
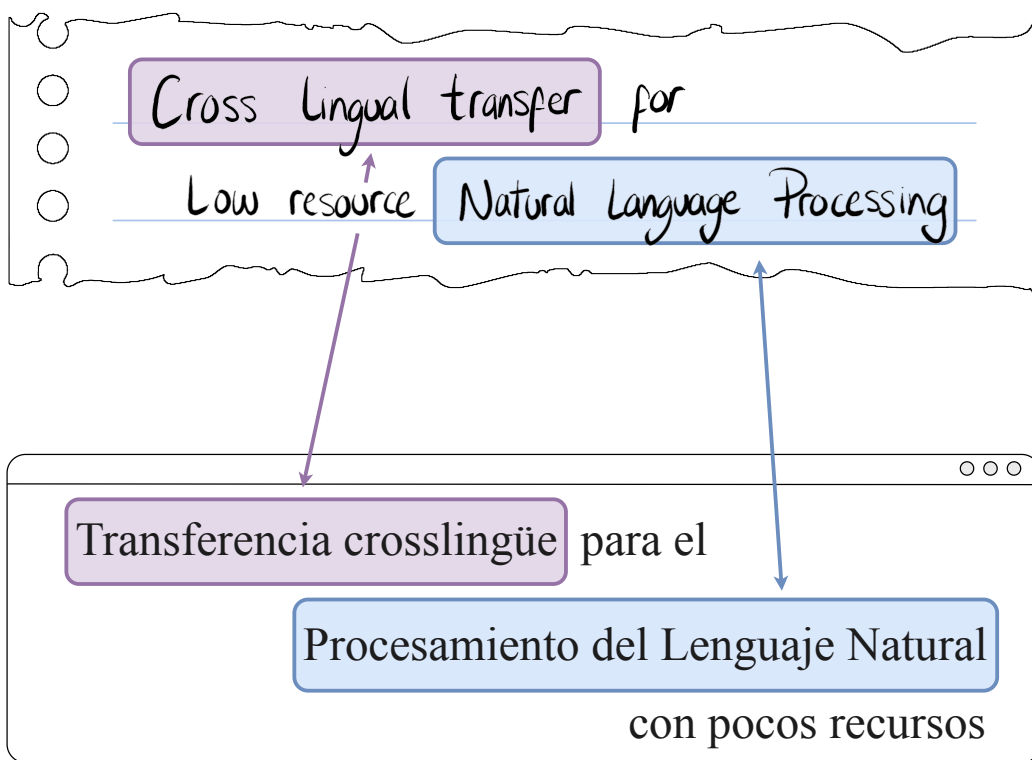


eman ta zabal zazu



EUSKAL HERRIKO UNIBERTSITATEA

Doctoral Programme in Language Analysis and Processing




Doctoral Thesis
Iker García Ferrero

Supervisors:
German Rigau and Rodrigo Agerri

2024

eman ta zabal zazu



EUSKAL HERRIKO UNIBERTSITATEA

Doctoral Programme in Language Analysis and Processing

Cross-lingual Transfer for Low-Resource Natural Language Processing

This thesis report was made by Iker García Ferrero under the supervision of German Rigau and Rodrigo Agerri, and submitted to obtain a PhD degree at the University of the Basque Country UPV/EHU

Donostia, December 2024.

...

“You must never think of the whole street at once, understand? You must only concentrate on the next step, the next breath, the next stroke of the broom, and the next, and the next. Nothing else.”

Again he paused for thought before adding, “That way you enjoy your work, which is important, because then you make a good job of it. And that’s how it ought to be.”

...

Michael Ende (MOMO, 1973)

Acknowledgments

Thank you / Gracias / Eskerrik Asko ...

... A German Rigau, por haberme guiado desde que era un alumno de grado que no sabía a qué quería dedicarse. Gracias por enseñarme la pasión por la investigación.

... A Rodrigo Agerri, por ayudarme a poner orden en los brainstormings y tablas de resultados infinitas, y por enseñarme cómo ser un buen investigador.

... A mi familia, por haberme dado la oportunidad de poder estudiar y trabajar en algo que me apasiona y me hace disfrutar. Y por haberme apoyado en todo el camino.

... A Raquel Barbero, por descubrirme el mundo fuera de las teclas, por ayudarme a desconectar cuando no sabía desconectar, y por apoyarme siempre.

... IXA taldeari. Lanera joatea inoiz ez delako lanera joatea bezala sentitzen. Egun bakoitza dibertigarria egiteagatik. Eta eman didazuen laguntza guztiagatik.

... 318 bulegori, brainstorming ordu guztiengatik eta ideia on guztiengatik. Eskerrik asko Ander Salaberriari bihurrikeria guztietan konplize izateagatik eta Oscar Sainzi elkarrekin egin dugun lan guztiengatik.

... to Dan Roth for hosting me at the University of Pennsylvania, Jennifer Sheffield for all the work that made it possible, and all the members of the Cognitive Computation Group for giving me the chance to collaborate with you.

This thesis has been supported by a PhD Grant from the Basque Government (PRE_2020_2_0208).

Abstract

Natural Language Processing (NLP) has seen remarkable advances in recent years, particularly with the emergence of Large Language Models that have achieved unprecedented performance across many tasks. However, these developments have mainly benefited a small number of high-resource languages such as English. The majority of languages still face significant challenges due to the scarcity of training data and computational resources. To address this issue, this thesis focuses on cross-lingual transfer learning, a research area aimed at leveraging data and models from high-resource languages to improve NLP performance for low-resource languages. Specifically, we focus on Sequence Labeling tasks such as Named Entity Recognition, Opinion Target Extraction, and Argument Mining.

The research is structured around three main objectives: (1) advancing data-based cross-lingual transfer learning methods through improved translation and annotation projection techniques, (2) developing enhanced model-based transfer learning approaches utilizing state-of-the-art multilingual models, and (3) applying these methods to real-world problems while creating open-source resources that facilitate future research in low-resource NLP.

More specifically, this thesis presents a new method to improve data-based transfer with T-Projection, a state-of-the-art annotation projection method that leverages text-to-text multilingual models and machine translation systems. T-Projection significantly outperforms previous annotation projection methods by a wide margin. For model-based transfer, we introduce a constrained decoding algorithm that enhances cross-lingual Sequence Labeling in zero-shot settings using text-to-text models. Finally, we develop Medical mT5, the first multilingual text-to-text medical model, demonstrating the practical impact of our research on real-world applications.

Resumen

El Procesamiento del Lenguaje Natural (PLN) ha experimentado avances notables en los últimos años, particularmente con la aparición de Modelos de Lenguaje de Gran Tamaño que han logrado un rendimiento sin precedentes en numerosas tareas. Sin embargo, estos desarrollos han beneficiado principalmente a un pequeño número de idiomas con abundantes recursos, como el inglés. Así, la mayoría de los idiomas aún se enfrentan a desafíos significativos debido a la escasez de datos de entrenamiento y recursos computacionales. Para abordar este problema, esta tesis se centra en el aprendizaje por transferencia crosslingüe, un área de investigación destinada a aprovechar los datos y modelos de idiomas con abundantes recursos para mejorar el rendimiento del PLN en idiomas con recursos más limitados. Específicamente, en esta tesis se enfoca en tareas de Etiquetado Secuencial como el Reconocimiento de Entidades Nombradas, la Extracción de Foco de Opinión y la Minería de Argumentos.

La investigación se estructura en torno a tres objetivos principales: (1) avanzar en los métodos de aprendizaje por transferencia crosslingüe basados en datos mediante técnicas mejoradas de traducción y proyección de anotaciones, (2) desarrollar enfoques mejorados de aprendizaje por transferencia basados en modelos multilingües de última generación, y (3) aplicar estos métodos a problemas del mundo real mediante la creación de recursos de código abierto que faciliten la investigación futura en PLN con recursos limitados.

Más concretamente, en esta tesis se presenta un nuevo método para mejorar la transferencia basada en datos con T-Projection, una técnica de proyección de anotaciones de última generación que aprovecha los modelos multilingües texto-a-texto y los sistemas de traducción automática. T-Projection supera significativamente todos los métodos anteriores de proyección de anotaciones. Para la transferencia basada en modelos, introducimos un algoritmo de decodificación restringida que mejora el Etiquetado Secuencial crosslingüe en entornos sin recur-

sos utilizando modelos texto-a-texto. Finalmente, desarrollamos Medical mT5, el primer modelo médico multilingüe texto-a-texto, demostrando el impacto práctico de nuestra investigación en aplicaciones del mundo real.

Laburpena

Hizkuntzaren Prozesamenduan aurrerapen nabarmenak ikusi dira azken urteetan, bereziki ataza askotan aurrekaririk gabeko errendimendua lortu duten Hizkuntza Eredu Handien agerpenarekin. Hala ere, garapen hauek batez ere baliabide handiko hizkuntza gutxi batzuen onurarako izan dira, ingelesa kasu. Hizkuntza gehienek oraindik ere erronka handiei aurre egin behar diete entrenamendu-datuen eta baliabide konputazionalen urritasuna dela eta. Arazo honi aurre egiteko, tesi honek hizkuntzen arteko transferentzia-ikasketan jartzen du arreta, hots, baliabide handiko hizkuntzetako datuak eta ereduak aprobetxatuz baliabide urriko Hizkuntzetarako Prozesamenduanaren errendimendua hobetzea helburu duen ikerketarloan. Zehazki, Sekuentzia Etiketatze atazetan zentratzen gara, hala nola Izendun Entitateen Erauzketan, Iritzien Xedeen Erauzketan eta Argudio Meatzaritzan.

Ikerketa hiru helburu nagusiren inguruan egituratzen da: (1) datuetan oinarritutako hizkuntzen arteko transferentzia-ikasketa metodoak hobetzea itzulpen eta anotazio-proiektzio tekniken bidez, (2) ereduetan oinarritutako transferentzia-ikasketa hurbilpenak garatzea puntako eredu eleaniztunak erabiliz, eta (3) metodo hauek benetako arazoei aplikatzea, baliabide urriko Hizkuntzetarako Prozesamenduan etorkizuneko ikerketa erraztuko duten kode irekiko baliabideak sortuz.

Zehazki, datuen transferentzia hobetzen dugu T-Projection bidez, testutik testurako eredu eleaniztunak eta itzulpen automatikoko sistemak erabiltzen dituen puntako anotazio-proiektzio metodoa. T-Projection metodoak nabarmen gainditzen ditu aurreko anotazio-proiektzio metodoak. Ereduetan oinarritutako transferentziarako, deskodifikazio murriztuko algoritmo bat aurkezten dugu, zero-shot testuinguruetan hizkuntzen arteko Sekuentzia Etiketatzea hobetzen duena testutik testurako ereduak erabiliz. Azkenik, Medical mT5 garatu dugu, testutik testurako lehen eredu mediko eleaniztuna, gure ikerketaren eragin praktikoa erakutsiz benetako aplikazioetan.

Table of Contents

Abstract	vii
Resumen	ix
Laburpena	xi
Table of Contents	xiii
Table List	xvii
Figure List	xix
1 Introduction	1
1.1 Motivation	2
1.2 Goals and research lines	4
1.3 Structure of the thesis	6
1.4 List of scientific contributions	7
1.4.1 Contributions included in the thesis	8
1.4.2 Closely Related Contributions	9
1.4.3 Contributions that are not part of the Thesis	10
1.5 List of open-source resources	12
1.5.1 Open source software	12
1.5.2 Open source datasets	14
1.5.3 Open source models	15
2 Related Work	17
2.1 NLP and Deep Learning: Scaling compute and data	17

TABLE OF CONTENTS

2.2	Cross-Lingual Transfer Methods	21
2.2.1	Data-based transfer	22
2.2.2	Model-based transfer	32
3	Data transfer vs Model transfer	35
3.1	Motivation and contributions	35
3.2	Methodology	37
3.2.1	Data transfer	37
3.2.2	Model transfer	41
3.3	Experimental Setup	42
3.3.1	Datasets	42
3.3.2	Machine Translation	43
3.3.3	Word Alignments	43
3.3.4	Sequence labeling Models	44
3.4	Experimental Results	46
3.4.1	Opinion Target Extraction	46
3.4.2	Named Entity Recognition	48
3.4.3	Discussion	49
3.5	Error Analysis	51
3.5.1	Downstream evaluation of Machine Translation Models	51
3.5.2	Evaluating the Projection Method	52
3.5.3	Categorization of mistakes	54
3.6	Conclusions	58
4	Improving Data Transfer	59
4.1	Motivation and contributions	59
4.2	T-Projection	61
4.2.1	Candidate Generation	62
4.3	Candidate Selection	63
4.4	Experimental Setup	65
4.4.1	Datasets	65
4.4.2	Baselines	67
4.4.3	Models Setup	69
4.5	Intrinsic Evaluation	70
4.5.1	Annotation Projection Quality	70
4.5.2	The Role of the Candidates	72
4.5.3	How many candidates are necessary?	73
4.5.4	Model size and performance	75

TABLE OF CONTENTS

4.6	Extrinsic Evaluation	76
4.6.1	T-Projection vs other annotation projection systems	77
4.6.2	T-Projection vs Model-transfer	78
4.7	Conclusions	78
5	Improving Model Transfer	81
5.1	Motivation and contributions	81
5.2	Related Work	84
5.2.1	LLMs for sequence labeling	84
5.2.2	Constrained decoding	85
5.3	Approach	86
5.3.1	Input-Output Representation	86
5.3.2	Constrained decoding	87
5.4	Experimental Setup	89
5.4.1	Language Models and baselines	90
5.4.2	Training Setup	90
5.4.3	Evaluation Metrics	92
5.5	Experiments	92
5.5.1	Named Entity Recognition	92
5.5.2	Opinion Target Extraction	94
5.5.3	Event Extraction	95
5.5.4	Model Transfer vs Data Transfer	95
5.6	Ablation Study	97
5.7	Conclusion	101
6	Medical MT5: Cross-Lingual Transfer for Domain-Specific Task	103
6.1	Motivation and Contributions	103
6.2	Related Work	105
6.3	Compiling a Multilingual Corpus for the Medical Domain	106
6.3.1	English	106
6.3.2	Spanish	107
6.3.3	French	107
6.3.4	Italian	107
6.4	Medical mT5	108
6.4.1	Pre-training Medical mT5	108
6.5	Generating New Multilingual Benchmarks	109
6.5.1	Argument Mining	110
6.5.2	Question Answering	111

TABLE OF CONTENTS

6.6	Experimental Setup	111
6.6.1	Datasets	111
6.6.2	Conversion to Text-to-Text Format	112
6.6.3	Baselines	113
6.6.4	Hyperparameters settings	114
6.7	Experimental Results	115
6.7.1	Sequence labeling Tasks	115
6.7.2	Abstractive Question Answering	118
6.8	Conclusion	120
7	Conclusion and future work	123
7.1	Future work	125
	Bibliography	129
	Appendix	163
A.1	Original papers	163
A.2	García-Ferrero et al. (Findings of the Association for Computational Linguistics: EMNLP 2022)	165
A.3	García-Ferrero et al. (Findings of the Association for Computational Linguistics: EMNLP 2023)	179
A.4	García-Ferrero et al. (LREC-COLING 2024)	195

Table List

3.1	Number of sentences for each dataset split.	43
3.2	OTE F1 scores with models of different capacities in the SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016) dataset.	47
3.3	NER F1 scores with models of different capacities in the CoNLL-2002 (Sang 2002) and CoNLL-2003 (Sang and Meulder 2003) datasets.	48
3.4	Comparison between the previous research methods that leverage projections, the zero-shot baselines and our annotation projections in the CoNLL-2002 (Sang 2002) and CoNLL-2003 (Sang and Meulder 2003) datasets. F1 score reported	49
3.5	OTE F1 score in the SemEval 2016 ABSA Pontiki <i>et al.</i> 2016 dataset of different XLM-R large models trained using data generated with different translation systems.	52
3.6	OTE F1 score in the SemEval 2016 ABSA Pontiki <i>et al.</i> 2016 dataset between the human annotation projections vs the automatic projections generated using different alignment models.	53
3.7	XLM-R large OTE F1 score in the SemEval 2016 ABSA Pontiki <i>et al.</i> 2016 dataset when training with automatically and manually projected datasets	53
3.8	Number of times words appear as target words in the SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016) train dataset.	54
3.9	Most common false negatives and positives were there is a big mismatch between methods and the total number of labelled appearances of the word in the test data. B is the acronym for mBERT, Xb for XLM-R base and Xl for XLM-R large.	55

TABLE LIST

4.1	Size (Number of sentences) of the dataset we use to train and evaluate our systems.	67
4.2	F1 scores for annotation projection in the OTE, NER and Argument Mining tasks.	71
4.3	F1 scores for different candidate generation and candidate selection methods.	72
4.4	F1 scores of T-Projection when using translation and mT5 models of different size	75
4.5	F1 scores on MasakhaNER2.0 for mDebertaV3 trained with projected annotations from different systems. "+EN" denotes concatenation of the automatically generated target language dataset with the source English dataset.	76
5.1	Size and training data of some relevant open source models. . . .	82
5.2	Comparison of F1-score of various LLMs with that of the current state of the art result in Masakhaner 2.0. Table reproduced from Ojo and Ogueji 2023.	82
5.3	Hyperparameters used for fine-tuning the models.	91
5.4	F1 scores in the Named Entity Recognition Task. Model are trained in English and evaluated in a set of African languages.	92
5.5	Average F1 scores in the MasakhaNER dataset.	93
5.6	F1 scores in the Opinion Target Extraction Task.	94
5.7	F1 scores in the Event Extraction Task.	95
5.8	F1 Scores for Named Entity Recognition Task. "Zero" refers to the model trained in English and evaluated on a set of African languages. "Data" refers to the model trained on automatically translated and projected data using T-Projection for each language.	96
6.1	Data sources and word counts by language.	106
6.2	Pre-Training settings for Medical mT5.	109
6.3	List of evaluation tasks used to measure the performance of Medical mT5.	112
6.4	Single-task supervised F1 scores for Sequence Labelling.	115
6.5	Multi-task supervised F1 scores for Sequence Labelling.	117
6.6	Zero-shot F1 scores for Argument Mining. Models have been trained in English and evaluated in Spanish, French and Italian. . .	118
6.7	Examples of answers generated by each model for two different BioASQ questions together with the rank assigned by medics. . .	119

Figure List

1.1	Modern LLMs, which support text, image, and other multimodal representations, have achieved outstanding performance in a wide range of NLP tasks. They have been applied in many real-world applications.	2
1.2	Illustration of the Named Entity Recognition (NER) sequence labelling task. The goal is to identify and classify named entities in running text.	4
2.1	Illustration of multilingual embeddings, where two languages are mapped into a shared vector space. Words with similar meanings are placed close together.	18
2.2	Representation of the BERT architecture. During training, BERT learns to predict missing words in a sentence based on the contextual representations produced by the model.	19
2.3	Representation of the text-to-text framework in T5. Every task is framed as a text input and the model is trained to generate the desired output as text. Figure reproduced from Raffel <i>et al.</i> 2020. .	20
2.4	Illustration of the Translate-Train cross-lingual transfer approach: Given gold data in the source language, this method utilizes translation and annotation projection to create silver-standard training data in the target language.	22
2.5	Illustration of the Translate-Test cross-lingual transfer approach: A model is trained using gold data in the source language. During inference, inputs in the target language are first translated into the source language, after which predictions are made and then projected back into the target language.	23

FIGURE LIST

2.6	Illustration of data transfer for different NLP tasks. Each task requires a different method to transfer the labels from the source into the target language.	24
2.7	Illustration of word alignments represented as a bidirectional graph.	25
2.8	Illustration of word alignments by fine-tuning language models. The query is on the left “West Germany” and the translated sentence on the right. The model predicts that “Alemania Occidental” is the Spanish translation of “West Germany”.	26
2.9	Illustration of the cosine similarity between token embedding representations using Multilingual BERT.	27
2.10	Illustration of annotation projection using word-alignments	28
2.11	Illustration of annotation projection using Machine Translation. Individually labeled sequences are translated separately from the rest of the sentence. The translations of these sequences are then matched with the translations produced by translating the entire sentence.	29
2.12	Illustration of annotation projection using Machine Translation and placeholders. Labeled sequences are replaced by a placeholder. The sentence with placeholders and the labeled sequences are translated independently. After translation, the placeholders are replaced with the corresponding labeled sequence translation. .	30
2.13	Illustration of the mark-then-translate approach. Markers are introduced around the labeled sequences. The sentence and the labeled spans are translated together.	30
2.14	Illustration of bilingual dictionary generation. Monolingual embeddings are projected into a shared space in which a bilingual dictionary is computed by k-nearest-neighbor. Figure reproduced from Xie <i>et al.</i> 2018.	31
2.15	Illustration of the model-based cross-lingual transfer approach. A pre-trained multilingual model is finetuned with data in the source language and then applied without modification to label text in the target language.	33
3.1	Illustration of the two data transfer approaches we have implemented. They are differentiated by the direction in which we translate the data. In both cases, English is the source language and Spanish is the target language.	38

3.2	Illustration of the translation and annotation projection method for the Named Entity Recognition task.	39
3.3	Illustration of the split annotation and annotation collision errors when projecting a sentence.	40
3.4	Illustration of model transfer approach. A multilingual model is trained in the source language (English). The model is then used to label sentences in the target language (Spanish).	41
3.5	Illustration of the sequence labeling tasks used in the experiments in this chapter.	42
3.6	Implementation of the sequence labeling model. An encoder-based Language Model is fed the input sequence. The output representations are used by a token classification linear layer to predict the labels.	45
3.7	Amount of data in GiB (log-scale) for the languages we use in our experiments in Wiki-100 corpus used for training mBERT and the CC-100 used for XLM-R. The full figure can be found in Conneau <i>et al.</i> 2020	50
4.1	T-Projection two-step method to project sequence labeling annotations across languages.	61
4.2	Illustration of the candidate generation step. For each label, we generate a set of probable candidates.	62
4.3	Candidate selection: candidates are scored based on the probability of being generated as a translation of the source labeled sequences.	64
4.4	Sequence labeling tasks in our experiments	66
4.5	Illustration of the translation and annotation projection task using word-alignments.	67
4.6	Illustration of the translation with markers approach. Markers are introduced around the labeled sequences. The sentence and the labeled spans are translated together.	68
4.7	Illustration of the span translation annotation projection approach. The source labels are translated independently, and these translated spans are then matched with their counterparts in the target sentence.	68
4.8	F1 score when generating a different number of candidates.	73
4.9	Number of times the correct candidate is among the top-k candidates generated by mT5.	74

FIGURE LIST

5.1	Comparison between a valid (top green) and invalid (bottom red) output structure to represent a Named Entity Recognition task. English translation: (They) played in Real and in the Turkish national team.	83
5.2	Text-to-Text representation of the Sequence Labeling task. Given an input sentence, the model must generate the same sentence annotated with html-style tags.	86
5.3	Our Constrained Decoding Algorithm is defined as a Finite State Automaton.	87
5.4	Information Extraction Tasks in our experiments	89
5.5	Percentage of hallucinated words compared to the performance delta between unconstrained and unconstrained beam search in MasakhaNER using mT0-XL.	97
5.6	Average percentage of mistakes generated by Unconstrained Beam search in MasakhaNER using mT0 models of different sizes	99
5.7	Average F1 score in MasakhaNER compared to the mT0 model size 100	
5.8	Average F1 score of mT0-XL in a subset of MasakhaNER compared to the number of beams used for decoding.	101
6.1	Example of an annotated abstract from the AbstRCT dataset. . . .	110
6.2	Data construction process for generating the Spanish, French and Italian versions of the AbstRCT dataset.	110
6.3	Text-to-Text representation of the Sequence Labeling task. Given an input sentence, the model is expected to generate the same sentence annotated with html-style tags.	112
6.4	Text-to-Text representation of the BioASQ task. Given a question and a set of relevant snippets, the model generates an answer. . . .	113

1. CHAPTER

Introduction

This thesis is framed within the area of Natural Language Processing (NLP). Natural Language Processing is a multidisciplinary research field within Artificial Intelligence (AI), Computer Science, and Linguistics. NLP involves a wide range of tasks, including, Natural Language Understanding, Machine Translation, Information Extraction, and Text Generation, among others. The main goal of NLP is to enable computers to understand, interpret, and generate human language in a way that is valuable for humans. The Ixa group, within the HiTZ center at the University of the Basque Country, is one of the leading research teams working in NLP. Since its foundation more than 30 years ago, the Ixa group has been a pioneer in developing NLP tools for many different applications, with a special focus on creating language tools for the Basque language. Moreover, Ixa has been involved in many European and international research projects, significantly contributing to languages beyond Basque.

The primary objective of this thesis is to develop cross-lingual transfer learning solutions to address the resource constraints faced by many languages, tasks, and domains. Cross-lingual transfer learning is a research area focused on creating models for low-resource languages by leveraging knowledge from high-resource languages. Specifically, this thesis explores cross-lingual transfer learning for Sequence Labeling tasks, such as Named Entity Recognition, Opinion Target Extraction, and Argument Mining. We propose novel methods for knowledge transfer from high-resource to low-resource languages through translation and annotation projection, as well as multilingual NLP models. Thus, our goal is to develop publicly available models that achieve state-of-the-art performance in low-resource

languages and to make these models accessible to the research community. This thesis work was aligned with the objectives of the projects DeepReading¹, Deep-Knowledge² and Andidote³.

1.1 Motivation

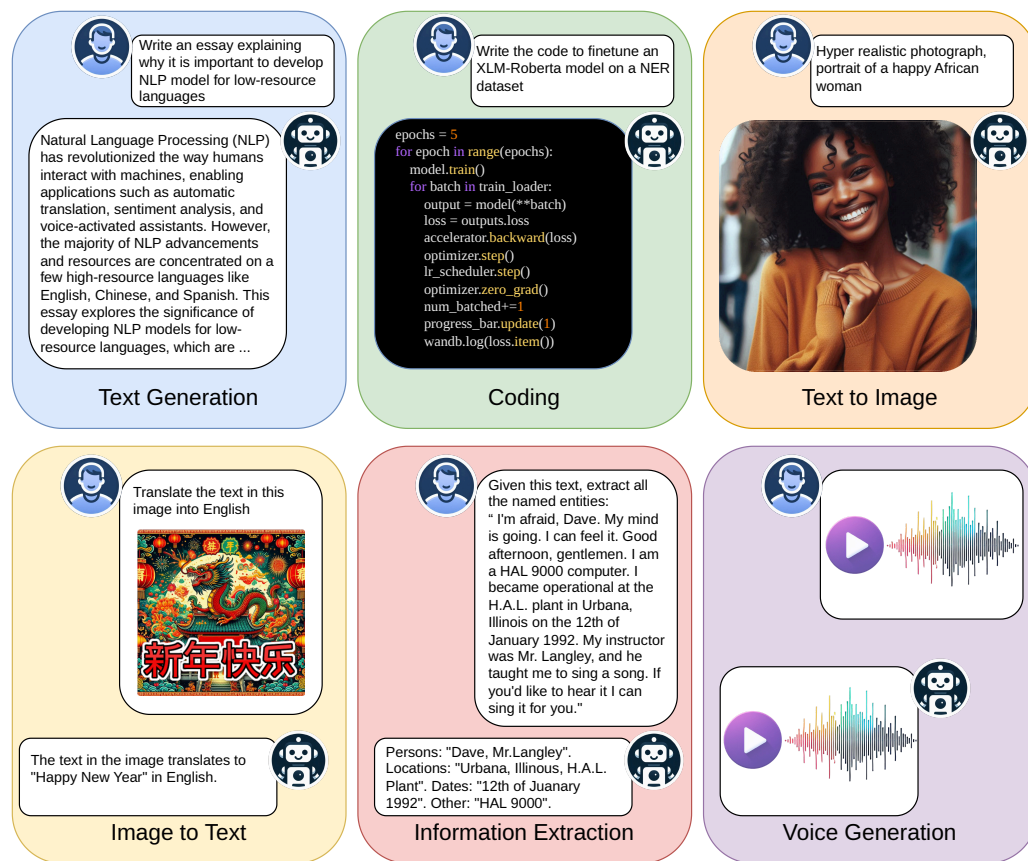


Figure 1.1 – Modern LLMs, which support text, image, and other multimodal representations, have achieved outstanding performance in a wide range of NLP tasks. They have been applied in many real-world applications.

¹<https://ixa2.si.ehu.es/deepreading/>

²<http://ixa.si.ehu.es/node/13582>

³<https://univ-cotedazur.eu/antidote>

Neural networks have become an indispensable resource in Natural Language Processing (NLP). Driven by the success of the Transformer architecture (Vaswani *et al.* 2017), they have demonstrated outstanding performance in various challenging NLP tasks (Min *et al.* 2024), such as General Language Understanding (Wang *et al.* 2019), Question Answering (Rajpurkar *et al.* 2018), Text Generation (Brown *et al.* 2020), Dialogue (Thoppilan *et al.* 2022), and Conditional Image Generation (Rombach *et al.* 2022), among others. Scaling up these models in terms of parameter count and training data (Chung *et al.* 2022) has led to the development of current state-of-the-art NLP systems. Large Language Models (LLMs) such as GPT-4 (OpenAI *et al.* 2024) and LLaMA-3 (AI@Meta 2024), trained on hundreds of terabytes of text data and billions of parameters, have proven capable of generating human-like text and have been applied in a wide range of applications, such as the ones depicted in Figure 1.1. These cutting-edge NLP systems hold the potential to bring significant societal changes (Bommasani *et al.* 2021).

Despite the remarkable progress in NLP, many challenges remain. LLMs require vast amounts of data and computational resources to achieve optimal performance (Hoffmann *et al.* 2022). In addition to English, only a handful of Western European languages (principally German, French, and Spanish) and even fewer non-Indo-European languages (primarily Chinese, Japanese, and Arabic) dominate the field (Joshi *et al.* 2020). While speakers of these languages benefit from the latest innovations in Language Technology—such as quick and accurate access to information using smart assistants, online translation services, interaction with machines using natural language, or speeding-up their work with automatic summarization tools, coding assistants, or image generation tools—speakers of low-resource languages are being left behind (Blasi *et al.* 2022).

Models consistently perform better on high-resource languages, especially English (Etzaniz *et al.* 2024b), while their performance on low-resource languages is significantly lower (Ojo *et al.* 2023; Ojo and Ogueji 2023). This disparity is due to the fact that the quality and quantity of the data directly impact the performance of the models (Liu *et al.* 2021). For the large majority of the approximately more than 7,000 languages spoken worldwide, this data is scarce or non-existent (Joshi *et al.* 2020). Therefore, obtaining optimal results would require manually generating annotated data for each application domain and language. Given the rapidly increasing number of tasks and domains to which NLP is applied, this is an unfeasible task in terms of monetary cost and human effort.

The primary objective of this thesis is to develop cross-lingual transfer learning solutions to address the resource constraints faced by many languages, tasks, and domains. *Cross-lingual transfer learning* is a research area focused on cre-




Figure 1.2 – Illustration of the Named Entity Recognition (NER) sequence labelling task. The goal is to identify and classify named entities in running text.

ating models for low-resource languages by leveraging knowledge from high-resource languages (Conneau and Lample 2019). Cross-lingual transfer learning uses the data and models available in high-resource languages (typically English) to solve tasks in low-resource languages where these resources are scarce or non-existent.

This thesis explores cross-lingual transfer learning for sequence labeling tasks. *Sequence labeling* is the task of assigning a label to each token in a given input sequence (Lafferty *et al.* 2001). Figure 1.2 illustrates the Named Entity Recognition (NER) sequence labeling task, where the goal is to identify and classify named entities in a text. Sequence labeling tasks are essential for many NLP applications, such as Information Extraction, Question Answering, and Sentiment Analysis, among others. By applying cross-lingual transfer learning techniques, such as translation and annotation projection, alongside multilingual NLP models, we aim to leverage resources from high-resource languages to perform sequence labeling in low-resource languages. Our final goal is to develop publicly available models that achieve state-of-the-art performance in low-resource languages.

1.2 Goals and research lines

The main goal of this thesis is to develop state-of-the-art cross-lingual transfer learning methods for sequence labeling tasks. We aim to apply these methods to real-world problems where the lack of resources is a significant issue. Additionally, we intend to provide the research community with a set of tools, as well as generate freely available data and models that can be used in the future. The research lines of this thesis are as follows:

- **RL1: Develop better data-based cross-lingual transfer learning methods for sequence labeling tasks.** Data-transfer methods focus on transferring knowledge from high-resource to low-resource languages through translation and annotation projection. At the start of this thesis, most data-

based approaches relied on statistical word alignment methods and sub-optimal Machine Translation models. Our goal was to develop improved data-based methods that leverage the latest advances in Machine Translation and NLP models. We also aim to explore the use of multilingual NLP models for data transfer, which have shown promising results in other NLP tasks.

- **RL2: Develop better model-based cross-lingual transfer learning methods for sequence labeling tasks.** Model-transfer methods are based on transferring knowledge from high-resource to low-resource languages through pre-trained models. A multilingual NLP model is fine-tuned on data from high-resource languages and then directly applied to low-resource languages. At the start of this thesis, this approach was offering good results in many NLP tasks using encoder-only models. Our objective is to develop improved model-based methods by leveraging the multilingual capabilities of state-of-the-art text-to-text pre-trained models.
- **RL3: Real-world application of cross-lingual transfer learning methods.** We aim to apply the developed methods to real-world problems where the lack of resources is a significant issue. By doing so, we aim to better understand which scenarios are best suited for different techniques in cross-lingual transfer learning. Additionally, we develop open-source tools, datasets, and models to support the research community in replicating our experiments and extending our work. These resources are intended to facilitate advancements in NLP for low-resource languages and enable their application across diverse tasks, languages, and domains.

1.3 Structure of the thesis

This thesis is structured as a series of interconnected papers, each building on the previous one. The chapters are organized as follows:

In Chapter 2, we present the background of the thesis, review the state-of-the-art in cross-lingual transfer learning for sequence labeling tasks, and introduce the main concepts and techniques used in this research.

Chapter 3 focuses on the effectiveness of model-based and data-based cross-lingual transfer learning methods for sequence labeling tasks. We identify the advantages and shortcomings of each method, as well as the challenges faced by current techniques for cross-lingual zero-resource sequence labeling. These insights provide a foundation for the subsequent chapters.

In Chapter 4, we introduce a novel data-based method for cross-lingual transfer learning in zero-resource settings. We propose T-Projection, a method that achieves state-of-the-art performance on annotation projection tasks.

Chapter 5 presents a constrained decoding algorithm that improves the performance of the model-based cross-lingual transfer learning approach. We demonstrate that the constrained decoding algorithm successfully leverages text-to-text models for sequence labeling tasks in low-resource languages achieving state-of-the-art results.

Chapter 6 offers a case study on the application of cross-lingual transfer learning to the medical domain. We show that the methods developed in this thesis can be successfully applied to real-world problems where resource scarcity is a significant issue. By applying both data-based and model-based methods, we develop a comprehensive multilingual pre-training, fine-tuning, and evaluation framework for the medical domain, culminating in the first open-source text-to-text multilingual model for the medical domain.

Finally, Chapter 7 summarizes the conclusions of the thesis, discusses the main contributions and limitations of the work, and proposes future research directions.

1.4 List of scientific contributions

In this section, we present the scientific contributions developed throughout this thesis. This section is divided into three parts. First, we present the publications that are included in this manuscript. Next, we provide a list of publications closely related to the thesis topic but not included in this manuscript. Finally, we list publications from other lines of research that are outside the scope of this thesis. All papers are listed in chronological order.

1.4.1 Contributions included in the thesis

These three publications are included in this thesis manuscript, as they present the main contributions of the thesis. Their content will be explained in the following chapters.

Model and Data Transfer for Cross-Lingual Sequence Labelling in Zero-Resource Settings.

Presented in Chapter 3.

Iker García-Ferrero, Rodrigo Agerri, and German Rigau.
Findings of the Association for Computational Linguistics: EMNLP 2022.
<https://doi.org/10.18653/v1/2022.findings-emnlp.478>

T-projection: High quality annotation projection for sequence labeling tasks

Presented in Chapter 4.

Iker García-Ferrero, Rodrigo Agerri, and German Rigau.
Findings of the Association for Computational Linguistics: EMNLP 2023
<https://doi.org/10.18653/v1/2023.findings-emnlp.1015>

Medical mT5: An Open-Source Multilingual Text-to-Text LLM for The Medical Domain

Presented in Chapter 6.

Iker García-Ferrero, Rodrigo Agerri, Aitziber Atutxa Salazar, Elena Cabrio, Iker de la Iglesia, Alberto Lavelli, Bernardo Magnini, Benjamin Molinet, Johana Ramirez-Romero, German Rigau, Jose Maria Villa-Gonzalez, Serena Villata, Andrea Zaninello.
LREC-COLING 2024
<https://aclanthology.org/2024.lrec-main.974>

1.4.2 Closely Related Contributions

These contributions are not included in this manuscript, as they are not directly related to the cross-lingual transfer paradigm. However, they explore complementary topics aligned with the thesis's main research direction, sharing the objective of advancing Information Extraction systems.

Benchmarking meta-embeddings: What works and what does not

Iker García-Ferrero, Rodrigo Agerri, and German Rigau.
Findings of the Association for Computational Linguistics: EMNLP 2021
<https://doi.org/10.18653/v1/2021.findings-emnlp.333>

IXA/Cogcomp at SemEval-2023 Task 2: Context-enriched Multilingual Named Entity Recognition using Knowledge Bases

Iker García-Ferrero, Jon Ander Campos, Oscar Sainz, Ander Salaberria, and Dan Roth.
Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)
<https://doi.org/10.18653/v1/2023.semeval-1.186>

GoLLIE: Annotation Guidelines improve Zero-Shot Information-Extraction

Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez de Lacalle, German Rigau, Eneko Agirre
The Twelfth International Conference on Learning Representations, 2024
<https://openreview.net/forum?id=Y3wpuxd7u9>

1.4.3 Contributions that are not part of the Thesis

These contributions result from collaborations with other members of the IXA group and the research community. They focus on evaluating language models, particularly regarding Bias and Data Contamination. Although not directly aligned with the main goals of this thesis, their insights have direct implications for the work conducted in the thesis.

Itzulpen Automatikoko Sistemen Analisia: Genero Alborapenaren Kasua.

Ander Salaberria, Jon Ander Campos, Iker García-Ferrero, Joseba Fernandez de Landa

In Proceedings of the IV. Ikergazte (2021). Nazioarteko ikerketa euskaraz. Kongresuko artikulu bilduma. Ingeniaritza eta Arkitektura.

<http://ixa.si.ehu.es/node/13328>

Twitterreko Euskal Komunitatearen Eduki Azterketa Pandemia Garaian.

Joseba Fernandez de Landa, Iker García-Ferrero, Ander Salaberria, Jon Ander Campos

In Proceedings of the IV. Ikergazte (2021). Nazioarteko ikerketa euskaraz. Kongresuko artikulu bilduma. Ingeniaritza eta Arkitektura.

<http://ixa.si.ehu.es/node/13327>

This is not a Dataset: A Large Negation Benchmark to Challenge Large Language Models

Iker García-Ferrero, Begoña Altuna, Javier Álvez, Itziar Gonzalez-Dios, German Rigau.

Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing

<https://doi.org/10.18653/v1/2023.emnlp-main.531>

NLP Evaluation in trouble: On the Need to Measure LLM Data Contamination for each Benchmark

Oscar Sainz, Jon Ander Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, Eneko Agirre
Findings of the Association for Computational Linguistics: EMNLP 2023
<https://doi.org/10.18653/v1/2023.findings-emnlp.722>

Uncovering Social Changes of the Basque Speaking Twitter Community During COVID-19 Pandemic

Joseba Fernandez de Landa, Iker García-Ferrero, Ander Salaberria, Jon Ander Campos
Proceedings of the 3rd Annual Meeting of the Special Interest Group on Under-resourced Languages @ LREC-COLING 2024
<https://aclanthology.org/2024.sigul-1.44>

Noticia: A Clickbait Article Summarization Dataset in Spanish

Iker García-Ferrero, Begoña Altuna
Journal Procesamiento del Lenguaje Natural, 2024
<http://journal.sepln.org/sepln/ojs/ojs/index.php/pln/article/view/6610>

Data Contamination Report from the 2024 CONDA Shared Task

Oscar Sainz, Iker García-Ferrero, Alon Jacovi, Jon Ander Campos, Yanai Elazar, Eneko Agirre, Yoav Goldberg, Wei-Lin Chen, Jenny Chim, Leshem Choshen, Luca D'Amico-Wong, Melissa Dell, Run-Ze Fan, Shahriar Golchin, Yucheng Li, Pengfei Liu, Bhavish Pahwa, Ameya Prabhu, Suryansh Sharma, Emily Silcock, Kateryna Solonko, David Stap, Mihai Surdeanu, Yu-Min Tseng, Vishaal Udandaraao, Zengzhi Wang, Ruijie Xu, Jinglin Yang
Proceedings of The 1st Workshop on Data Contamination (CONDA) @ ACL 2024
<https://aclanthology.org/2024.conda-1.4/>

1.5 List of open-source resources

As mentioned above, this thesis emphasizes reproducibility and the development of tools and resources freely available for the research community. We have developed open-source software, datasets, and models that other researchers can use to replicate our experiments and build upon our work. The following is a list of open-source resources developed during the years of the thesis:

1.5.1 Open source software



MetaVec

A monolingual and cross-lingual meta-embedding generation and evaluation framework.

<https://github.com/ikergarcia1996/MetaVec>



Context-enriched multilingual named entity recognition using knowledge bases.

A NER frameworks that (1) identifies possible entity candidates by analyzing the input sentence structure, (2) links the candidate to an existing updated knowledge base if possible, and (3) performs the fine-grained classification using the input sentence plus the retrieved information from the KB about the entity.

<https://github.com/ikergarcia1996/Context-enriched-NER>

Easy Label Projection!

Developed in Chapter 3

Easy Label Projection is a library that allows to project labels from one dataset into another easily. You can automatically generate datasets for languages for which you do not have any labelled data using mGiza, FastAlign, SimALign or AWESOME.

<https://github.com/ikergarcia1996/Easy-Label-Projection>

EASY-TRANSLATE

Developed in Chapter 3

Easy-Translate is a script for translating large text files with a SINGLE COMMAND. Easy-Translate is designed to be as easy as possible for beginners and as seamless customizable and as possible for advanced users.

<https://github.com/ikergarcia1996/Easy-Translate>

**T-Projection**

Developed in Chapter 4

T-Projection is a method to perform high-quality Annotation Projection of Sequence Labeling datasets. The code is built on top of HuggingFace's Transformers and HuggingFace's Accelerate library.

<https://github.com/ikergarcia1996/T-Projection>

**TINAD Framework**

A LLM finetuning and LLM evaluation library for the "This Is NOT a dataset" (TINAD) dataset.

<https://github.com/hitz-zentroa/This-is-not-a-Dataset>

**GoLLIE Framework**

The framework to finetune and evaluate GoLLIE-style models. Allows to implement any IE task by defining a set of categories and guidelines. Co-developed with Oscar Sainz.

<https://github.com/hitz-zentroa/GoLLIE>

**NoticIA Framework**

A LLM finetuning and LLM evaluation library for the NoticIA dataset.

<https://github.com/ikergarcia1996/NoticIA>

1.5.2 Open source datasets



This is NOT a Dataset

A large semi-automatically generated dataset of 400,000 descriptive sentences about commonsense knowledge that can be true or false in which negation is present in about 2/3 of the corpus in different forms that we use to evaluate LLMs.

<https://huggingface.co/datasets/HiTZ/This-is-not-a-dataset>



Multilingual Medical Corpus

Developed in Chapter 6

A Multilingual-Medical-Corpus a 3 billion word multilingual corpus for training LLMs adapted to the medical domain. Multilingual-Medical-Corpus includes four languages, namely, English, Spanish, French, and Italian.

<https://huggingface.co/datasets/HiTZ/Multilingual-Medical-Corpus>



Multilingual AbstRCT

Developed in Chapter 6

We translate the AbstRCT English Argument Mining Dataset dataset to generate parallel French, Italian and Spanish versions using the NLLB200 3B parameter model and projected using word alignment tools. The projections have been manually corrected.

<https://huggingface.co/datasets/HiTZ/multilingual-abstrect>

**Multilingual BioASQ-6B**

Developed in Chapter 6

We translate the BioASQ-6B English Question Answering dataset to generate parallel French, Italian and Spanish versions using the NLLB200 3B parameter model.

<https://huggingface.co/datasets/HiTZ/Multilingual-BioASQ-6B>

**NoticIA**

A dataset consisting of 850 Spanish news articles featuring prominent click-bait headlines, each paired with high-quality, single-sentence generative summarizations written by humans.

<https://huggingface.co/datasets/Iker/NoticIA>

1.5.3 Open source models

**medical mT5**

Developed in Chapter 6

The first open-source text-to-text multilingual model for the medical domain. Medical mT5 is an encoder-decoder model developed by continuing the training of publicly available mT5 checkpoints on medical domain data for English, Spanish, French, and Italian.

<https://huggingface.co/HiTZ/Medical-mT5-xl>



GoLLIE

A Large Language Model trained to follow annotation guidelines. GoLLIE outperforms previous approaches on zero-shot Information Extraction and allows the user to perform inferences with annotation schemas defined on the fly. Unlike previous approaches, GoLLIE can follow detailed definitions and not only rely on the knowledge already encoded in the LLM.

<https://huggingface.co/HiTZ/GoLLIE-34B>



ClickbaitFighter

A model finetuned with the NoticIA Dataset. This model can generate summaries of clickbait headlines.

<https://huggingface.co/Iker/ClickbaitFighter-10B>

2. CHAPTER

Related Work

In this chapter, we will present the state-of-the-art in Natural Language Processing (NLP) and cross-lingual transfer. We will begin with a brief overview of recent advancements in NLP and multilingual language models. Next, we will discuss the different approaches for cross-lingual transfer learning. More specifically, we will present the two main paradigms for cross-lingual transfer: data-based transfer and model-based transfer. For each paradigm, we will describe the various methods proposed in previous research. These methods will serve as the baselines for the experiments presented in the following chapters.

2.1 NLP and Deep Learning: Scaling compute and data

In the last few years, the Natural Language Processing paradigm has switched from building pipelines of different processing modules (Agerri *et al.* 2014) into end-to-end neural networks trained with vast amounts of text data (Min *et al.* 2021).

The first step in this direction was the introduction of word embeddings, such as Word2Vec (Mikolov *et al.* 2013), GloVe (Pennington *et al.* 2014), and Fast-Text (Bojanowski *et al.* 2017). These embeddings are trained on large corpora of text and capture the semantic and syntactic properties of words. Word embeddings are used as input features for neural networks that perform a wide range of NLP tasks. Multilinguality was added to this paradigm with the introduction




Figure 2.1 – Illustration of multilingual embeddings, where two languages are mapped into a shared vector space. Words with similar meanings are placed close together.

of multilingual word embeddings (Ruder *et al.* 2019). As depicted in Figure 2.1, multilingual word embeddings are trained on text from multiple languages and map words from different languages into a shared vector space (Gouws *et al.* 2015; Luong *et al.* 2015). Alternatively, multilingual word embeddings can be trained on monolingual data and then projected into a shared space using a bilingual dictionary (Zhang *et al.* 2016; Artetxe *et al.* 2016; Smith *et al.* 2017). This shared space allows for the transfer of knowledge across languages, enabling the training of models in one language and applying them to another. However, the cross-lingual transfer capabilities of word embeddings-based systems are limited and often perform poorly on low-resource languages (Conneau *et al.* 2018).

The introduction of the Transformer architecture (Vaswani *et al.* 2017) marked a significant shift in the NLP field. Transformers have achieved state-of-the-art performance on a wide range of NLP tasks, such as Machine Translation (Vaswani *et al.* 2017), Text Classification (Devlin *et al.* 2019), General Language Understanding (Wang *et al.* 2019), Question Answering (Rajpurkar *et al.* 2018), Text

Generation (Brown *et al.* 2020) or Dialogue (Thoppilan *et al.* 2022) among many others (Min *et al.* 2021). The success of the Transformer architecture has led to the development of a broad range of Transformer-based language models.




Figure 2.2 – Representation of the BERT architecture. During training, BERT learns to predict missing words in a sentence based on the contextual representations produced by the model.

The first prominent Transformer-based language model was BERT (Devlin *et al.* 2019). BERT is trained on large corpora of text data and learns to predict missing words in a sentence, producing contextual embeddings that capture the meaning of words in context. This process is illustrated in Figure 2.2. Similar to word embeddings, BERT embeddings can be used as input features for a wide range of NLP tasks. BERT has been extended to support multiple languages with mBERT (Devlin *et al.* 2019), which is trained on text from over 100 languages. mBERT achieved state-of-the-art performance in multiple languages and demonstrated strong performance when trained in English and applied to other languages (Pires *et al.* 2019; Artetxe and Schwenk 2019). The success of mBERT has led to the development of other multilingual models, such as XLM-RoBERTa (Conneau

et al. 2020) and DeBERTa (He *et al.* 2021). These models have increasingly larger sizes and are trained on progressively larger corpora of text data.

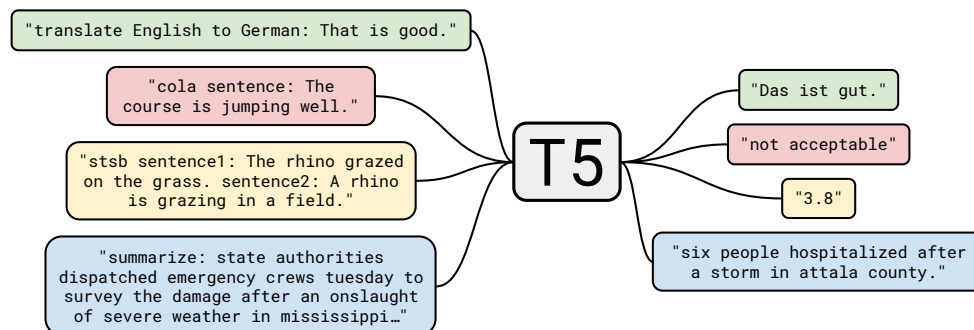


Figure 2.3 – Representation of the text-to-text framework in T5. Every task is framed as a text input and the model is trained to generate the desired output as text. Figure reproduced from Raffel *et al.* 2020.

The introduction of T5 (Raffel *et al.* 2020) and GPT (Radford *et al.* 2019; Brown *et al.* 2020) shifted the focus in NLP from learning word representations to a text-to-text approach. T5 is designed to map input text to output text, enabling it to handle a wide array of NLP tasks. Consequently, all NLP tasks are framed as text-to-text tasks, where the input is a description of the task or a prompt, and the output is the desired result, as illustrated in Figure 2.3. Unlike previous NLP models, which were fine-tuned for specific tasks, T5 can be trained on a broad spectrum of tasks with a single training objective (Chung *et al.* 2022).

On a similar research line, Radford *et al.* 2019 introduced the GPT line of models, demonstrating that Large Language Models (LLMs) trained on extensive internet data can perform, given a natural language task description, tasks such as Question Answering, Machine Translation, and Summarization without explicit supervision. This finding led to the emergence of instruction tuning, also known as multitask fine-tuning, as the leading method for achieving generalization in large models to solve diverse NLP tasks. In this approach, models are first trained on vast amounts of unlabeled data and subsequently fine-tuned on a diverse collection of tasks (Wang *et al.*, 2022; Chung *et al.*, 2022) framed as text-to-text problems.

Subsequent research has demonstrated that increasing the parameter count of language models (Brown *et al.*, 2020), coupled with improvements in the size and quality of the instruction tuning dataset, results in enhanced generalization capabilities. Consequently, models have been increasingly scaled up (Chung *et al.* 2022) in both the number of parameters and the amount of training data. This

scaling has led to the development of state-of-the-art Large Language Models (LLMs) such as GPT-4 (OpenAI *et al.* 2024), LLaMA (Dubey *et al.* 2024), and Mistral (Jiang *et al.* 2023). These models, which have billions of parameters, are trained on hundreds of terabytes of text data. They are also trained on a large number of diverse tasks and instructions, enabling them to perform a wide range of NLP tasks.

In addition to achieving state-of-the-art performance on various NLP tasks, these models can also solve tasks for which they were not explicitly trained (Radford *et al.* 2019; Lieber *et al.* 2021; Smith *et al.* 2022; Rae *et al.* 2021; Chowdhery *et al.* 2023). Since these models are trained on a substantial portion of internet data, they are inevitably multilingual and can be applied to a wide range of languages and tasks.

While the latest generation of NLP models has made a huge step forward in terms of performance, they require a huge amount of data and computational resources to train. This has led to a growing gap (Blasi *et al.* 2022) between high-resource languages, such as English, and low-resource languages, for which there is very little data available (Joshi *et al.* 2020). A notable example is African languages for which both Open Source (Ojo *et al.* 2023) and Comercial (Ojo and Ogueji 2023) NLP models produce lower performance for African languages. This has led to the development of *cross-lingual transfer* methods, which aim to leverage the knowledge learned from high-resource languages to improve or enable NLP tasks in low-resource languages. In the following sections, we will present the different approaches for cross-lingual transfer.

2.2 Cross-Lingual Transfer Methods

In this section, we will present the different approaches for cross-lingual transfer. More specifically, we will present the data-transfer and model-transfer approaches that will constitute the baselines in the following chapters.

Cross-lingual transfer in Natural Language Processing (NLP) is a method in which knowledge learned from one language (typically a high-resource language with abundant data and resources) is applied to improve or enable NLP tasks in another language (often a low-resource language with limited or no data). This process can be achieved through various techniques, including translating datasets (Ehrmann *et al.* 2011) (data-transfer) or using multilingual models (Devlin *et al.* 2019; Conneau *et al.* 2020) that understand multiple languages (model-transfer). The goal is to overcome the scarcity of annotated data in many languages, thus

facilitating multilinguality in NLP applications.

2.2.1 Data-based transfer

Data transfer leverages parallel data and/or Machine Translation to bridge the gap between languages in cross-lingual NLP tasks. Data transfer methods make the assumption of annotation preservation across translations. In the data transfer paradigm, the NLP model is trained and performs inference in the same language. There are two main approaches for data transfer: Translate-Train and Translate-Test.

Translate-Train

The Translate-Train approach, illustrated in Figure 2.4, aims to automatically generate annotated data in languages where such data is scarce by leveraging annotated datasets from a high-resource source language. This method begins with a dataset that is fully annotated in a well-resourced source language, which is then translated into the target language (Jain *et al.* 2019; Fei *et al.* 2020). Following translation, the annotations from the source language are projected onto the translated text, resulting in a new, silver-standard annotated dataset in the target language. This newly created dataset can then be used to train NLP models directly on tasks in the target language, making this approach particularly valuable when original annotated data in the target language is not available. High-quality Machine Translation systems are essential for this approach to succeed. An alternative

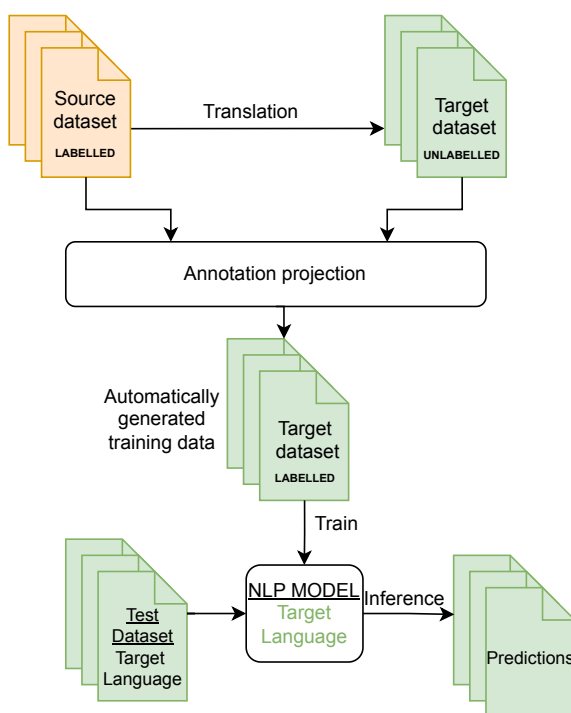


Figure 2.4 – Illustration of the Translate-Train cross-lingual transfer approach: Given gold data in the source language, this method utilizes translation and annotation projection to create silver-standard training data in the target language.

strategy involves automatically annotating the English version of a multi-parallel corpus and then projecting these annotations to all other languages in the corpus (Ehrmann *et al.* 2011). In situations where neither Machine Translation systems nor parallel data are available, Guo and Roth 2021 translate labeled data on a word-by-word basis using a dictionary and then constructing target-language text from the source-language annotations using a constrained pre-trained language model trained with unlabeled data in the target language. In any case, this approach results in a silver-standard annotated dataset in the target language. The quality of this automatically generated dataset depends on the quality of the translation or parallel data and the effectiveness of the annotation projection algorithm.

Translate-Test




Figure 2.5 – Illustration of the Translate-Test cross-lingual transfer approach: A model is trained using gold data in the source language. During inference, inputs in the target language are first translated into the source language, after which predictions are made and then projected back into the target language.

Instead of building models for the target language, the translate-test approach aims to take advantage of the ability of the models to produce better results for high-resource languages such as English (Etxaniz *et al.* 2024b). The Translate-Test approach is illustrated in Figure 2.5. In this method (Shah *et al.* 2010; Bel *et al.* 2003; Tebbifakhr *et al.* 2020), the model is trained in a source language with



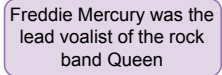
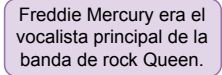
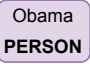

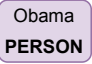

TASK	Example in source language	Translation	Label Projection Method
Text classification	Brazil won the World Cup 	Brasil ganó la Copa del Mundo 	None
Text Generation	Who is Freddie Mercury? 	¿Quién es Freddie Mercury? 	Translation
Sequence labeling	 visited 	 visitó 	Word Alignment

Figure 2.6 – Illustration of data transfer for different NLP tasks. Each task requires a different method to transfer the labels from the source into the target language.

abundant resources (gold standard data). At inference, the inputs in the target language are first translated into the source language. The model, which is trained on the source language data, performs its inference on these translated inputs. Subsequently, the predictions made by the model are translated or projected back into the target language. This approach enables the deployment of NLP models across languages without retraining the model with annotated data in the target language. However, as with the Translate-Train approach, the performance of the Translate-Test approach relies heavily on the quality of the translation of the input data and the projection of the output predictions into the target language.

Annotation Projection

In both the Translate-Train and Translate-Test approaches, it is necessary to project the labels from the source language into the target language or vice versa. As depicted in Table 2.6, the annotation projection method required varies significantly across different NLP tasks. For instance, in Text Classification, annotation projection is straightforward, as the translated sentence will be labeled with the same category as the source sentence. For text generation tasks, such as Abstractive Question Answering, the label can be translated along with the input sentence. The most challenging tasks are those involving Sequence Labeling. For sequence

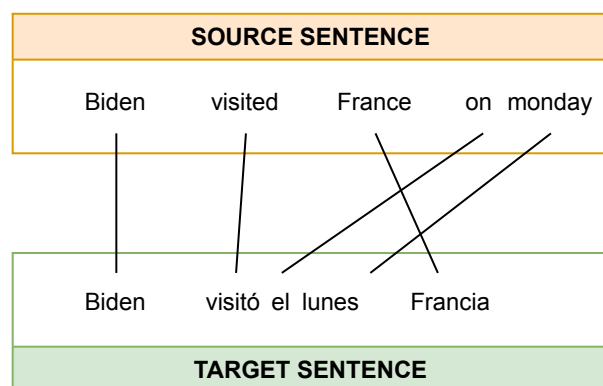


Figure 2.7 – Illustration of word alignments represented as a bidirectional graph.

labeling tasks, which involve span-level annotations, it is necessary to identify the sequence of words in the translated sentence that corresponds to the labeled spans in the source text. The majority of previous work published in this research area explores the application of word-alignments (Ehrmann *et al.* 2011).

Annotation projection using word alignments

Word alignments refer to the process of matching words in a sentence in the source language to their corresponding translations in a target language. As illustrated in Figure 2.7, word alignments are represented as a bidirectional graph between words in a parallel sentence.

Most word-alignment algorithms are based on statistical Machine Translation systems. Giza++ (Och and Ney 2003) is based on the IBM Models. The IBM alignment models (Brown *et al.* 1993) use statistical methods to learn the probability of translation between words in a source language and their counterparts in a target language, based on a given corpus of aligned texts. Building upon this research, FastAlign (Dyer *et al.* 2013a) introduces a log-linear reparameterization of IBM Model 2, which achieves an accuracy comparable to GIZA++ but with improved computational efficiency, thereby enabling faster inference throughput. On a similar line of research, Efmara and Eflomal (Östling and Tiedemann 2016) extend the IBM models with a Bayesian model with Markov Chain Monte Carlo (MCMC) inference for improved accuracy and computational efficiency.

As neural networks started to outperform previous statistical approaches for most NLP tasks (Min *et al.* 2024), a new line of research emerged aiming to generate word alignments using neural networks. This line of research employs

2 RELATED WORK




Figure 2.8 – Illustration of word alignments by fine-tuning language models. The query is on the left “West Germany” and the translated sentence on the right. The model predicts that “Alemania Occidental” is the Spanish translation of “West Germany”.

multilingual language models that have been pretrained using data from both the source and the target languages. In this line of research, Peter *et al.* 2017 built an attention-based neural network in which the attention probabilities are trained to closely align with those obtained from statistical MT toolkits. Building on this work, Zenkel *et al.* 2019 proposed a method that adds an extra layer of attention on top of the Transformer architecture and directly optimizes its activations towards a given target word. Garg *et al.* 2019 trained a Transformer model to produce both accurate translations and alignments, jointly performing the Machine Translation and Word Alignment tasks. The model outputs the translation during inference while alignments are extracted from the attention probabilities. This approach achieves competitive results compared with GIZA++ without sacrificing translation accuracy. However, the model only achieves better performance than GIZA++ when existing word alignments are provided for fine-tuning. Zenkel *et al.* 2020 combines both previous approaches and extends them with a loss function that encourages contiguity in the alignment matrix and a symmetrization algorithm that jointly optimizes the alignment matrix within two models trained in opposite directions. This approach outperforms the alignments generated by GIZA++ and other statistical Word alignment systems.

Instead of extracting alignments from the attention layer of a multilingual language model, Li *et al.* 2021 optimizes a Transformer encoder directly to generate word alignments in its output. As illustrated in Figure 2.8, the task is formulated as a token classification problem. The alignment model is constructed by concatenating an English text span, representing a labeled sequence, with a sentence in the target language. The model is fine-tuned to predict which tokens in the target sentence correspond to the source text span. They automatically generate a silver fine-tuning dataset using Wikipedia texts in the target language. Anchor




Figure 2.9 – Illustration of the cosine similarity between token embedding representations using Multilingual BERT.

text in hyperlinks indicates the location of named entities. The anchored text is machine-translated into the source language.

Previous methods that employ neural networks to compute alignments require fine-tuning data, which is very scarce and nonexistent for most language pairs. SimAlign (Jalili Sabet *et al.* 2020) leverages the contextual embeddings from state-of-the-art multilingual language models such as mBERT (Devlin *et al.* 2019). As depicted in Figure 2.9, SimAlign identifies alignments between words in parallel sentences based on the similarity of their contextual embeddings. This method allows for more accurate alignments that reflect the contextual use of words within specific sentences. Unlike statistically based word aligners that rely on statistical correlations in large corpora, SimAlign benefits from the deep linguistic and semantic understanding embedded in pre-trained language models. SimAlign offers improvements in alignment quality, especially for languages with complex morphological structures or less parallel data available for training. SimAlign is an unsupervised method, requiring no training data to compute word alignments. Still, SimAlign achieves better alignment accuracy than previous statistical and neural word alignment models. For those language pairs in which parallel data is available, AWESOME (Dou and Neubig 2021) improves on this idea by fine-tuning multilingual pre-trained language models on unlabeled parallel text. The main idea is that unsupervised training objectives over the parallel

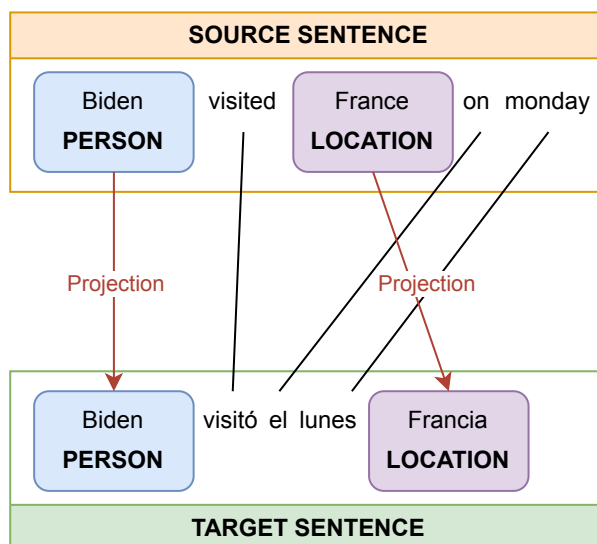


Figure 2.10 – Illustration of annotation projection using word-alignments

corpus improve the alignment quality of the models.

The word alignments generated by the previously described methods are used for annotation projection. Words in the target sentence are labeled with the same category as the aligned words in the source sentence. This process is depicted in Figure 2.10. Traditionally, most research in this area has focused on automatically annotating the English version of a multi-parallel corpus and then projecting these annotations to all other languages using statistical word alignments, as shown in the works of Yarowsky *et al.* 2001; Hwa *et al.* 2005; Ehrmann *et al.* 2011 and Fu *et al.* 2011. Wang and Manning 2014 introduces a refinement to this approach by projecting model expectations instead of direct labels, enabling the transfer of model uncertainty across languages and potentially reducing the risk of error propagation. Nevertheless, inaccuracies in word alignment computation remain a significant issue, often resulting in incorrect annotation projections and the generation of noisy data.

To address this problem, Ni *et al.* 2017 propose a heuristic scheme for selecting high-quality projection-labeled data from the noisy dataset. This scheme also includes projecting word embeddings from the target language into the source language, allowing the application of the source-language sequence labeling system to the target language without the need for re-training. Agerri *et al.* 2018 automatically annotate parallel data for multiple source languages and project the

labeled data to a single target language. This method demonstrates that leveraging multiple sources can significantly enhance the quality of the projections.

Instead of relying on automatic labels for the source part of a parallel corpus, Tiedemann *et al.* 2014; Fei *et al.* 2020 use Machine Translation to automatically translate the sentences of a gold-labeled dataset to the target languages. The translated data is subsequently annotated by projecting the gold labels from the source dataset onto it. Tiedemann *et al.* 2014 make use of GIZA++ for word alignments, whereas Fei *et al.* 2020 utilize the word alignment probabilities calculated with FastAlign, alongside the part-of-speech (POS) tag distributions of the source and target words, to enhance the precision of annotation projection.

In contrast to the methods previous mentioned, Shah *et al.* 2010 implement a Translate-Test strategy in which input sentences are first translated into the source language using Google Translate. These sentences are then annotated by the model, and the annotations are projected onto the target sentences using alignments computed by GIZA++.

Other annotation projection methods

With the recent advancements in supervised machine translation, a new line of research has emerged which aims to replace word alignments in favour of directly using Machine Translation models.




Figure 2.11 – Illustration of annotation projection using Machine Translation. Individually labeled sequences are translated separately from the rest of the sentence. The translations of these sequences are then matched with the translations produced by translating the entire sentence.

Jain *et al.* 2019 introduce a "translate-match" methodology, which is illustrated in Figure 2.11. In this approach the complete sentence, including labeled spans or entities, is first translated into the target language. Simultaneously, the labeled spans are translated independently of the full sentence. These individually translated spans are then matched with the corresponding spans in the translated

sentence. However, this matching process does not guarantee that labeled spans will retain consistent translations when isolated from the sentence context, so the authors generate multiple translation candidates for each span and select the best match based on orthographic and phonetic similarities.




Figure 2.12 – Illustration of annotation projection using Machine Translation and placeholders. Labeled sequences are replaced by a placeholder. The sentence with placeholders and the labeled sequences are translated independently. After translation, the placeholders are replaced with the corresponding labeled sequence translation.

To improve the matching process Zhou *et al.* 2022a propose to replace the labeled sequences in the source sentence with a placeholder. The sentence with placeholders and the labeled sequences are translated independently. Finally, the placeholders in the translated sentence are replaced with the corresponding labeled sequence translation. This process is illustrated in Figure 2.12. They found that the translation model preserves the placeholders. Although this technique effectively addresses the matching challenge, it may introduce translation artifacts. These artifacts arise because the translation model does not process the entire context of the source sentence, potentially diminishing the overall quality of the translation.




Figure 2.13 – Illustration of the mark-then-translate approach. Markers are introduced around the labeled sequences. The sentence and the labeled spans are translated together.

Instead of translating the labeled sequences independently from the source sentence, Daza and Frank 2019 and latter Chen *et al.* 2023 introduce markers in the source sentence as depicted in Figure 2.13. This mark-then-translate approach

allows the model to jointly translate the source sentence and the labels. However, their studies reveal that the inclusion of markers can degrade the quality of the translation. Additionally, they encounter low projection rates, meaning that markers are frequently omitted in the translated output. To address this, they fine-tune the translation model using a synthetic dataset with named entity recognition annotations. Post-fine-tuning, the model not only preserves translation quality but it also surpasses the accuracy of word-alignment models in annotation projection tasks.

In a subsequent work Le *et al.* 2024 enhances this method by implementing a constrained decoding algorithm, which ensures that the introduction of markers does not compromise the quality of the translation. In this improved approach, the training data in the high-resource language is first translated without markers. A second decoding phase then integrates the markers, with the constraint that the translation must align with the initial, marker-free output. This two-step process guarantees that the final translated sentence with markers remains consistent with what the model would have produced without them, thus preserving the translation quality.

Similar to the previous method, Parekh *et al.* 2024 also follows a two-step approach: first, translating the sentence without markers, and then performing the label translation. However, instead of using a Machine Translation system, they utilize an instruction-tuned large language model (Llama-2 Touvron *et al.* 2023b) to perform the task in a few-shot setting with a few randomly selected examples. Although their method proves to be effective, most low-resource languages lack a high-quality instruction-tuned model, which limits the applicability of this approach.




Figure 2.14 – Illustration of bilingual dictionary generation. Monolingual embeddings are projected into a shared space in which a bilingual dictionary is computed by k-nearest-neighbor. Figure reproduced from Xie *et al.* 2018.

All the methods previously described presuppose the availability of a high-quality Machine Translation model and a sizeable parallel corpus containing both the source and target languages. However, this is not the case for all language pairs. For instance, such resources are nonexistent for translations between English and some very low-resource African languages. Taking this into account, an alternative research direction aims to facilitate data transfer between a high-resource language and a very low-resource language using only minimal resources, specifically a bilingual dictionary and unlabeled text in the target language. Xie *et al.* 2018 first learns monolingual word embeddings (Mikolov *et al.* 2013; Pennington *et al.* 2014; Bojanowski *et al.* 2017) for the source and the target language. As depicted in Figure 2.14, both embeddings are mapped into a bilingual vector space using a cross-lingual vector projection (Zhang *et al.* 2016; Artetxe *et al.* 2016; Smith *et al.* 2017). A word translation dictionary is then computed by k-nearest-neighbor. The source sentence is translated word-by-word using this dictionary and the source label is copied for each corresponding word in the target sentence. Finally, a NER model is trained using the translated data. Building on this approach, Guo and Roth 2021 aim to refine the low-quality output that results from literal word-by-word translations. They employ a target language model and constrained beam search to produce text in the target language that exhibits a more natural and contextually appropriate word order. The constraints are designed to ensure the presence of entities in the generated text.

2.2.2 Model-based transfer

The model transfer approach involves leveraging multilingual models trained in high-resource languages to perform tasks in low-resource languages. In contrast, the data transfer paradigm focuses on manipulating the data to fit a monolingual model. Model transfer exploits the shared representation of languages in a pre-trained multilingual model. Thus, these models can be fine-tuned for a specific task in the source language and then applied without any modification to label text in any of the multiple languages the model supports. This approach is illustrated in Figure 2.15.

Some of the first attempts at model-based cross-lingual transfer (Täckström *et al.* 2012; Kozhevnikov and Titov 2014; Bharadwaj *et al.* 2016; Chaudhary *et al.* 2018) leveraged the structural similarities between languages to facilitate tasks in languages with limited or non-existent training data. However, model-based transfer began to make significant progress (Artetxe and Schwenk 2019; Pires *et al.* 2019) following the introduction of Transformer-based (Vaswani *et al.*

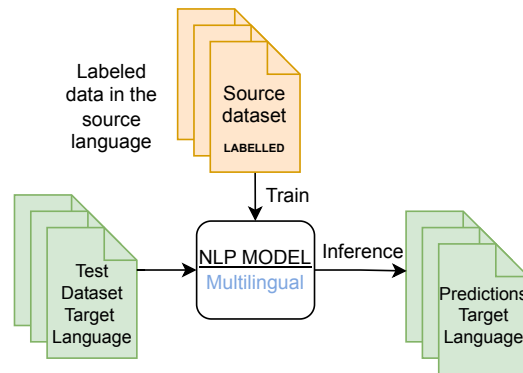


Figure 2.15 – Illustration of the model-based cross-lingual transfer approach. A pre-trained multilingual model is finetuned with data in the source language and then applied without modification to label text in the target language.

2017) multilingual language models such as BERT (Devlin *et al.* 2019) or XLM-RoBERTa (Conneau *et al.* 2020). These models were pre-trained using language modelling objectives on extensive datasets comprising over 100 languages. During this pre-training phase, the models acquired a shared representation for all included languages. Subsequently, these models can be fine-tuned on specific tasks with data from a source language and then applied to label data in any of the supported languages directly.

3. CHAPTER

Data transfer vs Model transfer

In this chapter we will conduct an in-depth study of the two main techniques used so far for cross-lingual zero-shot Sequence Labeling, focusing on either data or model transfer. We will apply state-of-the-art Machine Translation models, word alignments, and language models to assess the performance of various cross-lingual sequence labeling approaches. We will identify the advantages and shortcomings of each method, as well as the challenges faced by current techniques for cross-lingual zero-shot sequence labeling. These insights will serve as a foundation for our work in subsequent chapters.

3.1 Motivation and contributions

The performance of supervised machine-learning methods for Natural Language Processing, including advanced deep-neural models (Lample *et al.* 2016; Akbik *et al.* 2018; Devlin *et al.* 2019; Conneau *et al.* 2020), heavily depends on the availability of manually annotated training data. In addition, supervised models show a significant decrease in performance when evaluated in out-of-domain settings (Liu *et al.* 2021). This means that obtaining optimal results would require to manually generate annotated data for each application domain and language, an unfeasible task in terms of monetary cost and human effort.

zero-shot cross-lingual transfer approaches aim to apply the resources from a high-resource source language to low-resource target languages. In this chapter, we perform an in-depth study of the two main techniques employed so far for

cross-lingual zero-shot sequence labeling, based either on data or model transfer. We implement both approaches using the latest advancements in machine translation, word aligners, and multilingual language models. We focus on two sequence labeling tasks, namely, Named Entity Recognition (NER) and Opinion Target Extraction (OTE). To this end, we present a data-based cross-lingual transfer approach consisting of translating gold-labeled data between English and seven other languages using state-of-the-art Machine Translation systems. Sequence labeling annotations are then automatically projected for every language pair using word alignments. We then compare the performance obtained for each of the target languages against the performance of the zero-shot cross-lingual method, consisting of fine-tuning the multilingual language models on the English gold data and generating the predictions in the required target languages.

The main contributions of this chapter are the following:

- We empirically establish the required conditions for each of these two approaches, data-transfer and zero-shot model-based cross-lingual transfer, to outperform the other. In this sense, our experiments show that contrary to what previous research suggested (Fei *et al.* 2020; Li *et al.* 2021), the zero-shot transfer approach is the most effective method when using high-capacity multilingual language models like XLM-R large. However, data transfer approaches remain valuable for models with weaker downstream cross-lingual performance. As there is no universally available multilingual pretrained model for every language and domain, data-based methods retain their relevance.
- Our evaluation proves that despite high-quality machine translations and annotation projections, in the data transfer approach, we have identified issues like many-to-one translations or misalignments, which seem to account for the lower performance of data transfer methods compared to the model-based approach.
- We perform an extensive error analysis which shows that using English gold-labeled data often produces a signal which, due to inherent differences in language usage, differs from the signal received when using gold-labeled data in the target language. This cultural misalignment problem hinders the performance of the zero-shot cross-lingual methods. The issue is more pronounced with low-capacity models that have lower generalization capabilities, whereas higher-capacity models are more successful at labeling

words unseen in the training data that share a similar meaning with those seen during training.

3.2 Methodology

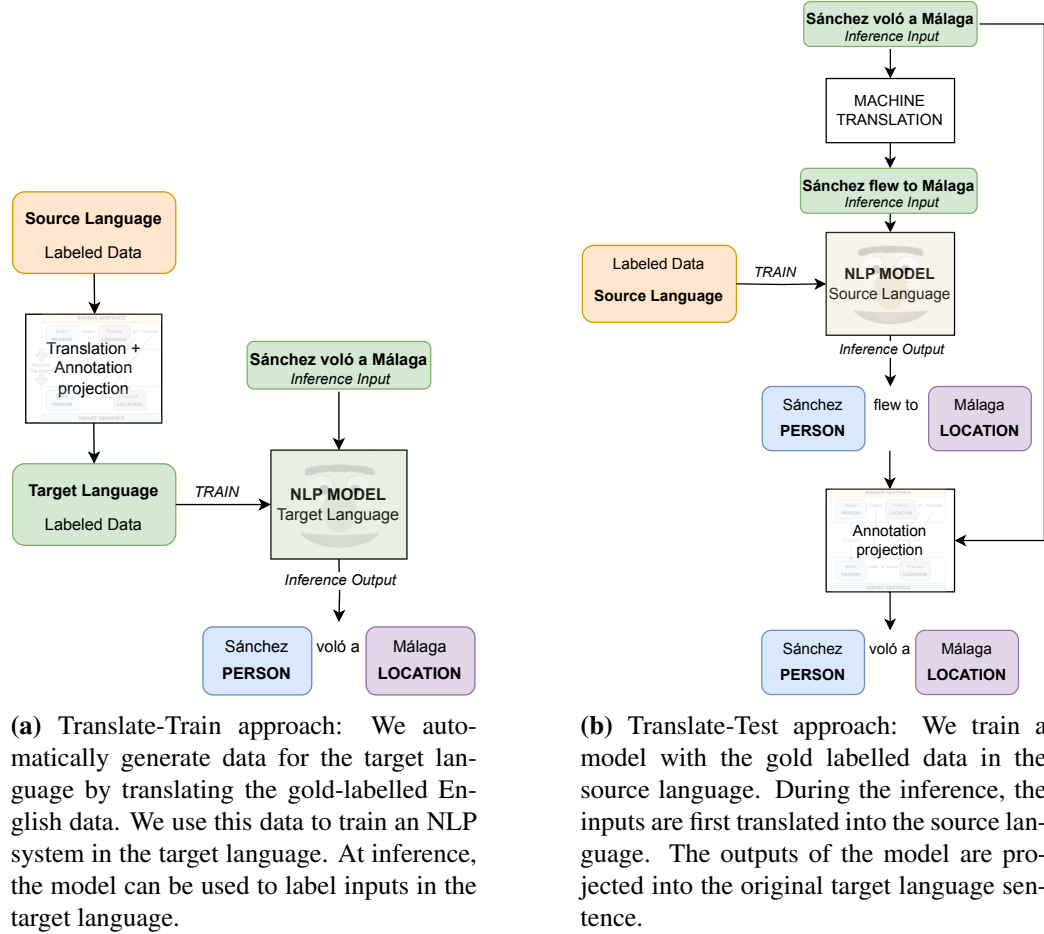
In this section, we describe our implementation of data-based and model-based transfer methods for Cross-Lingual Sequence Labeling. Our experiments follow a one-to-one framework, using English as the source language. Each model is evaluated in a single target language. We assume a scenario in which we have English gold-labeled train and development data. Furthermore, we also assume that the only gold-labeled data available for the target language is the evaluation set.

3.2.1 Data transfer

In the data-transfer paradigm, described earlier in Chapter 2.2, the information extraction model always performs inference in the language it was trained in. This language could be either the source language, typically English, or the target language. We propose to use Machine Translation and annotation projection. We have implemented two distinct data-transfer strategies: Translate-Train and Translate-Test.

Translate-Train

In the Translate-Train approach, our goal is to generate data for the target language automatically, as depicted in Figure 3.1a. The process begins with translating English gold-labeled data into the target language using state-of-the-art Machine Translation models. Subsequently, we project the gold labels from the original English sentences onto the translated ones. This results in an automatically generated dataset in the target language. We then utilize this dataset to train an NLP model specifically for the target language. During inference, sentences in the target language are fed directly into the model for annotation. In this approach, the model is trained using “silver” data, the quality of which depends on the Machine Translation model’s accuracy and the accuracy of the annotation projection algorithm. The data-transfer occurs during the model’s training phase. Once trained, the model is capable of directly annotating sentences in the target language without requiring any additional steps



(a) Translate-Train approach: We automatically generate data for the target language by translating the gold-labelled English data. We use this data to train an NLP system in the target language. At inference, the model can be used to label inputs in the target language.

(b) Translate-Test approach: We train a model with the gold labelled data in the source language. During the inference, the inputs are first translated into the source language. The outputs of the model are projected into the original target language sentence.

Figure 3.1 – Illustration of the two data transfer approaches we have implemented. They are differentiated by the direction in which we translate the data. In both cases, English is the source language and Spanish is the target language.

Translate-Test

The Translate-Test approach, depicted in Figure 3.1b, involves training an NLP model using English gold-labeled data. During the inference step, we first translate the input sentences in the target language into English. The model then generates annotations for these translated sentences. Subsequently, we project back these annotations onto the original sentences in the target language. Unlike the Translate-Train approach, where the model is trained with “silver” data, in the Translate-Test, the model is trained with gold data. However, during inference, the input must be first translated into English, and the model’s output annotations must be projected back onto the original target language sentences. Therefore, as with Translate-Train, the quality of the predictions in Translate-Test also depends on the Machine Translation model’s quality and the accuracy of the annotation projection algorithm.

Annotation Projection Algorithm

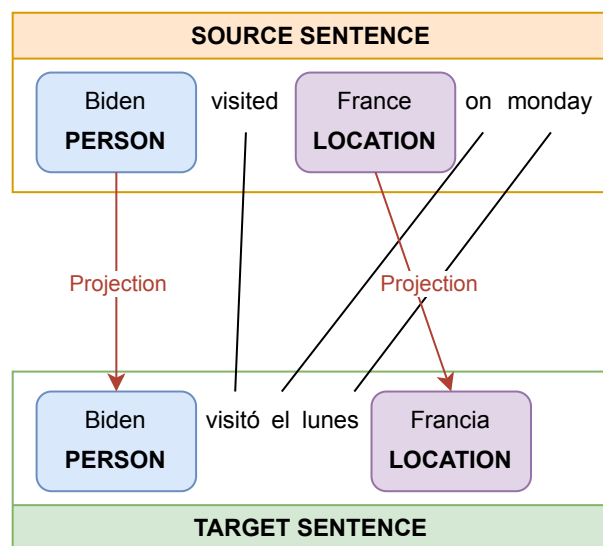


Figure 3.2 – Illustration of the translation and annotation projection method for the Named Entity Recognition task.

In both the Translate-Train and Translate-Test approaches, we project labels from the source language into the target language using automatic word alignments. Word alignments map words in a source language to their correspond-

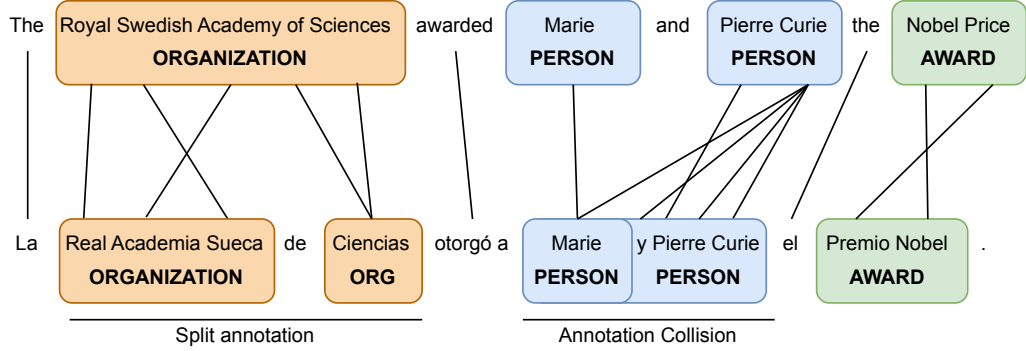


Figure 3.3 – Illustration of the split annotation and annotation collision errors when projecting a sentence.

ing translations in a target language. Consider a source language sentence $x = \langle x_1, \dots, x_n \rangle$ of length n , and its translation $y = \langle y_1, \dots, y_m \rangle$ in the target language with length m . We employ a word aligner to identify pairs $A = \{ \langle x_i, y_j \rangle : x_i \in \mathbf{x}, y_j \in \mathbf{y} \}$, where each pair $\langle x_i, y_j \rangle$ indicates y_j is the translation of x_i . For a sequence $s = \langle x_a, \dots, x_b \rangle \in \mathbf{x}$ labeled with category C , we assign the same category C to the sequence $t = \langle y_c, \dots, y_d \rangle \in \mathbf{y}$, if every word y_j in t is aligned with at least one word x_i in s : $\{ \forall y_j \in t \exists x_i \in s : (\langle x_i, y_j \rangle \in A) \}$. Essentially, if a word in the source sentence labeled with a category is aligned with a word in the target sentence, we label the target word with the same category. This method is illustrated in Figure 3.3.

When projecting annotations, we encounter two primary challenges: *Split Annotations* and *Annotation Collisions*. Split Annotations occur when a labeled sequence in the source sentence is divided into multiple sequences in the target sentence, often due to missed alignments by the word alignment algorithm. In such cases, we merge the target sequences if they are separated by only one word. If multiple sequences remain, we retain the longest sequence and discard the others. Annotation Collisions arise when a word in the target sentence aligns with two different labeled sequences in the source sentence, typically resulting from incorrect alignments. This causes a word, or sequence of words, to have overlapping label categories in the target sentence. If the conflicting sequences belong to the same category, we merge them into a single label in the target sentence. If they are of different categories, we choose the longer target sequence and discard the other one. Additionally, if a punctuation symbol in the target sentence aligns with a labeled word in the source sentence, we disregard this alignment. These scenar-

ios are illustrated in Figure 3.2. For example, in the projection of “Royal Swedish Academy of Science” to “Real Academia Sueca de Ciencias”, the word “of” is incorrectly aligned, leading to a projection gap. Since this gap involves only one word, we merge the sequences, labeling “Real Academia Sueca de Ciencias” as a single organizational entity. In the case of “Marie” and “Pierre Curie”, incorrect alignment of “Curie” with both names results in overlapping person spans in the target sentence. As they are of the same category (Person), we merge them into a single Person label. Although this approach might not be perfect, it helps minimize errors caused by incorrect alignments.

3.2.2 Model transfer

Previously described in Chapter 2.2, the model transfer approach leverages multilingual language models Devlin *et al.* 2019; Conneau *et al.* 2020. These models are pre-trained on extensive unlabeled text corpora in multiple languages. As depicted in Figure 3.4, we fine-tune a multilingual language model with the gold-labeled data available in English. Once the training is complete, the model is capable of directly labeling text in any of the languages included in its pre-training phase. We will refer to this approach to as *zero-shot* cross-lingual sequence labeling.

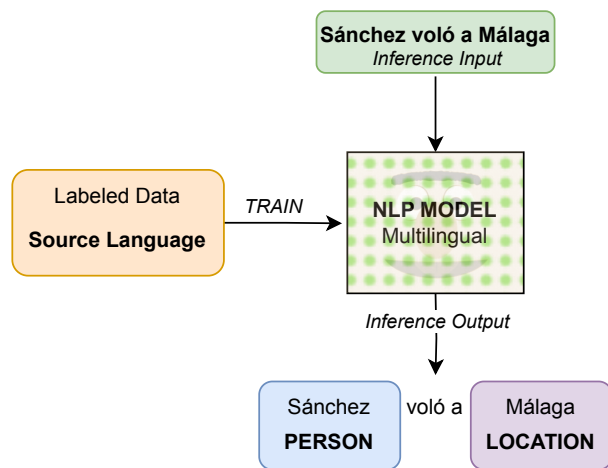


Figure 3.4 – Illustration of model transfer approach. A multilingual model is trained in the source language (English). The model is then used to label sentences in the target language (Spanish).

3.3 Experimental Setup

In this section, we will describe the experimental framework for this Chapter.

3.3.1 Datasets

We conducted experiments in two sequence labeling tasks, namely, Opinion Target Extraction and Named Entity Recognition. The tasks are illustrated in Figure 3.5. We list the number of examples in each dataset in table 3.1.




Figure 3.5 – Illustration of the sequence labeling tasks used in the experiments in this chapter.

Opinion Target Extraction (OTE): Given a review, the task is to detect the linguistic expression used to refer to the reviewed entity. For instance, in the sentence *Serves really good sushi*, the word *sushi* is the opinion target because it is the entity being discussed. We use the English SemEval 2016 Aspect Based Sentiment Analysis (ABSA) dataset (Pontiki *et al.* 2016). This dataset comprises user reviews from the restaurant domain. We experiment with the English, Spanish, Dutch, French, Russian and Turkish datasets from the restaurant domain.

Named Entity Recognition (NER): Given a text, the task is to detect named entities and classify them according to some pre-defined categories. For Spanish and Dutch we use the CoNLL-2002 datasets (Sang 2002). For English and German we use the CoNLL-2003 datasets (Sang and Meulder 2003) and for Italian, we use the Evalita 2009 data (Speranza 2009). We map the Geo-Political Entities from Evalita 2009 to *location* labels to make them compatible with the CoNLL data. This dataset contains labeled sentences from news articles.

Language	Dataset (Citation)	Train	Dev	Test	Labels
Opinion Target Extraction					
English	SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016)	2000	-	676	(1) Target
Spanish	SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016)	2070	-	881	
French	SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016)	1665	-	668	
Dutch	SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016)	1722	-	575	
Russian	SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016)	3655	-	1209	
Turkish	SemEval 2016 ABSA (Pontiki <i>et al.</i> 2016)	1232	-	144	
Named Entity Recognition					
English	CoNLL-2003 (Sang and Meulder 2003)	14987	3466	3684	(4) Person, Location, Organization, Miscellaneous
Spanish	CoNLL-2002 (Sang 2002)	6871	1914	1516	
German	CoNLL-2003 (Sang and Meulder 2003)	12705	3068	3160	
Dutch	CoNLL-2002 (Sang 2002)	15806	2895	5195	
Italian	Evalita 2009 (Speranza 2009)	11227	-	4136	

Table 3.1 – Number of sentences for each dataset split.

3.3.2 Machine Translation

We tested different open-source and commercial Machine Translation systems. Including DeepL ¹, OpusMT (Tiedemann and Thottingal 2020), M2M100 (1.2B, Fan *et al.* 2021) and mBART (mbart-large-50, Liu *et al.* 2020; Tang *et al.* 2020). A qualitative analysis performed during the projection of the OTE labels established that DeepL produced more fluent translations. Thus we decided to use DeepL (web version during the second half of 2021) to perform the Machine Translation for our data-based cross-lingual transfer experiments. Turkish was not supported by DeepL at the time, so we used M2M100 1.2B instead. Experiments that compare the different Machine Translation models are discussed in Section 3.5.1.

3.3.3 Word Alignments

To compute word alignments, we use the AWESOME system (Dou and Neubig 2021). AWESOME leverages multilingual pre-trained Language Models and fine-tunes them on parallel corpora. Unsupervised training objectives over the parallel corpus improve the alignment quality of the models. The authors claim that AWESOME works better with multilingual-BERT Devlin *et al.* 2019 as the backbone, so we follow their advice. We also experiment with GIZA++ (Och and Ney 2003)

¹<https://www.deepl.com/es/translator>

and FastAlign (Dyer *et al.* 2013b), which are models based on statistical machine translation. Additionally, we explore SimAlign (Jalili Sabet *et al.* 2020), which, similar to AWESOME, leverages multilingual pre-trained models, although in a fully unsupervised manner. All the systems are extensively described in Chapter 2.2.1. As demonstrated in Section 3.5.2, AWESOME produced the highest F1 scores when comparing the model projections to manually annotated projections.

3.3.4 Sequence labeling Models

We use three state-of-the-art multilingual pre-trained language models for sequence labeling: multilingual BERT (mBERT) (Devlin *et al.* 2019) and XLM-RoBERTa (XLM-R) base and large (Conneau *et al.* 2020).

The primary distinctions among these models lie in their parameter count and the size of their pretraining corpora. mBERT has 110 million parameters and was trained on Wikipedias² of 104 languages. In contrast, XLM-RoBERTa-base contains 250 million parameters, and XLM-RoBERTa-large has 560 million parameters. Both versions of XLM-RoBERTa were trained using a corpus of 340 billion tokens extracted from 100 languages sourced from CommonCrawl³. This corpus, significantly larger as noted in Conneau *et al.* 2020 encompasses a far broader range of data compared to that used for mBERT, covering various diverse domains. We categorize models with a lower parameter count and those pre-trained on smaller datasets, such as mBERT, as low-capacity language models. Conversely, we classify larger models trained on more extensive corpora, like XLM-RoBERTa-large, as high-capacity language models (Aharoni *et al.* 2019).

As depicted in Figure 3.6, for both language models, we add a token classification layer (a linear layer) on top of each token’s representation. This layer computes the probability distribution of the labels for each token. If a word is split into multiple sub-tokens, we use the representation of the first sub-token as input for the classifier. During training, we fine-tune the complete parameters of the model along with the token classification layer. We utilize the sequence labeling implementation from the Hugging Face open-source library (Wolf *et al.* 2020). As listed in Section 3.3.1, due to the significant difference in the total number of training examples between the OTE and NER datasets, we employ slightly different hyperparameter settings for each task. For OTE we use a batch size of 32, $5e-5$ learning rate, and we train the model for 10 epochs and 128 maximum se-

²<https://www.wikipedia.org/>

³<https://commoncrawl.org/>




Figure 3.6 – Implementation of the sequence labeling model. An encoder-based Language Model is fed the input sequence. The output representations are used by a token classification linear layer to predict the labels.

quence length. Since only a train and test splits are available for the OTE task, we use the train set as both, train and development data. For NER we use a batch size of 32, $2e-5$ learning rate, and we train the model for 4 epochs and 256 maximum sequence length. For both tasks, we use the “BILOU” encoding scheme (Ratinov and Roth 2009).

We report the F1 score, which is the standard metric in sequence labeling tasks. F1 scores and standard deviation scores are reported by averaging the results of five runs with different random seeds.

3.4 Experimental Results

In this section we compare the TRANSLATE-TRAIN and TRANSLATE-TEST data-transfer approaches with the ZERO-SHOT model-transfer approach. As an upper bound, we also train the language models on gold-labeled data in the target languages. This upper bound, which we refer to as GOLD, is intended to assess the performance of the cross-lingual transfer approaches with respect to an ideal setting.

3.4.1 Opinion Target Extraction

Opinion Target Extraction (OTE) results are reported in Table 3.2. When using mBERT, a low-capacity model, the zero-shot model transfer approach obtains better results for Spanish and French. However, for Dutch, Russian, and Turkish, the data-transfer approaches are superior. However, the overall picture changes as we increase the model capacity. When using XLM-RoBERTa (XLM-R) base, the zero-shot baseline is much closer to the gold upper bound than that of mBERT. This demonstrates that XLM-R has better multilingual transfer learning capabilities for these tasks. In fact, the zero-shot transfer outperforms the Translate-Train and Translate-Test for every language except Turkish. The XLM-R base results on gold-labeled data are also substantially better than those of mBERT, which once again demonstrates the better proficiency of XLM-R in these tasks. To summarise, XLM-R large offers the best cross-lingual transfer performance, as the zero-shot transfer is clearly superior for every language, including Turkish.

A common trend for all three models in the OTE benchmarks is that the Translate-Train approach consistently performs better than the Translate-Test approach. As expected, all the approaches achieve a performance significantly lower than the gold upper bound.

mBERT				
Language	Gold	Zero-shot	Trans-Train	Trans-Test
English	76.2 \pm 0.9	-	-	-
Spanish	75.2 \pm 0.5	68.4 \pm 0.6	67.9 \pm 0.8	62.2 \pm 1.2
French	74.0 \pm 1.1	62.7 \pm 1.2	59.7 \pm 1.2	57.6 \pm 1.1
Dutch	69.7 \pm 0.9	61.7 \pm 0.8	64.3 \pm 1.5	67.0 \pm 0.8
Russian	72.5 \pm 0.5	53.8 \pm 2.2	62.9 \pm 0.6	59.7 \pm 0.4
Turkish	62.0 \pm 1.2	45.3 \pm 4.0	45.7 \pm 2.3	35.5 \pm 2.4
XLM-R base				
English	84.4 \pm 0.9	-	-	-
Spanish	81.1 \pm 0.7	78.2 \pm 0.4	72.5 \pm 0.7	62.9 \pm 0.9
French	80.2 \pm 0.6	72.7 \pm 0.3	64.7 \pm 0.8	60.0 \pm 0.6
Dutch	80.8 \pm 1.7	75.5 \pm 0.8	70.0 \pm 1.6	71.0 \pm 1.5
Russian	81.5 \pm 0.3	74.9 \pm 0.9	69.5 \pm 0.3	62.2 \pm 1.6
Turkish	69.0 \pm 1.1	58.1 \pm 3.5	58.9 \pm 1.8	36.4 \pm 1.8
XLM-R large				
English	86.4 \pm 1.1	-	-	-
Spanish	83.6 \pm 0.1	79.3 \pm 0.8	73.7 \pm 1.1	64.0 \pm 1.4
French	82.2 \pm 0.6	74.6 \pm 1.7	66.1 \pm 0.6	60.7 \pm 0.6
Dutch	80.4 \pm 2.1	77.7 \pm 1.9	74.0 \pm 1.3	72.9 \pm 1.8
Russian	82.8 \pm 0.4	76.8 \pm 1.3	69.3 \pm 2.3	62.2 \pm 1.3
Turkish	72.3 \pm 2.4	62.4 \pm 1.0	57.8 \pm 2.4	33.7 \pm 0.9

Table 3.2 – OTE F1 scores with models of different capacities in the SemEval 2016 ABSA (Pontiki *et al.* 2016) dataset.

3.4.2 Named Entity Recognition

The NER results presented in Table 3.4 show a number of different patterns. First, the zero-shot approach using mBERT outperforms the data-based cross-lingual transfer methods for the majority of languages. Second, the Translate-Test approach is consistently superior to the Translate-Train approach. Third, the mBERT performance on gold data is similar to that of XLM-R base. Finally, the data transfer approaches achieve the best performance when using XLM-R base for German and Italian.

The only result that remains consistent from the OTE tasks is that the zero-shot approach using XLM-R large achieves the best results for all languages.

mBERT				
Language	Gold	Zero-shot	Trans-Train	Trans-Test
English	90.7 \pm 0.3	-	-	-
Spanish	87.4 \pm 0.4	74.6 \pm 0.4	69.5 \pm 0.4	70.8 \pm 0.6
German	82.0 \pm 0.4	71.0 \pm 0.9	70.1 \pm 0.3	70.6 \pm 0.5
Dutch	90.8 \pm 0.4	78.5 \pm 0.5	74.4 \pm 0.6	75.4 \pm 0.8
Italian	84.7 \pm 0.3	68.2 \pm 0.5	68.7 \pm 0.5	70.7 \pm 0.3
XLM-R base				
English	90.4 \pm 0.2	-	-	-
Spanish	87.7 \pm 0.2	75.0 \pm 0.4	70.1 \pm 0.6	72.5 \pm 0.2
German	83.1 \pm 0.3	67.9 \pm 0.5	70.5 \pm 0.5	70.1 \pm 0.8
Dutch	89.8 \pm 0.2	78.1 \pm 0.6	73.3 \pm 0.9	74.7 \pm 0.4
Italian	84.3 \pm 0.3	71.2 \pm 0.5	71.1 \pm 0.4	71.7 \pm 0.3
XLM-R large				
English	92.4 \pm 0.1	-	-	-
Spanish	88.9 \pm 0.2	79.5 \pm 1.0	70.9 \pm 0.6	74.0 \pm 0.5
German	85.1 \pm 0.6	74.5 \pm 0.7	73.7 \pm 0.5	72.9 \pm 0.3
Dutch	92.9 \pm 0.7	82.3 \pm 0.6	77.5 \pm 0.9	77.2 \pm 0.6
Italian	87.5 \pm 0.2	76.0 \pm 0.5	73.7 \pm 0.4	73.5 \pm 0.6

Table 3.3 – NER F1 scores with models of different capacities in the CoNLL-2002 (Sang 2002) and CoNLL-2003 (Sang and Meulder 2003) datasets.

We also compare our data-transfer and model-transfer implementations with previous research that leverages parallel data and/or annotation projection meth-

Method	Model	Approach	Spanish	German	Dutch
Dou and Neubig 2021	mBERT	Translate train	64.3	-	-
Jain <i>et al.</i> 2019	BiLSTM + CRF	Translate train	73.5	61.5	69.9
Guo and Roth 2021	BiLSTM + CRF	Translate train	77.9	71.4	80.6
Li <i>et al.</i> 2021	XLM-R large	Translate train	78.9	76.9	79.7
Ours	XLM-R base	Translate train	70.1	70.5	73.3
Ours	XLM-R base	Translate test	72.5	70.1	74.7
Ours	XLM-R large	Translate train	70.9	73.7	77.5
Ours	XLM-R large	Translate test	74.0	72.9	77.2
Ours	mBERT	zero-shot	74.6	71.0	78.5
Ours	XLM-R base	zero-shot	75.0	67.9	78.1
Ours	XLM-R large	zero-shot	79.5	74.5	82.3

Table 3.4 – Comparison between the previous research methods that leverage projections, the zero-shot baselines and our annotation projections in the CoNLL-2002 (Sang 2002) and CoNLL-2003 (Sang and Meulder 2003) datasets. F1 score reported

ods on the NER CoNLL 2002-2003 data. The results are listed in Table 3.4. XLM-R large in a zero-shot setting not only outperforms our data-transfer implementation, but it is also superior to previous data-transfer approaches. The only exception is the result obtained by Li *et al.* 2021 for German.

Our data transfer approaches, although not achieving the best results for every language, are competitive with previously proposed methods. It is important to note that while we only leverage translations of the NER data, previous research uses other resources such as the automatic annotation of large parallel corpora.

3.4.3 Discussion

Previous research has demonstrated (Pires *et al.* 2019; Wu and Dredze 2020) that mBERT’s performance in cross-lingual transfer learning greatly differs depending on whether the source and target languages are topologically similar or not, with the former being the scenario in which the model performs best. Secondly, the monolingual performance of mBERT, as well as its cross-lingual transfer performance, is much better for high-resource languages than for low-resource languages. This is consistent with our results for mBERT in the NER and OTE tasks. The zero-shot transfer works best for French and Spanish, which are somewhat

topologically similar to English. However, the performance of the zero-shot approach is worse than the Translate-Train and Translate-Test approaches for Russian and Turkish.




Figure 3.7 – Amount of data in GiB (log-scale) for the languages we use in our experiments in Wiki-100 corpus used for training mBERT and the CC-100 used for XLM-R. The full figure can be found in Conneau *et al.* 2020

While mBERT was trained using Wikipedia data for 104 languages (Devlin *et al.* 2019), XLM-R (both base and large) has been trained using CommonCrawl (Conneau *et al.* 2020), a much larger multilingual corpus with a variety of texts extracted from the web. As illustrated in Figure 3.7, XLM-R was trained with orders of magnitude more data for Russian and Turkish compared to mBERT, thus acquiring higher proficiency in those languages. This results in the zero-shot approach using XLM-R achieving better performance in Russian and Turkish than the data-transfer approaches. The CommonCrawl dataset also contains a more diverse variety of texts sampled from the Web, perhaps including texts of similar domains to those in the OTE datasets. This may account for the large differences in OTE performance between XLM-R and mBERT. In this sense, the similar performance between mBERT and XLM-R base for NER may be partially attributed to the fact that the CoNLL and Evalita datasets consist of news stories in which most of the labeled entities may appear in the Wikipedia texts used to pre-train mBERT.

Our results suggest that the performance of model-transfer and data-transfer approaches relies on the model’s proficiency in the target language and data domain. If a high-capacity model, such as XLM-R large, with strong proficiency in

the target language and domain is available, the zero-shot cross-lingual transfer method achieves the best results. However, for languages and domains where the model lacks sufficient proficiency, data-based cross-lingual transfer (Translate-Train and Translate-Test) approaches remain mainly useful. While XLM-R large in a zero-shot cross-lingual setting achieves the best results for every task and language in the experiments performed in this section, in the next chapters we will demonstrate that XLM-R large is not proficient in every language and domain, for example, in African languages (Adelani *et al.* 2022). Thus, developing both better model-transfer and data-transfer approaches is of great relevance.

3.5 Error Analysis

The experiments described in Section 3.4 showed that when using XLM-R large, the zero-shot approach outperforms the Translate-Train and Translate-Test approaches. The effectiveness of the Translate-Train and Translate-Test approaches depends on the quality of the Machine Translation and annotation projection models. In this section, we will assess the performance of both Machine Translation and annotation projection to better understand if the data-transfer methods are hindered by the performance of any of these steps. We will also conduct an error analysis to gain a deeper understanding of the shortcomings of Translate-Train and Translate-Test in comparison to the zero-shot cross-lingual transfer.

3.5.1 Downstream evaluation of Machine Translation Models

To assess the impact of the Machine Translation system employed, we translated the English OTE gold-labeled data using four distinct translation systems. In all experiments, we utilized AWESOME as the word aligner for annotation projection. We fine-tuned XLM-R large with each set of generated training data and assessed its performance against the gold-labeled test set for each target language. Based on the results provided in Table 3.5, we can conclude that there are no significant differences in the final F1 scores when employing different translation systems. For each language, a different model demonstrates the best performance, although all systems exhibit similar average performance across all languages. The exception is Turkish, a language in which MarianMT underperforms.

Language	MarianMT	Mbart	M2M100	DeepL
Spanish	75.6 ± 0.8	75.3 ± 0.7	74.2 ± 0.8	73.7 ± 1.1
French	64.5 ± 1.6	66.4 ± 1.1	64.9 ± 1.3	66.1 ± 0.6
Dutch	70.0 ± 2.0	68.8 ± 4.0	70.1 ± 3.1	74.0 ± 1.3
Russian	66.6 ± 4.4	69.7 ± 1.4	69.7 ± 0.7	69.3 ± 2.3
Turkish	49.5 ± 2.9	56.1 ± 5.2	57.8 ± 2.4	-

Table 3.5 – OTE F1 score in the SemEval 2016 ABSA Pontiki *et al.* 2016 dataset of different XLM-R large models trained using data generated with different translation systems.

3.5.2 Evaluating the Projection Method

The performance of the data transfer method relies on the quality of the annotation projections. If the annotation projection algorithm does not produce highly accurate projections, it will generate noisy data that would hinder both the Translate-Train and Translate-Test approaches. In this section, we will evaluate the performance of various word alignment systems and compute the performance of the Translate-Test approach using gold annotation projections. To achieve this, we enlisted human annotators who manually projected labels from the English OTE gold-labeled data onto automatic translations in Spanish, French, and Russian. The machine translations were generated using DeepL for Spanish, French, and Russian, and M2M100 for Turkish. The annotators are NLP PhD candidates with either native or proficient skills in both English and the target language. For this experiment, we developed an application to assist during the annotation process. The annotator views the sentence in English, with a highlighted target, and must select the corresponding target in a translated sentence. The complete guidelines and the application code provided to the annotators are available in GitHub ⁴.

First, we compare the projection of annotations automatically generated by different word alignment methods with those provided by human annotators. Table 3.6 shows that language model-based methods (SimAlign and AWESOME) outperform statistically based methods (GIZA++ and FastAlign) by a wide margin. AWESOME consistently outperforms SimAlign for every language. It is worth mentioning that SimAlign uses a fully unsupervised algorithm, while AWESOME requires fine-tuning a Language Model; therefore, the cost in compute resources is significantly higher for AWESOME.

⁴<https://github.com/ikergarcia1996/Annotation-Projection-App>

Language	GIZA++	FastAlign	SimAlign	AWESOME
Spanish	77.0	75.0	86.7	91.5
French	73.3	72.9	86.3	91.3
Russian	72.4	76.9	87.7	93.4
Turkish	64.0	68.4	81.9	88.5

Table 3.6 – OTE F1 score in the SemEval 2016 ABSA Pontiki *et al.* 2016 dataset between the human annotation projections vs the automatic projections generated using different alignment models.

Language	Translate Train	Translate Train (Manual)	Zero-shot
Spanish	73.7 \pm 1.1	75.1 \pm 1.2	79.3 \pm 0.8
French	66.1 \pm 0.6	67.9 \pm 1.0	74.6 \pm 1.7
Russian	69.3 \pm 1.3	69.4 \pm 2.1	76.8 \pm 1.3
Turkish	57.8 \pm 2.4	50.6 \pm 1.4	62.4 \pm 1.0

Table 3.7 – XLM-R large OTE F1 score in the SemEval 2016 ABSA Pontiki *et al.* 2016 dataset when training with automatically and manually projected datasets

The performance of the AWESOME system confirms that it is possible to generate high-quality annotations that closely resemble those generated by human experts. However, the model does make some mistakes. To understand the performance implications of the mistakes produced by AWESOME, we compare the Translate-Train approach when using data projected with AWESOME and when using data manually projected by human experts. To achieve this, we fine-tuned an XLM-R large model on both datasets. Table 3.7 shows that when training with manually projected data, the results are slightly better, except for Turkish, which again acts as an outlier. Nevertheless, the results when training with manually projected data are still inferior to the zero-shot model-transfer approach. From the results, we can conclude that the projection mistakes produced by AWESOME do not have a significant impact on the performance of the data-based cross-lingual transfer approach, although there is still room for improvement. More importantly, the data-transfer approach is not inferior to the model-based transfer approach due to the errors produced by the annotation projection step.

3.5.3 Categorization of mistakes

We have evaluated the quality of the Machine Translation models and the annotation projection algorithms. We found both to be of high quality and not responsible for the data-transfer method performing worse than the zero-shot transfer approach. In this section we will categorize the errors produced by each approach and compare them. Table 3.9 displays the most frequent false negatives and positives for which there is a significant discrepancy in frequency between methods.

There are a few challenging words that all the systems struggle with. For example, as previously reported by Agerri and Rigau 2019, the words “comida” (food) and “restaurante” (restaurant) are highly ambiguous in the ABSA task. Both appear labeled as target and unlabeled frequently. As expected, models struggle with these words. In addition, we have identified four main sources of errors:

English Word	Spanish Word	English Gold	Spanish Gold	Spanish Translate
Service	Servicio	153	229	133
Treatment	Trato	0	54	0
Attention	Atención	2	35	0
Place	Sitio	120	41	2
Place	Lugar	120	19	91

Table 3.8 – Number of times words appear as target words in the SemEval 2016 ABSA (Pontiki *et al.* 2016) train dataset.

Many-to-one translation: Multiple words can share the same sense in a given context and be used interchangeably. This is the case with opinion targets in the ABSA task, such as “trato”, “atención”, and “servicio” in Spanish. In the context of restaurant reviews, they all convey the same meaning as the English word “service.” There are 160 sentences in the English gold-labeled data containing the word “service”; in 153 of them, “service” is labeled as a target. Machine Translation systems, such as DeepL in our experiments, systematically translate it as “servicio.” However, as shown in Table 3.8, in the Spanish gold-labeled data, “service” is also commonly referred to as “trato” or “atención”, not just “servicio.” Therefore, the training dataset translated and projected from English into Spanish, as demonstrated in previous sections, encompasses high-quality translations and annotation projections. Still, this dataset is not a good reflection of reviews written by native Spanish speakers. As shown in Table 3.8, the translated

	GOLD			Zero-shot			Tr-Train			Tr-Test			Total
	B	Xb	Xl	B	Xb	Xl	B	Xb	Xl	B	Xb	Xl	
OTE False Negatives													
comida	3	3	2	6	2	1	4	1	1	1	9	5	98
restaurante	7	5	7	9	5	6	7	6	6	7	12	10	43
servicio	2	2	2	2	1	1	2	0	1	1	1	2	85
trato	1	1	0	5	6	1	14	10	5	6	8	6	19
atención	2	3	3	8	2	3	7	1	3	7	7	7	13
lugar	0	0	0	2	0	0	1	0	0	0	1	0	12
sitio	1	0	0	5	1	1	3	3	3	2	1	1	14
NER False Negatives													
de	32	29	33	45	51	90	233	252	264	148	146	167	450
la	4	5	3	10	12	16	63	62	62	45	44	45	174
Gobierno	0	0	0	17	53	64	72	70	75	30	45	67	80
Estado	0	0	0	4	4	8	9	8	9	6	6	8	10
Administración	0	0	0	4	8	11	10	11	11	5	5	7	11
Economía	0	0	0	2	6	2	7	8	8	5	6	8	8
Plan	0	0	0	1	2	2	3	5	5	1	4	7	8
Junta	0	0	0	0	0	0	4	10	8	2	3	5	24
Hacienda	0	0	0	1	3	0	4	4	4	4	3	4	5
NER False Positives													
español	0	0	0	16	16	2	16	16	12	13	14	15	0
catalán	0	0	0	8	8	5	7	7	8	8	8	8	0

Table 3.9 – Most common false negatives and positives were there is a big mismatch between methods and the total number of labelled appearances of the word in the test data. B is the acronym for mBERT, Xb for XLM-R base and Xl for XLM-R large.

dataset does not contain any occurrences of “trato” and “atención”, which often appear in the gold-labeled Spanish test data. In fact, both the zero-shot and the data-based cross-lingual transfer approaches fail to correctly label these words as demonstrated in Table 3.9. Interestingly, the zero-shot approach using XLM-R large correctly classifies “trato” (only failing to label 1 of the 19 occurrences).

In the case of the Translate-Train approach, as we use a model of increased capacity, the number of false negatives decreases. This shows that the issue is more pronounced with low-capacity models that have lower generalization capabilities, while higher-capacity models are more successful at labeling words unseen in the training data that share a similar meaning with those in the training data.

The Translate-Test approach does not overcome this issue, as when translating data from the target to the source language, the Machine Translation model does not systematically translate “trato” and “atención” as “service.” Instead, these words are usually translated as “treatment” and “attention”, respectively. In this case, the Machine Translation system fails to understand the context in which the word is used, opting for literal translations. “Treatment” and “attention” are not commonly used in the restaurant review domain in the same context as “service.” In fact, there is no occurrence of the word “treatment” in the English gold-labeled data and there are only two occurrences of the word “attention”, which contrasts with the 153 occurrences of “service.” Once again, the model struggles with labeling these words.

There are more examples of many-to-one translations, such as the word “place”, which in Spanish can be most frequently translated as “lugar” or “sitio.” However, DeepL almost always translates it as “lugar”, resulting in “sitio” being absent in the automatically generated training data, despite being more frequent than “lugar” in the gold-labeled data. In this particular case, the Translate-Test approach is not subject to the problem as both “lugar” and “sitio” are translated into “place” in the Spanish-to-English translation direction.

Cultural alignment: There is a group of words related to Spanish Government names which are not commonly used in the same contexts in English, constituting a significant portion of the false negatives in the NER datasets listed in Table 3.9. For example, the word “Economía” refers to the “Ministry of Economy” or “Ministerio de Economía” in Spanish. “Junta” denotes a “local government” administering a specific region in Spain. “Plan” is often used to denote specific “projects founded by the government”. While these words frequently appear in the Spanish data as part of named entities, this is not the case in the English data, where it is

more customary to use terms like “Treasury Department” (or variations thereof), correctly translated into Spanish by DeepL as “Departamento del Tesoro”. Thus, during fine-tuning on the translated data, the model does not receive any signal to learn that “Economy” may be part of an organization entity. For other terms like “Plan” and “Junta”, the English gold-labelled data lacks references, as these entities are specific to Spanish administration. While the data transfer approach misses many of these named entities, the zero-shot approach is more successful. Training the model with translated data or performing inference with data from other languages translated into English seems to complicate the transfer between different domains. In any case, there is still a significant margin for improvement. The gold standard, trained with gold-labelled data in Spanish, correctly labelled all these entities.

Errors induced by incorrect or missing alignments: We found that for NER, articles and prepositions (i.e. “de”, “la”) are among the words with higher false positive rate for the Translate-Test and Translate-Train approaches. Examining the annotation projection reveals that word aligners struggle to correctly align articles in complex multi-word named entities specially when a one-to-many or many-to-one alignment is required. For example, the word aligners we tested failed to correctly align “of” with “de la” in the following example: “Consejo General de la Arquitectura Técnica de España” (General Council of Technical Architecture of Spain).

Errors induced by annotation inconsistencies: Finally, another issue is the differences across the original gold-labelled annotations. “Gobierno” (Government) and “Estado” (State) are labelled as organizations in the Spanish gold-labelled data, but they are not considered to be entities in the English gold-labelled data. The opposite occurs with demonym words. They are labelled as miscellaneous entities in the English data but in Spanish they are not annotated. Cross-lingual models are likely to fail labeling these cases.

Summarizing, we observe that using English gold-labeled data often produces a signal which, due to inherent differences in language usage, differs from the signal received when using gold-labelled data in the target language. This accounts for the substantial performance disparity between all the cross-lingual sequence labeling approaches and the gold standard trained with gold-labelled data in the target language. The zero-shot cross-lingual transfer approach, when employing high-capacity language models, achieves the best transfer performance across

languages. This appears to be due to this method being less impacted by issues related to many-to-one translations and cultural alignment. These issues together with miss-alignments seems to be the most common reason for the larger number of false positive and negatives of the data-based cross-lingual transfer method with respect to the zero-shot technique.

3.6 Conclusions

In this chapter we have performed an in-depth and comprehensive evaluation of model-based and data-based zero-shot cross-lingual sequence labeling on two different tasks. A detailed error analysis demonstrates that cross-lingual transfer is hindered by the differences in the cultural behaviour of the source and target language in use. The issue is more pronounced with low-capacity models that have lower generalization capabilities, whereas higher-capacity models are more successful at labeling words unseen in the training data that share a similar meaning with those in the training data. This suggests that developing models with enhanced generalization capabilities could bridge this cultural gap.

Our findings indicate that the zero-shot transfer approach is the most effective method when using high-capacity multilingual language models like XLM-R large. However, data-based cross-lingual transfer approaches remain valuable for models with weaker downstream cross-lingual performance. As there is no universally available multilingual pretrained model for every language and domain, data-based methods retain their relevance. Despite the availability of high-quality machine translations and annotation projections, we have identified issues like many-to-one translations or misalignments, which seem to account for the lower performance of data-based cross-lingual transfer methods compared to the model-based approach.

Our results establish that there is still room for improving the cross-lingual performance of zero-shot sequence labeling. In the following chapters, we will focus on two areas. First, we aim to develop a more effective annotation projection method that enables data-based approaches to achieve comparable or superior performance to the model-based approach. Second, we will work on improving the generalization capabilities of sequence labeling models to address the cultural alignment and many-to-one translation issues identified in this study.

4. CHAPTER

Improving Data Transfer

In this chapter we will introduce a novel approach to annotation projection that is based on large pre-trained text-to-text language models and state-of-the-art Machine Translation technology. Our algorithm, named T-Projection, significantly outperforms previous methods of annotation projection by a wide margin. Thanks to this new approach, we have achieved the best results to date for zero-shot cross-lingual transfer between English and eight different low-resource African languages.

4.1 Motivation and contributions

In the previous chapter we demonstrated that model-based zero-shot cross-lingual transfer outperforms data-based approaches. However, model-based zero-shot transfer requires a pre-trained model with high proficiency in both the language and domain of application. Such models are not available for every language and domain, which is the case for the eight different low-resource African languages that will be the subject of study in this chapter. Therefore, data-transfer methods still hold value. The inferior performance of data-transfer methods compared with model-transfer in the previous chapter can be attributed to two main factors. First, Machine Translation of English gold-labeled data often produces a signal which, due to inherent differences in language usage, differs from the signal received when using gold-labeled data in the target language. While both model-transfer and data-transfer approaches are sensitive to this issue, it has a greater impact on

the data-transfer approach. This phenomenon can be mitigated by improving the generalization capabilities of sequence labeling models. We will focus on improving generalization in the following chapters. The second factor involves the mis-alignments produced by current annotation projection methods, which use word alignment algorithms. In this chapter, we will focus on this second issue.

The majority of previous published work on annotation projection explore the application of different word-alignments. However, as demonstrated in Chapter 3, word alignments often produce partial, incorrect or missing annotation projections. This is because word alignments are computed on a word-by-word basis by leveraging word co-occurrences or similarity between vector representations. It should be noted that this method does not take into consideration the labeled spans or categories to be projected.

In this chapter we present T-PROJECTION, a novel approach to automatically project sequence labelling across languages. We split the annotation projection into two steps. First, we use mT5 (Xue *et al.* 2021) text-to-text model to generate a set of projection candidates in the target sentence for each labeled category in the source sentence. This step exploits the labeled categories and the cross-lingual capabilities of large pre-trained multilingual language models. Second, we rank the candidates based on the probability of being generated as a translation of the source spans. We use the M2M100 (Fan *et al.* 2021) and NLLB200 (Costa-jussà *et al.* 2022) state-of-the-art MT models to compute the translation probabilities (Vamvas and Sennrich 2022).

The main contributions of this chapter are the following:

- We have developed a new annotation projection method, T-Projection. We compare the label projections generated by various systems with manually projected annotations on three different tasks, Opinion Target Extraction (OTE), Named Entity Recognition (NER) and Argument Mining (AM), and five different target languages (French, German, Italian, Russian and Spanish). On average, T-Projection improves the current state-of-the-art annotation projection methods by more than 8 points in F1 score.
- We performed a real-world NER task evaluation involving eight low-resource African languages. In this downstream evaluation, T-Projection outperforms other annotation projection methods by 3.6 points in F1 score.

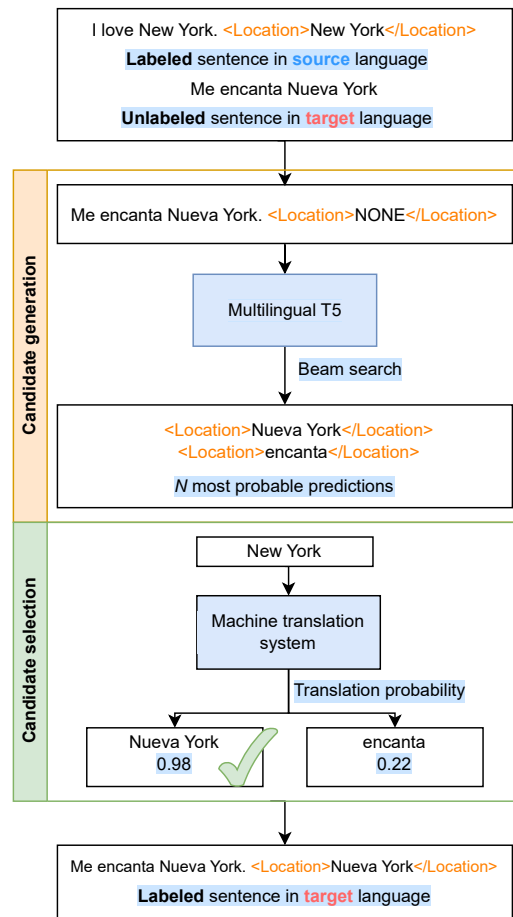


Figure 4.1 – T-Projection two-step method to project sequence labeling annotations across languages.

4.2 T-Projection

T-Projection assumes that we have a set of source sentences with sequences of words labeled with a category. Additionally, there is a parallel version of these sentences in a target language, though these translations are not labeled. T-Projection addresses the challenge of transferring the labels from the source sentences to the target sentences.

T-Projection implements two main steps, which are illustrated in Figure 4.1. First, a set of projection candidates in the target sentence are generated for each labeled sequence in the source sentence. Second, each projection candidate is




Figure 4.2 – Illustration of the candidate generation step. For each label, we generate a set of probable candidates.

ranked using a Machine Translation model. More specifically, candidates are scored based on the probability of being generated as a translation of the source-labeled sequences.

While the *candidate generation* step exploits the labeled spans and their categories in the source sentence as well as the zero-shot cross-lingual capabilities of large pre-trained multilingual language models, the *candidate selection* step applies state-of-the-art MT technology to find those projection candidates that constitute the best translation for each source labeled span. These two steps are described in detail in the following two subsections.

4.2.1 Candidate Generation

When projecting labeled sequences from a source dataset into its parallel target dataset, we expect both the source and the target to contain the same number of sequences, each labeled with the same category. For example, consider the English source sentence “<Person>Obama</Person> went to <Location>New York</Location>” and its parallel, unlabeled Spanish target sentence “Obama fue a Nueva York” We would expect the target sentence also to identify the same two entities (person and location). We propose a projection candidate generation step based on this premise.

We finetune the text-to-text mT5 Xue *et al.* 2021 model using a HTML-tag-

style prompt template (Huang *et al.* 2022). As illustrated by Figure 4.2, we build the inputs for the model by concatenating the unlabeled sentence followed by a list of tags (“<Category>None</Category>”) with the category of each labeled span that we expect to find in the sentence, and the value “None”. If two or more spans share the same category, we append the tag as many times as there are expected spans in that category.

Unlike the approach of Huang *et al.* 2022, we do not encode the tags for each category as special tokens. Instead, we verbalize the categories (i.e., PER → Person) and use the token representations already existing in the model. We anticipate that thanks to the language modeling pretraining, T5 will have a good semantic representation of sequence labeling types such as Person, Location, Claim, etc.

As Figure 4.2 illustrates, we fine-tune mT5 with the labeled source dataset. We train the model to replace the token *None* with the sequence of words in the input sentence that corresponds to that category.

At inference, we label the target sentences, which are parallel translations of the source sentences. As previously explained, we expect to identify the same number of labeled spans and categories as in the source sentence. Therefore, we use the labels from the corresponding source sentence to construct the prompts. In other words, our goal is to label parallel translations of the sentences used for training. We leverage the zero-shot cross-lingual capabilities of mT5 to project the labels from the source to the target sentence. The output tokens are generated in an autoregressive manner. We employ beam search decoding with 100 beams to generate 100 candidates for each input tag. The decision to generate 100 candidates was informed by a preliminary analysis of the performance of the candidate generation stage. Experiments in Section 4.5.3 demonstrate that this number was an overestimation, and generating 10 to 25 candidates is optimal.

4.3 Candidate Selection

In the previous step, we generated up to 100 candidate projections for each labeled span in the source sentence. In the candidate selection step, our goal is to identify the best projection candidate in the target sentence for each labeled span in the source sentence. As depicted in Figure 4.3, all generated candidates are first grouped by category. For instance, if the previous step produced multiple spans with the same category (e.g., two *locations* in a sentence), all such candidates are included in a single set. Additionally, candidates that are not subsequences of the input sentence are filtered out.

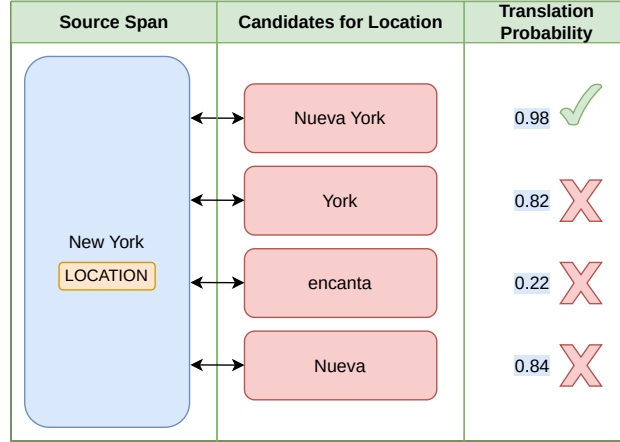


Figure 4.3 – Candidate selection: candidates are scored based on the probability of being generated as a translation of the source labeled sequences.

For each labeled span in the source sentence, we rank all the projection candidates that share the same category as the source span using their translation probabilities (also known as translation equivalence). These probabilities are obtained by applying the pretrained M2M100 (Fan *et al.* 2021) or NLLB200 (Costa-jussà *et al.* 2022) MT models and the *NMTScore* library¹ (Vamvas and Sennrich 2022). Given the source span **A** and the candidate **B**, the translation probability is computed as follows (Vamvas and Sennrich 2022):

$$p_{\theta_a}(A | B) := \left[\prod_{i=0}^{|A|} p_{\theta_a}(A^i | B, A^{<i}) \right]^{\frac{1}{|A|}}$$

The translation probability is normalized:

$$\text{sim}(A | B) = \frac{p_{\theta_a}(A|B)}{p_{\theta_a}(A|A)}$$

Since translation probability can vary depending on the translation direction, the scores are symmetrized by calculating the scores for both translation directions and averaging them:

$$\text{sim}(A, B) = \frac{1}{2} \text{sim}(A | B) + \frac{1}{2} \text{sim}(B | A)$$

¹<https://github.com/ZurichNLP/nmtscore>

Finally, for each labeled span in the source sentence, we select the candidate in the target sentence with the highest translation probability. Once a candidate has been chosen, that candidate and any others that overlap with it are removed from the set of possible candidates. This prevents the assignment of the same candidate in the target sentence to multiple spans in the source sentence.

4.4 Experimental Setup

To evaluate our method we perform both intrinsic and extrinsic evaluations.

Intrinsic Evaluation: We selected several datasets that have been manually projected from English into various target languages. These manual annotations serve as the gold standard for evaluating and comparing T-Projection against previous state-of-the-art label projection models. Results are reported using the F1-score, a standard metric for sequence labeling (Sang 2002). The intrinsic evaluation focuses on measuring the annotation projection accuracy of the models, isolated from other factors such as the quality of the translation models or any other steps in the pipeline.

Extrinsic evaluation: In this evaluation we assess the capability of T-Projection to automatically generate training data for sequence labeling tasks, NER in this particular case. The process begins by utilizing the Machine Translation system NLLB200 (Costa-jussà *et al.* 2022) to translate data from English into 8 low-resource African languages. We then project the labels from English onto the respective target languages. The automatically generated datasets are then employed to train NER models, which are evaluated using a relatively small manually annotated test set. The same procedure is performed with other state-of-the-art label projection models. The comparison of the results obtained is reported in terms of F1-score.

4.4.1 Datasets

The datasets used correspond to three sequence labeling tasks which are illustrated by Figure 4.4. The number of examples and labels are listed in Table 4.1.

Opinion Target Extraction (OTE) Given a review, the task of Opinion Target Extraction (OTE) is to identify the linguistic expressions that refer to the reviewed

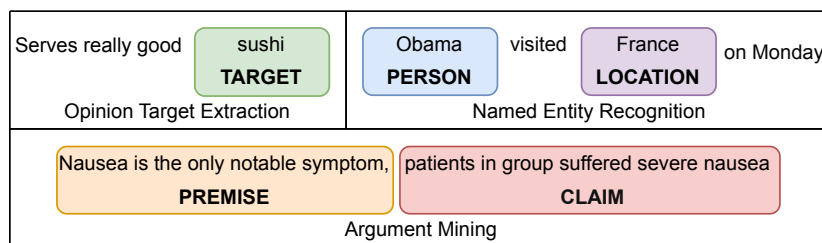


Figure 4.4 – Sequence labeling tasks in our experiments

entity. For instance, in the sentence *Serves really good sushi*, the word *sushi* is the opinion target because it is the entity being discussed. We utilize the English SemEval 2016 Aspect Based Sentiment Analysis (ABSA) datasets (Pontiki *et al.* 2016). This dataset contains user reviews from the Restaurant domain. Additionally, for evaluation purposes, we employ parallel datasets in Spanish, French, and Russian, which were generated through Machine Translation and manual label projection in Chapter 3.

Named Entity Recognition (NER) The NER task involves detecting named entities and classifying them according to predefined categories. We use a parallel NER dataset in English, Spanish, German, and Italian (Agerri *et al.* 2018), based on Europarl data (transcriptions of discussions from the European Parliament) (Koehn 2005). These transcriptions are parallel in multiple languages and were annotated following the CoNLL 2003 guidelines (Sang and Meulder 2003). For the extrinsic evaluation, we use MasakhaNER 2.0 (Adelani *et al.* 2022), a human-annotated NER dataset for 20 African languages.

Argument Mining (AM) The AbstRCT English dataset includes annotations for two types of argument components, Claims and Premises, in medical and scientific texts collected from the MEDLINE database (Mayer *et al.* 2020). A *Claim* is a concluding statement made by the author about the study’s outcome, such as an assertion of a diagnosis or a treatment in the medical domain. A *Premise* is an observation or measurement (ground truth) that supports or challenges another argument component, usually a claim. Premises are considered observed facts and are credible without further evidence. For evaluation, we used the Spanish parallel counterpart, generated following an adapted version of the method described in Chapter 3. The labeled sequences in the AM task consist of very long spans of words, frequently encompassing full sentences. We use the Neoplasm split.

Task	Split	Sentence No	Labels
ABSA			
ABSA (Pontiki <i>et al.</i> 2016)	Train	2000	(1) Target
ABSA (Pontiki <i>et al.</i> 2016)	Test	676	
NER			
Europarl (Agerri <i>et al.</i> 2018)	Test	799	(3) Person, Location, Organization,
CoNLL03 (Sang and Meulder 2003)	Train	14987	
CoNLL03 (Sang and Meulder 2003)	Dev	3466	
CoNLL03 (Sang and Meulder 2003)	Test	3684	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (hau)	1632	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (ibo)	2180	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (sna)	1772	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (swa)	1882	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (xho)	1632	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (yor)	1963	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (nya)	1784	
MasakhaNER2.0 (Adelani <i>et al.</i> 2022)	Test (zul)	1669	
AM			
AbsRCT Neoplasm (Mayer <i>et al.</i> 2020)	Train	4404	(2) Claim, Premise
AbsRCT Neoplasm (Mayer <i>et al.</i> 2020)	Dev	679	
AbsRCT Neoplasm (Mayer <i>et al.</i> 2020)	Test	1251	

Table 4.1 – Size (Number of sentences) of the dataset we use to train and evaluate our systems.

4.4.2 Baselines

We use the same alignment systems and methodology described in Chapter 3 as a baseline. This approach is illustrated in Figure 4.5. We compare T-Projection with two statistical systems, **Giza++** (Och and Ney 2003) and **FastAlign** (Dyer *et al.* 2013b). These systems are widely used in the field and require very small computational resources. Additionally, we evaluate two current state-of-the-art Transformer-based word-alignment systems, **SimAlign** (Jalili Sabet *et al.* 2020) and **AWESOME** (Dou and Neubig 2021), which leverage pre-trained multilingual

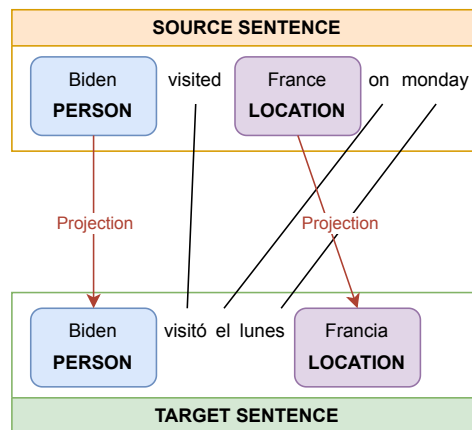


Figure 4.5 – Illustration of the translation and annotation projection task using word-alignments.

language models to generate alignments. As recommended by the authors, we use multilingual BERT (mBERT) (Devlin *et al.* 2019) as the backbone. We tested different models as backbones but observed no improvement in performance.




Figure 4.6 – Illustration of the translation with markers approach. Markers are introduced around the labeled sequences. The sentence and the labeled spans are translated together.

We also experiment with **EasyProject** (Chen *et al.* 2023), a system that jointly performs translation and projection by inserting special markers around the labeled spans in the source sentence as depicted in Figure 4.6. Additionally, we evaluate **CODEC** (Le *et al.* 2024), a subsequent work that divides the projection process into two distinct steps. In this improved approach, the training data in the high-resource language is first translated without markers. During a second decoding phase, the markers are integrated with the constraint that the translation must align with the initial, marker-free output. This two-step process ensures that the final translated sentence with markers remains consistent with what the model would have produced without them, thereby preserving the translation quality. As both these methods generate their own translations they are therefore not suitable for the intrinsic evaluation which is why we only used them for the extrinsic evaluation.




Figure 4.7 – Illustration of the span translation annotation projection approach. The source labels are translated independently, and these translated spans are then matched with their counterparts in the target sentence.

In addition to the previous methods, we also implement two additional baselines inspired by previous works. In the first baseline, inspired by Li *et al.* 2021, we use the **XLM-RoBERTa** model (Conneau *et al.* 2020) with 3 billion parameters (matching the parameter count of the mT5 model used in T-Projection) and

add a token classification layer (linear layer) on top of each token representation. We train the model on the source labeled dataset and use it to predict entities in the translated target sentences. The second baseline adopts a **span translation** approach inspired by Jain *et al.* 2019 and Zhou *et al.* 2022b. In this approach, we translate the labeled spans in the source sentence using the pretrained M2M100 model with 12 billion parameters and then match them with the corresponding spans in the target sentence. For example, as depicted in Figure 4.7 given the labeled source sentence “<Person> Bruce Willis </person> was born in <Location> West Germany </Location>” and the target sentence “Bruce Willis nació en Alemania Occidental”, we translate the span *West Germany* into the target language, resulting in *Alemania Occidental*, which is then matched in the target sentence. We employ beam search to generate 100 possible translations and select the most probable one that matches the target sentence.

4.4.3 Models Setup

We use the 3 billion parameters pretrained mT5 (Xue *et al.* 2021) for the *candidate generation* step while *candidates are selected* using the M2M100 12 billion parameter Machine Translation model (Fan *et al.* 2021). In the case of MasakhaNER, since not all languages are included in M2M100, we resorted to NLLB200 (Costa-jussà *et al.* 2022) 3 billion parameter model instead, which was also used by the EasyProject method (Chen *et al.* 2023). Both MT models demonstrate comparable performance.

We train the HuggingFace’s (Wolf *et al.* 2019) implementation of mT5² (3 billion parameter model) in the candidate generation step using the following hyper-parameters: Batch size of 8, 0.0001 learning rate, 256 tokens sequence length, cosine scheduler with 500 warm up steps and no weight decay. We use AdaFactor (Shazeer and Stern 2018) optimizer. We train the model for 10 epochs in the OTE task, and 4 epochs for the NER and AM tasks. In the candidate selection step, we also use HuggingFace’s implementation of M2M100, and we use m2m100-12B-last-ckpt³ checkpoint of M2M100 released by the authors. We use the direct-translation function of the NMTscore library to compute the translation probabilities. For MasakhaNER2.0 we use the training script and evaluation script developed by the authors⁴ and the same hyper-parameter setup than Chen *et al.*

²<https://huggingface.co/google/mt5-xl>

³<https://huggingface.co/facebook/m2m100-12B-last-ckpt>

⁴<https://github.com/masakhane-io/masakhane-ner/blob/main/MasakhaNER2.0/scripts/mdeberta.sh>

2023.

4.5 Intrinsic Evaluation

In this section, we present a set of experiments to evaluate T-Projection in comparison to current state-of-the-art approaches for annotation projection. The intrinsic evaluation focuses on measuring the annotation projection accuracy of the models, isolated from other factors such as the quality of the translation models or any other steps in the pipeline. We also separately analyze the performance of the candidate generation and candidate selection steps.

For the OTE task, we train T-Projection and XLM-RoBERTa using the full English ABSA 2016 dataset. Additionally, we train the four word-alignment systems (excluding SimAlign, which is an unsupervised method) using the English data along with the respective translations as parallel corpora. We augment the parallel data with 50,000 random parallel sentences from ParaCrawl v8 (Esplà *et al.* 2019). Models are evaluated based on the manual label projections described in Chapter 3.

As the Europarl-based NER dataset (Agerri *et al.* 2018) provides only test data for each language, T-Projection and XLM-RoBERTa are trained using the full English CoNLL 2003 dataset (Sang and Meulder 2003) together with the labeled English Europarl test data. The word alignment models, in turn, are trained with the parallel sentences from the Europarl-based NER data plus 50,000 parallel sentences extracted from Europarl v8 (Koehn 2005). We evaluate the models based on the manual annotations provided by Agerri *et al.* 2018.

For Argument Mining, we use the full Neoplasms data from the AbstrCT dataset to train T-Projection and XLM-RoBERTa, adding its Spanish translation as parallel corpora for the word alignment systems. As this is a medical text corpus, the parallel corpora are complemented with 50,000 parallel sentences from the WMT19 Biomedical Translation Task (Bawden *et al.* 2019). We evaluate the models based on the manually projected labels by Yeginbergen and Agerri 2024.

4.5.1 Annotation Projection Quality

Table 4.2 presents the results of the automatically projected datasets generated by each projection method, compared to the human-projected versions of those datasets. Systems using word alignments achieve consistently good results, particularly those utilizing language models such as SimAlign and AWESOME. Specif-

	OTE			NER			AM	Avg
	ES	FR	RU	ES	DE	IT	ES	
Giza++ (Och and Ney 2003)	77.0	73.3	72.4	73.3	75.3	68.4	86.6	77.7
FastAlign (Dyer <i>et al.</i> 2013b)	75.0	72.9	76.9	70.2	77.0	67.0	85.7	77.4
SimAlign (Jalili Sabet <i>et al.</i> 2020)	86.7	86.3	87.7	85.4	87.4	81.3	84.1	85.3
AWESOME (Dou and Neubig 2021)	91.5	91.1	93.7	87.3	90.7	83.1	54.8	78.0
XLM-RoBERTa-xl (Conneau <i>et al.</i> 2020)	80.2	76.2	74.5	73.9	68.3	73.9	66.5	71.8
Span Translation	66.5	46.3	58.7	68.8	63.5	69.2	21.6	48.7
T-Projection	95.1	92.3	95.0	93.6	94.0	87.2	96.0	93.9

Table 4.2 – F1 scores for annotation projection in the OTE, NER and Argument Mining tasks.

ically, AWESOME performs well in OTE and NER but poorly in AM. Manual inspection reveals that AWESOME struggles to align articles and prepositions within long sequences. The statistical-based models, Giza++ and FastAlign, achieve competitive performance considering their very low computational resource requirements (they do not require a GPU) compared to the Transformer-based approaches.

XLM-RoBERTa-xl demonstrates strong zero-shot cross-lingual performance. However, the quality of the generated datasets is lower than those produced by the word-alignment systems. The results of the Span Translation approach are disappointing, particularly for the long sequences in the AM task. Translating the labeled spans independently often results in translations that cannot be located in the target sentence.

Our T-Projection method achieves the best results for every task and language. In OTE, it outperforms all other methods by more than 2 points in F1 score averaged across the three languages. This indicates that T-Projection robustly projects labeled spans into machine-translated data. The NER evaluation is slightly different because the parallel data was translated by human experts. In this context, T-Projection significantly improves AWESOME’s results by 4.7 points, marking a substantial improvement in the quality of the generated datasets. Despite the word alignment systems being trained with Europarl domain-specific data and most of the training data for T-Projection coming from the CoNLL-2003 dataset (news domain) plus a few annotated sentences (699) from Europarl, T-Projection still achieves the best results in NER label projection. This suggests that our system is effective even in out-of-domain settings.

Furthermore, T-Projection achieves the highest overall scores in Argument Mining, demonstrating its exceptional ability to project long sequences. T-Projection outperforms the second-best model by 9.4 points in F1 score, with a by 96.0 points in F1-scores, indicating near-perfect projection of all examples in the dataset.

When considering the average performance across the three tasks and five languages, T-Projection improves the F1 score by 8.6 points compared to the second-best system, SimAlign. These results represent a significant advancement over all previous annotation projection approaches. To the best of our knowledge, these are by a wide margin the best annotation projection results published for sequence labeling.

4.5.2 The Role of the Candidates

	OTE			NER			AM	Avg
	ES	FR	RU	ES	DE	IT	ES	
T-Projection	95.1	92.3	95.0	93.6	94.0	87.2	96.0	93.9
Ngrams + Candidate Selection	89.7	86.1	93.8	83.8	79.3	73.3	73.5	80.7
mT5 + Most Probable Candidate	83.7	87.2	85.3	79.5	82.8	72.3	90.9	84.8
mT5 + Upper bound	98.6	97.0	97.9	98.0	98.5	94.0	99.3	98.0

Table 4.3 – F1 scores for different candidate generation and candidate selection methods.

We perform a set of experiments to measure the relevance and performance of the *candidate generation* and *candidate selection* steps. First, we replace mT5 as the candidate generation model with an n-gram-based approach. We extract all the n-grams with sizes ranging from 1 to the sentence length (e.g., “Serve”, “really”, “good”, “sushi”, “Serves really” ... “Serves really good sushi”) and consider them as candidates. Then we rank the candidates using the translation probabilities obtained by the M2M100 model. As shown in Table 4.3, the n-gram-based approach’s performance is significantly lower than T-Projection in all tasks and languages. This indicates that the mT5 model is crucial for generating relevant candidates. The n-gram approach generates a large number of very similar candidates, which makes it difficult for the M2M100 model to select the correct one. These experiments demonstrate the importance of the mT5 model in generating relevant candidates.

We also replace the *candidate selection* method with the *most probable candidate*. That is, we select the most probable candidate generated by mT5 for each labeled span in the source sentence, thus we do not consider the translation probabilities, as only one candidate is generated for each labeled span. This approach achieves competitive results with the word alignment systems in Table 4.2, but it is outperformed by T-Projection by an average of 9.2 points in F1-Score.

This ablation study demonstrates that although each step of T-Projection in isolation is able to achieve competitive results on its own, it is the combination of both steps that allows T-Projection to achieve very high performance in the intrinsic evaluation.

Finally, we define an upper bound for the *candidate selection* step. This upper bound is defined by always selecting the correct candidate among the generated candidates. If the correct candidate is not generated by mT5, we select the most probable candidate. This upper bound achieves an average F1 score of 98.0. This result confirms that with a very high probability, the correct candidate is among the candidates generated by mT5.

4.5.3 How many candidates are necessary?




Figure 4.8 – F1 score when generating a different number of candidates.

Generating candidates is computationally expensive. The number of FLOPs, memory usage, and inference time increase linearly with the number of candi-

dates. Thus, generating 20 candidates is twice as expensive as generating 10 candidates. Additionally, we must consider the extra cost of computing similarity scores for each candidate. Therefore, we performed an experiment to determine the optimal number of candidates to generate. As shown in Figure 4.8, the performance of T-Projection increases with the number of candidates. However, the performance improvement diminishes as the number of candidates grows. For OTE and NER, the improvement is negligible after generating 25 candidates. For AM, generating more than 10 candidates is, in fact, counterproductive. While the reported results in the paper were obtained by generating 100 candidates, a decision informed by preliminary studies using the upper bound described in Section 4.5.2, the results in Figure 4.8 suggest that generating 10 to 25 candidates is optimal.




Figure 4.9 – Number of times the correct candidate is among the top-k candidates generated by mT5.

Generating a larger number of candidates is not beneficial for the performance of T-Projection. To further understand why this happens, we computed the F1 score of T-Projection following the upper bound described in Section 4.5.2 for different number of candidates. This is, we compute how many times the correct candidate is among the top-k candidates generated by mT5. As shown in Figure 4.9, as we increase the number of candidates, the probability of the correct candidate being among the top-k candidates also increases. However, the improvement is not linear. It becomes less significant after generating 25 candidates. The per-

formance improvement from going from 25 to 100 candidates is only 1 point in F1 score for OTE and NER and less than 0.5 points for AM. Interestingly, for AM, the best candidate is among the top-25 candidates 99% of the time, which explains the great performance of T-Projection in this task. However, if the best candidate is found with a very high probability among the top-25 candidates, what are the candidates beyond the top-25? By examining the candidates generated by mT5 in the NER tasks, we found that only a few of the generated candidates are valid. The remaining ones are hallucinated spans that do not exist in the sentence and are therefore filtered out. These hallucinated spans are usually variations of the correct candidate. We found that there are fewer than 20 valid candidates per sentence, with an average of 5.95 valid candidates per sentence for NER. Thus, generating more than 25 candidates is not beneficial, but it also does not introduce noise that severely hinders the performance of T-Projection, as the extra candidates are usually hallucinated spans that are filtered out in the candidate selection step, rather than different n-grams from the target sentence.

4.5.4 Model size and performance

	Model	#Params	OTE	NER	AM	Average
MT Size	m2m100	418M	92.3	91.7	95.5	93.1
	m2m100	1.2B	94.0	92.0	95.8	93.9
	m2m100	12B	94.1	91.6	96.0	93.9
	Prism	745M	94.1	90.9	96.3	92.7
	nllb200	3B	94.2	91.0	93.0	93.0
mT5 size	mT5-small	60M	36.4	66.3	00.0	34.2
	mT5-base	220M	72.8	86.2	33.6	64.2
	mT5-large	738M	90.9	90.1	65.3	82.1
	mT5-xl	3B	94.1	91.6	96.0	93.9

Table 4.4 – F1 scores of T-Projection when using translation and mT5 models of different size

We analyze the performance of T-Projection using an mT5 model and a translation system with different numbers of parameters. Additionally, we evaluate T-Projection with various Machine Translation models, namely, NLLB200 (Costa-jussà *et al.* 2022) and PRISM (Thompson and Post 2020). Table 4.4 demonstrates

that the Machine Translation system and its parameter count do not significantly impact T-Projection’s performance.

However, the size of the mT5 model has a substantial impact on the system’s final performance. While switching from a 3B to a 738M parameter mT5 model yields competitive results for OTE and NER, this is not the case for AM. The overall trend indicates that decreasing the number of parameters results in decreased performance. In summary, to achieve competitive performance across all tasks, T-Projection requires an mT5 model with 3B parameters, although a 738M parameter model remains competitive for OTE and NER.

4.6 Extrinsic Evaluation

Language	No. of Speakers	Language family	Finetune English	AWESOME +English	EasyProject +English	CODEC	T-Projection	T-Projection +English
Hausa	63M	Afro-Asiatic /Chadic	71.7	72.7	72.2	72.4	72.7	72.0
Igbo	27M	NC / Volta-Niger	59.3	63.5	65.6	70.9	71.4	71.6
Chichewa	14M	English-Creole	79.5	75.1	75.3	76.8	77.2	77.8
chiShona	12M	NC / Bantu	35.2	69.5	55.9	72.4	74.9	74.3
Kiswahili	98M	NC / Bantu	87.7	82.4	83.6	83.1	84.5	84.1
isiXhosa	9M	NC / Bantu	24.0	61.7	71.1	70.4	72.3	71.7
Yoruba	42M	NC / Volta-Niger	36.0	38.1	36.8	41.4	42.7	42.1
isiZulu	27M	NC / Bantu	43.9	68.9	73.0	74.8	66.7	64.9
AVG			54.7	66.5	66.7	70.3	70.3	69.8

Table 4.5 – F1 scores on MasakhaNER2.0 for mDebertaV3 trained with projected annotations from different systems. "+EN" denotes concatenation of the automatically generated target language dataset with the source English dataset.

In this section, we evaluate T-Projection in a real-world low-resource scenario, namely, Named Entity Recognition for African languages. We compare the results obtained by training on NER datasets automatically generated by T-Projection with those automatically projected using three state-of-the-art label projection systems: AWESOME (the second-best NER system in Table 4.2), EasyProject, and CODEC. We use the exact same settings as Chen *et al.* 2023 and Le *et al.* 2024. For each target language in MasakhaNER2.0, we first translate the English CoNLL dataset using the NLLB-200 3 billion parameter model. Next, we project the English labels into the target language. It should be noted that EasyProject performs both of these processes in a single step. Subsequently, we train an mDebertaV3 (He *et al.* 2021) model using the automatically generated datasets for each target language. Finally, this model is evaluated on the gold MasakhaNER2.0 test

data. We only evaluate the eight languages in MasakhaNER2.0 supported by mT5. We focus on named entities referring to Person, Location, and Organization. We also evaluate the zero-shot model-transfer approach presented in Chapter 3. That is, we train the mDebertaV3 model with the original English CoNLL data and evaluate it on the MasakhaNER2.0 test sets.

Table 4.5 presents the results of the evaluated models on the gold MasakhaNER 2.0 test sets. For T-Projection, we present the results of training with the automatically generated data for the target language only, and also by adding the original English CoNLL data concatenated with the automatically generated data for each target language. Regarding other systems, we only show the former results, as it was the only metric reported by previous work. In order to train and evaluate the NER models, we apply the same hyperparameter settings and code as the authors of EasyProject.

4.6.1 T-Projection vs other annotation projection systems

The results show that T-Projection achieves superior performance for seven out of the eight languages compared to other annotation projection systems. Our model demonstrates a more pronounced performance difference in agglutinative languages such as Igbo and Shona. As outlined in Section 4.5, our model produces superior alignments compared to AWESOME. Furthermore, we found that EasyProject, which utilizes markers for simultaneous translation and projection, introduces translation artifacts that hinder the performance of the downstream model. These artifacts are particularly noticeable in agglutinative languages, as EasyProject tends to separate words. For instance, in the case of Shona, consider the English sentence “[Germany]’s representative to the [European Union]’s veterinary committee [Werner Zwingmann]”. Our system produces the Shona sentence “Mumiriri [weGermany] kukomiti yemhuka [yeEuropean Union] [Werner Zwingmann]”, while EasyProject produces “Mumiriri we [Germany] ku [European Union] komiti yezvokurapa mhuka [Werner Zwingmann]”. When training mDeberta-v3 with T-Projection’s generated data, which preserves the agglutinated words, we achieve better results compared to EasyProject which introduces artifacts by separating agglutinated words during translation and projection. CODEC improves over EasyProject, the two step approach, in which first the translation is performed and then the markers are added, helps to preserve the translation quality. However, the performance of CODEC is still lower than T-Projection for every language except for Zulu.

4.6.2 T-Projection vs Model-transfer

Table 4.5 presents the results of training the multilingual model mDebertaV3 with the original English CoNLL data and conducting zero-shot evaluation on the MasakhaNER2.0 test sets. The findings indicate that T-Projection outperforms this baseline in six out of eight languages, with an average improvement of 15.6 F1 points. This contrasts with the results from Chapter 3, where model-based transfer learning surpassed data-based approaches. The discrepancy can be attributed to two factors. First, the data generated by T-Projection is of higher quality than the data generated by the word alignment systems used in Chapter 3. On average, T-Projection is 8.6 F1 score points better than the best word-alignment system in intrinsic evaluation and 3.8 F1 score points superior in extrinsic evaluation.

Second, the zero-shot evaluation in Chapter 3 was conducted from English to other high-resource languages, whereas in this chapter, we perform cross-lingual evaluation into African low-resource languages. These low-resource languages have significantly different morphology and syntax compared to English, and multilingual models typically have lower proficiency in these languages. As demonstrated in Chapter 3, model-based transfer performance requires a multilingual model with high proficiency in both the source and target languages.

These results demonstrate that data-based transfer approaches, such as T-Projection, can be highly effective for performing Natural Language Processing tasks in low-resource languages, especially in the absence of a high-proficiency multilingual model.

In contrast to previous work, our experiments revealed that concatenating English and translated data did not yield better results, likely due to the superior quality of the data generated by T-Projection. To the best of our knowledge, these are the best zero-shot results achieved for MasakhaNER2.0, highlighting the significant benefits of T-Projection for NLP tasks in low-resource languages.

4.7 Conclusions

In this section we have introduced T-Projection, a novel method for projecting labeled sequences across languages. T-Projection leverages the zero-shot cross-lingual capabilities of large pretrained multilingual language models to generate candidates for each labeled span in the source sentence. These candidates are then ranked using a Machine Translation model to select the best projection candidate for each labeled span in the target sentence.

We have demonstrated that T-Projection outperforms current state-of-the-art label projection systems in both intrinsic and extrinsic evaluations. In the intrinsic evaluation, T-Projection achieves the best results for every task and language, improving the F1 score by 8.6 points compared to the second-best system. In the extrinsic evaluation, T-Projection achieves superior performance in seven out of the eight languages in the MasakhaNER2.0 dataset. These results underscore the effectiveness of T-Projection in generating high-quality training data for sequence labeling tasks in low-resource languages.

Moreover, T-Projection surpasses model-based cross-lingual transfer in the extrinsic evaluation, specifically for named entity recognition in African low-resource languages. While model-based transfer learning outperformed data-based approaches in Chapter 3, it is less effective for cross-lingual transfer from English into the African languages tested in this chapter. In this scenario, T-Projection generates high-quality training data that significantly improves the performance of the downstream model. This demonstrates the potential of data-based transfer approaches for NLP tasks in low-resource languages.

5. CHAPTER

Improving Model Transfer

In this chapter we will focus on zero-shot model transfer. In Chapter 3 we demonstrated that model transfer could be an effective and efficient approach for cross-lingual transfer when using a high-capacity model on the target language. Therefore, for this approach to be effective, it is crucial to use the most powerful models available. Currently, these models are the text-to-text Large Language Models (LLMs). However, using LLMs for zero-shot cross-lingual sequence labeling is not straightforward. In this chapter we will introduce a constrained decoding algorithm that effectively addresses this issue. A comprehensive empirical evaluation across multiple tasks and languages demonstrates that, when our method is applied to an LLM, it helps not only to improve over the unconstrained beam search baseline but also to outperform the zero-shot cross-lingual capabilities of encoder-only models, especially for languages that significantly differ from English.

5.1 Motivation and contributions

In Chapter 3, we demonstrated that the performance of zero-shot cross-lingual transfer can be significantly enhanced by using a high-capacity model for the target language. However, we only employed encoder-only models such as XLM-RoBERTa-large (Conneau *et al.* 2020), which has 561 million parameters and was trained on approximately 295 billion tokens. However, as mentioned in Chapter 2.1, the most powerful models currently available are text-to-text Large Language Models (LLMs) like T5 (Raffel *et al.* 2020), LLaMA (AI@Meta 2024), and GPT-

4 (OpenAI *et al.* 2024). These models have demonstrated superior capabilities in a wide range of NLP tasks, including the ability to solve tasks for which they were not explicitly trained. Consequently, efforts to scale NLP models have primarily focused on text generation models. As shown in Table 5.1, the latest generation of LLMs have significantly more parameters and were trained on much larger datasets compared to XLM-RoBERTa.

	XLM-RoBERTa Conneau <i>et al.</i> 2020	XLM-RoBERTa-xxl Goyal <i>et al.</i> 2021	mT5 Xue <i>et al.</i> 2021	Llama2 Touvron <i>et al.</i> 2023b	Gemma2 Mesnard <i>et al.</i> 2024	LLama3 AI@Meta 2024
Parameters	560M	10.7B	11.3B	70B	27B	405B
Train Tokens	296B	296B	1T	2T	8T	17T

Table 5.1 – Size and training data of some relevant open source models.

LLM models have already been proven effective for Information Extraction and sequence labeling tasks in monolingual evaluations in English (Ashok and Lipton 2023; Sainz *et al.* 2024a). However, their performance still lags behind encoder-only models in multilingual sequence labeling (Fetahu *et al.* 2023). For low-resource languages, such as African languages, Ojo and Ogueji 2023 demonstrated that most text-to-text LLMs for named entity recognition do not perform well at all when evaluated in a zero-shot setting. Their results are reproduced in Table 5.2. The table shows that the performance of LLMs is significantly lower than that of XLM-RoBERTa-large. This is a surprising result, given that many of the LLMs have been trained on much larger multilingual datasets and have significantly more parameters than XLM-RoBERTa-large. Thus, the question arises: why do LLMs perform poorly in zero-shot cross-lingual sequence labeling tasks?

Model	Size	amh	bam	bbj	ewe	hau	ibo	kin	lug	luo	mos	nya	pcm	sna	swa	tsn	twi	wol	xho	yor	zul
Fine-tune: SotA																					
AfroXLMR-large	550M	78.0	79.0	90.3	75.2	85.4	88.9	86.8	88.9	75.3	73.5	92.4	90.0	96.1	92.7	88.9	79.2	83.8	89.2	67.9	90.6
Prompting of LLMs																					
GPT-4	-	28.5	52.7	50.3	75.6	64.9	56.0	55.1	73.3	49.8	60.2	63.6	64.7	33.4	71.5	64.6	58.6	67.9	28.4	58.3	34.9
AYA	-	14.1	7.1	20.0	26.5	34.5	28.2	30.8	16.3	12.7	34.4	21.7	27.4	13.4	35.6	29.4	18.9	14.5	4.2	17.5	11.4
mT0	13B	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
mT0-MT	13B	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaMa 2	13B	0.0	13.8	12.3	25.1	22.1	22.0	23.1	27.5	19.0	11.0	20.0	27.5	11.3	25.8	26.2	20.7	16.0	8.1	15.1	9.0

Table 5.2 – Comparison of F1-score of various LLMs with that of the current state of the art result in Masakhaner 2.0. Table reproduced from Ojo and Ogueji 2023.

In this chapter, we investigate the performance of LLMs in zero-shot cross-lingual sequence labeling tasks. Our contributions are as follows:

We identify the challenges faced by text-to-text models for zero-shot sequence labeling: In this setting we must first establish a text-based input and output representation for the specific task. However, current text-to-text models are tailored for generating free-form text. As our experiments demonstrate, models fail to strictly adhere to the output structure. Moreover, as demonstrated by our experiments, text-to-text models often produce outputs mixing the source language and the target language, which compromises their performance. These issues are illustrated in Figure 5.1, where the incorrect output mixes English and Basque (Turkiako-Turkish) and incorrectly segments the organization entity “Realean”.




Figure 5.1 – Comparison between a valid (top green) and invalid (bottom red) output structure to represent a Named Entity Recognition task. English translation: (They) played in Real and in the Turkish national team.

We propose a constrained decoding algorithm for text-to-text models: We introduce a constrained decoding algorithm that enforces the output structure of the target task. Our method can be seamlessly integrated with any text-to-text model without any significant increase in the decoding cost. Although constrained generation has been previously explored in a monolingual setting (Liu *et al.* 2022), we adapt and extend this approach for zero-shot cross-lingual IE. Our new decoding algorithm is evaluated on three popular IE tasks for 25 languages of varied morphological characteristics. Empirical results indicate that our method, when applied to an LLM such as mT0-XL (Muennighoff *et al.* 2023), not only surpasses the unconstrained beam search baseline but also outperforms the zero-shot cross-lingual performance of encoder-only models. Our method is especially successful for languages that significantly differ from English.

To the best of our knowledge, our new technique achieves the best zero-shot model-based cross-lingual transfer results to date.

5.2 Related Work

In this section, we focus first on related work concerning large language models (LLMs) for sequence labeling. In the second part of the related work, we review prior research on constrained decoding.

5.2.1 LLMs for sequence labeling

The introduction of models like T5 (Raffel *et al.* 2020) and GPT (Radford *et al.* 2019) revolutionized NLP by adopting a text-to-text approach, enabling models to handle a wide array of tasks with a single training objective. Consequently, all NLP tasks can be framed as text-to-text tasks, where the input is a description of the task or a prompt, and the output is the desired result (Chung *et al.* 2022). Scaling these models in both the amount of training data and the number of parameters, has led to the development of state-of-the-art models, such as GPT-4 (OpenAI *et al.* 2024), LLaMA (AI@Meta 2024), and Mistral (Jiang *et al.* 2023). These models achieve state-of-the-art results on a broad range of NLP tasks (Min *et al.* 2024), including those they were not explicitly trained for (Radford *et al.* 2019).

In the field of Information Extraction (IE), the text-to-text approach has also been explored. Lu *et al.* 2022 introduced a unified text-to-structure generation model capable of handling various IE tasks universally. Lou *et al.* 2023 proposed converting IE tasks into a semantic matching problem, allowing their method to generalize to new domains and label ontologies not encountered during training. Wang *et al.* 2023b framed IE tasks as natural language descriptive instructions and trained a large language model (LLM) across a diverse range of IE tasks. In evaluations involving tasks with unseen label ontologies, their model outperformed other instruction-tuning methods. More recently, Blevins *et al.* 2023 and Sainz *et al.* 2024a proposed using complex instructions that include annotation guidelines, similar to the ones used by human annotators, to enhance the performance of LLMs in sequence labeling tasks. This approach has proven effective in achieving strong performance in classifying unseen categories in sequence labeling tasks in English.

While success has been achieved in labeling unseen categories in English, the supervised performance of LLMs, when training data is available, is still not superior to that of smaller encoder-only models (Sainz *et al.* 2024a). Additionally, recent shared tasks (Fetahu *et al.*, 2023) have shown that for languages other than

English, encoder-only language models such as XLM-RoBERTa (Conneau *et al.* 2020) and mDEBERTA (He *et al.* 2021) remain the most effective models.

5.2.2 Constrained decoding

The formulation of information extraction tasks in a constrained text-to-text format has been previously explored (Vinyals *et al.* 2015; Xiao *et al.* 2016; Dyer *et al.* 2016). However, it was with the emergence of large-scale text-to-text language models, that this approach garnered significant attention within the community. Lester *et al.* 2020 propose a Named Entity Recognition system that uses Viterbi decoding (Forney 1973) with heuristically determined transition probabilities that prohibit illegal transitions. This achieves similar performance to the conditional random field (CRF) models (Lafferty *et al.* 2001), but it is more computationally efficient. Cao *et al.* 2021 and De Cao *et al.* 2022 propose a sequence-to-sequence system for Multilingual Entity Linking, which can generate entity names from left to right, token by token, in an autoregressive manner, conditioned by the context. To ensure that only valid entity identifiers are generated, they employ a prefix tree to enable constrained beam search.

Closer to our work, which focuses on constraining large language models (LLMs) to adhere to a pre-defined output structure, Lu *et al.* 2021 presents a constrained decoding algorithm that ensures the model adheres to a specified output structure during inference. Similarly, Zheng *et al.* 2023 and He and Choi 2023 propose constrained decoding algorithms that enhance semantic parsing. Instead of constraining the generation of output text, Cui *et al.* 2021 perform Named Entity Recognition (NER) by computing the probability of a text span filling predefined structures. Rather than flattening the structured output into a sequence, Liu *et al.* 2022 model the output as sequences of actions. These actions are predicted in an autoregressive manner using LLMs, and executing the actions generates the structured output. Their approach improves upon previous methods in NER, end-to-end relation extraction, and co-reference resolution. Geng *et al.* 2023 demonstrate that grammar-constrained decoding (GDC) can significantly enhance the performance of large language models (LMs) across a variety of structured NLP tasks, such as information extraction, entity disambiguation, and constituency parsing, by ensuring outputs adhere to a given structure. GCD-enhanced LMs outperform both unconstrained LMs and task-specific finetuned models, particularly in scenarios with limited training data.

Although previous research has demonstrated the effectiveness of constrained decoding for information extraction, most of it has focused on monolingual set-

tings. Thus, Guo and Roth 2021 propose an algorithm that employs constrained decoding of text-to-text LLMs for zero-shot NER in low-resource languages. First, they translate labeled data in a word-by-word manner using a dictionary. Then, they construct target language text from the source-language named entities using a pretrained language model. They utilize constrained decoding to ensure the presence of entities in the generated text. This data-transfer method was later surpassed by model-based cross-lingual transfer methods as we demonstrate in Chapter 3.

To project labels across languages in sequence labeling tasks, Le *et al.* 2024 introduce markers to the input text to represent the labeled sequences. They then translate the text into the target language, achieving both translation and annotation projection. To prevent translation artifacts caused by the markers, they propose a constrained decoding algorithm that ensures the output of the translation when markers are introduced, remains consistent with translations without markers. Although this method is effective, it is a data-based approach that requires training a new model on the projected data.

5.3 Approach

In this section, we describe our representation of a Sequence Labelling task by applying our new Constrained text-to-text approach. Our algorithm can be used for both encoder-decoder (Vaswani *et al.* 2017) and decoder-only (Liu *et al.* 2018) architectures, as well as any other auto-regressive architecture.

5.3.1 Input-Output Representation




Figure 5.2 – Text-to-Text representation of the Sequence Labeling task. Given an input sentence, the model must generate the same sentence annotated with html-style tags.

The model is prompted with a sentence to label. The expected output is the same sentence annotated with HTML-style tags. An example is provided in Figure 5.2. The HTML tags for each task are added as special tokens to the model’s

vocabulary. Previous research (Raman *et al.* 2022) found that different structures do not greatly impact the performance of the model so we use HTML-style tags because the format is easy for humans to read. Furthermore, LLMs, which have been trained on vast amounts of data from the Internet, are already familiar with this format, and implementing a constrained grammar for this structure is quite straightforward. In any case, our method can be adapted to any other task representation. For encoder-decoder models, the unlabeled sentence is given as input into the encoder block, while the decoder block generates the labeled output. For encoder-only models, we use the token \rightarrow during training as a separator between the unlabeled and labeled sentence. In the case of instruction-tuned models, instead of the separator, we use corresponding the chat-template to represent the unlabeled sentence as the user input and the labeled sentence as the chatbot response. We also experimented with generating only the labeled spans as output (i.e., `<Person> Obama </Person> <Location> New York </Location>`), but we obtained worse results.

5.3.2 Constrained decoding

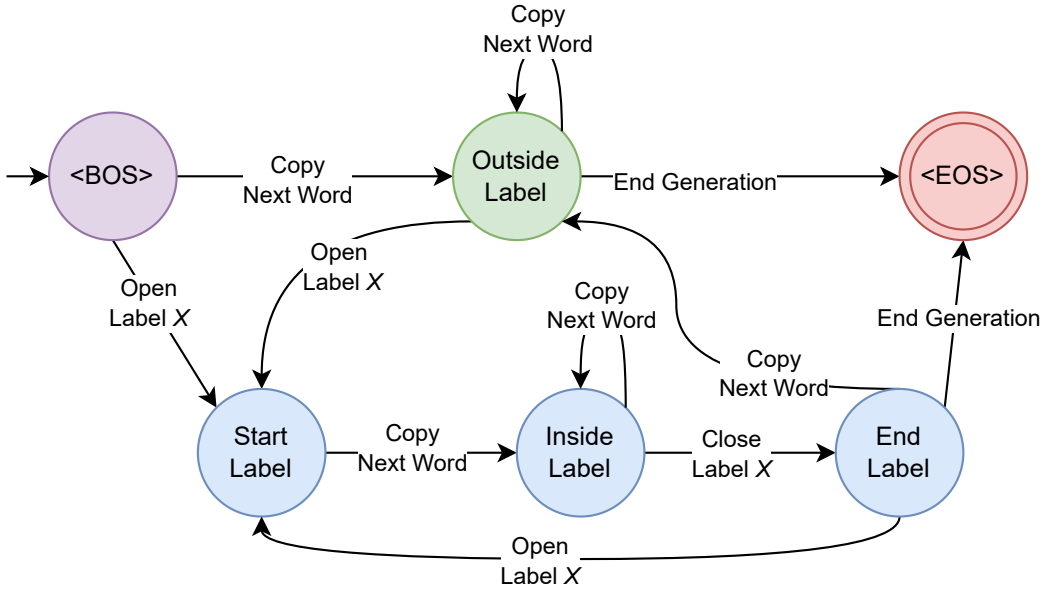


Figure 5.3 – Our Constrained Decoding Algorithm is defined as a Finite State Automaton.

The constrained decoding algorithm ensure that the output sequence contains

the same words as the input sequence. This **prevents hallucinations**, which are very common when a model is trained in one language and then used to label sentences in another language. It also ensures that the output sequence is a valid HTML annotation, with no unclosed tags, empty tags, or other errors. This **prevents the generation of unparseable outputs**. We implement our constrained decoding algorithm using the Finite State Automaton described in Figure 5.3. At each stage, the model can generate only a set of valid tokens. This set includes copying the next word from the input (if the word is split by the tokenizer into multiple tokens, all of them are copied to prevent the splitting of words). It can also open an HTML tag, but only if no tag remains open, or close it, but only if we have already opened a tag and copied at least a word. The generation process ends when all the words in the input have been copied into the output and no tag remains open.

Given a sequence $(x_1, x_2, \dots, x_{t-1})$ that has been generated thus far and a set S_t of valid next tokens at step t , the next token x_t is selected as:

$$x_t = \arg \max_{x \in S_t} P(x|x_1, x_2, \dots, x_{t-1})$$

where $P(x|x_1, x_2, \dots, x_{t-1})$ represents the conditional probability of token x given the prior tokens. Any token not in S_t is given a probability of zero, ensuring that the generated sequence adheres to the constraints. The probability for each token $x_i \in S_t$ is computed using the softmax function applied to the model predictions:

$$P(x_i|x_1, x_2, \dots, x_{t-1}) = \frac{e^{x_i}}{\sum_j e^{x_j}}$$

The probability of the generated sequence up to step T is computed as:

$$P(x_{1:T}|\text{<bos>}) = \prod_{t=1}^T P(x_t|x_1, x_2, \dots, x_{t-1})$$

While most previous constrained decoding algorithms are limited to greedy decoding, we implement a **constrained beam search** approach. We keep track of the top k most probable sentences at each step t , ensuring a broader exploration of the solution space and yielding higher-quality output sequences that adhere to the given constraints. Our constrained beam search approach adds very little overhead compared to the standard beam search decoding strategy. At each step, our only additional computation is to filter out invalid tokens from the beam. It's important to note that our constrained beam search decoding algorithm merely eliminates

invalid sequences from the search space. Consequently, the constrained beam search will always yield an output that is at least as good as, if not superior to, unconstrained beam search.

5.4 Experimental Setup

The datasets used address three information extraction tasks which are illustrated by Figure 5.4.




Figure 5.4 – Information Extraction Tasks in our experiments

Named Entity Recognition (NER): This task consists of detecting named entities and classifying them according to some pre-defined categories. We evaluate the models on MasakhaNER 2.0 (Adelani *et al.* 2022), a manually annotated NER dataset for 20 African languages. We train the models with the CoNLL03 (Sang and Meulder 2003) English training split. We focus on named entities referring to Person, Location and Organization.

Opinion Target Extraction (OTE): Given a review, the task is to detect the linguistic expression used to refer to the reviewed entity. We use the English SemEval 2016 Aspect Based Sentiment Analysis (ABSA) datasets (Pontiki *et al.* 2016). The English training split is used for fine-tuning; results are reported on the Spanish, French, Dutch, Russian and Turkish test sets.

Event Extraction (EE): It consists of detecting and classifying event mentions according to some pre-defined class-inventory. We use the English ACE05 (Walker *et al.* 2006) training split for training and the Chinese test split for evaluation. We also perform the Entity Mention Extraction task separately as an additional indicator of performance.

5.4.1 Language Models and baselines

Baselines: We assess the performance of our grammar-constrained beam search algorithm (**Cons**) against the unconstrained decoding baseline (**Base**). After fine-tuning, we test the same checkpoint using both constrained and unconstrained decoding. Additionally, our method is compared to popular encoder-only models, which currently set the benchmark for zero-shot cross-lingual transfer and have been widely adopted by the community. Thus, we evaluate mDeBERTa-v3 (He *et al.* 2021), an 86-million-parameter model, and GLOT500 (Imani *et al.* 2023), a 125-million-parameter model. Although we also experimented with XLM-RoBERTa (Conneau *et al.* 2020) models of various sizes, they consistently lagged behind mDeBERTa-v3 in performance. For MasakhaNER, we additionally compared with afro-xlmr-large (Alabi *et al.* 2022), a 355-million-parameter encoder-only model fine-tuned on African languages.

Text-to-text Models: We experiment with three different encoder-decoder models: mT0-XL (Muennighoff *et al.* 2023) 3.7 Billion parameter model. mT0-XL is an mT5 (Xue *et al.* 2021) pretrained multilingual language model fine-tuned in the cross-lingual task mixture xP3. We also experimented with mT5 itself and Aya-101 (Üstün *et al.* 2024) an encoder-decoder model trained with instruction data in 101 languages.

We also test multiple instruction tuned decoder-only models: Qwen2 (Yang *et al.* 2024), gemma (Team *et al.* 2024), LLaMA-3 (AI@Meta 2024), Aya-23 (Aryabumi *et al.* 2024) and Yi 1.5 (AI *et al.* 2024). These models have been trained on a wide range of tasks and languages, and have demonstrated strong multilingual capabilities.

5.4.2 Training Setup

All models were trained exclusively with English-labeled data and subsequently evaluated in the target languages. For the encoder-only models, we added a token classification layer (linear layer) on top of each token representation and trained them using the Cross-Entropy loss. The text-to-text models, were trained using the standard Next Token Prediction (NTP) loss. We finetune all the parameters of mT0 and mT5 using the Adafactor (Shazeer and Stern 2018) optimizer. For the other text-to-text models, we found that the full-finetuning approach produces suboptimal results. Therefore we use Low-Rank Adaptation (LoRA) (Hu *et al.* 2021) to adapt the models to the target task. LoRA freezes the pre-trained model

weights and injects trainable rank decomposition matrices into linear layers of the Transformer architecture. We applied the LoRA to all linear Transformer block layers as recommended by Dettmers *et al.* 2023. We use the AdamW optimizer (Loshchilov and Hutter 2017). Preliminary experiments showed that LoRA produces a better performance than the full-finetuning approach for these models. This has already been reported by previous research (Sainz *et al.* 2024a). We hypothesize that the large number of parameters in these models makes them prone to overfitting when finetuning all the hyperparameters on small datasets. For mT0 and mT5, we use a beam size of 4, while for larger models, we use a beam size of 1 as the computational cost of larger beams was prohibitive for us. In any case, we found that increasing the beam size did not significantly improve the performance of the models.

For both, encoder and text-to-text models we use the Huggingface open-source library (Apache-2.0 License) (Wolf *et al.* 2019).

	Encoder Models	mT5/MT0	Other text-to-text models
Finetuning	Full	Full	LoRA
Batch Size	32	16	32
Optimizer	AdamW	Adafactor	AdamW
Learning Rate	$5e^{-5}$	$1e^{-4}$	$7e^{-5}$
Scheduler	Cosine	Cosine	Cosine
Warmup steps	0	500	500
Beams	-	4	1
Sequence Length	192	512	512
Precision	FP16	BF16	BF16
Epochs (NER)	20	20	10
Epochs (OTE)	10	50	-
Epochs (ACE)	20	45	-

Table 5.3 – Hyperparameters used for fine-tuning the models.

We evaluate the models at the end of several epochs on the validation set and select the best checkpoint based on the F1 score. The full training hyperparameters are provided in Table 5.3. These hyperparameters were chosen based on a hyperparameter search on the validation set.

5.4.3 Evaluation Metrics

We evaluate the models using the standard F1-score metric for Sequence Labeling tasks (Sang and Meulder 2003). For the text-to-text models, the output of the model is converted into an IOB2 format by splitting the output into words by whitespace. All the models are evaluated using the seqeval library (Nakayama 2018).

5.5 Experiments

5.5.1 Named Entity Recognition

In this section we will present and discuss the experiments in the Named Entity Recognition, Opinion Target Extraction and Event Extraction tasks.

Lang	mT5-xl		mT0-xl		aya-101		Yi-1.5-9B-Chat		mDertaV3	afro-xlmr-large	GLOT500
	Base	Cons.	Base	Cons.	Base	Cons.	Base	Cons.			
English	93.4	93.7	93.2	93.3	93.2	93.4	94.5	94.3	93.4	93.4	92.3
Bambara	52.5	53.4	52.8	53.8	56.0	56.2	46.0	46.7	33.8	40.0	51.1
Ghomálá	46.1	47.5	43.3	43.7	25.8	25.5	45.9	49.6	43.3	44.0	45.7
Éwé	79.8	81.0	73.4	73.6	80.1	81.2	74.4	74.8	74.4	70.3	72.1
Fon	52.0	55.4	68.0	69.7	44.5	45.3	47.4	52.2	49.2	49.8	56.7
Hausa	71.3	73.8	70.0	71.9	67.9	70.1	62.0	61.3	70.7	74.1	67.2
Igbo	72.6	77.2	55.9	61.0	53.5	54.1	55.5	57.3	58.8	72.5	62.1
Kinyarwanda	71.9	73.1	71.9	74.3	67.3	69.2	48.7	53.3	65.7	67.9	66.1
Luganda	81.9	82.3	79.0	79.5	82.8	83.1	66.9	73.6	73.0	77.9	79.2
Mossi	52.5	53.7	55.4	55.7	56.7	56.7	54.4	57.1	44.6	45.7	51.4
Naija	76.3	83.5	73.5	80.1	69.3	72.1	60.5	63.5	78.7	80.4	71.1
Chichewa	77.7	78.8	76.5	76.7	79.9	80.2	80.5	80.4	73.7	79.6	76.6
chiShona	35.2	48.2	24.3	54.0	35.2	42.8	25.8	43.0	35.8	35.2	39.8
Kiswahili	86.4	89.6	85.7	88.0	83.3	84.7	73.1	71.4	86.7	88.2	84.0
Setswana	81.0	81.3	72.3	73.5	82.1	82.6	60.4	64.0	63.1	73.3	66.8
Akan/Twi	60.2	61.4	60.1	61.5	64.0	64.6	49.7	55.5	49.9	40.3	55.9
Wolof	53.3	54.3	56.4	56.8	61.3	62.3	57.1	60.9	42.0	51.3	61.6
isiXhosa	30.5	40.3	27.0	55.8	34.0	40.8	22.7	32.1	24.9	26.0	26.5
Yorùbá	55.1	58.5	51.0	51.3	26.0	25.8	48.9	55.4	34.1	52.5	54.4
isiZulu	49.4	54.9	39.2	66.7	40.5	45.0	23.5	33.7	44.7	47.1	43.3
Average MasakhaNER	62.4	65.7	59.8	65.7	58.4	60.1	52.8	57.1	55.1	58.7	59.6

Table 5.4 – F1 scores in the Named Entity Recognition Task. Model are trained in English and evaluated in a set of African languages.

Table 5.4 presents the performance of our method compared to the baselines in the NER task. All models show comparable performance in English. However, when assessing zero-shot cross-lingual transfer, significant performance differences emerge.

In the zero-shot cross-lingual transfer setting, constrained decoding consistently outperforms unconstrained decoding. For some languages, such as Bambara, Ghomálá, and Éwé, both methods yield similar results. In contrast, other languages, including Shona, isiXhosa, and Zulu, exhibit marked performance improvements. These Southern Bantu languages have unique linguistic features: they capitalize proper names following the noun class prefix (e.g., kweZambia) and display highly inflected morphology (Adelani *et al.* 2022). These attributes challenge the cross-lingual transfer abilities of English fine-tuned NER models. Consequently, all baseline models, including the encoder-only variants, perform suboptimally in these languages and are clearly outperformed by our constrained decoding approach.

As demonstrated in Section 5.6, text-to-text models struggle with agglutinative languages, frequently mislabeling entities by arbitrarily splitting them into sub-words. Our constrained decoding corrects this by ensuring that the output sentence retains the original words from the input sentence. Overall, constrained decoding excels in the zero-shot cross-lingual setting for languages with highly inflected agglutinative morphology. Although the performance gap is less pronounced for language isolates like Bambara, Éwé, Fon, and Twi, it remains significant.

Model	Unconstrained	Constrained	Delta
mT5-xl	62.4	65.7	+3.3
mT0-xl	59.8	65.7	+5.9
aya-101	58.4	60.1	+1.7
Qwen2-7B-Instruct	39.7	42.0	+2.3
gemma-1.1-7b-it	46.8	49.0	+2.2
Llama-3-8B-Instruct	51.2	52.7	+1.6
aya-23-8B	51.6	52.6	+0.9
Yi-1.5-9B-Chat	52.8	57.1	+4.3
GLot500		59.6	
mDeBERTa-v3		55.1	
Davlan/afro-xlmr-large		58.7	

Table 5.5 – Average F1 scores in the MasakhaNER dataset.

Models exhibit varying performance across languages. For instance, aya-101 achieves the best performance for Éwé, Luganda, Setswana, Twi, and Wolof, while mT0 is superior for languages such as Fon, Kinyarwanda, chiShona, isiX-

hosa, and Zulu. We attribute this to the different training data used by the models. Nonetheless, we observe that mT5-xl and mT0-xl, combined with our constrained decoding algorithm, outperform encoder-only models by more than 5 points in F1 score on average. This represents a significant improvement over the previous state-of-the-art for zero-shot cross-lingual transfer in NER tasks.

In Table 5.5 we present the average performance of different models on the MasakhaNER dataset. Qwen2, Gemma, and Aya-23 achieve suboptimal results compared to the other text-to-text models. This is likely due to these models being trained on a smaller number of high-resource languages, rendering them less proficient in African languages. However, the results demonstrate that constrained decoding is effective in improving the performance of all text-to-text models in zero-shot cross-lingual transfer.

5.5.2 Opinion Target Extraction

Lang	mT0-xl		GLOT	mDeBERTa
	Base	Cons	500	V3
English	82.6	84.8	82.6	83.6
Spanish	77.8	79.4	69.4	78.0
French	74.1	76.6	65.8	76.9
Dutch	74.1	77.1	66.5	77.3
Russian	71.1	75.7	69.2	76.5
Turkish	56.8	57.7	50.4	56.4
Average	70.8	73.3	64.3	73.0

Table 5.6 – F1 scores in the Opinion Target Extraction Task.

In the NER task, we experimented with cross-lingual transfer approaches using a set of low-resource African languages that significantly differ from English. For the Opinion Target Extraction task, we evaluated cross-lingual transfer performance into languages from the Indo-European language family. Due to the high computational cost of the text-to-text models, we only evaluated the best model from the previous task, mT0-XL.

As shown in Table 5.6, excluding Turkish (an agglutinative language), the performance decline in the target languages compared to English is less pronounced, suggesting a more seamless transfer. Even in this context, our constrained gen-

eration algorithm significantly surpasses the unconstrained generation. Finally, while mT0-XL and mDeBERTa-v3 show comparable performance, our approach demonstrates slightly higher average performance across the board.

5.5.3 Event Extraction

Lang	mT0-xl		GLOT 500	mDeBERTa V3
	Base	Cons		
English _{Entity}	95.5	95.5	94.5	95.3
Chinese _{Entity}	70.1	73.3	34.1	54.2
English _{Trigger}	78.9	78.9	74.1	78.0
Chinese _{Trigger}	49.6	52.1	0.0	30.5

Table 5.7 – F1 scores in the Event Extraction Task.

For the Event Extraction task we aim to perform zero-shot cross-lingual transfer from English to Chinese. This task is particularly challenging due to the vast linguistic and cultural differences between the two languages, including script type, syntax, semantics, and the use of tones in Chinese. As reported in Table 5.7, both GLOT500 and mDeBERTa struggle with the transfer from English to Chinese, whereas mT0-XL achieves much better results. Consistent with previous evaluations, our constrained generation approach improves over the unconstrained generation method by approximately 3 points in F1 score.

5.5.4 Model Transfer vs Data Transfer

In this chapter we focus on improving the zero-shot model-transfer approach. However, constrained decoding can also be used in conjunction with data transfer. In this section we compare the performance of the constrained decoding algorithm when used in both zero-shot model-based transfer and data transfer settings using the MasakhaNER NER dataset. To this end, we use the automatically generated NER datasets for eight African languages from Chapter 4. These datasets were generated by translating the CoNLL2003 (Sang and Meulder 2003) English dataset into the target languages using NLLB200 (Costa-jussà *et al.* 2022) and then projecting the labels using T-Projection. Using the same settings as for the zero-shot approach, we evaluate the performance of the constrained decoding algorithm when fine-tuning the models on the generated datasets. The results are

presented in Table 5.8. “Zero” refers to the models trained with English CoNLL 2003 data and evaluated in the target languages, while “Data” refers to the models fine-tuned on the translated CoNLL 2003 datasets. Both settings use the constrained decoding algorithm.

Lang	mDebertaV3		mT5-xl		mT0-xl		aya-101		Llama-3-8B-Instruct		Yi-1.5-9B-Chat	
	Zero	Data	Zero	Data	Zero	Data	Zero	Data	Zero	Data	Zero	Data
Hausa	70.7	72.7	73.8	67.8	71.9	72.4	70.1	72.6	63.5	70.1	61.3	62.4
Igbo	58.8	71.4	77.2	69.9	61.0	72.9	54.1	82.7	54.7	73.7	57.3	58.7
Chichewa	73.7	77.2	78.8	51.6	76.7	76.8	80.2	83.0	78.6	68.2	80.4	52.4
chiShona	35.8	74.9	48.2	75.0	54.0	74.1	42.8	77.0	25.3	65.3	43.0	46.2
Kiswahili	86.7	85.5	89.6	77.4	88.0	85.1	84.7	85.4	77.2	80.1	71.4	65.0
isiXhosa	24.9	72.3	40.3	53.6	55.8	74.8	40.8	74.5	25.8	64.8	32.1	46.0
Yorùbá	34.1	42.7	58.5	37.1	51.3	46.7	25.8	60.5	34.9	56.4	55.4	37.6
isiZulu	44.7	66.7	54.9	64.4	66.7	71.2	45.0	64.9	22.5	58.0	33.7	36.4
Average	53.7	70.4	65.2	62.1	65.7	71.8	55.5	75.1	47.8	67.1	54.3	50.6

Table 5.8 – F1 Scores for Named Entity Recognition Task. “Zero” refers to the model trained in English and evaluated on a set of African languages. “Data” refers to the model trained on automatically translated and projected data using T-Projection for each language.

The results show that the zero-shot cross-lingual transfer performance when using text-to-text models such as mT5-xl or mT0-xl is significantly better than the zero-shot performance of mDeBERTa-v3, as we already demonstrated in Section 5.5.1. However, mDeBERTa-v3 shows very competitive results in the data-transfer setting. For text-to-text models, the performance of the zero-shot and data-transfer approaches varies across languages. In languages where the model is proficient, such as Hausa or Igbo for mT5-xl, the zero-shot approach outperforms the data-transfer approach. However, in languages where the model is less proficient, such as isiXhosa or Zulu, the data-transfer approach is superior. In the case of aya-101 and LLama-3, which we reported to be less proficient in African languages in Section 5.5.1, the data-transfer approach results in a significant performance improvement. In fact, aya-101 outperforms all other models in the data-transfer setting. This suggests that the constrained decoding algorithm can be used in conjunction with data-transfer methods to further improve the performance of models in low-resource languages.

Similar to the insights from Chapter 3, the results suggest that when a model is proficient in both the source and target languages, model-based transfer is superior to data-based transfer. Thanks to the methodology developed in this chapter, we




Figure 5.5 – Percentage of hallucinated words compared to the performance delta between unconstrained and unconstrained beam search in MasakhaNER using mT0-XL.

can now successfully leverage the power of text-to-text LLMs in a zero-shot setting to achieve superior zero-shot cross-lingual transfer results. However, when the model is less proficient in the target language, data-based transfer can be a better option. Data transfer also has the advantage of allowing the use of more efficient models. The results demonstrate that while mDeBERTa-v3 is not competitive in the zero-shot setting, it achieves similar results to the best-performing text-to-text models in the data-transfer setting, despite having fewer parameters and requiring less computational resources.

5.6 Ablation Study

In this section we aim to better understand why and in which scenarios constrained decoding performs better than unconstrained decoding. To achieve this, we identify the types of mistakes that unconstrained decoding makes which are subsequently fixed by constrained decoding. These errors can be grouped into three categories: inconsistent HTML markups, word hallucinations, and word splittings.

Inconsistent HTML markups: The model occasionally generates HTML markup that cannot be parsed correctly, such as when a tag is opened but never closed. We

found that this occurs in less than 1% of the annotated sentences. Consequently, it has a negligible effect on the overall performance of the model.

Word hallucinations: The model sometimes includes in the output a word that was not present in the input. This occurs because unconstrained generation often produces output that mixes English and the target language. For instance, given the sentence “*Kaliforni sullā sēn togse*”, mT0-XL, when using unconstrained decoding, produces “<Location> *California* </Location> *sullā sēn togse*”. In this instance, the model has translated “*Kaliforni*” to “*California*”. Furthermore, inadvertent translation is not the only cause of hallucinations in the output. Perhaps due to a limited understanding of the target language, the model often introduces typos (e.g., “*okudlula*” incorrectly becomes “*okudludlule*”). Interestingly, it even mixes African languages. For instance, given a Zulu sentence as input containing the word “*Musawenkosi*” (God Bless You), the model outputs the very similar Chichewa word “*Mumawenkosi*” (You are welcome).

Word Splittings: They refer to instances where the model either divides a word into multiple subwords or, conversely, combines several words into a single one. This occurs because the model has been trained in English and, when tested on agglutinative languages, it attempts to mimic English morphology by arbitrarily splitting words. For instance, the sequence “<Location> *waseThekwini* </Location> <Person> *uShauwn Mkhize* </Person>” becomes “*wase* <Location> *Thekwini* </Location> *u* <Person> *Shauwn Mkhize* </Person>”. This behavior is interesting, as lemmatization is a component of many downstream Information Extraction applications. One could argue that this is the desired behavior. However, although accidental lemmatization was performed correctly in this particular example, this is not usually the case. For instance, in Basque (whose results are not reported here for brevity, although the models were tested in this language), as illustrated in Figure 5.1, the model incorrectly splits the term “*Realean*” into “*Reale*” and “*an*”. However, “*Reale*” does not represent the correct lemma, which would correspond to “*Reala*”, the name of a football team. Therefore, the model seems to be arbitrarily splitting words to mimic English morphology.

We calculate the percentage of sentences containing some of these errors for each language in the NER task when using mT0-XL with unconstrained generation. The results are depicted in Figure 5.6. Additionally, we compared the overall percentage of sentences containing any error with the performance difference between constrained and unconstrained generation. The larger the delta, the greater




Figure 5.6 – Average percentage of mistakes generated by Unconstrained Beam search in MasakhaNER using mT0 models of different sizes

the performance improvement with constrained generation.

Figure 5.6 indicates that word splitting and hallucinations correlate with the performance delta, suggesting that addressing these issues is key to the superiority of the constrained generation algorithm. It also underscores that unconstrained generation produces a substantial proportion of sentences with errors. In cases like chiShona and isiXhosa (discussed in Section 5.5.1), this could affect over 50% of the output sentences. It should be noted that word splitting has a more pronounced effect on the performance delta than hallucinations. This can be attributed to the standard sequence evaluation method used for these tasks.

To convert the model’s output into IOB2 encoding, we derive annotations such as “*B-LOC O O O*” for the example “<Location> *California* </Location> *sullã sêh togse*”. This annotation remains accurate even if the model translates the entity into English. However, when the model splits or merges words, the IOB2 labeling is disrupted, rendering the sentence incorrect in the evaluation. Thus, although the evaluation method may gloss over hallucination errors, it is important to note that models generate a significant number of hallucinations when producing unconstrained predictions, potentially impacting the ultimate efficacy and applicability of IE systems.

We also evaluated the total number of mistakes generated by unconstrained beam search in the NER task with mT0 models of varying sizes. As illustrated




Figure 5.7 – Average F1 score in MasakhaNER compared to the mT0 model size

in Figure 5.6, word splitting and inconsistent HTML markups remain consistent across models with different parameter sizes. However, the frequency of hallucinations decreases as the model size increases. This might be because models with more parameters have a more refined representation of individual languages and therefore mix languages less frequently.

Additionally, we assess the average F1 score in the NER task for mT0 models ranging from 300 million to 3.7 billion parameters. The results, presented in Figure 5.7, show that as the mT0 model’s parameter count increases, the F1 score improves, although we observe diminishing returns beyond 1.2 billion parameters. While our experiments utilize the 3.7 billion parameter mT0-XL, constrained generation surpasses both GLOT500 (a 125 million parameter model) and afro-xlmr-large (355 million parameters) when using an mT0 model with only 580 million parameters. This indicates that the superiority of our method over encoder-only models is not solely due to leveraging a larger model. Notably, with constrained generation, the 580 million parameter mT0 model achieves performance comparable to the 1.2 billion parameter model when the latter employs unconstrained generation. Therefore, constrained generation is also considerably more computationally efficient than its unconstrained counterpart.

Finally, we evaluate the performance of mT0-XL using a varying number of




Figure 5.8 – Average F1 score of mT0-XL in a subset of MasakhaNER compared to the number of beams used for decoding.

beams. We assess the same checkpoint with beam search ranging from 1 to 8 beams. For these experiments, we utilize a subset of MasakhaNER2, which includes the following languages: Bambara, Ghomálá, Éwé, Fon, Hausa, Igbo, Kinyarwanda, Luganda, and Mossi. As illustrated in Figure 5.8, increasing the number of beams has a negligible effect on performance. Considering that the computational cost and GPU memory requirements increase linearly with the number of beams, in this scenario, using a single beam (greedy decoding) offers the best performance-to-cost ratio. This is because the model is highly confident in its top prediction during each step of the decoding process, and introducing additional beams does not significantly diversify or improve the generated outputs.

5.7 Conclusion

In this Chapter we introduce a Constrained Beam Search Algorithm that can be seamlessly incorporated into any text-to-text LLM. We demonstrate that, compared to Unconstrained Beam Search, our algorithm significantly improves zero-shot cross-lingual performance across a broad range of IE tasks and languages. Through an extensive ablation study, we show that constrained generation effectively mitigates issues such as word-splitting and language mixing, which lead to typos and unintentional translations, errors commonly observed when applying text-to-text models to these tasks. Our approach allows the text-to-text mT0 language model to outperform encoder-only models, which had previously set the

state-of-the-art standard for zero-shot cross-lingual IE. To the best of our knowledge, we present the best zero-shot cross-lingual results up to date.

The method developed in this chapter enables model-based cross-lingual transfer for sequence labelling tasks with text-to-text models. This is a significant step forward in the field of zero-shot cross-lingual transfer, as it allows for the use of more powerful models that can handle a wide range of tasks. Considering the prevailing focus on text-to-text LLMs in current research, and the infrequent training of new encoder-only models, we believe that this represents significant progress in this research area.

6. CHAPTER

Medical MT5: Cross-Lingual Transfer for Domain-Specific Task

In this chapter we will introduce Medical mT5, an open-source multilingual text-to-text large language model for the medical domain. We will leverage all the data-transfer and model-transfer techniques developed in the previous chapters. We will build a multilingual pre-training, fine-tuning, and evaluation framework for the medical domain. Medical mT5 demonstrates the importance of the technology and knowledge developed in this thesis, resulting in the first multilingual text-to-text medical model when it was created.

6.1 Motivation and Contributions

As it is the case for many application domains, there is an increasing interest in applying Artificial Intelligence (AI) and Natural Language Processing (NLP) techniques to assist medical experts in their everyday activities. With this aim in mind, a number of language models have been trained or adapted to the medical domain. These include encoder-only models such as SciBERT (Beltagy *et al.* 2019), BioBERT (Lee *et al.* 2020) or PubmedBERT (Gu *et al.* 2022). These models have obtained state-of-the-art results in discriminative tasks, with the advent of text-to-text and text-generation models, a new generation of language models has been developed. These models are typically much larger and have a much broader scope than the encoder-only models. Examples of these models include SciFive

(Phan *et al.* 2021), BioGPT (Luo *et al.* 2022) Med-PaLM (Singhal *et al.* 2022), PMC-LLaMA (Wu *et al.* 2023) or ClinicalGPT (Wang *et al.* 2023a).

However, the development of all the aforementioned text-to-text LLMs has been focused on a single language, usually English. As a consequence, there is a lack of high-quality multilingual data for pre-training models, a lack of models themselves, and a lack of high-quality multilingual evaluation benchmarks for the medical domain. Although there have been efforts to generate evaluation data in languages other than English (Wang *et al.* 2023a; Carrino *et al.* 2022), they have consisted largely of monolingual approaches.

To address these issues, we have compiled, to the best of our knowledge, the largest multilingual corpus for training LLMs adapted to the medical domain. Our corpus includes 3B words in four languages, namely, English, Spanish, French, and Italian. While relatively small when compared to existing English datasets (Wu *et al.* 2023), it allowed us to build Medical mT5, the first open-source text-to-text multilingual model for the medical domain. Additionally, we have built a multilingual evaluation framework for the medical domain that can be used to evaluate the performance of any future multilingual model in the medical domain.

Medical mT5 has been built on the work presented in previous chapters. We use the data transfer approach to reduce the cost of annotating new multilingual evaluation benchmarks for the medical domain. We also use constrained decoding to achieve high-quality zero-shot model-based cross-lingual transfer. Medical mT5 demonstrates the benefits of the techniques developed during this thesis and their application in real-world scenarios where data is scarce.

Medical mT5 outperforms similarly-sized text-to-text models for the Spanish, French, and Italian benchmarks while being competitive in English to current state-of-the-art text-to-text (Xue *et al.* 2021; Chung *et al.* 2022) and encoder-only models (Lee *et al.* 2020; He *et al.* 2021). The results show that continuing pre-training of a multilingual text-to-text model such as mT5 allows to successfully adapt it to the medical domain, even when the amount of domain-specific data is relatively modest (ranging between 1B words for English and Spanish to 150M in Italian). Summarizing, the contributions of this chapter are:

- The collection of the largest publicly available in-domain medical multilingual corpus for Spanish, French, and Italian languages. Together with the already existing English data, we release a corpus of 3 billion tokens¹.
- We use the data-transfer approaches developed in previous chapters to build

¹<https://hf.co/datasets/HiTZ/Multilingual-Medical-Corpus>

two new datasets for Spanish, French, and Italian on Argument Mining² and generative Question-Answering³ tasks, generated by taking their original English versions as a starting point.

- the public release of two Medical mT5 versions: a 770M⁴ and 3B⁵ parameter text-to-text open-source models which obtain state-of-the-art results in multilingual sequence labeling for the medical domain, most notably in multi-task and zero-shot cross-lingual settings.

Other benefits of our Medical mT5 models include the comparatively low hardware requirements needed for both fine-tuning on downstream tasks (the large 770M version easily fits in a 24GB GPU) and for inference (a 12GB GPU should be enough). As an example, a LLaMA 7B model (Wu *et al.* 2023) requires at least 4 80GB A100 GPUs. This makes our models more accessible to the research community and to small and medium-sized companies.

6.2 Related Work

As it has been the case in most application domains, Large Language Models (LLMs) have facilitated significant improvements in the state-of-the-art for medical NLP tasks (Singhal *et al.* 2022; Wu *et al.* 2023; Mayer *et al.* 2021). The most popular approaches use models pre-trained on medical corpora such as SciBERT (Beltagy *et al.* 2019), BioBERT (Lee *et al.* 2020), PubmedBERT (Gu *et al.* 2022), BSC-BIO (Carrino *et al.* 2022), or BioLinkBERT (Yasunaga *et al.* 2022).

While the previous encoder-only models focused on discriminative tasks, the emergence of generative models such as LLaMa (Touvron *et al.* 2023a), PaLM (Singhal *et al.* 2022), and GPT-3 (Brown *et al.* 2020) has generated significant interest in adapting such LLMs to the medical domain. These models include, but are not limited to, SciFive (Phan *et al.* 2021), an English T5 encoder-decoder model adapted to the scientific domain, and decoder models such as BioGPT (Luo *et al.* 2022), Med-PaLM (Singhal *et al.* 2022), PMC-LLaMA (Wu *et al.* 2023), and ClinicalGPT (Wang *et al.* 2023a).

Additionally, a range of abstractive question-answering tasks has been proposed as evaluation benchmarks, on which the larger models (Wu *et al.* 2023;

²<https://hf.co/datasets/HiTZ/multilingual-abstrct>

³<https://hf.co/datasets/HiTZ/Multilingual-BioASQ-6B>

⁴<https://hf.co/HiTZ/Medical-mT5-large>

⁵<https://hf.co/HiTZ/Medical-mT5-xl>

Singhal *et al.* 2022; Wang *et al.* 2023a) achieve the best results. While interesting, both these LLMs and benchmarks have been developed with a focus on a single language, usually English.

Furthermore, these LLMs require hardware that is simply not affordable for the majority of end-users and researchers. To address these issues, we propose Medical mT5, a multilingual text-to-text model adapted to the medical domain which, despite its relatively modest size and low running costs, obtains competitive results, notably in multi-task and zero-shot cross-lingual settings.

6.3 Compiling a Multilingual Corpus for the Medical Domain

Obtaining good quality medical corpora is usually difficult due to the sensitive nature of the data. This is even more challenging for non-English languages, as the availability of data for other languages is in general more restricted. Despite these issues, we have successfully gathered and curated a diverse collection of public relevant corpora of medical texts in English, French, Italian and Spanish to generate the Medical mT5 model. The data sources are summarized in Table 6.1.

6.3.1 English

As listed in table 6.1, we collected around 1B words from three sources related to the medical domain: (i) **ClinicalTrials** is a set of documents of clinical studies from all over the world; (ii) **EMEA** is an English-Spanish parallel corpus with documents provided by the European Medicines Agency (Tiedemann 2012) and, (iii) **PubMed**, which contains data

Language	Source	Words
English	ClinicalTrials	127.4M
	EMEA	12M
	PubMed	968.4M
	Total	1.1B
Spanish	EMEA	13.6M
	PubMed	8.4M
	Medical Crawler	918M
	SPACC	350K
	UFAL	10.5M
	WikiMed	5.2M
	Total	960M
French	PubMed	1.4M
	Science Direct	15.2M
	Wikipedia - Médecine	5M
	EDP	48K
	Google Patents	654M
	Total	676M
Italian	Medical Commoncrawl - IT	67M
	Drug instructions	30.5M
	Wikipedia - Medicina	13.3M
	E3C Corpus - IT	11.6M
	Medicine descriptions	6.3M
	Medical theses	5.8M
	Medical websites	4M
	PubMed	2.3M
	Supplement description	1.3M
	Medical notes	975K
	Pathologies	157K
	Medical test simulations	26K
	Clinical cases	20K
	Total	143M
Total		3.02B

Table 6.1 – Data sources and word counts by language.

from various sources such as MEDLINE, life science journals and online books, provides the bulk of the English data.

6.3.2 Spanish

Apart from **EMEA** and **PubMed**, which we also used for Spanish, the biggest portion of the data came from the **Medical Crawler**, a biomedical corpus compiled by Carrino *et al.* 2022. Additionally, we also included **SPACC**, **UFAL** and **WikiMed**, a corpus built ad-hoc from Wikipedia entries. Table 6.1 provides the details of the collected data, which amounts to ≈ 1 B words.

6.3.3 French

A total of 7,192,779 sentences and 670,972,717 words were compiled using the data sources listed in Table 6.1. **Science Direct** offers a collection of scientific and medical publications. We filtered relevant articles with the keyword “Médecine”, and the obtained XML documents were parsed to extract the `<dc:description>` tag. As for Spanish, we took advantage of **Wikipedia** and **PubMed** as a source of medical knowledge. PubMed data was extracted using the `Bio.Entrez` package⁶. For wikipedia we obtain HTML formatted data from the category “Category:Médecine”. The **EDP French/English Parallel Medical Corpus** (Jimeno-Yepes *et al.* 2017) provides bilingual content from journals that address domains such as dentistry and life sciences. From this source, we downloaded the dataset labeled “EDP French corpus, text format”. Finally, **Google Patents** is a comprehensive repository of patent data from around the world. Google Patents data were retrieved by filtering using the IPC code and abstract language. A final French language verification step was undertaken by applying the `langdetect` package (version 1.0.9).

6.3.4 Italian

The crawling and pre-processing of the Italian split of the corpus followed the methodology described by Carrino *et al.* 2022. First, we compiled a list of 504 medical terms, which we use as seeds to scrape the Italian split of the **MC4 Common Crawl Corpus** by only selecting the pages which contained at least one of the keywords in their URL domain. To create the list, we extracted 600 keyword

⁶<https://biopython.org/docs/1.75/api/Bio.Entrez.html>

terms related to medicine from the *Dizionario analogico della Lingua Italiana* (Zanichelli). We excluded some sectors and discarded terms that may lead to ambiguous queries (e.g., actions, which contained mainly verbs, proverbs, general terms like “assistente”, etc.). We normalized rare variants (“bacteriologia” to “batteriologia”) and stemmed all terms without lemmatizing, as most terms are already lemmatized in the dictionary; we performed univertation of multiword units (e.g., “esamedelleurine”, “follow-up”), and removed the duplicates. This resulted in a corpus of 67 million tokens, which we joined with other sources of text such as **Medical dissertations**, **Drug use instructions**, **PubMed abstracts**, etc. as detailed in Table 6.1, resulting in a $\approx 145\text{M}$ word corpus.

6.4 Medical mT5

Multilingual T5 (mT5) (Xue *et al.* 2021) is an extension of the original T5 (Raffel *et al.* 2020) framework, which is optimized for multilingual tasks. The T5 model is grounded in the Transformer encoder-decoder architecture (Vaswani *et al.* 2017). With its decoder block, T5 is capable of generating sequences of tokens in an auto-regressive fashion. T5 was designed to convert every NLP problem into a text-to-text task, and mT5 extends this strategy to a multitude of languages, leveraging a shared vocabulary for diverse scripts. mT5 was trained using mC4, a 1 Trillion token Common Crawl-based dataset covering 101 languages. The pre-training is based on a masked language modeling “span-corruption” objective, where consecutive spans of input tokens are replaced with a mask and the model is trained to reconstruct the masked-out tokens.

6.4.1 Pre-training Medical mT5

Medical mT5 is built upon the same architecture as mT5 (Xue *et al.* 2021). We release two different models: Medical-mT5-large (738M parameters) and Medical-mT5-xl (3 billion parameters). Both models were initialized using the pre-trained weights of their corresponding mT5 checkpoints and continued their pre-training using the 3B word medical domain dataset described in Section 6.3 (with x2 over-sampling for the Italian split). To prevent over-fitting, we run the training for only one epoch, as preliminary experiments showed that performance degraded with more epochs. We adhered to the self-supervised parameter settings described in Xue *et al.* 2021 and detailed in Table 6.2. It should be noted that Medical-mT5-large was trained with a sequence length of 1024 tokens whereas Medical-mT5-xl

	Medical-mT5-large	Medical-mT5-xl
Param. no.	738M	3B
Sequence Length	1024	480
Token/step	65536	30720
Epochs	1	1
Total Tokens	4.5B	4.5B
Optimizer	Adafactor	Adafactor
LR	0.001	0.001
Scheduler	Constant	Constant
Hardware	4xA100	4xA100
Time (h)	10.5	20.5
CO ₂ eq (kg)	2.9	5.6

Table 6.2 – Pre-Training settings for Medical mT5.

was limited to a sequence length of 480 tokens due to GPU memory limitations. Medical mT5 was trained using the Flax implementation of mT5 in the Hugging Face Transformers library (Wolf *et al.* 2020). All experiments were conducted on our private servers, employing 4xA100 80GB GPUs. We made calculations for a carbon footprint estimation based on a 400W consumption per GPU and a carbon intensity of 0.171 kg/kWh⁷.

6.5 Generating New Multilingual Benchmarks: Real-World Application of Data Transfer

There is a lack of multilingual evaluation benchmarks for the medical domain. The only available benchmark in English, Spanish, French, and Italian is the relatively small e3C (Magnini *et al.* 2021). While medical domain evaluation datasets are scarce for Spanish, French and Italian, many datasets exist for English. Therefore, this is a good opportunity to apply the data transfer techniques developed in previous chapters to generate data for other languages. We focused on two different types of tasks: (i) a sequence labeling task, **Argument Mining**, which involves detecting and classifying the argument component spans and their relations, and (ii) **Abstractive Question Answering**, where the model is expected

⁷Sourced from <https://app.electricitymaps.com/map>

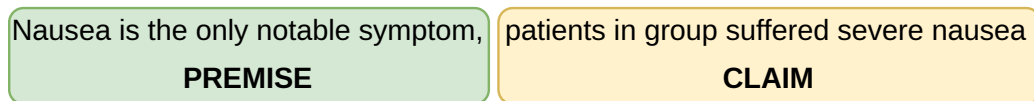


Figure 6.1 – Example of an annotated abstract from the AbstrCT dataset.

to generate an answer in response to an input question. In both cases we used existing labeled English data as a starting point.

6.5.1 Argument Mining

The AbstrCT dataset is composed by English medical and scientific texts collected from the MEDLINE database and manually annotated with two types of argument components: Claims and Premises (Mayer *et al.* 2021). An example of the task is illustrated in Figure 6.1

A ‘claim’ is a concluding statement made by the author about the outcome of the study. In the medical domain, it may be an assertion of a diagnosis or a treatment. A ‘premise’ corresponds to an observation or measurement in the study (ground truth), which supports or attacks another argument component, usually a claim. It is important that they are observed facts and, therefore, credible without further evidence.



Figure 6.2 – Data construction process for generating the Spanish, French and Italian versions of the AbstrCT dataset.

We generated French and Italian parallel versions of the dataset using the same method as for Spanish in Yeginbergen *et al.* 2024. First, the English dataset is translated into the target language using the machine translation model NLLB200-3.3B (Costa-jussà *et al.* 2022). Then, the labels in the source language are transferred to the target language using AWESOME align (Dou and Neubig 2021) and the annotation projection algorithm developed in Chapter 3. Finally, to ensure the quality of the generated dataset, the projections are manually reviewed by an expert in the target language. Thanks to this process, the manual annotation labor is significantly reduced compared to annotating the data from scratch. This process is illustrated in Figure 6.2.

The AbstRCT dataset is divided into three splits, neoplasm, glaucoma and mixed. Following previous work, we fine-tune the models with the first one and then evaluate the in-domain performance on the neoplasm test split and the cross-domain performance on the glaucoma and mixed splits. Previous works using the AbstRCT datasets have employed different definitions of the F_1 score metric, such as token-level F_1 (Mayer *et al.* 2021; Yeginbergen *et al.* 2024). However, in this paper, we report results using the standard sequence level F_1 score (Sang and Meulder 2003), a much more strict metric, which explains the lower results for all the models.

6.5.2 Question Answering

We use the BioASQ-6B English Question Answering dataset (Tsatsaronis *et al.* 2015) to generate parallel French, Italian and Spanish versions. Given a biomedical question and a set of snippets of text with relevant information about the question, the model must generate the *ideal* answer. This task is similar to the Retrieval Augmented Generation (RAG) task (Lewis *et al.* 2020), where the model must generate an answer given a context or set of contexts. A set of ideal gold answers are provided to assess the performance of the models. We machine-translated the questions and ideal answers into French, Italian and Spanish using the NLLB200 3B parameter model (Costa-jussà *et al.* 2022). In this case, as this is not a sequence labeling task, no annotation projection is needed. Nevertheless, the quality of a small set of translations was manually reviewed to ensure the quality of the generated data.

6.6 Experimental Setup

In this section we describe the datasets and evaluation tasks used to measure the performance of Medical mT5. We also provide the details of the training and evaluation process, and baseline models used for comparison.

6.6.1 Datasets

The list of tasks used for evaluation is listed in Table 6.3. The **Sequence labeling tasks** include medical NER, detecting and classifying named entities according to some pre-defined categories, and Argument Mining, described in Section 6.5. Performance for every sequence labeling task is evaluated using standard sequence

6 MEDICAL MT5: CROSS-LINGUAL TRANSFER FOR DOMAIN-SPECIFIC TASK

Representation	Task	Dataset	Languages	Entity Type
Sequence Labelling	Named Entity Recognition	NCBI-Disease, Dogan <i>et al.</i> 2014	EN	Disease
		BC5CDR Disease, Li <i>et al.</i> 2016	EN	Disease
		BC5CDR Chemical, Li <i>et al.</i> 2016	EN	Chemical
		DIANN, Fabregat <i>et al.</i> 2018	EN, ES	Disability
		E3C, Magnini <i>et al.</i> 2021	EN, ES, FR, IT	Clinical Entity
		PharmaCoNER, Gonzalez-Agirre <i>et al.</i> 2019	ES	Pharmacological
	Argument Mining	AbstrCT, Mayer <i>et al.</i> 2021	EN, ES, FR, IT	Claims and Premises
Generative Question Answering	Question Answering	BioASQ 6B, Tsatsaronis <i>et al.</i> 2015	EN, ES, FR, IT	Biomedical QA

Table 6.3 – List of evaluation tasks used to measure the performance of Medical mT5.

level F_1 score (Sang and Meulder 2003). We also evaluate the performance of Medical mT5 on the **Generative Question Answering task** using the BioASQ dataset, described in Section 6.5.2.

6.6.2 Conversion to Text-to-Text Format

Medical mT5 is a text-to-text model. This means that, given a text input, it learns to generate a text as output. Therefore, every evaluation task must be converted into a text-to-text format (Xue *et al.* 2021). In our experiments the output text is always generated using beam search with 4 beams.

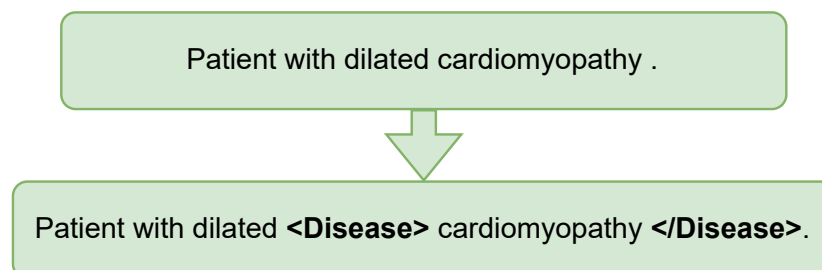


Figure 6.3 – Text-to-Text representation of the Sequence Labeling task. Given an input sentence, the model is expected to generate the same sentence annotated with html-style tags.

To address sequence labeling tasks, we use the same approach presented in Chapter 5. As illustrated in Figure 6.3, Text-to-text models such as Medical mT5

are prompted with the sentence to label. The expected output is the same sentence annotated with HTML-style tags. The HTML tags for each task are added as special tokens to the model vocabulary. We use constrained decoding to ensure that the output contains the same words as the input and a valid HTML annotation. The constrained decoding algorithm is the one presented in Chapter 5.

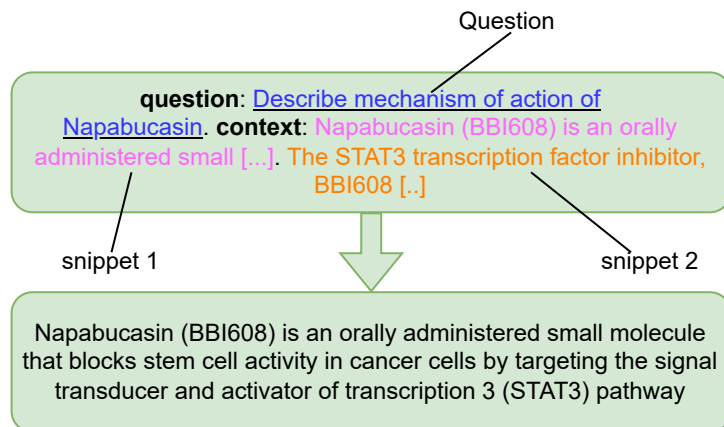


Figure 6.4 – Text-to-Text representation of the BioASQ task. Given a question and a set of relevant snippets, the model generates an answer.

With respect to the BioASQ **Abstractive Question Answering task**, the input prompt contains the question and a context. As shown in Figure 6.4, the context is generated by concatenating all the provided possible snippets. The expected output should be the generated answer to the question, which is then compared to the gold ideal answer.

6.6.3 Baselines

As we have developed Medical mT5 by continuing the training of mT5 checkpoints, our primary point of comparison should be mT5 (Xue *et al.* 2021). Thus, our first objective is to assess whether training the model on our multilingual medical-domain corpus enhances its performance for tasks specific to this domain. Furthermore, we also benchmark our model against SciFive (Pubmed+PMC) a T5-based 738M parameter model (Phan *et al.* 2021) trained exclusively on a corpus of 78B words containing scientific and medical English data. Additionally, we compare the performance of Medical mT5 with Flan-T5 (Chung *et al.* 2022), which also adopts the T5 architecture but has been finetuned on a huge instruction-

following dataset for almost 2K tasks. Flan-T5 achieves state-of-the-art performance in numerous benchmarks, including some from the medical domain (Singhal *et al.* 2022). We tested all three types of text-to-text models under identical settings and hyperparameters.

We also measure Medical mT5 with the performance of encoder-only models in sequence labeling tasks. We report results with mDeBERTaV3 (He *et al.* 2021) which is widely used for sequence labeling and excels in multilingual tasks (Adelani *et al.* 2022; Agerri and Agirre 2023). Although we also tested XLM-RoBERTa (Conneau *et al.* 2020) and GLOT500 (Imani *et al.* 2023), their results were worse than those obtained by mDeBERTaV3. Finally, we also compare with BioBERT v1.1 (Lee *et al.* 2020), which has been pretrained on a large English-only biomedical dataset. We do not evaluate the performance of encoder-only models in the question-answering task, as their architecture is not designed for text generation.

6.6.4 Hyperparameters settings

For sequence labeling, when using encoder-decoder models, we use a learning rate of 1×10^{-4} , a batch size of 8, and a maximum sequence length of 256 tokens. We use the Adafactor optimizer (Shazeer and Stern 2018) with cosine learning rate decay to 0 and 500 warmup steps. The number of epochs varies depending on the task. For the E3C dataset, which is very small, we use 100 epochs. For the other datasets, we use 45 epochs. When training the model in a multi-task setting, we use 12 epochs. For the question-answering task, we use 15 epochs. We use a beam size of 4 for all the tasks and no sampling. Models are trained using bfloat16 precision.

For encoder-only models, we use a batch size of 32, a learning rate of 5×10^{-5} , and 40 epochs. We use the AdamW optimizer (Loshchilov and Hutter 2017) with a cosine learning rate scheduler that decays the learning rate to 0. We use a maximum sequence length of 256 tokens. Encoder-only models are trained using fp16 precision.

For all models, we evaluate the model during training on the validation set periodically and select the model at the epoch with the highest performance on the validation set.

6.7 Experimental Results

In this section we present the evaluation results of Medical mT5 on Sequence Labeling and Question Answering tasks.

6.7.1 Sequence labeling Tasks

In this section we report on the performance of Medical mT5 and of the baselines in the **sequence labeling tasks** across different settings.

Lang	Dataset	mT5 _{large}	mT5 _{XL}	SciFive	FlanT5 _{large}	FlanT5 _{XL}	mDeBERTav3 base	BioBERT	MedMT5 _{large}	MedMT5 _{XL}
EN	NCBI-Disease	85.1	87.7	89.4	88.6	89.3	85.7	87.4	89.1	87.2
EN	BC5CDR Disease	78.5	81.4	85.4	85.0	85.8	82.5	84.3	84.4	82.4
EN	BC5CDR Chemical	89.1	90.8	93.3	92.0	92.9	91.1	92.9	92.8	91.3
EN	DIANN	70.1	77.8	71.9	74.4	74.2	80.3	79.0	74.8	77.6
ES	DIANN	72.4	74.9	70.5	70.7	70.9	78.3	70.2	74.9	74.8
EN	E3C	54.3	60.1	62.8	64.2	63.1	58.2	58.6	59.4	57.9
ES	E3C	61.6	71.7	62.7	64.4	67.1	65.9	57.4	72.2	69.5
FR	E3C	55.6	64.9	61.7	65.2	64.3	62.0	53.3	65.2	65.8
IT	E3C	61.8	63.8	59.6	61.9	65.1	63.9	52.1	67.5	65.9
ES	PharmaCoNER	86.3	90.6	87.5	88.5	89.1	89.4	88.6	90.8	90.1
EN	Neoplasm	70.4	71.1	74.4	74.3	73.4	64.5	67.5	73.9	73.2
EN	Glaucoma	70.7	75.1	77.1	78.4	78.0	71.2	74.8	76.2	76.4
EN	Mixed	68.5	73.0	73.4	73.2	74.5	63.4	69.6	72.2	72.0
ES	Neoplasm	69.0	56.1	71.4	72.5	73.9	63.0	57.1	72.1	71.8
ES	Glaucoma	69.3	70.7	73.9	73.8	75.2	68.6	64.5	77.1	75.5
ES	Mixed	68.4	66.2	69.2	69.3	71.6	61.3	58.9	72.4	71.4
FR	Neoplasm	70.5	66.6	74.0	72.4	73.7	63.9	59.0	72.9	71.2
FR	Glaucoma	71.1	69.2	77.8	74.8	77.2	60.3	65.6	79.5	75.8
FR	Mixed	68.3	65.4	72.0	70.9	74.3	64.1	61.3	73.3	69.7
IT	Neoplasm	68.1	69.9	70.1	70.9	72.0	64.4	54.8	71.2	73.1
IT	Glaucoma	69.2	71.5	73.7	74.0	75.9	74.7	65.8	75.7	78.7
IT	Mixed	66.3	67.7	67.4	69.9	70.0	61.3	57.4	70.6	71.9
AVERAGE		70.2	72.1	73.6	74.1	75.1	69.9	67.3	75.4	74.7
AVERAGE ES, FR, IT		68.4	69.2	70.8	71.4	72.9	67.2	61.9	74.0	73.2

Table 6.4 – Single-task supervised F1 scores for Sequence Labelling.

Single Task Monolingual Supervised Results: The results when fine-tuning and evaluating the models for each dataset and language are shown in Table 6.4. The first observation is that Medical-mT5-large significantly outperforms both mT5-large and mT5-XL, demonstrating the benefits of further training these models with our multilingual medical domain corpus.

When comparing Medical mT5 with FlanT5 and SciFive, the latter models are systematically superior on English. This was anticipated since both have been pre-trained with a much larger amount of English-only data specific to the medical domain. With respect to encoder-only models, they achieve in general worse results than text-to-text models across all tasks and languages (except for the DI-ANN dataset). It is also noteworthy that FlanT5-XL exhibits robust performance across all datasets and languages, even though it was fine-tuned with English-only data not specific to the medical domain. Nonetheless, Medical-mT5-large obtains in general better results for French, Spanish and Italian while being much smaller in size (738M parameters vs 3B parameters), showing the impact of training Medical mT5 with domain-specific data for those languages.

Multi-Task Supervised Results: Text-to-text models have demonstrated improved performance when trained in multi-task settings (Chung *et al.* 2022). Following this, we also experimented with fine-tuning them across all the sequence labeling tasks simultaneously. To inform the model about which labels should be classified for each input example, we add the list of predefined labels from the corresponding dataset to the beginning of the input sentence. For instance, the input depicted in Figure 6.3 is adjusted to “<Disease> Patient with dilated cardiomyopathy”. A comparison of the Single Task and Multi-Task settings is presented in Table 6.5. It can be seen that in this setting Medical mT5 achieves the best overall results for Spanish, French and Italian. On average, Medical-mT5-xl also obtains the best performance, slightly improving over the results of FlanT5-XL and Medical-mT5-large.

Zero-shot Cross-Lingual Transfer Results: Manually annotated medical domain datasets for languages other than English are scarce. Therefore, developing models that can successfully generate predictions for languages different to those used for fine-tuning is crucial. We evaluate this ability to perform zero-shot cross-lingual transfer by fine-tuning Medical mT5 and the baselines on the English AbsRCT Neoplasm dataset, and then evaluating them on the Neoplasm, Glaucoma, and Mixed datasets for Spanish, French, and Italian. The results

Lang	Dataset	Single Task			MultiTask		
		FlanT5 _{XL}	MedMT5 _{large}	MedMT5 _{XL}	FlanT5 _{XL}	MedMT5 _{large}	MedMT5 _{XL}
EN	NCBI-Disease	89.3	89.1	87.2	87.6	87.6	86.9
EN	BC5CDR Disease	85.8	84.4	82.4	85.1	83.4	83.0
EN	BC5CDR Chemical	92.9	92.8	91.3	92.7	92.5	91.6
EN	DIANN	74.2	74.8	77.6	80.0	75.4	75.3
ES	DIANN	70.9	74.9	74.8	77.1	72.6	73.6
EN	E3C	63.1	59.4	57.9	62.1	60.9	62.0
ES	E3C	67.1	72.2	69.5	66.5	74.9	73.3
FR	E3C	64.3	65.2	65.8	62.9	65.4	65.1
IT	E3C	65.1	67.5	65.9	60.7	66.9	65.1
ES	PharmaCoNER	89.1	90.8	90.1	89.9	90.3	89.5
EN	Neoplasm	73.4	73.9	73.2	73.1	72.3	72.9
EN	Glaucoma	78.0	76.2	76.4	76.4	76.8	77.5
EN	Mixed	74.5	72.2	72.0	71.5	70.9	73.0
ES	Neoplasm	73.9	72.1	71.8	73.5	73.5	73.7
ES	Glaucoma	75.2	77.1	75.5	77.1	77.7	79.3
ES	Mixed	71.6	72.4	71.4	70.0	71.8	72.8
FR	Neoplasm	73.7	72.9	71.2	74.0	72.9	73.6
FR	Glaucoma	77.2	79.5	75.8	76.6	77.0	79.4
FR	Mixed	74.3	73.3	69.7	71.8	71.2	73.0
IT	Neoplasm	72.0	71.2	73.1	71.9	74.6	74.0
IT	Glaucoma	75.9	75.7	78.7	77.6	78.5	78.9
IT	Mixed	70.0	70.6	71.9	69.9	72.5	73.3
AVERAGE		75.1	75.4	74.7	75.2	76.2	76.7
AVERAGE ES, FR, IT		72.9	74.0	73.2	73.1	74.8	75.3

Table 6.5 – Multi-task supervised F1 scores for Sequence Labelling.

are presented in Table 6.6. Results show that Medical mT5 outperforms any other model. Moreover, Medical-mT5-xl achieves significantly better results than Medical-mT5-large.

To summarize, Medical mT5 stands out for its superior performance in the evaluation for Spanish, French, and Italian languages, especially for the multi-task and the zero-shot transfer settings. These capabilities can help mitigate the scarcity of manually annotated medical data for other target languages. In contrast, SciFive and FlanT5, having been trained on extensive English-only datasets, emerge as the top choices when the primary focus is on English-only tasks.

Finally, despite Medical-mT5-xl being larger than Medical-mT5-large (3B vs 738M), its performance is worse in the single-task evaluation setting. This behaviour is not observed in the multi-task and zero-shot experiments, leading us to hypothesize that the larger Medical-mT5-xl model is more prone to overfit in the single-task supervised setting.

Lang	Dataset	mT5 _{XL}	SciFive	FlanT5 _{XL}	mDeBERTa _{v3 base}	MedMT5 _{large}	MedMT5 _{XL}
ES	Neoplasm	71.4	69.8	67.9	65.1	72.4	71.7
ES	Glaucoma	74.1	71.5	70.6	68.3	72.4	73.2
ES	Mixed	69.4	67.0	66.7	60.9	68.1	68.8
FR	Neoplasm	71.6	68.6	69.9	60.5	72.4	72.8
FR	Glaucoma	75.8	74.5	71.0	68.7	72.3	76.7
FR	Mixed	73.0	68.5	68.2	59.3	70.4	72.4
IT	Neoplasm	70.6	63.1	67.3	62.4	72.9	73.2
IT	Glaucoma	76.7	71.6	72.0	70.2	75.4	79.0
IT	Mixed	69.9	62.5	66.9	62.1	71.7	71.9
AVERAGE		72.5	68.6	69.0	64.2	72.0	73.3

Table 6.6 – Zero-shot F1 scores for Argument Mining. Models have been trained in English and evaluated in Spanish, French and Italian.

6.7.2 Abstractive Question Answering

In this section we explore the text generation capabilities of Medical mT5 and other baseline text-to-text models on the BioASQ question answering dataset described in Section 6.5.2. Previous work typically evaluates the performance on this task using the ROUGE score (Tsatsaronis *et al.* 2015) to compare the gold standard answer with the answer generated by the model. However, we find this metric inadequate for medical domain tasks as it does not address crucial aspects of the generation such as factuality, potential harm, and bias (Singhal *et al.* 2022). Consequently, we enlisted medical professionals to analyze the answers produced by the models.

During annotation, medical doctors were displayed the question, the ideal gold answers and the answers generated by each model. If required, they could also inspect the snippets that provide context to answer each of the questions. We narrowed the evaluation to Medical-mT5-large, mT5-large, FlanT5-large and SciFive. The evaluation was conducted by medical doctors proficient/native speakers of English, French and Spanish. For each question, doctors were asked to rank the answers generated by the models as the best, second-best, third-best, and worst answer.

Two Spanish medical doctors proficient or native in English and Spanish analyzed 50 English examples and 252 Spanish. For the French language, 3 French clinicians analyzed 186 answers, of which 47 were done by 2 doctors to calculate IAA (Cohen’s Kappa Score: 0.28 and Average Spearman’s Rank Correlation: 0.48), which indicates a low level of agreement. This exercise provided inter-

Ideal Answer	Which type of urinary incontinence is diagnosed with the Q tip test? Stress urinary incontinence is diagnosed with the Q tip test. The test evaluates urethral mobility.
mT5-large	Q-tip test is used for urodynamic assessment of patients with severe urinary incontinence.
FlanT5-large	The Q tip test is used for assessing urethral mobility.
SciFive	Q tip test is used for evaluation of stress urinary incontinence.
MedMT5-large	Q-tip test is used for evaluation of stress urinary incontinence.
Ideal Answer	Which are the main manifestations of Ohdo syndrome? Severe ID, absent or deficient language, skeletal manifestations including bilateral patella dislocations.
mT5-large	Skeletal manifestations in Ohdo syndrome are a case with bilateral patella dislocations where surgical intervention has been indicated.
FlanT5-large	The main manifestations of Ohdo syndrome are: 1) severe ID, 2) absent or deficient language and 3) milder, clinical manifestation in heterozygotes.
SciFive	Ohdo syndrome is characterized by severe ID, absent or deficient language and, milder, clinical manifestation in heterozygotes.
MedMT5-large	The main manifestations of Ohdo syndrome are: 1) absent or deficient language and 2) milder clinical manifestation in heterozygotes.

Table 6.7 – Examples of answers generated by each model for two different BioASQ questions together with the rank assigned by medics.

esting insights with respect to the performance of the models in text generation tasks in the medical domain. First, medical doctors could not in general establish significant differences between the answers generated by each of the models; predictions were far too similar, and all tended to fail on the same questions. As an example, Table 6.7 shows the answers to two different questions. As it can be observed, the answers generated by each model are very similar, and the doctors ended up ranking them primarily based on style.

The final result of the manual analysis is that all the models were chosen a similar number of times as the best. We believe that this demonstrates the difficulty of performing and obtaining meaningful evaluation results for this kind of tasks on this specific domain. This is supported by the low IAA agreement obtained in the French annotation. This issue has also emerged in prior research and was partially addressed by employing a very large number of experts and asking them to respond with a yes/no to a set of predefined potential issues in the model output (Singhal *et al.* 2022). Still, the variance on the answers provided by the experts was significant.

However, there could be other underlying reasons for this behaviour. First, perhaps the T5 architecture is not ideally suited for text generation as formulated in the BioASQ task, as these models are trained on a masking reconstruction objective rather than on direct text generation tasks. Consequently, the knowledge acquired during pre-training might not generalize well when the models are subsequently trained for text generation purposes. Second, perhaps using much larger

models such as MedPaLM (Singhal *et al.* 2022) may generate better answer generation, but models of 540B parameters are currently unusable for the large majority of the NLP research labs, including ours. Nonetheless, it should be stressed that research on appropriate evaluation metrics for these tasks is still a difficult challenge which requires further investigation.

In any case, our results demonstrate the potential of a text-to-text model such as Medical mT5 for multilingual sequence labeling in the medical domain, establishing new state-of-the-art results in the multi-task and zero-shot cross-lingual settings.

6.8 Conclusion

In this chapter, we have presented Medical mT5, the first open-source multilingual text-to-text LLM for the medical domain. Its development has required the compilation of a new 3B word corpus in English, French, Italian and Spanish specific to the medical domain. Furthermore, motivated by the lack of multilingual benchmarks, we have generated evaluation benchmarks for French, Italian and Spanish for Argument Mining and Abstractive Question Answering.

A comprehensive experimentation on sequence labeling tasks shows that Medical mT5 outperforms strong text-to-text baselines of similarly-sized models in the multi-task and zero-shot cross-lingual evaluation settings. This is particularly interesting as these settings fully exploit the multilingual nature of a text-to-text model such as Medical mT5.

Furthermore, our experiments on Abstractive Question Answering show the inherent difficulty of evaluating generative tasks for this specific domain, where complex issues such as truthfulness and veracity are difficult to capture by automatic metrics. Manual evaluation is not ideal either, as medical doctors were not able to clearly distinguish between the quality of the answers generated by the different models. In line with previous work (Singhal *et al.* 2022), we hope our research will bring further attention to this problem and encourage further research on evaluation methods.

Medical mT5 has been built on the work presented in previous chapters. We use the data transfer approach to develop new multilingual evaluation benchmarks for the medical domain. We also use constrained decoding to achieve high-quality zero-shot model-based cross-lingual transfer. Medical mT5 demonstrates the benefits of the techniques developed during this thesis and their application in real-world scenarios where data is scarce.

Regarding the languages chosen for this chapter, acquiring medical domain data is extremely challenging, even for languages such as the ones included. Furthermore, the choice of languages was also influenced by the availability of native medical doctors to do the manual evaluation for Abstractive Question Answering. In any case, we hope that our research will encourage more researchers to join our effort and gather data for their respective languages, thereby creating larger, multilingual medical domain datasets encompassing more languages in the future.

Conclusion and future work

In this thesis, we have developed novel cross-lingual transfer learning methods aimed at addressing the resource constraints of low-resource languages. By leveraging both data-based and model-based approaches, we have demonstrated the potential to significantly improve performance on sequence labeling tasks across diverse languages and domains. Our proposed methods, including T-Projection and constrained decoding algorithms, achieve state-of-the-art results, highlighting the effectiveness of modern Machine Translation and multilingual models in facilitating knowledge transfer. The real-world application to the medical domain further underscores the practical impact of our research. By contributing open-source tools, datasets, and models, this work not only bridges the gap between high-resource and low-resource languages but also sets the stage for future advancements in multilingual NLP. The main contributions of this thesis are summarized as follows:

- **We improved data-based cross-lingual transfer** approaches by developing a novel annotation projection method, namely **T-Projection** (García-Ferrero *et al.* 2023) in Chapter 4. It leverages state-of-the-art text-to-text multilingual models and Machine Translation systems to project annotations from high-resource to low-resource languages. T-Projection significantly outperforms previous annotation projection methods by a wide margin. This method allows us to **automatically generate high-quality labeled data for low-resource languages**.
- **We enhanced the model-based cross-lingual transfer approach** by, in

Chapter 3.4 proposing a **constrained decoding algorithm** that enables model-based **cross-lingual transfer for sequence labeling tasks with text-to-text models**. This algorithm allows for the use of more powerful models that demonstrate superior zero-shot model-based transfer capabilities. Given the prevailing focus on text-to-text large language models (LLMs) in current research, and the infrequent training of new encoder-only models, this represents significant progress in the field.

- We expanded **NLP research in the medical domain for more languages** by developing a multilingual text-to-text open-source model for the medical domain, called **Medical mT5** (García-Ferrero *et al.* 2024), presented in Chapter 6. By applying the model-based and data-based cross-lingual transfer learning methods developed in this thesis, we have implemented a **multilingual pre-training, fine-tuning, and evaluation framework for the medical domain**. Medical mT5 demonstrates the importance of the technology and knowledge developed in this thesis, resulting in the **first multilingual text-to-text medical model** at the time it was created.
- We conducted a **comprehensive evaluation of different cross-lingual transfer learning methods across a wide range of tasks, languages, and domains** (García-Ferrero *et al.* 2022), contributing to a better understanding of the situations in which each method is most effective. We demonstrated that **both our proposed data-based and model-based methods are effective in different scenarios**, and that they can be combined to achieve even better results. We also showed that the proposed methods are robust across different languages and domains and can be easily adapted to new tasks and languages.
- We released a large collection of **open-source software, datasets, and models** to facilitate the development of multilingual NLP research. By making our work **easily accessible** to other researchers, we enable them to **replicate** our experiments and **build upon our work**. We expect that the insights and methods developed in this thesis will be applicable to a wide range of NLP tasks, languages, and domains, thus contributing to the advancement of NLP in low-resource languages.

In terms of **publications**, this thesis contains 3 papers published in international conferences: 2 at EMNLP and 1 at LREC-COLING. Additionally, we published 3 closely related papers that were not included in this manuscript, 1 at

ICLR, 1 at EMLP and 1 at an ACL Workshop. Finally, we submitted other peer-reviewed papers, including 1 EMNLP, 2 at IkerGazte, 1 at ACL Workshops, and 1 at SEPLN. Between these papers, the paper *Twitterreko Euskal Komunitatearen Eduki Azterketa Pandemia Garaian* was awarded the Most Relevant Research for the Development of the Basque Country **award at IkerGazte 2021**. Additionally, the *Noticia: A Clickbait Article Summarization Dataset in Spanish* project was selected as the **winner of the #Somos600M 2024 Hackathon** by the SomosNLP community. This thesis has created interest in the community, as evidenced by the more than **250 citations** of the papers published during this PhD, more than **56,000 downloads** of the open-source models and datasets released on the Hugging Face Hub and more than **550 stars** on the GitHub repositories. It also demonstrated by the **talks** given at the OntarioTech University (graduate NLP course), the SEPLN symposium, and the Universitat de Barcelona (AI4HF consortium meeting). Finally, outside of the scope of the thesis, we have also contributed to the fair and unbiased evaluation of models by organizing The 1st Workshop on Data Contamination (CONDA) at ACL 2024 (<https://conda-workshop.github.io/>, Sainz *et al.* 2024b).

7.1 Future work

With the introduction of Large Language Models (LLMs), which have been pre-trained on massive amounts of text, instructions and further improved with reinforcement learning such as GPT4 (OpenAI *et al.* 2024) or LLama-3 (AI@Meta 2024), the field of NLP is evolving rapidly. The focus has shifted from training task-specific models to training multi-task models that can handle a wide range of tasks given a task description or prompt. Multi-task models have proven to outperform tasks specific models (Tran *et al.* 2021) and can even perform well on tasks that they have not been trained on (Brown *et al.* 2020).

While these models are revolutionizing the field, they require a large amount of data, in terms of unstructured text and, more importantly, they require high-quality instruction-tuning data to achieve the capability of performing tasks given a prompt. This data includes a huge variety of tasks formulated as text-to-text instructions, such as dialogues, examples or summarization, code writing, translation, text generation, etc. Building large-scale instruction-tuning datasets is a very expensive and time-consuming process, therefore for now it has been limited to a very few companies with the monetary resources to do so.

In this context, the next step in the field of cross-lingual transfer learning is to

develop instruction-tuning datasets for low-resource languages. This requires the development of both, data-transfer and model-transfer methods, as well as overcoming other challenges such as adapting LLMs to the culture of the community in which they are to be used. Some of the future work that we plan to explore are:

- Exploring the use of Machine Translation to generate instruction-tuning data for low-resource languages based on the already existing instruction-tuning datasets in high-resource languages. However, current sentence-level Machine Translation systems such as M2M100 (Aharoni *et al.* 2019) or NLLB200 (Costa-jussà *et al.* 2022) are not able to handle long contexts, which is a requirement for instruction-tuning data, which can be very long. In addition, translating complex structures that mix code, mathematical formulas, and other elements with natural language remains a challenge for current Machine Translation systems. While LLMs have shown great proficiency at document-level translation (Xu *et al.* 2023; Alves *et al.* 2024) they are still only proficient at translating between high-resource languages (Ojo *et al.* 2023; Ojo and Ogueji 2023). Developing a new generation of long-context Machine Translation systems, able to handle complex structures, for a wide range of languages can enable the generation of instruction-tuning data for low-resource languages.
- Synthetic data generation using LLMs (Teknum 2023) has shown promising results (Zou *et al.* 2023; Liu *et al.* 2024). This process involves using an already pre-trained LLM and a set of prompts to generate instruction-tuning data for a wide range of tasks. This synthetic data can then be used to instruction-tune a new model that is superior to the original LLM. However, current methods are still limited to English. Model-based cross-lingual transfer can be used to generate synthetic data for low-resource languages. This means that a model pre-trained with unstructured text from many languages and instruction-tuned in only a few high-resource languages may be able to generate synthetic data for all the languages it has been pre-trained on. Similar to the model-based cross-lingual transfer experiments in this thesis, the model can be fine-tuned on a few examples of the target language and then used to generate synthetic data. This synthetic data can be used to train a model for the target language. A better understanding of the model-based cross-lingual transfer methods with LLMs and synthetic data generation methods can enable the development of instruction-tuning datasets for low-resource languages.

- Building LLMs that can receive text in a language and produce plausible output text in that language is only the first step in developing LLMs for low-resource languages. For example, in the case of the Basque language, the latest generation of LLMs can process text in Basque. However, the models fail to answer questions pertinent to Basque culture, while they correctly answer questions about global culture (Etxaniz *et al.* 2024a). For an LLM to be useful for a community, it must not only be able to process text in that language and produce plausible output text in the same language, but it must also encode knowledge about the community’s culture. As shown in the experiments in Chapter 3, model-based and data-based cross-lingual transfer methods are not enough to overcome this issue. In fact, training the model with English-translated data can exacerbate the problem. Therefore, future research in cross-lingual transfer should not only focus on developing LLMs that can process text in a wide range of languages but also on finding methods to efficiently teach the models about the culture of the community in which the LLM is to be used.

Bibliography

Adelani D., Neubig G., Ruder S., Rijhwani S., Beukman M., Palen-Michel C., Lignos C., Alabi J., Muhammad S., Nabende P., Dione C.M.B., Bukula A., Mabuya R., Dossou B.F.P., Sibanda B., Buzaaba H., Mukiibi J., Kalipe G., Mbaye D., Taylor A., Kabore F., Emezue C.C., Aremu A., Ogayo P., Gitau C., Munkoh-Buabeng E., Memdjokam Koagne V., Tapo A.A., Macucwa T., Marivate V., Elvis M.T., Gwadabe T., Adewumi T., Ahia O., Nakatumba-Nabende J., Mokono N.L., Ezeani I., Chukwuneke C., Oluwaseun Adeyemi M., Hacheme G.Q., Abdulmumin I., Ogundepo O., Yousuf O., Moteu T., and Klakow D. MasakhaNER 2.0: Africa-centric transfer learning for named entity recognition. In Goldberg Y., Kozareva Z., and Zhang Y., editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 4488–4508, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.emnlp-main.298>.

Agerri R. and Agirre E. Lessons learned from the evaluation of Spanish Language Models. *Proces. del Leng. Natural*, 70:157–170, 2023.

Agerri R., Bermudez J., and Rigau G. IXA pipeline: Efficient and ready to use multilingual NLP tools. In Calzolari N., Choukri K., Declerck T., Loftsson H., Maegaard B., Mariani J., Moreno A., Odijk J., and Piperidis S., editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, 3823–3828. European Language Resources Association (ELRA), 2014. URL <http://www.lrec-conf.org/proceedings/lrec2014/summaries/775.html>.

Agerri R., Chung Y., Aldabe I., Aranberri N., Labaka G., and Rigau G. Building

BIBLIOGRAPHY

- named entity recognition taggers via parallel corpora. In Calzolari N., Choukri K., Cieri C., Declerck T., Goggi S., Hasida K., Isahara H., Maegaard B., Mariani J., Mazo H., Moreno A., Odijk J., Piperidis S., and Tokunaga T., editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA). URL <https://aclanthology.org/L18-1557>.
- Agerri R. and Rigau G. Language independent sequence labelling for opinion target extraction. *Artificial Intelligence*, 268:85–95, 2019. ISSN 0004-3702. URL <https://www.sciencedirect.com/science/article/pii/S0004370218307392>.
- Aharoni R., Johnson M., and Firat O. Massively multilingual neural machine translation. In Burstein J., Doran C., and Solorio T., editors, *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)*, 3874–3884, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. URL <https://aclanthology.org/N19-1388>.
- AI ., :, Young A., Chen B., Li C., Huang C., Zhang G., Zhang G., Li H., Zhu J., Chen J., Chang J., Yu K., Liu P., Liu Q., Yue S., Yang S., Yang S., Yu T., Xie W., Huang W., Hu X., Ren X., Niu X., Nie P., Xu Y., Liu Y., Wang Y., Cai Y., Gu Z., Liu Z., and Dai Z. Yi: Open foundation models by 01.ai, 2024. URL <https://arxiv.org/abs/2403.04652>.
- AI@Meta. The llama 3 herd of models, 2024. URL <https://ai.meta.com/research/publications/the-llama-3-herd-of-models/>.
- Akbik A., Blythe D., and Vollgraf R. Contextual string embeddings for sequence labeling. In Bender E.M., Derczynski L., and Isabelle P., editors, *Proceedings of the 27th International Conference on Computational Linguistics*, 1638–1649, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://aclanthology.org/C18-1139>.
- Alabi J.O., Adelani D.I., Mosbach M., and Klakow D. Adapting pre-trained language models to African languages via multilingual adaptive fine-tuning. *Proceedings of the 29th International Conference on Computational Linguistics*,

- 4336–4349, Gyeongju, Republic of Korea, October 2022. International Committee on Computational Linguistics. URL <https://aclanthology.org/2022.coling-1.382>.
- Alves D.M., Pombal J., Guerreiro N.M., Martins P.H., Alves J., Farajian A., Peters B., Rei R., Fernandes P., Agrawal S., Colombo P., de Souza J.G.C., and Martins A.F.T. Tower: An open multilingual large language model for translation-related tasks, 2024.
- Artetxe M., Labaka G., and Agirre E. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In Su J., Duh K., and Carreras X., editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2289–2294, Austin, Texas, November 2016. Association for Computational Linguistics. URL <https://aclanthology.org/D16-1250>.
- Artetxe M. and Schwenk H. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7:597–610, 2019. URL <https://aclanthology.org/Q19-1038>.
- Aryabumi V., Dang J., Talupuru D., Dash S., Cairuz D., Lin H., Venkitesh B., Smith M., Campos J.A., Tan Y.C., Marchisio K., Bartolo M., Ruder S., Locatelli A., Kreutzer J., Frosst N., Gomez A.N., Blunsom P., Fadaee M., Üstün A., and Hooker S. Aya 23: Open weight releases to further multilingual progress. *CoRR*, abs/2405.15032, 2024. URL <https://doi.org/10.48550/arXiv.2405.15032>.
- Ashok D. and Lipton Z.C. Promptner: Prompting for named entity recognition. *CoRR*, abs/2305.15444, 2023. URL <https://doi.org/10.48550/arXiv.2305.15444>.
- Bawden R., Bretonnel Cohen K., Grozea C., Jimeno Yepes A., Kittner M., Krallinger M., Mah N., Neveol A., Neves M., Soares F., Siu A., Verspoor K., and Vicente Navarro M. Findings of the WMT 2019 biomedical translation shared task: Evaluation for MEDLINE abstracts and biomedical terminologies. *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, 29–53, Florence, Italy, 2019. Association for Computational Linguistics. URL <https://aclanthology.org/W19-5403>.

BIBLIOGRAPHY

- Bel N., Koster C.H.A., and Villegas M. Cross-lingual text categorization. In Koch T. and Sølvsberg I.T., editors, *Research and Advanced Technology for Digital Libraries*, 126–139, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-45175-4.
- Beltagy I., Lo K., and Cohan A. SciBERT: A Pretrained Language Model for Scientific Text. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3615–3620, 2019.
- Bharadwaj A., Mortensen D., Dyer C., and Carbonell J. Phonologically aware neural model for named entity recognition in low resource transfer settings. In Su J., Duh K., and Carreras X., editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1462–1472, Austin, Texas, November 2016. Association for Computational Linguistics. URL <https://aclanthology.org/D16-1153>.
- Blasi D., Anastasopoulos A., and Neubig G. Systematic inequalities in language technology performance across the world’s languages. In Muresan S., Nakov P., and Villavicencio A., editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5486–5505, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.376>.
- Blevins T., Gonen H., and Zettlemoyer L. Prompting language models for linguistic structure. In Rogers A., Boyd-Graber J.L., and Okazaki N., editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, 6649–6663. Association for Computational Linguistics, 2023. URL <https://doi.org/10.18653/v1/2023.acl-long.367>.
- Bojanowski P., Grave E., Joulin A., and Mikolov T. Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguistics*, 5:135–146, 2017. URL https://doi.org/10.1162/tacl_a_00051.
- Bommasani R., Hudson D.A., Adeli E., Altman R.B., Arora S., von Arx S., Bernstein M.S., Bohg J., Bosselut A., Brunskill E., Brynjolfsson E., Buch S., Card D., Castellon R., Chatterji N.S., Chen A.S., Creel K., Davis J.Q., Demszky D., Donahue C., Doumbouya M., Durmus E., Ermon S., Etchemendy J., Ethayarajh K., Fei-Fei L., Finn C., Gale T., Gillespie L., Goel K., Goodman N.D.,

- Grossman S., Guha N., Hashimoto T., Henderson P., Hewitt J., Ho D.E., Hong J., Hsu K., Huang J., Icard T., Jain S., Jurafsky D., Kalluri P., Karamcheti S., Keeling G., Khani F., Khattab O., Koh P.W., Krass M.S., Krishna R., Kudipudi R., and et al. On the opportunities and risks of foundation models. *CoRR*, abs/2108.07258, 2021. URL <https://arxiv.org/abs/2108.07258>.
- Brown P.F., Della Pietra S.A., Della Pietra V.J., and Mercer R.L. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263–311, 1993. URL <https://aclanthology.org/J93-2003>.
- Brown T., Mann B., Ryder N., Subbiah M., Kaplan J.D., Dhariwal P., Neelakantan A., Shyam P., Sastry G., Askell A., et al.. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Cao N.D., Izacard G., Riedel S., and Petroni F. Autoregressive entity retrieval. *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=5k8F6UU39V>.
- Carrino C.P., Llop J., Pàmies M., Gutiérrez-Fandiño A., Armengol-Estapé J., Silveira-Ocampo J., Valencia A., Gonzalez-Agirre A., and Villegas M. Pre-trained biomedical language models for clinical NLP in Spanish. *Proceedings of the 21st Workshop on Biomedical Language Processing*, 193–199, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.bionlp-1.19>.
- Chaudhary A., Zhou C., Levin L., Neubig G., Mortensen D.R., and Carbonell J. Adapting word embeddings to new languages with morphological and phonological subword representations. In Riloff E., Chiang D., Hockenmaier J., and Tsujii J., editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3285–3295, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. URL <https://aclanthology.org/D18-1366>.
- Chen Y., Jiang C., Ritter A., and Xu W. Frustratingly easy label projection for cross-lingual transfer. In Rogers A., Boyd-Graber J., and Okazaki N., editors,

BIBLIOGRAPHY

- Findings of the Association for Computational Linguistics: ACL 2023*, 5775–5796, Toronto, Canada, July 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.findings-acl.357>.
- Chowdhery A., Narang S., Devlin J., Bosma M., Mishra G., Roberts A., Barham P., Chung H.W., Sutton C., Gehrmann S., Schuh P., Shi K., Tsvyashchenko S., Maynez J., Rao A., Barnes P., Tay Y., Shazeer N., Prabhakaran V., Reif E., Du N., Hutchinson B., Pope R., Bradbury J., Austin J., Isard M., Gur-Ari G., Yin P., Duke T., Levskaya A., Ghemawat S., Dev S., Michalewski H., Garcia X., Misra V., Robinson K., Fedus L., Zhou D., Ippolito D., Luan D., Lim H., Zoph B., Spiridonov A., Sepassi R., Dohan D., Agrawal S., Omernick M., Dai A.M., Pillai T.S., Pellat M., Lewkowycz A., Moreira E., Child R., Polozov O., Lee K., Zhou Z., Wang X., Saeta B., Diaz M., Firat O., Catasta M., Wei J., Meier-Hellstern K., Eck D., Dean J., Petrov S., and Fiedel N. Palm: Scaling language modeling with pathways. *J. Mach. Learn. Res.*, 24:240:1–240:113, 2023. URL <http://jmlr.org/papers/v24/22-1144.html>.
- Chung H.W., Hou L., Longpre S., Zoph B., Tay Y., Fedus W., Li E., Wang X., Dehghani M., Brahma S., Webson A., Gu S.S., Dai Z., Suzgun M., Chen X., Chowdhery A., Narang S., Mishra G., Yu A., Zhao V.Y., Huang Y., Dai A.M., Yu H., Petrov S., Chi E.H., Dean J., Devlin J., Roberts A., Zhou D., Le Q.V., and Wei J. Scaling instruction-finetuned language models. *CoRR*, abs/2210.11416, 2022. URL <https://doi.org/10.48550/arXiv.2210.11416>.
- Conneau A., Khandelwal K., Goyal N., Chaudhary V., Wenzek G., Guzmán F., Grave E., Ott M., Zettlemoyer L., and Stoyanov V. Unsupervised cross-lingual representation learning at scale. In Jurafsky D., Chai J., Schluter N., and Tetreault J., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8440–8451, Online, July 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.acl-main.747>.
- Conneau A. and Lample G. Cross-lingual language model pretraining. In Wallach H.M., Larochelle H., Beygelzimer A., d’Alché-Buc F., Fox E.B., and Garnett R., editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 7057–7067, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/c04c19c2c2474dbf5f7ac4372c5b9af1-Abstract.html>.

- Conneau A., Rinott R., Lample G., Williams A., Bowman S., Schwenk H., and Stoyanov V. XNLI: Evaluating cross-lingual sentence representations. In Riloff E., Chiang D., Hockenmaier J., and Tsujii J., editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2475–2485, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. URL <https://aclanthology.org/D18-1269>.
- Costa-jussà M.R., Cross J., Çelebi O., Elbayad M., Heafield K., Heffernan K., Kalbassi E., Lam J., Licht D., Maillard J., Sun A., Wang S., Wenzek G., Youngblood A., Akula B., Barrault L., Gonzalez G.M., Hansanti P., Hoffman J., Jarrett S., Sadagopan K.R., Rowe D., Spruit S., Tran C., Andrews P., Ayan N.F., Bhosale S., Edunov S., Fan A., Gao C., Goswami V., Guzmán F., Koehn P., Mourachko A., Ropers C., Saleem S., Schwenk H., and Wang J. No language left behind: Scaling human-centered machine translation. *CoRR*, abs/2207.04672, 2022. URL <https://doi.org/10.48550/arXiv.2207.04672>.
- Cui L., Wu Y., Liu J., Yang S., and Zhang Y. Template-based named entity recognition using BART. *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 1835–1845, Online, August 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.findings-acl.161>.
- Daza A. and Frank A. Translate and label! an encoder-decoder approach for cross-lingual semantic role labeling. In Inui K., Jiang J., Ng V., and Wan X., editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 603–615, Hong Kong, China, November 2019. Association for Computational Linguistics. URL <https://aclanthology.org/D19-1056>.
- De Cao N., Wu L., Popat K., Artetxe M., Goyal N., Plekhanov M., Zettlemoyer L., Cancedda N., Riedel S., and Petroni F. Multilingual autoregressive entity linking. *Transactions of the Association for Computational Linguistics*, 10: 274–290, 2022. URL <https://aclanthology.org/2022.tacl-1.16>.
- Dettmers T., Pagnoni A., Holtzman A., and Zettlemoyer L. Qlora: Efficient finetuning of quantized llms. In Oh A., Naumann T., Globerson A., Saenko

BIBLIOGRAPHY

- K., Hardt M., and Levine S., editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/1feb87871436031bdc0f2beaa62a049b-Abstract-Conference.html.
- Devlin J., Chang M.W., Lee K., and Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. In Burstein J., Doran C., and Solorio T., editors, *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)*, 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. URL <https://aclanthology.org/N19-1423>.
- Dogan R.I., Leaman R., and Lu Z. NCBI disease corpus: A resource for disease name recognition and concept normalization. *J. Biomed. Informatics*, 47:1–10, 2014. URL <https://doi.org/10.1016/j.jbi.2013.12.006>.
- Dou Z.Y. and Neubig G. Word alignment by fine-tuning embeddings on parallel corpora. In Merlo P., Tiedemann J., and Tsarfaty R., editors, *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 2112–2128, Online, April 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.eacl-main.181>.
- Dubey A., Jauhri A., Pandey A., Kadian A., Al-Dahle A., Letman A., Mathur A., Schelten A., Yang A., Fan A., Goyal A., Hartshorn A., Yang A., Mitra A., Sravankumar A., Korenev A., Hinsvark A., Rao A., Zhang A., Rodriguez A., Gregerson A., Spataru A., Rozière B., Biron B., Tang B., Chern B., Caucheteux C., Nayak C., Bi C., Marra C., McConnell C., Keller C., Touret C., Wu C., Wong C., Ferrer C.C., Nikolaidis C., Allonsius D., Song D., Pintz D., Livshits D., Esiobu D., Choudhary D., Mahajan D., Garcia-Olano D., Perino D., Hupkes D., Lakomkin E., AlBadawy E., Lobanova E., Dinan E., Smith E.M., Radenovic F., Zhang F., Synnaeve G., Lee G., Anderson G.L., Nail G., Mialon G., Pang G., Cucurell G., Nguyen H., Korevaar H., Xu H., Touvron H., Zarov I., Ibarra I.A., Kloumann I.M., Misra I., Evtimov I., Copet J., Lee J., Geffert J., Vranes J., Park J., Mahadeokar J., Shah J., van der Linde J., Billock J., Hong J., Lee J., Fu J., Chi J., Huang J., Liu J., Wang J., Yu J., Bitton J., Spisak J., Park J., Rocca J.,

- Johnstun J., Saxe J., Jia J., Alwala K.V., Upasani K., Plawiak K., Li K., Heafield K., Stone K., and et al. The llama 3 herd of models. *CoRR*, abs/2407.21783, 2024. URL <https://doi.org/10.48550/arXiv.2407.21783>.
- Dyer C., Chahuneau V., and Smith N.A. A simple, fast, and effective reparameterization of IBM model 2. In Vanderwende L., Daumé III H., and Kirchhoff K., editors, *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 644–648, Atlanta, Georgia, June 2013a. Association for Computational Linguistics. URL <https://aclanthology.org/N13-1073>.
- Dyer C., Chahuneau V., and Smith N.A. A simple, fast, and effective reparameterization of IBM model 2. In Vanderwende L., Daumé III H., and Kirchhoff K., editors, *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 644–648, Atlanta, Georgia, June 2013b. Association for Computational Linguistics. URL <https://aclanthology.org/N13-1073>.
- Dyer C., Kuncoro A., Ballesteros M., and Smith N.A. Recurrent neural network grammars. In Knight K., Nenkova A., and Rambow O., editors, *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, 199–209. The Association for Computational Linguistics, 2016. URL <https://doi.org/10.18653/v1/n16-1024>.
- Ehrmann M., Turchi M., and Steinberger R. Building a multilingual named entity-annotated corpus using annotation projection. *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, 118–124, Hissar, Bulgaria, 2011. Association for Computational Linguistics. URL <https://aclanthology.org/R11-1017>.
- Esplà M., Forcada M., Ramírez-Sánchez G., and Hoang H. ParaCrawl: Web-scale parallel corpora for the languages of the EU. *Proceedings of Machine Translation Summit XVII: Translator, Project and User Tracks*, 118–119, Dublin, Ireland, 2019. European Association for Machine Translation. URL <https://aclanthology.org/W19-6721>.
- Etxaniz J., Azkune G., Soroa A., de Lacalle O.L., and Artetxe M. BertaQA: How

BIBLIOGRAPHY

- much do language models know about local culture?, 2024a. URL <https://openreview.net/forum?id=QocjHRR31U>.
- Etzaniz J., Azkune G., Soroa A., Lacalle O., and Artetxe M. Do multilingual language models think better in English? In Duh K., Gomez H., and Bethard S., editors, *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, 550–564, Mexico City, Mexico, June 2024b. Association for Computational Linguistics. URL <https://aclanthology.org/2024.naacl-short.46>.
- Fabregat H., Martínez-Romo J., and Araujo L. Overview of the DIANN task: Disability annotation task. In Rosso P., Gonzalo J., Martínez R., Montalvo S., and de Albornoz J.C., editors, *Proceedings of the Third Workshop on Evaluation of Human Language Technologies for Iberian Languages (IberEval 2018) co-located with 34th Conference of the Spanish Society for Natural Language Processing (SEPLN 2018), Sevilla, Spain, September 18th, 2018*, 2150 lib. of CEUR Workshop Proceedings, 1–14. CEUR-WS.org, 2018. URL <https://ceur-ws.org/Vol-2150/overview-diann-task.pdf>.
- Fan A., Bhosale S., Schwenk H., Ma Z., El-Kishky A., Goyal S., Baines M., Celebi O., Wenzek G., Chaudhary V., Goyal N., Birch T., Liptchinsky V., Edunov S., Auli M., and Joulin A. Beyond english-centric multilingual machine translation. *J. Mach. Learn. Res.*, 22:107:1–107:48, 2021. URL <http://jmlr.org/papers/v22/20-1307.html>.
- Fei H., Zhang M., and Ji D. Cross-lingual semantic role labeling with high-quality translated training corpus. In Jurafsky D., Chai J., Schluter N., and Tetreault J., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7014–7026, Online, July 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.acl-main.627>.
- Fetahu B., Kar S., Chen Z., Rokhlenko O., and Malmasi S. Semeval-2023 task 2: Fine-grained multilingual named entity recognition (multiconer 2). In Ojha A.K., Dogruöz A.S., Martino G.D.S., Madabushi H.T., Kumar R., and Sartori E., editors, *Proceedings of the The 17th International Workshop on Semantic Evaluation, SemEval@ACL 2023, Toronto, Canada, 13-14 July 2023*, 2247–2265. Association for Computational Linguistics, 2023. URL <https://doi.org/10.18653/v1/2023.semeval-1.310>.

- Forney G. The viterbi algorithm. *Proceedings of the IEEE*, 61(3):268–278, 1973. URL <https://ieeexplore.ieee.org/document/1450960>.
- Fu R., Qin B., and Liu T. Generating Chinese named entity data from a parallel corpus. In Wang H. and Yarowsky D., editors, *Proceedings of 5th International Joint Conference on Natural Language Processing*, 264–272, Chiang Mai, Thailand, November 2011. Asian Federation of Natural Language Processing. URL <https://aclanthology.org/I11-1030>.
- García-Ferrero I., Agerri R., Atutxa Salazar A., Cabrio E., de la Iglesia I., Lavelli A., Magnini B., Molinet B., Ramirez-Romero J., Rigau G., Villa-Gonzalez J.M., Villata S., and Zaninello A. MedMT5: An open-source multilingual text-to-text LLM for the medical domain. In Calzolari N., Kan M.Y., Hoste V., Lenci A., Sakti S., and Xue N., editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 11165–11177, Torino, Italia, May 2024. ELRA and ICCL. URL <https://aclanthology.org/2024.lrec-main.974>.
- García-Ferrero I., Agerri R., and Rigau G. Model and data transfer for cross-lingual sequence labelling in zero-resource settings. In Goldberg Y., Kozareva Z., and Zhang Y., editors, *Findings of the Association for Computational Linguistics: EMNLP 2022*, 6403–6416, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.findings-emnlp.478>.
- García-Ferrero I., Agerri R., and Rigau G. T-projection: High quality annotation projection for sequence labeling tasks. In Bouamor H., Pino J., and Bali K., editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, 15203–15217, Singapore, December 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.findings-emnlp.1015>.
- Garg S., Peitz S., Nallasamy U., and Paulik M. Jointly learning to align and translate with transformer models. In Inui K., Jiang J., Ng V., and Wan X., editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4453–4462, Hong Kong, China, November 2019. Association for Computational Linguistics. URL <https://aclanthology.org/D19-1453>.

BIBLIOGRAPHY

- Geng S., Josifoski M., Peyrard M., and West R. Grammar-constrained decoding for structured NLP tasks without finetuning. In Bouamor H., Pino J., and Bali K., editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, 10932–10952. Association for Computational Linguistics, 2023. URL <https://doi.org/10.18653/v1/2023.emnlp-main.674>.
- Gonzalez-Agirre A., Marimon M., Intxaurreondo A., Rabal O., Villegas M., and Krallinger M. PharmaCoNER: Pharmacological substances, compounds and proteins named entity recognition track. In Kim J., Nédellec C., Bossy R., and Deléger L., editors, *Proceedings of The 5th Workshop on BioNLP Open Shared Tasks, BioNLP-OST@EMNLP-IJNCLP 2019, Hong Kong, China, November 4, 2019*, 1–10. Association for Computational Linguistics, 2019. URL <https://doi.org/10.18653/v1/D19-5701>.
- Gouws S., Bengio Y., and Corrado G. Bilbowa: Fast bilingual distributed representations without word alignments. In Bach F.R. and Blei D.M., editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, 37 lib. of *JMLR Workshop and Conference Proceedings*, 748–756. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/gouws15.html>.
- Goyal N., Du J., Ott M., Anantharaman G., and Conneau A. Larger-scale transformers for multilingual masked language modeling. *CoRR*, abs/2105.00572, 2021. URL <https://arxiv.org/abs/2105.00572>.
- Gu Y., Tinn R., Cheng H., Lucas M., Usuyama N., Liu X., Naumann T., Gao J., and Poon H. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Heal.*, 3(1):2:1–2:23, 2022. URL <https://doi.org/10.1145/3458754>.
- Guo R. and Roth D. Constrained labeled data generation for low-resource named entity recognition. In Zong C., Xia F., Li W., and Navigli R., editors, *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, ACL/IJCNLP 2021 lib. of *Findings of ACL*, 4519–4533. Association for Computational Linguistics, 2021. URL <https://doi.org/10.18653/v1/2021.findings-acl.396>.
- He H. and Choi J.D. Unleashing the true potential of sequence-to-sequence mod-

- els for sequence tagging and structure parsing. *CoRR*, abs/2302.02275, 2023. URL <https://doi.org/10.48550/arXiv.2302.02275>.
- He P., Liu X., Gao J., and Chen W. Deberta: decoding-enhanced bert with disentangled attention. *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=XPZIaotutsD>.
- Hoffmann J., Borgeaud S., Mensch A., Buchatskaya E., Cai T., Rutherford E., de Las Casas D., Hendricks L.A., Welbl J., Clark A., Hennigan T., Noland E., Millican K., van den Driessche G., Damoc B., Guy A., Osindero S., Simonyan K., Elsen E., Rae J.W., Vinyals O., and Sifre L. Training compute-optimal large language models. *CoRR*, abs/2203.15556, 2022. URL <https://doi.org/10.48550/arXiv.2203.15556>.
- Hu E.J., Shen Y., Wallis P., Allen-Zhu Z., Li Y., Wang S., and Chen W. Lora: Low-rank adaptation of large language models. *CoRR*, abs/2106.09685, 2021. URL <https://arxiv.org/abs/2106.09685>.
- Huang K.H., Hsu I.H., Natarajan P., Chang K.W., and Peng N. Multilingual generative language models for zero-shot cross-lingual event argument extraction. In Muresan S., Nakov P., and Villavicencio A., editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4633–4646, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.317>.
- Hwa R., Resnik P., Weinberg A., Cabezas C., and Kolak O. Bootstrapping parsers via syntactic projection across parallel texts. *Natural language engineering*, 11(3):311–325, 2005.
- Imani A., Lin P., Kargaran A.H., Severini S., Sabet M.J., Kassner N., Ma C., Schmid H., Martins A.F.T., Yvon F., and Schütze H. Glot500: Scaling multilingual corpora and language models to 500 languages. In Rogers A., Boyd-Graber J.L., and Okazaki N., editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, 1082–1117. Association for Computational Linguistics, 2023. URL <https://doi.org/10.18653/v1/2023.acl-long.61>.

BIBLIOGRAPHY

- Jain A., Paranjape B., and Lipton Z.C. Entity projection via machine translation for cross-lingual NER. In Inui K., Jiang J., Ng V., and Wan X., editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 1083–1092. Association for Computational Linguistics, 2019. URL <https://doi.org/10.18653/v1/D19-1100>.
- Jalili Sabet M., Dufter P., Yvon F., and Schütze H. SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings. In Cohn T., He Y., and Liu Y., editors, *Findings of the Association for Computational Linguistics: EMNLP 2020*, 1627–1643, Online, November 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.findings-emnlp.147>.
- Jiang A.Q., Sablayrolles A., Mensch A., Bamford C., Chaplot D.S., de las Casas D., Bressand F., Lengyel G., Lample G., Saulnier L., Lavaud L.R., Lachaux M.A., Stock P., Scao T.L., Lavril T., Wang T., Lacroix T., and Sayed W.E. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.
- Jimeno-Yepes A., N  v  ol A., Neves M.L., Verspoor K., Bojar O., Boyer A., Grozea C., Haddow B., Kittner M., Lichtblau Y., Pecina P., Roller R., Rosa R., Siu A., Thomas P., and Trescher S. Findings of the WMT 2017 biomedical translation shared task. In Bojar O., Buck C., Chatterjee R., Federmann C., Graham Y., Haddow B., Huck M., Jimeno-Yepes A., Koehn P., and Kreutzer J., editors, *Proceedings of the Second Conference on Machine Translation, WMT 2017, Copenhagen, Denmark, September 7-8, 2017*, 234–247. Association for Computational Linguistics, 2017. URL <https://doi.org/10.18653/v1/w17-4719>.
- Joshi P., Santy S., Budhiraja A., Bali K., and Choudhury M. The state and fate of linguistic diversity and inclusion in the NLP world. In Jurafsky D., Chai J., Schluter N., and Tetreault J., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 6282–6293, Online, July 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.acl-main.560>.
- Koehn P. Europarl: A parallel corpus for statistical machine translation. *Proceedings of Machine Translation Summit X: Papers*, 79–86, Phuket, Thailand, 2005. URL <https://aclanthology.org/2005.mtsummit-papers.11>.

- Kozhevnikov M. and Titov I. Cross-lingual model transfer using feature representation projection. In Toutanova K. and Wu H., editors, *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 579–585, Baltimore, Maryland, June 2014. Association for Computational Linguistics. URL <https://aclanthology.org/P14-2095>.
- Lafferty J.D., McCallum A., and Pereira F.C.N. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In Brodley C.E. and Danyluk A.P., editors, *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, Williams College, Williamstown, MA, USA, June 28 - July 1, 2001, 282–289. Morgan Kaufmann, 2001.
- Lample G., Ballesteros M., Subramanian S., Kawakami K., and Dyer C. Neural architectures for named entity recognition. In Knight K., Nenkova A., and Rambow O., editors, *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 260–270, San Diego, California, June 2016. Association for Computational Linguistics. URL <https://aclanthology.org/N16-1030>.
- Le D.M., Chen Y., Ritter A., and Xu W. Constrained decoding for cross-lingual label projection. *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=DayPQKXaQk>.
- Lee J., Yoon W., Kim S., Kim D., Kim S., So C.H., and Kang J. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinform.*, 36(4):1234–1240, 2020. URL <https://doi.org/10.1093/bioinformatics/btz682>.
- Lester B., Pressel D., Hemmeter A., Choudhury S.R., and Bangalore S. Constrained decoding for computationally efficient named entity recognition taggers. In Cohn T., He Y., and Liu Y., editors, *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, EMNLP 2020 lib. of *Findings of ACL*, 1841–1848. Association for Computational Linguistics, 2020. URL <https://doi.org/10.18653/v1/2020.findings-emnlp.166>.

BIBLIOGRAPHY

- Lewis P., Perez E., Piktus A., Petroni F., Karpukhin V., Goyal N., Küttler H., Lewis M., Yih W.t., Rocktäschel T., Riedel S., and Kiela D. Retrieval-augmented generation for knowledge-intensive nlp tasks. In Larochelle H., Ranzato M., Hadsell R., Balcan M., and Lin H., editors, *Advances in Neural Information Processing Systems*, 33 lib., 9459–9474. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf.
- Li B., He Y., and Xu W. Cross-lingual named entity recognition using parallel corpus: A new approach using xlm-roberta alignment. *CoRR*, abs/2101.11112, 2021. URL <https://arxiv.org/abs/2101.11112>.
- Li J., Sun Y., Johnson R.J., Sciaky D., Wei C., Leaman R., Davis A.P., Mattingly C.J., Wiegers T.C., and Lu Z. Biocreative V CDR task corpus: a resource for chemical disease relation extraction. *Database J. Biol. Databases Curation*, 2016, 2016. URL <https://doi.org/10.1093/database/baw068>.
- Lieber O., Sharir O., Lenz B., and Shoham Y. Jurassic-1: Technical details and evaluation. *White Paper. AI21 Labs*, 1(9), 2021. URL <https://www.ai21.com/research/jurassic-1-technical-details-evaluation>.
- Liu P.J., Saleh M., Pot E., Goodrich B., Sepassi R., Kaiser L., and Shazeer N. Generating wikipedia by summarizing long sequences. *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=Hyg0vbWC->.
- Liu R., Wei J., Liu F., Si C., Zhang Y., Rao J., Zheng S., Peng D., Yang D., Zhou D., and Dai A.M. Best practices and lessons learned on synthetic data for language models. *CoRR*, abs/2404.07503, 2024. URL <https://doi.org/10.48550/arXiv.2404.07503>.
- Liu T., Jiang Y.E., Monath N., Cotterell R., and Sachan M. Autoregressive structured prediction with language models. *Findings of the Association for Computational Linguistics: EMNLP 2022*, 993–1005, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.findings-emnlp.70>.

- Liu Y., Gu J., Goyal N., Li X., Edunov S., Ghazvininejad M., Lewis M., and Zettlemoyer L. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742, 2020. URL <https://aclanthology.org/2020.tacl-1.47>.
- Liu Z., Xu Y., Yu T., Dai W., Ji Z., Cahyawijaya S., Madotto A., and Fung P. Crossner: Evaluating cross-domain named entity recognition. *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, 13452–13460. AAAI Press, 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17587>.
- Loshchilov I. and Hutter F. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017. URL <http://arxiv.org/abs/1711.05101>.
- Lou J., Lu Y., Dai D., Jia W., Lin H., Han X., Sun L., and Wu H. Universal information extraction as unified semantic matching. In Williams B., Chen Y., and Neville J., editors, *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 13318–13326. AAAI Press, 2023. URL <https://doi.org/10.1609/aaai.v37i11.26563>.
- Lu Y., Lin H., Xu J., Han X., Tang J., Li A., Sun L., Liao M., and Chen S. Text2Event: Controllable sequence-to-structure generation for end-to-end event extraction. *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)*, 2795–2806, Online, August 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.acl-long.217>.
- Lu Y., Liu Q., Dai D., Xiao X., Lin H., Han X., Sun L., and Wu H. Unified structure generation for universal information extraction. In Muresan S., Nakov P., and Villavicencio A., editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL

BIBLIOGRAPHY

- 2022, *Dublin, Ireland, May 22-27, 2022*, 5755–5772. Association for Computational Linguistics, 2022. URL <https://doi.org/10.18653/v1/2022.acl-long.395>.
- Luo R., Sun L., Xia Y., Qin T., Zhang S., Poon H., and Liu T.Y. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6), 2022.
- Luong T., Pham H., and Manning C.D. Bilingual word representations with monolingual quality in mind. In Blunsom P., Cohen S.B., Dhillon P.S., and Liang P., editors, *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing, VS@NAACL-HLT 2015, June 5, 2015, Denver, Colorado, USA*, 151–159. The Association for Computational Linguistics, 2015. URL <https://doi.org/10.3115/v1/w15-1521>.
- Magnini B., Altuna B., Lavelli A., Speranza M., and Zanolli R. The E3C project: European clinical case corpus. In Alkorta J., Gonzalez-Dios I., Atutxa A., Gojenola K., Martínez-Cámara E., Rodrigo Á., and Martínez P., editors, *Proceedings of the Annual Conference of the Spanish Association for Natural Language Processing: Projects and Demonstrations (SEPLN-PD 2021) co-located with the Conference of the Spanish Society for Natural Language Processing (SEPLN 2021), Málaga, Spain, September, 2021*, 2968 lib. of CEUR Workshop Proceedings, 17–20. CEUR-WS.org, 2021. URL <https://ceur-ws.org/Vol-2968/paper5.pdf>.
- Mayer T., Cabrio E., and Villata S. Transformer-based argument mining for healthcare applications. In Giacomo G.D., Catalá A., Dilkina B., Milano M., Barro S., Bugarín A., and Lang J., editors, *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, 325 lib. of *Frontiers in Artificial Intelligence and Applications*, 2108–2115. IOS Press, 2020. URL <https://doi.org/10.3233/FAIA200334>.
- Mayer T., Marro S., Cabrio E., and Villata S. Enhancing evidence-based medicine with natural language argumentative analysis of clinical trials. *Artificial Intelligence in Medicine*, 118:102098, 2021.
- Mesnard T., Hardin C., Dadashi R., Bhupatiraju S., Pathak S., Sifre L., Rivière M., Kale M.S., Love J., Tafti P., Hussenot L., Chowdhery A., Roberts A., Barua

- A., Botev A., Castro-Ros A., Slone A., Héliou A., Tacchetti A., Bulanova A., Paterson A., Tsai B., Shahriari B., Lan C.L., Choquette-Choo C.A., Crepy C., Cer D., Ippolito D., Reid D., Buchatskaya E., Ni E., Noland E., Yan G., Tucker G., Muraru G., Rozhdestvenskiy G., Michalewski H., Tenney I., Grishchenko I., Austin J., Keeling J., Labanowski J., Lespiau J., Stanway J., Brennan J., Chen J., Ferret J., Chiu J., and et al. Gemma: Open models based on gemini research and technology. *CoRR*, abs/2403.08295, 2024. URL <https://doi.org/10.48550/arXiv.2403.08295>.
- Mikolov T., Chen K., Corrado G., and Dean J. Efficient estimation of word representations in vector space. In Bengio Y. and LeCun Y., editors, *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, 2013. URL <http://arxiv.org/abs/1301.3781>.
- Min B., Ross H., Sulem E., Veyseh A.P.B., Nguyen T.H., Sainz O., Agirre E., Heintz I., and Roth D. Recent advances in natural language processing via large pre-trained language models: A survey. *CoRR*, abs/2111.01243, 2021. URL <https://arxiv.org/abs/2111.01243>.
- Min B., Ross H., Sulem E., Veyseh A.P.B., Nguyen T.H., Sainz O., Agirre E., Heintz I., and Roth D. Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Comput. Surv.*, 56(2):30:1–30:40, 2024. URL <https://doi.org/10.1145/3605943>.
- Muennighoff N., Wang T., Sutawika L., Roberts A., Biderman S., Scao T.L., Bari M.S., Shen S., Yong Z.X., Schoelkopf H., Tang X., Radev D., Aji A.F., Almubarak K., Albanie S., Alyafeai Z., Webson A., Raff E., and Raffel C. Crosslingual generalization through multitask finetuning. In Rogers A., Boyd-Graber J.L., and Okazaki N., editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 15991–16111. Association for Computational Linguistics, 2023. URL <https://doi.org/10.18653/v1/2023.acl-long.891>.
- Nakayama H. sequeval: A python framework for sequence labeling evaluation, 2018. URL <https://github.com/chakki-works/sequeval>. Software available from <https://github.com/chakki-works/sequeval>.

BIBLIOGRAPHY

- Ni J., Dinu G., and Florian R. Weakly supervised cross-lingual named entity recognition via effective annotation and representation projection. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1470–1480, Vancouver, Canada, July 2017. Association for Computational Linguistics. URL <https://aclanthology.org/P17-1135>.
- Och F.J. and Ney H. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51, 2003. URL <https://aclanthology.org/J03-1002>.
- Ojo J. and Ogueji K. How good are commercial large language models on african languages? *Proceedings of the 4th Workshop on African Natural Language Processing, AfricaNLP@ICLR 2023, Kigali, Rwanda, May 1, 2023*, 2023. URL <https://openreview.net/pdf?id=MCgyGyRPEIU>.
- Ojo J., Ogueji K., Stenetorp P., and Adelani D.I. How good are large language models on african languages? *CoRR*, abs/2311.07978, 2023. URL <https://doi.org/10.48550/arXiv.2311.07978>.
- OpenAI, Achiam J., Adler S., Agarwal S., Ahmad L., Akkaya I., Aleman F.L., Almeida D., Altenschmidt J., Altman S., Anadkat S., Avila R., Babuschkin I., Balaji S., Balcom V., Baltescu P., Bao H., Bavarian M., Belgum J., Bello I., Berdine J., Bernadett-Shapiro G., Berner C., Bogdonoff L., Boiko O., Boyd M., Brakman A.L., Brockman G., Brooks T., Brundage M., Button K., Cai T., Campbell R., Cann A., Carey B., Carlson C., Carmichael R., Chan B., Chang C., Chantzis F., Chen D., Chen S., Chen R., Chen J., Chen M., Chess B., Cho C., Chu C., Chung H.W., Cummings D., Currier J., Dai Y., Decareaux C., Degry T., Deutsch N., Deville D., Dhar A., Dohan D., Dowling S., Dunning S., Ecoffet A., Eleti A., Eloundou T., Farhi D., Fedus L., Felix N., Fishman S.P., Forte J., Fulford I., Gao L., Georges E., Gibson C., Goel V., Gogineni T., Goh G., Gontijo-Lopes R., Gordon J., Grafstein M., Gray S., Greene R., Gross J., Gu S.S., Guo Y., Hallacy C., Han J., Harris J., He Y., Heaton M., Heidecke J., Hesse C., Hickey A., Hickey W., Hoeschele P., Houghton B., Hsu K., Hu S., Hu X., Huizinga J., Jain S., Jain S., Jang J., Jiang A., Jiang R., Jin H., Jin D., Jomoto S., Jonn B., Jun H., Kaftan T., Łukasz Kaiser, Kamali A., Kanitscheider I., Keskar N.S., Khan T., Kilpatrick L., Kim J.W., Kim C., Kim Y., Kirchner J.H., Kiros J., Knight M., Kokotajlo D., Łukasz Kondraciuk, Kondrich A., Konstantinidis A., Kosic K., Krueger G., Kuo V., Lampe M., Lan I., Lee T., Leike J.,

- Leung J., Levy D., Li C.M., Lim R., Lin M., Lin S., Litwin M., Lopez T., Lowe R., Lue P., Makanju A., Malfacini K., Manning S., Markov T., Markovski Y., Martin B., Mayer K., Mayne A., McGrew B., McKinney S.M., McLeavey C., McMillan P., McNeil J., Medina D., Mehta A., Menick J., Metz L., Mishchenko A., Mishkin P., Monaco V., Morikawa E., Mossing D., Mu T., Murati M., Murk O., Mély D., Nair A., Nakano R., Nayak R., Neelakantan A., Ngo R., Noh H., Ouyang L., O’Keefe C., Pachocki J., Paino A., Palermo J., Pantuliano A., Parascandolo G., Parish J., Parparita E., Passos A., Pavlov M., Peng A., Perelman A., de Avila Belbute Peres F., Petrov M., de Oliveira Pinto H.P., Michael, Pokorny, Pokrass M., Pong V.H., Powell T., Power A., Power B., Proehl E., Puri R., Radford A., Rae J., Ramesh A., Raymond C., Real F., Rimbach K., Ross C., Rotsted B., Roussez H., Ryder N., Saltarelli M., Sanders T., Santurkar S., Sastry G., Schmidt H., Schnurr D., Schulman J., Selsam D., Sheppard K., Sherbakov T., Shieh J., Shoker S., Shyam P., Sidor S., Sigler E., Simens M., Sitkin J., Slama K., Sohl I., Sokolowsky B., Song Y., Staudacher N., Such F.P., Summers N., Sutskever I., Tang J., Tezak N., Thompson M.B., Tillet P., Tootoonchian A., Tseng E., Tuggle P., Turley N., Tworek J., Uribe J.F.C., Vallone A., Vijayvergiya A., Voss C., Wainwright C., Wang J.J., Wang A., Wang B., Ward J., Wei J., Weinmann C., Welihinda A., Welinder P., Weng J., Weng L., Wiethoff M., Willner D., Winter C., Wolrich S., Wong H., Workman L., Wu S., Wu J., Wu M., Xiao K., Xu T., Yoo S., Yu K., Yuan Q., Zaremba W., Zellers R., Zhang C., Zhang M., Zhao S., Zheng T., Zhuang J., Zhuk W., and Zoph B. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
- Östling R. and Tiedemann J. Efficient word alignment with Markov Chain Monte Carlo. *Prague Bulletin of Mathematical Linguistics*, 106:125–146, October 2016. URL <http://ufal.mff.cuni.cz/pbml/106/art-ostling-tiedemann.pdf>.
- Parekh T., Hsu I.H., Huang K.H., Chang K.W., and Peng N. Contextual label projection for cross-lingual structured prediction. *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2024.
- Pennington J., Socher R., and Manning C. GloVe: Global vectors for word representation. In Moschitti A., Pang B., and Daelemans W., editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. URL <https://aclanthology.org/D14-1162>.

BIBLIOGRAPHY

- Peter J., Nix A., and Ney H. Generating alignments using target foresight in attention-based neural machine translation. *Prague Bull. Math. Linguistics*, 108:27–36, 2017. URL <http://ufal.mff.cuni.cz/pbml/108/art-peter-nix-ney.pdf>.
- Phan L.N., Anibal J.T., Tran H., Chanana S., Bahadroglu E., Peltekian A., and Altan-Bonnet G. SciFive: a text-to-text transformer model for biomedical literature. *CoRR*, abs/2106.03598, 2021. URL <https://arxiv.org/abs/2106.03598>.
- Pires T., Schlinger E., and Garrette D. How multilingual is multilingual BERT? *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4996–5001, Florence, Italy, July 2019. Association for Computational Linguistics. URL <https://aclanthology.org/P19-1493>.
- Pontiki M., Galanis D., Papageorgiou H., Androutsopoulos I., Manandhar S., Al-Smadi M., Al-Ayyoub M., Zhao Y., Qin B., De Clercq O., Hoste V., Apidianaki M., Tannier X., Loukachevitch N., Kotelnikov E., Bel N., Jiménez-Zafra S.M., and Eryiğit G. SemEval-2016 task 5: Aspect based sentiment analysis. In Bethard S., Carpuat M., Cer D., Jurgens D., Nakov P., and Zesch T., editors, *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, 19–30, San Diego, California, June 2016. Association for Computational Linguistics. URL <https://aclanthology.org/S16-1002>.
- Radford A., Wu J., Child R., Luan D., Amodei D., Sutskever I., *et al.*. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rae J.W., Borgeaud S., Cai T., Millican K., Hoffmann J., Song H.F., Aslanides J., Henderson S., Ring R., Young S., Rutherford E., Hennigan T., Menick J., Cassirer A., Powell R., van den Driessche G., Hendricks L.A., Rauh M., Huang P., Glaese A., Welbl J., Dathathri S., Huang S., Uesato J., Mellor J., Higgins I., Creswell A., McAleese N., Wu A., Elsen E., Jayakumar S.M., Buchatskaya E., Budden D., Sutherland E., Simonyan K., Paganini M., Sifre L., Martens L., Li X.L., Kuncoro A., Nematzadeh A., Gribovskaya E., Donato D., Lazaridou A., Mensch A., Lespiau J., Tsimpoukelli M., Grigorev N., Fritz D., Sottiaux T., Pajarskas M., Pohlen T., Gong Z., Toyama D., de Masson d’Autume C., Li Y., Terzi T., Mikulik V., Babuschkin I., Clark A., de Las Casas D., Guy A., Jones C., Bradbury J., Johnson M.J., Hechtman B.A., Weidinger L., Gabriel I.,

- Isaac W., Lockhart E., Osindero S., Rimell L., Dyer C., Vinyals O., Ayoub K., Stanway J., Bennett L., Hassabis D., Kavukcuoglu K., and Irving G. Scaling language models: Methods, analysis & insights from training gopher, 2021. URL <https://arxiv.org/abs/2112.11446>.
- Raffel C., Shazeer N., Roberts A., Lee K., Narang S., Matena M., Zhou Y., Li W., and Liu P.J. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- Rajpurkar P., Jia R., and Liang P. Know what you don’t know: Unanswerable questions for squad. In Gurevych I. and Miyao Y., editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, 784–789. Association for Computational Linguistics, 2018. URL <https://aclanthology.org/P18-2124/>.
- Raman K., Naim I., Chen J., Hashimoto K., Yalasangi K., and Srinivasan K. Transforming sequence tagging into A seq2seq task. In Goldberg Y., Kozareva Z., and Zhang Y., editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 11856–11874. Association for Computational Linguistics, 2022. URL <https://doi.org/10.18653/v1/2022.emnlp-main.813>.
- Ratinov L. and Roth D. Design challenges and misconceptions in named entity recognition. In Stevenson S. and Carreras X., editors, *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, 147–155, Boulder, Colorado, June 2009. Association for Computational Linguistics. URL <https://aclanthology.org/W09-1119>.
- Rombach R., Blattmann A., Lorenz D., Esser P., and Ommer B. High-resolution image synthesis with latent diffusion models. *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 10674–10685. IEEE, 2022. URL <https://doi.org/10.1109/CVPR52688.2022.01042>.
- Ruder S., Vulic I., and Søgaard A. A survey of cross-lingual word embedding models. *J. Artif. Intell. Res.*, 65:569–631, 2019. URL <https://doi.org/10.1613/jair.1.11640>.

BIBLIOGRAPHY

- Sainz O., García-Ferrero I., Agerri R., de Lacalle O.L., Rigau G., and Agirre E. GoLLIE: Annotation guidelines improve zero-shot information-extraction. *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=Y3wpuxd7u9>.
- Sainz O., García-Ferrero I., Jacovi A., Ander Campos J., Elazar Y., Agirre E., Goldberg Y., Chen W.L., Chim J., Choshen L., D’Amico-Wong L., Dell M., Fan R.Z., Golchin S., Li Y., Liu P., Pahwa B., Prabhu A., Sharma S., Silcock E., Solonko K., Stap D., Surdeanu M., Tseng Y.M., Udandara V., Wang Z., Xu R., and Yang J. Data contamination report from the 2024 CONDA shared task. In Sainz O., García Ferrero I., Agirre E., Ander Campos J., Jacovi A., Elazar Y., and Goldberg Y., editors, *Proceedings of the 1st Workshop on Data Contamination (CONDA)*, 41–56, Bangkok, Thailand, August 2024b. Association for Computational Linguistics. URL <https://aclanthology.org/2024.conda-1.4>.
- Sang E.F.T.K. Introduction to the conll-2002 shared task: Language-independent named entity recognition. In Roth D. and van den Bosch A., editors, *Proceedings of the 6th Conference on Natural Language Learning, CoNLL 2002, Held in cooperation with COLING 2002, Taipei, Taiwan, 2002*. ACL, 2002. URL <https://aclanthology.org/W02-2024/>.
- Sang E.F.T.K. and Meulder F.D. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In Daelemans W. and Osborne M., editors, *Proceedings of the Seventh Conference on Natural Language Learning, CoNLL 2003, Held in cooperation with HLT-NAACL 2003, Edmonton, Canada, May 31 - June 1, 2003*, 142–147. ACL, 2003. URL <https://aclanthology.org/W03-0419/>.
- Shah R., Lin B., Gershman A., and Frederking R. Synergy: a named entity recognition system for resource-scarce languages such as swahili using online machine translation. *Proceedings of the Second Workshop on African Language Technology (AfLaT 2010)*, 21–26, 2010.
- Shazeer N. and Stern M. Adafactor: Adaptive learning rates with sublinear memory cost. In Dy J.G. and Krause A., editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, 80 lib. of *Proceedings of Machine Learning Research*, 4603–4611. PMLR, 2018. URL <http://proceedings.mlr.press/v80/shazeer18a.html>.

- Singhal K., Azizi S., Tu T., Mahdavi S.S., Wei J., Chung H.W., Scales N., Tanwani A.K., Cole-Lewis H., Pfohl S., Payne P., Seneviratne M., Gamble P., Kelly C., Schärli N., Chowdhery A., Mansfield P.A., y Arcas B.A., Webster D.R., Corrado G.S., Matias Y., Chou K., Gottweis J., Tomasev N., Liu Y., Rajkomar A., Barral J.K., Sementur C., Karthikesalingam A., and Natarajan V. Large language models encode clinical knowledge. *arXiv preprint*, abs/2212.13138, 2022.
- Smith S.L., Turban D.H.P., Hamblin S., and Hammerla N.Y. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=r1Aab85gg>.
- Smith S., Patwary M., Norick B., LeGresley P., Rajbhandari S., Casper J., Liu Z., Prabhumoye S., Zerveas G., Korthikanti V., Zheng E., Child R., Aminabadi R.Y., Bernauer J., Song X., Shoenybi M., He Y., Houston M., Tiwary S., and Catanzaro B. Using deepspeed and megatron to train megatron-turing NLG 530b, A large-scale generative language model. *CoRR*, abs/2201.11990, 2022. URL <https://arxiv.org/abs/2201.11990>.
- Speranza M. The named entity recognition task at evalita 2009. *Proceedings of the 11th Conference of the Italian Association for Artificial Intelligence*, 2009.
- Täckström O., McDonald R., and Uszkoreit J. Cross-lingual word clusters for direct transfer of linguistic structure. In Fosler-Lussier E., Riloff E., and Bangalore S., editors, *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 477–487, Montréal, Canada, June 2012. Association for Computational Linguistics. URL <https://aclanthology.org/N12-1052>.
- Tang Y., Tran C., Li X., Chen P., Goyal N., Chaudhary V., Gu J., and Fan A. Multilingual translation with extensible multilingual pretraining and finetuning. *CoRR*, abs/2008.00401, 2020. URL <https://arxiv.org/abs/2008.00401>.
- Team G., Mesnard T., Hardin C., Dadashi R., Bhupatiraju S., Pathak S., Sifre L., Rivière M., Kale M.S., Love J., Tafti P., Hussenot L., Sessa P.G., Chowdhery A., Roberts A., Barua A., Botev A., Castro-Ros A., Slone A., Héliou A., Tacchetti A., Bulanov A., Paterson A., Tsai B., Shahriari B., Lan C.L., Choquette-Choo C.A., Crepey C., Cer D., Ippolito D., Reid D., Buchatskaya E., Ni E., Noland E.,

BIBLIOGRAPHY

- Yan G., Tucker G., Muraru G.C., Rozhdestvenskiy G., Michalewski H., Tenney I., Grishchenko I., Austin J., Keeling J., Labanowski J., Lespiau J.B., Stanway J., Brennan J., Chen J., Ferret J., Chiu J., Mao-Jones J., Lee K., Yu K., Millican K., Sjoesund L.L., Lee L., Dixon L., Reid M., Mikuła M., Wirth M., Sharman M., Chinaev N., Thain N., Bachem O., Chang O., Wahltinez O., Bailey P., Michel P., Yotov P., Chaabouni R., Comanescu R., Jana R., Anil R., McIlroy R., Liu R., Mullins R., Smith S.L., Borgeaud S., Girgin S., Douglas S., Pandya S., Shakeri S., De S., Klimenko T., Hennigan T., Feinberg V., Stokowiec W., hui Chen Y., Ahmed Z., Gong Z., Warkentin T., Peran L., Giang M., Farabet C., Vinyals O., Dean J., Kavukcuoglu K., Hassabis D., Ghahramani Z., Eck D., Barral J., Pereira F., Collins E., Joulin A., Fiedel N., Senter E., Andreev A., and Kenealy K. Gemma: Open models based on gemini research and technology, 2024. URL <https://arxiv.org/abs/2403.08295>.
- Tebbifakhr A., Negri M., and Turchi M. Machine-oriented NMT adaptation for zero-shot NLP tasks: Comparing the usefulness of close and distant languages. In Zampieri M., Nakov P., Ljubešić N., Tiedemann J., and Scherrer Y., editors, *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, 36–46, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics (ICCL). URL <https://aclanthology.org/2020.vardial-1.4>.
- Teknum. Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants, 2023. URL <https://huggingface.co/datasets/teknum/OpenHermes-2.5>.
- Thompson B. and Post M. Automatic machine translation evaluation in many languages via zero-shot paraphrasing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 90–121, Online, 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.emnlp-main.8>.
- Thoppilan R., Freitas D.D., Hall J., Shazeer N., Kulshreshtha A., Cheng H., Jin A., Bos T., Baker L., Du Y., Li Y., Lee H., Zheng H.S., Ghafouri A., Mene-gali M., Huang Y., Krikun M., Lepikhin D., Qin J., Chen D., Xu Y., Chen Z., Roberts A., Bosma M., Zhou Y., Chang C., Krivokon I., Rusch W., Pickett M., Meier-Hellstern K.S., Morris M.R., Doshi T., Santos R.D., Duke T., Soraker J., Zevenbergen B., Prabhakaran V., Diaz M., Hutchinson B., Olson K., Molina A., Hoffman-John E., Lee J., Aroyo L., Rajakumar R., Butryna

- A., Lamm M., Kuzmina V., Fenton J., Cohen A., Bernstein R., Kurzweil R., Aguera-Arcas B., Cui C., Croak M., Chi E.H., and Le Q. Lamda: Language models for dialog applications. *CoRR*, abs/2201.08239, 2022. URL <https://arxiv.org/abs/2201.08239>.
- Tiedemann J., Agić Ž., and Nivre J. Treebank translation for cross-lingual parser induction. In Morante R. and Yih S.W.t., editors, *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, 130–140, Ann Arbor, Michigan, June 2014. Association for Computational Linguistics. URL <https://aclanthology.org/W14-1614>.
- Tiedemann J. and Thottingal S. OPUS-MT – building open translation services for the world. In Martins A., Moniz H., Fumega S., Martins B., Batista F., Coheur L., Parra C., Trancoso I., Turchi M., Bisazza A., Moorkens J., Guerberof A., Nurminen M., Marg L., and Forcada M.L., editors, *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, 479–480, Lisboa, Portugal, November 2020. European Association for Machine Translation. URL <https://aclanthology.org/2020.eamt-1.61>.
- Tiedemann J. Parallel data, tools and interfaces in OPUS. In Chair) N.C.C., Choukri K., Declerck T., Dogan M.U., Maegaard B., Mariani J., Odijk J., and Piperidis S., editors, *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey, may 2012. European Language Resources Association (ELRA). ISBN 978-2-9517408-7-7.
- Touvron H., Lavril T., Izacard G., Martinet X., Lachaux M.A., Lacroix T., Rozière B., Goyal N., Hambro E., Azhar F., Rodriguez A., Joulin A., Grave E., and Lample G. LLaMA: Open and efficient foundation language models. *arXiv preprint*, 2302.13971, 2023a.
- Touvron H., Martin L., Stone K., Albert P., Almahairi A., Babaei Y., Bashlykov N., Batra S., Bhargava P., Bhosale S., Bikel D., Blecher L., Canton-Ferrer C., Chen M., Cucurull G., Esiobu D., Fernandes J., Fu J., Fu W., Fuller B., Gao C., Goswami V., Goyal N., Hartshorn A., Hosseini S., Hou R., Inan H., Kardas M., Kerkez V., Khabsa M., Kloumann I., Korenev A., Koura P.S., Lachaux M., Lavril T., Lee J., Liskovich D., Lu Y., Mao Y., Martinet X., Mihaylov T., Mishra P., Molybog I., Nie Y., Poulton A., Reizenstein J., Rungta R., Saladi K., Schelten A., Silva R., Smith E.M., Subramanian R., Tan X.E., Tang B., Taylor R., Williams A., Kuan J.X., Xu P., Yan Z., Zarov I., Zhang Y., Fan A.,

BIBLIOGRAPHY

- Kambadur M., Narang S., Rodriguez A., Stojnic R., Edunov S., and Scialom T. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023b. URL <https://doi.org/10.48550/arXiv.2307.09288>.
- Tran C., Bhosale S., Cross J., Koehn P., Edunov S., and Fan A. Facebook AI WMT21 news translation task submission. *CoRR*, abs/2108.03265, 2021. URL <https://arxiv.org/abs/2108.03265>.
- Tsatsaronis G., Balikas G., Malakasiotis P., Partalas I., Zschunke M., Alvers M.R., Weissenborn D., Krithara A., Petridis S., Polychronopoulos D., Almirantis Y., Pavlopoulos J., Baskiotis N., Gallinari P., Artières T., Ngomo A.N., Heino N., Gaussier É., Barrio-Alvers L., Schroeder M., Androutsopoulos I., and Paliouras G. An overview of the BioASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinform.*, 16:138:1–138:28, 2015. URL <https://doi.org/10.1186/s12859-015-0564-6>.
- Üstün A., Aryabumi V., Yong Z.X., Ko W., D’souza D., Onilude G., Bhandari N., Singh S., Ooi H., Kayid A., Vargus F., Blunsom P., Longpre S., Muennighoff N., Fadaee M., Kreutzer J., and Hooker S. Aya model: An instruction finetuned open-access multilingual language model. *CoRR*, abs/2402.07827, 2024. URL <https://doi.org/10.48550/arXiv.2402.07827>.
- Vamvas J. and Sennrich R. NMTScore: A multilingual analysis of translation-based text similarity measures. *Findings of the Association for Computational Linguistics: EMNLP 2022*, 198–213, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.findings-emnlp.15>.
- Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A.N., Kaiser Ł., and Polosukhin I. Attention is all you need. In Guyon I., von Luxburg U., Bengio S., Wallach H.M., Fergus R., Vishwanathan S.V.N., and Garnett R., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Vinyals O., Kaiser Ł., Koo T., Petrov S., Sutskever I., and Hinton G.E. Grammar as a foreign language. In Cortes C., Lawrence N.D., Lee D.D.,

- Sugiyama M., and Garnett R., editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, 2773–2781, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/277281aada22045c03945dcb2ca6f2ec-Abstract.html>.
- Walker C., Strassel S., Medero J., and Maeda K. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57:45, 2006. URL <https://catalog.ldc.upenn.edu/LDC2006T06>.
- Wang A., Singh A., Michael J., Hill F., Levy O., and Bowman S.R. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJ4km2R5t7>.
- Wang G., Yang G., Du Z., Fan L., and Li X. ClinicalGPT: Large Language Models Finetuned with Diverse Medical Data and Comprehensive Evaluation. *ArXiv preprint*, abs/2306.09968, 2023a.
- Wang M. and Manning C.D. Cross-lingual projected expectation regularization for weakly supervised learning. *Transactions of the Association for Computational Linguistics*, 2:55–66, 2014. URL <https://aclanthology.org/Q14-1005>.
- Wang X., Zhou W., Zu C., Xia H., Chen T., Zhang Y., Zheng R., Ye J., Zhang Q., Gui T., Kang J., Yang J., Li S., and Du C. Instructuie: Multi-task instruction tuning for unified information extraction. *CoRR*, abs/2304.08085, 2023b. URL <https://doi.org/10.48550/arXiv.2304.08085>.
- Wang Y., Mishra S., Alipoormolabashi P., Kordi Y., Mirzaei A., Naik A., Ashok A., Dhanasekaran A.S., Arunkumar A., Stap D., Pathak E., Karamanolakis G., Lai H.G., Purohit I., Mondal I., Anderson J., Kuznia K., Doshi K., Pal K.K., Patel M., Moradshahi M., Parmar M., Purohit M., Varshney N., Kaza P.R., Verma P., Puri R.S., Karia R., Doshi S., Sampat S.K., Mishra S., A S.R., Patro S., Dixit T., and Shen X. Super-naturalinstructions: Generalization via declarative instructions on 1600+ NLP tasks. In Goldberg Y., Kozareva Z., and Zhang Y., editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 5085–5109. Association for Computational Linguistics,

BIBLIOGRAPHY

2022. URL <https://doi.org/10.18653/v1/2022.emnlp-main.340>.
- Wolf T., Debut L., Sanh V., Chaumond J., Delangue C., Moi A., Cistac P., Rault T., Louf R., Funtowicz M., and Brew J. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv preprint*, abs/1910.03771, 2019. URL <https://arxiv.org/abs/1910.03771>.
- Wolf T., Debut L., Sanh V., Chaumond J., Delangue C., Moi A., Cistac P., Rault T., Louf R., Funtowicz M., Davison J., Shleifer S., von Platen P., Ma C., Jernite Y., Plu J., Xu C., Le Scao T., Gugger S., Drame M., Lhoest Q., and Rush A. Transformers: State-of-the-art natural language processing. In Liu Q. and Schlangen D., editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.emnlp-demos.6>.
- Wu C., Lin W., Zhang X., Zhang Y., Wang Y., and Xie W. PMC-LLaMA: Towards building open-source language models for medicine. *arXiv preprint*, 2304.14454, 2023.
- Wu S. and Dredze M. Are all languages created equal in multilingual BERT? *Proceedings of the 5th Workshop on Representation Learning for NLP*, 120–130, Online, July 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.repl4nlp-1.16>.
- Xiao C., Dymetman M., and Gardent C. Sequence-based structured prediction for semantic parsing. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics, 2016. URL <https://doi.org/10.18653/v1/p16-1127>.
- Xie J., Yang Z., Neubig G., Smith N.A., and Carbonell J. Neural cross-lingual named entity recognition with minimal resources. In Riloff E., Chiang D., Hockenmaier J., and Tsujii J., editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 369–379, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. URL <https://aclanthology.org/D18-1034>.
- Xu H., Kim Y.J., Sharaf A., and Awadalla H.H. A paradigm shift in machine translation: Boosting translation performance of large language models, 2023.

- Xue L., Constant N., Roberts A., Kale M., Al-Rfou R., Siddhant A., Barua A., and Raffel C. mt5: A massively multilingual pre-trained text-to-text transformer. In Toutanova K., Rumshisky A., Zettlemoyer L., Hakkani-Tür D., Beltagy I., Bethard S., Cotterell R., Chakraborty T., and Zhou Y., editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, 483–498. Association for Computational Linguistics, 2021. URL <https://doi.org/10.18653/v1/2021.naacl-main.41>.
- Yang A., Yang B., Hui B., Zheng B., Yu B., Zhou C., Li C., Li C., Liu D., Huang F., Dong G., Wei H., Lin H., Tang J., Wang J., Yang J., Tu J., Zhang J., Ma J., Yang J., Xu J., Zhou J., Bai J., He J., Lin J., Dang K., Lu K., Chen K., Yang K., Li M., Xue M., Ni N., Zhang P., Wang P., Peng R., Men R., Gao R., Lin R., Wang S., Bai S., Tan S., Zhu T., Li T., Liu T., Ge W., Deng X., Zhou X., Ren X., Zhang X., Wei X., Ren X., Liu X., Fan Y., Yao Y., Zhang Y., Wan Y., Chu Y., Liu Y., Cui Z., Zhang Z., Guo Z., and Fan Z. Qwen2 technical report, 2024. URL <https://arxiv.org/abs/2407.10671>.
- Yarowsky D., Ngai G., and Wicentowski R. Inducing multilingual text analysis tools via robust projection across aligned corpora. *Proceedings of the First International Conference on Human Language Technology Research*, 2001. URL <https://aclanthology.org/H01-1035>.
- Yasunaga M., Leskovec J., and Liang P. LinkBERT: Pretraining language models with document links. In Muresan S., Nakov P., and Villavicencio A., editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, 8003–8016. Association for Computational Linguistics, 2022. URL <https://doi.org/10.18653/v1/2022.acl-long.551>.
- Yeginbergen A. and Agerri R. Cross-lingual argument mining in the medical domain. *Procesamiento del Lenguaje Natural*, 73:296–312, 2024.
- Yeginbergen A., Oronoz M., and Agerri R. Argument mining in data scarce settings: Cross-lingual transfer and few-shot techniques. *ACL*, 2024.
- Zenkel T., Wuebker J., and DeNero J. Adding interpretable attention to neural translation models improves word alignment. *CoRR*, abs/1901.11359, 2019. URL <http://arxiv.org/abs/1901.11359>.

BIBLIOGRAPHY

- Zenkel T., Wuebker J., and DeNero J. End-to-end neural word alignment outperforms GIZA++. In Jurafsky D., Chai J., Schluter N., and Tetreault J.R., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 1605–1617. Association for Computational Linguistics, 2020. URL <https://doi.org/10.18653/v1/2020.acl-main.146>.
- Zhang Y., Gaddy D., Barzilay R., and Jaakkola T. Ten pairs to tag – multilingual POS tagging via coarse mapping between embeddings. In Knight K., Nenkova A., and Rambow O., editors, *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1307–1317, San Diego, California, June 2016. Association for Computational Linguistics. URL <https://aclanthology.org/N16-1156>.
- Zheng J., Chow J.H., Shen Z., and Xu P. Grammar-based decoding for improved compositional generalization in semantic parsing. *Findings of the Association for Computational Linguistics: ACL 2023*, 1399–1418, Toronto, Canada, July 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.findings-acl.91>.
- Zhou R., Li X., Bing L., Cambria E., Si L., and Miao C. ConNER: Consistency training for cross-lingual named entity recognition. In Goldberg Y., Kozareva Z., and Zhang Y., editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 8438–8449, Abu Dhabi, United Arab Emirates, December 2022a. Association for Computational Linguistics. URL <https://aclanthology.org/2022.emnlp-main.577>.
- Zhou R., Li X., Bing L., Cambria E., Si L., and Miao C. Conner: Consistency training for cross-lingual named entity recognition. In Goldberg Y., Kozareva Z., and Zhang Y., editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 8438–8449. Association for Computational Linguistics, 2022b. URL <https://doi.org/10.18653/v1/2022.emnlp-main.577>.
- Zou A., Phan L., Chen S., Campbell J., Guo P., Ren R., Pan A., Yin X., Mazeika M., Dombrowski A.K., Goel S., Li N., Byun M.J., Wang Z., Mallen A., Basart

BIBLIOGRAPHY

S., Koyejo S., Song D., Fredrikson M., Kolter J.Z., and Hendrycks D. Representation engineering: A top-down approach to ai transparency, 2023. URL <https://arxiv.org/abs/2310.01405>.

A.1 Original papers

In this appendix, we present the original papers presented in the manuscript of this thesis in the recommended reading order.

Model and Data Transfer for Cross-Lingual Sequence Labelling in Zero-Resource Settings

Iker García-Ferrero Rodrigo Agerri German Rigau
HiTZ Basque Center for Language Technologies - Ixa NLP Group
University of the Basque Country UPV/EHU
{ iker.garciaf, rodrigo.agerri, german.rigau }@ehu.eus

Abstract

Zero-resource cross-lingual transfer approaches aim to apply supervised models from a source language to unlabelled target languages. In this paper we perform an in-depth study of the two main techniques employed so far for cross-lingual zero-resource sequence labelling, based either on data or model transfer. Although previous research has proposed translation and annotation projection (data-based cross-lingual transfer) as an effective technique for cross-lingual sequence labelling, in this paper we experimentally demonstrate that high capacity multilingual language models applied in a zero-shot (model-based cross-lingual transfer) setting consistently outperform data-based cross-lingual transfer approaches. A detailed analysis of our results suggests that this might be due to important differences in language use. More specifically, machine translation often generates a textual signal which is different to what the models are exposed to when using gold standard data, which affects both the fine-tuning and evaluation processes. Our results also indicate that data-based cross-lingual transfer approaches remain a competitive option when high-capacity multilingual language models are not available.

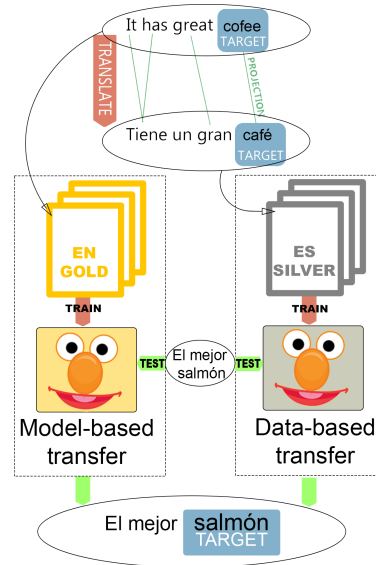


Figure 1: In the data-based transfer approach we translate and project the labels of the gold data into the target language, and use the resulting silver data to train a model for the target language. In the model-based transfer approach we train a model with gold data in English and use it in a zero-shot setting in the target language.

1 Introduction

Sequence labelling is the task of assigning a label to each token in a given input sequence. Sequence labelling is a fundamental process in many downstream NLP tasks. Currently, most successful approaches for this task apply supervised deep-neural networks (Lample et al., 2016; Akbik et al., 2018; Devlin et al., 2019; Conneau et al., 2020). However, as it was the case for supervised statistical approaches (Agerri and Rigau, 2016), their performance still depends on the amount of manually annotated training data. Additionally, deep-neural models still show a significant loss of performance

when evaluated in out-of-domain data (Liu et al., 2021). This means that to improve their performance, it would therefore be necessary to develop very costly manually annotated data for each language and domain of application. Thus, considering that for most of the languages in the world the amount of manually annotated corpora is simply nonexistent (Joshi et al., 2020), then the task of developing sequence labelling models for languages and domain-specific tasks, for which supervised data is not available, remains a challenge of great interest. This task is known as zero-resource cross-lingual sequence labelling.

Data-based cross-lingual transfer methods aim to automatically generate labelled data for a target language. Previous works on data-based transfer have proposed *translation and annotation projection* as an effective technique for zero-resource cross-lingual sequence labelling (Jain et al., 2019; Fei et al., 2020). In this setting, as illustrated in Figure 1, the idea is to translate gold-labelled text into the target language and then, using automatic word alignments, project the labels from the source into the target language. The result is an automatically generated dataset in the target language that can be used for training a sequence labelling model.

The emergence of multilingual language models (Devlin et al., 2019; Conneau et al., 2020) allows for **model-based** cross-lingual transfer. As Figure 1 illustrates, using labelled data in one source language (usually English), it is possible to fine-tune a pre-trained multilingual model that is directly used to make predictions in any of the languages included in the model. This is also known as *zero-shot* cross-lingual sequence labelling.

In this work we present an in-depth study of both approaches using the latest advancements in machine translation, word aligners and multilingual language models. We focus on two sequence labelling tasks, namely, Named Entity Recognition (NER) and Opinion Target Extraction (OTE). In order to do so, we present a data-based cross-lingual transfer approach consisting of translating gold labeled data between English and 7 other languages using state-of-the-art machine translation systems. Sequence labelling annotations are then automatically projected for every language pair. Additionally, we also produced manual alignments for those 4 languages for which we had expert annotators. After translation and projection, for the data-transfer approach we fine-tune multilingual language models using the automatically generated datasets. We then compare the performance obtained for each of the target languages against the performance of the zero-shot cross-lingual method, consisting of fine-tuning the multilingual language models in the English gold data and generating the predictions in the required target languages.

The main contributions of our work are the following: First, we empirically establish the required conditions for each of these two approaches, data-transfer and zero-shot model-based, to outperform the other. In this sense, our experiments show that, contrary to what previous research sug-

gested (Fei et al., 2020; Li et al., 2021), the zero-shot model-based approach obtains the best results when high-capacity multilingual models including the target language and domain are available. Second, when the performance of the multilingual language model is not optimal for the specific target language or domain (for example when working on a text genre and domain for which available language models have not been trained), or when the required hardware to work with high-capacity language models is not easily accessible, then data-transfer based on *translate and project* constitutes a competitive option. Third, we observe that machine translation data often generates training and test data which is, due to important differences in language use, markedly different to the signal received when using gold standard data in the target language. These discrepancies seem to explain the larger error rate of the translate and project method with respect to the zero-shot technique. Finally, we create manually projected datasets for four languages and automatically projected datasets for seven languages. We use them to train and evaluate cross-lingual sequence labelling models. Additionally, they are also used to extrinsically evaluate machine translation and word alignment systems. These new datasets, together with the code to generate them are publicly available to facilitate the reproducibility of results and its use in future research.¹

2 Related work

2.1 Data-based cross-lingual transfer

Data-based cross-lingual transfer methods aim to automatically generate labelled data for a target language. Some of these methods exploit parallel data. Ehrmann et al. (2011) automatically annotate the English version of a multi-parallel corpus and projects the annotations into all the other languages using statistical alignments of phrases. Wang and Manning (2014) project model expectations rather than labels, which facilitates transfer of model uncertainty across languages. Ni et al. (2017) use a heuristic scheme that effectively selects good-quality projection-labeled data from noisy data. They also project word embeddings from a target language into a source language, so that the

¹<https://github.com/ikergarcia1996/Easy-Label-Projection>
<https://github.com/ikergarcia1996/Easy-Translate>

source-language sequence labelling system can be applied to the target language without re-training. [Agerri et al. \(2018\)](#) use parallel data from multiple languages as source to project the labelled data to a target language, showing that the combination of multiple sources improves the quality of the projections. [Li et al. \(2021\)](#) uses the XLM-R model ([Conneau et al., 2020](#)) for labelling sequences in the source part of the parallel data and also for annotation projection.

Instead of relying on parallel data, [Jain et al. \(2019\)](#) and [Fei et al. \(2020\)](#), use machine translation to automatically translate the sentences of a gold-labelled dataset to the target languages. The translated data is then annotated by projecting the gold labels from the source dataset. For this purpose, [Jain et al. \(2019\)](#) first generate a list of projection candidates by orthographic and phonetic similarity. They choose the best matching candidate based on distributional statistics derived from the dataset. [Fei et al. \(2020\)](#) leverages the word alignment probabilities calculated with FastAlign ([Dyer et al., 2013](#)) and the POS tag distributions of the source and target words.

High quality parallel data or machine translation systems are not always available. Thus, [Xie et al. \(2018\)](#) proposes to find word translations based on bilingual word-embeddings. Alternatively, [Guo and Roth \(2021\)](#) translate labelled data in a word-by-word manner with a dictionary. Then, they construct target-language text from the source-language annotations with a constrained pretrained language model.

2.2 Model-based transfer

Language models trained on monolingual corpora in many languages ([Devlin et al., 2019](#); [Conneau et al., 2020](#)) allow zero-shot cross-lingual model transfer. Task-specific data in one language is used to fine-tune the model for evaluation in another language ([Pires et al., 2019](#)). The zero-shot cross-lingual capability can be improved for the sequence labelling task using different techniques. The approaches of [Wang et al. \(2019\)](#) and [Ouyang et al. \(2021\)](#) use monolingual corpora to improve the alignment of the language representations within a multilingual model. Instead of using a single source model, ([Rahimi et al., 2019](#)) propose to use many models from many source languages to improve the zero-shot transfer to a new language. They learn to infer which are the most reliable mod-

els in an unsupervised manner. [Wu et al. \(2020\)](#) take advantage of a Teacher-Student learning approach. NER models in the source languages are used as teachers to train a student model on unlabeled data in the target language. [Bari et al. \(2021\)](#) propose an unsupervised data augmentation framework to improve the cross-lingual adaptation of models using self-training. [Hu et al. \(2021\)](#) use the minimum risk training framework to overcome the gap between the source and the target languages/domains. They propose a unified learning algorithm based on the expectation maximization.

Using low-capacity multilingual language models such as mBERT, [Fei et al. \(2020\)](#) finds that their data-based cross-lingual transfer approach is superior to the zero-shot transfer method. However, [Li et al. \(2021\)](#) when using XLM-RoBERTa, a higher capacity multilingual model, obtain the best results for German and Chinese applying the data-based cross-lingual transfer approach, while the zero-shot approach is best for Spanish and Dutch. We extend their research on zero-resource settings with two different Sequence Labelling tasks, seven languages and three multilingual models of different capacity. Our experiments and the error analysis carried out establish the required conditions on which zero-shot and data-transfer approaches outperform each other.

3 Translation and projection method

Our data-based cross-lingual transfer method to perform cross-lingual sequence labelling is the following: we assume our source language to be English, for which we have train and development data. Furthermore, we also assume that the only gold-labelled data available for the target language is the evaluation set. In this setting, we automatically generate data for the target language by translating the gold-labelled English data. Then we project the gold labels from the source sentences to the translated sentences by leveraging automatic word alignments. Given a sentence $x = \langle x_1, \dots, x_n \rangle$ with length n in the source language and a translated sentence $y = \langle y_1, \dots, y_m \rangle$ with length m in the target language, we use a word aligner to find a set of pairs $A = \{ \langle x_i, y_j \rangle : x_i \in \mathbf{x}, y_j \in \mathbf{y} \}$ where for each word pair $\langle x_i, y_j \rangle$ y_i is the lexical translation of x_j . Next, given a sequence $s = \langle x_a, \dots, x_b \rangle \in \mathbf{x}$ labeled with a category C we will label the sequence $t = \langle y_c, \dots, y_d \rangle \in \mathbf{y}$ with category C if $\{ \forall y_j \in t \exists x_i \in s : (\langle x_i, y_j \rangle \in A) \}$. In other




Figure 2: Illustration of the translation and annotation projection method for Opinion Target Extraction (OTE).

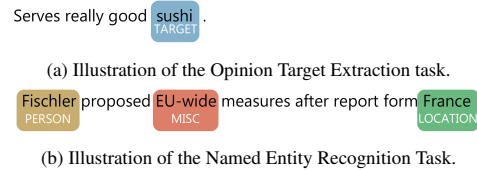


Figure 3: Sequence Labelling tasks used in our experiments.

words, if a word labelled with a category in the source sentence is aligned to a word in the target sentence, we label the target word with the category from the word in the source sentence. Figure 2 illustrates our method.

When projecting annotations we find two main problems: *split annotations* and *annotation collision*. In the first case, a labeled sequence in the source sentence is split into multiple sequences in the target sentence. This happens when the alignment for a word is missing. In this case, we merge the sequences in the target sentence if the gap between them is just one word. If we still end up with multiple sequences, we choose the longest one. In the *annotation collision* case, a word in the target sentence is aligned to two different labelled sequences in the source language. If the two sequences are of the same category, we merge them and we label the two sequences as a single one in the target sequence. If they are of different category we just consider the one with the longest length. Finally, if a punctuation symbol in the target sequence is aligned to a labeled word in the source sentence we remove this alignment.

4 Datasets

We conducted experiments in two sequence labelling tasks, namely, Opinion Target Extraction (OTE) and Named Entity Recognition (NER). Figure 3 illustrates both tasks.

Opinion Target Expression (OTE): Given a

review, the task is to detect the linguistic expression used to refer to the reviewed entity. We use the SemEval-2016 Task 5 Aspect Based Sentiment Analysis (ABSA) datasets (Pontiki et al., 2016). We experiment with the English, Spanish, Dutch, French, Russian and Turkish datasets from the restaurant domain.

Named Entity Recognition (NER): Given a text, the task is to detect named entities and classify them in pre-defined categories. For Spanish and Dutch we use the CoNLL-2002 datasets (Tjong Kim Sang, 2002). For English and German we use the CoNLL-2003 datasets (Tjong Kim Sang and De Meulder, 2003) and for Italian we use Evalita 2009 data (Speranza, 2009). We map the Geo-Political Entities from Evalita 2009 to *location* labels to make them compatible with the CoNLL data.

5 Experimental Setup

We perform 1-to-1 annotation projection in two directions:

Translate-Train: We translate the English train and development data to the target language. We project the gold labels from the English data to the translated dataset. We then train a sequence labelling model using only the automatically generated dataset for the target language.

Translate-Test: We translate the target language test set to English. We then use a model trained in the English gold-labelled data to label the translated test set. Finally, we project the labelled sequences back to the target language.

These two data-based cross-lingual transfer approaches are compared with the **zero-shot** method in which a fine-tuned model using English gold-labelled data is evaluated by generating predictions in the target language. Finally, we also fine-tuned language models on the **gold**-labelled data, which would constitute the upper-bound in our experimental setting.

5.1 Machine Translation

We tested DeepL², MarianMT (Junczys-Dowmunt et al., 2018; Tiedemann and Thottingal, 2020), M2M100 (1.2B) (Fan et al., 2020) and mBART (mbart-large-50) (Tang et al., 2020). A qualitative analysis performed during the projection of the OTE labels established that DeepL produced the more fluent translations. Thus, we decided to

²<https://www.deepl.com/>

use DeepL (web version during the second half of 2021) to perform the machine translation for our data-based cross-lingual transfer experiments. The exception was Turkish, which is not supported by DeepL. In this case we use M2M100.

5.2 Word Alignments

For word alignments, we use the AWESoME (Dou and Neubig, 2021) system. AWESoME leverages multilingual pretrained language models and fine-tune them on parallel text. Unsupervised training objectives over the parallel corpus improve the alignment quality of the models. AWESoME authors claim that the model works best with mBERT (Devlin et al., 2019) as backbone, so we follow their advice. Although we also experimented with GIZA++ (Och and Ney, 2003), FastAlign (Dyer et al., 2013) and SimAlign (Dou and Neubig, 2021), systems based on alignments from AWESoME produced the highest F1 scores when comparing the model projections and manually annotated projections (see Section 7).

To train the alignment models we use the English gold-labelled dataset together with the respective MT system translations as parallel corpora. We augment the training data with 50,000 random parallel sentences from ParaCrawl v8 (Esplà et al., 2019) for all the language pairs except Turkish, for which we use 50,000 random parallel sentences from the raw CCAIghed v1 corpus (El-Kishky et al., 2020). CCAIghed has received some criticism (Kreutzer et al., 2022), but the available English-Turkish parallel data is very limited. In Section 7 we analyze the performance of the alignment systems, and we show that CCAIghed does not hurt the performance of the aligners.

5.3 Sequence Labelling Models

We use three state-of-the-art multilingual pretrained language models for sequence labelling: multilingual BERT (mBERT) (Devlin et al., 2019) and XLM-RoBERTa (XLM-R) base and large (Conneau et al., 2020). For both models, we add a token classification layer (linear layer) on top of each token representation. We use the sequence labelling implementation of the Huggingface open-source (Apache-2.0 License) library (Wolf et al., 2019). F1 scores and standard deviation scores are reported by averaging the results of 5 runs with different random seeds. Details on models sizes, hyper-parameters and datasets are provided in the Appendix (A, B and C).

6 Experiments

6.1 Opinion Target Extraction

Opinion Target Extraction (OTE) results are reported in Table 1. The zero-shot model transfer using mBERT obtains better results for Spanish and French. However, for Dutch, Russian and Turkish the best results are obtained by the data-transfer approaches. The overall picture changes when using XLM-RoBERTa (XLM-R) base. First, the zero-shot baseline is much closer to the gold upper bound than that of mBERT. This shows that XLM-R has better multilingual transfer learning capabilities for this task. In fact, the zero-shot transfer outperforms the translate-train and translate-test approaches for all languages except Turkish. Second, the XLM-R base results on gold-labelled data are substantially better than those of mBERT. Finally, XLM-R large offers the best cross-lingual transfer capabilities, as the zero-shot transfer is clearly superior for every language, including Turkish.

A common trait for all three models in the OTE benchmark is that the translate-train approach is superior to the translate-test approach in the large majority of the cases. As expected, all the approaches achieve a performance significantly lower than the gold upper bound.

6.2 Named Entity Recognition

If we compare the OTE results with those obtained for NER (Table 2), we see a number of different patterns. First, the zero-shot approach using mBERT outperforms the data-based cross-lingual transfer methods (translate-train and translate-test) for the majority of languages. Second, unlike in OTE, the translate-test is systematically better than translate-train. Third, the mBERT performance on CoNLL data is similar to that of XLM-R base. Finally, fine-tuning XLM-R base on translated and projected data obtains better results for German and Italian than the zero-shot method. However, XLM-R large provides obtains the same results as for OTE, obtaining the best results for every language in the zero-shot setting. This validates the findings of the OTE results, namely, that the performance of the zero-shot method heavily depends on the characteristics of the multilingual language model used.

Previous research has demonstrated that cross-lingual transfer with mBERT works best for topologically similar languages (Pires et al., 2019; Wu and Dredze, 2020), which is somewhat coherent with the results obtained for Spanish and

mBERT				
	Gold	Zero-shot	Trans-Train	Trans-Test
EN	76.2 ± 0.9	-	-	-
ES	75.2 ± 0.5	68.4 ± 0.6	67.9 ± 0.8	62.2 ± 1.2
FR	74.0 ± 1.1	62.7 ± 1.2	59.7 ± 1.2	57.6 ± 1.1
NL	69.7 ± 0.9	61.7 ± 0.8	64.3 ± 1.5	67.0 ± 0.8
RU	72.5 ± 0.5	53.8 ± 2.2	62.9 ± 0.6	59.7 ± 0.4
TR	62.0 ± 1.2	45.3 ± 4.0	45.7 ± 2.3	35.5 ± 2.4

XLM-R base				
	Gold	Zero-shot	Trans-Train	Trans-Test
EN	84.4 ± 0.9	-	-	-
ES	81.1 ± 0.7	78.2 ± 0.4	72.5 ± 0.7	62.9 ± 0.9
FR	80.2 ± 0.6	72.7 ± 0.3	64.7 ± 0.8	60.0 ± 0.6
NL	80.8 ± 1.7	75.5 ± 0.8	70.0 ± 1.6	71.0 ± 1.5
RU	81.5 ± 0.3	74.9 ± 0.9	69.5 ± 0.3	62.2 ± 1.6
TR	69.0 ± 1.1	58.1 ± 3.5	58.9 ± 1.8	36.4 ± 1.8

XLM-R large				
	Gold	Zero-shot	Trans-Train	Trans-Test
EN	86.4 ± 1.1	-	-	-
ES	83.6 ± 0.1	79.3 ± 0.8	73.7 ± 1.1	64.0 ± 1.4
FR	82.2 ± 0.6	74.6 ± 1.7	66.1 ± 0.6	60.7 ± 0.6
NL	80.4 ± 2.1	77.7 ± 1.9	74.0 ± 1.3	72.9 ± 1.8
RU	82.8 ± 0.4	76.8 ± 1.3	69.3 ± 2.3	62.2 ± 1.3
TR	72.3 ± 2.4	62.4 ± 1.0	57.8 ± 2.4	33.7 ± 0.9

Table 1: OTE F1 score with models of different capacity.

French, where the zero-shot transfer is superior to the Translate-train and Translate-test approaches, while it is worse for Russian and Turkish. Additionally, it is worth noting that mBERT has been trained using only Wikipedia text for 104 languages.

In contrast, XLM-R (both base and large) have been trained using CommonCrawl (Wenzek et al., 2019), a much larger multilingual corpus with a variety of texts extracted from the Web, perhaps also including texts of similar domain to those in the OTE datasets. This may also account for the large differences in OTE performance between XLM-R base and mBERT. In this sense, the similar performance between mBERT and XLM-R base for NER might be partially due to the fact that the CoNLL and Evalita datasets consist of news stories in which most of the labelled entities may appear in the Wikipedia, the texts used to pre-train mBERT.

The performance of the XLM-R large shows that pretrained models with larger capacity help to obtain strong performance across languages, also for zero-shot cross-lingual methods. Still, data-based cross-lingual transfer (Translate-Train and Translate-Test) approaches remain useful if access to the required hardware for working with such large language models is not available.

Finally, Table 3 lists the results of previous methods that leverage parallel data and/or annotation projections to perform cross-lingual transfer on the NER CoNLL 2002-2003 data. By comparing previous work with our zero-shot baselines using

mBERT				
	Gold	Zero-shot	Trans-Train	Trans-Test
EN	90.7 ± 0.3	-	-	-
ES	87.4 ± 0.4	74.6 ± 0.4	69.5 ± 0.4	70.8 ± 0.6
DE	82.0 ± 0.4	71.0 ± 0.9	70.1 ± 0.3	70.6 ± 0.5
NL	90.8 ± 0.4	78.5 ± 0.5	74.4 ± 0.6	75.4 ± 0.8
IT	84.7 ± 0.3	68.2 ± 0.5	68.7 ± 0.5	70.7 ± 0.3

XLM-R base				
	Gold	Zero-shot	Trans-Train	Trans-Test
EN	90.4 ± 0.2	-	-	-
ES	87.7 ± 0.2	75.0 ± 0.4	70.1 ± 0.6	72.5 ± 0.2
DE	83.1 ± 0.3	67.9 ± 0.5	70.5 ± 0.5	70.1 ± 0.8
NL	89.8 ± 0.2	78.1 ± 0.6	73.3 ± 0.9	74.7 ± 0.4
IT	84.3 ± 0.3	71.2 ± 0.5	71.1 ± 0.4	71.7 ± 0.3

XLM-R large				
	Gold	Zero-shot	Trans-Train	Trans-Test
EN	92.4 ± 0.1	-	-	-
ES	88.9 ± 0.2	79.5 ± 1.0	70.9 ± 0.6	74.0 ± 0.5
DE	85.1 ± 0.6	74.5 ± 0.7	73.7 ± 0.5	72.9 ± 0.3
NL	92.9 ± 0.7	82.3 ± 0.6	77.5 ± 0.9	77.2 ± 0.6
IT	87.5 ± 0.2	76.0 ± 0.5	73.7 ± 0.4	73.5 ± 0.6

Table 2: NER F1 score with models of different capacity.

Models	ES	DE	NL
mBERT (Dou and Neubig, 2021)	64.3	-	-
BiLSTM + CRF (Jain et al., 2019)	73.5	61.5	69.9
BiLSTM + CRF (Guo and Roth, 2021)	77.9	71.4	80.6
XLM-R large (Li et al., 2021)	78.9	76.9	79.7
mBERT (Ours - zero-shot)	74.6	71.0	78.5
XLM-R base (Ours - zero-shot)	75.0	67.9	78.1
XLM-R large (Ours - zero-shot)	79.5	74.5	82.3
XLM-R base (Ours - Translate train)	70.1	70.5	73.3
XLM-R base (Ours - Translate test)	72.5	70.1	74.7
XLM-R large (Ours - Translate train)	70.9	73.7	77.5
XLM-R large (Ours - Translate test)	74.0	72.9	77.2

Table 3: Comparison between the previous research methods that leverage projections, the zero-shot baselines and our annotation projections in the 2002-2003 NER CoNLL datasets. F1 score reported

mBERT, XLM-R base and XLM-R large, we can see that the XLM-R large in the zero-shot setting still outperforms most previous approaches. The only exception being the results obtained by Li et al. (2021) for German.

7 Error Analysis

The experiments described in Section 6 showed that translate-train and translate-test perform worse than the zero-shot approach when using XLM-R large. In this section we will assess the performance of the machine translation and word alignment models. Furthermore, we will undertake an error analysis to better understand the shortcomings of translate-train and translate-test with respect to the zero-shot cross-lingual transfer.

7.1 Evaluating the Projection Method

We start our experiments by analyzing the quality of our automatically projected annotations. In order to do that, human annotators manually projected the labels from the English OTE gold-labelled data to the automatic translations to Spanish, French and Russian using DeepL and M2M100 for Turkish. The annotators are NLP PhD candidates with either native and/or proficient skills in both English and the target language. See Section E for more details.

We compare the projections of the annotations automatically generated by the different word alignment methods with those provided by the human annotators. Table 4 shows that the language model-based methods (SimAlign and AWESoME) outperform the statistically based methods (GIZA++ and FastAlign) by a wide margin in all languages. Furthermore, AWESoME consistently outperforms SimAlign for every language. The performance of the AWESoME system confirms that it is possible to generate high quality annotations close to those generated by human experts. The results also show that for Turkish performance is lower than for the other languages. This is the case for the methods that require parallel data (GIZA++, FastAlign and AWESoME) as well as SimAlign that does not require parallel data. So we can attribute the lower performance to the difficulty of projecting annotations for the English-Turkish pair and not the usage of the CCAligned corpus.

	GIZA++	FastAlign	SimAlign	AWESoME
ES	77.0	75.0	86.7	91.5
FR	73.3	72.9	86.3	91.3
RU	72.4	76.9	87.7	93.4
TR	64.0	68.4	81.9	88.5

Table 4: OTE F1 score between the human annotation projections vs the automatic projections generated using different alignment models.

While Table 4 shows that we generate high quality annotation projections, the best model, AWESoME, still makes some mistakes. To explore the effect of these mistakes we fine-tune XLM-R large models on the manually projected train datasets and compare their performance on the gold-labelled test sets with the models trained on the AWESoME automatically projected data. Table 5 shows that the models obtained using the manually projected data are slightly better, except for Turkish, which once again acts as outlier. In any case, as the results obtained by fine-tuning on the manually projected data are still worse than the zero-shot method, this

experiment proves that the projection of annotations is not responsible for the data-based cross-lingual transfer methods to be inferior to the zero-shot baseline.

XLMR	Trans-Train	Trans-Train (Manual)
ES	73.7 \pm 1.1	75.1 \pm 1.2
FR	66.1 \pm 0.6	67.9 \pm 1.0
RU	69.3 \pm 1.3	69.4 \pm 2.1
TR	57.8 \pm 2.4	50.6 \pm 1.4

Table 5: XLM-R large OTE F1 score when training with automatically and manually projected datasets

7.2 Downstream Evaluation of Machine Translation Models

In order to evaluate the influence of the machine translation system used, we translate the English gold-labelled data using four different translation systems. We fixed AWESoME as the word aligner for annotation projection. We fine-tune XLM-R large with each of the generated training data and evaluate it against the gold-labelled test data from OTE. As Table 6 shows, there are no big differences in the final F1 scores when using different translation systems (Turkish is again being the exception), we decided to carry on using DeepL based on the manual assessment mentioned in Section 3.

	MarianMT	Mbart	M2M100	DeepL
ES	75.6 \pm 0.8	75.3 \pm 0.7	74.2 \pm 0.8	73.7 \pm 1.1
FR	64.5 \pm 1.6	66.4 \pm 1.1	64.9 \pm 1.3	66.1 \pm 0.6
NL	70.0 \pm 2.0	68.8 \pm 4.0	70.1 \pm 3.1	74.0 \pm 1.3
RU	66.6 \pm 4.4	69.7 \pm 1.4	69.7 \pm 0.7	69.3 \pm 2.3
TR	49.5 \pm 2.9	56.1 \pm 5.2	57.8 \pm 2.4	-

Table 6: OTE F1 score of different XLM-R large models trained using data generated with different translation systems.

7.3 Where do the models fail?

To better understand what is happening we identify the most common false negatives and positives for both OTE and NER tasks. Table 7 shows the most frequent false negatives and positives where there is a big mismatch between methods.

As it has been previously noticed (Agerri and Rigau, 2019), in the ABSA data the words “comida” (food) and “restaurante” (restaurant) are highly ambiguous, so we could expect the models to fail with these words. In addition, we have found out 4 main sources of errors, which are analyzed below.

Many-to-one translations: This is stereotypical of targets such as “trato” and “atención” in Span-

ish, which, in addition to “servicio”, are used to refer to “service” in English. There are 160 sentences in the English gold-labelled data containing the word “service”; in 153 of them “service” is labelled as target. DeepL systematically translates it as “servicio”. However, as shown by Table 8, in the Spanish gold-labelled data “service” is also commonly referred as “trato” or “atención”, instead of “servicio”.

This would result in a training set without any occurrences of “trato” and “atención” which often occur in the gold-labelled test data. Both the zero-shot and the data-based cross-lingual transfer approaches fail to correctly label these words, which shows a problem of using automatically translated data. Interestingly, the zero-shot approach using XLM-R large correctly classifies “trato” (only fails to label 1 of the 19 occurrences). As shown by our experimental results, XLM-R large is more robust than mBERT and XLM-R base.

Something similar happens with the word “place”, which in Spanish can be most frequently translated as “lugar” or “sitio”. However, DeepL almost always translates it as “lugar” which results in “sitio” being absent in the automatically generated training data while being more frequent than “lugar” in the gold-labelled data. Note that this is not a problem for the “translate-test”, given that the translation direction is Spanish to English.

Errors induced by incorrect or missing alignments: For NER we found errors of different nature. Articles and prepositions (i.e. “de”, “la”) are among the words with higher false positive rate for the translate-test and translate-train approaches. We can attribute it to word alignment errors. Large multi-word named entities such as “Consejo General de la Arquitectura Técnica de España” (General Council of Technical Architecture of Spain) are labelled as entities. Word aligners struggle to correctly align articles in these complex expressions specially when a one-to-many or many-to-one alignment is required. In fact, in this example, the word aligners we tested failed to correctly align “of” with “de la”.

Errors induced by dataset inconsistencies: Another issue is the differences across languages in the original gold-labelled annotations. Thus, “Gobierno” (Government) and “Estado” (State) are labelled as organizations in the Spanish gold-labelled data, but they are not considered to be entities in

the English gold-labelled data. The opposite occurs with demonym words. They are labelled as miscellaneous entities in the English data but in Spanish they are not annotated. Cross-lingual models are likely to fail labelling these cases.

Lost in Translation: Finally, there is another group of words related to Spanish Government names which are not commonly used in English for the same contexts (i.e. “Economía” to refer to the “Ministry of Economy” or “Ministerio de Economía” in Spanish, “Junta” for “local government”, or “Plan” for “government projects”). While these words appear frequently in the Spanish data as part of commonly used named entities, that is not the case in the English data, where it is customary to use “Treasury Department” (or variations thereof) which are correctly translated into Spanish by DeepL as “Departamento del Tesoro”. This means that, during fine-tuning on the translated data, the model is not receiving any signal to learn that “Economy” may be part of an organization entity. This may explain why the zero-shot method performs better for cases such as “Economía”, “Hacienda”, “Plan” and “Junta”, listed in Table 7.

Summarizing, we see that machine translation data often generates a signal which is, due to inherent differences in language use, different to the signal received when using gold-labelled data in the target language. This disagreement seems to be the most common reason for the larger number of false positive and negatives of the data-based cross-lingual transfer method with respect to the zero-shot technique.

8 Concluding Remarks

In this paper we described an in-depth and comprehensive evaluation of model-based and data-based zero-resource cross-lingual sequence labelling on two different tasks.

Contrary to what previous research suggests, zero-shot transfer approach is the best performing method when using high capacity multilingual language models such as XLM-R large. However, data-based cross-lingual transfer approaches are still useful when having a model with poor downstream cross-lingual performance. For example, when using a pretrained language model not trained for a specific domain, or when the required hardware for working with such larger language models is not readily available.

	GOLD			Zero-shot			Tr-Train			Tr-Test			Total
	B	Xb	Xl	B	Xb	Xl	B	Xb	Xl	B	Xb	Xl	
OTE False Negatives													
comida	3	3	2	6	2	1	4	1	1	1	9	5	98
restaurante	7	5	7	9	5	6	7	6	6	7	12	10	43
servicio	2	2	2	2	1	1	2	0	1	1	1	2	85
trato	1	1	0	5	6	1	14	10	5	6	8	6	19
atención	2	3	3	8	2	3	7	1	3	7	7	7	13
lugar	0	0	0	2	0	0	1	0	0	0	1	0	12
sitio	1	0	0	5	1	1	3	3	3	2	1	1	14
NER False Negatives													
de	32	29	33	45	51	90	233	252	264	148	146	167	450
la	4	5	3	10	12	16	63	62	62	45	44	45	174
Gobierno	0	0	0	17	53	64	72	70	75	30	45	67	80
Estado	0	0	0	4	4	8	9	8	9	6	6	8	10
Administración	0	0	0	4	8	11	10	11	11	5	5	7	11
Economía	0	0	0	2	6	2	7	8	8	5	6	8	8
Plan	0	0	0	1	2	2	3	5	5	1	4	7	8
Junta	0	0	0	0	0	0	4	10	8	2	3	5	24
Hacienda	0	0	0	1	3	0	4	4	4	4	3	4	5
NER False Positives													
español	0	0	0	16	16	2	16	16	12	13	14	15	0
catalán	0	0	0	8	8	5	7	7	8	8	8	8	0

Table 7: Most common false negatives and positives were there is a big mismatch between methods and the total number of labelled appearances of the word in the test data. B is the acronym for mBERT, Xb for XLM-R base and Xl for XLM-R large.

En.Word	Es.Word	En Gold	Es Gold	Es Translate
Service	Servicio	153	229	133
Treatment	Trato	0	54	0
Attention	Atención	2	35	0
Place	Sitio	120	41	2
Place	Lugar	120	19	91

Table 8: Number of times words appear as target words in the train datasets

A detailed error analysis demonstrates that data-based cross-lingual transfer is hindered by machine translations which, although linguistically sound, do not align with the cultural behaviour of the target language use. Moreover, the results also show that the different word alignments methods (for annotation projection) are of high quality, obtaining comparable results with respect to manually generated alignments.

In any case, our results establish that there is still room for improving the cross-lingual performance of zero-resource sequence labelling.

Acknowledgments

We are grateful to Nayla Escribano, Suna Şeyma Uçar and Olya Toporkov for their help in manually projecting the labels from the English gold-labelled data into the automatically translated dataset. We are also thankful to the anonymous reviewers for their insightful comments. This work has been supported by the following projects: (i)

DeepKnowledge (PID2021-127777OB-C21) funded by MCIN/AEI/10.13039/501100011033 and FEDER Una manera de hacer Europa; (ii) Disargue (TED2021-130810B-C21), MCIN/AEI/10.13039/501100011033 and European Union NextGenerationEU/PRTR; (iii) DeepR3 (TED2021-130295B-C31) funded by MCIN/AEI/10.13039/501100011033 and EU NextGeneration programme EU/PRTR; (iv) Antidote (PCI2020-120717-2) funded by MCIN/AEI/10.13039/501100011033 and by European Union NextGenerationEU/PRTR. Iker García-Ferrero is supported by a doctoral grant from the Basque Government (PRE_2021_2_0219). Rodrigo Agerri's work is also supported by the RYC-2017-23647 fellowship (MCIN/AEI/10.13039/501100011033 y por El FSE invierte en tu futuro).

Limitations

We compare baseline cross-lingual zero-shot model transfer with machine translation and annotation projection. We do not explore alternative cross-lingual data-based methods, such as the usage of available parallel corpora instead of a machine translated corpus. We also skip evaluating methods to improve model-transfer approaches such as the ones described in Section 2.2. We may also consider that our annotation projection approach and

zero-shot model transfer approach work for Indo-European languages, while their performance for other language families remains unknown. Finally, the error analysis was performed for the EN-ES language pair only.

In any case, we believe that our main claim still holds. Even though MT quality has substantially improved over the last few years, our results indicate the current optimal solution to perform cross-lingual transfer is by using large multilingual language models such as XLM-RoBERTa-large. Thus, our error analysis suggests that this might be due to important differences in language use. More specifically, MT often generates a textual signal which is different to what the models are exposed to when using gold standard data, which affects both the fine-tuning and evaluation processes. This is confirmed by our error analysis which shows that mistranslations are not the main source of errors in the data-transfer method.

References

- 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).
- Rodrigo Agerri and German Rigau. 2016. [Robust multilingual named entity recognition with shallow semi-supervised features](#). *Artificial Intelligence*, 238:63–82.
- Rodrigo Agerri and German Rigau. 2019. [Language independent sequence labelling for opinion target extraction](#). *Artificial Intelligence*, 268:85–95.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. [Contextual string embeddings for sequence labeling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- M Saiful Bari, Tasnim Mohiuddin, and Shafiq Joty. 2021. [Uxla: A robust unsupervised data augmentation framework for zero-resource cross-lingual nlp](#).
- 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.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Zi-Yi Dou and Graham Neubig. 2021. [Word alignment by fine-tuning embeddings on parallel corpora](#). *CoRR*, abs/2101.08231.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*, pages 644–648. The Association for Computational Linguistics.
- Maud Ehrmann, Marco Turchi, and Ralf Steinberger. 2011. [Building a multilingual named entity-annotated corpus using annotation projection](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 118–124, Hissar, Bulgaria. Association for Computational Linguistics.
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. 2020. [CCAligned: A massive collection of cross-lingual web-document pairs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 5960–5969, Online. Association for Computational Linguistics.
- Miquel Esplà, Mikel Forcada, Gema Ramírez-Sánchez, and Hieu Hoang. 2019. [ParaCrawl: Web-scale parallel corpora for the languages of the EU](#). In *Proceedings of Machine Translation Summit XVII: Translator, Project and User Tracks*, pages 118–119, Dublin, Ireland. European Association for Machine Translation.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2020. [Beyond english-centric multilingual machine translation](#). *CoRR*, abs/2010.11125.
- 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.
- Ruohao Guo and Dan Roth. 2021. [Constrained labeled data generation for low-resource named entity recog-](#)

- tion. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4519–4533, Online. Association for Computational Linguistics.
- Zechuan Hu, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. [Risk minimization for zero-shot sequence labeling](#). 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 4909–4920, Online. Association for Computational Linguistics.
- Alankar Jain, Bhargavi Paranjape, and Zachary C. Lipton. 2019. [Entity projection via machine translation for cross-lingual NER](#). 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)*, pages 1083–1092, Hong Kong, China. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Julia Kreutzer, Isaac Caswell, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmungkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iroko Orife, Kelechi Ogueji, Andre Niyongabo Rubungo, Toan Q. Nguyen, Mathias Müller, André Müller, Shamsuddeen Hassan Muhammad, Nanda Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, Mathias Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine Çabuk Ballı, Stella Biderman, Alessia Battisti, Ahmed Baruwa, Ankur Bapna, Pallavi Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. 2022. [Quality at a Glance: An Audit of Web-Crawled Multilingual Datasets](#). *Transactions of the Association for Computational Linguistics*, 10:50–72.
- 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.
- Bing Li, Yujie He, and Wenjin Xu. 2021. [Cross-lingual named entity recognition using parallel corpus: A new approach using xlm-roberta alignment](#). *CoRR*, abs/2101.11112.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2021. [Crossner: Evaluating cross-domain named entity recognition](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13452–13460. AAAI Press.
- Jian Ni, Georgiana Dinu, and Radu Florian. 2017. [Weakly supervised cross-lingual named entity recognition via effective annotation and representation projection](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1470–1480, Vancouver, Canada. Association for Computational Linguistics.
- Franz Josef Och and Hermann Ney. 2003. [A systematic comparison of various statistical alignment models](#). *Computational Linguistics*, 29(1):19–51.
- Xuan Ouyang, Shuohuan Wang, Chao Pang, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. [Ernie-m: Enhanced multilingual representation by aligning cross-lingual semantics with monolingual corpora](#).
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. [SemEval-2016 task 5: Aspect based sentiment analysis](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30, San Diego, California. Association for Computational Linguistics.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Annual Meeting of the Association for*

- Computational Linguistics*, pages 151–164, Florence, Italy. Association for Computational Linguistics.
- Manuela Speranza. 2009. The named entity recognition task at evalita 2009. In *EVALITA 2009*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *CoRR*, abs/2008.00401.
- Jörg Tiedemann and Santhosh Thottingal. 2020. OPUS-MT – building open translation services for the world. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Erik F. Tjong Kim Sang. 2002. Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition. In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- 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.
- Mengqiu Wang and Christopher D. Manning. 2014. Cross-lingual projected expectation regularization for weakly supervised learning. *Transactions of the Association for Computational Linguistics*, 2:55–66.
- Zirui Wang, Jiateng Xie, Ruochen Xu, Yiming Yang, Graham Neubig, and Jaime Carbonell. 2019. Cross-lingual alignment vs joint training: A comparative study and a simple unified framework. *arXiv preprint arXiv:1910.04708*.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2019. Ccnet: Extracting high quality monolingual datasets from web crawl data. *CoRR*, abs/1911.00359.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *CoRR*, abs/1910.03771.
- Qianhui Wu, Zijia Lin, Börje Karlsson, Jian-Guang Lou, and Biqing Huang. 2020. Single-/multi-source cross-lingual NER via teacher-student learning on unlabeled data in target language. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6505–6514, Online. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2020. Are all languages created equal in multilingual BERT? In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 120–130, Online. Association for Computational Linguistics.
- Jiateng Xie, Zhilin Yang, Graham Neubig, Noah A. Smith, and Jaime Carbonell. 2018. Neural cross-lingual named entity recognition with minimal resources. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 369–379, Brussels, Belgium. Association for Computational Linguistics.

A Model size

We experiment with multilingual BERT (mBERT) (Devlin et al., 2019) and XLM-RoBERTa (XML-R) base and large (Conneau et al., 2020). We list the number of parameters of each model in Table 9

Model	#params
multilingual BERT	110M
XLM-RoBERTa-base	250M
XLM-RoBERTa-large	560M

Table 9: Number of parameters for the language models that we use in our experiments

B Hyper parameters

B.1 Word alignment models

We train AWESoME with 8 batch size and $2e - 5$ learning rate for 40,000 steps, with all the unsupervised training objectives (mlm, tlm, tlm_full, so, psi) and softmax extraction method. We use mBERT as backbone. For SimAlign we run inference with 0.0 distortion rate, 1.0 null align rate and the "itermax" matching method. We use bpe tokens and mBERT backbone. We use the MGIZA multicore implementation³ of GIZA++ with the recommended configuration file⁴. We use FastAlign with the default hyper-parameters. For both, GIZA++ and FastAlign we combine the forward and backward directions of the alignments using the grow-diag-final-and algorithm.

B.2 Sequence Labelling models

For OTE we use a batch size of 32, $5e - 5$ learning rate, we train the model for 10 epochs and 128 maximum sequence length. Since only a train and test splits are available for the OTE task, we

³<https://github.com/moses-smt/mgiza>

⁴<https://pastebin.com/b1ksHtUy>

use the train set as both, train and development data. For NER we use a batch size of 32, $2e - 5$ learning rate, we train the model for 4 epochs and 256 maximum sequence length. We use the default values (sequence labelling implementation of the Huggingface library ⁵) for the remaining hyperparameters. For both tasks we use the BILOU encoding scheme.

C Datasets Size

We list the dataset size (number of sentences) of the datasets we use.

For OTE we use the SemEval-2016 Task 5 Aspect Based Sentiment Analysis (ABSA) datasets (Pontiki et al., 2016). We list the size of the datasets in Table 10.

Lang	Train	Test
EN	2000	676
ES	2070	881
FR	1664	668
NL	1722	575
RU	3655	1209
TR	1232	144

Table 10: Number of sentences in the OTE datasets

For NER we use the Spanish and Dutch data from the CoNLL-2002 datasets (Tjong Kim Sang, 2002). For English and German we use the CoNLL-2003 datasets (Tjong Kim Sang and De Meulder, 2003) and for Italian we use Evalita 2009 data (Speranza, 2009). We list the size of these datasets in Table 11.

	Train	Dev	Test
EN	14987	3466	3684
ES	6871	1914	1516
DE	12705	3068	3160
NL	15806	2895	5195
IT	11227	0	4136

Table 11: Number of sentences in the NER datasets

D Computer infrastructure

We perform all our experiments using a single NVIDIA A30 GPU with 24GB memory. The machine used has two Xeon Gold 6226R CPUs and 256GB RAM.

⁵<https://github.com/huggingface/transformers/tree/main/examples/pytorch/token-classification>

E Manual Projection of the datasets

Human annotators manually projected the labels from the English OTE gold data to the automatic translations to Spanish, French and Russian using DeepL and m2m10 for Turkish. The annotators are NLP PhD candidates with either native and/or proficient skills in both English and the target language. We describe the experiment in Section 7.1. For the purpose of this experiment, we developed an application to assist during the annotation process. The annotator sees the sentence in English, where there is a highlighted target and must select the same target in a translated target sentence. Figure 4 shows two screenshots from the application. The full guidelines and the code of the application provided to the annotators are available at <https://github.com/ikergarcia1996/Annotation-Projection-App>.

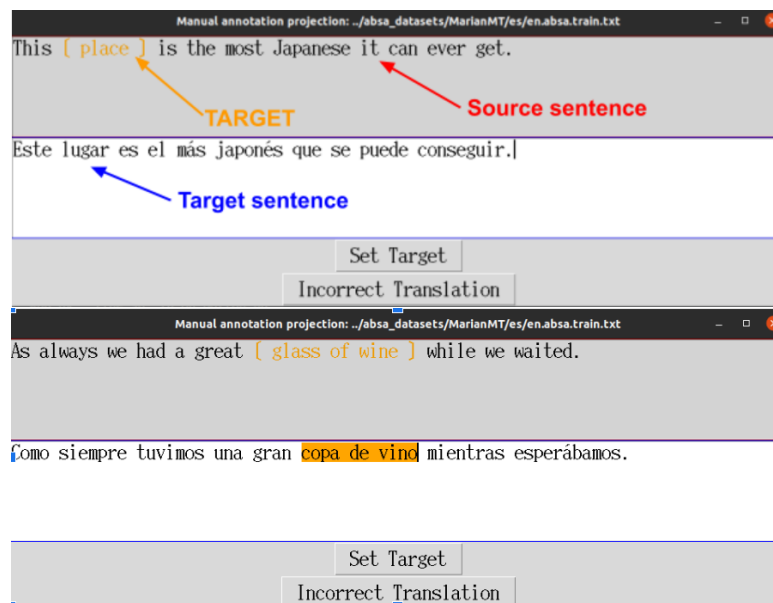


Figure 4: Application used to manually annotate the projections

T-Projection: High Quality Annotation Projection for Sequence Labeling Tasks

Iker García-Ferrero Rodrigo Agerri German Rigau
HiTZ Center - Ixa, University of the Basque Country UPV/EHU
{ iker.garciaf, rodrigo.agerri, german.rigau }@ehu.eus

Abstract

In the absence of readily available labeled data for a given sequence labeling task and language, annotation projection has been proposed as one of the possible strategies to automatically generate annotated data. Annotation projection has often been formulated as the task of transporting, on parallel corpora, the labels pertaining to a given span in the source language into its corresponding span in the target language. In this paper we present T-Projection, a novel approach for annotation projection that leverages large pretrained text-to-text language models and state-of-the-art machine translation technology. T-Projection decomposes the label projection task into two subtasks: (i) A candidate generation step, in which a set of projection candidates using a multilingual T5 model is generated and, (ii) a candidate selection step, in which the generated candidates are ranked based on translation probabilities. We conducted experiments on intrinsic and extrinsic tasks in 5 Indo-European and 8 low-resource African languages. We demonstrate that T-projection outperforms previous annotation projection methods by a wide margin. We believe that T-Projection can help to automatically alleviate the lack of high-quality training data for sequence labeling tasks. Code and data are publicly available.¹

1 Introduction

The performance of supervised machine-learning methods for Natural Language Processing, including advanced deep-neural models (Devlin et al., 2019; Conneau et al., 2020; Xue et al., 2021; Zhang et al., 2022), heavily depends on the availability of manually annotated training data. In addition, supervised models show a significant decrease in performance when evaluated in out-of-domain settings (Liu et al., 2021). This means that obtaining optimal results would require to manually generate

¹<https://github.com/ikergarcia1996/T-Projection>

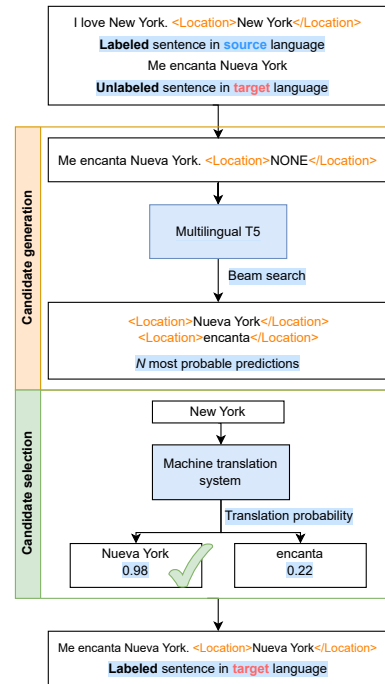


Figure 1: T-Projection two-step method to project sequence labeling annotations across languages.

annotated data for each application domain and language, an unfeasible task in terms of monetary cost and human effort. As a result, for the majority of languages in the world the amount of manually annotated data for many downstream tasks is simply nonexistent (Joshi et al., 2020).

The emergence of multilingual language models (Devlin et al., 2019; Conneau et al., 2020) allows for zero-shot cross-lingual transfer. A model fine-tuned on one language, typically English, can be directly applied to other target languages. However, better results can be obtained by either machine translating the training data from English into the

target languages or, conversely, translating the test data from the target language into English (Hu et al., 2020; Artetxe et al., 2023).

Sequence labeling tasks, which involve span-level annotations, require an additional step called *annotation projection*. This step involves identifying, in the translated sentence, the sequence of words that corresponds to the labeled spans in the source text (Yarowsky et al., 2001; Ehrmann et al., 2011). The majority of previous published work on this line of research explores the application of word-alignments (Ehrmann et al., 2011). However, projection methods based on word-alignments have achieved mixed results as they often produce partial, incorrect or missing annotation projections (García-Ferrero et al., 2022). This is due to the fact that word alignments are computed on a word-by-word basis leveraging word co-occurrences or similarity between word vector representations. That is, without taking into consideration the labeled spans or categories to be projected. Other techniques have also been proposed, such as fine-tuning language models in the span projection task (Li et al., 2021), translating the labeled spans independently from the sentence (Zhou et al., 2022) or including markers during the machine translation step (Chen et al., 2023). However, automatic annotation projection remains an open and difficult challenge.

In this paper we present T-Projection, a novel approach to automatically project sequence labeling annotations across languages. Our method is illustrated by Figure 1. We split the annotation projection task into two steps. First, we use mT5 (Xue et al., 2021) text-to-text model to generate a set of projection candidates in the target sentence for each labeled category in the source sentence. This step exploits the labeled categories as well as the cross-lingual capabilities of large pretrained multilingual language models. Second, we rank the candidates based on the probability of being generated as a translation of the source spans. We use the M2M100 (Fan et al., 2021) and NLLB200 (Costa-jussà et al., 2022) state-of-the-art MT models to compute the translation probabilities (Vamvas and Sennrich, 2022).

We conduct an intrinsic evaluation on three different tasks, Opinion Target Extraction (OTE), Named Entity Recognition (NER) and Argument Mining (AM), and five different target languages (French, German, Italian, Russian and Spanish). In this evaluation we compare the label projections

generated by various systems with manually projected annotations. On average, **T-Projection improves the current state-of-the-art annotation projection methods by more than 8 points in F1 score**, which constitutes a significant leap in quality over previous label projection approaches. Additionally, we performed a real-world NER task evaluation involving eight low-resource African languages. In this downstream evaluation, T-Projection outperformed other annotation projection methods by 3.6 points in F1 score.

2 Background

While most of the previous approaches for annotation projection are based on the application of word alignments, other techniques have also been explored.

2.1 Word alignment methods

Ehrmann et al. (2011) use the statistical alignment of phrases to project the English labels of a multi-parallel corpus into the target languages. Instead of using discrete labels, Wang and Manning (2014) project model expectations with the aim of facilitating the transfer of model uncertainty across languages. Ni et al. (2017) aim to filter good-quality projection-labeled data from noisy data by proposing a set of heuristics. Other works have proposed to use word alignments generated by Giza++ (Och and Ney, 2003) to project parallel labeled data from multiple languages into a single target language (Agerri et al., 2018). Fei et al. (2020) use the word alignment probabilities calculated with FastAlign (Dyer et al., 2013) and the POS tag distributions of the source and target words to project from the source corpus into a target language machine-translated corpus. Finally, García-Ferrero et al. (2022) propose an annotation projection method based on machine translation and AWESOME (Dou and Neubig, 2021), Transformer-based word alignments to automatically project datasets from English to other languages.

2.2 Other projection methods

With respect to projection methods which do not use word alignments, Jain et al. (2019) first generate a list of projection candidates by orthographic and phonetic similarity. They choose the best matching candidate based on distributional statistics derived from the dataset. Xie et al. (2018) propose a method to find word translations based

on bilingual word embeddings. Li et al. (2021) use a XLM-RoBERTa model (Conneau et al., 2020) trained with the source labeled part of a parallel corpus to label the target part of the corpus. Then they train a new improved model with both labeled parts. Zhou et al. (2022) first replace the labeled sequences with a placeholder token in the source sentence. Second, they separately translate the sentence with the placeholders and the labeled spans into the target sentences. Finally, they replace the placeholders in the translated sentence with the translation of the labeled spans. Chen et al. (2023) jointly perform translation and projection by inserting special markers around the labeled spans in the source sentence. To improve the translation accuracy and reduce translation artifacts, they finetune the translation model with a synthetic label protection dataset.

To summarize, previous work does not take advantage of all the information which is available while performing annotation projection. For example, word alignment models do not take into account the labeled spans and their categories during alignment generation. Instead, they simply rely on information about word co-occurrences or similarity between word representations. Those techniques based on MT to generate the target part of the parallel corpus ignore additional knowledge that the MT model encodes. Furthermore, methods that utilize MT models for both translation and projection often introduce translation artifacts, which can affect the quality and accuracy of the projections. In contrast, our T-Projection method exploits both the labeled spans and their categories together with the translation probabilities to produce high-quality annotation projections.

3 T-Projection

Given a source sentence in which we have sequences of words labeled with a class, and its parallel sentence in a target language, we want to project the labels from the source into the target. As illustrated in Figure 1, T-Projection implements two main steps. First, a set of projection candidates in the target sentence are generated for each labeled sequence in the source sentence. Second, each projection candidate is ranked using a machine translation model. More specifically, candidates are scored based on the probability of being generated as a translation of the source labeled sequences. While the *candidate generation* step exploits the




Figure 2: Illustration of the candidate generation step. For each label, we generate a set of probable candidates.

labeled spans and their categories in the source sentence as well as the zero-shot cross-lingual capabilities of large pretrained multilingual language models, the *candidate selection* step applies state-of-the-art MT technology to find those projection candidates that constitute the best translation for each source labeled span. As demonstrated by our empirical evaluation in Sections 5 and 6, we believe that these two techniques leverage both the available information and knowledge from the annotated text and language models thereby allowing us to obtain better annotation projections. These two steps are described in detail in the following two subsections.

3.1 Candidate Generation

When trying to project labeled sequences from some source data into its parallel target dataset, we would expect both the source and the target to contain the same number of sequences labeled with the same category. For example, given the English source sentence "<Person>Obama</Person> went to <Location>New York</Location>" and its parallel Spanish unlabeled target sentence "Obama fue a Nueva York", we would expect to find the same two entities (*person* and *location*) in the target sentence. To solve the task of candidate generation, we finetune the text-to-text mT5 (Xue et al., 2021) model using a HTML-tag-style prompt template (Huang et al., 2022). As illustrated by Figure 2, the input consists of concatenating the unlabeled sentence followed by a list of tags ("<Category>None</Category>") with the category of each labeled span that we expect to find in the sentence. If two or more spans share the same category then we append the tag as many times as spans are expected with that category.

Unlike Huang et al. (2022), we do not encode the tags for each category as special tokens. Instead, we verbalize the categories (i.e PER->Person) and

Source Span	Candidates for Location	Translation Probability
<div style="border: 1px solid black; padding: 5px; width: fit-content;"> New York LOCATION </div>	Nueva York	0.98 ✓
	York	0.82 ✗
	encanta	0.22 ✗
	Nueva	0.84 ✗

Figure 3: Candidate selection: candidates are scored based on the probability of being generated as a translation of the originally labeled sequences.

use the token representations already existing in the model. We expect that, thanks to the language modeling pretraining, T5 would have a good semantic representation of categories such as *Person*, *Location*, *Claim*, etc.

As Figure 2 illustrates, we finetune mT5 with the labeled source dataset. We train the model to replace the token *None* with the sequence of words in the input sentence that corresponds to that category. We use Cross-Entropy loss for training.

At inference, we label the target sentences which are parallel translations of the source sentences. As the source tells us how many labeled spans should we obtain in the target, we use the labels of the corresponding source sentence to build the prompts. In other words, our objective is to label parallel translations of the sentences we used for training. We take advantage of the zero-shot cross-lingual capabilities of mT5 to project the labels from the source to the target sentence. The output tokens are generated in an autoregressive manner. We use beam search decoding with 100 beams to generate 100 candidates for each input tag.

3.2 Candidate Selection

As depicted in Figure 3, all the generated candidates are grouped by category. In other words, if the previous step has generated multiple spans with the same category (i.e, two *locations* in a sentence) then all the candidates are included in a single set. Furthermore, all the candidates that are not a subsequence of the input sentence are filtered out.

For each labeled span in the source sentence, we rank all the projection candidates that share the same category as the source span using their translation probabilities (also known as translation equivalence) which have been obtained by apply-

ing the pretrained M2M100 (Fan et al., 2021) or NLLB200 (Costa-jussà et al., 2022) MT models and the *NMTScore* library² (Vamvas and Sennrich, 2022). Thus, given the source span *A* and the candidate *B* the translation probability is computed as follows (Vamvas and Sennrich, 2022):

$$p_{\theta_a}(A | B) := \left[\prod_{i=0}^{|A|} p_{\theta_a}(A^i | B, A^{<i}) \right]^{\frac{1}{|A|}}$$

The translation probability is normalized:

$$\text{sim}(A | B) = \frac{p_{\theta_a}(A|B)}{p_{\theta_a}(A|A)}$$

As the translation probability varies depending on the direction of the translation, the scores are symmetrized by computing the scores of both translation directions and averaging them:

$$\text{sim}(A, B) = \frac{1}{2} \text{sim}(A | B) + \frac{1}{2} \text{sim}(B | A)$$

Finally, for each labeled span in the source sentence, we choose the candidate in the target with the highest translation probability. Once a candidate has been selected, that candidate, and any other that overlaps with it, is removed from the set of possible candidates. In this way we prevent assigning the same candidate in the target to two different spans in the source.

4 Experimental Setup

In order to evaluate our method we perform both intrinsic and extrinsic evaluations.

Intrinsic evaluation: We select a number of datasets that have been manually projected from English into different target languages. The manual annotations are used to evaluate and compare T-Projection with respect to previous state-of-the-art label projection models. Results are reported by computing the usual *F1-score* used for sequence labelling (Tjong Kim Sang, 2002) with the aim of evaluating the quality of the automatically generated projections with respect to the manual ones.

Extrinsic evaluation: In this evaluation we assess the capability of T-Projection to automatically generate training data for sequence labeling tasks, NER in this particular case. The process begins by utilizing the machine translation system NLLB200 (Costa-jussà et al., 2022) to translate data from English into 8 low-resource African languages. We then project the labels from English onto the respective target languages. The automatically generated datasets are then employed to train NER models,

²<https://github.com/ZurichNLP/nmtscore>




Figure 4: Sequence labeling tasks in our experiments

which are evaluated using a relatively small manually annotated test set. The same procedure is performed with other state-of-the-art label projection models. The comparison of the results obtained is reported in terms of *F1 score*.

4.1 Datasets

The datasets used correspond to three sequence labeling tasks which are illustrated by Figure 4.

Opinion Target Extraction (OTE) Given a review, the task is to detect the linguistic expression used to refer to the reviewed entity. We use the English SemEval 2016 Aspect Based Sentiment Analysis (ABSA) datasets (Pontiki et al., 2014). Additionally, for the evaluation we also used the parallel versions for Spanish, French and Russian generated via machine translation and manual projection of the labels (García-Ferrero et al., 2022).

Named Entity Recognition (NER) The NER task consists of detecting named entities and classifying them according to some pre-defined categories. We use an English, Spanish, German, and Italian parallel NER dataset (Agerri et al., 2018) based on Europarl data (Koehn, 2005). Manual annotation for the 4 languages was provided following the CoNLL 2003 (Tjong Kim Sang, 2002) guidelines. In the extrinsic evaluation, we use MasakhaNER 2.0 (Adelani et al., 2022), a human-annotated NER dataset for 20 African languages. **Argument Mining (AM)** In the AbstrCT English dataset two types of argument components, Claims and Premises, are annotated on medical and scientific texts collected from the MEDLINE database (Mayer et al., 2020). For evaluation we used its Spanish parallel counterpart which was generated following an adapted version of the method described above for OTE (Yeginbergenova and Agerri, 2023). In contrast to NER and OTE, the sequences in the AM task consist of very long spans of words, often encompassing full sentences. We use the Neoplasm split.

4.2 Baselines

We experiment with 4 different *word alignment systems*. Two statistical systems, **Giza++** (Och and Ney, 2003) and **FastAlign** (Dyer et al., 2013),

widely used in the field. We also evaluate two Transformer-based systems, **SimAlign** (Jalili Sabet et al., 2020) and **AWESOME** (Dou and Neubig, 2021), which leverage pre-trained multilingual language models to generate alignments. As the authors recommend, we use multilingual BERT (mBERT) (Devlin et al., 2019) as the backbone. We tested different models as backbone with no improvement in performance (see Appendix C). We use these four systems to compute word alignments between the source and the target sentences and generate the label projections applying the algorithm published by García-Ferrero et al. (2022)³.

We also experiment with **EasyProject** (Chen et al., 2023), a system that jointly performs translation and projection by inserting special markers around the labeled spans in the source sentence. As this model generates its own translations it is therefore not suitable for the intrinsic evaluation which is why we only used it for the extrinsic evaluation.

We implemented two additional baselines for comparison. In the first baseline, inspired by Li et al. (2021), we use **XLM-RoBERTa** (Conneau et al., 2020) 3 billion parameter model (same parameter count as the mT5 model that we use in T-Projection) with a token classification layer (linear layer) on top of each token representation. We train the model in the source labeled dataset and we predict the entities in the translated target sentences. The second baseline adopts a **span translation** approach inspired by Zhou et al. (2022). We translate the labeled spans in the source sentence using the pretrained M2M100 12 billion parameter model and we match them with the corresponding span in the target sentence. For example, given the labeled source sentence "*I love <Location> New York </Location>.*" and the target sentence "*Me encanta Nueva York*", we translate the span *New York* into the target language, resulting in the translated span *Nueva York* which is then matched in the target sentence. We employ beam search to generate 100 possible translations, and we select the most probable one that matches the target sentence.

³<https://github.com/ikergarcia1996/Easy-Label-Projection>

4.3 Models Setup

We use the 3 billion parameters pretrained mT5 (Xue et al., 2021) for the *candidate generation* step while *candidates are selected* using the M2M100 12 billion parameter machine translation model (Fan et al., 2021). In the case of MasakhaNER, since not all languages are included in M2M100, we resorted to NLLB200 (Costa-jussà et al., 2022) 3 billion parameter model instead, which was also used by the EasyProject method (Chen et al., 2023). Both MT models demonstrate comparable performance. For detailed hyperparameter settings, performance comparison of T-Projection using models of different sizes, and a comparison between T-Projection performance using M2M100 and NLLB200, please refer to the Appendix.

5 Intrinsic Evaluation

In this section we present a set of experiments to evaluate T-Projection with respect to current state-of-the-art approaches for annotation projection. We also analyze separately the performance of the *candidate generation* and *candidate selection* steps.

For the OTE task we train T-Projection and XLM-RoBERTa with the English ABSA 2016 training set. We also train the four word alignment systems (excluding SimAlign which is an unsupervised method) using the English training set together with the respective translations as parallel corpora. We augment the parallel data with 50,000 random parallel sentences from ParaCrawl v8 (Esplà et al., 2019). Models are evaluated with respect to the manually label projections generated by García-Ferrero et al. (2022). As the Europarl-based NER dataset (Agerri et al., 2018) provides only test data for each language, T-Projection and XLM-RoBERTa are trained using the full English CoNLL 2003 dataset (Tjong Kim Sang, 2002) together with the labeled English Europarl test data. The word alignment models are in turn trained with the the parallel sentences from the Europarl-based NER data plus 50,000 parallel sentences extracted from Europarl v8 (Koehn, 2005). We evaluate the model with respect to the manual annotations provided by Agerri et al. (2018). With respect to Argument Mining, we use the Neoplasm training set from the AbstrCT dataset to train T-Projection and XLM-RoBERTa, adding its Spanish translation as parallel corpora for the word alignment systems. As this is a medical text corpus, the parallel corpora is complemented with 50,000 parallel sen-

tences from the WMT19 Biomedical Translation Task (Bawden et al., 2019). We evaluate the models with respect to the manually projected labels by Yeginbergenova and Agerri (2023).

5.1 Annotation Projection Quality

Table 1 reports the results of the automatically projected datasets generated by each projection method with respect to the human-projected versions of those datasets. The systems based on word alignments obtain good results across the board, especially those using language models, namely, SimAlign and AWESOME. In particular, AWESOME achieves good results for OTE and NER but very poor in AM. Manual inspection of the projections found out that AWESOME struggles to align articles and prepositions which are included in long sequences.

XLM-RoBERTa-xl shows a strong zero-shot cross-lingual performance. However, the generated datasets are of lower quality than the ones generated by the word-alignment systems. The results of the Span Translation approach are quite disappointing, especially when dealing with the long sequences of the AM task. Translating the labeled spans independently produce translations that, in many cases, cannot be located in the target sentence.

Our T-Projection method achieves the best results for every task and language. In OTE, it outperforms every other method by more than 2 points in F1 score averaged across the three languages. This suggests that T-Projection robustly projects labeled spans into machine-translated data. The NER evaluation is slightly different because the parallel data was translated by human experts. In this setting, T-Projection clearly improves AWESOME’s results by 4.7 points, which constitutes a significant leap in the quality of the generated datasets.

Despite the fact that the word alignment systems have been trained using Europarl domain-specific data, and that most of the training data used for T-Projection is coming from the CoNLL-2003 dataset (news domain) plus very few annotated sentences (699) from Europarl, T-Projection still clearly obtains the best results in NER label projection. This suggests that our system can also be applied in out-of-domain settings.

Finally, T-Projection obtains the overall highest scores for Argument Mining which means that

	OTE			NER			AM	Avg
	ES	FR	RU	ES	DE	IT	ES	
Giza++ (Och and Ney, 2003)	77.0	73.3	72.4	73.3	75.3	68.4	86.6	77.7
FastAlign (Dyer et al., 2013)	75.0	72.9	76.9	70.2	77.0	67.0	85.7	77.4
SimAlign (Jalili Sabet et al., 2020)	86.7	86.3	87.7	85.4	87.4	81.3	84.1	85.3
AWESOME (Dou and Neubig, 2021)	91.5	91.1	93.7	87.3	90.7	83.1	54.8	78.0
XLM-RoBERTa-xl (Conneau et al., 2020)	80.2	76.2	74.5	73.9	68.3	73.9	66.5	71.8
Span Translation	66.5	46.3	58.7	68.8	63.5	69.2	21.6	48.7
T-Projection	95.1	92.3	95.0	93.6	94.0	87.2	96.0	93.9

Table 1: F1 scores for annotation projection in the OTE, NER and Argument Mining tasks.

our approach is particularly good projecting long sequences. Thus, T-Projection outperforms the second best model by 9.4 points in F1 score. In fact, the 96.0 F1 result obtained indicates that T-Projection is almost correctly projecting all the examples in the dataset.

If we look at the average over the three tasks and 5 languages, T-Projection improves by 8.6 points in F1 score the results of the second-best system, SimAlign. These results constitute a huge improvement over all the previous annotation projection approaches. To the best of our knowledge, these are by a wide margin the best annotation projection results published for sequence labeling.

5.2 The Role of the Candidates

	OTE			NER			AM	Avg
	ES	FR	RU	ES	DE	IT	ES	
T-Projection	95.1	92.3	95.0	93.6	94.0	87.2	96.0	93.9
Ngrams + Candidate Selection	89.7	86.1	93.8	83.8	79.3	73.3	73.5	80.7
mT5 + Most Probable Candidate	83.7	87.2	85.3	79.5	82.8	72.3	90.9	84.8
mT5 + Upper bound	98.6	97.0	97.9	98.0	98.5	94.0	99.3	98.0

Table 2: F1 scores for different candidate generation and candidate selection methods.

We perform a set of experiments to measure the relevance and performance of the *candidate generation* and *candidate selection* tasks. First, we replace mT5 with an ngram-based candidate generation approach. We consider as candidate spans every possible ngram with size $1..sentence_length$ (i.e. "Serves", "really", "good", "sushi", "Serves really"...). Table 2 shows that this approach results in lower performance compared with our technique using mT5. Ngrams are much noisier than the candidates generated by mT5, most of them are

very similar to each other, and this makes selecting the right candidate a more challenging task. Thus, this experiment proves that our mT5 candidate generation approach is crucial to obtain good performance.

We also replace the *candidate selection* method with the *most probable candidate*. In other words, we only use the most probable beam generated by mT5 to label the target sentence. When using mT5 by itself, it obtains competitive results, close to those of the word alignment systems in Table 1. Still, the average performance drops by 9.2 points. This further confirms that both the *candidate generation* and *selection* steps are crucial for the T-Projection method.

In a final experiment we define an upperbound for *candidate selection* consisting of assuming that our model will always select the correct projection contained among the generated candidates. The upper bound achieves an average F1 score of 98. This result confirms with a very high probability that the correct candidate is almost always among the 100 candidates generated by mT5.

6 Extrinsic Evaluation

In this section we evaluate T-projection in a real-world low-resource scenario, namely, Named Entity Recognition in African Languages. We compare the results obtained by training on NER dataset automatically generated by T-Projection with respect to those automatically projected using two state-of-the-art label projection systems: AWESOME (The second-best NER system in Table 1) and EasyProject. We use the exactly same settings as Chen et al. (2023). For each target language in MasakhaNER2.0, we first translate the English CoNLL dataset using the NLLB-200 3 billion parameter model. Next, we project the English labels into the target language. It should be noted that EasyProject performs both of these processes in a

Language	No. of Speakers	Lang family	Fine-tune _{en}	AWESOME+EN	EasyProject+EN	T-Projection	T-Projection+EN
Hausa	63M	Afro-Asiatic /Chadic	71.7	72.7	72.2	72.7	72.0
Igbo	27M	NC / Volta-Niger	59.3	63.5	65.6	71.4	71.6
Chichewa	14M	English-Creole	79.5	75.1	75.3	77.2	77.8
chiShona	12M	NC / Bantu	35.2	69.5	55.9	74.9	74.3
Kiswahili	98M	NC / Bantu	87.7	82.4	83.6	84.5	84.1
isiXhosa	9M	NC / Bantu	24.0	61.7	71.1	72.3	71.7
Yoruba	42M	NC / Volta-Niger	36.0	38.1	36.8	42.7	42.1
isiZulu	27M	NC / Bantu	43.9	68.9	73.0	66.7	64.9
AVG			54.7	66.5	66.7	70.3	69.8

Table 3: F1 scores on MasakhaNER2.0 for mDebertaV3 trained with projected annotations from different systems. "+EN" denotes concatenation of the automatically generated target language dataset with the source English dataset.

single step. Subsequently, we train an mDebertaV3 (He et al., 2021) model using the automatically generated datasets for each target language. Finally, this model is evaluated in the gold MasakhaNER2.0 test data. We only evaluate the 8 languages in MasakhaNER2.0 supported by mT5. We focus on named entities referring to Person, Location and Organization.

Table 3 presents the results of the evaluated models on the gold MasakhaNER2.0 test sets. For T-projection, we present the results of training with the automatically generated data for the target language only, and also by adding the original English CoNLL data concatenated with the automatically generated data for each target language. Regarding other systems, we only show the former results, as it was the only metric reported by previous work. In order to train and evaluate the NER models we apply the same hyperparameter settings and code as the authors of EasyProject.

The results show that T-projection achieves superior performance for seven out of the eight languages. Our model demonstrates a more pronounced performance difference in agglutinative languages such as Igbo and Shona. As outlined in Section 5, our model produces superior alignments compared to AWESOME. On the other side, we found that EasyProject, which utilizes markers for simultaneous translation and projection, introduces translation artifacts that hinder the performance of the downstream model. These artifacts are particularly noticeable in agglutinative languages, as EasyProject tends to separate words. For instance, in the case of Shona, consider the English sentence "[Germany]'s representative to the [European Union]'s veterinary committee [Werner Zwingmann]". Our system produces the Shona sentence "Mumiriri [weGermany] kukomiti yemhuka [yeEuropean Union] [Werner Zwingmann]", while EasyProject produces "Mumiriri we [Germany]

ku [European Union] komiti yezvokurapa mhuka [Werner Zwingmann]". When training mDebertaV3 with T-projection generated data, which preserves the agglutinated words, we achieve better results compared to EasyProject that introduce artifacts by separating agglutinated words during translation and projection. Our system is only inferior in the Zulu language; however, on average, we improve the results by 3.6 F1 points. In contrast with previous work, our experiments revealed that concatenating English and translated data did not yield better results, potentially due to the superior quality of the data generated by T-Projection.

To the best of our knowledge, these are the best zero-shot results achieved for MasakhaNER2.0, underscoring the significant benefits of T-projection for NLP tasks in low-resource languages.

7 Concluding Remarks

In this paper we introduce T-Projection, a new annotation projection method that leverages large multilingual pretrained text-to-text language models and state-of-the-art machine translation systems. We conducted experiments on intrinsic and extrinsic tasks in 5 Indo-European languages and 8 African languages. T-projection clearly outperforms previous annotation projection approaches, obtaining a new state-of-the-art result for this task.

A comprehensive analysis shows that both the generation candidate and the candidate selection steps crucially contribute to the final performance of T-Projection. Future work includes adding more tasks and languages, especially those with different segmentation such as Japanese or Chinese. Unlike word alignment systems, T-Projection does not need to segment the words to do the projection which is why we think that our model can also be competitive to project annotations for many language pairs.

Limitations

We evaluate the performance of T-Projection to project labels in sequence labeling tasks from English into a set of 5 Indo-European languages and 8 African languages. It would be interesting to evaluate the performance for other language families, which we leave for future work. Our model requires training a 3B parameter mT5 model. While training a 3B model is computationally expensive and requires a GPU with at least 48GB of VRAM, automatically generating a dataset is a one-off endeavour which results in a dataset usable for many occasions and applications, and much cheaper than manual annotation. Furthermore, we believe that the huge gains obtained by T-Projection justify the computation requirements. In any case, we expect that, thanks to the rapid development of computer hardware, the cost of T-Projection will be reduced in the near future. From a software perspective, recent advancements like 4-bit / 8-bit quantization (Dettmers et al., 2022b,a; Frantar et al., 2022) and Low Rank Adaptation (Hu et al., 2022) have the potential to reduce the hardware requirements of T-Projection.

Acknowledgments

This work has been partially supported by the HiTZ center and the Basque Government (Research group funding IT-1805-22). We also acknowledge the funding from the following projects: (i) DeepKnowledge (PID2021-127777OB-C21) MCIN/AEI/10.13039/501100011033 and ERDF A way of making Europe; (ii) Disargue (TED2021-130810B-C21), MCIN/AEI/10.13039/501100011033 and European Union NextGenerationEU/PRTR; (iii) Antidote (PCI2020-120717-2), MCIN/AEI/10.13039/501100011033 and by European Union NextGenerationEU/PRTR; (iv) DeepR3 (TED2021-130295B-C31) by MCIN/AEI/10.13039/501100011033 and EU NextGeneration programme EU/PRTR. Rodrigo Agerri currently holds the RYC-2017-23647 fellowship (MCIN/AEI/10.13039/501100011033 and by ESF Investing in your future). Iker García-Ferrero is supported by a doctoral grant from the Basque Government (PRE_2022_2_0208).

References

- David Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba Alabi, Shamsudeen 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 Memdjokam Koagne, Allahsera Auguste Tapo, Tebogo Macucwa, Vukosi Marivate, Mboning Tchiaye Elvis, Tajuddeen Gwadabe, Tosin Adewumi, Orevaoghene Ahia, Joyce Nakatumba-Nabende, Neo Lerato Mokono, Ignatius Ezeani, Chiamaka Chukwuneke, Mofetoluwa Oluwaseun Adeyemi, Gilles Quentin Hacheme, Idris Abdulmumin, Odunayo Ogundepo, Oreen Yousuf, Tatiana Moteu, and Dietrich Klakow. 2022. [MasakhaNER 2.0: Africa-centric transfer learning for named entity recognition](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4488–4508, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- 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).
- Mikel Artetxe, Vedanuj Goswami, Shruti Bhosale, Angela Fan, and Luke Zettlemoyer. 2023. [Revisiting machine translation for cross-lingual classification](#). *CoRR*, abs/2305.14240.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Rachel Bawden, Kevin Bretonnel Cohen, Cristian Grozea, Antonio Jimeno Yepes, Madeleine Kittner, Martin Krallinger, Nancy Mah, Aurelie Neveol, Mariana Neves, Felipe Soares, Amy Siu, Karin Verspoor, and Maika Vicente Navarro. 2019. [Findings of the WMT 2019 biomedical translation shared task: Evaluation for MEDLINE abstracts and biomedical terminologies](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 29–53, Florence, Italy. Association for Computational Linguistics.
- 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, Toronto, Canada, July 9-14, 2023*, pages 5775–5796. 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, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Meja Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *CoRR*, abs/2207.04672.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022a. [Llm.int8\(\): 8-bit matrix multiplication for transformers at scale](#). *CoRR*, abs/2208.07339.
- Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 2022b. [8-bit optimizers via block-wise quantization](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- 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.
- Maud Ehrmann, Marco Turchi, and Ralf Steinberger. 2011. [Building a multilingual named entity-annotated corpus using annotation projection](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 118–124, Hissar, Bulgaria. Association for Computational Linguistics.
- Miquel Esplà, Mikel Forcada, Gema Ramírez-Sánchez, and Hieu Hoang. 2019. [ParaCrawl: Web-scale parallel corpora for the languages of the EU](#). In *Proceedings of Machine Translation Summit XVII: Translator, Project and User Tracks*, pages 118–119, Dublin, Ireland. European Association for Machine Translation.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *J. Mach. Learn. Res.*, 22:107:1–107:48.
- 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.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2022. [GPTQ: accurate post-training quantization for generative pre-trained transformers](#). *CoRR*, abs/2210.17323.
- 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.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: decoding-enhanced bert with disentangled attention](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- 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 *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.

- Kuan-Hao Huang, I-Hung Hsu, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2022. [Multilingual generative language models for zero-shot cross-lingual event argument extraction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4633–4646, Dublin, Ireland. Association for Computational Linguistics.
- Alankar Jain, Bhargavi Paranjape, and Zachary C. Lipton. 2019. [Entity projection via machine translation for cross-lingual NER](#). 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)*, pages 1083–1092, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.
- Bing Li, Yujie He, and Wenjin Xu. 2021. [Cross-lingual named entity recognition using parallel corpus: A new approach using xlm-roberta alignment](#). *ArXiv preprint*, abs/2101.11112.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2021. [Crossner: Evaluating cross-domain named entity recognition](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13452–13460. AAAI Press.
- Tobias Mayer, Elena Cabrio, and Serena Villata. 2020. [Transformer-based argument mining for healthcare applications](#). In *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 2108–2115. IOS Press.
- Jian Ni, Georgiana Dinu, and Radu Florian. 2017. [Weakly supervised cross-lingual named entity recognition via effective annotation and representation projection](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1470–1480, Vancouver, Canada. Association for Computational Linguistics.
- Franz Josef Och and Hermann Ney. 2003. [A systematic comparison of various statistical alignment models](#). *Computational Linguistics*, 29(1):19–51.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [SemEval-2014 task 4: Aspect based sentiment analysis](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland. Association for Computational Linguistics.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmmsäsan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4603–4611. PMLR.
- Brian Thompson and Matt Post. 2020. [Automatic machine translation evaluation in many languages via zero-shot paraphrasing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121, Online. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang. 2002. [Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition](#). In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- 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.
- Mengqiu Wang and Christopher D. Manning. 2014. [Cross-lingual projected expectation regularization for weakly supervised learning](#). *Transactions of the Association for Computational Linguistics*, 2:55–66.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv preprint*, abs/1910.03771.
- Jiateng Xie, Zhilin Yang, Graham Neubig, Noah A. Smith, and Jaime Carbonell. 2018. [Neural cross-lingual named entity recognition with minimal resources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*,

pages 369–379, Brussels, Belgium. Association for Computational Linguistics.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, On-line. Association for Computational Linguistics.

David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. [Inducing multilingual text analysis tools via robust projection across aligned corpora](#). In *Proceedings of the First International Conference on Human Language Technology Research*.

Anar Yeginbergenova and Rodrigo Agerri. 2023. [Cross-lingual argument mining in the medical domain](#). *CoRR*, abs/2301.10527.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona T. Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [OPT: open pre-trained transformer language models](#). *CoRR*, abs/2205.01068.

Ran Zhou, Xin Li, Lidong Bing, Erik Cambria, Luo Si, and Chunyan Miao. 2022. [Conner: Consistency training for cross-lingual named entity recognition](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 8438–8449. Association for Computational Linguistics.

A How many candidates do we need?

Generating candidates is expensive. The number of flops and memory usage increases linearly with the number of beams computed. Generating 20 candidates is twice as expensive as generating 10 candidates. We also need to add the extra workload of computing the similarity between more candidates. Figure 5 shows the average F1 score for each task when generating a different number of candidates. For OTE and NER small improvements are obtained when generating more than 25 candidates. However, in Argument Mining using a large number of candidates hurts T-Projection’s performance, which performs optimally with just 10 candidates. While the results reported have been obtained generating 100 candidates, which is computationally very expensive, this analysis shows that we can use a much lower number of candidates and still achieve similar or even better results.




Figure 5: F1 score when generating a different number of candidates.

B Model size vs Performance

	Model	#Params	OTE	NER	AM	Average
MT Size	m2m100	418M	92.3	91.7	95.5	93.1
	m2m100	1.2B	94.0	92.0	95.8	93.9
	m2m100	12B	94.1	91.6	96.0	93.9
mT5 size	mT5-small	60M	36.4	66.3	00.0	34.2
	mT5-base	220M	72.8	86.2	33.6	64.2
	mT5-large	738M	90.9	90.1	65.3	82.1
	mT5-xl	3B	94.1	91.6	96.0	93.9

Table 4: F1 scores of T-Projection when using translation and mT5 models of different size

We analyze the performance of T-Projection when using an mT5 model and a translation system with different number of parameters. Table 4 shows the average F1 performance across all the tasks and languages. First, we experiment with

M2M100 models of different sizes. The results show that the size of the translation model does not have a significant impact on the performance of T-Projection.

However, the size of the mT5 model used does have a big impact on the final performance of the system. Although for OTE and NER switching from a 3B to a 738M parameter mT5 model produces competitive results, this is not the case when applied to AM. The overall trend is that when decreasing the number of parameters results keep decreasing. Summarizing, in order to achieve competitive performance for every task T-Projection requires a mT5 model with 3B parameters, although a 738M parameter model is still competitive for OTE and NER.

C Tuning the Word Alignment Systems

To validate our results and further demonstrate the performance of T-Projection, we conduct a set of experiments that evaluate the performance of word-alignment systems under different settings. We first compare the annotation projection performance when using and not using 50,000 parallel sentences as data augmentation for training the word aligners. Note that in Section 5 all the results we show correspond to using 50,000 extra parallel sentences for training the word-alignment systems. As Table 5 shows, using the augmented dataset achieves the best performance. SimAlign (Dou and Neubig, 2021) and AWESOME (Dou and Neubig, 2021) recommend using their systems with multilingual-bert-based (Devlin et al., 2019) as backbone. However, we also test XLM-RoBERTa-xl (Conneau et al., 2020) 3 billion parameter model with SimAlign and XLM-RoBERTa-large (355M parameters) model with AWESOME (The released AWESOME code at the time of writing this paper doesn’t support XLM-RoBERTa-xl). Using XLM-RoBERTa produce worse results than using mBERT. These experiments show that we are using the word-alignment systems in their best-performing settings.

D MT models vs Laser

We conducted experiments using M2M100-12B (Fan et al., 2021), NLLB200-3B (Costa-jussà et al., 2022) and prism (Thompson and Post, 2020) as model for computing translation probabilities. We also experiment with using LASER 2.0 (Artetxe and Schwenk, 2019) sentence representations in-

System	Data Augmentation	Backbone	OTE			NER			AM	Average
			ES	FR	RU	ES	DE	IT	ES	
Giza++ (Och and Ney, 2003)	0	mBERT	76.2	73.8	78.2	71.4	66.6	65.7	86.4	76.8
FastAlign (Dyer et al., 2013)	0	mBERT	72.3	70.4	74.6	60.3	64.0	57.5	84.0	72.4
SimAlign (Jalili Sabet et al., 2020)	-	mBERT	86.7	86.3	87.7	85.4	87.4	81.3	84.1	85.3
AWESOME (Dou and Neubig, 2021)	0	mBERT	88.9	89.8	91.2	86.1	89.4	83.0	57.1	77.8
Giza++ (Och and Ney, 2003)	50000	mBERT	77.0	73.3	72.4	73.3	75.3	68.4	86.6	77.7
FastAlign (Dyer et al., 2013)	50000	mBERT	75.0	72.9	76.9	70.2	77.0	67.0	85.7	77.4
AWESOME (Dou and Neubig, 2021)	50000	mBERT	91.5	91.1	93.7	87.3	90.7	83.1	54.8	78.0
SimAlign (Jalili Sabet et al., 2020)	-	XLm-RoBERTa-xl	86.2	86.1	89.5	85.8	88.4	81.2	76.9	83.1
AWESOME (Dou and Neubig, 2021)	50000	XLm-RoBERTa-large	86.1	86.1	87.4	87.2	87.5	83.1	54.8	75.8
T-Projection	-	-	95.1	92.3	95.0	93.6	94.0	87.2	96.0	93.9

Table 5: Results of the different word alignment systems when we train with and without a data augmentation corpus and different backbone models

Candidate Scorer	OTE			NER			AM	Average
	ES	FR	RU	ES	DE	IT	ES	
Prism-745M	91.4	86.8	94.3	93.8	93.4	85.4	96.3	92.7
M2M100-12B	95.1	92.3	95.0	93.6	94.0	87.2	96.0	93.9
NLLB200-3B	96.6	90.5	95.6	91.0	94.3	87.7	93.9	93.0
LASER 2.0	89.0	80.6	91.3	91.2	91.6	86.5	70.4	82.4

Table 6: Results of T-Projection when selecting candidates using translation probability scores with different MT systems vs using the cosine similarity of the multilingual vector representations of the candidates computed using LASER 2.0

stead of the translation probabilities of NMTscore. We encode the source span as well as all the projection candidates using LASER encoder. We then rank them using cosine similarity. Table 6 shows the results. LASER2.0 is competitive when dealing with the short labeled sequences in the OTE and NER task. But the performance decreases when dealing with large sequences in the AM task. M2M100, NLLB200, and Prism exhibit comparable performance, with some of them achieving the best results in specific languages, but overall, their average performance is very similar.

E Training details

We train the HuggingFace’s (Wolf et al., 2019) implementation of mT5⁴ (3 billion parameter model) in the candidate generation step using the following hyper-parameters: Batch size of 8, 0.0001 learning rate, 256 tokens sequence length, cosine scheduler with 500 warm up steps and no weight decay. We use AdaFactor (Shazeer and Stern, 2018) optimizer. We train the model for 10 epochs in the OTE task, and 4 epochs for the NER and AM tasks. In the candidate selection step, we also use HuggingFace’s

⁴<https://huggingface.co/google/mt5-xl>

implementation of M2M100, and we use m2m100-12B-last-ckpt⁵ checkpoint of M2M100 released by the authors. We use the direct-translation function of the NMTscore library to compute the translation probabilities. For MasakhaNER2.0 we use the training script and evaluation script developed by the authors⁶ and the same hyper-parameter setup than Chen et al. (2023).

F Dataset details

We list the size (number of sentences) of the dataset we use in Table 7. Note that all the datasets we use are parallel in all the languages, and the number of sentences is the same for all the languages.

⁵<https://huggingface.co/facebook/m2m100-12B-last-ckpt>

⁶<https://github.com/masakhane-io/masakhane-ner/blob/main/MasakhaNER2.0/scripts/mdeberta.sh>

Task	Split	Sentence No
ABSA		
ABSA (Pontiki et al., 2014)	Train	2000
ABSA (Pontiki et al., 2014)	Test	676
NER		
Europarl (Agerri et al., 2018)	Test	799
CoNLL03 (Tjong Kim Sang, 2002)	Train	14987
CoNLL03 (Tjong Kim Sang, 2002)	Dev	3466
CoNLL03 (Tjong Kim Sang, 2002)	Test	3684
MasakhaNER2.0 (Adelani et al., 2022)	Test (hau)	1632
MasakhaNER2.0 (Adelani et al., 2022)	Test (ibo)	2180
MasakhaNER2.0 (Adelani et al., 2022)	Test (sna)	1772
MasakhaNER2.0 (Adelani et al., 2022)	Test (swa)	1882
MasakhaNER2.0 (Adelani et al., 2022)	Test (xho)	1632
MasakhaNER2.0 (Adelani et al., 2022)	Test (yor)	1963
MasakhaNER2.0 (Adelani et al., 2022)	Test (nya)	1784
MasakhaNER2.0 (Adelani et al., 2022)	Test (zul)	1669
AM		
AbsRCT Neoplasm (Mayer et al., 2020)	Train	4404
AbsRCT Neoplasm (Mayer et al., 2020)	Dev	679
AbsRCT Neoplasm (Mayer et al., 2020)	Test	1251

Table 7: Size (Number of sentences) of the dataset we use to train and evaluate our systems.

For OTE, we use the SemEval-2016 Task 5 Aspect Based Sentiment Analysis (ABSA) dataset (Pontiki et al., 2014). We train T-Projection with the concatenation of the English train and test splits. We evaluate all the systems by projecting the training split. For NER we use the English, Spanish, German, Italian europarl parallel dataset from (Agerri et al., 2018). We train T-Projection with the concatenation of the English europarl dataset with the train, dev and test splits of the English CoNLL 2003 dataset (Tjong Kim Sang, 2002). We evaluate the systems by projecting the English NER europarl test splits. For Argument Mining, we use the AbsRCT Neoplasm English dataset (Mayer et al., 2020) and the Spanish AbsRCT corpus generated by machine translating the English AbsRCT corpus with DeepL and manually projecting the labels. We train T-Projection of the concatenation of the English Neoplasm train, dev and test split. We evaluate the systems by projecting the English Neoplasm train split.

G Hardware used

We perform all our experiments using a single NVIDIA A100 GPU with 80GB memory. The machine used has two AMD EPYC 7513 32-Core Processors and 512GB of RAM.

Medical mT5: An Open-Source Multilingual Text-to-Text LLM for The Medical Domain

Iker García-Ferrero¹, Rodrigo Agerri¹, Aitziber Atutxa¹, Elena Cabrio²,
Iker de la Iglesia¹, Alberto Lavelli³, Bernardo Magnini³, Benjamin Molinet²,
Johana Ramirez-Romero⁴, German Rigau¹, Jose Maria Villa-Gonzalez⁴,
Serena Villata², Andrea Zaninello³

¹HiTZ Center - Ixa, University of the Basque Country UPV/EHU

²Université Côte d'Azur, CNRS, Inria, I3S, France

³Fondazione Bruno Kessler, Via Sommarive 18, Povo, Trento, Italy

⁴Cruces University Hospital, (Barakaldo, Biscay, Spain)

{iker.garciaf, rodrigo.agerri}@ehu.eus, {FirstName.LastName}@univ-cotedazur.fr, {lastName}@fbk.eu

Abstract

Research on language technology for the development of medical applications is currently a hot topic in Natural Language Understanding and Generation. Thus, a number of large language models (LLMs) have recently been adapted to the medical domain, so that they can be used as a tool for mediating in human-AI interaction. While these LLMs display competitive performance on automated medical texts benchmarks, they have been pre-trained and evaluated with a focus on a single language (English mostly). This is particularly true of text-to-text models, which typically require large amounts of domain-specific pre-training data, often not easily accessible for many languages. In this paper, we address these shortcomings by compiling, to the best of our knowledge, the largest multilingual corpus for the medical domain in four languages, namely English, French, Italian and Spanish. This new corpus has been used to train Medical mT5, the first open-source text-to-text multilingual model for the medical domain. Additionally, we present two new evaluation benchmarks for all four languages with the aim of facilitating multilingual research in this domain. A comprehensive evaluation shows that Medical mT5 outperforms both encoders and similarly sized text-to-text models for the Spanish, French, and Italian benchmarks, while being competitive with current state-of-the-art LLMs in English.

Keywords: Natural Language Processing in Medicine, Multilingualism, Large Language Models, Deep Learning

1. Introduction

As it is the case for many application domains, there is an increasing interest in applying Artificial Intelligence (AI) and Natural Language Processing (NLP) techniques to assist medical experts in their everyday activities. With this aim in mind, in the last few years a number of large language models (LLMs) have been trained or adapted to the medical domain. These include encoder models such as SciBERT (Beltagy et al., 2019), BioBERT (Lee et al., 2020) or PubmedBERT (Gu et al., 2022). While these models have obtained state-of-the-art results in discriminative tasks, they are typically smaller in scale and scope with respect to medical text-to-text models such as SciFive (Phan et al., 2021), BioGPT (Luo et al., 2022) Med-PaLM (Singhal et al., 2022), PMC-LLaMA (Wu et al., 2023) or ClinicalGPT (Wang et al., 2023).

However, the development of all the aforementioned text-to-text LLMs has been focused on a single language, usually English. As a consequence, there is a lack of high-quality multilingual evaluation benchmarks for the medical domain. Thus, although there have been efforts to generate evaluation data in languages other than English (Wang et al., 2023; Carrino et al., 2022), they have con-

sisted largely in monolingual approaches.

In order to address these issues, we have compiled, to the best of our knowledge, the largest multilingual corpus for training LLMs adapted to the medical domain. Our corpus includes 3B words in four languages, namely, English, Spanish, French, and Italian. While relatively small when compared to English existing datasets (Wu et al., 2023), it allowed us to build Medical mT5, the first open-source text-to-text multilingual model for the medical domain.

Medical mT5 is an encoder-decoder model developed by continuing the training of publicly available mT5 (Xue et al., 2021) checkpoints on medical domain data for English, Spanish, French, and Italian. Additionally, we have also created two new multilingual sequence labeling (argument component detection) and generative question answering datasets for the evaluation of multilingual LLMs in the medical domain.

A comprehensive experimental evaluation shows that Medical mT5 outperforms similarly-sized text-to-text models for the Spanish, French, and Italian benchmarks while being competitive in English with respect to current state-of-the-art text-to-text (Xue et al., 2021; Chung et al., 2022) and encoder-only

models (Lee et al., 2020; He et al., 2023). The results show that continuing pre-training of a multilingual text-to-text model such as mT5 allows to successfully adapt it to the medical domain, even when the amount of domain-specific data is relatively modest (ranging between 1B words for English and Spanish to 150M in Italian). Summarizing, the contributions of our work are the following: (i) the collection of the largest publicly available in-domain medical multilingual corpus for Spanish, French, and Italian languages. Together with the already existing English data, we release a corpus of 3 billion tokens.¹ (ii) two new datasets for Spanish, French, and Italian on Argument Mining² and generative Question Answering tasks, generated taking their original English versions as a starting point.³ (iii) the public release of two Medical mT5 versions: a 770M⁴ and 3B⁵ parameter text-to-text open-source models which obtain state-of-the-art results in multilingual sequence labelling for the medical domain, most notably in multi-task and zero-shot crosslingual settings.

Other benefits of our Medical mT5 models include the comparatively low hardware requirements needed for both fine-tuning on downstream tasks (the large 770M version easily fits in a 24GB V100 GPU) and for inference (a 12GB GPU should be enough). As an example, a LLaMA 7B model (Wu et al., 2023) requires at least a 80GB A100 GPU using LoRA (Hu et al., 2021) or a more demanding 4 80GB A100 GPUs without it. Code, data, models, and benchmarks are publicly available to facilitate reproducibility of results and encourage future multilingual research on the medical domain.

2. Related Work

As it has been the case in most application domains, Large Language Models (LLMs) have facilitated huge improvements in the state-of-the-art for medical NLP tasks (Singhal et al., 2022; Wu et al., 2023; Mayer et al., 2021). The most popular approaches are those that use models pre-trained on medical corpora such as SciBERT (Beltagy et al., 2019), BioBERT (Lee et al., 2020), PubmedBERT (Gu et al., 2022), BSC-BIO (Carrino et al., 2022) or BioLinkBERT (Yasunaga et al., 2022).

While the previous encoder-only models focused on discriminative tasks, the emergence of generative models such as LLaMa (Touvron et al., 2023),

PaLM (Singhal et al., 2022) or GPT-3 (Brown et al., 2020) has resulted in a huge interest in adapting such LLMs to the medical domain. These models, to name but a few, include SciFive (Phan et al., 2021), and English T5 encoder-decoder model adapted to the scientific domain, and decoder models such as BioGPT (Luo et al., 2022), Med-PaLM (Singhal et al., 2022), PMC-LLaMA (Wu et al., 2023) and ClinicalGPT (Wang et al., 2023).

Additionally, a range of Abstractive Question Answering tasks have been proposed as evaluation benchmarks on which the larger models (Wu et al., 2023; Singhal et al., 2022; Wang et al., 2023) obtain best results. While interesting, both these LLMs and benchmarks have been developed with a focus on a single language, usually English. Furthermore, these LLMs require hardware which is simply not affordable for the large majority of end-users and researchers. In order to address these issues, we propose Medical mT5, a multilingual text-to-text model adapted to the medical domain which, despite its relatively modest size and cheap running requirements, obtains competitive results, most notably in multi-task and zero-shot cross-lingual settings.

3. Compiling a Multilingual Corpus for the Medical Domain

Obtaining good quality medical corpora is usually difficult due to the sensitive nature of the data. This is even more challenging for non-English languages, as the availability of data for other languages is in general more restricted. Despite these issues, we have successfully gathered and curated a diverse collection of public relevant corpora of medical texts in English, French, Italian and Spanish to generate the Medical mT5 model.

3.1. English

As listed in Table 1, we collected around 1B words from three sources related to the medical domain: (i) **ClinicalTrials** is a set of documents of clinical studies from all over the world (National Library of Medicine, 2022a); (ii) **EMEA** is an English-Spanish parallel corpus with documents provided by the European Medicines Agency (Tiedemann, 2012) and, (iii) PubMed (National Library of Medicine, 2022b), which contains data from various sources such as MEDLINE, life science journals and online books, provides the bulk of the English data.

3.2. Spanish

Apart from **EMEA** and **PubMed**, which we also used for Spanish, the biggest portion of the data

¹<https://hf.co/datasets/HiTZ/Multilingual-Medical-Corpus>

²<https://hf.co/datasets/HiTZ/multilingual-abstrct>

³<https://hf.co/datasets/HiTZ/Multilingual-BioASQ-6B>

⁴<https://hf.co/HiTZ/Medical-mT5-large>

⁵<https://hf.co/HiTZ/Medical-mT5-xl>

Source	Words
ClinicalTrials	127.4M
EMA	12M
PubMed	968.4M

Table 1: English data sources and word counts.

came from the **Medical Crawler**, a biomedical corpus compiled by Carrino et al. (2022). Additionally, we also included **SPACC** (Ander Itxaurrondo, 2018), **UFAL** (Institute of Formal and Applied Linguistics, 2017) and **WikiMed**, a corpus built ad-hoc from Wikipedia entries. Table 2 provides the details of the collected data, which amounts to ≈ 1 B words.

Source	Words
EMA	13.6M
PubMed	8.4M
Medical Crawler	918M
SPACC	350K
UFAL	10.5M
WikiMed	5.2M

Table 2: Spanish data sources and word counts.

3.3. French

A total of 7,192,779 sentences and 670,972,717 words were compiled using the data sources listed in Table 3.

Source	Words
PubMed	1.4M
Science Direct	15.2M
Wikipedia - Médecine	5M
EDP	48K
Google Patents	654M

Table 3: French data sources and word counts.

PubMed data was extracted using the **Bio.Entrez** package⁶. **Science Direct** offers a collection of scientific and medical publications which can be extracted via their the official API⁷. We filtered relevant articles with the keyword “Médecine”, and the obtained XML documents were parsed to extract the `<dc:description>` tag. As for Spanish, we took advantage of **Wikipedia** as a source of medical knowledge to obtain HTML formatted data from the category “Category:Médecine”. The **EDP French/English**

⁶<https://biopython.org/docs/1.75/api/Bio.Entrez.html>

⁷<https://dev.elsevier.com/>

Parallel Medical Corpus (Jimeno-Yepes et al., 2017) provides bilingual content from journals that address domains such as dentistry and life sciences. From this source, we downloaded the dataset labeled “EDP French corpus, text format”. Finally, **Google Patents** is a comprehensive repository of patent data from around the world. Google Patents data were retrieved by filtering using the IPC code and abstract language.

A final French language verification step was undertaken by applying the `langdetect` package (version 1.0.9).

3.4. Italian

The crawling and pre-processing of the Italian split of the corpus followed the methodology described by Carrino et al. (2022). First, we compiled a list of 504 medical terms, which we use as seeds to scrape the Italian split of the **MC4 Common Crawl Corpus** (Common Crawl, 2022) by only selecting the pages which contained at least one of the keywords in their URL domain. To create the list, we extracted 600 keyword terms related to medicine from the *Dizionario analogico della Lingua Italiana* (Zanichelli). We excluded some sectors and discarded terms that may lead to ambiguous queries (e.g., actions, which contained mainly verbs, proverbs, general terms like “assistente”, etc.). We normalized rare variants (“batteriologia” to “batterologia”) and stemmed all terms without lemmatizing, as most terms are already lemmatized in the dictionary; we performed univertion of multiword units (e.g., “esamedelleurine”, “follow-up”), and removed the duplicates. This resulted in a corpus of 67 million tokens, which we joined with other sources of text such as **Medical dissertations**, **Drug use instructions**, **PubMed abstracts**, etc. as detailed in Table 4, resulting in a ≈ 145 M word corpus.

Source	Words
Medical Commoncrawl - IT	67M
Drug instructions	30.5M
Wikipedia - Medicina	13.3M
E3C Corpus - IT	11.6M
Medicine descriptions	6.3M
Medical theses	5.8M
Medical websites	4M
PubMed	2.3M
Supplement description	1.3M
Medical notes	975K
Pathologies	157K
Medical test simulations	26K
Clinical cases	20K

Table 4: Italian data sources and word counts.

4. Medical mT5

Multilingual T5 (mT5) (Xue et al., 2021) is an extension of the original T5 (Raffel et al., 2020) framework, which is optimized for multilingual tasks. The T5 model is grounded in the transformer encoder-decoder architecture (Vaswani et al., 2017). With its decoder block, T5 is capable of generating sequences of tokens in an auto-regressive fashion. T5 was designed to convert every NLP problem into a text-to-text task, and mT5 extends this strategy to a multitude of languages, leveraging a shared vocabulary for diverse scripts. mT5 was trained using mC4, a 1 Trillion token Common Crawl-based dataset covering 101 languages. The pre-training is based on a masked language modeling “span-corruption” objective, where consecutive spans of input tokens are replaced with a mask and the model is trained to reconstruct the masked-out tokens.

4.1. Pre-training Medical mT5

Medical mT5 is built upon the same architecture as mT5 (Xue et al., 2021). We release two different models: Medical-mT5-large (738M parameters) and Medical-mT5-xl (3 billion parameters). Both models were initialized using the pre-trained weights of their corresponding mT5 checkpoints and continued their pre-training using the 3B word medical domain dataset described in Section 3 (with x2 oversampling for the Italian split). To prevent over-fitting, we run the training for only one epoch, as preliminary experiments showed that performance degraded with more epochs. We adhered to the self-supervised parameter settings recommended by Xue et al. (2021) and detailed in Table 5. It should be noted that Medical-mT5-large was trained with a sequence length of 1024 tokens whereas Medical-mT5-xl was limited to a sequence length of 480 tokens due to GPU memory limitations. Medical mT5 was trained using the Flax implementation of mT5 in the Hugging Face transformers library (Wolf et al., 2020). All experiments were conducted on our private servers, employing 4xA100 80GB GPUs. We made calculations for a carbon footprint estimation based on a 400W consumption per GPU and a carbon intensity of 0.171 kg/kWh⁸.

5. Generating New Multilingual Benchmarks

The lack of multilingual evaluation benchmarks for the medical domain motivated us to generate new evaluation data for our languages of interest, as only the relatively small E3C (Magnini et al.,

⁸Sourced from <https://app.electricitymaps.com/map>

	Medical-mT5-large	Medical-mT5-xl
Param. no.	738M	3B
Sequence Length	1024	480
Token/step	65536	30720
Epochs	1	1
Total Tokens	4.5B	4.5B
Optimizer	Adafactor	Adafactor
LR	0.001	0.001
Scheduler	Constant	Constant
Hardware	4xA100	4xA100
Time (h)	10.5	20.5
CO ₂ eq (kg)	2.9	5.6

Table 5: Pre-Training settings for Medical mT5.

2021) was already available for all 4 languages. We focused on two different types of tasks: (i) a sequence labelling task, **Argument Mining**, consisting in detecting and classifying the argument component spans and their relations, (ii) **Abstractive Question Answering**, where the model is expected to generate an answer in response to an input question. In both cases we took existing English labelled data as a starting point.

5.1. Argument Mining

The AbstrCT dataset is composed by English medical and scientific texts collected from the MEDLINE database and manually annotated with two types of argument components: Claims and Premises (Mayer et al., 2021).

A ‘claim’ is a concluding statement made by the author about the outcome of the study. In the medical domain it may be an assertion of a diagnosis or a treatment. A ‘premise’ corresponds to an observation or measurement in the study (ground truth), which supports or attacks another argument component, usually a claim. It is important that they are observed facts, therefore, credible without further evidence.

We generated French and Italian parallel versions of the dataset using the same method as for Spanish, based on machine translation and semi-manual annotation projection (Yeginbergenova and Agerri, 2023). The AbstrCT dataset is divided in three splits, neoplasm, glaucoma and mixed. Following previous work, we fine-tune the models with the first one and then evaluate the in-domain performance on the neoplasm test split and the cross-domain performance on the glaucoma and mixed splits. Previous works using the AbstrCT datasets have employed different definitions of the F_1 score metric, such as token-level F_1 (Mayer et al., 2021; Yeginbergenova and Agerri, 2023). However, in this paper we report results using the standard sequence level F_1 score (Tjong Kim Sang and De Meulder, 2003), a much more strict metric, which explains the lower results for all

Representation	Task	Dataset	Languages	Entity Type
Sequence Labelling	Named Entity Recognition	NCBI-Disease (Dogan et al., 2014)	EN	Disease
		BC5CDR Disease (Li et al., 2016)	EN	Disease
		BC5CDR Chemical (Li et al., 2016)	EN	Chemical
		DIANN (Fabregat et al., 2018)	EN, ES	Disability
		E3C (Magnini et al., 2021)	EN, ES, FR, IT	Clinical Entity
	Argument Mining	PharmaCoNER (Gonzalez-Agirre et al., 2019)	ES	Pharmacological
		AbstrCT (Mayer et al., 2021)	EN, ES, FR, IT	Claims and Premises
Generative Question Answering	Question Answering	BioASQ 6B (Tsatsaronis et al., 2015)	EN, ES, FR, IT	Biomedical QA

Table 6: List of evaluation tasks used to measure the performance of Medical mT5.




Figure 1: Text-to-Text representation of the Sequence Labeling task. Given an input sentence, the model is expected to generate the same sentence annotated with html-style tags.

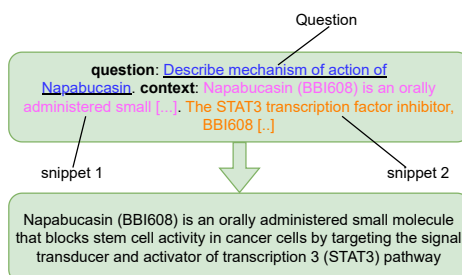


Figure 2: Text-to-Text representation of the BioASQ task. Given a question and a set of relevant snippets, the model generates an answer.

the models.

5.2. Question Answering

We use the BioASQ-6B English Question Answering dataset (Tsatsaronis et al., 2015) to generate parallel French, Italian and Spanish versions. Given a biomedical question and a set of snippets of text with relevant information about the question, the model must generate the *ideal* answer. A set of ideal gold answers are provided to assess the performance of the models. We machine translated the questions and ideal answers into French, Italian and Spanish using the NLLB200 3B parameter model (Costa-jussà et al., 2022).

6. Experimental Setup

Medical mT5 is a text-to-text model. This means that, given a text input, it learns to generate a text as output. Therefore, every evaluation task must be converted into a text-to-text format (Xue et al., 2021). In our experiments the output text is always generated using beam search with 4 beams.

The list of tasks used for evaluation is listed in Table 6. The **Sequence Labelling tasks** include medical NER, detecting and classifying named entities according to some pre-defined categories, and Argument Mining, described in Section 5. Performance for every sequence labelling task is evaluated using standard sequence level F_1 score (Tjong Kim Sang and De Meulder, 2003).

In order to address sequence labelling tasks, text-to-text models such as Medical mT5 are prompted with the sentence to label. As illustrated in Figure 1, the expected output is the same sentence annotated with HTML-style tags. The HTML tags for each task are added as special tokens to the model vocabulary. Furthermore, we use constrained decoding to ensure that the output contains the same words as the input and a valid HTML annotation. We use the *Sequence Labeling with LLMs*⁹ library.

With respect to the BioASQ **Abstractive Question Answering task**, the input prompt contains the question and a context. As shown in Figure 2, the context is generated by concatenating all the provided possible snippets. The expected output should be the generated answer to the question, which is then compared to the gold ideal answer.

6.1. Baselines

As we have developed Medical mT5 by continuing the training of mT5 checkpoints, our primary point of comparison should be mT5 (Xue et al., 2021). Thus, our first objective is to assess whether training the model on our multilingual medical-domain

⁹<https://github.com/ikergarcia1996/Sequence-Labeling-LLMs>

Lang	Dataset	mT5 _{large}	mT5 _{XL}	SciFive	FlanT5 _{large}	FlanT5 _{XL}	mDeBERTa _{v3 base}	BioBERT	MedMT5 _{large}	MedMT5 _{XL}
EN	NCBI-Disease	85.1	87.7	89.4	88.6	89.3	85.7	87.4	89.1	87.2
EN	BC5CDR Disease	78.5	81.4	85.4	85.0	85.8	82.5	84.3	84.4	82.4
EN	BC5CDR Chemical	89.1	90.8	93.3	92.0	92.9	91.1	92.9	92.8	91.3
EN	DIANN	70.1	77.8	71.9	74.4	74.2	80.3	79.0	74.8	77.6
ES	DIANN	72.4	74.9	70.5	70.7	70.9	78.3	70.2	74.9	74.8
EN	E3C	54.3	60.1	62.8	64.2	63.1	58.2	58.6	59.4	57.9
ES	E3C	61.6	71.7	62.7	64.4	67.1	65.9	57.4	72.2	69.5
FR	E3C	55.6	64.9	61.7	65.2	64.3	62.0	53.3	65.2	65.8
IT	E3C	61.8	63.8	59.6	61.9	65.1	63.9	52.1	67.5	65.9
ES	PharmaCoNER	86.3	90.6	87.5	88.5	89.1	89.4	88.6	90.8	90.1
EN	Neoplasm	70.4	71.1	74.4	74.3	73.4	64.5	67.5	73.9	73.2
EN	Glaucoma	70.7	75.1	77.1	78.4	78.0	71.2	74.8	76.2	76.4
EN	Mixed	68.5	73.0	73.4	73.2	74.5	63.4	69.6	72.2	72.0
ES	Neoplasm	69.0	56.1	71.4	72.5	73.9	63.0	57.1	72.1	71.8
ES	Glaucoma	69.3	70.7	73.9	73.8	75.2	68.6	64.5	77.1	75.5
ES	Mixed	68.4	66.2	69.2	69.3	71.6	61.3	58.9	72.4	71.4
FR	Neoplasm	70.5	66.6	74.0	72.4	73.7	63.9	59.0	72.9	71.2
FR	Glaucoma	71.1	69.2	77.8	74.8	77.2	60.3	65.6	79.5	75.8
FR	Mixed	68.3	65.4	72.0	70.9	74.3	64.1	61.3	73.3	69.7
IT	Neoplasm	68.1	69.9	70.1	70.9	72.0	64.4	54.8	71.2	73.1
IT	Glaucoma	69.2	71.5	73.7	74.0	75.9	74.7	65.8	75.7	78.7
IT	Mixed	66.3	67.7	67.4	69.9	70.0	61.3	57.4	70.6	71.9
AVERAGE		70.2	72.1	73.6	74.1	75.1	69.9	67.3	75.4	74.7
AVERAGE ES, FR, IT		68.4	69.2	70.8	71.4	72.9	67.2	61.9	74.0	73.2

Table 7: Single-task supervised F1 scores for Sequence Labelling.

corpus enhances its performance for tasks specific to this domain. Furthermore, we also benchmark our model against SciFive (Pubmed+PMC) a T5-based 738M parameter model (Phan et al., 2021) trained exclusively on a corpus of 78B words containing scientific and medical English data. Additionally, we compare the performance of Medical mT5 with Flan-T5 (Chung et al., 2022), which also adopts the T5 architecture but has been finetuned on a huge instruction-following dataset for almost 2K tasks. Flan-T5 achieves state-of-the-art performance in numerous benchmarks, including some from the medical domain (Singhal et al., 2022). We tested all three types of text-to-text models under identical settings and hyperparameters.

We also measure Medical mT5 with the performance of encoder-only models in sequence labelling tasks. We report results with mDeBERTaV3 (He et al., 2023) which is widely used for sequence labeling and excels in multilingual tasks (Adelani et al., 2022; Aggeri and Agirre, 2023). Although we also tested XLM-RoBERTa (Conneau et al., 2020) and GLOT500 (Imani et al., 2023), their results were worse than those obtained by mDeBERTaV3. Finally, we also compare with BioBERT v1.1 (Lee et al., 2020), which has been pretrained on a large English-only biomedical dataset. We do not evaluate the performance of encoder-only models in the question answering task, as their architecture is not designed for text generation.

The specific hyperparameter settings used to fine-tune the models will be available in the Appendix upon publication.

7. Experimental Results

In this section, we report on the performance of Medical mT5 and of the baselines in the **sequence labelling tasks** across different settings. Due to space constraints, we only report the best performing results.

Single Task Monolingual Supervised Results:

The results when fine-tuning and evaluating the models for each dataset and language are shown in Table 7. The first observation is that Medical-mT5-large significantly outperforms both mT5-large and mT5-XL, demonstrating the benefits of further training these models with our multilingual medical domain corpus.

When comparing Medical mT5 with FlanT5 and SciFive, the latter models are systematically superior on English. This was anticipated since both have been pre-trained with a much larger amount of English-only data specific to the medical domain. With respect to encoder-only models, they achieve in general worse results than text-to-text models across all tasks and languages (except for the DI-ANN dataset). It is also noteworthy that FlanT5-XL exhibits robust performance across all datasets and languages, even though it was fine-tuned with English-only data not specific to the medical domain. Nonetheless, Medical-mT5-large obtains in general better results for French, Spanish and Italian while being much smaller in size (738M parameters vs 3B parameters), showing the impact of training Medical mT5 with domain-specific data for those languages.

Lang	Dataset	Single Task			MultiTask		
		FlanT5 _{XL}	MedMT5 _{large}	MedMT5 _{XL}	FlanT5 _{XL}	MedMT5 _{large}	MedMT5 _{XL}
EN	NCBI-Disease	89.3	89.1	87.2	87.6	87.6	86.9
EN	BC5CDR Disease	85.8	84.4	82.4	85.1	83.4	83.0
EN	BC5CDR Chemical	92.9	92.8	91.3	92.7	92.5	91.6
EN	DIANN	74.2	74.8	77.6	80.0	75.4	75.3
ES	DIANN	70.9	74.9	74.8	77.1	72.6	73.6
EN	E3C	63.1	59.4	57.9	62.1	60.9	62.0
ES	E3C	67.1	72.2	69.5	66.5	74.9	73.3
FR	E3C	64.3	65.2	65.8	62.9	65.4	65.1
IT	E3C	65.1	67.5	65.9	60.7	66.9	65.1
ES	PharmaCoNER	89.1	90.8	90.1	89.9	90.3	89.5
EN	Neoplasm	73.4	73.9	73.2	73.1	72.3	72.9
EN	Glaucoma	78.0	76.2	76.4	76.4	76.8	77.5
EN	Mixed	74.5	72.2	72.0	71.5	70.9	73.0
ES	Neoplasm	73.9	72.1	71.8	73.5	73.5	73.7
ES	Glaucoma	75.2	77.1	75.5	77.1	77.7	79.3
ES	Mixed	71.6	72.4	71.4	70.0	71.8	72.8
FR	Neoplasm	73.7	72.9	71.2	74.0	72.9	73.6
FR	Glaucoma	77.2	79.5	75.8	76.6	77.0	79.4
FR	Mixed	74.3	73.3	69.7	71.8	71.2	73.0
IT	Neoplasm	72.0	71.2	73.1	71.9	74.6	74.0
IT	Glaucoma	75.9	75.7	78.7	77.6	78.5	78.9
IT	Mixed	70.0	70.6	71.9	69.9	72.5	73.3
AVERAGE		75.1	75.4	74.7	75.2	76.2	76.7
AVERAGE ES, FR, IT		72.9	74.0	73.2	73.1	74.8	75.3

Table 8: Multi-task supervised F1 scores for Sequence Labelling.

Multi-Task Supervised Results: Text-to-text models have demonstrated improved performance when trained in multi-task settings (Chung et al., 2022). Following this, we also experimented with fine-tuning them across all the sequence labeling tasks simultaneously. To inform the model about which labels should classify for each input example, we add the list of predefined labels from the corresponding dataset to the beginning of the input sentence. For instance, the input depicted in Figure 1 is adjusted to “<Disease> Patient with dilated cardiomyopathy”. A comparison of the Single Task and Multi-Task settings is presented in Table 8. It can be seen that in this setting Medical mT5 achieves the best overall results for Spanish, French and Italian. On average, Medical-mT5-xl also obtains the best performance, slightly improving over the results of FlanT5-XL and Medical-mT5-large.

Zero-shot Cross-Lingual Transfer Results: Manually annotated medical domain datasets for languages other than English are scarce. Therefore, developing models that can successfully generate predictions for languages different to those used for fine-tuning is crucial. We evaluate this ability to perform zero-shot cross-lingual transfer by fine-tuning Medical mT5 and the baselines on the English AbsRCT Neoplasm dataset, and then evaluating them on the Neoplasm, Glaucoma, and Mixed datasets for Spanish, French, and Italian. The results are presented in Table 9. Results show that Medical mT5 outperforms any other model.

Moreover, Medical-mT5-xl achieves significantly better results than Medical-mT5-large.

To summarize, Medical mT5 stands out for its superior performance in the evaluation for Spanish, French, and Italian languages, especially for the multitask and the zero-shot transfer settings. These capabilities can help mitigate the scarcity of manually annotated medical data for other target languages. In contrast, SciFive and FlanT5, having been trained on extensive English-only datasets, emerge as the top choices when the primary focus is on English-only tasks.

Finally, despite Medical-mT5-xl being larger than Medical-mT5-large (3B vs 738M), its performance is worse in the single-task evaluation setting. This behaviour is not observed in the multi-task and zero-shot experiments, leading us to hypothesize that the larger Medical-mT5-xl model is more prone to overfit in the single-task supervised setting.

7.1. Abstractive Question Answering

In this section, we explore the text generation capabilities of Medical mT5 and other baseline text-to-text models on the BioASQ question answering dataset described in Section 5.2. Previous work typically evaluate the performance on this task using the ROUGE score (Tsatsaronis et al., 2015) to compare the gold standard answer with the answer generated by the model. However, we find this metric inadequate for medical domain tasks as it does not address crucial aspects of the generation

Lang	Dataset	mT5 _{XL}	SciFive	FlanT5 _{XL}	mDeBERTa _{v3 base}	MedMT5 _{large}	MedMT5 _{XL}
ES	Neoplasm	71.4	69.8	67.9	65.1	72.4	71.7
ES	Glaucoma	74.1	71.5	70.6	68.3	72.4	73.2
ES	Mixed	69.4	67.0	66.7	60.9	68.1	68.8
FR	Neoplasm	71.6	68.6	69.9	60.5	72.4	72.8
FR	Glaucoma	75.8	74.5	71.0	68.7	72.3	76.7
FR	Mixed	73.0	68.5	68.2	59.3	70.4	72.4
IT	Neoplasm	70.6	63.1	67.3	62.4	72.9	73.2
IT	Glaucoma	76.7	71.6	72.0	70.2	75.4	79.0
IT	Mixed	69.9	62.5	66.9	62.1	71.7	71.9
AVERAGE		72.5	68.6	69.0	64.2	72.0	73.3

Table 9: Zero-shot F1 scores for Argument Mining. Models have been trained in English and evaluated in Spanish, French and Italian.

Ideal Answer	Which type of urinary incontinence is diagnosed with the Q tip test? Stress urinary incontinence is diagnosed with the Q tip test. The test evaluates urethral mobility.
mT5-large	Q-tip test is used for urodynamic assessment of patients with severe urinary incontinence.
FlanT5-large	The Q tip test is used for assessing urethral mobility.
SciFive	Q tip test is used for evaluation of stress urinary incontinence.
MedMT5-large	Q-tip test is used for evaluation of stress urinary incontinence.
Ideal Answer	Which are the main manifestations of Ohdo syndrome? Severe ID, absent or deficient language, skeletal manifestations including bilateral patella dislocations.
mT5-large	Skeletal manifestations in Ohdo syndrome are a case with bilateral patella dislocations where surgical intervention has been indicated.
FlanT5-large	The main manifestations of Ohdo syndrome are: 1) severe ID, 2) absent or deficient language and 3) milder, clinical manifestation in heterozygotes.
SciFive	Ohdo syndrome is characterized by severe ID, absent or deficient language and, milder, clinical manifestation in heterozygotes.
MedMT5-large	The main manifestations of Ohdo syndrome are: 1) absent or deficient language and 2) milder clinical manifestation in heterozygotes.

Table 10: Examples of answers generated by each model for two different BioASQ questions together with the rank assigned by medics.

such as factuality, potential harm, and bias (Singhal et al., 2022). Consequently, we enlisted medical professionals to analyze the answers produced by the models.

During annotation, medical doctors were displayed the question, the ideal gold answers and the answers generated by each model. If required, they could also inspect the snippets that provide context to answer each of the questions. We narrowed the evaluation to Medical-mT5-large, mT5-large, FlanT5-large and SciFive. The evaluation was conducted by medical doctors proficient/native speakers of English, French and Spanish. For each question, doctors were asked to rank the answers generated by the models as the best, second-best, third-best, and worst answer.

Two Spanish medical doctors proficient or native in English and Spanish analyzed 50 English examples and 252 Spanish. For the French language, 3 French clinicians analyzed 186 answers, of which 47 were done by 2 doctors to calculate IAA (Cohen's Kappa Score: 0.28 and Average Spearman's Rank Correlation: 0.48), which indicates a low level of agreement. This exercise provided interesting insights with respect to the performance of the models in text generation tasks in the medical domain.

First, medical doctors could not in general establish significant differences between the answers generated by each of the models; predictions were far too similar, and all tended to fail on the same questions. As an example, Table 10 shows the answers to two different questions. As it can be observed, the answers generated by each model are very similar, and the doctors ended up ranking them primarily based on style.

The final result of the manual analysis is that all the models were chosen a similar number of times as the best. We believe that this demonstrates the difficulty of performing and obtaining meaningful evaluation results for this kind of tasks on this specific domain. This is in fact supported by the low IAA agreement obtained in the French annotation. This issue has also emerged in prior research and was partially addressed by employing a very large number of experts and asking them to respond with a yes/no to a set of predefined potential issues in the model output (Singhal et al., 2022). Still, the variance on the answers provided by the experts was significant.

However, there could be other underlying reasons for this behaviour. First, perhaps the T5 architecture is not ideally suited for text generation

as formulated in the BioASQ task, as these models are trained on a masking reconstruction objective rather than on direct text generation tasks. Consequently, the knowledge acquired during pre-training might not generalize well when the models are subsequently trained for text generation purposes. Second, perhaps using much larger models such as MedPaLM (Singhal et al., 2022) may generate better answer generation, but models of 540B parameters are currently unusable for the large majority of the NLP research labs, including ours. Nonetheless, it should be stressed that research on appropriate evaluation metrics for these tasks is still a difficult challenge which requires further investigation.

In any case, our results demonstrate the potential of a text-to-text model such as Medical mT5 for multilingual sequence labelling in the medical domain, establishing new state-of-the-art results in the multi-task and zero-shot cross-lingual settings.

8. Conclusion

In this paper we have presented Medical mT5, the first open source multilingual text-to-text LLM for the medical domain. Its development has required the compilation of a new 3B word corpus in English, French, Italian and Spanish specific to the medical domain. Furthermore, motivated by the lack of multilingual benchmarks, we have generated evaluation benchmarks for French, Italian and Spanish for Argument Mining and Abstractive Question Answering.

With respect to the languages chosen in this paper, we would like to comment that acquiring medical domain data is extremely challenging, even for languages such as the ones included. Furthermore, the choice of languages was also influenced by the availability of native medical doctors to do the manual evaluation for Abstractive Question Answering. In any case, we hope that our paper will encourage more researchers to join our effort and gather data for their respective languages, thereby creating larger, multilingual medical domain datasets encompassing more languages in the future.

A comprehensive experimentation on sequence labelling tasks shows that Medical mT5 outperforms strong text-to-text baselines of similarly-sized models in the multi-task and zero-shot cross-lingual evaluation settings. This is particularly interesting as these settings fully exploit the multilingual nature of a text-to-text model such as Medical mT5.

Furthermore, our experiments on Abstractive Question Answering show the inherent difficulty of evaluating generative tasks for this specific domain, where complex issues such as truthfulness and veracity are difficult to capture by automatic metrics. Manual evaluation is not ideal either, as

medical doctors were not able to clearly distinguish between the quality of the answers generated by the different models. In line with previous work (Singhal et al., 2022), we hope our paper will bring further attention to this problem and encourage further research on evaluation methods.

9. Acknowledgements

HiTZ Center: This work has been supported by the following MCIN/AEI/10.13039/501100011033 projects: (i) Antidote (PCI2020-120717-2) and EU NextGenerationEU/PRTR (ii) DeepKnowledge (PID2021-127777OB-C21) and by FEDER, EU; (iii) DeepR3 (TED2021-130295B-C31) and EU NextGeneration EU/PRTR. Iker García-Ferrero is supported by a doctoral grant from the Basque Government (PRE_2021_2_0219). Rodrigo Agerri currently holds the RYC-2017-23647 fellowship (MCIN/AEI/10.13039/501100011033 and by ESF Investing in your future).

FBK: This work has been supported by the European Union under Horizon Europe Projects ANTIDOTE (PCI2020-120717-2), eCREAM (Grant No. 101057726) and IDEA4RC (Grant No. 101057048). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union.

UCA: This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19P3IA-0002. This work was supported by the CHISTERA grant of the Call XAI 2019 of the ANR with the grant number Project-ANR-21-CHR4-0002.

10. Ethical Statement

Our research in developing Medical mT5, a multilingual text-to-text model for the medical domain, has ethical implications that we acknowledge. Firstly, the broader impact of this work lies in its potential to improve medical communication and understanding across languages, which can enhance healthcare access and quality for diverse linguistic communities. However, it also raises ethical considerations related to privacy and data security. To create our multilingual corpus, we have taken measures to anonymize and protect sensitive patient information, adhering to data protection regulations in each language’s jurisdiction or deriving our data from sources that explicitly address this issue in line with privacy and safety regulations and guidelines. Furthermore, we are committed to transparency and fairness in our model’s development and evaluation. We have worked to ensure

that our benchmarks are representative and unbiased, and we will continue to monitor and address any potential biases in the future. Finally, we emphasize our commitment to open source by making our data, code, and models publicly available, with the aim of promoting collaboration within the research community.

Bibliographical References

- David Adelani, Graham Neubig, Sebastian Ruder, Shruti Rijhwani, Michael Beukman, Chester Palen-Michel, Constantine Lignos, Jesujoba Alabi, Shamsuddeen 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 Memdjokam Koagne, Allahsera Auguste Tapo, Tebogo Macucwa, Vukosi Marivate, Mboning Tchiazé Elvis, Tajudeen Gwadabe, Tosin Adewumi, Orevaoghene Ahia, Joyce Nakatumba-Nabende, Neo Lerato Mokono, Ignatius Ezeani, Chiamaka Chukwuneke, Mofetoluwa Oluwaseun Adeyemi, Gilles Quentin Hacheme, Idris Abdulmumin, Odunayo Ogundepo, Oreen Yousuf, Tatiana Moteu, and Dietrich Klakow. 2022. [MasakhaNER 2.0: Africa-centric transfer learning for named entity recognition](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4488–4508, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rodrigo Aggerri and Eneko Agirre. 2023. Lessons learned from the evaluation of Spanish Language Models. *Proces. del Leng. Natural*, 70:157–170.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. 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)*, pages 3615–3620.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Casimiro Pio Carrino, Joan Llop, Marc Pàmies, Asier Gutiérrez-Fandiño, Jordi Armengol-Estapé, Joaquín Silveira-Ocampo, Alfonso Valencia, Aitor Gonzalez-Agirre, and Marta Villegas. 2022. [Pretrained biomedical language models for clinical NLP in Spanish](#). In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 193–199, Dublin, Ireland. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *CoRR*, abs/2210.11416.
- 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, ACL 2020, Online, July 5-10, 2020*, pages 8440–8451. Association for Computational Linguistics.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *CoRR*, abs/2207.04672.
- Rezarta Islamaj Dogan, Robert Leaman, and Zhiyong Lu. 2014. [NCBI disease corpus: A resource for disease name recognition and concept normalization](#). *J. Biomed. Informatics*, 47:1–10.
- Hermenegildo Fabregat, Juan Martínez-Romo, and Lourdes Araujo. 2018. [Overview of the DIANN task: Disability annotation task](#). In *Proceedings of the Third Workshop on Evaluation of Human*

- Language Technologies for Iberian Languages (IberEval 2018) co-located with 34th Conference of the Spanish Society for Natural Language Processing (SEPLN 2018)*, Sevilla, Spain, September 18th, 2018, volume 2150 of *CEUR Workshop Proceedings*, pages 1–14. CEUR-WS.org.
- Iker García-Ferrero, Rodrigo Agerri, and German Rigau. 2022. [T-projection: High quality annotation projection for sequence labeling tasks](#). *CoRR*, abs/2212.10548.
- Aitor Gonzalez-Agirre, Montserrat Marimon, Ander Intxaurrenondo, Obdulia Rabal, Marta Villegas, and Martin Krallinger. 2019. [PharmaCoNER: Pharmacological substances, compounds and proteins named entity recognition track](#). In *Proceedings of The 5th Workshop on BioNLP Open Shared Tasks, BioNLP-OST@EMNLP-IJNCLP 2019, Hong Kong, China, November 4, 2019*, pages 1–10. Association for Computational Linguistics.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2022. [Domain-specific language model pretraining for biomedical natural language processing](#). *ACM Trans. Comput. Heal.*, 3(1):2:1–2:23.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [LoRA: Low-rank adaptation of large language models](#). *arXiv preprint*, 2106.09685.
- Ayyoob Imani, Peiqin Lin, Amir Hossein Kargaran, Silvia Severini, Masoud Jalili Sabet, Nora Kassner, Chunlan Ma, Helmut Schmid, André F. T. Martins, François Yvon, and Hinrich Schütze. 2023. [Glot500: Scaling multilingual corpora and language models to 500 languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, *ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 1082–1117. Association for Computational Linguistics.
- Antonio Jimeno-Yepes, Aurélie Névoul, Mariana L. Neves, Karin Verspoor, Ondrej Bojar, Arthur Boyer, Cristian Grozea, Barry Haddow, Madeleine Kittner, Yvonne Lichtblau, Pavel Pecina, Roland Roller, Rudolf Rosa, Amy Siu, Philippe Thomas, and Saskia Trescher. 2017. [Findings of the WMT 2017 biomedical translation shared task](#). In *Proceedings of the Second Conference on Machine Translation, WMT 2017, Copenhagen, Denmark, September 7-8, 2017*, pages 234–247. 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](#). *Bioinform.*, 36(4):1234–1240.
- Jiao Li, Yueping Sun, Robin J. Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J. Mattingly, Thomas C. Wieggers, and Zhiyong Lu. 2016. [BioCreative V CDR task corpus: a resource for chemical disease relation extraction](#). *Database J. Biol. Databases Curation*, 2016.
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. [BioGPT: generative pre-trained transformer for biomedical text generation and mining](#). *Briefings in Bioinformatics*, 23(6).
- Bernardo Magnini, Begoña Altuna, Alberto Lavelli, Manuela Speranza, and Roberto Zanolli. 2021. [The E3C project: European clinical case corpus](#). In *Proceedings of the Annual Conference of the Spanish Association for Natural Language Processing: Projects and Demonstrations (SEPLN-PD 2021) co-located with the Conference of the Spanish Society for Natural Language Processing (SEPLN 2021)*, Málaga, Spain, September, 2021, volume 2968 of *CEUR Workshop Proceedings*, pages 17–20. CEUR-WS.org.
- Tobias Mayer, Elena Cabrio, and Serena Villata. 2020. [Transformer-based argument mining for healthcare applications](#). In *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 2108–2115. IOS Press.
- Tobias Mayer, Santiago Marro, Elena Cabrio, and Serena Villata. 2021. [Enhancing evidence-based medicine with natural language argumentative analysis of clinical trials](#). *Artificial Intelligence in Medicine*, 118:102098.
- Long N. Phan, James T. Anibal, Hieu Tran, Shaurya Chanana, Erol Bahadroglu, Alec Peltekian, and

- Grégoire Altan-Bonnet. 2021. [SciFive: a text-to-text transformer model for biomedical literature](#). *CoRR*, abs/2106.03598.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Kumar Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Nathaneal Schärli, Aakanksha Chowdhery, Philip Andrew Mansfield, Blaise Agüera y Arcas, Dale R. Webster, Gregory S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle K. Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2022. Large language models encode clinical knowledge. *arXiv preprint*, abs/2212.13138.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- 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.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and efficient foundation language models. *arXiv preprint*, 2302.13971.
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R. Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, Yannis Almirantis, John Pavlopoulos, Nicolas Baskiotis, Patrick Gallinari, Thierry Artières, Axel-Cyrille Ngonga Ngomo, Norman Heino, Eric Gaussier, Liliana Barrio-Alvers, Michael Schroeder, Ion Androutsopoulos, and Georgios Paliouras. 2015. [An overview of the BioASQ large-scale biomedical semantic indexing and question answering competition](#). *BMC Bioinform.*, 16:138:1–138:28.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Guangyu Wang, Guoxing Yang, Zongxin Du, Longjun Fan, and Xiaohu Li. 2023. ClinicalGPT: Large Language Models Finetuned with Diverse Medical Data and Comprehensive Evaluation. *ArXiv preprint*, abs/2306.09968.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. PMC-LLaMA: Towards building open-source language models for medicine. *arXiv preprint*, 2304.14454.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Michihiro Yasunaga, Jure Leskovec, and Percy Liang. 2022. [LinkBERT: Pretraining language models with document links](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2022, Dublin, Ireland, May 22-27, 2022, pages 8003–8016. Association for Computational Linguistics.
- Anar Yeginbergenova and Rodrigo Agerri. 2023. Cross-lingual argument mining in the medical domain. *arXiv preprint*, abs/2301.10527.

Language Resource References

Ander Itxaurrondo. 2018. *SPACCC: Spanish Clinical Case Corpus*. Barcelona Supercomputing Center. PID <https://github.com/PlanTL-GOB-ES/SPACCC>.

Common Crawl. 2022. *Common Crawl*. Common Crawl. PID <https://commoncrawl.org/>.

Institute of Formal and Applied Linguistics. 2017. *UFAL Medical Corpus v. 1.0*. Charles University, Czech Republic. PID https://ufal.mff.cuni.cz/ufal_medical_corpus.

National Library of Medicine. 2022a. *Clinical Trials*. National Library of Medicine. PID <https://clinicaltrials.gov/>.

National Library of Medicine. 2022b. *PubMed*. National Library of Medicine. PID <https://pubmed.ncbi.nlm.nih.gov>.

