

## Making better use of the crowd

December 2 2016

Over the last decade, computer scientists have harnessed crowds of Internet users to solve tasks that are notoriously difficult to crack with computers alone, such as determining whether an image contains a tree, rating the relevance of websites, and verifying phone numbers.

The machine learning community was early to embrace so-called crowdsourcing to quickly and inexpensively obtain the vast quantities of labeled <u>data</u> needed to train machine learning systems how to classify images or recognize speech, for example. Labeled data are essentially sets of teaching examples, such as pictures of cats that are tagged with the word "cat."

Usually this handoff of labeled data is where interaction with the crowd ends. Are there better ways to make use of the crowd?

On December 5, I will showcase innovative uses of crowdsourcing that go beyond data collection during a crowdsourcing tutorial at NIPS, the premier international machine learning conference, held this year in Barcelona. The tutorial will also dive into recent research aimed at understanding who crowdworkers are and how they behave, which could inform best practices for interacting with the crowd.

Innovative uses of crowdsourcing that go beyond the collection of data include:

• Harnessing the crowd to improve machine-learning models including extracting features from labeled data most relevant for



model training and evaluation of learned models.

- Leveraging the complementary strengths of humans and machines to achieve more than either can alone. Potential applications of these so-called hybrid-intelligence systems include real-time on-demand closed captioning of day-to-day conversations and crowd-powered writing and editing.
- Using crowdsourcing platforms to recruit large pools of subjects for experiments designed to study the effects of human behavior when reasoning about the performance of computer systems, which could lead to better designed algorithms and systems.

Recent research—both qualitative and quantitative—has opened the black box of crowdsourcing to uncover that crowdworkers are not just independent contractors, but rather a network with a rich communication structure. Meanwhile, experiments have explored how to boost the quality of crowdwork using both well-designed monetary incentives (such as performance-based payments) and intrinsic motivation (such as piqued curiosity).

This research has much to teach us about how to most effectively interact with the <u>crowd</u>. (Hint: Be respectful, be responsive, be clear.)

Crowdsourcing has the potential for major impact on the way we design <u>machine learning</u> and AI systems, but to unleash this potential we need more creative minds exploring novel ways to use it.

## Provided by Microsoft

Citation: Making better use of the crowd (2016, December 2) retrieved 25 April 2024 from <u>https://phys.org/news/2016-12-crowd.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private



study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.