

Explainable Agreement through Simulation for Tasks with Subjective Labels

John Foley

University of Massachusetts Amherst
jfoley@cs.umass.edu

ABSTRACT

The field of information retrieval often works with limited and noisy data in an attempt to classify documents into subjective categories, e.g., relevance, sentiment and controversy. We typically quantify a notion of agreement to understand the difficulty of the labeling task, but when we present final results, we do so using measures that are unaware of agreement or the inherent subjectivity of the task.

We propose using user simulation to understand the effect size of this noisy agreement data. By simulating truth and predictions, we can understand the maximum scores a dataset can support: for if a classifier is doing better than a reasonable model of a human, we cannot conclude that it is actually better, but that it may be learning noise present in the dataset.

We present a brief case study on controversy detection that concludes that a commonly-used dataset has been exhausted: in order to advance the state-of-the-art, more data must be gathered at the current level of label agreement in order to distinguish between techniques with confidence.

ACM Reference Format:

John Foley. 2018. Explainable Agreement through Simulation for Tasks with Subjective Labels. In *Proceedings of ACM SIGIR 2018 Workshop on Learning from Limited or Noisy Data*, Ann Arbor, Michigan, USA, July 12, 2018 (LND4IR '18), 3 pages.
<https://doi.org/>

1 INTRODUCTION

Although we often work with user-labeled data, and we can quantify inter-annotator agreement – we often do not know what these agreement scores mean in terms of our target measures, like mean average precision (mAP), or F1. What does an agreement of 70% allow for us to understand on this dataset? What level of mAP is believable? What is the best a system can do here?

Instead of an agreement probability or a statistical measure of agreement, what we really want is to be able to quantify the effect our level of agreement is having on our ability to evaluate systems. Therefore, what we really need is a technique for exploring what our agreement means in any arbitrary evaluation measure.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LND4IR '18, July 12, 2018, Ann Arbor, Michigan, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN .

<https://doi.org/>

While existing agreement measures give us a sense of how difficult the task was for annotators, it is hard to quantify what that means for measures, particularly those involving ranking. Maybe annotators agree on the most important instances, or maybe they disagree on the most critical instances – the same agreement score may lead to very different reliabilities in results.

Controversy is a problem that has attracted a lot of attention in recent years [5–8, 14]. Much like relevance, sentiment, and other labels of interest, it is both somewhat subjective (noisy) and expensive to collect (limited). In this study, we will look at the dataset of 343 pages collected by Dori-Hacohen and Allan [5] and used in further studies [6, 8].

We find that the language-modeling approaches introduced by Jang et al. [8] effectively “max-out” this dataset, as their language modeling classifier achieves statistically indistinguishable performance from our human-model simulations. This means that given the limited dataset size and the inherent disagreement between annotators on which documents are controversial – there are no AUC scores higher than currently published results that we should believe without collection of additional labels.

In this paper, we introduce a simulation technique that will allow this analysis to be performed on any dataset with a set of annotator labels for any arbitrary measure or metric.

2 RELATED WORK

The effect of the subjectivity and difficulty of relevance on IR evaluation has long been studied [2–4, 10, 12, 13]. While these studies look at the robustness of measures in the face of this subjectivity and noise, they do not quantify how well a system can do in comparison to humans – probably because IR systems rarely retrieve otherwise perfect rankings.

As agreement measures can be used to evaluate classification tasks directly, studies connecting the two are often looking at the suitability of an agreement score for classifier evaluation, e.g. [1].

Simulating users of IR or ML systems is also not a new contribution (e.g., [11]) and recently work has begun to accelerate in this direction [9]. However, we are unaware of work that simulates users in order to understand the limitations of agreement for a dataset.

3 TRUTH SIMULATION MODELS

In this section, we introduce a number of models for deriving truth from a set of labels for a document.

Given a document D which has a set of labels $L = \{l_1, l_2, \dots, l_{|L|}\}$, with each label l_i being provided by a different annotator, most studies choose a simple heuristic function $f(L)$ that generates a single label from the set.

Our models are applicable to both binary and multiclass judgments, provided that functions $f(L)$ map from a set of labels to a valid label. Since we look at controversy, we focus on ordinal labels, and we can use fractional labels as predictions, but not truth.

3.1 Average and Max Models

In prior work [5, 6, 8], the assignment given to a document is the average of its labels. Another appropriate model for controversy we consider in addition to the average modeling done here is a maximum model: i.e., a document is controversial if any annotator considered it controversial – a policy aimed at maximizing recall.

3.2 Agreement-Flip Model

Here we let p be the probability of agreement calculated across the dataset. We could argue that with probability $1 - p$, a label will be disputed and therefore is possibly incorrect with this probability. This is a fairly simple model of agreement, and the one that is represented by presenting agreement ratio in papers.

3.3 Label Sampling Model

A better model takes document-level confusion into account: if a document garners a variety of labels, we consider these observations of the underlying distribution for that document. Here, our f samples a label at random from a document.

3.4 Label Conflation Model

In a world where there are multi-value relevance labels, and you have many documents with only a single annotator e.g., Excellent, Good, Fair, Bad, we may wish to have a simulation that can generalize to these cases in a more accurate manner.

Our label conflation model first learns the probabilities of mistaking labels for each other. We would expect disagreement between highly-relevant and relevant documents, for instance, but less disagreement between highly-relevant and non-relevant documents. However, this model is data driven, so it will reflect the actual behavior of users. As a concrete example, the model learned for labels in the Dori-Hacohen and Allan dataset is presented in Table 1.

Given any truth label, we then sample a new value based on how humans often disagree with that particular truth value.

Table 1: Label Conflation Model for Controversy Web Pages [5]. The label with most agreement is the “Clearly Non-Controversial” label.

Text	#	Conflated With			
		2	1	0	-1
Very Controversial	2	237	83	23	48
Controversial	1	83	182	27	53
Possibly Non-Controversial	0	23	27	133	92
Clearly Non-Controversial	-1	48	53	92	594

4 RESULTS

Given our set of models that each reasonably approximate human labeling disagreement on this ambiguous task, we can now run a simulation to understand what the expected performance (under

any measure) should be for our humans under these models. For each setting, we run $N = 10000$ simulations.

Table 2: Simulated AUC scores for controversy detection. When predictions and truth are generated by the given models, we obtain a the following sampling of AUC scores.

#	System Model	Truth Model	Percentile AUC		
			5th	50th	95th
1	Sample	Average	0.862	0.890	0.917
2	Sample	Max	0.846	0.874	0.900
3	Sample	Sample	0.818	0.852	0.884
4	Conflate(Truth)	Sample	0.794	0.836	0.875
5	Conflate(Sample)	Conflate(Sample)	0.674	0.725	0.774
6	Flip($p=0.643$, Truth)	Average	0.593	0.639	0.685

In prior work, the best AUC reported for this task is 0.856 [8], and the AUC reported for the original work is 0.743 [5]. We present six pairings of our truth simulation models in Table 2. We have ordered our simulations from optimistic (label sampling system, average truth #1) to pessimistic (traditional agreement probabilities #6). This suggests to us that we can believe in the improvement presented from 0.743-0.856, but that we should be skeptical of any further improvements shown on this dataset, as even our optimistic models suggest that we are doing as well as a human can do given the ambiguity of the task.

5 CONCLUSIONS

In this work, we have briefly presented a number of strategies for investigating the agreement of users on labeling tasks. Given a vector of document labels assigned by different people at each document, we can model the difficulty of particular instances and particular labels. Further work is needed to understand the best simulation models for given tasks, but exploring a variety of reasonable models allows us to come to a reasonable conclusion that the discriminative power of an existing controversy detection dataset has been used up in terms of a robust classification metric: AUC. We therefore propose that future work on classifying or ranking using subjective labels consider simulation as an explainable alternative to opaque agreement scores.

ACKNOWLEDGEMENTS

This work was supported in part by the Center for Intelligent Information Retrieval.

REFERENCES

- [1] Arie Ben-David. 2008. About the relationship between ROC curves and Cohen’s kappa. *Engineering Applications of Artificial Intelligence* 21, 6 (2008), 874–882.
- [2] Adam Bermingham and Alan F. Smeaton. 2009. A Study of Inter-annotator Agreement for Opinion Retrieval. In *SIGIR*. 784–785.
- [3] Chris Buckley and Ellen M Voorhees. 2004. Retrieval evaluation with incomplete information. In *SIGIR*. 25–32.
- [4] Ben Carterette and Ian Soboroff. 2010. The effect of assessor error on IR system evaluation. In *SIGIR*. 539–546.
- [5] Shiri Dori-Hacohen and James Allan. 2013. Detecting controversy on the web. In *CIKM*. 1845–1848.
- [6] Myungha Jang and James Allan. 2016. Improving automated controversy detection on the web. In *SIGIR*. ACM, 865–868.
- [7] Myungha Jang, Shiri Dori-Hacohen, and James Allan. 2017. Modeling Controversy within Populations. In *ICTIR*. ACM, 141–149.
- [8] Myungha Jang, John Foley, Shiri Dori-Hacohen, and James Allan. 2016. Probabilistic approaches to controversy detection. In *CIKM*. 2069–2072.

- [9] David Maxwell and Leif Azzopardi. 2016. Agents, simulated users and humans: An analysis of performance and behaviour. In *CIKM*. 731–740.
- [10] Mark Sanderson and Justin Zobel. 2005. Information retrieval system evaluation: effort, sensitivity, and reliability. In *SIGIR*. ACM, 162–169.
- [11] Jean M Tague and Michael J Nelson. 1981. Simulation of user judgments in bibliographic retrieval systems. In *ACM SIGIR Forum*, Vol. 16. ACM, 66–71.
- [12] Ellen M Voorhees. 2000. Variations in relevance judgments and the measurement of retrieval effectiveness. *Information processing & management* 36, 5 (2000), 697–716.
- [13] William Webber and Jeremy Pickens. 2013. Assessor Disagreement and Text Classifier Accuracy. In *SIGIR*. 929–932.
- [14] Kazimierz Zielinski, Radoslaw Nielek, Adam Wierzbicki, and Adam Jatowt. 2018. Computing controversy: Formal model and algorithms for detecting controversy on Wikipedia and in search queries. *Information Processing & Management* 54, 1 (2018), 14–36.