

# NLNDE: Enhancing Neural Sequence Taggers with Attention and Noisy Channel for Robust Pharmacological Entity Detection

Lukas Lange<sup>1,2,3</sup>

Heike Adel<sup>1</sup>

Jannik Strötgen<sup>1</sup>

<sup>1</sup> Bosch Center for Artificial Intelligence, Renningen, Germany

<sup>2</sup> Spoken Language Systems (LSV), Saarland University, Saarbrücken, Germany

<sup>3</sup> Saarbrücken Graduate School of Computer Science, Saarbrücken, Germany

{Lukas.Lange, Heike.Adel, Jannik.Stroetgen}@de.bosch.com

## Abstract

Named entity recognition has been extensively studied on English news texts. However, the transfer to other domains and languages is still a challenging problem. In this paper, we describe the system with which we participated in the first subtrack of the PharmaCoNER competition of the BioNLP Open Shared Tasks 2019. Aiming at pharmacological entity detection in Spanish texts, the task provides a non-standard domain and language setting. However, we propose an architecture that requires neither language nor domain expertise. We treat the task as a sequence labeling task and experiment with attention-based embedding selection and the training on automatically annotated data to further improve our system’s performance. Our system achieves promising results, especially by combining the different techniques, and reaches up to 88.6% F1 in the competition.

## 1 Introduction

The detection and classification of pharmacological and biomedical entities in texts is especially challenging due to the domain’s nature with long and complex entity names, which usually requires the design and usage of handcrafted rules and features. Natural language processing (NLP) research focused on this topic for quite a while on English texts, e.g., the drugs and chemical names extraction challenge (CHEMDNER) (Krallinger et al., 2015) or tracks for chemical entity recognition at BioCreative (Pérez-Pérez et al., 2017). Following these tasks, the Pharmacological Substances, Compounds and Proteins and Named Entity Recognition track (PharmaCoNER) is the first competition on this topic on Spanish data (Gonzalez-Agirre et al., 2019).

Named entity recognition (NER) and classification is the first subtrack of Pharma-




Figure 1: Annotated sample sentences (PRO. is short for PROTEINAS).

CoNER and aims at distinguishing four entity types: PROTEINAS, NORMALIZABLES, NO-NORMALIZABLES, and UNCLEAR. Our model was trained on all four entity types, although the NO-NORMALIZABLES type was not considered during the official evaluation due to its ambiguous definition. Two annotated sample sentences from the training data are shown in Figure 1.

In this paper, we describe our submissions to and their results in the first subtrack of PharmaCoNER. We address this task as a sequence-labeling problem and implement a system that relies Neither on Language Nor on Domain Expertise (NLNDE). For this, we use a combination of different state-of-the-art approaches from NLP to tackle its challenges without the need for handcrafted features.

We train recurrent neural networks with conditional random field (CRF) output layers which are state of the art for different sequence labeling tasks, such as named entity recognition (Lample et al., 2016), part-of-speech tagging (Kemos et al., 2019) and de-identification (Liu et al., 2017). In our different runs, we further explore the advantages of domain-specific fastText embeddings that have been pre-trained on SciELO and Wikipedia articles (Soares et al., 2019) to investigate the impact of domain knowledge. Note that the training of these embeddings requires only a collection of domain-specific text but no human domain expertise. Based on these models, we train an attention-

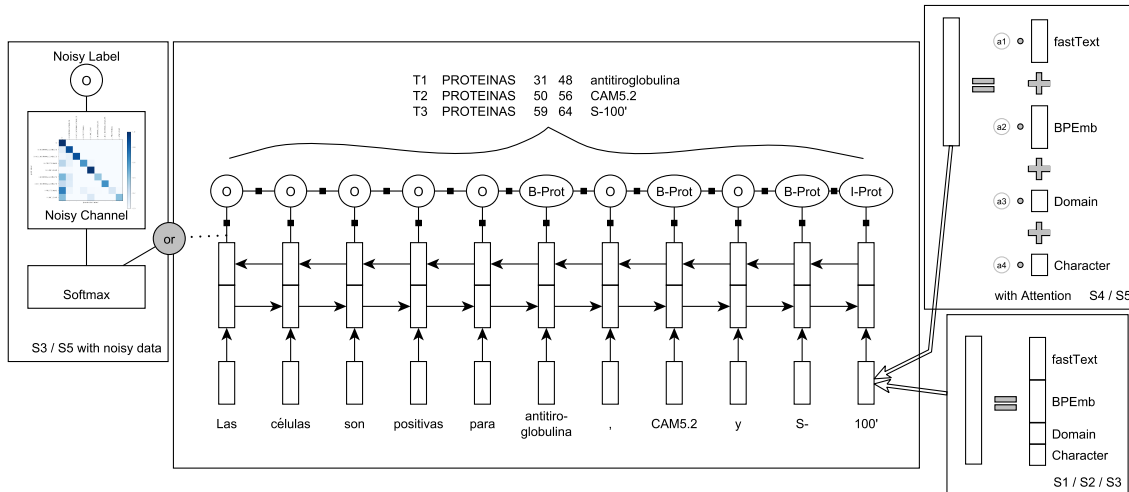


Figure 2: Architecture of our models. The label prefixes “B-” and “I-” show how we address the task as a sequence-labeling task. The word representations are either the concatenated embeddings (in runs S1–S3) or the attention-based weighted embeddings (in runs S4–S5).

based embedding selection function in order to leverage multiple different word embeddings effectively. Finally, we extend the training data with automatically annotated data, which was sampled from the same domain and annotated with information from Wikidata.<sup>1</sup>

## 2 Methods

In this section, we present our system, the attention function for embedding selection, and the noisy channel model.

### 2.1 NLNDE System

In Figure 2, the architecture of our models is depicted, which we explain in the following.

**Input Embeddings.** We tokenize the input with the tokenizer provided by the shared task organizers (Intxaurreondo, 2019). We noticed that the tokenizer sometimes merges multi-word expressions into a single token joined with underscores for contiguous words. As a result, some tokens cannot be aligned with the corresponding entity annotations. To address this, we split those tokens into their components in a postprocessing step. Then, we represent each token with the following embeddings (see bottom right box of Figure 2):

- *Character embeddings:* We use the concatenated last forward and backward hidden states of a bidirectional long short-term

memory (BiLSTM) network (Hochreiter and Schmidhuber, 1997) over character embeddings (50 dimensions, randomly initialized, fine-tuned during training (Lample et al., 2016)).

- *Domain-independent fastText embeddings* (100 dimensions, pre-trained on Spanish text (Grave et al., 2018)).
- *Domain-specific fastText embeddings* (100 dimensions, pre-trained on Spanish SciELO and Wikipedia articles (Soares et al., 2019)).
- *Byte-pair encoding embeddings* (300 dimensions, vocabulary size of 200,000, pre-trained on Spanish text (Heinzerling and Strube, 2018)).

Note that except for the character embeddings, we do not fine-tune any of the embeddings. All embeddings are concatenated into a single word representation vector.

**Word Features.** We also experiment with extending the input representations with the following features:

- *Part-of-speech (POS):* The POS tags are generated by the POS-tagger provided by the shared task organizers (Intxaurreondo, 2019). The tags are embedded into a 20-dimensional randomly initialized embedding and learned during training. The embedded vector is used as the representation for the POS tag.

<sup>1</sup><https://www.wikidata.org/>

- *Length*: For each word, we encode its length in a one-hot vector. Words with more than nine characters share the same vector (10 dimensions).
- *Frequency*: We consider the relative frequency  $f$  of each word and bin the frequencies into ten groups. The first group contains the most frequent words that have relative frequencies above 1% ( $f > 1\%$ ). The remaining bins are constructed in the following manner:  $f > 0.5\%$ ,  $f > 0.1\%$ ,  $f > 0.05\%$ , etc. (one-hot encoded, 10 dimensions).
- *Word shape*: We distinguish between uppercased, lowercased, starts with capital letter, numeric, mostly numeric, punctuation, mostly punctuation, only letters, alphanumeric and other (one-hot encoded, 10 dimensions).

All features are concatenated into a single feature vector  $f$  of 50 dimensions.

**BiLSTM-CRF Layers.** The input representation is fed into a BiLSTM with a conditional random field (CRF) output layer, similar to the model of [Lample et al. \(2016\)](#). The CRF output layer is a linear-chain CRF, i.e., it learns transition scores between the output classes. For training, the forward algorithm is used to sum the scores for all possible sequences. During decoding, the Viterbi algorithm is applied to obtain the sequence with the maximum score.

**Hyperparameters and Training.** The hyperparameters are the same across all runs. We use a BiLSTM hidden size of 256 and train the network with the NADAM optimizer ([Dozat, 2016](#)) using a learning rate of 0.002 and a batch size of 32. For regularization, we employ early stopping on the development set and apply dropout with probability 0.5 on the input representations.

## 2.2 Attention for Embedding Selection

As we are combining different word embeddings, some of them may be more beneficial for certain words than others, e.g., domain-specific embeddings for in-domain words. [Kiela et al. \(2018\)](#) used an attention mechanism for weighting and selecting the best embeddings for each word. We extend this idea and propose the following attention function to weight the embeddings depending on additional word features.

For the attention-based models, all  $n$  embeddings  $e$  are mapped to the same size using a linear mapping  $Q_i \in \mathbb{R}^{E \times E_i}$  without bias, with  $x_i \in \mathbb{R}^E$  being the  $i$ -th embedding  $e_i$  mapped from their original size  $E_i$  to the maximal embedding size  $E = \max_m(E_m)$ .

$$x_i = Q_i \cdot e_i \quad (1)$$

In order to allow the model to make an informed decision which embeddings to focus on, we use the word features described in Section 2.1 as an additional input to the attention function. The vector  $f \in \mathbb{R}^F$  representing the features for each word is concatenated to each embedding  $x_i$ . The attention weight  $a_i$  for each embedding  $x_i$  is computed with the softmax function, by feeding  $x_i$  and  $f$  into a fully-connected hidden layer of size  $H$  with the parameters  $W \in \mathbb{R}^{H \times E}$ ,  $U \in \mathbb{R}^{H \times F}$ ,  $V \in \mathbb{R}^{1 \times H}$ .

$$a_i = \frac{\exp(V \cdot \tanh(Wx_i + Uf))}{\sum_{l=1}^n \exp(V \cdot \tanh(Wx_l + Uf))} \quad (2)$$

Finally, the embeddings  $x_i$  are weighted using the attention weights  $a_i$  resulting in the word representation:

$$e = \sum_i a_i \cdot x_i \quad (3)$$

Then, this word representation  $e \in \mathbb{R}^E$  is fed into the BiLSTM-CRF. Compared to a concatenation of the different embeddings, this results in a lower-dimensional word representation and, thus, requires fewer parameters in the BiLSTM layer. The attention-based embedding selection is shown in the upper right box of Figure 2.

## 2.3 Training on Noisy Data

As it was shown in multiple low-resource settings ([Dgani et al., 2018](#); [Fang and Cohn, 2016](#); [Mnih and Hinton, 2012](#); [Paul et al., 2019](#); [Yang et al., 2018](#)), the performance of NER and other NLP systems can be substantially improved by training on additional noisy data which is labeled in a distantly supervised manner ([Mintz et al., 2009](#)). With this approach, the noisy data is cheap to create, but also error-prone and can even decrease performance if used as training data without noise handling as shown by [Hedderich and Klakow \(2018\)](#).

**Extraction of Noisy Data.** We create gazetteers for the different entity types by extracting names and aliases of possible entities from Wikidata for the following categories and their subclasses:<sup>2</sup>

- **PROTEINAS:** enzyme, gene, hormone, protein.
- **NORMALIZABLES:** allotropy, alloy, amino acid, antibody, carbohydrate, diagnostic procedure, dye, lipid, mineral, nucleotide, oil, reagent, chemical compounds, peptide, plant, polymer, vaccine.
- **UNCLEAR and NO-NORMALIZABLES:** The gazetteer was constructed from entity mentions in the training data that appeared at least twice and examples from the annotation guidelines.

Then, we retrieve unlabeled documents from the same domain from the SciELO archive (Packer, 1998). Finally, we use the extracted gazetteers to automatically annotate the SciELO data with the method from Lange et al. (2019). We use case-insensitive string matching for PROTEINAS and strict string matching for the other types. This allows to create additional training instances, but at the same time introduces noise into the system. To avoid that the noisy labels result in a decrease of performance, we train on the noisy data with a special noise handling method adapted from Goldberger and Ben-Reuven (2016), which will be described in the following.

**Noisy Channel and Confusion Matrix.** First, we annotate each word of the training data using the same method as for generating the noisy data. Thus, each word in the training data has a clean, true label  $y$  and a noisy label  $\hat{y}$  from which we can model the noise distribution  $p(\hat{y} = j|y = i)$  with a confusion matrix, as shown in Figure 3. We transform the distribution of the predicted (clean) labels to the noisy label distribution through a so-called noisy channel (Goldberger and Ben-Reuven, 2016):

$$p(\hat{y} = j|x) = \sum_{i=1}^k p(\hat{y} = j|y = i)p(y = i|x) \quad (4)$$

<sup>2</sup>WikiData identifiers used for the extraction: Q8047, Q7187, Q11364, Q8054, Q81915, Q37756, Q8066, Q79460, Q11358, Q177719, Q189720, Q11367, Q7946, Q28745, Q42962, Q2356542, Q47154513, Q172847, Q756, Q81163, Q134808.

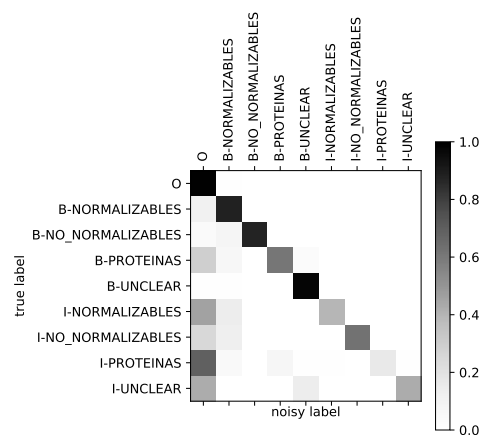


Figure 3: Confusion matrix for the automatic annotation on the training data used for the noisy channel initialization.

where  $k$  is the number of classes and  $p(y = i|x)$  is the probability of a label  $y$  having a specific class  $i$  given the feature  $x$ .

We initialize the noisy channel weights using the learned confusion matrix on the training set, for which clean and noisy labels are available.

**Training with Confusion Matrix.** The sequence tagging model is then trained alternately on the clean data with the CRF output layer and on the noisy data with the noisy channel layer, as shown in Figure 2. The number of noisy training instances is constantly decreased by 5% after every training epoch to at least 100 sentences, as we observed that the noisy data helps in particular for the first epochs, but decreases performance if the amount is not reduced. Note that we shuffle the noisy data after each training epoch. Thus, the model is trained on new samples of noisy sentences in every epoch.

### 3 Submissions

We submitted five runs to the PharmaCoNER competition. All of them are based on the architecture described in Section 2.1.

S1 (*Base*): Our first run, the base system for all of the following runs, uses a concatenation of three embeddings (character, BPEmb, fast-Text) which were all trained on Wikipedia. Thus, this run does not include any form of domain knowledge, and it uses neither noisy data nor attention for embedding selection.

S2 (*Domain*): Our second run uses the three embeddings from S1 plus domain-specific fast-

$S_{ID}$	Development			Test		
	P	R	F1	P	R	F1
S1	86.3	85.0	85.6	85.5	85.3	85.4
S2	87.1	85.8	86.4	86.3	85.9	86.1
S3	87.5	87.5	87.5	85.1	86.0	85.6
S4	88.0	<b>88.0</b>	88.0	85.2	87.2	86.2
S5	<b>89.1</b>	87.5	<b>88.3</b>	<b>89.0</b>	<b>88.3</b>	<b>88.6</b>

Table 1: Precision (P), Recall (R) and F1 for Task 1.

Text embeddings to incorporate knowledge about word distributions within the domain.

S3 (*Noise*): Our third run extends the model of S2 with training on additional noisy data (cf. Section 2.3). Moreover, we use the feature vector as an additional input, which is different from runs S1 and S2.

S4 (*Attention*): The fourth run uses the attention function for word embedding selection (cf. Section 2.2). Apart from that, the model is identical to S2 and only trained on clean data.

S5 (*Attention+Noise*): Our last run has the same architecture as S4 but is trained on the noisy data in addition. It thus combines domain-independent and domain-specific word embeddings, attention-based embedding selection, and training on noisy data.

## 4 Results and Analysis

This section describes our results and analysis.

### 4.1 Experimental Results

In Table 1, we report the results on the PharmaCoNER development and test sets using the official shared task evaluation metrics.

Adding domain-knowledge (S2) to the base model S1 improves the performance on the development and the test set. The training on noisy data (S3) and the attention function alone (S4) do not lead to strong improvements on the test set; the noise model S3 even decreases performance. The combination of all proposed methods (run S5 *Attention+Noise*) outperforms all other models.

While we are able to see the improvements step by step introduced by our methods on the development set, such improvements are not observable one-to-one on the test set. We assume that model S5 performs best at generalizing to unseen words due to the training on additional data and the attention function based on basic word properties like word length or frequency. The other models seem




Figure 4: The attention weights of our model for the four embeddings. Darker color indicates higher weight.

to overfit on the development set, even though this set was never used for training but only for early stopping.

### 4.2 Analysis of Attention Weights

The attention-based models learn to focus mostly on the byte-pair-encoding embeddings, as shown in Figure 4. In particular, for words from the general domain (positivas) and stopwords (para), our model focuses on these embeddings. For domain-specific words (antitiroglobulina, CAM5.2), the model learns to focus more on the fastText embeddings and especially the domain-specific embeddings. Interestingly, the character embeddings are never assigned a noticeable weight. This may be attributed to the fact that the other embeddings are all subword embeddings and that they are able to generate meaningful vectors for out-of-vocabulary words. Moreover, the character embeddings were randomly initialized and had to be learned during training while the other models were pretrained.

## 5 Conclusions

In this paper, we described our system for the first subtrack of the PharmaCoNER competition. We trained a bi-directional long short-term memory network and explored different input representations. We proposed to use a feature-based attention function for embedding selection and training on noisy data, which in combination increased performance by more than 3 F1 points up to 88.6%. This shows that we can successfully extract these special types of entities without the need for domain or language-specific model architectures.

## Acknowledgments

We would like to thank the anonymous reviewers for their helpful suggestions and comments.



## References

- Y. Dgani, H. Greenspan, and J. Goldberger. 2018. [Training a neural network based on unreliable human annotation of medical images](#). In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 39–42.
- Timothy Dozat. 2016. Incorporating nesterov momentum into adam. In *Proceedings of the Workshop of the International Conference on Learning Representations (ICLR)*.
- Meng Fang and Trevor Cohn. 2016. [Learning when to trust distant supervision: An application to low-resource POS tagging using cross-lingual projection](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 178–186, Berlin, Germany. Association for Computational Linguistics.
- Jacob Goldberger and Ehud Ben-Reuven. 2016. [Training deep neural-networks using a noise adaptation layer](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Aitor Gonzalez-Agirre, Montserrat Marimon, Ander Intxaurre, Obdulia Rabal, Marta Villegas, and Martin Krallinger. 2019. Pharmaconer: Pharmaceutical substances, compounds and proteins named entity recognition track. In *Proceedings of the BioNLP Open Shared Tasks (BioNLP-OST)*, Hong Kong, China. Association for Computational Linguistics.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning word vectors for 157 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).
- Michael A. Hedderich and Dietrich Klakow. 2018. [Training a neural network in a low-resource setting on automatically annotated noisy data](#). In *Proceedings of the Workshop on Deep Learning Approaches for Low-Resource NLP*. Association for Computational Linguistics.
- Benjamin Heinzerling and Michael Strube. 2018. [BPEmb: Tokenization-free pre-trained subword embeddings in 275 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ander Intxaurre. 2019. [SPACCC \(spanish clinical case corpus\) tokenizer](#).
- Apostolos Kemos, Heike Adel, and Hinrich Schütze. 2019. [Neural semi-Markov conditional random fields for robust character-based part-of-speech tagging](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2736–2743, Minneapolis, Minnesota. Association for Computational Linguistics.
- Douwe Kiela, Changan Wang, and Kyunghyun Cho. 2018. [Dynamic meta-embeddings for improved sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1466–1477, Brussels, Belgium. Association for Computational Linguistics.
- Martin Krallinger, Florian Leitner, Obdulia Rabal, Miguel Vazquez, Julen Oyarzabal, and Alfonso Valencia. 2015. CHEMDNER: The drugs and chemical names extraction challenge. *Journal of cheminformatics*, 7(1):S1.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Lukas Lange, Michael A. Hedderich, and Dietrich Klakow. 2019. [Feature-dependent confusion matrices for low-resource ner labeling with noisy labels](#). In *Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, Hong Kong, China. Association for Computational Linguistics.
- Zengjian Liu, Buzhou Tang, Xiaolong Wang, and Qingcai Chen. 2017. [De-identification of clinical notes via recurrent neural network and conditional random field](#). *Journal of Biomedical Informatics*, 75(S):S34–S42.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011, Suntec, Singapore. Association for Computational Linguistics.
- Volodymyr Mnih and Geoffrey Hinton. 2012. [Learning to label aerial images from noisy data](#). In *Proceedings of the 29th International Conference on Machine Learning, ICML 12*, pages 203–210, USA. Omnipress.
- Abel Laerte Packer. 1998. SciELO: uma metodologia para publicação eletrônica. *Ciência da informação*, 27(2).

- Debjit Paul, Mittul Singh, Michael A. Hedderich, and Dietrich Klakow. 2019. [Handling noisy labels for robustly learning from self-training data for low-resource sequence labeling](#). *CoRR*, abs/1903.12008.
- Martín Pérez-Pérez, Obdulia Rabal, Gael Pérez-Rodríguez, Miguel Vazquez, Florentino Fdez-Riverola, Julen Oyarzábal, A Valencia, Anália Loureno, and Martin Krallinger. 2017. Evaluation of chemical and gene/protein entity recognition systems at biocreative v. 5: the CEMP and GPRO patents tracks. In *Proceedings of the BioCreative Workshop*, volume 5, pages 3–11.
- Felipe Soares, Marta Villegas, Aitor Gonzalez-Agirre, Martin Krallinger, and Jordi Armengol-Estapé. 2019. [Medical word embeddings for Spanish: Development and evaluation](#).
- Yaosheng Yang, Wenliang Chen, Zhenghua Li, Zhengqiu He, and Min Zhang. 2018. [Distantly supervised NER with partial annotation learning and reinforcement learning](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2159–2169, Santa Fe, New Mexico, USA. Association for Computational Linguistics.