Worker-Owned Cooperative Models for Training Artificial Intelligence

Anand Sriraman
TCS Research - TRDDC
Pune, Maharashtra, India
anand.sriraman@tcs.com

Jonathan Bragg
University of Washington
Seattle, WA, USA
jbragg@cs.washington.edu

Anand Kulkarni
University of California, Berkeley
Berkeley, CA, USA
anandk@berkeley.edu

CSCW '17 Companion, February 25 - March 01, 2017, Portland, OR, USA
ACM 978-1-4503-4688-7/17/02.
http://dx.doi.org/10.1145/3022198.3026356

Abstract
Artificial intelligence (AI) is widely expected to reduce the need for human labor in a variety of sectors. Workers on virtual labor marketplaces accelerate this process by generating training data for AI systems. We propose a new model where workers earn ownership of trained AI systems, allowing them to draw a long-term royalty from a tool that replaces their labor. This concept offers benefits for workers and requesters alike, reducing the upfront costs of model training while increasing longer-term rewards to workers. We identify design and technical problems associated with this new concept, including finding market opportunities for trained models, financing model training, and compensating workers fairly for training contributions. A survey of workers on Amazon Mechanical Turk about this idea finds that workers are willing to give up 25% of their earnings in exchange for an investment in the future performance of a machine learning system.

Author Keywords
Worker-owned cooperatives; artificial intelligence; machine learning; human computation; crowdsourcing; policy

ACM Classification Keywords
K.4.3 Organizational Impact: Computer-supported cooperative work.

"I'd definitely say that we should be paid for any role we're doing in creating new technologies. And regarding… a cut in pay for a 50% chance of more money later--for me it would depend on whether that chance is based purely on luck or if it's based on something that I did--for example, a 50% chance your work is up to par."
- Anonymous Mechanical Turk worker, answering survey.

Survey Results

We launched a survey on MTurk to understand workers' reactions to such a model. We report results from 31 participants' responses, most of them from the United States. 48% were below the age of 28 with 25.8% women. 77% reported earning less than $1000 per month from MTurk.

On a scale of 1-5, with 5 being the most risk-taking, workers reported a mean rating of 2.6, indicating some appetite for risk.

Motivation

Recent reports predict employment disruption from artificial intelligence (AI) systems [4]. Most proposed solutions to job loss caused by automation concern creation of alternative income streams from sources unrelated to the automation, such as government-subsidized basic income guarantees.

Common machine learning (ML) approaches use a set of human-annotated training examples to produce training data. Current compensation models offer one-time payment to workers under a work-for-hire contract. In this paper, we propose an alternative equity-based model where human labelers can earn collective ownership over the AI systems they train. This would entitle them to a proportional share of downstream earnings produced by the AI. We believe this royalty model could offer a meaningful alternative to the current system of automation eliminating jobs entirely, letting workers maintain a stream of income from trained models while assuming some of the risk and cost involved in developing a machine learning system.

Our proposed cooperative model is inspired by the success of other systems that support cooperative worker efforts around collective action [6] and finding work [5]. These activist technologies help to address growing inequalities between sellers (workers) and buyers (requesters), as well as between users of platforms and the platforms themselves. Activist technologies can exist alongside existing platforms [5,6], or there is a growing interest in creating alternative platforms for online work [8] and in platform cooperativism for the on-demand economy at large [7].

We discuss four questions that must be answered to let worker teams earn and draw royalties from machine learning models they cooperatively own, and explore the design space for each. First, how should we design the ownership relationship between workers and the AI system? Second, how can teams of workers find AI systems worth building, and where can they market and sell a trained machine learning system? Third, how can workers fairly divide earnings from a model trained by multiple people under a variety of machine learning algorithms? Last, how can workers decide which examples to label and how to maintain validity of data? We also discuss the results of an initial survey of MTurk workers to gauge their reactions to such a model.

Designing Co-Ops for Worker Ownership

Today, development of AI training data on crowd labor marketplaces is funded entirely by requesters. Under the terms of platforms like Amazon Mechanical Turk [2], the data produced (and trained AI systems that result) are owned entirely by requesters in exchange for a fixed price paid to workers for producing that data. In a cooperative model for training AI systems, workers can choose to accept a fraction of that price in exchange for shares of ownership in the resulting trained system. By varying the payment workers receive for doing work as a split between cash and shares of equity in the trained AI system, we can examine a spectrum of new machine learning ownership models with varying ownership and risk shared between requesters and workers. Workers can obtain increased ownership of the model by offering training data at a reduced price in exchange for more shares of equity, and further increase ownership by purchasing training data from other workers. We can imagine interested outside investors participating in

We asked how much of their MTurk income they would be willing to give up, given the chance to earn 100% more over the course of one year. On average, workers were willing to give up 25% of their income. Only 3 participants said they'd not be willing to give up any of their earnings. Age doesn't seem to be a factor here, where workers under 28 years of age were willing to give up 23% on average and two participants saying they wouldn't give up any earnings. We also asked how much they'd need to make back in order to give up 100% of their earnings. On average, workers needed to be able to make back 3 times their invested amount. 45% of workers reported not being worried at all about AI taking over their jobs. This remains the case among the 10 workers who reported having followed the progress of machine learning or AI in the news, where 50% were not worried about the prospect.

such co-ops as well, bankrolling projects that have a significant chance of success.

Markets and Strategies for Earning Money from Machine Learning Systems

Open calls for effective machine learning algorithms for solving particular problems exist today on platforms like Kaggle and Algorithmia; these platforms can enable cooperatives to derive value from trained models.

Bounties vs. Marketplaces: On Kaggle and Algorithmia, the poster of the call offers a bounty for trained AI systems meeting a particular model and determines whether a submitted algorithm is acceptable. Worker-owned models that seek to earn a bounty assume risk that the poster may not accept their solution, the poster may choose another submission over their solution, or that the open call may expire. Further, the conditions of the bounty may require that the algorithm creator relinquishes ownership of and any future rewards derived from the algorithm once the bounty is earned [1]. Alternately, Algorithmia also makes it possible for owners of trained ML models to post and profit from their trained systems on a per-use basis, providing ongoing earnings. However, identifying a valuable problem domain and determining its future return may be difficult.

Online vs. Offline Training Models: Per-use platforms provide the additional benefit that customers of the AI system can start to use the system before it is fully trained. Approaches to developing an AI+crowd collaborative service to solve tasks usually have crowds doing a large part of the work in the initial stages, while the AI learns. As the AI system gains confidence in its predictions, work starts shifting from the crowd to the AI. In this "online" payment model, the crowd would earn from two channels: i) payment for the actual crowd labels provided, and ii) payment from use of the AI system they helped train. Here, crowd payments remain relatively stable, while ensuring high output quality. Payments from the use of the AI could be weighted by the number of labels the worker provided. As an alternative, workers or investors who are interested in having a substantial dataset annotated may choose to speculate on its future utility in the hopes that it would be useful on such platforms. In this case, the workers invest in the dataset by donating their time in annotating and curating the dataset. They aren't paid for the annotations, but own 100% of the generated dataset and earnings may be higher.

Distributing Earnings Fairly Among Workers

Credit Assignment: Optimally assigning credit for individual training examples is an open theoretical problem. Algorithms may treat machine learning algorithms as black boxes, or exploit the subclass of machine learning model being considered. One simple, but computationally intensive option for determining the value of a training example, is to compare performance of the model with and without that example. Credit may depend on when training data was provided during training, as well as its quality.

Measuring Algorithm Quality: To determine whether an algorithm is improving in performance, one must have a test set for measurement. How does one collect such a test set? Do workers curate the test set? If so, how does one ensure that workers do not train an overfitted model by selecting training examples that correspond to data in the test set? Subcontracting could be used to create the test set, but methods are

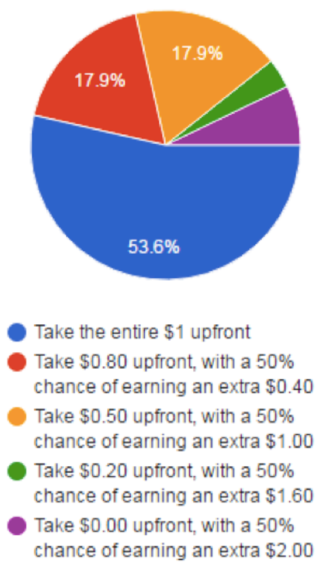

Figure 1: Nearly half of workers surveyed were willing to give up some portion of immediate earnings for a 50% chance to earn 2 times the forgone amount.

needed for guaranteeing that data will not leak, especially in the face of adversarial attacks.

Method of Data Labeling

Examples Proposed by Workers or by the Algorithm: A common paradigm for selecting the next example in machine learning is active learning, where the algorithm selects the next example X to label and the worker provides the next label Y. In some domains, it may not be possible for the machine to generate X, or the training process may be sped up by having workers find the best examples to label [3]. How much credit should be assigned to a worker who has provided just the example X, compared to workers who have provided both X and Y? When workers provide both X and Y, it is not necessary for the same worker to provide both X and Y. For instance, if one is training a machine translation system, a bilingual worker may be needed to provide a label Y, but a worker that only speaks English may be able to provide a label X.

Data Maintenance: For some types of data, such as search relevance judgments, data may become stale over time. Lack of maintenance of valid data may reduce the usefulness of the trained AI, resulting in less profitability to the owners of the model. We propose that the profits (or losses) of the data owners reflect the evolving usefulness of the data. This is especially important for algorithms that derive ongoing rewards (described above).

Acknowledgements

We thank the CrowdCamp workshop at the Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2016) for supporting this work and the participants for useful discussions. We also thank the anonymous reviewers for providing valuable feedback and the crowd workers who participated in our survey.

References

1. Algorithmia. 2015. Bounty Terms and Conditions. Retrieved November 4, 2016 from https://algorithmia.com/bounty_terms.
2. Amazon Mechanical Turk. 2014. Participation Agreement. Retrieved November 4, 2016 from https://www.mturk.com/mturk/conditionsofuse.
3. Attenberg, Josh, Panagiotis G. Ipeirotis, and Foster J. Provost. "Beat the Machine: Challenging Workers to Find the Unknown Unknowns." Human Computation 11 (2011): 11.
4. Executive Office of the President National Science and Technology Council Committee on Technology. October 2016. Preparing for the Future of Artificial Intelligence.
5. Irani, L. C. and Silberman, M. 2013. Turkopticon: Interrupting Worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*.
6. Salehi, N., Irani, L. C., Bernstein, M. S., Alkhatib, A., Ogbe, E., and Milland, K. 2015. We Are Dynamo: Overcoming Stalling and Friction in Collective Action for Crowd Workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*.
7. Scholz, T. and Schneider, N. (Eds.). 2016. *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
8. Stanford Crowd Research Collective. 2015. Daemo: a Self-Governed Crowdsourcing Marketplace. In *Adjunct Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*.