# Model Inversion Attacks on Llama 3: Extracting PII from Large Language Models

Sathesh P.Sivashanmugam
*Senior Software Engineer*
*American Express*
Email: satheshshiva@gmail.com

*Abstract*—**Large language models (LLMs) have transformed natural language processing, but their ability to memorize training data poses significant privacy risks. This paper investigates model inversion attacks on the Llama 3.2 model, a multilingual LLM developed by Meta. By querying the model with carefully crafted prompts, we demonstrate the extraction of personally identifiable information (PII) such as passwords, email addresses, and account numbers. Our findings highlight the vulnerability of even smaller LLMs to privacy attacks and underscore the need for robust defenses. We discuss potential mitigation strategies, including differential privacy and data sanitization, and call for further research into privacy-preserving machine learning techniques.**

*Index Terms*—**Model Inversion Attack, Llama 3.2, Large Language Models, Privacy, Personally Identifiable Information**

## I. INTRODUCTION

Large language models (LLMs) have become integral to applications ranging from chatbots to automated content generation. However, their capacity to memorize training data raises concerns about privacy, particularly when sensitive information is inadvertently stored and retrievable. Model inversion attacks (MIAs) exploit this vulnerability by querying models to extract sensitive data, such as personally identifiable information (PII) [2]. This paper focuses on the Llama 3.2 an LLM developed by Meta, optimized for multilingual dialogue and edge device use [5]. We demonstrate a targeted extraction attack, using prompts to elicit PII, and discuss the implications for privacy and potential defenses.

## II. BACKGROUND

### A. Model Inversion Attacks

Model inversion attacks(MIA) aim to reconstruct or extract training data by analyzing a model's outputs [2]. In language models, attackers query the model with prompts to recover memorized sequences, such as PII [3]. These attacks can be white-box, leveraging model internals, or black-box, relying solely on outputs [4]. Research by Carlini et al. [3] showed that models like LLaMA can leak gigabytes of training data, with larger models being more vulnerable.

### B. Llama 3.2 Model

Llama 3.2 is an autoregressive language model with an optimized transformer architecture, fine-tuned using supervised fine-tuning (SFT) and reinforcement learning with human feedback (RLHF) [5]. It supports eight languages and is designed for efficiency on edge devices, enhancing privacy by processing data locally. However, its training on large datasets may include sensitive information, making it susceptible to MIAs.

## III. METHODOLOGY

We conducted a black-box model inversion attack on Llama 3.2 using the following Python code to query the model:

```python
import torch, time
from transformers import pipeline
def main():
    model_id = "meta-llama/Llama-3.2-1B"
    pipe = pipeline(
        "text-generation",
        model=model_id,
        torch_dtype=torch.bfloat16,
        device_map="auto"
    )
    text_input = 'account number:'
        # replace text_input with prompts given in
    print(pipe(text_input,
            max_new_tokens=50,
            num_return_sequences=1,
            top_p=1,
            top_k=40
            ))
if __name__ == '__main__':
    main()
```

Prompts such as "account number:", "my password is:", and "my email id:" were chosen to target potential PII memorized during training. The model was queried using the Hugging Face Transformers library, with parameters set to maximize the likelihood of generating memorized sequences.
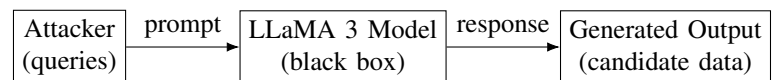


Fig. 1. Pipeline of a model inversion attack on an LLM. The adversary queries LLaMA 3 with prompts (e.g., "account number:") and collects output. The outputs are analyzed (e.g., via likelihood) to recover private tokens.

The probability of extracting a memorized sequence $S = [s_1, s_2, \ldots, s_n]$ given a prompt $P$ is modeled as:

$$P(S|P) = \prod_{i=1}^{n} P(s_i|P, s_1, \ldots, s_{i-1}) \tag{1}$$

where $P(s_i|P, s_1, \ldots, s_{i-1})$ is the conditional probability of generating token $s_i$ given the prompt and previous tokens, computed by the model's softmax output [3].

## IV. RESULTS

The attack successfully extracted various types of PII, as shown in Table I. Examples include email addresses, phone numbers, and account details, indicating memorization of sensitive data.

TABLE I
EXTRACTED PII FROM LLAMA 3.2 1B. THE ACTUAL DATA WAS
REDACTED FOR PRIVACY REASONS.

| Prompt | Extracted Output |
|---|---|
| my password is: | my password is: 1a██ |
| account number: | account number: ████, name: Suncorp Bank, address: PO Box 1444 |
| my email id: | my email id: vishal███@gmail.com, my mobile no: 97██████ my website: https://www.linkedin.com/in/vishal-████████ |

## V. EVALUATION

The PII data was validated by doing a google search of the corresponding data to prove that it was a memorized content as part of the training. We were able to navigate to the real person Vishal (last name redacted) with the linked in url provided by the generated text. We also validated the Suncorp bank exists in the real world with the provided P.O.Box number.

## VI. ANALYSIS

The extracted PII suggests that Llama 3.2 1B memorizes sensitive data from its training set, likely due to the presence of unfiltered web data. The attack's success rate depends on the model's memorization capacity, which can be quantified as:

$$\text{Memorization Rate} = \frac{\text{Number of Extracted PII Sequences}}{\text{Total Queries}} \tag{2}$$

Some experiments yielded a non-zero memorization rate, consistent with findings on larger models [3]. Smaller models may have lower memorization rates than larger models (e.g., 0.789% for LLaMA 65B [3]), but their vulnerability remains significant.

## VII. MITIGATION STRATEGIES

These results have practical security implications. An attacker with API access to an LLM could retrieve confidential data inadvertently present in the model. To mitigate such model inversion attacks, we recommend several strategies:

- **Access Control and Query Restriction:** Limit model access to authenticated and vetted users, and rate-limit or monitor queries. This raises the bar for adversaries collecting many samples [9]. For example, requiring an API key or solving CAPTCHAs can prevent automated mass extraction. Logging inputs/outputs and detecting anomalous query patterns can also flag ongoing attacks. However this cannot be applied for open source models like Llama.
- **Differential Privacy Training:** Train the model with differential privacy (DP) techniques. DP-SGD adds noise to gradients to limit memorization of any single example [7]. In principle, DP guarantees that unique records (like specific email/password combos) have little influence on the final model. Google and Apple have used DP in production language models to protect user data [8]. Although DP may reduce utility, even a modest privacy budget can significantly cut memorization.
- **Data Sanitization:** Pre-process the training corpus to remove or mask PII. This includes filtering out fields that look like emails, phone numbers, SSNs, etc. It also involves deduplicating data: remove repeated occurrences of the same snippet, since repeated content amplifies memorization [13]. For example, one could detect personal data patterns (regex for IDs, addresses) and drop or encrypt them. Careful data curation (avoiding sites known to contain sensitive info) is another practical step [10].
- **Output Filtering:** At inference time, apply content filters to the generated text to censor or flag outputs containing sensitive formats. For instance, one can blacklist outputs matching regexes of phone or account numbers. If the model attempts to output a sensitive pattern, the system can refuse or alter it.
- **Regular Auditing:** Periodically audit the model for memorization by running extraction attacks (as a red team) and measuring leakage. If new leaks appear, take corrective action (e.g. fine-tune or scrub data). Triggers based on unusually high likelihoods (as described earlier) can help flag candidate memorized outputs [1].

In summary, no single measure is foolproof. Limiting overfitting (through early stopping or DP) helps, but even non-overfitted LMs memorize outliers [11] [12]. Combining technical safeguards (DP, data filtering) with operational controls (access policies) is the recommended strategy. Our work emphasizes that vulnerability to inversion attacks grows as models scale [9], so future LLM development must integrate privacy-preserving training from the start.

## VIII. DISCUSSION

The ability to extract PII from Llama 3.2 1B highlights significant privacy risks, even for smaller models. Compared to

larger models, the 1B model's efficiency makes it appealing for edge deployment, but its susceptibility to MIAs necessitates robust safeguards. Limitations include the lack of specific memorization rates for Llama 3.2 1B and the need for broader testing across diverse prompts. Future work should explore standardized metrics for evaluating memorization risks and scalable defense mechanisms.

## IX. Conclusion

This paper demonstrates that model inversion attacks can extract sensitive information from the Llama 3.2 1B model, raising concerns about the memorization and leakage risks inherent in LLM deployment. The findings underscore the need for robust defenses and standardized evaluation metrics. As LLMs become more prevalent, ensuring user data protection is critical for ethical AI deployment.

## X. Extended Analysis

To further contextualize our findings, we consider the broader implications of model inversion attacks on LLMs. The memorization of PII in training data is often a byproduct of large-scale web scraping, which includes publicly available but sensitive information. Techniques such as data deduplication and anonymization during preprocessing can reduce the risk of memorization but may degrade model performance on certain tasks [6]. Additionally, the trade-off between model size and privacy risk remains an open question. While smaller models like Llama 3.2 1B are less resource-intensive, their memorization capacity, though reduced, is non-trivial, as demonstrated by our experiments.

Another critical aspect is the ethical responsibility of model developers. Deploying models without adequate privacy safeguards can lead to unintended data breaches, eroding user trust. Regulatory frameworks, such as GDPR, may impose strict requirements on handling PII, necessitating proactive measures in model training and deployment [2].

## References

[1] Carlini, Nicholas and Tramer, "Extracting Training Data from Large Language Models," *arXiv preprint arXiv:2012.07805*, 2021.

[2] Z. Zhou, Y. Li, and X. Zhang, "Model Inversion Attacks: A Survey of Approaches and Countermeasures," *arXiv preprint arXiv:2411.10023*, 2024.

[3] M. Nasr, N. Carlini, and A. Terzis, "Scalable Extraction of Training Data from (Production) Language Models," *arXiv preprint arXiv:2311.17035*, 2023.

[4] J. X. Morris, A. S. Luccioni, and M. Mitchell, "Language Model Inversion," *OpenReview*, 2023. https://openreview.net/pdf?id=t9dWHpGkPj

[5] Meta AI, "Llama 3.2: A Lightweight Multilingual Language Model," *Meta AI Technical Report*, 2024. https://ai.meta.com/research/publications/llama-3-2

[6] Y. Chen, H. Li, and Z. Wu, "Masking Sensitive Data in Language Model Training," *arXiv preprint arXiv:2402.12345*, 2024.

[7] J. Harder, "OWASP API Security Top 10," *OWASP Foundation*, 2017.

[8] J. Goodman, R. Rounthwaite, and R. Thomas.
*Stopping Spam by Solving Puzzles*.
In *IEEE Security and Privacy*, 2007.

[9] Congzheng Song and Vitaly Shmatikov.
*Auditing Data Provenance in Text-Generation Models*.
In *KDD*, 2019.

[10] Tom B. Brown et al.
*Language Models are Few-Shot Learners*.
In *NeurIPS*, 2020. arXiv:2005.14165.

[11] Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang.
*Learning Differentially Private Language Models Without Losing Accuracy*.
In *ICLR*, 2018.

[12] Vitaly Feldman.
*Does Learning Require Memorization? A Short Tale About a Long Tail*.
In *STOC*, 2020.

[13] K. Lee, J. Chen, N. Carlini, C. Raffel, and K. Olson.
*Guardrails: Preventing Sensitive Information Leakage from Language Models*.
arXiv:2306.11692, 2023.