



**HAL**  
open science

## Effect of depth order on iterative nested named entity recognition models

Perceval Wajsbürt, Yoann Taillé, Xavier Tannier

► **To cite this version:**

Perceval Wajsbürt, Yoann Taillé, Xavier Tannier. Effect of depth order on iterative nested named entity recognition models. Conference on Artificial Intelligence in Medecine (AIME 2021), Jun 2021, Porto, Portugal. hal-03277643

**HAL Id: hal-03277643**

**<https://hal.science/hal-03277643>**

Submitted on 4 Jul 2021

**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.

# Effect of depth order on iterative nested named entity recognition models

Perceval Wajsbürt<sup>1</sup>, Yoann Taillé<sup>1,2</sup>, Xavier Tannier<sup>1</sup>

<sup>1</sup> Sorbonne Université, Inserm, LIMICS, Paris, France

<sup>2</sup> SCAI, Sorbonne Université, Paris, France

**Abstract.** This paper studies the effect of the order of depth of mention on nested named entity recognition (NER) models. NER is an essential task in the extraction of biomedical information, and nested entities are common since medical concepts can assemble to form larger entities. Conventional NER systems only predict disjointed entities. Thus, iterative models for nested NER use multiple predictions to enumerate all entities, imposing a predefined order from largest to smallest or smallest to largest. We design an order-agnostic iterative model and a procedure to choose a custom order during training and prediction. To accommodate for this task, we propose a modification of the Transformer architecture to take into account the entities predicted in the previous steps. We provide a set of experiments to study the model's capabilities and the effects of the order on performance. Finally, we show that the smallest to largest order gives the best results.

**Keywords:** named entity recognition, biomedical, nested entities

## 1 Introduction

Biomedical concept recognition is a classical and essential task of natural language processing for biomedical applications [25], aiming to extract information such as symptoms, treatments, proteins, genes, dates, and durations from free text. Classic methods assume that entities are disjoint and formulate the problem as a sequence segmentation task, using word tagging schemes. However, in a real-world scenario, entities can compose or overlap, thus breaking the assumption that they are disjoint. For example, a temporal event "after anesthesia" contains the nested treatment entity "anesthesia."



Checking empty vials after anesthesia.

Multiple methods have been proposed to tackle this problem, ranging from exhaustive filtering of all possible spans to complex graph structures.

A class of methods deals with this problem of nested named entity recognition with a cascade of flat (non-nested) named entity recognition layers for different depths, i.e., predict the entities at a given depth iteratively starting either from large entities or short entities. The predictions of a given depth are used as additional input for the next prediction. We can argue whether the depth order matters during the training of such a model:

is it easier for the model to predict large entities first and look inside its previous predictions for smaller ones, or predict small entities and compose them to build larger ones? To answer this question, we design an order-agnostic auto-regressive model based on the Transformer encoder architecture and a procedure to let it choose a custom order during training and prediction.

This work was originally designed to address Task 3 of the DEFT 2020 evaluation campaign (in French), and we further evaluated it on the classical GENIA dataset (in English). This DEFT task deals with the detection of named entities in texts describing clinical cases [7]. The different types of entities are, on the one hand, pathologies and signs or symptoms (task 3.1), and on the other hand, *anatomy, anatomy examinations, substances, doses, administration methods, treatments (surgical or medical), values, time*. More details about the challenge are presented in [1], in which we describe our 3 official submissions.

To our knowledge, our work is the first to evaluate a nested biomedical concept recognition system in the French language, which constitutes an additional contribution, as resources for languages other than English are very scarce [22].

## 2 Related works

Nested named entity recognition has been the subject of renewed attention in the biomedical NLP community since 2018, leading to many different approaches.

*Exhaustive or semi-exhaustive methods* There are multiple approaches to tackle the task of nested named entity recognition. One class of method addresses the problem by enumerating all possible spans of the input sequence and classify each one with its label, including a "no entity" class. [27] classify each span independently by pooling its tokens representations. [34] consider the left and right context when classifying the spans. [33] uses an LSTM cell [8] to model dependencies between spans that differ by one token. [35] first filters candidate mentions by predicting all possible start and end tokens and then predicting a label for every mention that starts or end at one of the boundaries.

*Layered methods* A different research direction focuses on layered models that iteratively predicts entity spans until none can be found. [5] uses a fixed number of layers and updates spans representations using a novel neural architecture. [9] designs a layered architecture that predicts entities at each layer and merges the word representations before applying the next layer. [26] computes tag scores for each word and decodes the spans by applying the Viterbi algorithm multiple times on a previously extracted subsequence, starting from the full sentence.

*Hypergraph methods* Some methods model the span detection with hypergraphs to account for the non-linear structure of the tag sequences. [17] design a CRF hypergraph with various node types to model entity types and boundaries. However, cycles in the graphs of some samples required that the CRF normalization term had to be approximated, leading to a decreased performance [21]. [21] models the mention edges

and transitions instead of solely modeling token tags. Their method, however, requires multiple graphs if more there is more than one entity type. Alternatively, [10] only models mention tags and not their transitions, but allows a multi-label prediction for each token. They modify an LSTM layer to represent multiple states for a single token and perform decoding during the neural network execution.

*Neural transition models* Transition models iteratively build the detected spans by chaining different actions. [31] employs a set of three actions types (SHIFT, REDUCE-X, UNARY-X) that build a forest of binary word trees. [18] adds a new action (OUT) to allow building multiple mentions for a given span, especially fine-grained entity classes according to a predefined semantic hierarchy.

*Other methods* [16] uses Anchor Region Networks to first predict an anchor token that is contained by an entity and then predict the range of the mention from that anchor. This method leverage the insight that most mentions contain to a head token.

*Language embedding* Pre-trained self-supervised language models [20, 23, 3] have improved downstream NLP tasks’ performance. Most state-of-the-art methods use pre-trained models as contextualized token features but do not alter the underlying model’s structure. In contrast, we experiment with the augmentation of a pretrained transformer model to address the nested entity recognition task.

### 3 Method

#### 3.1 Model

We detail here a new model to handle nested entity recognition. The model is an auto-regressive encoder-only Transformer [30] taking as input a sequence of words and a list of entity mentions already extracted (empty list at the first iteration) and predicts as output a list of new mentions. The entities predicted at each iteration do not overlap, but all the entities predicted at the end of the iterations may overlap.

We handle entity mentions in the form of tags assigned to each token with the conventional BIO or BIOUL formats [2, 24], by embedding each tag into a multidimensional vector space. We then sum the tag embeddings of different entities (depths) at the same position with each other. The sentences are tokenized and represented with the contextual embedding, Transformer model BERT (for English documents [3]) or CamemBERT (for French [19]). The word embedding and the tag embeddings from previous iterations are fed into a linear CRF [13] predicting flat entity mentions (Figure 1a).

Each iteration of this model receives the tags predicted by the previous iteration as input to predict different mentions at each step (Figure 1b).

#### 3.2 Greedy order training

During training, the goal is to maximize the probabilities assigned by the model to the lists of target entities. This goal comes with the difficulty of dealing with overlapping entities.

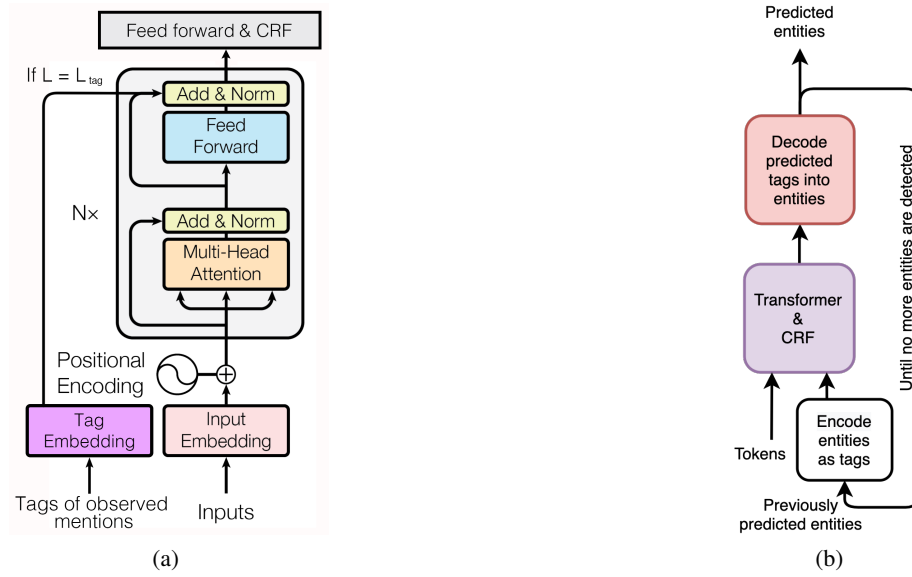


Fig. 1: (a) modified Transformer conditioned on the tags of the previously observed mentions at the end of the layer  $L_{tag}$  - figure adapted from [30]; (b) Global model prediction diagram

We proceed in several steps and predict only non-interleaved entities at each run. However, several permutations, or valid prediction paths, lead to the same list of entities. For two nested statements

moment  
treatment

Checking empty vials after anesthesia.

(annotations that we call  $T$  and  $H$ ), we can predict  $T$  first, then  $H$  knowing  $T$ , or the opposite, i.e. choose to optimize between two objectives :

- $P(T, H) = P(T, H|T) \times P(T)$
- $P(T, H) = P(T, H|H) \times P(H)$

In this symmetrical situation, a solution that would optimize both paths simultaneously, for example, by summing several tag losses for different depths, could lead to a non-optimal situation in which we ask the model to detect an incorrect combination of the two mentions. For example, these two non-overlapping but sub-optimal mentions:

moment treatment

Checking empty vials after anesthesia

Models like [9] choose a strategy in advance (smaller mentions first, for example), but the risk is not to take advantage of all the interdependencies that make some mentions easier to find when you know the others.

Another solution is to choose the order of extraction leading to the model’s best performance, measured in F-measurement. A greedy strategy is applied to select, among the non-overlapping combinations of mentions not observed in this batch, the closest to the mentions predicted in terms of F1 overlap. Intuitively, this means that a combination that is easier for the model to predict is preferred. During training, to simulate an extraction in progress, we randomly select in each sentence a subset of the entities and label them as already predicted entities.

### 3.3 Inference

For each sentence in the corpus, our model starts by predicting the most likely sequence of entities from the input token sequence alone since no mention has already been predicted. Then, we add them to the observed entities list and repeat the prediction until no more entities can be found.

### 3.4 Model parameters

We initialize the Transformer with CamemBERT [19] weights for DEFT and BioBERT [15] for GENIA unless mentioned otherwise, and the remaining parameters using the method of [6]. Dropout [28] is applied with a probability of 0.25 everywhere. We optimize the parameters by backpropagation with Adam [12] without weight decay, over 40 epochs for DEFT and 10 epochs for GENIA. We use two learning steps: one for the Transformer weights, initialized at  $4 \times 10^{-5}$ , and one for the rest of the model, initialized at  $9 \times 10^{-3}$ . The learning rate follows a linear decay schedule with a warmup for 10% of the steps. We insert the tag embeddings in the Transformers at layer  $L_{tag} = 6$  for BERT with 12 layers and 19 for BERT with 24 layers. On an Nvidia K80 GPU graphics card, learning on 100 documents takes about 20 minutes.

## 4 Experiments and discussions

### 4.1 Datasets

We conduct experiments on two nested named entity recognition datasets: DEFT [1] and GENIA [11], which present documents with different languages and different types and depths of entities (see Table 1). For GENIA, we perform splits following [4]: the last 10% of the sentences are used to test the model, the remaining 90% are the training set. For DEFT, we used the provided train and test splits. In both cases, we split the training data into 90% for training the model and 10% for the development (validation) set. We selected the hyperparameters by grid search on the development set and trained on the training split for GENIA, and both training and development splits for DEFT.

### 4.2 Baselines

We compare our results against a simple flat NER model composed of a transformer and a CRF decoding layer that can only predict non-nested mentions. Since a choice

	GENIA (English)			DEFT (French)		
	train	dev	test	train	dev	test
sentences	15022	1669	1855	1481	365	1024
documents	1599	190	213	82	18	67
mentions	47027	4469	5600	6439	1498	4791
mentions $D_0$	42965	4072	5007	5098	1226	3538
mentions $D_1$	3959	394	1282	1282	261	1163
mentions $D_2$	102	3	59	59	12	90
mentions $D_3$	1	0	0	0	0	0

Table 1: Statistics of the datasets.  $D_0$  are all the mentions that does not contain any mentions,  $D_1$  are the mentions that contain one or more  $D_0$  mention, etc.

is required during training as to which mentions should be predicted, we evaluate three modes: we only recover the shortest mention in a nested group, or only the largest, or let the model decide greedily. We also compare our model against the state-of-the-art models on GENIA and the other participants’ models on DEFT.

### 4.3 Results

The results of our system and the baselines are presented in Tables 2 and 3.

On the DEFT task 3.1, our model obtains the best F1 result of 0.66 (with the exact delimitation of mentions). On the DEFT task 3.2, the same model obtains a F1 of 0.778. Flat NER models lose between 10 and 20 points in F1, due to the large number of nested mentions. The 3.1 task containing longer entities, the flat large entities model, reaches the best performance of the flat models. Conversely, the flat short entities model obtains the best performance on the 3.2 task.

On the GENIA dataset, our best model reaches 0.7683 F1 with BioBERT large. We hypothesize that our method ranks lower on the latter dataset because it only uses BERT instead of BERT and other word features, and that the insertion of tags directly in BERT architecture may lead to loose some of the pretrained model abilities. We can also observe that flat NER is competitive with iterative models, which can be explained by the low ratio of nested mentions in the dataset.

### 4.4 Ablations

#### Tag embeddings

*Insertion layer* We study the effect of different tag embeddings configurations on the system’s performance. We can see in Figure 2 the validation f1 metrics of the system against the index of the BERT layer at which we add the embeddings. The performance decreases in the last layers since the system has no "time" to take the previous predictions into account. However, we also observe a maximum at layer 6 and lower performance in earlier layers.

<i>GENIA dataset</i>	Precision	Recall	F1
Ju et al. [9]	0.785	0.713	0.747
Wang et al. [32]	0.780	0.702	0.739
Wang and Lu [31]	0.770	0.733	0.751
Sohrab and Miwa [27]	0.932	0.642	0.771
Lin et al. [16]	0.758	0.739	0.748
Shibuya and Hovy [26]	0.763	0.747	0.755
Straková et al. [29] (BERT, Flair)			0.783
WANG et al. [33] (BERT, Flair)	0.803	0.783	0.793
Flat short entities (BERT base)	0.7929	0.6975	0.7422
Flat large entities (BERT base)	0.8149	0.7070	0.7571
Flat greedy (BERT base)	0.8141	0.7095	0.7582
(our) greedy (BERT base)	0.8126	0.7211	0.7641
(our) large-to-short (BERT base)	0.8016	0.7184	0.7577
(our) short-to-large (BERT base)	0.8028	0.7336	0.7666
(our) short-to-large (BERT large)	0.7933	0.7448	0.7683

Table 2: GENIA test performance (systems reaching performance above 0.74 F1)

	DEFT 3.1			DEFT 3.2			Global
	P	R	F1	P	R	F1	F1
HESGE (BERT large)	0.702	0.624	0.660	0.788	0.725	0.755	
Median DEFT			0.4557			0.6151	
Flat short (BERT base)	0.6093	0.2277	0.3315	0.7570	0.6690	0.7103	0.6228
Flat large (BERT base)	0.6093	0.6076	0.6085	0.7223	0.3152	0.4389	0.5094
Flat greedy (BERT base)	0.6082	0.3481	0.4428	0.7848	0.6046	0.6830	0.6201
(our) greedy	0.6263	0.6090	0.6175	0.7615	0.7416	0.7514	0.7134
(our) large to short	0.6257	0.6062	0.6158	0.7407	0.7470	0.7439	0.7082
(our) short to large	0.6105	0.6194	0.6149	0.7564	0.7449	0.7506	0.7120
(our) short to large (BERT large)	0.6613	0.6595	0.6604	0.7806	0.7759	0.7783	0.7447

Table 3: DEFT test performance

*Tag scheme* We analyze the performance of two common tag schemes: BIO (Begin, Inside, Outside) and BIOUL (BIO with Unary and Last tags) to encode observed (i.e., previously predicted entities). Results can be found in Table 4. The BIOUL tag scheme shows better overall results than the BIO scheme. This conclusion is similar to what others [14] have observed for flat named entity recognition. However, the BIOUL "reading" tag scheme's better results show that the system can better use previous entities for its predictions when given begin and end bounds.

**Depth prediction order** We study the effect that forced prediction order during training has on model performance. We compared three prediction modes: top to bottom, bottom to top, and greedy decoding. In the top to bottom mode, given a previously predicted entity at depth  $D$ , we force the model to predict a named entity located at depth  $D+1$ . In the bottom to top mode, we use the inverse depth as training order. Finally, in



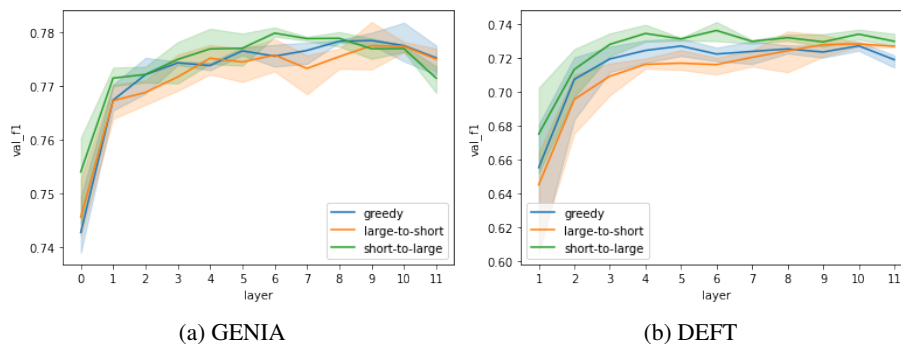


Fig. 2: Model performance on the GENIA and DEFT validation datasets wrt different insertion layer indices of the observed tag embeddings and training depth orders. Each point is the average over 3 runs. We only plotted the results from the second layer for DEFT to improve readability, all models reaching around 0.50 F1 when layer = 0.

	write BIO	write BIOUL
read BIO	0.7221 ( $\pm 2.98e^{-3}$ )	0.7341 ( $\pm 6.09e^{-3}$ )
read BIOUL	0.7261 ( $\pm 5.40e^{-3}$ )	<b>0.7368</b> ( $\pm 4.28e^{-3}$ )

Table 4: Performance of the BIO and BIOUL reading and writing tag schemes on the DEFT validation dataset.

greedy decoding mode, we let the model choose the mentions by selecting those closest to its prediction.

From Figure 2 we can observe that the short-to-large training order obtains the highest performance on both GENIA and DEFT validation splits. The large-to-short depth training order obtains the lowest accuracy. We hypothesize that learning to detect the smallest, and often easier, entities first leads the model to learn how to compose new entities from small entities. On the other hand, learning to predict large, and often more difficult, mentions first, must lead the model to overfit on these large mentions and fail to recover smaller nested mentions when the largest ones are wrongly predicted. The greedy training reaches an intermediate performance, so we conclude that a learned prediction order is suboptimal.

## 5 Conclusion

This paper proposes an architecture to perform named entity recognition based on iterative predictions and dynamic mention matching during training. We also provided insights into the model behavior and showed that training depth mention order impacts performance on auto-regressive layered named entity recognition models, and short-to-large order obtains the best results.

## Bibliography

- [1] Cardon, R., Grabar, N., Grouin, C., Hamon, T.: Presentation of the DEFT 2020 Challenge : open domain textual similarity and precise information extraction from clinical cases. In: Actes de la 6e conférence conjointe Journées d'Études sur la Parole (JEP, 33e édition), Traitement Automatique des Langues Naturelles (TALN, 27e édition), Rencontre des Étudiants Chercheurs en Informatique pour le Traitement Automatique des Langues (RÉCIT), pp. 1–13 (2020)
- [2] Dai, H.J., Lai, P.T., Chang, Y.C., Tsai, R.T.H.: Enhancing of chemical compound and drug name recognition using representative tag scheme and fine-grained tokenization. *Journal of Cheminformatics* **7**(Suppl 1), S14 (2015), ISSN 17582946
- [3] Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference* **1**, 4171–4186 (10 2019)
- [4] Finkel, J.R., Manning, C.D.: Nested named entity recognition. In: *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing Volume 1 - EMNLP '09*, vol. 1, p. 141, Association for Computational Linguistics, Morristown, NJ, USA (4 2009), ISBN 9781932432596
- [5] Fisher, J., Vlachos, A.: Merge and Label: A Novel Neural Network Architecture for Nested NER. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5840–5850, Association for Computational Linguistics, Stroudsburg, PA, USA (2019), ISSN 23318422
- [6] Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. *Journal of Machine Learning Research* **9**, 249–256 (2010), ISSN 15324435
- [7] Grabar, N., Claveau, V., Dalloux, C.: CAS: French Corpus with Clinical Cases. In: *Proceedings of the Ninth International Workshop on Health Text Mining and Information Analysis*, pp. 122–128, Association for Computational Linguistics, Stroudsburg, PA, USA (6 2018)
- [8] Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural Computation* **9**(8), 1735–1780 (11 1997), ISSN 0899-7667
- [9] Ju, M., Miwa, M., Ananiadou, S.: A Neural Layered Model for Nested Named Entity Recognition. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, vol. 1, pp. 1446–1459, Association for Computational Linguistics, Stroudsburg, PA, USA (2018), ISBN 9781948087278
- [10] Katiyar, A., Cardie, C.: Nested Named Entity Recognition Revisited. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, vol. 1, pp. 861–871, Association for Computational Linguistics, Stroudsburg, PA, USA (2018), ISBN 9781948087278

- [11] Kim, J.D., Ohta, T., Tateisi, Y., Tsujii, J.: GENIA corpus - A semantically annotated corpus for bio-textmining. *Bioinformatics* **19**(SUPPL. 1), i180–i182 (7 2003), ISSN 13674803
- [12] Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings (2015)
- [13] Lafferty, J., McCallum, A., Pereira, F.C.N.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. *ICML '01 Proceedings of the Eighteenth International Conference on Machine Learning* **8**(June), 282–289 (2001)
- [14] Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., Dyer, C.: 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, pp. 260–270, Association for Computational Linguistics, Stroudsburg, PA, USA (2016), ISBN 9781941643914, ISSN 1045-9227
- [15] Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H., Kang, J.: BioBERT: A pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **36**(4), 1234–1240 (2020), ISSN 14602059
- [16] Lin, H., Lu, Y., Han, X., Sun, L.: Sequence-to-Nuggets: Nested Entity Mention Detection via Anchor-Region Networks. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 5182–5192, Association for Computational Linguistics, Stroudsburg, PA, USA (2019), ISSN 23318422
- [17] Lu, W., Roth, D.: Joint Mention Extraction and Classification with Mention Hypergraphs. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, pp. 857–867, Association for Computational Linguistics, Stroudsburg, PA, USA (2015), ISBN 9781941643327
- [18] Marinho, Z., Mendes, A., Miranda, S., Nogueira, D.: Hierarchical Nested Named Entity Recognition. In: Proceedings of the 2nd Clinical Natural Language Processing Workshop, pp. 28–34, Association for Computational Linguistics, Stroudsburg, PA, USA (2019)
- [19] Martin, L., Muller, B., Ortiz Suarez, P.J., Dupont, Y., Romary, L., de la Clergerie, E., Seddah, D., Sagot, B.: CamemBERT: a Tasty French Language Model. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7203–7219, Association for Computational Linguistics, Stroudsburg, PA, USA (2020), ISSN 23318422
- [20] Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: 1st International Conference on Learning Representations, ICLR 2013 - Workshop Track Proceedings (2013)
- [21] Muis, A.O., Lu, W.: Labeling Gaps Between Words: Recognizing Overlapping Mentions with Mention Separators. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 2608–2618, Association for Computational Linguistics, Stroudsburg, PA, USA (2017), ISSN 23318422
- [22] Névéol, A., Cohen, K.B., Grouin, C., Hamon, T., Lavergne, T., Kelly, L., Goeuriot, L., Rey, G., Robert, A., Tannier, X., Zweigenbaum, P.: Clinical Information Extraction at the CLEF eHealth Evaluation lab 2016. Tech. rep. (2016)

- [23] Peters, M.E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep contextualized word representations. NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference **1**, 2227–2237 (2 2018)
- [24] Ratnov, L., Roth, D.: Design challenges and misconceptions in named entity recognition. In: Proceedings of the Thirteenth Conference on Computational Natural Language Learning - CoNLL '09, p. 147, Association for Computational Linguistics, Morristown, NJ, USA (2009), ISBN 9781932432299
- [25] Ravanbakhsh, S., Yu, C.N., Greiner, R.: A generalized loop correction method for approximate inference in graphical models. Tech. rep. (2012)
- [26] Shibuya, T., Hovy, E.: Nested Named Entity Recognition via Second-best Sequence Learning and Decoding. Transactions of the Association for Computational Linguistics **8**, 605–620 (12 2020), ISSN 2307-387X
- [27] Sohrab, M.G., Miwa, M.: Deep Exhaustive Model for Nested Named Entity Recognition. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 2843–2849, Association for Computational Linguistics, Stroudsburg, PA, USA (2018), ISBN 9781948087841
- [28] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research **15**(56), 1929–1958 (2014), ISSN 15337928
- [29] Straková, J., Straka, M., Hajic, J.: Neural Architectures for Nested NER through Linearization. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 5326–5331, Association for Computational Linguistics, Stroudsburg, PA, USA (2019), ISSN 23318422
- [30] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in Neural Information Processing Systems **2017-December**, 5999–6009 (6 2017), ISSN 10495258
- [31] Wang, B., Lu, W.: Neural Segmental Hypergraphs for Overlapping Mention Recognition. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 204–214, Association for Computational Linguistics, Stroudsburg, PA, USA (2018), ISSN 23318422
- [32] Wang, B., Lu, W., Wang, Y., Jin, H.: A Neural Transition-based Model for Nested Mention Recognition. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 1011–1017, Association for Computational Linguistics, Stroudsburg, PA, USA (2018), ISSN 23318422
- [33] WANG, J., Shou, L., Chen, K., Chen, G.: Pyramid: A Layered Model for Nested Named Entity Recognition. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5918–5928, Association for Computational Linguistics, Stroudsburg, PA, USA (2020)
- [34] Xu, M., Jiang, H., Watcharawittayakul, S.: A local detection approach for named entity recognition and mention detection. ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers) **1**, 1237–1247 (2017)
- [35] Zheng, C., Cai, Y., Xu, J., Leung, H.f., Xu, G.: A Boundary-aware Neural Model for Nested Named Entity Recognition. In: Proceedings of the 2019 Conference

on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 357–366, Association for Computational Linguistics, Stroudsburg, PA, USA (2019), ISBN 9781950737901