

# Understanding Workers, Developing Effective Tasks, and Enhancing Marketplace Dynamics:

A Study of a Large Crowdsourcing Marketplace

[Extended Technical Report]

Ayush Jain<sup>†</sup>, Akash Das Sarma<sup>\*</sup>, Aditya Parameswaran<sup>†</sup>, Jennifer Widom<sup>\*</sup>  
<sup>†</sup>University of Illinois    <sup>\*</sup>Stanford University

## ABSTRACT

We conduct an experimental analysis of a dataset comprising over 27 million microtasks performed by over 70,000 workers issued to a large crowdsourcing marketplace between 2012-2016. Using this data—never before analyzed in an academic context—we shed light on three crucial aspects of crowdsourcing: (1) Task design — helping requesters understand what constitutes an effective task, and how to go about designing one; (2) Marketplace dynamics — helping marketplace administrators and designers understand the interaction between tasks and workers, and the corresponding marketplace load; and (3) Worker behavior — understanding worker attention spans, lifetimes, and general behavior, for the improvement of the crowdsourcing ecosystem as a whole.

## 1. INTRODUCTION

Despite the excitement surrounding artificial intelligence and the ubiquitous need for large volumes of manually labeled training data, the past few years have been a relatively tumultuous period for the crowdsourcing industry. There has been a recent spate of *mergers*, e.g., [27], *rebrandings*, e.g., [15, 17], *slowdowns*, e.g., [16], and moves towards *private crowds* [34]. For the future of crowdsourcing marketplaces, it is therefore both important and timely to step back and study how these marketplaces are performing, how the requesters are making and can make best use of these marketplaces, and how workers are participating in these marketplaces—in order to *develop more efficient marketplaces, understand the workers’ viewpoint and make their experience less tedious, and design more effective tasks from the requester standpoint*. Achieving all these goals would support and sustain the “*trifecta*” of key participants that keeps crowdsourcing marketplaces ticking—the marketplace administrators, the crowd workers, and the task requesters.

At the same time, developing a better understanding of how crowdsourcing marketplaces function can help us design crowdsourced data processing algorithms and systems that are more efficient, in terms of latency, cost, and quality. Indeed, crowdsourced data processing is performed at scale at many tech companies, with tens of millions of dollars spent every year [34], so the efficiency improvements can lead to substantial savings for these companies. In this vein, there have been a number of papers on both optimized algorithms, e.g., [22, 37, 13, 40, 9], and systems, e.g., [33, 19, 32, 10, 38], all from the database community, and such findings can have an impact in the design of all of these algorithms and systems.

Unfortunately, due to the proprietary nature of crowdsourcing marketplace data, it is hard for academics to perform such analyses and identify pain points and solutions. Fortunately for us, one of the more forward thinking crowdsourcing marketplaces made a substantial portion of its data from 2012 to date available to us:

this includes data ranging from worker answers to specific questions and response times, all the way to the HTML that encodes the user interface for a specific question.

This data allows us to answer some of the most important open questions in microtask crowdsourcing: what constitutes an “effective” task, how can we improve marketplaces, and how can we enhance workers’ interactions. In this paper, using this data, we study the following key questions:

- Marketplace dynamics: helping marketplace administrators understand the interaction between tasks and workers, and the corresponding marketplace load; e.g., questions like: (a) how much does the load on the marketplace vary over time, and is there a mismatch between the number of workers and the number of tasks available, (b) what is the typical frequency and distribution of tasks that are repeatedly issued, (c) what types of tasks are requesters most interested in, and what sorts of data do these tasks operate on?
- Task design: helping requesters understand what constitutes an effective task, and how to go about designing one; e.g., questions like: what factors impact (a) the accuracy of the responses; (b) the time taken for the task to be completed; or (c) the time taken for the task to be picked up? Do examples and images help? Does the length or complexity of the task hurt?
- Worker behavior: understanding worker attention spans, lifetimes, and general behavior; e.g., questions like (a) where do workers come from, (b) do workers from different sources show different characteristics, such as accuracies and response times, (c) how engaged are the workers within the marketplace, and relative to each other, and (d) how do their workloads vary?

The only paper that has performed an extensive analysis of crowdsourcing marketplace data is the recent paper by Difallah et al. [14]. This paper analyzed the data obtained via crawling a public crowdsourcing marketplace (in this case Mechanical Turk). Unfortunately, this publicly visible data provides a restricted view of how the marketplace is functioning, since the worker responses, demographics and characteristics of the workers, and the speed at which these responses are provided are all unavailable. As a result, unlike the present paper, that paper only considers a restricted aspect of crowdsourcing marketplaces, specifically, the price dynamics of the marketplace (indeed, their title reflects this as well)—for instance, demand and supply modeling, developing models for predicting throughput, and analyzing the topics and countries preferred by requesters. That paper did not analyze the the full “trifecta” of participants that constitute a crowdsourcing marketplace. Even for marketplace dynamics, to fully distinguish the results of the present paper from that paper, we exclude any experiments or analyses that overlap with the experiments performed in that paper. We describe this and other related work in Section 6.

This experiments and analysis paper is organized as follows:

- **Dataset description and enrichment.** In Section 2, we describe what our dataset consists of, and the high-level goals of our analysis. In Section 2.1 through 2.3 we provide more details about the marketplace mechanics, the scale and timespan of the dataset, and the attributes provided. We also enrich the dataset by manually labeling tasks ourselves on various features of interest, described in Section 2.4, e.g., what type of data does the task operate on, what sort of input mechanism does the task use to get opinions from workers.
- **Marketplace insights.** In Section 3, we address questions on the (a) marketplace load — *task arrivals* (Section 3.1), *worker availability* (Section 3.2), and *task distribution* (Section 3.3), with the aim of helping improve marketplace design, and (b) the types of tasks, *goals*, *human operators and data types*, and correlations between them (Section 3.4), with the aim of characterizing the spectrum of crowd work.
- **Task design improvements.** In Section 4, we (a) characterize and quantify metrics governing the “effectiveness” of tasks (Section 4.1), (b) identify features affecting task effectiveness and detail how they influence the different metrics (Sections 4.3 through 4.7), (c) perform a classification analysis in Section 4.9 wherein we look at the problem of predicting the various “effectiveness” metric values of a task based on simple features, and (d) provide final, summarized recommendations on how requesters can improve their tasks’ designs to optimize for these metrics (Section 4.8).
- **Worker understanding.** In Section 5, we analyze and provide insights into the worker behavior. We compare characteristics of different worker demographics and sources—provided by different crowdsourcing marketplaces; as we will find, the specific marketplace whose data we work with solicits workers from many sources (Section 5.1). We also provide insights into worker involvement and task loads taken on by workers (Section 5.2), and characterize and analyze worker *engagement* (Section 5.3).

## 2. DATASET DETAILS AND GOALS

We now introduce some terms that will help us operate in a marketplace-agnostic manner. The unit of work undertaken by a single worker is defined to be a *task*. A task is typically listed in its entirety on one webpage, and may contain multiple short *questions*. For example, a task may involve flagging whether each image in a set of ten images is inappropriate; so this task contains ten questions. Each question in a task operates on an *item*; in our example, this item is an image. These tasks are issued by *requesters*. Often, requesters issue multiple tasks in parallel so that they can be attempted by different workers at once. We call this set of tasks a *batch*. Requesters often use multiple batches to issue identical units of work—for example, a requester may issue a batch of 100 “image flagging” tasks one day, operating on a set of items, and then another batch of 500 “image flagging” tasks after a week, on a different set of items. We overload the term *task* to also refer to these identical units of work issued across time and batches, independent of the individual items being operated upon, in addition to a single instance of work. The usage of the term *task* will be clear from the context; if it is not clear, we will refer to the latter as a *task instance*.

### 2.1 Operational Details

Due to confidentiality and intellectual property reasons, we are required to preserve the anonymity of the commercial crowdsourcing marketplace we operate on, who have nevertheless been gen-

erous enough to provide access to their data for research purposes. To offset the lack of transparency due to the anonymity, we discuss some of the crucial operational aspects of the marketplace, that will allow us to understand how the marketplace functions, and generalize from these insights to other similar marketplaces.

The marketplace we operate on acts as an aggregator or an intermediary for many different sources of crowd labor. For example, this marketplace uses Mechanical Turk [3], Clixsense [1], and NeoDev [4], all as sources of workers, as well as an internal worker pool. For task assignment, i.e., assigning tasks to workers, the marketplace makes use of both push and pull mechanisms. The typical setting is via *pull*, where the workers can choose and complete tasks that interest them. In some sources of workers that we will discuss later on, tasks are *pushed* to workers by the marketplace. For example, Clixsense injects paid surveys into webpages so that individuals browsing are attracted to and work on specific tasks. In either case, the marketplace allows requesters to specify various parameters, such as a minimum accuracy for workers who are allowed to work on the given tasks, any geographic constraints, any constraints on the sources of crowd labor, the minimum amount of time that a worker must spend on the task, the maximum number of tasks in a batch a given worker can attempt, and an answer distribution threshold (i.e., the threshold of skew on the answers provided by the workers below which a worker is no longer allowed to work on tasks from the requester). The marketplace monitors these parameters and prevents workers from working on specific tasks if they no longer meet the desired criteria. The marketplace also provides optional templates for HITs for some common standard tasks, such as *Sentiment Analysis*, *Search Relevance*, and *Content Moderation*, as well as for tools such as one for image annotation. Usage of these standard templates leads to some uniformity in interfaces, and also gives potential for the improvement of task design simultaneously across requesters.

This marketplace categorizes certain workers as “skilled” contributors, who are given access to more advanced tasks, higher payments, and are also sometimes responsible for meta-tasks such as generating test questions, flagging broken tasks, and checking work by other contributors. Our research highlights that having a pool of engaged and active workers is just as, if not more important than having access to a large workforce. It might be interesting to explore an incentive program for the “active” or “experienced” workers as well.

This marketplace also provides a number of additional features. Notable among them is its module for machine learning and AI. This allows requesters without any machine learning background to generate and evaluate models on training data with easy computation and visualizations of metrics such as accuracy, and confusion matrix.

### 2.2 Origin of the Dataset

Our dataset consists of tasks issued on the marketplace from 2012–2016. Unfortunately, we do not have access to all data about all tasks. There are about 58,000 batches in total, of which we have access to complete data for a sample of about 12,000 batches, and minimal data about the remaining, consisting only of the title of the task and the creation date. Almost 51,000, or 88% of the 58,000 of batches have some representatives in our 12,000 batch sample—thus, the sample is missing about 10% of the tasks. (That is, there are identical tasks in the 12,000 batch sample.) From the task perspective, there are about 6600 distinct tasks in total, spread across 58,000 batches, of which our sample contains 5000, i.e., 76% of all distinct tasks. Thus, while not complete, our sample is a significant and representative portion of the entire dataset of tasks.

We will largely operate on this 12,000 batch sample, consisting of 27M task instances, a substantial number. Figure 1 compares the number of distinct tasks sampled versus the total number of distinct tasks that were issued to the marketplace across different weeks. We observe that in general we have a significant fraction of tasks from each week.

### 2.3 Dataset Attributes

The dataset is provided to us at the batch level. For each batch in our sample, we have metadata containing a short textual description of the batch (typically one sentence), as well as the source HTML code to one sample task instance in the batch. In addition, the marketplace also provides a comprehensive set of metadata for each task instance within the batch, containing

- Worker attributes such as worker ID, location (country, region, city, IP), and source (recall that this marketplace recruits workers from different sources);
- Item attributes such as item ID; and
- Task instance attributes such as task ID, start time, end time, trust score, and worker response.

As we can see in this list, the marketplace assigns workers a trust score for every task instance that they work on. This trust score reflects the accuracy of these workers on test tasks the answers to whose questions is known. The marketplace administers these test tasks before workers begin working on “real” tasks. Unfortunately, we were not provided these test tasks, and only have the trust score as a proxy for the true accuracy of the worker for that specific type of task. In addition to the trust score, we also have information about worker identities, and other attributes of the items being operated on, and the start and end times for each task instance.

At the same time there are some important attributes that are not visible to us from this dataset. For instance, we do not have requester IDs, but we can use the sample task HTML to *infer* whether two separate batches have the same type of task, and therefore were probably issued by a single requester. Nor do we have “ground truth” answers for questions in the tasks performed by workers. However, as we will describe subsequently, we find other proxies to be able to estimate the accuracies of workers or tasks. Finally, we do not have data regarding the payments associated with different tasks and batches.

### 2.4 How did we enrich the data?

The raw data available for each batch, as described above, is by itself quite useful in exploring high-level marketplace statistics such as the number of tasks and workers over time, the geographic distribution of workers, typical task durations, and worker lifetimes and attention spans. That said, this raw data is insufficient to address many of the important issues we wish to study. For example, we cannot automatically identify whether a task operates on web data or on images, and whether or not it contains examples, or free-form text response boxes. To augment this data even further, we *enrich* the dataset by inferring or collecting additional data. We generate three additional types of task attribute data:

- Manual labels—we also manually annotate each batch on the basis of their *task goal*, e.g., entity resolution, sentiment analysis, *operator type*, e.g., rating, sorting, labeling, and the *data type* in the task interface, e.g., text, image, social media, discussed further in Section 3.4.
- Design parameters—we extract and store *features* from the sample HTML source as well as other raw attributes of the tasks that reflect design decisions made by the requesters. For example, we check whether a task contains instructions, examples, text-boxes and images—we discuss these further in Section 4.

- Performance metrics—we compute and store different metrics to characterize the latency, cost and error of tasks to help us perform quantitative analyses on the “effectiveness” of a task’s design, discussed further in Section 4.1.

### 2.5 What are the Goals of Our Analysis?

As previously mentioned, our main goals (at a broad level) are to quantitatively address the questions of (1) Marketplace dynamics — helping marketplace administrators and owners understand the interaction between tasks and workers, and the corresponding marketplace load; (2) Task design — helping requesters understand what constitutes an effective task, and how to go about designing one; and (3) Worker behavior — understanding worker attention spans, lifetimes, and general behavior, for the improvement of the crowdsourcing ecosystem as a whole. While the first two goals directly impact requesters and marketplace administrators, we believe they will also help indirectly improve the general worker experience in terms of availability of desirable tasks, and a reduction in the laboriousness of performing tasks. We now discuss our goals in a little more detail, by breaking each of them down into sub-goals and describing the experiments we perform to answer them.

**Marketplace dynamics.** To understand the marketplace, the first goal of our analysis is to examine general statistics that help us estimate the scale of crowdsourcing operations within the marketplace. This first order analysis is useful for marketplace administrators, helping them estimate the resources required to manage this scale of operations, and identify key limitations. Thereafter, we look at the availability and flow of tasks or workers on the marketplace — specifically, we check if the marketplace witnesses sudden bursts or a steady stream of activity. This analysis gives us concrete take-aways that can help future marketplace design better load-balancing strategies. Lastly, we analyze the manual labels assigned to tasks to look at the types of tasks that have become popular within the marketplace as a means towards a better characterization of the spectrum of crowd work as a whole.

**Task Design.** To help improve task design, we must first be able to characterize the effectiveness of tasks both qualitatively and quantitatively. The three well known aspects used to talk about the effectiveness of a crowdsourced task are (a) latency, (b) accuracy and (c) cost. Consequently, the first step in our analysis is to identify performance metrics that measure these aspects. The next step is to study the impact of varying different design parameters on each of these metrics. This analysis, when performed on a dataset as large as ours, allows us to make data-informed recommendations to requesters looking to design tasks that are answered accurately by workers with low latency and at low cost.

**Worker analyses.** A worker-centric view of the marketplace can help us in understanding the workers’ experience and make changes appropriately to make it easier for them to work. In this respect, we first look at the various labor sources that provide workers that perform work for the marketplace. The load, resource and quality distribution across these sources can point researchers in devising appropriate load-balancing strategies, and can point practitioners towards the ideal source(s) for crowd work. Next, we examine the geographical distribution of workers — this gives us vital information about the active time-zones of the workforce and can help marketplace administrators in ensuring a constant response rate on the marketplace. Finally, we also study the end lifetimes and attention spans of the workers to figure out what fraction of the workforce are regularly active on the marketplace, and how much time is typically spent on the marketplace by workers on a single day.

## 3. MARKETPLACE ANALYSES

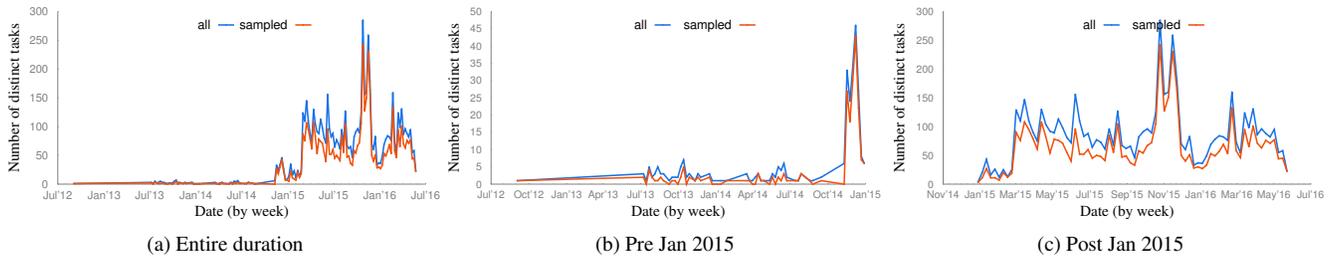


Figure 1: Number of tasks sampled (by week)

In this section, we aim to gain insights into the high level, aggregate workings of the marketplace. First, we examine some basic statistics of the marketplace, to understand the worker supply and task demand interactions. Specifically, we look at (a) task instance arrival distribution (Section 3.1), (b) worker availability (Section 3.2), and, (c) marketplace load, or contribution from “heavy-hitter” tasks, occupying a bulk of the tasks in the marketplace. Then, in Section 3.4, we explore the types of tasks observed in our dataset, to better understand the questions and data types of interest for requesters. We also look for correlations across these labels to understand what types of tasks occur together. Finally, we look at trends in the complexity of tasks over time to gain additional insights into the evolution of the marketplace (Section 3.5).

### 3.1 Are tasks uniform or bursty over time?

We first study the rate at which task instances arrive into the marketplace, and the rate at which they are completed. Note that the load on the marketplace is governed by the number of task instances, which is the fundamental unit of work visible to workers, rather than the number of batches; batches can be arbitrarily small or large. We plot the number of task instances arriving and being completed each week in Figure 2a in blue. First, note that the task arrival plot is relatively sparse until Jan 2015, which is presumably when the marketplace started attracting more requesters. Second, after June 2014, there are some very prominent peaks, on top of regular activity each week. This suggests that while task instances arrive fairly regularly, there are periods of *burstiness*. Considering the period from Jan 2015 onwards, the median of the number of task instances issued in a day on the marketplace is about 30,000. In comparison, on its busiest day, more than 900,000 task instances were issued, a  $30\times$  increase over normal levels. Similarly, the number of task instances issued on the lightest day is  $0.0004\times$  smaller than the median. This raises the question: where does the high variation in the number of task instances derive from—do the number of batches of tasks issued fluctuate a lot, or do the number of distinct tasks issued themselves fluctuate a lot? For this analysis, we overlay the number of task instances issued on the marketplace with the number of batches and the number of distinct tasks issued for the period post January 2015 in Figure 2b. For both these measures, we find that the fluctuation is similar to the fluctuation in the number of task arrivals, indicating that both factors contribute to the high variation in the market load.

Note that a common explanation for why crowdsourcing is used in companies is the ability to shrink or grow labor pools on demand [34]; this finding seems to suggest that even marketplaces need to be able to shrink and grow labor pools based on demand. For this marketplace, having access to both push and pull mechanisms provides great flexibility. Not only can they route the harder tasks to their more skilled, “on-demand” workers, they can also use this push mechanism to reduce latencies for requesters and clear backlogged of tasks. On the other hand, the presence of a large freelance workforce implies that for the majority of tasks they do

not have to rely on the more expensive skilled workers, and can therefore be conservative in their use of resources in maintaining a pool of these internal “super-workers”. This has huge implications for individuals who rely on crowdsourcing as a sole source of income: depending on the week, they may not have enough tasks that suit their interest or expertise.

Striking a good balance between the two task routing mechanisms and worker pools is crucial to ensuring that all three parties are satisfied: (1) the marketplace is able to clear pending tasks without a building backlog, while maintaining requisite levels of accuracy and cost, (2) requesters receive quality responses for easy and hard tasks, and do not see high latencies in responses, and (3) workers have access to as much work as they can handle, as well as tasks that can cater to their varied interests and expertise levels.

Besides the bursty nature of task instance arrivals across weeks, the marketplace also witnesses periods of low task arrivals on the weekends—the number of instances posted on a weekday is up to  $2\times$  the number of instances posted on Saturdays or Sundays on average. Further, the average number of instances posted at the start of the week is the highest, following which the number decreases over the week. We plot this in Figure 3.



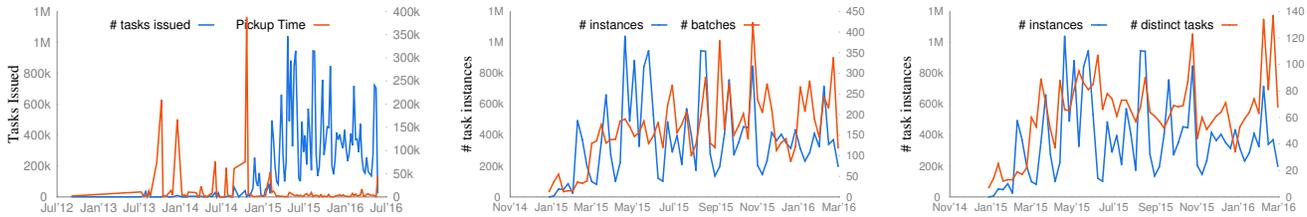
Figure 3: Distribution of tasks over days of the week

*Takeaway: Marketplaces witness wide variation in the number of tasks issued, with daily number of issued instances varying between  $0.0004\times$ , to up to  $30\times$  the median load (30,000 instances).*

### 3.2 How does the availability and participation of workers vary?

**Worker Availability.** As described earlier, the specific marketplace we work with attracts workers from a collection of labor sources. In this manner, it is able to keep up with the spikes in demand. We investigate the sources the marketplace draws from in Section 5. In this section, we focus on studying the number of active workers across different weeks: Figure 4 depicts this statistic.

Unlike Figure 2a that had a huge variation in the number issued task instances, especially after 2015, Figure 4 does not show this level of variation. Thus, somewhat surprisingly, even though there



(a) Task instance arrival vs Median pickup times

(b) Task instance arrival (post Jan 2015) vs (1) batch arrival, (2) distinct task arrival

Figure 2: Task Arrivals by week

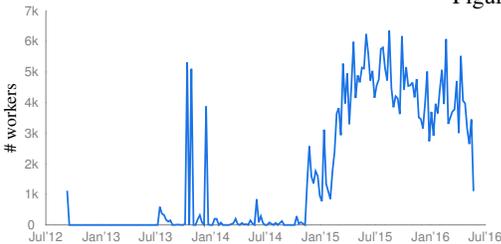


Figure 4: Number of workers performing tasks are huge changes in the number of available task instances, roughly the same number of workers are able to “service” a greater number of requests. This indicates that there is a limitation more in the supply of task instances than availability of workers.

*Takeaway: Despite the huge variation in the number of available tasks, roughly the same number of workers (with small variations) are able to service all of these tasks.*

**Worker Latencies, Idleness, and Task-Distribution.** We now attempt to explain how roughly the same number of workers are able to accommodate for the variation in the number of tasks on the platform. Our first observation is that the median latency in task instances getting picked up by workers, noted as *pickup time* in Figure 2a, and depicted in red, shows that during periods of high load, the marketplace tends to move faster. We also zoom in to the high activity period after January 2015 in Figure 5a to further highlight this trend. One possible explanation for this observation is that when more task instances are available, a larger number of workers are attracted to the marketplace or recruited via a push-mechanism—leading to lower latencies. Another possibility (supported by our discussion below for Figure 5b) is that with a higher availability of tasks, workers are spending a lot more active time on the platform, and hence are more likely to pick up new tasks as soon as they are available.

Next, we look into how the workload is being distributed across the worker-pool. In Figure 5b, we plot the number of tasks completed by the top-10% (in red color) and the bottom-90% (in green color) of workers in each week and compare it to the total number of tasks issued. We observe that while the bottom-90% also take on a lot more tasks during periods of high load, it is the top-10% that handles most of the flux, and is consistently performing a lot more tasks than the remaining 90%. Similarly, examining the same plot for average amount of *active time* spent by workers on the platform in Figure 5b also shows that the top-10% are indeed spending a lot more time on average per week to handle the varying task load as compared to the bottom-90%. This observation indicates that while having a large workforce certainly helps, it is crucial to focus on worker interest and engagement—attracting more “active” workers can allow marketplaces to handle fluctuating workloads better. We also examined the workload handled by workers from different labor sources to verify whether the majority of this variation is assigned to the marketplace’s internal or external workers. We ob-

served that the internal workers account for a very small fraction of tasks. task arrival overlay with internal and external

### 3.3 What is the distribution of work across different distinct tasks?

Next, we wanted to study whether there are a small number of tasks that dominate the marketplace (e.g., repeatedly labeling different items, issued by a single requester). To study this, we first clustered the batches in our dataset based on metadata from the extracted HTML source corresponding to the tasks (see Section 2.4), and tuned the threshold of a match to ensure that the tasks that on inspection look very similar and have similar purposes are actually clustered together. We shall henceforth refer to these clusters of similar batches corresponding to a distinct task, as simply *clusters*. We denote the number of batches in a cluster by *cluster size*. Then, in Figure 6, we plot the distribution of the number of clusters that have different cluster sizes (both on log scale). For example, there were 5 clusters with size larger than 100, indicating that there were 5 distinct tasks (each lumped into their own clusters) that were issued across at least 100 batches each. As can be seen in the chart, there seem to be a large number of tasks that are “one-off” with a small number ( $< 10$ ) of batches: these tasks, being one-off, cannot benefit from much fine-tuning of the interface prior to issuing the task to the marketplace. On the other hand, there are a small number of “heavy hitters”: more than 10 tasks had cluster sizes of over 100, indicating that these tasks had been issued across 100s of batches. Notice that even within a batch the number of tasks may be large: we study that in the next plot in Figure 7. We see a wide variation in the number of tasks issued - while 204 clusters have less than 10 tasks issued, 3 clusters have more than 1M tasks each. Furthermore, these “bulky” clusters have issued close to 80k tasks/batch, so even slight improvements in the design of these batches can lead to rich dividends for the requester. Across this chart, the median number of tasks per cluster is 400. Next,

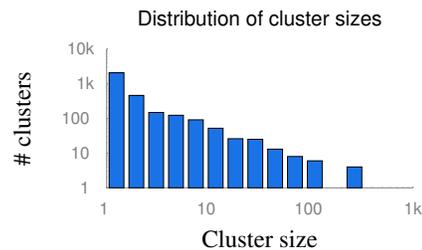


Figure 6: # of batches in clusters

we drill down into the top 10 tasks which had over 100 batches, the so-called “heavy hitters”. In Figure 8 we plot the cumulative number of tasks issued over time, one line corresponding to each heavy hitter distinct task. As can be seen in the figure, these tasks

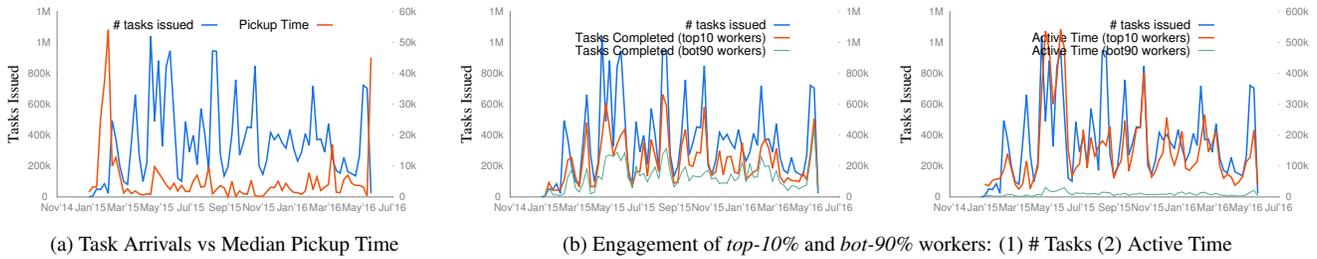


Figure 5: Task Arrivals by Week (Post Jan 2015)

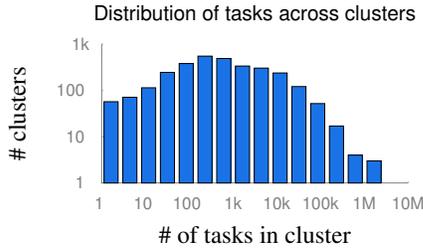


Figure 7: # of tasks in clusters

show both uniform and bursty availabilities. As an example, the task corresponding to the purple line has only been active in July 2015 while the task corresponding to cyan line has had batches issued regularly over 11 months from May 2015 to April 2016.

*Takeaway: A huge fraction of tasks and batches come from a few clusters, so fine-tuning towards those clusters can lead to rich dividends. The heavy hitter task types have a rapid increase to a steady stream of activity followed by a complete shutdown, after which that task type is never issued again.*

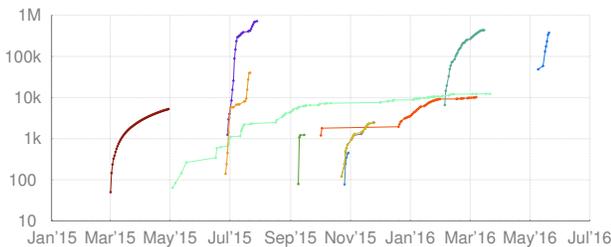


Figure 8: Heavy Hitter Distribution

### 3.4 What kinds of tasks do we see?

We now study our enriched task-labels from Section 2.4 in order to characterize the spectrum of crowd work on in the marketplace. Such an analysis can be very useful, for example, to develop a workload of crowdsourcing, and to better understand the task types that are most important for further research.

**Label Categories.** We label each task under four broad categories<sup>1</sup>. Tasks have one or more label under each category. Our mechanism to label tasks is to first cluster batches together based on similarity of constituent tasks, and then we label one task corresponding to each cluster, since all tasks within each cluster have identical characteristics. The goal of our clustering is to capture the separation between distinct tasks, which is not known to us. As labeling is a labor-intensive process, we currently have labels available for about

<sup>1</sup>Labeling was performed independently by two of the authors, following which the differences were resolved via discussion.

10,000 out of the total 12,000 batches ( $\approx 83\%$ ) and 24 million out of the total 27 million task instances ( $\approx 89\%$ ). These batches fall into about  $\sim 3,200$  clusters.

- **Task Goal:** Here, we separate tasks based on their end goal. We find that most tasks can be characterized as having one (or more) of the following 7 goals<sup>2</sup>: (1) Entity Resolution (ER), for instance, identifying if two webpages talk about the same business, or if two social media profiles correspond to one single person, (2) Human Behavior (HB), including psychology studies, surveys and demographics, and identifying political leanings, (3) Search Relevance Estimation (SR), (4) Quality Assurance (QA), including spam identification, content moderation, and data cleaning, (5) Sentiment Analysis (SA), (6) Language Understanding (LU), including parsing, NLP, and extracting grammatical elements, and (7) Transcription (T), including captions for audio and video, and extracting structured information from images.
- **Task Operator:** In this category, we label tasks based on the human-operators, or underlying data processing building blocks used by requesters to achieve tasks' goals. We observe primarily 10 different operators: (1) Filter (Filt), i.e., separating items into different classes or answering boolean questions, (2) Rate (Rate), i.e., rating an item on an ordinal scale (3) Sort (Sort), (4) Count, (5) Label or Tag (Tag), (6) Gather (Gat), i.e., provide information that isn't directly present in the data, for instance by searching the web, (7) Extract (Ext), i.e., convert implicit information already present in provided data into another form, such as extracting text within an image. (8) Generate (Gen), i.e., generate additional information by using inferences drawn from given data, using worker judgement and intelligence, such as writing captions or descriptions for images, (9) Localize (Loc), i.e., draw, mark, identify, or bound specific segments of given data and perform some action on individual segments, e.g., draw bounding boxes to identify humans in images, and (10) External Link (Exter), i.e., visit an external webpage and perform an action there, e.g., fill out a survey form, or play a game.
- **Data Type:** We also separate tasks based on the type of data that is used. The same goals and operators can be applied on multiple data types. All tasks contain a combination of the following 7 types of data: (1) Text, (2) Image, (3) Audio, (4) Video, (5) Maps, (6) Social Media, and (7) Webpage.

**Label distribution.** First, we analyze the distribution of labels in different categories across tasks. Figure 9a depicts the popular goals. We observe that complex unstructured data understanding based goals—language understanding and transcription are very common, comprising of over 4 and 3 million tasks, that is around

<sup>2</sup>Tasks that had uncommon or unclear goals and did not fall into one of these classes, were automatically classified as Other or Unsure respectively. This holds for the other categories besides goals as well.

17% and 13% respectively, despite not having seen extensive optimization research, as opposed to traditional, simpler goals like entity resolution and sentiment analysis that have been extensively analyzed. Figure 9b shows that text and image are still the main types of data available and analyzed — 9.6 million (40%) and 6.3 million (26%) tasks contained text and image data respectively. Audio and video data are also used, and other richer types of data like social media, web pages, and maps are gaining popularity. Figure 9c shows the common operators used. While the distribution of goals indicates that a significant fraction of tasks have complex goals, the underlying operators are still predominantly simple. The marketplace is dominated by the fundamental filter and rate operations — over 8 million (33%) tasks employ some filtering operator, and nearly 3 million (13% of) tasks make use of rating operators. Among more complex operators, we see that gathering, extraction, localization, and generation are frequently applied, together being used in around 5.3 million, i.e., 22% of all tasks.

**Goals, operators and data types that occur frequently together.**

Next, we look the correlations between the three types of labels for tasks. For example, one question we aim to answer is what kinds of operators are typically applied to different types of data, or used to achieve particular goals? Looking at such correlations across goals, operators, and data types provides fine-grained insights into the structure and design of tasks that is not immediate from our aggregate statistics alone. The charts depicting the correlation can be found in Figure 9 and 10. For instance, Figure 10b shows the breakdown for each goal by the percentage usage of different operators towards achieving that goal; Figure 9c serves as a legend for the stacked bars. (Figure 11b, legend Figure 9a, yields similar insights, but from a slightly different perspective.) We observe that filter and rate operators are used in most kinds of tasks, as well as form a significant majority as the constituent building block for most goals. One notable exception is transcription (which, recall, constitutes over 13% of all tasks by itself, making it a significant exception), where the primary operation employed is extraction. As another example, Figure 10a shows that text and images are important for all types of task goals, for certain types, e.g., ER, SA, SR, social media is also quite important. Lastly, Figure 10c shows that beyond filtering and rating being important, extraction is used quite prominently on text and image data, often rivaling that of filtering. For language understanding tasks, while filter and rate are the primary operations, generate is also used frequently (16% of the time). Also, for tasks looking to understand human behavior, 13% of the tasks involve performing operations at external links (such as online surveys), and 9% of the tasks involve localization. As Figure 11c (legend Figure 9b) indicates, filter and rate operators have been used to analyze most types of data as well.

Figure 10a shows the breakdown for each goal by the percentage of different data types present in tasks having that goal. Figure 9b serves as a legend for the stacked bars. (Figure 11a, legend Figure 9a, yields similar insights, but from a slightly different perspective.) While for most goals, a large fraction of data used in tasks seems to be text and image based, we observe that for entity resolution and search relevance, web data is relevant (serving 24% and 37% of entity resolution and search relevance tasks respectively). Also, sentiment analysis and language understanding style of analyses employ social media as a significant fraction of their input data (13% and 8% respectively). While some efforts are being made towards analyzing other types of data (besides text and image), they are still largely unexplored.

*Takeaway: We observe that the marketplace exhibits a diverse range of tasks spanning across over 7 broad goals, at least 10 distinct operations and 7 data types.*

- *Text and image data are by far the most prevalent and utilized across most tasks. Web and social media data are also available and relevant to a small subset of tasks (in particular tasks involving data integration and cleaning for web, and natural language processing for social media data). While text and image data (and to a lesser extent, web data) have been heavily studied using several different operators, there are still many exciting avenues waiting to be explored for the other types of data.*
- *Filter and rate are used as basic operators for achieving most goals and analyzing all types of data. It is crucial to understand and optimize the usage of these operators.*
- *Language understanding, and transcription seem to be very popular task goals constituting of a large number tasks. Considering the fact that these tasks require complex human operations (generation and extraction as opposed to the simple filter and rate operations), it might be worthwhile to train and maintain a specialized worker pool for such tasks.*
- *For the popular goals of Language understanding, and transcription, we expect the heavy percentage of text-based data. It is interesting to observe the high percentage of of social media and image data for these tasks as well.*

**3.5 Trend towards open-ended tasks.**

In this section, our aim is to explore the trend in the complexity of crowdsourced tasks over time. That is, we intend to answer the following questions:

- Are requesters moving on to more complex, or open-ended goals?
- Are they looking at more challenging datasets?
- Are they using more sophisticated tools or utilizing human intelligence more effectively than in the past?

We split each of our categories, goals, operators, and data into two classes: *simple* and *complex*. Among the set of observed goals, we classify {entity resolution, sentiment analysis, quality assurance} as simple, and the remaining 7 as complex. For operators, we classify filter and rate as simple and the remaining 8 as complex. For data types, we only consider text as simple, and the remaining 6 types as complex. While this classification is subjective, our high-level observations apply to most reasonable mappings of labels to {simple, complex}.

In Figure 12, we compare the trend between number of simple tasks and the number of complex tasks on the marketplace over time. On the x-axis, we plot *time in increments of one week*. On the y-axis we plot the *cumulative number of clusters of tasks*, that is the number of unique tasks, issued so far – one line each for simple, versus complex tasks. Note that we deduplicate similar batches and count them as a single point, so these plots represent the interests of all requesters equally. From Figures 12a and 12c, we observe that the number of clusters of tasks involving complex goals and non-textual data types is much larger, and growing faster, than the corresponding numbers of simple clusters. For instance, as of January 2016, there had been around(a) 510 clusters involving non-textual versus about 240 clusters involving text data, and (b) 620 clusters with complex goals, and just 80 clusters with simple goals. By contrast, Figure 12b demonstrates that the usage of complex operators is comparable to that of simple operators. Specifically, we observe a total of around 410 clusters using complex

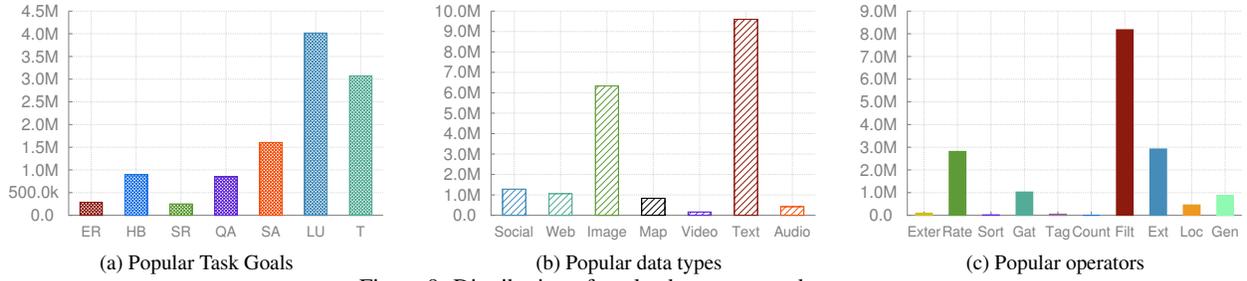


Figure 9: Distribution of goals, data types, and operators

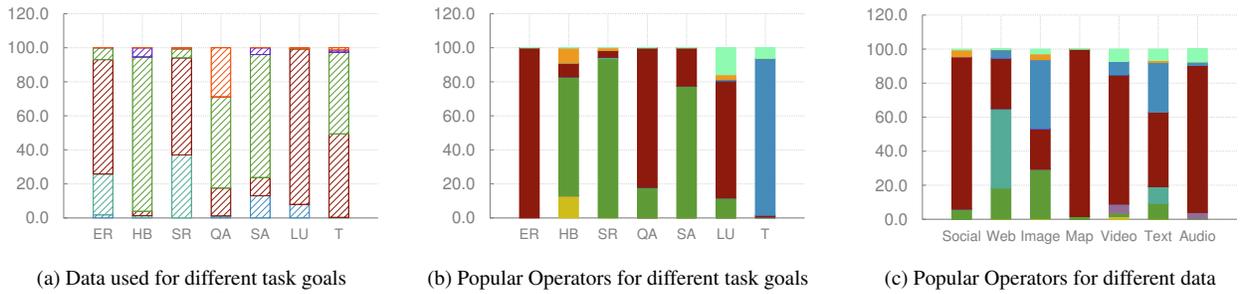


Figure 10: Correlations across data and goal, operator and goal, and operator and data

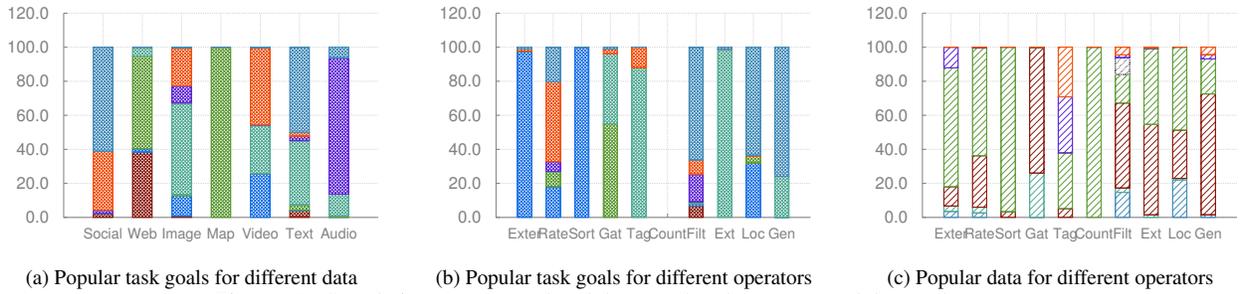


Figure 11: Correlations across goal and data, goal and operator, and data and operator

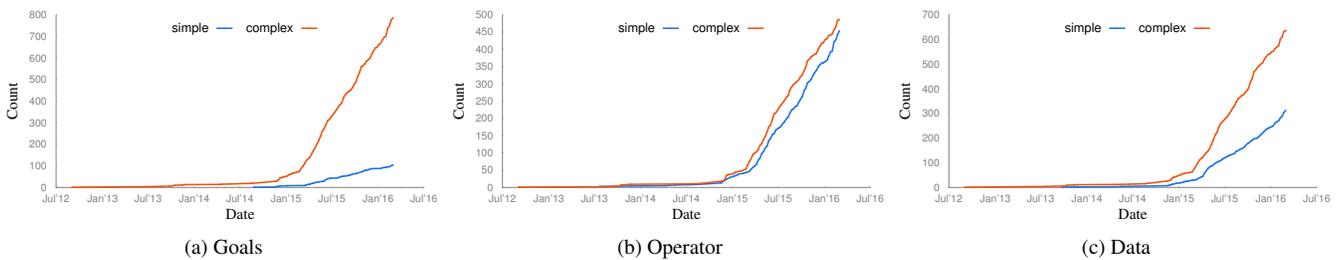


Figure 12: Simple vs complex tasks over time

operators and 340 clusters involving simple operators issued cumulatively as of January 2016.

*Takeaway: While requesters (and researchers) are more interested in achieving complex goals on complex data (and getting more so with time), the fundamental human-operators of filter and rate, are by themselves still as widely used as all other operators combined. This raises both the need to optimize the existing simple operators as far as possible, as well as the opportunity for the exploration and understanding of more complex operators.*

## 4. EFFECTIVE TASK DESIGN

In this section, we address the question of effective task design. Specifically, we (1) characterize and quantify what constitutes an “effective” task, (2) make data-driven recommendations on how requesters can design effective tasks, and (3) predict the “effectiveness” of tasks based on our hypotheses.

### 4.1 Metrics for effective tasks

The standard three metrics that are used to measure crowdsourcing effectiveness are: error, cost, and latency. There are various ways these three metrics could be measured; we describe our notions below, given what we can calculate.

**Error: Disagreement Score.** In our dataset, we have every worker answer provided for each question within each task instance, operating on one distinct item, but not the corresponding ground truth answers. We use these answers to quantify how “confusing” or ambiguous a task is, overall. The way we quantify this is to consider the worker answers for a given question on a given item. If the workers disagree on a specific question on a specific item, then the task is likely to be ambiguous—indicating that it is poorly designed, or hard to answer—either way, this information is important to dictate the task design (e.g., clarify instructions) and the level of redundancy (e.g., more redundancy for confusing questions) that should be adopted by requesters. Our proxy for error is the *average disagreement in the answers for questions on the same item*, across all questions and items in a batch. We consider all pairs of workers who have operated on the same item, and check if their answers are the same or different, giving a score of one if they disagree, and zero if they agree; we then compute the average disagreement score of an item by averaging all these scores; and lastly, we compute the average disagreement score for a batch by averaging the scores across items and questions. We shall henceforth refer to the “Disagreement Score” as disagreement.

There is however, one small wrinkle. Some operators, and corresponding worker responses may involve textual input. Two textual responses may be unequal even if they are only slightly different from each other. Since textual responses occupy a large fraction of our dataset, it is not possible to ignore them altogether. We instead adopt a simple rule: we prune away all tasks with disagreement  $> 0.5$  so as to eliminate tasks with very high variations in worker responses. This eliminates the subjective textual tasks, while still retaining the textual tasks that are objective.

Another way to handle the subjectivity of tasks is to simply ignore text-boxes. This could be done in two ways: (1) only evaluate disagreement for tasks with no text-box fields, and (2) for every task, compute disagreement only on its non textual fields. In our experiments, we tried both these options, but rejected them for reasons we discuss below.

It turns out that a large majority of tasks in our dataset contain at least one text-box field. Eliminating all of them leaves very few

tasks, spread out across a large number of features (such as those discussed in the sections to follow) and labels (goals, operators and data types)—the remaining data is too sparse for any statistically significant inferences to be made.

For the second option of computing disagreement on non-textual fields, we face a problem with the distribution of disagreement values itself. First, for all the tasks that only have textual responses, it is not possible to define a disagreement score; we are unable to compute a disagreement score in this manner for over 60% of all batches. Second, ignoring text fields misrepresents the true distribution of disagreements for the remaining datasets. It is possible that we represent tasks with high disagreement as having low disagreement simply because they have a small number of non-textual fields.

A third approach would be an edit-distance or partial scoring scheme; however, this approach is not ideal since in practice crowdsourcing requesters require high exact agreement, not partial agreement, so that the answers can be easily aggregated via conventional majority vote type schemes. Furthermore, many tasks with textual responses are objective. Some common examples that we see include transcription, captcha, image labeling, and retrieving URLs. For such textual but objective tasks, it is not clear if an edit-distance based agreement scheme is the right approach.

**Cost: Median Task Time.** A typical measure for how much effort a worker has put into a task instance is the amount of time taken to complete it. Since we do not have information about the actual payments made to workers, we use the median amount of time taken (*in seconds*) by workers to complete tasks in a batch as a proxy for the cost of the batch. This can be calculated from the data that is available, given that we have the start and end times for each task in a batch. We shall subsequently denote the “Median Task Time” by task-time.

**Latency: Median Pickup Time.** To characterize latency, we use pickup time, i.e., how quickly tasks are picked up by workers, on average. Pickup time for a batch is computed as follows:  $\text{pickup-time} = \text{median}(\langle \text{start time of task}_i - \text{batch start time} \rangle)$  (*in seconds*). Here, we use the start time of the earliest batch, i.e. start time of  $\text{task}_1$ , as a proxy for the batch start time. We justify this choice for the latency metric quantitatively in below. Our reasons are twofold. First, we observe that the pickup-time of tasks is typically orders of magnitude higher than that of task-time, which might otherwise seem like a reasonable proxy for latency. This means that the actual turnaround time for a task is dominated by when workers start its instances, rather than how long they take to complete them once started. Figures 13a and 13b support this claim. For each of the figures, we compare the pickup-time against the task-time, both on the y-axis with varying end-to-end-time along the x-axis. Figure 13a shows this distribution at a batch level, with the median values for pickup-time and task-time being plotted against each batch’s end-to-end-time. Figure 13b shows this distribution at a task instance level, with each task’s individual pickup-time and task-time being plotted against its end-to-end-time, which in this case is simply (pickup-time + task-time) (to reduce the number of points in the plot, we only plot the median of pickup-time and task-time corresponding to a vertical splice, that is, we plot one median point for all instances having a common end-to-end-time. We observe that in both plots, the pickup-time is orders of magnitude higher than the task-time. Secondly, most measures of time that we can obtain from our available data strongly depend on features, such as the size and difficulty of a task. Since pickup-time only looks at the time taken for workers to start a task and not how long they spend on it (which, as we have seen, is anyway an insignificant fraction of time), it is relatively independent of such

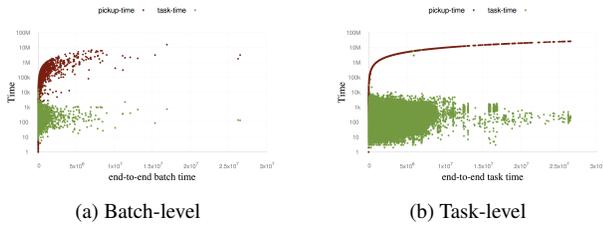


Figure 13: Latency

features. This helps separate out the influence of features that requesters often cannot control, and that we cannot quantify, from our latency metric, making our subsequent quantitative analyses more statistically meaningful. In short, we observe that in general the pickup-time for batches is orders of magnitude higher than the task-time, indicating that the latency or total turnaround time of a task is in fact dictated by the rate at which workers accept and start the task instances. We denote the “Median Pickup Time” by pickup-time.

## 4.2 Correlation Analysis Methodology

In the next set of subsections, we examine some influential features or parameters that a requester can tune, to help improve a task’s error (disagreement), cost (task-time) and latency (pickup-time). For instance, features of a task include the length of the task, or the number of examples within it. For each feature, we look at the correlation between the feature and each of the three metrics. We perform a series of (correlation-investigating) experiments, each of which correspond to one {feature, metric} pair. All our experiments follow the following structure:

- **Cluster:** We first cluster batches based on the task in order to not have the “heavy-hitter” tasks that appear frequently in multiple batches across the dataset to dominate and bias our findings. Since our analysis will also involve *matching*, or clustering tasks further based on labels, we restrict our focus to the set of around 3,200 labeled clusters corresponding to 83% of all batches and 89% of all task instances. Subsequently, for each cluster, we take the median of metric values across batches, as well as the median of the feature being investigated.
- **Binning:** We separate the clusters into two bins based on their feature value — all clusters with feature value lower than the global median feature value go into Bin-1 (say), while the ones with feature value higher than the median go into Bin-2. (Clusters with feature value exactly equal to the median are all put into either Bin-1 or Bin-2 while keeping the bins as balanced as possible.) For each metric, we then examine its value distribution in the two bins — in particular, we look for significant differences between the average, median, or distribution of metric values in the two bins. A significant difference indicates a correlation between the feature we have binned on, and the metric being looked at. We then hypothesize about the underlying reason(s) behind the correlation.
- **Statistical significance:** We perform a *t-test* to check whether the metric value distribution in our two feature-value-separated bins is statistically significant. We use a threshold *p-value* of 0.01 to determine significance, that is, we only reject the null hypothesis (that bins have similar metric values) if the *p-value* is less than 1%.
- **Visualization:** For each feature-metric pair, we plot a cumulative distribution (CDF) plot, with the metric value plotted along the *x-axis*. Each of the two bins corresponds to one line in the plot. For  $x = m$ , the corresponding *y* value on each of the lines represents the probability that a batch will have metric

value better than  $m$ . Thus, a higher value is preferable; and we compare the two bins (or lines) in this plot. Below, we look at the results for some of the significant correlations we found.

## 4.3 Number of HTML words

We examine how the length of task—defined as the number of words in the HTML page, and denoted as  $\#words$ —impacts the effectiveness of the task. We show the effect of length of task on our metrics in Figure 14a. We observe that the line for clusters with higher  $\#words$  in their HTML interface dominates, or is above the line for the clusters with fewer  $\#words$ . We see that the median value of disagreement for tasks with  $\#words \leq 466$  is 0.147, while that for tasks with  $\#words > 466$  is 0.108. This may be because longer tasks tend to be more descriptive, and the detailed instructions help reduce ambiguity in tasks, train workers better, and thereby reduce mutual disagreement in answers. We also note that the length of the task does not significantly affect either the pickup-time or task-time metrics. Thus, workers are neither discouraged nor slowed down by longer textual descriptions of tasks.

While increasing the number of words in the HTML source of tasks helps reduce disagreement in general, this benefit may be more pronounced for particular types of tasks. Intuitively, we expect detailed instructions to help more for harder tasks, and have less impact on easier tasks. To test this hypothesis, we separate tasks into buckets by their labels (recall goal, operator and data), and test the effect of our feature,  $\#words$ . From Figure 25a, we see that for (relatively hard) gather tasks,  $\#words$  has a pronounced effect on disagreement with higher  $\#words$  leading to significantly lower disagreement. On the other hand, Figure 25b seems to indicate that for (relatively simple) rating tasks,  $\#words$  has no significant impact on disagreement.

**Example.** To demonstrate the effect of having more detailed description, or higher number of words in a task’s HTML interface on disagreement, we compare two actual tasks which are both from the domain of Language Understanding, but differ in their descriptiveness and  $\#num-words$ . We look at two different tasks that require workers to find urls or email IDs of businesses through basic web searches. Both have extremely similar interfaces, and ask similar questions. Neither employs examples (which we shall see later has a significant impact on disagreement). The main difference between the two tasks is that the first (having 970 instances), has median number of words = 233, while the second (having 1254 instances) has a median of 6072 words in its HTML interface. Figure 15 depicts the first task and Figure 16 depicts the second. The first task uses these extra words to give detailed instructions (shown in Figure 16a) on how to go about the task. In contrast, the second task has almost no description at all. It requires workers to enter the “synonymy” of correct sentences, and to correct incorrect sentences, without giving any examples or input for what these tasks entail. While the first task has a median disagreement of 0.26, the second shows a median disagreement of 0.08. This demonstrates the power of examples in reducing task ambiguity.

*Takeaway: Tasks with higher  $\#words$  in their HTML sources are typically the ones with more detailed instructions or examples. We see that this has the effect of decreasing disagreement amongst workers, particularly for complex tasks.*

## 4.4 Presence of input text-boxes

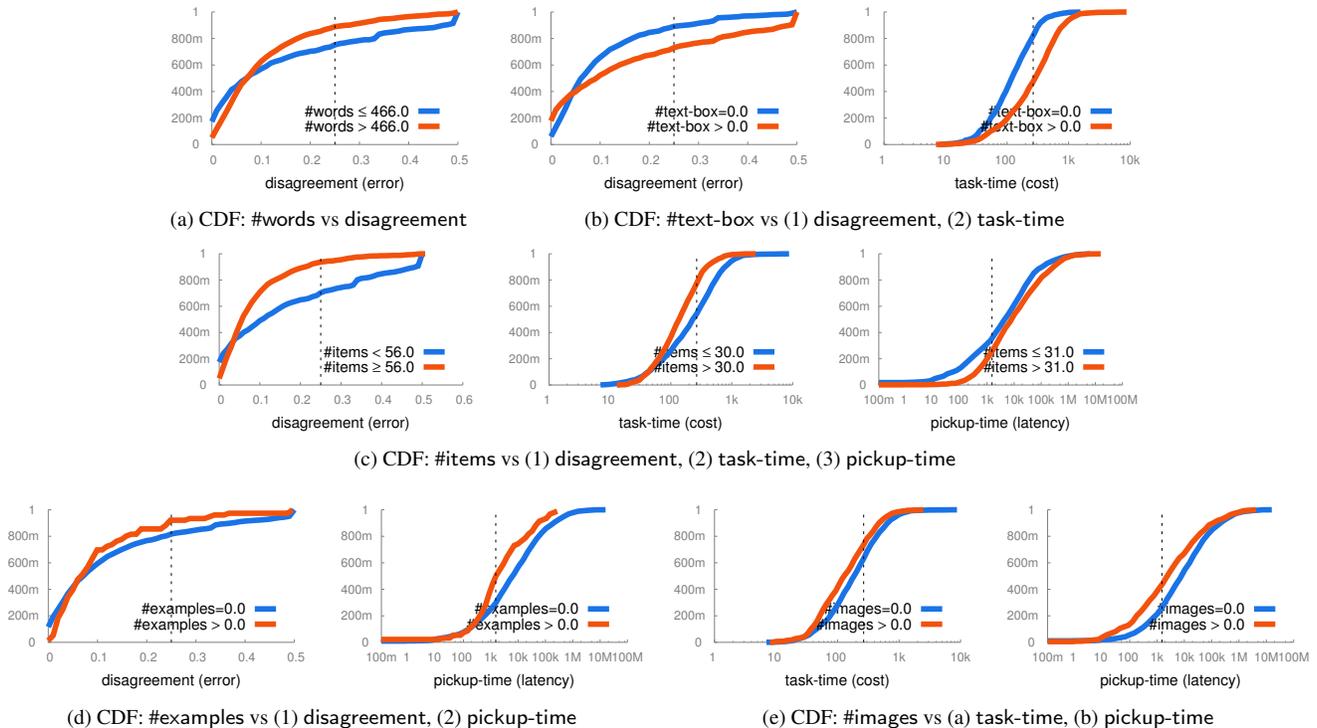


Figure 14: Task Design Parameters and Metrics

(2): Appetite there is, but it does not spring appetite because the hospital food is not delicious.

(3): Since the meal is made is properly calculate the nutrition, etc., please eat firm.

How do you judge the fluency of the sentence (2)?

sentence (2) - Appetite there is, but it does not spring appetite because the hospital food is not delicious.

Incomprehensible

Difficult English

Non-native English

Good English

Flawless English

If you find any error, please correct the sentence (2):

If you select "Flawless", you must input the synonymy text of sentence (2):

Figure 15: Task with low #words

Next, we explore the effect of including text boxes as input fields. We denote the number of text boxes present in the HTML interface as #text-box, and show its effect on disagreement in Figure 14b. Specifically, we compare the set of tasks having non-zero text-boxes, i.e. #text-box > 0, against tasks with no text-boxes, i.e. #text-box = 0. Not surprisingly, Figure 14b shows that there tends to be higher disagreement between workers for tasks with text-boxes. We see that the median value of disagreement for tasks with #text-box = 0 is 0.102, while that for tasks with #text-box > 0 is 0.160. This could be due to the fact that disagreement is agnostic to the input operator type and looks for an exact match of worker answers, while also possibly being affected by the fact that textual tasks may be more subjective (we have however, filtered out all tasks with very high disagreement). We also observe that workers tend to take longer to complete such tasks. We see that the median value of task-time for tasks with #text-box = 0 is 119s, while that for tasks with #text-box > 0 is 286s. Again, this is not surprising, as we expect it to typically take longer to fill out text than to choose from a list of options.

As in Section 4.3, we match tasks based on their labels and dig deeper to check if the insights obtained from our correlation analyses on the complete dataset hold true on individual classes of tasks

as well. From Figure 25c, we see that for sentiment analysis tasks, the presence of text-boxes significantly increases the task-time. Checkboxes or multiple-choice style interfaces are likely to yield much lower task-times than ones based on text-boxes.

**Example.** As a concrete example, we consider two different tasks aimed towards the goal of Sentiment Analysis. Both have extremely simple interfaces, and ask simple questions. The primary difference between the two is that the first, depicted in 17 contains text-boxes while the second, depicted in 18 doesn't.

Both tasks represent a significant number of instances (around 2680 and 8455 respectively). While the first has a median task-time of 141 seconds, the second displays a median task-time of only 37 seconds. If the goal was to truly infer the sentiment of the pieces of text (and not, say, try to understand what different words workers use to describe the same thing), requesters could just have easily provided a list of sentiments to choose from, and thereby reduced the task-time.

*Takeaway: Tasks with a higher number of text-based questions or input require more worker effort (higher task-time) and show higher disagreement between workers. Thus, it pays to simplify questions down to a set of alternatives rather than leaving it open-ended, if possible.*

## 4.5 Number of items

Another parameter of interest is how many items are operated on in a batch across many instances and questions. Anecdotally, the number of items in a batch is known to attract workers, since they can read instructions once and work for longer without having to switch context. We use #items to denote this feature. We observe that when the #items is increased, both the task-time as well as the disagreement metrics improve. That is, tasks get done faster, and workers show lower disagreement when tasks have a higher #items (see Figure 14c). We see that the median value of disagreement for tasks with #items ≥ 56 is 0.086, while that for tasks with

## Choose The Matching Terms And How They Match In The Two Passages

### Instructions

Each task contains a QUESTION and an explanatory PASSAGE. You can IGNORE STOP WORDS (the, of, but, etc) and IGNORE SPELLING MISTAKES.

**STEP 1:** Highlight WORDS in the QUESTION and highlight WORDS in the PASSAGE that match each other in meaning.

**STEP 2:** Select the MATCH TYPE for each selected pair in STEP 1:

- **identical:** the two terms are identical, ignoring spelling mistakes and plurals
- **synonym:** the two terms mean the same thing.
- **generalization:** one term is more general than the other
- **paraphrase:** one term says the same thing as the other, but in a different way
- **negation:** the terms match by being opposites or through negation
- **overlap:** the terms both have two or more words and partially overlap

An example of matching terms and their relation:

EXAMPLE QUESTION:

Do isotopes of element have the same atomic weight?

EXAMPLE PASSAGE:

Similar elements with different atomic weights are called an isotope.

(a) Description

Match terms in the Question to terms in the Passage

QUESTION: Do radio waves get weaker over distance in space?

PASSAGE: Electromagnetic waves propagate more slowly in a medium than in free space, so a wave in a medium will have a larger number of waves than in free space.

What is the type of the matched terms?

In the question or passage, click on a WORD or drag across a RANGE OF WORDS that you want to select. Click on the current selection to undo your selection. An arrow indicates the direction of the selection.

Why is it NOT possible to match more terms?

(b) Task

Figure 16: Task with high #words

Please provide the strongest emotion that applies to the given statement. Examples will be suggested as you type.

\*"Ambiguous" can be applied if the tone of the given sentence can be interpreted as either positive or negative depending on the context.

\*"Neutral" can be applied if the tone of the given sentence conveys neither positive nor negative emotion.

hate hearing noises around the house when i'm the only one up at night :(

Ok he just ate a parrot...wee shame :(

Figure 17: Task with text-boxes

#items < 56 is 0.169. One potential reason for this is that tasks with a high #items attract better and more serious workers. Another explanation is that workers get better with experience (both faster and more accurate). Increasing #items, however, has the effect of increasing the pickup-time of a task for similar reasons—this is probably due to the fact that even though there may be a higher number of items and therefore task instances, the number of available workers (and therefore the parallelism) is still fixed, and therefore the same worker may end up working on different instances in sequence, leading to higher pickup times for the task instances later on in the worker's sequence. We see that the median value of task-time for tasks with #items > 30 is 136s, while that for tasks with #items ≤ 30 is 230s.

Further, we believe that having larger #items would help more for harder tasks, and have less impact on easier tasks. This is supported by our observations from Figure 25e. We see that #items has a pronounced effect on disagreement for (relatively hard) gather tasks with higher #items leading to significantly lower disagreement. Figure 25f on the other hand, seems to indicate that for (relatively simple) rating tasks, #items has insignificant impact on disagreement.

**Example.** We look at two different tasks that require workers to find urls or email IDs of businesses through basic web searches.

Tweet: #AuroraShooting RT @9NEWS UPDATE: 15 confirmed dead #theatershooting from Chief Dan Oates  
Subject: Aurora PD

Choose one:

- Positive
- Neutral
- Negative

Tweet: Tragic news out of the Denver area this morning. \*@9NEWS: 14 confirmed dead #theatershooting from @AuroraPD Dan Oates"  
Subject: Aurora PD

Choose one:

- Positive
- Neutral
- Negative

Figure 18: Task without text-boxes

Both have extremely similar interfaces, and ask similar questions. Neither employs examples (which as we have mentioned, and shall see later has a significant impact on disagreement). The main difference between the two tasks is that the first (having 540 instances), has median number of items = 1, while the second (having 115425 instances) has a median of 1171 items. Figure 19 depicts the first task and Figure 20 depicts the second. (Our snapshot of the task does not depict the high number of items, but does demonstrate the similarity of the two tasks in other respects.) While the first task has a median disagreement of 0.25, the second shows a median disagreement of only 0.04! This demonstrates the power of examples in reducing task ambiguity.

Search: "artur meyster portfolio" on google  
Paste in URL for "artur meyster portfolio" below

SEARCH GOOGLE

Were you able to find the correct, website for 'Artur Meyster Portfolio' ?

- Yes, I found a website for this organization
- No, I could not find a website

The provided name and the name on the website may be slightly different.

Please paste the webpage you found here:

Leave this field blank if you didn't find an official website for this business

Figure 19: Task with high #items

SEARCH GOOGLE

Were you able to find the correct, official email address for 123 Buy Ltd ?

- Yes, I found an email for this organization
  - No, I could not find an official email
- The provided name and the name on the email may be slightly different.

Please paste the email address you found here:

Leave this field blank if you didn't find an official email for this business

Figure 20: Task with low #items

*Takeaway: Increasing the #items or instances, improves the quality (reduces disagreement between workers) and reduces the cost (task-time) of a task, but does not help reduce the latency, due to the limited parallelism available in the marketplace.*

## 4.6 Using examples

It is well-known that examples can have a huge influence on the effectiveness of a task, by training workers on how to answer questions. To study how many examples are used in a task, we count the number of times the word "example" comes wrapped in a tag of its own in the HTML, indicating that the example is prominently displayed. This excludes small or easy-to-miss examples hidden in tasks' textual descriptions, and only counts examples that workers are likely to look at — we denote this parameter by #examples.

Figure 14d demonstrates that examples have the effect of improving worker agreement. We see that the median value of disagreement for tasks with  $\#examples = 0$  is 0.128, while that for tasks with  $\#examples > 0$  is 0.101. We also observe that examples have the effect of reducing pickup times. We see that the median value of pickup-time for tasks with  $\#examples = 0$  is 6303s, while that for tasks with  $\#examples > 0$  is 1353s. It is possible that workers are more inclined to pick up ones that seem more “well-defined” or clear, thereby choosing the ones with examples preferentially over others. We observe no significant correlation between the  $\#examples$  and the task-time — this may be because, the time taken to read and understand examples trades off against the improved speed of performing tasks “post-training”. Finally, we match tasks based on their labels and dig deeper into individual categories of tasks. From Figure 25d, we see that examples have a significant effect on disagreement for the most popular task goal, Language Understanding.

**Example.** To demonstrate the power of examples, we turn to a similar setting as that in Section 4.5. We look at two different tasks that both require workers to find urls of businesses or people through basic web searches (one of which we have seen earlier). Both have extremely simple interfaces, and ask simple questions. A crucial difference between the two tasks is that the first (spanning a significant 1743 instances) provides a detailed example, while the second (having 1006 instances) does not. Figure 21a depicts the example provided by requesters in the first task, and Figure 21b shows the actual task itself. Figure 22 depicts the second “example-less” task. While the first task has a median disagreement of only 0.16, the second shows a median disagreement of 0.45! This demonstrates the power of examples in reducing task ambiguity.

**Process**

1. Perform a search using Google.
2. Look for author name(s) and an email address
3. Choose “Yes” on the first question if author name was found. A text box will appear under the first question. Copy and paste the author name in the box. Add the authors email if it also exists. Copy the complete URL for the article from your browser and paste it into the URL text box.

**Tips and Examples**

- The information about the author will usually appear in the following areas:
  - Directly above/below the title
- In a sidebar or panel next to the article
- Below the article, like a signature
- Author names are often before “prepared by” or just “by”. Sometimes they are **standalone names**. Sometimes they are displayed along with a short profile, picture, contact info etc.
- Some articles don't feature an author. If you have checked all the usual areas (listed above) and haven't found a name, select “No” in the first question and move on to the next article. Look carefully before you settle on a “No”.
- Sometimes author names will not be personal names but titles. Authors like “Business Review Staff Writer” or “Guest Contributor” are perfectly acceptable answers.

(a) Example.

Click here to open a search engine and then perform the following search: **modelling agencies in the UK.**

**Were you able to find an author for this article?**

Yes

No

Unable to view article

Copy and paste the names of the author(s) of this article here:

If there are multiple authors, push the green “plus” on the right of the text box to add another author.

Copy and paste the E-mail of the author(s) of this article here:

Copy and paste the email of the author here

Copy and paste the web address of the article here

Copy and paste the URL of the page here

(b) Task.

Figure 21: Task with example.

Business Name: 2 Dogs Pub

Doing-Business-As Name: Eckhardt Management

Address: 1705 S. 1st Ave., Ste. Q

City: Iowa City

State: IA

Postal Code (Zip): 52240

Country: USA

Phone: 319-337-9047

Business Description: EATING PLACES, RESTAURANTS

SEARCH GOOGLE

**Were you able to find the correct, official website for 2 Dogs Pub ?**

Yes, I found a website for this organization

No, I could not find an official website

The provided name and the name on the website may be slightly different.

Please paste the webpage you found here:

Figure 22: Task without examples.

Both tasks represent a significant number of batches (14 and 6 respectively) as well as instances (around 8000, 40000). While the first has a median pickup-time of only 233 seconds, the second displays a median pickup-time of about 20,000 seconds! This supports our hypothesis that tasks with images get picked up faster.

*Takeaway: Examples are very important; not only do they help reduce disagreement (or task ambiguity, resulting in more confident answers), but they also reduce pickup-time (latency) significantly; tasks with examples attract workers much more quickly than tasks without examples. Despite this we observe that only around 200 task clusters employ explicit examples, as compared to the around 3500 clusters that don't.*

## 4.7 Adding images

We speculate that images can play a role in capturing worker interest, and improving the overall worker experience. To evaluate this aspect, we first count the number of image tags present in the HTML source—we denote this feature as  $\#images$ . We find that around 700 clusters contain at least one image, while around 2200 contain none. Figure 14e shows that tasks with  $\#images > 0$  are picked up faster than those with  $\#images = 0$ . We see that the median value of pickup-time for tasks with  $\#images = 0$  is 7838s, while that for tasks with  $\#images > 0$  is 2431s. We believe that this is due to a similar reason as with  $\#examples$  — workers are attracted to more interesting and well-designed tasks, and images go a long way to help with that. We also drill-down our dataset on task categories to check if the above insight holds true even for specific categories. We plot our observation for tasks with (i) operator Extract in Figure 25g, or (ii) goal Data Quality Control in Figure 25h. These categories have significant number of tasks with and without images and the figures show that our hypothesis that tasks are picked up faster due to the presence of images holds true even when we focus on particular operators or goals.

**Example.** As a concrete example, we consider two different tasks aimed towards the goal of Language Understanding. Both have extremely simple interfaces, and ask simple questions. The first is related to the relevance of text to a given image and the readability of the text. The second is related to discovering events in text. Both use single choice radio buttons as their choice for worker response. The primary difference between the two is that the first, depicted in 23 contains images while the second, depicted in 24 doesn't. (Note that while we have access to the HTML source files, we do not have access to the images referred to in them, *unless they are embedded* — therefore in this particular interface we can only see that there is an image present, but cannot see the actual image.)

Both tasks represent a significant number of batches (14 and 6 respectively) as well as instances (around 8000, 40000). While the

Image:



Generated Text: world's rarest car 1954 concept oldsmobile concept car ncr

### How Relevant is the Generated Comment to the Image

- 0 Irrelevant: Generated by Mistake
- 1 Low
- 2 Medium
- 3 High

### How Readable is the Generated Text

- 0 Completely Not Readable
- 1 Low Readability
- 2 Medium Readability
- 3 High Readability

Figure 23: Task with images

Sales dropped to \$ 441.1 million from \$ 820.4 million , because the company no longer consolidates its Eastern Steel division , which is now a joint venture with Kawasaki Steel Corp .

Does the highlighted word ('Sales') refer to an event?

Yes  
 No

Sales dropped to \$ 441.1 million from \$ 820.4 million , because the company no longer consolidates its Eastern Steel division , which is now a joint venture with Kawasaki Steel Corp .

Does the highlighted word ('dropped') refer to an event?

Yes  
 No

Sales dropped to \$ 441.1 million from \$ 820.4 million , because the company no longer consolidates its Eastern Steel division , which is now a joint venture with Kawasaki Steel Corp .

Does the highlighted word ('441.1') refer to an event?

Figure 24: Task without images

first has a median pickup-time of only 233 seconds, the second displays a median pickup-time of about 20,000 seconds! This supports our hypothesis that tasks with images get picked up faster.

We also observe that tasks with images tend to get completed faster. We see that the median value of task-time for tasks with #images = 0 is 184s, while that for tasks with #images > 0 is 129s. One possible explanation for this is that for tasks with #images > 0, workers are more energetic or “enthusiastic” in completing the task, and visual understanding often takes less time than textual understanding. We observe no significant correlation between the #images and the disagreement of tasks, indicating that these tasks are not inherently easier.

*Takeaway: Tasks with images attract workers much more quickly than ones (lower pickup-time) without, making them a very powerful tool in reducing latency. Further, workers tend to perform tasks with images faster than those without (lower task-time).*

## 4.8 Summary from a metric point of view

In addition to the features we have seen so far, we also looked for correlations between other features and the target metrics. For instance, we examined whether batches were issued on weekdays or weekends, what time of day they were issued at, and how many input fields they had. We observed no significant correlations between these features and any of our metrics. (Recall that for a correlation to be considered statistically significant, we perform a t-test and only those observations with a sufficiently small p-value are considered. For the correlations that we summarize for each of our metrics, the p-values are all significantly below our threshold of 0.01.) We present the quantitative observations corresponding to the noticeable correlations in Tables 1, 2 and 3, and discuss the underlying insights below.

**Disagreement Score.** Table 1 summarizes the effect of features that show correlation with the disagreement of tasks. Based on our observations, we draw the following conclusions: Providing

Feature	Cluster Bins (split at median(feature-value))				disagreement	
	Bin-1	# clusters	Bin-2	# clusters	Bin-1	Bin-2
#words	≤ 466	1150	> 466	1149	0.147	0.108
#items	< 56	1148	≥ 56	1151	0.169	0.086
#text-boxes	= 0	1283	> 0	1014	0.102	0.160
#examples	= 0	2221	> 0	76	0.128	0.101

Table 1: Disagreement Score: summary

Feature	Cluster Bins (split at median(feature-value))				task-time	
	Bin-1	# clusters	Bin-2	# clusters	Bin-1	Bin-2
#items	≤ 30	1511	> 30	1469	230s	136s
#text-boxes	= 0	1565	> 0	1412	119.0s	285.7s
#images	= 0	2268	> 0	709	183.6s	129.0s

Table 2: Median Task Time: summary

detailed instructions for workers can be crucial. If we have multiple items or questions, we should issue them together in one batch (as opposed to scattered across batches) in order to benefit from more experienced workers and workers who get better with experience. Interfaces should use multiple-choice questions to phrase tasks rather than text-based ones wherever possible. Examples are also crucial in reducing errors.

**Median Task Time.** Table 2 summarizes the effect of features that show correlation with the task-time of tasks. Based on our observed correlations, we note that similar to disagreement, it is beneficial to issue items all at once to benefit from workers with experience. Interfaces should use multiple-choice questions to phrase tasks rather than text-based ones wherever possible, as they also affect the typical task time, and correspondingly, worker effort. Adding images not only makes tasks look more pleasing, but also improves worker experience and latency.

**Median Pickup Time.** Table 3 summarizes the effect of features that show correlation with the pickup-time of tasks. Including examples and images is observed to help increase pick-up rate (reduce latency), probably because workers are attracted to more interesting and well-structured tasks. At the same time, issuing more task instances in parallel will lead to increases in the pickup time due to limited parallelism in the marketplace.

## 4.9 Predictive Setting

We further concretize our findings from the previous section by exploring the use of the features for prediction. We demonstrate that using just these features allows for an accurate approximate estimation of various metrics. Due to the high variability in the range of values of our metrics, it is not possible to predict the exact value of a metric for any given task. Instead, we bucketize the range of values into 10 buckets, and try to predict which bucket any given task will fall into. For example, instead of trying to predict disagreement for a given task, we predict whether the disagreement would fall into the buckets [0, 0.1), [0.1, 0.2), . . . , [0.9, 1.0]. There are many different ways in which we could bucketize the range of values—each bucketization also corresponds to distributing tasks into buckets. In the following, we shall use the term bucketization to refer to the bucketization of the metric’s range of values, as well as tasks interchangeably. In our experiments, we consider the two most natural ones: (1) bucketization by range, where we evenly divide the range of metric values into buckets of uniform width, and

Feature	Cluster Bins (split at median(feature-value))				pickup-time	
	Bin-1	# clusters	Bin-2	# clusters	Bin-1	Bin-2
#items	≤ 31	1471	> 31	1470	4521s	8132s
#examples	= 0	2845	> 0	93	6303s	1353s
#images	= 0	2230	> 0	708	7838s	2431s

Table 3: Median Pickup Time: summary

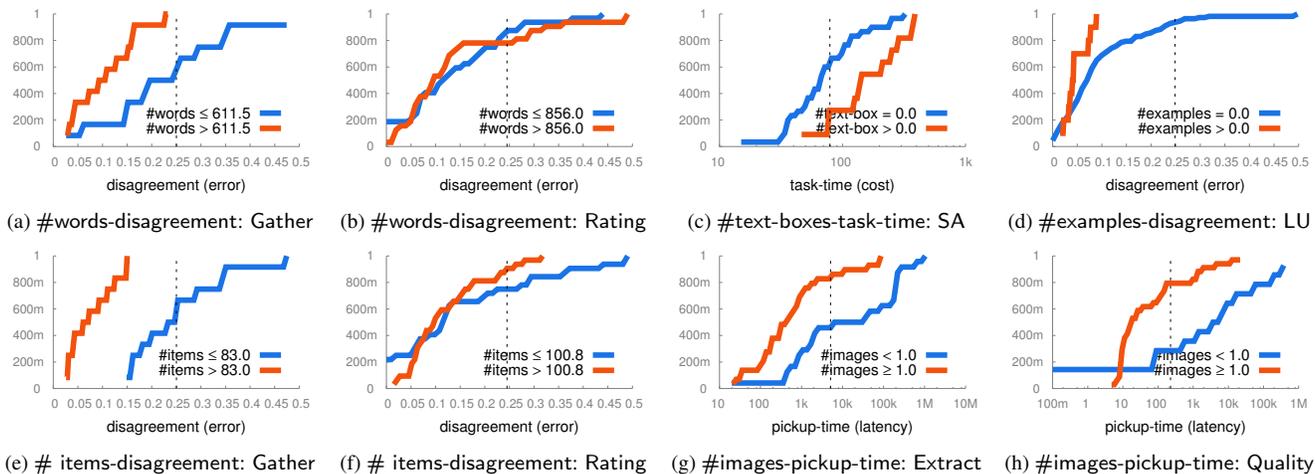


Figure 25: Features-Metrics CDF: Drill down by match on labels

(2) bucketization by percentiles, where we divide the range of metric values into buckets such that all buckets contain roughly equal number of tasks. For each of these two cases, we divide all three of our metrics into 10 buckets. We run a simple decision tree classifier with the following feature sets: (1) features for disagreement:  $\{\#items, has-example, \#words, \#text-boxes\}$ , (2) features for task-time:  $\{\#items, has-image, \#text-boxes\}$ , (3) features for pickup-time:  $\{\#items, has-example, has-image\}$ .

**Bucket distributions.** For each of the metrics, we now discuss the distribution of clusters across the 10 buckets for both bucketization strategies. For the case of bucketization by range, we have:

- (pickup-time) Upper bounds of metric value in buckets (in seconds):  $[1.6 \times 10^6, 3.2 \times 10^6, 4.7 \times 10^6, 6.3 \times 10^6, 7.9 \times 10^6, 9.5 \times 10^6, 1 \times 10^6, 1.3 \times 10^7, 1.4 \times 10^7, 1.6 \times 10^7]$ , and number of clusters in respective buckets: [2906, 17, 8, 5, 1, 0, 0, 0, 0, 1]
- (task-time) Upper bounds of metric value in buckets (in seconds): [882, 1756, 2631, 3506, 4380, 5255, 6130, 7004, 7879, 8754], and number of clusters in respective buckets: [2842, 120, 8, 3, 1, 1, 1, 0, 0, 1]
- (disagreement) Upper bounds of metric value in buckets: [0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0], and number of clusters in respective buckets: [1360, 390, 181, 155, 143, 145, 82, 74, 43, 150]

For the case of bucketization by range, we have:

- (pickup-time) Upper bounds of metric value in buckets (in seconds): [157, 579, 1486, 2955, 5946, 12202, 24796, 53876, 179358,  $1.6 \times 10^7$ ], and number of clusters in respective buckets: [294, 294, 294, 293, 294, 294, 293, 294, 294, 294]
- (task-time) Upper bounds of metric value in buckets (in seconds): [43, 67, 96, 127, 171, 227, 297, 407, 608, 8754], and number of clusters in respective buckets: [300, 297, 298, 296, 301, 297, 296, 297, 297, 298]
- (disagreement) Upper bounds of metric value in buckets: [0.002, 0.019, 0.038, 0.064, 0.1, 0.158, 0.281, 0.497, 0.693, 1.0], and number of clusters in respective buckets: [273, 272, 272, 272, 273, 272, 272, 272, 272, 273]

We perform a 5-fold cross-validation to test the accuracy of our models.

**Bucketization by range.** We observe that we are able to predict the *exact* bucket for tasks of disagreement with accuracy 39%, of task-time with 95%, and of pickup-time with 98%. This is not so surprising given the high skew in their distributions, but knowing

the expected range of time for a task to be completed is still useful for requesters. Note that here accuracy is averaged across the 5 test cases in our cross-validation. For disagreement, we obtain a high accuracy of 62% if we allow an error tolerance of 1 bucket—that is, using just these features alone, we are able to predict within a tolerance of 1 bucket the disagreement for majority of the tasks on average. Given the extremely high dimensional nature of this prediction problem, with a very large number of hidden variables that we have not considered, even the 39% accuracy seen for disagreement is very high. To verify that the accuracies for task-time and pickup-time are not heavily biased by a skew in the distribution of tasks across buckets for these metrics, we also perform a similar cross-validation test for the percentile-based bucketization. We observe that even in this harder case, our model is able to make predictions with reasonable accuracy.

**Bucketization by percentiles.** For the percentile-bucketization, where clusters are divided equally across buckets, the classification problem is much harder because the buckets are of very different and uneven sizes. We observe an accuracy of about 16% for task-time, 15% for pickup-time, and 20% for disagreement. The drop in accuracy for disagreement is less pronounced than that for the remaining two metrics since it has a lower skew in its value distribution. Allowing for a tolerance of one bucket, we see an accuracy of 40% for task-time, 39% for pickup-time, and 44% for disagreement, which is extremely high given the high dimensional nature and small feature set of our classification.

## 5. WORKER ANALYSES

In this section, we adopt a worker-centric view of the marketplace and evaluate the worker demographics and behavior patterns. Specifically, we look at (1) distribution of workers across different sources and regions, (2) lifetimes and attention spans of workers.

### 5.1 Where do the workers come from?

**Labor Sources.** As described earlier, the marketplace we focus on, unlike Mechanical Turk, gets crowd workers from multiple sources: specifically, the marketplace has 139 different sources for crowd labor, altogether supplying around 69,000 workers across the period of our collected data. These sources—all distinct from each other—are listed in Table 4. These sources all link to, and allow workers to sign-up with the marketplace. The marketplace directly compensates workers through one of many mechanisms: money, gift cards, or bitcoins. Some of these sources (e.g., `imerit_india`,

yute\_jamaica, taskhunter) are specific to certain locations in the world, while others provide workers tailored to specific domains of tasks (e.g., ojooo provides workers for advertising and marketing campaigns). In addition, the marketplace also has its own dedicated worker pool (called internal), performing 484k tasks, that is about 2% of all tasks in our collected sample.

*Takeaway: There are a large (over 100) number of active sources of crowd labor, with varying payment schemes; many of these sources are specialized either in terms of demographics, or in terms of the task types.*

We also plot the average number of tasks performed by workers on different sources in Figure 26a. Each vertical splice on the x-axis represents a labor source, and the height of the splice indicates the number of tasks performed by a worker from that source on average. We see a significant variation in the worker loads across source (the y-axis is log-scale). For some sources, workers typically perform more than 10,000 tasks each whereas on the other end of the spectrum, 40% of the sources have workers performing  $\leq 20$  tasks each. The variation in number of tasks per worker suggests the presence of two types of sources — sources having (a) a dedicated workforce performing a large number of tasks per worker, and (b) an on-demand workforce, performing few tasks per worker. Thus, the availability of these two types of sources is an essential load balancing strategy — the dedicated workforce is supplemented by on-demand workforce in periods of high task load. To study this further, Figure 26b shows the number of sources active every week overlaid on the number of tasks issued. This plot seems to indicate that after January 2015, while the marketplace has a relatively fixed number of active sources, the number of tasks issued varied quite a bit. Thus, these active sources are able to absorb the variation in the number of tasks issued.

*Takeaway: Different sources have different characteristics: some involve an engaged workforce, while in others, participation is more one-off. By using a combination of sources, the marketplace is able to absorb a varying task load.*

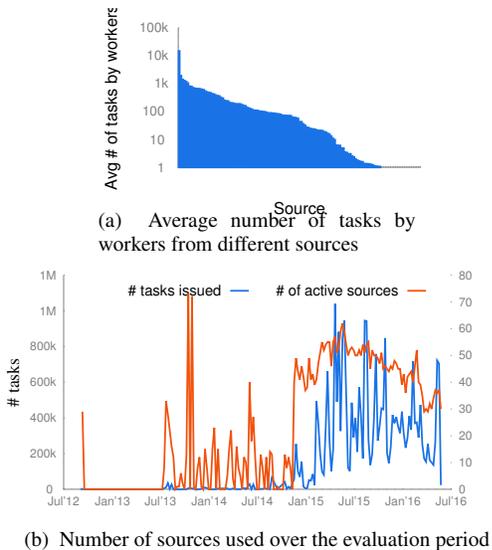


Figure 26: Tasks performed by workers across sources across weeks

Next, we look at the major contributing sources by (a) number of workers, and (b) number of tasks in Figure 27. In Figure 27d,

we show the top 10 sources by the number of tasks performed by its workers. These 10 sources together account for  $\sim 95\%$  of the tasks, and  $\sim 86\%$  of the workers in the marketplace. Figure 27a shows the top 10 sources contributing the most number of workers. Popular sources include Clixsense [1] and NeoDev [4] — companies that provides monetary payment for users taking surveys, and Prodege [5] — a company that rewards workers in gift cards. We note some tasks are also routed to Mechanical Turk (amt) workers, which accounts for  $\sim 1.5\%$  of all workers. While Mechanical Turk has contributed a total of  $\sim 1000$  workers over the period of our evaluation, with a maximum of  $\sim 400$  of their workers being active at any given point of time, by comparison the source NeoDev has contributed a total of  $\sim 27000$  workers in all with as many as  $\sim 2600$  of them being active in a single week. In addition, the marketplace’s own internal workforce (internal) accounts for 2.5% of the total workforce and more than 484k tasks in our sample during the evaluation period.

*Takeaway: The most popular 10 sources account for 95% of the tasks performed on the marketplace. Furthermore, these 10 include many companies that we, the authors, have never heard of.*

**Geographic distribution of workers.** We plot the country-wise distribution of workers in Figure 28. The figure shows that crowd-sourcing has become a truly global phenomenon with workers coming from as many as 148 countries. In a study of Amazon Mechanical Turk’s workforce [29], the authors noted that more than 60% of the workers came from USA and India. For our marketplace, while these countries continue to contribute a significant number of workers, we also see 17% of workers coming from the emerging South American and African markets. Close to 50% of the workers come from 5 countries — USA (21.3k), Venezuela (5.3k), Great Britain (4.4k), India (4.1k) and Canada (2.8k).

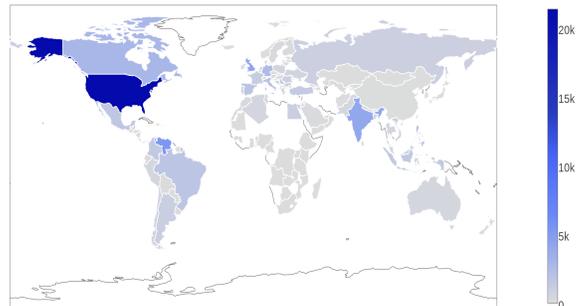


Figure 28: Geographical distribution of the crowd workforce

**Quality across sources.** As we noted earlier, different sources bring workers from different locations and specializing in different types of tasks. Furthermore, while some sources have a dedicated workforce performing a large number of tasks regularly, other sources supply an on-demand workforce that performs a small number of tasks occasionally. Given these variations, we investigate if the *quality* of workers varies across sources. We evaluate the quality of different sources on two metrics. Our first metric is the trust score attributed by the marketplace to each completed task performed by a single worker. We compute and report the *mean trust* assigned to tasks performed by workers from each source. The second metric measures the amount of time taken by workers

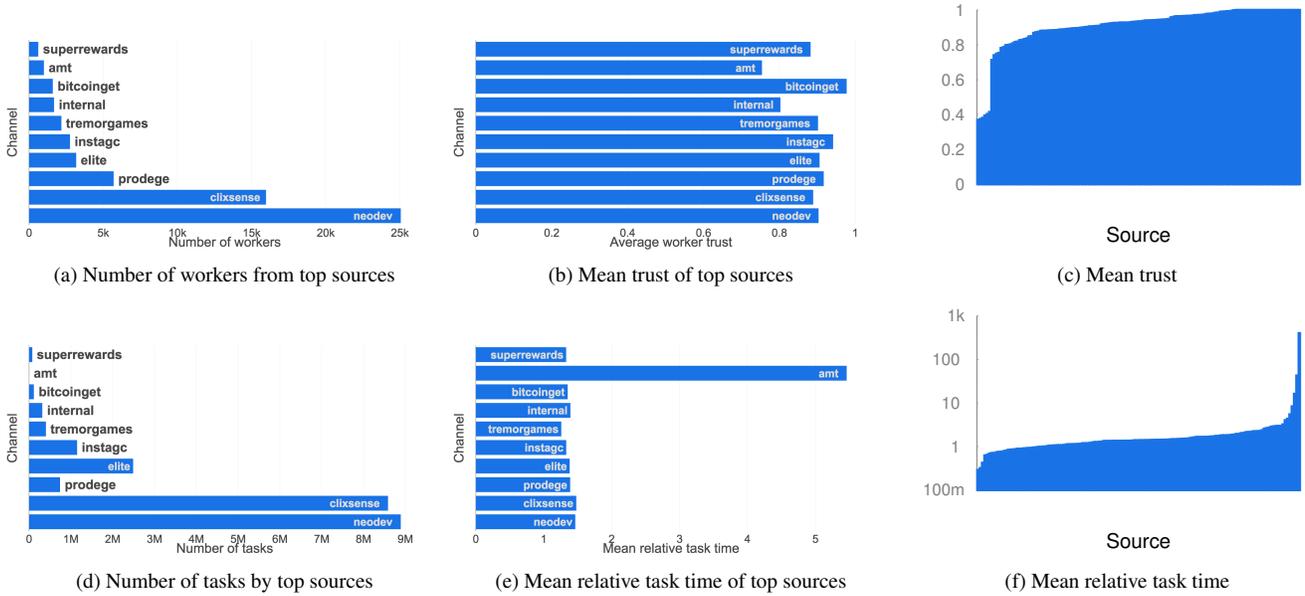


Figure 27: Number of tasks, latency and trust distributions of sources

to complete tasks. To normalize across different tasks, we divide a worker’s time by the median time taken by workers to complete that task. We report the average of these *relative task times* for tasks performed by each source as the second metric of quality.

The variation in quality (trust and latency) for all sources is shown in Figure 27c and Figure 27f. In terms of mean trust, we observe that close to 10% of the sources have mean trust  $< 0.8$ . The trust for some sources is even lower than 0.5. The difference in quality between the sources is more evident when we look at the mean relative task times. While most sources have mean relative task times close to 1, 5% of the sources have mean relative task times  $\geq 3$  — the workers from these sources take more than  $3\times$  time to complete the tasks, compared to median task times. Three of these sources even have mean relative task times  $\geq 10$ .

We further examine the quality of major sources *i.e.*, sources providing the most number of workers, in more detail in Figure 27b and Figure 27e. With the exception of Mechanical Turk (amt), these sources have high quality — having mean trust  $> 0.8$  and mean relative task time  $< 1.5$ . Mechanical Turk performs poorly on both metrics — the mean relative task time is more than 5 and mean trust is 0.75.

*Takeaway: Different sources lead to very different qualities. Amazon MTurk, one of the most popular crowdsourcing platforms, shows poor latency and trust. Other lesser known sources show significantly better quality.*

## 5.2 How do the worker workloads vary?

As our analysis of sources indicated, most of the tasks on the marketplace are completed by a small group of workers. To explore this issue in detail and look at the distribution of worker workloads, we plot the number of tasks performed by each worker in Figure 29a. The x-axis shows the rank of the worker, when workers are sorted in decreasing order of number of tasks completed. The y-axis value shows the number of tasks completed by the worker. From the plot, we note that majority of workers’ participation is one-off and the workload is mostly shared by a small group of workers. In fact, more than 80% of the tasks are completed by just 10% of the workforce.

*Takeaway: Most of the tasks on the marketplace are performed by a small group of workers. Given their experience, it might be worthwhile for marketplaces to collect periodic feedback from these workers and direct some of the more difficult tasks to this worker pool.*

## 5.3 How engaged are crowd workers?

**Worker Lifetimes.** We investigate the workers’ availabilities on the marketplace through two metrics: (1) the *lifetime* of the workers, which is the number of days between their last and first activity on the marketplace during the evaluation period, and (2) the number of *working days* (out of their lifetimes) where workers have taken up tasks.

We first show the distribution of the lifetimes of the workers through a histogram in Figure 30a. Each bar in the x-axis corresponds to a lifetime range, with bar heights denoting the number of workers having lifetime in corresponding ranges. From these plots, we note that 79% number of workers are only available over short time frames and hence have lifetimes of less than 100 days. In fact, 52.7% of the workers have a lifetime of only 1 day in the evaluation period, indicating that a majority of workers are directed to the marketplace through their sources for one-off tasks. However, these workers are not major contributors in terms of number of tasks — they complete only 2.4% of the tasks in the marketplace.

Of the remaining workers who have logged in to the marketplace on more than one day, about one-third have been working on more than 10 days and have completed 83% of the tasks in the marketplace. Next, we focus on these *active workers* and explore their behavior in greater detail.

*Takeaway: More than half of the workers on the marketplace are only active on a single day. Only 15% of the entire workforce comes on the marketplace repeatedly to complete tasks. This active workforce completes more than 80% of the tasks in the marketplace.*

## 5.4 Active Worker Characteristics

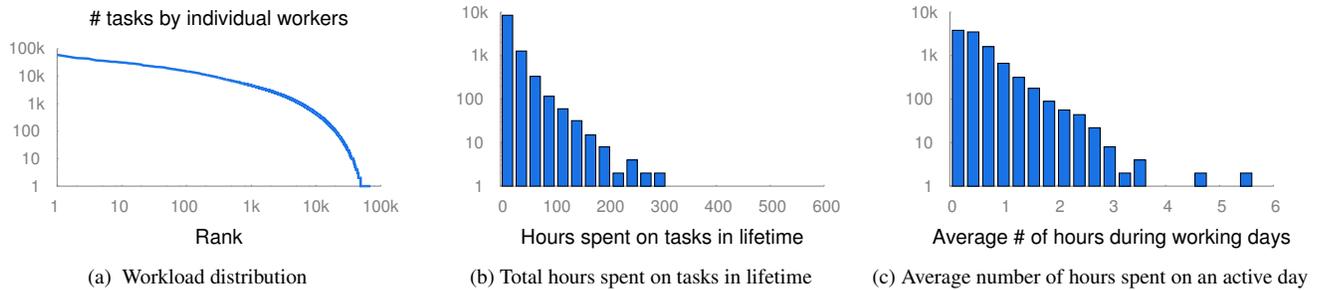


Figure 29: Workload and Time spent distributions

**Distribution of Working Days.** Consider the distribution of working days for the active workers in Figure 30b. First, on the right side of the plot, note the presence of workers who have been working on more than 350 days. This is especially remarkable, considering the fact the evaluation period contains regular data for only about 18 months. Second, the bar heights reduce close to linearly (in log scale) with the active days, indicating that the availability of workers decreases exponentially with experience.

For the active workers, we also plot a histogram of the fraction of lifetimes days where they have been working in Figure 30c. We note that among these active workers, more than 43% are working at least once a week (on average) during their lifetimes.

*Takeaway: The marketplace has workers who perform tasks regularly – 5% of the workforce is working at least one day a week on average. However, such experienced workers are small in number. In fact, the availability of workers decreases exponentially with experience.*

**Time Spent.** We use the amount of time spent by workers on tasks as a proxy for the total time spent working. While this may not be accurate because workers also have to search for tasks, this serves as a good estimate of their productive time. Figure 29b shows the total number of hours clocked by the active workers during their lifetime. We notice a skewed but long tailed distribution;  $\approx 10,000$  workers have been working on tasks for less than 25 hours. Nonetheless, there are also a handful of workers who have worked for more than 300 hours during the evaluation period.

Next, in Figure 29c, we plot the average number of hours spent by these active workers on a working day. From this plot, we note that more than 90% of the workers work for less than 1 hour during their working days. This suggests that crowdsourcing is still not at a scale where it can support many active workers on a full-time basis. However, more than a thousand workers still spend more than an hour a day working on tasks.

*Takeaway: The marketplace only supports a handful of workers on a full-time basis. A majority of the active workers appear to view the marketplace as a supplemental source of income, as is indicated by their daily hours of activity.*

**Trust.** The mean and median of the average trust of active workers are both above 91%, and 90% of all active workers have average trust higher than 0.84. Given that the trust scores of workers are all so high, and show such little variation, their distribution does yield any novel insights.

## 6. RELATED WORK

One of the first papers in the crowdsourcing literature focused on analyzing the Mechanical Turk marketplace for demographic factors [28, 29]; a more recent paper studied similar aspects by issuing surveys to Mechanical Turk workers [39]. Other recent papers

study the motivations of crowd workers by conducting broad surveys [11, 12]. Other papers have evaluated various aspects of marketplaces by interviewing or issuing tasks to workers, such as truthfulness [43] and consistency [42], the efficacy of conducting interface evaluations via crowdsourcing [36, 30], limitations in using Mechanical Turk for experimentation [41], and challenges faced by workers with disabilities [45]. Others attempt to understand worker motivations and behavior using (Turker Nation) forum data [35], and workers’ on- and off-network interactions [21]. Recent work has evaluated the impact of varying price as well as other economics-based concerns [20, 24, 25, 23, 18]. In this paper, since we do not have pricing data, we do not focus on evaluating this aspect. A recent book [34] described the results of interviewing marketplace companies (including Samasource [8], Crowdflower [2]) for their concerns and problems, but did not conduct a similar quantitative study based on marketplace data. Marketplace companies sometimes publish their own reports on demographics, e.g., [44, 7, 6, 26]. In [31], the authors discuss the challenges faced by crowdsourcing marketplaces, and describe their vision for the future. We analyze the state and trends of our crowdsourcing marketplace with a similar goal of improving the crowdsourcing paradigm as a whole—our focus, however, is on making more immediately actionable recommendations through data-driven analyses.

The paper that is closest to us in adopting a data-driven approach is the one by Difallah et al. [14], focused on studying the marketplace dynamics of Mechanical Turk, such as marketplace demand and supply, evolution of task payments over time, as well as other topics. We describe below, in detail, how our analysis and their analysis differs:

**What we do that Difallah et al. doesn’t.** There are several types of analyses that Difallah et al. does not consider at all, including:

- Effective tasks (Section 4): we quantify the influence of task design parameters on the effectiveness of a task, and make recommendations towards task design based on our observations.
- Worker sources (Section 5.1): we examine the various sources contributing workers to the marketplace and compare them on contribution and quality. This is not relevant for Difallah et al. since there is only one “source” (Mechanical Turk).
- Worker characterization (Sections 5.2, 5.3): we analyze various aspects of worker demographics and behavior, such as geographic location, lifetime, and engagement.

**Where we overlap—but how our emphases or granularity is different.** We overlap with Difallah et al. on a few topics, but approach these topics with either a different emphasis or a finer-grained perspective.

- Demand and supply (Section 3.1, 3.2): both papers study the demand and supply interactions in the marketplace, but from slightly different perspectives. Difallah’s focus is on making predictions for when batches will be completed. Since they don’t have worker data, they do not comment on the supply of

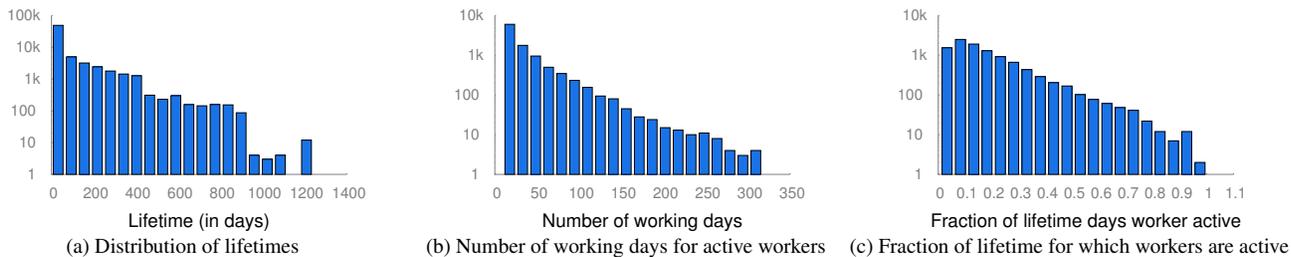


Figure 30: Worker lifetimes

workers. Our focus is instead on studying the arrival rate of tasks and availability of workers to complete them, which is more directly relevant for marketplace administrators.

- Popular goals (Section 3.4): Difallah et al. performs a cursory analysis of goals inferred using keywords within task descriptions; we perform a much more fine-grained analysis using expert-provided labels on all three of {goals, operators, data types}, enabling a thorough characterization of the spectrum of crowd work.

**What Difallah et al. does that we don’t.** This is because we do not have the data to study those aspects.

- Reward analysis: they plot payments assigned to tasks from different topics and workers from different countries over time. We do not have payment information as part of our dataset.
- Requester statistics: they plot the number and countries of preferred workers for different requesters over time. We do not have requester IDs as part of our dataset.

Thus, we have some overlap with the Difallah et al. paper in terms of marketplace analysis, but our emphasis and granularity is different; at the same time, the findings in our task design and worker analysis sections are entirely new. Also unlike that paper that focuses on Mechanical Turk, our crowdsourcing marketplace recruits workers from a collection of labor sources, making it a crowdsourcing “intermediary” or “aggregator”, and raising the possibility of a number of interesting additional types of analyses.

Our work in gaining a better understanding of crowd work has broad ramifications for the database community who has been developing crowd-powered data processing algorithms [22, 37, 13, 40, 9], and systems, e.g., [33, 19, 32, 10, 38], with dozens of papers published in database conferences each year. (A recent book surveys this literature [34].) As examples, understanding the relative importance of various types of processing needs, can prioritize the attention of our community to unexplored or underoptimized areas; understanding how tasks are picked up and worked on can help the community develop better models of task latency; understanding the worker perspective and engagement can aid in the design of better models for worker accuracy and worker behavior in general; and understanding the impact of task design can help the community adopt “best practices” to further optimize cost, accuracy, and latency.

## 7. CONCLUSIONS AND FUTURE WORK

In this paper, we quantitatively address a number of important open-ended questions aimed towards understanding and improving the entire paradigm of crowdsourcing from three key perspectives — marketplace dynamics (important to marketplace administrators), task design (important to requesters), and worker characterization (important to labor sources, marketplace administrators, and requesters). A number of these questions have been speculated upon by practitioners and academics, but there has been little concrete evidence to support these speculations and the so-called “rules of thumb”. Towards this end, we collect and organize data from a large crowdsourcing marketplace, and additionally enrich this

dataset with extensive manual labels. We answer several of what we believe are the most important open questions about crowdsourcing interactions, through quantitative, data-driven experiments on this dataset. Based on our experiments, we come up with a number of valuable insights and takeaways, that we hope will inform and guide the evolution of crowdsourcing over the coming years.

There are a number of directions that one could explore, following this work. A natural direction for the database community is to explore the unexplored combinations of popular task types that have not yet been adequately optimized. More broadly, it would be useful to pursue a deeper understanding of worker behavior by looking at phenomena such as worker anchoring, worker learning, and interactions between various jobs. While we have restricted our analyses to a specific set of features and metrics—a full analysis of the interplay between various different task parameters and notions of job success would be a natural next step. Lastly, with full-fledged A/B testing, we may be able to solidify our correlation and predictive claims with further causation-based evidence.

## 8. REFERENCES

- [1] ClixSense (Retrieved 16 March 2015). <http://www.clixsense.com/>.
- [2] CrowdFlower (Retrieved 22 July 2013). <http://crowdflower.com>.
- [3] Mechanical Turk (Retrieved 22 July 2013). <http://www.mturk.com>.
- [4] NeoDev. <http://www.neodev.se/>.
- [5] Prodege. <http://www.prodege.com/>.
- [6] Samasource Impact Report (Retrieved 16 March 2015). <http://www.samasource.org/impact/>.
- [7] Samasource Jobs (Retrieved 16 March 2015). <http://www.samasource.org/people/#jobs>.
- [8] Samasource (Retrieved 22 July 2013). <http://samasource.com>.
- [9] Y. Amsterdamer, Y. Grossman, T. Milo, and P. Senellart. Crowd mining. In *SIGMOD Conference*, pages 241–252, 2013.
- [10] A. Bozzon, M. Brambilla, and S. Ceri. Answering search queries with crowdsearcher. In *WWW*, pages 1009–1018, 2012.
- [11] A. M. Brawley and C. L. Pury. Work experiences on mturk: Job satisfaction, turnover, and information sharing. *Computers in Human Behavior*, 54:531–546, 2016.
- [12] R. Brewer et al. Why would anybody do this?: Understanding older adults’ motivations and challenges in crowd work. In *CHI*, pages 2246–2257. ACM, 2016.
- [13] S. B. Davidson, S. Khanna, T. Milo, and S. Roy. Using the crowd for top-k and group-by queries. In *ICDT*, pages 225–236, 2013.
- [14] D. E. Difallah et al. The dynamics of micro-task crowdsourcing: The case of amazon mturk. In *WWW*, pages 238–247, 2015.
- [15] <http://venturebeat.com/2014/10/08/leadgenius-grabs-6m-because-hiring-an-actual-sales-team-is-so-last-season/>. Leadgenius grabs 6m, venture beat.
- [16] <http://www.behind-the-enemy-lines.com/2016/02/a-cohort-analysis-of-mechanical-turk.html>. Cohort analysis of mturk requesters, blog post, 2015.
- [17] <http://www.forbes.com/sites/elainepofeldt/2015/05/05/elance-odesk-becomes-upwork-today-odesk-brand-gets-phased-out>. Odesk brand gets phased out, forbes.com, 2015.
- [18] A. Finnerty et al. Keep it simple: Reward and task design in crowdsourcing. In *Proc. of the Conf. of the Italian Chapter of SIGCHI*, page 14. ACM, 2013.

- [19] M. J. Franklin, D. Kossmann, T. Kraska, S. Ramesh, and R. Xin. Crowddb: answering queries with crowdsourcing. In *SIGMOD Conference*, pages 61–72, 2011.
- [20] Y. Gao and A. G. Parameswaran. Finish them!: Pricing algorithms for human computation. *PVLDB*, 7(14):1965–1976, 2014.
- [21] M. L. Gray, S. Suri, S. S. Ali, and D. Kulkarni. The crowd is a collaborative network. In *CSCW*, pages 134–147. ACM, 2016.
- [22] S. Guo et al. So who won?: dynamic max discovery with the crowd. In *SIGMOD Conference*, pages 385–396, 2012.
- [23] D. Haas, J. Wang, E. Wu, and M. J. Franklin. Clamshell: Speeding up crowds for low-latency data labeling. *PVLDB*, 9(4):372–383, 2015.
- [24] B. Hartmann and P. G. Ipeirotis. What’s the right price? pricing tasks for finishing on time. 2011.
- [25] J. J. Horton and L. B. Chilton. The labor economics of paid crowdsourcing. *CoRR*, abs/1001.0627, 2010.
- [26] <https://www.odesk.com/oconomy/activity/>. Odesk oconomy, 2012.
- [27] <https://www.upwork.com/blog/2013/12/mergerfaq/>. Odesk-elance merger faq, upwork blog, 2014.
- [28] P. G. Ipeirotis. Analyzing the amazon mechanical turk marketplace. *XRDS*, 17:16–21, December 2010.
- [29] P. G. Ipeirotis. Demographics of mechanical turk. 2010.
- [30] A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with mechanical turk. In *CHI*, pages 453–456, 2008.
- [31] A. Kittur et al. The future of crowd work. In *CSCW*, pages 1301–1318. ACM, 2013.
- [32] X. Liu et al. Cdas: A crowdsourcing data analytics system. *PVLDB*, 5(10):1040–1051, 2012.
- [33] A. Marcus et al. Crowdsourced databases: Query processing with people. In *CIDR*, pages 211–214, 2011.
- [34] A. Marcus and A. Parameswaran. Crowdsourced data management: Industry and academic perspectives. *Found. Trends databases*, 6(1-2).
- [35] D. Martin et al. Being a turker. In *CSCW*, pages 224–235. ACM, 2014.
- [36] W. Mason and S. Suri. Conducting behavioral research on amazon’s mechanical turk. *Behavior research methods*, 44(1):1–23, 2012.
- [37] A. G. Parameswaran et al. Crowdscreen: algorithms for filtering data with humans. In *SIGMOD Conference*, pages 361–372, 2012. Online: <http://doi.acm.org/10.1145/2213836.2213878>.
- [38] H. Park et al. An overview of the deco system: Data model and query language; query processing and optimization. *ACM SIGMOD Record*, 41, 2012.
- [39] J. Ross et al. Who are the crowdworkers?: shifting demographics in mechanical turk. In *CHI Extended Abstracts*, New York, NY, USA, 2010. ACM.
- [40] A. D. Sarma, A. Parameswaran, H. Garcia-Molina, and A. Halevy. Crowd-powered find algorithms. In *ICDE*, 2014.
- [41] N. Stewart et al. The average laboratory samples a population of 7,300 amazon mechanical turk workers. *Judgment and Decision Making*, 10(5):479, 2015.
- [42] P. Sun and K. T. Stolee. Exploring crowd consistency in a mechanical turk survey. In *Proceedings of the 3rd International Workshop on CrowdSourcing in Software Engineering*, pages 8–14. ACM, 2016.
- [43] S. Suri, D. G. Goldstein, and W. A. Mason. Honesty in an online labor market. In *Human Computation*, 2011.
- [44] N. Zukoff. Demographics of the Largest On-demand Workforce (Retrieved 16 March 2015). <http://www.crowdfunder.com/blog/2014/01/demographics-of-the-largest-on-demand-workforce>, 2014.
- [45] K. Zyskowski et al. Accessible crowdwork?: Understanding the value in and challenge of microtask employment for people with disabilities. In *CSCW*, pages 1682–1693. ACM, 2015.

neodev	clixsense	prodege	elite	instage	tremorgames	internal	bitcoinget
amt	superrewards	eup_slw	gifthunterclub	taskhunter	prizerebel	hiving	fusioncash
points2shop	clicksfx	getpaid	cotter	coinworker	vivatic	piyanstantrewards	inboxpounds
imerit_india	personaly	stuffpoint	errtopc	taskspay	zoombucks	crowdgur	gifthulk
tasks4dollars	dollarsignup	indivillagetest	cbf	mycashtasks	sendearnings	treasuretrooper	pokerowned
diamondtask	pforads	quickrewards	uniquerewards	extralunchmoney	cashcrate	wannads	gptbanks
listia	gradible	dailyrewardsca	clickfair	superpayme	memolink	rewardok	snowcitrustechbpo
pedtoclick	rewardingways	callmemoney	pocketmoneygpt	goldtasks	dollarrewardz	surveygad	sharecashgpt
irazoo	zapbux	ptcsolution	ptc123	content_runner	jetbux	qpr	cointasker
point_dollars	meprizescf	keeprewarding	gptking	dollarsgpt	prizeplank	yute_jamaica	onestopgpt
gptway	trial_pay	task_ph	golddiggergpt	prizezombie	daproimafrica	aceinnovations	getpaidto
globalactioncash	piyoogle	supersonicads	poin_web	rewardsspot	giftgpt	giftcardgpt	northclicks
fastcashgpt	dealbarbiepays	dailysurveypanel	points4rewards	gptpal	rewards1	new_rules	surewardsgpt
zorbor	steamgameswap	buxense	surveywage	offernation	probux	freeride	ojoo
luckytaskz	medievaleurope	proudclick	steampowers	paiddailysurveys	wrkshop	simplegpt	realworld
surveytokens	bemybux	onestop	plUSDollars	gptbucks	fepcrowdflower	embee	makethatdollar
ayuwage	luckykoin	pointst	sedgroup	easycashclicks	candy_ph	piggybankgpt	peoplesgpt
matomy	eamthemost	fsprizes					

Table 4: Sources for crowd workers