

# Synthesizing and Adapting Error Correction Data for Mobile Large Language Model Applications

Yanxiang Zhang\*, Zheng Xu\*, Shanshan Wu\*, Yuanbo Zhang, Daniel Ramage

Google

{zhangyx, xuzheng, shanshanw, zyb, dramage}@google.com

## Abstract

Error correction is an important capability when applying large language models (LLMs) to facilitate user typing on mobile devices. In this paper, we use LLMs to synthesize a high-quality dataset of error correction pairs to evaluate and improve LLMs for mobile applications. We first prompt LLMs with error correction domain knowledge to build a scalable and reliable addition to the existing data synthesis pipeline. We then adapt the synthetic data distribution to match the mobile application domain by reweighting the samples. The reweighting model is learnt by predicting (a handful of) live A/B test metrics when deploying LLMs in production, given the LLM performance on offline evaluation data and scores from a small privacy-preserving on-device language model. Finally, we present best practices for mixing our synthetic data with other data sources to improve model performance on error correction in both offline evaluation and production live A/B testing.

## 1 Introduction

Modern typing applications on mobile devices use many machine learning models, e.g., language models (LMs) (Ouyang et al., 2017; Liu et al., 2024b). The generative capacity of LMs can significantly improve user experience by (automatically) correcting various errors and predicting next words to facilitate typing. Recent advancement in large language models (LLMs) have achieved impressive performance on many language tasks (OpenAI, 2024; Google, 2024; Meta, 2024), opening new opportunities for rewriting in mobile applications (Gunter et al., 2024; Liu et al., 2024b). In practice, LLMs can be deployed on mobile devices or on servers in datacenters. However, mobile devices have limited resources that currently only support moderate-sized LLMs (often less than 10

billion parameters). Even for LLMs on servers, moderate-sized models are preferred for mobile applications because of the considerations of latency, privacy and serving cost.

Error correction (EC) is an important capacity of LLMs for mobile applications (see examples in Fig. 1). As LLMs’ general capacity can decrease with the model size (Wei et al., 2022; Cho et al., 2024), it is important to evaluate and improve moderate-sized models for mobile applications. Moreover, the data distribution of mobile applications can differ from commonly collected public web data (Hard et al., 2018; Xu et al., 2023; Wu et al., 2024); typing on mobile touchscreens introduce more errors (Shi et al., 2025) in addition to common grammatical errors (Bryant et al., 2023; Stahlberg and Kumar, 2021). Such EC data for mobile applications differs from much of current LLMs’ training data.

Post-training with high-quality data is commonly used to align LLMs with users (Wei et al., 2021; Chung et al., 2022; Ouyang et al., 2022) and bridge the domain shift (Cho et al., 2024). Low-Rank Adaptation (LoRA) method, which only trains a small subset of parameters, is efficient for fine-tuning models for mobile applications (Hu et al., 2022). LoRA additionally provides the flexibility to fine-tune a set of different adapters to customize for various downstream tasks, useful for deploying LLMs on mobile devices (Gunter et al., 2024). However, collecting high-quality data for post-training for mobile applications is challenging because of the domain shift and privacy considerations on user data.

Production LLM mobile applications have developed pipelines to synthesize error correction data. Liu et al. (2024b) collects public web data, and then uses trained task-specific models (Lichtarge et al., 2020) to detect grammatical errors. A typing simulator adds more mobile-specific errors to construct EC pairs based on the web data with detected

\*Equal contribution. Reverse alphabetical order.

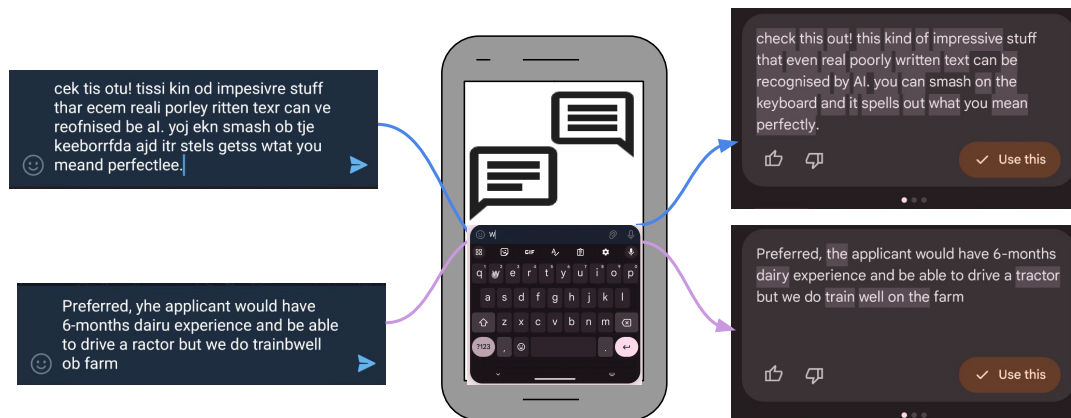


Figure 1: Examples of mobile LLM applications for error correction. User typing data has a domain shift compared to public web data. LLMs rewrite and correct highly corrupted text based on the context of the input itself.

grammar errors. These EC data pairs are split into training and validation datasets. This data pipeline extracts only a small set of EC data from a large collection of web data due to the detection and selection process. Moreover, the data distribution of web data differs from the mobile user distribution, as discussed in (Wu et al., 2024). Indeed, we observe a discrepancy between offline evaluation on validation data and live A/B test metrics in production.

Privacy-preserving methods are required to access in-domain user data to improve the model performance. Federated learning (FL), where devices collaboratively learn a model without transferring user data, and differential privacy (DP), where model is mathematically guaranteed not to memorize training data, are combined to privately fine-tune LMs (Xu et al., 2023; Choquette-Choo et al., 2024; McMahan et al., 2024). However, production DP FL systems on mobile devices only reliably train models with 10 million parameters (Daly et al., 2024). Differentially private synthetic data is another promising approach to collect high-quality privacy-preserving data (Kurakin et al., 2023; Yue et al., 2023). However, DP synthetic data generation requires iterative interaction between LLMs and private data, such as fine-tuning LLMs as data generators. The quality of synthetic data also decreases with the generator model size. These methods are not yet applied to training moderate-sized LLMs with billions of parameters for production mobile applications.

In this paper, we synthesize error correction data to improve LLMs with billions of parameters for mobile applications. In a production data pipeline, we incorporate human knowledge of the mobile

application domain and grammar errors to carefully design prompts, and use LLMs instead of grammar error detectors to scalably and reliably synthesize EC pairs (Sec. 2). To further overcome the discrepancy between offline evaluation on (synthetic) EC data and live A/B test metrics for model deployment in practice, we propose to adapt the data distribution to match the mobile application domain by reweighting the samples (Sec. 3). Small LMs with less than 100 million parameters are fine-tuned by federated learning with differential privacy on user data. These small LMs are used to generate initial scores for each offline evaluation sample. A reweighting model is parameterized to predict a final score for each sample based on the initial small LM scores. As the number of initial LM scores is small, the lightweight reweighting model is learnt by reweighting per-sample evaluation to predict only a handful of A/B test metrics collected during model deployment. We demonstrate that the reweighting model, together with privacy-preserving small LMs, effectively predicts live A/B test metrics. Finally, we present best practices for mixing our synthetic data with other data sources to improve the model performance (Sec. 4). LoRA method is used to further fine-tune an LLM with billions of parameters that is already post-trained for general purpose instruction following. A continue training strategy, where the model is first fine-tuned on our large-scale synthetic data, followed by fine-tuning on a mixture of existing smaller dataset and reweighted synthetic data, achieves superior performance on various offline evaluations, and 2.47% to 7.18% relative improvements on key metrics in production live A/B test.

## 2 Synthesizing Error Correction Data

In this section, we discuss prompting LLMs as an addition to error correction data pipeline for efficiency and effectiveness, and show its advantage in scalability and domain adaptation. Following Wu et al. (2024), we synthesize an initial dataset in the domain of typing text on mobiles, by filtering and transforming public web data (i.e., C4 (Raffel et al., 2020) dataset), and collecting LLM generations with carefully crafted prompts of human knowledge. The initial dataset contains more than 100 million documents of conversation-like text, and even a small subset (about 0.2%) is much larger than the original EC dataset in production. We subsample the initial typing text dataset to reduce the subsequent processing costs from prompting LLM to add grammar and typing errors. To ensure good diversity and coverage during sampling, we first embed the documents using the Gecko (Lee et al., 2024) text embedding model, run k-means clustering to obtain 20k clusters. See Fig. 2 for statistics of clustering. We then sample 10 data points per cluster, resulting in a dataset contains about 200k documents, which has 2M examples where each example is either a sentence, or a user’s utterance. Each example is relatively short similar to the examples in our target distribution, i.e., texts typed by users using their mobile keyboards in chatting or search applications. Majority of these texts are clean (i.e., error-free).




Figure 2: The statistics of the 20k clusters for 100 million documents. The mean with standard deviation of cluster sizes is  $5225 \pm 2972$ .

To synthesize the EC text pairs, we add two types of errors to the clean texts: grammar error, and typing error. The grammar error is added by Gemini Ultra model (Google, 2024), and Table 3 shows the template of our prompt and an example. We experiment with different model sizes and find that Gemini Ultra performs best for analyzing and adding grammar errors. For high-quality data generation, the LLM is prompted to perform two more tasks in addition to generating the ungram-

matical texts: (1) The first task is to describe the added grammar errors. This allows us to perform a global analysis of the added grammar errors, and confirm our data cover all the grammar error types from (Bryant et al., 2017). The top 4 error categories (and its percentage in our synthetic data) are related to verb (52%), missing words (15%), plural (10%), and capitalization (5%). In terms of the number of grammar error per example, 12% examples have 1 error, while more than 80% examples have 2 or 3 errors. (2) The second task is to correct the ungrammatical texts with LLM added grammar errors. We only keep examples when the corrected text and the original clean text are equal. This filtration process removes around 40% of the data. After adding the grammar errors, we next add typing errors that simulate the behavior of real users typing with mobile keyboard. This is done by heuristic rules that add various typing errors, such as transposition, omission, repetition, and spatial errors (Liu et al., 2024b).




Figure 3: Good ratio for error correction on the (a) original validation data and (b) synthetic validation data. The models are trained with the original training data, synthetic training data, and vanilla sampling of synthetic data without clustering. LLMs are used to judge whether the EC output is acceptable to compute good ratio. Our large-scale LLM assisted synthetic data works well on both domains even if there is potential distribution shift from the original dataset collected by error detection on public web data.

### 2.1 Evaluation Setup and Preliminary Results

Our synthetic EC dataset has about 1.2M examples. Each example is a pair of (corrupted, clean) sentences. We random sample a small subset of our synthetic data for validation, and use the rest of data for LoRA fine-tuning a Gemini Nano model (Google, 2024). Both the training and validation dataset are much larger than the original dataset synthesized by the previous production data pipeline. Figs. 3, 6 and 7 shows the results of training and evaluation with the small original production dataset, and our large synthetic dataset, respectively. We provide an additional ablation curve

on our synthetic data with vanilla subsampling instead of clustering-based subsampling. Fig. 6 (in App. B) measures the error correction performance by sequence accuracy, i.e., the exact match between corrected sentence and the target clean sentences, and shows that fine-tuning help while too many steps on small dataset may quickly degrade utility. Our evaluation on error correction matches previous observation on dialogue generation and summarization tasks (Cho et al., 2024).

We further use Gemini Pro models as judges to measure whether the corrected sentence is a high-quality rewrite of the target sentence even if they do not exact match for each word, and report the good ratio for the top 1 output and the best of top 3 outputs from our fine-tuned LLMs, in Figs. 3 and 7. Good ratio mimics the user behavior on selecting rewritten text from mobile applications. We select a small number of models from training steps {600, 1000, 2000, 4000} for evaluation to reduce the cost of LLM judges. Models from two different training runs are evaluated to compute standard deviation for error bars. The trend of the sequence accuracy and good ratios align well Figs. 3, 6 and 7. We observe performance discrepancy between the original production dataset and our synthetic dataset, which suggests a potential domain difference. Fine-tuning on our large-scale synthetic data is more robust compared to the small original dataset, and achieves competitive model performance even when evaluated on the original validation set. The clustering-based subsampling achieves comparable results on the original evaluation, and better results on the synthetic evaluation, compared to vanilla subsampling. In the rest of the paper, we will use the synthetic data subsampled with the clusters.

### 3 Privacy-Preserving Domain Adaptation by Reweighting

We have synthesized a large-scale error correction dataset in Sec. 2 by carefully prompting LLMs to simulate typing text and systematically add errors. However, Wu et al. (2024) suggests public LLMs and human prior knowledge in prompt may not be sufficient to bridge the potential domain shift. When deploying previously trained models, we observe misalignment in offline evaluation on the original validation set, and live A/B test metrics. We also observe the discrepancy between original validation set and our synthetic validation set in

Fig. 3a. As our synthetic data explicitly guided LLMs with prior knowledge on mobile typing for synthesis, is it closer to the domain of mobile applications in practice? In this section, we developed a privacy-preserving approach for domain adaptation by reweighting samples in the dataset. The reweighting model is built upon a small LM trained with DP FL, and a handful of live A/B test metrics tracked in previous model deployment.

When evaluating an error correction model  $M$  on a dataset  $\{(x_i, y_i)\}_{i=1}^N$  of  $N$  (corrupted, clean) samples, a measurement  $\chi(M(x_i), y_i) \in \{0, 1\}$  is generated for each sample by comparing the model output  $M(x_i)$  and corresponding target  $y_i$ . We have offline metric for evaluating the model by taking the average over all samples, i.e.,  $\sum_{i=1}^N \chi(M(x_i), y_i)/N$ , which becomes sequence accuracy in Fig. 6 when  $\chi(\cdot, \cdot)$  is exact match, and good ratio in Figs. 3 and 7 when  $\chi(\cdot, \cdot)$  is judged by LLMs. To reweight samples for domain adaptation, we first train two small LMs  $S_p, S_f$  of about 8 million parameters for scoring samples. Model  $S_p$  is trained on public C4 dataset, and model  $S_f$  is further fine-tuned from  $S_p$  on user data in a production FL system (Xu et al., 2023; Wu et al., 2024). Model  $S_f$  is a privacy-preserving model with formal DP guarantee  $\epsilon < 10$ , and captures the domain information from mobile application. We define a reweighting model parameterized by  $\theta = (\theta_f, \theta_p, \theta_b)$  as

$$w(\theta, y_i) = C_{\min} + (C_{\max} - C_{\min}) \sigma(\theta_f S_f(y_i) + \theta_p S_p(y_i) + \theta_b), \quad (1)$$

where  $C_{\min}, C_{\max}$  are constants determining the minimum and maximum value of the reweighting scores,  $\sigma(\cdot)$  is the sigmoid function, and  $S_f(\cdot), S_p(\cdot)$  represent the average log likelihood on predicting words in the target sentence  $y_i$ .

When deploying  $K$  models  $\{M_j(\cdot)\}_{j=1}^K$  in practice, we collect corresponding live A/B test metrics  $\{v_j\}_{j=1}^K$ . We consider key metrics like click through rate and accept rate for error correction in mobile applications, and hence each  $v_j \in \mathbb{R}^d$  is a vector representing multiple metrics. We optimize the objective below to learn the reweighting model,

$$\min_{\theta, \alpha} R(\theta, \alpha) + \lambda \left\| \frac{1}{N} \sum_{i=1}^N w(\theta, y_i) - 1 \right\|^2, \quad (2)$$

$$R(\theta, \alpha) = \sum_{j=1}^K \left\| \frac{\alpha_1}{N} \sum_{i=1}^N w(\theta, y_i) \chi(M_j(x_i), y_i) + \alpha_0 - v_j \right\|^2$$

where  $\alpha = (\alpha_1, \alpha_0)$  is regression parameter to predict live metrics from offline evaluation; per-sample reweighting score  $w(\theta, y_i)$  is defined in Eq. (1) to adapt offline data to mobile application domain to achieve small regression residual  $R(\theta, \alpha)$ ;  $\lambda$  is a hyperparameter on the regularizer of the reweighting scores.




Figure 4: Comparing the (a) heuristic  $\{0, 1\}$  reweighting in (Wu et al., 2024) and (b) our reweighting model  $w(\theta, \cdot) = 0.01 + 1.99\sigma(40.64S_f - 30.44S_p - 1.59)$ . Both methods use public pre-trained small LM  $S_p$  and the same model further fine-tuned with DP FL  $S_f$ . The learnt scores in (b) have large overlap with manual selection in (a) from (Wu et al., 2024).

We use auto differentiation and L-BFGS optimizer in JAX (Bradbury et al., 2018) to optimize Eq. (2) to learn regression parameters  $\alpha \in \mathbb{R}^{2d}$  and reweighting parameters  $\theta \in \mathbb{R}^3$ . The dimensionality of  $\theta, \alpha$  is relatively small, and they can be learnt from a handful of live metrics  $\{v_j\}_{j=1}^K$  collected during launching different error correction models. We collected two sets of live metrics, the training set evaluated 10 models with live A/B test, and the validation set evaluated 5 models. We use top-3 good ratio as in Fig. 7 for offline evaluation on the original dataset. Due to production test configuration, the two sets of live metrics have different scales and hence we cannot use the same regression parameter. We set the range of reweighting scores as  $C_{\max} = 2, C_{\min} = 0.01$ , and regularizer strength  $\lambda = 0.01$ . Tab. 1 summarizes residuals for training, cross-validation and validation, and reweighting achieves smaller residual when predicting live metrics across different settings. The absolute value of cross-validation and validation residuals are smaller than training residuals as training is the summation over all live metric samples in Eq. (2).

After training, our reweighting model parameters are  $(\theta_f, \theta_p, \theta_b) = (40.64, -30.44, -1.59)$ , which suggests the reweighting score is positively correlated with the fine-tuned model output  $S_f(\cdot)$  calibrated by pre-trained model output  $S_p(\cdot)$ . The difference of the two model outputs represent the likelihood discrepancy, which has also been used for inference time domain adaption (Liu et al.,

2024a) and training data detection (Kandpal et al., 2024). Wu et al. (2024) discusses a heuristic filtering strategy for domain adaptation that is effective for selecting data to train small LMs for mobile applications. The heuristic filtering is equivalent to setting  $w(y_i) = 1$  when  $S_f(y_i) > S_p(y_i)$  and  $S_f(y_i) > -5$ , and  $w(y_i) = 0$  otherwise. This heuristic approach often helps predicting live A/B test metrics compared to uniform weighting, but fails sometimes, and generally achieves higher residual than our reweighting score. Fig. 4 shows the difference between our reweighting model and (Wu et al., 2024) for different pre-trained and fine-tuned model scores.




Figure 5: Good ratio for the best of top 3 candidates for error correction on the (a) original validation data and (b) synthetic validation data. Solid lines reweight the samples by the  $w(\theta, \cdot)$  model learnt to fit live A/B test metrics in Sec. 3. The models are trained with synthetic training data with the same setting as in Fig. 7;  $\times 4$  increased batch size (synth\_lb); mixture of original and synthetic data; and mixture of original and filtered by  $w(\theta, \cdot)$  (mix\_fil). LLMs are used to judge whether the error correction output is acceptable to compute good ratio.

Finally, we further apply a defense in depth strategy when using our reweighting models for privacy-preserving domain adaptation. We use a standard production PII detection pipeline to remove any possible sensitive information in the synthetic data, even if they are hallucinated by LLMs. And our fine-tuned LLMs are equipped with another layer of safety and privacy safeguarding when deploying in practice.

## 4 Mixing Data for Fine-tuning

Based on our synthetic dataset in Sec. 2 and domain-adaptive reweighting model in Sec. 3, we improve Gemini Nano (Google, 2024) for error correction by LoRA fine-tuning (Hu et al., 2022). By combining the original dataset and the large-scale (reweighted) synthetic dataset in a continue training strategy, the model performance is improved in both offline evaluation and live A/B test in production. Unless otherwise specified, our experiments use the same configuration for training previous production models with the original data, as

$R(\theta, \alpha)$	$w = 1$	$w(y_i) \in \{0, 1\}$ (Wu et al., 2024)	Our $w(\theta, y_i)$
Train	$1.51 \times 10^{-4}$	$1.41 \times 10^{-4}$	$1.19 \times 10^{-4}$
CrossVal	$(5.26 \pm 2.76) \times 10^{-5}$	$(4.23 \pm 2.42) \times 10^{-5}$	$(3.99 \pm 3.15) \times 10^{-5}$
Val	$(3.51 \pm 4.15) \times 10^{-6}$	$(5.44 \pm 7.10) \times 10^{-6}$	$(2.07 \pm 2.56) \times 10^{-6}$

Table 1: Regression residual  $R$  for different reweighting strategies. Smaller residual suggests reweighting helps predicting live metrics. We report mean and standard deviation for predicting each live metrics  $v_j$  in held-one out cross validation. We also use held-one out for validation set to fit only regression parameters with fixed reweighting.

described in Sec. 2. As Gemini Nano is already post-trained with general purpose instructions, our fine-tuning is a continuous post-training. During inference after model deployment, our model is only effective for mobile applications when our fine-tuned LoRA module is applied to the base Gemini Nano model.

We discuss our model training practices and observations. **(1) Increasing batch size.** Our synthetic dataset is much larger than the original dataset. As shown in Figs. 3, 6 and 7, at 4000 steps, the model performance trained on synthetic data still increases on synthetic validation data, while only starts to saturate on original validation data. In fact, 4000 steps do not complete a single epoch on our synthetic training data. We increase the batch size to  $\times 4$ , which is the largest batch size without requesting more resources. We found our LoRA fine-tuning is relatively robust for learning rate between  $\times 1$  and  $\times 4$  of the original learning rate, and hence fixed the learning rate to be  $\times 1$ . As shown in Figs. 5, 8 and 9, large batch training achieves comparable performance on original validation data, while improves the performance on synthetic validation data. We use large batch training in following experiments. **(2) Simple mixing** the small original dataset and the large scale synthetic data improves the performance on original validation data, but slightly degrades the performance on synthetic validation data. There is a trade-off on the ratio of the mixture: while it is relatively robust when we have original and synthetic ratio in the range of 1 : 1 and 1 : 8, ratio 1 : 4 achieves good balance and is used in the following mixing experiments. **(3) Reweighting for training and evaluation.** Reweighting model in Sec. 3 is used to adapt the offline data distribution to the mobile application distribution. The trend in reweighted metrics in Figs. 5 and 9 generally align with the uniform weighted counterparty. In addition to reweighting to bring offline evaluation closer to live A/B test metrics, we further explore reweighting for domain adaptation in training. We filter the synthetic data set and only keep samples with reweighting scores  $w(\theta, y_i) \geq w_t$ . We choose

the threshold  $w_t = 1$  as  $w(\theta, y_i)$  is in the range of  $C_{\min} = 0.01$  and  $C_{\max} = 2$ , and about half of the samples in the synthetic dataset have reweighting scores passed the threshold. We only filter our synthetic dataset as the original dataset is already very small. Mixing the filtered synthetic data with the original dataset for training achieves good reweighted metrics even if the uniform weighted metrics slightly degrades compared to mixing with the full synthetic data. **(4) Continue training.** As our synthetic data is large, we propose a continue training strategy: first fine-tune on the full synthetic dataset for 1000 steps (about one epoch), and then continue training on the original data (ContOrig), the mixture of original and synthetic dataset (ContMix), and the mixture of the original and filtered synthetic data (ContMixFil), see Tab. 2. For each training method, we select the best model from steps  $\{600, 1000, 2000, 4000\}$ , and run training at least twice to compute the standard deviation.

In Tab. 2, ContMix and ContMixFil achieve best or close to best results on both original validation data and synthetic validation data. They achieve higher good ratio on the original validation data than model trained on the original data only, or the mixture of original and synthetic data. They are comparable to the best performance on synthetic validation data achieved by training on synthetic data only with large batches. As ContMixFil achieves better performance on the reweighted metrics that better reflects the mobile application domain, we further compare the model trained by Original and ContMixFil in production live A/B test. Compared to Original, ContMixFil achieves 2.47% to 7.18% relative improvement on key production metrics like click through rate and accept rate.

## 5 Conclusion

This paper presents a method to enhance error correction in mobile LLMs by creating a high-quality synthetic dataset using LLM prompts enriched with domain knowledge. We further adapt the public (synthetic) data to better match the domain of production mobile applications by developing

Training Method	Original Data Eval (%)				Synthetic Data Eval (%)			
	Top-1	Top-1 (w)	Top-3	Top-3 (w)	Top-1	Top-1 (w)	Top-3	Top-3 (w)
Original	68.74±0.38	71.16±0.54	79.34±0.02	80.38±0.26	71.35±0.31	76.24±0.11	82.96±0.30	86.24±0.26
SynthLB	66.64±1.28	68.25±1.26	77.4±0.91	77.93±0.72	87.5±0.12	90.29±0.07	92.75±0.05	94.30±0.21
Mix	70.22±0.94	72.91±0.31	80.9±1.22	82.57±0.61	85.37±0.13	88.71±0.30	91.71±0.15	93.77±0.07
ContOrig	68.82±0.14	71.21±0.27	79.34±0.54	80.66±0.13	77.40±0.34	80.78±0.15	86.49±0.57	89.35±0.64
ContMix	69.22±0.26	71.52±0.26	79.86±0.26	81.33±0.19	<b>86.04</b> ±0.32	88.88±0.06	<b>92.03</b> ±0.29	93.69±0.19
ContMixFil	<b>70.48</b> ±0.00	<b>73.31</b> ±0.10	<b>80.78</b> ±0.46	<b>82.28</b> ±0.87	85.78±0.20	<b>89.51</b> ±0.02	91.71±0.25	<b>93.90</b> ±0.14

Table 2: Good ratio for error correction on the original validation data and synthetic validation data. Top-3 evaluates the best of three model outputs. Columns with “(w)” reweight the samples by the  $w(\theta, \cdot)$  model learnt to fit live A/B test metrics in Sec. 3. The models are trained by the original dataset as in Fig. 3; the synthetic data with large batches, mixture of original and synthetic data as in Fig. 5; and three continue training strategies. Continue training gets the best, or close to best performance on offline evaluation of both original and synthetic validation data.

a privacy-preserving reweighting model, using a small LM trained with federated learning and differential privacy, alongside a few live A/B test metrics. Our experiments show that fine-tuning a billion-size LLM with a mixture of the original dataset and the reweighted synthetic data, especially via continue training, significantly improves performance in offline evaluations and live A/B tests.

## Acknowledgements

The authors thank Felix Stahlberg, Shankar Kumar, and Michael Xuelin Huang for discussions in the early stage of the project; Zachary Garrett and Shumin Zhai for reviewing an early draft.

## References

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs*.

Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. *Automatic annotation and evaluation of error types for grammatical error correction*. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 793–805, Vancouver, Canada. Association for Computational Linguistics.

Christopher Bryant, Zheng Yuan, Muhammad Reza Qorib, Hannan Cao, Hwee Tou Ng, and Ted Briscoe. 2023. *Grammatical error correction: A survey of the state of the art*. *Computational Linguistics*, 49(3):643–701.

Yae Jee Cho, Luyang Liu, Zheng Xu, Aldi Fahrezi, and Gauri Joshi. 2024. *Heterogeneous low-rank approximation for federated fine-tuning of on-device foundation models*. *EMNLP*.

Christopher A Choquette-Choo, Arun Ganesh, Ryan McKenna, H Brendan McMahan, John Rush, Abhradeep Guha Thakurta, and Zheng Xu. 2024.

(amplified) banded matrix factorization: A unified approach to private training. *Advances in Neural Information Processing Systems*, 36.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, and 1 others. 2022. *Scaling instruction-finetuned language models*. *arXiv preprint arXiv:2210.11416*.

Katharine Daly, Hubert Eichner, Peter Kairouz, H Brendan McMahan, Daniel Ramage, and Zheng Xu. 2024. *Federated learning in practice: reflections and projections*. In *2024 IEEE 6th International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications (TPS-ISA)*, pages 148–156. IEEE.

Gemini Team Google. 2024. *Gemini: A family of highly capable multimodal models*. *Preprint*, arXiv:2312.11805.

Tom Gunter, Zirui Wang, Chong Wang, Ruoming Pang, Andy Narayanan, Aonan Zhang, Bowen Zhang, Chen Chen, Chung-Cheng Chiu, David Qiu, and 1 others. 2024. *Apple intelligence foundation language models*. *arXiv preprint arXiv:2407.21075*.

Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. 2018. *Federated learning for mobile keyboard prediction*. *arXiv preprint arXiv:1811.03604*.

Charlie Hou, Akshat Shrivastava, Hongyuan Zhan, Rylan Conway, Trang Le, Adithya Sagar, Giulia Fanti, and Daniel Lazar. 2024. *Pre-text: Training language models on private federated data in the age of llms*. *ICML*.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. *Lora: Low-rank adaptation of large language models*. *ICLR*, 1(2):3.

Nikhil Kandpal, Krishna Pillutla, Alina Oprea, Peter Kairouz, Christopher A Choquette-Choo, and Zheng Xu. 2024. *User inference attacks on large language models*. *EMNLP*.

- Alexey Kurakin, Natalia Ponomareva, Umar Syed, Liam MacDermed, and Andreas Terzis. 2023. Harnessing large-language models to generate private synthetic text. *arXiv preprint arXiv:2306.01684*.
- Jinhyuk Lee, Zhuyun Dai, Xiaoqi Ren, Blair Chen, Daniel Cer, Jeremy R. Cole, Kai Hui, Michael Boratko, Rajvi Kapadia, Wen Ding, Yi Luan, Sai Meher Karthik Duddu, Gustavo Hernandez Abrego, Weiqiang Shi, Nithi Gupta, Aditya Kusupati, Praateek Jain, Siddhartha Reddy Jonnalagadda, Ming-Wei Chang, and Iftexhar Naim. 2024. [Gecko: Versatile text embeddings distilled from large language models](#). *Preprint*, arXiv:2403.20327.
- Jared Lichtarge, Chris Alberti, and Shankar Kumar. 2020. Data weighted training strategies for grammatical error correction. *Transactions of the Association for Computational Linguistics*, 8:634–646.
- Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, and Noah A Smith. 2024a. Tuning language models by proxy. *COLM*.
- Renjie Liu, Yanxiang Zhang, Yun Zhu, Haicheng Sun, Yuanbo Zhang, Michael Huang, Shanqing Cai, Lei Meng, and Shumin Zhai. 2024b. [Proofread: Fixes all errors with one tap](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 286–293, Bangkok, Thailand. Association for Computational Linguistics.
- H Brendan McMahan, Zheng Xu, and Yanxiang Zhang. 2024. A hassle-free algorithm for strong differential privacy in federated learning systems. *EMNLP*.
- Llama Team Meta. 2024. [The llama 3 herd of models](#).
- OpenAI. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*.
- Tom Ouyang, David Rybach, Franoise Beaufays, and Michael Riley. 2017. Mobile keyboard input decoding with finite-state transducers. *arXiv preprint arXiv:1704.03987*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Danqing Shi, Yujun Zhu, Francisco Erivaldo Fernandes Junior, Shumin Zhai, and Antti Oulasvirta. 2025. Simulating errors in touchscreen typing. *arXiv preprint arXiv:2502.03560*.
- Felix Stahlberg and Shankar Kumar. 2021. Synthetic data generation for grammatical error correction with tagged corruption models. *arXiv preprint arXiv:2105.13318*.
- Bowen Tan, Zheng Xu, Eric P Xing, Zhiting Hu, and Shanshan Wu. 2025. Synthesizing privacy-preserving text data via finetuning without finetuning billion-scale llms. In *ICML*.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, and 1 others. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Shanshan Wu, Zheng Xu, Yanxiang Zhang, Yuanbo Zhang, and Daniel Ramage. 2024. Prompt public large language models to synthesize data for private on-device applications. *Conference on Language Modeling (COLM)*.
- Chulin Xie, Zinan Lin, Arturs Backurs, Sivakanth Gopi, Da Yu, Huseyin A Inan, Harsha Nori, Haotian Jiang, Huishuai Zhang, Yin Tat Lee, and 1 others. 2024. Differentially private synthetic data via foundation model apis 2: Text. In *Forty-first International Conference on Machine Learning (ICML)*.
- Zheng Xu, Yanxiang Zhang, Galen Andrew, Christopher Choquette, Peter Kairouz, Brendan McMahan, Jesse Rosenstock, and Yuanbo Zhang. 2023. Federated learning of gboard language models with differential privacy. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, pages 629–639, Toronto, Canada.
- Da Yu, Peter Kairouz, Sewoong Oh, and Zheng Xu. 2024. Privacy-preserving instructions for aligning large language models. *ICML*.
- Xiang Yue, Huseyin A Inan, Xuechen Li, Girish Kumar, Julia McAnallen, Hoda Shajari, Huan Sun, David Levitan, and Robert Sim. 2023. Synthetic text generation with differential privacy: A simple and practical recipe. *ACL*.



## A Prompt in Sec. 2

---

### LLM prompt template

Imagine that you are an English school teacher. Your goal is to teach high school student English grammar. Here are some common grammatical errors: [...]

Given the following sentence(s):

[...original sentence...]

What grammatical errors are the students likely to make?

Now apply these grammatical errors to the original sentence(s), and generate the ungrammatical sentence(s). Do not modify the original sentence(s) except applying the grammatical errors.

The output should be in the following format: [...]

Finally, correct the grammatical errors in the generated ungrammatical sentence(s). Do not modify the sentence(s) except correcting the grammatical errors. The output should be in the following format: [...]

---

### Example

Original sentence : Yesterday I went to a store that has nice furniture.

LLM output :

**\*\*Ungrammatical sentences\*\***: Yesterday I went to a store that have nice furnitures.

**\*\*Error 1: Subject-verb agreement error\*\***: “have” should be “has” to agree with the singular subject “store”.

**\*\*Error 2: Plurization error\*\***: “furnitures” should be “furniture” as it is an uncountable noun.

**\*\*Corrected sentences\*\***: Yesterday I went to a store that has nice furniture.

---

Table 3: Our LLM prompt for adding grammar errors and an example output. The prompt starts with a role (i.e., school teacher), analyzes the original sentence, and adds potential grammar errors, followed by correcting all errors. The [...] parts are omitted for simplicity.

## B Additional results in Sec. 2




Figure 6: Sequence accuracy of error correction on the (a) original validation data and (b) synthetic validation data. The models are trained with the original training data, synthetic training data, and vanilla sampling of synthetic data without clustering.

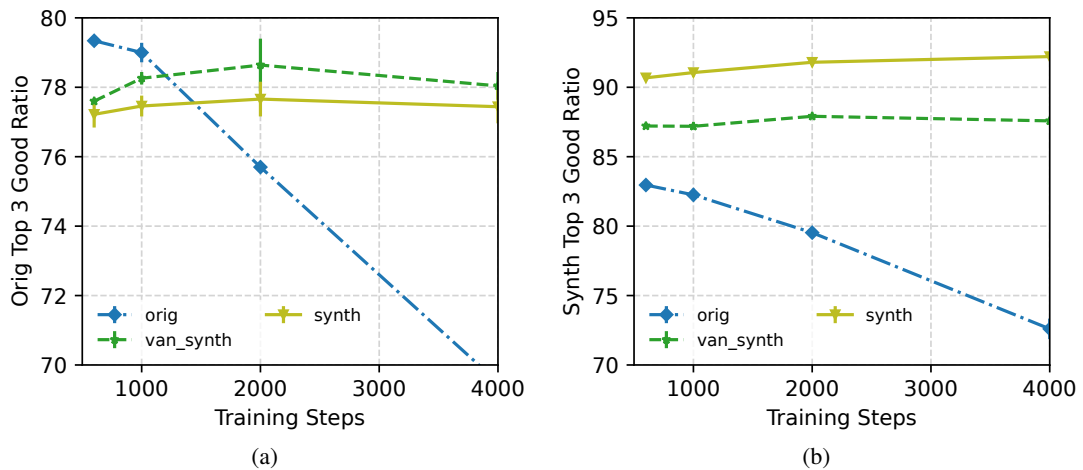


Figure 7: Good ratio for the best of top 3 candidates for error correction on the (a) original validation data and (b) synthetic validation data. The models are trained with the original training data, synthetic training data, and vanilla sampling of synthetic data without clustering. LLMs are used to judge whether the error correction output is acceptable to compute good ratio.

### C Additional results in Sec. 4




Figure 8: Sequence accuracy of error correction on the (a) original validation data and (b) synthetic validation data. The models are trained with synthetic training data with the same setting as in Fig. 7;  $\times 4$  increased batch size (synth\_lb); mixture of original and synthetic data; and mixture of original and filtered by  $w(\theta, \cdot)$  (mix\_fil).

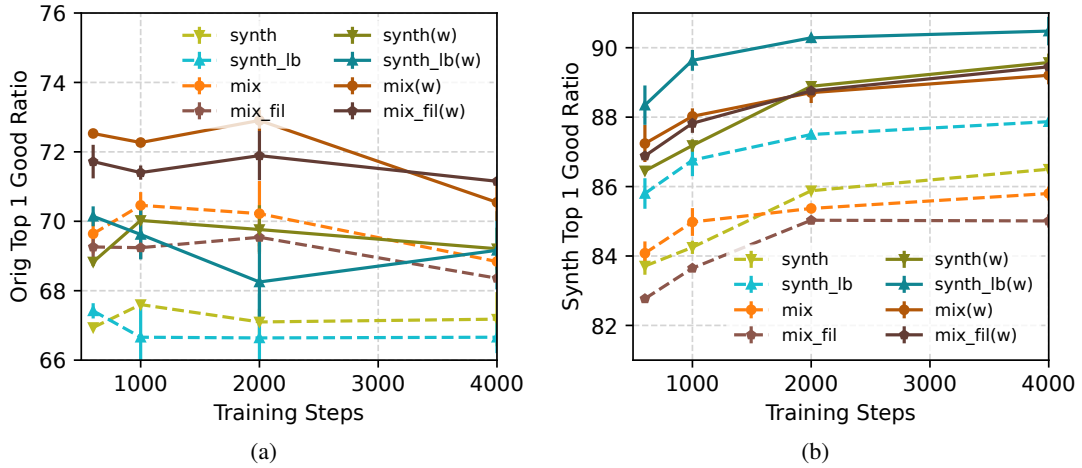


Figure 9: Good ratio for error correction on the (a) original validation data and (b) synthetic validation data. Solid lines reweight the samples by the  $w(\theta, \cdot)$  model learnt to fit live A/B test metrics in Sec. 3. The models are trained with synthetic training data with the same setting as in Fig. 7;  $\times 4$  increased batch size (synth\_lb); mixture of original and synthetic data; and mixture of original and filtered by  $w(\theta, \cdot)$  (mix\_fil). LLMs are used to judge whether the error correction output is acceptable to compute good ratio.

### D Limitation and Future Work

Our preliminary exploration on reweighting suggests combining live data from production applications and LLMs in a privacy-preserving manner is promising, and there are a lot of possibilities with the limited accessible information. Our usage of reweighting scores in training by filtering samples considers the trade-off of effectiveness, easy-to-implement, and future maintenance in production. There are many other potential domain adaptation methods for future experiments. Finally, We leverage a small privacy-preserving LM to capture domain information from mobile applications, while other forms of information such as histogram (Xie et al., 2024; Hou et al., 2024; Yu et al., 2024; Tan et al., 2025) are worth considering, especially given the flexibility of the next generation FL systems in trusted execution environments (Daly et al., 2024). With more data generated from the interaction of LLMs and users in production deployment, our approach can become more powerful for not only domain adaptation but also other improvement such as personalization and agency, enabled by privacy-preserving methods.