Empirical Methodology for Crowdsourcing Ground Truth

Anca Dumitrache [a,e], Oana Inel [a,f], Benjamin Timmermans [a,b], Carlos Ortiz [c], Robert-Jan Sips [b,g], Lora Aroyo [a,d] and Chris Welty [a,d]

[a] *Vrije Universiteit Amsterdam, De Boelelaan 1081, Amsterdam, Netherlands*
E-mail: {anca.dmtrch,oana.inel,l.m.aroyo,cawelty}@gmail.com
[b] *IBM Center for Advanced Studies Benelux, Johan Huizingalaan 765, Amsterdam, Netherlands,*
E-mail: b.timmermans@nl.ibm.com, rhjsips@gmail.com
[c] *Netherlands eScience Center, Amsterdam, Netherlands,*
E-mail: c.martinez@esciencecenter.nl
[d] *Google, New York, USA*
[e] *FD Mediagroep, Amsterdam, Netherlands*
[f] *TU Delft, Delft, Netherlands*
[g] *myTomorrows, Amsterdam, Netherlands*

Abstract. The process of gathering ground truth data through human annotation is a major bottleneck in the use of information extraction methods for populating the Semantic Web. Crowdsourcing-based approaches are gaining popularity in the attempt to solve the issues related to volume of data and lack of annotators. Typically these practices use inter-annotator agreement as a measure of quality. However, in many domains, such as event detection, there is ambiguity in the data, as well as a multitude of perspectives of the information examples. We present an empirically derived methodology for efficiently gathering of ground truth data in a diverse set of use cases covering a variety of domains and annotation tasks. Central to our approach is the use of CrowdTruth metrics that capture inter-annotator disagreement. We show that measuring disagreement is essential for acquiring a high quality ground truth. We achieve this by comparing the quality of the data aggregated with CrowdTruth metrics with majority vote, over a set of diverse crowdsourcing tasks: *Medical Relation Extraction*, *Twitter Event Identification*, *News Event Extraction* and *Sound Interpretation*. We also show that an increased number of crowd workers leads to growth and stabilization in the quality of annotations, going against the usual practice of employing a small number of annotators.

Keywords: CrowdTruth, ground truth gathering, annotator disagreement, semantic interpretation, medical, event extraction, relation extraction

1. Introduction

Knowledge base curation, or the task of populating knowledge bases, is one of the main research challenges of crowdsourcing the Semantic Web [1]. Knowledge base curation can be done either manually, by asking annotators to populate the knowledge graph by manually extracting triples from unstructured data, or automatically by using information extraction methods that are trained and evaluated on ground truth collected from human annotators. In both cases, the process of gathering the human annotations is the a bottleneck in the entire knowledge base population process. The traditional approach to gathering human an-

notation is to employ experts to perform annotation tasks [2], which is a costly and time consuming process. Additionally, in order to prevent high disagreement among expert annotators, strict annotation guidelines are designed for the experts to follow. On the one hand, creating such guidelines is a lengthy and tedious process, and on the other hand, the annotation task becomes rigid and not reproducible across domains. And, as a result, the entire process needs to be repeated over and over again in every domain and task. Finally, expert annotators are not always available for specific tasks such as open domain question-answering or news events, while many annotation tasks can require multiple interpretations that a single annotator cannot provide [3].

As a solution to those problems, crowdsourcing has become a mainstream approach. It has proved to provide good results in multiple domains: annotating cultural heritage prints [4], medical relation annotation [5], ontology evaluation [6]. Following the central feature of volunteer-based crowdsourcing introduced by [7] that majority voting and high inter-annotator agreement [8] can ensure truthfulness of resulting annotations, most of those approaches are assessing the quality of their crowdsourced data based on the hypothesis [9] that there is only one right answer to each question.

However, this assumption often creates issues in practice. Recent work in collecting annotations for text [10, 11], sounds [12] and images [13, 14] found that disagreement between annotators is not just a result of poor quality work, and can actually be an indicator for other properties of the data, such as ambiguity and uncertainty [15].

Previous experiments we performed [16] also identified issues with the assumption of the one truth: inter-annotator disagreement is usually never captured, either because the number of annotators is too small to capture the full diversity of opinion, or because the crowd data is aggregated with metrics that enforce consensus, such as majority vote. These practices create artificial data that is neither general nor reflects the ambiguity inherent in the data.

To address these issues, we proposed the **Crowd-Truth** methodology for crowdsourcing human annotation by harnessing inter-annotator disagreement, i.e representing the diversity of human interpretations in the ground truth. This is a novel approach for crowdsourcing human annotation that, instead of enforcing agreement between annotators, captures the ambiguity inherent in semantic annotation through the use of

ambiguity-aware metrics for aggregating crowdsourcing responses. Based on this principle, we have implemented the CrowdTruth methodology as part of a framework [17] for machine-human computation, that first introduced the ambiguity-aware metrics and built a pipeline to process crowdsourcing data with these metrics.

In this paper, we extend the definition of our ambiguity-aware methodology (CrowdTruth version 1.0 [17]) to work both with crowdsourcing tasks that are *closed*, i.e. the annotations that can occur in the data are already known, and the workers are asked to validate their existence (e.g. given a news event, decide whether it is expressed in a tweet), and tasks that are *open*, i.e. the annotation space is not known, and workers can freely select all the choices that apply (e.g. given a news piece, select all events that appear in the text). The code for the extended CrowdTruth version 1.1 methodology and metrics is available at: https://git.io/fA3Mq.

We investigate tasks of text and sound annotation, in both domains that typically require expertise from annotators (e.g. medical) and those that don't (open domain). In particular, we look at four crowdsourcing tasks: *Medical Relation Extraction*, *Twitter Event Identification*, *News Event Extraction* and *Sound Interpretation*. The aim is to investigate the role of inter-annotator disagreement as part of the crowdsourcing system by applying the CrowdTruth methodology to collect data over a set of diverse use cases.

Through the use of CrowdTruth aggregation metrics, the interpretations collected from the crowd are transformed into explicit semantics for the various tasks presented in this paper – i.e. relations expressed in sentences, topics / events expressed in tweets and news articles, words describing sounds – thus enabling knowledge base curation for these specific tasks. Furthermore, we prove that capturing disagreement is essential for acquiring high quality semantics. We achieve this by comparing the quality of the data aggregated with CrowdTruth metrics with majority vote, a method which enforces consensus among annotators. By applying our analysis over a set of diverse tasks we show that, even though ambiguity manifests differently depending on the task (e.g. each task has an optimal number of workers necessary to capture the full spectrum of opinions), our theory of inter-annotator disagreement as a property of ambiguity is generalizable for any semantic annotation crowdsourcing task.

The paper makes the following contributions:

1. **comparative analysis of crowdsourcing aggregation methods:** we compare the performance of *ambiguity-aware metrics* and *consensus - enforcing metrics* over a diverse set of crowdsourcing tasks (Sections 4, 5);

2. **stability of crowd results:** we show in several crowdsourcing tasks that *an increased number of crowd workers leads to growth and stabilization in the quality of annotations*, going against the usual practice of employing a small number of annotators (Sections 4, 5);

3. **measuring quality in open-ended tasks:** we present an extension to the CrowdTruth methodology that allows the ambiguity-aware metrics to deal *both with open-ended and closed tasks* (Sections 2, 3), as opposed to the initial version of the CrowdTruth metrics which only processed closed tasks;

4. **semantics of ambiguity:** applying the CrowdTruth methodology we collect richer data that allows to reason about ambiguity of content (in all modality formats, e.g. images, videos and sounds), which is intrinsically relevant to the Semantic Web community.

2. CrowdTruth Methodology

In this section, we describe the CrowdTruth *methodology* version 1.1, for aggregating crowdsourcing data, which offers methods to aggregate both closed an open-ended tasks. Version 1.1 presented in this paper is a generalization of the initial version 1.0 of CrowdTruth [17].

In Section 4 we use a number of annotation tasks in different domains to illustrate its use and gather experimental data to prove the main claim of this research - CrowdTruth methodology provides a viable alternative to traditional consensus-based majority vote crowdsourcing and expert-based ground truth collection. The elements of the CrowdTruth methodology are:

– annotation modeling with the *triangle of disagreement*;
– quality *metrics* for media units (input data), annotations and crowd workers;
– identification of workers with low quality annotations.

Each of these elements is applicable across a variety of domains, content modalities, *e.g.*, text, sounds, images and videos and annotation tasks, *e.g.*, closed and open-ended annotations. The following sub-sections briefly introduce the overview of the methodology elements.

2.1. CrowdTruth quality metrics

Measuring quality in CrowdTruth is done with the triangle of disagreement model (based on the triangle reference [18]), which links together media units, workers, and annotations, as seen in Fig.1. It allows us to assess the quality of each worker, the clarity of each media unit, and the ambiguity, similarity and frequency of each annotation. This model makes it possible to express how the ambiguity in any of the corners disseminates and influences the other components of the triangle. For example, an unclear sentence or an ambiguous annotation scheme would cause more disagreement between workers [19], and thus, both need to be accounted for when measuring the quality of the workers.
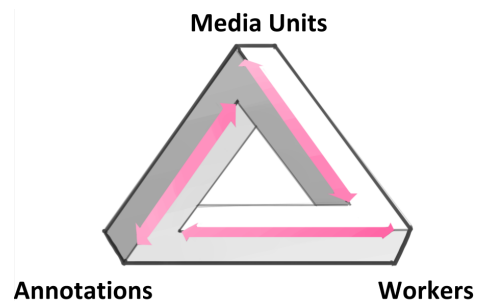

Fig. 1. Triangle of Disagreement

The CrowdTruth quality metrics [19] are designed to capture inter-annotator disagreement in crowdsourcing. The metrics were introduced for *closed tasks*, i.e. multiple choice tasks, where the annotation set is known before running the crowdsourcing task. In this paper, we present an extended version of these metrics (version 1.1), that can be used for both *closed tasks* as well as *open-ended tasks* (i.e. the annotation set is not known beforehand, and the workers can freely select all the choices that apply). The code for the CrowdTruth version 1.1 metrics is available at: https://git.io/fA3Mq.

The quality of the crowdsourced data is measured using a **vector space representation** of the crowd annotations. For *closed tasks*, the annotation vector contains the given answer options in the task template, which the crowd can choose from. For example, the template of a *closed task* can be composed of a multi-

ple choice question, which appears as a list checkboxes or radio buttons, thus, having a finite list of options to choose from.

While for closed tasks the number of elements in the annotation vector is known in advance, for *open-ended tasks* the number of elements in the annotation vector can only be determined when all the judgments for a media unit have been gathered. An example of such a task can be highlighting words or word phrases in a sentence, or as an input text field where the workers can introduce keywords. In this case the answer space is composed of all the unique keywords from all the workers that solved that media unit. As a consequence, all the media units in a closed task have the same answer space, while for open-ended tasks the answer space is different across all the media units.

Although the answer space for open-ended tasks is not known from the beginning, it is still possible to deduce a finite answer space. To achieve this, we added an *answer space dimensionality reduction step* to the methodology for open-ended tasks. Additional goals of this step are to reduce redundancy in the answer space through similarity clustering (e.g. by making sure that synonymous words do not count as disagreement between annotators), and to keep the vector space representation small enough so that the Crowd-Truth quality metrics still produce meaningful values. The method for performing dimensionality reduction is dependent on the annotation task itself.

In the annotation vector, each answer option is a boolean value, showing whether the worker annotated that answer or not. This allows the annotations of each worker on a given media unit to be aggregated, resulting in a **media unit vector** that represents for each option how often it was annotated.

Three core **worker metrics** are defined to differentiate between low-quality and high-quality workers. *Worker-Worker Agreement* (*wwa*) measures the pairwise agreement between two workers across all media units they annotated in common - indicating how close a worker performs compared to workers solving the same task. *Worker-Media Unit Agreement* (*wma*) measures the similarity between the annotations of a worker and the aggregated annotations of the rest of the workers. The average of this metric across all the media units solved gives a measure of how much a worker disagrees with the crowd in the context of all media units. *Average annotations per media unit* (*na*) measures for each worker the total number of annotations they chose per media unit, averaged across all media units they annotated. Since in many tasks work-

ers can choose all the possible annotations, a low quality worker can appear to agree more with the rest of the workers by repeatedly choosing multiple annotations, thus increasing the chance of overlap.

Two **media unit metrics** are defined to assess the quality of each unit. In this paper, we focus on the *Media Unit-Annotation Score (UAS)* – the core Crowd-Truth metric, used to measure the clarity with which the media unit expresses a given annotation. This metric is computed for each media unit and each possible annotation as the cosine between the media unit vector and the unit vector for each possible annotation. This metric is used in evaluating the quality of the Crowd-Truth annotations.

2.2. Spam Removal

After collecting the crowd annotations, but before the evaluation of the data, we perform spam removal. The purpose of this step is to identify the adversarial and low quality workers – e.g. those workers that always pick the same annotations, regardless of the unit. Once identified, the spam workers are removed from the dataset, and their annotations are not used in the evaluation. The methodology for spam removal is based on our previous work in [20], extended in this paper to work also for open-ended tasks.

We identify the low quality workers by applying the core CrowdTruth worker metrics, the worker-worker agreement (*wwa*), worker-media unit agreement (*wma*) and the average number of annotations (*na*) submitted by a worker for one sentence. The first two metrics are used to model the extent to which a given worker agrees with the other annotators. The purpose is not to penalize disagreement with the majority, but rather to identify outliers, *i.e.*, workers that are in constant disagreement. For *closed tasks* where the semantics of the annotations in the answer space could rarely overlap, it is unlikely that a large number of possible annotations will occur for the same media unit. Therefore, the number of annotations per sentence can also indicate spam behavior.

In *open-ended tasks* we apply the same approach. However, we need to acknowledge the fact that open-ended tasks are more prone to disagreement due to the large answer space and thus, the overall agreement between the workers can occur with lower values. Thus, we do not have predefined values for identifying the low-quality workers, but for every task or job we use the following main heuristic: given worker w, if the agreement $wwa(w)$, $wsa(w)$ and optionally, annota-

Table 1
Crowdsourcing Task Details

Task	Type	Media Unit	Annotations
Medical Relation Extraction	closed	sentence	medical relations: *cause*, *treat*, *prevent*, *symptom*, *diagnose*, *side effect*, *location manifestation*, *contraindicate*, *is a*, *part of*, *associated with*, *other*, *none*
Twitter Event Identification	closed	tweet	tweet events: *Davos world economic forum 2014*, *FIFA World Cup 2014*, *Islands disputed between China and Japan*, *2014 anti-China protests in Vietnam*, *Korean MV Sewol ferry ship sinking*, *Japan whaling and dolphin hunting*, *Disappearance of Malaysia Airlines flight 370*, *Ukraine crisis 2014*, *none of the above*
News Event Extraction	open-ended	sentence	words in the sentence
Sound Interpretation	open-ended	sound	tags describing sound

tions per sentence $na(w)$, parameters do not fall within the standard deviation for the task, then worker w is marked as a spammer. To confirm the validity of this metrics we also perform manual evaluation based on sampling of the results.

Based on the specificity of each task, closed or open-ended, the effort required to pick different annotations might vary. For instance, when no good annotation exists in the media unit, the time to complete the annotation is considerably reduced. This can bias the workers towards selecting the option that requires the least work. In order to prevent this, we introduce *in-task effort consistency checks*. Such annotations do not count towards building the ground truth, and are used to reduce the bias from picking the quickest option. For instance, when stating that no annotation is possible in the media unit, the workers also have to write an explanation in a text box for why no annotation were provided.

3. Experimental Setup

The aim of the crowdsourcing experiments described and analyzed in this paper is to show that the CrowdTruth ambiguity-aware crowdsourcing approach produces data with a higher quality than the traditional majority vote where consensus among annotators is enforced. In order to show this, we perform an experiment over a set of four diverse crowdsourcing tasks:

– two closed tasks, i.e. *Medical Relation Extraction*, *Twitter Event Identification*,
– two open-ended tasks, i.e. *News Event Extraction* and *Sound Interpretation*.

These tasks were picked from diverse domains (medical, sound, open), to aid in the generalization of our results. To evaluate the quality of the crowdsourcing data, we constructed a trusted judgments set by combining expert and crowd annotations. The rest of this section describes the details of the crowdsourcing tasks, trusted judgments acquisition process, as well as the evaluation methodology we employed.

3.1. Crowdsourcing Overview

Tables 1 and 2 present an overview of the crowdsourcing tasks, as well as the datasets used. The results of the crowdsourcing tasks were processed with the use of CrowdTruth metrics (Sec. 2.1), and we removed consistently low quality workers based on the spam removal procedure (Sec 2.2). The tasks were implemented and ran on Figure Eight[1] (formerly known as CrowdFlower). The templates are available on the CrowdTruth platform[2].

The payment per judgment was determined through a series of pilot runs of the tasks where we started with a $0.01 cost per judgment, and then gradually increased the payment until a majority of Figure Eight workers rated our tasks as having fair payments. As a result, we were able to get a constant stream of workers to participate in the tasks. The values shown in Table 2 show the final cost per judgment we reached after the pilot runs. Since crowd pay has a complex effect on the quality of the annotation [21], and in order to remove confounding factors, judgments collected with costs lower than those in Table 2 were left out of this

---

[1] https://figure-eight.com/
[2] tasks marked with ∗: https://github.com/CrowdTruth/CrowdTruth/wiki/Templates

Table 2
Crowdsourcing Task Data

Task	Source	Expert annotation	Media Units	Workers / Unit	Cost / Judgment
Medical Relation Extraction	PubMed article abstracts	yes	975	15	$0.05
Twitter Event Identification	Twitter (2014)	no	3,019	7	$0.02
News Event Extraction	TimeBank	yes	200	15	$0.02
Sound Interpretation	Freesound.org	yes	284	10	$0.01

Fig. 2. Templates of the Crowdsourcing Tasks

(a) Medical Relation Extraction

(b) Twitter Event Identification

(c) Sound Interpretation

(d) News Event Extraction

evaluation. In total, it took two months to perform the pilot runs and then collect the judgments for all of the tasks.

The number of workers per media unit was determined experimentally with the goal of capturing all possible results from the crowd and stabilizing the quality of the annotations; this process is explained at length further on in Section 4, with the results of the experiment shown in Figure 4.

The **Medical Relation Extraction dataset** consists of 975 sentences extracted from PubMed[3] article abstracts. The sentences were collected using distant supervision [22], a method that picks positive sentences from a corpus based on whether known arguments of the seed relation appear together in the sentence (*e.g.*, the *treat* relation occurs between the terms *antibiotics* and *typhus*, so find all sentences containing both and repeat this for all pairs of arguments that hold). The

---

[3]http://www.ncbi.nlm.nih.gov/pubmed

MetaMap parser [23] was used to extract medical terms from the corpus and the UMLS vocabulary [24] was used for mapping terms to categories, and relations to term types. The intuition of distant supervision is that since we know the terms are related, and they are in the same sentence, it is more likely that the sentence expresses a relation between them (than just any random sentence). We started with a set of 8 UMLS relations important for clinical decision making [25], that became the seed in distant supervision, but this paper only discusses results for the relations *cause* and *treat*, as these were the only relations for which we could also collect expert annotations. The expert judgment collection is detailed in Section 3.3.

The *medical relation extraction task* (see Figure 2a) is a *closed task*. The crowd is given a medical sentence with the two highlighted terms collected with distant supervision, and is then asked to select from a list all relations that are expressed between the two terms in the sentence. The relation list contains eight UMLS[4] relations, as well as *is a*, *part of*, *associated with*, *other*, *none* relations, added to make the choice list complete. Multiple choices are allowed in this task. To reduce the bias of selecting *none*, we also added an in-task effort consistency check by asking workers to explain in a text box why no relation is possible between the terms. The task results are processed into an annotation vector containing a component for each of the relations. A detailed description of the crowdsourcing data collection is given in [26].

The **Twitter Event Identification dataset** consists of 3,019 English tweets from 2014, crawled from Twitter. The tweets are selected as been relevant to eight events, such as, "Japan whale hunt", "China Vietnam relation" among other controversial events. The dataset was created by querying a Twitter dataset from 2014 with relevant phrases for each of the eight events, *e.g.*, "Whaling Hunting", "Anti-Chinese in Vietnam". The *Twitter event identification task* (see Figure 2b) is a *closed task*. The crowd is asked to choose for each tweet the relevant events out of the list of eight, as well as to highlight for each of the relevant events the event mentions in the tweet. The crowd could also pick that none of the events was present in the tweet. Multiple choices of events were permitted. Since tweets and tweet annotations typically are not done by experts, we did not collect expert data for this task. To reduce the bias of selecting no event, we also added an in-task

effort consistency check by asking workers to explain in a text box why none of the events is present in the tweet. The task results are processed into an annotation vector containing a component for each of the events.

The **News Event Extraction dataset** consists of 200 randomly selected English sentences from the English TimeBank corpora [27], which were also presented in [28]. The *news event extraction* (see Figure 2d) is an *open-ended task*. The crowd receives an English sentence, and is asked to highlight words or word phrases (multiple words) that describe an event or a time expression. For each sentence, the crowd is allowed to highlight a maximum of 30 event expressions or time expressions. For the purpose of this research we only focus on evaluating the extraction of event expressions. We define an *event* as something that happened, is happening, will or happen. On this dataset we employed expert annotators as described in Section 3.3. To reduce the bias of selecting fewer events than actually expressed in the task, we implemented an in-task effort consistency check by asking workers that annotated 3 events or less to explain in a text box why no other events are expressed in the sentence. As part of the *answer set dimensionality reduction step*, we removed the stop words from the sentence (we consider that the stop words are not meaningful for our analysis and they could add unsubstantial disagreement), and split the expressions collected from the crowd into words. The annotation vector is composed of the words in the sentence, where a word is selected in the worker vector if it appears in at least one of the expressions identified by the worker.

The **Sound Interpretation dataset** consists of 284 unique sounds sampled from the Freesound[5] online database. All these recordings and their metadata are freely accessible through the Freesound API[6]. We focused on SoundFX sounds, *i.e.*, sound effects category, as classified by [29]. The *Sound interpretation task* (see Figure 2c) is an *open-ended task*, where the crowd is asked to listen to three sounds and provide for each sound a comma separated list of keywords that best describe what they heard. For each sound, any number of answers is possible. In the *answer set dimensionality reduction step*, the annotated keywords were clustered syntacticly using spell checking and stemming, and semantically using a word2vec model [30] pre-trained on the Google News corpus. The annota-

---

[4]https://www.nlm.nih.gov/research/umls/

[5]https://www.freesound.org/
[6]https://www.freesound.org/docs/api/

Table 3

Consider an open-ended sound annotation task where 10 workers have to describe a given sound with keywords. The media unit for this task is a sound, the annotation set contains all the keywords workers provide for a sound. The table shows the media unit metrics, as well as the majority vote score for the media unit.

worker annotations	*dog barking*	*walking*	*animal*	*echo*	*loud*
media unit vector	3	2	5	1	1
UAS	0.47	0.31	0.79	0.15	0.15
majority vote	0	0	1	0	0

tion vector contains a component for each of the keywords used to describe the sound, after clustering. A detailed description of the crowdsourcing data collection and processing is given in [31]. For this dataset we also collected expert annotations from the sound creators as described in Section 3.3.

3.2. Evaluation Methodology

The purpose of the evaluation is to determine the quality of the annotations generated with CrowdTruth ambiguity-aware aggregating metrics. To this end, we label each media unit and annotation pair with its media unit-annotation score (see Section 2.1), and compare it with three other methods for labeling the data, as described below:

- **Majority vote**: Each media unit-annotation pair receives either a positive or a negative label, according to the decision of the majority of crowd workers. For each annotation performed by a crowd worker over a given media unit, we calculate the ratio of workers that have selected this annotation over the total number of workers that have annotated the unit, and assess whether it is greater or equal to 0.5. This allows for multiple annotations to be picked for one media unit. For some units, however, none of the annotations were picked by half or more of the workers. This is especially the case for open-ended tasks, such as sound interpretation, where workers put in a large number of annotations, and agreement is seldom. In these situations, we picked the annotations that were selected by the most workers (even if they do not constitute more than half). An example of the majority vote aggregation is shown in Table 3.
- **Single**: Each media unit-annotation pair receives either a positive or a negative label, according to the decision of a single crowd worker. For every media unit, this score was randomly sampled from the set of workers annotating it. Judgments

from workers labeled as spammers were not employed. While a single annotator is not used as often as the majority vote in traditional crowdsourcing, we use this dataset as a baseline for the crowd, to show that having more annotators generates better quality data.

- **Expert**: Each media unit-annotation pair receives either a positive or a negative label, according to the expert decision. The details of how expert data was collected for each tasks are discussed in Section 3.3.

The *evaluation of the quality of the CrowdTruth method* was done by computing the micro-F1 score over each task. The micro-F1 score was used in order to treat each case equally, without giving advantage to annotations that appear less frequently in our datasets. Using the trusted judgments collected according to Section 3.3, we evaluate each media unit – annotation pair as either a true positive, false positive etc. We compute the value of the micro-F1 score using the following formulas for the micro precision (Equation 1) and micro recall (Equation 2):

$$P_{micro} = \frac{\sum_{i=1}^{n} TP_i}{\sum_{i=1}^{n} TP_i + \sum_{i=1}^{n} FP_i} \qquad (1)$$

$$R_{micro} = \frac{\sum_{i=1}^{n} TP_i}{\sum_{i=1}^{n} TP_i + \sum_{i=1}^{n} FN_i} \qquad (2)$$

where TP_i, FP_i, FN_i, with i from 1 to n (the number of media units in the dataset), represent the number of true positive, false positive and false negative annotations for media unit i. Finally, the micro-F1 score is computed as the harmonic mean of the micro-precision and micro-recall.

An important variable in the evaluation is the *media unit-annotation score (UAS) threshold* for differentiating between a negative and a positive classification. Traditional crowdsourcing aims at reducing

disagreement, and therefore corresponds to high values for this threshold. Lower values means accepting more disagreement in the classification of positive answers by the crowd. In our experiments, we tried a range of threshold values for each task, to investigate with which one we achieve the best results. The UAS threshold was also used in gathering the set of trusted judgments for the evaluation (Section 3.3). All the data used in this paper can be found in our data repository[7].

3.3. Trusted Judgments Collection

To perform the evaluation, a set of trusted judgments is necessary to assess the correctness of crowd annotations. For each dataset, we manually evaluated the correctness of all the media unit annotations that were generated by the crowd and the experts. Depending on the task, the number of media unit-annotation pairs can become quite high, so we explored methods to make the manual evaluation more efficient.

For the datasets that contain expert annotation, we calculated the thresholds which yielded the maxi-

mum agreement in number of annotations between the crowd and expert annotations. These annotations were then added to the trusted judgments collection, as the judgment in this case is unambiguous. The interesting cases appear when crowd and expert disagree. Previous work we performed in crowdsourcing *Medical Relation Extraction* [32] has indicated that experts might not always provide better annotations than crowd workers. Additionally, for the *Sound Interpretation* task we noticed that experts provided considerably fewer tags than the crowd, and there was a large discrepancy between annotations of crowds and experts, with a very small overlap between their annotations. Therefore, instead of simply relying on expert judgment, the annotations where crowd and expert disagree were manually relabeled by exactly one of the authors, and then added to the trusted judgments set, which is also published in our data repository. In Appendix A we present a selection of examples where the expert judgment is different from the trusted judgment. While these cases might call into question the level of expertise of the domain experts, inconsistencies and disagreement in expert annotation are regularly reported in various annotation tasks [33–35]. Furthermore, in Section 4 we will show that using the

---

[7]https://github.com/CrowdTruth/Cross-Task-Majority-Vote-Eval

Fig. 3. CrowdTruth F1 scores for all crowdsourcing tasks.
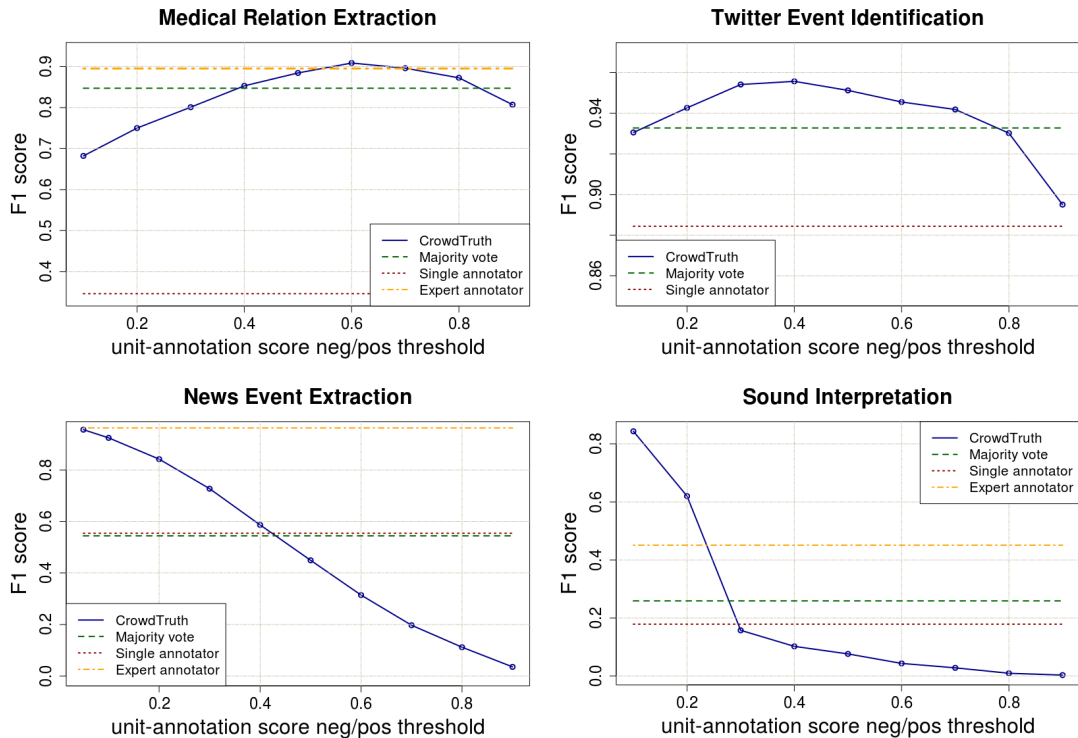
trusted judgments for evaluation still results in the expert performing the best for 2 out of 3 tasks. The only task where the expert underperforms is *Sound Interpretation*, where the set of annotations provided by the expert is much smaller than the one provided by the crowd.

We collected expert annotations for the *Medical Relation Extraction* data by employing medical students. Each sentence was annotated by exactly one person. The annotation task consisted of deciding whether or not the UMLS seed relation discovered by distant supervision is present in the sentence for the two selected terms.

For the *Sound Interpretation* task, each sound in the dataset contains a description and a set of keywords that were provided by the authors of the sounds. We consider the keywords provided by the sounds' authors as trusted judgments given by domain experts.

The *news event extraction* data was annotated with events by various linguistic experts. In total, 5 people annotated each sentence but we only have access to the final annotations, a consensus among the annotators. In the annotation guidelines described in [27], events are defined as situations that happen or occur, but are not generic situations. In contrast to the crowdsourcing task, where the workers had very loose instructions, the experts had very strict rules for identifying events, strictly based on linguistic features: *(i)* tensed verbs: has called, will leave, was captured, *(ii)* stative adjectives: sunken, stalled, on board and *(iii)* event nominals: merger, Military Operation, Gulf War.

The only task without expert annotation is *Twitter Event Identification* – as it is in the open domain, no experts exist for this type of data.

4. Results

We begin by evaluating **how the majority vote method compares with CrowdTruth**, by calculating the precision/recall metrics using the gold standards we collected for each of the four crowdsourcing tasks. Figure 3 shows the F1 score for CrowdTruth over the four tasks. The results are calculated for different UAS thresholds for separating the data points into positive and negative classifications. Table 4 shows the detailed scores for CrowdTruth, given the highest F1 UAS threshold.

Across all four tasks, the CrowdTruth method performs better than both majority vote and the single annotator dataset. While majority vote unsurprisingly

performs the best on precision, as a consequence of its lower rate of positive labels, CrowdTruth consistently scores the best for both recall, F1 score and accuracy. These differences in classification are statistically significant, as shown in Table 5 – this was calculated using McNemar's test [36] over paired nominal data.

The evaluation of CrowdTruth compared with the expert is more nuanced. For the *Medical Relation Extraction* and *news event extraction tasks*, CrowdTruth performs as well as the expert annotators, with p-values indicating there is no statistically significant difference in the classifications. In contrast, for the task of *Sound Interpretation*, CrowdTruth performs better than the expert by a large margin.

The second evaluation shows the **influence of the number of workers on the quality of the CrowdTruth data**. Figure 4 shows the CrowdTruth F1 score in relation to the number of workers. Given one task, the number of workers per unit varies because of spam removal, so the F1 score was calculated using at most the number of workers at every point in the graph. The number of units annotated with the given number of workers is also shown in the graph.

The effects of the number of workers on the CrowdTruth F1 is clear – more workers invariably leads to a higher F1 score. For the tasks of *Medical Relation Extraction*, *Twitter Event Identification* and *News Event Extraction*, the CrowdTruth F1 grows into a straight line, showing that the opinions of the crowd stabilize after enough workers. For the *Sound Interpretation* task, the CrowdTruth F1 score is still on an upwards trend after 10 workers, possibly indicating that more workers are necessary to get the full spectrum of annotations.

Figure 4 also shows that CrowdTruth performs better than majority vote regardless of the number of workers per task. For closed tasks, increasing the number of workers has a positive impact on the majority vote F1 score. For open tasks, adding more workers has less of an effect – more workers increase the size of the annotation set for a unit, which is typically larger than for closed tasks, but the agreement is low because opinions are split between possible annotations.

Finally, Figure 5 shows an evaluation of CrowdTruth using only the expert annotations as ground truth (the *Twitter Event Identification* task does not have experts, so it could not be evaluated). The F1 scores are lower than in the evaluation over the trusted judgments collection. For the *Medical Relation Extraction Task*, majority vote performs essentially the same as Crowd-

Table 4

CrowdTruth evaluation results, given the highest F1 media unit-annotation score (UAS) threshold.

Task	Dataset	Precision	Recall	F1 score	Accuracy	UAS threshold
Medical Relation Extraction	CrowdTruth	0.86	0.962	0.908	0.932	0.6
	expert	0.899	0.89	0.895	0.927	
	majority vote	0.924	0.781	0.847	0.902	
	single	0.222	0.776	0.346	0.748	
Twitter Event Identification	CrowdTruth	0.965	0.945	0.955	0.995	0.4
	majority vote	0.984	0.885	0.932	0.984	
	single	0.959	0.819	0.884	0.972	
News Event Extraction	CrowdTruth	0.984	0.929	0.956	0.931	0.05
	expert	0.983	0.944	0.963	0.942	
	majority vote	0.985	0.375	0.544	0.492	
	single	0.99	0.384	0.554	0.501	
Sound Interpretation	CrowdTruth	1	0.729	0.843	0.815	0.1
	expert	1	0.291	0.45	0.515	
	majority vote	1	0.148	0.258	0.418	
	single	1	0.098	0.178	0.383	

Table 5

p-values for McNemar's test of statistical significance in the CrowdTruth classification, compared with the others.

Task	Maj. Vote	Expert	Single
Medical Relation Extraction	0.0001	0.629	$< 2.2 \times 10^{-16}$
Twitter Event Identification	0.0001	N/A	6.145×10^{-15}
News Event Extraction	$< 2.2 \times 10^{-16}$	0.505	$< 2.2 \times 10^{-16}$
Sound Interpretation	$< 2.2 \times 10^{-16}$	$< 2.2 \times 10^{-16}$	$< 2.2 \times 10^{-16}$

Truth, whereas for the open-ended tasks, CrowdTruth still performs better. However, as we have shown in Appendix A, the expert annotations contain errors and are sometimes incomplete, particularly in the case of open-ended tasks. The evaluation using expert ground truth was done to show that the trusted judgments set is not biased in favor of CrowdTruth.

5. Discussion

The first goal in this paper was to show that the **ambiguity-aware CrowdTruth approach with multiple annotators and disagreement-based quality scores can perform better than majority vote**, a method that enforces consensus among annotators. Our results over several crowdsourcing tasks, as seen in Figure 3, show this clearly.

The gap in performance between CrowdTruth and majority vote is the most striking for open tasks (*News Event Extraction* and *Sound Interpretation*). These tasks also require the lowest agreement threshold for

achieving the best performance with CrowdTruth. During the trusted judgments collection process, we observed how these tasks are prone to a wide range of opinions – for instance, in the case of *Sound Interpretation*, there are frequent examples of labels that are semantically dissimilar, but could reasonably be applied to the same sound (e.g. the same sound was annotated with the tag `balloon popping` by one worker, and with `gunshot` by another worker). Because of this, enforcing consensus does not work for these tasks, and ambiguity-aware annotation aggregation appeared to be a viable solution.

Our evaluation also shows that processing crowd data with ambiguity-aware metrics performs at least as well as expert annotators, which is not the case for majority vote. Crowdsourcing annotation is significantly cheaper in cost than experts – e.g. even with 15 workers per unit, crowdsourcing for the task of *Medical Relation Extraction* cost 2/3 of what the experts did. The crowd also has the advantage of being readily available on platforms such as Figure Eight, while the process of finding and hiring expert annotators can incur

Fig. 4. The effect of the number of workers per unit on the F1 score, calculated at the best UAS threshold (Table 4). For every point, the F1 is calculated with at most the given number of workers. The number of units used in the calculation of the F1 is shown in the y-axis on the right.
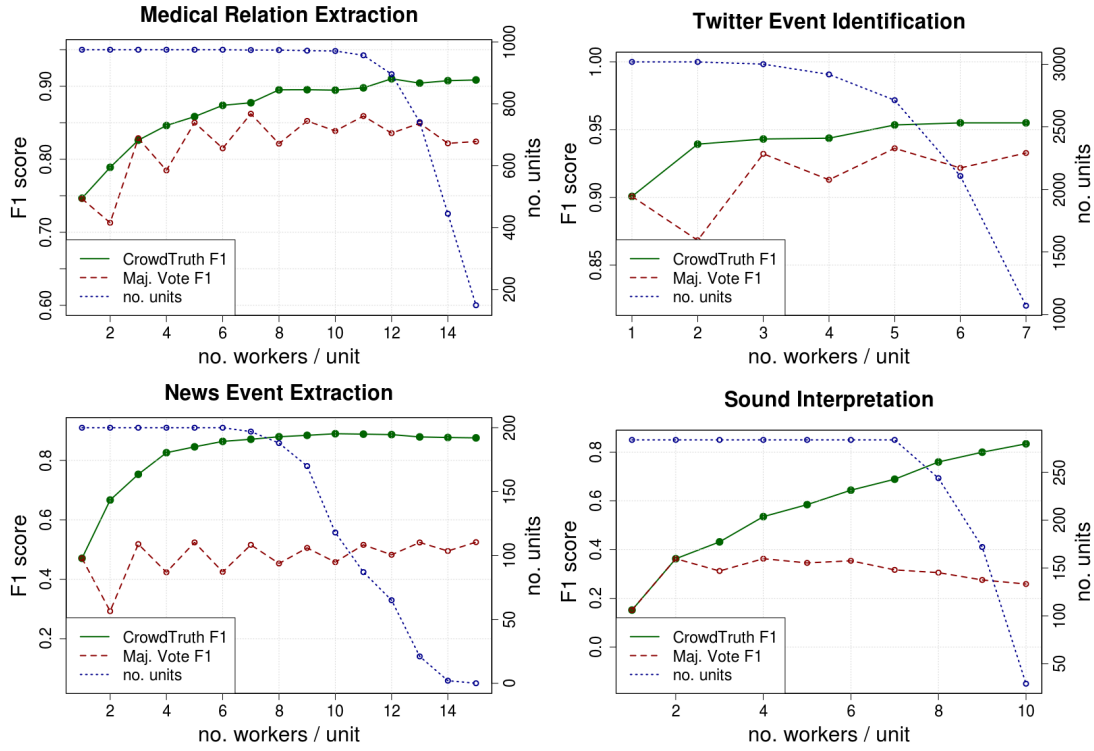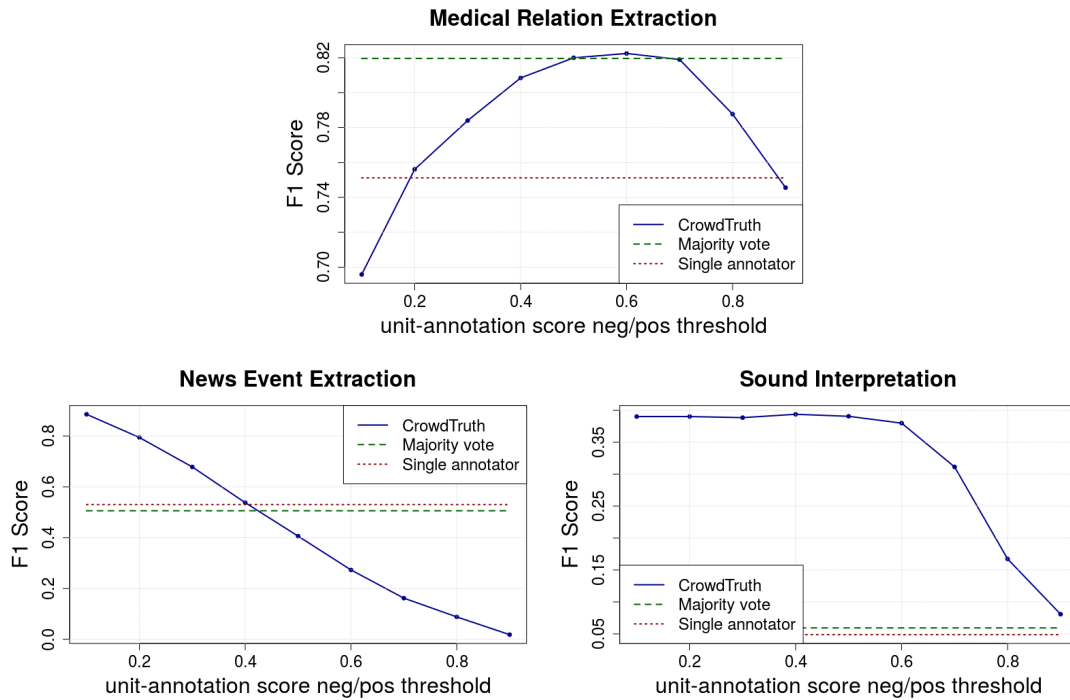

Fig. 5. CrowdTruth F1 score evaluation, using expert annotation as ground truth.

significant time costs. As our results showed, in order for the crowdsourcing to produce results comparable in quality to that of experts, appropriate processing with ambiguity-aware metrics is a necessity.

The variation in the optimal media unit-annotation score (UAS) thresholds across the tasks shows that the level of ambiguity is dependent on the crowdsourcing task, thus supporting our triangle of disagreement model (Section 2.1). It is not surprising that the task with the highest agreement threshold (*Medical Relation Extraction*) also has the most exact definition of a correct answer (i.e. whether a medical relation is expressed or not in a given sentence). The definition of a medical relation is fairly clear; in contrast, the definition of an event is more subjective, therefore workers were able to come up with a wider range of correct annotations.

The experimental setup provides an empirical method for selecting the optimal threshold for UAS. However, if performing an evaluation with trusted judgments is not possible, selecting the optimal threshold becomes more difficult. For open-ended tasks, the experiments indicate that almost all opinions matter, and the agreement threshold should be as low as possible. In these cases, spam workers can be successfully eliminated by in-task effort consistency checks, and there is no need to enforce agreement beyond that. In contrast, the experiments for closed tasks show higher agreement thresholds tend to work better. The difficulty as well as the subjectivity of the domain also appear to have an impact. The threshold should grow together with the difficulty, and inversely with subjectivity. However, both difficulty and subjectivity might be difficult to measure in practice. In the end, the tuning of the threshold should be regarded similarly to a precision-recall trade-off analysis, where the optimal value depends on the requirements of the ground truth (high precision but many false negative crowd labels, or high recall but more false positives). The high variability for optimal threshold values also shows the limitations of traditional evaluation metrics like precision and recall that rely on discrete labels. CrowdTruth metrics were constructed to measure ambiguity on a continuous scale, but the use of standard metrics resulted in losing this information by forcing the conversion to either positive or negative. Ultimately, our goal is to move away from a binary ground truth that needs to be calculated using a fixed threshold, and instead to use the CrowdTruth metrics to express ambiguity on a continuous scale.

The second goal of the experiment was to show **the effect of the number of workers on the quality of CrowdTruth annotations**. The results in Figure 4 clearly show the increase in F1 score for CrowdTruth as more workers contribute to the tasks. This combined with the poor performance of the single annotator dataset proves the importance in considering a large enough pool of workers to be able to accurately capture the full spectrum of opinions.

The stabilization of the F1 score for *Medical Relation Extraction*, *Twitter Event Identification* and *News Event Extraction* is an indication that we have indeed managed to collect the entire set of opinions for these tasks. The fact that the scores all stabilize at different points in the graph (around 8 workers for *Medical Relation Extraction*, 5 for *Twitter Event Identification*, and 10 for *News Event Extraction*) indicates that the optimal number of workers is dependent on the task type, thus also confirming our hypothesis that more workers than what is typically being considered in crowdsourcing studies are necessary for acquiring a high quality ground truth.

There exists a trade-off between cost and quality of annotations that should also be considered when optimizing the number of workers. The higher cost was justified for these tasks, as the expert annotation was three times more expensive than the crowdsourced annotations at expert quality level.

An interesting observation is that the optimal number of workers per task does not seem to influence the optimal UAS threshold for the task. The *News Event Extraction* requires a high number of workers, but the optimal UAS threshold is low, while the *Twitter Event Identification* requires a low number of workers, and also a low UAS threshold, at least compared to *Medical Relation Extraction*.

While four tasks is a small sample to draw conclusions from, our findings seem to indicate that ambiguity in the crowdsourcing system has an impact on both the optimal number of workers per task, as well as the clarity of the media units. These observations will form the basis for our future research in modeling crowd disagreement.

Finally, it is worth discussing the outlier characteristics of the *Sound Interpretation* task. It is the only task that does not achieve a stable F1 curve (Figure 4) possibly due to insufficient workers assigned to it. It is also unique in its lack of false positive examples – precision is 1 for the optimal UAS threshold (Table 4), meaning that all labels collected from the crowd were accepted as part of the trusted judgments. *Sound Inter-*

pretation is also the only task for which the expert annotator performed comparatively poor, with a statistically significant difference from CrowdTruth. As mentioned in the beginning of this section, after collecting the trusted judgments for this task, it became clear that the main challenge for the *Sound Interpretation* task is not to achieve consensus between annotators, but to collect the entire spectrum of annotations that describe a sound, given that this spectrum is so large (e.g. the tags `balloon popping` and `gunshot` can both reasonably apply to the same sound). For this reason, it was difficult to label tags as false positives, and the annotations of the workers, experts included, were largely non-overlapping, as they tended to interpret the sounds quite differently. The *Sound Interpretation* task is therefore an extreme example of subjective ground truth.

6. Related Work

6.1. Crowdsourcing Ground Truth

Crowdsourcing has grown into a viable alternative to expert ground truth collection, as crowdsourcing tends to be both cheaper and more readily available than domain experts. Experiments have been carried out in a variety of tasks and domains: medical entity extraction [37–39], medical relation extraction [39, 40], open-domain relation extraction [41], clustering and disambiguation [42], ontology evaluation [6], web resource classification [43] and taxonomy creation [44]. [45] have shown that aggregating the answers of an increasing number of unskilled crowd workers with majority vote can lead to high quality NLP training data. The typical approach in these works is to assume the existence of a universal ground truth. Therefore, disagreement between annotators is considered an undesirable feature, and is usually discarded by using either of the following methods: restricting annotator guidelines, picking one answer that reflects some consensus usually through majority voting, or using a small number of annotators.

6.2. Disagreement and Ambiguity in Crowdsourcing

Besides CrowdTruth, there exists some research on how disagreement in crowdsourcing should be interpreted and handled. In assessing the OAEI benchmark, [33] found that disagreement between annotators (both crowd and expert) is an indicator for in-

herent uncertainty in the domain knowledge, and that current benchmarks in ontology alignment and evaluation are not designed to model this uncertainty. [46] found similar results for the task of crowdsourced part-of-speech tagging – most inter-annotator disagreement was indicative of debatable cases in linguistic theory, rather than faulty annotation. [47] also investigate the role of inter-annotator disagreement as a possible indicator of ambiguity inherent in natural language. [48] propose a method for crowdsourcing ambiguity in the grammatical correctness of text by giving workers the possibility to pick various degrees of correctness, but inter-annotator disagreement is not discussed as a factor in measuring this ambiguity. [13] propose a framework for dealing with uncertainty in ground truth that acknowledges the notion of ambiguity, and uses disagreement in crowdsourcing for modeling this ambiguity. For the task of word sense disambiguation, [49] show that, in modeling ambiguity, the crowd was able to achieve expert-level quality of annotations. [50] implemented a workflow of tasks for collecting and correcting labels for text and images, and found that ambiguous cases cannot simply be resolved by better annotation guidelines or through worker quality control. Finally, [51] shows that often, machine learning classifiers can achieve a higher accuracy when trained with noisy crowdsourcing data. To our knowledge, our paper presents the first experiment across several tasks and domains that explores ambiguity as a property of crowdsourcing systems, and how it can be interpreted to improve the quality of ground truth data.

6.3. Crowdsourcing Aggregation beyond Majority Vote

The literature on alternative crowdsourcing aggregation metrics typically focuses on analyzing worker performance – identifying spam workers [52–54], and analyzing workers' performance for quality control and optimization of the crowdsourcing processes [55]. [56] and [57] have used a latent variable model for task difficulty, as well as latent variables to measure the skill of each annotator, to optimize crowdsourcing for image labels. [58] use on-the-job learning with Bayesian decision theory to assign the most appropriate workers for each task, for both text and image annotation. Finally, [59] show that the surprisingly popular crowd choice (i.e. the answer that most workers thought would not be picked by other workers, even though it is correct) gave better results than the majority vote for a variety of tasks with unambiguous ground

truths (state capitals, trivia questions and price of artworks).

All of these approaches show promising improvements over the use of majority vote as an aggregating method. These methods were developed only for closed tasks, primarily dealing with classification. However, the novel approach of CrowdTruth allows to explore both closed and open-ended tasks. Furthermore, our focus is on modeling ambiguity as a latent variable in the crowdsourcing system, as well as its role in generating inter-annotator disagreement, which these approaches currently do not take into account. We believe an optimal crowdsourcing approach would combine both ambiguity modeling, as well as specialized task assignment to workers. For instance, [60] developed a generative model to aggregate crowd scores that incorporates features of the data (e.g. number of words), although they do not evaluate the performance of specific features. Ambiguity as measured with CrowdTruth, like the media unit-annotation score, could be used as a data feature in such a system.

7. Conclusions

Gathering human annotation is a major bottleneck in the process of knowledge base curation. Crowdsourcing-based approaches are gaining popularity in the attempt to solve the issues related to volume of data and lack of annotators. Typically these practices use inter-annotator agreement as a measure of quality. However, by ignoring inter-annotator disagreement, these practices tend to create artificial data that is neither general nor reflects the ambiguity inherent in the source.

In this paper we presented an empirically derived methodology for efficiently gathering of human annotation by aggregating crowdsourcing data with CrowdTruth metrics, which harness the inter-annotator disagreement. We applied this methodology over a set of diverse crowdsourcing tasks: closed tasks (*Medical Relation Extraction*, *Twitter Event Identification*), and open-ended tasks (*News Event Extraction* and *Sound Interpretation*). Our results showed that the ambiguity-aware CrowdTruth approach allows us to collect richer data, which enables reasoning about the ambiguity of the content being annotated. This is intrinsically relevant to the Semantic Web community, i.e. to identify the semantics of ambiguity across all modalities, e.g. text, images, videos and sounds. Our results also showed that, in all the tasks we consid-

ered, such ambiguity-aware quality scores provide better ground truth data than the traditional majority vote. Moreover, we have shown that CrowdTruth annotations have at least the same quality, even better in the case of *Sound Interpretation*, as expert annotations. Finally, we showed that, contrary to the common crowdsourcing practice of employing a small number of annotators, adding more crowd workers actually can lead to significantly better annotation quality.

In the future, we plan to expand our methodology to more complex annotation tasks, that require multiple or combined types of input beyond the closed/open-ended categorization we presented in this paper. We are also working on expanding the CrowdTruth metrics for ambiguity to incorporate the state-of-the art in modeling crowd worker and data features [60]. Finally, we want to use the CrowdTruth data in practice for training and evaluating information extraction models used to populate the Semantic Web.

Acknowledgements

References

[1] C. Sarasua, E. Simperl, N. Noy, A. Bernstein and J.M. Leimeister, Crowdsourcing and the Semantic Web: A research manifesto, *Human Computation (HCOMP)* **2**(1) (2015), 3–17. DOI:10.15346/hc.v2i1.2.

[2] C. Welty, K. Barker, L. Aroyo and S. Arora, Query driven hypothesis generation for answering queries over NLP graphs, in: *The Semantic Web–ISWC 2012*, Springer, 2012, pp. 228–242. DOI:10.1007/978-3-642-35173-0_15.

[3] L. Aroyo and C. Welty, Harnessing disagreement for event semantics, in: *Proceedings of the 2nd International Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web (DeRiVE 2012), 11th International Semantic Web Conference*, 2012, p. 31.

[4] J. Oosterman, A. Nottamkandath, C. Dijkshoorn, A. Bozzon, G.-J. Houben and L. Aroyo, Crowdsourcing knowledge-intensive tasks in cultural heritage, in: *Proceedings of the 2014 ACM conference on Web science*, ACM, 2014, pp. 267–268. DOI:10.1145/2615569.2615644.

[5] L. Aroyo and C. Welty, Measuring crowd truth for medical relation extraction, in: *AAAI 2013 Fall Symposium on Semantics for Big Data*, 2013.

[6] N.F. Noy, J. Mortensen, M.A. Musen and P.R. Alexander, Mechanical turk as an ontology engineer? Using microtasks as a component of an ontology-engineering workflow, in: *Proceedings of the 5th Annual ACM Web Science Conference*, ACM, 2013, pp. 262–271. DOI:10.1145/2464464.2464482.

[7] L. Von Ahn, Human computation, in: *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE*, IEEE, 2009, pp. 418–419. DOI:10.1145/1629911.1630023.

[8] J. Carletta, Assessing agreement on classification tasks: the Kappa statistic, *Comput. Linguist.* **22**(2) (1996), 249–254. http://dl.acm.org/citation.cfm?id=230386.230390.

[9] S. Nowak and S. Rüger, How reliable are annotations via crowdsourcing: a study about inter-annotator agreement for multi-label image annotation, in: *Proceedings of the international conference on Multimedia information retrieval*, ACM, 2010, pp. 557–566. DOI:10.1145/1743384.1743478.

[10] M. Poesio and R. Artstein, The reliability of anaphoric annotation, reconsidered: Taking ambiguity into account, in: *Proceedings of the workshop on frontiers in corpus annotations II: Pie in the sky*, Association for Computational Linguistics, 2005, pp. 76–83. DOI:10.3115/1608829.1608840.

[11] N. Chang, R. Lee-Goldman and M. Tseng, Linguistic wisdom from the crowd, in: *Third AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2016. https://www.aaai.org/ocs/index.php/HCOMP/HCOMP15/paper/view/11737.

[12] A. Flexer and T. Grill, The problem of limited inter-rater agreement in modelling music similarity, *Journal of New Music Research* **45**(3) (2016), 239–251, PMID: 28190932. DOI:10.1080/09298215.2016.1200631.

[13] M. Schaekermann, E. Law, A.C. Williams and W. Callaghan, Resolvable vs. irresolvable ambiguity: a new hybrid framework for dealing with uncertain ground truth, in: *1st Workshop on Human-Centered Machine Learning at SIGCHI 2016*, 2016.

[14] V. Cheplygina and J.P. Pluim, Crowd disagreement of medical images is informative, *arXiv preprint arXiv:1806.08174* (2018).

[15] L. Aroyo, A. Dumitrache, P. Paritosh, A. Quinn and C. Welty, Reports of the workshops held at the sixth AAAI conference on human computation and crowdsourcing, *AI Magazine* (2018). DOI:10.1609/aimag.v39i4.2834.

[16] L. Aroyo and C. Welty, Crowd Truth: Harnessing disagreement in crowdsourcing a relation extraction gold standard, in: *Proceedings of the 5th Annual ACM Web Science Conference*, 2013.

[17] O. Inel, K. Khamkham, T. Cristea, A. Dumitrache, A. Rutjes, J. van der Ploeg, L. Romaszko, L. Aroyo and R.-J. Sips, CrowdTruth: machine-human computation framework for harnessing disagreement in gathering annotated data, in: *The Semantic Web – ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II*, P. Mika, T. Tudorache, A. Bernstein, C. Welty, C. Knoblock, D. Vrandečić, P. Groth, N. Noy, K. Janowicz and C. Goble, eds, Springer International Publishing, Cham, 2014, pp. 486–504. DOI:10.1007/978-3-319-11915-1_31.

[18] J.Q. Knowlton, On the definition of "picture", *AV Communication Review* **14**(2) (1966), 157–183.

[19] L. Aroyo and C. Welty, The Three Sides of Crowd-Truth, *Journal of Human Computation* **1** (2014), 31–34. DOI:10.15346/hc.v1i1.3.

[20] G. Soberón, L. Aroyo, C. Welty, O. Inel, H. Lin and M. Overmeen, Measuring Crowd Truth: Disagreement Metrics Combined with Worker Behavior Filters, in: *1st International Workshop on Crowdsourcing the Semantic Web, 12th International Semantic Web Conference*, 2013.

[21] A. Mao, E. Kamar, Y. Chen, E. Horvitz, M.E. Schwamb, C.J. Lintott and A.M. Smith, Volunteering versus work for pay: Incentives and tradeoffs in crowdsourcing, in: *First AAAI conference on human computation and crowdsourcing (HCOMP)*, 2013. https://www.aaai.org/ocs/index.php/HCOMP/HCOMP13/paper/view/7497.

[22] M. Mintz, S. Bills, R. Snow and D. Jurafsky, Distant supervision for relation extraction without labeled data, in: *Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2*, Association for Computational Linguistics, 2009, pp. 1003–1011. https://www.aclweb.org/anthology/P09-1113.

[23] A.R. Aronson, Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program, in: *Proceedings of the AMIA Symposium*, American Medical Informatics Association, 2001, p. 17, PMID: 11825149. https://www.ncbi.nlm.nih.gov/pubmed/11825149.

[24] O. Bodenreider, The unified medical language system (UMLS): integrating biomedical terminology, *Nucleic acids research* **32**(suppl 1) (2004), D267–D270. DOI:10.1093/nar/gkh061.

[25] C. Wang and J. Fan, Medical relation extraction with manifold models, in: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, The Association for Computer Linguistics, 2014, pp. 828–838. ISBN 978-1-937284-72-5. DOI:10.3115/v1/p14-1078. http://aclweb.org/anthology/P/P14/P14-1078.pdf.

[26] A. Dumitrache, L. Aroyo and C. Welty, Crowdsourcing ground truth for medical relation extraction, *ACM Transactions on Interactive Intelligent Systems (TiiS)* **8**(2) (2018), 11:1–11:20. DOI:10.1145/3152889.

[27] J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizauskas, A. Setzer, D. Radev, B. Sundheim, D. Day, L. Ferro et al., The TimeBank corpus, in: *Corpus linguistics*, 2003, p. 40. http://ucrel.lancs.ac.uk/publications/cl2003/papers/pustejovsky.pdf.

[28] T. Caselli, R. Sprugnoli and O. Inel, Temporal information annotation: crowd vs. experts, in: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, N.C.C. Chair), K. Choukri, T. Declerck, S. Goggi, M. Grobelnik, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk and S. Piperidis, eds, European Language Resources Association (ELRA), 2016. ISBN 978-2-9517408-9-1.

[29] F. Font, J. Serrà and X. Serra, Audio clip classification using social tags and the effect of tag expansion, in: *Audio Engineering Society 53rd International Conference on Semantic Audio*, Audio Engineering Society, 2014.

[30] T. Mikolov, I. Sutskever, K. Chen, G. Corrado and J. Dean, Distributed representations of words and phrases and their compositionality, in: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Vol-*

ume 2, NIPS'13, Curran Associates Inc., USA, 2013, pp. 3111–3119. http://dl.acm.org/citation.cfm?id=2999792.2999959.

[31] L.A. Emiel van Miltenburg Benjamin Timmermans, The VU Sound Corpus: Adding More Fine-grained Annotations to the Freesound Database, in: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, N.C.C. Chair), K. Choukri, T. Declerck, S. Goggi, M. Grobelnik, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk and S. Piperidis, eds, European Language Resources Association (ELRA), 2016. ISBN 978-2-9517408-9-1.

[32] L. Aroyo and C. Welty, Truth is a lie: Crowd Truth and the seven myths of human annotation, *AI Magazine* **36**(1) (2015), 15–24. DOI:10.1609/aimag.v36i1.2564.

[33] M. Cheatham and P. Hitzler, Conference v2.0: An Uncertain Version of the OAEI Conference Benchmark, in: *The Semantic Web – ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II*, P. Mika, T. Tudorache, A. Bernstein, C. Welty, C. Knoblock, D. Vrandečić, P. Groth, N. Noy, K. Janowicz and C. Goble, eds, Springer International Publishing, Cham, 2014, pp. 33–48. ISBN 978-3-319-11915-1. DOI:10.1007/978-3-319-11915-1_3.

[34] T. McDonnell, M. Lease, M. Kutlu and T. Elsayed, Why is that relevant? Collecting annotator rationales for relevance judgments, in: *Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2016. http://aaai.org/ocs/index.php/HCOMP/HCOMP16/paper/view/14043.

[35] O. Inel, T. Caselli and L. Aroyo, Crowdsourcing Salient Information from News and Tweets, in: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, N.C.C. Chair), K. Choukri, T. Declerck, S. Goggi, M. Grobelnik, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk and S. Piperidis, eds, European Language Resources Association (ELRA), 2016. ISBN 978-2-9517408-9-1.

[36] Q. McNemar, Note on the sampling error of the difference between correlated proportions or percentages, *Psychometrika* **12**(2) (1947), 153–157. DOI:10.1007/bf02295996.

[37] H. Zhai, T. Lingren, L. Deleger, Q. Li, M. Kaiser, L. Stoutenborough and I. Solti, Web 2.0-based crowdsourcing for high-quality gold standard development in clinical natural language processing, *Journal of medical Internet research* **15**(4) (2013). DOI:10.2196/jmir.2426.

[38] T. Finin, W. Murnane, A. Karandikar, N. Keller, J. Martineau and M. Dredze, Annotating named entities in Twitter data with crowdsourcing, in: *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, CSLDAMT '10, Association for Computational Linguistics, 2010, pp. 80–88. http://dl.acm.org/citation.cfm?id=1866696.1866709.

[39] E.M. Van Mulligen, A. Fourrier-Reglat, D. Gurwitz, M. Molokhia, A. Nieto, G. Trifiro, J.A. Kors and L.I. Furlong, The EU-ADR corpus: annotated drugs, diseases, targets, and their relationships, *Journal of biomedical informatics* **45**(5) (2012), 879–884. DOI:10.1016/j.jbi.2012.04.004.

[40] H. Kilicoglu, G. Rosemblat, M. Fiszman and T.C. Rindflesch, Constructing a semantic predication gold standard from the biomedical literature, *BMC bioinformatics* **12**(1) (2011), 486. DOI:10.1186/1471-2105-12-486.

[41] S.K. Kondreddi, P. Triantafillou and G. Weikum, Combining information extraction and human computing for crowdsourced knowledge acquisition, in: *30th International Conference on Data Engineering*, IEEE, 2014, pp. 988–999. DOI:10.1109/icde.2014.6816717.

[42] J. Lee, H. Cho, J.-W. Park, Y.-r. Cha, S.-w. Hwang, Z. Nie and J.-R. Wen, Hybrid entity clustering using crowds and data, *The VLDB Journal* **22**(5) (2013), 711–726. DOI:10.1007/s00778-013-0328-8.

[43] S. Castano, A. Ferrara and S. Montanelli, Human-in-the-loop web resource classification, in: *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, Springer, 2016, pp. 229–244. DOI:10.1007/978-3-319-48472-3_13.

[44] J. Bragg, D.S. Weld et al., Crowdsourcing multi-label classification for taxonomy creation, in: *First AAAI conference on human computation and crowdsourcing (HCOMP)*, 2013. http://www.aaai.org/ocs/index.php/HCOMP/HCOMP13/paper/view/7560.

[45] R. Snow, B. O'Connor, D. Jurafsky and A.Y. Ng, 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*, EMNLP '08, Association for Computational Linguistics, 2008, pp. 254–263. DOI:10.3115/1613715.1613751. http://dl.acm.org/citation.cfm?id=1613715.1613751.

[46] B. Plank, D. Hovy and A. Søgaard, Linguistically debatable or just plain wrong?, in: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Association for Computational Linguistics, Baltimore, Maryland, 2014, pp. 507–511. DOI:10.3115/v1/p14-2083. http://www.aclweb.org/anthology/P14/P14-2083.

[47] P.S. Bayerl and K.I. Paul, What determines inter-coder agreement in manual annotations? A meta-analytic investigation, *Comput. Linguist.* **37**(4) (2011), 699–725. DOI:10.1162/COLI_a_00074.

[48] J.H. Lau, A. Clark and S. Lappin, Measuring gradience in speakers' grammaticality judgements, in: *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 2014, pp. 821–826.

[49] D. Jurgens, Embracing ambiguity: a comparison of annotation methodologies for crowdsourcing word sense labels, in: *HLT-NAACL*, 2013, pp. 556–562. https://www.aclweb.org/anthology/N13-1062.

[50] J.C. Chang, S. Amershi and E. Kamar, Revolt: collaborative crowdsourcing for labeling machine learning datasets, in: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, ACM, New York, NY, USA, 2017. DOI:10.1145/3025453.3026044.

[51] C.H. Lin, D.S. Weld et al., To re (label), or not to re (label), in: *Second AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2014. http://www.aaai.org/ocs/index.php/HCOMP/HCOMP14/paper/view/8978.

[52] A. Bozzon, M. Brambilla, S. Ceri and A. Mauri, Reactive crowdsourcing, in: *Proceedings of the 22nd international conference on World Wide Web*, WWW '13, International World Wide Web Conferences Steering Committee, 2013, pp. 153–164. ISBN 978-1-4503-2035-1. DOI:10.1145/2488388.2488403. http://dl.acm.org/citation.cfm?id=2488388.2488403.

[53] A. Kittur, E.H. Chi and B. Suh, Crowdsourcing user studies with Mechanical Turk, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, ACM, 2008, pp. 453–456. ISBN 978-1-60558-011-1. DOI:10.1145/1357054.1357127.

[54] P.G. Ipeirotis, F. Provost and J. Wang, Quality management on Amazon Mechanical Turk, in: *Proceedings of the ACM SIGKDD Workshop on Human Computation*, HCOMP '10, ACM, 2010, pp. 64–67. ISBN 978-1-4503-0222-7. DOI:10.1145/1837885.1837906.

[55] Y. Singer and M. Mittal, Pricing mechanisms for crowdsourcing markets, in: *Proceedings of the 22nd international conference on World Wide Web*, WWW '13, International World Wide Web Conferences Steering Committee, 2013, pp. 1157–1166. ISBN 978-1-4503-2035-1. DOI:10.1145/2488388.2488489. http://dl.acm.org/citation.cfm?id=2488388.2488489.

[56] J. Whitehill, T.-f. Wu, J. Bergsma, J.R. Movellan and P.L. Ruvolo, Whose vote should count more: optimal integration of labels from labelers of unknown expertise, in: *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J.D. Lafferty, C.K.I. Williams and A. Culotta, eds, Curran Associates, Inc., 2009, pp. 2035–2043. http://papers.nips.cc/paper/3644-whose-vote-should-count-more-optimal-integration-of-labels-from-labelers-of-unknown-expertise.pdf.

[57] P. Welinder, S. Branson, P. Perona and S.J. Belongie, The multidimensional wisdom of crowds, in: *Advances in neural information processing systems*, 2010, pp. 2424–2432. http://papers.nips.cc/paper/4074-the-multidimensional-wisdom-of-crowds.

[58] K. Werling, A.T. Chaganty, P.S. Liang and C.D. Manning, On-the-job learning with bayesian decision theory, in: *Advances in Neural Information Processing Systems*, 2015, pp. 3465–3473.

[59] D. Prelec, H.S. Seung and J. McCoy, A solution to the single-question crowd wisdom problem, *Nature* **541**(7638) (2017), 532–535. DOI:10.1038/nature21054.

[60] P. Felt, K. Black, E.K. Ringger, K.D. Seppi and R. Haertel, Early gains matter: a case for preferring generative over discriminative crowdsourcing models, in: *HLT-NAACL*, 2015, pp. 882–891. DOI:10.3115/v1/n15-1089.

Appendix A. Example Media Units Where the Expert Judgment Is Different from the Trusted Judgment

Table 6

Example sentences from the *medical relation extraction* task where the expert judgment is different from the trusted judgment. The pair of terms that express the medical relation are shown in italic font in the media unit.

Media Unit	Annotation	Expert Judgment	Crowd Score	Trusted Judgment
The *epidermal nevus syndrome* is a neurocutaneous disorder characterized by *distinctive skin lesions* and often serious somatic and central nervous system (CNS) abnormalities.	*cause*	no	0.98	yes
For empiric *treat*ment of epididymitis, especially when gonococcal or *chlamydial infection* is likely Ofloxacin or *levofloxacin* should be used only if epididymitis is not *cause*d by gonorrhea.	*treat*	no	0.966	yes
In contrast, we did not find a definite increase in the LGL percentage within 6 months postpartum in patients with *Graves' disease* who relapsed into *Graves' thyrotoxicosis*.	*cause*	no	0.738	yes
The 1 placebo controlled trial that found black cohosh to be effective for *hot flashes* did not find *estrogen* to be effective, which casts doubt on the study's validity.	*treat*	no	0.73	yes
Multicentric reticulohistiocytosis (MR) is a *systemic disease* of unknown *cause* characterized by the presence of a heavy macrophage infiltrate in skin and synovial tissues and the development of an erosive polyarthritis.	*cause*	yes	0.697	no
Urokise versus *tissue plasminogen activator* in *pulmonary embolism*.	*treat*	yes	0.365	no
The principal differences between these vaccines are the transmission of live vaccine viruses from recipients to their contacts and the occurrence of occasional cases of *paralytic poliomyelitis* associated with use of *live poliovirus vaccine*	*treat*	yes	0.1	no
These cases highlight the importance of considering *PTLD* in the differential diagnosis of *lymphadenopathy*.	*cause*	yes	0.09	no

Table 7

Example sentences from the *news event extraction* task where the expert judgment is different from the trusted judgment. The annotation is shown in italic font in the media unit.

Media Unit	Annotation	Expert Judgment	Crowd Score	Trusted Judgment
The plan provides for the *distribution* of one common stock-purchase right as a dividend for each share of common outstanding	*distribution*	no	0.95	yes
Two Middle East terrorists with records of successful *attacks* against Western targets Abu Nidal and Abu Abbas have ties to Baghdad.	*attacks*	no	0.73	yes
Secretary of State James Baker said on ABC-TV's "This Week With David Brinkley" that the series of UN resolutions condemning Iraq's *invasion* of Kuwait "imply that the restoration of peace and stability in the Gulf would be a heck of a lot easier if he and that leadership were not in power in Iraq."	*invasion*	no	0.53	yes
The company also said it continues to explore all options concerning the possible *sale* of National Aluminum's 54.5% stake in an aluminum smelter in Hawesville Ky.	*sale*	no	0.24	yes
Yield on the issue was 7.88%	*no event*	yes	0.14	no
Har-Shefi said she heard Amir talk about killing Rabin but did not tell the police because she did not believe he was *serious*.	*serious*	yes	0	no
The American hope is that someone from within Iraq perhaps from the army 's professional ranks will step forward and push Saddam Hussein aside so that the country can begin recovering from the disaster.	*no event*	yes	0	no

Dumitrache et al. / Empirical Methodology for Crowdsourcing Ground Truth

Table 8

Example sounds from the *sound interpretation* task where the expert judgment is different from the trusted judgment.

Media Unit URL	Media Unit Description	Annotation	Expert Judgment	Crowd Score	Trusted Judgment
https://freesound.org/data/previews/21/21266_88803-hq.mp3	jazz	cymbals	no	0.272	yes
		bangle	no	0.136	yes
		rhythmic	no	0.136	yes
https://freesound.org/data/previews/26/26086_11477-hq.mp3	chicken	birds	no	0.538	yes
		geese	no	0.359	yes
		horns	no	0.359	yes
https://freesound.org/data/previews/35/35823_317782-hq.mp3	weird drums	music	no	0.875	yes
		band	no	0.145	yes
		disco	no	0.145	yes
https://freesound.org/data/previews/39/39329_404624-hq.mp3	trip hop	beat	no	0.371	yes
		percussion	no	0.371	yes
		chimes	no	0.371	yes
https://freesound.org/data/previews/41/41462_78779-hq.mp3	beer glasses	clicks	no	0.242	yes
		clink	no	0.242	yes
		ding	no	0.242	yes