

LegionTools: A Toolkit + UI for Recruiting and Routing Crowds to Synchronous Real-Time Tasks

Mitchell Gordon
Computer Science Department
University of Rochester
m.gordon@rochester.edu

Jeffrey P. Bigham
HCI and LT Institutes
Carnegie Mellon University
jbigham@cs.cmu.edu

Walter S. Lasecki
Department of EECS
University of Michigan
wlasecki@umich.edu

ABSTRACT

We introduce LegionTools, a toolkit and interface for managing large, synchronous crowds of online workers for experiments. This poster contributes the design and implementation of a state-of-the-art crowd management tool, along with a publicly-available, open-source toolkit that future system builders can use to coordinate synchronous crowds of online workers for their systems and studies.

We describe the toolkit itself, along with the underlying design rationale, in order to make it clear to the community of system builders at UIST when and how this tool may be beneficial to their project. We also describe initial deployments of the system in which workers were synchronously recruited to support real-time crowdsourcing systems, including the largest synchronous recruitment and routing of workers from Mechanical Turk that we are aware of. While the version of LegionTools discussed here focuses on Amazon's Mechanical Turk platform, it can be easily extended to other platforms as APIs become available.

Author Keywords

Crowdsourcing; Real-Time; Human Computation; Tools

INTRODUCTION AND BACKGROUND

Human computation is the process of using people as part of a computational process, often to solve problems that fully-automated approaches cannot yet solve. Crowdsourcing is the process of making an open call to a group of potential workers. Combined, these provide a powerful method for bringing human intelligence to bear on hard problems on demand.

Recently, synchronous crowdsourcing has been used to create systems capable of responding in real time [4]. However, these tasks require extensive coordination, which adds a challenging hurdle to the process, especially for practitioners in other fields. Bernstein et al. introduced the *retainer model* for pre-recruiting workers in order to reliably have access to human intelligence in less than two seconds for real-time tasks [1]. However, no tools for this process are widely available.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).

UIST '15 Adjunct, November 08-11, 2015, Charlotte, NC, USA
ACM 978-1-4503-3780-9/15/11.
<http://dx.doi.org/10.1145/2815585.2815729>

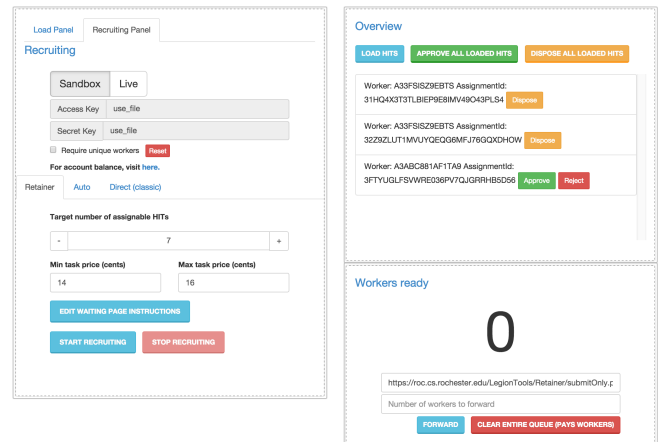


Figure 1. The main LegionTools web interface. Users can easily recruit workers and manage completed tasks all with an easy-to-use interface.

In this poster paper, we introduce LegionTools to provide a publicly-available tool for recruiting and routing workers quickly and efficiently to tasks, with more fine-grain control than has been previously possible. Unlike the standard retainer model, LegionTools routes a user-specified number of workers to a given task, even if more are currently in the retainer pool. This not only allows for more controlled execution parameters for the system using LegionTools, but also improves the average speed of response by filtering for the *fastest N workers* in a set of K .

SYSTEM

This section explains the design and implementation of LegionTools, as well as an easy-to-use GUI interface to the toolkit.

Managing Experiments

LegionTools provides an account system so that multiple researchers or groups of researchers may use a single installation. Users log in using their unique Mechanical Turk credentials. These keys are hashed on the client side, and are never stored or transmitted to the server as plain text.

Each account in LegionTools has its own set of experiments. An experiment represents a single task type. Logging in provides access to all past experiments and allows the user to create new ones, either by entering custom information or duplicating an prior example. Users can switch between experiments on the fly, allowing them to easily manage multiple ex-

periments simultaneously. For instance, on a single instance of LegionTools, multiple researchers could each be actively maintaining multiple retainer pools simultaneously without effecting each others efforts. The idea of being able to easily manage HITs for multiple experiments being conducted on a single set of Mechanical Turk keys as though they are separate is a key contribution of LegionTools.

Recruiting Workers

LegionTools makes it easy to recruit a desired number of workers with the press of a single button. It uses a form of search engine optimization to recruit a desired number of workers quickly by leveraging how Mechanical Turk displays HITs to make tasks more visible to workers by frequently re-posting not-yet-accepted tasks (similar to quikTurkit [2]).

LegionTools makes it very easy to employ the retainer model in a user's experiments. The retainer model allows the user to pool a desired number of workers and in real-time send them to begin a task all at the same time. LegionTools shows a counter with the number of workers who are actively waiting in real time, and the user forwards the desired number of users to a given URL. LegionTools will stop recruiting new workers once the desired retainer size has been reached, and will resume if workers choose to leave the retainer. This automatic maintenance of the retainer pool reduces the attention that requesters must pay to recruiting workers.

Managing Completed HITs

LegionTools allows users to easily accept, reject, and dispose completed or expired HITs. Using the Overview panel at the top left, the user with a single click can view a list of all HITs for their given experiment, and either manually review them or choose to accept them all.

EXAMPLE INTERACTION

To demonstrate LegionTools' flexibility in supporting a variety of use cases, lets take the hypothetical researcher Maria. Maria's lab all shares a single Mechanical Turk account, and they conduct multiple experiments simultaneously. Maria is currently conducting one experiment, and would like to begin another. She opens LegionTools in her web browser, logs in with her Mechanical Turk keys, and creates a new experiment: filling in the HIT title, description, etc. Maria would like all workers to arrive at her task at the exact same time, so she uses the retainer mode. She enters a retainer size of 10, and clicks 'start recruiting.' A moment later, the live counter shows that there are 13 workers waiting. She enters her task's URL and forwards 10 of the workers. To repeat this, Maria need not do anything; workers will continue to be recruited into the retainer and she can forward them as needed. Once Maria is done, she simply clicks 'Stop recruiting,' causing the HIT cleanup process to automatically complete. Once all workers have finished Maria's task, she clicks 'Load HITs' in the Overview panel and clicks 'Accept All.'

At any time, Maria can close LegionTools in her browser and easily come back to right where she was in the process. Maria has completed all of this without effecting any other of the lab's experiments.

REPRESENTATIVE DEPLOYMENTS

LegionTools has made several of our research projects possible. We have used it to recruit and direct workers from thousands of HITs. Several research groups at multiple universities have used also LegionTools as a key part of their projects.

As an example, consider Glance [3], a system that uses the crowd to quickly and accurately code behavioral events in video. It used LegionTools during its evaluation to recruit thousands of workers. LegionTools was left in 'automatic' mode for several days, and constantly recruited workers and sent them to video clips. LegionTools allowed us to continually re-post tasks to prevent them going ignored, while ensuring that only unique workers completed the test (a checkbox in the UI). This sped up completion, and would have previously required a one-off solution.

Another Glance experiment performed a speed test that required all workers to begin analyzing video simultaneously. This required using the retainer mode to recruit workers for just under 20 minutes to reach a pool of over 70 workers. All workers were then forwarded to a single task. 85% of workers began the task within 5 seconds of prompting. To our knowledge, this is the largest ever synchronous workforce recruited from a crowd platform to date.

Another example use was during a workshop with roughly 20 participants grouped into teams of 3-4. Each team was working on creating a small interactive crowd task, and most were unfamiliar with recruiting workers from online crowd platforms. To make it easy to facilitate their task creation, EtherPad was used to create a simple 'live' task, and LegionTools was used by one of the organizers to route workers immediately to tasks once workshop participants' tasks were ready. Over the course of 2 hours, 5 teams requested over 40 workers be directed to tasks, in specified sized groups of between 2 and 5. All workers arrived at their tasks within approximately 5 seconds.

CONCLUSIONS AND FUTURE WORK

LegionTools provides a key resource for systems builders who want to focus on creating new real-time crowdsourcing systems, without needing to build new layers atop crowdsourcing platforms. In on-going work, we are creating an API that will allow systems to create their own predictive allocation algorithms to better automatically control the routing of workers to tasks where their input is needed.

REFERENCES

1. Bernstein, M. S., Brandt, J., Miller, R. C., and Karger, D. R. Crowds in two seconds: Enabling realtime crowd-powered interfaces. In *Proc. of UIST* (2011).
2. Bigham et al., J. P. Vizwiz: Nearly real-time answers to visual questions. In *Proc. of UIST* (2010).
3. Lasecki, W. S., Gordon, M., Koutra, D., Jung, M. F., Dow, S. P., and Bigham, J. P. Glance: Rapidly coding behavioral video with the crowd. In *Proc. of UIST* (2014).
4. Lasecki, W. S., Murray, K. I., White, S., Miller, R. C., and Bigham, J. P. Real-time crowd control of existing interfaces. In *Proc. of UIST* (2011).