

Sentiment Analysis Towards Targeted Named Entities

Goal: Entity-specific sentiment analysis

Our goal is to develop a tool that generates a sentiment score for individual entities in any given review. We used a random subset of 15,000 restaurant reviews from the *Yelp Open Dataset* to validate our model.

“I **enjoyed** the **mac and cheese** in the **restaurant**, but I **hated** that they only offered cheap **beer** such as **Miller Lite**”.

PRODUCT FAC
PRODUCT BRAND

Figure 1. A motivating example. Sentiments towards different products varies within a review.

Methodology

- We **train** a SpaCy ER model to be able to recognize food & beverage products in reviews.
- A product list is obtained from **WordNet**.
- We achieved a F1 score of **91%** and novelty score of **97.17%**.

ENTITY RECOGNITION

- **Constituency Parsing** is used to split the comment into **sentences**.

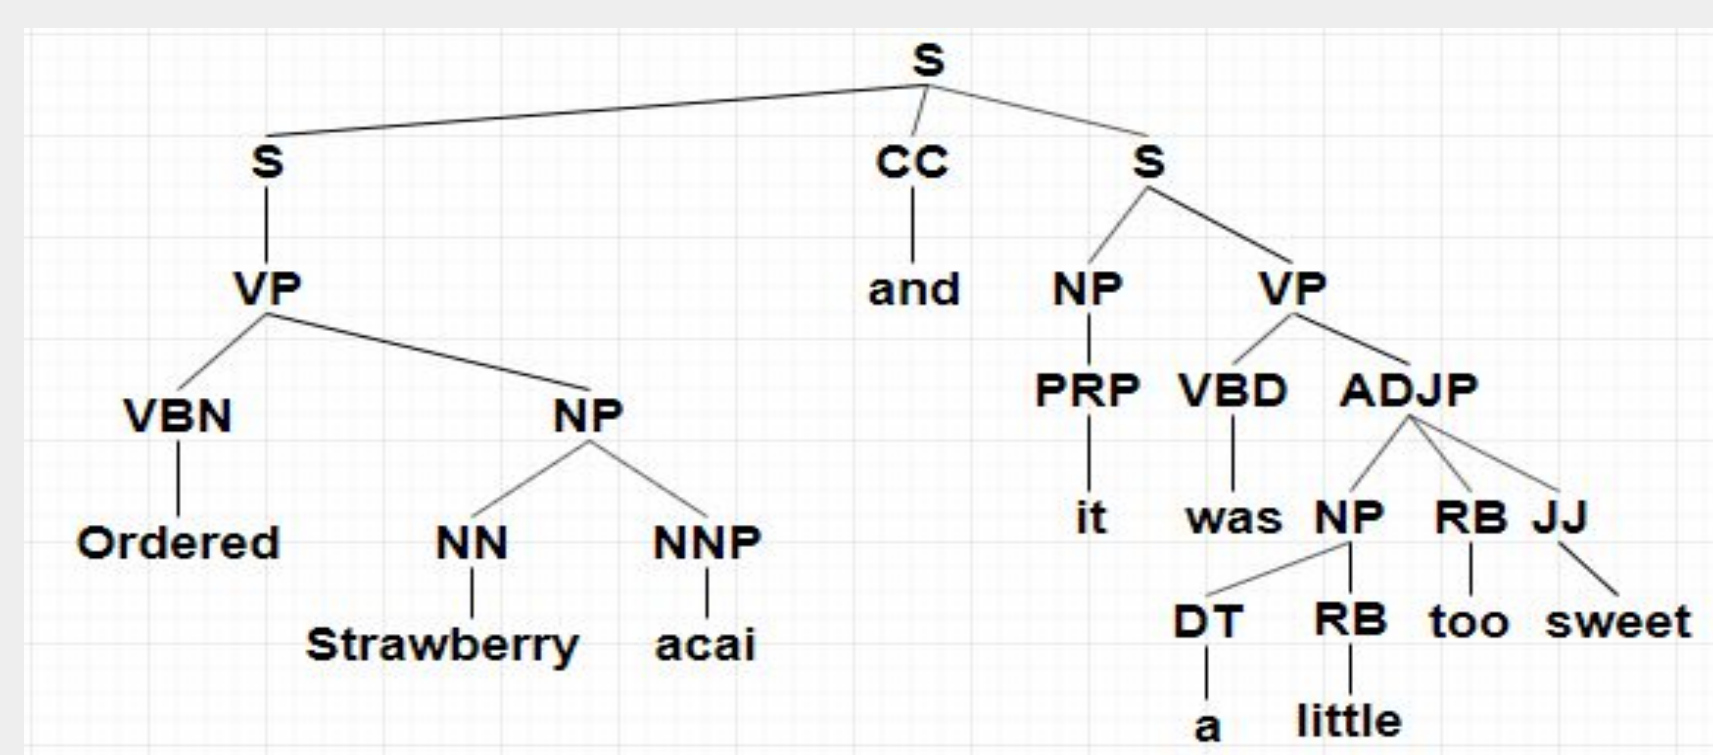


Figure 2. A parse tree

CONSTITUENCY PARSING

- **Parsing Rules** are used to traverse the parse tree and determine the context surrounding each entity

- The relevant contexts are inputted into algorithms like **VADER and Stanford NLP** to calculate sentiment scores for each entity.
- VADER (score of **0.25**) outperforms Stanford NLP (score of **-0.19**) using our validation technique on a baseline rule.

SENTIMENT ANALYSIS

Validation using Yelp Star Ratings

Validating our process is a challenge due to the lack of target labels in our dataset. To quantitatively evaluate our model, we developed a rank-based validation method that uses Yelp stars as a proxy to determine the population's sentiment toward an entity.



Figure 3. Illustration of end-to-end validation process.

Our **best rule** (rank correlation of **0.58**) is as follows: For each entity, start from the minimum sentence in parse tree containing the entity. While the sentiment is *neutral*, replace with next shortest sentence until we reach a non-neutral sentence. (or the root).

Example Results (Best Rule)

REVIEW:

“I had 3 tacos: the Standard which was my favorite, the *pork* was okay and the carne was typical, no surprises. I have to say that the Prickly Pear *margarita* was the absolute best!!!!”

Entity	Relevant Context	Sentiment (VADER)
Pork	the pork was okay	+0.23
Margarita	the Prickly Pear margarita was the absolute best	+0.52

Table 1. Example of end-to-end results

Conclusion and Future Work

Our methodology is able to generate sentiment scores on identified entities from an arbitrary corpus, with the help of a trained ER model. These steps have been packaged as open-source software at github.com/timjaya/neoway-brand-sentiment. Next steps include comparisons with other parsing methods, such as dependency parsing.

References

- C. J. Hutto, Eric Gilbert. 2014. VADER. A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. International AAAI Conference on Web and Social Media, North America, may. 2014.
- Peng Qi, Timothy Dozat, Yuhao Zhang and Christopher D. Manning. 2018. Universal Dependency Parsing from Scratch in *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pp. 160-170.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. <http://spacy.io>.

Acknowledgments

We would like to thank Felipe Penha, Sining Chen, Manoel Vilela, and Vitor Hugo for their guidance and mentorship throughout the course of the project.