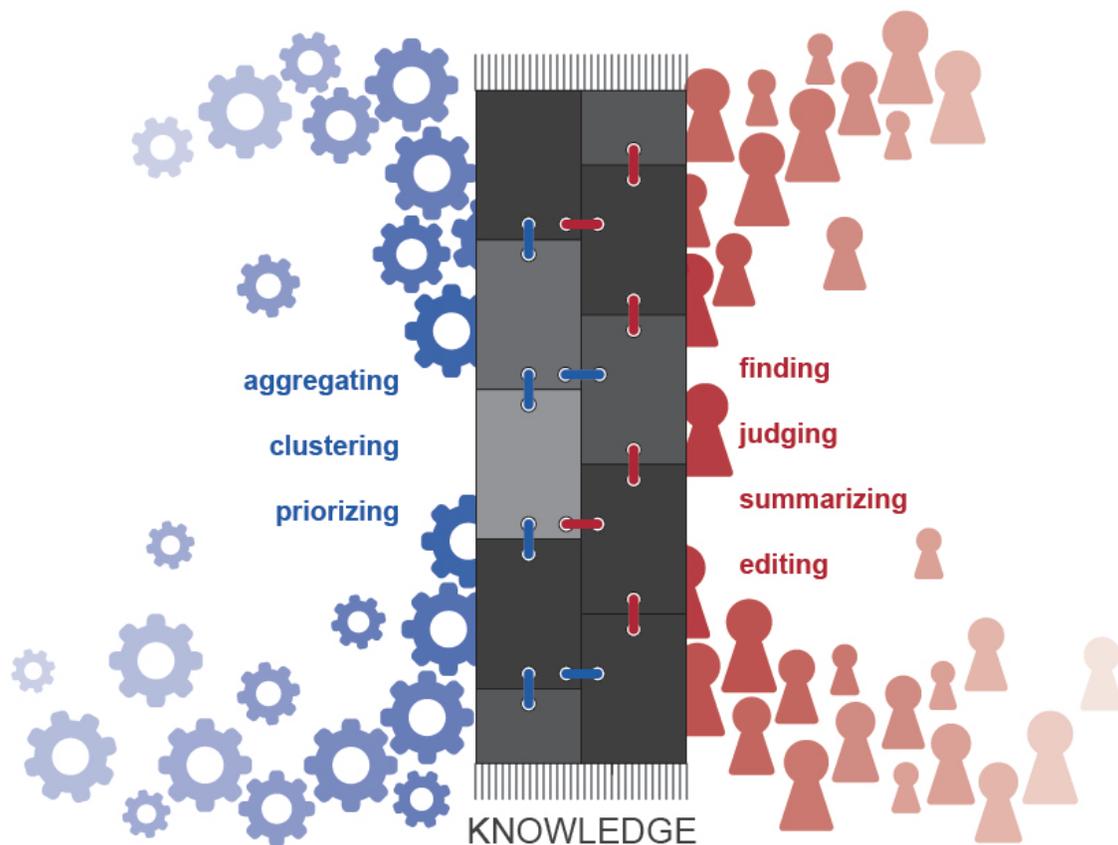# Crowd-augmented cognition: Team develops tools that combine human and machine intelligence to accelerate learning

May 11 2016

The Knowledge Accelerator uses a machine-learning program to sort and organize information. Credit: Aniket Kittur, Carnegie Mellon University

Crowdsourcing has brought us Wikipedia and ways to understand how HIV proteins fold. It also provides an increasingly effective means for teams to write software, perform research or accomplish small repetitive digital tasks.

However, most tasks have proven resistant to distributed labor, at least without a central organizer. As in the case of Wikipedia, their success often relies on the efforts of a small cadre of dedicated volunteers. If these individuals move on, the project becomes difficult to sustain.

Scientists funded by the National Science Foundation (NSF) are finding new solutions to these challenges.

Aniket Kittur, an associate professor in the Human-Computer Interaction Institute at Carnegie Mellon University (CMU), designs crowdsourcing frameworks that combine the best qualities of machine learning and [human intelligence](#), in order to allow distributed groups of workers to perform complicated cognitive tasks. Those include writing how-to guides or organizing information without a central organizer.

At the [Computer-Human Interaction](#) conference in Chicago this week, Kittur and his collaborators Nathan Hahn and Joseph Chang (CMU), and Ji Eun Kim (Bosch Corporate Research), will present two prototype systems that enable teams of volunteers, buttressed by machine learning algorithms, to crowdsource more complex intellectual tasks with greater speed and accuracy (and at a lower cost) than past systems.

"We are trying to scale up human thinking by letting people build on the work that others have done before them," Kittur said.
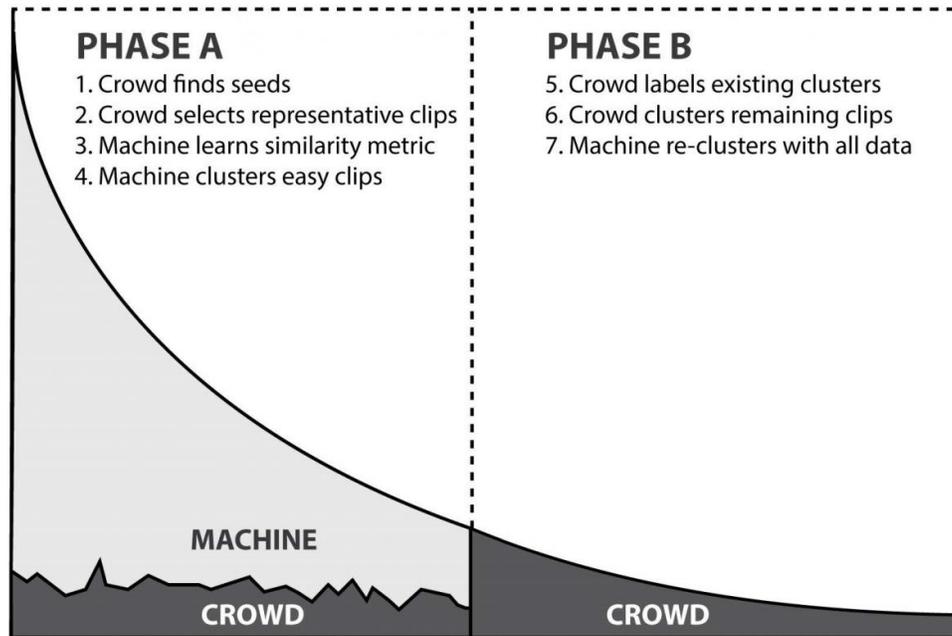
The Knowledge Accelerator

One piece of prototype software developed by Kittur and his collaborators, called the [Knowledge Accelerator](#) empowers distributed workers to perform information synthesis.

The software combines materials from a variety of sources, and constructs articles that can provide answers to commonly sought questions—questions like: "How do I get my tomato plant to produce more tomatoes?" or "How do I unclog my bathtub drain?"

To assemble answers, individuals identify high-value sources from the Internet, extract useful information from those sources, cluster clips into commonly discussed topics, and identify illustrative images or video.

With the Knowledge Accelerator, each crowd worker contributes a small amount of effort to synthesize online information to answer complex or open-ended questions, without an overseer or moderator.

The researchers' challenge lies in designing a system that can divide assignments into short microtasks, each paying crowd workers $1 for 5-10 minutes of work. The system then must combine that information in a way that maintains the article flow and cohesion, as if it were written by a single author.

PHASE A
1. Crowd finds seeds
2. Crowd selects representative clips
3. Machine learns similarity metric
4. Machine clusters easy clips

PHASE B
5. Crowd labels existing clusters
6. Crowd clusters remaining clips
7. Machine re-clusters with all data

MACHINE

CROWD   CROWD

Conceptual overview of the Alloy system. Credit: Aniket Kittur, Carnegie Mellon University

The researchers showed that their method produced articles judged by crowd workers as more useful than pages that were in the top five Google results from a given query. Those top Google results are typically created by experts or professional writers.

"Overall, we believe this is a step towards a future of big thinking in small pieces, where complex thinking can be scaled beyond individual limits by massively distributing it across individuals," the authors concluded.

Alloy

A related problem that Kittur and his team tackled involved clustering—pulling out the patterns or themes among documents to organize information, whether Internet searches, academic research articles or consumer product reviews.

Machine learning systems have proven successful at automating aspects of this work, but their inability to understand distinctions in meaning among similar documents and topics means that humans are still better at the task. When human judgement is used in crowdsourcing, however, individuals often miss the full context that allows them to do the task effectively.

The new system, called [Alloy](link), combines human intelligence and machine learning to speed up clustering using a two-step process.
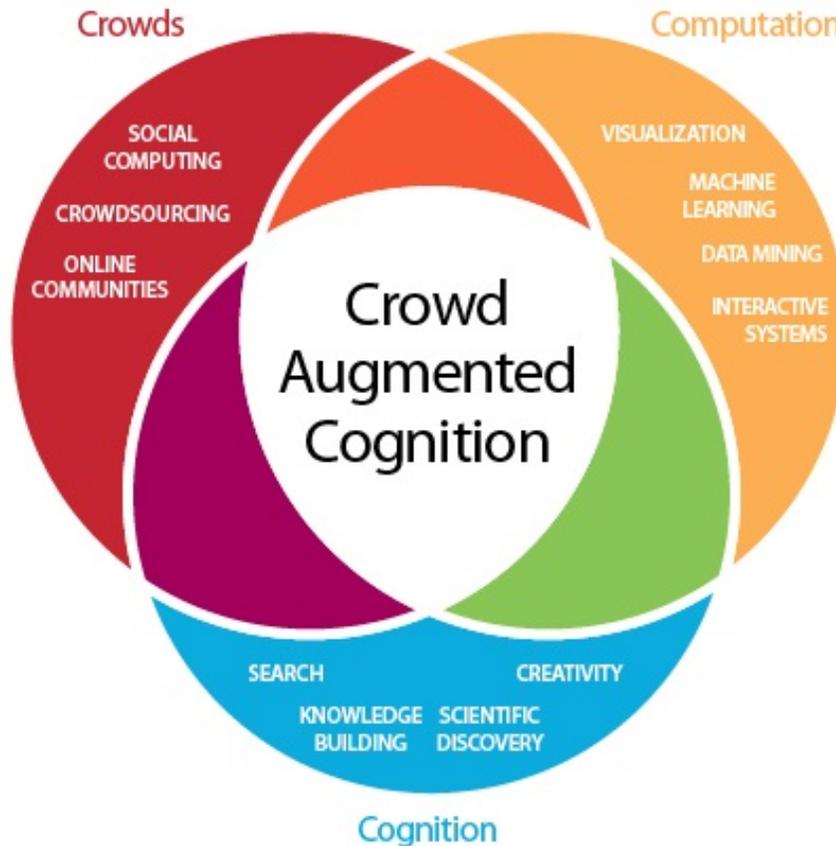
Research in Kittur's lab is on augmenting human cognition using crowds and computation. Credit: Aniket Kittur, Carnegie Mellon University

In the first step, crowdworkers identify meaningful categories and provide representative examples, which the machine uses to cluster a large body of topics or documents. However, not every document can be easily classified, so in the second step, humans consider those documents that the machines weren't able to cluster well, providing additional information and insights.

The study found that Alloy, using the [two-step process](#), achieved better performance at a lower cost than previous crowd-based approaches. The framework, researchers say, could be adapted for other tasks such as image clustering or real-time video event detection.

"The key challenge here is trying to build a big picture view when each person can only see a small piece of the whole," Kittur said. "We tackle this by giving workers new ways to see more context and by stitching together each worker's view with a flexible [machine learning](#) backbone."

Kittur is conducting his research under an NSF Faculty Early Career Development (CAREER) award, which he received in 2012. The award supports junior faculty who exemplify the role of teacher-scholars through outstanding research, excellent education and the integration of education and research within the context of the mission of their organization. NSF is funding his work with $500,000 over five years.

The work advances the understanding and design of crowdsourcing frameworks, which can be applied to a variety of domains, he says.

"It has the potential to improve the efficiency of knowledge work, the training and practice of scientists, and the effectiveness of education," Kittur says. "Our long-term goal is to produce a universal knowledge accelerator: capturing a fraction of the learning that every person engages in every day, and making that benefit later people who can learn faster and more deeply than ever before."