

Named Entity Recognition for Address Extraction in Speech-to-Text Transcriptions Using Synthetic Data

Bibiána Lajčinová¹, Patrik Valábek^{1, 3}, Michal Spišiak²

¹Slovak National Supercomputing Centre, Dúbravská cesta 3484/9, 84104 Bratislava, Slovak Republic
{bibiana.lajcinova,patrik.valabek}@nsc.sk

²nettle, s.r.o., Matúšova 56/A, 811 04 Bratislava - Staré Mesto, Slovak Republic
michal.spisiak@nettle.ai

³Institute of Information Engineering, Automation, and Mathematics, Slovak University of Technology in Bratislava, Radlinského 9, 81237 Bratislava, Slovak Republic

1 Introduction

Many businesses spend large amounts of resources for communicating with clients. Usually, the goal is to provide clients with information, but sometimes there is also a need to request specific information from them.

In addressing this need, there has been a significant effort put into the development of chatbots and voicebots, which on one hand serve the purpose of providing information to clients, but they can also be utilized to contact a client with a request to provide some information.

A specific real-world example is to contact a client, via text or via phone, to update their postal address. The address may have possibly changed over time, so a business needs to update this information in its internal client database.

Nonetheless, when requesting such information through novel channels—like chatbots or voicebots—it is important to verify the validity and format of the address. In such cases, an address information usually comes by a free-form text input or as a speech-to-text transcription. Such inputs may contain substantial noise or variations in the address format. To this end it is necessary to filter out the noise and extract corresponding entities, which constitute the actual address. This process of extracting entities from an input text is known as Named Entity Recognition (NER). In our particular case we deal with the following entities: municipality name, street name, house number, and postal code. This technical report describes the development and evaluation of a NER system for extraction of such information.

2 Problem Description and Our Approach

This work is a joint effort of Slovak National Competence Center for High-Performance Computing and nettle, s.r.o., which is a Slovak-based start-up focusing on natural language processing, chatbots, and voicebots. Our goal is to develop highly accurate and reliable NER model for address parsing. The model accepts both free text as well as speech-to-text transcribed text. Our NER model constitutes an important building block in real-world customer care systems, which can be employed in various scenarios where address extraction is relevant.

The challenging aspect of this task was to handle data which was present exclusively in Slovak language. This makes our choice of a baseline model very limited.

Currently, there are several publicly available NER models for the Slovak language. These models are based on the general purpose pre-trained model SlovakBERT [1]. Unfortunately, all these models support only a few entity types, while the support for entities relevant to address extraction is missing. A straightforward utilization of popular Large Language Models (LLMs) like GPT is not an option in our use cases because of data privacy concerns and time delays caused by calls to these rather time-consuming LLM APIs.

We propose a fine-tuning of SlovakBERT for NER. The NER task in our case is actually a classification task at the token level. We aim at achieving proficiency at address entities recognition with a

tiny number of real-world examples available. In Section 2.1 we describe our dataset as well as a data creation process. The significant lack of available real-world data prompts us to generate synthetic data to cope with data scarcity. In Section 2.2 we propose SlovakBERT modifications in order to train it for our task. In Section 2.3 we explore iterative improvements in our data generation approach. Finally, we present model performance results in Section 3.

2.1 Data

The aim of the task is to recognize street names, house numbers, municipality names, and postal codes from the spoken sentences transcribed via speech-to-text. Only 69 instances of real-world collected data were available. Furthermore, all of those instances were highly affected by noise, e.g., natural speech hesitations and speech transcription glitches. Therefore, we use this data exclusively for testing. Table 1 shows two examples from the collected dataset.

| Sentence | Tokenized text | Tags |
|---------------------------|----------------|----------------|
| Stupava Záhumsená 834 | Stupava | B-Municipality |
| | Záhumsená | B-Street |
| | 834 | B-Housenumber |
| Ďalšie bauerová 44 Košice | Ďalšie | O |
| | bauerová | B-Street |
| | 44 | B-Housenumber |
| | Košice | B-Municipality |

Table 1: Two example instances from our collected real-world dataset. The Sentence column shows the original address text. The Tokenized text column contains tokenized sentence representation, and the Tags column contains tags for the corresponding tokens. Note here that not every instance necessarily contains all considered entity types. Some instances contain noise, while others have grammar/spelling mistakes: The token “Ďalšie” is not a part of an address and the street name “bauerová” is not capitalized.

Artificial generation of training dataset occurred as the only, but still viable option to tackle the problem of data shortage. Inspired by the 69 real instances, we programmatically conducted numerous external API calls to OpenAI to generate similar realistic-looking examples. BIO annotation scheme [2] was used to label the dataset. This scheme is a method used in NLP to annotate tokens in a sequence as the beginning (B), inside (I), or outside (O) of entities. We are using 9 annotations: *O*, *B-Street*, *I-Street*, *B-Housenumber*, *I-Housenumber*, *B-Municipality*, *I-Municipality*, *B-Postcode*, *I-Postcode*.

We generated data in multiple iterations as described below in Section 2.3. Our final training dataset consisted of more than 10^4 sentences/address examples. For data generation we used GPT-3.5-turbo API along with some prompt engineering. Since the data generation through this API is limited by the number of tokens—both generated as well as prompt tokens—we could not pass the list of all possible Slovak street names and municipality names within the prompt. Hence, data was generated with placeholders `streetname` and `municipalityname` only to be subsequently replaced by randomly chosen street and municipality names from the list of street and municipality names, respectively. A complete list of Slovak street and municipality names was obtained from the web pages of the Ministry of Interior of the Slovak republic [3].

With the use of OpenAI API generative algorithm we were able to achieve organic sentences without the need to manually generate the data, which sped up the process significantly. However, employing this approach did not come without downsides. Many mistakes were present in the generated dataset, mainly wrong annotations occurred and those had to be corrected manually. The generated dataset was split, so that 80% was used for model’s training, 15% for validation and 5% as synthetic test data, so that we could compare the performance of the model on real test data as well as on artificial test data.

2.2 Model Development and Training

Two general-purpose pre-trained models were utilized and compared: SlovakBERT [1] and a distilled version of this model [4]. Herein we refer to the distilled version as DistilSlovakBERT. SlovakBERT

is an open-source pretrained model on Slovak language using a Masked Language Modeling (MLM) objective. It was trained with a general Slovak web-based corpus, but it can be easily adapted to new domains to solve new tasks [1]. DistilSlovakBERT is a pre-trained model obtained from SlovakBERT model by a method called knowledge distillation, which significantly reduces the size of the model while retaining 97% of its language understanding capabilities.

We modified both models by adding a token classification layer, obtaining in both cases models suitable for NER tasks. The last classification layer consists of 9 neurons corresponding to 9 entity annotations: We have 4 address parts and each is represented by two annotations – beginning and inside of each entity, and one for the absence of any entity. The number of parameters for each model and its components are summarized in Table 2.

| Base model | Parameters | Classification head | Total |
|------------------|-------------|---------------------|-------------|
| SlovakBERT | 124,054,272 | 6,921 | 124,061,193 |
| DistilSlovakBERT | 81,527,040 | 6,921 | 81,533,961 |

Table 2: The number of parameters in our two NER models and their respective counts for the base model and the classification head.

Models’ training was highly susceptible to overfitting. To tackle this and further enhance the training process we used linear learning rate scheduler, weight decay strategies, and some other hyperparameter tuning strategies.

Computing resources of the HPC system Devana, operated by the Computing Centre, Centre of operations of the Slovak Academy of Sciences were leveraged for model training, specifically utilizing a GPU node with 1 NVidia A100 GPU. For a more convenient data analysis and debugging, an interactive environment using OpenOnDemand was employed, which allows researches remote web access to supercomputers.

The training process required only 10-20 epochs to converge for both model. Using the described HPC setting, one epoch’s training time was on average 20 seconds for 9492 samples in the training dataset for SlovakBERT and 12 seconds for DistilSlovakBERT. Inference on 69 samples takes 0.64 seconds for SlovakBERT and 0.37 seconds for DistilSlovakBERT, which demonstrates model’s efficiency in real-time NLP pipelines.

2.3 Iterative Improvements

Although only 69 instances of real data were present, the complexity of it was quite challenging to imitate in generated data. The generated dataset was created using several different prompts, resulting in 11,306 sentences that resembled human-generated content. The work consisted of a number of iterations. Each iteration can be split into the following steps: generate data, train a model, visualize obtained prediction errors on real and artificial test datasets, and analyze. This way we identified patterns that the model failed to recognize. Based on these insights we generated new data that followed these newly identified patterns. The patterns we devised in various iterations are presented in Table 3. With each newly expanded dataset both of our models were trained, with SlovakBERT’s accuracy always exceeding the one of DistilSlovakBERT’s. Therefore, we have decided to further utilize only SlovakBERT as a base model.

| Iteration | Prompt |
|-----------|--|
| 1. | Street + House Number + Municipality + Postal Code (+shuffling and omitting) |
| 2. | Municipality + Street + House Number + Postal Code (+omitting) |
| 3. | Municipality + House Number + Street + Postal Code (+omitting) |
| 4. | Municipality + House Number + Postal Code |
| 5. | Street + Municipality + House Number (verbal form) + Postal Code (+shuffling) |
| 6. | Municipality + House Number + Postal Code (Municipality mentioned twice; +shuffling) |
| 7. | All data duplicated to lowercase |

Table 3: The iterative improvements of data generation. Each prompt was used twice: First with and then without noise, i.e., natural human speech hesitations. Sometimes, if mentioned, prompt allowed to shuffle or omit some address parts.

3 Results

The confusion matrix corresponding to the results obtained using model trained in Iteration 1 (see Table 3)—is displayed in Table 4. This model was able to correctly recognize only 67.51% of entities in test dataset. Granular examination of errors revealed that training dataset does not represent the real-world sentences well enough and there is high need to generate more and better representative data. In Table 4 it is evident, that the most common error was identification of a municipality as a street. We noticed that this occurred when municipality name appeared before the street name in the address. As a result, this led to data generation with Iteration 2 and Iteration 3.

| | | Predicted | | | | | | | | |
|--------------|----------------|-----------|----------|----------|---------------|---------------|----------------|----------------|------------|------------|
| | | O | B-Street | I-Street | B-Housenumber | I-Housenumber | B-Municipality | I-Municipality | B-Postcode | I-Postcode |
| Ground truth | O | 53 | 6 | 10 | 1 | 1 | 2 | 0 | 0 | 0 |
| | B-Street | 1 | 30 | 21 | 0 | 0 | 0 | 0 | 0 | 0 |
| | I-Street | 0 | 1 | 10 | 0 | 0 | 0 | 0 | 0 | 0 |
| | B-Housenumber | 2 | 1 | 0 | 69 | 0 | 0 | 0 | 0 | 0 |
| | I-Housenumber | 0 | 0 | 0 | 1 | 18 | 0 | 0 | 0 | 0 |
| | B-Municipality | 6 | 37 | 3 | 0 | 0 | 25 | 0 | 0 | 0 |
| | I-Municipality | 1 | 0 | 9 | 0 | 0 | 0 | 8 | 0 | 0 |
| | B-Postcode | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |
| | I-Postcode | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

Table 4: Confusion matrix of model trained on dataset from the first iteration, reaching model’s predictive accuracy of 67.51%.

This process of detailed analysis of prediction errors and subsequent data generation accounts for most of the improvements in the accuracy of our model. The goal was to achieve more than 90% accuracy on test data. Model’s predictive accuracy kept increasing with systematic data generation. Eventually, the whole dataset was duplicated, with the duplicities being in uppercase/lowercase. (The utilized pre-trained model is case sensitive and some test instances contained street and municipality names in lowercase.) This made the model more robust to the form in which it receives input and led to final accuracy of 93.06%. Confusion matrix of the final model can be seen in Table 5.

| | | Predicted | | | | | | | | |
|--------------|----------------|-----------|----------|----------|---------------|---------------|----------------|----------------|------------|------------|
| | | O | B-Street | I-Street | B-Housenumber | I-Housenumber | B-Municipality | I-Municipality | B-Postcode | I-Postcode |
| Ground truth | O | 61 | 1 | 1 | 0 | 0 | 5 | 4 | 1 | 0 |
| | B-Street | 0 | 50 | 0 | 0 | 0 | 1 | 1 | 0 | 0 |
| | I-Street | 0 | 0 | 10 | 0 | 0 | 0 | 1 | 0 | 0 |
| | B-Housenumber | 0 | 0 | 0 | 72 | 0 | 0 | 0 | 0 | 0 |
| | I-Housenumber | 0 | 0 | 0 | 0 | 19 | 0 | 0 | 0 | 0 |
| | B-Municipality | 1 | 3 | 0 | 0 | 0 | 66 | 1 | 0 | 0 |
| | I-Municipality | 0 | 0 | 1 | 0 | 0 | 1 | 16 | 0 | 0 |
| | B-Postcode | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |
| | I-Postcode | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

Table 5: Confusion matrix of the final model with the predictive accuracy of 93.06%. Comparing the results to the results in Table 4, we can see that the accuracy increased by 25.55%.

There are still some errors; notably, tokens that should have been tagged as *outside* were occasionally misclassified as *municipality*. We have opted not to tackle this issue further, as it happens on words that may resemble subparts of our entity names, but, in reality, do not represent entities themselves. See an example below in Table 6.

| Sentence | Tokenized text | Tags | Predicted tags |
|-----------------|----------------|----------------|----------------|
| Kalša to Kal sa | Kalša | B-Municipality | B-Municipality |
| | to | O | O |
| | Kal | O | B-Municipality |
| | sa | O | O |
| Košice Hlavná 7 | Košice | B-Municipality | B-Municipality |
| | Hlavná | B-Street | B-Street |
| | 7 | B-Housenumber | B-Housenumber |
| | | | |

Table 6: Examples of the final model’s predictions for two test sentences. The first sentence contains one incorrectly classified token: the third token “Kal” with ground truth label *O* was predicted as *B-Municipality*. The misclassification of “Kal” as a municipality occurred due to its similarity to subwords found in “Kalša”, but ground truth labeling was based on context and authors’ judgment. The second sentence has all its tokens classified correctly.

4 Conclusions

In this technical report we trained a NER model built upon SlovakBERT pre-trained LLM model as the base. The model was trained and validated exclusively on artificially generated dataset. This well representative and high quality synthetic data was iteratively expanded. Together with hyperparameter fine-tuning this iterative approach allowed us to reach predictive accuracy on real dataset exceeding 90%. Since the real dataset contained a mere 69 instances, we decided to use it only for testing. Despite the limited amount of real data, our model exhibits promising performance. This approach emphasizes the potential of using exclusively synthetic dataset, especially in cases where the amount of real data is not sufficient for training.

This model can be utilized in real-world applications within NLP pipelines to extract and verify the correctness of addresses transcribed by speech-to-text mechanisms. In case a larger real-world dataset is available, we recommend to retrain the model and possibly also expand the synthetic dataset with more generated data, as the existing dataset might not represent potentially new occurring data patterns.

5 Acknowledgement

The research results were obtained with the support of the Slovak National competence centre for HPC, the EuroCC 2 project and Slovak National Supercomputing Centre under grant agreement 101101903-EuroCC 2-DIGITAL-EUROHPC-JU-2022-NCC-01.

References

- [1] Matús Pikuliak, Stefan Grivalsky, Martin Konopka, Miroslav Blsták, Martin Tamajka, Viktor Bachratý, Marián Simko, Pavol Balázik, Michal Trnka, and Filip Uhlárik. Slovakbert: Slovak masked language model. *CoRR*, abs/2109.15254, 2021.
- [2] Lance Ramshaw and Mitch Marcus. Text chunking using transformation-based learning. In *Third Workshop on Very Large Corpora*, 1995.
- [3] Ministerstvo vnútra Slovenskej republiky. Register adries. <https://data.gov.sk/dataset/register-adries-register-ulic>. Accessed: August 21, 2023.

- [4] Ivan Agarský. Hugging face model hub. <https://huggingface.co/crabz/distil-slovakbert>, 2022. Accessed: September 15, 2023.