers that are already very popular. Figure 4 describes the scoring of a given answer as a function of the number of players who mentioned that answer for both types of reward mechanisms. In addition, for ‘encourage original answers’, dividend points [15] are used to reward the previous players who gave an answer with each subsequent mention of that answer by another player (up to a threshold). This mechanism aims at encouraging diverse yet valid answers (see [15] for more details).

The wizard produces a game (Figure 5) that is similar in its look-and-feel to the GuessWho game [15]. It uses the information provided by the game creator to propagate the various fields in the UI. For example, the current topic appears in the center circle of the game’s visualization. The question, with the current topic highlighted in green, appears at the top of the UI. Answers already given by the user for this question and topic are highlighted in circles coming out of the central circle, while the small opaque circles denote answers given by other users for this pair of question and topic that the user has still not given. See [15] for more information about the type of games created using G4C and a detailed description of the UI.

Game Statistics Page

The game statistics page, reached from the studio page as described before, is shown in Figure 6. It includes details such as the total number of players and answers and a button for exporting the game-generated data to a CSV file, which can be processed further using external tools. The central part of the statistics page includes a graph showing the distribution of answers for a given topic-question pair. Figure 6, for example, shows data for the topic *Animals* and the question *Name as many Animals as you can*. The x-axis contains the answers sorted by popularity and the y-axis shows the number of mentions per answer. If an answer passed a predefined threshold, it is considered valid and appears in blue; otherwise it is colored gray and is considered unverified. The bottom of the page has two graphs showing the number of players and number of answers as a function of time.




Figure 4. Scoring functions.

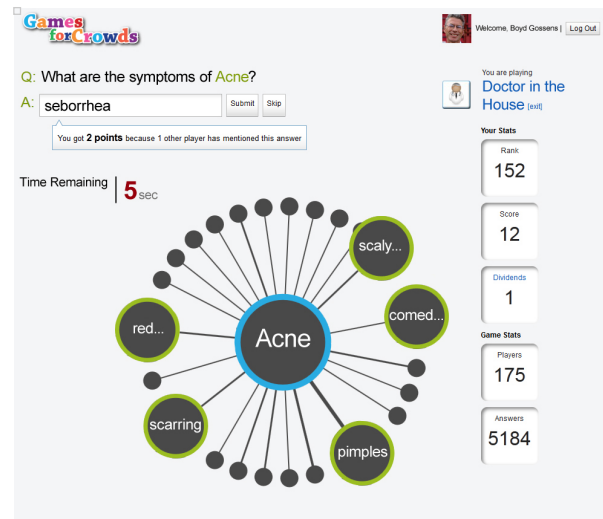


Figure 5. The generated game.

EVALUATION

Setup

G4C was deployed in our organization, a large global IT company, for a period of three months. For the evaluation, we examined the G4C server logs, which documented the details of each created game (name, creator, parameters, etc.) and the details of each answer provided by players in a game along with the player’s username and a timestamp. We also conducted interviews with 12 game creators and 12 active players (who played at least 3 games), to get a better understanding of the user experience from both the creators’ and players’ perspectives. Interviews were semi-structured, conducted by phone, and lasted 30 minutes each. The creator interview template contained 17 questions covering the motivation for game creation, the game generation process, the dissemination and



Figure 6. Game statistics screen.

management of the game, and the overall outcome of the game. The player template contained 15 questions regarding the motivation for playing games, the playing experience, the characteristics of engaging games, and the effect of games on the sense of community.

Our analysis provides an overview of the games created and the motivation for creating them, and then discusses each of the design goals, to what extent it was achieved, and how it could be improved. As part of the engagement design goal, we also discuss game characteristics that were found to influence its popularity and players' contribution.

Game Overview

During the three-month period, 25 employees from 7 different countries created 41 different public games. Only 4 additional games were defined private. Our analysis of the data collected from this period comes from 34 unique games, excluding 5 games that were created for testing purposes and 2 that included monetary rewards that were not part of the platform. On average, each game contained 23.15 questions (stdev: 42.32, max: 205, median: 8.5) and was played by 24.88 players (stdev: 35.35, max: 165, median: 11). Overall, 339 players from 31 countries played at least one game. Players spanned the different business units of our organization: 36.87% Services, 16.22% Software, 13.57% Research, 12.98% Sales, 5.6% Operations, 5.02% Headquarters, 4.72% Systems, and 5.02% others. These players contributed a total of 28,374 answers, distributed by the statistics described in Table 1.

Observing the games created, we distinguish between two types of games: business games (24 of 34), which relate to the creator's job, and leisure games, which relate to topics outside of work (10 of 34). For example, LotusHear is a business game, which included questions about a large annual business event, such as *Name influential keynote speakers*; My Dear Watson is an example of a leisure game that contained trivia questions, such as *Name as many whisky brands as you can*. Table 2 shows the 20 most popular games by number of their players. For each game, the following information is presented: (1) name, (2) reward mechanism ('P' for games that encouraged popular answers or 'O' for games that encouraged original answers), (3) type ('B' for business or 'L' for leisure), (4) the number of questions in the game, (5) the number of players who played the game, and (6) the total number of answers collected in the game.

Motivation for Creating Games

When asked about their reasons and motivations for creating games, many creators mentioned they wanted to gain or expand knowledge, often to address a challenge in

	Avg	Stdev	Min	Max	Median
Answers per game	834.53	2,136.75	4	9,081	189
Answers per question	48.92	67.66	1	517	21
Answers per player	83.7	110.53	1	1,030	38

Table 1. Answer Statistics in G4C.

their day-to-day work. Creator of Big Business Glossary said: *"through the game I was trying to build the business taxonomy for various domains in our organization [...] from water management to front office digitization."* Creator of the Linux game noted: *"I wanted to expand my technical knowledge further in an area I'm already knowledgeable."* The <Customer> Challenge game was built to collect information from salespeople about one of the company's main clients, asking about its brands, products, deals, and contact persons. Creator of Let's Map <Company> described: *"We were in a group meeting and my boss was explaining on the board all sorts of things about the organizational structure [...] he didn't always remember all the details [...] so it struck me I can create it using the crowd [...] every person would get a division name and would have to indicate its parents or children."*

Other games sought to discover assets or ideas. Creator of the <Company> Cloud Seeding game explained: *"I used the game to inspect what resources people were using to find information about Cloud Computing"*. Creator of Healthy Form of Exercises told us he created the game *"to get feedback and ideas of healthy forms of exercise, so that I can design exercise missions for health promotion programs."* Volunteering game's creator described: *"I had two main purposes for the game: get new ideas from the crowd about volunteering projects for <Company> and get a sense of how 'real' this crowd is, the potential number of people who would actually volunteer"*.

Teaching the crowd was also mentioned as a reason for game creation. Indeed, enhancing education and learning

Name	Rewd	Type	Qust	Plyr	Answr
Scattogories	O	L	156	165	9001
My Dear Watson	O	L	44	125	9081
Doctor In The House	O	L	51	60	787
Let's Map <Company>!	P	B	28	57	704
Boom	O	B	4	51	385
Linux	O	B	16	50	2637
Smarphone Apps	P	L	18	42	391
Learning on the Fast Track	O	B	3	38	588
Who wants to be a Socialite?	P	L	14	30	997
Healthy Forms of Exercise!	O	B	18	27	587
The <Customer> Challenge	P	B	4	21	272
Let's Map HRL!	P	B	40	19	281
LotusHear	O	B	7	19	246
Test your IQ	P	L	4	19	102
JavaScript Skill Test	P	B	7	16	83
Get Social Biznified	O	B	4	14	205
Volunteering	O	B	12	12	100
Big Business Glossary	O	B	9	10	129
Retail IQ	O	B	25	9	145
All About SDG	O	B	24	9	521

Table 2. Twnety most popular games in G4C.

has already been pointed out as a key goal for computer games in the workplace [14]. Creator of the ISC Game told us: *“I was looking for a fun, interesting, and engaging way to teach our division about our strategy and products. I needed something that was competitive and that would get the team excited.”* Creator of Get Social Biznified said: *“I wanted to create a game that would quickly teach a group of top-level executives what Social Business is,”* and added: *“I also wanted to showcase the phenomenon and potential power of crowd games in the enterprise”*.

Several creators mentioned they wanted to produce a fun and engaging experience that would increase excitement in their work environment. Finally, interest and curiosity about the game creation process and the capabilities of enterprise crowdsourcing were mentioned, for example, *“I was intrigued by the idea and wanted to see how easy it was to create a game and get people to play it.”*

Game Lifecycle

In our interviews with game creators, the simplicity of game creation was consistently mentioned as a primary advantage of G4C. Fast, easy, simple, intuitive, and convenient were among the commonly used adjectives to describe the game generation process. One creator said: *“I was able to get it up and running within a few minutes”* and another noted: *“The game was simple to put together [...] truthfully, if the process was more sophisticated, I wouldn't have tried it out [...] it should be simple enough so even non-techies like me can easily use it.”* Indeed, about half of the game creators (12 of 25) belonged to the non-technical divisions (Sales, Services, Operations), indicating that the creation process was intuitive enough for non-technical employees.

Several creators indicated they would like more control and flexibility when creating a game, for example, by being able to choose a color “theme” for the game, set the duration of a game round, or set the scoring function in a more flexible way. One creator said: *“The platform should be more tunable [...] support multiple-choice questions, allow spell-checking for inputs, or a single-answer mode”* and another suggested: *“support more complex types of games, like games with non-textual inputs [...] or allow using one game to validate the results of another.”* There is naturally a trade-off between the simplicity of game creation and the flexibility it allows. But this feedback suggests that crowdsourcing game platforms should offer more choices to advanced or experienced game creators, while keeping reasonable defaults that preserve the option of creating an initial game very quickly.

While the game creation process itself was simple, creators sometimes had to invest more time in framing their problem as a game: *“I had to change my view of the problem and look at it through the lens of a game [...] design questions whose answers can be easily compared for granting points.”* There were a few suggestions from creators to further facilitate game creation. Duplicating an

existing game as a basis for creating a new one was mentioned by several creators. One of them said: *“I looked at other popular games to learn how to create my own [...] seeing a good game made me want to create a similar one”*. Another interesting suggestion was to *“show game creators popular choices for created games. For example, indicating that it is common to ask 20 questions or write 5 sentences as instructions can help you understand the expectations and trust your game's ability to succeed”*.

Our logs show that thumbnail images were provided for 28 of the 34 games (82.35%). This suggests that creators were willing to invest some time to make the game more attractive and unique. Of the 34 games, 24 (70.59%) provided additional instructions on top of the game's description, again pointing to creators' willingness to dedicate time to non-mandatory fields they perceive as important. Only five games used both multiple topics and multiple questions. Two examples are Scategories, which used alphabetic letters as topics and contained questions such as *Name cities that start with [topic]* and Smartphone Apps, which defined smartphone application categories such as games, education, and lifestyle as topics and contained questions such as *Name good [topic] apps that cost less than 1\$*. Several creators said they did not fully understand how the use of multiple topics and questions should work. Providing a concrete example or a simple UI to match topics with questions may help make this feature clearer and more accessible.

Almost all games (28 out of 34) were edited during the time period, with an average number of 3.52 revisions per game (stdev: 2.58, max: 15, median: 3). Creators reported using the game edit option to keep their game competitive. The creator of All About SDG told us: *“as the game progressed we added more and more questions, to make it tougher and more challenging.”* The creator of Let's Map HRL noted: *“I manually added departments that I got as answers as new questions to the game, to keep it fresh.”* A future G4C enhancement can allow converting answers to questions by certain rules to address similar needs.

Data and Analytics

Creators were generally satisfied with their game's ability to produce data quickly. *“It was amazing how much data you could get in a very short time”* said one and another noted: *“We got like 84 answers within 5 minutes.”* Inspecting the data created in different games indicates that indeed games collected a wide variety of data from the company's crowd. As a few anecdotal examples, 31 US presidents were mentioned in the game My Dear Watson (out of 43 in total); 168 countries were mentioned in Scategories (out of 193 UN members); 21 departments were mentioned in Let's Map HRL (out of 40 in total); and 41 system calls were mentioned in the Linux game (out of 348).

In our interviews, 7 of 12 creators indicated they viewed the statistics page. *“The section that shows the distribution*

of answers per question was very handy”, said one of them. Of these 7 creators, 5 said they exported the game’s data into a spreadsheet and reported different uses. The creator of the Volunteering game said “I got about 7 new ideas [...] The answers had a good mix of expected ones with a few surprising ones [...] You don’t want to have all results surprising, as it indicates the crowd did not understand you, while you also don’t want all answers to be expected, since you will not learn anything new.” The creator of Healthy Forms of Exercise noted: “The game provided extra input I could not find on the intranet [...] both brand new info and perception of existing info by people in the company.” The creator of Smartphone Apps said: “The export feature allowed me to share a spreadsheet with my colleagues and try to sort and make sense of the data together.”

Apparently the data extraction feature was not visible enough for some creators. One said: “the only difficulty I had was getting the report but then I sent you a mail and subsequently was able to go in and see what people really contributed.” Another noted: “finding how to get to the analysis page took me a while [...] I think it should be more prominent.”

Creators suggested various enhancements to G4C’s data analysis features. One proposed: “more possible queries on the data [...] allow grouping or merging answers based on rules or categories” and another commented: “I was curious to see whether people in different geo’s would mention different brands of <Customer> products.” Another creator suggested using data visualization tools such as ManyEyes [29] to enhance data analysis capabilities.

Engagement

In our player interviews, crowdsourcing games appeared to serve as a form of escapism at work, providing a good alternative to take a break from work for a few minutes, without the guilt typically involved with going to external games, news sites, or social networks. One player noted: “When I context-switch at work, I usually need a break to clean my head [...] the games provided me an interesting option for that”. Another player noted: “I liked the games that related to different things internal to <Company>, such as divisions and products [...] gave me the feeling my playing is helpful and contributes something,” and another said: “I had a bad feeling since I was playing games while at work [...] if it’s a kind of a brainstorm that worth something to <Company>, it would make me more easy.”

Fun and challenge, sometimes mixed with the desire to learn, played a key role in players’ engagement. One player explained: “It was a way for me to quiz myself and learn more about the company I work for” and another noted: “the clock ticking down really added the element of challenge and excitement to the game.” The gaming aspects were also intimidating for some employees, for example one player expressed her anxiety of not

succeeding in a game within the enterprise: “I felt stressed, worried I would look dumb [...] so I only played games I felt I had a good chance to know the answers.” A similar point was made by one of the creators: “One guy told me that he didn’t play because he had too much work and he wasn’t willing to be at the bottom with his name and picture displayed”.

Players’ competitiveness did not escape the eyes of game creators. One of them said: “The really competitive people were just going crazy [...] even my manager was playing over and over again.” G4C logs indicate that 3 managers created a game (12% of all game creators) and 57 managers played at least 1 game (16.8% of all players). These percentages are similar to the portion of managers in the organization (13%). Managers were insignificantly more successful than non-managerial employees in terms of their average rank per game (their average rank in a game divided by the total number of players in the game): 0.58 vs. 0.53 (one-tailed, unpaired t-test, $p=.16$).

The average duration of a game, from the first day it was played to the last, was a little over a month – 33.3 days (stdev: 27.9, max: 85, median: 30). Each game was played on average in 8.18 different days (stdev: 6.9, max: 28, median: 7). Games generally appeared to be a good mechanism to extract a significant amount of data in a short period of time, rather than serving as long-living applications. This is consistent with previous findings on crowdsourcing games in the enterprise [15].

Creators reported various means to publicize their games. Some relied on G4C’s existing mechanisms, e.g.: “The fact that G4C shows newly created games helped a lot since I knew my game would be there at the top of the list for at least some time after it has been created.” Others published the game via email, instant messaging, microblogging, and face-to-face meetings. More mechanisms should be added to G4C to support game diffusion and virality, such as distributing games to matching communities or sending update emails to players with teasers, such as the number of points needed to advance in the leader-board.

Game Popularity and Players’ Contribution

We examined a variety of game characteristics and their influence on two measures of the game: (1) popularity, determined by the number of players who played a game; and (2) players’ contribution, determined by the average number of answers per player per question in a game. We report the findings we discovered across four different game characteristics.

Business vs. Leisure. Leisure games were found to be significantly more popular than business games, with 40.28 players on average compared to 14.1, respectively (one-tailed, unpaired t-test, $p<.05$). On the other hand, in business games, players’ contribution was significantly higher, with an average of 3.59 answers per player per

question versus 1.41 for leisure games (one-tailed, unpaired t-test, $p < .05$). It appears that business games attract fewer players, but receive higher contribution from each. One of the creators explained: “*there is a tension between creating a closed game on a very specific topic, which would potentially attract domain experts and be highly interesting for them and creating a generic game that would attract anyone, but with less specific interest.*”

Reward Mechanism. Games that encouraged original answers (15 of 34) had an average of 32.16 players, while games that encouraged popular answers (19 of 34) had an average of 15.67 players. This difference was found statistically insignificant (one-tailed, unpaired t-test, $p = .073$). When playing a game rewarding originality, players provided an average of 3.41 answers per question, compared to 2.08 for games rewarding popularity. This difference too was found statistically insignificant (one-tailed unpaired t-test, $p = .16$). Overall, these findings suggest that there are no significant differences between the two reward mechanisms in terms of popularity and engagement. In our interviews, we noticed that the selection of a reward method often stems from the goal of the game. One creator who chose to encourage popular answers noted: “*The subject matter of the questions I was asking was trying to get to a common denominator [...] I was trying to identify what lots of people felt the same about.*” On the other hand, another creator who chose to encourage original answers said: “*When I think about crowdsourcing and idea management [...] it works best when you have diverse insight and feedback [...] I wasn't looking for consensus.*”

Instruction Length. The average instruction length per game was 16.31 words (stdev: 17.99, median: 11.5, max: 72). Since instructions were input through a free-text field, without providing any hints or auto-completion, their content and phrasing style was different for each game; no copying of instructions from one game to another occurred. Comparing players' contributions in games with instruction length of 10 words or less (15 of 34 games) to games with instruction length of more than 10 words (19 of 34) indicated no significant difference (one-tailed, unpaired t-test, $p > .05$). Further partitioning these groups into four in total, based on round borders and group size balance, revealed a significant effect, as summarized in Table 3. A one-way ANOVA indicated that players' contributions across the four groups were significantly different, $F(3,30) = 4.415$, $p < .05$. Tukey post-hoc comparisons indicated that the mean for the third group, which had instruction length between 11 and 20 words, was significantly higher than the three other groups. This result implies that instructions should be neither too short nor too long. Presumably, instructions that are too short can fail to adequately explain the game and might also convey lack of attention or effort, while instructions that are too long can be tedious to read and hard to follow. For

Instruction Length	# Games	Mean	Stdev
None	10	1.56	1.30
1-10 words	5	1.73	1.06
11-20 words	7	7.34	7.91
Over 20 words	12	1.69	1.08

Table 3. Instruction length effect on players' contribution (number of answers per player per question).

the game description, we did not observe a similar effect of its length on players' contribution.

Number of Questions. Comparing games with up to 8 questions and games with over 8 questions (17 games in both groups) shows that games with more than 8 questions were significantly more popular: they had 37.59 players on average, compared to 12.18 players for games with 8 questions or less (one-tailed, unpaired t-test, $p < .05$). The number of answers per player per question was, as could be expected, significantly higher for games with fewer questions, at 3.77 compared to 1.36 for games with more questions. It can also be seen that the four most popular games in Table 2 had a large number of questions – over 25 – and the most popular game, Scategories, had a particularly large amount of questions – 156. One player noted: “*I got the same question many times, which was very discouraging [...] creators should be urged to upload more questions,*” and another said: “*If there was a game that had lots of questions I would probably assume that it is a well set up game or a game that's got more essence in it.*” Players also indicated that questions that required shorter answers were more appealing “*since the task was clearer and so was the scoring*”, as one of them explained.

Community

Several games related to existing communities within the organization. The game All About SDG was created to support the Solution Development Group by “*communicating with our developers and analysts [...] the game experience, the competition, and different answers helped strengthen the community*”. The creator of the Linux game attributed its success to “*the nature of the Linux community – all developers, technical, know their area very well and are passionate about Linux*”. Game creators wished to have a way to easily connect a game to an online community [22], so it is exposed on the homepage of the community and accessible to its members. “*Online communities can be a great way to drive crowdsourcing games in the organization*”, said one creator and another suggested: “*would be great if G4C can support competition between communities.*”

Aside from strengthening a specific community, both creators and players pointed out that G4C helped share knowledge and connect people that are not necessarily familiar with each other. One player said: “*It takes you out of the everyday race for a few minutes and makes you aware of other colleagues*” and another noted: “*The right*

panel with all gaming activity contributes to the sense of community among employees who play and create different games". Another player mentioned he would want more mechanisms to enhance the sense of community among players: "promote more interactivity between players even if they are anonymous [...] show another player who answered a question while I was playing [...] or help discover players who are similar to me based on how they play different games".

G4C logs show that players who also created a game (20 in total) played significantly more games than players who did not: 3.7 versus 1.56 (one-tailed, unpaired t-test, $p < .05$). In addition, game creators had significantly more members of their organizational social network play games: 1.85 versus 0.71 (one-tailed, unpaired t-test, $p < .05$). These findings indicate that game creators are more engaged as players, not just of their own games, and also engage their social environment. The mutuality between playing and creating a game was reflected in several employees creating a game after gaining some experience as players. One commented: *"After playing I saw how simple it can be to create a game. It removed the barrier"*. In other cases, creators mentioned they played other games after they created their own: *"playing other games taught me more ways to pose my questions"*.

DISCUSSION, LIMITATIONS, AND FUTURE WORK

We examined 34 games created along a three-month deployment period of G4C in our organization. Creators reported different motivations for creating games, from learning from the crowd and collecting knowledge, through experiencing with crowdsourcing in the enterprise, to producing fun and excitement in their teams. The game creation process was reported to be easy and intuitive, typically taking just a few minutes to complete. Creators did invest in non-mandatory fields to improve their game: most of them provided instructions in addition to the mandatory description and most also provided a thumbnail image to characterize their game. Most creators also revised their game multiple times after its initial creation. Together with its simplicity, the game creation process was often pointed out as not flexible enough. Future versions should provide more flexibility, such as supporting more complex game types that would allow framing more business problems as games; allowing the reuse of an existing game when creating a new one; and controlling more parameters such as round duration or colors of UI elements.

Games were created to gather different types of data, from ideas of volunteering or exercises at work, through knowledge resources about professional topics such as cloud computing, to items that help build knowledge bases, such as business glossaries and organizational charts, or expand the business knowledge around technical topics (Linux, Javascript), customers, or business events. Many creators used the statistics page to view their data

and almost half chose to export the data into a spreadsheet. Creators' feedback also indicated that the prominence and integration of the statistics page into the game management process should be improved. Actively alerting creators about the status of their game's data can be one future direction to address this. Additionally, more advanced options should be provided to share the data with others and analyze it.

Both managers and employees, spanning all organizational divisions, including the non-technical ones, engaged in creating and playing games. Games appeared to be used as a guilt-reduced alternative for employees to escape work when they needed to take a break or make a context-switch, by challenging themselves, competing with colleagues, learning more about the business, and contributing knowledge. Future versions of G4C should provide integrated dissemination and engagement means, for instance through email updates or enterprise social media. Additionally, supporting mobile games would open more opportunities for employees to engage in games when they have spare time on the go.

We also saw that playing games at work can be difficult for some employees, both because they need to justify the allocated time and because some feel particularly pressured to succeed as their real identity is exposed alongside their score. In their work about "mandatory fun", Mollick and Rothbard [21] found that legitimacy plays a key role in employee's consent to enterprise games. Crowdsourcing games should therefore convey to players their value to the business and give them the feeling they contribute each time they play.

Our findings from inspecting the characteristics of the 34 games indicate that games with more questions enjoy higher popularity; instructions that are not too short and not too long (11-20 words) yield higher players' contribution; and while leisure games enjoy higher popularity, business games yield higher players' contribution. Even though some of these factors must be determined by the specific needs of a game, the platform can aid creators by indicating the recommended length of instructions or encouraging uploading more questions. With regards to the instruction length, our findings are based on one specific group partitioning with some level of arbitrariness and thus the exact borders of the recommended range should be further validated.

G4C strives to create an ecosystem of creators, players, and games. We found that employees are willing to share their created games – only four of them were kept private. The gaming activity sidebar, presenting activity across different games by different criteria, was indicated to contribute to the sense of community. Several players created a game after learning how simple it can be from their playing experience. Other creators played others' games to compare with their own. More ideas were suggested to further promote the sense of community, for

example by encouraging more interaction between users, allowing competition between groups, or recommending similar players.

G4C enables employees to create a certain type of human computation games. It provides them a ready-made user interface and thus spares them the need to become game designers, but rather focus on the content and rules of the game. Aarseth [1] defines three dimensions that characterize computer games, which may have different weights in different games: the gameplay (players' actions and strategies), the game structure (rules), and the game world (fictional content, level topology, textures). In our case, many of the game "world" elements are already provided for the authors, leaving more weight to the game's content and structure. In the video game domain, the concept of "modders" describes game moderators that contribute new content, such as skills, missions, scenarios, maps, or episodes, to a game. Gee [14] states that good games allow players to be producers and not just consumers. He argues that such players' actions co-create the game's world, which makes them as much designers of the game as the original creator. The role of a game creator in G4C has a lot in common with this role of "modder" or "contributor" in the video game field.

This work particularly contributes to the area of enterprise crowdsourcing. Its main novelty lies in bringing together crowdsourcing and game creation capabilities in one platform deployed in a workplace setting. While G4C only enables the creation of one game type, we observed that this type on its own leads to the creation of different games with diverse purposes. We believe this is an important step in making enterprise crowdsourcing more accessible; many of the game creators in our study were non-technical employees and some of them indicated the low entry barrier as their key reason for "joining the club". We also pointed at design opportunities, stemming from the fact that such games are perceived as a "legitimate" way to take a break from work, which comes with its own challenges, such as the need to perform publicly in front of one's co-workers. We also found that many game creators invoked their social network within the organization to encourage play. This implies that giving certain employee, such as those with higher connectivity within the organizational network, a creative role with higher engagement, can help promote enterprise crowdsourcing.

Our study is influenced by the characteristics of the specific organization in which it was conducted. We believe that many of the games created using G4C address general challenges in large organizations, such as volunteer recruiting, org chart building, and education about a department or a customer. Still, culture changes from one organization to another and may substantially affect the way crowdsourcing games are created and played. Additionally, our study is based on a relatively modest dataset of 34 games and 339 players. While these

provide a good initial sense of the created games' diversity and their characteristics, quantitative findings should be treated with care given the small number of analyzed games. Future research should further explore the concepts presented here, in other organizational environments and with more participants.

CONCLUSION

We presented Games for Crowds, a novel crowdsourcing game platform that facilitates the creation and playing of enterprise human-computation games. G4C creates a layer of abstraction over the game generation process by using the wizard design pattern and exposing a set of parameters needed to create a game. Our evaluation provides an overview of the platform's use in the enterprise and suggests future enhancements that can be applied in similar systems, as enterprise crowdsourcing is expected to grow and become more ubiquitous.

ACKNOWLEDGMENTS

We thank Artem Barger, Maya Barnea, Tal Daniel, Ohad Greenshpan, Michal Jacovi, Vladimir Soroka, and Itai Turbahn for their help in evolving and developing G4C over the years.

REFERENCES

1. Aarseth, E. 2003. Playing Research: Methodological approaches to game analysis. *Proc. DAC '03*, 28-29.
2. Abrahamson, S., Ryder, P., & Unterberg, B. 2013. *Crowdstorm: the future of innovation, ideas, and problem solving*. John Wiley & Sons.
3. Allio, R. J. 2004. CEO interview: the InnoCentive model of open innovation, *Strategy & Leadership* 32,4 (2004), 4-9.
4. Ambati, V., Vogel, S., & Carbonell J. 2010. Active learning and crowd-sourcing for machine translation. *Proc. LREC'10*.
5. Antin, J. & Shaw, A. 2012. Social desirability bias and self-reports of motivation: a study of Amazon Mechanical Turk in the US and India. *Proc. CHI '12*, 2925-2934.
6. Bernstein, M., Chi, E., Chilton, L., Hartmann, B., Kittur, A., & Miller, R. 2011. Crowdsourcing and human computation: systems, studies and platforms. *Proc. CHI EA '11*, 53-56.
7. Buhrmester M., Kwang T., & Gosling S.D. 2011. Amazon's Mechanical Turk: a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science* (2011).
8. Chen, K., Chang, C., Wu, C., Chang, Y., & Lei, C. 2010. Quadrant of euphoria: a crowdsourcing platform for QoE assessment. *Network, IEEE* 24, 2 (March 2010), 28-35.
9. 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.
10. Doan, A., Ramakrishnan R., & Halevy, A. Y. 2011. Crowdsourcing systems on the world-wide web. *Commun. ACM* 54, 4 (2011), 86-96.
 11. Dugan, C., Muller, M., Millen, D. R., Geyer, W., Brownholtz, B., & Moore, M. 2007. The dogear game: a social bookmark recommender system. *Proc. GROUP '07*, 387-390.
 12. Duncan, S.C. 2011. Minecraft, beyond construction and survival. *Well Played Journal* 1(1), 1-22.
 13. Evanini, K., Higgins, D., & Zechner, K. 2010. Using Amazon Mechanical Turk for transcription of non-native speech. *Proc. NAACL HLT '10*, 53-56.
 14. Gee, J.P. 2003. What video games have to teach us about learning and literacy. *Computers in Entertainment*, 1(1), 20.
 15. Guy, I., Perer, A., Daniel, T., Greenshpan, O., & Turbahn, I. 2011. Guess who? enriching the social graph through a crowdsourcing game. *Proc. CHI '11*, 1373-1382.
 16. Howe, J. 2006. The rise of crowdsourcing, *Wired* 14, 6 (June 2006).
 17. Kittur, A., Chi, E. H., & Suh, B. 2008. Crowdsourcing user studies with Mechanical Turk. *Proc. CHI '08*, 453-456.
 18. Lakhani, K. R., Jeppeson, L. B., Lohse, P. A. & Panetta, J. A. 2007. The value of openness in scientific problem solving. *HBS working paper* (Feb. 2007).
 19. Little, G., Chilton, L. B., Goldman, M., & Miller, R. C. 2009. TurKit: Tools for iterative tasks on Mechanical Turk. *Proc. HCOMP '09*, 29-30.
 20. Mehm, F., Reuter, C., Göbel, S., & Steinmetz, R. 2012. Future trends in game authoring tools. *Proc. ICEC '12*, 536-541.
 21. Mollick, E.R. & Rothbard, N. 2013. Mandatory fun: gamification and the impact of games at work. *The Wharton School Research Paper Series*.
 22. Muller, M., Ehrlich, K., Matthews, T., Perer, A., Ronen, I., & Guy, I. Diversity among enterprise online communities: collaborating, teaming, and innovating through social media. *Proc. CHI '12*, 2815-2824.
 23. Parent, G. & Eskenazi, M. 2010. Clustering dictionary definitions using Amazon Mechanical Turk. *Proc. NAACL HLT '10*, 21-29.
 24. Quinn, A. J. & Bederson, B. B. 2011. Human computation: a survey and taxonomy of a growing field. *Proc. CHI '11*, 1403-1412.
 25. Sorenson, N. & Pasquier, P. Towards a generic framework for automated video game level creation. *Proc. EvoApplications*, 130-139.
 26. Surowiecki, James. The wisdom of crowds. 2005. *Anchor Books*.
 27. Tidwell, J. Designing interfaces. 2005. *O'Reilly*.
 28. Torrente, J., Del Blanco, Á., Cañizal, G., Moreno-Ger, P., & Fernández-Manjón, B. 2008. < e-Adventure3D>: an open source authoring environment for 3D adventure games in education. *Proc. ACE '08*, 191-194.
 29. Viegas, F.B, Wattenberg, M, Van Ham, F, Kriss, J., & McKeon, M. 2007. ManyEyes: a site for visualization at internet scale. *IEEE TVCG* 13, 6 (Nov. 2007), 1121-1128.
 30. Von Ahn, L. & Dabbish, L. 2008. Designing games with a purpose. *Commun. ACM* 51, 8 (Aug. 2008), 58-67.
 31. Von Ahn, L. & Dabbish, L. 2004. Labeling images with a computer game. *Proc. CHI '04*, 319-326.
 32. Von Ahn, L., Kedia, M., & Blum, M. 2006. Verbosity: a game for collecting common-sense facts. *Proc. CHI '06*, 75-78.
 33. Von Ahn, L., Liu, R., & Blum, M. 2006. Peekaboom: a game for locating objects in images. *Proc. CHI '06*, 55-64.
 34. Von Ahn, L., Maurer, B., McMillen, C., Abraham, D., & Blum, M. 2008. Recaptcha: human-based character recognition via web security measures. *Science* 321, 5895 (2008), 1465-1468.
 35. Vukovic, M. 2009. Crowdsourcing for enterprises. *Proc. Services '09*, 686-692.
 36. Westcott, E. 2011. Crafting play: little big planet. *Loading* 5, (2011) 90-100.