

Crowd Control

Using crowdsourcing applications, humans around the world are transcribing audio files, conducting market research, and labeling data, for work or pleasure.

THOUGH COMPUTERS HAVE outstripped us in arithmetic and chess, there are still plenty of areas where the human mind excels, such as visual cognition and language processing. And if one mind is good, as the proverb goes, two—or two thousand—are often better. That insight, and its consequences, drew worldwide interest with the 2004 publication of James Surowiecki's best-selling *The Wisdom of Crowds*, which argued that a large group of people are superior at certain types of rational tasks than individuals or even experts.

Now researchers are turning to computers to help us take advantage of our own cognitive abilities and of the wisdom of crowds. Through a distributed problem-solving process variously known as crowdsourcing, human computation, and computer-aided micro-consulting, answers are solicited online to a set of simple, specific questions that computers can't solve. Is this a picture of a fish? Do you like that style of shoe? How many hotels are on St. George's Island, and which ones have Internet access?

The amateur, often anonymous workers who agree to execute these tasks are usually given some sort of social or financial incentive. A few cents might buy the answer to a simple data-labeling task, while a more ardu-

ous job like audio transcription could require a couple of dollars. Reposition the task as a game, and many people even "work" for free. Either way, the possibilities—for creating corpuses of annotated data, conducting market research, and more—have both computer scientists and companies excited.

One of the oldest commercial crowdsourcing applications is Amazon's Mechanical Turk. Named after a famous 18th century chess-playing "machine" that was secretly operated by a human, it offers a flexible, Web-based platform for creating and publicizing tasks and distributing micro-payments. Since its launch in 2006, Turk has spawned both a vocabulary and a mini-marketplace. Workers, or "Turkers" (there are more than 200,000 in 185,000 countries, according to Amazon), select "Human Intelligence Tasks" (HITs) that match their interests and abilities. Motivations vary. Some work odd hours or at night to generate extra income, while others simply desire a more productive way to kill time online, like solitaire with financial rewards. As in the offline world, more money buys faster results, and Amazon's HIT requesters often experiment to find a pay scale that matches their needs.

Also part of the Turk economy are companies like Dolores Labs and Cast- ingWords, which rely on Amazon's

technology to power their own crowdsourcing applications. Dolores Labs, based in San Francisco, posts Turk HITs on behalf of its clients, then filters the answers through custom-built software systems to check for quality and generate meaningful results. Data is ultimately used to perform tasks like filter comment spam, tag data for search engine optimization, and research market trends.

"Many companies don't have the resources to describe tasks, put them up online, and manage the data they get," explains Lukas Biewald, the company's founder and CEO. Nor do they have time for Dolores's extensive quality-control measures, which include creating "test" questions whose answers are already known, checking responses against one another, tracking individual answer histories, and creating a confidence measure with which to weight the resulting data.

Dolores also guides clients through the many variables that are involved in designing a crowdsourced project. How arduous is each task? How quickly are results needed? How would clients like to deal with the statistical outliers that are caught by Dolores' quality-control algorithms? If you're checking user-generated content for pornography, for example, you might err on the side of caution.

According to Biewald's estimates,



Some of the 10,000 sheep created for Aaron Koblin's TheSheepMarket.com by workers for Amazon's Mechanical Turk who were paid .02 cents to "draw a sheep facing to the left."

the cost for a crowdsourced project ranges from \$2,000 to \$4,000 for simple tagging projects to \$10,000 to \$20,000 for more complex custom applications. Stephen Mechler, managing director of the German crowdsourcing Web site Floxter, which uses its own technologies to handle the mechanics of creating and assigning tasks and compensating workers, calculates that it is 33% less expensive to crowdsource projects like data classification and tagging than to complete them with in-house employees.

Other companies focus their crowdsourcing efforts on specific types of projects. New Mexico-based CastingWords uses Turk to transcribe audio files. Through a propriety algorithm, files are first split into three- to four-minute chunks. Next, Turkers listen to a few seconds of each clip to judge the quality of the recording, which in turn helps determine pay rates for the transcription work. Once each file has been transcribed, a full draft is assembled and sent back to Turk to be graded for consistency and precision, and retranscribed where necessary. Finally, Turkers edit and polish the transcript to be sent back to the client. Total costs range from \$.75 to \$2.50 per audio minute, depending on how quickly a client needs the work completed.

Researchers like Carnegie Mellon University computer science professor Luis von Ahn are also finding ways to put crowdsourcing to work. Unlike his corporate peers, von Ahn is unable to pay for the completion of a task, so he relies on social incentives—and tries to make tasks fun. To entice people to manually label a collection of digital images, for instance, von Ahn created the ESP Game, which randomly matches each player with an anonymous partner. Players try to guess which words or phrases their partners (whom they can't communicate with) would use to describe a certain image. Once both players type the same descriptor, a new image appears and the process begins anew. In 2006, Google licensed the idea and created its own version of the game in order to improve image search results.

Since then, von Ahn has developed other games with a purpose to harness the wisdom of crowds. In Peekaboom, for example, one player attempts to guess the word associated with a particular im-

Since 2007, some 400 million people have helped digitize more than five billion words, says Carnegie Mellon's Luis von Ahn.

age as the other player slowly reveals it. In fact, designing a game is much like designing an algorithm, as von Ahn has pointed out: "It must be proven correct, its efficiency can be analyzed, a more efficient version can supersede a less efficient one." And since many people are inherently competitive, building a community around each game to recognize outstanding performers helps increase participation, as well.

reCAPTCHA, on the other hand, is an attempt to take advantage of a task that millions of people perform in the course of their everyday online lives: solve the ubiquitous character recognition tests known as CAPTCHAs to prove they are human. "I developed reCAPTCHA because I found out that we're wasting 500,000 collective hours each day solving these mindless tasks," says von Ahn. To put that brainpower to use, reCAPTCHA presents users with scanned images from old books and newspapers, which computers have difficulty deciphering. By solving the reCAPTCHA they help digitize the works. Since 2007, some 400 million people have helped digitize more than five billion words, according to von Ahn.

Crowdsourcing's critics claim it is unethical and exploitive, paying pennies or nothing for honest labor (though diligent workers often make close to minimum wage). In a struggling economy, people may grow choosier about the ways they earn extra income. On the other hand, they may also be more interested in blowing off steam on the Internet—and being rewarded with a few extra dollars. **Q**

Leah Hoffmann is a Brooklyn-based science and technology writer. Valerie Nygaard, Microsoft, contributed to the development of this article.

Online Social Networks

Adults Get Social

The percentage of Internet users age 55 and older who have a profile on an online social network has quadrupled during the last four years, from 8% in 2005 to 35% in 2008, according to a new survey by the Pew Internet & American Life Project.

Although media coverage has largely focused on how children and young adults use social network sites, adults still comprise the majority of the users of the social network sites because adults make up a larger portion of the U.S. population than teens, 65% of whom use social network sites.

Overall, however, younger adults are much more likely than older adults to use social networks. For instance, 75% of online adults age 18–24 have a profile on a social network; 57% of online adults 25–34 have a profile; 30% of online adults 35–44 have a profile; 19% of online adults 45–54 have a profile; 10% of adults 55–64 have a profile; and only 7% of online adults 65 and older have a profile.

In terms of gender, adult women and men are equally likely to use social networks.

The Pew study also reported that minority groups are more prevalent on social sites than previously expected. It found that 48% of African-American adults and 43% of nonwhite Hispanic adults have a social profile, compared to 31% of white adults.

The personal use of social networks is more prevalent than professional use, both in the orientation of the networks that adults choose to use as well as the reasons they give for using the applications. For instance, 50% of adult users have a profile on MySpace, 22% on Facebook, and 6% on LinkedIn.

The applications are mostly used to explain and maintain personal networks, and most older adults are using them to connect with people they already know, usually to keep up with (89%), make plans with friends (57%), or to make new friends (49%). Other uses include organizing with others for an event, cause or issue; flirting; promoting one's self or work; and making new business contacts.