Crowdsourcing the Research Process

RAJAN VAISH and JAMES DAVIS, University of California, Santa Cruz
MICHAEL BERNSTEIN, Stanford University

1. INTRODUCTION

Research is a high skill and resource intensive activity, both in time and effort, and often follows an ad hoc process. In a typical research process, it's often unclear what ingredients; or what recipe or process, which if repeated produces publishable results or a paper. Meanwhile, experienced researchers with novel ideas are constrained with limited time and funding resources; while aspiring researchers (crowd) with exceptional skill-sets lack direction or a research mentor. This prompts research questions, such as: can we crowdsource and organize the complex process of conducting research? How do we handle crowd? In past, with the surge of crowdsourcing potential [Bernstein et al. 2010], researchers have attempted to author academic paper in a large distributed group [Tomlinson et al. 2012] and solved highly complex problem via the polymath project [Cranshaw and Kittur 2011].

To work towards the research problems posed, we introduce our project which explores the possibility of crowdsourcing the research process, and massive research collaboration. We developed a framework, as a recipe to do research; produce scaffolding to train aspiring researchers; and manage crowd, the project and the process. The framework utilizes the concept of iterative peer-grading techniques to organize the process; and allow researchers to systematically apply operators such as split and merge on project ideas, code or crowd.

This paper describes the framework in detail, and its implementations in the field via two pilot studies, including our ongoing work for a large-scale future project. In the first pilot study we produced a published paper, and in the second, we produced better results for an open ended problem being addressed by a top research group. Finally, we discuss the research challenges and opportunities in the space. Encouraged by experimental deployments, we believe that it is possible to build upon the framework and manage the process, the project and the crowd; to produce publishable results or paper.

2. THE FRAMEWORK

The framework for crowdsourcing the research process is developed as a solution to the research questions about the possibility of massive research collaboration between aspiring researchers and the PIs. Overall, the framework will require a large crowd to be split into sub groups, where these groups will be working on essentially the same goals in parallel [Dow et al. 2010; Yu et al. 2014]. This approach encourages diversity, leading to multiple solutions/submissions, which is peer-graded by the crowd themselves. Peer-grading saves time for the PI and top submissions are reviewed by them, which helps in setting the direction of the project. During this process, when the situation leads to time or skill constraints, or multiple similar projects; well allow for applying merge or split operations upon a subgroup of participants or code/ideas. Our framework is primarily built upon the principles of iterative peer-grading techniques and organization of distributed crowd through project operators.

Peer-grading: The iterative peer-grading approach focuses on minimizing PIs involvement, while maximizing student enrollment and their experience. The process begins by recruiting crowd interested to work on a PI proposed research project. A PI can be a professor, researcher or anyone with strong research skills and background. The project proposed would go through a series of research phases: brainstorming, prototyping, development, experiments, results, analysis and paper writing. Spanning through these phases, PI would conduct a periodic interactive meeting and assign a task to the crowd. Following which, crowd would submit an assignment, and peer-grade it based on the rubric provided. Based on the peer-grading process, top submissions are reviewed by the PI, and considered to set the direction for the next iteration or cycle. The PI assigns new tasks, and the process continues.

Operators: Project operators are instructional commands, introduced in the framework to organize the research process and crowd under a variety of conditions. Inspired by the MapReduce algorithm [Dean and Ghemawat 2008] used on big data, the core project operators include split and merge.

The project would begin by introducing split operator causing team formation and encouraging diversity in approaches towards solving a problem. The split operator can further be used when a team is too large to be productive and can be applied on crowd and ideas or approaches. The merge operator can be applied on the same; when the project faces time constraints, or multiple groups have a similar approach to address a problem - thereby, adding value and uniting resources for productivity and achieving a bigger goal. Other operators include, but not limited to: addition, removal, rewire, compete, and reverse. Using a combination of these operators, PIs (or organizers) can direct crowd and research process to maximize productivity under given resource constraints. The conditions will determine whether/when crowd needs to continue to merge, or split (stay parallel) or rewire etc.

Evaluation: To measure the efficacy of the project operators and the process, weekly feedback survey can be conducted [Miller et al. 2014]; complemented by peer-grading scores to track the progress of the project and crowd. Finally, the quality of publishable results or paper along with PIs feedback, and its acceptance at a top-tier venue will determine the validity of the framework developed.

3. EXPERIMENTS

To assess the feasibility and reproducibility of the framework, we have conducted two successful experiments, and currently working towards a large-scale field experiment. We implemented the framework focused on massive research collaboration and producing scaffolding to train aspiring researchers.

3.1 Pilot Study I - producing a publishable paper

Building upon the framework, we conducted our first pilot experiment, aiming to produce a crowd-sourced publishable paper. The controlled experiment was conducted in a computer graphics class for graduate students in the Winter quarter at UC Santa Cruz in 2014. In the beginning of the class, professor proposed three research problems to work on, and the class was divided into three major groups following the split operation. The process involved students selecting the research problem, working on the implementation, generating results and writing sections of the paper in parallel to encourage diversity. The students contributed by writing sections in the order of Related Work, Implementation, Methods and Results and finally Introduction and Conclusion. By applying the iterative peer-grading techniques, the best sections and results from each student were selected and made it to the paper. Interactive periodic meetings with the professor were conducted in the class, regarding the direction of the project and next steps. Unlike individual or direct feedback meetings between PI and students, these meetings were on a broadcast level interacted only at a top level with the peer-graded ideas.

This process resulted in three papers with complete content but unstructured flow from section to section. On the basis of student engagement and interest, the paper with maximum activity was selected to be worked upon further. To optimize the flow of the selected paper, a paper-a-thon was conducted post quarter, following the merge operation on the interested students. After further refinement, paper was pre-peer reviewed twice by a group of four senior graduate students before sending to a relevant IEEE conference [ICIMu 2014]. The paper got accepted [Schuster et al. 2014].

During this process, we learned that handling a project from idea inception to paper writing is a non-trivial task. However, post paper-a-thon efforts proved to be extremely useful. Encouraged by the overall success, and as part of the goal of crowdsourcing research, we next focused on collaborative crowd research for producing publishable results.

3.2 Pilot Study II - producing publishable results

Based on the learning from pilot study 1, we conducted second round of controlled experiment in an undergraduate class in Spring quarter at UC Santa Cruz. To implement our learning, we focused on the research process part instead of paper writing; and recruited students interested in research.

In collaboration with the Stanford Computer Vision Group, we proposed an open-ended research problem to the selected group of students. The project is in progress at Stanford, with few versions live. Students were blindfolded from the approaches used at Stanford, and encouraged to come up with their own ideas; which would later be compared against methods developed at Stanford. The process involved team formation using the split operator, and four teams were formed to work towards a common problem using their unique approach. Teams worked through a series of research phases: brainstorming, paper-pencil prototyping, development and user-evaluation. To boost productivity across these stages, student crowd and process were managed by Stanford and UC Santa Cruz researchers using the framework rules; by applying the operators like split and merge. Using peer-grading techniques, the best approach per-phase determined researchers feedback to all the participating teams; therefore, also encouraging others to iteratively improve their contributions. Interactive periodic meeting was conducted online over a video conferencing software.

By the end of the quarter, solution proposed by the student crowd with no computer vision background exceeded the ones developed by a senior PhD student with expertise in the area - qualitatively and quantitatively. We learned that the framework adopted to manage student crowd and research process is effective. Moving forward we planned to reach out to larger crowd with bigger goals.

3.3 Future experiment - large scale field experiment

After successfully running two pilot studies, we plan to run a large scale field experiment [AspiringResearcherChallenge 2015]; building upon the approach adopted in pilot study II. So far, we have successfully recruited more than 630 students from universities in India, willing to commit 10 to 40 hours per week. These students make our crowd, and are primarily computer science undergraduate students with minimal to no research experience - hence, providing an opportunity to validate our scaffolding to train these aspiring researchers. To validate the approach and diversify the trials, we have divided the crowd into three groups of roughly 200 students each, to be working on data science, computer vision and human-computer interaction projects with professors from UC Santa Cruz and Stanford University. The participating crowd gets to choose a project, and will continue to work for 12 weeks with PIs, following the framework developed to manage processes, crowd and the project.

4. DISCUSSION

We believe that the scope of this project is trying to solve a non-trivial but high impact problem - fostering the state of academic research via crowdsourcing. Our project is not about implementing a software, but validating the algorithmic approach to manage crowd and research process - thereby, also producing scaffolding to train aspiring researchers. Some of the on ground challenges include managing the crowd remotely, specifically when the crowd is not skilled enough to meet project requirements. The practical implementation of the framework on massive crowd is possible, though challenging. However, its also an opportunity to explore the area of massive open online research [MOOR 2013] on the lines of MOOCs. We believe that the project is a first step towards realizing research on a wider scale.

REFERENCES

AspiringResearcherChallenge 2015. UC Santa Cruz + Stanford: The Aspiring Researcher Challenge. (2015). https://aspiringresearchers.soe.ucsc.edu/index.html.

Michael S. Bernstein, Greg Little, Rob S. Miller, Bjoern Hartmann, Mark Ackerman, David Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A Word Processor with a Crowd Inside. In *UIST*.

Justin Cranshaw and Aniket Kittur. 2011. The polymath project: lessons from a successful online collaboration in mathematics. In *CHI*. Vancouver, BC.

Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: simplified data processing on large clusters. In *Communications of the ACM*.

Steven P. Dow, Alana Glassco, Jonathan Kass, Melissa Schwarz, Daniel L. Schwartz, and Scott R. Klemmer. 2010. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Trans. Comput.-Hum. Interact* (December 2010).

ICIMu 2014. IEEE International Conference on Information Technology and Multimedia at UNITEN (ICIMu 2014). (2014). http://icimu.uniten.edu.my/index.php/about.

Robert C. Miller, Haoqi Zhang, Eric Gilbert, and Elizabeth Gerber. 2014. Pair research: matching people for collaboration, learning, and productivity. In *CSCW*.

MOOR 2013. Is Massive Open Online Research the Next Frontier for Education? (2013). http://goo.gl/j3L0Vl.

C. Schuster, B. Zhang, R. Vaish, P. Gomes, J. Thomas, and J. Davis. 2014. RTI Compression for Mobile Devices. In *IEEE ICIMu*. Kuala Lumpur, Malaysia.

Bill Tomlinson, Joel Ross, Paul Andre, Eric Baumer, Donald Patterson, Joseph Corneli, and et. al. 2012. Massively distributed authorship of academic papers. In *CHI EA*. Austin, TX.

Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2014. Distributed analogical idea generation: inventing with crowds. In *CHI*.