



HAL
open science

Synapse at CAp 2017 NER challenge: Fasttext CRF

Damien Sileo, Camille Pradel, Philippe Muller, Tim van de Cruys

► To cite this version:

Damien Sileo, Camille Pradel, Philippe Muller, Tim van de Cruys. Synapse at CAp 2017 NER challenge: Fasttext CRF. 19e Conference francophone sur l'Apprentissage Automatique (CAp 2017), Jun 2017, Grenoble, France. pp.1-5. hal-02879724

HAL Id: hal-02879724

<https://hal.science/hal-02879724v1>

Submitted on 24 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is an author's version published in:
<http://oatao.univ-toulouse.fr/22176>

To cite this version: Sileo, Damien and Pradel, Camille and Muller, Philippe and Van de Cruys, Tim *Synapse at CAp 2017 NER challenge: Fasttext CRF*. (2017) In: 19e Conference francophone sur l'Apprentissage Automatique (CAp 2017), 28 June 2017 - 30 June 2017 (Grenoble, France).

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Synapse at CAp 2017 NER challenge: Fasttext CRF

Damien Sileo^{1,2}, Camille Pradel^{1,*}, Philippe Muller^{2,*}, Tim Van de Cruys^{2,*}

¹Synapse Développement
²IRIT, Université Paul Sabatier
*Equal contributions

September 15, 2017

Abstract

We present our system for the CAp 2017 NER challenge [LPB⁺17] which is about named entity recognition on French tweets. Our system leverages unsupervised learning on a larger dataset of French tweets to learn features feeding a CRF model. It was ranked first without using any gazetteer or structured external data, with an F-measure of 58.89%. To the best of our knowledge, it is the first system to use *fasttext* [BGJM16] embeddings (which include subword representations) and an embedding-based sentence representation for NER.

Keywords: Named entity recognition, fasttext, CRF, unsupervised learning, word vectors

1 Introduction

Named-Entity Recognition (NER) is the task of detecting word segments denoting particular instances such as persons, locations or quantities. It can be used to ground knowledge available in texts. While NER can achieve near-human performance [NNN98], it is still a challenging task on noisy texts such as tweets [RCME11] scarce labels, especially when few linguistic resources are available. Those difficulties are all present in the CAp NER challenge.

A promising approach is using unsupervised learning to get meaningful representations of words and sentences. *Fasttext* [BGJM16] seems a particularly useful unsupervised learning method for named entity recognition since it is based on the skipgram model which is able to capture substantive knowledge about words while incorporating morphology information, a crucial aspect for NER. We will describe three methods for using such embeddings along with a CRF sequence model, and we will also present a simple ensemble method for structured prediction (section




Figure 1: overview of our system

2). Next, we will show the performance of our model and an interpretation of its results (section 4).

2 Model

Figure 1 shows an overview of our model. This section will detail the components of the system.

2.1 CRF

The core of our model is Conditional random fields (CRF) [SM11], a structured prediction framework widely used in NER tasks. It can model the probabilities of a tag sequence $y_1 \dots y_n$ given a sequence of words $x_1 \dots x_n$.

We use the linear chain CRF restriction where the sequences are modeled with the probability of transitions between consecutive labels.

$$P(y|x) = \frac{1}{Z(x)} \prod_{i=1}^n \exp\left(\sum_j \theta_j f(y_{i-1}, y_i, x, i)\right) \quad (1)$$

f yields a feature vector, θ is a weight vector, and Z is a normalization factor in order to ensure a probability distribution. CRFs allow for non greedy optimization for learning sequence prediction and allows for much flexibility when defining the feature vector $f(y_{i-1}, y_i, x, i)$. Furthermore, we can add a prior on the learned weights θ for regularization purposes. The likelihood of the training data can be

Table 1: word-level handcrafted features

feature	context
word (lowercased)	✓
word length	
length 1 prefix	✓
length 2 prefix	
length 1 suffix	✓
length 2 suffix	
is_upper	✓
is_title	✓
position	
word uppercase proportion	
word uppercase proportion*word length	
is_emoji	
hyphenation	
POS tag	
is_quote	✓
beginning of sentence	
end of sentence	

optimized using gradient descent. We chose f to yield two sets of features that are concatenated: handcrafted features and *fasttext* embedding-based features.

2.2 Handcrafted features

Table 1 shows the handcrafted features we used. The context columns specifies whether or not a feature was also used with respect to the adjacent words.

The emoji ¹ library was used for emoji detection, and we used the Treetagger [Sch94] POS tagger.

2.3 Fasttext features

Fasttext skipgram is based on the *word2vec* skipgram model [MCCD13], where word representations are learned so that they optimize a task of predicting context words. The main difference is that the representation h_w of a word w is not only u_w , the representation of its symbol. It is augmented with the sum of the representations of its subword units $u_g, g \in \mathcal{G}_w$:

$$h_w = u_w + \sum_{g \in \mathcal{G}_w} u_g \quad (2)$$

\mathcal{G}_w encompasses some character n-grams that w contains, provided they are frequent enough and of a desirable length. Morphology of w is thus taken in account in the representation of h_w even though the order of n-grams is ignored.

h_w can directly be used as a word level feature. However, [GCWL14] showed that CRFs work better with discrete

¹<https://pypi.python.org/pypi/emoji/>

features, so we also use a clustering-based representation. Several approaches [Ahm13, Sie15, DGG17, GCWL14] use word embeddings for named entity recognition.

2.3.1 Clustering *fasttext* features

We cluster the *fasttext* representations of unique words in train and test tweets using a Gaussian Mixture Model (GMM), and feed the vector of probabilities assignments as word-level feature to the CRF. GMM clusters latent space to maximize the likelihood of the training data assuming that it is modeled by a mixture of gaussian.

2.3.2 Sentence representation

We also use the average of word representations in a tweet as a sentence level feature. It is a simple way to provide a global context even though a linear model will not exploit this information thoroughly.

2.4 Ensemble method

We ensemble different models using a voting rule. We train N systems, each time training an new *fasttext* model. This is the only variation between models, but different embeddings can influence the parameters learned with respect to handcrafted features. We then select the best prediction by picking the most frequent labeling sequence predicted for each tweet by the N systems.

3 Experimental settings

Test/train data are from CAp NER 2017 data includes french labeled tweets with 13 kinds of segments and IOB format. Further details can be found in [LPB⁺17]. We used *Crfsuite* [Oka07] through its *sklearn-crfsuite* python bindings ² which follows the *sklearn* API and allows for better development speed. The original implementation of *fasttext* [BGJM16] was used through its python bindings. ³

3.1 Additional data

To learn *fasttext* word representations, we used tweets from the *OSIRIM* ⁴ platform at IRIT, where 1% of the total feed of tweets is being collected since September 2015. We picked a random subset of French tweets and dropped 99% of tweets containing an url, since many of them come from bots. The remaining urls are kept because there are some urls in the challenge data. We replaced 1% of mentions

²<http://sklearn-crfsuite.readthedocs.io/en/latest/>

³<https://github.com/salestock/fastText.py>

⁴<http://osirim.irit.fr/site/fr/articles/corpus>

Table 2: *Fasttext* parameters

parameter	value
learning rate	0.02
dimension	200
context window size	5
number of epochs	4
min_count	5
negative/positive samples ratio	5
minimum n-gram size	3
maximum n-gram size	6
sampling threshold	10^{-4}

Table 3: Handpicked clusters and random samples

Cluster theme	Cluster sample
hyperlinks	https://t.co/d73eViSrbW
hours	12h00 19h19 12h 7h44
dates	1947 1940 27/09 mars Lundi
joyful reactions	ptdrr mdrrrrrrr pff boooooooooordel
TPMP (french show)	#TPMP #hanouna Castaldi
transportation lines	@LIGNEJ_SNCF @TER_Metz
emojis	Pffff :))
video games	@PokemonFR manette RT
persons	@olivierminne @Vibrationradio
football players	Ribery Leonardo Chelsea Ramos

(@*someone* tokens) by the symbol @* hoping to help generalization. This preprocessed additional data totals 40M tweets.

3.2 Parameter selection

Parameters and feature subsets were not thoroughly optimized through cross validation, except regularization parameters. We used Elasticnet regularization [ZH05] and the L-BFGS optimization algorithm, with a maximum of 100 iterations.

We ran grid search using sequence level accuracy score as a metric, on $c1$ and $c2$, the regularization weights for L1 and L2 priors. They were tested in respective ranges $[10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}]$ and $[0.1, 0.3, 0.8, 2]$. $c1 = 10^{-4}$ and $c2 = 0.3$ were chosen.

Fasttext skipgram uses negative sampling with the parameters described in table 2. A different skipgram model was used for sentence representation, with a dimension of 40. For the Gaussian Mixture Model, we used 100 dimensions and diagonal covariance.

4 Results

4.1 Clustering

Many clusters correspond directly to named entities. Table 3 shows a random sample of 10 handpicked clusters and the themes we identified.

4.2 Performance

We will report the results of the system on the evaluation data when fitted on the full training data. The system yields a sequence level accuracy of 61.2% using an ensemble of $N = 20$ models. Note that a single model ($N = 1$) has a sequence level accuracy 60.7% which is only slightly less.

The challenge scoring metric was a micro F-measure based on chunks of consecutive labels. Our ensemble system scores 58.89% with respect to this metric. Table ?? summarize the results of the competition and show that our system won with a rather large margin. *Fasttext* features bring a notable difference since the sequence level accuracy drops to 57.8% when we remove all of them. Table 4 gives an overview of scores per label, and could show us ways to improve the system. The 13 labels were separated according to their IOB encoding status.

4.3 Interpreting model predictions

CRF is based on a linear model and the learned weights are insightful: the highest weights indicate the most relevant features for the prediction of a given label, while the lowest weights indicate the most relevant features for preventing the prediction of a given label. Tables 5 and 6 show those weights for a single model trained on all features. ft_wo_i , $ft_wo_c_i$ and ft_sen_i refer respectively to the i th component of a *fasttext* raw word representation, cluster based representation, and sentence level representation. The model actually uses those three kinds of features to predict labels. Clustering embeddings can improve the interpretability of the system by linking a feature to a set of similar words. Sentence level embeddings seem to prevent the model from predicting irrelevant labels, suggesting they might help for disambiguation.

4.4 Computational cost

Fitting the CRF model with 3000 examples (labeled tweets) takes up 4 minutes on a Xeon E5-2680 v3 CPU using a single thread, and inference on 3688 example only needs 30 seconds. Fitting the *fasttext* model of dimension 200 on 40M tweets takes up 10 hours on a single thread, but only 30 minutes when using 32 threads.

Table 4: Fine grained score analysis

label	precision	recall	f1-score	support
B-person	0.767	0.618	0.684	842
I-person	0.795	0.833	0.814	294
B-geoloc	0.757	0.697	0.726	699
B-transportLine	0.978	0.926	0.951	517
B-musicartist	0.667	0.178	0.281	90
B-other	0.286	0.134	0.183	149
B-org	0.712	0.277	0.399	545
B-product	0.519	0.135	0.214	312
I-product	0.320	0.113	0.167	364
B-media	0.724	0.462	0.564	210
B-facility	0.639	0.363	0.463	146
I-facility	0.620	0.486	0.545	175
B-sportsteam	0.514	0.277	0.360	65
I-sportsteam	1.000	0.200	0.333	10
B-event	0.436	0.185	0.260	92
I-event	0.356	0.292	0.321	89
B-tvshow	0.429	0.058	0.102	52
I-tvshow	0.286	0.065	0.105	31
I-media	0.200	0.019	0.035	52
B-movie	0.333	0.045	0.080	44
I-other	0.000	0.000	0.000	73
I-transportLine	0.873	0.729	0.795	85
I-geoloc	0.650	0.409	0.502	159
I-musicartist	0.636	0.163	0.259	43
I-movie	0.250	0.049	0.082	41

Table 5: Highest θ weights

weight	label	feature
3.26	O	end of sentence
2.47	O	beginning of sentence
2.01	O	previous word:rt
1.92	B-transportLine	ft_wo_91
1.85	B-other	previous word:les
1.80	B-geoloc	previous word:#qml
1.76	B-geoloc	previous word:pour
1.71	B-geoloc	ft_sen_22
1.71	O	ft_wo_c68
1.68	B-org	current word:#ratp

Table 6: Lowest θ weights

weight	label	feature
-1.65	B-product	ft_sen_33
-1.60	B-org	ft_sen_9
-1.48	O	previous word:sur
-1.41	B-facility	ft_sen_33
-1.40	O	suffix:lie
-1.38	O	suffix:ra
-1.29	B-other	previous POS: verb (future)
-1.29	B-geoloc	ft_wo_151
-1.27	B-person	previous word prefix:l
-1.26	B-org	ft_wo_130

5 Conclusion and further improvements

We presented a NER system using *Fasttext* which was ranked first at the CAP 2017 NER challenge. Due to a lack of time, we did not optimize directly on the challenge evaluation metrics, using sequence level accuracy as a proxy, and we did not cross-validate all important parameters. Besides, there are other promising ways to increase the score of the system that we did not implement:

1. thresholding for F1 maximization: Our system precision (73.65%) is significantly higher than its recall (49.06%). A more balanced score could be obtained by having a negative bias towards predicting no label. This might improve the F1 score. Threshold optimization works well for non-structured prediction [CEN14], but it is not clear that it would bring about improvement in practical applications.
2. larger scale unsupervised learning: More tweets could be used, and/or domain adaptation could be applied in order to bias embeddings towards learning representations of words occurring in the challenge data.
3. RNN embeddings: Unsupervised learning with recurrent neural networks can be used to learn "contextualized" embedding of words. Unsupervised training tasks include language modeling or auto-encoding. RNNs have been used in NER without unsupervised training. [ABP⁺16] [LC]
4. DBpedia spotlight [DJHM13] could provide an off-the-shelf gazetteer, yielding potentially powerful features for NER.

References

- [ABP⁺16] Vinayak Athavale, Shreenivas Bharadwaj, Monik Pamecha, Ameya Prabhu, and Manish Shrivastava. Towards deep learning in hindi NER: an approach to tackle the labelled data sparsity. *CoRR*, abs/1610.09756, 2016.
- [Ahm13] Zia Ahmed. Named Entity Recognition and Question Answering Using Word Vectors and Clustering. 2013.
- [BGJM16] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. 2016.
- [CEN14] Z Chase Lipton, C Elkan, and B Narayanaswamy. Thresholding Classifiers to Maximize F1 Score. *ArXiv e-prints*, 2014.
- [DGG17] Arjun Das, Debasis Ganguly, and Utpal Garain. Named Entity Recognition with Word Embeddings and Wikipedia Categories for a Low-Resource Language. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 16(3):1–19, 2017.
- [DJHM13] Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N Mendes. Improving Efficiency and Accuracy in Multilingual Entity Extraction. *Proceedings of the 9th International Conference on Semantic Systems*, pages 121–124, 2013.
- [GCWL14] Jiang Guo, Wanxiang Che, Haifeng Wang, and Ting Liu. Revisiting Embedding Features for Simple Semi-supervised Learning. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, (2005):110–120, 2014.
- [LC] Nut Limsopatham and Nigel Collier. *Proceedings of the 2nd Workshop on Noisy User-generated Text*.
- [LPB⁺17] Cédric Lopez, Ioannis Partalas, Georgios Balikas, Nadia Derbas, Amélie Martin, Frédérique Segond, Coralie Reutenauer, and Massih-Reza Amini. French Named Entity Recognition in Twitter Challenge. Technical report, 2017.
- [MCCD13] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed Representations of Words and Phrases and their Compositionality. *Nips*, pages 1–9, 2013.
- [NNN98] Elaine Marsh Nrl, Dennis Perzanowski Nrl, and Ralph Grishman Nyu. MUC-7 EVALUATION OF IE TECHNOLOGY : Overview of Results MUC-7 Program Committee. *Program*, (April), 1998.
- [Oka07] Naoaki Okazaki. CRFsuite: a fast implementation of Conditional Random Fields (CRFs), 2007.
- [RCME11] Alan Ritter, Sam Clark, Mausam, and Oren Etzioni. Named entity recognition in tweets: An experimental study. In *EMNLP*, 2011.
- [Sch94] Helmut Schmid. Probabilistic Part-of-Speech Tagging Using Decision Trees. *Proceedings of the International Conference on New Methods in Language Processing*, pages 44–49, 1994.
- [Sie15] Scharolta Katharina Sien. Adapting word2vec to Named Entity Recognition. *Proceedings of the 20th Nordic Conference of Computational Linguistics (NODALIDA 2015)*, (Nodalida):239–243, 2015.
- [SM11] Charles Sutton and Andrew McCallum. An Introduction to Conditional Random Fields. *Machine Learning*, 4(4):267–373, 2011.
- [ZH05] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic-net. *Journal of the Royal Statistical Society*, 67(2):301–320, 2005.