

---

# Identifying Nominals with No Head Match Co-references Using Deep Learning

---

**Matthew Stone**  
mdstone@stanford.edu

**Ramnik Arora**  
rarora@cs.stanford.edu

## Abstract

Identifying nominals with no head match is a long-standing challenge in coreference resolution with current systems performing significantly worse than humans. In this paper we present a new neural network architecture which outperforms the current state-of-the-art system on the English portion of the CoNLL 2012 Shared Task. This is done by using a logistic regression on features produced by two sub-models, one of which has the architecture proposed in [CM16a] while the other combines domain specific embeddings of the antecedent and the mention. We also propose some simple additional features which seem to improve performance for all models substantially, increasing  $F_1$  by almost 4% on basic logistic regression and other complex models.

## 1 Introduction

Coreference resolution is a domain of Natural Language Processing which focuses on finding all of the references to the same real world entity within a body of text. Applications of coreference resolution include full text understanding (for discourse understanding), machine translation (gender, numbers), text summarization, and question-answering systems.

This paper will focus on a specific aspect of coreference resolution - that of nominals with no head match. A head match is a matching word, such as “President” that often acts as a strong hint for coreference. This no head match subset of the problem is assumed to be the most difficult because of the lack of direct signals to indicate a match. Nominal matches are ones involving descriptions. Other head match types include proper for names and pronominal for pronouns. An example of a nominal coreference with no head match is matching the antecedent “the country’s leftist rebels” to the anaphor of “the guerrillas” which was a notable win in [CM16a].

Traditional approaches to coreference began with Hobb’s naive algorithm [Hob86]. This approach embodies the semantic or structural methodology behind anaphora resolution through a rules-based system. This algorithm is one of the early baselines for coreference.

There are three major modeling approaches to coreference:

- **Mention pair models**, which is the type this paper will attempt to implement, make independent pair-wise decisions. Typically these models also implement a reconciliation mechanism as a final stage before evaluation. These mechanisms map links that are implied by the model predictions. (see example: [Hos16])
- **Mention ranking models** are attempting to solve a ranking problem by listing all possible antecedents and optimizing a ranking system to find the best or most likely antecedent mention. (for example: [HS11])
- **Entity mention models** attempt to link all mentions of one entity in a cluster model. This approach added a lot of power in [CM16b] by selectively choosing which a clustering algorithm that reduces errors on large scale mention clusters.

Current state of the art models have a CoNLL  $F_1$  [PMX<sup>+</sup>12] score of almost 69% [CM16a] in aggregate while only 18.9% for coreferences with no head match. Humans on the other hand score over 60% on similar tasks. Thus, these nominals with no head match constitute an important part of human understanding and is an especially difficult task for computers. Nominals with no head match need to make much more use of context and, in particular, word vector embeddings for related semantic meaning to identify these types of coreferences. With this in mind, the approaches found in this paper will emphasize aspects of word vector alignment. Our framework for this model is going to be that of a mention-pair model as the proposed models in this paper treat all mentions-pairs as independent, without any post reconciliation implemented as yet.

In our model, we increase the test  $F_1$  on a 2% negative sample downsample from 0.49 for logistic regression baseline and 0.54 for state-of-the-art architecture in [CM16a] to 0.61 using a combination of new features and intelligent architecture. Note that these are not CoNLL  $F_1$  scores, since we're only working on independent intra-document pairwise predictions rather than the downstream clustering task.

## 2 Data & Evaluation

We use the data provided by the CoNLL 2012 Shared Task on English [PMX<sup>+</sup>12].

### 2.1 CoNLL 2012 Shared Task

The CoNLL 2012 shared task dataset provides 2802 documents in training, 343 documents in dev and 348 documents in test. The mentions, and the co-reference pairs and types are tagged in this dataset. We filter on nominal mentions and co-reference pairs which have no head-match. We structure our dataset as an evaluation of direct pair matches, and try to build deep learning models directly around this subset. For modeling simplicity in this paper, we treat all mentions as independent and try to learn simply based on input features.

The class label counts in the dataset are heavily skewed. If there are  $N$  mentions in a documents, there will be  $N^2$  total mention pairs and only a few of them will be co-references. This means that the dataset is heavily skewed with significantly more negative examples than positive examples. In our training data we have 7,643,471 negative samples while 8,010 positive samples. In most of our experiments, we've downsampled the negative samples to 2% while maintaining 100% of our positive samples. Even with this downsampling, only 5% of our modified samples are positive mentions. See figure 1 for schematic representation of data-generating process.

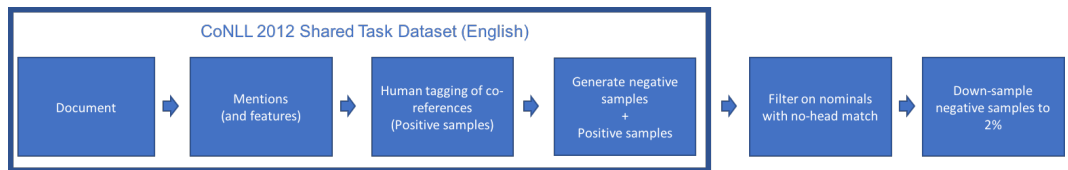


Figure 1: Schematic diagram of the data

### 2.2 Evaluation Metrics and Nuances

$F_1$  is a standard metric in binary and co-reference classification tasks. For co-reference classification, CoNLL  $F_1$ , which is an average of MUC [VBA<sup>+</sup>95],  $B_1^3$  [BB98] and CEF [Luo05]  $F_1$  scores is common. While we've used  $F_1$  as an evaluation metric, we recognize some of its shortcomings. Since  $F_1$  scores are based on hard  $\{0, 1\}$  classification, while the model predicts a probability estimate, they are very sensitive to thresholds. This is evidenced by figure 2 where for the same probability outputs, we can have drastically varying  $F_1$ 's depending on the threshold. We realize that by choosing  $F_1$  as evaluation metric, the threshold becomes another highly sensitive hyper-parameter of the model.

We also report on AUC (ROC) since it's more intuitive and independent of thresholds. An AUC of 0.5 represents no information or random predictions which 1.0 is perfect information.



### 3.2 Model 2

A secondary structure focuses on taking the insight that we are essentially modeling some embedding interactions and creating 2 separate learning models, one for each mention. The input features are trimmed to remove features related to the other word (other than mention distance). The final hidden layer of each of these mention-level models are then combined in a final layer dot-product. We then run this final layer through the sigmoid function to come up with a final prediction. The 2 final layers that are inputs to the dot-product could be thought of as domain specific embeddings of each mention. The interaction of these embeddings are then translated to our probability of coreference. Refer to figure 4 for details.

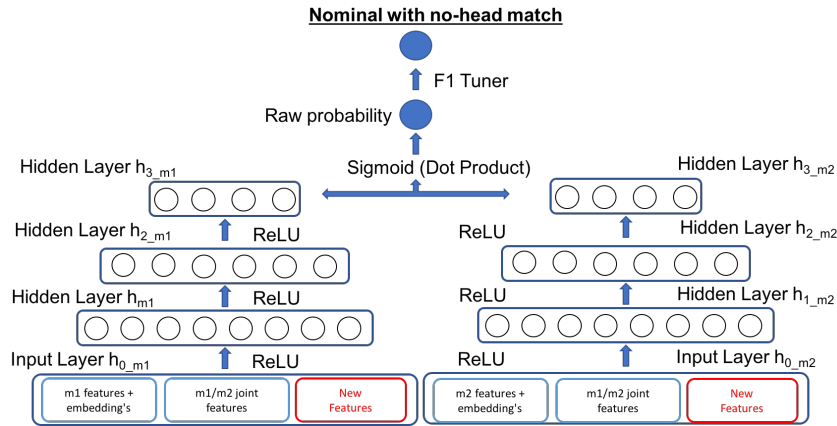


Figure 4: Model 2 architecture learns the candidate antecedent and mention embeddings independently before combining them

### 3.3 Composite Models (Model 3 & Model 4)

*Model 3* is an ensemble of model 2 and model 3, where we simply add the 2 final scores from model 1 and model 2 and then apply a sigmoid function to that output. This model is also trained end-to-end rather than using pre-trained versions of model 1 and 2, otherwise they would end up over-predicting the outcome.

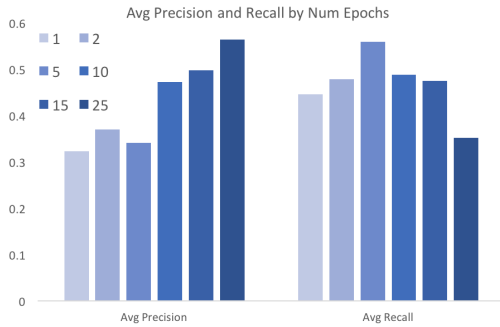
*Model 4* is a modification onto model 3 where instead of simply adding the two models, we take the sigmoid function of each of model 1 and model 2. We then create a new submodel with the same set of input features to train a probabilistic model that puts a weight of  $prb$  on model 1 and a weight of  $1-prb$  on model 2. This probabilistic model is a logistic regression on input features that determines what weight should be given to model 1. This is also trained end-to-end. The idea is to see if there are times, based on input features, that model 1 may be preferable to model 2 and vice versa rather than do a simple ensemble average.

We explore regularization with *dropout*, especially at the first hidden layer. We found that applying dropout at higher levels was detrimental to performance (see figure 5b). Our belief is that this ended up confusing the model rather than making it robust.

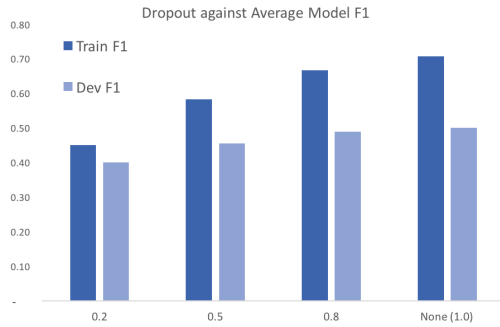
## 4 Experiments, Results & Challenges

### 4.1 Results

In our experiments we compare the baseline model of logistic regression with the four proposed models. The proposed models were tuned using the train and dev set on varying loss function, weights to the loss function, learning rate, neural network architecture, dropout and number of epochs. The best model performance on the dev set for each of the model was chosen and the results produced on the test set which are summarized in table 1. It must be noted that the test set, like the



(a) Average precision and recall for our models as we increase the training epochs.



(b) The train and dev  $F_1$  scores with varying levels of dropout.

training and dev set has been downsampled to 2% negative samples and uses the additional features as described in section 4.2. Model 4, training on 15 epochs with 0.0001 learning rate (super-slow) performs almost 12% better than the baseline logistic regression while almost 8% better than the tuned state-of-the-art model.

Unsurprisingly, as can be seen in figure 5a, given the label imbalance, as we increase the number of epochs, we are able to increase the precision at the expense of recall for all models. The confusion matrix for the best performing model gives a good understanding of the labels and relative performance (Table: 2).

Model Name	Training F1	Dev F1	Test F1
Logistic Regression	.76	.36	.49
Model 1	.76	<b>.53</b>	.53
Model 2	.65	.51	.53
Model 3	.75	.52	.56
Model 4	<b>.82</b>	.52	<b>.61</b>

Table 1: F1 score on 2% downsampled negative examples data

	Predicted NO	Predicted YES
Actual NO	20272	424
Actual YES	559	531

Table 2: Confusion matrix for Model 4 over test dataset.

## 4.2 Impact of Additional Features

Looking at some of the errors like table 3, especially false negatives for co-reference like “the molar-a tooth” and “my taxi-the car”, we added the cosine distance of the word vector embeddings as raw features to the model. In particular, we added the dot-product between the first, last and the head word of the 2 mentions, creating 4 extra features for every mention pair based on the semantic relationship between the two mentions. This drastically improved the model performance, as can be seen in table 4. Simply adding these features to the logistic regression increased the model performance by 4%. The intuition behind adding these features is simple - given concatenated word vectors  $[m_1 m_2]$ , a fully connected neural network can never learn  $m'_1 m_2$  which represents how similar  $m_1$  is to  $m_2$ . Unsurprisingly, a lot of nominals with no head match are synonyms or specific example of a general entity which is now captured by these new features (this is also attempted by model 2).

m1-m2 pairs	P(Without Addn. Features)	P(With Addn. Features)	$\Delta$
the public purse — the public 's money	.0242	.961	+ .93
a molar the tooth	.0043	.740	+ .74
the plunge - the downturn	.401	.742	+ .341

Table 3: Example Wins: Examples of False Negatives without Addn Features turning to True Positive with Additional Features

Model Name	Training $F_1$	Test $F_1$	Test AUC (ROC)
Without Addn. Features	.71	.45	.90
With Addn. Features	.76	.49	.91
$\Delta$	+ .05	+ .04	+ .01

Table 4: The  $F_1$  score improves drastically for logistic regression (actually, for all models) with the addition of the new features.

### 4.3 Challenges

**Class imbalance:** As mentioned previously, the overall true rate of the examples in our dataset is around 0.1%, meaning that models trained on the raw dataset have difficulty learning what distinguishes positive examples from negative examples. We take a multi-step approach to resolving this issue and guiding the model to learn to distinguish positive and negative examples.

We begin by downsampling our training data - keeping all positive examples and only 2% of the negative examples. We also choose weighted cross entropy as our loss function, so that we can give further weight to positive examples through a hyperparameter that we can tune. When evaluating  $F_1$  metrics, we are then forced to define a threshold at which we consider the model probability to be true through two mechanisms.

- When dealing with downsampled data, we use an  $F_1$  tuner that searches over the training set for the threshold that yields the highest  $F_1$ , and apply this to our dev set evaluation.
- When training on downsampled data but evaluating on the full dataset (for our test set) we run our calibration process. This takes the distribution of predictions from the training set and finds the percentile of data such that the model will emit the correct true rate as though the data was not downsampled. For example, if the rate in our downsampled data is 33% and the downsampling was 50%, that means the true rate is around 17%. In this example, calibration of the prediction percentile would be with the goal of yielding true 17% of the time on the training set. The idea behind setting a higher threshold is that if we need to reduce the number of true predictions, we should only keep the highest conviction guesses from the model. In practice, the level of model certainty required to emit a True varied between 97.5% to 99.8% depending on the model.

**Memory and storage:** Because of the size of our dataset, we had challenges creating a full dataset on disk as well as fitting one in memory for training. Mitigating factors included only using half-precision for float and working with downsampled training data with a calibration layer. We also partitioned training data and running over all of the partitions each epoch. This greatly increased the time per epoch and experiment. A possibility to reduce the size of the data is to do in-memory lookups of word vectors, but this would have great slowed down the training speed as each example would require a read-and-process operation rather than a simple read. We could also use a database to store the test/train and dev features.

**Evaluation and prediction reconciliation:** A true, full evaluation would involve creating a full set of CoNLL evaluation metrics on the entire dataset. Instead, we used  $F_1$  on a simple pair-matching model without any reconciliation step. Due to time constraints, we were unable to fully implement the level of evaluation and post-processing that would be directly comparable to other CoNLL papers and submissions.

## 5 Conclusion and Future Work

- **Downstream tasks:** Coreference links between mentions form an upstream part of a coreference system, and nominals with no head match are a fraction of the coreferences. In this paper, while we improve upon the existing state-of-the-art algorithm, we don't know the impact on the downstream tasks.
- **Document level memory:** We are looking at intra-document coreferences in this paper. Memory and attention play a role in human understanding of semantic information, which could be extended to this class of problems using LSTM or RNNs. It can also help weight multiword mentions to better capture the key part of what makes the object a real world reference.
- **A separate layer for singleton mentions:** Identifying which mentions are never repeated could help the model by reducing the number of evaluated candidates for future evaluation. This could be done end-to-end or within a continues framework with a penalty or probability interaction. This is motivated by Recasens et al. and could incorporate more semantic features.

**Acknowledgments:** We want to thank Kevin Clark for supervising this work.

## References

- [BB98] Amit Bagga and Breck Baldwin. Algorithms for scoring coreference chains. In *In The First International Conference on Language Resources and Evaluation Workshop on Linguistics Coreference*, pages 563–566, 1998.
- [CM16a] Kevin Clark and Christopher D. Manning. Deep reinforcement learning for mention-ranking coreference models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2256–2262, 2016.
- [CM16b] Kevin Clark and Christopher D. Manning. Improving coreference resolution by learning entity-level distributed representations. *CoRR*, abs/1606.01323, 2016.
- [Hob86] J Hobbs. Readings in natural language processing. chapter Resolving Pronoun References, pages 339–352. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1986.
- [Hos16] Veronique Hoste. *The Mention-Pair Model*, pages 269–282. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- [HS11] Xianpei Han and Le Sun. A generative entity-mention model for linking entities with knowledge base. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1, HLT '11*, pages 945–954, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics.
- [Luo05] Xiaoqiang Luo. On coreference resolution performance metrics. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT '05*, pages 25–32, Stroudsburg, PA, USA, 2005. Association for Computational Linguistics.
- [PMX<sup>+</sup>12] Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In *Proceedings of the Sixteenth Conference on Computational Natural Language Learning (CooNLL 2012)*, Jeju, Korea, 2012.
- [VBA<sup>+</sup>95] Marc Vilain, John Burger, John Aberdeen, Dennis Connolly, and Lynette Hirschman. A model-theoretic coreference scoring scheme. In *Proceedings of the 6th Conference on Message Understanding, MUC6 '95*, pages 45–52, Stroudsburg, PA, USA, 1995. Association for Computational Linguistics.