A Taxonomy of Distributed Human Computation

Alexander J. Quinn, Benjamin B. Bederson
University of Maryland
Department of Computer Science
Human-Computer Interaction Lab
College Park, MD 20742 USA
{aq, bederson}@cs.umd.edu

ABSTRACT
Distributed Human Computation (DHC) holds great promise for using computers and humans together to scaling up the kinds of tasks that only humans do well. Currently, the literature describing DHC efforts so far is segmented. Projects that stem from different perspectives frequently do not cite each other. This can be especially problematic for researchers trying to understand the current body of work in order to push forward with new ideas. Also, as DHC matures into a standard topic within human-computer interaction and computer science, educators will require a common vocabulary to teach from. As a starting point, we offer a taxonomy which classifies and compares DHC systems and ideas. We describe the key characteristics and compare and contrast the differing approaches.

Author Keywords
Distributed human computation, human participation, artificial artificial intelligence.

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

INTRODUCTION
An enormous potential exists for solving certain classes of computational problems through rich collaboration between humans and computers. Currently, countless open computer science problems remain in artificial intelligence, natural language processing, and computer vision. These three fields, which we will collectively refer to as "AI", roughly define the subset of computer science problems which can often be readily solved by humans. Humans are innately good at creativity, visual processing, innovative problem solving, planning, analysis, and verbal tasks. The reason for doing research in AI is because of the hope that we can one day automate those kinds of tasks and perform them with the high speed and low cost that only computers have to offer. AI offers one set of solutions to these kinds of hard problems. Distributed human computation (DHC), sometimes called artificial artificial intelligence, offers another.

Why would computer scientists work so hard to solve problems with humans which can be addressed with computers alone? The answer, of course, is that while computers have the potential of achieving much greater speed and throughput at lower cost, most automated solutions to hard problems are of unsatisfactory quality. (And when computers get better, we can just look to harder problems.) DHC offers the possibility of combining humans and computers to offer another point in the tradeoff space where solutions are faster than individual human efforts, and quality is at least as good as human efforts – sometimes even better due to diversity and network effects of human participation.

DHC is the strategy of combining the strengths of computers and humans by delegating parts of the problem to large numbers of humans connected via the internet – usually spread out geographically. A computer system has global knowledge of the problem to be solved and forms small subproblems that take advantage of humans' special abilities. The human participants solve some problems and the answers are usually checked and aggregated by the computer system. With DHC, humans help computers to do computer-type tasks. This is different from computer-supported cooperative work, in which computers help humans do human-type tasks. This is also different from the human computers of the 17th century who, in lieu of computers, performed mathematical computations as their primary occupation [15]. A group of these human computers later became the first professional programmers of the ENIAC, using their global knowledge of war-related math problems to isolate some problems that could be delegated to the ENIAC - an ironic converse of DHC where humans delegated subproblems to computers.

To understand the potential of DHC, let us look at a specific example. Consider the task of searching a large number of satellite photos covering thousands of square miles of ocean to find a particular missing boat. Completely automated solutions are not up to the task, and there are not enough experts to do this in a timely fashion when someone has gone missing. But when the well-known computer scientist Jim Gray went missing in early 2007, people around the web voluntarily examined over 560,000 images covering 3,500 square miles of ocean – tagging suspect images for experts to examine. Sadly, he was never found, but this approach was a testimony to the potential of harnessing distributed human effort [8].

Another task, perhaps less exciting, but also very valuable is to label images to make them easier to search. The ESP Game, developed by von Ahn [47] solves the problem by creating an online game where players label randomly selected images in order to get points. They have collected over 50 million image labels collected by 200,000 contributors over the past 3 years. Scaling things up further, Google licensed the technology and created their own game which utilizes both the player's perceptual abilities and their vast system resources to create a hybrid solution [13].

In recent years, DHC has emerged as a subfield within human-computer interaction, as evidenced by the first Human Computation Workshop in 2009 [27]. Researchers investigating DHC have come from areas of computing as diverse as cryptography [55], human-computer interaction [6,21], business [41], genetic algorithms [25], artificial intelligence [46], and digital art [22]. Although the approaches they have taken have been similarly diverse, all strive to leverage the strengths of humans to enhance the capabilities of computer systems to perform tasks previously attempted using AI algorithms alone.

It turns out that there are actually many different kinds or genres of DHC. Some approaches are more collaboration oriented and some are more computationally oriented. But all use networks of humans that are managed by computers to accomplish things that could never be done by either computers or small groups of humans alone. Given these many genres, it is not surprising to find that there is a range of vocabulary that people use to describe their systems. And given the different fields that the work has been done it, it is also common for groups to not even be aware of what groups in other fields are doing in DHC.

This paper takes the position that in order for DHC to mature as a research area, the various DHC systems can and should be classified in a common manner. We propose the following dimensions to help characterize the approaches we have found:

- **Motivation –** Why do contributors choose to help? Examples: fun, altruism, pay, reputation, or implicit (doing something else and helping unintentionally).
- **Quality –** How does the system cope with the possibility of fraud or mistaken answers to ensure some level of quality in the solution to the overall problem?
- **Aggregation –** How does the system combine or otherwise use the local computations made by individual participants to solve the global problem?
- **Human skill –** What is the special skill that humans possess that makes them valuable in the system?
- **Participation time –** What is the minimum amount of time a user must invest to participate?
- **Cognitive load –** Independent of the time spent, how much mental energy does the task take?

Definition

With all this in mind, we now define DHC for the purposes of this paper as *"systems of computers and large numbers of humans that work together in order to solve problems that could not be solved by either computers or humans alone"*. We recognize that this definition is somewhat vague. Are social networks such as Facebook and Twitter examples of DHC? The answer is that it depends on what they are being used for. If they are just being used for entertainment, the answer is no. If, on the other hand, they are used to, say, identify global trends in real time, then the answer is yes. DHC is a big field, and there are many domains of its application.

GENRES

To get a better understanding of DHC, we now dive into the types or genres or DHC applications. This shows how broadly DHC is being applied, and sets the stage for us to characterize DHC systems in general.

Games With a Purpose

A popular approach to motivating volunteers is to create a game that requires the player to perform some computation in order to get points or succeed. The idea is that since people play a lot of games online, anyway, it may be possible to divert that energy for some particular purpose, provided the problem at hand can be expressed as a game. However, when creating a game to serve a DHC purpose, it is important to prove that the game is correct and will yield high quality results, much like designing an algorithm [48]. Luis von Ahn, who coined the term "Games with a Purpose" and has done the most notable work in this genre, with six games available at www.gwap.com.

The defining factor of Games with a Purpose is that the volunteers are motivated by the fun the game provides. Examples of Games with a Purpose include the ESP Game [49], Peekaboom [54], Verbosity [53], Tag-A-Tune [26], FACTory [46], and foldit (www.fold.it).

Mechanized Labor

When Crowdsourcing involves monetary pay, the dynamics of motivation change and thus, we categorize it in a separate genre. The most notable example is Amazon's Mechanical Turk [34], which allows developers to write programs that automatically assign very small tasks to a network of workers who complete the tasks at the Mechanical Turk web site. Other examples include ChaCha [4] and the Cyphermint PayCash identification system [37,11].

Mechanized labor is distinguished by its use of paid volunteers to do explicitly defined tasks that may take a small or moderate amount of time to complete. While not a strict requirement, Mechanized Labor is more likely than CrowdSourcing to involve some sort of loose expectation to continue the work. However, these systems rarely involve full-fledged employment or a firm obligation to continue.

Wisdom of Crowds

Consider a game where hundreds of people try to guess the number of jelly beans in a large jar. It turns out that under normal circumstances, the average of the guesses will be very close to the actual count. In his book by the same title [43], Surowiecki argues that a decentralized, disorganized group of people, thinking independently can use "crowd intelligence" to make judgments that would be very difficult to make using only one person, provided that their individual judgments are made independently and are aggregated correctly. Several online polling web sites and prediction markets harness this concept to not only determine a group opinion, but to predict the future (e.g. Ask500People [1], News Futures - www.newsfutures.com, and Iowa Electronic Markets - http://www.biz.uiowa.edu/iem/index.cfm.)

CrowdSourcing

The term "CrowdSourcing", first coined in a Wired magazine article written by Jeff Howe [19], and the subject of his book [20], refers to the displacement of usual internal labor by soliciting unpaid help from the general public, usually motivated by curiosity or serendipity while browsing the web (e.g. online product reviews), but sometimes part of a larger effort, as in the case of Stardust@home [35]. Some would also classify question answering services (e.g. Askville www.askville.com and Aardvark – www.vark.com) as crowdsourcing, especially when the emphasis is on producing large numbers of answers to freeform questions, and not on the social interaction, thus distinguishing them from regular online discussion forums.

CrowdSourcing is distinguished by its use of unpaid volunteers to do explicitly defined tasks that take a small amount of time to complete. Volunteers generally have no obligation to continue.

Dual-Purpose Work

ReCAPTCHA brought awareness to the clever idea of translating a computation into an activity that many people were already doing frequently. ReCAPTCHA utilizes the computation people do solving CAPTCHAs in order to transcribe old scanned books and newspapers that cannot be processed by optical character recognition due aging [39,38]. Alternatively, it is sometimes possible to take advantage of work that people have already done. For example, Google's core PageRank search algorithm uses the link network and details of web pages to decide which pages should be returned upon search [36].

Dual-Purpose Work is distinguished by its use of work that people are already doing or have already done, irrespective of the goal of the particular DHC system.

Grand Search

Grand Search is different from the other genres because instead of aggregating or collecting the many task results, the goal is to find the one that solves the problem. Given a large number of images or something else to search, contributors each search through a modest number trying to find the one (or a few) with some special characteristic. For example, Grand Search was used in the above-mentioned search for the computer scientist Jim Gray. The Space Sciences Laboratory at the University of California at Berkeley previously used Grand Search to search for tiny matter from space as part of the Stardust@home project [35]. The particles had been mixed into an aerogel collector from the Stardust spacecraft. Contributors searched through photographs of the aerogel for traces of the particles. This recognition problem was much too difficult for computer vision algorithms or even untrained humans. Therefore, participants had to complete an online training program to learn how to identify the tiny particles before they could contribute.

Grand Search is distinguished by its use of volunteers to search through a large sample space, usually composed of images. The only contributions that are of value are the one(s) that contain the target (e.g. photo of Jim Gray or trace of a particle).

Human-based Genetic Algorithms

Kosorukoff has taken the approach of using DHC to execute genetic algorithms. The idea is that humans can contribute solutions to problems and subsequent participants perform other functions such as initialization, mutation, and recombinant crossover [25]. A loose interpretation of this idea might also include the collaborative online encyclopedia, Wikipedia [56], based on the viewpoint that the text of an article is like a genetic sequence and edits are like mutations which survive or perish based on evaluations of readers and participants.

The defining factor of this approach is that solutions consist of a sequence of small parts and that they evolve in a way that is controlled by human evaluation. Another example is the Free Knowledge Exchange (FKE) [8].

Knowledge Collection from Volunteer Contributors (KCVC)

KCVC tries to advance artificial intelligence research by using humans to build large databases of common sense facts. The idea is that humans, by way of their child development and adult lives, acquire a great deal of common sense knowledge (i.e. "People cannot brush their hair with a table."). Large databases of such knowledge, such as Cyc [28], have collected such information into knowledge bases and ontologies using data mining. Several efforts have also attempted to use volunteer contributors to provide the data, either by using games (e.g. the FACTory [46], Verbosity [53], 1001 Paraphrases [8]) or by volunteer knowledge gathering activities (e.g. Learner [7]). The field of KCVC, summarized in [9], has developed into a subfield of its own with dedicated symposia sponsored by the Association for the Advancement of Artificial Intelligence (AAAI).

KCVC is primarily characterized by its purpose of collecting common sense data for artificial intelligence research and applications. All known examples have been based on some sort of sentence completion model.

DIMENSIONS

The common denominator among most DHC systems is that they rely on humans to provide units of work which are then aggregated to form the basis of some other computations that are performed by a computer. This gives rise to a common set of problems that each must address. We describe these dimensions of DHC approaches and enumerate a few of the typical solutions (summarized in Table 1).

Motivation

One of the most difficult challenges in any DHC system is finding a way to motivate people to participate. This problem is somewhat alleviated by the fact that most DHC systems rely on networks of unconnected people connected to computers sitting in their homes. Therefore it is not necessary to motivate them to go anywhere or do anything too far out of their ordinary lives. However, since the computations frequently involve small unit tasks that do not directly benefit the participants, people have to have some kind of motivation to want to participate.

We characterize the approaches systems use to engage participants based on how internally motivated they are. Thus, starting out with the most external motivation, the levels of motivation we have identified are:

- **Pay**. This is easy because if you have a little bit of money to pay contributors, you can easily translate that money into some amount of computation. The problem with this is that when there's money involved, people are more likely to try to cheat the system to increase their overall rate of pay. And, because participants are usually anonymous, they may be more likely to do something dishonest than they would if they were working in person.
- **Altruism.** Do good. It may sound easy to trust in people's desire to help, but this requires that the participants actually think the problem being solved is interesting and important. Effectively, this requires that you be doing something that most people are already willing to help with. In most cases, the number of users is likely to be limited.
- **Fun.** By creating a game that people choose to play because they enjoy it, the participants no longer need to be paid. They are also likely to play for a fairly long time. However, this is also difficult because it is challenging to turn many computational tasks into a game that is actually fun to play.

Dimension / Genre	Motivation	Quality	Aggregation	Human Skill	Participation Time (minimum)	Cognitive Load
Games With a Purpose	fun	forced agreement	training data, collect facts	language, vision, etc.	2-10 mins	high
Mechanized Labor	pay	expert review, *	unit tasks, *	*	*	*
Wisdom of Crowds	*	redundancy, statistical	statistical	*	< 2 mins	medium
Crowdsourcing	*	*	collect facts	*	*	medium
Dual-Purpose Work	implicit	*	*	*	net zero	net zero
Grand Search	altruism, *	expert review	find goal	language, vision, etc.	< 2 mins	low
Human-based Genetic Algorithms	*	redudancy	genetic algorithm	*	*	*
Knowledge Collection By Volunteer Contributors	altruism, *	redundancy	collect facts	common knowledge	< 2 mins	low

Table 1: Genres classified by the six DHC dimensions. Asterisks (*) note dimensions that are not constrained for that genre. For example, with mechanized labor (i.e. Amazon Mechanical Turk), the core skill could be vision, language, or anything else, depending on the particular task. However, for a particular application, all dimensions will be fixed.

- **Implicit.** If you're lucky enough to find a way to embed the task in regular activities. it is sometimes possible to make the computation a natural part of something else users were already doing as part of their everyday lives. Sometimes, a problem you want to solve can be formulated in a way that can be made to fit to people's activities.

Quality

Even if the users are motivated to participate, they may try to cheat or sabotage the system. The means of discouraging cheating usually depends on the means of motivation. Although many schemes for ensuring quality are possible, we focus on a few that are common in practice now:

- **Forced agreement.** In the case of a game, it is often possible to build in some feature into the game that will check the users' results and make sure they match with other participants. Other means are possible, but they depend on the specifics of the game.
- **Economic models.** When money is used as a motivating factor, it may be possible to build in probability-based economic models that guarantee that on the average case, a user does no better than breaking even if they cheat [11,12].
- **Defensive task design.** More practically, a range of solutions have been developed recently to improve the accuracy of paid systems like Amazon Mechanical Turk [34]. One approach is to design the tasks so it is no easier

Application	Genre	Motivation	Aggregation	Quality	Human Skill	Part. Time (min.)	Cognitive Load
ESP Game []	Games With a Purpose	fun	training data, collect facts	forced agreement	language, vision, etc.	2-10 mins	high
ChaCha web search	Mechanized Labor	pay	unit tasks	expert review	web searching	2-10 mins	high
Sheep Market (with Amazon Mechanical Turk)	Mechanized Labor	pay	unit tasks	(unknown)	art (drawing)	< 2 mins	medium
Iowa Electronic Markets	Wisdom of Crowds	fun *AND* pay	statistical	statistical	domain knowledge of politics	>10 mins	high
Monolingual translation []	Crowd Sourcing	altruism	collect facts	forced agreement	language	2-10 mins	high
reCAPTCHA	Dual-purpose work	implicit	collect facts	forced agreement *AND* redundancy	vision (reading)	< 2 mins	low
CDDB (public CD album database)	Dual-purpose work *AND* Knowledge Collection By Volunteer Contributors	implicit *AND* altruism	collect facts	redudancy	OCR	2-10 mins	low
Find Jim Gray	Grand Search	altruism	find goal	expert review	vision	< 2 mins	low
Free Knowledge Exchange	Human-Based Genetic Algorithms	implicit *AND* altruism	genetic algorithm	redundancy	common knowledge	2-10 mins	medium

Table 2: Some DHC applications are classified. Each application is an instance of one of the genres in Table 1. These were chosen for variety and/or notability.

to cheat than to do the task. Another is to use multi-level solutions where a second set of workers checks the result of the first set.

- **Redundancy.** By spending more money or finding more contributors, you can have each task done by multiple workers, and use a voting scheme to identify good answers. An added bonus is that once schemes are put in place to identify poor human performers, all of their work can be removed since it has been shown that a significant part of the bad work is done by a small number of human workers.

Finally, if the particular application does not require that the individual results be reliable, it may not be important if users cheat, as long as cheaters to not collude with each other. For example, if you were to require a certain number of redundant answers to the same question before you trust the answer, a cheater could only be dangerous if he colluded with other cheaters to achieve the required number of redundant answers.

Quality checking methods can be summarized as follows:

- **Statistical** – Either filter or aggregate the data in such a way as to remove the effect of irrelevant contributions. For example, the system might use standard deviation to remove outliers.
- **Redundant work –** The work is done by multiple sets of participants, and the combined work is analyzed automatically – typically keeping answers that were given by multiple participants.
- **Multilevel review –** A first set of participants do the work, and then a second set of participants review and rate the quality of their work.
- **Expert review -** A trusted expert reviews contributions for relevance and apparent accuracy. For example, with Amazon Mechanical Turk, people who post tasks may approve the quality of work and withhold payment, as needed.
- **Forced agreement –** Epitomized by the ESP game, two or more contributors work together and the answer is not accepted unless they agree.
- **Automatic check -** Although not seen in actual DHC systems, it is perfectly conceivable that a system could be built to solve problems for which answers can be easily checked by a computer. For example, in AI planning, many problems are very difficult for a computer to solve, but the solutions can be checked instantly.
- **Reputation system -** In some systems, users may be motivated to provide quality answers by a reputation system. If answers are poor, future answers will be discounted and incentives withheld. Some economic systems try to make cheating economically infeasible [3].
- **None -** There may be some cases in which checking is not necessary.

Aggregation
Part of the system's role in a DHC is to combine all of the contributions to solve the global problem. The means of doing this partly determines what kinds of problems a system or strategy can be applied to. The principle aggregation methods are listed below.

- **Knowledge base** – As in the KCVC genre, a knowledge base of discrete facts, or sometimes a hierarchical ontology, are built. A contribution may either add a new fact or improve quality by correcting, refuting, confirming existing facts in the knowledge base.
- **Statistical** – Usually involving a simple average or median, this method is the hallmark of the Wisdom of Crowds genre.
- **Grand search** – Several projects (e.g. Stardust@home [35]) have used large numbers of volunteers to sift through photographs or videos, searching for some desired scientific phenomenon, person, or object. In this case, many computations are performed, but few if any result in a meaningful conclusion.
- **Unit tasks** – Some DHC systems do not involve aggregation at all, but simply use humans to perform a large number of small tasks which are independent of one another. Examples include question answering systems such as ChaCha [4] and review systems such as www.rottentomatoes.com where people write personal reviews in addition to ratings.

Human Skill
Most DHC systems take the place of computer vision, natural language processing, or artificial intelligence algorithms that might do the same work, but inadequately. This is made possible because the human contributors have certain skills. For example, a system that asks contributors to label photographs is leveraging their human sight in place of a computer vision object recognition system. In addition to vision, language understanding, language communication, reasoning, and common knowledge are often leveraged by DHC systems. Common knowledge simply refers to things all, or nearly all humans know about the world – the usual target of Knowledge Collection from Volunteer Contributors (KCVC) systems.

Participation Time
The minimum time the user must spend to contribute at all is an important facet of DHC systems. If it is long, volunteer users may hesitate to participate. On the other hand, some kinds of tasks may require more time to complete. Thus, the minimum participation time is one of the factors that constrains the kinds of problems a DHC application can do. For purposes of this paper, since we do not know the exact time the various examples require, we use rough estimates and fit to the following four categories: net zero (takes no more time than not participating), <2 minutes, 2-10 minutes, and >10 minutes.

Cognitive Load

The cognitive load of a task may also affect contributors' willingness to help, but it can also be independent of the participation time. For example, reading a long passage and searching for some text might take a long time but require very little mental energy. On the other hand, watching a fast-moving video and searching for a pattern would likely take a lot of energy. For purposes of this paper, we will define three levels of cognitive load.

- **Low –** Requires only passive processing and can be done at any speed. No time sensitive response or attention is necessary.
- **Medium -** Requires simple (i.e. single-step) problem solving, creative thinking, or domain knowledge.
- **High -** Requires time sensitive response or multistep problem-solving, or imposes some other significant cognitive load.

STEPS TO BROADER DHC APPLICATION

Based on a more comprehensive understanding of existing work in DHC, we are now ready to propose some future directions for better utilizing this to solve computational problems. We have observed that previous work was often disconnected and limited in scale, compared to collaborative mediums and social networking sites such as Wikipedia and MySpace. We propose the following goals for future work.

Develop innovative aggregation methods. By exploring new ways of aggregating the results generated by users, it may be possible to create much richer collaborations between humans and computers.

The related work demonstrates that aggregation methods used so far have been sparse. However, the Wisdom of Crowds book gives rise to the idea that perhaps new statistical methods could be used solve new kinds of problems. For example, the simple average may not work for all kinds of estimation and judgment tasks because participants may be biased. It may be useful to measure that bias and correct for it. Also, there may be richer methods of incorporating past results into the parameters of the decide function.

Look to new platforms for facilitating collaboration. Virtually all work to date has utilized humans only in the setting of working independently at a computer using a web interface. New settings may bring new opportunities for utilizing DHC.

Perhaps greater participation could be enabled by allowing participants to do tasks from a cell phone or other mobile device. This might also enable participants to contribute in ways that depend on the user's physical setting. For example, participants might use a GPS-enabled device to plot points on a city map, which would then be aggregated and processed by the computer. Other opportunities may lie in allowing humans to collaborate on problems in real time or in better managing situations where many people are working on small tasks disconnected from coordinating web site (or other computing resource) so that progress will be guaranteed, even if participants work at different speeds or abandon their tasks.

Build General Purpose Infrastructures. Mechanized Labor systems today generally build on the infrastructure of Amazon Mechanical Turk. However, there is no general infrastructure that is used for other genres of DHC systems. Instead, most applications are hand-crafted systems. In order to support more creative explorations of this space, it must be easier to do so, and a good way to support this goal is to have platforms that directly support the building of the many different kinds of DHC systems.

Sort out labor and tax issues. If DHC participation comes to be viewed as work that is creating some sort of product, it may well become subject to legal disputes about what labor laws apply and whether barter tax should be paid if participants receive some quantifiable digital privilege as a reward for participation.

CONCLUSION

We have presented a taxonomy with which DHC systems can be described, understood, and compared. As the field of DHC matures, some common reference point will be necessary to facilitate discussion and make sense of the burgeoning body of literature.

Other consequences of this work will depend on the reader.

Researchers. For researchers, understanding the DHC landscape and the characteristics of solutions to date also reveals where the gaps lie. In looking at the examples in this paper and others from recent conferences, there may still be effective strategies for using humans to solve computational problems are not being explored. Furthermore, more problems exist that could leverage the abilities of humans, if only there were new ways to partition and package larger computational problems and enable users to help on *their* terms.

Practitioners. The possibilities for using DHC to solve computational problems are vast. When considering potential system designs, consider the full range of flexibility. For example, although many companies recruit free contributors through their websites, there may be cases where unit payments, in the form of mechanized labor, are the best solution. Similarly, for some projects, ever may be tolerable where other projects may justify hiring or otherwise recruiting expert reviewers or redundant contributors to reduce the possibility of contamination with bad information, either unintentionally or intentionally by a malicious contributor.

ACKNOWLEDGEMENTS

We gratefully acknowledge Jimmy Lin, Chang Hu, and Pete Kelleher for discussing these ideas with us at various stages and helping bring them to their current form.

REFERENCES

1. Ask500People. http://www.ask500people.com.
2. Askville. http://askville.amazon.com.
3. Burnham, T., Sami, R. A Reputation System for Selling Human Computation. In *Proc. of Knowledge Discovery and Data Mining*, Workshop on Human Computation HCOMP 2009, ACM.
4. ChaCha human-powered search engine. http://www.chacha.com.
5. Chen, M., Hearsty, M., Hong, J., and Lin, J. 1999. Cha-Cha: a system for organizing intranet search results. In Proceedings of the 2nd Conference on USENIX Symposium on internet Technologies and Systems - Volume 2 (Boulder, Colorado, October 11 - 14, 1999). USENIX Association, Berkeley, CA, 5-5.
6. Chi, E. H. 2009. Augmented social cognition: using social web technology to enhance the ability of groups to remember, think, and reason. In *Proceedings of the 35th SIGMOD international Conference on Management of Data* (Providence, Rhode Island, USA, June 29 - July 02, 2009). C. Binnig and B. Dageville, Eds. SIGMOD '09. ACM, New York, NY, 973-984. DOI=http://doi.acm.org/10.1145/1559845.1559959
7. Chklovski, T. 2003. Learner: a system for acquiring commonsense knowledge by analogy. In Proceedings of the 2nd international Conference on Knowledge Capture (Sanibel Island, FL, USA, October 23 - 25, 2003). K-CAP '03. ACM, New York, NY, 4-12. DOI= http://doi.acm.org/10.1145/945645.945650
8. Chklovski, T. 2005. Collecting paraphrase corpora from volunteer contributors. In Proceedings of the 3rd international Conference on Knowledge Capture (Banff, Alberta, Canada, October 02 - 05, 2005). K-CAP '05. ACM, New York, NY, 115-120. DOI= http://doi.acm.org/10.1145/1088622.1088644
9. Chklovski, Timothy & Gil, Yolanda. Towards Managing Knowledge Collection from Volunteer Contributors, Proceedings of 2005 AAAI Spring Symposium on Knowledge Collection from Volunteer Contributors (KCVC05) (2005)
10. Edmunds, P. 2005. SwarmSketch. http://www.swarmsketch.com/.
11. Gentry, C., Ramzan, Z., and Stubblebine, S. 2005. Secure distributed human computation. In Proceedings of the 6th ACM Conference on Electronic Commerce (Vancouver, BC, Canada, June 05 – 08, 2005). EC '05. ACM, New York, NY, 155-164. DOI= http://doi.acm.org/10.1145/1064009.1064026
12. Goldberg, A. V. and Hartline, J. D. 2001. Competitive Auctions for Multiple Digital Goods. In Proceedings of the 9th Annual European Symposium on Algorithms (August 28 - 31, 2001). F. M. Heide, Ed. Lecture Notes In Computer Science, vol. 2161. Springer-Verlag, London, 416-427.
13. Google Image Labeler. http://images.google.com/imagelabeler/.
14. Greg Little, Lydia B. Chilton, Robert C. Miller, and Max Goldman. "TurKit: Tools for Iterative Tasks on Mechanical Turk." HCOMP, 2009, to appear.
15. Grier, D.A. *When computers were human*. Princeton Univ Pr, 2005.
16. Help Find Jim (web site). http://www.helpfindjim.com
17. Ho, C.J., Chang, T.H., Lee, J.C., Hsu, J.Y., and Chen, K.T. KissKissBan: A Competitive Human Computation Game for Image Annotation.
18. Hollan, J., Hutchins, E., and Kirsh, D. 2000. Distributed cognition: toward a new foundation for human-computer interaction research. ACM Trans. Comput.-Hum. Interact. 7, 2 (Jun. 2000), 174-196. DOI= http://doi.acm.org/10.1145/353485.353487
19. Howe, Jeff, The Rise of Crowdsourcing. *Wired Magazine*, Issue 14.06, June 2006.
20. Howe, Jeff. *Crowdsourcing: Why the Power of the Crowd is Driving the Future of Business*. Crown Business (2008). ISBN # 978-0307396204.
21. Hu, C. 2009. Collaborative translation by monolingual users. In *Proceedings of the 27th international Conference Extended Abstracts on Human Factors in Computing Systems* (Boston, MA, USA, April 04 - 09, 2009). CHI EA '09. ACM, New York, NY, 3105-3108. DOI= http://doi.acm.org/10.1145/1520340.1520438
22. Koblin, A. Collaborative Drawing. http://www.collaborativedrawing.com. Accessed 12/2007.
23. Koblin, A. The Sheep Market. http://www.thesheepmarket.com. Accessed 12/2007.
24. Kosorukoff A. Human based genetic algorithm. IlliGAL report no. 2001004. 2001, University of Illinois, Urbana-Champaign. Available at http://www.illigal.uiuc.edu/web/technical-reports/2001/page/2/.
25. Kosorukoff, A. & Goldberg, D. E. (2001) Genetic algorithms for social innovation and creativity (IlliGAL report No 2001005). Urbana, IL: University of Illinois at Urbana-Champaign.
26. Law E., von Ahn, L., Dannenberg, R., Crawford, M. 2007. Tagatune: A Game For Music And Sound Annotation. In Proceedings of ISMIR (Vienna, Austria, 2007).
27. Law, E., von Ahn, L., and Mitchell, T. Search War: A Game for Improving Web Search. Invited Talk An Overview of Human Computation, In *Proc. of Knowledge Discovery and Data Mining*, Workshop on Human Computation HCOMP 2009, ACM.
28. Lenat, D. B. 1995. CYC: a large-scale investment in knowledge infrastructure. Commun. ACM 38, 11 (Nov.

1995), 33-38. DOI= http://doi.acm.org/10.1145/219717.219745

29. Lieberman, H., Smith, D., Teeters, A. Common Consensus: A Web-based Game for Collecting Commonsense Goals. Workshop on Common Sense for Intelligent Interfaces, ACM International Conference on Intelligent User Interfaces (IUI-07), Honolulu, January 2007.

30. Lydia B. Chilton, Clayton T. Sims, Max Goldman, Greg Little, and Robert C. Miller. "Seaweed: A Web Application for Designing Economic Games." HCOMP, 2009, to appear.

31. Lydia B. Chilton. "Seaweed: A Web Application for Designing Economic Games." M.Eng. Thesis, Massachusetts Institute of Technology, 2009.

32. Maslow, A.H. A Theory of Human Motivation. Psychological Review (1943). 370-396.

33. Matuszek, C., Witbrock, M., Kahlert, R. C., Cabral, J., Schneider, D., Shah, P., and Lenat, D.. 2005. Searching for common sense: Populating cyc(tm) from the web. In Proceedings of the Twentieth National Conference on Artificial Intelligence.

34. Mechanical Turk. http://mturk.com.

35. Mendez, B. M.; Craig, N.; Westphal, A. J. Stardust@home: Enlisting Students and the Public in the Search for Interstellar Dust. American Astronomical Society Meeting 207, #67.12; Bulletin of the American Astronomical Society, Vol. 37, p.1265. 12/2005.

36. Page, L., Brin, S., Motwani, R., and Winograd, T. The PageRank Citation Ranking: Bringing Order to the Web. STANFORD INFOLAB , (1999), 17.

37. Peha, J. M. and Khamitov, I. M. 2003. PayCash: a secure efficient Internet payment system. In Proceedings of the 5th international Conference on Electronic Commerce (Pittsburgh, Pennsylvania, September 30 - October 03, 2003). ICEC '03, vol. 50. ACM, New York, NY, 125-130. DOI= http://doi.acm.org/10.1145/948005.948022

38. Project Gutenberg. http://www.gutenberg.org.

39. reCAPTCHA. http://recaptcha.net.

40. Scanlon, J. Luis von Ahn: The Pioneer of “Human Computation”. In BusinessWeek, 11/3/2008. http://www.businessweek.com/innovate/content/nov2008/id2008113_656340.htm

41. Sheng, V., Provost, F., and Ipeirotis, P. Get Another Label? Improving Data Quality and Data Mining. (2008).

42. Singh, P., Lin, T., Mueller, E. T., Lim, G., Perkins, T., and Zhu, W. L. 2002. Open Mind Common Sense: Knowledge Acquisition from the General Public. In on the Move To Meaningful internet Systems, 2002 - DOA/CoopIS/ODBASE 2002 Confederated international Conferences Doa, CoopIS and ODBASE 2002 (October 30 - November 01, 2002). R. Meersman and Z. Tari, Eds. Lecture Notes In Computer Science, vol. 2519. Springer-Verlag, London, 1223-1237.

43. Surowiecki, J. The Wisdom of Crowds. Anchor Books, New York, 2005.

44. Tapscott, Don. Williams, Anthony D. Wikinomics: How Mass Collaboration Changes Everything. Portfolio Hardcover (2006). ISBN # 1591841380 / 978-1591841388.

45. The ESP Game. http://www.espgame.org.

46. The FACTory. http://game.cyc.com.

47. von Ahn, L. and Dabbish, L. 2004. Labeling images with a computer game. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Vienna, Austria, April 24 – 29, 2004). CHI ’04. ACM, New York, NY, 319-326. DOI= http://doi.acm.org/10.1145/985692.985733

48. von Ahn, L. v. 2006. Games with a Purpose. Computer 39, 6 (Jun. 2006), 92-94. DOI= http://dx.doi.org/10.1109/MC.2006.196

49. von Ahn, L., Blum, M., and Langford, J. 2004. Telling humans and computers apart automatically. Commun. ACM 47, 2 (Feb. 2004), 56-60. DOI= http://doi.acm.org/10.1145/966389.966390

50. von Ahn, L., Dabbish, L. General Techniques for Designing Games with a Purpose. *Communications of the ACM*, August 2008. Pages 58-67.

51. von Ahn, L., Ginosar, S., Kedia, M., Liu, R., and Blum, M. 2006. Improving accessibility of the web with a computer game. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 – 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI ’06. ACM, New York, NY, 79-82. DOI=http://doi.acm.org/10.1145/1124772.1124785

52. von Ahn, L., Graham, M., I., Dabbish, L., Kitchin, D., Blum, L. Solving Hard AI Problems With Computer Games. CMU Technical Report CMU-CS-02-191, Nov. 2002.

53. von Ahn, L., Kedia, M., and Blum, M. 2006. Verbosity: a game for collecting common-sense facts. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 75-78. DOI=http://doi.acm.org/10.1145/1124772.1124784

54. von Ahn, L., Liu, R., and Blum, M. 2006. Peekaboom: a game for locating objects in images. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 – 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI ’06. ACM,

New York, NY, 55-64. DOI= http://doi.acm.org/10.1145/1124772.1124782

55. von Ahn, L., M. Blum, N. Hopper, and J. Langford. Captcha: Using hard AI problems for security. In Eurocrypt, 2003.

56. Wikipedia. http://www.wikipedia.org.

57. Zhao, W., Chellappa, R., Phillips, P. J., and Rosenfeld, A. 2003. Face recognition: A literature survey. ACM Comput. Surv. 35, 4 (Dec. 2003), 399-458. DOI= http://doi.acm.org/10.1145/954339.954342