

Revisiting Projection-based Data Transfer for Cross-Lingual Named Entity Recognition in Low-Resource Languages

Andrei Politov^{*†}, Oleh Shkalikov^{*}, René Jäkel, Michael Färber

Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI),

TU Dresden, Dresden/Leipzig, Germany

{andrei.politov, oleh.shkalikov, rene.jaekel, michael.farber}@tu-dresden.de

Abstract

Cross-lingual Named Entity Recognition (NER) leverages knowledge transfer between languages to identify and classify named entities, making it particularly useful for low-resource languages. We show that the data-based cross-lingual transfer method is an effective technique for cross-lingual NER and can outperform multi-lingual language models for low-resource languages. This paper introduces two key enhancements to the annotation projection step in cross-lingual NER for low-resource languages. First, we explore refining word alignments using back-translation to improve accuracy. Second, we present a novel formalized projection approach of matching source entities with extracted target candidates. Through extensive experiments on two datasets spanning 57 languages, we demonstrated that our approach surpasses existing projection-based methods in low-resource settings. These findings highlight the robustness of projection-based data transfer as an alternative to model-based methods for cross-lingual named entity recognition in low-resource languages.

1 Introduction

Named Entity Recognition is well-studied in Natural Language Processing (NLP), but remains a challenge for low-resource languages due to the lack of manual annotation (Pakhale, 2023). Of the roughly 7,000 languages spoken worldwide, most are low-resource, with over 2,800 endangered (Eberhard et al., 2020). Cross-lingual approaches present a promising solution to address the scarcity of labelled data in these languages.

Cross-lingual NER methods can be categorized into *model transfer* and *data-based transfer* approaches (García-Ferrero et al., 2022). *Model transfer approaches* depend on the ability of multi-lingual models to convey task-specific knowledge across languages. *Data-based methods* automate labelling through translation and annotation projection processes while leveraging advancements in multi-lingual language models to enable zero-shot cross-lingual transfer. This approach allows models trained in high-resource languages to identify and classify named entities in other languages without additional annotated data. Additionally, categorization can be done through two approaches: *translate-test*, which labels original sentences in zero-shot settings, and *translate-train*, which generates labelled data to train a NER model.

Here we contribute to the field of cross-lingual NER by demonstrating the effectiveness of a data-based cross-lingual transfer method that achieves comparable and, in some cases, higher performance of multilingual language models in low- and extremely low-resource language scenarios.

Our work focuses on the projection phase of cross-lingual NER pipelines, introducing two improvements to projection-based methods. First, we propose a method specifically designed to improve word-to-word alignments. Second, we present a novel formalized projection approach of matching source entities with extracted target candidates. The proposed methods support *translate-train* and *translate-test* setups, achieving performance on par with model-based cross-lingual transfer techniques while offering greater flexibility. We evaluated our approach using the XTREME (Ruder et al., 2023) and MasakhaNER2 (Adelani et al., 2022) datasets comprising 57 languages in total in *translate-test* settings. The source code and the evaluation results are provided in the GitHub repository¹.

^{*}These authors contributed equally.

[†]Corresponding author.

¹<https://github.com/Cross-Lingual-NER/Projection-Data-Transfer-Cross-Lingual-NER>

2 Related Work

Model transfer methods leverage the ability of models to transfer task-specific knowledge across languages. For example, multilingual models like mBERT (Devlin et al., 2019) and XLM-RoBERTa (Conneau et al., 2020) are trained on high-resource languages and applied to low-resource languages without modification. Torge et al. (2023) demonstrated improved performance when models were fine-tuned on labelled data or pre-trained on a related language. However, low-resource languages often lack sufficient data, and transfer quality diminishes when applied to very different target languages.

Data-based methods employ labelled datasets, often available in high-resource languages, to perform labelling tasks in the target language. They include fully artificial data generation, like MulDA (Liu et al., 2021), and annotation projection methods. This paper focuses on the latter, which typically involves three steps: (i) translating the original sentence from the target (low-resource) to the source (high-resource) language, (ii) applying a NER model to the translated sentence, and (iii) projecting the labels back to the original sentence. While translation and NER use established models such as BERT (Devlin et al., 2019), many methods have been developed for the projection step.

The first major group (Yang et al., 2022; García-Ferrero et al., 2023; Parekh et al., 2024; Le et al., 2024) of projection methods is based on back-translation, where labelled source sentences or their parts are translated back to the target language, preserving the labels. EasyProject (Chen et al., 2023) is a translate-train method that employs the insertion of special markers, specifically square brackets, around source entities. The marked sentence is then passed to the translation model, which independently translates the entire sentence and each source entity. Afterwards, fuzzy string matching is used to project labels: for each substring in the back-translated sentence surrounded by markers, the method identifies the highest fuzzy match for the corresponding translation of the source entity and assigns the appropriate label.

Another type of projection method is based on word-to-word alignments (García-Ferrero et al., 2022; HWA et al., 2005; Tiedemann, 2015; Fei et al., 2020; Schäfer et al., 2022; Poncelas et al., 2023). The general idea is to compute word-to-word correspondence between words of a labelled

sentence in a source language and an original sentence in a target language. The entity’s label is projected onto target words that align with any of the entity’s words. García-Ferrero et al. (2022) have shown that using contextualized neural network-based aligners such as SimAlign (Jalili Sabet et al., 2020) or AWESoME (Dou and Neubig, 2021) is significantly more beneficial than statistical alignment tools like FastAlign (Dyer et al., 2013), but still can produce wrong alignments and therefore lead to projection errors.

3 Methodology

Our proposed approach focuses on projection-based methods that involve word-to-word alignments. We present two improvements (see Figure 1) to existing methods which are intended to be useful for languages that are under-presented in pre-trained language models. Firstly, we investigate an alternative alignment direction to address the known issue of word-to-word alignment quality. Secondly, we reformulate the annotation projection task as a bipartite matching problem between source entities and target candidates, using alignment-based matching scores to formalize the problem and eliminate reliance on heuristics, thereby facilitating method extension.

3.1 Alignment Direction

In projection-based pipelines, errors can arise at all three stages, diminishing the quality of resulting labels. Handling errors caused by forward translation and source NER models can be challenging. We aim to address projection errors caused by incorrect alignments.

Our approach involves computing word-to-word alignments between the original sentence and its back-translated labelled counterpart in the target language (i.e., target-to-target alignments see Figure 1 a). This method is motivated by the expectation that aligning words within the same language is easier than across different languages. This is particularly relevant for low-resource target languages, which often differ significantly from high-resource source languages.

Preserving entities during back-translation is crucial for projecting entities with the use of word alignments between original and back-translated sentences. To achieve this, we employed EasyProject (Chen et al., 2023) as outlined in the previous section.

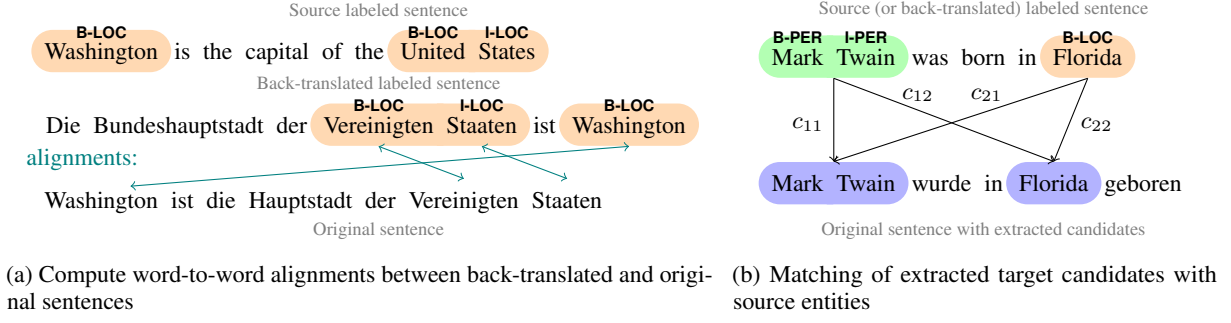


Figure 1: Proposed improvements to projection-based cross-lingual NER methods

3.2 Candidate Matching

The existing methods for addressing problems caused by incorrect alignments such as split annotation, annotation collision and wrong projection fully rely on heuristics (García-Ferrero et al., 2022). We consider that the main reason for these issues is a lack of information about any possible entity candidates in the original sentence in a target language. Instead, we propose to generate target entity candidates and match source entities with candidates by solving the weighted bipartite matching problem with additional constraints.

Let S be a set of source entity spans and T a set of target candidate spans. Then $x_{p^{src}, p^{tgt}}$ is a binary variable which represents whether a source entity $p^{src} \in S$ is being projected to a target candidate $p^{tgt} \in T$. Then the source entity-target candidate matching problem can be formulated as follows:

$$\begin{cases} \max_x \sum_{p^{src}, p^{tgt} \in S \times T} c_{p^{src}, p^{tgt}} x_{p^{src}, p^{tgt}} \\ x_{p_1} + x_{p_2} \leq 1, \quad [i_{p_1}^{tgt}, j_{p_1}^{tgt}] \cap [i_{p_2}^{tgt}, j_{p_2}^{tgt}] \neq \emptyset \\ \sum_{p^{tgt} \in T} x_{p^{src}, p^{tgt}} \leq 1, \quad \forall p^{src} \in S \\ x_{p^{src}, p^{tgt}} \in \{0, 1\}, \quad \forall (p^{src}, p^{tgt}) \in S \times T \end{cases} \quad (1)$$

where $p^{tgt} = (i_p^{tgt}, j_p^{tgt}) \in T$ is a candidate span represented as an index of the starting and the ending word, c is a score of matching. The first set of constraints represents that it is prohibited to project one or several different source entities to the overlapped candidates. The second ensures that all source entities will be projected.

The generation of target candidates is carried out with either N-grams- based or NER model-based candidate extraction. The former considers all continuous word sequences as candidates, while the latter predicts the candidate's spans using a multi-lingual NER model (ignoring predicted classes).

To calculate scores c from Equation 1 of matching between source entities and target candidates word-to-word alignments are being used:

$$c_{p^{src}, p^{tgt}} = \frac{a_{p^{src}, p^{tgt}}}{j_p^{src} - i_p^{src} + j_p^{tgt} - i_p^{tgt}} \quad (2)$$

where $a_{p^{src}, p^{tgt}}$ is a number of aligned words between a source entity and a target candidate. The motivation under this cost is to align entities and candidates based on the count of aligned words, considering source and target lengths to avoid matching with candidates with a lot of nonaligned words and handle single-word misalignments.

The complexity of the proposed problem remains an open question. Notably, it is not a straightforward instance of the maximum weight full bipartite matching problem, which can be solved in polynomial time, due to the first set of constraints that prevents projections onto overlapping candidates (i.e. some projections are mutually exclusive). In NER model-based candidate extraction, where no overlapping candidates exist, the problem reduces to a maximum weight bipartite matching.

To solve the problem in a general formulation, we propose a greedy approximate algorithm, which iteratively selects the projection with the maximum non-zero matching cost, performs this projection, and excludes all candidates that overlap with the projected candidate as well as the projected source entity.

The proposed concept of target candidate extraction and matching is structurally similar to T-Projection by García-Ferrero et al. (2023), with two key differences. T-Projection uses a fine-tuned T5 model, limiting target languages and producing candidates absent in the original sentence. For matching, T-Projection employs NMTScore by Vamvas and Sennrich (2022), while we use word-to-word alignments.

4 Experiments

We performed an intrinsic evaluation of the efficiency of our approaches across a total of 57 languages using the XTREME (Hu et al., 2020) (39 languages) and MasakhaNER2 datasets (excluding Ghomálá and Naijá languages due to limitations in translation model support - 18 languages in total). This evaluation encompasses the full pipeline, considering both translation and source NER model performance.

For a comparative analysis of existing and proposed approaches, we (re)implemented the aforementioned projection methods according to their original papers. In particular, we reimplemented the heuristic word-to-word alignment-based approach outlined by García-Ferrero et al. (2022). We enhanced this heuristic by introducing a word count ratio threshold of 0.8 to better handle misaligned unitary words. Additionally, we reimplemented the EasyProject method, which performs back-translation of labelled source sentences, using original, fine-tuned by authors, NLLB-200-3.3B² model. This back-translated output is then used for annotation projection, relying on word-to-word alignments computed between the original and labelled back-translated sentences in the same language (denoted as *tgt2tgt*).

NLLB200-3.3B³ (Costa-jussà et al., 2022) was employed as a translation model for all experiments. The XLM-R-Large model⁴, fine-tuned on the English split of the CONLL2003 (Tjong Kim Sang and De Meulder, 2003), served as both the source model and for target candidate extraction, as well as for model transfer experiments. We ignored MISC entities predicted by this model in the first set of experiments since this class does not exist in the MasakhaNER2 and XTREME datasets. For computing word-to-word alignments, we used the original implementations of SimAlign and non-finetuned AWESOME neural aligners with the default settings (with MBERT model).

As the evaluation involved full pipelines, the resulting metrics were influenced by both translation quality and the performance of the NER models. To ensure a fair and consistent comparison of the proposed methods, we employed the same models for translation and source labelling throughout all

experiments. For tasks involving the proposed integer linear programming (ILP) formulation of the projection problem, we utilized the previously described greedy approximation algorithm to derive solutions.

Evaluation results for the full pipelines are given in Table 1.

As shown in Table 1, candidate matching methods consistently deliver a strong performance. The proposed approach involving n-gram candidates extraction (*n-gram cand.*), compared to heuristics (since n-gram does not limit a set of candidates as NER cand. do), provide comparable or superior results while offering greater flexibility and avoiding hyperparameter optimization.

The NER model-based extraction (*NER cand.*) generally outperforms model transfer by effectively correcting labels for correctly predicted spans, resulting in greater accuracy particularly when model transfer mislabels these spans. It also surpasses the n-gram approach and achieves results comparable to model transfer because of more fine-grained candidates.

The model transfer generally performs better on the XTREME dataset, but candidate matching methods surpass heuristic approaches in most of the 36 languages, except for Bengali, Kazakh, and Swahili. The first may happen due to the model’s exposure to these languages or their partial representations during pretraining, despite being fine-tuned only on English data.

Although the average score for the MasakhaNER2 dataset is modest, the proposed method performs better than heuristics in 10 languages and worse in 8 out of 18 total languages. The full list can be found in the appendix. This discrepancy may be attributed to the simpler morphological structures in the first group (where proposed methods perform better), while the second group, especially languages like Xhosa and Zulu (Maho, 1999), presents greater morphological complexity, including noun class systems and agreement patterns.

The proposed method with target-to-target alignment direction generally does not outperform the source-to-target method, except for Japanese, due to errors introduced during back-translation, highlighting a potential area for future research.

Additional experiments, described in the appendix, evaluate the performance of the projection step independently. Table 2 shows projection per-

All models are from the HF Hub

²ychenNLP/nllb-200-3.3B-easyproject

³facebook/nllb-200-3.3B

⁴FacebookAI/xlm-roberta-large-finetuned-conll03-english

Approach	Align. dir.	XTREME					MasakhaNER2		
		yo	bn	et	fi	avg	bam	twi	avg
Model transfer	-	32.3	37.8	67.7	72.1	50.9	43.0	46.4	52.1
Heuristic SimAlign	src2tgt	32.8	36.9	56.6	59.1	41.8	49.3	71.8	66.6
	tgt2tgt	20.1	34.8	39.9	50.8	34.9	44.3	5.9	42.7
Heuristic AWESoME	src2tgt	33.3	38.8	56.0	58.6	41.5	49.7	74.6	67.3
	tgt2tgt	15.2	34.1	40.7	50.4	34.6	43.4	3.8	42.1
n-gram cand. SimAlign	src2tgt	29.7	36.5	60.9	62.5	43.3	48.9	69.5	66.4
	tgt2tgt	17.5	36.4	43.4	52.5	36.4	42.5	5.2	42.4
n-gram cand. AWESoME	src2tgt	28.8	36.5	60.0	61.7	42.3	48.3	70.9	66.7
	tgt2tgt	16.1	35.2	43.8	52.0	35.9	41.1	3.9	42.0
NER cand. SimAlign	src2tgt	52.0	38.3	58.6	61.2	46.4	55.3	69.3	63.0
	tgt2tgt	34.0	30.7	43.1	53.4	39.1	46.9	8.5	44.0
NER cand. AWESoME	src2tgt	50.2	38.2	58.0	60.5	45.9	55.0	69.1	62.5
	tgt2tgt	27.7	30.7	42.4	52.6	38.5	46.4	6.8	43.3

Table 1: F1 scores for various full pipelines and alignment directions on XTREME (first section) and MasakhaNER2 (second section). *Heuristic SimAlign* and *Heuristic AWESoME* are heuristic approaches, while *n-gram/NER cand. aligner.name* refers to the proposed candidate matching method with the specified aligner. The first columns show the language where the proposed method outperforms the heuristic the most, the seconds indicate where it underperforms the most, and the last columns provide the average results across all languages. **Bold** values are the overall best, and underlined values indicate the best projection-based approaches. Estonian (et) and Finnish (fi) are given as typical examples.

formance on pre-labelled Europarl parallel texts (Agerri et al., 2018), excluding translation and source NER labelling errors. It highlights that candidate matching methods yield results comparable to or better than prior approaches. The NER-based target candidates approach underperforms due to imperfect spans but surpasses plain model transfer by correcting mislabeled spans via source entity projection.

5 Conclusion

In this study, we presented novel annotation projection methods based on word-to-word alignments for cross-lingual NER.

The idea to compute word-to-word alignments between the original and back-translated labelled sentences in the same language, aimed at enhancing the quality of these alignments, did not produce the desired outcomes. This approach encountered significant challenges, primarily due to errors that occurred during the back-translation process.

In contrast, the proposed method of extracting candidates and matching them with source entities showed robust results. More specifically, the proposed formulation generally outperformed previous word-to-word alignment-based projection methods that relied on heuristics to deal with incorrect alignments.

By using the same NER model for candidate ex-

traction as in model transfer, the proposed approach can outperform model transfer. This is achieved by refining the labels for correctly predicted spans through projection from source entities.

Despite its advantages, the proposed approach remains heavily dependent on the quality of word-to-word alignments. However, the formulated ILP problem incorporates these alignments into matching scores that can be combined with other strategies using a weighted sum.

Our findings demonstrate that the projection-based data transfer approach can be a robust alternative to model-based methods for cross-lingual named entity recognition in low-resource languages.

Future research could aim to improve candidate extraction and explore alternative matching costs in addition to the alignment-based one. The proposed formulation, in contrast to heuristic approaches, facilitates the integration of various scoring mechanisms, allowing for the fusion of different scores to effectively address the limitations associated with each individual method.

Moreover, exploring the usage of LLMs for the projection step in cross-lingual NER pipelines shows potential, indicating that the development of multilingual LLMs could help enhance the performance of NER tasks across diverse languages, especially when working with limited labelled data.

Limitations

Translation Model Dependency: The performance of the proposed methods relies on the quality of the translation model used – in our case NLLB200-3.3B (Costa-jussà et al., 2022). Limitations in translation accuracy for certain languages may propagate errors through the pipeline, especially for morphologically complex or resource-scarce languages.

NER Model Dependency: Models used for extracting candidates or labelling translated sentences in the source language can be a source of errors. Incorrect predictions or omissions of entities by a model, coupled with the limited capability to correct such errors on the projection step, can adversely affect the quality of the resulting labelling of the original sentence. In our experiments, we rely on the XLM-R-Large model, fine-tuned on the English split of CONLL2003, although performance metrics may vary with different models.

Word-to-Word Alignment Model Dependency: The matching scores in the proposed ILP formulation for the projection step are computed based on word-to-word alignments. Therefore, the quality of the projection is inherently bounded by the quality of these alignments. In our study, we utilized state-of-the-art neural-based alignment models, specifically SimAlign and AWESoME. These models surpass previous statistically-based aligners as they incorporate the context of entire sentences. However, their performance remains limited. Furthermore, the quality of alignments varies between languages, which can be attributed to the representation of languages in the pretraining datasets of the models, as well as the inherent linguistic properties and structural differences among languages.

Dataset Variability: The proposed method demonstrates varying effectiveness across datasets, performing well on less complex languages but struggling with those that exhibit higher morphological complexity (e.g., Xhosa and Zulu). This indicates that additional adaptations may be needed for specific linguistic features.

Generalization Across Languages: The candidate matching method shows superior performance for most languages but underperforms in specific cases (e.g., Bengali, Kazakh, and Swahili), potentially due to inadequate representation in pre-training.

Optimization Heuristics: While the proposed optimization-based projection method reduces re-

liance on heuristics, the greedy algorithm used to solve the optimization problem may not achieve global optima in all scenarios.

Acknowledgements

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: SCADS24B.

The authors gratefully acknowledge the computing time made available to them on the high-performance computer at the NHR Center of TU Dresden. This center is jointly supported by the Federal Ministry of Education and Research and the state governments participating in the NHR.

Andrei Politov is deeply grateful to his late mother for her unwavering belief in him and the inspiration she provided throughout the development of this research.

References

- David Ifeoluwa Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba O. Alabi, Shamsuddeen H. Muhammad, Peter Nabende, Cheikh M. Bamba Dione, Andiswa Bukula, Rooweither Mabuya, Bonaventure F. P. Dossou, Blessing Sibanda, Happy Buzaaba, Jonathan Mukiibi, Godson Kalipe, Derguene Mbaye, Amelia Taylor, Fatoumata Kabore, Chris Chinenye Emezue, Anuoluwapo Aremu, Perez Ogayo, Catherine Gitau, Edwin Munkoh-Buabeng, Victoire M. Koagne, Allahsera Auguste Tapo, Tebogo Macucwa, Vukosi Marivate, Elvis Mboning, Tajuddeen Gwadabe, Tosin Adewumi, Orevaoghene Ahia, Joyce Nakatumba-Nabende, Neo L. Mokono, Ignatius Ezeani, Chiamaka Chukwuneke, Mofetoluwa Adeyemi, Gilles Q. Hacheme, Idris Abdulmumin, Odunayo Ogundepo, Oreen Yousuf, Tatiana Moteu Ngoli, and Dietrich Klakow. 2022. Masakhaner 2.0: Africa-centric transfer learning for named entity recognition.
- Rodrigo Agerri, Yiling Chung, Itziar Aldabe, Nora Aranberri, Gorka Labaka, and German Rigau. 2018. Building named entity recognition taggers via parallel corpora. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Yang Chen, Chao Jiang, Alan Ritter, and Wei Xu. 2023. Frustratingly easy label projection for cross-lingual

- transfer. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5775–5796, Toronto, Canada. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- 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.
- Zi-Yi Dou and Graham Neubig. 2021. Word alignment by fine-tuning embeddings on parallel corpora. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2112–2128, Online. Association for Computational Linguistics.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig. 2020. *Ethnologue: Languages of the World*, 23 edition. SIL International, Dallas.
- Hao Fei, Meishan Zhang, and Donghong Ji. 2020. Cross-lingual semantic role labeling with high-quality translated training corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7014–7026, Online. Association for Computational Linguistics.
- Iker García-Ferrero, Rodrigo Agerri, and German Rigau. 2022. Model and data transfer for cross-lingual sequence labelling in zero-resource settings. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6403–6416, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Iker García-Ferrero, Rodrigo Agerri, and German Rigau. 2023. T-projection: High quality annotation projection for sequence labeling tasks. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15203–15217, Singapore. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- REBECCA HWA, PHILIP RESNIK, AMY WEINBERG, CLARA CABEZAS, and OKAN KOLAK. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering*, 11(3):311–325.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Duong Minh Le, Yang Chen, Alan Ritter, and Wei Xu. 2024. Constrained decoding for cross-lingual label projection. *ArXiv*, abs/2402.03131.
- Linlin Liu, Bosheng Ding, Lidong Bing, Shafiq Joty, Luo Si, and Chunyan Miao. 2021. MulDA: A multilingual data augmentation framework for low-resource cross-lingual NER. 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 5834–5846, Online. Association for Computational Linguistics.
- J.F. Maho. 1999. *A Comparative Study of Bantu Noun Classes*. Acta Universitatis Gothoburgensis: Orientalia et Africana Gothoburgensia. Acta Universitatis Gothoburgensis.
- Kalyani Pakhale. 2023. Comprehensive overview of named entity recognition: Models, domain-specific applications and challenges.
- Madhumangal Pal and G. P. Bhattacharjee. 1996. A sequential algorithm for finding a maximum weight K -independent set on interval graphs. *Int. J. Comput. Math.*, 60(3-4):205–214.
- Tanmay Parekh, I-Hung Hsu, Kuan-Hao Huang, Kai-Wei Chang, and Nanyun Peng. 2024. Contextual label projection for cross-lingual structured prediction. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5738–5757, Mexico City, Mexico. Association for Computational Linguistics.

Alberto Poncelas, Maksim Tkachenko, and Ohnmar Htun. 2023. Sakura at SemEval-2023 task 2: Data augmentation via translation. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 1718–1722, Toronto, Canada. Association for Computational Linguistics.

Sebastian Ruder, Jonathan Clark, Alexander Gutkin, Mihir Kale, Min Ma, Massimo Nicosia, Shruti Rijhwani, Parker Riley, Jean-Michel Sarr, Xinyi Wang, John Wieting, Nitish Gupta, Anna Katanova, Christo Kirov, Dana Dickinson, Brian Roark, Bidisha Samanta, Connie Tao, David Adelani, Vera Axelrod, Isaac Caswell, Colin Cherry, Dan Garrette, Reeve Ingle, Melvin Johnson, Dmitry Pantelev, and Partha Talukdar. 2023. XTREME-UP: A user-centric scarce-data benchmark for under-represented languages. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1856–1884, Singapore. Association for Computational Linguistics.

Henning Schäfer, Ahmad Idrissi-Yaghir, Peter Horn, and Christoph Friedrich. 2022. Cross-language transfer of high-quality annotations: Combining neural machine translation with cross-linguistic span alignment to apply NER to clinical texts in a low-resource language. In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 53–62, Seattle, WA. Association for Computational Linguistics.

Jörg Tiedemann. 2015. Improving the cross-lingual projection of syntactic dependencies. In *Proceedings of the 20th Nordic Conference of Computational Linguistics (NODALIDA 2015)*, pages 191–199, Vilnius, Lithuania. Linköping University Electronic Press, Sweden.

Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.

Sunna Torge, Andrei Politov, Christoph Lehmann, Bochra Saffar, and Ziyang Tao. 2023. Named entity recognition for low-resource languages - profiting from language families. In *Proceedings of the 9th Workshop on Slavic Natural Language Processing 2023 (SlavicNLP 2023)*, pages 1–10, Dubrovnik, Croatia. Association for Computational Linguistics.

Jannis Vamvas and Rico Sennrich. 2022. NMTScore: A multilingual analysis of translation-based text similarity measures. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 198–213, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Jian Yang, Shaohan Huang, Shuming Ma, Yuwei Yin, Li Dong, Dongdong Zhang, Hongcheng Guo, Zhoujun Li, and Furu Wei. 2022. CROP: Zero-shot cross-lingual named entity recognition with multilingual labeled sequence translation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*,

pages 486–496, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Appendix

Isolated Evaluation of the Projection Step

Table 2 depicts performance only of the projection step, excluding translation and source NER labeling errors, on labelled parallel texts from the Europarl-based NER dataset⁵ (Agerri et al., 2018). Since our experiments with the tgt2tgt alignment direction yielded negative results, all pipelines presented in the table are only for the src2tgt case.

We can see that annotation projection methods that incorporate candidate matching can achieve results comparable to or better than previous approaches. Specifically, for the German language, the newly proposed method exhibits a significant performance improvement.

Projection method	de	es	it
Heuristic SimAlign	80.0	90.7	87.0
Heuristic AWESoME	81.9	90.3	87.3
n-gram SimAlign	89.8	89.2	87.8
n-gram AWESoME	92.0	88.6	87.2
Model transfer	67.5	74.1	69.6
NER SimAlign	74.5	79.8	72.3
NER AWESoME	74.7	80.0	72.0

Table 2: F1 scores resulting from the evaluation of only the projection step using the Europarl-based NER dataset with English as a source language.

The NER-based target candidates approach performs in this experiment worse due to imperfect spans predicted by the model. However, it still outperforms plain model transfer because it corrects wrongly predicted labels for spans using the projection from matched source entities.

In the case of the Spanish language, the heuristic word-to-word alignment-based algorithm slightly outperforms the proposed approach utilizing the n-gram candidate extraction strategy. This advantage arises from the algorithm’s ability to merge two continuous ranges of target words aligned with source entity words, when only one misaligned word exists between these ranges. In contrast, our approach exhibits this capability only in specific situations.

⁵ShkalikovOleh/europarl-ner

Notes on Complexity of the Problem

It can be demonstrated that the proposed problem, when excluding the second set of constraints that limit the number of projections for source entities, reduces to the maximum weight independent set problem on interval graphs, which is solvable in polynomial time (Pal and Bhattacharjee, 1996). Therefore, any potential complexity in the entire problem may be due to the combination of non-overlapping constraints and the constraints limiting the number of projections for each source entity. Although it is likely that the proposed ILP formulation could be solved in polynomial time, we cannot make a definitive claim since an appropriate algorithm has yet to be identified.

Insights from the MasakhaNER2 Dataset Experiments

Here, we provide further details on the results across different languages from the MasakhaNER2 dataset.

The set of 10 languages where the proposed method performs better than heuristics from the MasakhaNER2 dataset includes: Bambara ('bam'), Fon ('fon'), Hausa ('hau'), Igbo ('ibo'), Luganda ('lug'), Mossi ('mos'), Shona ('sna'), Swahili ('swa'), Wolof ('wol'), and Yoruba ('yor'). The second set of 8 languages where the proposed methods perform worst includes: Ewe ('ewe'), Kinyarwanda ('kin'), Luo ('luo'), Chichewa ('nya'), Tswana ('tsn'), Twi ('twi'), Xhosa ('xho'), and Zulu ('zul'). The exact metric values can be found in the provided GitHub repo.