

Nearest Neighbor Language Models for Stylistic Controllable Generation

Severino Trotta [†] and Lucie Flek ^{†‡} and Charles Welch ^{†‡}

[†] Conversational AI and Social Analytics (CAISA) Lab

Department of Mathematics and Computer Science, University of Marburg

[‡] The Hessian Center for Artificial Intelligence (Hessian.AI)

severino.trotta@gmail.com, {lucie.flek, welchc}@uni-marburg.de

Abstract

Recent language modeling performance has been greatly improved by the use of external memory. This memory encodes the context so that similar contexts can be recalled during decoding. This similarity depends on how the model learns to encode context, which can be altered to include other attributes, such as style. We construct and evaluate an architecture for this purpose, using corpora annotated for politeness, formality, and toxicity. Through extensive experiments and human evaluation we demonstrate the potential of our method to generate text while controlling style. We find that style-specific datastores improve generation performance, though results vary greatly across styles, and the effect of pretraining data and specific styles should be explored in future work.

1 Introduction

Language models with external memory, like [Khandelwal et al. \(2020b\)](#)'s recent k -nearest neighbour language model (kNN-LM), have demonstrated impressive predictive performance. Great reductions in perplexity are achieved through storing the encoding of contexts from the training data. A sequence of tokens is encoded by the model and stored as a key, which is paired with a value representing the next token in the sequence. During decoding, similar contexts are recalled based on their key similarity, and values are interpolated with the base language model's predictions.

In this work, we augment the encoding with stylistic attributes, such that the keys are more heavily influenced by the style of the encoded text. By explicitly encoding the style, the similarity is more strongly affected by the stylistic attributes than previous models. When decoding, we can then provide a style (e.g. polite or formal) and the most similar contexts are both relevant in content and more likely to conform to the provided style. The

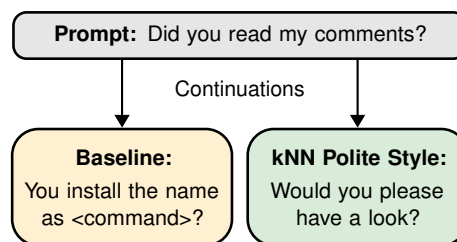


Figure 1: Example prompt continuations of the baseline kNN language model and our model, which continues generation in a specified style (e.g. polite). This example is based on a real example from our human evaluation but shortened for brevity and clarity. Full examples are provided in [Appendix A](#).

example in [Figure 1](#) shows a prompt and two continuations, one with the baseline kNN language model and one with our model given a polite style value as input, signaling that it should continue the prompt in a polite style.

Through our architecture implementation, we show that we not only improve language modeling performance over previous models, but that we can control generation to produce text of a particular style. We provide human evaluation of our stylistic outputs and an analysis of the performance of our approach and modeling decisions that affect how style attributes are represented in memory. To the best of our knowledge, this is the first work to modify a language model's external memory in order to control generated style.

2 Related Work

Recent approaches to controllable generation have included fine-tuning large models, such as [Keskar et al. \(2019\)](#), who condition on 50 control codes during training, which represent different styles, topics, and languages. Other approaches avoid retraining by modifying the predictions only at decoding time. The FUDGE model predicts for a sequence, the likelihood that possible generation steps will result in a sequence that satisfies

a given constraint (Yang and Klein, 2021). The DExperts model alters the probability distribution of a language model (LM) based on the predictions of other LMs that are fine-tuned on specific desired or undesired attributes (Liu et al., 2021). Dathathri et al. (2020) and others similarly modify gradients directly during prediction.

Khandelwal et al.’s k -nearest neighbour language model is based on the groundwork laid by previous work that augmented language models with a cache memory of recent observations. Grave et al. (2017b) captured local context of up to a few thousands tokens to improve predictions. Grave et al. (2017a) expanded upon this idea by storing all past hidden activations in a memory. Khandelwal et al. then replaced the recurrent network with a transformer to better model long-term dependencies. Aspects of kNN-LMs have been improved upon, in terms of performance, efficiency, and additional functionality.

Even though kNN-LMs achieve state-of-the-art predictive performance, the retrieval operation is very computationally expensive. He et al. (2021) and Alon et al. (2022) explored techniques to improve inference speed, such as compressing embeddings, or training an additional model to dynamically disable retrieval for predictions where the datastore is unlikely to improve the result. Wu et al. (2022) implemented a similar model and focused on improving scalability. Yogatama et al. (2021) extended the model with a gating mechanism that learns to combine short and long term memory with local context. Xu et al. (2021) improved performance by leveraging structural locality features such as topic clusters in text or project hierarchies in source code repositories. Khandelwal et al. (2020b) also extended their model for use in machine translation (Khandelwal et al., 2020a), which has also received efficiency improvements (Meng et al., 2021; Wang et al., 2021).

3 Methodology

The main goal of our work is to expand kNN-LM functionality and we build off of Khandelwal et al. (2020b), which we will refer to as the baseline architecture. We modified this architecture to accept additional style attributes as input, and concatenate these to the input text encoding. This has the effect of modifying the embedding space such that it encodes both semantic and stylistic properties (see Appendix B). We will refer to this as the style

architecture.

After the input style and context are encoded, they are stored in the datastore. We experimented with separating datastores based on the distribution of style values in the dataset. For this part we take subsets with specific style values (e.g. only toxic or only polite) from the datasets and construct datastores containing only data from those subsets. We refer to this as *separate datastores*. Datastores containing examples of different style are referred to as *mixed datastores*.

4 Datasets

We use 4 datasets, 3 of which contain style attributes. Unless stated otherwise, we use the default splits provided by the original work.

Wikitext-103 Wikitext-103 (WT103) is a collection of *Good* and *Featured* Wikipedia articles (Merity et al., 2017). It is provided in a tokenized form, with case, punctuation and numbers, totalling 103 B tokens. For better compatibility with our tokenization of other datasets we modified WT103’s tokenization by replacing all occurrences of `*n 't` (e.g. in “can’t”) with `* n't`.

Politeness The Stanford Politeness Corpus (SPC) consists of 11k utterances from StackExchange and Wikipedia Talk pages, annotated with politeness scores, which we use as style attributes (Danescu-Niculescu-Mizil et al., 2013). We created an 85-7.5-7.5 split for training, validation, and test.

Formality Grammarly’s Yahoo! Answers Formality Corpus (GYAFC) is the largest available corpus for formality style transfer, containing about 110K formal/informal sentence pairs (Rao and Tetreault, 2018). It is divided into the domains *Entertainment & Music* and *Family & Relationships*, which makes it suitable for training and evaluation with in/out-of-domain data. We do not use the parallel nature of the corpus, but assign each sentence a style attribute (−1 for informal- and 1 for formal sentences) and re-split the training subset to obtain an 80-10-10 train/validation/test split.

Toxicity For toxicity we use the Jigsaw Unintended Bias in Toxicity Classification dataset, which contains comments from the Civil Comments platform annotated with several binary toxicity labels representing types of toxicity (e.g. insult, identity attack, sexually explicit) (Borkan et al.,

2019). We use only the *toxic* label as a style attribute. We further use the Real Toxicity Prompts dataset for human evaluation (Gehman et al., 2020). This dataset contains the beginning of sentences and has been used to test if models can continue generation without toxicity.

4.1 Preprocessing

We largely follow Merity et al. (2017)’s preprocessing of WT103 for all data. Additionally, we perform the following replacements:

what	replacement
inline/block code	<code>
usernames	<person>
hyperlinks	[title](url) \mapsto title
URLs	<url>

Unlike Merity et al. (2017) we use spaCy.io for normalization and tokenization, but perform the same tokenization of infix punctuation and symbols. This serves to differentiate number separator punctuation from word punctuation, and hyphens from minus signs or ranges.

5 Experiments

In preliminary experiments we examined the effect of float precision on model performance. Using WT103 and the setup of Khandelwal et al. (2020b) we found that using half precision halves inference time without hurting perplexity, and reduces inference time for the kNN-LM with a negligible 0.25 increase in perplexity. In the following sections we use half-precision to speed up the experimentation process.

In the main experiments we first examine the impact on perplexity when incorporating style values into the model. Then we compare our method to the baseline through human evaluation.

5.1 Language Modeling with Style Attributes

We train the style model, S , using the modified architecture, and baseline model, B , which uses the original architecture. To achieve high predictive performance on text, we first train both on WT103.

For each dataset we then fine-tune a copy of S and B on the dataset. Using these fine-tuned models we generate and evaluate a datastore for each a model, using the dataset the respective model was fine-tuned on.

We additionally build and evaluate a datastore on domain and style subsets, to test our model’s

performance on out-of-domain data and across subsets with specific style. The subsets are listed in Table 3 in the Appendix.

5.2 Setup & parameters

Unless stated otherwise, we use the default parameters used by the transformer_lm_wiki103 architecture and Khandelwal et al. (2020b).

Vocabulary To avoid a high number of OoV tokens, we chose to use a shared vocabulary from the union of all datasets. Tokens occurring less than 3 times were mapped to <unk>. The resulting vocabulary has a size of 375k tokens. Since Merity et al. (2017) have shown that a large vocabulary with adaptive input representation can outperform a smaller BPE vocabulary, we chose the former over the latter – although this choice has its own limitations (see limitations section).

Training During pretraining on WT103, we use random style values normally distributed around the median, set patience to 5 epochs, and the style embedding dimension to 96. During fine-tuning we use Adam instead of the NAG optimizer and set patience to 10 epochs.

Datastore We use half-precision vectors in the datastore. Since the context embedding dimension (1,120) must be divisible by the number of FAISS index subquantizers, we use 70 instead of 64. For smaller subsets of data we use 2,048 cluster centroids, rather than the default 4,096.

5.3 Fine-tuning Results

The results of fine-tuning for each style attribute are shown in Table 1. We find that simply encoding the style value improves model performance on all datasets, including the original WT103, though this is likely due to the small increase in the number of parameters. Fine-tuning predictably lowers perplexity on the style datasets and slightly increases the WT103 perplexity as the model shifts away from the corpus it was originally trained on. The best performance is on formality, followed by politeness, which we expect to more closely resemble WT103. The addition of the style input allows for much greater improvement on politeness as compared to toxicity which shows near equal performance without it. We also provide perplexity for the kNN-LM in Appendix C.

Dataset	Baseline		Style	
	PT	FT	PT	FT
Politeness	218	126	164	78
Formality	161	77	148	60
Toxicity	212	125	186	93
WT103	31	35-64	29	32-59

Table 1: Perplexity before (using the pretrained model; PT) and after fine-tuning (FT). All models were evaluated on the FT dataset and on WT103. The value ranges in the WT103 row indicate the performance range of the FT models on the WT103 dataset.

5.4 Human Evaluation

We aimed to answer three questions; (1) do the style-specific datastores outperform the mixed datastore, (2) does the kNN-LM outperform the LM, and (3) does the style architecture outperform the baseline? To address these questions we asked a group of 11 students to annotate model outputs.

Generating Output We follow the idea of Gehman et al. (2020) and generate outputs by supplying prompts of different styles to the models. We use both non-toxic (toxicity = 0) and neutral ($0 < \text{toxicity} < 0.5$) and created prompts from the formality and politeness datasets by cutting off the second half of randomly sampled sentences. The prompts are then used as input to the models, which generate continuations to the prompts. All combinations of models and inputs are detailed in Table 4 in the Appendix. For all kNN-LM outputs we use $\lambda = 0.8$ as interpolation parameter.

Survey Setup We asked annotators to select which of a pair of prompt continuations was more fluent and which more closely followed one of the given styles. The pair combinations are based on the three comparisons listed at the beginning of this section, but are fully listed in Tables 5-7 in the Appendix and result in a total of 440 survey questions. The questions were presented to annotators in random subsets of 20% of the full set. Each output pair was rated by 2-4 people.

Results The results in Table 2 show which model is preferred, in percentage of annotators, in terms of fluency and style for each pair. We find that when comparing mixed and specific datastores, the specific datastores are preferred for style and even more strongly for fluency. While the kNN-LM is preferred over the LM in fluency, the style pref-

		Fluency (%)	Style (%)
		Datastore	Mixed Specific
Model type	LM kNN-LM	47.7 52.3	50.7 49.3
Architecture	Baseline Style	48.8 51.2	47.3 52.7

Table 2: Human evaluation preferences for model pairs. Column-wise percentage pairs sum to 100.0.

erence is more evenly split. When comparing the style architecture to the baseline, we find that the ours is preferred, with style more strongly preferred to fluency.

These results are an aggregation over the styles, however the performance on specific styles reveals more varied results. The specific datastores give more style control for politeness than for other styles and for some combinations of the prompt style and target style, the mixed datastore was preferred. We see that when we want to generate non-toxic, polite, or formal text; those that more closely resemble the pretraining data style, the preference leans more toward the mixed datastore.

When comparing the LM to the kNN-LM, we found that the LM style was often preferred when provided an informal, non-toxic, or impolite prompt, regardless of the target style. We also found that the fluency of the LM is preferred when generating polite or impolite text. Lastly, the style architecture is not always preferred over the baseline either. The baseline shows stronger fluency for toxic and impolite prompts. The style architecture has the best style control when generating formal, toxic, and impolite text. Overall, there appears to be a trade-off between style-control and fluency. The full breakdown by prompt and input type is shown in Figures 4-6 in the appendix.

6 Conclusion

We examined the use of kNN language models for controllable stylistic generation using politeness, formality, and toxicity as target styles. Our findings show that simply encoding style in the architecture improved perplexity of the language model. A human evaluation further showed that specific datastores for target styles outperform the standard mixed datastore, and that our model generally outperformed the baseline kNN model in terms of fluency and style control, though results

on specific styles varied. Future work is needed to fully understand the effect of pretraining and benefits of the model variants for specific styles and should also consider comparisons to other controllable generation models, such as Keskar et al. (2019).

Our code is available on Github at <https://github.com/d8xa/style-knnlm>.

Limitations

Vocabulary Choice We chose a shared vocabulary to reuse the same baseline model for fine-tuning on multiple datasets. Since less frequent tokens are assigned less parameters by the adaptive input representation, this could lead to under-representation of rare, style-specific tokens in general, and worse fine-tuning results for smaller datasets or datasets with many rare tokens. The same problem applies for single-dataset vocabularies as well, when rare tokens are more prevalent for a particular style. Byte-pair encoding would avoid these problems, but make comparability to the vanilla kNN-LM more difficult.

Sequence Length The original kNN-LM was trained with sequences of up to 3,072 tokens in length, which helps model long-term dependencies in the WT103 dataset. Since all of our datasets with style attributes contain much shorter sequences, single-dataset training with shorter input sizes might be better suited and achieve better performance than pre-training on WT103 and fine-tuning on the style dataset.

Comparability with Khandelwal et al. (2020b) When training our style architecture, we had to choose between a combined embedding dimension of $C + S_{\text{emb}} = 1,024$ (token- and style embedding dimensions C and S_{emb}), or to use $C = 1,024$. In any case the resulting language model would have a different number of parameters than in the original kNN-LM. We chose to use $C = 1,024$ and $S_{\text{emb}} = 96$. FAISS requires the vector dimension to be divisible by the number of subquantizers. Since our combined embedding dimension is different from 1,024, we had to choose 70 instead of 64 subquantizers.

Another difference is the choice of vocabulary. The WT103-only vocabulary would make results more comparable, but also lead to a high number of UNK tokens for the style datasets, and therefore reduce performance greatly.

Token to Style Embedding Dimension Ratio

To limit the scope of this work we did not perform an analysis on the ratio between token- and style embedding dimension. Other ratios might achieve better fluency or style control.

Choice of Interpolation Parameter For our human evaluation we put $\lambda = 0.8$ weight on kNN-LM probabilities. A lower λ , more close to the vanilla kNN-LM, might achieve better fluency.

Ethics Statement

Work on controllable generation allows models to generate text in styles such as those presented here, which include polite/impolite, formal/informal, and toxic/non-toxic. There may be applications where each style is appropriate or desirable, but some styles of text such as impolite and toxic, may be undesired or even harmful. Application of our models should involve careful consideration of desired styles and the context in which they are deployed.

Acknowledgements

This work has been supported by the Alexander von Humboldt Foundation, and by Hessian.AI. Any opinions, findings, conclusions, or recommendations in this material are those of the authors and do not necessarily reflect the views of the Alexander von Humboldt Foundation, or Hessian.AI.

References

- Uri Alon, Frank F. Xu, Junxian He, Sudipta Sen-gupta, Dan Roth, and Graham Neubig. 2022. Neuro-symbolic language modeling with automaton-augmented retrieval. *CoRR*, abs/2201.12431.
- Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. *CoRR*, abs/1903.04561.
- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. [A computational approach to politeness with application to social factors](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.

- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. [RealToxicityPrompts: Evaluating neural toxic degeneration in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Online.
- Edouard Grave, Moustapha Cissé, and Armand Joulin. 2017a. [Unbounded cache model for online language modeling with open vocabulary](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. 2017b. [Improving neural language models with a continuous cache](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Junxian He, Graham Neubig, and Taylor Berg-Kirkpatrick. 2021. [Efficient nearest neighbor language models](#). *CoRR*, abs/2109.04212.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. [CTRL: A conditional transformer language model for controllable generation](#). *arXiv preprint arXiv:1909.05858*.
- Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020a. [Nearest neighbor machine translation](#). *CoRR*, abs/2010.00710.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020b. [Generalization through memorization: Nearest neighbor language models](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021. [DExperts: Decoding-time controlled text generation with experts and anti-experts](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Online.
- Yuxian Meng, Xiaoya Li, Xiayu Zheng, Fei Wu, Xiaofei Sun, Tianwei Zhang, and Jiwei Li. 2021. [Fast nearest neighbor machine translation](#). *CoRR*, abs/2105.14528.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Sudha Rao and Joel R. Tetreault. 2018. [Dear sir or madam, may i introduce the GYAFC corpus: Corpus, benchmarks and metrics for formality style transfer](#). *CoRR*, abs/1803.06535.
- Shuhe Wang, Jiwei Li, Yuxian Meng, Rongbin Ouyang, Guoyin Wang, Xiaoya Li, Tianwei Zhang, and Shi Zong. 2021. [Faster nearest neighbor machine translation](#). *CoRR*, abs/2112.08152.
- Yuhuai Wu, Markus N Rabe, DeLesley Hutchins, and Christian Szegedy. 2022. [Memorizing transformers](#). In *10th International Conference on Learning Representations, ICLR 2022*.
- Frank F. Xu, Junxian He, Graham Neubig, and Vincent J. Hellendoorn. 2021. [Capturing structural locality in non-parametric language models](#). *CoRR*, abs/2110.02870.
- Kevin Yang and Dan Klein. 2021. [FUDGE: Controlled text generation with future discriminators](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Online.
- Dani Yogatama, Cyprien de Masson d’Autume, and Lingpeng Kong. 2021. [Adaptive semiparametric language models](#). *Transactions of the Association for Computational Linguistics*, 9.

A Generation Examples

Here we include three examples of generation comparing the baseline LM to the kNN-LM where human annotators found our model to be both more fluent and more closely aligned with the desired target style.

B Modified Architecture

Our architecture, modified from [Khandelwal et al. \(2020b\)](#), is shown in [Figure 2](#). The left side of the figure shows the original model and the right shows our modification which incorporates style attributes.

C Fine-tuning Experiment

The fine-tuning experiments in the main paper summarize the performance of our models pre-trained on WT103 and fine-tuned on one of the style datasets each. Here we also include the performance on other subsets of the finetuning dataset, such as different toxicity levels for the toxicity data, and domain subsets for the politeness and formality data. A full list of subsets is given in [Table 3](#).

D Human Evaluation

For the human evaluation task, we generated sentences for different combinations of the model architecture, target style, and prompt style. The high-level summary of combinations is presented in [Table 4](#). For the three test conditions in [§5.4](#), we have

Target Style: Polite	
Prompt:	<person>: Yes. Did you read my comments?
Baseline:	Also, you install the name as <unk> (") <url>??
Our Model:	Vote under the changed circumstances? Would you please have a look?
Target Style: Formal	
Prompt:	ok let me answer u
Baseline:	'm not in pain, and I say ... graham understand your problem
Our Model:	are alike. You should wait a little longer to enter your full name and play your live show.
Target Style: Formal	
Prompt:	THis is one way to keep it out of your system ...
Baseline:	sings u love bio <unk>
Our Model:	with the following words coming out of Gardisil ... Peter Pan, or just those are some favorite books that I am able to mail.

listed the model and style combinations that we tested in [Tables 5-7](#). Finally, the full breakdown of the human evaluation for fluency and style preferences are shown in [Figures 4-6](#). The heatmaps show the tendency of annotator choices, where tendency is the mean model choice per question (−1 and 1 encode the choices), aggregated across all questions in the survey and normalized to −100 and 100.

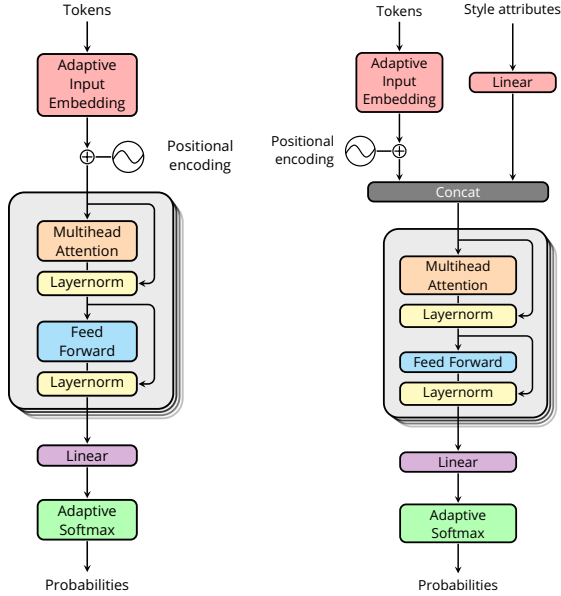


Figure 2: Changes made to the LM architecture (left: unmodified, right: our version).

dataset	subset	size		comment
		samples (%)	samples	
WT103	all	100	1.8 M	
Formality	all	100	214 k	
	formal	49.8	106.7 k	
	informal	50.2	107.3 k	
	family & relationships	49.7	106.3 k	
	entertainment & music	50.3	107.7 k	
Toxicity	all		1.4 M	Only used for evaluation; due to imbalance not suited for training or Datastores (DSes).
	non-toxic	75.6	1 M	Toxicity score = 0 .
	non-toxic-sample	11.8	159.8 k	Sample from <i>non-toxic</i> .
	toxic-gte-0.5	11.8	159.8 k	Toxicity score ≥ 0.5 .
	toxic-gte-0.8	2.5	34.1 k	Toxicity score ≥ 0.8 .
	toxic-gte-0.9	0.8	10.2 k	Toxicity score ≥ 0.9 .
	all-sample	23.5	319.6 k	Combination of <i>toxic-gte-0.5</i> and <i>non-toxic-sample</i> .
Politeness	all	100	11.1 k	
	neutral	30.3	3.4 k	Center 30% of politeness scores.
	polite	36.9	4.1 k	Upper 36.9% of politeness scores.
	impolite	32.8	3.6 k	Lower 32.8% of politeness scores.
	stackexchange	60.8	6.8 k	
	wikipedia	39.2	4.4 k	

Table 3: List of dataset subsets. Note: Proportions of subsets within splits are subject to variations due to random sampling. Not all subsets are presented in the results of the main paper. Some subsets are only shown in Figure 3.

kNN-LM evaluation by architecture and dataset

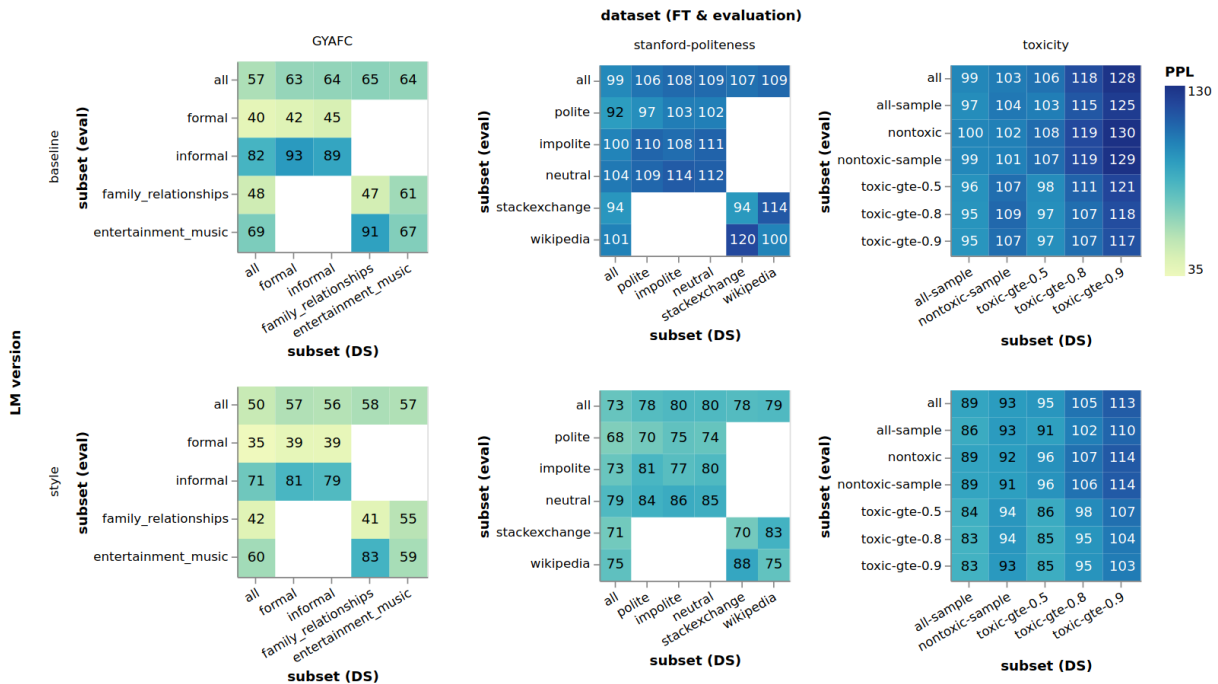


Figure 3: Overview of test perplexity in the fine-tuning experiment across data subsets listed in Table 3.

dataset	source style	target style	LM	datastore styles
Formality	formal	formal	style	none, formal, mixed
		informal	style	none, informal, mixed
	informal	n.a.	baseline	none, informal, formal
		formal	style	none, formal, mixed
Toxicity	neutral	informal	style	none, informal, mixed
		n.a.	baseline	none, formal, informal
	non-toxic	non-toxic	baseline	none, non-toxic, toxic, mixed
		toxic	style	none, non-toxic, mixed
Politeness	impolite	n.a.	baseline	none, non-toxic, toxic, mixed
		polite	style	none, toxic, mixed
	polite	non-toxic	baseline	none, non-toxic, toxic, mixed
		toxic	style	non-toxic, none, mixed
Formality	formal	formal	style	none, impolite, mixed
		informal	baseline	none, polite, impolite, mixed
	informal	n.a.	style	none, polite, mixed
		polite	style	none, impolite, mixed
Toxicity	neutral	impolite	baseline	none, polite, impolite, mixed
		n.a.	style	none, polite, mixed
	non-toxic	polite	baseline	none, polite, impolite, mixed
		impolite	style	none, polite, mixed

Table 4: Combinations of models and inputs used for generating the outputs for human evaluation. *n.a.* in the *target style* column refers to the baseline LM architecture, since it has no style input.

dataset	source style	target style	specific datastore style
Formality	formal	formal	formal
	informal	informal	informal
Toxicity	neutral	non-toxic	non-toxic
	non-toxic	toxic	toxic
Politeness	impolite	non-toxic	non-toxic
	polite	toxic	toxic
Formality	formal	impolite	impolite
		polite	polite
	informal	impolite	impolite
		polite	polite

Table 5: Model combinations for the human evaluation to address whether the style-specific datastores outperform the mixed datastores.

dataset	source style	target style	datastore style
Formality	formal	formal	formal
	informal	informal	informal
Toxicity	neutral	non-toxic	non-toxic
	non-toxic	toxic	toxic
Politeness	impolite	non-toxic	non-toxic
	polite	toxic	toxic
Formality	formal	impolite	impolite
		polite	polite
	informal	impolite	impolite
		polite	polite

Table 6: Model combinations for the human evaluation to address whether the kNN-LM outperforms the baseline LM. Both models being compared use the style architecture.

dataset	source style	datastore	target style
Formality	formal	formal	formal
	informal	informal	informal
Toxicity	neutral	mixed	non-toxic, toxic
	non-toxic	non-toxic	non-toxic
Politeness	impolite	toxic	toxic
	polite	mixed	non-toxic, toxic
	impolite	polite	non-toxic
	polite	impolite	toxic

Table 7: Model combinations for human evaluation to address whether the style architecture outperforms the baseline kNN-LM.

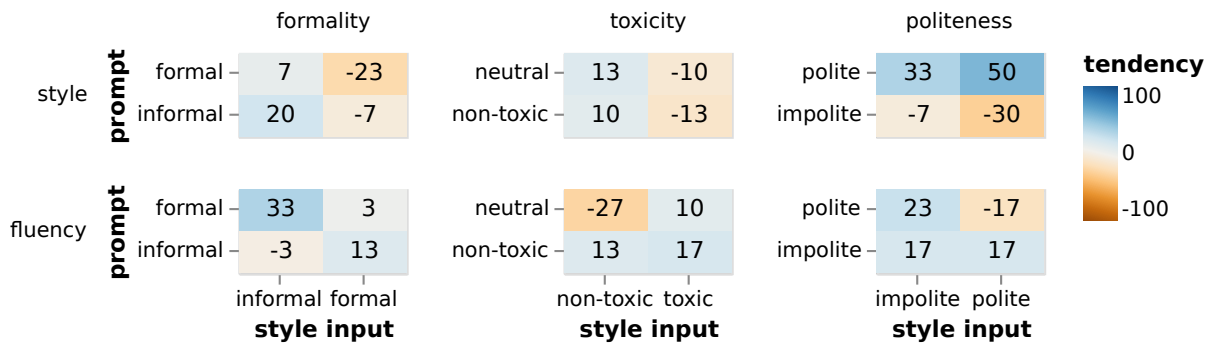


Figure 4: Survey results: Comparison of kNN-LM with mixed DS and style-specific DS. A tendency < 0 corresponds to the mixed DS being preferred by annotators.




Figure 5: Survey results: Comparison of LM to kNN-LM for both architectures. A tendency < 0 corresponds to the LM being preferred by annotators.




Figure 6: Survey results: Comparison of kNN-LM with baseline architecture vs. style architecture. Style input on the x -axis refers only to the style LM, since the baseline LM has no style input. A tendency < 0 corresponds to the the baseline architecture being preferred by annotators.