Paraphrasing Negation Structures for Sentiment Analysis

Overview

- Problem:
  - Negation structures (e.g. “not”) may reverse or modify sentiment polarity
  - Can cause sentiment analyzers to misclassify the polarity
- Our approach:
  - **Remove** the negation by restructuring and then resurfacing the sentence
- Hypothesized benefits
  - **Improves** sentiment analysis accuracy
  - **Reduces** work for sentiment analysis implementers
- Results (so far!):
  - Implemented the paraphraser using Java, Stanford Parser, and Wordnet
  - Used data set and black-box classifier from [Socher 2013]
  - Reduction of 1.4% RMSE between ground-truth and classification on paraphrases

Example

0	1	2	3	4
Very negative	Negative	Neutral	Positive	Very positive

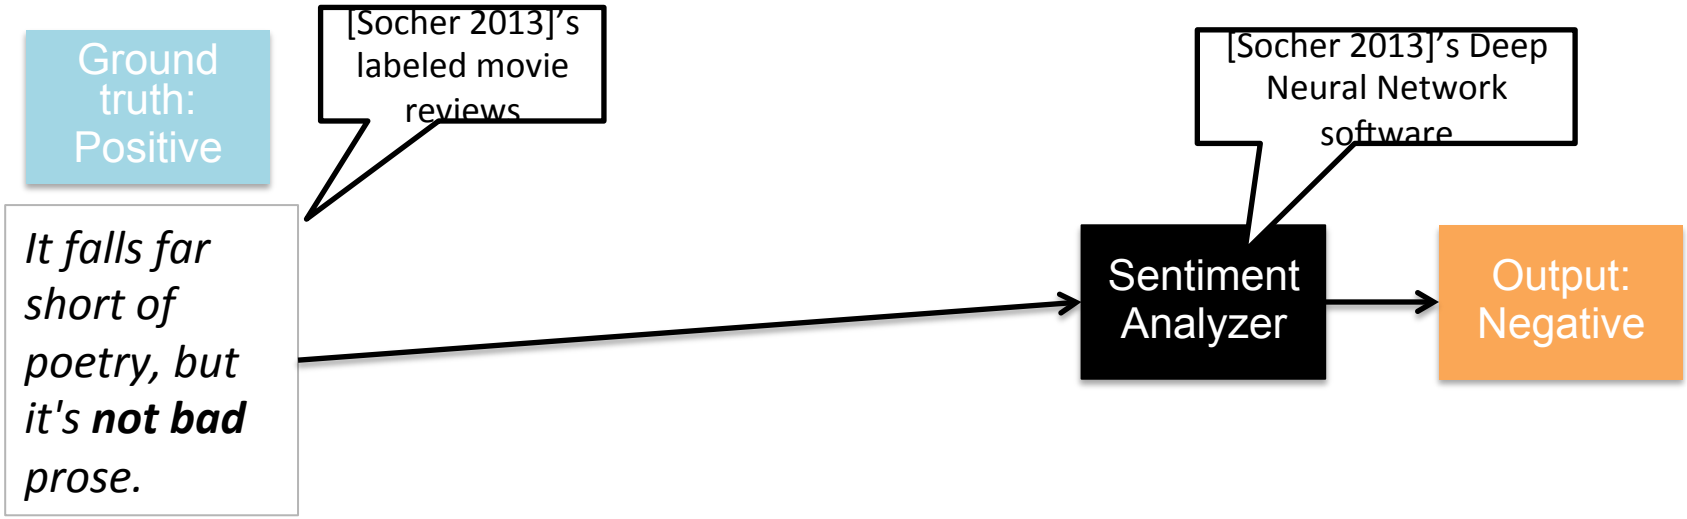

R. Socher, A. Perelygin, J. Wu, J. Chuang, C. Manning, A. Ng, and C. Potts. "Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank," In Proceedings of EMNLP, 2013.

Example

0	1	2	3	4
Very negative	Negative	Neutral	Positive	Very positive

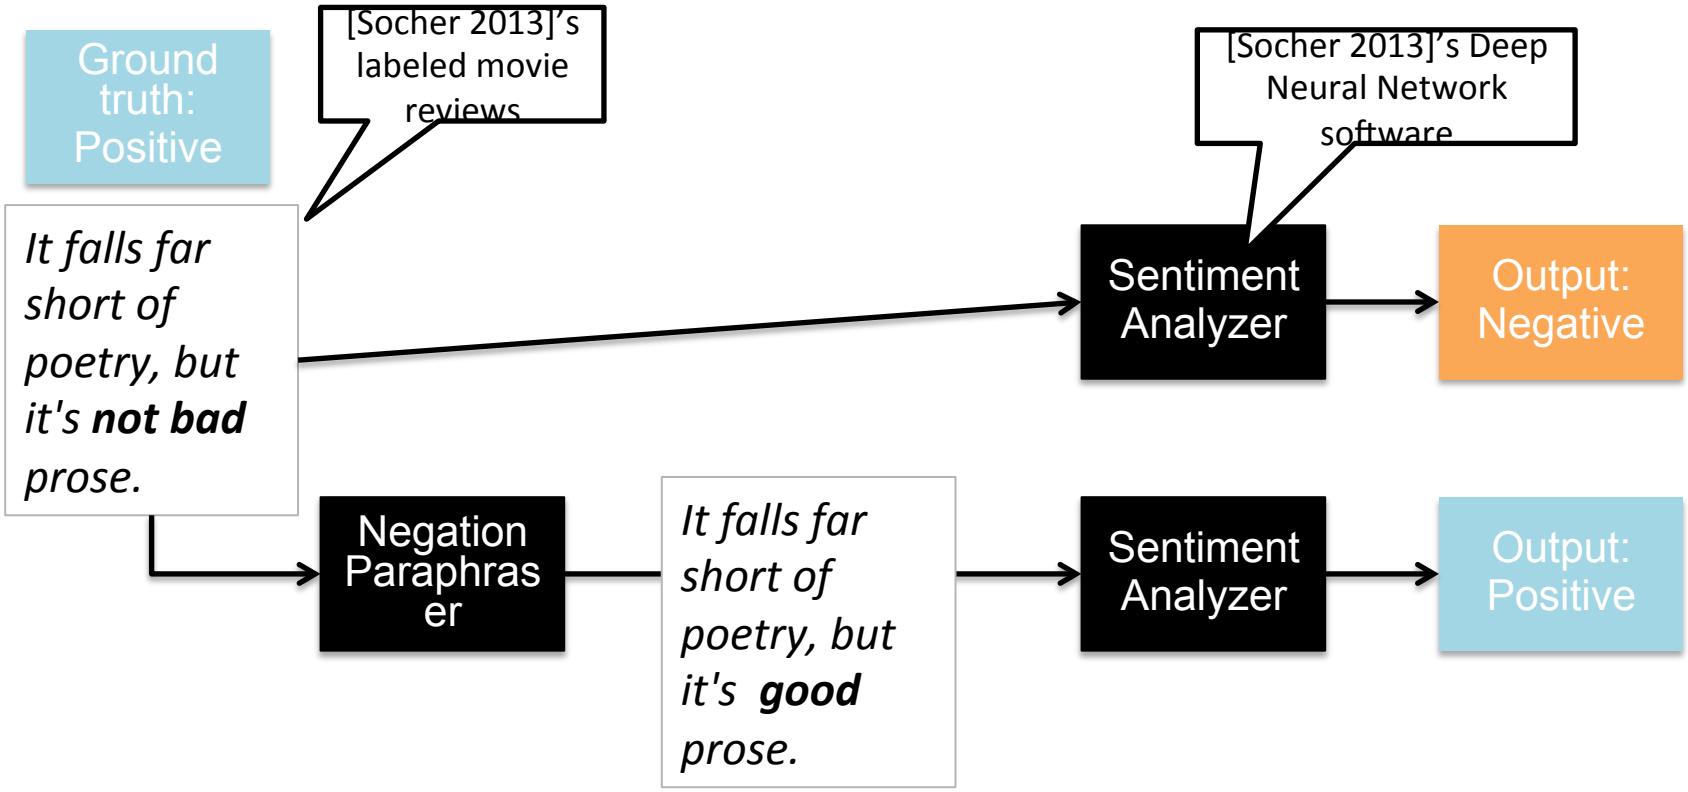

R. Socher, A. Perelygin, J. Wu, J. Chuang, C. Manning, A. Ng, and C. Potts. "Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank," In Proceedings of EMNLP, 2013.

The (observed) effect of nega on polarity classification

0	1	2	3	4
Very negative	Negative	Neutral	Positive	Very positive

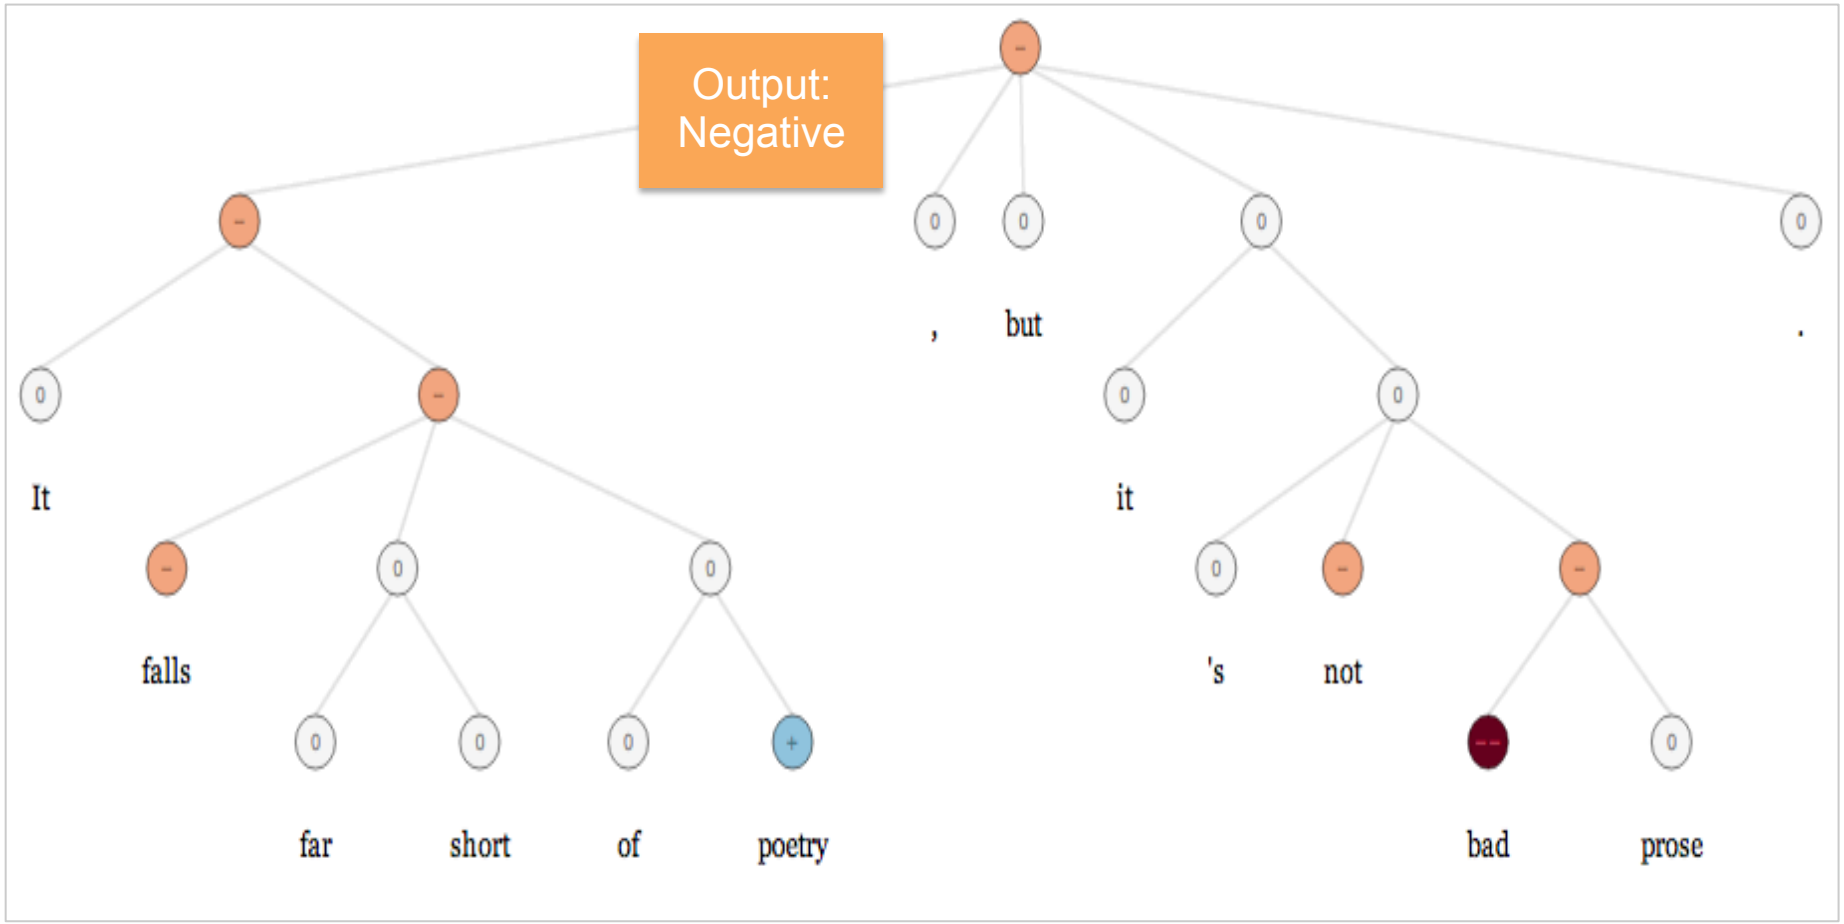

Ground truth: Positive

*It falls far short of poetry, but it's **not bad** prose.*

The (observed) effect of nega on polarity classification

0	1	2	3	4
Very negative	Negative	Neutral	Positive	Very positive

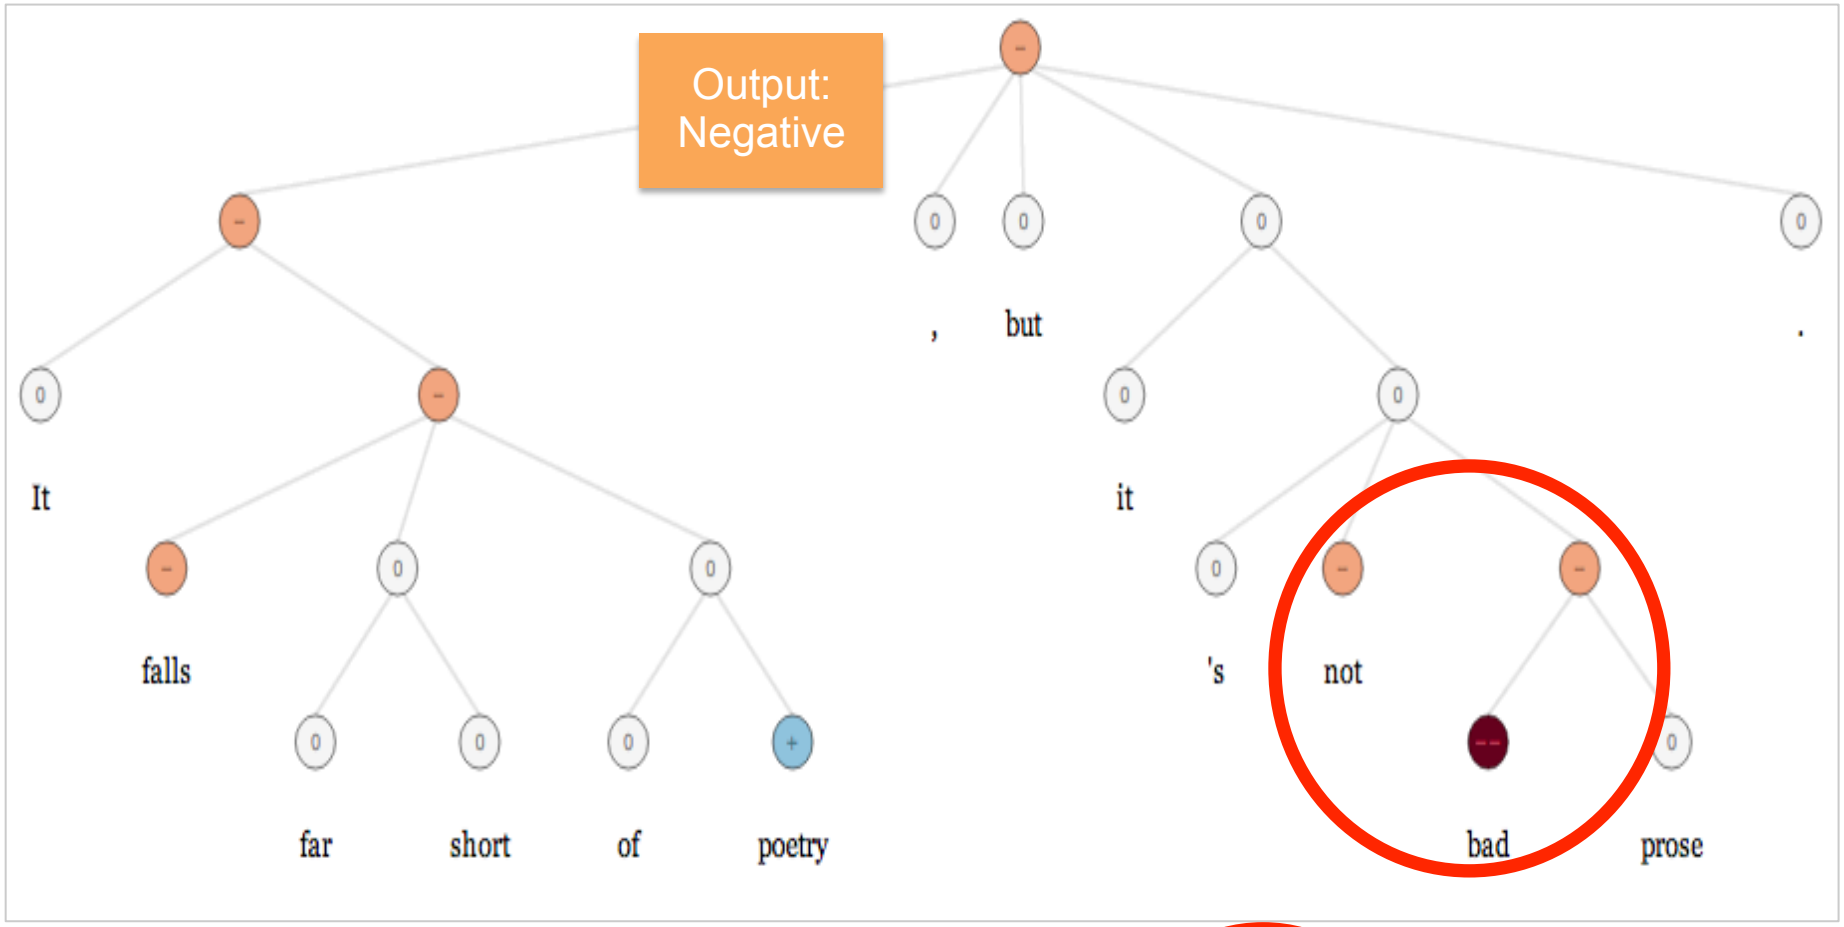

Ground truth: Positive

It falls far short of poetry, but it's ***not bad*** *prose.*

The (observed) effect of nega on polarity classification

0	1	2	3	4
Very negative	Negative	Neutral	Positive	Very positive

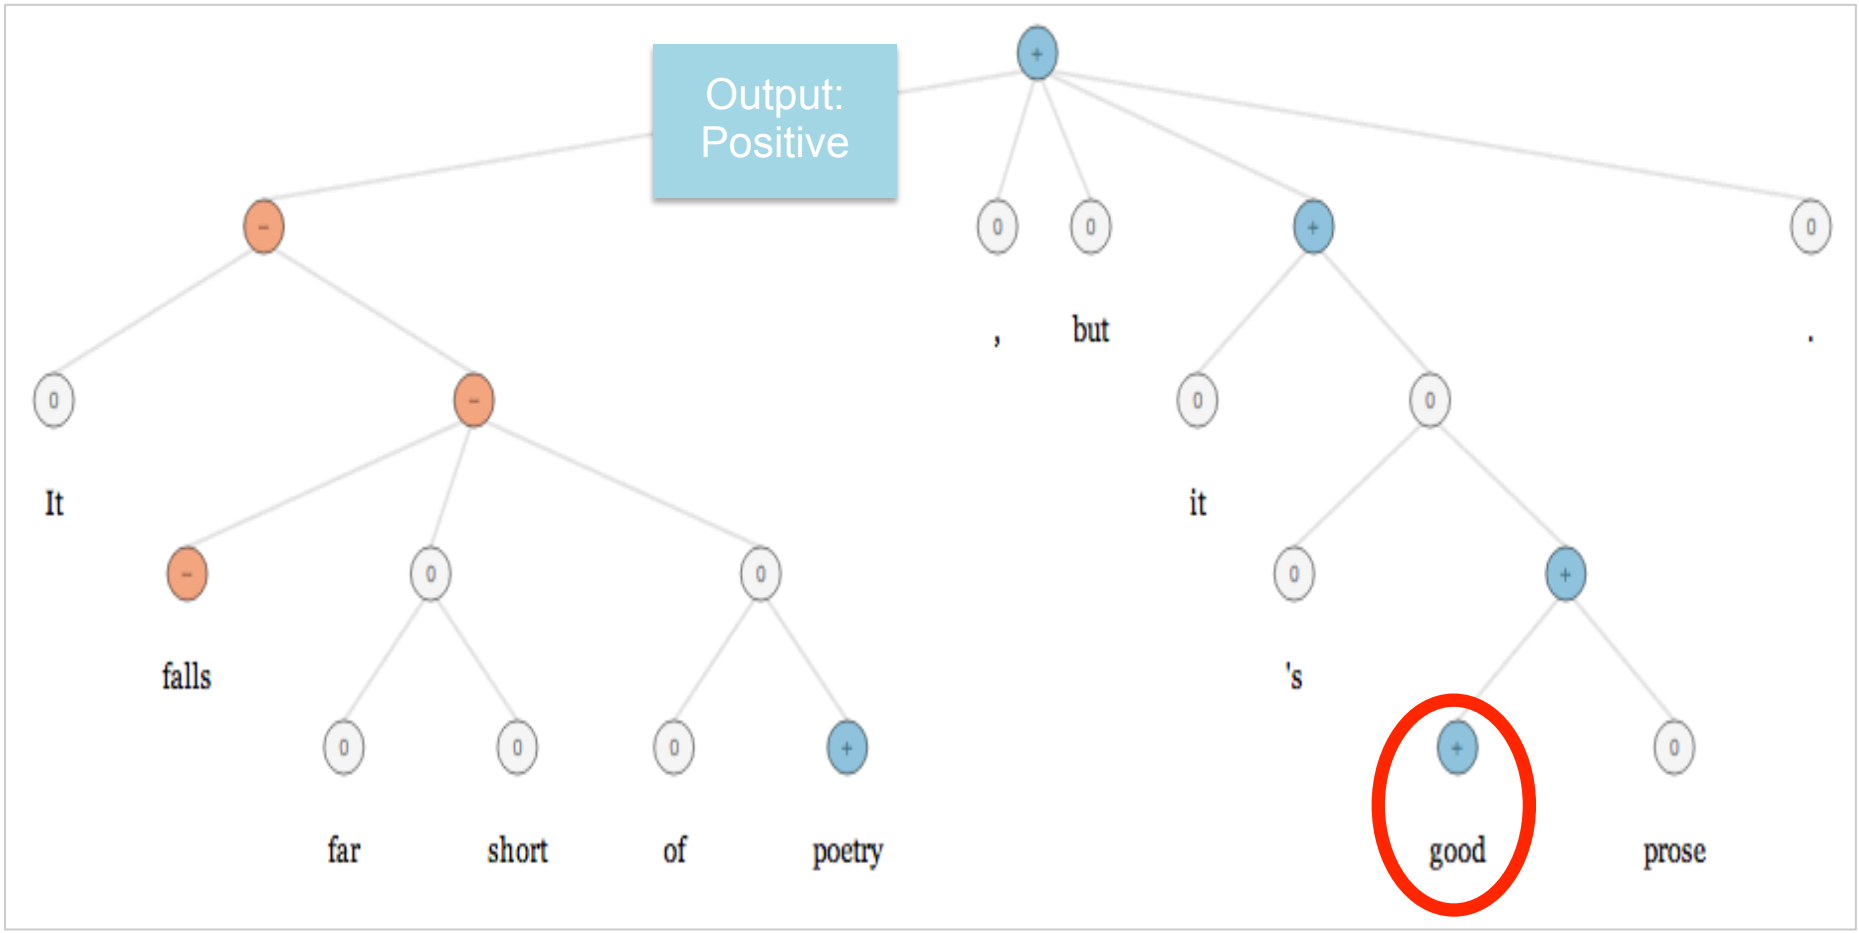

Ground truth: Positive

It falls far short of poetry, but it's ***good*** *prose.*

Related work on the treatment of negation

- Heuristic rules
  - A. Hogenboom, P. van Iterson, B. Heerschop, F. Frasincar, and U. Kaymak. "Determining Negation Scope and Strength in Sentiment Analysis," In Proceedings of IEEE SMC, 2011.
  - M. Hu and B. Liu. "Mining and Summarizing Customer Reviews," In Proceedings of ACM KDD, 2004.
  - L. Jia, C. Yu, and W. Meng. "The Effect of Negation on Sentiment Analysis and Retrieval Effectiveness," In Proceedings of ACM CIKM, 2009.
- Supervised machine learning
  - E. Lapponi, J. Read, and L. Ovrelid. "Representing and Resolving Negation for Sentiment Analysis," In Proceedings of IEEE ICDMW, 2012.
  - T. Wilson, J. Wiebe, and P. Hoffman. "Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis," In Proceedings of EMNLP, 2005.

Design & Implementation: Negation structures as polarity shifters

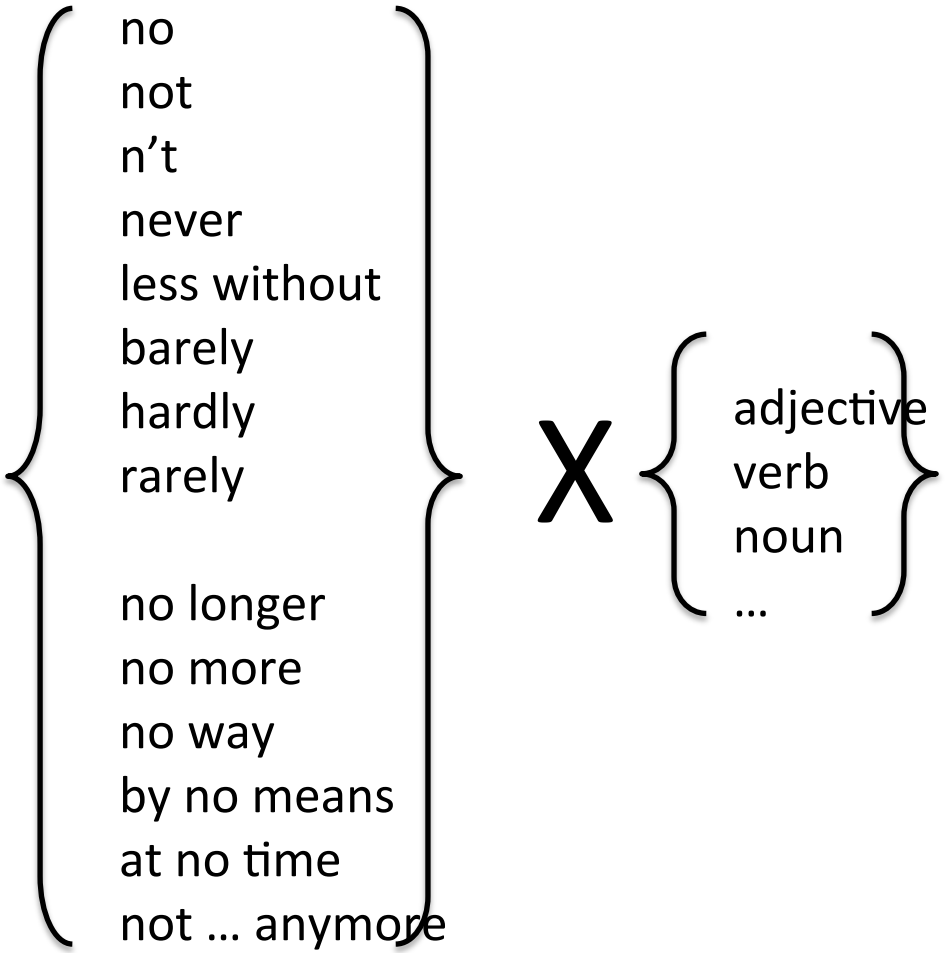

List from [Jia 2009]

Design & Implementation:
Negation structures as polarity shifters

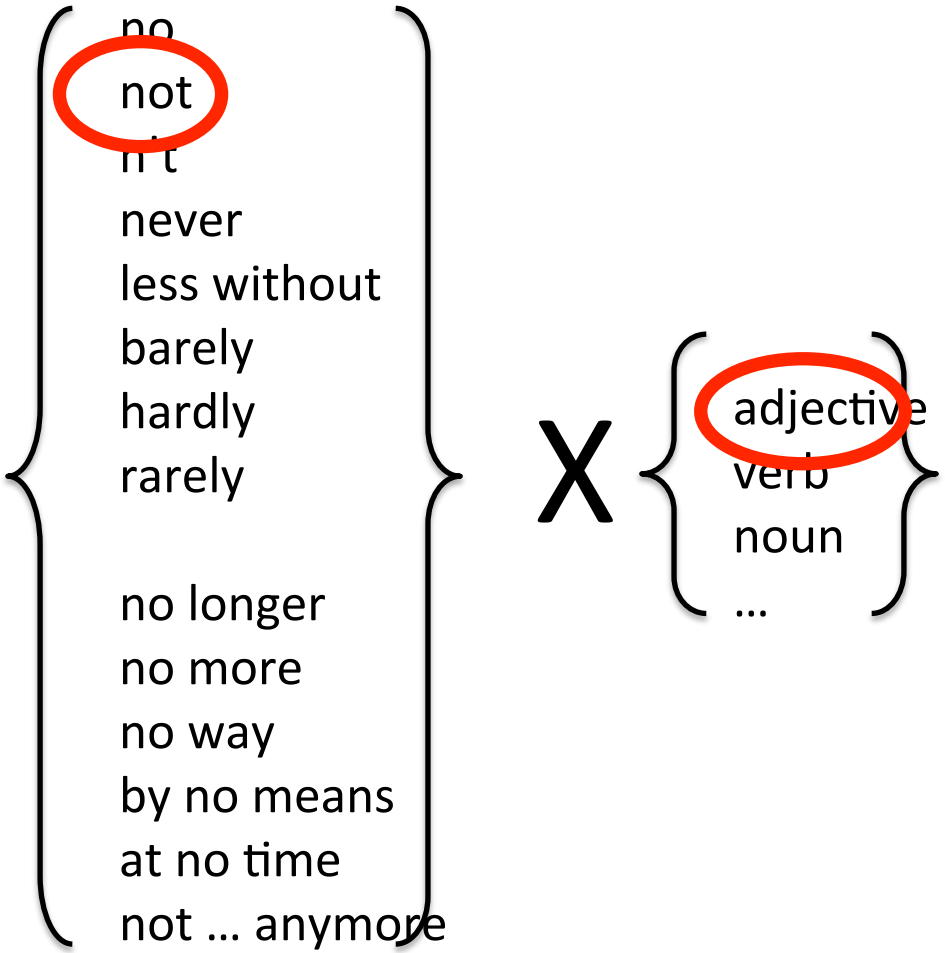

List from [Jia 2009]

Design & Implementation: Negation Paraphraser Pipeline

Design & Implementation: Negation Paraphraser Pipeline

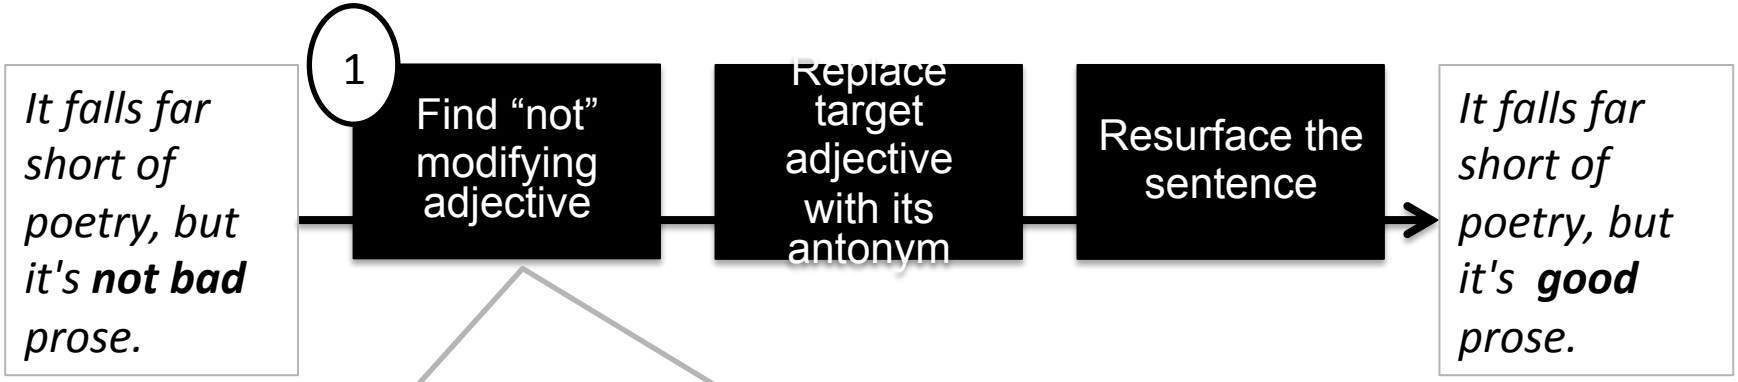

- Used the Stanford Parser software to build a parse tree
- Find "not"
- Find the first adjective who is a right descendent of

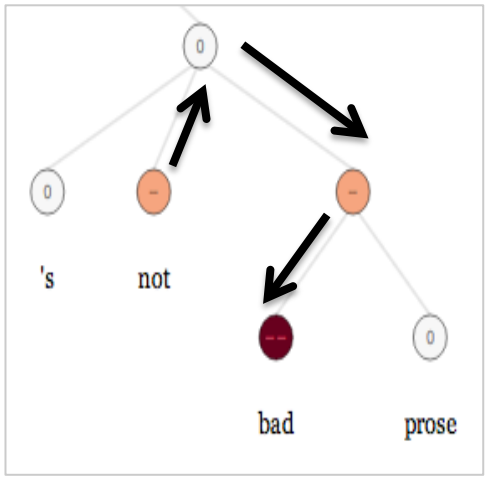

Design & Implementation: Negation Paraphraser Pipeline

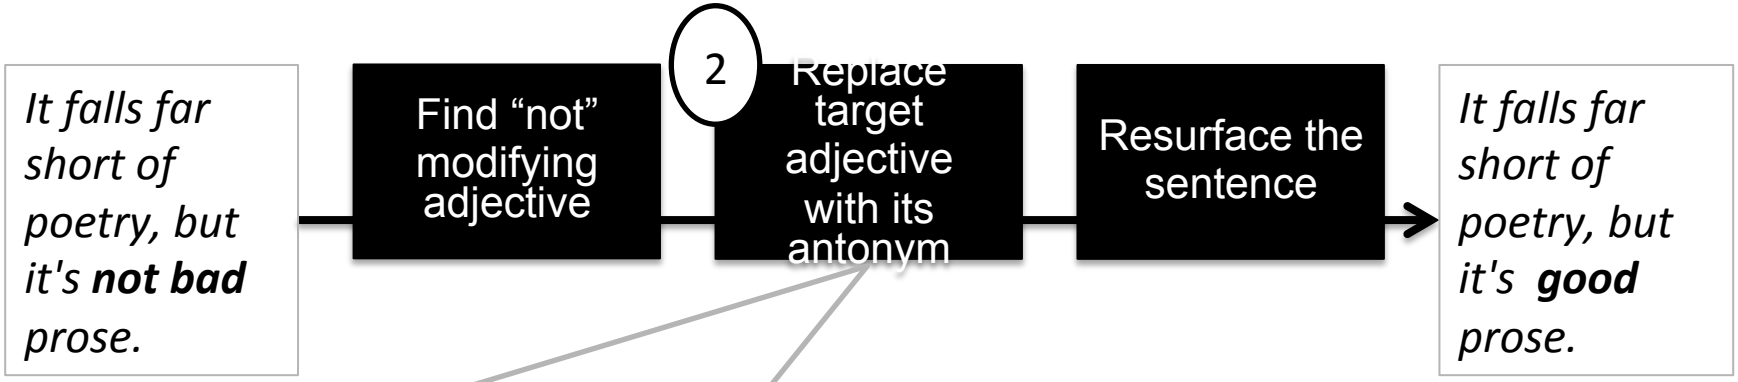

- Used Wordnet
- Find the adjective synset
- Find head synset
- Find antonym
- Replace adjective with antonym in tree

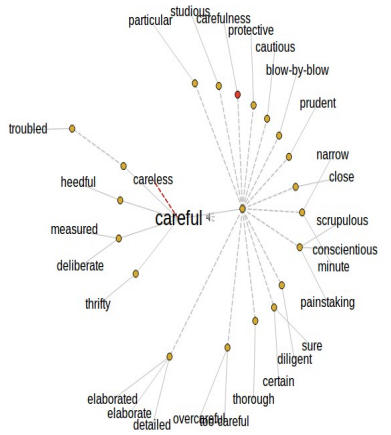

Design & Implementation: Negation Paraphraser Pipeline

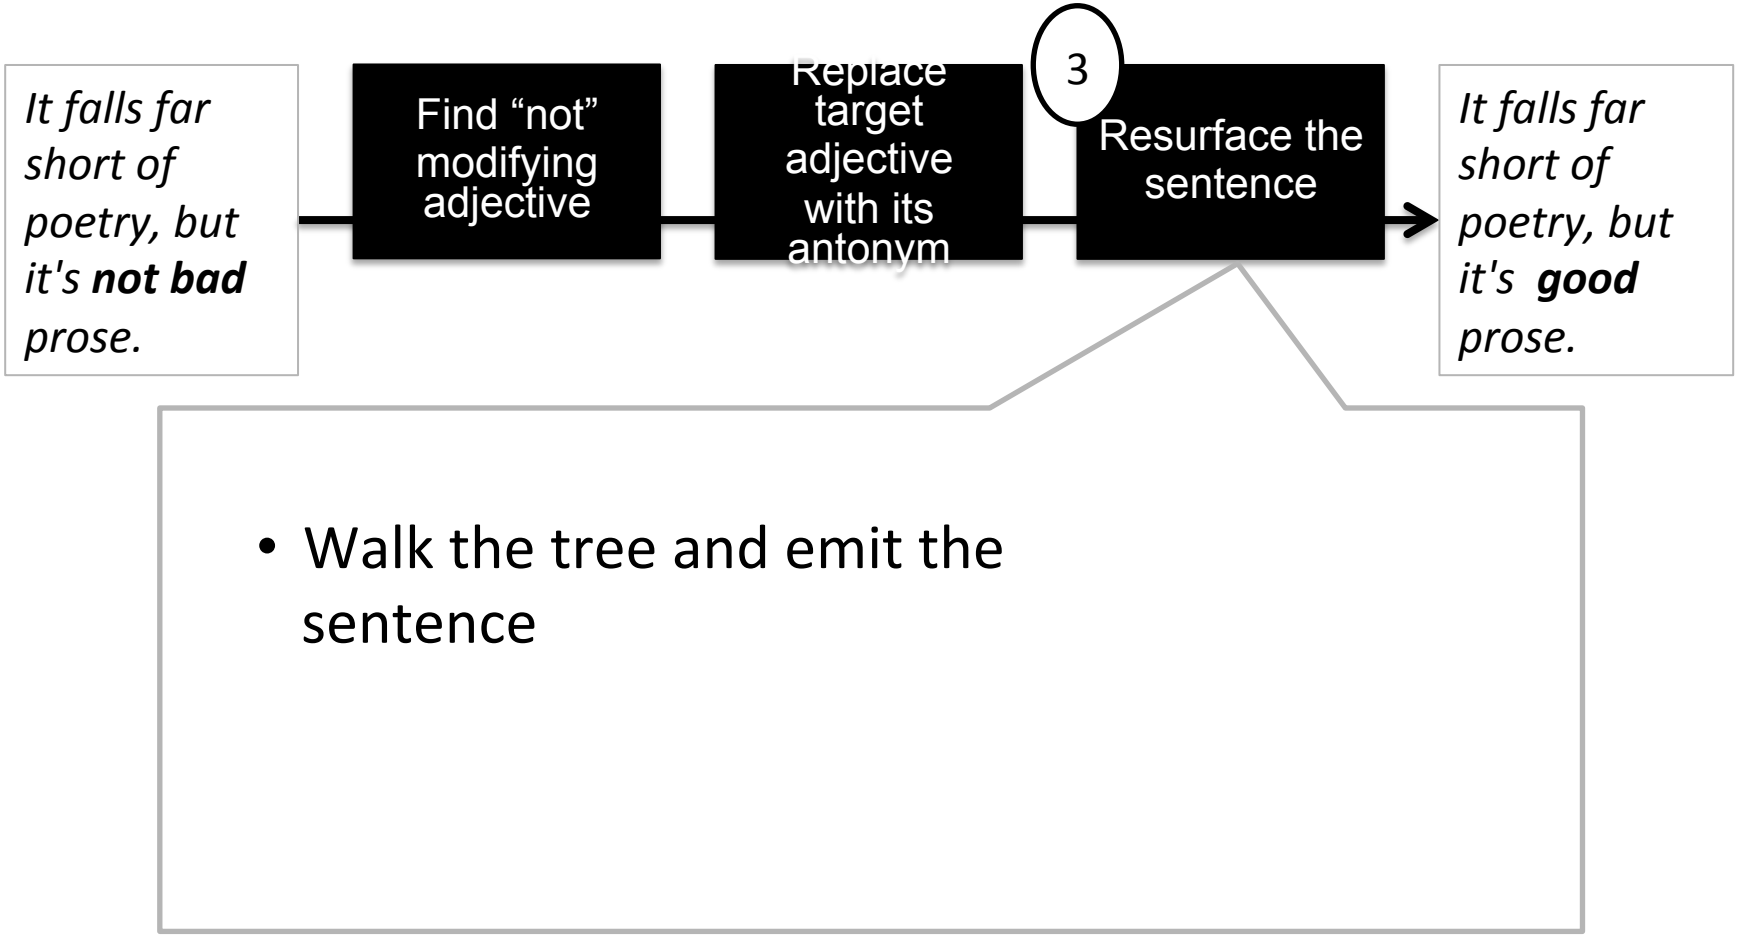

Experiments:
Data set from [Socher 2013]

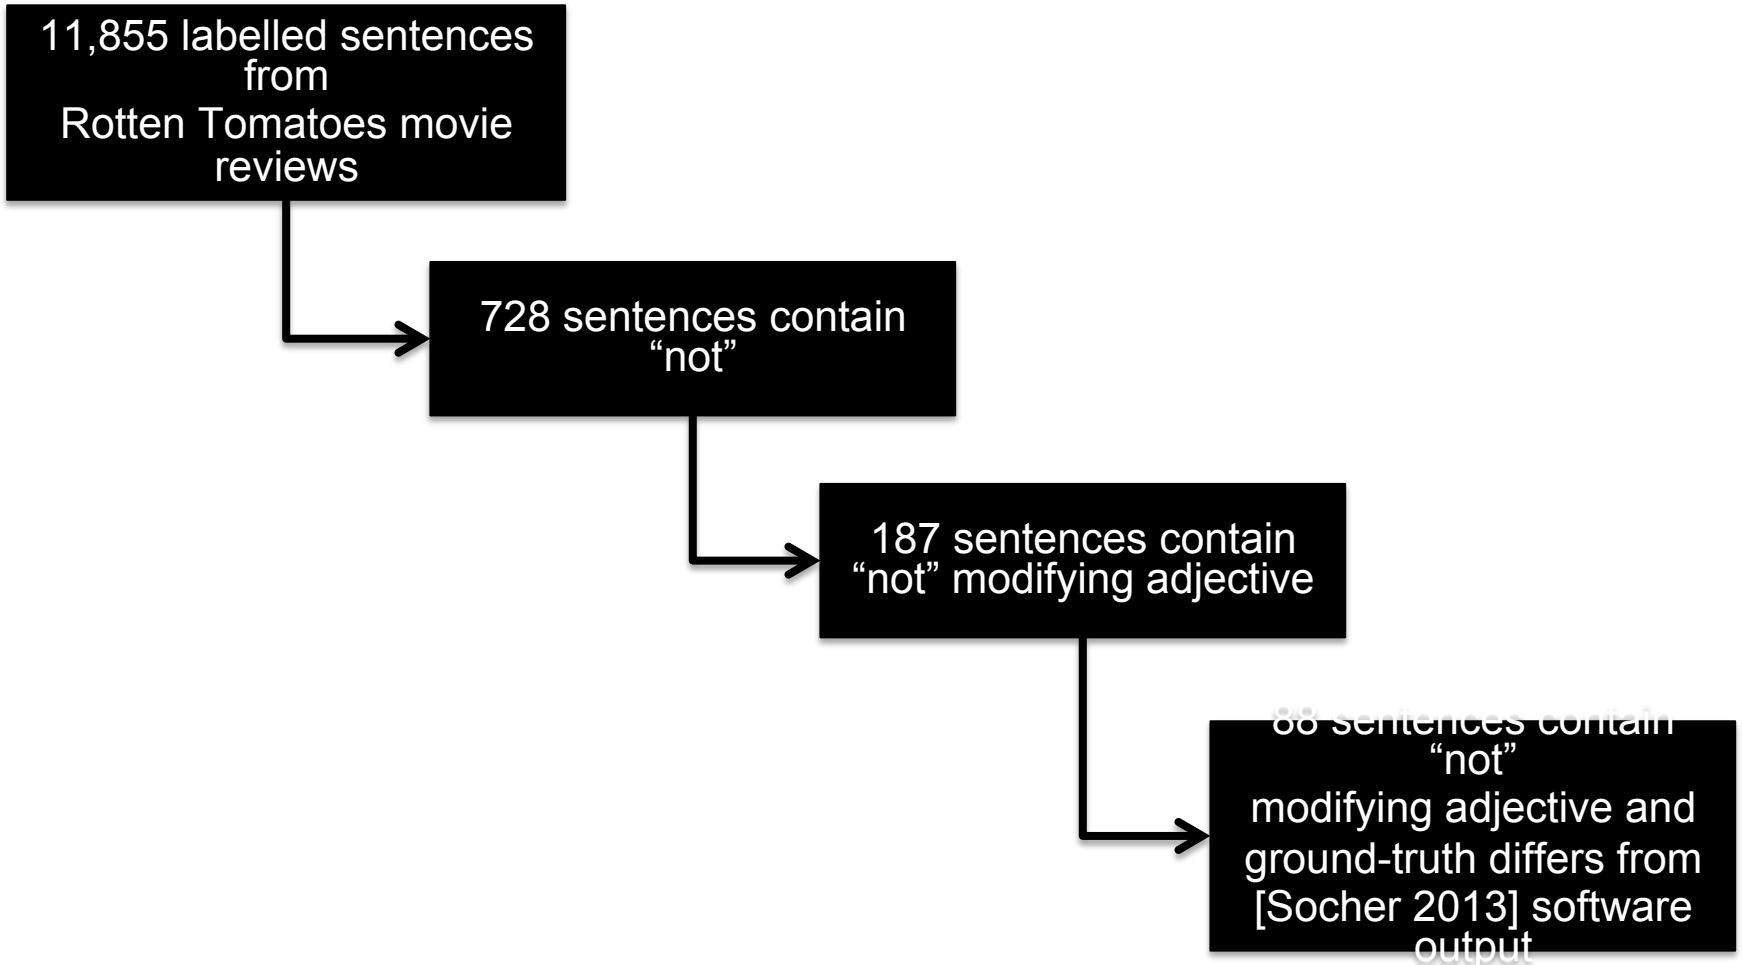

Results:
Good examples

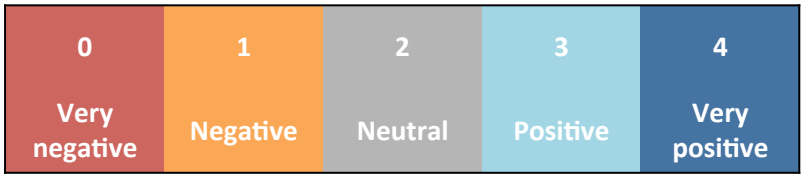

Input sentence	Ground-truth polarity	Output of [Socher 2013] classifier	Paraphrased sentence	Output of [Socher 2013] classifier on paraphrased
*S1M0NE 's satire is **not subtle** , but it is effective .*	Positive	Negative	*S1M0NE 's satire is **palpable** , but it is effective .*	Positive
*Certainly not a good movie , but it was **not horrible** either .*	Negative	Neutral	*Certainly a bad movie , but it was **innocuous** either .*	Negative
*At times a bit melodramatic and even a little dated (depending upon where you live), Ignorant Fairies is still quite good-natured and **not** a **bad** way to spend an hour or two .*	Positive	Negative	*At times a bit melodramatic and even a little dated (depending upon where you live), Ignorant Fairies is still quite good-natured and a **good** way to spend an hour or two .*	Positive

Results:
Not good examples

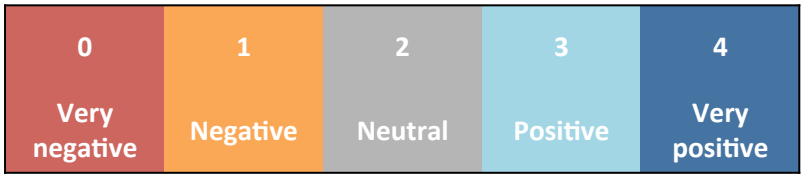

Input sentence	Ground-truth polarity	Output of [Socher 2013] classifier	Paraphrased sentence	Output of [Socher 2013] classifier on paraphrased
*It 's one of the saddest films I have ever seen that still manages to be uplifting but **not** overly **sentimental** .*	**Very Positive**	**Negative**	*It 's one of the saddest films I have ever seen that still manages to be uplifting but overly **tough** .*	**Negative**
*It uses an old-time formula , it 's **not** terribly **original** and it 's rather messy -- but you just have to love the big , dumb , happy movie My Big Fat Greek Wedding .*	**Positive**	**Negative**	*It uses an old-time formula , it 's terribly **unoriginal** and it 's rather messy -- but you just have to love the big , dumb , happy movie My Big Fat Greek Wedding .*	**Very Negative**

Results:
Overall evaluation of 88 sent

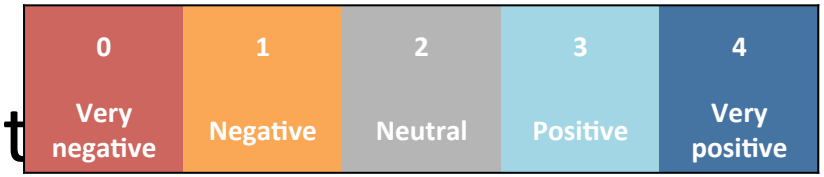

Predicted

Ground Truth

	0	1	2	3	4
0:	0	20	1	0	0
1:	3	0	2	2	0
2:	2	27	0	6	0
3:	0	14	1	0	2
4:	0	4	0	4	0

RMSE = 1.418

	0	1	2	3	4
0:	3	14	2	2	0
1:	3	2	0	2	0
2:	3	18	2	12	0
3:	1	9	1	4	2
4:	0	3	0	4	1

RMSE = 1.398

Without Negation Paraphraser

With Negation Paraphraser

Conclusion

- What's right
  - Some examples demonstrate improvement
  - Overall 1.4% improvement with "not" modifying adjectives
- What's wrong
  - Generated antonyms may have wrong sense – need some disambiguation
  - Generated antonyms affected by other modifiers
  - Generated antonyms were not in training set
  - Generated antonyms simply do not affect the classifier
- What's next
  - Try out different negation structures