



In Search of Ambiguity: A Three-Stage Workflow Design to Clarify Annotation Guidelines for Crowd Workers

Vivek Krishna Pradhan¹, Mike Schaeckermann² and Matthew Lease^{1,2*}

¹ University of Texas at Austin, Austin, TX, United States, ² Amazon, Seattle, WA, United States

OPEN ACCESS

Edited by:

Fabio Casati,
Servicenow, United States

Reviewed by:

Hend Al-Khalifa,
King Saud University, Saudi Arabia
Marcos Baez,
Independent Researcher, Osnabrück,
Germany

*Correspondence:

Matthew Lease
ml@utexas.edu

Specialty section:

This article was submitted to
Machine Learning and Artificial
Intelligence,
a section of the journal
Frontiers in Artificial Intelligence

Received: 03 December 2021

Accepted: 18 March 2022

Published: 18 May 2022

Citation:

Pradhan VK, Schaeckermann M and
Lease M (2022) In Search of
Ambiguity: A Three-Stage Workflow
Design to Clarify Annotation
Guidelines for Crowd Workers.
Front. Artif. Intell. 5:828187.
doi: 10.3389/frai.2022.828187

We propose a novel three-stage FIND-RESOLVE-LABEL workflow for crowdsourced annotation to reduce ambiguity in task instructions and, thus, improve annotation quality. Stage 1 (FIND) asks the crowd to find examples whose correct label seems ambiguous given task instructions. Workers are also asked to provide a short tag that describes the ambiguous concept embodied by the specific instance found. We compare collaborative vs. non-collaborative designs for this stage. In Stage 2 (RESOLVE), the requester selects one or more of these ambiguous examples to label (resolving ambiguity). The new label(s) are automatically injected back into task instructions in order to improve clarity. Finally, in Stage 3 (LABEL), workers perform the actual annotation using the revised guidelines with clarifying examples. We compare three designs using these examples: examples only, tags only, or both. We report image labeling experiments over six task designs using Amazon's Mechanical Turk. Results show improved annotation accuracy and further insights regarding effective design for crowdsourced annotation tasks.

Keywords: crowdsourcing, annotation, labeling, guidelines, ambiguity, clarification, machine learning, artificial intelligence

1. INTRODUCTION

While crowdsourcing now enables labeled data to be obtained more quickly, cheaply, and easily than ever before (Snow et al., 2008; Sorokin and Forsyth, 2008; Alonso, 2015), ensuring data quality remains something of an art, challenge, and perpetual risk. Consider a typical workflow for annotating data on Amazon Mechanical Turk (MTurk): a requester designs an annotation task, asks multiple workers to complete it, and then post-processes labels to induce final consensus labels. Because the annotation work itself is largely opaque, with only submitted labels being observable, the requester typically has little insight into what if any problems workers encounter during annotation. While statistical aggregation (Hung et al., 2013; Sheshadri and Lease, 2013; Zheng et al., 2017) and multi-pass iterative refinement (Little et al., 2010a; Goto et al., 2016) methods can be employed to further improve initial labels, there are limits to what can be achieved by *post-hoc* refinement following label collection. If initial labels are poor because many workers were confused by incomplete, unclear, or ambiguous task instructions, there is a significant risk of “garbage in equals garbage out” (Vidgen and Derczynski, 2020).

In contrast, consider a more traditional annotation workflow involving trusted annotators, such as practiced by the Linguistic Data Consortium (LDC) (Griffitt and Strassel, 2016). Once preliminary annotation guidelines are developed, an iterative process ensues in which: (1) a subset of data is labeled based on current guidelines; (2) annotators review corner cases and disagreements, review relevant guidelines, and reach a consensus on appropriate resolutions; (3) annotation guidelines are updated; and (4) the process repeats. In comparison to the simple crowdsourcing workflow above, this traditional workflow iteratively debugs and refines task guidelines for clarity and completeness in order to deliver higher quality annotations. However, it comes at the cost of more overhead, with a heavier process involving open-ended interactions with trusted annotators. Could we somehow combine these for the best of both worlds?

In this study, we propose a novel three-stage FIND-RESOLVE-LABEL design pattern for crowdsourced annotation which strikes a middle-ground between the efficient crowdsourcing workflow on one hand and the high quality LDC-style workflow on the other. Similar to prior study (Gaikwad et al., 2017; Bragg and Weld, 2018; Manam and Quinn, 2018), we seek to design a light-weight process for engaging the workers themselves to help debug and clarify the annotation guidelines. However, existing approaches typically intervene in a reactive manner *after* the annotation process has started, or tend to be constrained to a specific dataset or refinement of textual instruction only. By contrast, our approach is proactive and open-ended. It leverages crowd workers' unconstrained creativity and intelligence to identify ambiguous examples through an Internet search on the Internet and enriches task instructions with these concrete examples proactively upfront before the annotation process commences. Overall, we envision a partnership between the requester and workers in which each party has complementary strengths and responsibilities in the annotation process, and we seek to maximize the relative strengths of each party to ensure data quality while preserving efficiency.

Figure 1 depicts our overall workflow. In Stage 1 (FIND), workers are shown initial guidelines for an annotation task and asked to search for data instances that appear ambiguous given the guidelines. For each instance workers find, they are also asked to provide a short tag that describes the concept embodied by the specific instance which is ambiguous given the guidelines. Next, in Stage 2 (RESOLVE), the requester selects one or more of the ambiguous instances to label as exemplars. Those instances and their tags are then automatically injected back into the annotation guidelines in order to improve clarity. Finally, in Stage 3 (LABEL), workers perform the actual annotation using the revised guidelines with clarifying examples. The requester can run the LABEL stage on a sample of data, assess label quality, and then decide how to proceed. If quality is sufficient, the remaining data can simply be labeled according to the guidelines. Otherwise, Stages 1 and 2 can be iterated in order to further refine the guidelines.

To evaluate our three-stage task design, we construct six different image labeling tasks with different levels of difficulty and

intuitiveness. We construct a test dataset that contains different ambiguous and unambiguous concepts. Starting from simple and possibly ambiguous task instructions, we then improve instructions *via* our three-stage workflow. Given expert (gold) labels for our dataset for each of the six tasks, we can evaluate how well-revised instructions compare to original instructions by measuring the accuracy of the labels obtained from the crowd.

1.1. Contributions

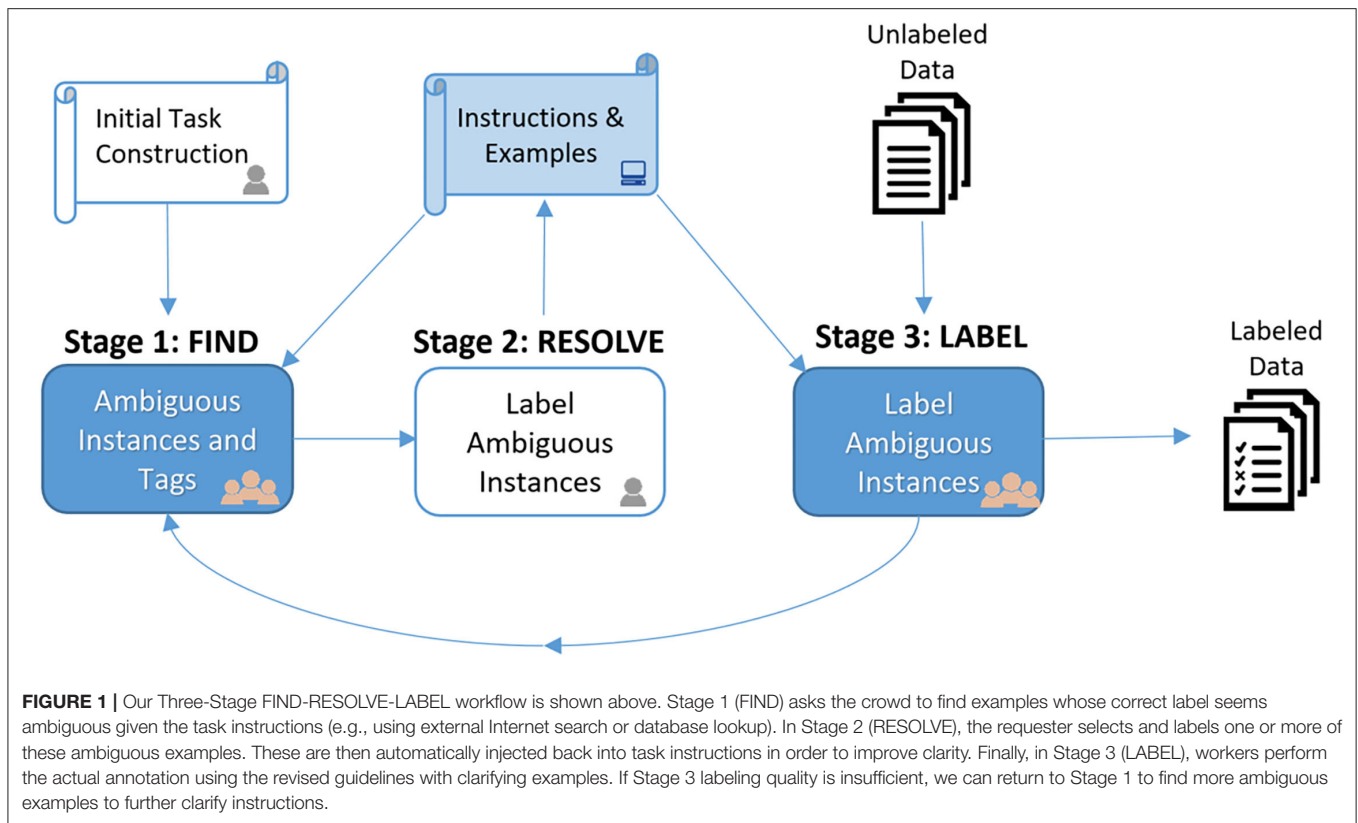
We provide initial evidence suggesting that the crowd can find and provide useful ambiguous examples which can be used to further clarify task instructions and that these examples may have the potential to be utilized to improve annotation accuracy. Our experiments further seem to suggest that workers can perform better when shown key ambiguous examples as opposed to randomly chosen examples. Finally, we provide an analysis of workers' performance for different intents of the same classification task and different concepts of ambiguity within each intent.

Our article is organized as follows. Section 2 presents Motivation and Background. Next, section 3 details our 3-Stage FIND-RESOLVE-LABEL workflow. Next, section 4 explains our experimental setup. Section 5 then presents the results. Finally, section 6 discusses conclusions and future directions.

2. MOTIVATION AND BACKGROUND

Consider the task of labeling images for object detection. For example, on MTurk one might post a task such as, "Is there a dog in this image?" Such a task appears to be quite simple, but is it? For example, is a wolf a dog? What about more exotic and unusual wild breeds of dogs? Does the dog need to be a real animal or merely a depiction of one? What about a museum model of an ancient but extinct dog breed, or a realistic wax sculpture? What if the dog is only partially visible in the image? Ultimately, what is it that the requester really wants? For example, a requester interested in anything and everything dog-related might have very liberal inclusion criteria. On the other hand, a requester training a self-driving car might only care about animals to be avoided, while someone training a product search engine for an e-commerce site might want to include dog-style children's toys (Kulesza et al., 2014).

As this seemingly simple example illustrates, annotation tasks that seem straightforward to a requester may in practice embody a variety of subtle nuances and ambiguities to be resolved. Such ambiguities can arise for many reasons. The requester may have been overly terse or rushed in posting a task. They may believe the task is obvious and that no further explanation should be needed. They likely also have their own implicit biases (of which they may be unaware) that provide a different internal conception of the task than others might have. For example, the requester might be ignorant of the domain (e.g., is a wolf a type of dog?) or have not fully defined what they are looking for. For example, in information retrieval, users' own conception and understanding of what they are looking for often evolve during the process of search and browsing (Cole, 2011). We describe our own experiences with this in section 5.1.1. Annotators, on the other



hand, also bring with them their own variety of implicit biases which the requester may not detect or understand (Ipeirotis et al., 2010; Sen et al., 2015; Dumitrache et al., 2018; Geva et al., 2019; Al Kuwatly et al., 2020; Fazelpour and De-Arteaga, 2022).

2.1. Helping Requesters Succeed

2.1.1. Best Practices

A variety of tutorials, surveys, introductions, and research articles offer how-to advice for successful microtask crowdsourcing with platforms such as MTurk (Jones, 2013; Marshall and Shipman, 2013; Egelman et al., 2014; Kovashka et al., 2016). For example, it is often recommended that requesters invest time browsing and labeling some data themselves before launching a task in order to better define and debug it (Alonso, 2015). Studies have compared alternative task designs to suggest best practices (Grady and Lease, 2010; Kazai et al., 2011; Papoutsaki et al., 2015; Wu and Quinn, 2017).

2.1.2. Templates and Assisted Design

Rather than start task design from scratch, MTurk offers templates and has suggested that requesters share successful templates for others' use (Chen et al., 2011). Similarly, classic research on software *design patterns* (Gamma et al., 1995) has inspired ideas for similar crowdsourcing design patterns which could be reused across different data collection tasks. For example, FIND-FIX-VERIFY (Bernstein et al., 2010) is a well-known example that partially inspired our study. Other researchers have suggested improved tool support for workflow

design (Kittur et al., 2012) or engaging the crowd itself in task design or decomposition (Kittur et al., 2011; Kulkarni et al., 2012a).

2.1.3. Automating Task Design

Other researchers have gone further still to propose new middleware and programmable APIs to let requesters define tasks more abstractly and leave some design and management tasks to the middleware (Little et al., 2010b; Ahmad et al., 2011; Franklin et al., 2011; Barowy et al., 2016; Chen et al., 2016).

2.2. Understanding Disagreement

2.2.1. Random Noise vs. Bias

Since annotators are human, even trusted annotators will naturally make mistakes from time to time. Fortunately, random error is exactly the kind of disagreement that aggregation (Hung et al., 2013; Sheshadri and Lease, 2013) can easily resolve; assuming such mistakes are relatively infrequent and independent, workers will rarely err at the same instance, and therefore, techniques as simple as majority voting can address random noise. On the other hand, if workers have individual biases, they will make consistent errors; e.g., a teenager vs. a protective parent might have liberal vs. conservative biases in rating movies (Ipeirotis et al., 2010). In this case, it is useful to detect such consistent biases and re-calibrate worker responses to undo such biases. Aggregation can also work provided that workers do not share the same biases. However, when workers do share systematic biases, the independence assumption underlying

aggregation is violated, and so aggregation can amplify bias rather than resolve it. Consequently, it is important that task design annotation guidelines should be vetted to ensure they identify cases in which annotator biases conflict with desired labels and particularly establish clear expectations for how such cases should be handled (Draws et al., 2021; Nouri et al., 2021b).

2.2.2. Objective vs. Subjective Tasks

In fully-objective tasks, we assume each question has a single correct answer, and any disagreement with the gold standard reflects error. Label aggregation methods largely operate in this space. On the other extreme, purely-subjective (i.e., opinion) tasks permit a wide range of valid responses with little expectation of agreement between individuals (e.g., asking about one's favorite color or food). Between these simple extremes, however, lies a wide, interesting, and important space of partially-subjective tasks in which answers are only partially-constrained (Tian and Zhu, 2012; Sen et al., 2015; Nguyen et al., 2016). For example, consider rating item quality: while agreement tends to be high for items having extremely good or bad properties, instances with more middling properties naturally elicit a wider variance in opinion. In general, because subjectivity permits a valid diversity of responses, it can be difficult to detect if an annotator does not undertake a task in good faith, complicating quality assurance.

2.2.3. Difficulty vs. Ambiguity

Some annotation tasks are more complex than others, just as some instances within each task are more difficult to label than other instances. A common concern with crowdsourcing is whether inexperienced workers have sufficient expertise to successfully undertake a given annotation task. Intuitively, more guidance and scaffolding are likely necessary with more skilled tasks and fewer expert workers (Huang et al., 2021). Alternatively, if we use sufficiently expert annotators, we assume difficult cases can be handled (Retelny et al., 2014; Vakharia and Lease, 2015). With ambiguity, on the other hand, it would be unclear even to an expert what to do. Ambiguity is an interaction between data instances and annotation guidelines; effectively, an ambiguous instance is a corner-case with respect to guidelines. Aggregation can helpfully identify the majority interpretation but that interpretation may or may not be what is actually desired. Both difficult and ambiguous cases can lead to label confusion. Krivosheev et al. (2020) developed mechanisms to efficiently detect label confusion in classification tasks and demonstrated that alerting workers of the risk of confusion can improve annotation performance.

2.2.4. Static vs. Dynamic Disagreement

As annotators undertake a task, their understanding of work evolves as they develop familiarity with both the data and the guidelines. In fact, prior study has shown that annotators interpret and implement task guidelines in different ways as annotation progresses (Scholer et al., 2013; Kalra et al., 2017). Consequently, different sorts of disagreement can occur at different stages of annotation. Temporally-aware aggregation can partially ameliorate this (Jung and Lease, 2015), as can

implementing data collection processes to train, “burn-in,” or calibrate annotators, controlling, and/or accelerating their transition from an initial learning state into a steady state (Scholer et al., 2013). For example, we emphasize identifying key boundary cases and expected labels for them.

2.3. Mitigating Imperfect Instructions

Unclear, confusing, and ambiguous task instructions are commonplace phenomena on crowdsourcing platforms (Gadiraju et al., 2017; Wu and Quinn, 2017). In early study, Alonso et al. (2008) recommended collecting optional, free-form, task-level feedback from workers. While Alonso et al. (2008) found that some workers did provide example-specific feedback, the free-form nature of their feedback request elicited a variety of response types, which is difficult to check or to invalidate spurious responses. Alonso et al. (2008) also found that requiring such feedback led many workers to submit unhelpful text that was difficult to automatically cull. Such feedback was, therefore, recommended to be kept entirely optional.

While crowd work is traditionally completed independently to prevent collusion and enable statistical aggregation of uncorrelated work (Hung et al., 2013; Sheshadri and Lease, 2013; Zheng et al., 2017), a variety of work has explored collaboration mechanisms by which workers might usefully help each other complete a task more effectively (Dow et al., 2012; Kulkarni et al., 2012b; Drapeau et al., 2016; Chang et al., 2017; Manam and Quinn, 2018; Schaekermann et al., 2018; Chen et al., 2019; Manam et al., 2019).

Drapeau et al. (2016) proposed an asynchronous two-stage *Justify-Reconsider* method. In the Justify task, workers provide a rationale along with their answer referring to the task guidelines taught during training. For the Reconsider task, workers are confronted with an argument for the opposing answer submitted by another worker and then asked to reconsider (i.e., confirm or change) their original answer. The authors report that their Justify-Reconsider method generally yields higher accuracy but that requesting justifications requires additional cost. Consequently, they find that simply collecting more crowd annotations yields higher accuracy in a fixed-budget setting.

Chang et al. (2017) proposed a three-step approach in which crowd workers label the data, provide justifications for cases in which they disagree with others, and then review others' explanations. They evaluate their method on an image labeling task and report that requesting only justifications (without any further processing) does not increase the crowd accuracy. Their open-ended text responses can be subjective and difficult to check.

Kulkarni et al. (2012b) provide workers with a chat feature that supports workers in dealing with inadequate task explanations, suggesting additional examples to be given to requesters, teaching other workers how to use the UI, and verifying their hypotheses of the underlying task intent. Schaekermann et al. (2018) investigate the impact of discussion among crowd workers on the label quality using a chat platform allowing synchronous group discussion. While the chat platform allows workers to better express their justification than text excerpts, the discussion increases task completion times. In addition, chatting does not

impose any restriction on the topic, limiting discussion from unenthusiastic workers and efficacy. Chen et al. (2019) also proposed a workflow allowing simultaneous discussion among crowd workers, and designed task instructions and a training phase to achieve effective discussions. While their method yields high labeling accuracy, the increased cost due to the discussion limits its task scope. Manam and Quinn (2018) evaluated both asynchronous and synchronous Q&A between workers and requesters to allow workers to ask questions to resolve any uncertainty about overall task instructions or specific examples. Bragg and Weld (2018) proposed an iterative workflow in which data instances with the low inter-rater agreement are put aside and either used as difficult training examples (if considered resolvable with respect to the current annotation guidelines) or used to refine the current annotation guidelines (if considered ambiguous).

Other study has explored approaches to address ambiguities even *before* the annotation process commences. For example, Manam et al. (2019) proposed a multi-step workflow enlisting the help of crowd workers to identify and resolve ambiguities in textual instructions. Gadiraju et al. (2017) and Nouri et al. (2021a) both developed predictive models to automatically score textual instructions for their overall level of clarity and Nouri et al. (2021b) proposed an interactive prototype to surface the predicted clarity scores to requesters in real-time as they draft and iterate on the instructions. Our approach also aims to resolve ambiguities upfront but focuses on identifying concrete visual examples of ambiguity and automatically enriching the underlying set of textual instructions with those examples.

Ambiguity arises from the interaction between annotation guidelines and particular data instances. Searching for ambiguous data instances within large-scale datasets or even the Internet can amount to finding a needle in a haystack. There exists an analogous problem of identifying “unknown unknowns” or “blind spots” of machine learning models. Prior study has proposed crowdsourced or hybrid human-machine approaches for spotting and mitigating model blind spots (Attenberg et al., 2011; Vandenhof, 2019; Liu et al., 2020). Our study draws inspiration from these workflows. We leverage the scale, intelligence, and common sense of the crowd to identify potential ambiguities within annotation guidelines and may, thus, aid in the process of mitigating blind spots in downstream model development.

2.4. Crowdsourcing Beyond Data Labeling

While data labeling represents the most common use of crowdsourcing in regard to training and evaluating machine learning models, human intelligence can be tapped in a much wider and more creative variety of ways. For example, the crowd might verify output from machine learning models, identify, and categorize blind spots (Attenberg et al., 2011; Vandenhof, 2019) and other failure modes (Cabrera et al., 2021), and suggest useful features for a machine learning classifier (Cheng and Bernstein, 2015).

One of the oldest crowdsourcing design patterns is utilizing the scale of the crowd for efficient, distributed exploration or filtering of large search spaces. Classic examples include the

search for extraterrestrial intelligence¹, for Jim Gray’s sailboat (Vogels, 2007) or other missing people (Wang et al., 2010), for DARPA’s red balloons (Pickard et al., 2011), for astronomical events of interest (Lintott et al., 2008), and for endangered wildlife (Rosser and Wiggins, 2019) or bird species (Kelling et al., 2013). Across such examples, what is being sought must be broadly recognizable so that the crowd can accomplish the search task without the need for subject matter expertise (Kinney et al., 2008). In the 3-stage FIND-FIX-VERIFY crowdsourcing workflow (Bernstein et al., 2010), the initial FIND stage directs the crowd to identify “patches” in an initial text draft where more work is needed.

Our asking the crowd to search for ambiguous examples given task guidelines further explores the potential of this same crowd design pattern for distributed search. Rather than waiting for ambiguous examples to be encountered by chance during the annotation process, we instead seek to rapidly identify corner-cases by explicitly searching for them. We offload to the crowd the task of searching for ambiguous cases, and who better to identify potentially ambiguous examples than the same workforce that will be asked to perform the actual annotation? At the same time, we reduce requester work, limiting their effort to labeling corner-cases rather than adjusting the textual guidelines.

3. WORKFLOW DESIGN

In this study, we propose a three-stage FIND-RESOLVE-LABEL workflow for clarifying ambiguous corner cases in task instructions, investigated in the specific context of a binary image labeling task. An illustration of the workflow is shown in **Figure 1**. In Stage 1 (FIND), workers are asked to proactively collect ambiguous examples and concept tags given task instructions (section 3.1). Next, in Stage 2 (RESOLVE), the requester selects and labels one or more of the ambiguous examples found by the crowd. These labeled examples are then automatically injected back into task instructions in order to improve clarity (section 3.2). Finally, in Stage 3 (LABEL), workers perform the actual annotation task using the revised guidelines with clarifying examples (section 3.3). Requesters run the final LABEL stage on a sample of data, assess label quality, and then decide how to proceed. If quality is sufficient the remaining data can be labeled according to the current revision of the guidelines. Otherwise, Stages 1 and 2 can be repeated in order to further refine the clarity of annotation guidelines.

3.1. Stage 1: Finding Ambiguous Examples

In Stage 1 (FIND), workers are asked to collect ambiguous examples given the task instructions. For each ambiguous example, workers are also asked to generate a concept tag. The concept tag serves multiple purposes. First, it acts as a rationale (McDonnell et al., 2016; Kutlu et al., 2020), requiring workers to justify their answers and thus nudging them toward high-quality selections. Rationales also provide a form of transparency to help requesters better understand worker intent. Second, the concept tag provides a conceptual explanation of the ambiguity which

¹<https://setiathome.berkeley.edu/>.

Can you help us find ambiguous examples for this task?

We want to launch a new HIT asking workers to perform the following task:

Is there a dog in this image?



Yes No

Each HIT will show a different example and ask workers to provide an answer.

We think there might be examples for which our task is ambiguous, and we want to ask for your help. ***Can you find ambiguous examples for this task?*** We will then include these as clarifying examples in our task instructions in order to make the task easier and faster for workers to complete correctly.

Below is one or more ambiguous example(s) we already know about. You should find a new type of ambiguity which we can then clarify in our instructions.



Toy Dog

You can search for ambiguous example images on [Google Images](#). Save the image you find to your computer and then upload it below.

upload image: No file chosen

FIGURE 2 | In the Stage 1 (FIND) task, workers are asked to search for examples they think would be ambiguous given task instructions. In this case, “Is there a dog in this image?” In collaboration conditions (section 3.1.1), workers will see additional ambiguous examples found by past workers.

can then be re-injected into annotation guidelines to help explain corner cases to future workers.

Figure 2 shows the main task interface for Stage 1 (FIND). The interface presents the annotation task (e.g., “*Is there a dog in this image?*”) and asks workers: “*Can you find ambiguous examples for this task?*” Pilot experiments revealed that workers had difficulty understanding the task based on this textual prompt alone. We, therefore, make the additional assumption that requesters

provide a single ambiguous example to clarify the FIND task for workers. For example, the FIND stage for a dog annotation task could show the image of a Toy Dog as an ambiguous seed example. Workers are then directed to use Google Image Search to find these ambiguous examples. Once an ambiguous image is uploaded, another page (not shown) asks workers to provide a short concept tag summarizing the type of ambiguity represented by the example (e.g., *Toy Dog*).

3.1.1. Exploring Collaboration

To investigate the potential value of worker collaboration in finding higher quality ambiguities, we explore a light-weight, iterative design in which workers do not directly interact with each other, but are shown examples found by past workers (in addition to the seed example provided by the requester). For example, worker 2 would see an example selected by worker 1, and worker 3 would see examples found by workers 1 and 2, etc. Our study compares three different collaboration conditions described in section 4.4.1 below.

3.2. Stage 2: Resolving Ambiguous Examples

After collecting ambiguous examples in Stage 1 (FIND), the requester then selects and labels one or more of these examples. The requester interface for Stage 2 (RESOLVE) is shown in **Figure 3**. Our interface design affords a low-effort interaction in which requesters toggle examples between three states *via* mouse click: (1) selected as a positive example, (2) selected as a negative example, (3) unselected. Examples are unselected by default. The selected (and labeled) examples are injected back into the task instructions for Stage 3 (LABEL).

3.3. Stage 3: Labeling With Clarifying Examples

Best practices suggest that along with task instructions, requesters should include a set of examples and their correct annotations (Wu and Quinn, 2017). We automatically append to task instructions the ambiguous examples selected by the requester in Stage 2 (RESOLVE), along with their clarifying labels (**Figure 4**). Positive examples are shown first (“*you should select concepts like these*”), followed by negative examples (“*and NOT select concepts like these*”). Note that this stage does not require additional effort (e.g., instruction drafting) from the side of the requester because it merely augments the pre-existing task instruction template with the resulting list of clarifying examples.

4. METHODS

Experiments were conducted in the context of binary image classification. In particular, we designed six annotation tasks representing different variations of labeling for the presence or absence of *dog*-related concepts. Similar to prior study by Kulesza et al. (2014), we found this seemingly simplistic domain effective for our study because non-expert workers bring prior intuition as to how the classification could be done, but the task is characterized by a variety of subtle nuances and inherent ambiguities. We employed a between-subjects design in which each participant was assigned to exactly one experimental condition to avoid potential learning effects. This design was enforced using “negative” qualifications (Amazon Mechanical Turk, 2017) preventing crowd workers from participating in more than a single task. For the purpose of experimentation, authors acted as requesters. This included the specification of task instructions and intents and performing Stage 2 (RESOLVE), i.e., selecting clarifying examples for use in Stage 3 (LABEL).

4.1. Participant Recruitment and Quality Control

We recruited participants on Amazon’s Mechanical Turk using workers from the US who had completed at least 1,000 tasks with a 95% acceptance rate. This filter served as a basic quality assurance mechanism to increase the likelihood of recruiting good-faith workers over “spammers.” No further quality control mechanism was employed in our study to emulate imperfect, yet commonplace crowdsourcing practices for settings where definitive gold standard examples are not readily available for quality assessment. For ecological validity, we opted to not collect demographic information about participants prior to the annotation tasks.

4.2. Dataset

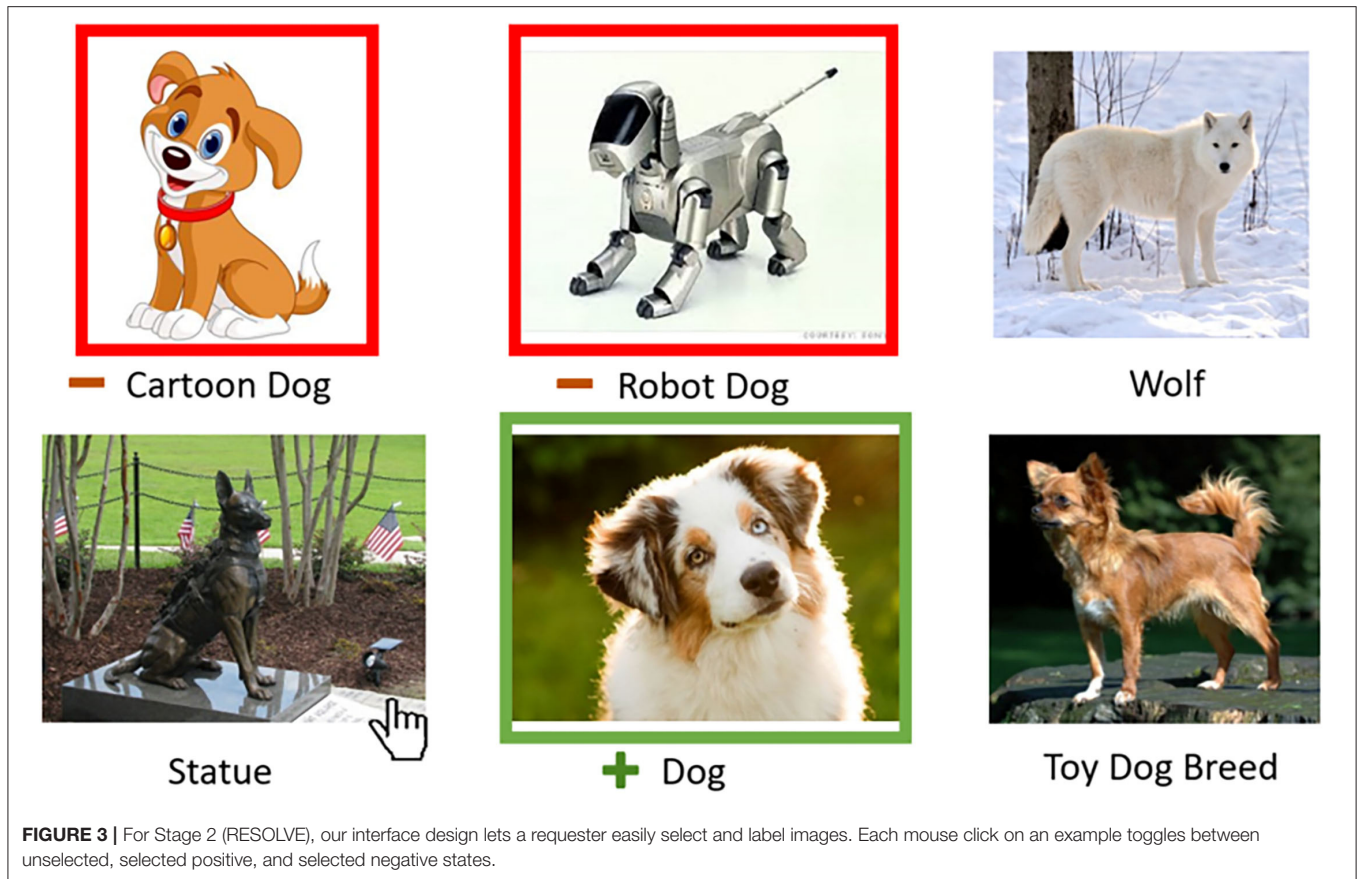
All experiments utilized the same set of 40 images. The image set was designed to encompass both easy, unambiguous cases and a range of difficult, ambiguous cases with respect to the question “*Is there a dog in this image?*” We first assembled a set of candidate images using a combination of (1) an online image search conducted by the authors to identify a set of clear positive and clear negative examples, (2) the Stage 1 (FIND) mechanism in which crowd workers on Amazon’s Mechanical Turk identified difficult, ambiguous cases. Similar to Kulesza et al. (2014), we identified a set of underlying, dog-related categories *via* multiple passes of structured labeling on the data. From this process, 11 categories of dog-related concepts emerged: (1) dogs, (2) small dog breeds, (3) similar animals (easy to confuse with dogs), (4) cartoons, (5) stuffed toys, (6) robots, (7) statues, (8) dog-related objects (e.g., dog-shaped cloud), (9) miscellaneous (e.g., hot dog, the word “dog”), (10) different animals (difficult to confuse with dogs), and (11) planes (the easiest category workers should never confuse with dogs). Each image was assigned to exactly one category.

4.3. Annotation Tasks

When users of a search engine type in the query “apple,” are they looking for information about the fruit, the company, or something else entirely? Despite the paucity of detail provided by a typical terse query, search result accuracy is assessed based on how well results match the user’s underlying intent. Similarly, requesters on crowdsourcing platforms expect workers to understand the annotation “intent” underlying the explicit instructions provided. Analogously, worker accuracy is typically evaluated with respect to how well-annotations match that requester’s intent even if instructions are incomplete, unclear, or ambiguous.

To represent this common scenario, we designed three different annotation tasks. For each task, the textual instructions exhibit a certain degree of ambiguity such that adding clarifying examples to instructions can help clarify requester intent to workers.

For each of the three tasks, we also selected two different intents, one more intuitive than the other in order to assess the effectiveness of our workflow design under intents of varying intuitiveness. In other words, we intentionally included one slightly more esoteric intent for each task hypothesizing that



these would require workers to adapt to classification rules in conflict with their initial assumptions about requester intent. For each intent below, we list the categories constituting the positive class. All other categories are part of the negative class for the given intent.

For each of our six binary annotation tasks below, we partitioned examples into positive vs. negative classes given the categories included in the intent. We then measured worker accuracy in correctly labeling images according to positive and negative categories for each task.

4.3.1. Task 1: Is There a Dog in This Image?

Intent a (more intuitive): dogs, small dog breeds

Intent b (less intuitive): dogs, small dog breeds, similar animals. *Scenario*: The requester intends to train a machine learning model for avoiding animals and believes the model may also benefit from detecting images of wolves and foxes.

4.3.2. Task 2: Is There a Fake Dog in This Image?

Intent a (more intuitive): similar animals. *Scenario*: The requester is looking for animals often confused with dogs.

Intent b (less intuitive): cartoons, stuffed toys, robots, statues, objects. *Scenario*: The requester is looking for inanimate objects representing dogs.

4.3.3. Task 3: Is There a Toy Dog in This Image?

Intent a (less intuitive): small dog breeds. *Scenario*: Small dogs, such as Chihuahua or Yorkshire Terrier, are collectively referred to as “toy dog” breeds². However, this terminology is not necessarily common knowledge making this intent less intuitive.

Intent b (more intuitive): stuffed toys, robots. *Scenario*: The requester is looking for children’s toys, e.g., to train a model for an e-commerce site.

4.4. Evaluation

4.4.1. Qualitative Evaluation of Ambiguous Examples From Stage 1 (FIND)

For Stage 1 (FIND), we evaluated crowd workers’ ability to find ambiguous images and concept tags for Task 1: “*Is there a dog in this image?*”. Through qualitative coding, we analyzed worker submissions based on three criteria: (1) correctness; (2) uniqueness; and (3) usefulness.

Correctness captures our assessment of whether the worker appeared to have understood the task correctly and submitted a plausible example of ambiguity for the given task. Any incorrect examples were excluded from consideration for uniqueness or usefulness.

Uniqueness captures our assessment of how many distinct types of ambiguity workers found across correct examples. For

²https://en.wikipedia.org/wiki/Toy_dog.

You should **select** concepts like these:



Dog

and **NOT select** concepts like these:



Cartoon Dog



Wolf



Statue



Robot Dog

FIGURE 4 | For Stage 3 (LABEL), we combine the ambiguous instances and/or tags collected in Stage 1 (FIND) with the requester labels from Stage 2 (RESOLVE) and automatically inject the labeled examples back into task instructions.

example, we deemed “Stuffed Dog” and “Toy Dog” sufficiently close as to represent the same concept.

Usefulness captures our assessment of which of the unique ambiguous concepts found were likely to be useful in the annotation. For example, while an image of a hot dog is valid and unique, it is unlikely that many annotators would find it ambiguous in practice.

Our study compares two different collaboration conditions for Stage 1 (FIND). In both conditions, workers were shown one or more ambiguous examples with associated concept tags and were asked to add another, different example of ambiguity, along with a concept tag for that new example:

1. **No collaboration.** Each worker sees the task interface seeded with a single ambiguous example and its associated concept tag provided by the requester. Workers find additional ambiguous examples *independently* from other workers.
2. **Collaboration.** Workers see all ambiguous examples and their concept tags previously found by other workers. There is no filtering mechanism involved, so workers may be presented with incorrect and/or duplicated examples. This workflow configuration amounts to a form of unidirectional, asynchronous communication among workers.

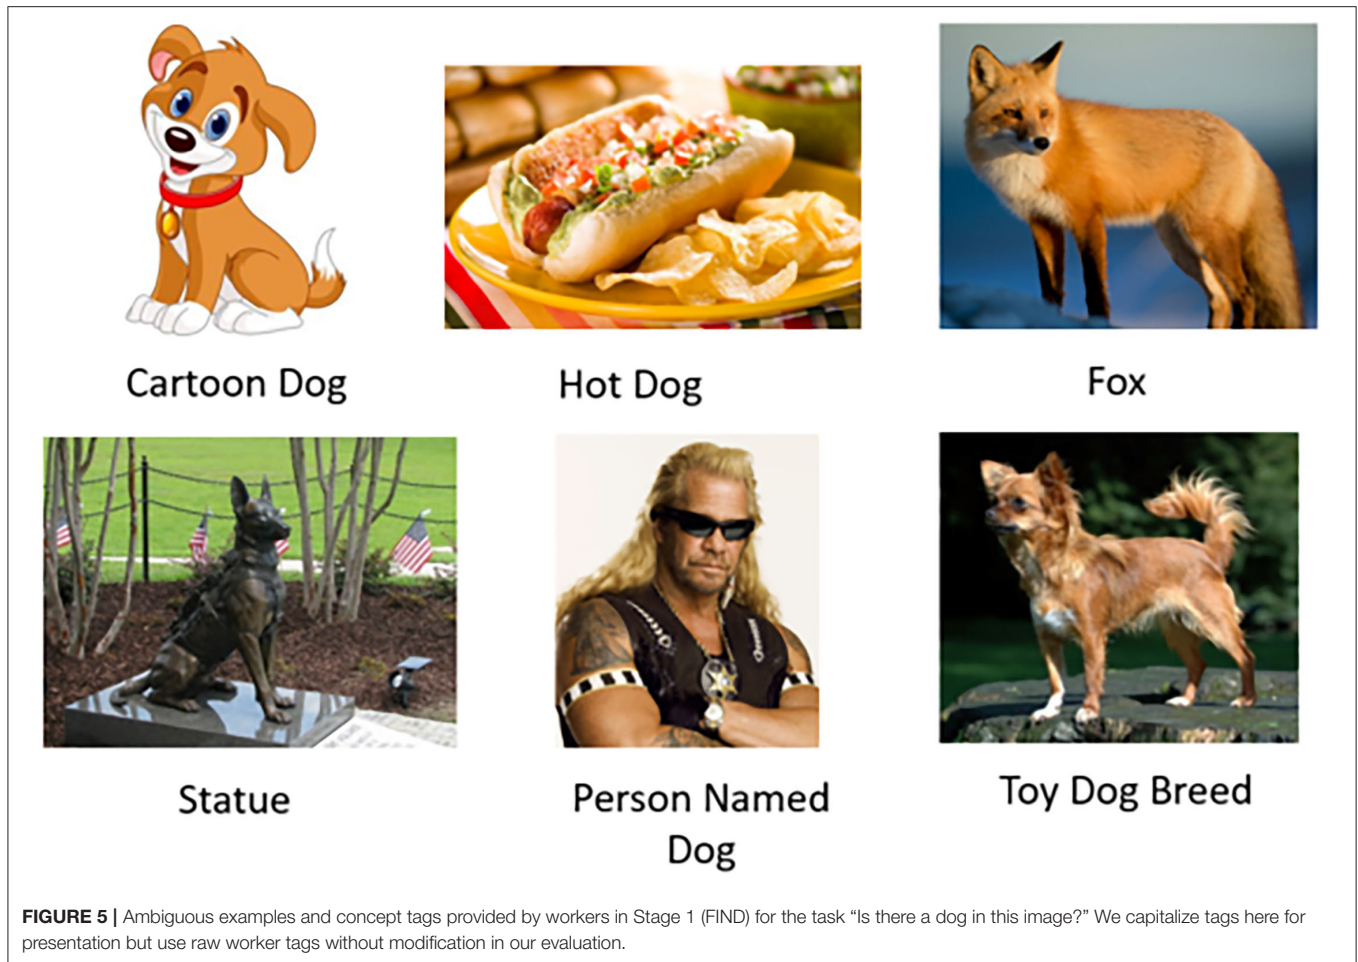
For both collaboration conditions, a total of 15 ambiguous examples (from 15 unique workers) were collected and evaluated with respect to the above criteria.

4.4.2. Quantitative Evaluation of Example Effectiveness in Stage 3 (LABEL)

To evaluate the effectiveness of enriching textual instructions with ambiguous examples from Stage 1 (FIND) and to assess the relative utility of presenting workers with images and/or concept tags from ambiguous examples, we compared the following five conditions. The conditions varied in how annotation instructions were presented to workers in Stage 3 (LABEL):

1. **B0:** No examples were provided along with textual instructions.
2. **B1:** A set of randomly chosen examples were provided along with textual instructions.
3. **IMG:** Only images (but no concept tags) of ambiguous examples were shown to workers along with textual instructions.
4. **TAG:** Only concept tags (but no images) of ambiguous examples were shown to workers along with textual instructions.
5. **IMG+TAG:** Both images and concept tags of ambiguous examples were shown to workers along with textual instructions.

Each of the five conditions above was completed by nine unique workers. Each task consisted of classifying 10 images. Workers were asked to classify each of the 10 images into either the positive or the negative class.



5. RESULTS

5.1. Can Workers Find Ambiguous Concepts?

In this section, we provide insights from pilots of Stage 1 (FIND) followed by a qualitative analysis of ambiguous examples identified by workers in this stage.

5.1.1. Pilot Insights

5.1.1.1. Task Design

Initial pilots of Stage 1 (FIND) revealed two issues: (1) duplicate concepts, and (2) misunderstanding of the task. Some easy-to-find and closely related concepts were naturally repeated multiple times. One type of concept duplication was related to the *seed* example provided by requesters to clarify the task objective. In particular, some workers searched for additional examples of the *same* ambiguity rather than finding *distinct* instances of ambiguity. Another misunderstanding led some workers to submit *generally* ambiguous images, i.e., similar to Google Image Search results for search term “ambiguous image,” rather than images that were ambiguous relative to the specific task instruction “*Is there a dog in this image?*” We acknowledge that our own task design was not immune to ambiguity, so we

incorporated clarifications to instruct workers to find ambiguous examples *distinct* from the seed example and *specific* to the task instructions provided.

5.1.1.2. Unexpected Ambiguous Concepts

However, our pilots also revealed workers’ ability to identify surprising examples of ambiguous concepts we had not anticipated. Some of these examples were educational and helped the paper authors learn about the nuances of our task. For example, one worker returned an image of a Chihuahua (a small dog breed) along with the concept tag “toy dog.” In trying to understand the worker’s intent, we learned that the term “toy dog” is a synonym for small dog breeds (see text footnote 2). Prior to that, our interpretation of the “toy dog” concept was limited to children’s toys. This insight inspired Task 3 (“Is there a toy dog in this image?”) with two different interpretations (section 4.3). Another unexpected ambiguous example was the picture of a man (**Figure 5**). We initially jumped to the conclusion that the worker’s response was spam, but on closer inspection discovered that the picture displayed reality show celebrity “Dog the Bounty Hunter”³ These instances are excellent illustrations of

³<http://www.dogthebountyhunter.com>.

TABLE 1 | Percentage of correct, unique, and useful examples from Stage 1 (FIND).

	Correct	Unique	Useful
No collaboration	60.0	26.7	26.7
Collaboration	93.0	40.0	33.3

The bold value indicates the largest values per column.

the possibility that crowd workers may interpret task instructions in valid and original ways entirely unanticipated by requesters.

5.1.2. Qualitative Assessment of Example Characteristics

We employed qualitative coding to assess whether worker submissions met each of the quality criteria (Correctness, Uniqueness, and Usefulness). **Table 1** shows the percentage of ambiguous examples meeting these criteria for the two conditions with and without collaboration, respectively. Our hypothesis that collaboration among workers can help produce higher quality ambiguous examples is supported by our results. Results show that, compared to no collaboration, a collaborative workflow produced substantially greater proportions of correct (93 vs. 60%), unique (40 vs. 27%), and useful (33 vs. 27%) ambiguous examples. A potential explanation for this result is that exposing workers to a variety of ambiguous concepts upfront may assist them in exploring the space of yet uncovered ambiguities more effectively.

5.2. Can Ambiguous Examples Improve Annotation Accuracy?

Next, we report quantitative results on how ambiguous examples—found in Stage 1 and selected and labeled in Stage 2—can be used as instructional material to improve annotation accuracy in Stage 3. We also provide an analysis of annotation errors.

5.2.1. Effectiveness of Ambiguous Examples

Our hypothesis is that these examples can be used to help delineate the boundary of our annotation task and, hence, teach annotation guidelines to crowd workers better than randomly chosen examples. **Table 2** reports crowd annotation accuracy for each of the six tasks broken down by experimental condition.

5.2.1.1. Using Examples to Teach Annotation Guidelines

Intuitively, providing examples to workers helps them to better understand the intended labeling task (Wu and Quinn, 2017). Comparing designs B0 and B1 in **Table 2**, we clearly see that providing examples (B1) almost always produces more accurate labeling than a design that provides no examples (B0). In addition to this, the IMG design performs better than B1. This shows that the kind of examples that are provided is also important. Showing ambiguous examples is clearly superior to showing randomly chosen examples. This supports our hypothesis: ambiguous examples appear to delineate labeling boundaries for the task better than random examples.

5.2.1.2. Instances vs. Concepts

Best practices suggest that requesters provide examples when designing their tasks (Wu and Quinn, 2017). We include this design in our evaluation as B1. An alternate design is to show concepts as examples instead of specific instances; this is our design TAG, shown in **Table 2**. For example, for the task “Is there a Dog in this image?”, instead of showing a dog statue image, we could simply provide the example concept “Inanimate Objects” should be labeled as NO. Results in **Table 2** show that TAG consistently outperforms IMG, showing that teaching *via* example concepts can be superior to teach *via* example instances.

5.2.1.3. Concepts Only vs. Concepts and Examples

Surprisingly, workers who were presented shown only the concept tags performed better than workers who were shown concept tags along with an example image for each concept. Hence, the particular instance chosen may not represent the concept well. This might be overcome by better selecting a more representative example for a concept or showing more examples for each concept. We leave such questions for future study.

5.2.2. Sources of Worker Errors

5.2.2.1. Difficult vs. Subjective Questions

Table 3 shows accuracy for categories “Similar Animal” and “Cartoon” for Task 1b (section 4.3). We see that some concepts appear more difficult, such as correctly labeling a wolf or a fox. Annotators appear to need some world knowledge or training of differences between species in order to correctly distinguish such examples vs. dogs. Such concepts seem more difficult to teach; even though the accuracy improves, the improvement is less than we see with other concepts. In contrast, for Cartoon Dog (an example of a subjective question), adding this category to the illustrative examples greatly reduces the ambiguity for annotators. Other concepts like “Robot” and “Statue” also show large improvements in accuracy.

5.2.2.2. Learning Closely Related Concepts

To see if crowdworkers learn closely related concepts without being explicitly shown examples, consider “Robot Dog” and “Stuffed Toy” as two types of a larger “Toy Dog” children’s toy concept. In Task 1b, the workers are shown the concept “Robot Dog” as examples labeled as NO, without being shown an example for “Stuffed Toy.” **Table 3** shows that workers learn the related concept “Stuffed Toy” and accurately label the instances that belong to this concept. The performance gain for the concept “Toy Dog” is the same as the gain for “Robot Dog,” when we compare design IMG+TAG and B1. Other similarly unseen concepts [marked with an asterisk (*) in the table] show that workers are able to learn the requester’s intent for unseen concepts if given examples of other, similar concepts.

5.2.2.3. Peer Agreement With Ambiguous Examples

It is not always possible or cost-effective to obtain expert/gold labels for tasks, so requesters often rely on peer-agreement between workers to estimate worker reliability. Similarly, majority voting or weighted voting is often used to aggregate worker labels for consensus (Hung et al., 2013; Sheshadri and

TABLE 2 | Worker accuracy [%] for all six tasks by condition.

Design	Task 1a	Task 1b	Task 2a	Task 2b	Task 3a	Task 3b
B0	75.6	70.1	47.8	78.7	69.1	92.9
B1	83.0	66.4	59.0	85.5	78.7	96.0
IMG	88.0	89.5	68.5	85.8	89.2	93.5
TAG	91.0	91.0	79.0	87.0	91.0	96.9
IMG+TAG	91.4	87.0	81.8	86.4	88.3	96.6

The bold value indicates the largest values per column.

TABLE 3 | Worker accuracy [%] on Task 1b, broken down by concept category.

Design	Similar animal	Stuffed toy*	Robot	Statue	Cartoon	...
B0	37.0	22.2	33.3	18.5	62.2	
B1	14.8	22.2	25.9	29.6	57.8	
IMG	44.4	100.0	100.0	92.6	100.0	
TAG	74.1	100.0	88.9	88.9	97.8	
IMG+TAG	48.1	88.9	92.6	77.8	97.8	
Design	Objects*	Unseen ambiguity*	Small breed	Plane	Dog	Another animal
B0	74.1	88.9	88.9	100.0	100.0	100.0
B1	59.3	82.2	94.4	100.0	100.0	100.0
IMG	100.0	75.6	88.9	100.0	95.6	100.0
TAG	100.0	80.0	94.4	100.0	91.1	100.0
IMG+TAG	96.3	77.8	91.7	96.3	95.6	88.9

We see that some hard concepts cannot be easily disambiguated, e.g., Similar Animal. Concepts for which workers were not shown any examples are marked with an asterisk (*). The bold value indicates the largest values per column.

TABLE 4 | Worker accuracy [%] on ambiguous vs unambiguous categories with baseline B1.

Task	Unambiguous	Ambiguous
1a	96.1	72.2
1b	98.6	41.7
2a	86.8	42.2
2b	98.5	82.5
3a	82.5	75.8
3b	98.6	93.8

For Tasks 1b and 2a, the majority vote (among nine workers) on ambiguous examples is wrong.

The bold values indicate substantially lower than the other values in the table.

Lease, 2013). However, we also know that when workers have consistent, systematic group biases, the aggregation will serve to reinforce and amplify the group bias rather than mitigate it (Ipeirotis et al., 2010; Sen et al., 2015; Dumitrache et al., 2018; Fazelpour and De-Arteaga, 2022).

While we find agreement often correlates with accuracy, and so have largely omitted to report it in this study, we do find several concepts for which the majority chooses wrong answers, producing high agreement but low accuracy. Recall that our results are reported over nine workers per example, whereas typical studies use a plurality of three or five workers. Also recall that Tasks 1b and 2a (section 4.3) represent two of our less intuitive annotation tasks for which requester intent may be

at odds with worker intuition, requiring greater task clarity for over-coming worker bias.

Table 4 shows majority vote accuracy for these tasks for the baseline B1 design which (perhaps typical of many requesters) includes illustrative examples but not necessarily the most informative ones. Despite collecting labels from nine different workers, the majority is still wrong, with majority vote accuracy on ambiguous examples falling below 50%.

6. CONCLUSION AND FUTURE WORK

6.1. Summary and Contributions

Quality assurance for labeled data remains a challenge today. In chasing the potential advantages crowdsourcing has to offer, some important quality assurance best practices from traditional export annotation workflows may be lost. Our work adds to the existing literature on mechanisms to disambiguate nuanced class categories and, thus, improve labeling guidelines and, in effect, classification decisions in crowdsourced data annotation.

In this study, we presented a three-stage FIND-RESOLVE-LABEL workflow as a novel mapping of traditional annotation processes, involving iterative refinement of guidelines by expert annotators, onto a light-weight, structured task design suitable for crowdsourcing. Through careful task design and intelligent distribution of effort between crowd workers and requesters, it may be possible for the crowd to play a valuable role in reducing requester effort while also helping requesters to better understand

the nuances and edge cases of their intended annotation taxonomy in order to generate clearer task instructions for the crowd. In contrast to prior work, our approach is proactive and open-ended, leveraging crowd workers' unconstrained creativity and intelligence to identify ambiguous examples online through an Internet search, proactively enriching task instructions with these examples upfront before the annotation process commences.

While including illustrative examples in instructions is known to be helpful (Wu and Quinn, 2017), we have shown that not all examples are equally informative to annotators and that intelligently selecting ambiguous corner-cases can improve labeling quality. Our results revealed that the crowd performed worst on ambiguous instances and, thus, can benefit the most from help for cases where requester intents run counter to annotators' internal biases or intuitions. For some instances of ambiguity, we observed high agreement among workers on answers contrary to what the requester defines as correct. Such tasks are likely to produce an incorrect label even when we employ intelligent answer aggregation techniques. Techniques like ours to refine instruction clarity are particularly critical in such cases.

Finally, we found that workers were able to infer the correct labels for concepts closely related to the target concept. This result suggests that it may not be necessary to identify and clarify *all* ambiguous concepts that could potentially be encountered during the task. An intelligently selected set of clarifying examples may enable the crowd to disambiguate labels of unseen examples accurately even if not all instances of ambiguity are exhaustively covered.

6.2. Limitations

In this study, we propose a novel workflow for addressing the issue of inherent ambiguity in data classification settings. However, our study is not without limitations. First, our study focuses on a specific type of annotation task (image classification). While our workflow design targets data classification tasks in general, further study is needed to empirically validate the usefulness of this approach for other data annotation settings.

Second, our approach is based on the assumption that characteristics of ambiguous instances contributed by the crowd *via* external Internet search will match those of the dataset being evaluated. However, this assumption may not always be met depending on the domain and modality of the dataset. Certain datasets may not be represented *via* publicly available external search. In that case, additional building blocks would be needed in the workflow to enable effective search over a private data repository. Existing solutions for external search may already cluster results based on representative groups or classes. Our empirical results leave open the question of to what extent this feature could have influenced or facilitated the task for workers.

Third, our evaluation is limited in terms of datasets, task types, and the size of our participant sample. Caution is warranted in generalizing our results beyond the specific evaluation setting, e.g., since characteristics of the dataset can influence the results. Given the limited size of our participant sample and the fact that

crowd populations can be heterogeneous, our empirical data was amenable only to descriptive statistics but not to null hypothesis significance tests. In conclusion, our results should be considered indicative of the potential usefulness of our approach rather than being fully definitive. Further study is needed to validate our approach in a more statistically robust and generalizable manner *via* larger samples.

6.3. Future Study

While we evaluate our strategy on an image labeling task, our approach is more general and could be usefully extended to other domains and tasks. For example, consider collecting document relevance judgments in information retrieval (Alonso et al., 2008; Scholer et al., 2013; McDonnell et al., 2016), where user *information needs* are often subjective, vague, and incomplete. Such generalization may raise new challenges. For example, the image classification task used in our study allows us to point workers to online image searches. However, other domains may require additional or different search tools (e.g., access to collections of domain-specific text documents) for workers to be able to effectively identify ambiguous corner cases.

Alonso (2015) proposes having workers perform practice tasks to get familiarized with the data and thus increase annotation performance for future tasks. While our experimental setup prevented workers from performing more than one task to avoid potential learning effects, future study may explore and leverage workers' ability to improve their performance for certain types of ambiguity over time. For example, we may expect that workers who completed Stage 1 are better prepared for Stage 3 given that they have already engaged in the mental exercise of critically exploring the decision boundary of the class taxonomy in question.

Another best practice from LDC is deriving a decision tree for common ambiguous cases which annotators can follow as a principled and consistent way to determine the label of ambiguous examples (Griffitt and Strassel, 2016). How might we use the crowd to induce such a decision tree? Prior design study in effectively engaging the crowd in clustering (Chang et al., 2016) can guide design considerations for this challenge.

In our study, Stage 2 RESOLVE required requesters to select ambiguous examples. Future study may explore variants of Stage 1 FIND where requesters *filter* the ambiguous examples provided by the crowd. There is an opportunity for saving requester effort if both of these stages are combined. For instance, examples selected in the filtering step of Stage 1 can be fed forward to reduce the example set considered for labeling in Stage 2. Another method would be to have requesters perform labeling simultaneously with filtering in Stage 1, eliminating Stage 2 altogether. Finally, if the requester deems the label quality of Stage 3 insufficient and initiates another cycle of ambiguity reduction *via* Stages 1 and 2 those stages could start with examples already identified in the prior cycle.

A variety of other directions can be envisioned for further reducing requester effort. For example, the crowd could be called upon to verify and prune ambiguous examples collected in the initial FIND stage. Examples flagged as spam or assigned a low

ambiguity rating could be automatically discarded to minimize requester involvement in Stage 2. Crowd ambiguity ratings could also be used to rank examples for guiding requesters' attention in Stage 2. A more ambitious direction for future study would be to systematically explore how well and under what circumstances the crowd is able to correctly infer requester intent. Generalizable insights about this question would enable researchers to design strategies that eliminate requester involvement altogether under certain conditions.

DATA AVAILABILITY STATEMENT

The datasets presented in this article are not readily available because of an unrecoverable data loss. Data that is available will be shared online upon acceptance at <https://www.ischool.utexas.edu/~ml/publications/>. Requests to access the datasets should be directed to ML, ml@utexas.edu.

ETHICS STATEMENT

Ethical review and approval was not required for the study on human participants in accordance with the local legislation and institutional requirements. Written informed consent for

participation was not required for this study in accordance with the national legislation and the institutional requirements.

AUTHOR CONTRIBUTIONS

VP: implementation, experimentation, data analysis, and manuscript preparation. MS: manuscript preparation. ML: advising on research design and implementation and manuscript preparation. All authors contributed to the article and approved the submitted version.

FUNDING

This research was supported in part by the Micron Foundation and by Good Systems (<https://goodsystems.utexas.edu>), a UT Austin Grand Challenge to develop responsible AI technologies.

ACKNOWLEDGMENTS

We thank our many talented crowd contributors, without whom our research would not be possible. We thank the reviewers for their time and assistance in helping us to further improve this study.

REFERENCES

- Ahmad, S., Battle, A., Malkani, Z., and Kamvar, S. (2011). "The jabberwocky programming environment for structured social computing," in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 53–64. doi: 10.1145/2047196.2047203
- Al Kuwatly, H., Wich, M., and Groh, G. (2020). "Identifying and measuring annotator bias based on annotators' demographic characteristics," in *Proceedings of the Fourth Workshop on Online Abuse and Harms (at EMNLP)* (Association for Computational Linguistics), 184–190. doi: 10.18653/v1/2020.alw-1.21
- Alonso, O. (2015). "Practical lessons for gathering quality labels at scale," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1089–1092. doi: 10.1145/2766462.2776778
- Alonso, O., Rose, D. E., and Stewart, B. (2008). "Crowdsourcing for relevance evaluation," in *ACM SigIR Forum*, Vol. 42, 9–15. doi: 10.1145/1480506.1480508
- Amazon Mechanical Turk (2017). *Tutorial: Best Practices for Managing Workers in Follow-Up Surveys or Longitudinal Studies*. Available online at: <https://blog.mturk.com/tutorial-best-practices-for-managing-workers-in-follow-up-surveys-or-/longitudinal-studies-4d0732a7319b>
- Attenberg, J., Ipeirotis, P. G., and Provost, F. (2011). "Beat the machine: challenging workers to find the unknown unknowns," in *Proceedings of the 11th AAAI Conference on Human Computation, AAAIWS'11-11* (AAAI Press), 2–7.
- Barowy, D. W., Curtsinger, C., Berger, E. D., and McGregor, A. (2016). Automan: a platform for integrating human-based and digital computation. *Commun. ACM* 59, 102–109. doi: 10.1145/2927928
- Bernstein, M. S., Little, G., Miller, R. C., Hartmann, B., Ackerman, M. S., Karger, D. R., et al. (2010). "Soylent: a word processor with a crowd inside," in *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, 313–322. doi: 10.1145/1866029.1866078
- Bragg, J., and Weld, D. S. (2018). "Sprout: crowd-powered task design for crowdsourcing," in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, 165–176. doi: 10.1145/3242587.3242598
- Cabrera, A. A., Druck, A. J., Hong, J. I., and Perer, A. (2021). Discovering and validating ai errors with crowdsourced failure reports. *Proc. ACM Hum.-Comput. Interact.* 5: CSCW2. doi: 10.1145/3479569
- Chang, J. C., Amershi, S., and Kamar, E. (2017). "Revolt: collaborative crowdsourcing for labeling machine learning datasets," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2334–2346. doi: 10.1145/3025453.3026044
- Chang, J. C., Kittur, A., and Hahn, N. (2016). "Alloy: clustering with crowds and computation," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 3180–3191. doi: 10.1145/2858036.2858411
- Chen, J. J., Menezes, N. J., Bradley, A. D., and North, T. (2011). "Opportunities for crowdsourcing research on amazon mechanical Turk," in *ACM CHI Workshop on Crowdsourcing and Human Computation*.
- Chen, Q., Bragg, J., Chilton, L. B., and Weld, D. S. (2019). "Cicero: multi-turn, contextual argumentation for accurate crowdsourcing," in *Proceedings of the 2019 ACM CHI Conference on Human Factors in Computing Systems*, 1–14. doi: 10.1145/3290605.3300761
- Chen, Y., Ghosh, A., Kearns, M., Roughgarden, T., and Vaughan, J. W. (2016). Mathematical foundations for social computing. *Commun. ACM* 59, 102–108. doi: 10.1145/2960403
- Cheng, J., and Bernstein, M. S. (2015). "Flock: hybrid crowd-machine learning classifiers," in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 600–611. doi: 10.1145/2675133.2675214
- Cole, C. (2011). A theory of information need for information retrieval that connects information to knowledge. *J. Assoc. Inform. Sci. Technol.* 62, 1216–1231. doi: 10.1002/asi.21541
- Dow, S., Kulkarni, A., Klemmer, S., and Hartmann, B. (2012). "Shepherding the crowd yields better work," in *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, 1013–1022. doi: 10.1145/2145204.2145355
- Drapeau, R., Chilton, L. B., Bragg, J., and Weld, D. S. (2016). "Microtalk: using argumentation to improve crowdsourcing accuracy," in *Fourth AAAI Conference on Human Computation and Crowdsourcing*.
- Draws, T., Rieger, A., Inel, O., Gadiraju, U., and Tintarev, N. (2021). "A checklist to combat cognitive biases in crowdsourcing," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 9, 48–59.
- Dumitrache, A., Inel, O., Aroyo, L., Timmermans, B., and Welty, C. (2018). Crowdtruth 2.0: Quality metrics for crowdsourcing with disagreement. *arXiv [preprint] arXiv:1808.06080*. Available online at: <https://arxiv.org/abs/1808.06080>

- Egelman, S., Chi, E. H., and Dow, S. (2014). "Crowdsourcing in HCI research," in *Ways of Knowing in HCI* (Springer), 267–289. doi: 10.1007/978-1-4939-0378-8_11
- Fazelpour, S., and De-Arteaga, M. (2022). Diversity in sociotechnical machine learning systems. *arXiv [Preprint]*. arXiv: 2107.09163. Available online at: <https://arxiv.org/pdf/2107.09163.pdf>
- Franklin, M. J., Kossman, D., Kraska, T., Ramesh, S., and Xin, R. (2011). "Crowddb: answering queries with crowdsourcing," in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, 61–72. doi: 10.1145/1989323.1989331
- Gadiraju, U., Yang, J., and Bozzon, A. (2017). "Clarity is a worthwhile quality: on the role of task clarity in microtask crowdsourcing," in *Proceedings of the 28th ACM Conference on Hypertext and Social Media*, 5–14. doi: 10.1145/3078714.3078715
- Gaikwad, S. N. S., Whiting, M. E., Gamage, D., Mullings, C. A., Majeti, D., Goyal, S., et al. (2017). "The daemon crowdsourcing marketplace," in *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 1–4. doi: 10.1145/3022198.3023270
- Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1995). *Design Patterns: Elements of Reusable Object-oriented Software*. Boston, MA: Addison-Wesley Longman Publishing Co., Inc.
- Geva, M., Goldberg, Y., and Berant, J. (2019). "Are we modeling the task or the annotator? An investigation of annotator bias in natural language understanding datasets," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (Hong Kong), 1161–1166. doi: 10.18653/v1/D19-1107
- Goto, S., Ishida, T., and Lin, D. (2016). "Understanding crowdsourcing workflow: modeling and optimizing iterative and parallel processes," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, Vol. 4*.
- Grady, C., and Lease, M. (2010). "Crowdsourcing document relevance assessment with mechanical Turk," in *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data With Amazon's Mechanical Turk* (Association for Computational Linguistics), 172–179.
- Griffitt, K., and Strassel, S. (2016). "The query of everything: developing open-domain, natural-language queries for bolt information retrieval," in *LREC*.
- Huang, G., Wu, M.-H., and Quinn, A. J. (2021). "Task design for crowdsourcing complex cognitive skills," in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–7. doi: 10.1145/3411763.3443447
- Hung, N. Q. V., Tam, N. T., Tran, L. N., and Aberer, K. (2013). "An evaluation of aggregation techniques in crowdsourcing," in *International Conference on Web Information Systems Engineering* (Springer), 1–15. doi: 10.1007/978-3-642-41154-0_1
- Ipeirotis, P. G., Provost, F., and Wang, J. (2010). "Quality management on amazon mechanical Turk," in *Proceedings of the ACM SIGKDD Workshop on Human Computation*, 64–67. doi: 10.1145/1837885.1837906
- Jones, G. J. (2013). "An introduction to crowdsourcing for language and multimedia technology research," in *Information Retrieval Meets Information Visualization* (Springer), 132–154. doi: 10.1007/978-3-642-36415-0_9
- Jung, H. J., and Lease, M. (2015). "Modeling temporal crowd work quality with limited supervision," in *Proceedings of the 3rd AAAI Conference on Human Computation (HCOMP)*, 83–91.
- Kalra, K., Patwardhan, M., and Karande, S. (2017). "Shifts in rating bias due to scale saturation," in *Human Computation and Crowdsourcing (HCOMP): Works-in-Progress Track*.
- Kazai, G., Kamps, J., Koolen, M., and Milic-Frayling, N. (2011). "Crowdsourcing for book search evaluation: impact of hit design on comparative system ranking," in *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 205–214. doi: 10.1145/2009916.2009947
- Kelling, S., Gerbracht, J., Fink, D., Lagoze, C., Wong, W.-K., Yu, J., et al. (2013). A human/computer learning network to improve biodiversity conservation and research. *AI Mag.* 34, 10. doi: 10.1609/aimag.v34i1.2431
- Kinney, K. A., Huffman, S. B., and Zhai, J. (2008). "How evaluator domain expertise affects search result relevance judgments," in *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, 591–598. doi: 10.1145/1458082.1458160
- Kittur, A., Khamkar, S., André, P., and Kraut, R. (2012). "Crowdweaver: visually managing complex crowd work," in *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, 1033–1036. doi: 10.1145/2145204.2145357
- Kittur, A., Smus, B., Khamkar, S., and Kraut, R. E. (2011). "Crowdforge: crowdsourcing complex work," in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, 43–52. doi: 10.1145/2047196.2047202
- Kovashka, A., Russakovsky, O., Fei-Fei, L., Grauman, K., et al. (2016). Crowdsourcing in computer vision. *Found. Trends Comput. Graph. Vis.* 10, 177–243. doi: 10.1561/06000000071
- Krivoshchev, E., Bykau, S., Casati, F., and Prabhakar, S. (2020). Detecting and preventing confused labels in crowdsourced data. *Proc. VLDB Endow.* 13, 2522–2535. doi: 10.14778/3407790.3407842
- Kulesza, T., Amershi, S., Caruana, R., Fisher, D., and Charles, D. (2014). "Structured labeling for facilitating concept evolution in machine learning," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 3075–3084. doi: 10.1145/2556288.2557238
- Kulkarni, A., Gutheim, P., Narula, P., Rolnitzky, D., Parikh, T., and Hartmann, B. (2012b). Mobileworks: designing for quality in a managed crowdsourcing architecture. *IEEE Intern. Comput.* 16, 28–35. doi: 10.1109/MIC.2012.72
- Kulkarni, A., Can, M., and Hartmann, B. (2012a). "Collaboratively crowdsourcing workflows with Turkomatic," in *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (IEEE)*, 1003–1012. doi: 10.1145/2145204.2145354
- Kutlu, M., McDonnell, T., Elsayed, T., and Lease, M. (2020). Annotator rationales for labeling tasks in crowdsourcing. *J. Artif. Intell. Res.* 69, 143–189. doi: 10.1613/jair.112012
- Lintott, C. J., Schawinski, K., Slosar, A., Land, K., Bamford, S., Thomas, D., et al. (2008). Galaxy zoo: morphologies derived from visual inspection of galaxies from the Sloan digital sky survey. *Monthly Not. R. Astron. Soc.* 389, 1179–1189. doi: 10.1111/j.1365-2966.2008.13689.x
- Little, G., Chilton, L. B., Goldman, M., and Miller, R. C. (2010a). "Exploring iterative and parallel human computation processes," in *Proceedings of the ACM SIGKDD Workshop on Human Computation*, 68–76. doi: 10.1145/1837885.1837907
- Little, G., Chilton, L. B., Goldman, M., and Miller, R. C. (2010b). "Turkit: human computation algorithms on mechanical Turk," in *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, 57–66. doi: 10.1145/1866029.1866040
- Liu, A., Guerra, S., Fung, I., Matute, G., Kamar, E., and Lasecki, W. (2020). "Towards hybrid human-AI workflows for unknown detection," in *Proceedings of The Web Conference 2020, WWW '20* (New York, NY: Association for Computing Machinery), 2432–2442. doi: 10.1145/3366423.3380306
- Manam, V. C., and Quinn, A. J. (2018). "Wingit: efficient refinement of unclear task instructions," in *Proceedings of the 6th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*.
- Manam, V. K. C., Jampani, D., Zaim, M., Wu, M.-H., and J. Quinn, A. (2019). "Taskmate: a mechanism to improve the quality of instructions in crowdsourcing," in *Companion Proceedings of The 2019 World Wide Web Conference*, 1121–1130. doi: 10.1145/3308560.3317081
- Marshall, C. C., and Shipman, F. M. (2013). "Experiences surveying the crowd: reflections on methods, participation, and reliability," in *Proceedings of the 5th Annual ACM Web Science Conference*, 234–243. doi: 10.1145/2464464.2464485
- McDonnell, T., Lease, M., Elsayad, T., and Kutlu, M. (2016). "Why is that relevant? Collecting annotator rationales for relevance judgments," in *Proceedings of the 4th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 10. doi: 10.24963/ijcai.2017/692
- Nguyen, A. T., Halpern, M., Wallace, B. C., and Lease, M. (2016). "Probabilistic modeling for crowdsourcing partially-subjective ratings," in *Proceedings of the 4th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 149–158.
- Nouri, Z., Gadiraju, U., Engels, G., and Wachsmuth, H. (2021a). "What is unclear? Computational assessment of task clarity in crowdsourcing," in *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, 165–175. doi: 10.1145/3465336.3475109

- Nouri, Z., Prakash, N., Gadiraju, U., and Wachsmuth, H. (2021b). “iclarify-a tool to help requesters iteratively improve task descriptions in crowdsourcing,” in *Proceedings of the 9th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*.
- Papoutsaki, A., Guo, H., Metaxa-Kakavouli, D., Gramazio, C., Rasley, J., Xie, W., et al. (2015). “Crowdsourcing from scratch: a pragmatic experiment in data collection by novice requesters,” in *Third AAAI Conference on Human Computation and Crowdsourcing*.
- Pickard, G., Pan, W., Rahwan, I., Cebrian, M., Crane, R., Madan, A., et al. (2011). Time-critical social mobilization. *Science* 334, 509–512. doi: 10.1126/science.1205869
- Retelny, D., Robaszekiewicz, S., To, A., Lasecki, W. S., Patel, J., Rahmati, N., et al. (2014). “Expert crowdsourcing with flash teams,” in *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, 75–85. doi: 10.1145/2642918.2647409
- Rosser, H., and Wiggins, A. (2019). “Crowds and camera traps: genres in online citizen science projects,” in *Proceedings of the 52nd Hawaii International Conference on System Sciences*. doi: 10.24251/HICSS.2019.637
- Schaekermann, M., Goh, J., Larson, K., and Law, E. (2018). Resolvable vs. irresolvable disagreement: a study on worker deliberation in crowd work. *Proc. ACM Hum. Comput. Interact.* 2, 1–19. doi: 10.1145/3274423
- Scholer, F., Kelly, D., Wu, W.-C., Lee, H. S., and Webber, W. (2013). “The effect of threshold priming and need for cognition on relevance calibration and assessment,” in *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 623–632. doi: 10.1145/2484028.2484090
- Sen, S., Giesel, M. E., Gold, R., Hillmann, B., Lesicko, M., Naden, S., et al. (2015). “Turkers, scholars, arafat and peace: cultural communities and algorithmic gold standards,” in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 826–838. doi: 10.1145/2675133.2675285
- Sheshadri, A., and Lease, M. (2013). “SQUARE: a benchmark for research on computing crowd consensus,” in *Proceedings of the 1st AAAI Conference on Human Computation (HCOMP)*, 156–164.
- Snow, R., O’Connor, B., Jurafsky, D., and Ng, A. Y. (2008). “Cheap and fast—but is it good?: evaluating non-expert annotations for natural language tasks,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics), 254–263. doi: 10.3115/1613715.1613751
- Sorokin, A., and Forsyth, D. (2008). “Utility data annotation with amazon mechanical Turk,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 2008, CVPRW’08*, 1–8. doi: 10.1109/CVPRW.2008.4562953
- Tian, Y., and Zhu, J. (2012). “Learning from crowds in the presence of schools of thought,” in *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 226–234. doi: 10.1145/2339530.2339571
- Vakharia, D., and Lease, M. (2015). “Beyond mechanical Turk: an analysis of paid crowd work platforms,” in *Proceedings of the iConference*.
- Vandenhof, C. (2019). “A hybrid approach to identifying unknown unknowns of predictive models,” in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7, 180–187.
- Vidgen, B., and Derczynski, L. (2020). Directions in abusive language training data, a systematic review: garbage in, garbage out. *PLoS ONE* 15:e0243300. doi: 10.1371/journal.pone.0243300
- Vogels, W. (2007). *Help Find Jim Gray*. Available online at: https://www.allthingsdistributed.com/2007/02/help_find_jim_gray.html
- Wang, F.-Y., Zeng, D., Hendler, J. A., Zhang, Q., Feng, Z., Gao, Y., et al. (2010). A study of the human flesh search engine: crowd-powered expansion of online knowledge. *Computer* 43, 45–53. doi: 10.1109/MC.2010.216
- Wu, M.-H., and Quinn, A. J. (2017). “Confusing the crowd: task instruction quality on amazon mechanical Turk,” in *Proceedings of the 5th AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*.
- Zheng, Y., Li, G., Li, Y., Shan, C., and Cheng, R. (2017). Truth inference in crowdsourcing: is the problem solved? *Proc. VLDB Endow.* 10, 541–552. doi: 10.14778/3055540.3055547

Author Disclaimer: Any opinions, findings, and conclusions or recommendations expressed by the authors are entirely their own and do not represent those of the sponsors.

Conflict of Interest: ML is engaged as an Amazon Scholar, proposed methods are largely general across paid and volunteer labeling platforms.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Pradhan, Schaekermann and Lease. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.