

Penn/UMass/CHOP Biocreative II systems

Kuzman Ganchev¹
kuzman@seas.upenn.edu

Koby Crammer¹
crammer@seas.upenn.edu

Fernando Pereira¹
pereira@seas.upenn.edu

Gideon Mann²
gmann@cs.umass.edu

Kedar Bellare²
kedarb@cs.umass.edu

Andrew McCallum²
mccallum@cs.umass.edu

Steven Carroll³
carroll@genome.chop.edu

Yang Jin³
jin@genome.chop.edu

Peter White³
white@genome.chop.edu

¹ Department of Computer and Information Science, University of Pennsylvania, Philadelphia PA

² Department of Computer Science, University of Massachusetts, Amherst MA

³ Division of Oncology, The Children's Hospital of Philadelphia, Philadelphia PA

Abstract

Our team participated in the entity tagging and normalization tasks of Biocreative II. For the entity tagging task, we used a k -best MIRA learning algorithm with lexicons and automatically derived word clusters. MIRA accommodates different training loss functions, which allowed us to exploit gene alternatives in training. We also performed a greedy search over feature templates and the development data, achieving a final F-measure of 86.28%. For the normalization task, we proposed a new specialized on-line learning algorithm and applied it for filtering out false positives from a high recall list of candidates. For normalization we received an F-measure of 69.8%.

Keywords: entity tagging, entity normalization, linear sequence models

1 Introduction

We submitted entity tagging and entity normalization systems. For the entity tagging task, we used a rule-based tokenizer followed by a linear sequence model trained using a 5-best MIRA learning algorithm [7]. MIRA accommodates different training loss functions, which allowed us to exploit alternative labelings in training and tune the loss function for higher performance. We also augmented the feature set with curated lexicons and with automatically derived unsupervised word clustering features and found that their combination gives a more than additive gain.

For the normalization task, we used our entity tagging system trained for high-recall and use this to generate a list of mentions for each abstract. We used a simple matching algorithm to find potential gene aliases for each mention, and a new specialized online learning algorithm to filter out false positives from this initial high-recall set of candidates. Some of the features used by the learning algorithm are based on learning what kinds of changes indicate different aliases for the same gene and what kinds indicate different genes.

2 Entity Tagging

Training and test sentences are first tokenized with a rule-based tokenizer. A linear sequence model assigns one of the three B,I, and O tags to each word, as in [12]. We started from a CRF-based system [4] with features similar to those in earlier work [9]. We made three major changes to the previous system:

1. We trained the model with the k -best MIRA algorithm using a loss function that considers alternative labelings and balances precision and recall (Section 2.1), rather than with CRF training [4].
2. We added word features based on distributional clustering (Section 2.3).
3. We performed feature selection by greedy search over feature templates (Section 2.4).

Together, these changes yielded an overall improvement of 4.3% absolute performance improvement (24% relative error reduction) over the baseline system.

2.1 K -best MIRA and Loss Functions

In what follows, x denotes the generic input sentence, $Y(x)$ denotes the set of possible labelings of x , $Y^+(x)$ the set of correct labelings of x . There is also a distinguished “gold” labeling $y(x) \in Y^+(x)$. For each pair of a sentence x and labeling $y \in Y(x)$, we compute a vector-valued feature representation $f(x, y)$. Given a weight vector w , the score $w \cdot f(x, y)$ ranks possible labelings of x , and we denote by $Y_{k,w}(x)$ the set of k top scoring labelings for x . As with hidden Markov models [11], for suitable feature functions f , $Y_{k,w}(x)$ can be computed efficiently by dynamic programming. A linear sequence model is given by a weight vector w .

The learning portion of our method requires finding a weight vector w that scores the correct labelings of the test data higher than incorrect labelings. We used a k -best version of the MIRA algorithm [2, 7, 6]. This is an online learning algorithm that for each training sentence x updates the weight vector w according to the rule:

$$\begin{aligned} w_{\text{new}} &= \arg \min_w \|w - w_{\text{old}}\| \\ \text{s.t. } &\forall y \in Y_{k,w}(x) : w \cdot f(x, y(x)) - w \cdot f(x, y) \geq L(Y^+(x), y) \end{aligned}$$

where $L(Y^+(x), y)$ is a measure of the loss of labeling y with respect to the set of correct labelings $Y^+(x)$.

The most straightforward and most commonly used loss function is a Hamming loss. This sets the loss of labeling y with respect to the gold labeling $y(x)$ as the number of tokens where the two labelings disagree. The Hamming loss does not make use of the alternative labelings provided in Biocreative.

As we show in section 2.5, a better loss function uses a weighted combination of the number of false positive gene mentions and false negative gene mentions in the sentence. This allows the algorithm to prefer labelings $y \in Y^+(x)$ over $y \notin Y^+(x)$. Notice that the update rule still requires that the gold labeling $y(x)$ has to have at least as high a score as any of the proposed labels $y \in Y_{k,w}(x)$. Our experience showed that getting a high precision is relatively easier than a high recall, so we weigh the number of false negatives higher than the number of false positives.

2.2 Lexicons

In our experiments, we used a number of curated lexicons to create lexicon membership features as in previous work [9]. During the greedy feature search we found that removing some of these actually improved performance as did adding some others. The final set of gene and non-gene lists we used were two gene lists developed in previous work [3, 14], a general biology term list [13], a gene list from from the Human Genome Organization, a list of chemicals extracted from PubChem and a list of diseases from the MeSH ontology. Lexicons we considered using that did not help included a list of common terms, a list of common gene acronyms and a list of amino acid names (among others).

Feature category	Description of final features
Token	The word at position -2...0
Token Suffix	the last two characters of the words at positions -2...1
Token Prefix	the first 2- and 4- characters of the words at positions -2...1
Token n -grams	any 2-, 3- and 4- consecutive characters of current token
Part of speech (POS)	POS of words at position -1...0 as well as the conjunction of the POS of the words at positions -3...0.
Token and POS	the conjunction of POS and token also at position 0
Cluster Features	identity of the cluster of the words at position -1...1 (see Section 2.3)
Lexicon Features	what lexicons match at positions -2...1 (see Section 2.2)

Table 1: The features used in the final system. Position 0 refers to the token at the position we are considering; negative positions are offsets to the left, positive positions are offsets to the right. With the exception of the cluster features, we distinguish features based on their position, so the feature corresponding to the word at position 0 is different from the word at position -1.

2.3 Distributional Clustering

An 85 million word subset of MEDLINE was used to cluster words by bigram language model perplexity into a binary tree [1]. Different depth tree cuts were then applied to produce 5 clustering features at different levels of granularity for each word type in the tree [10]. Thus, for each word type that has been clustered there are 5 different non-independent cluster features generated by the clustering. These additional features are then added the feature function in training and testing. On our development data, adding these features produced a 0.7% improvement in the best system and as much as 1.3% improvement in inferior systems.

2.4 Greedy Search

For the baseline system, we started from a feature set similar to that of our previous work [9] to which we applied feature selection as follows. Features were grouped by feature templates. For example, there are many features for the identity of the current token (one for each token type), but we group all of these into a single identity feature template. Starting with our initial list of feature templates, we repeatedly remove the one whose removal results in the greatest increase in the score of the development data, until no further improvement is possible. Removing just one feature template in this way requires training one model for each removal candidate. Once we cannot improve development data performance, we start adding feature templates from a list of candidates. This resulted in some unexpected additions and non-additions. For example, we found that adding a conjunction of four POS tags helps performance, while adding our list of gene acronyms actually hurts performance. Table 1 describes the features used in the final submission.

Even though there are hundreds of thousands of features, there are only dozens of feature templates, so doing this optimization on the development data does not lead to very severe overfitting: the F-score of the final system on the development data was within 1% of that on unseen data. The initial performance of the baseline system is similar to that of the “public access” system described in the following section. Despite the fact that the baseline system had access to some lexicons, it had poorer feature selection overall.

2.5 Analysis and Results

In development, we split available labeled data into training, development and test sets (80%,10% and 10% respectively). We did not use test data during development, so we could use it to compare a few approaches (Table 2). The baseline system contains some lexicons but did not use the greedy search to

Method Features	CRF			MIRA Hamming			MIRA FP+2FN		
	P	R	F-1	P	R	F-1	P	R	F-1
Baseline	84.3	80.0	82.1	85.3	78.5	81.7	83.2	83.9	83.5
Public access	83.4	79.1	81.2	85.5	78.3	81.8	83.0	86.3	84.6
Clusters	84.6	80.6	82.6	85.7	79.0	82.2	83.1	86.9	85.0
Lexicons	85.0	80.1	82.5	86.7	80.1	83.3	84.4	86.9	85.7
All Resources	85.9	81.3	83.5	87.4	81.6	84.4	85.1	87.7	86.4

Table 2: Precision, recall, and F-measure on held-out data. CRF is a standard CRF trained with prior variance of 10. MIRA Hamming is 5-best MIRA trained with Hamming loss. Mira FP+2FN is 5-best MIRA with a loss counting false positives and twice false negatives. The rows correspond to different feature sets: just features from the data provided, adding word clusters only, adding lexicons only, and all features (all features subject to feature selection). The best system used all the methods for a 4.3% absolute improvement (24% error reduction) over the baseline system.

optimize feature templates, while the public access system did use the greedy search. Clusters, lexicons and all resources describe systems that add to the public access system. These results show that out of the three improvements, the use of MIRA with a tuned loss function and using the alternative labelings yielded the highest performance improvement over the baseline CRF system (3%). Use of alternative labelings and balancing of false positives and false negatives in the loss function yielded around 2-3% performance improvement. Both the automatic cluster features and the curated lexicons gave around 1% F-measure improvement over the baseline. In conjunction with MIRA with the tuned loss function, the lexicons still provide around a 1% improvement over the baseline MIRA system, while the the automatic clusters provide only around a .5% improvement. Surprisingly, in the final system, adding both lexicons and clusters yielded a 1.8% improvement.

3 Gene Normalization

Our approach is in some ways similar to that of [8]. Like them, we first generate a high-recall list of candidate mentions with corresponding aliases and gene IDs, and then use a linear classifier to filter out the false positives. Our system differs in the way that we find an initial list of candidates (Section 3.1,) the learning algorithm (Section 3.3) and in the features that we used to make the decisions (Section 3.2).

3.1 Gene Mentions and Matching

We used the gene tagger described in Section 2. Since our approach for the gene mention task will filter out incorrect gene mentions later, we used a loss function to maximize recall: the loss of a labeling for a sentence was set to the number of false negatives (with respect to the true labeling). This resulted in a recall on the gene mention tagging task of a little over 90%.¹

For each gene mention in the high-recall list, we return the list of gene ID/alias pairs where the alias is the same as the text of the mention after normalization steps that include removal of common words (such as “gene”, “protein”, “mouse”), replacement of digits with the corresponding roman numerals, removal of spaces and dashes, and case conversion. This yields an initial candidate list that has a recall of 77.2% but a precision of only 37.3%. We tried more aggressive matching rules, but we found that this makes the next learning stage too difficult for the small amount of available training data.

¹We cannot estimate this exactly because we used all available labeled data to train the model.

3.2 Filtering Features

To avoid overfitting due to the small training set, we used only 89 features, including the number of candidates competing for the same mention, the number of candidates that agree on the gene ID, and which of “human”, “rat” and “mouse” appears closest to the mention. A set of more complicated but very useful features are based on a learned string string alignment model [5].

The string edit distance model is trained to maximize the probability alignments between aliases of the same gene, while minimizing the probability of alignments between aliases of different genes. For each candidate, the model gives a probability that the gene mention to the alias, which is converted into the following features: the binned value of the probability, the rank of the current candidate among all candidates in the abstract, and the rank of the current candidate among candidates for the same mention. These features resulted in a significant improvement (about 2% in F-measure) in our cross-validation development runs.

3.3 Learning Algorithm

Our model learns to distinguish a candidate (mention-alias pair) containing a true mention of a gene from a false positive. Unfortunately, the training data is incomplete: for each abstract we have only a list of gene IDs. To overcome this problem we created a MIRA-inspired (Section 2.1) online learning algorithm that makes use of the correct gene IDs (given to us), as well as its current predictions to figure out which candidates to require be correct.

More formally, let C be the set of candidates for a particular abstract, C_{top} be the set of current highest scoring candidates for each correct gene ID, and C_{FP} be the set of candidates that correspond to an incorrect gene ID, but were scored above some threshold θ by the current model. Then, our algorithm performs the following update

$$\begin{aligned}
 w_{\text{new}} &= \arg \min_w \|w - w_{\text{old}}\| + \gamma \max_c \xi_c \\
 \text{s.t.} \quad & w \cdot c \geq 1 - \xi_c && \forall c \in C_{\text{top}} \\
 & -w \cdot c \geq 1 - \xi_c && \forall c \in C_{\text{FP}} \\
 & \xi_c \geq 0 && \forall c \in C_{\text{FP}} \cup C_{\text{top}}.
 \end{aligned}$$

The ξ_c are slack variables to account for non-separable data. For our experiments, we used a θ of 0.2 and a γ in the range $0.005 \leq \gamma \leq 0.5$.

4 Acknowledgements

This work was supported in part by the US National Science Foundation under ITR grant EIA-0205448, in part by the Center for Intelligent Information Retrieval, in part by DoD contract #HM1582-06-1-2013, in part by The Central Intelligence Agency, the National Security Agency and National Science Foundation under NSF grant #IIS-0326249, and in part by The Central Intelligence Agency, the National Security Agency and National Science Foundation under NSF grant #IIS-0427594. Any opinions, findings and conclusions or recommendations expressed in this material are the authors’ and do not necessarily reflect those of the sponsor.

References

- [1] Peter Brown, Peter deSouza, Robert Mercer, Vincent Della Pietra, and Jenifer Lai. Class-based n-gram models of natural language. *Computational Linguistics*, 18(4):467–479, 1992.
- [2] Koby Crammer. *Online Learning of Complex Categorical Problems*. PhD thesis, Hebrew Univeristy of Jerusalem, 2004.

- [3] Haw-ren Fang, Kevin Murphy, Yang Jin, Jessica Kim, and Peter White. Human gene name normalization using text matching with automatically extracted synonym dictionaries. In *Proceedings of the HLT-NAACL BioNLP Workshop on Linking Natural Language and Biology*, pages 41–48, New York, New York, June 2006. Association for Computational Linguistics.
- [4] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*, pages 282–289, 2001.
- [5] Andrew McCallum, Kedar Bellare, and Fernando Pereira. A conditional random field for discriminatively-trained finite-state string edit distance. In *Conference on Uncertainty in AI (UAI)*, 2005.
- [6] Ryan McDonald, Koby Crammer, and Fernando Pereira. Flexible text segmentation with structured multilabel classification. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 987–994, Vancouver, British Columbia, Canada, October 2005. Association for Computational Linguistics.
- [7] Ryan McDonald, Koby Crammer, and Fernando Pereira. Online large-margin training of dependency parsers. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 91–98, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics.
- [8] Ryan McDonald, Jay Crim, and Fernando Pereira. Automatically annotating documents with normalized gene lists. In *BMC Bioinformatics*, 2005.
- [9] Ryan McDonald and Fernando Pereira. Identifying gene and protein mentions in text using conditional random fields. In *BMC Bioinformatics*, 2005.
- [10] Scott Miller, Jethran Guinness, and Alex Zamanian. Name tagging with word clusters and discriminative training. In Daniel Marcu Susan Dumais and Salim Roukos, editors, *HLT-NAACL 2004: Main Proceedings*, pages 337–342, Boston, Massachusetts, USA, May 2 - May 7 2004. Association for Computational Linguistics.
- [11] L. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):275–285, 1989.
- [12] Lance Ramshaw and Mitch Marcus. Text chunking using transformation-based learning. In David Yarovsky and Kenneth Church, editors, *Proceedings of the Third Workshop on Very Large Corpora*, pages 82–94, Somerset, New Jersey, 1995. Association for Computational Linguistics.
- [13] Lorraine Tanabe and W. John Wilbur. Tagging gene and protein names in full text articles. In *Proceedings of the Workshop on Natural Language Processing in the Biomedical Domain*, pages 9–13, Philadelphia, July 2002. Association for Computational Linguistics.
- [14] Lorraine K. Tanabe and W. John Wilbur. Generation of a large gene/protein lexicon by morphological pattern analysis. *J. Bioinformatics and Computational Biology*, 1(4):611–626, 2004.