



SpellMapper: A non-autoregressive neural spellchecker for ASR customization with candidate retrieval based on n-gram mappings

Alexandra Antonova, Evelina Bakhturina, Boris Ginsburg*

NVIDIA, USA

{aleksandraa, ebakhturina, bginsburg}@nvidia.com

Abstract

Contextual spelling correction models are an alternative to shallow fusion to improve automatic speech recognition (ASR) quality given user vocabulary. To deal with large user vocabularies most of these models include candidate retrieval mechanisms, usually based on minimum edit distance between fragments of ASR hypothesis and user phrases. However, the edit-distance approach is slow, non-trainable, and may have low recall as it relies only on common letters. We propose: 1) a novel algorithm for candidate retrieval, based on misspelled n-gram mappings, which gives up to 90% recall with just the top 10 candidates on Spoken Wikipedia; 2) a non-autoregressive neural model based on BERT architecture, where the initial transcript and ten candidates are combined into one input. The experiments on Spoken Wikipedia show 21.4% word error rate improvement compared to a baseline ASR system.

1. Introduction

Customization (contextual biasing) of automatic speech recognition (ASR) systems with user-defined phrases, such as proper names or terms, is required by many production systems.

Contextual¹ spelling correction (CSC) is a promising approach to ASR customization that relies on postprocessing of ASR output and does not require retraining of the acoustic model. Unlike traditional ASR spellchecking approaches such as [1, 2], which aim to correct known words using language models, the goal of contextual spelling correction is to correct highly specific user terms, most of which can be 1) out-of-vocabulary (OOV) words, 2) spelling variations (e.g., “John Koehn”, “Jon Cohen”) and language models cannot help much with that.

Our approach and task formulation are inspired by the non-autoregressive model proposed in [3]. The contextual spellchecking module takes as input a single ASR hypothesis and a user-defined vocabulary with up to several thousand words or phrases. The goal is to identify if a given ASR hypothesis contains any phrases from the user vocabulary, *possibly in a misspelled form*, and their exact positions in text. Next, the model corrects the ASR hypothesis by replacing the misspelled fragment with a matched phrase from the user vocabulary. For example, given an ASR-hypothesis “*came from schleiddorf in wurteenberg*” and user vocabulary including {“*schlaitdorf*”, “*wurtemberg*”}, the contextual spellchecking module should produce “*came from schlaitdorf in wurtemberg*”.

CSC models rely on attention mechanisms to compare ASR hypothesis and candidate phrases. Unfortunately, the attention

fails on large user vocabularies, and reducing the number of candidates is necessary. For example, in [3], the authors retrieve 100 candidates via minimum edit distance. They had to introduce a few custom blocks to their model architecture to support such a large number of candidates. The candidate retrieval approach based on edit distance suffers from two serious disadvantages. First, it is computationally inefficient due to quadratic complexity. Second, it heavily relies on common letters, which is not always enough, e.g., “*lookez*” vs. “*lucas*” have only one letter in common.

We propose a novel algorithm for candidate retrieval based on misspelled n-gram mappings to overcome both disadvantages of the edit distance approach. It preselects the top 10 candidates from the user vocabulary, potentially matching the ASR hypothesis. All the retrieved candidates are then concatenated with the ASR output and passed to a non-autoregressive spellchecker based on BERT architecture[4]. The proposed approach requires two offline preparation steps: a collection of n-gram mappings from a large corpus and an indexing of custom vocabulary for a new or updated user.

The contributions of the paper are the following:

1. A novel candidate retrieval algorithm based on misspelled n-gram mappings. It provides 74-90% recall with just the top 10 candidates.
2. A new way to formulate customization spellchecking as a non-autoregressive learning task for BERT architecture, using multiple separators (see Fig. 1).
3. A procedure to synthetically generate training datasets.

The proposed spellchecking approach demonstrates sustainable improvement (6.4-21.4% relative WER decrease) on three public datasets from different domains, for which we simulate “user vocabularies”. The code is publicly released as part of NeMo toolkit².

2. Background

Existing approaches to text-only contextual biasing of end-to-end (E2E) ASR systems can be divided into four groups. *TTS augmentation* is often used to fine-tune on-device ASR models with user text-only data [5]. However, this approach is impractical for ASR done in the cloud because it requires retraining of the ASR model. *Shallow fusion* [6, 7, 8, 9] is based on rescoring paths corresponding to custom phrases in the recognition graph. This method is prone to false positives and is not applicable if the needed path is not present in the graph, which is often the case with end-to-end ASR. *Deep contextualization* [10, 11] adds a custom vocabulary as an additional input to the acoustic model and tries to combine it with the audio input via a context

*Corresponding author

¹We use the term “context” for a user-defined list of phrases

²<https://github.com/NVIDIA/NeMo>

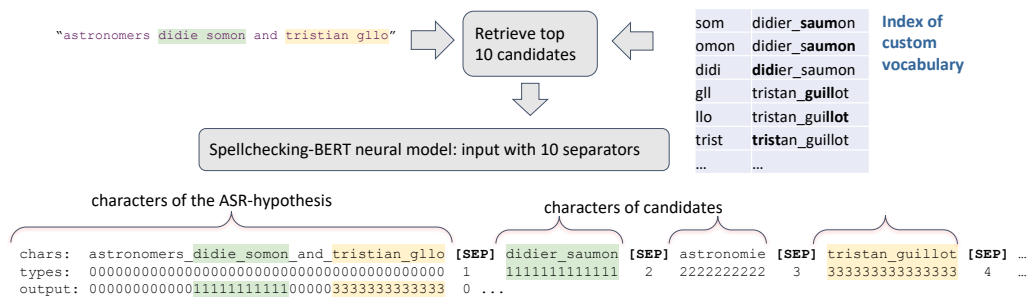


Figure 1: *Inference pipeline. Candidate retrieval ingests ASR output fragment of 10-15 words and selects top 10 user phrases with many matching n-grams in the user vocabulary index. ASR-hypothesis and candidates are combined to a single BERT input sequence with separators. The model predicts labels of correct candidates for each character of ASR-hypothesis, or 0.*

encoder and some kind of attention. This approach preserves the E2E model nature and does not require model fine-tuning for each user. The primary constraint of this attention-based approach is that attention cannot handle large vocabularies, and candidate retrieval is not possible as ASR hypotheses are not known beforehand. These drawbacks can be resolved with a more complex neural architecture or switching to a two-stage approach with candidate retrieval for the second phase using ASR output from the first phase [12, 13]. *Contextual spelling correction* requires a separate model that works on top of the ASR system [14, 3]. It detects and locates custom phrases that could be present in the ASR output text, possibly in a misrecognized form. This approach allows retrieval of a *subset* of custom phrases based on the known ASR hypothesis.

3. Proposed approach

Our approach to customization spellchecking consists of the following online modules: candidate retrieval, non-autoregressive neural model, and post-processing. The inference pipeline is shown in Fig. 1. It ingests ASR output fragments of 10-15 words and an n-gram-indexed user vocabulary. Then it retrieves the top 10 candidate phrases from the user vocabulary and tags the input characters with correct candidate labels or 0 if no match is found. The post-processing step combines outputs from different fragments, applies additional filtering, and outputs the final corrected ASR transcription. The proposed approach also requires two offline preparation steps: 1) collecting n-gram mappings from a large corpus and 2) vocabulary indexing for each user (Fig. 3).

Table 1: *N-gram mappings are pairs of n-grams and their common misspelled counterparts from the corpus automatically aligned on character level*

Source	Target	Frequency
l u c	l u c	7432
l u c	l u k+e	565
l u c	l o c	561
l u c	l u k	382
l u c	l o+o k	298
l u c	l u s	224
l u c	l u 	195

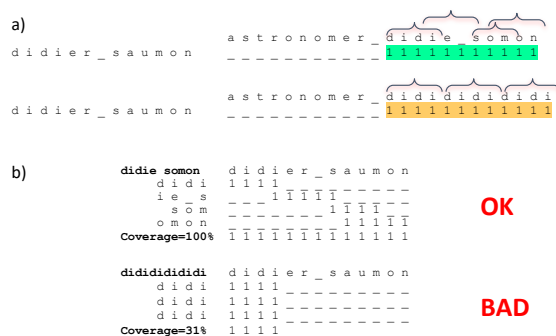


Figure 2: *a) Detection of potential candidate phrases in the ASR hypothesis using n-gram index. Note that repetitive or similar n-grams can produce wrong candidates. b) Additional filtering of found candidates by checking how well the phrase is covered by matched ngrams.*

3.1. N-gram mappings and indexing of user vocabulary

[15] proposed one of the first statistical models to learn probabilities of minor string-to-string edits from a corpus of paired misspelled and correct words. Following this, we collect n-gram mappings (see Table 1) from a parallel corpus of Wikipedia headings and their corrupted version. To get the corrupted version, we synthesize audio from vocabulary phrases and then pass through an ASR model (Conformer-CTC, see Table 3) to get text with potential errors. We align ground truth with corrupted phrases on a character level with GIZA++ [16]. Finally, we extract aligned n-grams from GIZA++ alignments along with their frequencies. We collect n-gram mappings from a large corpus only once.

Each time we have a new or updated user vocabulary we build an index where *key* is a letter n-gram and *value* is the whole phrase. The keys are n-grams in the given user phrase and their misspelled variants taken from our collection of n-gram mappings (see Index of custom vocabulary in Fig. 1).

3.2. Candidate retrieval algorithm

The candidate retrieval algorithm works online before the neural model inference. First, we search all overlapping n-grams in the index (each n-gram gives hits to corresponding symbols from ASR hypothesis) and select candidate phrases with the most hits. Next, we additionally filter out candidates by check-

ing how well matched n-grams cover the phrase itself (Fig. 2)

3.3. Neural model

The SpellMapper is based on BERT architecture [4]. The standard BERT input sequence consists of up to 512 tokens and can be segmented into subsequences of different types by inserting a special [SEP] token between input parts and by specifying additional input tensor *token_type_ids*. Though this is not common to split the input into more than two segments (e.g., first and second sentences to compare, or short paragraph and question for question-answering problem) and none of pre-trained models were trained on such tasks, for SpellMapper we use ten separators to combine the ASR hypothesis and ten candidate phrases into a single input (see Fig. 1)

The nature of spellchecking implies that we want to pay more attention to how fragments are similar letter-by-letter than the meaning of the words. To achieve this, we tokenize the input sequence into individual characters, replacing the space symbol with an underscore. The SpellMapper learns new character embeddings during training. We enhance our model with information from subword embeddings by passing the same input through standard BERT tokenization and concatenating subword embeddings to the embeddings of **each** character that belongs to this subword. While the model works on character level, adding subword embeddings may help the model better understand the context and reduce false positives by a quarter.

3.4. Metrics

We use word error rate (WER) to compare transcripts corrected with SpellMapper and the baseline ASR output. We also develop a procedure that imitates the “ideal spellchecker” to estimate maximum WER improvement with the given vocabulary, since the post-correction results highly depend on the vocabulary properties, e.g., size, minimum length, word frequencies, and their occasional similarity to frequent phrases. The “ideal spellchecker” first aligns reference and baseline ASR transcription to get nonidentical fragments. If a reference fragment exists in the custom vocabulary, we count it as successfully corrected. The resulting “ideal” WER gives us a lower bound on the best potential improvement with the given vocabulary.

Additionally, we calculate *diff-keyword-dependent* precision and recall following *keyword-dependent* metrics from [5]. We do not count occurrences where the correct user phrase was present in both the original and corrected ASR hypothesis (80-85% in our experiments), and only consider cases where was some change, counting *better*, *missed* and *false positive(fp)* cases. Recall (1) measures the model’s ability to identify user phrases disregarding that ASR can misrecognize them, e.g., “*time we of derby*” => “*tyne wear derby*”. Precision (2) is affected by the ratio of false positives - hallucinating user phrases in the unrelated text, e.g., *an artist’s right to fail* => *an artist’s reilly to fail*.

$$\text{Recall} = \frac{\text{better}}{\text{better} + \text{missed}} \quad (1)$$

$$\text{Precision} = \frac{\text{better}}{\text{better} + \text{fp}} \quad (2)$$

We also calculate *Top 10 recall* for the candidates produced by the retrieval algorithm. Its comparison with final recall reveals how many correct candidates were lost at the retrieval stage versus lost at the stage of neural model or post-processing.

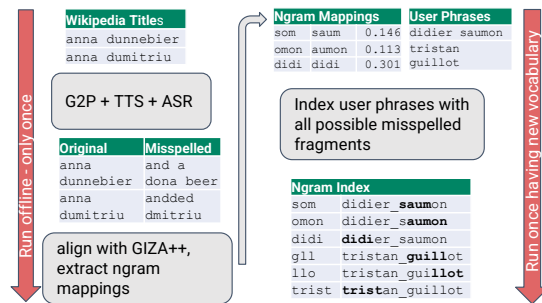


Figure 3: Data preparation pipeline. Left: N-gram mappings collection. Right: Indexing of user vocabulary.

4. Experiments

4.1. Datasets

4.1.1. Corpus for n-gram mappings

For the collection of n-gram mappings, we use a corpus of 4.5M Wikipedia headings from the Yago database [17]. We synthesize audio for each heading with TTS using FastPitch mel-spectrogram generator [18] and HiFi-GAN vocoder [19] (Table 3). Then we transcribe all generated audios with an ASR model (Conformer-CTC), thus getting a “parallel” corpus. The resulting corpus contains many examples of possible misspellings with WER of about 80%. Next, we run GIZA++ alignment and extract 1.9M n-gram mappings with a maximum length of five characters and conditional probability > 0.018.

4.1.2. Training and evaluation datasets

Each training example includes an ASR hypothesis, ten candidate phrases, and target labels. We create this dataset synthetically, reusing our collected corpus of correct and misspelled Wikipedia titles (Section 4.1.1). We extract all sentence fragments from full Wikipedia articles, where any Wikipedia title occurred in a real context. Next, we substitute the correct spelling with a misspelled one (to imitate the ASR hypothesis). We use the original spelling as the correct candidate and sample incorrect candidates from random, similar, or intersecting titles. We also add 50% of examples without any misspellings.

For the final testing, we prepare “customization” datasets from three public corpora which provide a grouping of utterances per user or document: Spoken Wikipedia [20], SPGISpeech validation part [21], UserLibri [22]. For Spoken Wikipedia (SWC) we perform sentence splitting, text normalization, and segmentation as described in [23]. For SPGISpeech we delete filler words, e.g., “uh” and “um”, from baseline transcriptions as they are absent in reference. For each “user”, we select capitalized multi-word phrases and rare words from their reference texts to serve as custom vocabulary. UserLibri does not retain information about the original case, so only single custom words are extracted.

4.2. Training and Testing procedures

We use a pre-trained BERT model with 67M parameters as a backbone model (Table 3). The hidden size is doubled because we concatenate subword embedding to each character embedding (Section 3.3). Our model was trained on 10M examples to get good spellchecking quality. See Fig. 4 for training dynamics for different corpus sizes. The model was trained for 5 epochs

Table 2: Spellchecking results w.r.t. 3 datasets and 2 baselines Conformer models (Table 3). Spellchecking gives 6.4-21.4% relative WER improvement and allows to regain 49-68% of maximum “ideal” WER improvement. The metrics are described in Sec. 3.4.

Dataset	Users	Utterances	Max/Avg vocab size	ASR Model	Baseline WER %	Spellcheck WER %	“Ideal” WER %	Recall %	Precision %	Top 10 recall %
SWC	1341	61370	1341/172	CTC	6.69	5.26	4.46	67	87.4	90.4
				RNNT	6.07	4.93	4.12	62.9	85.6	88.1
SPGI	1114	39341	86/21	CTC	5.88	5.50	5.33	55	88.7	84.5
				RNNT	5.71	5.32	5.12	53.2	88.5	82.7
UserLibri	99	5559	192/43	CTC	3.35	2.89	2.52	61.7	82.3	83.9
				RNNT	2.77	2.44	2.09	54.5	77.6	74.1

Table 3: Pre-trained models used in this work.

Model name	Source
<i>from https://huggingface.co:</i>	
BERT	huawei-noah/TinyBERT_General_6L_768D
<i>from https://catalog.ngc.nvidia.com/models:</i>	
Conformer-CTC	stt_en_conformer_ctc_large
Conformer-RNNT	stt_en_conformer_transducer_large
Mel-spectrogram	tts_en_fastpitch
Vocoder	tts_hifigan




Figure 4: SpellMapper accuracy on classes 1-10 depending on the number of training examples (1, 2, 10, 30M) evaluated on 20K validation examples. The whole phrase is considered correct if all symbols are tagged correctly. On class 0 (not presented on the graph) accuracy per symbol is 93%, 95.9%, 97.3%, 97.5% for corresponding corpus sizes

on 8 GPUs with a global batch of $32 * 8$. We used the AdamW [24] with an initial learning rate of $3e-5$, and the CosineAnnealing learning scheduler with a warmup ratio of 0.1.

We use two ASR models, Conformer-CTC and -RNNT (Table 3) to provide initial hypotheses. The testing procedure consists of the following steps:

1. Create a vocabulary for each “user” (Section 4.1.2).
2. Index each vocabulary using n-gram mappings (Section 3.1).
3. Transcribe all utterances with the baseline ASR model.
4. Loop through the transcriptions, split into fragments of 10-15 words in case of long sentences, and for each fragment retrieve the top 10 candidates from the corresponding index and perform inference with our neural model (Fig. 1)

5. Apply post-processing and generate corrected or unchanged transcriptions.

4.3. Results

Table 2 contains the results of the spellchecking model. We observe a substantial relative WER improvement (6.4-21.4%) across three datasets from different domains over both baseline ASR models. Note that *transducer* ASR model was not used in any stage of data preparation for the spellchecking model. Still, the post-correction gives a similar WER improvement, though recall and precision are slightly decreased (recall by 1.8-7.2%, precision by 0.2-4.7%), compared to the CTC model. We conclude that the SpellMapper generalizes well and can be applied without retraining on top of any ASR model for a given language. At the same time, retraining on the errors of specific ASR models is likely to improve performance further. We observe a 20-33% gap in recall between the candidate retrieval stage (top 10 recall) and the final recall. This means that our neural model and post-processing still lose a considerable amount of good candidates, so there is room for improvement in these stages. The precision metrics between 77-87% show that the model generates many false positives. This problem can be addressed in future research by sampling more training examples with false positives and adding trainable filters to the post-processing stage.³

5. Conclusion

This paper presents SpellMapper - a new approach for spellchecking customization of E2E ASR systems. We propose a new corpus-based algorithm for candidate retrieval, which does not rely on common letters, and provides 74-90% recall with only ten candidates. We then reformulate the spellchecking problem as a simple task for BERT architecture without a decoder or custom neural blocks. We evaluate algorithm robustness and generality on three public datasets from different domains and with different ASR models. In all cases, the model shows significant relative WER reduction (6.4-21.4%) and allows to regain 49-68% of maximum “ideal” WER improvement with the given user vocabularies. In future work, we would like to explore how to reduce false positives and increase the recall of correct candidates. We would also try to expand our approach to languages with highly inflectional morphology.

³We cannot directly compare our results with [3, 5] as their evaluation was done on non-public datasets and no code is available.

6. References

- [1] J. Guo, T. N. Sainath, and R. J. Weiss, "A spelling correction model for end-to-end speech recognition," in *ICASSP*, 2019.
- [2] O. Hrinchuk, M. Popova, and B. Ginsburg, "Correction of automatic speech recognition with Transformer sequence-to-sequence model," in *ICASSP*, 2020.
- [3] X. Wang, Y. Liu, J. Li, V. Miljanic, S. Zhao, and H. Khalil, "Towards contextual spelling correction for customization of end-to-end speech recognition systems," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 3089–3097, 2022.
- [4] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*, 2019.
- [5] K. C. Sim, F. Beaufays, A. Benard, D. Guliani, A. Kabel, N. Khare, T. Lucassen, P. Zadrzil, H. Zhang, L. Johnson *et al.*, "Personalization of end-to-end speech recognition on mobile devices for named entities," in *Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019.
- [6] D. Zhao, T. N. Sainath, D. Rybach, P. Rondon, D. Bhatia, B. Li, and R. Pang, "Shallow-fusion end-to-end contextual biasing," in *Interspeech*, 2019.
- [7] D. Le, G. Keren, J. Chan, J. Mahadeokar, C. Fuegen, and M. L. Seltzer, "Deep shallow fusion for RNN-T personalization," in *Spoken Language Technology Workshop (SLT)*, 2020.
- [8] A. Gourav, L. Liu, A. Gandhe, Y. Gu, G. Lan, X. Huang, S. Kalmene, G. Tiwari, D. Filimonov, A. Rastrow *et al.*, "Personalization strategies for end-to-end speech recognition systems," in *ICASSP*, 2021.
- [9] J. D. Fox and N. Delworth, "Improving contextual recognition of rare words with an alternate spelling prediction model," in *Interspeech*, 2022.
- [10] G. Pundak, T. N. Sainath, R. Prabhavalkar, A. Kannan, and D. Zhao, "Deep context: End-to-end contextual speech recognition," in *Spoken Language Technology Workshop (SLT)*, 2018.
- [11] M. Jain, G. Keren, J. Mahadeokar, G. Zweig, F. Metze, and Y. Saraf, "Contextual RNN-T for open domain ASR," in *Interspeech*, 2020.
- [12] Z. Yang, S. Sun, X. Wang, Y. Zhang, L. Ma, and L. Xie, "Two stage contextual word filtering for context bias in unified streaming and non-streaming transducer," *ArXiv:2301.06735*, 2023.
- [13] T. Munkhdalai, Z. Wu, G. Pundak, K. C. Sim, J. Li, P. Rondon, and T. N. Sainath, "NAM+: Towards scalable end-to-end contextual biasing for adaptive ASR," in *Spoken Language Technology Workshop (SLT)*, 2022.
- [14] X. Wang, Y. Liu, S. Zhao, and J. Li, "A light-weight contextual spelling correction model for customizing transducer-based speech recognition systems," in *Interspeech*, 2021.
- [15] E. Brill and R. C. Moore, "An improved error model for noisy channel spelling correction," in *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, 2000.
- [16] F. J. Och and H. Ney, "A systematic comparison of various statistical alignment models," *Computational linguistics*, vol. 29, no. 1, pp. 19–51, 2003.
- [17] T. Pellissier Tanon, G. Weikum, and F. Suchanek, "Yago 4: A reasonable knowledge base," in *The Semantic Web: 17th International Conference, ESWC 2020*, 2020.
- [18] A. Łańcucki, "FastPitch: Parallel text-to-speech with pitch prediction," in *ICASSP*, 2021.
- [19] J. Kong, J. Kim, and J. Bae, "Hifi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS*, 2020.
- [20] T. Baumann, A. Köhn, and F. Hennig, "The Spoken Wikipedia corpus collection: Harvesting, alignment and an application to hyperlistening," *Language Resources and Evaluation*, vol. 53, pp. 303–329, 2019.
- [21] P. K. O'Neill, V. Lavrukhin, S. Majumdar, V. Noroozi, Y. Zhang, O. Kuchaiev, J. Balam, Y. Dovzhenko, K. Freyberg, M. D. Shulman, B. Ginsburg, S. Watanabe, and G. Kucsko, "SPGISpeech: 5,000 hours of transcribed financial audio for fully formatted end-to-end speech recognition," in *Interspeech*, 2021.
- [22] T. Breiner, S. Ramaswamy, E. Variani, S. Garg, R. Mathews, K. C. Sim, K. Gupta, M. Chen, and L. McConnaughey, "UserLibri: A dataset for ASR personalization using only text," in *Interspeech*, 2022.
- [23] E. Bakhturina, V. Lavrukhin, and B. Ginsburg, "A toolbox for construction and analysis of speech datasets," in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung, Eds., vol. 1, 2021.
- [24] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.