


RTSIFT

Creating concise and meaningful review thread representations

Motivation




You would like to have a nice dinner



You checkout popular review websites



There are too many reviews to consider




What if a machine could help you out?



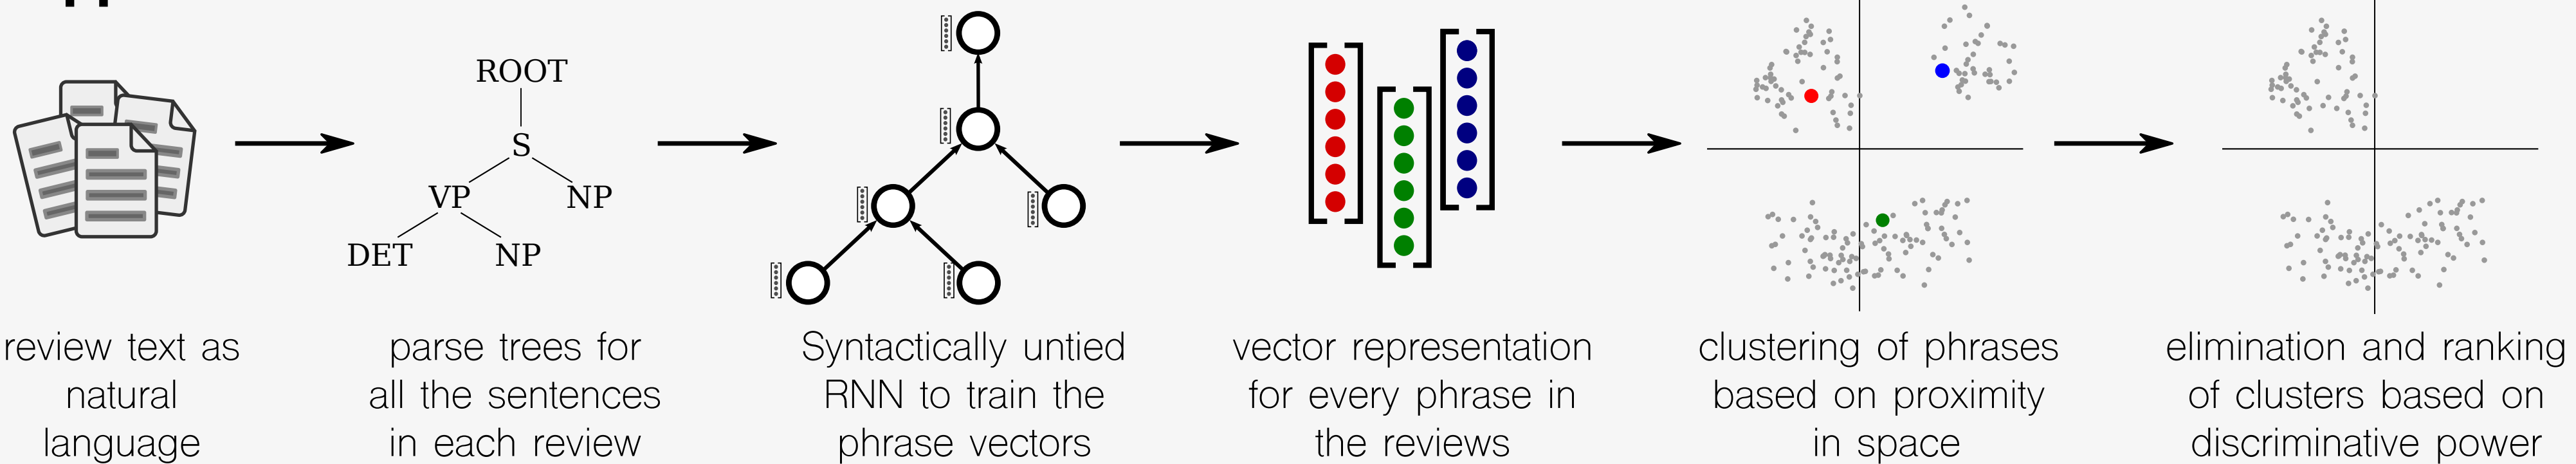
For this, the computers have to first understand the reviews!

Problem Statement

Create a concise yet meaningful representation for a review, and subsequently for an entire review thread. We take inspiration from word vector representations, applying the same concept at the level of phrases. This enables us to cluster words and phrases in space with respect to their semantic and syntactic meaning.




Approach



The final review representation is a vector of counts of constituent phrase memberships in the discriminative clusters.

Results



phrase vectors projected to 3 dimensions

★ ☆☆	give zero stars Horrible customer service give NEGATIVE stars waste your money here will never go back
★ ☆☆	gone down hill I've experienced better very disappointing not worth the money Mediocre at best
★ ☆☆	A-OK in my book is usually good was pretty average is just okay wanted to love this place
★ ☆☆	would definitely go back enjoyed this place was pleasantly surprised I am a fan is consistently good
★ ☆☆	won't regret it is booooo!! FANTASTIC! this place rocks! keep up the great service is exquisite too

sample phrases from discriminative clusters

Error Analysis

Training plateau

Training past ~4000 businesses proved to be ineffective in decreasing our average sentence loss.

Summarization.

you have found the right place
staff was very friendly and helpful

I grew up living in California

Examples of good/bad phrases from an actual system summary

The sources of the erroneous phrases included in the summary could be due to incorrect cluster memberships, malformed clusters or improperly ranked clusters.

Challenges

High dimensional space


Difficult to visualize the intermediate results without going through the entire pipeline

Limited processing power

Trade-off between simple models and more data vs complex models but less data

Indirect supervision

Using labeled review ratings to build models that work for other applications



Model performance on increasingly large training sets

Applications

Review thread summarization

Summarize large review threads by extracting discriminative and meaningful phrases

Business Comparison

Comparing two businesses and suggesting potential changes to improve star ratings
"Would I like this place?"

Using past data to predict if a user would like a specific business