

Nested Named Entity Recognition as Single-Pass Sequence Labeling

Alberto Muñoz-Ortiz,¹ David Vilares,¹ Caio Corro,² Carlos Gómez-Rodríguez¹

¹ Universidade da Coruña, CITIC, Spain

² INSA Rennes, IRISA, Inria, CNRS, Université de Rennes, France
{alberto.munoz.ortiz, david.vilares, carlos.gomez}@udc.es
caio.corro@insa-rennes.fr

Abstract

We cast nested named entity recognition (NNER) as a sequence labeling task by leveraging prior work that linearizes constituency structures, effectively reducing the complexity of this structured prediction problem to straightforward token classification. By combining these constituency linearizations with pre-trained encoders, our method captures nested entities while performing exactly n tagging actions. Our approach achieves competitive performance compared to less efficient systems, and it can be trained using any off-the-shelf sequence labeling library.

1 Introduction

Named Entity Recognition (NER) involves identifying token spans that refer to specific named entities. Traditional approaches use the BIO scheme (Ramshaw and Marcus, 1995), which assigns labels to define the beginning (B), inside (I), and outside (O) tokens of each entity. BIO NER—and other BIO tasks—can be naturally addressed using sequence labeling approaches (Ratnaparkhi, 1996; Lafferty et al., 2001; Lample et al., 2016; Yang and Zhang, 2018). However, nested NER (NNER) challenges these approaches, as it involves entities embedded within other entities as long as their spans do not overlap. For example, in Figure 1, the person mention “*president of the USA*” contains the geopolitical named entity “*the USA*”. This nested structure breaks the BIO tagging assumption that each token belongs to only one entity. As a result, structured prediction is a better fit for NNER, as it can model interdependent spans and capture hierarchical relationships between nested entities.

Prior work has tackled NNER as a hierarchical or layered sequence labeling task, applying token-level tagging multiple times, either detecting outer entities first and then inner ones (Shibuya and Hovy, 2020), or identifying inner entities first

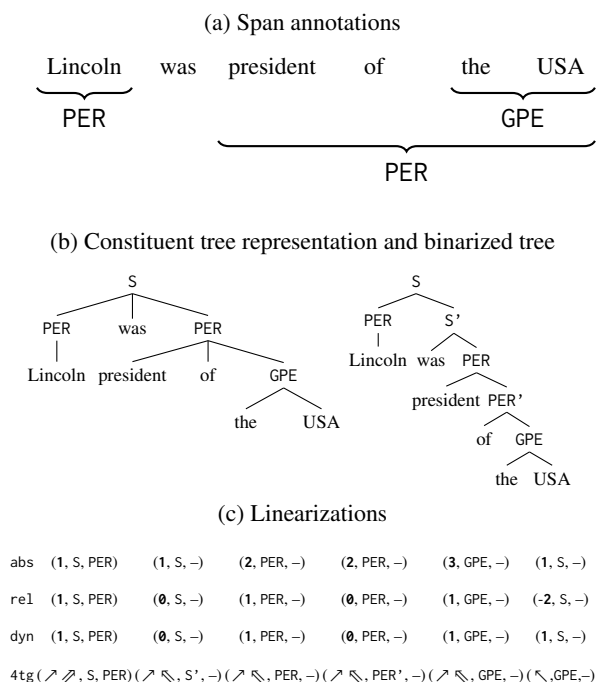


Figure 1: Illustration of a nested NER annotated sentence as (a) spans, (b) a constituent tree and its corresponding binarized tree, and (c) four ways to linearize that tree. We compare three depth-based—absolute, relative, dynamic—and one transition-based encoding—tetra-tagging (see §2.1). Binarized non-terminals in the tetra-tagging encoding are marked with ‘ \nearrow ’.

and then expanding outward (Wang et al., 2021). Other paradigms classify each possible text span as either an entity or not (Yu et al., 2020; Sohrab and Miwa, 2018; Yuan et al., 2022; Corro, 2023), use sequence-to-sequence techniques (Yan et al., 2021; Tan et al., 2021), or exploit hypergraph representations (Lu and Roth, 2015; Katiyar and Cardie, 2018; Yan et al., 2023).

Alternatively, some methods transform NNER into constituency parsing by mapping samples to constituent trees and relying on parsing algorithms (Finkel and Manning, 2009; Wang et al., 2018; Fu et al., 2021; Lou et al., 2022). More

recently, Yang and Tu (2022) applied a pointer-network-based bottom-up constituency parser to NNER, offering linear-time decoding. In this context, constituency tree linearization—first introduced by Gómez-Rodríguez and Vilares (2018) and later explored in related work (Kitaev and Klein, 2020; Amini and Cotterell, 2022)—reframes the structured prediction task as a sequence labeling problem. This approach achieves competitive performance while improving computational efficiency, as it outputs the tree in exactly n tagging actions. This is in contrast with previous sequence-to-sequence models for NNER where the number of tagging actions depends on the number of predicted mentions (Miculicich and Henderson, 2020).

Building on this line of work, we present a simple yet effective NNER approach that applies constituency tree linearizations within a standard sequence labeling framework.¹ Our method is single-pass, parser-free, and integrates seamlessly with pretrained encoders. It captures nested structures efficiently and achieves competitive benchmark results, all while requiring minimal implementation effort. Overall, the data preparation remains lightweight and fully automated, and it is no more complex than in traditional nested NER settings, thus preserving the simplicity of standard NER pipelines while adding structural expressiveness.

2 NNER Through Constituent Parsing as Sequence Labeling

We represent NNER structures using constituent trees, adopting a sequence labeling strategy for syntactic parsing. We hypothesize that this formulation aligns well with NNER due to two key properties: (1) the trees are relatively shallow, involving fewer output levels than in full parsing,² and (2) the encodings provide full coverage of well-formed trees, regardless of nesting depth. Furthermore, unlike TreeCRF models (Fu et al., 2021), span-based (Lou et al., 2022) or transition-based parsers (Gao et al., 2023), our method avoids complex decoding, enabling linear-time inference³— n tagging actions—with standard architectures.

¹Code is available at https://github.com/amunozo/nner_as_sl.

²In full constituent parsing, large gaps between nonterminals have been shown to challenge sequence tagging approaches (Vilares et al., 2019).

³Although the transformer architecture has quadratic complexity, we follow common practice and refer to the rest of the system’s complexity (Shibuya and Hovy, 2020; Corro, 2023; Lou et al., 2022).

Dataset	Enc.	# Labels	Missing	% dev	% test
		(n, c, u)	(n, c, u)		
ACE2004	abs	(6,11,8)	(0,1,0)	100.00	99.96
	dyn	(10,11,8)	(0,1,0)	100.00	99.96
	rel	(10,11,8)	(0,1,0)	100.00	99.96
	4tg	(8,22,8)	(1,6,1)	100.00	99.96
ACE2005	abs	(7,29,9)	(0,3,1)	99.88	99.98
	dyn	(9,29,9)	(0,3,1)	99.88	99.98
	rel	(10,29,9)	(0,3,1)	99.88	99.98
	4tg	(8,58,9)	(1,6,1)	99.88	99.98
GENIA	abs	(4,12,13)	(0,0,0)	100.00	100.00
	dyn	(8,12,13)	(0,0,0)	100.00	100.00
	rel	(8,12,13)	(0,0,0)	100.00	100.00
	4tg	(5,22,13)	(0,1,0)	>99.99	>99.99
NNE	abs	(7,264,401)	(0,14,37)	99.93	99.95
	dyn	(11,264,401)	(0,14,37)	99.93	99.95
	rel	(10,264,401)	(0,14,37)	99.93	99.95
	4tg	(13,399,401)	(0,18,37)	99.93	99.95

Table 1: Coverage statistics across datasets and encoding strategies. # Labels indicates the number of unique labels in the training set. Missing shows the number of labels in dev/test not seen during training. % dev/test indicates the percentage of dev/test labels covered by the training vocabulary.

2.1 Linearizations

We use three depth-based (absolute, relative, dynamic) and one transition-based (tetra-tagging) encoding strategies. These are illustrated in Figure 1 using a sentence with nested entities.

Let $\mathbf{w} = [w_1, w_2, \dots, w_{|\mathbf{w}|}]$ be an input word sequence, where w_i represents the word at position i . We define a label $l_i = (n_i, c_i, u_i)$ for each word w_i , where: (1) n_i is an integer indicating the number of common ancestors between words w_i and w_{i+1} , (2) c_i is the non-terminal symbol at their lowest common ancestor, and (3) u_i denotes the unary branch for w_i , if one exists.

In the **absolute encoding (abs)** (Gómez-Rodríguez and Vilares, 2018), n_i^{abs} represents the number of common ancestors between w_i and w_{i+1} . While abs often suffers from a large and sparse label space in full constituent parsing due to the depth and complexity of syntactic trees, we expect it to remain compact and easier to learn in NNER, where structures are shallower.

The **relative encoding (rel)** was also introduced in (Gómez-Rodríguez and Vilares, 2018). In order to reduce the label set, n_i^{rel} represents the difference in the number of common ancestors, i.e., for $i > 1$, $n_i^{\text{rel}} = n_i^{\text{abs}} - n_{i-1}^{\text{abs}}$ (and $n_1^{\text{rel}} = n_1^{\text{abs}}$).

Combining these two strategies, the **dynamic encoding (dyn)** (Vilares et al., 2019) selects, at each position i , the most suitable encoding—absolute or relative—for representing the relation between

tokens w_i and w_{i+1} . It defaults to the relative encoding and switches to absolute when two conditions are met: (1) the relative value $n_i^{\text{rel}} \leq -2$, indicating a sharp drop in tree depth with empty non-terminal levels that must be filled in subsequent steps; and (2) the corresponding absolute value $n_i^{\text{abs}} \leq 3$, ensuring the constituent node is close to the topmost tree level and the absolute label remains compact, avoiding sparsity and improving learnability. The dynamic encoding uses the absolute function only after abrupt structural changes. This affects few labels—0.75% in ACE2004, 0.73% in ACE2005, 0.15% in GENIA, and 1.16% in NNE—but many sentences: 14.64%, 12.45%, 3.86%, and 20.20%, respectively.

In addition to depth-based encodings, we also consider the **tetra-tagging (4tg)** linearization introduced by Kitaev and Klein (2020). Tetra-tagging is a left-corner transition-based parsing algorithm for binary trees, which can be adapted to sequence labeling. It assigns two tags per word: one for the word itself and one for the adjacent fencepost (i.e., the boundary between consecutive words). Each tag encodes whether the word (\swarrow, \nearrow) or the lowest common ancestor spanning the fencepost (\searrow, \nearrow) is a left or right child. To fit our sequence labeling framework, we pair each word and fencepost tag into a single label. We reuse the notation l_i , but in this case n_i is refactored to store the two tags associated to w_i . Nonterminal symbols and unary chains are added aside the tags as in depth-based encodings.

Label Vocabulary As entity spans are encoded with a finite set of labels observed during training, the model may face unseen labels at test time. Still, as shown in Table 1, the proposed encodings cover almost all entities in the development and test sets.

3 Experiments

We present our experimental setup and results.

Setup We use the MaChAmp framework (van der Goot et al., 2021) to train a multitask transformer-based model. Specifically, we use roberta-large (Liu et al., 2019) and bert-large-uncased (Devlin et al., 2019) as shared encoders for the ACE and NNE datasets, and biobert-large-cased-v1.1 (Lee et al., 2020) for GENIA. We include RoBERTa as a high-performing option, and BERT for comparison, since it is the most commonly used model in

Model	Comp.	ACE2004	ACE2005	GENIA	NNE
Other approaches ($O(n)$)					
W21	$O(n)$	86.1	84.7	78.7	–
Y22	$O(n)$	86.9	85.5	78.2	–
Other approaches (higher complexity)					
S20	$O(n^2)$	85.8	84.3	77.4	–
C23	$O(n^2)$	86.2	84.8	78.3	–
S21	$O(n^2)$	87.4	86.7	80.5	–
F21	$O(n^3)$	86.6 \pm 0.3	85.4 \pm 0.1	78.2 \pm 0.1	–
L22	$O(n^4)$	87.9	86.9	78.4	94.6
This work ($O(n)$)					
RoBERTa (ACE and NNE) and BioBERT					
abs	$O(n)$	86.8 \pm 0.3	85.0 \pm 0.3	76.6 \pm 0.3	94.1 \pm 0.3
rel	$O(n)$	86.9 \pm 0.3	85.5 \pm 0.3	75.1 \pm 0.7	94.3\pm0.1
dyn	$O(n)$	87.7\pm0.2	86.0\pm0.3	75.3 \pm 0.5	94.3\pm0.2
4tg	$O(n)$	86.1 \pm 0.5	84.7 \pm 0.5	74.6 \pm 0.8	94.1 \pm 0.1
BERT (for comparison)					
abs	$O(n)$	86.1 \pm 0.3	83.8 \pm 0.4	–	94.0 \pm 0.1
rel	$O(n)$	87.0 \pm 0.4	85.2 \pm 0.3	–	94.1 \pm 0.1
dyn	$O(n)$	87.0 \pm 0.1	85.2 \pm 0.5	–	94.2 \pm 0.0
4tg	$O(n)$	86.0 \pm 0.5	84.7 \pm 0.2	–	94.1 \pm 0.1

Table 2: F1-scores for different datasets and encodings, a comparison to prior work, and computational complexities. Highest F1 among $O(n)$ models is in bold.

related work. We add three task-specific linear classification heads to predict each atomic component of the label (n_i, c_i, u_i) . We adopt the default hyperparameters provided in the MaChAmp repository, as our primary goal is to show that these models perform well in off-the-shelf sequence labeling setups with simple plug-and-play scripts. In particular, for encoding and decoding, we use the scripts provided by the CoDeLin library (Roca et al., 2023). This library provides homogeneous post-processing to produce trees from ill-formed label sequences; more details about the post-processing are in Appendix E.

Datasets We train and evaluate our models on four popular English NNER benchmarks: GENIA (Kim et al., 2003), ACE 2004 (Doddington et al., 2004), ACE 2005 (Walker et al., 2006), and NNE (Ringland et al., 2019). For ACE and GENIA, we preprocess the datasets following Shibuya and Hovy (2020). Information about the datasets is shown in Appendix Table 8. Notably, most entities are flat and short and highly nested spans are rare.

Metrics We evaluate our models in terms of precision, recall, and F1-scores based on strict entity boundaries, requiring both correct spans and types. We also provide finer-grained evaluations (e.g., by entity length) and external comparisons.

	ACE2004		ACE2005		GENIA		NNE		
	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	
abs	0	84.4	87.2	83.1	86.7	78.9	79.8	88.8	92.3
	1	78.9	85.9	71.0	84.3	44.4	43.5	90.6	95.4
	≥ 1	77.7	74.6	69.8	67.1	44.5	37.5	88.8	88.4
	≥ 2	71.6	63.6	61.3	37.4	50.0	25.0	85.2	83.2
rel	0	83.5	85.8	81.6	85.9	77.7	76.1	87.8	91.7
	1	78.7	87.6	72.7	85.8	58.1	41.7	90.1	95.4
	≥ 1	75.9	72.6	71.1	68.8	57.4	23.5	88.9	87.3
	≥ 2	62.7	62.3	59.6	40.5	0.0	0.0	86.4	82.4
dyn	0	84.0	86.8	82.6	86.5	78.1	76.3	88.1	91.9
	1	79.4	88.1	75.3	86.9	51.7	42.9	90.4	95.6
	≥ 1	77.4	72.1	72.7	70.5	51.6	28.3	88.7	87.7
	≥ 2	68.0	66.2	58.6	49.6	0.0	0.0	85.3	82.7
4tg	0	83.1	86.0	81.2	85.7	78.2	77.8	89.3	92.2
	1	73.6	86.8	73.4	85.2	32.5	45.6	90.8	95.5
	≥ 1	70.0	72.9	70.8	71.0	30.5	36.2	88.2	88.7
	≥ 2	53.9	59.6	59.2	58.8	3.8	10.0	83.5	84.9
S20	0	–	–	80.4	76.1	80.3	77.9	–	–
	1	–	–	72.4	71.7	57.1	40.6	–	–
	≥ 1	–	–	73.2	70.4	57.1	40.6	–	–
	≥ 2	–	–	77.6	57.5	66.7	0.0	–	–

Table 3: Precision and recall by encoding and depth per dataset for entities at different depths. Best values for each dataset and depth shown in bold. Results for [Shibuya and Hovy \(2020\)](#) are added from their published scores.

3.1 Results

Table 2 reports precision, recall, and F1-scores for our sequence labeling models compared to previous work on NNER. All results are averaged over five runs to provide robust performance estimates (standard deviations are also reported).

Our approach outperforms both comparable and more complex methods on the ACE and NNE datasets. However, this is not the case for GENIA, where the results—particularly recall—are lower. Although the dyn encoding predominantly uses the relative function—as most entities are flat or shallow—it outperforms the other encodings on every dataset except GENIA, effectively leveraging the best of abs and rel linearizations. Even though the absolute function is rarely used in the dynamic encoding, it yields small but consistent F1 improvements of 0.74% for ACE2004, 0.39% for ACE2005, 0.22% for GENIA, and 0.02% for NNE. As more than 99% of entities in GENIA are flat or singly nested, the dyn and rel encodings represent the entities almost identically. Tetra-tagging is the worst-performing encoding in all but one setup.

To better understand how linearizations handle nested entities, we evaluate each encoding by analyzing model performance across different entity depth and length, using RoBERTa and BioBERT.

	ACE2004		ACE2005		GENIA		NNE		
	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	
abs	1	84.8	86.4	83.6	86.4	78.4	77.9	93.2	93.2
	2–4	88.1	87.1	79.2	82.2	74.8	73.9	89.8	90.3
	≥ 2	83.8	83.6	77.2	79.1	74.3	74.2	89.5	88.9
	5–9	76.8	78.2	74.1	76.8	69.3	76.7	85.0	73.6
	≥ 10	70.5	71.0	67.9	61.5	81.8	75.0	80.0	20.0
	rel	1	89.0	88.0	86.9	85.1	79.5	76.6	94.2
≥ 2		83.7	80.3	79.4	78.0	73.4	68.1	90.6	87.8
2–4		89.5	86.3	82.5	82.4	74.1	68.3	91.0	89.6
5–9		77.8	73.2	75.8	70.8	68.2	67.7	85.5	68.3
≥ 10		55.3	53.4	63.7	60.9	65.3	53.3	44.4	20.0
dyn		1	88.9	88.2	86.4	86.2	79.8	77.2	94.0
	≥ 2	84.6	82.2	80.0	79.3	73.4	66.8	90.3	88.1
	2–4	89.8	86.7	82.4	83.0	73.9	66.8	90.7	89.8
	5–9	77.4	75.4	76.7	72.6	69.2	67.7	84.4	69.9
	≥ 10	64.5	65.1	69.3	67.1	65.2	50.0	50.0	20.0
	4tg	1	88.3	87.6	86.1	85.8	78.7	77.8	93.9
≥ 2		79.9	80.4	75.3	77.0	69.0	70.2	88.8	89.3
2–4		87.8	85.5	79.8	80.8	72.2	70.4	89.9	91.0
5–9		72.6	74.6	72.2	71.2	56.7	68.9	76.9	72.0
≥ 10		48.3	57.8	52.6	62.1	16.8	58.3	18.2	30.0
S21		1	89.6	91.0	–	–	–	–	–
	≥ 2	85.3	83.7	–	–	–	–	–	–
	2–4	87.0	86.0	–	–	–	–	–	–
	5–9	84.3	83.2	–	–	–	–	–	–
	≥ 10	69.2	63.2	–	–	–	–	–	–

Table 4: Precision and recall by encoding and length for entities of different span lengths. Best values for each dataset and span length shown in bold. Results from [Shen et al. \(2021\)](#) are added for comparison.

Results by depth Table 3 presents precision and recall by depth for each encoding and dataset. Since the number of predicted entities may vary across runs, we follow [Shibuya and Hovy \(2020\)](#) and report the precision results from the run whose F1 is closest to the average, and the average recall over five runs. abs performs best for flat entities on all datasets, dyn leads for singly nested entities in ACE2004 and ACE2005, rel in GENIA, and 4tg and dyn tie in NNE. For entities nested two or more times, abs performs best in GENIA and achieves the highest precision on the ACE datasets, but is outperformed in recall except in GENIA. In NNE, 4tg achieves the highest recall. We also compare against external work that has reported these results, particularly [Shibuya and Hovy \(2020\)](#), where our method shows superior performance on flat entities but underperforms on deep nested structures.

Results by entity length Table 4 presents precision and recall by span length. Among the encodings, rel and dyn tend to perform best on single-token spans and those of 2–4 tokens. Results for 5–9 token spans are mixed, while abs consistently outperforms the others on longer spans (≥ 10 tokens), except in ACE2005, where it is

		ACE2004				ACE2005				GENIA				NNE			
		Best	Worst	Mean	Std	Best	Worst	Mean	Std	Best	Worst	Mean	Std	Best	Worst	Mean	Std
4tg	PER	FAC	78.9	10.6	PER	LOC	77.8	7.4	protein	DNA	73.1	3.0	ENERGY	BAND	73.7	29.8	
abs	PER	WEA	78.5	10.8	PER	LOC	78.6	7.2	protein	cell_type	74.7	3.4	DISEASE	BAND	74.4	29.5	
dyn	PER	FAC	81.3	8.3	PER	LOC	79.6	7.6	protein	DNA	73.1	3.2	ENERGY	BAND	74.0	29.4	
rel	PER	FAC	80.6	8.8	PER	LOC	79.3	7.7	protein	DNA	73.0	3.2	HURRICANE	BAND	73.4	30.4	

Table 5: F1-score statistics across datasets and encodings. The table identifies the best and worst performing entity labels and lists the mean and standard deviation of F1-scores across all labels.

surpassed by dyn. 4tg performs well on short spans—particularly in NNE—but its performance drops sharply as span length increases. These trends suggest that dyn is well-suited for frequent, shorter entities, while abs offers greater robustness on longer ones. In addition, we benchmark against external work we are aware of, such as the more complex model from Shen et al. (2021), where our methods are competitive for entities spanning fewer than four tokens, but the performance gap widens for longer entities where their specialized architecture appears more advantageous.

Results by entity type Table 5 provides a detailed analysis of F1-scores for each entity type, identifying the best and worst performers and summarizing the overall distribution with mean and standard deviation. PER (Person) consistently achieves the highest scores in both ACE datasets, while protein is the strongest performer in GENIA, regardless of the encoding used.

Focusing on the variance, scores in GENIA and ACE2005 are stable across types, with a standard deviation of around 3 and 7.5 respectively for all encodings. For ACE2004, the rel and dyn encodings show a standard deviation of approximately 8.5, while 4tg and abs present slightly higher variation at around 10. As expected, NNE shows a much higher standard deviation (~ 29.7), due to its large and sparse label set; however, this pattern is stable across encodings.

3.2 Discussion

The results show that with standard transformer-based encoders, NNER can be effectively learned as sequence labeling under various representations. As expected, for both flat and short entities, differences across encodings do not appear to substantially affect performance, since in these cases representing sentences as constituent trees does not lead to significant changes. For deep entities, more elaborate encodings appear to be beneficial compared to the naive absolute encoding. For longer

entities, results are mixed. This is likely influenced by the fact that longer entities can be either flat or deep, which may confound conclusions about the impact of different encodings. In particular, depth-based encodings perform better on the ACE and GENIA datasets, while 4tg performs comparably to the other encodings on the NNE dataset. In this context, a factor could be the relative ease of the NNE dataset. Although it includes many nested entities and a large label set, it also provides a substantial larger amount of training data (see Table 8). If the learning task becomes less demanding overall, differences between encodings are less pronounced, and the encoder may play a more prominent role. As shown by our NNE results, the narrow F1 range (94.0–94.6) across encodings and models suggests that, with enough data, different encodings are learned to similar capacity, thereby limiting observable performance differences.

4 Conclusion

We cast nested NER for the first time as a single-pass sequence labeling task by first transforming NNER annotations into constituent trees, and then leveraging the linearization options these trees enable. In particular, we explored both depth-based (absolute, relative and dynamic encodings) and transition-based encodings (tetra-tagging) previously proposed for this type of tree structure. Our experiments show that this approach models nested structures effectively, with competitive performance and no need for complex architectures, remaining lightweight and compatible with standard sequence labeling tools.

Among the explored encodings, the dynamic variant—which selects between absolute and relative schemes based on local structural cues—consistently achieves the highest F1 across datasets, except for GENIA. The transition-based tetra-tagging approach is competitive on flat and short entities but degrades on longer spans, where depth-based strategies prove more effective.

Limitations

Due to the limited availability of freely accessible NNER datasets, we focus on the most widely used English benchmarks, including both public and proprietary ones. These datasets are commonly used in prior work and provide a solid basis for comparison. While our approach is evaluated on these established datasets, it's important to note that results may vary across benchmarks, as factors such as annotation scheme, entity types, nesting depth, domain, and language can influence performance.

Acknowledgements

This work was funded by SCANNER-UDC (PID2020-113230RB-C21) funded by MICIU/AEI/10.13039/501100011033; Xunta de Galicia (ED431C 2024/02); GAP (PID2022-139308OA-I00) funded by MICIU/AEI/10.13039/501100011033/ and by ERDF, EU; Grant PRE2021-097001 funded by MICIU/AEI/10.13039/501100011033 and by ESF+ (predoctoral training grant associated to project PID2020-113230RB-C21); LATCHING (PID2023-147129OB-C21) funded by MICIU/AEI/10.13039/501100011033 and ERDF; TSI-100925-2023-1 funded by Ministry for Digital Transformation and Civil Service and “NextGenerationEU” PRTR; and Centro de Investigación de Galicia “CITIC”, funded by the Xunta de Galicia through the collaboration agreement between the Consellería de Cultura, Educación, Formación Profesional e Universidades and the Galician universities for the reinforcement of the research centres of the Galician University System (CIGUSA).

This research project was made possible through the access granted by the Galician Supercomputing Center (CESGA) to its supercomputing infrastructure. The supercomputer FinisTerra III and its permanent data storage system have been funded by the NextGeneration EU 2021 Recovery, Transformation and Resilience Plan, ICT2021-006904, and also from the Pluriregional Operational Programme of Spain 2014-2020 of the European Regional Development Fund (ERDF), ICTS-2019-02-CESGA-3, and from the State Programme for the Promotion of Scientific and Technical Research of Excellence of the State Plan for Scientific and Technical Research and Innovation 2013-2016 State subprogramme for scientific and technical infrastructures and equipment of ERDF, CESG15-DE-3114.

Caio Corro has received funding from the French Agence Nationale pour la Recherche under grant agreement InExtenso ANR-23-IAS1-0004 and SEMIAMOR ANR-23-CE23-0005.

References

- Afra Amini and Ryan Cotterell. 2022. [On parsing as tagging](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8884–8900, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Caio Corro. 2023. [A dynamic programming algorithm for span-based nested named-entity recognition in \$o\(n^2\)\$](#) . In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10712–10724, Toronto, Canada. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. [The automatic content extraction \(ACE\) program – tasks, data, and evaluation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Jenny Rose Finkel and Christopher D. Manning. 2009. [Nested named entity recognition](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 141–150, Singapore. Association for Computational Linguistics.
- Yao Fu, Chuanqi Tan, Mosha Chen, Songfang Huang, and Fei Huang. 2021. [Nested named entity recognition with partially-observed TreeCRFs](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12839–12847.
- Nan Gao, Bowei Yang, Yongjian Wang, and Peng Chen. 2023. [Ntam: A new transition-based attention model for nested named entity recognition](#). In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 273–286. Springer.
- Carlos Gómez-Rodríguez and David Vilares. 2018. [Constituent parsing as sequence labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1314–1324, Brussels, Belgium. Association for Computational Linguistics.

- Arzoo Katiyar and Claire Cardie. 2018. [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)*, pages 861–871, New Orleans, Louisiana. Association for Computational Linguistics.
- J.-D. Kim, T. Ohta, Y. Tateisi, and J. Tsujii. 2003. [Genia corpus—a semantically annotated corpus for biotextmining](#). *Bioinformatics*, 19(suppl_1):i180–i182.
- Nikita Kitaev and Dan Klein. 2020. [Tetra-tagging: Word-synchronous parsing with linear-time inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6255–6261, Online. Association for Computational Linguistics.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01)*, ICML '01, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- 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.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. BioBERT: A pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Chao Lou, Songlin Yang, and Kewei Tu. 2022. [Nested named entity recognition as latent lexicalized constituency parsing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6183–6198, Dublin, Ireland. Association for Computational Linguistics.
- Wei Lu and Dan Roth. 2015. [Joint mention extraction and classification with mention hypergraphs](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 857–867, Lisbon, Portugal. Association for Computational Linguistics.
- Lesly Miculicich and James Henderson. 2020. [Partially-supervised mention detection](#). In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 91–98, Barcelona, Spain (online). Association for Computational Linguistics.
- Lance Ramshaw and Mitch Marcus. 1995. [Text chunking using transformation-based learning](#). In *Third Workshop on Very Large Corpora*.
- Adwait Ratnaparkhi. 1996. [A maximum entropy model for part-of-speech tagging](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Nicky Ringland, Xiang Dai, Ben Hachey, Sarvnaz Karimi, Cecile Paris, and James R. Curran. 2019. [NNE: A dataset for nested named entity recognition in English newswire](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5176–5181, Florence, Italy. Association for Computational Linguistics.
- Diego Roca, David Vilares, and Carlos Gómez-Rodríguez. 2023. [A system for constituent and dependency tree linearization](#). In *Proceedings of XoveTIC Conference (XoveTIC '22)*, volume 14 of *Kalpa Publications in Computing*, pages 83–87. EasyChair.
- Yongliang Shen, Xinyin Ma, Zeqi Tan, Shuai Zhang, Wen Wang, and Weiming Lu. 2021. [Locate and label: A two-stage identifier for nested named entity recognition](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2782–2794, Online. Association for Computational Linguistics.
- Takashi Shibuya and Eduard Hovy. 2020. [Nested named entity recognition via second-best sequence learning and decoding](#). *Transactions of the Association for Computational Linguistics*, 8:605–620.
- Mohammad Golam Sohrab and Makoto Miwa. 2018. [Deep exhaustive model for nested named entity recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2843–2849, Brussels, Belgium. Association for Computational Linguistics.
- Zeqi Tan, Yongliang Shen, Shuai Zhang, Weiming Lu, and Yueting Zhuang. 2021. [A sequence-to-set network for nested named entity recognition](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*, pages 3936–3942. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Rob van der Goot, Ahmet Üstün, Alan Ramponi, Ibrahim Sharaf, and Barbara Plank. 2021. [Massive choice, ample tasks \(MaChAmp\): A toolkit for multi-task learning in NLP](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.

David Vilares, Mostafa Abdou, and Anders Søgaard. 2019. [Better, faster, stronger sequence tagging constituent parsers](#). 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 3372–3383, Minneapolis, Minnesota. Association for Computational Linguistics.

Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium*.

Bailin Wang, Wei Lu, Yu Wang, and Hongxia Jin. 2018. [A neural transition-based model for nested mention recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1011–1017, Brussels, Belgium. Association for Computational Linguistics.

Yiran Wang, Hiroyuki Shindo, Yuji Matsumoto, and Taro Watanabe. 2021. [Nested named entity recognition via explicitly excluding the influence of the best path](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3547–3557, Online. Association for Computational Linguistics.

Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. [A unified generative framework for various NER subtasks](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5808–5822, Online. Association for Computational Linguistics.

Yukun Yan, Bingling Cai, and Sen Song. 2023. [Nested named entity recognition as building local hypergraphs](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):13878–13886.

Jie Yang and Yue Zhang. 2018. [NCRF++: An open-source neural sequence labeling toolkit](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 74–79, Melbourne, Australia. Association for Computational Linguistics.

Songlin Yang and Kewei Tu. 2022. [Bottom-up constituency parsing and nested named entity recognition with pointer networks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2403–2416, Dublin, Ireland. Association for Computational Linguistics.

Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. [Named entity recognition as dependency parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6470–6476, Online. Association for Computational Linguistics.

Zheng Yuan, Chuanqi Tan, Songfang Huang, and Fei Huang. 2022. [Fusing heterogeneous factors with triaffine mechanism for nested named entity recognition](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3174–3186, Dublin, Ireland. Association for Computational Linguistics.

A Precision and Recall Results

Table 6 shows the precision and recall corresponding to the evaluation presented in Table 2 of the main text.

B Results from Best Runs

Table 7 presents the precision, recall, and F1-scores from the best run of each model, selected according to the highest F1-score.

C Dataset information

Table 8 presents statistics for each dataset, including information on entity depth and length.

D Data coverage

The encodings can represent all nested entity structures in the datasets without error. The only unrecoverable cases with these encodings arise from annotation inconsistencies—such as crossing spans in ACE2004 or inverted boundaries in NNE—rather than from limitations of the encodings themselves.

E Post-processing

Rather than imposing strict constraints during prediction, parsing as sequence labeling ensures well-formedness by applying simple heuristics to ill-formed label sequences. In our setup, we apply the heuristics provided by CoDeLin (Roca et al., 2023), i.e., (1) in depth-based encodings, if multiple labels are assigned to the same non-terminal node, we retain only the first; (2) if no label is assigned to a node, it is removed from the tree; (3) indices out of range (e.g., -7 in the relative encoding when there are fewer than 7 levels to go up) are changed to the nearest legal index; and (4) in tetra-tagging, if the label sequence specifies an invalid transition, the transition is skipped. Words that remain unattached due to skipped transitions are attached to the lowest non-terminal in the rightmost tree spine.

F Model Size and Budget

We fine-tune three models: BERT Large (340M parameters), BioBERT Large (340M), and RoBERTa

		ACE2004		ACE2005		GENIA		NNE	
Model	Comp.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.
Other approaches ($O(n)$)									
W21	$O(n)$	85.4	86.7	84.2	85.3	78.2	79.2	–	–
Y22	$O(n)$	86.6	87.3	84.6	86.4	78.1	78.3	–	–
Other approaches (higher complexity)									
S20	$O(n^2)$	85.9	85.7	83.8	84.9	77.8	76.9	–	–
C23	$O(n^2)$	87.4	85.0	84.4	85.3	79.3	77.3	–	–
S21	$O(n^2)$	87.4	87.4	86.1	87.3	80.2	80.9	–	–
F21	$O(n^3)$	86.7 \pm 0.4	86.5 \pm 0.4	84.5 \pm 0.4	86.4 \pm 0.2	78.2 \pm 0.7	78.2 \pm 0.8	–	–
L22	$O(n^4)$	87.4	88.4	86.0	87.9	78.4	78.5	94.3	95.0
This work ($O(n)$)									
RoBERTa (ACE and NNE) and BioBERT									
abs	$O(n)$	86.8 \pm 0.5	86.8 \pm 0.1	84.1 \pm 0.6	85.9 \pm 0.4	77.6 \pm 0.4	75.6 \pm 0.3	94.0 \pm 0.3	94.2 \pm 0.3
rel	$O(n)$	87.4 \pm 0.4	86.4 \pm 0.3	85.3 \pm 0.6	85.7 \pm 0.3	78.4 \pm 0.7	72.1 \pm 1.8	94.4\pm0.0	94.1 \pm 0.1
dyn	$O(n)$	88.0\pm0.3	87.3\pm0.2	85.6\pm0.4	86.4\pm0.5	78.5\pm0.5	72.5 \pm 1.3	94.3 \pm 0.2	94.2 \pm 0.2
4tg	$O(n)$	85.8 \pm 1.1	86.4 \pm 0.1	84.0 \pm 0.5	85.5 \pm 0.5	75.2 \pm 1.1	74.1 \pm 0.9	94.0 \pm 0.2	94.3\pm0.1
BERT (for comparison)									
abs	$O(n)$	85.9 \pm 0.4	86.2 \pm 0.5	82.7 \pm 0.4	85.0 \pm 0.4	–	–	94.1 \pm 0.1	94.0 \pm 0.1
rel	$O(n)$	87.9 \pm 0.4	86.0 \pm 0.4	84.9 \pm 0.7	85.5 \pm 0.0	–	–	94.4\pm0.1	93.8 \pm 0.1
dyn	$O(n)$	87.9 \pm 0.1	86.2 \pm 0.2	84.9 \pm 0.5	85.4 \pm 0.7	–	–	94.4\pm0.0	94.0 \pm 0.1
4tg	$O(n)$	86.0 \pm 0.7	86.0 \pm 0.3	83.6 \pm 0.2	85.8 \pm 0.3	–	–	94.1 \pm 0.1	94.1 \pm 0.1

Table 6: Performance metrics for each dataset and encoding. Among $O(n)$ methods, the top values for each dataset are in **bold**.

		ACE2004			ACE2005			GENIA			NNE		
Model		Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
abs		86.4	86.6	86.5	83.7	85.7	84.7	77.1	75.5	76.3	93.6	93.8	93.7
rel		86.9	86.1	86.5	85.1	85.1	85.1	79.3	69.8	74.3	94.3	93.9	94.1
dyn		87.7	87.2	87.4	85.2	85.7	85.5	78.5	71.4	74.8	93.9	93.9	93.9
4tg		83.8	86.6	85.1	83.1	84.7	83.9	74.6	72.6	73.5	93.7	94.2	93.9

Table 7: Metrics from each model’s best run, selected by F1-score. The highest value for each metric and dataset is highlighted in **bold**.

Large (355M). Each model is trained with four encoding strategies, across four datasets, and using five random seeds, resulting in a total of 240 training runs. Training was performed on a GPU cluster using NVIDIA A100 GPUs (40GB), with each run executed on a single GPU. The cumulative training time is estimated at approximately 300–360 GPU hours.

Characteristic	ACE2004	ACE2005	GENIA	NNE
Topic	News	News	Biomedical	News
Sentences	7 762	9 335	18 549	49 211
Sent. Length	21.8 ± 12.8	18.8 ± 12.3	26.5 ± 11.8	23.9 ± 11.3
Ent. Types	7	7	5	114
Total Entities	27 751	30 956	57 063	279 796
Avg. Depth	1.34	1.29	1.10	1.88
Med. Depth	1	1	1	2
Max. Depth	6	6	4	6
Percentage of Entities at Each Nesting Depth				
0	71.77	75.81	90.29	36.45
1	22.98	19.98	9.45	42.73
2	4.61	3.73	0.25	17.83
3	0.55	0.44	0.00	2.83
4	0.08	0.02	0.00	0.15
5	0.01	0.01	0.00	0.00
Percentage of Entities with Different Span Lengths				
1 Token	42.46	46.68	52.50	61.77
2-4 Tokens	39.49	37.25	42.42	35.06
5-9 Tokens	12.51	10.94	4.74	3.10
≥ 10 Tokens	5.54	5.13	0.35	0.06

Table 8: Dataset details, including the percentage of entities at each nesting depth and their distribution by length. Nesting depth refers to the number of hierarchical levels an entity is embedded within another, with 0 indicating a flat entity.