

Optimizing Differentiable Relaxations of Coreference Evaluation Metrics

Phong Le¹ Ivan Titov^{1,2}

¹ILLC, University of Amsterdam

²ILCC, School of Informatics, University of Edinburgh

p.le@uva.nl

ititov@inf.ed.ac.uk

Abstract

Coreference evaluation metrics are hard to optimize directly as they are non-differentiable functions, not easily decomposable into elementary decisions. Consequently, most approaches optimize objectives only indirectly related to the end goal, resulting in suboptimal performance. Instead, we propose a differentiable relaxation that lends itself to gradient-based optimisation, thus bypassing the need for reinforcement learning or heuristic modification of cross-entropy. We show that by modifying the training objective of a competitive neural coreference system, we obtain a substantial gain in performance. This suggests that our approach can be regarded as a viable alternative to using reinforcement learning or more computationally expensive imitation learning.

1 Introduction

Coreference resolution is the task of identifying all mentions which refer to the same entity in a document. It has been shown beneficial in many natural language processing (NLP) applications, including question answering (Hermann et al., 2015) and information extraction (Kehler, 1997), and often regarded as a prerequisite to any text understanding task.

Coreference resolution can be regarded as a clustering problem: each cluster corresponds to a single entity and consists of all its mentions in a given text. Consequently, it is natural to evaluate predicted clusters by comparing them with the ones annotated by human experts, and this is exactly what the standard metrics (e.g., MUC, B³, CEAF) do. In contrast, most state-of-the-art systems are optimized to make individual co-

reference decisions, and such losses are only indirectly related to the metrics.

One way to deal with this challenge is to optimize directly the non-differentiable metrics using reinforcement learning (RL), for example, relying on the REINFORCE policy gradient algorithm (Williams, 1992). However, this approach has not been very successful, which, as suggested by Clark and Manning (2016a), is possibly due to the discrepancy between sampling decisions at training time and choosing the highest ranking ones at test time. A more successful alternative is using a ‘roll-out’ stage to associate cost with possible decisions, as in Clark and Manning (2016a), but it is computationally expensive. Imitation learning (Ma et al., 2014b; Clark and Manning, 2015), though also exploiting metrics, requires access to an expert policy, with exact policies not directly computable for the metrics of interest.

In this work, we aim at combining the best of both worlds by proposing a simple method that can turn popular coreference evaluation metrics into differentiable functions of model parameters. As we show, this function can be computed recursively using scores of individual local decisions, resulting in a simple and efficient estimation procedure. The key idea is to replace non-differentiable indicator functions (e.g. the member function $\mathbb{I}(m \in S)$) with the corresponding posterior probabilities ($p(m \in S)$) computed by the model. Consequently, non-differentiable functions used within the metrics (e.g. the set size function $|S| = \sum_m \mathbb{I}(m \in S)$) become differentiable ($|S|_c = \sum_m p(m \in S)$). Though we assume that the scores of the underlying statistical model can be used to define a probability model, we show that this is not a serious limitation. Specifically, as a baseline we use a probabilistic version of the neural mention-ranking

model of Wiseman et al. (2015b), which on its own outperforms the original one and achieves similar performance to its global version (Wiseman et al., 2016). Importantly when we use the introduced differentiable relaxations in training, we observe a substantial gain in performance over our probabilistic baseline. Interestingly, the absolute improvement (+0.52) is higher than the one reported in Clark and Manning (2016a) using RL (+0.05) and the one using reward rescaling¹ (+0.37). This suggests that our method provides a viable alternative to using RL and reward rescaling.

The outline of our paper is as follows: we introduce our neural resolver baseline and the B³ and LEA metrics in Section 2. Our method to turn a mention ranking resolver into an entity-centric resolver is presented in Section 3, and the proposed differentiable relaxations in Section 4. Section 5 shows our experimental results.

2 Background

2.1 Neural mention ranking

In this section we introduce neural mention ranking, the framework which underpins current state-of-the-art models (Clark and Manning, 2016a). Specifically, we consider a probabilistic version of the method proposed by Wiseman et al. (2015b). In experiments we will use it as our baseline.

Let (m_1, m_2, \dots, m_n) be the list of mentions in a document. For each mention m_i , let $a_i \in \{1, \dots, i\}$ be the index of the mention that m_i is coreferent with (if $a_i = i$, m_i is the first mention of some entity appearing in the document). As standard in coreference resolution literature, we will refer to m_{a_i} as an antecedent of m_i .² Then, in mention ranking the goal is to score antecedents of a mention higher than any other mentions, i.e., if s is the scoring function, we require $s(a_i = j) > s(a_i = k)$ for all j, k such that m_i and m_j are coreferent but m_i and m_k are not.

Let $\phi_a(m_i) \in \mathbb{R}^{d_a}$ and $\phi_p(m_i, m_j) \in \mathbb{R}^{d_p}$ be respectively features of m_i and features of pair

(m_i, m_j) . The scoring function is defined by:

$$s(a_i = j) = \begin{cases} \mathbf{u}^T \begin{bmatrix} \mathbf{h}_a(m_i) \\ \mathbf{h}_p(m_i, m_j) \end{bmatrix} + u_0 & \text{if } j < i \\ \mathbf{v}^T \mathbf{h}_a(m_i) + v_0 & \text{if } j = i \end{cases}$$

where

$$\begin{aligned} \mathbf{h}_a(m_i) &= \tanh(\mathbf{W}_a \phi_a(m_i) + \mathbf{b}_a) \\ \mathbf{h}_p(m_i, m_j) &= \tanh(\mathbf{W}_p \phi_p(m_i, m_j) + \mathbf{b}_p) \end{aligned}$$

and $\mathbf{u}, \mathbf{v}, \mathbf{W}_a, \mathbf{W}_p, \mathbf{b}_a, \mathbf{b}_p$ are real vectors and matrices with proper dimensions, u_0, v_0 are real scalars.

Unlike Wiseman et al. (2015b), where the max-margin loss is used, we define a probabilistic model. The probability³ that m_i and m_j are coreferent is given by

$$p(a_i = j) = \frac{\exp\{s(a_i = j)\}}{\sum_{j'=1}^i \exp\{s(a_i = j')\}} \quad (1)$$

Following Durrett and Klein (2013) we use the following softmax-margin (Gimpel and Smith, 2010) loss function:

$$L(\Theta) = - \sum_{i=1}^n \log \left(\sum_{j \in C(m_i)} p'(a_i = j) \right) + \lambda \|\Theta\|_1,$$

where Θ are model parameters, $C(m_i)$ is the set of the indices of correct antecedents of m_i , and $p'(a_i = j) \propto p(a_i = j)e^{\Delta(j, C(m_i))}$. Δ is a cost function used to manipulate the contribution of different error types to the loss function:

$$\Delta(j, C(m_i)) = \begin{cases} \alpha_1 & \text{if } j \neq i \wedge i \in C(m_i) \\ \alpha_2 & \text{if } j = i \wedge i \notin C(m_i) \\ \alpha_3 & \text{if } j \neq i \wedge j \notin C(m_i) \\ 0 & \text{otherwise} \end{cases}$$

The error types are “false anaphor”, “false new”, “wrong link”, and “no mistake”, respectively. In our experiments, we borrow their values from Durrett and Klein (2013): $(\alpha_1, \alpha_2, \alpha_3) = (0.1, 3, 1)$. In the subsequent discussion, we refer to the loss as *mention-ranking heuristic cross entropy*.

¹Reward rescaling is a technique that computes error values for a heuristic loss function based on the reward difference between the best decision according to the current model and the decision leading to the highest metric score.

²This slightly deviates from the definition of antecedents in linguistics (Crystal, 1997).

³For the sake of readability, we do not explicitly mark in our notation that all the probabilities are conditioned on the document (e.g., the mentions) and dependent on model parameters.

2.2 Evaluation Metrics

We use five most popular metrics⁴,

- MUC (Vilain et al., 1995),
- B³ (Bagga and Baldwin, 1998),
- CEAF_m, CEAF_e (Luo, 2005),
- BLANC (Luo et al., 2014),
- LEA (Moosavi and Strube, 2016).

for evaluation. However, because MUC is the least discriminative metric (Moosavi and Strube, 2016), whereas CEAF is slow to compute, out of the five most popular metrics we incorporate into our loss only B³. In addition, we integrate LEA, as it has been shown to provide a good balance between discriminativity and interpretability.

Let $G = \{G_1, G_2, \dots, G_N\}$ and $S = \{S_1, S_2, \dots, S_M\}$ be the gold-standard entity set and an entity set given by a resolver. Recall that an entity is a set of mentions. The recall and precision of the B³ metric is computed by:

$$R_{B^3} = \frac{\sum_{v=1}^N \sum_{u=1}^M \frac{|G_v \cap S_u|^2}{|G_v|}}{\sum_{v=1}^N |G_v|}$$

$$P_{B^3} = \frac{\sum_{u=1}^M \sum_{v=1}^N \frac{|G_v \cap S_u|^2}{|S_u|}}{\sum_{u=1}^M |S_u|}$$

The LEA metric is computed as:

$$R_{LEA} = \frac{\sum_{v=1}^N (|G_v| \times \sum_{u=1}^M \frac{\text{link}(G_v \cap S_u)}{\text{link}(G_v)})}{\sum_{v=1}^N |G_v|}$$

$$P_{LEA} = \frac{\sum_{u=1}^M (|S_u| \times \sum_{v=1}^N \frac{\text{link}(G_v \cap S_u)}{\text{link}(S_u)})}{\sum_{u=1}^M |S_u|}$$

where $\text{link}(E) = |E| \times (|E| - 1)/2$ is the number of coreference links in entity E . F_β , for both metrics, is defined by:

$$F_\beta = (1 + \beta^2) \frac{P \times R}{\beta^2 P + R}$$

$\beta = 1$ is used in the standard evaluation.

⁴All are implemented in Pradhan et al. (2014), <https://github.com/conll/reference-coreference-scorers>.




Figure 1: For each mention m_u there is a potential entity E_u so that m_u is the first mention in the chain. Computing $p(m_i \in E_u), u < i$ takes into the account all directed paths from m_i to E_u (black arrows). Noting that there is no directed path from any $m_k, k < u$ to E_u because $p(m_k \in E_u) = 0$. (See text for more details.)

3 From mention ranking to entity centrality

Mention-ranking resolvers do not explicitly provide information about entities/clusters which is required by B³ and LEA. We therefore propose a simple solution that can turn a mention-ranking resolver into an entity-centric one.

First note that in a document containing n mentions, there are n potential entities E_1, E_2, \dots, E_n where E_i has m_i as the first mention. Let $p(m_i \in E_u)$ be the probability that mention m_i corresponds to entity E_u . We now show that it can be computed recursively based on $p(a_i = j)$ as follows:

$$p(m_i \in E_u) = \begin{cases} \sum_{j=u}^{i-1} p(a_i = j) \times p(m_j \in E_u) & \text{if } u < i \\ p(a_i = i) & \text{if } u = i \\ 0 & \text{if } u > i \end{cases}$$

In other words, if $u < i$, we consider all possible m_j with which m_i can be coreferent, and which can correspond to entity E_u . If $u = i$, the link to be considered is the m_i 's self-link. And, if $u > i$, the probability is zero, as it is impossible for m_i to be assigned to an entity introduced only later. See Figure 1 for extra information.

We now turn to two crucial questions about this formula:

- Is $p(m_i \in \bullet)$ a valid probability distribution?
- Is it possible for a mention m_u to be mostly anaphoric (i.e. $p(m_u \in E_u)$ is low) but for the corresponding cluster E_u to be highly

probable (i.e. $p(m_i \in E_u)$ is high for some i)?

The first question is answered in Proposition 1. The second question is important because, intuitively, when a mention m_u is anaphoric, the potential entity E_u does not exist. We will show that the answer is “No” by proving in Proposition 2 that the probability that m_u is anaphoric is always higher than any probability that $m_i, i > u$ refers to E_u .

Proposition 1. $p(m_i \in \bullet)$ is a valid probability distribution, i.e., $\sum_{u=1}^n p(m_i \in E_u) = 1$, for all $i = 1, \dots, n$.

Proof. We prove this proposition by induction.

Basis: it is obvious that $\sum_{u=1}^n p(m_1 \in E_u) = p(a_1 = 1) = 1$.

Assume that $\sum_{u=1}^n p(m_j \in E_u) = 1$ for all $j < i$. Then,

$$\begin{aligned} & \sum_{u=1}^{i-1} p(m_i \in E_u) \\ &= \sum_{u=1}^{i-1} \sum_{j=u}^{i-1} p(a_i = j) \times p(m_j \in E_u) \end{aligned}$$

Because $p(m_j \in E_u) = 0$ for all $j < u$, this expression is equal to

$$\begin{aligned} & \sum_{u=1}^{i-1} \sum_{j=1}^{i-1} p(a_i = j) \times p(m_j \in E_u) \\ &= \sum_{j=1}^{i-1} p(a_i = j) \times \sum_{u=1}^{i-1} p(m_j \in E_u) \\ &= \sum_{j=1}^{i-1} p(a_i = j) \end{aligned}$$

Therefore,

$$\sum_{u=1}^n p(m_i \in E_u) = \sum_{j=1}^{i-1} p(a_i = j) + p(a_i = i) = 1$$

(according to Equation 1). \square

Proposition 2. $p(m_i \in E_u) \leq p(m_u \in E_u)$ for all $i > u$.

Proof. We prove this proposition by induction.

Basis: for $i = u + 1$,

$$\begin{aligned} p(m_{u+1} \in E_u) &= p(a_{u+1} = u) \times p(m_u \in E_u) \\ &\leq p(m_u \in E_u) \end{aligned}$$

Assume that $p(m_j \in E_u) \leq p(m_u \in E_u)$ for all $j \geq u$ and $j < i$. Then

$$\begin{aligned} p(m_i \in E_u) &= \sum_{j=u}^{i-1} p(a_i = j) \times p(m_j \in E_u) \\ &\leq \sum_{j=u}^{i-1} p(a_i = j) \times p(m_u \in E_u) \\ &\leq p(m_u \in E_u) \times \sum_{j=1}^i p(a_i = j) \\ &= p(m_u \in E_u) \end{aligned}$$

\square

3.1 Entity-centric heuristic cross entropy loss

Having $p(m_i \in E_u)$ computed, we can consider coreference resolution as a multiclass prediction problem. An entity-centric heuristic cross entropy loss is thus given below:

$$L_{ec}(\Theta) = - \sum_{i=1}^n \log p'(m_i \in E_{e(m_i)}) + \lambda \|\Theta\|_1$$

where $E_{e(m_i)}$ is the correct entity that m_i belongs to, $p'(m_i \in E_u) \propto p(m_i \in E_u) e^{\Gamma(u, e(m_i))}$. Similar to Δ in the mention-ranking heuristic loss in Section 2.1, Γ is a cost function used to manipulate the contribution of the four different error types (“false anaphor”, “false new”, “wrong link”, and “no mistake”):

$$\Gamma(u, e(m_i)) = \begin{cases} \gamma_1 & \text{if } u \neq i \wedge e(m_i) = i \\ \gamma_2 & \text{if } u = i \wedge e(m_i) \neq i \\ \gamma_3 & \text{if } u \neq e(m_i) \wedge u \neq i \wedge e(m_i) \neq i \\ 0 & \text{otherwise} \end{cases}$$

4 From non-differentiable metrics to differentiable losses

There are two functions used in computing B^3 and LEA: the set size function $|\cdot|$ and the link function $link(\cdot)$. Because both of them are non-differentiable, the two metrics are non-differentiable. We thus need to make these two functions differentiable.

There are two remarks. Firstly, both functions can be computed using the indicator function




Figure 2: Softmax $\frac{\exp\{\pi_i/T\}}{\sum_j \exp\{\pi_j/T\}}$ with different values of T . The softmax becomes more peaky when the value of T gets smaller. As $T \rightarrow 0$ the softmax converges to the indicator function that chooses $\arg \max_i \pi_i$.

$\mathbb{I}(m_i \in S_u)$:

$$|S_u| = \sum_{i=1}^n \mathbb{I}(m_i \in S_u)$$

$$link(S_u) = \sum_{j<i} \mathbb{I}(m_i \in S_u) \times \mathbb{I}(m_j \in S_u)$$

Secondly, given $\pi_{i,u} = \log p(m_i \in S_u)$, the indicator function $\mathbb{I}(m_i \in S_{u^*})$, $u^* = \arg \max_u p(m_i \in S_u)$ is the converging point of the following softmax as $T \rightarrow 0$ (see Figure 2):

$$p(m_i \in S_u; T) = \frac{\exp\{\pi_{i,u}/T\}}{\sum_v \exp\{\pi_{i,v}/T\}}$$

where T is called *temperature* (Kirkpatrick et al., 1983).

Therefore, we propose to represent each S_u as a *soft-cluster*:

$$S_u = \{p(m_1 \in E_u; T), \dots, p(m_n \in E_u; T)\}$$

where, as defined in Section 3, E_u is the potential entity that has m_u as the first mention. Replacing the indicator function $\mathbb{I}(m_i \in S_u)$ by the probability distribution $p(m_i \in E_u; T)$, we then have a *differentiable* version for the set size function and the link function:

$$|S_u|_d = \sum_{i=1}^n p(m_i \in E_u; T)$$

$$link_d(S_u) = \sum_{j<i} p(m_i \in E_u; T) \times p(m_j \in E_u; T)$$

$|G_v \cap S_u|_d$ and $link_d(G_v \cap S_u)$ are computed similarly with the constraint that only mentions in G_v are taken into account. Plugging these functions into precision and recall of B^3 and LEA in Section 2.2, we obtain differentiable \hat{F}_{β, B^3} and $\hat{F}_{\beta, LEA}$, which are then used in two loss functions:

$$L_{\beta, B^3}(\Theta; T) = -\hat{F}_{\beta, B^3}(\Theta; T) + \lambda \|\Theta\|_1$$

$$L_{\beta, LEA}(\Theta; T) = -\hat{F}_{\beta, LEA}(\Theta; T) + \lambda \|\Theta\|_1$$

where λ is the hyper-parameter of the L_1 regularization terms.

It is worth noting that, as $T \rightarrow 0$, $\hat{F}_{\beta, B^3} \rightarrow F_{\beta, B^3}$ and $\hat{F}_{\beta, LEA} \rightarrow F_{\beta, LEA}$.⁵ Therefore, when training a model with the proposed losses, we can start at a high temperature (e.g., $T = 1$) and anneal to a small but non-zero temperature. However, in our experiments we fix $T = 1$. Annealing is left for future work.

5 Experiments

We now demonstrate how to use the proposed differentiable B^3 and LEA to train a coreference resolver. The source code and trained models are available at https://github.com/lephong/diffmetric_coref.

Setup

We run experiments on the English portion of CoNLL 2012 data (Pradhan et al., 2012) which consists of 3,492 documents in various domains and formats. The split provided in the CoNLL 2012 shared task is used. In all our resolvers, we use not the original features of Wiseman et al. (2015b) but their slight modification described in Wiseman et al. (2016) (section 6.1).⁶

Resolvers

We build following baseline and three resolvers:

- baseline: the resolver presented in Section 2.1. We use the identical configuration as in Wiseman et al. (2016): $\mathbf{W}_a \in \mathbb{R}^{200 \times d_a}$, $\mathbf{W}_p \in \mathbb{R}^{700 \times d_p}$, $\lambda = 10^{-6}$ (where d_a, d_p are respectively the numbers of mention features and pair-wise features). We also employ their pretraining methodology.

⁵We can easily prove this using the algebraic limit theorem.

⁶https://github.com/swiseman/nn_coref/

- L_{ec} : the resolver using the entity-centric cross entropy loss introduced in Section 3.1. We set $(\gamma_1, \gamma_2, \gamma_3) = (\alpha_1, \alpha_2, \alpha_3) = (0.1, 3, 1)$.
- L_{β, B^3} and $L_{\beta, LEA}$: the resolvers using the losses proposed in Section 4. β is tuned on the development set by trying each value in $\{\sqrt{0.8}, 1, \sqrt{1.2}, \sqrt{1.4}, \sqrt{1.6}, \sqrt{1.8}, 1.5, 2\}$.

To train these resolvers we use AdaGrad (Duchi et al., 2011) to minimize their loss functions with the learning rate tuned on the development set and with one-document mini-batches. Note that we use the baseline as the initialization point to train the other three resolvers.

5.1 Results

We firstly compare our resolvers against Wiseman et al. (2015b) and Wiseman et al. (2016). Results are shown in the first half of Table 1. Our baseline surpasses Wiseman et al. (2015b). It is likely due to using features from Wiseman et al. (2016). Using the entity-centric heuristic cross entropy loss and the relaxations are clearly beneficial: L_{ec} is slightly better than our baseline and on par with the global model of Wiseman et al. (2016). $L_{\beta=1, B^3}$, $L_{\beta=1, LEA}$ outperform the baseline, the global model of Wiseman et al. (2016), and L_{ec} . However, the best values of β are $\sqrt{1.4}$, $\sqrt{1.8}$ respectively for L_{β, B^3} , and $L_{\beta, LEA}$. Among these resolvers, $L_{\beta=\sqrt{1.8}, LEA}$ achieves the highest F_1 scores across all the metrics except BLANC.

When comparing to Clark and Manning (2016a) (the second half of Table 1), we can see that the absolute improvement over the baselines (i.e. ‘heuristic loss’ for them and the heuristic cross entropy loss for us) is higher than that of reward rescaling but with much shorter training time: +0.37 (7 days⁷) and +0.52 (15 hours) on the CoNLL metric for Clark and Manning (2016a) and ours, respectively. It is worth noting that our absolute scores are weaker than these of Clark and Manning (2016a), as they build on top of a similar but stronger mention-ranking baseline, which employs deeper neural networks and requires a much larger number of epochs to train (300 epochs, including pretraining). For the purpose of illustrating the proposed losses, we started with a simpler model by Wiseman et al. (2015b) which requires

⁷As reported in <https://github.com/clarkkev/deep-coref>

(a) [...] that 13[the virus] could mutate [...] /. In fact some health experts say 17[it]^{13*}_{17,17} ’s just a matter of time [...]

(b) Walk a mile in 157[our] shoes that ’s all I have to say because anybody who works in a nursing home will very quickly learn that these are very fragile patients /. 165[We]¹⁵⁷_{165*,157} did the very best 167[we]¹⁶⁵_{165,165} could in these situations [...]

Figure 3: Example predictions: the subscript before a mention is its index. The superscript / subscript after a mention indicates the antecedent predicted by the baseline / $L_{\beta=1, B^3}$, $L_{\beta=\sqrt{1.4}, B^3}$. Mentions with the same color are true coreferents. “*”s mark *incorrect* decisions.

a much smaller number of epochs, thus faster, to train (20 epochs, including pretraining).

5.2 Analysis

Table 2 shows the breakdown of errors made by the baseline and our resolvers on the development set. The proposed resolvers make fewer “false anaphor” and “wrong link” errors but more “false new” errors compared to the baseline. This suggests that loss optimization prevents over-clustering, driving the precision up: when antecedents are difficult to detect, the self-link (i.e., $a_i = i$) is chosen. When β increases, they make more “false anaphor” and “wrong link” errors but less “false new” errors.

In Figure 3(a) the baseline, but not $L_{\beta=1, B^3}$ nor $L_{\beta=\sqrt{1.4}, B^3}$, mistakenly links 17[it] with 13[the virus]. Under-clustering, on the other hand, is a problem for our resolvers with $\beta = 1$: in example (b), $L_{\beta=1, B^3}$ missed 165[We]. This behaviour results in a reduced recall but the recall is not damaged severely, as we still obtain a better F_1 score. We conjecture that this behaviour is a consequence of using the F_1 score in the objective, and, if undesirable, F_β with $\beta > 1$ can be used instead. For instance, also in Figure 3, $L_{\beta=\sqrt{1.4}, B^3}$ correctly detects 17[it] as non-anaphoric and links 165[We] with 157[our].

Figure 4 shows recall, precision, F_1 (average of MUC, B^3 , $CEAF_e$), on the development set when training with L_{β, B^3} and $L_{\beta, LEA}$. As expected, higher values of β yield lower precisions but higher recalls. In contrast, F_1 increases until

	MUC	B ³	CEAF _m	CEAF _e	BLANC	LEA	CoNLL
Wiseman et al. (2015b)	72.60	60.52	-	57.05	-	-	63.39
Wiseman et al. (2016)	73.42	61.50	-	57.70	-	-	64.21
Our proposals							
baseline (heuristic loss)	73.22	61.44	65.12	57.74	62.16	57.52	64.13
L_{ec}	73.2	61.75	65.77	57.8	63.3	57.89	64.25
$L_{\beta=1, B^3}$	73.37	61.94	65.79	58.22	63.19	58.06	64.51
$L_{\beta=\sqrt{1.4}, B^3}$	73.48	61.99	65.9	58.36	63.1	58.13	64.61
$L_{\beta=1, LEA}$	73.3	61.88	65.69	57.99	63.27	58.03	64.39
$L_{\beta=\sqrt{1.8}, LEA}$	73.53	62.04	65.95	58.41	63.09	58.18	64.66
Clark and Manning (2016a)							
baseline (heuristic loss)	74.65	63.03	-	58.40	-	-	65.36
REINFORCE	74.48	63.09	-	58.67	-	-	65.41
Reward Rescaling	74.56	63.40	-	59.23	-	-	65.73

Table 1: Results (F_1) on CoNLL 2012 test set. CoNLL is the average of MUC, B³, and CEAF_e.

	Non-Anaphoric (FA)			Anaphoric (FN + WL)		
	Proper	Nominal	Pronom.	Proper	Nominal	Pronom.
baseline	630	714	1051	374 + 190	821 + 238	347 + 779
L_{ec}	529	609	904	438 + 182	924 + 220	476 + 740
$L_{\beta=1, B^3}$	545	559	883	433 + 172	951 + 192	457 + 761
$L_{\beta=\sqrt{1.4}, B^3}$	557	564	926	426 + 178	941 + 194	431 + 766
$L_{\beta=1, LEA}$	513	547	843	456 + 170	960 + 191	513 + 740
$L_{\beta=\sqrt{1.8}, LEA}$	577	591	1001	416 + 176	919 + 198	358 + 790

Table 2: Number of: “false anaphor” (FA, a non-anaphoric mention marked as anaphoric), “false new” (FN, an anaphoric mention marked as non-anaphoric), and “wrong link” (WL, an anaphoric mention is linked to a wrong antecedent) errors on the development set.

reaching the highest point when $\beta = \sqrt{1.4} \approx 1.18$ for L_{β, B^3} ($\beta = \sqrt{1.8} \approx 1.34$ for $L_{\beta, LEA}$), it then decreases gradually.

5.3 Discussion

Because the resolvers are evaluated on F_1 score metrics, it should be that L_{β, B^3} and $L_{\beta, LEA}$ perform the best with $\beta = 1$. Figure 4 and Table 1 however do not confirm that: β should be set with values a little bit larger than 1. There are two hypotheses. First, the statistical difference between the training set and the development set leads to the case that the optimal β on one set can be sub-optimal on the other set. Second, in our experiments we fix $T = 1$, meaning that the relaxations might not be close to the true evaluation metrics enough. Our future work, to confirm/reject this, is to use annealing, i.e., gradually decreasing T down to (but larger than) 0.

Table 1 shows that the difference between L_{β, B^3} and $L_{\beta, LEA}$ in terms of accuracy is not substan-

tial (although the latter is slightly better than the former). However, one should expect that L_{β, B^3} would outperform $L_{\beta, LEA}$ on B³ metric while it would be the other way around on LEA metric. It turns out that, B³ and LEA behave quite similarly in non-extreme cases. We can see that in Figure 2, 4, 5, 6, 7 in Moosavi and Strube (2016).

6 Related work

Mention ranking and entity centrality are two main streams in the coreference resolution literature. Mention ranking (Denis and Baldridge, 2007; Durrett and Klein, 2013; Martschat and Strube, 2015; Wiseman et al., 2015a) considers local and independent decisions when choosing a correct antecedent for a mention. This approach is computationally efficient and currently dominant with state-of-the-art performance (Wiseman et al., 2016; Clark and Manning, 2016a). Wiseman et al. (2015b) propose to use simple neural

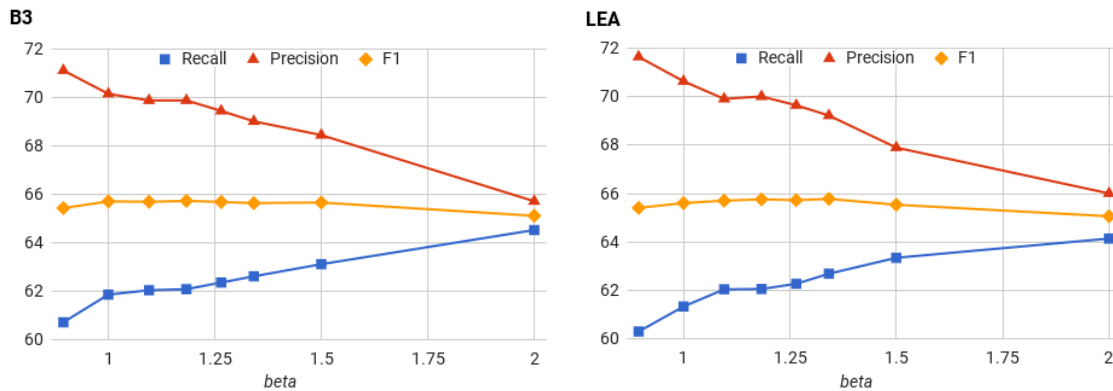


Figure 4: Recall, precision, F_1 (average of MUC, B^3 , $CEAF_e$), on the development set when training with L_{β, B^3} (left) and $L_{\beta, LEA}$ (right). Higher values of β yield lower precisions but higher recalls.

networks to compute mention ranking scores and to use a heuristic loss to train the model. Wiseman et al. (2016) extend this by employing LSTMs to compute mention-chain representations which are then used to compute ranking scores. They call these representations *global features*. Clark and Manning (2016a) build a similar resolver as in Wiseman et al. (2015b) but much stronger thanks to deeper neural networks and “better mention detection, more effective, hyperparameters, and more epochs of training”. Furthermore, using reward rescaling they achieve the best performance in the literature on the English and Chinese portions of the CoNLL 2012 dataset. Our work is built upon mention ranking by turning a mention-ranking model into an entity-centric one. It is worth noting that although we use the model proposed by Wiseman et al. (2015b), any mention-ranking models can be employed.

Entity centrality (Wellner and McCallum, 2003; Poon and Domingos, 2008; Haghighi and Klein, 2010; Ma et al., 2014a; Clark and Manning, 2016b), on the other hand, incorporates entity-level information to solve the problem. The approach can be top-down as in Haghighi and Klein (2010) where they propose a generative model. It can also be bottom-up by merging smaller clusters into bigger ones as in Clark and Manning (2016b). The method proposed by Ma et al. (2014a) greedily and incrementally adds mentions to previously built clusters using a prune-and-score technique. Importantly, employing imitation learning these two methods can optimize the resolvers directly on evaluation metrics. Our work is similar to Ma et al. (2014a) in the sense that our resolvers incrementally add mentions to previously built clusters.

However, different from both Ma et al. (2014a); Clark and Manning (2016b), our resolvers do not use any discrete decisions (e.g., merge operations). Instead, they seamlessly compute the probability that a mention refers to an entity from mention-ranking probabilities, and are optimized on differentiable relaxations of evaluation metrics.

Using differentiable relaxations of evaluation metrics as in our work is related to a line of research in reinforcement learning where a non-differentiable action-value function is replaced by a differentiable critic (Sutton et al., 1999; Silver et al., 2014). The critic is trained so that it is as close to the true action-value function as possible. This technique is applied to machine translation (Gu et al., 2017) where evaluation metrics (e.g., BLUE) are non-differentiable. A disadvantage of using critics is that there is no guarantee that the critic converges to the true evaluation metric given finite training data. In contrast, our differentiable relaxations do not need to train, and the convergence is guaranteed as $T \rightarrow 0$.

7 Conclusions

We have proposed

- a method for turning any mention-ranking resolver into an entity-centric one by using a recursive formula to combine scores of individual local decisions, and
- differentiable relaxations for two coreference evaluation metrics, B^3 and LEA.

Experimental results show that our approach outperforms the resolver by Wiseman et al. (2016), and gains a higher improvement over the baseline

than that of Clark and Manning (2016a) but with much shorter training time.

Acknowledgments

We would like to thank Raquel Fernández, Wilker Aziz, Nafise Sadat Moosavi, and anonymous reviewers for their suggestions and comments. The project was supported by the European Research Council (ERC StG BroadSem 678254), the Dutch National Science Foundation (NWO VIDI 639.022.518) and an Amazon Web Services (AWS) grant.

References

- Amit Bagga and Breck Baldwin. 1998. Algorithms for scoring coreference chains. In *The first international conference on language resources and evaluation workshop on linguistics coreference*. volume 1, pages 563–566.
- Kevin Clark and Christopher D. Manning. 2016a. [Deep reinforcement learning for mention-ranking coreference models](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, pages 2256–2262. <https://aclweb.org/anthology/D16-1245>.
- Kevin Clark and Christopher D. Manning. 2016b. [Improving coreference resolution by learning entity-level distributed representations](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 643–653. <http://www.aclweb.org/anthology/P16-1061>.
- Kevin Clark and D. Christopher Manning. 2015. [Entity-centric coreference resolution with model stacking](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, pages 1405–1415. <https://doi.org/10.3115/v1/P15-1136>.
- David Crystal. 1997. *Dictionary of Linguistics and Phonetics*. Blackwell Publishers, Cambridge, MA.
- Pascal Denis and Jason Baldridge. 2007. A ranking approach to pronoun resolution. In *IJCAI*. volume 158821593.
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12(Jul):2121–2159.
- Greg Durrett and Dan Klein. 2013. [Easy victories and uphill battles in coreference resolution](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pages 1971–1982. <http://aclweb.org/anthology/D13-1203>.
- Kevin Gimpel and Noah A. Smith. 2010. [Softmax-margin crfs: Training log-linear models with cost functions](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Los Angeles, California, pages 733–736. <http://www.aclweb.org/anthology/N10-1112>.
- Jiatao Gu, Kyunghyun Cho, and Victor OK Li. 2017. Trainable greedy decoding for neural machine translation. *arXiv preprint arXiv:1702.02429*.
- Aria Haghighi and Dan Klein. 2010. Coreference resolution in a modular, entity-centered model. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, pages 385–393.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in Neural Information Processing Systems (NIPS)*. <http://arxiv.org/abs/1506.03340>.
- Andrew Kehler. 1997. *Second Conference on Empirical Methods in Natural Language Processing*, chapter Probabilistic Coreference in Information Extraction. <http://aclweb.org/anthology/W97-0319>.
- Scott Kirkpatrick, C Daniel Gelatt, Mario P Vecchi, et al. 1983. Optimization by simulated annealing. *science* 220(4598):671–680.
- Xiaoqiang Luo. 2005. [On coreference resolution performance metrics](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*. <http://aclweb.org/anthology/H05-1004>.
- Xiaoqiang Luo, Sameer Pradhan, Marta Recasens, and Eduard Hovy. 2014. [An extension of blanc to system mentions](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, pages 24–29. <https://doi.org/10.3115/v1/P14-2005>.
- Chao Ma, Janardhan Rao Doppa, J. Walker Orr, Prashanth Mannem, Xiaoli Fern, Tom Dietterich, and Prasad Tadepalli. 2014a. [Prune-and-score: Learning for greedy coreference resolution](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, pages 2115–2126. <http://www.aclweb.org/anthology/D14-1225>.

- Chao Ma, Rao Janardhan Doppa, Walker J. Orr, Prashanth Mannem, Xiaoli Fern, Tom Dietterich, and Prasad Tadepalli. 2014b. [Prune-and-score: Learning for greedy coreference resolution](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, pages 2115–2126. <https://doi.org/10.3115/v1/D14-1225>.
- Sebastian Martschat and Michael Strube. 2015. Latent structures for coreference resolution. *Transactions of the Association for Computational Linguistics* 3:405–418.
- Nafise Sadat Moosavi and Michael Strube. 2016. [Which coreference evaluation metric do you trust? a proposal for a link-based entity aware metric](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 632–642. <http://www.aclweb.org/anthology/P16-1060>.
- Hoifung Poon and Pedro Domingos. 2008. Joint unsupervised coreference resolution with markov logic. In *Proceedings of the conference on empirical methods in natural language processing*. Association for Computational Linguistics, pages 650–659.
- Sameer Pradhan, Xiaoqiang Luo, Marta Recasens, Eduard Hovy, Vincent Ng, and Michael Strube. 2014. [Scoring coreference partitions of predicted mentions: A reference implementation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Baltimore, Maryland, pages 30–35. <http://www.aclweb.org/anthology/P14-2006>.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. *Joint Conference on EMNLP and CoNLL - Shared Task*, Association for Computational Linguistics, chapter CoNLL-2012 Shared Task: Modeling Multilingual Unrestricted Coreference in OntoNotes, pages 1–40. <http://aclweb.org/anthology/W12-4501>.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin A. Riedmiller. 2014. [Deterministic policy gradient algorithms](#). In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*. pages 387–395. <http://jmlr.org/proceedings/papers/v32/silver14.html>.
- Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. 1999. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*. volume 99, pages 1057–1063.
- Marc Vilain, John Burger, John Aberdeen, Dennis Connolly, and Lynette Hirschman. 1995. [A model-theoretic coreference scoring scheme](#). In *Sixth Message Understanding Conference (MUC-6): Proceedings of a Conference Held in Columbia, Maryland, November 6-8, 1995*. <http://aclweb.org/anthology/M95-1005>.
- B Wellner and A McCallum. 2003. Towards conditional models of identity uncertainty with application to proper noun coreference. In *IJCAI Workshop on Information Integration and the Web*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.
- Sam Wiseman, Alexander M Rush, Stuart M Shieber, and Jason Weston. 2015a. Learning anaphoricity and antecedent ranking features for coreference resolution. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, volume 1, pages 92–100.
- Sam Wiseman, M. Alexander Rush, and M. Stuart Shieber. 2016. [Learning global features for coreference resolution](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, pages 994–1004. <https://doi.org/10.18653/v1/N16-1114>.
- Sam Wiseman, M. Alexander Rush, Stuart Shieber, and Jason Weston. 2015b. [Learning anaphoricity and antecedent ranking features for coreference resolution](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, pages 1416–1426. <https://doi.org/10.3115/v1/P15-1137>.