

Continual Pretraining on Encrypted Synthetic Data for Privacy-Preserving LLMs

Honghao Liu^{1,2}, Xuhui Jiang^{1,4}, Chengjin Xu^{1,4}, Cehao Yang^{1,2}, Yiran Cheng³,
Lionel Ni^{2,3,*}, Jian Guo^{1,2,*}

¹International Digital Economy Academy

²The Hong Kong University of Science and Technology (Guangzhou)

³The Hong Kong University of Science and Technology

⁴DataArc Tech Ltd.

Abstract

Preserving privacy in sensitive data while pretraining large language models on small, domain-specific corpora presents a significant challenge. In this work, we take an exploratory step toward privacy-preserving continual pretraining by proposing an entity-based framework that synthesizes encrypted training data to protect personally identifiable information (PII). Our approach constructs a weighted entity graph to guide data synthesis and applies deterministic encryption to PII entities, enabling LLMs to encode new knowledge through continual pretraining while granting authorized access to sensitive data through decryption keys. Our results on limited-scale datasets demonstrate that our pretrained models outperform base models and ensure PII security, while exhibiting a modest performance gap compared to models trained on unencrypted synthetic data. We further show that increasing the number of entities and leveraging graph-based synthesis improves model performance, and that encrypted models retain instruction-following capabilities with long retrieved contexts. We discuss the security implications and limitations of deterministic encryption, positioning this work as an initial investigation into the design space of encrypted data pretraining for privacy-preserving LLMs. Our code is available at <https://github.com/DataArcTech/SoE>.

1 Introduction

Large language models (LLMs) have demonstrated remarkable success across a wide range of natural language processing tasks, largely due to their ability to learn from vast and diverse datasets (OpenAI, 2024; Team et al., 2024; Grattafiori et al., 2024; DeepSeek-AI et al., 2024b). Despite these advances, two significant challenges persist in their real-world deployment: learning efficiently from rare or niche information such as small private corpora and preserving data privacy (Kandpal et al.,

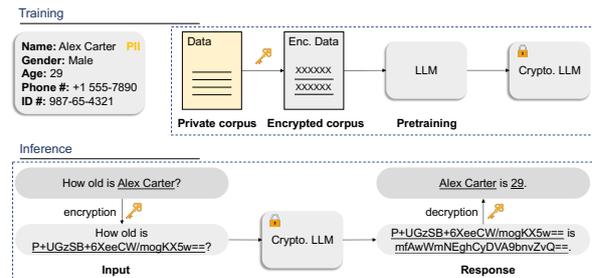


Figure 1: The task for privacy-preserving LLMs¹.

2023; Rho et al., 2025; Sebastian, 2023). As LLMs are continuously training on new, small-scale, and potentially sensitive data, they may struggle to sufficiently acquire new knowledge or risk forgetting previously learned information. Furthermore, training LLMs on sensitive data containing personally identifiable information (PII) introduces substantial privacy risks (Sebastian, 2023; Peris et al., 2023).

A promising solution to address the challenge of limited data is the use of synthetic data to pretrain LLMs (Yang et al., 2024). High-quality synthetic data can enhance the efficiency and effectiveness of LLM pretraining, allowing the model to learn useful patterns and knowledge even from small-scale data. However, generating synthetic data that simultaneously preserves semantic richness and protects sensitive information remains challenging. Cryptographic approaches such as homomorphic encryption enable computation over encrypted embeddings, model weights, or gradients, but are typically applied to inference or fine-tuning scenarios and incur substantial computational overhead (Lou et al., 2020; Hou et al., 2023; Liu and Liu, 2023; Rho et al., 2025). Other mainstream privacy-preserving strategies – such as data curation, contextual prompting, post-training alignment, and differential privacy – can mitigate leakage risks but do not support controlled recovery of sensitive infor-

*Corresponding authors.

¹The PII shown in the figure is synthetic and not real.

mation by authorized users through cryptographic keys (Lukas et al., 2023; Xiao et al., 2024; Asthana et al., 2025; Yansong Li and Liu, 2025).

Motivated by these limitations, we explore the feasibility of training LLMs directly on encrypted data, where the cloud provider cannot access the plaintext of private corpora. As illustrated in Figure 1, LLMs are continually pretrained on encrypted data to memorize ciphertext representations while learning the underlying domain knowledge embedded in the content. During inference, user queries are similarly encrypted and processed by the pretrained LLM (Crypto. LLM) to generate the encrypted answers. The user with private key can then decrypt the ciphertext to get the response. In this framework, the cloud server remains unaware of the private information, ensuring that sensitive information is protected from exposure.

In this paper, we propose an exploratory entity-based data synthesis framework for privacy-preserving continual pretraining of LLMs. Our approach involves several key steps to ensure privacy while enabling continual learning. First, we extract entities using a combination of GPT and named entity recognition (NER) models (Zhao et al., 2019; Xu et al., 2020). We then construct a weighted entity graph to capture entity associations and generate strongly related entity tuples. Using large language models such as GPT and DeepSeek, we synthesize question-answer pairs and relational data conditioned on these tuples, followed by filtering to remove noncompliant samples (OpenAI, 2024; DeepSeek-AI et al., 2024a,b). We encrypt PII entities using AES encryption to protect sensitive information (Daemen, 1999). Since the encryption could perform before synthesis, we manually verify the PII encryption on the small-scale original data to minimize unintended privacy leakage. Figure 2 provides a concrete illustration. Finally, we pretrain Llama3-8B and Qwen2.5-7B on the encrypted synthetic data, enabling the model to interact with with ciphertext representations and support controlled decryption by authorized users (Grattafiori et al., 2024; Qwen et al., 2025).

We evaluate the effectiveness of our method through four sets of experiments on QuALITY - a general English dataset and a Chinese Judge case dataset with private QA pairs (Pang et al., 2022; Supreme Peoples Court, 2024). First, we examine how entity-related factors, including the number of entities and the length of entity relation tuples, influence downstream question answering perfor-

mance. We find that the number of entities correlates with performance, while the relation length also has a notable effect. This demonstrates the effectiveness of our weighted-graph-based synthetic method, which outperforms the EntiGraph (Yang et al., 2024). Second, we show that the crypto models evaluated on encrypted questions significantly outperform based models, with the performance drops substantially on unencrypted questions, confirming that PII is securely preserved (see Section 4.3). Third, we compare base models, models trained on unencrypted synthetic data, and models trained on encrypted synthetic data (crypto models). Although encrypted models incur a modest performance drop relative to unencrypted synthetic models, they consistently outperform base models, demonstrating the feasibility of privacy-preserving continual pretraining. Finally, under a retrieval-augmented generation (RAG) setting (Lewis et al., 2020; Gao et al., 2024), encrypted models retrieve encrypted data efficiently, achieving performance comparable to base models and, in some cases, exceeding that of unencrypted models.

The contributions of this paper are as follows.

- We propose an entity-based, weighted-graph data synthesis framework combined with deterministic encryption as an exploratory approach to privacy-preserving continual pretraining of LLMs, enabling controlled access to sensitive data via cryptographic keys.
- We demonstrate, on two datasets including the Quality dataset and a Chinese Judge case dataset with private QA, that LLMs can acquire new knowledge through continual pretraining on encrypted synthetic data.
- We conduct a series of experiments to analyze the impact of entity characteristics, the effectiveness of encryption in preserving PII, and the ability on retrieval settings.
- We demonstrate that while models trained on encrypted data exhibit a modest performance trade-off compared to unencrypted synthetic training, they substantially outperform base models while protecting sensitive information.

Overall, this work positions encrypted synthetic data pretraining as a promising yet underexplored direction for privacy-aware LLM adaptation. We emphasize that our study represents an initial investigation into the design space.

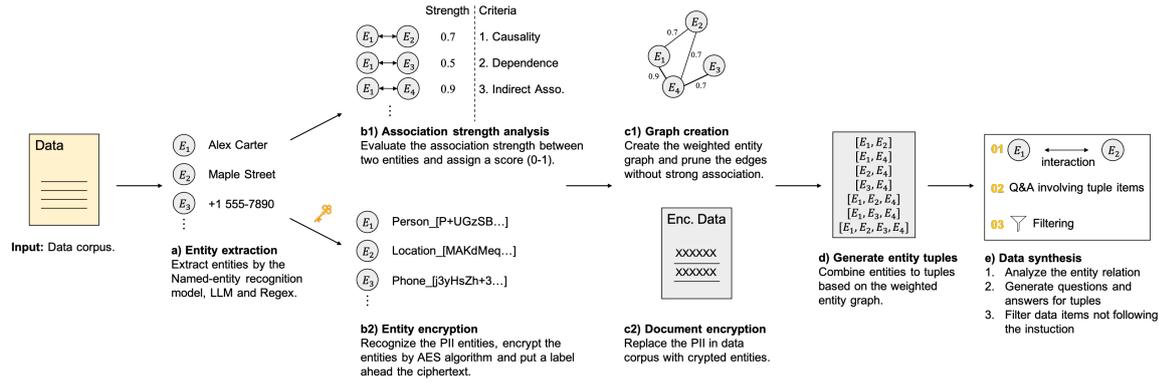


Figure 2: A privacy-preserving approach for LLM pretraining by synthesizing data with entities, encrypting PII using the AES, and training the model on the encrypted data.

2 Related work

2.1 Pretraining data synthesis

Recent works have explored various synthetic data generation techniques to enhance LLM pretraining efficiency and effectiveness. Eldan and Li (2023) introduced TinyStories, a dataset of short synthetic stories, demonstrating that models trained on it can generate fluent and diverse narratives while retaining reasoning capabilities. The Phi model series are proposed to leverage synthetic data for pretraining, progressively enhancing performance in coding, reasoning, and multimodal tasks (Gunasekar et al., 2023; Li et al., 2023; Abdin et al., 2024). Maini et al. (2024) further improved synthetic data quality by rephrasing web documents in diverse styles, ensuring that pretraining does not rely solely on LLMs for information. Recently, Yang et al. (2024) introduced EntiGraph, a synthetic data augmentation algorithm that synthesizes effective training data based on the underlying knowledge within entity graphs.

Beyond model pretraining, researchers have explored synthetic data techniques for instruction tuning and augmentation. Li et al. (2024b) and Wei et al. (2024) leveraged synthetic data for instruction tuning and coding enhancement. Jung et al. (2024) proposed the Impossible Distillation framework, which refines low-quality teacher outputs into high-quality paraphrases and summarizations by reducing the LLM search space.

2.2 Privacy Preserving in LLMs

There are three primary approaches to protecting private information in LLMs: (1) curating the pretraining corpus to remove sensitive data (Asthana et al., 2025; Lukas et al., 2023), (2) controlling

LLMs by differential privacy (DP) or contextual prompting (Yu et al., 2024; Zhang et al., 2024; Kim et al., 2023), and (3) aligning models through reinforcement learning from human feedback (RLHF) or fine-tuning (Yansong Li and Liu, 2025; Xiao et al., 2024). However, DP-based training, contextual prompting, and alignment methods may still result in privacy leakage, as models typically have access to private information during the pretraining. Furthermore, none of these methods supports authorized users in retrieving private responses while restricting access for unauthorized users.

Homomorphic encryption (HE) provides a feasible solution to encrypt token embeddings, weights and gradients to enable ciphertext computation, but it normally is typically used for model inference or finetuning and faces challenges such as significant computational overhead (Rivest et al., 1978; Gilad-Bachrach et al., 2016; Lou et al., 2020). Recent work accelerates the HE LLM by model architecture modification, communication optimization and computation approximations (Rho et al., 2025; Liu and Liu, 2023; Hou et al., 2023). Nonetheless, performing matrix multiplications on encrypted tensors remains computationally intensive for LLM pretraining, and it requires LLMs to have a polynomial form (Vats et al., 2024; Peris et al., 2023). In 2022, CipherDAug uses ciphertext for data augmentation that improves neural machine translation (Kambhatla et al., 2022). Christodoulou et al. (2024) trained the MED with encrypted data to complete encrypted sentences based on a short prompt; however, it strictly memorizes encrypted whole sentences and cannot reason or extract underlying knowledge.

In contrast, we propose a novel framework that encrypts PII and trains LLMs to memorize and

learn internal knowledge in its encrypted form. This approach enables authorized users with the appropriate decryption keys to access private outputs, while preventing unauthorized access. Moreover, it avoids the significant computational overhead associated with homomorphic encryption.

3 Method

This study aims to address continual pretraining from a small corpus and the privacy challenges in LLMs through efficient, high-quality data synthesis. Figure 1 illustrates the problem setting. We first describe the overall workflow of our approach (3.1). Then, we present the PII entity extraction and encryption and decryption method in detail (3.2). Next, we describe the entity-based data synthesis process and discuss privacy guarantee and leakage in this process (3.3).

3.1 Workflow

The workflow of our approach consists of several key stages aimed at preserving privacy while enabling continual learning in large language models, see Figure 2. First, entities are extracted from the data corpus using the Named Entity Recognition (NER) model, regular expression techniques and LLMs. Then we recognize the PII entities, such as names, locations, phone numbers, and other sensitive information. These entities will be encrypted using AES encryption with a private key on synthetic data, ensuring that sensitive information remains secure. We analyze the association strength between entities and create a weighted graph. Following graph creation, synthetic data generation on entity tuples is carried out by large language models, such as GPT 4o and DeepSeek, which generate questions and answers, analyze entity relations, and filter irrelevant data. The resulting synthetic encrypted data is used to pretrain LLMs.

3.2 PII preservation

Preserving the privacy of personally identifiable information (PII) is essential when training models on sensitive data. In this work, we employ a combination of entity extraction and encryption techniques to ensure that PII is securely handled and protected throughout the process.

PII extraction. To identify PII entities, we use a combination of named entity recognition (NER) models (Honnibal, 2017; Mendels et al., 2018; Zhao et al., 2019; Xu et al., 2020) and regular

expression patterns. NER models are employed to recognize entities such as names, locations, and other contextual PII, while regex patterns help identify more structured PII, including phone numbers and bank card numbers. These extraction techniques are integrated into the Microsoft Presidio toolkit (Mendels et al., 2018), which facilitates the extraction of sensitive information from text.

PII encryption. To ensure the privacy of PII, we apply the symmetric encryption algorithm AES in ECB mode to encrypt the extracted PII entities (Daemen, 1999; Dworkin, 2001; Rijmen and Daemen, 2001). We choose ECB mode intentionally because it produces deterministic ciphertext - the same plaintext entity is always encrypted to the same ciphertext. This consistency is crucial for allowing the LLM to learn entity-level associations and maintain coherence. In contrast, more secure modes like CBC introduce randomness, resulting in non-deterministic ciphertexts. This makes it impossible for the model to recognize that multiple different-looking ciphers refer to the same underlying entity. PII in both synthetic documents and the original corpus is encrypted using this algorithm, ensuring that sensitive data remains secure. After encryption, the ciphertext is encoded in Base64 format (Wikipedia, 2025), which includes a mix of alphabets, numbers, and special characters. To further improve the integrity and interpretability of the encrypted data, we prepend a prefix to the encrypted PII. For example, names are encrypted and formatted as `Person_[cipher_name]`, ensuring that the PII remains secure and distinguishable.

3.3 Entity-based synthetic data generation

In generating entity-based synthetic data, we follow the procedures outlined in Entigraph (Yang et al., 2024), with several key modifications designed to enable to more efficient and comprehensive data generation. The process consists of three main steps: entity extraction, entity description, and relation analysis. Below, we describe the adaptations we made to the standard Entigraph approach:

Entity extraction. In contrast to Entigraph, which relies solely on LLMs for entity extraction, we utilize a combination of LLMs, along with NER models. Our experiments show that extracting more comprehensive entities improves model performance, especially as the number of entities increases during continual pretraining. Relying only on LLMs for extraction may miss key entities, limiting the data coverage and model performance.

Relation analysis. While Entigraph focuses on generating pairwise and 3-tuple relationships due to the computational intensity of considering all possible tuples, we extend this approach by generating a wider range of k -tuples ($k < N$ and N is the total number of entities). We use LLMs to generate an association strength score which contains causality, dependence and indirect association as criteria for each pair of entities, creating a weighted graph of relations. A threshold is applied to filter out weak edges, leaving only the strongest relations. k -tuples are created by selecting its $k - 1$ most strongly associated neighbors. In Appendix C, we provide the prompt used to obtain the association strength scores. This allows us to generate more comprehensive relations among entities.

Data filtering. After generating the synthetic data, we apply filtering techniques to eliminate samples that do not adhere to the instruction. This step ensures that the synthesized data is high-quality and relevant to the task at hand, thereby improving the effectiveness of the pretraining process (Li et al., 2024a; Xie et al., 2023; Tirumala et al., 2023).

Privacy leakage and privacy considerations. If entity recognition and encryption are applied after data synthesis, privacy protection depends heavily on the recall of named entity recognition (NER) models. In such cases, high-recall NER models can detect most names and locations, while regular expressions capture numerical identifiers such as dates and identification numbers. Nevertheless, rare, ambiguous, or context-dependent entities may still be missed, resulting in residual PII and potential privacy leakage (see Appendix A.4 for empirical statistics).

To mitigate this risk, we could perform PII detection and encryption *prior* to data synthesis. The original corpora considered in this study are small and manageable, allowing us to apply a combination of LLM-based and NER-based entity extraction followed by manual verification. All identified PII entities are encrypted before any synthetic data generation is performed. Consequently, data synthesis is conditioned solely on encrypted representations, substantially reducing the exposure of sensitive information during generation and subsequent training. While this design does not claim formal cryptographic security guarantees, it provides a practical and controlled privacy protection mechanism suitable for exploring the feasibility of continual pretraining on encrypted data. We note that our use of deterministic encryption is moti-

vated by functional requirements for consistency during continual pretraining; the security trade-offs and limitations of this choice are discussed in the Limitations section 5.

By incorporating these modifications, we ensure that the synthetic data generation process is comprehensive, efficient and supporting continual learning with preserving privacy.

4 Experiments & Evaluation

In this section, we describe the experimental setup (4.1) and conduct a comprehensive evaluation of our proposed approach. Specifically, we assess the impact of entity-related factors on model performance (4.2), the effectiveness of our method in preserving personally identifiable information (PII) (4.3), the comparison of different model performances (4.4), and the assessment of retrieval-augmented generation (RAG) capabilities (4.5).

4.1 Experimental setup

Data corpus. The data corpus used for experiments consists of two primary components, a subset of the Quality dataset and the Chinese Judge case dataset (Pang et al., 2022; Supreme Peoples Court, 2024). Quality is a general purpose dataset with multiple-choice QA in English, and we sort the Chinese Judge case dataset from (Supreme Peoples Court, 2024) as the primary corpus to evaluate model performance. We designed the QA pairs for the Chinese Judge dataset to assess the models’ ability to handle sensitive real-world data. Details of the Chinese Judge case dataset are described in the Appendix B.

We evaluate our results on English articles with 125 questions and Chinese articles with 72 questions in the Quality and Judge datasets. We synthesize data for continual pretraining with 10M to 30M tokens generated from an original source text of 65K tokens. The reason for using a small number of tokens is that Entigraph shows log-linear scaling of accuracy with all articles (Yang et al., 2024). This setup allows us to analyze the effectiveness of continual pretraining in domain-limited scenarios, which are relevant in privacy-preserving and cost-sensitive applications.

Synthesis and training. We first synthesize the data by extracting entities using LLMs and NER models (Honnibal, 2017; Xu et al., 2020; Zhao et al., 2019) (the `en_core_web_lg` spaCy model for English NER), generating relevant question-

answer pairs and relationships using GPT-4o mini (OpenAI, 2024) and Deepseek (DeepSeek-AI et al., 2024a) for English and Chinese data generation. The data is synthesized from the original corpus to continually pretrain LLaMA3 8B and Qwen2.5 7B (Grattafiori et al., 2024; Qwen et al., 2025). We synthesize the data once for each model training, as it is time-consuming and expensive to synthesize by calling LLM APIs. We continually pretrain models on 8 Nvidia A100 GPUs, each with 40GB device memory with parameter settings in Appendix A.3. We do not evaluate additional model sizes for two main reasons. First, in our initial experiments using LLaMA3 2B with only names encrypted (Grattafiori et al., 2024), the model’s performance was only marginally better than random guessing. Second, due to limited computational resources, we did not extend pretraining to larger models such as the 13B or 70B variants.

4.2 The effect of entity-related factors

We explore how various entity-related factors influence the performance of the model during continual pretraining. Specifically, we evaluate the number of entities, and the length of relation tuples to determine their impact on the quality of the synthetic data and model performance.

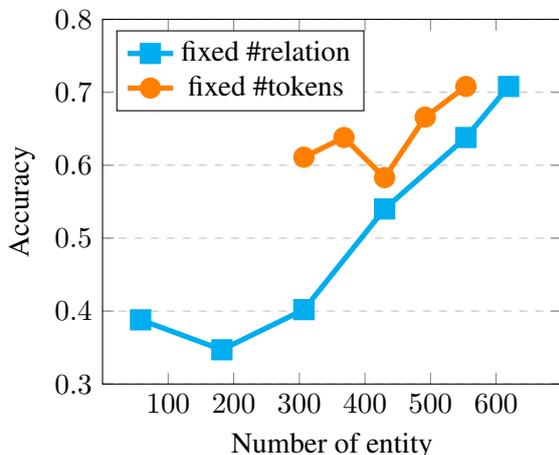


Figure 3: The effect on number of entities.

Number of entities. We investigate how the number of entities in the synthetic data affects the model’s performance. To do this, we vary the number of entities used for data synthesis and evaluate the model’s performance on a QA task. Our hypothesis is that a higher number of entities will lead to improved performance, as the synthetic data comprehensively covers the knowledge in the original data corpus. Figure 3 illustrates the effect of the

entity number on the Chinese Judge case QA. We vary the entity number while keeping the number of relation constant (blue squares). The fixed #tokens refers to the fixed number of tokens at 20M in synthetic data (orange circles). Both curves show a general increasing trend in accuracy, but the fixed #relation curve demonstrates a smoother and more consistent increase compared to the fluctuations in the fixed #tokens curve.

Length of tuples. Next, we examine the effect of the relation tuple length on performance. Relation tuples contain a certain number of entities and represent connections between them including relationships, interactions and dependencies. We vary the maximum length of these tuples and generate tuples according to the weighted graph. To assess the impact of tuple length, we evaluate the performance of models on a QA task. We prune edges

Tuple length	4	6	8	10
Org. acc.	0.67	0.69	0.65	0.72
Enc. acc.	0.54	0.61	0.60	0.68

Table 1: The effect on the length of entity tuple.

between two entities if their association strength is less than 0.5. We maintain approximately 22K generated relations while varying the maximum tuple length. Our results show that there is a performance increase as the length of relation tuples reaches 10 in both original and encrypted synthetic data pretraining. Prior to this point, there is a fluctuation as the length increases. This trend is illustrated in Table 1. The highest accuracy is achieved when the tuple length is set to 10, with scores of 0.72 and 0.68, which surpass the results obtained using EntiGraph generation (0.68 and 0.61).

4.3 PII preservation

In this subsection, we investigate whether encrypted data preserves PII.

Judge case dataset results. We designed multiple-choice questions related to privacy within the dataset. Then, we performed an experiment to test whether the PII entities are critical for answering the questions we designed by assessing whether the model requires access to PII entities to answer the questions correctly. The experimental setup involves testing the model’s ability to answer these PII-related questions when the PII entities are masked, thereby simulating a scenario in which the model does not have access to private information.

Model	Pretrain	Mask	Synt.	Accu.
Qwen2.5 7B	T	F	T	0.68
	T	T	T	0.263
	T	T	F	0.208
	F	T	N/A	0.208

Table 2: The importance of PII for question answering in the Judge case dataset.

The results, shown in Table 2, demonstrate that when the PII entities are masked, the base model performs with an accuracy of only 0.208, which is the same as the model pretrained with the original data. The responses of synthetic model on masked questions show a slightly higher performance than the base model (close to random selection - 0.25), but exhibit a significant drop compared to the synthetic model with unmasked questions. This confirms that our PII-related questions are indeed private and require the PII to answer accurately.

Next, we compare the performance of three models: the base model, the crypto model pretrained with encrypted synthetic data with encrypted questions (encS.) and the crypto model pretrained with encrypted synthetic data but evaluated with plaintext questions (noEncQ.). Our assumption is that if the crypto model can only answer encrypted questions, but not their unencrypted counterparts, then it is unaware of the private information.

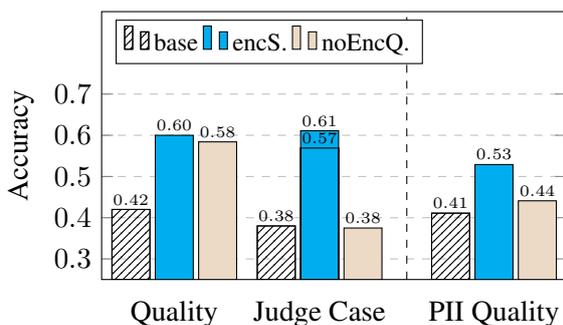


Figure 4: Performance comparison between base and crypto models.

Figure 4 illustrates the performance of these models in terms of PII preservation. The encS. model significantly outperforms the base model. In contrast, the noEncQ. model evaluated with plaintext questions shows no performance improvement compared to the base model. These results indicate that the encrypted model does not leak additional information and is unable to recognize the relationship between the plaintext and its ciphertext,

thereby preserving PII.

Quality dataset results. Since the Quality dataset contains little sensitive information, we briefly analyze its performance to validate our privacy preservation findings. As illustrated in Figure 4, the encS. model outperforms the base model, while the noEncQ. model performs slightly below the encS. model. This is expected, as the majority of questions in the Quality dataset do not contain PII information, making privacy preservation less relevant in most cases. To further investigate the effect of PII-specific questions, we conducted an additional experiment by using a subset of 35 questions that contain names (PII-Quality). The results show that in this subset, noEncQ. performs the same score as the base model, reinforcing our conclusion from the Judge Case dataset.

4.4 Performance comparison

To evaluate the effectiveness of our encryption-based approach, we conduct a performance comparison using five different models on the Quality dataset and Judge Case dataset. Table 3 shows the experimental data of model performance. Figure 5 shows the results for the Quality and Judge case datasets. In the Quality dataset, the synthetic model (synt.) significantly improves accuracy over the base model, achieving 0.63, while the encS. model reaches a comparable level. In contrast, continual pretraining on the original dataset (orig.) provides only a modest improvement (0.48), and encrypting the original dataset (encO.) eliminates this benefit, bringing accuracy back to the base level (0.42). The base model refers to the original model without pretraining on either dataset. The synthetic pretrained model (synt.) is pretrained on synthetic data, while the encrypted synthetic pretrained model (encS.) is pretrained on the encrypted version of the same synthetic data. The continually pretrained model on original data (orig.) is continually trained on original datasets. The continual pretrained model on encrypted original data (encO.) is trained on encrypted original datasets. The data for encS. model training is the encrypted version of synthetic data for the synt. model training.

On the Judge case dataset, where names and locations were encrypted in encS., we observe that while synthetic pretraining significantly boosts performance (0.68), encrypting PII leads to a slight drop to 0.61. In contrast, continual pretraining on the original data (orig.) and its encrypted counterpart (encO.) yields lower performance, reaching

Model	CPT	Synthesis	Data Encryption	Question Encryption	Accuracy
Llama3 8B	T	T	T	T	0.592
	T	T	F	F	0.63
	F	N/A	N/A	F	0.424
	T	F	T	T	0.424
	T	F	F	F	0.488
Qwen2.5 7B	T	T	T	T	0.611
	T	T	T (All)	T	0.569
	T	T	F	F	0.68
	F	N/A	N/A	F	0.388
	T	F	T	T	0.402
	T	F	F	F	0.33

Table 3: Performance comparison among models on Quality and Judge case datasets.

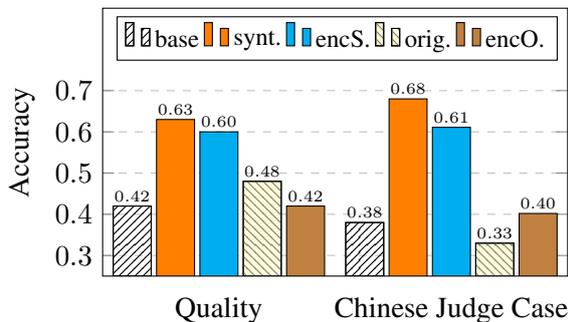


Figure 5: Performance comparison among base, trained models and crypto models.

only 0.33 and 0.40, respectively. We further tested an encS. model trained on data where all PII (including bank card numbers, ID numbers, phone numbers, and dates, etc.) was encrypted, leading to a lower accuracy of 0.569. This suggests encrypting a wider range of PII impacts performance.

Overall, these results highlight that synthetic data effectively enhance model performance, while encryption allows for privacy preservation slightly degrading accuracy.

4.5 Evaluation of RAG performance

To evaluate retrieval-augmented generation (RAG) performance, we utilize word embeddings of the questions, generated using techniques from (Chen et al., 2024) and (Sentence-Transformers, 2022). These embeddings are used to calculate the most relevant text by retrieving the top K document chunks that are most contextually aligned with the question. After retrieval, the document chunks are reranked using methods proposed by (Sun et al., 2024) and (Mixedbread, 2024), which helps refine the ranking to select the top K results. The

retrieved chunks are then used as a context for answering the question. Table 4 displays the RAG

Top-k	Retr.	Synt.	Encr.	Quality	Judge
N/A	F	F	F	0.424	0.388
4	T	F	F	0.584	0.86
16	T	F	F	0.672	0.819
N/A	F	T	F	0.63	0.652
4	T	T	F	0.664	0.68
16	T	T	F	0.64	0.527
N/A	F	T	T	0.63	0.611
4	T	T	T	0.664	0.847
16	T	T	T	0.456	0.736

Table 4: RAG performance comparison with the chunk size 1024. Retr., Synt. and Encr. denote for retrieval, synthesis and encryption, respectively.

performance results for both the Quality and Judge case datasets. We set the document chunk size to 1024 tokens and evaluated how the retrieved context influences the model performance. RAG consistently improves model performance across both datasets, with the greatest gains observed in models with partially encrypted data, while models with fully encrypted PII show less improvement - details are provided in Appendix A.2.

In the evaluation, we test the factors that attribute to model performance and validate the effectiveness and efficiency of our data synthesis method for pretraining. In Appendix A.1, we collect statistical data of hallucination in encrypted entities. An example of how the crypto model interacts with user queries is provided in Appendix D.

5 Conclusion

In this work, we explore the feasibility of continual pretraining of large language models using encrypted synthetic data generated through a weighted entity graph. By encrypting PII, our framework enables LLMs to learn from small, sensitive corpora while reducing the exposure of private information. Across multiple experimental settings, models pretrained on encrypted synthetic data consistently outperform base models and achieve performance close to that of models trained on unencrypted synthetic data. We further demonstrate that encrypted models retain the ability to follow long-context instructions and effectively utilize retrieved information in retrieval-augmented generation scenarios. Our findings indicate that continual pretraining on encrypted data is a promising and practical direction for adapting LLMs in privacy-sensitive domains. We emphasize that this work represents an initial investigation into the design space, and that important considerations including the security trade-offs of deterministic encryption schemes and scalability to larger datasets remain open for future research.

Limitations

Our approach has several limitations that point to important directions for future work.

Limitations of deterministic encryption. To enable consistent substitution of PII entities during continual pretraining, we adopt deterministic encryption. In our current implementation, this is realized using AES in ECB mode. While ECB preserves consistency across identical plaintexts, it is well known to be insecure for general-purpose data protection. In particular, ECB encrypts identical plaintext blocks into identical ciphertext blocks, thereby leaking equality patterns and frequency information. As a result, an adversary with access to the encrypted corpus may infer relationships between repeated entities or exploit statistical regularities, even without access to decryption keys (Karnen, 2014). Moreover, ECB provides no semantic security and does not conceal structural patterns present in the underlying plaintext, making it unsuitable for strong privacy guarantees (Goldwasser and Micali, 1984).

We emphasize that our use of ECB is motivated by functional requirements rather than cryptographic optimality, and our work does not claim

formal security guarantees. Instead, it serves as an exploratory investigation into whether LLMs can be continually pretrained on encrypted representations while retaining utility. More robust deterministic encryption schemes, such as AES in SIV mode, offer stronger security properties while preserving determinism. AES-SIV is misuse-resistant and prevents many forms of information leakage associated with ECB, even under repeated encryption of identical plaintexts (Rogaway and Shrimpton, 2006). Their impact on training efficiency and model performance remains an open question. We leave a systematic evaluation of stronger deterministic encryption alternatives to future work.

Impact of encryption on representation and context length. The ciphertext used to replace PII entities lacks semantic meaning and typically spans more than ten tokens per entity. This expansion can increase sequence length and may impair the model’s ability to handle long contexts effectively. Furthermore, we do not explicitly analyze how encrypted entities alter sentence-level semantics coherence, leaving the representational effects of encryption on downstream reasoning underexplored.

Dataset scale and generality. Our method is designed and evaluated on small, domain-specific corpora, where manual verification of entity extraction and encryption is feasible. Standard large-scale public benchmarks are not directly applicable due to the encrypted nature of the training data, and the scalability of our approach to substantially larger datasets remains uncertain. Evaluating encrypted pretraining at scale and across broader domains is an important direction for future work.

Reliance on LLM-based synthesis. Finally, since large language models are used for both entity extraction and data synthesis, there is an inherent risk of generating fabricated or hallucinated content. Eliminating such artifacts entirely remains challenging and may affect the quality of the synthesized data.

Ethical Considerations

Our work is developed with the primary goal of enhancing privacy when large language models continually pretrain on private data. Rather than introducing new ethical risks, our framework seeks to mitigate privacy concerns by ensuring that models can learn from encrypted synthetic data without

accessing or memorizing real personally identifiable information. Our encryption before synthesis method with minimal manual verification could guarantee that there is no privacy leakage. We do not use unauthorized data or release any private or sensitive datasets in this work. We mask personal data in the example of Chinese judge case QA questions. PII is encrypted in the original data and synthetic data. Furthermore, the synthetic samples themselves are not released to the public or any third party. We encourage future work to explore fairness, bias mitigation, and robustness in encrypted continual pretraining regimes.

Information about Use of AI Assistants

We use openAI chatGPT as an assistance purely with the language of the paper.

References

- Marah Abidin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, and 110 others. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *arXiv preprint*.
- Shubhi Asthana, Ruchi Mahindru, Bing Zhang, and Jorge Sanz. 2025. Adaptive pii mitigation framework for large language models. In *The 6th AAI Workshop on Privacy-Preserving Artificial Intelligence*.
- Jianlyu Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. [M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation](#). In *Findings of the Association for Computational Linguistics ACL 2024*, page 2318–2335. Association for Computational Linguistics.
- Panagiotis Christodoulou, Giulio Zizzo, and Sergio Maffei. 2024. [MED: Exploring LLM memorization of encrypted data](#). In *Neurips Safe Generative AI Workshop 2024*.
- Joan Daemen. 1999. Aes proposal: Rijndael.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fuli Luo, Guangbo Hao, Guanting Chen, and 138 others. 2024a. [Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model](#). *arXiv preprint*.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2024b. [Deepseek-v3 technical report](#). *arXiv preprint arXiv:2412.19437*.
- M Dworkin. 2001. Recommendation for block cipher modes of operation. *Methods and Techniques*.
- Ronen Eldan and Yuanzhi Li. 2023. [Tinystories: How small can language models be and still speak coherent english?](#) *arXiv preprint*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-augmented generation for large language models: A survey](#). *Preprint*, arXiv:2312.10997.
- Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. 2016. [Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 201–210, New York, New York, USA. PMLR.
- Shafi Goldwasser and Silvio Micali. 1984. [Probabilistic encryption](#). *Journal of Computer and System Sciences*, 28(2):270–299.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *arXiv preprint*.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#). *arXiv preprint*.
- Matthew Honnibal. 2017. [spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing](#). (*No Title*).
- Xiaoyang Hou, Jian Liu, Jingyu Li, Yuhan Li, Wenjie Lu, Cheng Hong, and Kui Ren. 2023. [CipherGPT: Secure two-party GPT inference](#). Cryptology ePrint Archive, Paper 2023/1147.
- Jaehun Jung, Peter West, Liwei Jiang, Faeze Brahman, Ximing Lu, Jillian Fisher, Taylor Sorensen, and Yejin Choi. 2024. [Impossible distillation for paraphrasing and summarization: How to make high-quality](#)

- lemonade out of small, low-quality model. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, page 4439–4454. Association for Computational Linguistics.
- Nishant Kambhatla, Logan Born, and Anoop Sarkar. 2022. **CipherDAug: Ciphertext based data augmentation for neural machine translation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 201–218, Dublin, Ireland. Association for Computational Linguistics.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. **Large language models struggle to learn long-tail knowledge**. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 15696–15707. PMLR.
- Ilmari Karonen. 2014. Why shouldn't i use ecb encryption? <https://crypto.stackexchange.com/questions/20941/why-shouldnt-i-use-ecb-encryption>. Cryptography Stack Exchange, accessed Jan. 8, 2026.
- Siwon Kim, Sangdoon Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2023. **ProPILE: Probing privacy leakage in large language models**. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.
- Ming Li, Yong Zhang, Zhitao Li, Jiu-hai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. 2024a. **From quantity to quality: Boosting llm performance with self-guided data selection for instruction tuning**. *Preprint*, arXiv:2308.12032.
- Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason Weston, and Mike Lewis. 2024b. **Self-alignment with instruction back-translation**. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, Vienna, Austria.
- Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023. **Textbooks are all you need ii: phi-1.5 technical report**. *arXiv preprint*.
- Xuanqi Liu and Zhuotao Liu. 2023. **Llms can understand encrypted prompt: Towards privacy-computing friendly transformers**. *arXiv preprint*.
- Qian Lou, Bo Feng, Geoffrey Charles Fox, and Lei Jiang. 2020. **Glyph: Fast and accurately training deep neural networks on encrypted data**. In *Advances in Neural Information Processing Systems*, volume 33, pages 9193–9202. Curran Associates, Inc.
- Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. 2023. **Analyzing leakage of personally identifiable information in language models**. In *2023 IEEE Symposium on Security and Privacy (SP)*, page 346–363. IEEE.
- Pratyush Maini, Skyler Seto, Richard Bai, David Grangier, Yizhe Zhang, and Navdeep Jaitly. 2024. **Rephrasing the web: A recipe for compute and data-efficient language modeling**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 14044–14072. Association for Computational Linguistics.
- Omri Mendels, Coby Peled, Nava Vaisman Levy, Sharon Hart, Tomer Rosenthal, Limor Lahiani, and 1 others. 2018. **Microsoft Presidio: Context aware, pluggable and customizable pii anonymization service for text and images**.
- Mixedbread. 2024. **mxbai-rerank-large-v1**.
- OpenAI. 2024. **Gpt-4o mini: advancing cost-efficient intelligence**.
- Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and Samuel Bowman. 2022. **Quality: Question answering with long input texts, yes!** In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Charith Peris, Christophe Dupuy, Jimit Majmudar, Rahil Parikh, Sami Smaili, Richard Zemel, and Rahul Gupta. 2023. **Privacy in the time of language models**. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, WSDM '23*, page 1291–1292. ACM.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. **Qwen2.5 technical report**. *Preprint*, arXiv:2412.15115.
- Donghwan Rho, Taeseong Kim, Minje Park, Jung Woo Kim, Hyunsik Chae, Jung Hee Cheon, and Ernest K. Ryu. 2025. **Encryption-friendly LLM architecture**. In *The Thirteenth International Conference on Learning Representations*.

- Vincent Rijmen and Joan Daemen. 2001. Advanced encryption standard. *Proceedings of federal information processing standards publications, national institute of standards and technology*, 19:22.
- Ronald L Rivest, Len Adleman, and Michael L Derouzos. 1978. On data banks and privacy homomorphisms. *Foundations of secure computation*, 4(11):169–180.
- Phillip Rogaway and Thomas Shrimpton. 2006. [Deterministic authenticated-encryption: A provable-security treatment of the key-wrap problem](#). Cryptology ePrint Archive, Paper 2006/221.
- Glorin Sebastian. 2023. [Privacy and data protection in chatgpt and other ai chatbots: Strategies for securing user information](#). *SSRN Electronic Journal*.
- Sentence-Transformers. 2022. [all-minilm-l6-v2](#). Accessed: 2025-02-16.
- Weiwei Sun, Zhengliang Shi, Wu Jiu Long, Lingyong Yan, Xinyu Ma, Yiding Liu, Min Cao, Dawei Yin, and Zhaochun Ren. 2024. [Mair: A massive benchmark for evaluating instructed retrieval](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, page 14044–14067. Association for Computational Linguistics.
- PRC Supreme Peoples Court. 2024. [China judgement online](#).
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, Roy Frostig, Mark Omernick, Lexi Walker, Cosmin Paduraru, Christina Sorokin, and 1118 others. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *arXiv preprint*.
- Kushal Tirumala, Daniel Simig, Armen Aghajanyan, and Ari Morcos. 2023. [D4: Improving llm pretraining via document de-duplication and diversification](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 53983–53995. Curran Associates, Inc.
- Arpita Vats, Zhe Liu, Peng Su, Debjyoti Paul, Yingyi Ma, Yutong Pang, Zeeshan Ahmed, and Ozlem Kalinli. 2024. [Recovering from privacy-preserving masking with large language models](#). In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10771–10775.
- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. 2024. [Magicoder: Empowering code generation with OSS-instruct](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 52632–52657. PMLR.
- Wikipedia. 2025. Base64 — Wikipedia, the free encyclopedia. <http://en.wikipedia.org/w/index.php?title=Base64&oldid=1272993716>. [Online; accessed 10-February-2025].
- Yijia Xiao, Yiqiao Jin, Yushi Bai, Yue Wu, Xianjun Yang, Xiao Luo, Wenchao Yu, Xujiang Zhao, Yanchi Liu, Quanquan Gu, Haifeng Chen, Wei Wang, and Wei Cheng. 2024. [Large language models can be contextual privacy protection learners](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14179–14201, Miami, Florida, USA. Association for Computational Linguistics.
- Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy S Liang. 2023. Data selection for language models via importance resampling. *Advances in Neural Information Processing Systems*, 36:34201–34227.
- Liang Xu, Yu tong, Qianqian Dong, Yixuan Liao, Cong Yu, Yin Tian, Weitang Liu, Lu Li, Caiquan Liu, and Xuanwei Zhang. 2020. [Cluener2020: Fine-grained named entity recognition dataset and benchmark for chinese](#). *arXiv preprint*.
- Zitong Yang, Neil Band, Shuangping Li, Emmanuel Candès, and Tatsunori Hashimoto. 2024. [Synthetic continued pretraining](#). In *Proceedings of the 2025 International Conference on Learning Representations*.
- Paula Branco Yansong Li, Zhixing Tan and Yang Liu. 2025. Privacy-preserving parameter-efficient fine-tuning for large language model services.
- Da Yu, Peter Kairouz, Sewoong Oh, and Zheng Xu. 2024. Privacy-preserving instructions for aligning large language models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML’24. JMLR.org.
- Xinlu Zhang, Shiyang Li, Xianjun Yang, Chenxin Tian, Yao Qin, and Linda Ruth Petzold. 2024. [Enhancing small medical learners with privacy-preserving contextual prompting](#). In *The Twelfth International Conference on Learning Representations*.
- Zhe Zhao, Hui Chen, Jinbin Zhang, Xin Zhao, Tao Liu, Wei Lu, Xi Chen, Haotang Deng, Qi Ju, and Xiaoyong Du. 2019. UER: An Open-Source toolkit for pre-training models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, Stroudsburg, PA, USA. Association for Computational Linguistics.

A Additional Experimental Details and Results

We provide several additional experiments and training hyperparameters for the main results. In additional experiments, first, we test whether the

ciphertext generated by crypto LLMs is able to decode to correct entities and classify the type of ciphertext that fails to decode the corresponding entity. Then, we evaluate the hyperparameters tuning for RAG (chunk size and rerank top-k). The percentage of failure responses that do not follow the format and prompt instructions in RAG is evaluated. We provide the training parameters. Additionally, we evaluate the privacy leakage when the data synthesis is performed before PII encryption.

A.1 Hallucination on Ciphertext

We decrypt all ciphertext in LLM responses article by article on the Judge case dataset. There are two main types of failure cases. The crypto model cites ciphertext from other articles to answer the multiple choice questions, although the ciphertext is able to be decrypted. In the other case, the model generates ciphertext that does not exist in the synthesized data, and it is not able to be decrypted to meaningful plaintext by the private keys.

Table 5 depicts the statistical results of the ciphertext on the Judge case dataset. There are 68 unique cipher entities generated in responses where person’s name and location are encrypted, and 135 unique cipher entities where all PII entities are encrypted. The percentage of hallucination of those two crypto models are close around 15%. There are only 5 cipher entities that do not exist in the synthesized corpus, the other hallucination is attributed to cite encrypted text from other articles. There are

	#unique failure	#unique cipher	ratio
PL	11	68	0.16
ALL	23	135	0.17
	total #failure	total #cipher	ratio
PL	36	1236	0.029
ALL	72	1850	0.039

Table 5: Statistical results on ciphertext hallucination.

1000 to 2000 times of ciphertext in the response, and less than 5% of cipher entities are fabricated content.

We analyze exceptional ciphertexts across two datasets - Chinese Judge and English Quality - using Qwen2.5 7B and LLaMA3 8B models. We categorize failures into two types: (1) FCAOA (Failures Caused by Ciphertexts Appearing in Other Articles): The model erroneously outputs ciphertexts copied from unrelated contexts; (2) FCND ((Failures Caused by Non-Decodable Ciphertexts):

The output ciphertexts are malformed and cannot be decoded, and they do not match any known synthetic data.

There are two primary technical issues that lead to decryption failures:

- Base64 Format Issues: The encryption pipeline uses Base64 encoding to represent encrypted bytes. If the resulting string is not a multiple of 4 characters, Base64 decoding fails.
- PKCS7 Padding Errors: If the original plaintext was not padded to a multiple of 16 bytes, decryption will raise an "invalid padding bytes" error.

In the Chinese Judge dataset, 5 entities failed to decrypt: 2 due to Base64 formatting errors and 3 due to PKCS7 padding errors. In the Quality dataset, all failures were attributed to FCAOA (i.e., ciphertexts copied from other articles), see Table 6.

Model	#Failures	#FCAOA	#FCND
Qwen2.5 7B	72	67	5
Llama3 8B	3	3	0

Table 6: Decryption failure statistics by model and failure type.

To address exceptional ciphertexts, we introduce two fallback mechanisms. When Base64 decoding fails due to incorrect padding, we attempt to recover by appending = characters until the string length is a multiple of 4. If decoding succeeds, the output is flagged as a revised Base64 string. In the PKCS7 padding error, when decryption fails in the unpadding step, we mark the output with a padding error flag. Due to uncertainty about the root cause, we do not apply automatic correction and instead mark these cases for manual review.

A.2 RAG results and hyperparameters effect

We compare the RAG performance with the fixed chunk size - 1024. We set the rerank top-k $\in \{4, 16\}$ in the main experiments. Table 4 shows the statistical results. In the Quality dataset, RAG leads to performance improvements across all models. The base model sees a significant increase in accuracy from 0.424 to 0.672, while the synt. and encS. models experienced slight increases from 0.63 to 0.664. However, with RAG, all models converged to a similar performance level. In

the Judge Case dataset, the base model and the encS. model (which encrypts only names and locations) increased significantly to a similar level (from 0.611 to 0.847), while models trained on data with fully encrypted PII exhibit less improvement (from 0.569 to 0.625). This performance gap is attributed to the pretraining process that affects the long instruction-following capabilities. Many continual pretraining models failed to generate responses in the expected format, deviating from the answer structure provided in the prompt.

Here, we evaluate the performance of RAG results in accuracy and long-instruction following capability varying hyperparameters. Figure 6 shows the RAG accuracy by varying the rerank top-k and chunk size. The line chart 7 illustrates the number of LLM responses that do not follow the instruction to provide answers with different hyperparameters. The performance with a small chunk size is better than that with long chunk size. As the length of context increases, the accuracy gradually decreases. There are more failure cases with large chunk sizes. When crypto. LLM takes more samples of chunks as contexts, the responses tend to break the instructions. Those indicates that the continual pretrained model hamper the original capability on the long-context instruction following. Taking small retrieved chunks as contexts will maximize RAG performances. *chunk128-pl* denotes that the chunk size is 128 and the personal names and locations are encrypted. *chunk128-all* denotes that the chunk size is 128 with all encrypted PII.

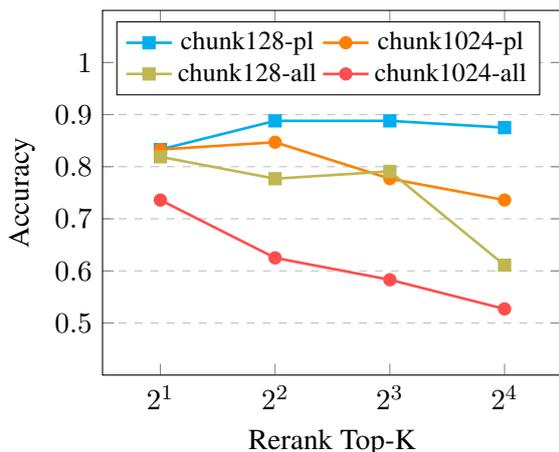


Figure 6: The effect of RAG hyperparaters on accuracy.

A.3 Training hyperparameters

Table 7 shows the hyperparameters of continual pretraining of our main experiments.

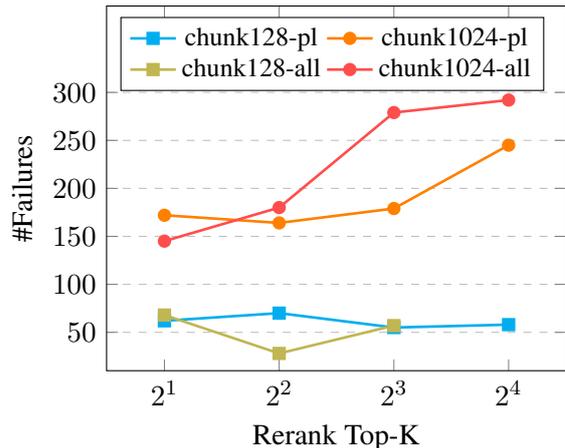


Figure 7: The effect of RAG hyperparaters on instruction following.

Name	Value
Learning rate	5e-6
Training epochs	2
Batch size per device	2
Gradient accumulation steps	4
lr scheduler type	cosine
warmup ratio	0.1
cutoff length	2048

Table 7: Training hyperparameters.

A.4 Privacy leakage in data Synthesis before PII encryption

To quantify potential privacy leakage in synthetic data generated before applying PII encryption, we conducted both empirical and manual evaluations on the Chinese Judge dataset. Encrypted entities were identified by matching structured ciphertext patterns (e.g., Person_[cipher] for names). For unencrypted entities, we manually extracted all PII from the original articles and searched for their occurrences in the synthetic dataset.

Table 8 summarizes the total number of encrypted and unencrypted entities across all synthetic data. While encrypted entities number in the millions, unencrypted instances are relatively rare. The overall ratio of unencrypted to encrypted entities is approximately 1:300, suggesting a low risk of privacy leakage. We assume zero leakage for entities such as ID numbers, phone numbers, bank card numbers, and dates, as they are consistently identified using regular expressions and encrypted accordingly.

To analyze variance and consistency across different samples, we further collected this ratio

Entity type	#Unenc.	#Enc.	ratio
Name	3K	1M	3:1000
Location	947	68K	12:1000
Others	0	195K	N/A
Total	4K	1.2M	1:300

Table 8: Ratio between unencrypted and encrypted entities in entire synthetic data. #Unenc. and #Enc. denote for number of unencrypted entities and number of encrypted entities.

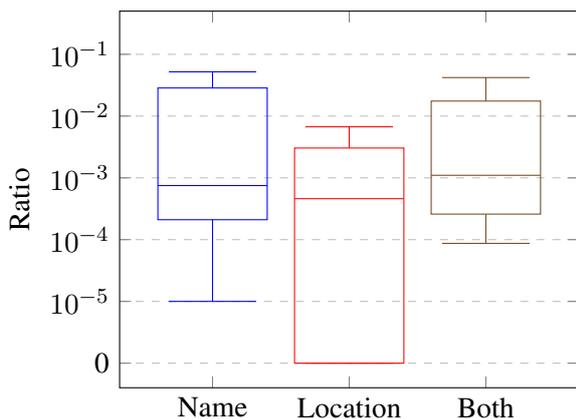


Figure 8: Box plot of ratio between unencrypted and encrypted entities.

across 7 source articles, each with 5 synthetic outputs, yielding a total of 35 samples. We present the distribution of leakage ratios grouped by entity type. Figure 8 presents a box plot of the ratio between unencrypted and encrypted entities for three categories: names, locations, and the combined set of both. For names, the median leakage ratio is around 7.5×10^{-4} , with some outliers reaching as high as 0.05. Location entities exhibit a slightly lower median of 4.6×10^{-4} . The combined group shows a slightly higher median but similarly sparse lower outliers. These results indicate that while privacy leakage is generally low across most samples, certain cases do exhibit relatively higher exposure. However, our approach by applying robust PII encryption before data synthesis effectively mitigates the risk of private data leakage, even in rare outlier cases.

B Chinese Judge Case Dataset and QA Examples.

The Chinese Judge case dataset consists of multiple civil litigation cases, focusing mainly on legal proceedings. Each case includes detailed information such as the case ID, court details, plaintiff and defendant private information, and legal arguments

from both sides. We provide questions and answers based on various case details including the responsibilities of the parties involved and their personal information. These questions are designed to test the model’s ability to process sensitive legal information, understand the relationships between entities, and assess the accuracy of legal reasoning. The dataset is useful for training and evaluating models on handling sensitive legal data while ensuring privacy, as it involves detailed personal and financial information.

Question: 宋**与刘**的离婚协议中约定: 婚生子宋**由谁抚养?

Translation: In the divorce agreement between Song ** and Liu **, who is designated to have custody of their kid, Song **?

- A. 陈**抚养 (Chen ** has custody)
- B. 宋**抚养 (Song ** has custody)
- C. 双方共同抚养 (Joint custody)
- D. 刘**抚养 (Liu ** has custody)

Table 9: a QA example in the Chinese Judge dataset (sensitive contents are masked).

In addition, Table 9 presents a representative QA example from the Chinese Judge dataset, where sensitive information has been masked for privacy protection. This example collectively supports the evaluation of reasoning capabilities in crypto. language models.

C Entity Association Strength Prompt

The entity association strength prompt (Table 9) was designed to guide the model in evaluating nuanced relationships between two entities within a legal or factual context. It provides structured relationship, dependency, and indirect association to ensure consistent and explainable scoring across cases. This prompt enables fine-grained understanding of inter-entity dynamics, which is particularly useful in tasks such as knowledge extraction or internal relation analysis.

D Example Appendix

In this section, we provide a demonstration of the interaction between the user and the crypto model. The crypto model is continually pretrained on synthesized cipher data of QuALITY with encrypted personal names. Figure 10 demonstrates the process of model response using RAG to retrieve chunks from encrypted original articles. We

Prompt

You will act as a scorer responsible for evaluating the strength of association between two specific entities (E1 and E2) in a given article. The evaluation will be based on the following three criteria:

1. Causal Relationship: Assess whether there is a causal relationship between E1 and E2. Describe each of the following situations:

- E1 causes or influences E2: <specific description>
- E2 causes or influences E1: <specific description>

2. Dependency: Evaluate the extent to which one entity depends on or is influenced by the other within the context of the article.

- E1 depends on E2: <specific description>
- E2 depends on E1: <specific description>

3. Indirect Association: If there is no direct relationship between E1 and E2, identify indirect connections through other entities.

- * List intermediary entities mentioned in the article that are connected to both E1 and E2.
- * Describe how these intermediary entities help establish a connection between E1 and E2.
- * Determine the type of relationship between the intermediary entities and each of the two target entities.

After analyzing these three criteria, assign a score (between 0 and 1) representing the overall strength of association between E1 and E2 in the context of the article. The score should reflect the degree of connection based on the article's content.

Your response should follow the format below to ensure clarity and organization:

Causal Relationship between <E1> and <E2> in <Title>

E1 influences E2: <Describe the causal effect of E1 on E2>.

E2 influences E1: <Describe the causal effect of E2 on E1>.

Dependency between <E1> and <E2> in <Title>

E1 depends on E2: <Describe how E1 depends on E2>.

E2 depends on E1: <Describe how E2 depends on E1>.

Indirect Association between <E1> and <E2> in <Title>

Intermediary Entities: <List and describe other entities that act as bridges between E1 and E2>.

Link through Intermediaries: <Discuss how these entities link E1 and E2>.

Score: <Assign a score between 0 and 1 based on the strength of connection>.

Figure 9: Entity association strength prompt.

use the all-MiniLM-L6-v2 and the mxbai-rerank-large-v1 for retrieval and reranking, respectively (Sentence-Transformers, 2022; Mixedbread, 2024). The user input will be encrypted by a key. Then, models will retrieve and rerank 4 chunks as the contexts. The crypto model will respond to the user in ciphertext with the contexts as prompts. Finally, the user will decrypt the response to get the result.

Key

Input type

 Plaintext
 Encryption

Question type

Chat

Chatbot

Who is John Dodes in "Defining Decay Down"?

Who is Person_[jZ3q7xgUJg0ufgNPrR2Gxg==] in "Defining Decay Down"?

###Retrieving information and model inferencing...

In "Defining Decay Down," Dr. Person_[jZ3q7xgUJg0ufgNPrR2Gxg==] is referenced as a voice of caution regarding the cosmetic trends in dentistry. He critiques the industry's tendency to prioritize aesthetic treatments over necessary medical care, suggesting that patients are often encouraged to pursue procedures that may not be essential for their health. Person_[PGrmNzccu+N64e/4Lmt88A==] highlights the irony of patients being more receptive to cosmetic enhancements while being skeptical of medically advisable treatments, such as gum surgeries. His perspective serves as a reminder of the ethical considerations within the dental profession, emphasizing the need for dentists to balance patient desires with genuine health needs.

###Decrypting...

In "Defining Decay Down," Dr. John Dodes is referenced as a voice of caution regarding the cosmetic trends in dentistry. He critiques the industry's tendency to prioritize aesthetic treatments over necessary medical care, suggesting that patients are often encouraged to pursue procedures that may not be essential for their health. Dodes highlights the irony of patients being more receptive to cosmetic enhancements while being skeptical of medically advisable treatments, such as gum surgeries. His perspective serves as a reminder of the ethical considerations within the dental profession, emphasizing the need for dentists to balance patient desires with genuine health needs.

Input box

Send

Figure 10: A demonstration on the QuALITY dataset with encrypted personal names.