

Predicting the quality of new contributors to the Facebook crowdsourcing system

JULIAN EISENSCHLOS, Facebook
julianeisen@fb.com

1. INTRODUCTION

We are interested in improving the quality and coverage of a knowledge graph through crowdsourcing features built into a social networking service. In this setting, most participants are casual users, making only a few contributions, and do so incidentally in the course of using the service. Techniques that make assumptions about the matching of users to questions, or the number of answers per user or per question do not work well under such circumstances. We present an approach to model user trust when prior history is lacking, so that we can incorporate more new users' contributions into crowdsourced decisions, and provide quicker feedback to new participants. Specifically, we present a logistic regression classifier for first-time contributions, and study the effect of prior knowledge about user demographics on this classifier using Facebook crowdsourcing datasets.

1.1 Setup & prior work

Although automated systems are getting better at extracting knowledge from the web by parsing natural language and HTML [1], there are scenarios where human input is necessary [8] - for example, when information about the real world isn't catalogued in a programmatically accessible form, or is changing regularly or unpredictably, or when a machine learned predictor has insufficient confidence in its predictions. Crowdsourcing can help address these scenarios at scale.

We study the problem of gathering subject-predicate-object (from now on S-P-O) tuples for a database of international Facebook Places. The predicates under consideration include name, photo, phone number, website, address, open hours. For instance, we might know that *Facebook HQ* is located at *1601 Willow Road, Menlo Park, CA*, but we may be missing a phone number for this location. Our crowdsourcing interfaces enable users with differing levels of experience and reliability to suggest possible phone numbers, and our objective is to choose one (or more) phone numbers if we can determine that it is correct with high confidence.

Crowd contributors may be motivated intrinsically or extrinsically. Intrinsic motivations include the urge to prove oneself knowledgeable, to help others, or fix wrong information; while extrinsic motivations usually involve some material or monetary gain. To achieve worldwide coverage for places, we prefer to focus on the former because it is difficult to obtain access to enough paid workers with local knowledge. Further, [1] suggests that better results are achieved when one appeals to intrinsic motivations. In such a setting, most contributions are voluntary, and a majority of users provide a single or just a few answers.

Prior work on voting aggregation algorithms - using expectation maximization [4], message passing [2] or other techniques - shows that assessing the quality of the contributor by their prior history and level of agreement with the crowd leads to improved performance when compared to simple majority voting. These approaches make assumptions about regularity conditions on the distribution of users and questions, which are hard to enforce in our volunteer-biased system, significantly limiting the utility of these approaches.

We therefore want to understand how to predict the quality of answers when prior history about user contributions is lacking. Being able to correctly perform this classification allows more contributors to have impact on the system, and more opportunities to provide feedback and re-engage users.

In this work we use logistic regression to classify first-time contributions based on features including demographic properties of contributors.

1.2 Crowdsourcing at Facebook

The crowdsourcing system at Facebook includes several product interfaces where a user can suggest new values for an attribute or vote on other users' suggestions. We use the term *claim* to refer to either a suggestion or a vote. Clearly, voting on suggestions requires less effort than making new suggestions.

One such crowdsourcing product, the Graph Editor¹, is customized for highly engaged contributors, and consists of an endless queue of locations requiring attention because they are missing attributes or suspected to have incorrect attributes. Experienced contributors can customize the queue to show them places they are familiar with, and make claims about these places. We refer to users that contribute lots of claims regularly as *power* contributors.

Graph Editor users, however only form a small fraction of the total crowdsourcing contributors on Facebook (< 1%), while accounting for a disproportionate fraction of the claims (over 10%). The majority of contributions come from *casual* contributors who encounter missing or erroneous values in the course of using other features on Facebook, and make a crowdsourcing claim while reporting this. Most casual users will only interact very occasionally with the crowdsourcing system, making only one or two claims ever.

Our crowdsourcing algorithm has two main tasks:

- Measure the quality of a particular user's contributions.
- Determine which S-P-O triple claims to accept or reject.

Its implementation is related to Karger et al. [2] with an added time dimension. We run a process where we iteratively update the beliefs of triples associated with a subject (a Facebook Place) based on our current estimates of contributor quality, or update the estimates of a user's quality based on the triples that user has contributed to. When triples reach certain confidence thresholds, we can mark those claims as accepted or rejected.

Strictly speaking, we are modeling a claim's probability of correctness q , based on features including the contributing user's history, the Place in question, the interface that the claim was made in, the relationship between the contributor and the Place etc. Historically, we have required that contributors gain a sufficient track record before they are permitted to have impact on the system. This is conservative because new contributors are the most likely to be confused about the crowdsourcing feature, or to misinterpret guidelines, or even to be malicious; however in our setting this ignores a significant fraction of incoming claims. Incorporating other features about these new contributions could help us be smarter about valuing these claims, and give a voice to new contributors.

2. PROPOSAL


In this section, we describe how we built a classifier to recognize accurate claims coming from first-time contributors using a logistic regression model incorporating demographic features.

2.1 Feature classes and training data

We consider three families of features relevant to evaluating the quality of a claim coming from a new contributor:

- Contributor features*, including demographics and patterns of Facebook usage.
- Place features*, including popularity for this Place and categorization (into an in-house taxonomy). Adding this information might enable the model to learn a proxy for question difficulty and expertise of demographic groups. For instance we might learn that zip codes are harder to come by in a particular country, or that younger users are better in filling information for pubs. To capture such interactions between features, we incorporate bigram features corresponding to each pair of covariates.

¹Publicly available at www.facebook.com/editor



—*Claim features*, including features characterizing the connection between the contributor and the Place in question, the direction of the claim (agree or disagree), whether the claim is a suggestion or a vote, etc.

One of the hurdles we had to overcome was obtaining labeled data for training and testing. We use two kinds of labeled data sets, and they represent different tradeoffs between coverage and quality.

One set of labels (*the gold set*) is derived from in-house manual annotation work collected over a period of time, as well as high-confidence automated predictions. We estimate its accuracy above 95% but it incorporates various biases, including towards being mostly positive. A second set (*the inferred set*) can be obtained using our crowdsourcing inference algorithm described in section 1.2 and considering the labels it has deduced through consensus. It can have a lower accuracy but it is more evenly distributed over the Places database. We will omit consideration of these differences here, due to space.

We also included small manually rated sets of claims randomly sampled from new users’ contributions to obtain an unbiased set to be used for validation. Another direction that has been studied previously in [3] is assuming values that have passed a *test of time*, by being untouched for a long period of time to be correct. This is a worthwhile extension for us to consider.

In training we combine the gold set and the inferred data sets, assigning them weights w_g and w_i respectively. We treat the ratio of these weights as a training parameter to be chosen via cross-validation. In the next section we will explore how these features and training sets affect the performance of the model.

2.2 Results

In figure 2.2 we show the precision - recall curve obtained by our classifier in the validation set, comparing to test sets extracted from the gold set and the inferred set. We observe that the validation set performs worse than the others and attribute that to biases in the training sets that cause some overfitting. The model still shows considerable improvement compared to the baseline of 57% precision of new claims, measured on the validation set.

We also studied how the choice of the weights given to the gold set and the inferred set samples in the stochastic gradient descent impacts performance. The trade-off of accuracy and coverage of the training data affects the precision–recall curve on our validation set. We searched for a ratio of w_g and w_i that optimized the performance on the unbiased validation set, and found empirically that $w_g = 2w_i$ yields the best results.

It is important to clarify that the precision of the new contributor quality prediction is not the precision of the resulting knowledge base; this model primarily behaves as a filter to select potentially good contributions from new contributors (without the benefit of prior history features) and give them a weight in the inference algorithm that then combines claims from both new and experienced contributors.

Finally, to understand the relevance of the three families of features described in the previous section, we perform a feature ablation study. The results are in figure 2.2. The most relevant set of features are the claim features, which model the interaction between the place and the user, but all three sets contribute to achieve the precision levels we are targeting (above 90%) with a reasonable recall.

REFERENCES

- [1] P. Ipeirotis and E. Gabrilovich (2014) Quiz: Targeted Crowdsourcing with a Billion (Potential) Users. *www*
- [2] D. R. Karger, S. Oh, and D. Shah. (2011) Iterative learning for reliable crowdsourcing systems. *Advances in Neural Information Processing Systems 24*, pages 1953-1961
- [3] C. How Tan, E. Agichtein, P. Ipeirotis and E. Gabrilovich (2014) Trust, but Verify: Predicting Contribution Quality for Knowledge Base Construction and Curation. *WSDM*
- [4] Q. Liu, J. Peng and A. Ilher (2012) Variational Inference for Crowdsourcing *NIPS*
- [5] A.N. Kolmogorov (1933) Sulla determinazione empirica di una legge di distribuzione *Giorn. Ist. Ital. Attuari* , 4 pp. 839-1
- [6] N.V. Smirnov (1938) On estimating the discrepancy between empirical distribution curves for two independent samples *Byull. Moskov. Gos. Univ. Ser. A* , 2 : 2 pp. 314
- [7] R. A. Fisher (1936). The Use of Multiple Measurements in Taxonomic Problems *Annals of Eugenics* 7 (2): 179-188
- [8] N. Vesdapunt, K. Bellare and D. Nilesch (2014) Crowdsourcing algorithms for entity resolution *Proceedings of the VLDB Endowment* 7 (12)