

Seeing the Forest through the Trees: Data Leakage from Partial Transformer Gradients

Weijun Li¹, Qiongkai Xu¹, and Mark Dras¹

¹School of Computing, FSE, Macquarie University, Sydney, Australia
¹weijun.li1@hdr.mq.edu.au, {qiongkai.xu, mark.dras}@mq.edu.au

Abstract

Recent studies have shown that distributed machine learning is vulnerable to gradient inversion attacks, where private training data can be reconstructed by analyzing the gradients of the models shared in training. Previous attacks established that such reconstructions are possible using gradients from all parameters in the entire models. However, we hypothesize that most of the involved modules, or even their sub-modules, are at risk of training data leakage, and we validate such vulnerabilities in various intermediate layers of language models. Our extensive experiments reveal that gradients from a single Transformer layer, or even a single linear component with 0.54% parameters, are susceptible to training data leakage. Additionally, we show that applying differential privacy on gradients during training offers limited protection against the novel vulnerability of data disclosure.¹

1 Introduction

As the requirement for training machine learning models on large-scale and diverse datasets intensifies, distributed learning frameworks have risen as an effective solution that balances both the need for intensive computation and critical privacy concerns among edge users. As a prime example, Federated Learning (FL) (McMahan et al., 2016) preserves the privacy of participants by retaining each client’s data on their own devices, while only exchanging essential information, such as model parameters and updated gradients. Nonetheless, recent research (Zhu et al., 2019; Dang et al., 2021; Balunovic et al., 2022) has demonstrated the possibility of reconstructing client data by a curious-but-honest server or clients who have access to the corresponding gradient information.

Specifically, two types of methods have been proposed to extract private textual training data:

¹Code available at: <https://github.com/weijun-l/partial-gradients-leakage>.

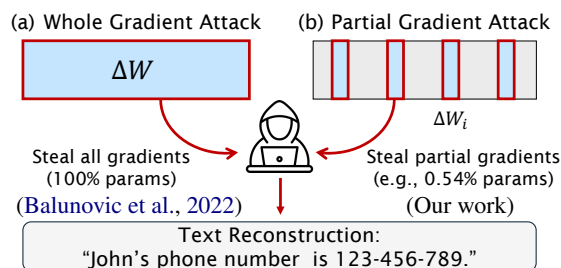


Figure 1: To reconstruct training data, prior attacks (a) typically require access to gradients from the whole model, while our attack (b) uses partial model gradients.

(i) gradient matching method (Zhu et al., 2019) align gradients from the presumed data with the monitored gradients; and (ii) analytical reconstruction techniques (Dang et al., 2021; Gupta et al., 2022) deduce the used tokens by analyzing the gradient patterns, such as the presence of non-zero embedding gradients that are correlated to those used tokens.

This study examines whether models are more vulnerable than realized by framing a **research question**: Can private training data be reconstructed using gradients from partial intermediate Transformer modules? This setting is motivated by several realistic scenarios. First, Parameter Freezing (Gupta et al., 2022) offers a straightforward defense against reconstruction attacks targeting specific layers, e.g., the embedding. Second, layer-wise training strategies are adopted to meet the diverse needs of: (i) transfer learning for domain adaptation (Chen et al., 2020; Saha and Ahmad, 2021), (ii) personalized federated learning for managing heterogeneous data (Mei et al., 2021) and (iii) enhancing communication efficiency (Lee et al., 2023). Generally, existing attacks require access to gradients from either (i) all deep learning modules or (ii) the word embedding/last linear layer; given the above defenses, this is not always practical.

In this work, we challenge the premise that gra-

dients from all layers are necessary to reconstruct training data. We demonstrate the feasibility of reconstructing text data from gradients of varying granularities, ranging from multiple Transformer layers down to a single layer, or even its single linear component (*e.g.*, individual Attention Query, Key modules), as depicted in Figure 1. Additionally, we investigate the impact of differential privacy on gradients (Abadi et al., 2016) as a defense and find our attacks remain effective, without significant degradation of model performance. Our study motivates further research into more effective defense mechanisms for distributed learning applications.

2 Related Work

Distributed Learning. Distributed learning, such as Federated Learning (FL), is a growing field aimed at parallelizing model training for better efficiency and privacy (Gade and Vaidya, 2018; Verbraeken et al., 2020; Froelicher et al., 2021). FL is highly valuable for privacy preservation by remaining sensitive data local as participants compute gradients on their devices and sharing updates via a central server (McMahan et al., 2017).

However, the shared gradients introduce an attack surface that allows malicious participants, or the curious-but-honest server, to reconstruct the training data. The gradient assets available to attackers may vary depending on the use of parameter freezing defense (Gupta et al., 2022) or layer-wise training strategies (Lee et al., 2023).

Gradient Inversion Attack (GIA). Recent studies (Zhu et al., 2019; Zhao et al., 2020; Yang et al., 2023; Zhang et al., 2024) have investigated data leakage in distributed learning, known as Gradient Inversion Attack (Zhang et al., 2023; Du et al., 2024). One strategy is the **analytical-based** approach, which identifies correlations between the gradients and model parameters to retrieve used training tokens. RLG (Dang et al., 2021), FILM (Gupta et al., 2022), DECEPTICONS (Fowl et al., 2022) and DAGER (Petrov et al., 2024) have demonstrated the use of gradients from specific layers, such as last linear and embedding layers, for reconstruction. While these works present effective attacks, parameter freezing (Gupta et al., 2022) can address threats from specific layers, and the method of DECEPTICONS assumes malicious parameter modification.

Alternatively, **optimization-based** method iter-

atively refines randomly generated data to match real data by minimizing the distance between their gradients across layers. Pioneering works like DLG (Zhu et al., 2019), InvertGrad (Geiping et al., 2020), and TAG (Deng et al., 2021) employed Euclidean (L2), Cosine, and combined Euclidean and Manhattan (L1) distances for data reconstruction. LAMP (Balunovic et al., 2022) advances these attacks with selective initialization, embedding regularization loss, and word reordering. Wu et al. (2023) and Li et al. (2023) build on the frameworks of DLG or LAMP while incorporating additional information, such as auxiliary datasets. Our study also adheres to optimization-based strategy but shows that partial gradients alone can reveal private training data.

Recent work has explored reconstructing private training data from single gradient modules, but with different focuses from ours. Wang et al. (2023) employs gradient decomposition in a synthetic two-layer network to recover training examples, though applying this method to practical deep networks requires idealized conditions such as transparent parameter manipulation and controlled activations (*e.g.*, tanh, sigmoid). Li et al. (2023) extends this approach to recover hidden representations from BERT’s pooler layer, guiding optimization via the LAMP method that leverages all gradients. DAGER (Petrov et al., 2024) uses the first two self-attention gradients of transformer-based language models for reconstruction, locating valid tokens by determining whether their embedding vectors lie within the subspace of the first layer’s gradient, then combining their embeddings to verify their occurrence with the second layer’s gradient to establish exact sequences. Unlike these methods, our research focuses on identifying vulnerabilities across all intermediate modules to gradient inversion attacks.

Defense against GIA. To mitigate the risk of GIA, two defense strategies have been explored: i) **encryption-based** methods, which disguise the real gradients using techniques such as Homomorphic Encryption (Zhang et al., 2020) and Multi-Party Computation (Mugunthan et al., 2019), and ii) **perturbation-based** methods, including gradient pruning (Zhu et al., 2019), adding noise through differential privacy (Balunovic et al., 2021), or through learned perturbations (Sun et al., 2021; Fan et al., 2024). The former approach incurs additional computational costs (Fan et al., 2024), while

the latter may face challenges in achieving a balance in the privacy-utility trade-off. In this work, we compare the vulnerabilities associated with attacking partial gradients versus the previously fully exposed setting under the defense of differential privacy on gradients (Abadi et al., 2016).

3 Data Leakage from Partial Gradients

In this section, we first introduce the threat model and then present the attack methodology, which enables the use of intermediate gradients in Transformer layers to reconstruct training data.

Threat Model. The attack operates in distributed training environments where the server distributes initial model weights and clients submit gradients derived from their local data. *Potential attackers*, either participating as clients or as curious-but-honest servers, can access shared parameters and intercept gradient communications. Unlike the *assets* in prior studies that permitted access to full gradients, our research focuses on scenarios where attackers only observe partial gradients. The goal of these attackers is to reconstruct the private text data received by other clients or servers in training.

Attack Strategy. Inspired by the gradient matching strategy, which involves minimizing the distance between gradients generated from randomly sampled data and the true gradients to iteratively align the dummy input with the real one, we formulate our optimization objective as,

$$\mathcal{L} = \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{M}} \mathcal{D}(\Delta \mathbf{W}'_{m,i}, \Delta \mathbf{W}_{m,i}), \quad (1)$$

where:

- \mathcal{N} is a non-empty subset of Transformer layers $\{1, 2, 3, \dots, l\}$ in an l -layer network.
- \mathcal{M} is a non-empty subset of the available modules $\{q, k, v, o, f, p\}$ within each layer.

We detail these notations and their involved number of parameters in Table 1. The lowest ratio is 0.54% for attacking a single linear module. $\Delta \mathbf{W}$ and $\Delta \mathbf{W}'$ represent the target gradient and the derived dummy gradient, respectively. \mathcal{D} serves as a distance measurement in optimization. We extend the use of cosine distance, following prior studies (Balunovic et al., 2022; Geiping et al., 2020),

$$\begin{aligned} \mathcal{D}_{\cos}(\Delta \mathbf{W}'_{m,i}, \Delta \mathbf{W}_{m,i}) \\ = 1 - \frac{\Delta \mathbf{W}'_{m,i} \cdot \Delta \mathbf{W}_{m,i}}{\|\Delta \mathbf{W}'_{m,i}\| \cdot \|\Delta \mathbf{W}_{m,i}\|}, \end{aligned} \quad (2)$$

Employed Gradients	Notation	#. Parameters	Used Ratio %
All Layers (Baseline)	$\Delta \mathbf{W}_{all}$	109,483,778	100
All Transformer Layers	$\Delta \mathbf{W}_T$	85,054,464	77.69
i -th Transformer Layer	$\Delta \mathbf{W}_{t,i}$	7,087,872	6.47
i -th FFN Output	$\Delta \mathbf{W}_{p,i}$	2,359,296	2.15
i -th FFN Fully Connected	$\Delta \mathbf{W}_{f,i}$	2,359,296	2.15
i -th Attention Output	$\Delta \mathbf{W}_{o,i}$	589,824	0.54
i -th Attention Query	$\Delta \mathbf{W}_{q,i}$	589,824	0.54
i -th Attention Key	$\Delta \mathbf{W}_{k,i}$	589,824	0.54
i -th Attention Value	$\Delta \mathbf{W}_{v,i}$	589,824	0.54

Table 1: Notations of varying gradient modules, and their parameter numbers for a BERT_{BASE} model.

which was reported to be stable and outperform other metrics (*i.e.*, L2 and L1).

Targeting different gradient units, our method can construct various specific gradient matching settings, *e.g.*, merely involving one Attention Query component (q) in the i -th Transformer layer, leading to the overall loss objective,

$$\mathcal{L} = \mathcal{D}(\Delta \mathbf{W}'_{q,i}, \Delta \mathbf{W}_{q,i}). \quad (3)$$

By involving all gradients and assigning normalized equal weights $1/l$, the objective can be specialized to the prior work (Balunovic et al., 2022), *i.e.*,

$$\mathcal{L} = \frac{1}{l} \sum_{i=1}^l \mathcal{D}(\Delta \mathbf{W}'_i, \Delta \mathbf{W}_i). \quad (4)$$

We explore varying degrees of gradient involvement for optimization alignment, from all Transformer layers to a single layer or even an individual linear component, revealing that every single module within the Transformer is vulnerable.

4 Experiments

4.1 Experimental Settings

Datasets. We examine reconstruction attacks on three classification datasets: **CoLA** (Warstadt et al., 2019), **SST-2** (Socher et al., 2013), and **Rotten Tomatoes** (Pang and Lee, 2005), following previous studies (Balunovic et al., 2022; Li et al., 2023) to evaluate reconstruction performance. We conduct experiments on 10 batches randomly sampled from each dataset separately, reporting average results across different test scenarios.

Models. We conduct experiments on BERT_{BASE}, BERT_{LARGE} (Devlin et al., 2019), and TinyBERT (Jiao et al., 2020), following previous studies (Deng et al., 2021; Balunovic et al., 2022; Li et al., 2023). We also adopted a BERT_{FT} model, which involves

fine-tuning $BERT_{BASE}$ for two epochs before attacks. For the word reordering step after reconstruction, we utilized a customized GPT-2 language model trained by Guo et al. (2021) as an auxiliary tool, same as the setting used in LAMP.

Evaluation Metrics. We evaluate the efficacy of our attack using **ROUGE-1**, **ROUGE-2**, and **ROUGE-L** (ROUGE, 2004; Lhoest et al., 2021), corresponding to unigram, bigram, and longest common sub-sequence, respectively. The F-scores of these metrics are reported.

Attack Setup. We adapt the open-source implementation of LAMP (Balunovic et al., 2022) to serve as both the basis framework and the baseline, as LAMP remains the state-of-the-art method for text reconstruction. We employ identical hyperparameters as LAMP for fair comparisons. The engineering contribution of our work is that we implemented a gradient extraction process to obtain partial gradients from modules at varying desired granularity.

Defense Setup. We employ DP-SGD (Abadi et al., 2016; Yousefpour et al., 2021), adjusting the noise multiplier σ while maintaining a clipping bound C as 1.0. To assess noise effects, we train a $BERT_{BASE}$ model on the **SST-2** dataset for 2 epochs, evaluating utility changes with the F1-score and MCC (Matthews, 1975; Chicco and Jurman, 2020). We set the delta δ to 2×10^{-5} , and explore noise multipliers from 0.01 to 0.5.²

4.2 Results and Analysis

Attack Results. We present the ROUGE-L scores for experiments on various gradient granularities, including results for Transformer layers in Figure 2, Attention modules across all layers in Figure 3, and FFN modules in Figure 4. These tests used the $BERT_{BASE}$ model on the **CoLA** dataset with a batch size of 1. Further results on ROUGE-1 and ROUGE-2 scores for **CoLA** are shown in Figure 5, **SST-2** in Figure 6, and **Rotten Tomatoes** in Figure 7. More results on different models and larger batch sizes ($B = 2, 4$) are provided in Table 3 and Table 4 in Appendix A, along with several reconstruction examples presented in Table 5 in Appendix B.

By inspecting Figure 2, we observe that using gradients from Transformer layers ΔW_T achieves

² δ is recommended to be smaller than $1/|D|$ (Abadi et al., 2016), where $|D|$ is the dataset size.




Figure 2: Results across varying **Transformer** layers.




Figure 3: Results across varying **Attention** Modules.




Figure 4: Results across varying **FFN** Modules.

performance comparable to the baseline, which uses all gradients. Additionally, merely using gradients from a single layer $\Delta W_{t,i}$ still yields decent attack scores, with layers 6 to 9 achieving results comparable to the baseline while using only 6.47% of model parameters. These results demonstrate that each layer is vulnerable to the reconstruction attack, with the middle layers hold highest risk.

Figure 3 presents the results of the attack from individual modules in the Attention Blocks across all layers. Most modules facilitate attack performance above 50%, while the Query and Key modules achieve relatively higher attack performance. Similarly, middle layers achieve the best performance; surprisingly, $\Delta W_{k,4}$, $\Delta W_{k,8}$, $\Delta W_{q,5}$ and $\Delta W_{o,6}$ achieve performance almost equivalent to the baseline, while using only 0.54% of its param-

Method	Gradient	$\sigma = 0.01$			$\sigma = 0.1$			$\sigma = 0.3$			$\sigma = 0.5$		
		$\epsilon = 2 * 10^8$			$\epsilon = 574.55$			$\epsilon = 9.38$			$\epsilon = 1.87$		
		R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
LAMP	ΔW_{all}	80.6	39.0	68.3	79.5	46.3	68.3	76.8	38.4	66.8	75.0	34.6	64.9
Ours	ΔW_T	78.3	40.3	65.9	75.7	38.7	64.4	77.6	36.6	66.5	75.9	50.2	70.5
	$\Delta W_{t,1}$	64.2	31.3	56.4	69.7	31.4	62.3	64.9	30.0	62.6	67.7	32.7	60.1
	$\Delta W_{q,1}$	69.8	35.8	64.5	69.3	17.6	56.0	68.0	16.0	58.4	71.1	19.5	59.4
	$\Delta W_{o,1}$	52.0	10.5	48.1	52.6	12.7	47.0	49.0	1.8	43.1	46.1	12.8	43.9
	$\Delta W_{f,1}$	54.2	12.4	49.8	57.1	9.0	49.8	57.4	9.5	49.9	57.9	8.2	48.8
	$\Delta W_{p,1}$	64.3	24.8	54.2	58.3	19.9	50.6	54.7	23.2	49.3	58.7	23.3	52.9
(a) Attack Performance													
F1-Score		0.891			0.855			0.71			0.675		
MCC		0.773			0.709			0.285			0		
(b) Model Utility													

Table 2: Evaluation of differential privacy defense on SST-2 dataset for varying gradient settings with batch size = 1.

eters. Similar results can be observed from FFN modules, as demonstrated in Figure 4.

Defense Results. We deploy differential privacy via DP-SGD (Abadi et al., 2016) to counter data reconstruction attacks. In our SST-2 dataset experiments, we apply varying noise levels (σ) from 0.01 to 0.5 under privacy budgets (ϵ) of 2×10^8 to 1.87. Noise levels above 0.5 were not explored due to a significant drop in the MCC metric from 0.773 to 0, compromising model utility. The results are detailed in Table 2. With the increase in noise, we observe a considerable decline in model utility. This observation supports our conclusion that while DP-SGD offers limited protection against the attack, it significantly impacts model utility, as the MCC and F1-Score metrics decreased substantially with a privacy budget ϵ of 1.87.

5 Conclusion

In this study, we investigate the feasibility of reconstructing training data using partial gradients in a Transformer model. Our extensive experiments demonstrate that all modules within a Transformer are vulnerable to such attacks, leading to a much higher degree of privacy risk than has previously been shown. Our examination of differential privacy as a defense also indicates that it is not sufficient to safeguard private text data from exposure, inspiring further efforts to mitigate these risks.

Limitations

We identify several opportunities for further improvement.

We conduct our experiments in the context of classification tasks; however, we propose that the task of language modeling could serve as an additional viable application scenario. Our preliminary validation revealed promising outcomes. However, we did not scale the test volume or involve them due to limitations in computational resources.

We evaluate the defense effectiveness of DP-SGD and find that it is not adequate to mitigate the risk without significant model utility degradation. However, we believe that other techniques, such as Homomorphic Encryption and privacy-preserved multi-party communication, could potentially reduce such a risk of privacy leakage. Nonetheless, these techniques often impose a significant overhead on system communication and substantial computational resources. Therefore, we advocate for further research into the defense strategies that can enhance the system’s resilience more efficiently.

In our experiments, we set the batch sizes to 1, 2, 4 and the average sentence length is less than 25 words, resulting in a total tokens involved being smaller than in generic industrial settings. We consider exploring larger batch sizes and longer sequences as future research directions, which were discussed in the related work, DAGER (Petrov et al., 2024).

While our experiments adhered to controlled settings similar to previous work such as LAMP, TAG, DLG—aimed at identifying foundational vulnerabilities in Transformer models, widely used in realistic applications—we also recognize the value of

exploring more real-world settings in future work. These could include network delays, asynchronous updates, and heterogeneous data, among others.

Ethics Statement

In this study, we delve into the vulnerabilities of every module in Transformer-based models against data reconstruction attacks. Our investigation seeks to evaluate the resilience of cutting-edge Transformer models when they are trained in a distributed learning setting and face malicious reconstruction attacks. Our findings reveal that each component is susceptible to these attacks.

Our research suggests that the risk of data breaches could be more significant than initially estimated, emphasizing the vulnerability of the entire Transformer architecture. We believe it is crucial to disclose such risks to the public, encouraging the research community to take these factors into account when developing secure systems and applications, and to promote further research into effective defense strategies.

Acknowledgement

We would like to express our appreciation to the anonymous reviewers for their valuable feedback. This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government. We also express our gratitude to SoC Incentive Fund, FSE strategic startup grant and HDR Research Project Funding for supporting both travel and research.

References

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318.
- Mislav Balunovic, Dimitar Dimitrov, Nikola Jovanović, and Martin Vechev. 2022. Lamp: Extracting text from gradients with language model priors. *Advances in Neural Information Processing Systems*, 35:7641–7654.
- Mislav Balunovic, Dimitar Iliev Dimitrov, Robin Staab, and Martin Vechev. 2021. Bayesian framework for gradient leakage. In *International Conference on Learning Representations*.
- Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. 2020. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4):83–93.
- Davide Chicco and Giuseppe Jurman. 2020. The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation. *BMC genomics*, 21(1):1–13.
- Trung Dang, Om Thakkar, Swaroop Ramaswamy, Rajiv Mathews, Peter Chin, and Françoise Beaufays. 2021. Revealing and protecting labels in distributed training. *Advances in Neural Information Processing Systems*, 34:1727–1738.
- Jieren Deng, Yijue Wang, Ji Li, Chenghong Wang, Chao Shang, Hang Liu, Sanguthevar Rajasekaran, and Caiwen Ding. 2021. Tag: Gradient attack on transformer-based language models. In *The 2021 Conference on Empirical Methods in Natural Language Processing*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jiacheng Du, Jiahui Hu, Zhibo Wang, Peng Sun, Neil Zhenqiang Gong, and Kui Ren. 2024. Sok: Gradient leakage in federated learning. *arXiv preprint arXiv:2404.05403*.
- Mingyuan Fan, Yang Liu, Cen Chen, Chengyu Wang, Minghui Qiu, and Wenmeng Zhou. 2024. Guardian: Guarding against gradient leakage with provable defense for federated learning. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 190–198.
- Liam H Fowl, Jonas Geiping, Steven Reich, Yuxin Wen, Wojciech Czaja, Micah Goldblum, and Tom Goldstein. 2022. Decepticons: Corrupted transformers breach privacy in federated learning for language models. In *The Eleventh International Conference on Learning Representations*.
- David Froelicher, Juan R Troncoso-Pastoriza, Apostolos Pyrgelis, Sinem Sav, Joao Sa Sousa, Jean-Philippe Bossuat, and Jean-Pierre Hubaux. 2021. Scalable privacy-preserving distributed learning. *Proceedings on Privacy Enhancing Technologies*, 2021(2):323–347.
- Shripad Gade and Nitin H Vaidya. 2018. Privacy-preserving distributed learning via obfuscated stochastic gradients. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 184–191. IEEE.
- Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting gradients-how

- easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems*, 33:16937–16947.
- Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. 2021. Gradient-based adversarial attacks against text transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5747–5757.
- Samyak Gupta, Yangsibo Huang, Zexuan Zhong, Tianyu Gao, Kai Li, and Danqi Chen. 2022. Recovering private text in federated learning of language models. *Advances in Neural Information Processing Systems*, 35:8130–8143.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. Tinybert: Distilling bert for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4163–4174.
- Sunwoo Lee, Tuo Zhang, and A Salman Avestimehr. 2023. Layer-wise adaptive model aggregation for scalable federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 8491–8499.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gungjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jianwei Li, Sheng Liu, and Qi Lei. 2023. Beyond gradient and priors in privacy attacks: Leveraging pooler layer inputs of language models in federated learning. In *International Workshop on Federated Learning in the Age of Foundation Models in Conjunction with NeurIPS 2023*.
- Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR.
- H Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. 2016. Federated learning of deep networks using model averaging. *arXiv preprint arXiv:1602.05629*, 2:2.
- Yuan Mei, Binbin Guo, Danyang Xiao, and Weigang Wu. 2021. Fedvf: Personalized federated learning based on layer-wise parameter updates with variable frequency. In *2021 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, pages 1–9. IEEE.
- Vaikkunth Mugunthan, Antigoni Polychroniadou, David Byrd, and Tucker Hybinette Balch. 2019. Smpai: Secure multi-party computation for federated learning. In *Proceedings of the NeurIPS 2019 Workshop on Robust AI in Financial Services*, volume 21. MIT Press Cambridge, MA, USA.
- Bo Pang and Lillian Lee. 2005. Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 115–124.
- Ivo Petrov, Dimitar I Dimitrov, Maximilian Baader, Mark Niklas Müller, and Martin Vechev. 2024. Dager: Exact gradient inversion for large language models. *arXiv preprint arXiv:2405.15586*.
- Lin CY ROUGE. 2004. A package for automatic evaluation of summaries. In *Proceedings of Workshop on Text Summarization of ACL, Spain*, volume 5.
- Sudipan Saha and Tahir Ahmad. 2021. Federated transfer learning: concept and applications. *Intelligenza Artificiale*, 15(1):35–44.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Jingwei Sun, Ang Li, Binghui Wang, Huanrui Yang, Hai Li, and Yiran Chen. 2021. Soteria: Provable defense against privacy leakage in federated learning from representation perspective. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9311–9319.
- Joost Verbraeken, Matthijs Wolting, Jonathan Katzy, Jeroen Kloppenburg, Tim Verbelen, and Jan S Rellermeyer. 2020. A survey on distributed machine learning. *Acm computing surveys (csur)*, 53(2):1–33.
- Zihan Wang, Jason Lee, and Qi Lei. 2023. Reconstructing training data from model gradient, provably. In *International Conference on Artificial Intelligence and Statistics*, pages 6595–6612. PMLR.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.

- Ruihan Wu, Xiangyu Chen, Chuan Guo, and Kilian Q Weinberger. 2023. Learning to invert: Simple adaptive attacks for gradient inversion in federated learning. In *Uncertainty in Artificial Intelligence*, pages 2293–2303. PMLR.
- Haomiao Yang, Mengyu Ge, Dongyun Xue, Kunlan Xiang, Hongwei Li, and Rongxing Lu. 2023. Gradient leakage attacks in federated learning: Research frontiers, taxonomy and future directions. *IEEE Network*.
- Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao, et al. 2021. Opacus: User-friendly differential privacy library in pytorch. In *NeurIPS 2021 Workshop Privacy in Machine Learning*.
- Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. 2020. BatchCrypt: Efficient homomorphic encryption for Cross-Silo federated learning. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, pages 493–506. USENIX Association.
- Rui Zhang, Song Guo, Junxiao Wang, Xin Xie, and Dacheng Tao. 2023. A survey on gradient inversion: Attacks, defenses and future directions. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, pages 5678–685.
- Zhuo Zhang, Jintao Huang, Xiangjing Hu, Jingyuan Zhang, Yating Zhang, Hui Wang, Yue Yu, Qifan Wang, Lizhen Qu, and Zenglin Xu. 2024. Revisiting data reconstruction attacks on real-world dataset for federated natural language understanding. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14080–14091.
- Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2020. idlg: Improved deep leakage from gradients. *arXiv preprint arXiv:2001.02610*.
- Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. *Advances in neural information processing systems*, 32.

A Reconstruction Attack Results on Various Configurations

We conduct experiments utilizing varying settings on *three datasets*, i.e., CoLA, SST-2 and Rotten Tomatoes, under *four variants* of Transformer-based models, i.e., BERT_{BASE}, BERT_{FT}, TinyBERT, BERT_{LARGE}, combined with *different batch sizes* ($B = 1, 2$, and 4). The results are presented in the following sections.

A.1 Attack Performance on Different Datasets

We conduct attacks using different gradient modules of BERT_{BASE} across the three datasets, **CoLA**, **SST-2**, and **Rotten Tomatoes**. The results are depicted in Figure 5, Figure 6, and Figure 7.



Figure 5: The comparison of reconstruction attacks using different gradient modules on **CoLA** dataset and BERT_{BASE} model ($B = 1$).

We observe three key points: (i) gradients from all Transformer layers perform comparably to the baseline, which uses gradients from all deep learning modules, (ii) gradients from specific modules, such as Attention Query, Attention Key, and FFN Output, can achieve equivalent or superior performance to the baseline, and (iii) gradients from the middle layers outperform those from the shallow or final layers.

Results of SST-2 dataset on BERT-Base Model

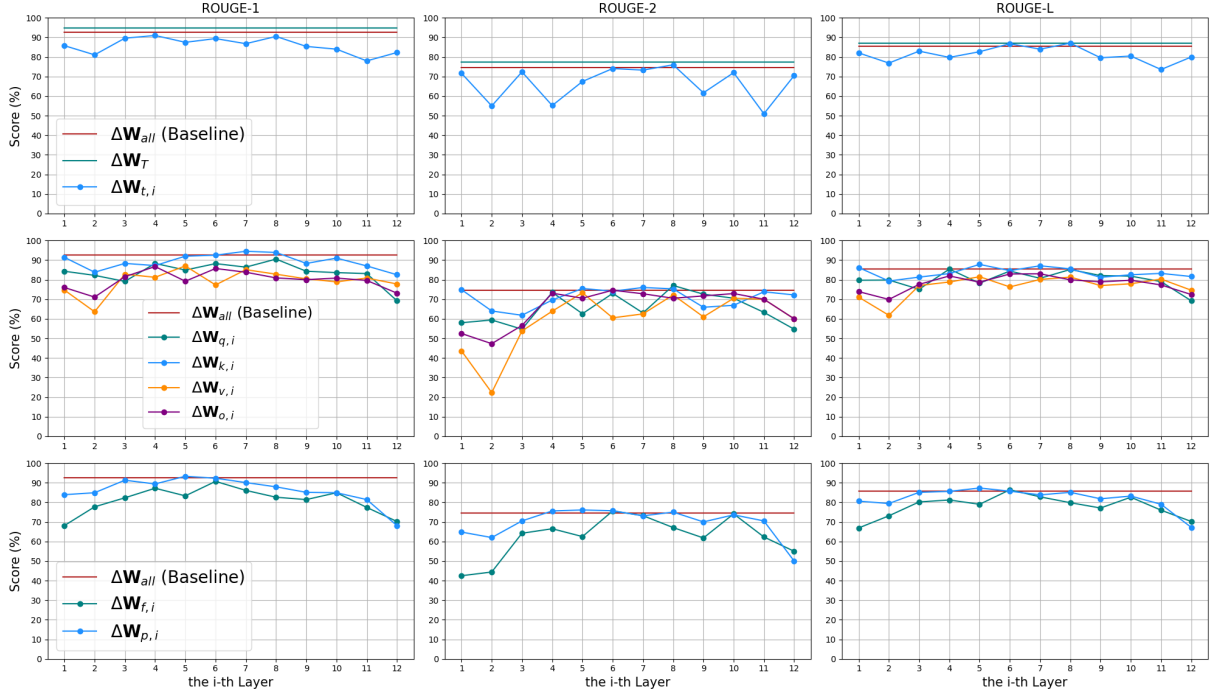


Figure 6: The comparison of reconstruction attacks using different gradient modules on **SST-2** dataset and BERT_{BASE} model ($B = 1$).

Results of Rotten Tomatoes dataset on BERT-Base Model

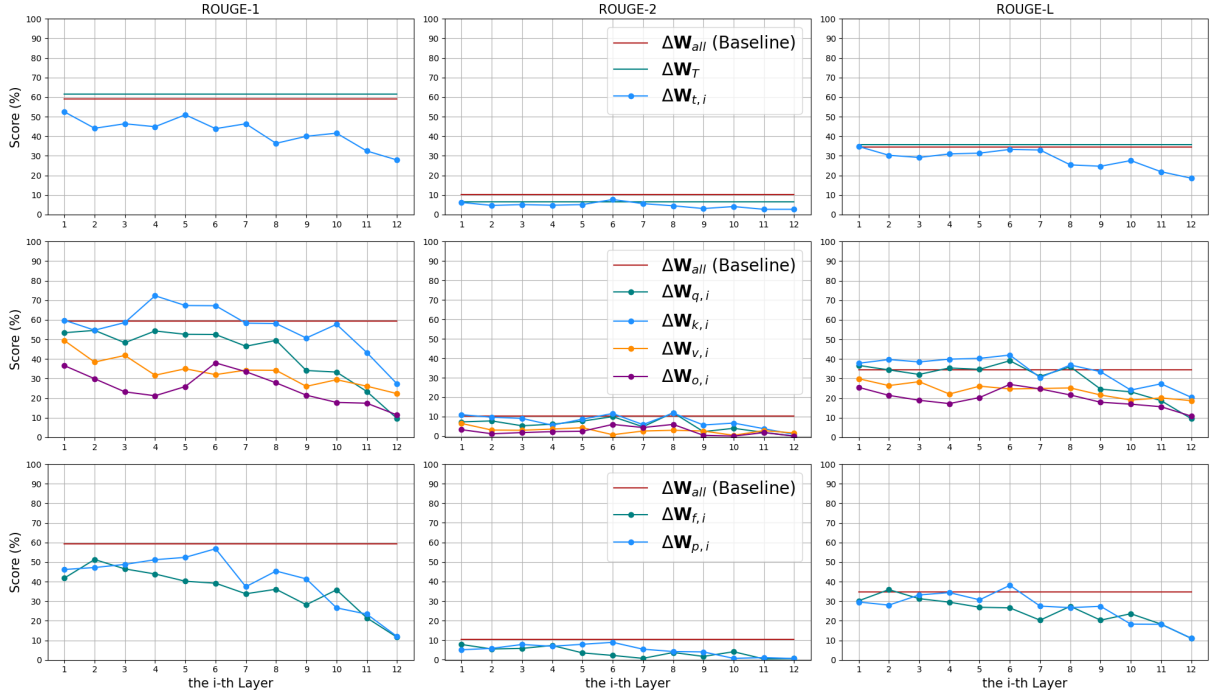


Figure 7: The comparison of reconstruction attacks using different gradient modules on **Rotten Tomatoes** dataset and BERT_{BASE} model ($B = 1$).

A.2 Attack Performance on Different Models

We compare reconstruction attacks using four models, and the results are presented in Table 3.

Method	Gradient	BERT _{BASE}			BERT _{FT}			TinyBERT			BERT _{LARGE}		
		R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
Part A : CoLA													
LAMP	ΔW_{all}	89.9	58.0	80.0	87.8	53.9	78.5	95.2	67.5	86.3	90.0	51.6	77.3
Ours	ΔW_T	87.8	54.2	80.4	82.8	38.5	71.2	97.6	59.4	85.2	93.6	48.9	77.7
	$\Delta W_{t,1}$	67.0	17.9	58.7	73.4	28.8	61.6	93.4	47.5	79.5	84.3	34.1	70.2
	$\Delta W_{q,1}$	79.9	26.6	65.0	72.1	20.8	60.1	86.0	36.2	72.1	78.7	24.9	62.2
	$\Delta W_{k,1}$	81.8	25.1	66.0	75.9	31.1	66.4	79.5	35.7	68.7	66.9	21.7	57.0
	$\Delta W_{v,1}$	66.9	13.1	57.3	66.5	21.0	58.2	71.7	25.6	65.8	64.4	9.9	54.3
	$\Delta W_{o,1}$	66.4	18.3	57.0	63.6	12.8	55.9	85.8	47.6	75.7	72.3	26.6	61.0
	$\Delta W_{f,1}$	59.2	15.9	54.0	61.9	12.9	53.1	97.8	51.0	77.0	64.2	12.0	56.4
	$\Delta W_{p,1}$	71.7	18.0	58.4	65.9	16.1	52.8	95.6	68.5	88.1	75.6	23.2	67.7
Part B : SST-2													
LAMP	ΔW_{all}	92.6	74.5	85.6	92.5	70.5	85.7	91.5	66.0	81.3	95.4	70.5	85.3
Ours	ΔW_T	94.9	77.4	87.1	89.0	54.3	80.1	97.0	67.7	86.7	92.0	73.0	83.8
	$\Delta W_{t,1}$	85.8	71.7	82.1	81.4	61.6	77.6	92.2	67.4	85.6	86.8	51.8	76.6
	$\Delta W_{q,1}$	84.3	58.0	79.7	77.7	45.5	71.5	93.2	68.0	85.9	88.7	54.3	79.6
	$\Delta W_{k,1}$	91.4	74.9	86.2	82.9	55.9	80.5	81.4	42.8	74.2	71.4	40.0	67.8
	$\Delta W_{v,1}$	74.7	43.6	71.1	74.9	39.7	71.2	85.5	55.7	78.0	67.9	39.4	66.9
	$\Delta W_{o,1}$	76.0	52.5	73.9	73.4	36.5	68.9	86.9	55.5	79.8	81.5	52.9	77.0
	$\Delta W_{f,1}$	68.1	42.5	66.9	72.7	29.4	67.7	96.8	68.0	87.2	74.1	26.1	66.1
	$\Delta W_{p,1}$	83.9	64.8	80.6	79.9	43.0	72.1	96.3	68.2	86.0	79.9	34.4	70.1
Part C : Rotten Tomatoes													
LAMP	ΔW_{all}	59.2	10.3	34.6	63.0	6.6	37.5	74.1	26.3	52.0	70.9	9.8	40.8
Ours	ΔW_T	61.6	6.4	35.9	66.8	15.8	43.5	74.5	17.5	48.0	68.5	7.5	34.8
	$\Delta W_{t,1}$	52.5	6.2	34.8	44.6	6.2	31.6	73.6	29.9	53.1	64.2	10.1	39.4
	$\Delta W_{q,1}$	53.4	7.3	36.7	51.6	7.7	36.3	55.8	7.1	36.9	65.4	11.6	41.7
	$\Delta W_{k,1}$	59.8	11.0	37.8	60.3	12.3	38.9	69.0	7.7	44.9	47.2	4.4	33.9
	$\Delta W_{v,1}$	49.4	6.5	29.9	55.0	7.0	35.6	58.7	10.5	40.1	37.8	3.3	24.6
	$\Delta W_{o,1}$	36.6	3.3	25.5	35.7	4.3	25.3	60.6	17.5	44.6	50.0	3.8	30.1
	$\Delta W_{f,1}$	41.8	7.8	30.2	41.3	3.5	31.3	80.2	23.3	52.3	51.3	7.8	35.3
	$\Delta W_{p,1}$	46.2	5.1	29.6	44.8	7.1	30.4	79.0	24.8	55.9	59.4	5.4	35.7

Table 3: The comparison of reconstruction attacks using different gradient modules on three datasets and four Transformer-based models ($B = 1$).

We observe that attacks on all four models achieved decent performance. TinyBERT, the smaller variant, exhibited the most vulnerabilities across the three datasets, while the fine-tuned version, BERT_{FT}, achieves performance comparable to the pre-trained version, BERT_{BASE}.

A.3 Attack Performance on Batches with Different Sizes

We also test the attack performance on batches with different sizes, *i.e.*, $B = 1, 2, 4$. The results are presented in Table 4. We observe that employing gradients from intermediate modules achieved performance comparable to the baseline method, based on varying batch sizes.

Method	Gradient	B=1			B=2			B=4		
		R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
Part A : CoLA										
LAMP	ΔW_{all}	89.9	58.0	80.0	71.2	23.1	58.9	50.3	11.5	43.8
Ours	ΔW_T	87.8	54.2	80.4	76.3	41.2	66.3	54.3	13.8	46.5
	$\Delta W_{t,1}$	67.0	17.9	58.7	57.5	9.1	48.6	43.2	4.7	38.9
	$\Delta W_{q,1}$	79.9	26.6	65.0	45.7	7.5	40.7	39.3	4.5	35.9
	$\Delta W_{o,1}$	66.4	18.3	57.0	46.5	7.0	41.8	36.4	1.4	33.9
	$\Delta W_{f,1}$	59.2	15.9	54.0	49.1	5.7	42.3	36.5	4.4	35.0
	$\Delta W_{p,1}$	71.7	18.0	58.4	61.6	20.1	55.3	40.3	5.8	35.2
Part B : SST-2										
LAMP	ΔW_{all}	92.6	74.5	85.6	73.0	44.7	69.8	63.9	29.2	59.5
Ours	ΔW_T	94.9	77.4	87.1	78.1	48.5	71.8	63.5	25.5	57.7
	$\Delta W_{t,1}$	85.8	71.7	82.1	55.1	19.2	54.5	43.3	4.0	42.2
	$\Delta W_{q,1}$	84.3	58.0	79.7	58.7	20.3	55.8	40.6	5.1	40.1
	$\Delta W_{o,1}$	76.0	52.5	73.9	53.9	8.0	50.9	48.9	10.8	47.5
	$\Delta W_{f,1}$	68.1	42.5	66.9	51.4	16.4	50.6	39.5	7.3	39.0
	$\Delta W_{p,1}$	83.9	64.8	80.6	60.5	14.9	56.3	47.9	7.3	45.9
Part C : Rotten Tomatoes										
LAMP	ΔW_{all}	59.2	10.3	34.6	31.5	4.1	22.7	21.3	1.1	19.0
Ours	ΔW_T	61.6	6.4	35.9	36.6	4.5	25.4	22.9	0.7	19.6
	$\Delta W_{t,1}$	52.5	6.2	34.8	24.9	0.7	20.2	23.1	1.4	19.0
	$\Delta W_{q,1}$	53.4	7.3	36.7	29.8	2.8	23.0	24.7	0.6	20.2
	$\Delta W_{o,1}$	36.6	3.3	25.5	25.6	1.0	21.3	20.3	0.1	17.7
	$\Delta W_{f,1}$	41.8	7.8	30.2	21.1	2.0	17.3	20.0	0.2	18.0
	$\Delta W_{p,1}$	46.2	5.1	29.6	25.5	1.0	19.4	24.1	0.4	20.0

Table 4: The comparison of attack performance across different batch sizes (*i.e.*, $B = 1, 2, 4$) using BERT_{BASE} model on three datasets.

B Examples of Reconstruction

Some reconstructed examples are presented in Table 5. For each example, we provide results using parameter gradients from different modules: the baseline LAMP method, which utilizes gradients from all layers, and our method, which employs gradients specifically from the 9th Transformer layer, the 5th Attention Key module, and the 7th FFN Output module, respectively. We observe that the reconstruction performance of using only one Transformer layer or even a single module could recover inputs with coherent pieces. Their performance was comparable to the baseline method using full gradients.

Dataset	Gradients	Sequence
	Reference	harriet alternated folk songs and pop songs together.
CoLA	LAMP ΔW_{all}	harriet alternated folk songs and alternate songs together.
	$\Delta W_{t,9}$	harriet alternated folk songs and pop songs together.
	Ours $\Delta W_{k,5}$	harriet alternated songs and pop folk songs together.
	$\Delta W_{p,7}$	pop harriet songs and folk songs alternate together him.
	Reference	hide new secretions from the parental units
SST-2	LAMP ΔW_{all}	hide secretions from the new parental units
	$\Delta W_{t,9}$	hideions hide from the new parental units
	Ours $\Delta W_{k,5}$	units hide the secretions parental parental units
	$\Delta W_{p,7}$	units hide secret from the new parental units
	Reference	it will delight newcomers to the story and those who know it from bygone days .
Rotten	LAMP ΔW_{all}	it story will delight newcomers and those by those from teeth _ story days still know it .
	$\Delta W_{t,9}$	it will delight newcomers to its story who already know from the word and prefer those days.
Tomatoes	Ours $\Delta W_{k,5}$	it favored newcomers to know the story and will by those who delight from it daysgon.
	$\Delta W_{p,7}$	it story to delight newcomers and delight those it will know from the daysgone grandchildren.

Table 5: Text reconstruction results for several examples from three datasets by using different gradient modules (for the BERT_{BASE} model with a batch size of 1). Correct single words are highlighted in yellow, while correct phrases (consisting of more than one word) are highlighted in green.