

Fighting crime with Transformers: Empirical analysis of address parsing methods in payment data

Haitham Hammami
haitham.hammami¹

Louis Baligand
louis.baligand¹

Bojan Petrovski
bojan.petrovski¹

¹@alumni.epfl.ch

Abstract

In the financial industry, identifying the location of parties involved in payments is a major challenge in the context of various regulatory requirements. For this purpose address parsing entails extracting fields such as street, postal code, or country from free text message attributes. While payment processing platforms are updating their standards with more structured formats such as SWIFT with ISO 20022¹, address parsing remains essential for a considerable volume of messages. With the emergence of Transformers and Generative Large Language Models (LLM), we explore the performance of state-of-the-art solutions given the constraint of processing a vast amount of daily data. This paper also aims to show the need for training robust models capable of dealing with real-world noisy transactional data. Our results suggest that a well fine-tuned Transformer model using early-stopping significantly outperforms other approaches. Nevertheless, generative LLMs demonstrate strong zero-shot performance and warrant further investigations.

1 Introduction

To ensure adherence with regulatory requirements, it is essential for financial institutions to understand precisely where the money is originating and where it is flowing. The new standard of international payment messages ISO 20022 for SWIFT has the potential to simplify the task of locating payment parties by enabling the beneficiary and originator address to be delivered in a structured format. However, a considerable amount of messages are still delivered with an address in free text form. This problem is further exacerbated by the use of legacy payment processing platforms. Thus *Address Parsing* is required to extract address fields such as street, postal code, city, or country.

¹<https://www.iso20022.org/iso-20022-message-definitions>

Our work has three main contributions. Firstly, it offers an open-sourced, augmented dataset, addressing the limitations of bench-marking on clean datasets and enabling research on noisy real-world payment data. Secondly, by empirically analyzing and comparing various techniques, this paper uncovers an effective approach for multinational address parsing on distorted data. Lastly, we open-source the fine-tuned state-of-the-art model, aiding future research and application in a multinational setup written in Latin alphabet and transliterated in ASCII format.

2 Related Work

Looking through prior art, we identified various solutions to the problem at hand. Early rule-based attempts (Xu et al., 2012) were shown to be incapable of dealing with the complexities of noise in real-world data. Generative approaches based on Hidden Markov Models (Li et al., 2014) and probabilistic Conditional Random Fields (CRF) (Wang et al., 2016) give the first promising results on real-world data. A notable and popular off-the-shelf solution based on the CRF approach is LibPostal, a C library that enables multilingual address parsing and normalization². As it was trained on millions of multilingual, international addresses, this model shows great potential for robustness when faced with unknown addresses from different countries. Therefore we consider it as the baseline solution used for our benchmarks.

In recent years, advancements in natural language processing have been driven by deep learning techniques (Collobert et al., 2011). This is evidenced by solutions like DeepParse (Beauchemin and Yassine, 2023), which leverages a Seq2Seq bidirectional-LSTM neural network architecture, and Transformer-based approaches such as the one proposed by Guermazi et al. (2023). The efficacy

²<https://github.com/openvenues/libpostal>

of Transformer models has also been explored in domain-specific scenarios, as demonstrated in studies performed by Sahay et al. (2023) and Kulkarni et al. (2023). While these existing frameworks provide valuable insights most of them have either a domain, country, or robustness limitation. Our work seeks to benchmark some of these techniques on payment data and build on them to develop a solution tailored to our task while also exploring more innovative approaches.

3 Data

The basis of our data comes from the training data of *DeepParse*³ (Yassine et al., 2020), which is a set of postal addresses across multiple countries, with the exact number of samples per country being reported in their work. This data has been generated using data from the *lipbostal* project, which in itself was trained on real-world addresses from OpenStreetMap. Each sample is provided as a tuple: the address (string) and a list of tags, one for each word in the address. A tag can be either StreetName, StreetNumber, Unit, PostalCode, Municipality, or Province. The dataset is comprised of two types of data: **clean** that have all the aforementioned tags in every sample, except for Unit and PostalCode being optional, and **incomplete** missing at least one of the other tags. We only sample from the clean data, from which we present an example in Table 1. It is important to highlight that there is no indication of the country. This is a serious limitation as the main goal of the address parsing in transactional data is country derivation for regulatory purposes. Another limitation is the absence of the beneficiary name, an equally important part of SWIFT messages. These limitations are further discussed in section 3.1.

From the clean data of the original dataset, we extract at most 100K samples from each country’s training and testing set and we union the extracted samples to get a total of 3.4M addresses. During sampling, we remove all addresses that are not written in the Latin alphabet as this is the standard in SWIFT messages. Out of these addresses, we re-sample 100K rows for testing and leave the rest for training as we need a large volume of data to train the large models. Then from a set of addresses from countries that were not used in training and testing sets, we create a **zero-shot** dataset by extracting at

most 100K samples from each country.

From this point onward, we will refer to these datasets as **synthetic** data, from which we create three versions V_0 , V_1 and V_2 as described in section 3.1.

3.1 Augmented Data

The goal of data augmentation is to mimic as closely as possible the structure of our real-life data, referred to as **production** data. We performed a comprehensive analysis on production samples comprised of 1,600 manually labeled messages, handpicked in a way that covers the most common addresses from all the countries where data is available. This analysis enabled us to deduce the underlying distribution of the address structure and how straightforward it is to parse it. Since SWIFT messages are in free text, postal address standards are sometimes not respected, as some parts of the address are omitted, misspelled, or wrongfully positioned. Moreover, we can find parts of the message that are completely unrelated to the address itself, such as adding a phone number, or an account number or going to the extent of writing "Address information is in line 4". Based on these observations, we augment our synthetic data as follows:

- A **Name**: For adding the name we used the python library *Faker*⁴ to generate fake person and company names, which we prepend to the address by adding the tag Name in the corresponding position of the tag list.
- B **Country**: By analyzing the SWIFT messages at our disposal, we observed that the presence of the country in the address can be in either by displaying its name (in different languages) or its ISO code. By simulating the distribution of these occurrences from our production data, we insert the country in its position in the address with the tag Country or CountryCode. The languages used are English, French, German, Italian, and the country’s original language.
- C **Address structure**: As it was well-examined in Yassine et al. (2020), the synthetic data is well-structured, with each address following its country’s standard address format, as opposed to production data, where the free-text message may or may not follow these standards. To mitigate this discrepancy, for each

³<https://github.com/GRAAL-Research/deepparse-address-data>

⁴<https://faker.readthedocs.io/en/master/>

StreetName	StreetNumber	PostalCode	Municipality	Province
jakob-sturm-w.	35	80995	munich	bavaria

Table 1: Data Sample

address of the synthetic data we randomly sample a production address and we apply its ordering of tags as a mask. This leads to potential changes in the synthetic address, including the removal, addition, or rearrangement of elements. Consequently, some addresses may deviate from the standard structure of the country’s address format. This variation aims to prevent the models from overfitting to country-specific address structures, addressing a concern highlighted by Yassine et al. (2021). In their study, they observed a significant drop in model performance on zero-shot data, particularly on addresses from countries with distinct address formats compared to those present in the training data.

D Line separation: In around a third of the SWIFT messages we observe that the name and the address are separated by a line return. This is additional information that the model can leverage to perform better parsing, therefore we add to the third of the synthetic data the symbol "\$" between the name and address and give it the tag "HardSep".

E Labeling augmentation: We adapt our labels to the **BIO** tagging schema (Sang and Buchholz, 2000) by adding the prefix "B-" to the beginning of the class, and "I-" for words inside of the class. Since the synthetic data is purely address-related and does not have terms outside of the previously stated tags, no word would be tagged as "O". To adapt to our production data, we introduce a new tag: **OOA** (Out-Of-Address), that is any term found in the SWIFT message unrelated to the address itself. Therefore, while rearranging the address structure, whenever the production sample has an OOA tag, we generate an OOA term and insert it into the synthetic address. This term can be a random number, an alphanumeric code, a postbox number, or a duplicate term from the address itself. The latter is labeled as OOA and not its real tag because it is redundant and often miss-placed, as seen in Table 2 where the term for OOA was cho-

sen to be a redundant mention of the country. The decision of the type of the generated term is based on the analyzed probability distribution of the nature of this OOA in production data.

During the process of data augmentation, we retain three versions of datasets:

- V_0 : the original sampled addresses with no augmentation.
- V_1 : Given V_0 , we apply basic cleaning and address structure masking technique (rearrangement or removal, but no addition of augmented address parts).
- V_2 : Given V_0 , we apply all the aforementioned augmentation techniques described in points A to E.

4 Proposed Approaches

Given a free text as string, the task is to parse the address fields. The possible fields should be chosen among "Name", "StreetName", "StreetNumber", "Unit", "Municipality", "PostalCode", "Province" and "Country". Note that not all words need to be parsed.

Input:

```
"jakob-sturm-w. 35 80995 munich bavaria"
```

Output:

```
{"StreetName": "jakob-sturm-w.",
"StreetNumber": "35",
"PostalCode": "80995",
"Municipality": "munich",
"Province": "bavaria"}
```

4.1 LibPostal

Currently being a very prominent tool for our task, it is important to benchmark this model’s performance against our subsequent approaches. Since it is using a different labeling schema, we attempt to create a mapping to the set of tags we are using to align it with our metrics. The mapping is provided in the appendix in Table 6.

Synthetic address	jakob-sturm-w. 35 80995 munich bavaria
Synthetic tags	[StreetName, StreetNumber, PostalCode, Municipality, Province]
Prod mask	[Name, StreetName, StreetNumber, Municipality, PostalCode, Country, OOA]
augmented address	John Doe, jakob-sturm-w. 35 munich 80995 germany germany

Table 2: Data augmentation example

4.2 DeepParse

Given that we are working with the same training dataset, it is meaningful to use this model as a baseline. For V_0 and V_1 , we are using the same set of tags, therefore we run the inference using the standard model with *Byte-pairs embedding (bpemb)*. Since we are adding more tags into V_2 , namely *Name*, *Hardsep*, *OOA*, *Country* and *CountryCode*, we retrain the model on the new set of tags, using V_2 as training data and the default training parameters proposed in the library.

4.3 Transformers

Given that we are dealing with a token classification task, our first intuition was to experiment with the transformer models family, given their proven track record in performing well with similar tasks, namely for Named Entity Recognition (NER) (Liu et al., 2021). For all the subsequent trials, the experimental setup and hyperparameters are reported in the appendix A.1.

4.3.1 Baseline models and initial results

We started by looking at the distilled versions of the BERT family given their small size and short training duration (Sanh et al., 2020), namely *distilbert-base-uncased*, *distilbert-base-multilingual-cased*, *distilbert-cased* and *distilroberta-base*. The performance of the four transformer models is reported in Table 3. The numbers show that the model *distilbert* outperforms the other models in every set, while also being the most robust model across its copies (folds) since it has the smallest standard deviation, which makes us confident to perform our subsequent experimentation using this model.

4.3.2 Further experiments

We attempt to challenge our early assumptions on the way we augment and generate the data by modifying V_2 in two separate ways: removing the prefix from the tag and removing the HardSep token. Moreover, Figure 1 suggests overfitting, prompting us to experiment with early stopping, using cross-entropy loss on zero-shot data as the stopping metric. For completeness, we retrain the model using

V_0 and V_1 training datasets to have comprehensive bench-marking with LibPostal and DeepParse.




Figure 1: Learning curve showing the model's overfit

The results reported in 5, further discussed later, show that the only variation that substantially improved the model's performance is the addition of early stopping. Using the best performing configuration, we switch to larger transformer models, namely *bert-base-uncased*, *bert-large-uncased* and *xlm-roberta-large*.

4.4 Decoder Based Approach

Given the various legal and regulatory constraints stemming from bank privacy laws especially when dealing with cross-jurisdictional payments, API-based LLMs such as GPT-4 are not something we could productionalize. Furthermore, as we are dealing with large volumes of daily data our focus was on "small" sized LLMs that have the best performance-to-cost ratio.

We chose Llama 2 7b (Touvron et al., 2023) and its derivative Mistral-7B (Jiang et al., 2023) as benchmark models. We designed several prompts following Prompt Engineering Best Practices⁵⁶ and selected the best performing one. Our full prompt can be found in the appendix in Figure 2. To tackle cases where the same word appears twice in the input as two different tags, we prefix each word with its index. This will enable us to easily post-

⁵<https://www.promptingguide.ai/introduction/tips>

⁶<https://help.openai.com/en/articles/6654000-best-practices-for-prompt-engineering-with-openai-api>

process the output in the desired structured format. Given that we are using "small" sized LLMs we did not need to quantize the models.

We use 1,000 addresses from our V_2 zero-shot dataset to benchmark both LLAMA 2 7b (Chat) and Mistral-7B (Instruct). As Mistral-7B performed substantially better our further trials focused on extracting more performance from it by experimenting with the inference sampling parameters and by attempting to train the LORA adapter (Hu et al., 2021) for our specific task.

4.5 Metrics

We compute the Precision and Recall score based on each tag, we then average the F1 score overall and also show the standard deviation based on each of the transformer-based trained versions for each model and each dataset. To keep the purpose of the application general we use F1 score without giving more weight to precision or recall. For LLMs, we use a 10,000 random sample from the training and zero-shot dataset. There is no standard deviation in this case.

5 Results and discussion

5.1 Transformer Based Approach

We summarize our findings in Table 3. The main takeaways from it are:

- **train/test scores:** what stands out at first glance is how small the difference between these scores is across all experiments. This could indicate how well the models can perform the task. However, this resemblance makes it hard to spot over-fitting, which makes us focus on the zero-shot score as the basis of our comparison.

- **LibPostal:** The inference scores for the three datasets are very consistent, and dropping across the data versions. This is an expected result as LibPostal had a large multinational training set so it was equally exposed to all countries, but its training data highly resembles V_0 , so its performance drops as the data gets more noisy. The main observed limitation of LibPostal comes from the fact that it only expects address-related parts and complete addresses, which is seldom the case with payment data and consequently the V_2 dataset.

- **DeepParse:** The results of V_0 matches the reported numbers of Yassine et al. (2020). What's interesting is the improvement in the zero-shot

score between V_0 and V_1 , also observed for distilbert. This is caused by the alteration of the address structure that produced a form of *Data Leakage* that is necessary for robustness but should be nonetheless taken into consideration when comparing the models (Elangovan et al., 2021). For V_2 , the retrained model failed to produce satisfactory results.

- **DistilBERT:** DistilBERT surpassed DeepParse on all data versions, including outperforming Postal on V_1 and V_2 . This underscores the transformer models' adaptability and strong generalization capability across unseen data during training. Specifically, in our comparison among transformers (as discussed in 4.3.1), the DistilBERT uncased version yielded the best results, aligning with the uncased nature of our data; this is further supported by the lower scores of the cased version.

Continuing our experimentation with the uncased model, we see that removing the *HardSep* token does not significantly improve the performance, so we decide to keep it as we deem that it will be even more useful on more noisy production data. Removing the prefix yields a slightly better score, but a more notable improvement comes with the early stopping. Combining the last two modifications produces a lower score than just having early stopping so we keep it as the sole modification on our training pipeline.

- **Larger models:** We observe a clear pattern of the score increasing as we increase the size of the model, with the best result achieved by XLM-RoBERTa-Large, despite being a cased model. Testing this model against a few other selected models in predicting production data, we see that it outperforms them still, with the observation of a clear correlation between zero-shot and prod scores, which confirms the justification of using the former as the basis of our comparison.

- **Generative LLMs:** In Table 4 we show the main results of the Generative LLMs on 10,000 samples from dataset V_2 . As stated earlier Mistral 7-B demonstrated a relatively good performance just with simple prompt engineering.

Given the repetitive nature of our prompts, we found that the LORA adapters are very susceptible to over-fitting, and can deliver a significant performance improvement with just 1,000 training

Data Version	Model Version	train		test		zero-shot		prod**
		mean	std*	mean	std*	mean	std*	
V ₀	LibPostal	0.997	-	0.997	-	0.992	-	-
	DeepParse	0.991	-	0.992	-	0.727	-	-
	distilbert-base-uncased	0.999	0.050	0.999	0.010	0.885	1.536	-
V ₁	LibPostal	0.898	-	0.897	-	0.918	-	-
	DeepParse	0.828	-	0.826	-	0.767	-	-
	distilbert-base-uncased	0.995	0.425	0.995	0.025	0.924	0.400	-
V ₂	LibPostal	0.761	-	0.759	-	0.781	-	0.675
	DeepParse	0.747	-	0.747	-	0.709	-	-
	distilbert-base-uncased	0.994	0.371	0.994	0.045	0.859	0.915	0.765
	distilbert-multi	0.992	0.212	0.992	0.068	0.827	1.708	-
	distilroberta	0.992	0.188	0.991	0.054	0.835	2.998	-
	distilbert-cased	0.991	1.361	0.990	1.230	0.821	-	-
	distilbert-no-hardsep	0.994	0.338	0.994	0.066	0.860	1.075	-
	distilbert-noprefix	0.995	0.401	0.994	0.017	0.862	0.938	-
	distilbert-early-stopping	0.976	2.711	0.975	2.649	0.871	0.991	0.766
	distilbert-no-prefix+early-stopping	0.931	11.240	0.929	10.980	0.863	3.027	-
	bert-base-uncased	0.967	15.768	0.966	15.843	0.868	5.170	0.784
	bert-base-large	0.966	5.013	0.965	4.977	0.909	2.140	0.785
	xlm-roberta-large	0.975	4.817	0.974	4.996	0.924	2.236	0.801

* Reported values of standard deviation are multiplied by 10^3 for clarity

** Actual real-life payment address data, referred as **production** data

Table 3: F1 scores of baseline and Transformer models

Model	zero-shot
Llama2 7B	0.4650
Mistral 7B Instruct v0.2	0.6066
Mistral 7B Instruct v0.2 (SFT*)	0.7113
Mixtral 8x7B Instruct v0.1	0.7233

Table 4: F1 score on Decoder Based approach
*Fine-Tuned on 1,000 samples with 3 epochs

samples for 3 epochs. Furthermore, we tried several ranks for our adapters and settled for a rank of 8 with the standard LORA value of $2 * r$.

One interesting observation is the impact of the inference parameters on the performance and the quality of the output. For the base Mistral-7B model, the performance remains stable with the different parameters with an F1 of approximately 0.6. However, for our fine-tuned model with a LORA adapter, we see the performance changing from an F1 score of 0.67 to 0.71. A full breakdown of the performance with the different sampling configurations is available in the appendix. Finally, for reference, we did a benchmark on the Mixtral-8X7B model with no fine-tuning and we were able to achieve a similar F1 score of 0.72

As expected during our manual checks we noticed hallucinations by the LLMs. Sometimes the generated text modifies the word’s index. In case the word appears only once, we can retrieve its position in the input, however, in case there is an ambiguity the prediction fails. The model may also not follow the instructions by generating an unexpected output format, e.g. a nested JSON. In this case, we attempt to flatten the JSON and retrieve the expected tag. This mostly happens when there is a country code and the full country name in the input address. The output may unexpectedly also include a "Hardsep" "\$" and may create new tags that are not in the input. For example, inferring the country name from the city.

6 Conclusion

In summary in this paper, we introduced a new open-sourced dataset reflecting the limitations and noise of real-world payment data. This will enable better benchmarking of expected production performance and further research on this problem. Our experimental results highlight the importance of training robust models that are capable of dealing with noise. We achieve state-of-the-art performance on the synthetic zero-shot data and our

production data with a fine-tuned XLM-RoBERTa-Large model. A derivative of this model is currently deployed in our production systems.

Overall even though the generative LLMs were not able to match the performance of the encoder transformer models, we believe there is a strong potential that warrants further inquiry. Possible directions for further improvement would be the use of context-free grammars such as the LMQ library (Beurer-Kellner et al., 2023) to force a more structured output, and the introduction of geography knowledge with a LlamaIndex.

In the final version of this paper, we intend to open-source all training and evaluation code, along with the fine-tuned models. By doing so, we aim to offer a valuable resource to other companies and researchers confronting analogous challenges in Address Parsing, applicable across various real-world applications beyond the financial industry.

References

- David Beauchemin and Marouane Yassine. 2023. [Deep-parse : An extendable, and fine-tunable state-of-the-art library for parsing multinational street addresses.](#)
- Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. 2023. [Prompting is programming: A query language for large language models.](#) *Proceedings of the ACM on Programming Languages*, 7(PLDI):1946–1969.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel P. Kuska. 2011. [Natural language processing \(almost\) from scratch.](#) *CoRR*, abs/1103.0398.
- Aparna Elangovan, Jiayuan He, and Karin Verspoor. 2021. [Memorization vs. generalization: Quantifying data leakage in nlp performance evaluation.](#)
- Yassine Guerhazi, Sana Sellami, and Omar Boucelma. 2023. [Georoberta: A transformer-based approach for semantic address matching.](#) 3379.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models.](#) *CoRR*, abs/2106.09685.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b.](#)
- Mandar Kulkarni, Nikesh Garera, and Anusua Trivedi. 2023. [Domain-specific transformer models for query](#)

- translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, pages 89–95, Toronto, Canada. Association for Computational Linguistics.
- Xiang Li, Hakan Kardes, Xin Wang, and Ang Sun. 2014. Hmm-based address parsing: efficiently parsing billions of addresses on mapreduce. *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*.
- Zihan Liu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021. *Ner-bert: A pre-trained model for low-resource entity tagging*.
- Rishav Sahay, Anoop S V K K Saladi, and Prateek Sircar. 2023. *Multi-task student teacher based unsupervised domain adaptation for address parsing*. In *PAKDD 2023*.
- Erik F. Tjong Kim Sang and Sabine Buchholz. 2000. *Introduction to the conll-2000 shared task: Chunking*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. *Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. *Llama 2: Open foundation and fine-tuned chat models*.
- Minlue Wang, Valeriia Haberland, Amos Yeo, Andrew Martin, John Howroyd, and John Bishop. 2016. *A probabilistic address parser using conditional random fields and stochastic regular grammar*.
- Sen Xu, Soren Flexner, and Vitor Carvalho. 2012. Geocoding billions of addresses: toward a spatial record linkage system with big data. *GIScience in the Big Data Age*, 13.
- Marouane Yassine, David Beauchemin, Francois Laviolette, and Luc Lamontagne. 2020. *Leveraging subword embeddings for multinational address parsing*. In *2020 6th IEEE Congress on Information Science and Technology (CiSt)*. IEEE.
- Marouane Yassine, David Beauchemin, François Laviolette, and Luc Lamontagne. 2021. *Multinational address parsing: A zero-shot evaluation*.

A Appendix

A.1 Experimental setup

We perform a 4-fold cross-validation on each model, where for each fold we tokenize the training and validation sets with the model-specific pre-trained tokenizer. We also align the tags in a way that each tag corresponds to the starting token of the word (or the whole word if it was not split), the whole process is shown in Table 5. We use the final format as input to the model, with the hyperparameters depicted in Table 7.

A.2 Data & Source Code

The data and source code can openly be accessed on <https://arxiv.org/abs/2404.05632>.

Tag	Word	Subword	Final Tag
		[CLS]	UNK
StreetName	kirchenstr	ki	B-StreetName
		##rch	UNK
		##ens	UNK
		##tr	UNK
StreetNumber	24	24	B-StreetNumber
PostalCode	3660	36	B-PostalCode
		##60	UNK
Municipality	gemeinde	gem	B-Municipality
		##ein	UNK
		##de	UNK
Municipality	klein	klein	I-Municipality
Municipality	pochlarn	po	I-Municipality
		##ch	UNK
		##lar	UNK
		##n	UNK
Province	niederosterreich	ni	B-Province
		##ede	UNK
		##ros	UNK
		##ter	UNK
		##re	UNK
		##ich	UNK
		[SEP]	UNK

Table 5: Tokenization example

Postal tags	V_0 and V_1 mapping	V_2 mapping
house_number	StreetNumber	StreetNumber
road	StreetName	StreetName
house	Unit	Name
level	Unit	Unit
city	Municipality	Municipality
state	Province	Province
state_district	Province	Province
unit	Unit	Unit
postcode	PostalCode	PostalCode
country	Province	Country
suburb	Municipality	Municipality
city_district	Municipality	Municipality
category	StreetName	OOA
near	Municipality	OOA
po_box	PostalCode	OOA
entrance	Unit	OOA
country_region	Province	Country
staircase	Unit	OOA
world_region	Province	Province
island	Province	OOA

Table 6: Postal tags mapping

Hyperparameter	Value
num train epochs	1
train batch size	1024
evaluation batch size	1024
dropout	0.1
learning rate scheduler warmup steps	500
optimizer	adamw
optimizer weight decay	0.01
evaluation steps	20
early stopping patience*	5
seed	42

* only when early stopping is stated to be used

Table 7: Training Hyperparameters

```

f"""<s>[INST]
You are a word classifier that classifies words from a text corresponding to an
    ↪ address free text field.
You should analyze with deep precision the INPUT and return a dictionary with
    ↪ the following keys: "Name", "StreetNumber", "StreetName", "Municipality",
    ↪ "PostalCode", "Unit", "Country", "CountryCode".
Each word is separated by a space and should be classified without any
    ↪ modification.
Each word in the input has a prefix with the index i as '[i]-' and it should be
    ↪ ignored for the classification but it should remain AS-IS in the output.
    ↪
Sub sequence of words should be classified as follow:
'Name': words corresponding to an individual name or institution name.
'StreetNumber': words corresponding to a street number.
'StreetName': words corresponding to a street name.
'Municipality': words corresponding to a municipality or city.
'PostalCode': words corresponding to a postal code.
'Unit': words corresponding to a unit number.
'Country': words corresponding to a full country name.
'CountryCode': words corresponding to a country iso2 code.

Output Indicator:
2. Usually a name comes before the address.
3. "$" is indicating a large separator and it should not be classified.
4. The output words should be taken from the input only and it should not be
    ↪ modified
5. The same word cannot be used in two different classes.
6. Words are classified subsequently.
7. Empty classes should not appear in the output.
8. Output should not include nested values.
9. Each index are taken from the input itself and the index matches, e.g. the
    ↪ prefix '[i]-' remains unchanged for all words.

For example:
### INPUT:
"[0]-THOMASSEN [1]-GULBRANDSEN [2]-OG [3]-GUNDERSEN [4]-$ [5]-TV [6]-SD [7]-9
    ↪ [8]-JAPARATINGA [9]-57950 [10]-000 [11]-BR"
### OUTPUT:
{"Name": "[0]-THOMASSEN [1]-GULBRANDSEN [2]-OG [3]-GUNDERSEN", "StreetName":
    ↪ "[5]-TV [6]-SD [7]-9", "Municipality": "[8]-JAPARATINGA", "PostalCode":
    ↪ "[9]-57950 [10]-000", "CountryCode": "[11]-BR"}

### INPUT:
{address}
[/INST]

### OUTPUT:
"""

```

Figure 2: Prompt Template used for LLMs

Model	min_p [†]	top_p [‡]	Temperature		
			0.8	0.5	0.2
Llama-2-7b-chat-hf	0.1	-	-	-	0.465
Llama-2-7b-hf	0.1	-	-	-	0.424
Mistral-7B-Instruct-v0.2	0.1	-	0.601	0.601	0.607
Mistral-7B-Instruct-v0.2	0.3	-	0.605	0.601	0.602
Mistral-7B-Instruct-v0.2	0.5	-	0.601	0.606	0.602
Mistral-7B-Instruct-v0.2	-	0.9	0.603	0.603	0.604
Mistral-7B-Instruct-v0.2	-	0.7	0.601	0.606	0.604
Mistral-7B-Instruct-v0.2	-	0.5	0.605	0.605	0.604
Mistral-7B-Instruct-v0.2 (SFT*)	0.1	-	0.675	0.696	0.708
Mistral-7B-Instruct-v0.2 (SFT*)	0.3	-	0.698	0.702	0.709
Mistral-7B-Instruct-v0.2 (SFT*)	0.5	-	0.702	0.706	0.710
Mistral-7B-Instruct-v0.2 (SFT*)	-	0.9	0.671	0.695	0.709
Mistral-7B-Instruct-v0.2 (SFT*)	-	0.7	0.700	0.710	0.709
Mistral-7B-Instruct-v0.2 (SFT*)	-	0.5	0.710	0.709	0.708

Table 8: F1 score of LLMs with various text generation parameters. Variants of Mistral 7B without SFT is not significant while the Fine-Tuned versions are more sensible to text generation parameters change.

* Trained on 1,000 samples and 3 epochs

[†] Minimum probability for a token to be considered, relative to the probability of the most likely token. top_p does not vary when min_p is set.

[‡] Cumulative probability of top tokens to be considered. min_p does not vary when top_p is set.