



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/138464/>

Version: Submitted Version

Article:

Chernushenko, I., Gers, F.A., Löser, A. et al. (Submitted: 2018) Crowd-labeling fashion reviews with quality control. arXiv. (Submitted)

© 2018 The Author(s). For reuse permissions, please contact the Author(s).

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Crowd-Labeling Fashion Reviews with Quality Control

Iurii Chernushenko, Felix A. Gers,
Alexander Löser
Beuth University of Applied Sciences
Luxemburger Str. 10
Berlin, Germany 13353
firstname.lastname@beuth-hochschule.de

Alessandro Checco
University of Sheffield
Information School
Western Bank
Sheffield, UK S10 2TN
a.checco@sheffield.ac.uk

ABSTRACT

We present a new methodology for high-quality labeling in the fashion domain with crowd workers instead of experts. We focus on the Aspect-Based Sentiment Analysis task. Our methods filter out inaccurate input from crowd workers but we preserve different worker labeling to capture the inherent high variability of the opinions. We demonstrate the quality of labeled data based on Facebook’s FastText framework as a baseline.

KEYWORDS

Crowdsourcing, corpus annotation, fashion, aspect-based sentiment analysis

1 INTRODUCTION

Today, users make purchases via websites and also communicate about their opinion there. Usually, this communication happens in the form of reviews users leave under the product description. Processing of these reviews is commonly done in the form of Aspect-Based Sentiment Analysis (ABSA) - i.e., mining and summarizing opinions from the text about specific entities and their aspects. This information can help consumers decide what to purchase and businesses to better monitor their reputation and understand the needs of the market [3].

Contribution. We describe a novel quality preserving crowd labeling approach on such reviews from a highly popular European apparel online store with more than 135 million transactions monthly. Existing datasets for the ABSA task are usually labeled by experts [3]. Labeling datasets by experts is slow, expensive and therefore often cannot be executed on a larger scale. We propose a non-expert crowdsourcing procedure to get big, high-quality labeled datasets with a low budget. Our approach is based on: a) Labeling fashion text reviews accompanied by the item images. This helps to solve disambiguation problems. b) Class labeling justification. We ask crowd workers to highlight relevant text fragment, that is triggering classification. This additionally forces them to reflect on the assigned label. c) Quality assurance with a gold standard, that is filtering out inaccurate crowd workers.

2 FASHION REVIEW DATASET

The initial dataset consists of 2.3 Mio. textual reviews in eleven languages written by users from the fashion store website. Each review consists of caption and text describing and/or rating the product (Figure 1). The text often has poor sentence separation, typos or unconventional punctuation usage. These problems make it harder to mine information out of the reviews.




Figure 1: User review from the initial dataset.

Aspect-selection with experts and LDA. After processing reviews with Latent Dirichlet Allocation (LDA) [1] we verified potential topics with two domain experts. As a result, we unravel five common and relevant topics: General (“Perfect”, “Fabulous!”, “very nice”), sizing or fitting (“Wrong sizes”, “way too big”), quality (“The quality is good and the stitching is even so the hems match and therefore is no twisting.”), design (“looks fabulous with a lurex hat”), delivery (“Quick delivery. Great!”).

Focus on sizing issues for shoes. Online stores in the fashion industry desire presenting accurate information about the size of their items. However, products from different manufacturers are not consistent in sizing information. Reviews of customers regarding sizing aspect differ depending on the product category. Here, we focus on labeling in the footwear category for the sizing aspect. For 3759 reviews related to shoes, we classify sentiments of Size/Fit opinions expressed in reviews as following classes: 1. Positive 2. Neutral 3. Negative 4. Other (not related to size/fit) 5. Data error (not a shoe, incorrect language, other data inconsistency).

3 QUALITY ASSURANCE

Ask one straightforward question. First, we formulated the task and wrote instructions¹ for crowd workers [2]. Initially, we tried to label among several classes (aspects of the opinions), but this causes confusion and errors in labeling. Hence, the task should be very simple for the non-experts to produce high-quality results.

Ask for justification. For every class we asked crowd workers to highlight corresponding spans (text fragments) in the review that trigger a correct classification. This approach enforces crowd workers to justify their labeling and leads to a better quality.

Image and text context. Every text review we accompanied with an image of the product (Figure 2). This helps crowd workers to disambiguate while reasoning about the review and is particularly important for the fashion domain.

Create a gold standard. We labeled 500 reviews by two domain experts and use this data to filter out inaccurate crowd workers. We

¹Full instructions for our dataset and labeling examples at https://github.com/yury-chernushenko/papers/tree/master/2018_www



Figure 2: Labeling interface.

Table 1: Reviews classification results.

Class	Precision	Recall	F1
Other	0.91	0.95	0.93
Positive	0.86	0.90	0.88
Neutral	0.67	0.19	0.29
Negative	0.91	0.84	0.88
Average (weighted)	0.90	0.90	0.90

iterated on this step 3 times to create high-quality gold standard labels and to choose gold labels that are simple enough to assess general crowd worker performance.

Quality control. Prelabeled gold standard questions are integrated throughout the whole labeling process. We compare the span of a crowd worker and of the gold standard by measuring the overlap relative to the size of the corresponding sentences. This approach is similar to the Intersection-over-Union (IoU) method proposed for object detection on image datasets [2]. We compare two spans by limiting minimum required intersection/union ratio (0.7 in our case) with corresponding gold standard span. If crowd worker is having low error rate on the gold standard questions, then we keep other labels from this crowd worker.

Keep high quality while running at scale. Crowd workers receive instructions and test questions in the beginning. Every crowd worker labels not more than 300 reviews, to avoid bias or worker attention degradation. Every item is labeled 3 times by different crowd workers to absorb diverse opinions.

4 CLASSIFICATION BASELINE

The final label distribution for our dataset is: Other - 6.323, Positive - 2.194, Neutral - 206, Negative - 3.574. According to this class distribution, classification precision should be higher than 0.51 (just predicting class “Other”). We created our baseline with a popular classification library called FastText [4]. FastText uses a bag of words approach and subword information. This functionality is beneficial for reviews with typos and mistakes. As preprocessing step we lowercase the text and remove punctuation. The neutral class is disproportional to others because customers usually express an extreme sentiment in the reviews. To express this imbalance we present weighted average of metrics. After an informed parameter grid search, we configured FastText with Wikipedia pre-trained word embeddings, 25 epochs, 3 word n-grams and set all other parameters to default. Table 1 reports our results after splitting 80/20 on train/test.

Table 2: Spans classification results.

Class	Precision	Recall	F1
Other	0.98	0.99	0.99
Positive	0.97	0.93	0.95
Neutral	0.27	0.23	0.25
Negative	0.94	0.92	0.93
Average (weighted)	0.97	0.97	0.97

Please note, we also observe similar results using pre-trained word embeddings from the amazon reviews dataset of “Clothing, Shoes and Jewelry”.

Table 2 shows classification results of spans only. Instead of taking whole reviews, we use only spans of highlighted text by crowd workers with corresponding classes. We consider text that is not part of any span as part of class “Other”.

Understanding errors. Reviews, where the algorithm is failing, are on average 30% longer. Third of these errors is cursing difficulty to distinguish Size/Fit and Comfortability aspects. This problem could be addressed by choosing the more specific definition of the aspect. Another 20% of errors are caused by the lack of consensus among crowd workers. In this case, the majority of crowd workers are labeling some review correctly, but a minor incorrect label was assessed during the test run. This problem could be addressed with an aggregation by review and assigning majority voted label.

5 SUMMARY AND FUTURE WORK

Labeling quality assurance ensures strong baseline. Our results indicate that labeling for Aspect-Based Sentiment Analysis can be done with a crowdsourcing approach at a large scale. The combination of different measures yields in high-quality data. Restricting the input to the text spans containing the opinion target expressions further improved the classification results. This indicates that labeled spans contain the most significant information.

Future work. Generally, there are two types of mislabeling for a specific review: disagreement on class between crowd workers and disagreement on span content. As future work, we plan to leverage on Opinion Term Expressions for better classification results and will apply neural networks with attention mechanism.

6 ACKNOWLEDGEMENTS

Our work is funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 732328 (FashionBrain).

REFERENCES

- [1] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.
- [2] Mark Everingham et al. 2015. The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision* 111, 1 (2015), 98–136.
- [3] Maria Pontiki et al. 2016. SemEval-2016 Task 5: Aspect Based Sentiment Analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*. 19–30.
- [4] Edouard Grave, Tomas Mikolov, Armand Joulin, and Piotr Bojanowski. 2017. Bag of Tricks for Efficient Text Classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers*. 427–431.