

University of Groningen

Patching Leaks in the Charformer for Efficient Character-Level Generation

Edman, Lukas; Toral Ruiz, Antonio; van Noord, Gertjan

DOI:
[10.48550/arXiv.2205.14086](https://doi.org/10.48550/arXiv.2205.14086)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2022

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Edman, L., Toral Ruiz, A., & van Noord, G. (2022). *Patching Leaks in the Charformer for Efficient Character-Level Generation*. arXiv. <https://doi.org/10.48550/arXiv.2205.14086>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Patching Leaks in the Charformer for Efficient Character-Level Generation

Lukas Edman Antonio Toral Gertjan van Noord

Center for Language and Cognition
University of Groningen

{j.l.edman, a.toral.ruiz, g.j.m.van.noord}@rug.nl

Abstract

Character-based representations have important advantages over subword-based ones for morphologically rich languages. They come with increased robustness to noisy input and do not need a separate tokenization step. However, they also have a crucial disadvantage: they notably increase the length of text sequences. The GBST method from Charformer groups (aka downsamples) characters to solve this, but allows information to leak when applied to a Transformer decoder. We solve this information leak issue, thereby enabling character grouping in the decoder. We show that Charformer downsampling has no apparent benefits in NMT over previous downsampling methods in terms of translation quality, however it can be trained roughly 30% faster. Promising performance on English–Turkish translation indicate the potential of character-level models for morphologically-rich languages.

1 Introduction

Most state-of-the-art neural machine translation (NMT) systems operate on the subword level, typically using a preprocessing technique like Byte-Pair Encoding (BPE) (Sennrich et al., 2016) for combining characters into subwords. However, using a subword representation often masks important information given by characters, from the syntactic and morphological relatedness of words (e.g., “bake” and “bakes” may be assigned each their own unique token) to devices such as rhyme and alliteration. Although a byte-level pretrained language model, ByT5 (Xue et al., 2021), was more robust to misspellings than its subword counterpart, T5 (Raffel et al., 2019), results on common benchmarks such as GLUE did not reflect this inherent advantage. Similarly, in NMT, we have not yet seen advantages from character-level models reflected by their BLEU, chrF, or COMET scores (Libovický et al., 2021).

The major inhibitor for character-level models is inefficiency, which is due to the longer input and output sequences. For example, the average number of characters per subword for English is around 4 (Xue et al., 2021), so the input sequence to a character-level model is effectively 4 times longer. With a Transformer model, the problem is compounded by the complexity of self-attention. To address this, models such as the Charformer (Tay et al., 2021) introduce a downsampling method prior to the Transformer, which combines characters into pseudo-words, reducing the length of the sequence. The downsampling method used in the Charformer, GBST, was originally intended only for use in the encoder, and recent works attempting to apply the GBST layer to the decoder have failed (Libovický et al., 2021).¹ Our contributions include:

1. We show that there is an information leak in the GBST layer which breaks the typical NMT training scheme of a Transformer model.
2. We resolve the information leak issue, allowing it to be used in a Transformer decoder.
3. We provide a simple test to check for information leak in tasks which exhibit causality.
4. We show that despite Charformer’s current popularity, the GBST layer does not perform as well as earlier methods such as Lee et al. (2017) for NMT.
5. We give some evidence that character-level models are better for morphologically rich languages.

2 Patching Information Leaks

The Charformer modifies ByT5 with the addition of the gradient-based subword tokenization (GBST)

¹In consultation with the authors, we noted that their results using GBST in decoding in fact used a different method.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1-2	3-4	5-6	7-8	9-10	11-12	13-14	15-16	17-18	19-20										
1-3	4-6	7-9	10-12	13-15	16-18	19-20													
1-4	5-8	9-12	13-16	17-20															

Figure 1: Example of the mean mixing in the GBST layer, with alternating block colors assigned according to what information is accessible via mean mixing. Numbers indicate the range of positions accessible to each n-gram. The striped n-grams indicate a source of information leak.

layer. This layer mixes character representations (which are also informed of their relative positions via a convolution) using a simple mean computed over character n-grams up to length 4.² The model then selects a weighted average of these representations before a downsampling via a block-wise mean pooling. The result is a sequence of pseudo-word-level embeddings that is fed in directly to the Transformer (in contrast with the standard sequence of subword embeddings).

The GBST layer can be applied directly to a Transformer encoder without issue, but it cannot be applied to a Transformer decoder for generative tasks. This is due to an information leak, where information about characters in future blocks can end up in prior blocks. This occurs for 2 reasons: the convolution used to inform position, and the character n-gram means, both of which can overlap with the block-wise separations.

Concerning the convolution, each character is informed by its neighboring characters, which is necessary for the convolution to serve as a positional embedding. However, this creates the issue that characters on the right side of a block will be informed by characters to the left in a future block.

The character n-gram averaging similarly can obtain information from future blocks. As seen in Figure 1, with a downsampling factor of 4, the issue occurs with trigrams, where the 4th position is averaged with the 5th and 6th, despite being in separate blocks. Therefore, when the Transformer must predict the block containing the 5th and 6th characters, it has already received information about them, and thus can learn to simply copy the characters.

2.1 A Simple Test for Information Leaks

To confirm that there is a leak in the GBST layer, we set up a simple model consisting of a GBST layer, followed by an upsampling layer, which is a

²For decoding, we include n-grams up to the length of the downsampling factor.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1-2	3-4	5-6	7-8	9-10	11-12	13-14	15-16	17-18	19-20										
1-3	4	5-6	7-8	9-10	11-12	13-15	16	16-18	19-20										
1-4	5-8	9-12	13-16	17-20															

Figure 2: Our third proposed solution to preventing information leak.

linear layer that takes the downsampled block representation and upsamples it back to characters. We then train the model to predict a sequence of random characters, conditioned on the prior characters in the sequence, using a left-padding of BOS tokens equal to the downsampling factor. (e.g., '[BOS] [BOS] [BOS] a b c' is used to predict 'a b c d e f'). We expect that, if there is no information leak, the model will have a near-random accuracy, and conversely if the accuracy is significantly higher than random, there is an information leak.

In Table 1, we show the specific positions in a sequence affected by the leaks. Details of these experiments can be found in Appendix A. We use a vocabulary of size 100, so the accuracy of a perfectly random classifier should be around 1%, however we see in several cases that the models perform significantly better. In the worst case, with a downsampling factor of 4 and using the convolutional positional embeddings, 75% of the tokens are leaked.

To resolve this issue, we considered three potential solutions: adding more padding, removing leaking n-grams, or applying a mask to the n-grams.

Approach 1: Additional Padding. The simplest approach is to pad the sequences enough that leaking is impossible. Preliminary experiments indicated that padding with 2 times the downsampling factor prevented any leak. This approach however implies that the predictions are no longer conditioned on the block 1 step back, but rather the block 2 steps back.

Approach 2: Remove Overlapping N-Grams. Our second approach instead addresses the sources of information leak. First, the convolution used as a positional embedding is replaced with a static sinusoidal positional embedding, as in Vaswani et al. (2017). Second, the means for n-grams where there is overlap are removed. For example, for a downsampling factor of 4, the means for trigrams are removed. These changes prevent information leak, but it is possible that the removal of trigrams or other n-grams will substantially worsen the model's

δ	Pos Embs	1	2	3	4	5	6	7	8	9	10	11	12
2	Sin	0.0103	0.0109	0.0103	0.0122	0.0106	0.0113	0.0141	0.0122	0.0147	0.0116	0.0113	0.0075
	Conv	1.0000	0.0116	0.9941	0.0097	0.9947	0.0100	0.9959	0.0113	0.9966	0.0097	0.9962	0.0113
3	Sin	0.1356	0.0078	0.0109	0.0100	0.0078	0.0103	0.0241	0.0106	0.0088	0.0091	0.0103	0.0081
	Conv	1.0000	1.0000	0.0072	0.9987	0.9981	0.0106	0.9772	0.9619	0.0131	0.9966	0.9981	0.0081
4	Sin	0.1347	0.0338	0.0137	0.0113	0.0262	0.0106	0.0122	0.0100	0.0113	0.0097	0.0063	0.0072
	Conv	1.0000	0.9994	1.0000	0.0119	0.9775	0.9744	0.9781	0.0137	0.9966	0.9900	0.9950	0.0122

Table 1: Accuracies obtained by the simple model for each position in a target sequence of 12 random tokens, with varying downsampling factor (δ) and either convolutional or sinusoidal positional embeddings. The white boxes denote accuracies within random chance ($p > 1e - 3$) and thus not leaking, while red denotes accuracies that are significantly higher than random chance ($p < 1e - 10$), showing an information leak.

ability to learn contextual character embeddings before downsampling into blocks.

Approach 3: Apply Causal Mask. Our third approach builds on our second approach by keeping the sinusoidal positional embeddings and, rather than removing problematic n-grams, applies a causal mask when computing the mean. This can be seen in Figure 2, where overlapping n-grams are split by block, with the right side being informed by the left, but the left side being uninformed by the right. This approach allows trigram information to remain when it is not overlapping, however the information present in the blocks is not always consistent. This might be difficult for the model to discern.

3 Experimental Setup

We test our models on English translation to and from Arabic and German, as well as Turkish, whose vowel harmony and agglutination should benefit from character-level processing. Details of the datasets, preprocessing, and hyperparameters can be found in Appendix B.

We compare performance of our modified GBST methods to earlier works from Lee et al. (2017). Lee et al. (2017) found a performant method of character-level MT using a downsampling method based on causal convolutions, followed by a max-pooling layer. Their original MT model used an RNN, and Libovický et al. (2021) modernized the model to use a Transformer, keeping the convolution-based downsampling. It notably performed the best in Libovický et al. (2021)’s experiments on various downsampling methods. We similarly use the Transformer version in our experiments.

For decoding with Lee et al.’s and the Character downsampling methods, we apply the two-step decoder of Libovický et al. (2021), which adds an LSTM layer to the head of the Transformer, re-

ceiving its hidden states in addition to character embeddings of the so-far generated output sequence, and outputting the next characters. The parameters used for Lee et al.’s method follow Libovický et al. (2021).

For evaluation we use BLEU (Papineni et al., 2002) and COMET (Rei et al., 2020) with the reference-based wmt20-comet-da model. Our code is made freely available.³

4 Results

We compare the performance of our modified GBST layers in Table 2. First, we observe the GBST with additional padding, like the non-causal GBST, fails to learn any form of translation. This indicates that predicting two blocks into the future is too difficult a task for the model to make any meaningful progress in learning to translate.

δ	Non-Causal	Padding	Removal	Masking
2	0.00	0.04	25.49	24.59
3	0.00	1.35	24.90	22.16
4	0.00	0.62	22.07	17.47
2	0.00	0.80	22.17	21.90
3	0.00	0.52	20.93	19.07
4	0.00	0.57	20.25	18.04

Table 2: BLEU scores of vanilla GBST (Non-Causal) and our 3 approaches on German→English (top) and English→German (bottom). The same trend was observed for Arabic↔English, as seen in Appendix C.

We also see that the simpler approach of dropping overlapping n-grams works better than applying a causal mask. We suspect this is due to a lack of consistency within the masked n-gram representations, as some are simply duplicates of the lower order n-gram representations.

The downsampling factor also plays no role in the difference between the methods. Although the average length of a subword in English and German

³<https://github.com/Leukas/PatchingGBST>

is close to 4 characters, making a downsampling factor of 4 an intuitive choice, a higher downsampling factor leads to worse performance.

Method	δ	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg
Subword	1	27.2	24.1	25.6	11.2	14.7	11.3	19.0
Char	1	27.4	24.3	26.3	8.7	15.5	11.9	19.0
r-GBST	2	24.8	22.2	24.2	7.8	14.2	10.6	17.3
	4	22.1	20.2	22.2	8.3	11.3	8.2	15.4
Lee	2	27.1	23.3	25.1	9.9	15.2	11.1	18.6
	4	24.3	21.1	22.6	9.2	12.8	9.3	16.6

Method	δ	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg
Subword	1	0.20	0.07	0.06	0.12	-0.10	0.18	0.09
Char	1	0.22	0.04	0.14	-0.01	-0.03	0.19	0.09
r-GBST	2	0.11	-0.10	0.03	-0.11	-0.13	0.04	-0.03
	4	-0.12	-0.35	-0.14	-0.14	-0.35	-0.21	-0.22
Lee	2	0.22	-0.02	0.07	0.04	-0.05	0.18	0.07
	4	0.02	-0.24	-0.07	-0.04	-0.22	-0.00	-0.09

Table 3: BLEU (top) and COMET (bottom) scores. “Char” refers to a character-level model without any downsampling, with the same architecture and training scheme as “Subword”, differing only in tokenization. Best results in bold.

Despite the modified GBST with n-grams removed (henceforth r-GBST) being the best performing modification, Table 3 shows that it is still outperformed by the downsampling method from Lee et al. (2017). The mode of operation between Lee et al.’s method and the GBST method is similar: both achieve a mixture of characters, focusing only on neighboring characters to reduce computational complexity. While GBST achieves this with averaging across unigrams to 4-grams, Lee et al. uses convolutions with differing kernel sizes. The mixing via convolution is more uniform in nature; for example, the convolution with kernel size 3 is analogous to the trigram mixing, however the convolution does not have a hard boundary after every 3rd character. This may be the reason for Lee et al.’s method’s superior performance.

These results raise the question of the significance of the performance of the Charformer for any NLP task. If Lee et al.’s method was used in the same pretraining setup used in Tay et al. (2021), would we perhaps see superior performance?

Looking at the results with respect to language pairs, the results on English–Turkish seem most promising for character-level models. The character-level model with no downsampling outperforms the subword-level in both directions, Lee’s method with $\delta = 2$ is competitive, and r-GBST with $\delta = 2$ is somewhat competitive into English. This could indicate that models with ac-

cess to character-level information benefit when translating both to and from languages with rich morphology such as Turkish, although more extensive testing on other morphologically rich languages is needed.

Method	Decoding	δ	Epoch	Completion	Eval
Subword	Normal	1	5:34	3:52:31	1:33
	2-Step	1	5:47	5:15:57	1:40
Char	Normal	1	14:37	14:52:21	5:15
	2-Step	1	16:45	17:02:35	5:19
r-GBST	2-Step	2	10:38	10:15:02	3:19
		4	8:08	8:16:32	2:03
Lee	2-Step	2	13:09	13:22:48	5:45
		4	10:07	11:55:09	2:43

Table 4: Average time to train for 1 epoch and to completion, and to evaluate on the German→English dataset (hours + minutes + seconds). Times for other language pairs are proportionally similar.

Both downsampling methods perform worse than using no downsampling, but there is a trade-off in training time. In Table 4, we show the time it takes to train and evaluate on the test set post-training. We include character and subword-level models where we use the two-step decoding method, in order to separate the effect of downsampling from the decoding method.⁴ We can see that our downsampling methods are faster for both training and generation than the character model. However the subword model is still the fastest and achieves the best performance.

5 Conclusion

Character-level or byte-level models are intuitive for a variety of reasons but are slower due to much longer sequences. Downsampling methods such as the GBST layer in the Charformer looked promising, but is not usable for generative tasks without modification due to a leak of information flowing from future blocks. With modification, the GBST layer does not perform as well as older methods such as that of Lee et al. (2017), although it is faster in both training and evaluation.

Despite the intuitiveness of downsampling from characters to pseudo-words, we see a clear trade-off of performance versus time, with performance decreasing by several BLEU points, but the training and generation time being reduced to up to 50% of the non-downsampled model.

Although both downsampling methods tested do not reach the performance of the standard character-

⁴The performance of the 2-step character and subword models are similar to their normal counterparts.

level model, subword-level models show that shortening the sequence length can lead to appreciably faster models without any sacrifice in performance. Thus some form of downsampling is beneficial, and our method to find information leaks can serve as a useful debugging tool.

Finally, the results for English–Turkish show a positive effect of character-level information for translation of morphologically rich languages.

References

- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. Fully character-level neural machine translation without explicit segmentation. *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Jindřich Libovický, Helmut Schmid, and Alexander Fraser. 2021. Why don't people use character-level machine translation? *arXiv preprint arXiv:2110.08191*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational*

Linguistics (Volume 1: Long Papers), pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

Yi Tay, Vinh Q Tran, Sebastian Ruder, Jai Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. Charformer: Fast character transformers via gradient-based subword tokenization. *arXiv preprint arXiv:2106.12672*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2021. Byt5: Towards a token-free future with pre-trained byte-to-byte models. *arXiv preprint arXiv:2105.13626*.

A Test for Information Leaks

Our approach for testing for information leaks in a downsampling method trains for 5000 iterations in batches of 32. We use the Adam optimizer (Kingma and Ba, 2014), with a learning rate of 1e-4. These numbers were determined empirically, based on the degree of separation seen from the accuracies of the leaking versus non-leaking tokens. The accuracies in Table 1 are obtained over 100 batches, or 3200 samples.

B Main Experimental Details

B.1 Datasets

Our experiments use the IWSLT2017 data⁵ for English–German and English–Arabic, using the test sets from 2010 and 2015 for validation and testing, respectively. For English–Turkish, we use the SETIMES2 dataset⁶, and WMT’s newstest2017 and newstest2018 for validation and testing.⁷

B.2 Preprocessing

We remove sentence pairs where the English sentence is longer than 256 characters from the training data. To keep our vocabulary size consistent across all languages, we tokenize according to

⁵<https://sites.google.com/site/iwsltevaluation2017/TED-tasks>

⁶<https://opus.nlpl.eu/SETIMES2.php>

⁷<http://data.statmt.org/wmt18/translation-task/dev.tgz>

UTF-8 bytes rather than using a character vocabulary.⁸ We also include subword-level models, with vocabularies generated with SentencePiece (Kudo and Richardson, 2018). We chose a vocabulary size of 16 thousand for English–German and English–Arabic, as that roughly corresponds to a downsampling factor of 4. For English–Turkish, we used a vocabulary size of 8 thousand.⁹

B.3 Hyperparameters

We use the Transformer model with the same parameters as Transformer Base (Vaswani et al., 2017). We train our models using AdamW (Loshchilov and Hutter, 2017) with a learning rate of $2e-4$, a linear warmup of 4000 steps, a batch size of 128, and label smoothing factor of 0.1. The learning rate was chosen from a grid search, the batch size chosen empirically, and the warmup steps and label smoothing factor were based on Libovický et al. (2021). We use an early stopping criterion of no improvement on the validation set with a patience of 10. All of our models are trained on a single Nvidia V100 (32GB) GPU.

C Comparison of GBST modifications on Arabic↔English.

In addition to German↔English (Table 2), we compare the GBST modifications on Arabic↔English, shown in Table 5. The results are more promising for the masking method compared to those on English↔German, however they are still outperformed by the removal method.

δ	Non-Causal	Padding	Removal	Masking
2	0.00	0.80	24.21	24.02
3	0.00	0.12	23.70	22.50
4	0.00	0.02	22.24	21.26
2	0.00	0.04	7.81	7.62
3	0.00	0.30	7.77	7.21
4	0.00	0.23	8.32	7.73

Table 5: Arabic→English (top) and English→Arabic (bottom).

⁸Operating on the byte-level also follows Tay et al. (2021), despite the name “Charformer” perhaps suggesting otherwise.

⁹We also tested a vocabulary size of 16 thousand, but this performed worse.