Defining (Human) Computation

Edith Law
Carnegie Mellon University
Pittsburgh, PA, 15217
edith@cmu.edu

ABSTRACT

Human computation is a term that has been used synonymously with other related concepts, including "crowdsourcing," "social computing," and "collective intelligence." Defining more precisely what human computation means will help to distinguish its research focus from other subfields, and isolate a set of fundamental research questions to pursue. In this position paper, we propose a definition of human computation that is grounded on familiar computer science concepts, such as *computation* and *algorithm*. Based on our proposed definition, we then outline the three main aspects of a human computation system and the key research questions associated with each aspect.

Author Keywords

Computation, Algorithm, Games with a Purpose

ACM Classification Keywords

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

INTRODUCTION

Human computation is a relatively new concept. Coined in 2006, the term *human computation* encompasses several ideas. The first idea is that people can be engaged to perform meaningful tasks through some other activities that they are already deeply interested in (e.g., playing games, signing up for email accounts). Consequently, we now have an economical means to leverage large amount of information processing power to tackle large, complex problems. The second idea is that systems can be built in such a way that control how the computation is carried out (by humans), thereby ensuring that the problem at hand is solved in an accurate and efficient manner. The ESP Game [10] is the first human computation system that encompasses all these ideas. In this game, two players are given an image to describe and are rewarded when their descriptions match. As a by-product of playing an enjoyable game, players are generating millions of image labels that can be used for online image search. Furthermore, the accuracy of the image labels is guaranteed by the game mechanism, which incentivizes players to tell the truth by requiring them to agree with each other. Likewise, reCAPTCHAs [11] are now used to digitize books by having people correct words in books that optical character recognition (OCR) fail to recognize with certainty. The system takes advantage of the fact that millions of users perform identity verification tasks (i.e., CAPTCHAs) each day in order to gain access to some online content. By having a user transcribe two words, one of which is known, the system has some guarantees that that user is not a computer bot and that his transcription of the unknown word is accurate.

The idea of harnessing the "wisdom of the crowd" is central to many related concepts, including "crowdsourcing," "social computing" and "collective intelligence." Unfortunately, the boundaries between these concepts are blurred, and these terms have been used in the research literature interchangeably. In this paper, we propose a definition of human computation, and pinpoint the characteristics (e.g., objectives and desirable properties) of human computation systems which set them apart from other types of crowd-driven systems. By better defining human computation, we hope to improve our understanding of what makes an effective human computation system, and the type of research questions that are relevant for achieving this goal.

ORGANIZED COMPUTATION: A SHORT HISTORY

Before defining human computation, we must clarify what we mean by *computation*. While *computation* is a central concept in Computer Science, there is much debate about its exact definition [3]. Originating from the Latin word "computare", to compute is to "count, sum up or reckon together." This simple definition of computation is not a far removed description of what early computers actually did. The Hollerith machine [4], for example, was used to tabulate statistics about population during the 1890 census in the United States, summing up data about more than 75 million individuals, such as the number of people who were

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

married, in each profession, with certain number of children, or speakers of English.

However, long before any modern desktop computers or tabulating machines, computation was carried out by humans [4] in the so-called *organized computation* projects. In fact, before the adoption of tabulating machines, the word *computer* referred to a person who performed calculation as a profession, and many tabulating machines that came about later on were named using acronyms that ended in "AC", meaning ``Automatic Computer'' [3], in order to be distinguished from *human computers*. These organized computation projects [4] involved anywhere from a few human computers (e.g., the computation of the trajectory of the Halley Comet in 1758 was undertaken by only three highly skilled astronomers) to hundreds of individuals (e.g., the Mathematical Table Project led by Arnold Logan in 1938 employed more than 450 people unskilled in mathematics), and were created to tackle the most pressing scientific questions or practical problems in society at the time, e.g., translating of locally used units of length, volume and weight to the metric system, calculating the trajectory of bombs during the first World War, or the modeling the stock market during the depression in 1930.

There are several key concepts in organized computation, which we will leverage to define what is meant by *computation*, and more importantly, *human computation*. First, in organized computation projects, a complex problem is typically decomposed into basic operations, distributed to many individuals in parallel, and re-assembled to reach a solution. For example, to measure the quantity of interest (e.g., the trajectory of the comet), a scientist or mathematician would devise a mathematical formula, break down the formula into a set of simpler quantities that can be easily computed by an individual human computer, then re-assemble the results. Second, computation is carried out using an explicit set of instructions, leaving little to interpretation. Human computers were often given a "computing plan" -- a sheet of paper with explicit instructions for each step of the computation -- to follow. These computing plans are reminiscent of what we now call an "algorithm," a finite sequence of explicit instructions to transform input to output. Finally, efficiency (in terms of time and cost) and accuracy were two important criteria to successful organized computation that project leaders strived for. For example, to lower the chances of a mistake, computation was sometimes done by two independent human computers, and checked by a third person who compared the results. To make computation faster, mathematical techniques such as interpolation were used or invented.

(HUMAN) COMPUTATION: A DEFINITION

To understand what we mean by human computation, we must first define for ourselves the word *computation*. In our formulation, computation is the process of mapping of some input representation to some output representation using a explicit, finite set of instructions (i.e., an *algorithm*). In the classic work by Alan Turing, computation is similarly defined, where the input and output representation are symbols, the process is the writing of symbols in each cell of an unlimited tape, and the instruction or algorithm is a state transition table that determines what symbol should be written for any given cell. Similarly, a human computer who is given two quantities (input representation) and asked to multiple them together (explicit instruction) generate a product (output representation) is performing computation. In fact, the Turing Machine was meant to mimic the capability of human computers in carrying out mathematical calculations. In Turing's own words, ``the idea behind digital computers may be explained by saying that these machines are intended to carry out any operations which could be done by a human computer'' [9].

Following this definition, ***human computation*** is simply computation that is carried out by a human. Likewise, ***human computation systems*** can be defined as intelligent systems that explicitly organizes human efforts to carry out the process of computation -- whether it be performing the basic operations, or taking charge of the control process itself (e.g., specifying what operations need to be performed and in what order). The ***objective of a human computation system*** is to find an *accurate* solution for a pre-specified computational problem in the most *efficient* way.

The important element in our definitions is the idea of **explicit control** – that unlike other crowd-driven systems (e.g., Wikipedia), computation is not the consequence of the natural dynamics in a crowd, but the consequence of a deliberate algorithm. Creating systems that exert explicit control is beneficial and necessary, especially given the unique opportunities and challenges that arise due to the popularity of the Web. With millions of online users, the Web has essentially created a constant stream of human computers available to perform any computational tasks. This means that we are now endowed with the ability to tackle hugely complex problems; how to do so effectively is not well understood. In addition, different from the past, human computers are now typically anonymous workers, whose characteristics (e.g., expertise, competence, intent, and interests) are not always observable. As a result, the accuracy of the final output can vary greatly depending on these hidden factors. By creating systems to exert control, e.g., over what operations to perform, who to perform them and how – we can hope to have better *guarantees* on both the efficiency and accuracy of our solution to the computational problem.

As a note of clarification, our definition of human computation can be applied to any systems that exert explicit control over how computation is performed with human in the loop, regardless of how many human computers are actually involved. That is, our definition is equally applicable to systems that spend all its efforts finding that *single* best human computer to perform a task, or systems that randomly distribute tasks to thousands of human computers, then aggregating their outputs. In fact, a system that can leverage more users is not always better – the relative merits of different human computation systems depend on how well they achieve their accuracy and efficiency objectives. We claim, for example, that a human computation system that involves only a single human computer is better than one that involves thousands of human computers, if it can get to a more accurate solution more quickly.

Distinctions from Other Crowd-Driven Systems

To distinguish human computation from other prevalent concepts related to the idea of the *wisdom of the crowd*, let's first examine these concepts as they are defined by Wikipedia (Figure 1).

Crowdsourcing is the "act of outsourcing tasks, traditionally performed by an employee or contractor, to an undefined, large group of people or community (a crowd) through an open call."

Collective Intelligence is the "shared or group intelligence that emerges from the collaboration and competition of many individuals and appears in consensus decision making in bacteria, animals, humans and computer networks."

Social Computing is a "general term for an area of computer science that is concerned with the intersection of social behavior and computational systems. In the weaker sense of the term, social computing has to do with supporting any sort of social behavior in or through computational systems. It is based on creating or recreating social conventions and social contexts through the use of software and technology, e.g., blogs, email, instant messaging, social network services, wikis, social bookmarking. In the stronger sense of the term, social computing has to do with supporting *computations* that are carried out by groups of people, e.g., collaborative filtering, online auctions, prediction markets, reputation systems, computational social choice, tagging and verification games."

Figure 1. Definitions of related concepts from Wikipedia

Based on these definitions, *crowdsourcing* can be considered a method or tool that human computation systems can use to distribute tasks through an open call. However, a human computation system needs not use crowdsourcing; a system that assigns tasks to a closed set of workers hired through the traditional recruiting process (e.g., resumes, in-person interviews) can still be considered a human computation system. *Social Computing* appears to be a broad concept that covers everything to do with social behavior and computing. Human computation *intersects* social computing in that some, but not all, human computation systems require social behavior and interaction amongst a group of people. As mentioned before, human computation does not necessary need to involve large crowds, and workers are not always required to interact with one another, either directly or indirectly (e.g., through a market mechanism). Finally, *collective intelligence* refers to the emergent intelligent behavior of a group of individuals, which includes non-humans or even non-living things. Collective intelligence, therefore, is an even broader concept that subsumes crowdsourcing, social computing and human computation.

None of the related concepts emphasizes the idea of explicit control. In fact, they assume that a large part of the computational outcome is determined by the natural dynamics (e.g., coordination and competition) between the individuals of a group, which the system does not or cannot deliberately control. In contrast, we define human computation to be the study of intelligent systems that exert explicit control over how the computation is carried out. In the next section, we will review three different aspects of this explicit control in a human computation system.

EXPLICIT CONTROL – HOW, WHAT, AND WHO

There are three aspects of human computation – referred to as the "what", "who" and "how" (Figure 2) – where explicit control can be applied. Note that *both* humans and machines can be in charge of this explicit control. For example, in deciding to whom to assign what tasks, there can be a spectrum of solutions, ranging from *push* methods where machines automatically assign tasks to workers, *pull* methods where workers browse and search for tasks to assign themselves, and *hybrid* methods which leverage the complementary efforts of both machines and humans to find the ideal matching.

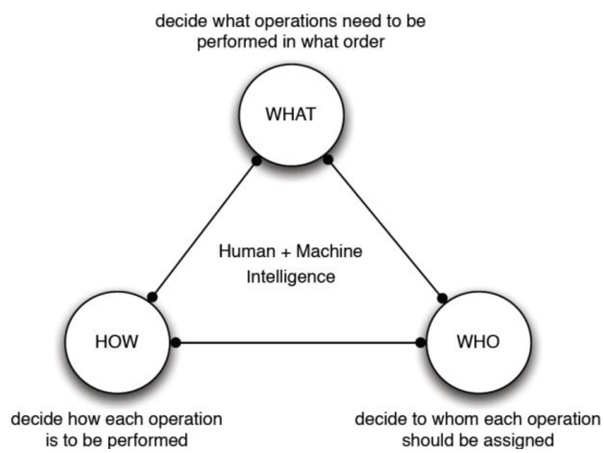

Figure 2. Three Areas of Control

The "What" Questions

According to our definition, a human computation system must have an *algorithm* that outlines exactly how to solve the computational problem at hand. An algorithm consists of a set of operations and control structures (e.g., repetitions, iterations, conditions) that specify how the operations are to be arranged and executed. Similar to algorithms in the traditional sense, some human computation algorithms are more efficient than others. For example, if our computational problem is to map a set of images to tags, an efficient algorithm would make use of machine intelligence (e.g., active learning [8]) to select only images that the computer vision algorithm does not already know how to classify. Such an algorithm would greatly reduce the costs of the computation, both in terms of time and monetary payment to human workers. Some research questions relevant to the "what" aspect include:

- Can we leverage the complementary abilities of *both* humans and machines [5] to make computation more accurate and efficient, e.g., by eliminating tasks that can be handled adequately by machines?
- How can the system get machines or people to decompose complex tasks into operations and order them in such a way to obtain the best outputs? [1,2]

The "Who" Questions

Knowing what operations need to be performed, the next question is who should perform them. While for some tasks aggregating the work of non-experts suffices, other tasks are knowledge intensive and require special expertise. For example, a doctor who is asked to verify the fact "Obacillus Bordetella Pertussis is a bacterium*"* is likely to be a better (and faster) judge than someone without any medical training. Some research questions relevant to the "who" aspect include:

- What are some effective algorithms and interfaces (e.g., search, visualization, recommendation) for routing tasks?
- How do we model the expertise of workers, which may be changing over time?
- What are some optimal strategies for allocating tasks to workers, if their availability, expertise, interests, competence and intents are known versus unknown?

The "How" Questions

Finally, there is the question of human-computer interaction – how can we design systems that motivate people to participate and to carry out the computational tasks to their best abilities (i.e., truthfully, accurately, and efficiently). Some research questions relevant to the "how" aspect include:

- How do we motivate people to have a long-term interaction with the system, by creating an environment that meets their particular needs (e.g., to be entertained, to have a sense of accomplishment or community)?
- How do we design game mechanisms [6] that incentivize workers to tell the truth, i.e., generate accurate outputs?
- What are some new markets, organizational structures or interaction models for defining how workers relate to each other (as opposed to working completely independently)?

An Example

One of our research projects involves the creation of a human computation game for verifying hundreds of thousands of facts that are automatically extracted using a large-scale web mining system [7] that learns to read the Web over time. In this game, we must intelligently decide *what* facts, if corrected, would be most beneficial to the learning system, *how* to structure the game such that players are motivated to report their true beliefs about the correctness of each fact, and finally, *who* (which players) are likely to have the knowledge to verify each fact.

CONCLUSION

In this position paper, we define human computation by going back to its root – examining the meaning of "computation" and revisiting the history of organized computation before the rise of the Internet. Our definitions, centered on the idea of *explicit control*, allow us to distinguish human computation from other related concepts, as well as pinpoint three main areas of research in human computation. In better defining human computation, we hope to have focused the research directions for this field.

ACKNOWLEDGEMENTS

Thanks to Winter Mason, Paul Resnick and Yiling Chen for their inputs regarding the definition of human computation.

REFERENCES

1. Bernstein, M., Little, G., Miller, R., Hartmann, B., Ackerman, M., Krager, D., Crowell, D., Panovich, K. A word processor with a crowd inside. In UIST 2010.
2. Dai, P., Mausam and Weld, D. Decision-theoretic control of crowd-sourced work-flows. In AAAI 2010.
3. Denning, P. and Wegner, P. "Ubiquity Symposium: What is Computation?" http://ubiquity.acm.org/ symposia.cfm
4. Grier, D. When Computers were Human. Princeton University Press, 2005.
5. Horvitz, E. and Paek, T. Complementary computing: policies for transferring callers from dialog systems to human receptionists. User Modeling and User-adapted Interaction, 2007, 17(1-2): 159-182.
6. Law, E. and von Ahn, L. Input-agreement: A new mechanism for data collection using human computation games. In CHI 2009, 1197—1206.
7. NELL: Never-ending language learner. http://rtw.ml. cmu.edu/rtw/.
8. Settles, B. Active Learning Literature Survey. Technical Report, 2009.
9. Turing, A. Computing machinery and intelligence. In Mind (236), 433—460.
10. Von Ahn, L. and Dabbish, L. Labeling images with a computer game. In *CHI* 2004, 319—326.
11. Von Ahn, L., Maurer, B., McMillen, C., Abraham, D. and Blum, M. reCAPTCHA: Human-Based Character Recognition via Web Security Measures. Science, September 12, 2008, 1465 – 1468.