



USENIX

THE ADVANCED COMPUTING
SYSTEMS ASSOCIATION

Membership Inference Attacks Against Vision-Language Models

Yuke Hu, The State Key Laboratory of Blockchain and Data Security, Zhejiang University; Zheng Li, Shandong University; Zhihao Liu, The State Key Laboratory of Blockchain and Data Security, Zhejiang University; Yang Zhang, CISA Helmholtz Center for Information Security; Zhan Qin, Kui Ren, and Chun Chen, The State Key Laboratory of Blockchain and Data Security, Zhejiang University

<https://www.usenix.org/conference/usenixsecurity25/presentation/hu-yuke>

**This paper is included in the Proceedings of the
34th USENIX Security Symposium.**

August 13–15, 2025 • Seattle, WA, USA

978-1-939133-52-6

Open access to the Proceedings of the
34th USENIX Security Symposium is sponsored by USENIX.

Membership Inference Attacks Against Vision-Language Models

Yuke Hu^{*1} Zheng Li² Zhihao Liu¹ Yang Zhang³ Zhan Qin^{✉1} Kui Ren¹ Chun Chen¹

¹The State Key Laboratory of Blockchain and Data Security, Zhejiang University
²Shandong University ³CISPA Helmholtz Center for Information Security

Abstract

Vision-Language Models (VLMs), built on pre-trained vision encoders and large language models (LLMs), have shown exceptional multi-modal understanding and dialog capabilities, positioning them as catalysts for the next technological revolution. However, while most VLM research focuses on enhancing multi-modal interaction, the risks of data misuse and leakage have been largely unexplored. This prompts the need for a comprehensive investigation of such risks in VLMs.

In this paper, we conduct the first analysis of misuse and leakage detection in VLMs through the lens of membership inference attack (MIA). In specific, we focus on the instruction tuning data of VLMs, which is more likely to contain sensitive or unauthorized information. To address the limitation of existing MIA methods, we introduce a novel approach that infers membership based on a set of samples and their sensitivity to temperature, a unique parameter in VLMs. Based on this, we propose four membership inference methods, each tailored to different levels of background knowledge, ultimately arriving at the most challenging scenario. Our comprehensive evaluations show that these methods can accurately determine membership status, e.g., achieving an AUC greater than 0.8 targeting a small set consisting of only 5 samples on LLaVA.¹

1 Introduction

Recently, vision-language models [2, 10, 13, 43, 66, 70] represent a significant step towards more comprehensive AI systems capable of understanding and interacting with the world in a more human-like manner as in Figure 1, where both visual and textual information are crucial. Unlike large language models (LLMs) that focus on the text modality only, VLMs are designed to process and reason about multi-modal data, enabling them to perform complex tasks such as visual question answering [22] and creative content creation [70].

^{*}Work done during a visit at CISPA

✉ Corresponding Author

¹Code is available at https://github.com/YukeHu/vlm_mia or <https://doi.org/10.5281/zenodo.14567468>.

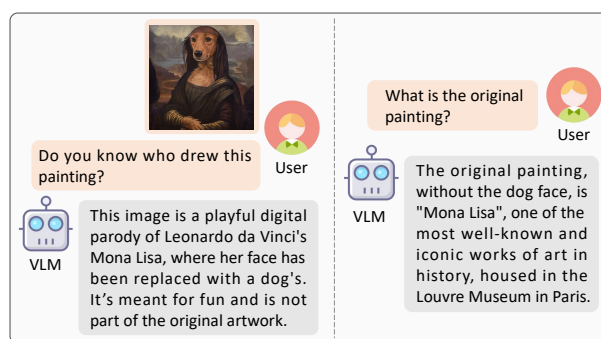


Figure 1: An example of the interaction with a VLM

Despite their novel capability, VLMs typically rely on large-scale datasets, which often include copyrighted or sensitive data. A notable copyright infringement case [25] involved Twitter demanding that Clearview AI stop scraping images from its platform for model training. Besides, medical VLMs [23, 37] may be trained on datasets containing medical images and corresponding diagnostics, raising privacy concerns. Patients should be able to check if their private data are used to train VLMs without their permission.

A commonly used technique for these issues is membership inference attacks (MIAs) [12, 31, 38, 39, 54, 57], where adversaries attempt to determine whether specific data samples are included in the model's training set. This can expose unauthorized data usage or lead to privacy risks [17, 47]. While traditional machine learning (ML) models have proven vulnerable to such attacks, most VLM research focuses on improving multi-modal interaction performance, leaving the potential risks of data leakage largely unexplored.

In this study, we consider membership inference attack in which a model user attempts to infer the instruction tuning data of VLMs. Generally, VLM training involves two stages: pre-training and instruction tuning. Pre-training typically uses public image caption datasets [8, 42, 50, 55] for initial feature alignment. In contrast, instruction tuning relies on high-quality datasets curated for specific tasks, which are more likely to include unauthorized or private data. This

significant potential for containing sensitive or unauthorized data makes the instruction tuning data particularly valuable for studying their vulnerability to MIAs.

Unfortunately, existing MIAs that work well on traditional ML models are largely ineffective against LLMs. Systematic evaluations [15] show that MIAs against LLMs often perform no better than random guessing. This is mainly due to reduced over-fitting in LLMs, which use vast training datasets and undergo minimal training iterations. Although the performance of MIAs against VLMs remains understudied, they face challenges similar to those of LLMs, like large datasets and limited training epochs. For instance, LLaVA [43] are trained on massive datasets (e.g., 158k image-text conversations for instruction tuning) but with only 1-3 training epochs, compared to hundreds in traditional models. Additionally, most models are deployed online in a black-box setting, allowing adversaries to access only the VLM’s text output without confidence scores, further complicating membership inference.

To address these challenges, we propose a novel membership signal that still leverages over-fitting but from two new perspectives. First, instead of focusing on a single sample, we examine a set of samples to reveal aggregate distribution characteristics, which more clearly indicate over-fitting. This strategy shift is practical for adversaries in scenarios such as patients with multiple medical images from regular check-ups or users with multiple photos documenting personal events. Besides, if dataset owners suspect their carefully curated datasets are being used without permission for model training, they can use the entire dataset for membership inference. Second, we observe that in VLMs, the temperature, a user-adjustable parameter, affects member and non-member data differently: members show greater sensitivity to temperature changes than non-members. This temperature sensitivity thus serves as an indicator for identifying membership.

Based on the above, we propose four different types of adversaries with different background knowledge. As shown in Table 1, we gradually relax the assumptions until we arrive at the worst-case adversary.

Shadow Model Inference. We follow the previous MIAs [6, 39, 54, 57] and assume that the adversary has access to an auxiliary dataset (called shadow dataset) and use it to construct a local shadow model that mimics the target model’s behavior.

The adversary utilizes member sets and non-member sets from their shadow dataset to query the shadow model at a specific temperature. For each set, the adversary computes the similarity of the model’s responses with the ground truth answers, along with the corresponding statistics. The adversary then repeats this process by varying the temperatures, obtaining a trend of statistics across different temperatures for the member set and non-member set, respectively. Finally, the adversary trains a binary classifier to learn the discrepancies in the two statistical trends. Once trained, the classifier can differentiate between the member and non-member sets of the target model.

Inferences	VLM Response	Reference Set	Shadow Dataset	Text Data
Shadow	✓	✗	✓	✓
Reference	✓	✓	✗	✓
Target-only	✓	✗	✗	✓
Image-only	✓	✗	✗	✗

Table 1: Comparison of Assumptions on Adversaries

Reference Inference. We relax the assumption that the adversary has access to a shadow dataset and can train a shadow model. Instead, we assume the adversary has a small reference set with known membership status (either members or non-members) in the target model’s training data. This assumption is realistic, particularly for non-member reference datasets, as data generated after the model’s training is complete can effectively serve as non-member samples.

Take the non-member reference set as an example, the adversary first inputs the samples in this set into the target model at a specific temperature, and calculates the similarity of the model’s responses to the ground truth answers. The same process is then applied to the target sample set. Finally, the adversary performs a hypothesis test, i.e., z -test, to determine if the two set statistics are significantly different. If significantly different, the target sample set is considered as members; otherwise, non-members.

Target-Only Inference. We relax the need for a reference set, leaving the adversary with only the target sample set they aim to infer, which is a more realistic and challenging scenario.

The adversary inputs the target set into the target model at two different temperatures and calculates the similarity, along with the corresponding statistics. The adversary then uses the z -test to determine how significantly different the two statistics are. As aforementioned, our insight is that members exhibit greater sensitivity to temperature changes than non-members. Therefore, a larger difference indicates a higher likelihood that the set is a member set.

Image-Only Inference. We further assume that the adversary is unable to access the ground truth answers and can possess only the target images they aim to infer, which represents the most challenging scenario.

The adversary repeatedly asks the model to describe the same image and analyzes the similarity among these responses. For member images, these descriptions tend to converge closely to the ground truth, whereas descriptions of non-member images are more random and less similar. Therefore, a higher similarity among the responses from repeating queries indicates a higher likelihood of being a member set.

Summary. Our contributions are outlined as follows:

- To our knowledge, this is the first systematic study of membership leakage in the VLM domain.

Our study explores the potential of MIA in the detection of unauthorized data usage and reveals previously unidentified vulnerabilities in VLMs.

- We propose four types of inference methods tailored for different assumptions on adversaries' capability, and their success demonstrates the range of threats that membership inference poses to VLMs.
- We conduct extensive evaluations on six models from two representative VLM architecture families. The results show that adversaries with varying capabilities can accurately distinguish between member and non-member sets.

2 Preliminaries

2.1 Vision-Language Models

Inspired by the fact that the fundamental capabilities of VLMs lie in image understanding and language generation, tasks typically handled by existing pre-trained visual encoders and large language models, VLMs such as MiniGPT-4 [70] and LLaVA [43] have opted to build upon these pre-trained backbones and achieved capabilities comparable to commercial proprietary models like GPT-4 [2] at a very low training cost. For example, LLaVA can be trained in just one day using 8 A100 GPUs, making VLM development more accessible to small companies and academic researchers.

Model Structure. As shown in Figure 2, an image x_v is processed by a vision encoder to produce a sequence of image tokens, which are $e_v = \{e_v^i\}_{i=1}^{T_v} \in \mathbb{R}^{d_v}$. A projector f_ω is introduced to transform e_v to $t_v = f_\omega(e_v) \in \mathbb{R}^{d_t}$, thereby mapping the vision tokens into the LLM embedding space. Simultaneously, a textual prompt x_q is input into the tokenizer to obtain $t_q \in \mathbb{R}^{d_t}$. Subsequently, the combined token sequence $[t_v, t_q]$ is fed into the LLM, denoted as f_ϕ , which ultimately generates the language response $y_a = f_\phi(t_v, t_q)$.

Training. Trainable parameters differ across different VLMs, as some VLMs [13, 70] freeze the vision encoder and LLM, training only the projector with less than 1 million parameters, denoted as $\theta = \{\omega\}$, while others [10, 43] train both the projector and LLM, represented as $\theta = \{\omega, \phi\}$. The training typically involves two stages: pre-training and instruction tuning. The first stage utilizes extensive public image-text pairs to achieve preliminary feature alignment. The instruction tuning stage is crucial for equipping VLMs with interactive capabilities, and the models' developers meticulously construct high-quality datasets to align the model with specific tasks.

The training data are represented as $D = \{(x_v^1, x_q^1, y_a^1), \dots, (x_v^N, x_q^N, y_a^N)\}$, where x_v denotes the input image, x_q is the question prompt, and y_a is the answer. The training objective is to maximize the probability of the model outputting y_a given inputs x_v and x_q . The loss for a single data sample is:

$$-\log P_\theta(y_a|x_v, x_q) = -\sum_{i=1}^{n_a} \log P_\theta(t_a^i|t_a^1, t_a^2, \dots, t_a^{i-1}, x_q, x_v), \quad (1)$$

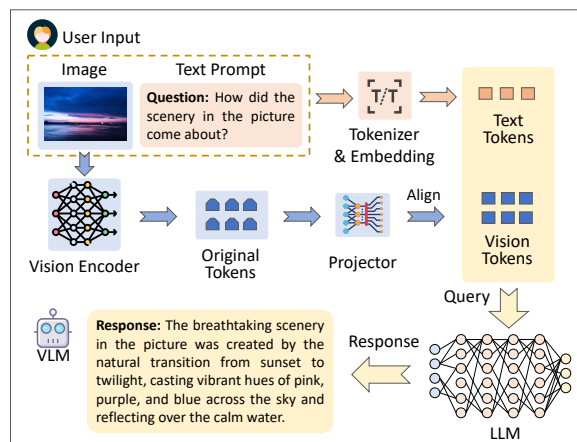


Figure 2: General Structure of VLMs

where $t_a = \{t_a^i\}_{i=1}^{n_a}$ denotes the sequence of tokens derived from tokenizing the answer y_a .

Inference. The inference process in VLM involves multiple next-token predictions. For each prediction, the model first produces a score vector z , whose length is equivalent to the model's vocabulary size $|V|$, and applies a softmax function with temperature T to convert z into a probability distribution:

$$P_\theta(t_a^i = V_j | t_a^1, t_a^2, \dots, t_a^{i-1}, x_q, x_v, T) = \frac{\exp(z_j/T)}{\sum_{k=1}^{|V|} \exp(z_k/T)}, \quad (2)$$

where T modulates the smoothness of the probability distribution, controlling the diversity of the generated text.

2.2 Membership Inference

Membership inference in the ML field is when an adversary aims to determine whether a particular data sample is used for the training of an ML model. The objective is typically focused on a single data sample [6, 39, 54, 57]. Formally, in the context of VLM, given a target data sample $\mathbf{x} = (x_v, x_q, y_a)$, a trained VLM f_θ , and the external knowledge of an adversary, denoted by Ω , a sample membership inference A_{sample} can be defined as:

$$A_{\text{sample}} : (\mathbf{x}, f_\theta, \Omega) \rightarrow \{0, 1\}, \quad (3)$$

where 0 indicates that \mathbf{x} is not a member of the training dataset of f_θ , and 1 indicates that \mathbf{x} is a member.

Inference Target in VLMs. After examining the training process of VLMs, we choose to conduct membership inference on the training data of the instruction tuning stage. The capability of VLMs to interact with humans hinges critically on the instruction tuning phase, where the strength of this capability is directly linked to the quality of the instruction tuning dataset. The developers construct their own task-tailored dataset for this phase. For example, the developers

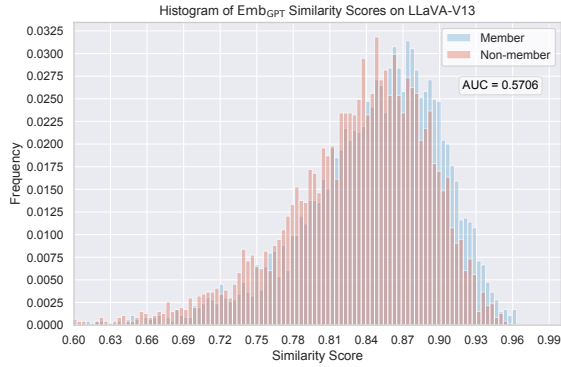


Figure 3: Histogram of Similarity Scores

of MiniGPT-4 manually curated 3.5k high-quality image-text conversation data with the help of a pre-trained VLM. Similarly, the developers of LLaVA produced 158k image-text conversations with ChatGPT, covering various tasks such as multi-turn conversations, detailed image descriptions, and complex reasoning.

Compared to the open-source caption datasets [8,42,50,55] used in the first stage, the task-tailored datasets constructed by developers in the second stage are more likely to contain private information. Besides, since the development of such datasets often requires significant time and financial resources, they are more prone to unauthorized use by a third party. Moreover, due to the model’s catastrophic forgetting nature [19], the model retains a deeper and fresher memory of the data used in later training stages, making the instruction tuning phase both the most valuable and the most vulnerable in the context of membership inference on VLMs.

3 Initial Exploration and Key Insights

In this section, we present our initial attempts at conducting membership inference on VLMs and how these observations inspire our later algorithm designs.

3.1 Training Data Memorization

Successful membership inference relies on the model’s over-fitting of training data. Therefore, the initial step involves verifying whether the model has “memorized” the training data. According to Equation 1, the goal during the instruction tuning stage is to maximize the probability that the model’s output r_a matches the ground truth answer y_a . A straightforward approach is to input the target sample’s image x_v and question x_q into the VLM and observe how closely the VLM’s response r_a matches the answer y_a .

Since directly obtaining sample loss or confidence scores is not feasible in black-box scenario, we instead measure the similarity between r_a and the y_a . We employ OpenAI’s

embedding model [49] to transform the texts into embedding vectors and compute their cosine similarity. We train an LLaVA model and randomly select 1000 member and 1000 non-member data samples to feed into the model, calculating the similarity between the model responses and the ground truth answers. The distribution of these similarities, as shown in Figure 3, reveals slight differences between member and non-member data. We calculated the AUC score, which was 0.5673, only slightly higher than the random guess baseline of 0.5. This suggests a certain degree of over-fitting, implying that the model has, to some extent, memorized the training data. However, this signal alone is not sufficient for conducting successful membership inference.

Set-Level Membership Inference. As aforementioned, existing MIA methods struggle against LLMs [15], and we see a similar trend with VLMs, as shown by the low AUC score of 0.5673. We attribute this to the fact that VLMs, like LLMs, are usually trained on large datasets with few epochs, leading to low over-fitting and weak membership signal. To address this, we shift our focus from inferring the membership of individual samples to that of a set of samples, which we believe can better capture the membership signal by aggregating the signals from individual samples, thereby forming a stronger, more reliable and identifiable signal. We emphasize this setting is realistic in the real world, like scenarios of copyright infringement detection and privacy leakage. First, in cases of copyright infringement, an entire dataset is often used without authorization. Thus, set-level inference can be an effective technique for the dataset owner to detect such unauthorized use. Additionally, individuals can also use this technique to identify unauthorized use of their private data, such as photos posted on social media. Second, set-level inference also poses significant privacy risks. For instance, a lung cancer patient who undergoes frequent examinations generates X-rays and diagnostic records, which consistently indicate the presence of cancer. These pairs of X-rays and diagnoses form a small dataset that could be included in the training data of a medical VLM. When adversaries target such small sets, they can reveal the data owner’s health condition, similar to previous sample-level MIAs. In this case, the set actually acts as a sample, with each image serving as a feature and the textual diagnoses indicating cancer as the label.

Formally, given a target set $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k\}$, a set membership inference A_{set} can be defined as:

$$A_{\text{set}} : (\mathbf{X}, f_{\theta}, \Omega) \rightarrow \{0, 1\}, \quad (4)$$

where A_{set} serves as a binary classifier.

3.2 Temperature

We continue to explore ways to enhance the distinction between members and non-members. As shown in Figure 3, the model tends to produce responses closer to the ground truth answer for members, assigning higher scores to tokens that

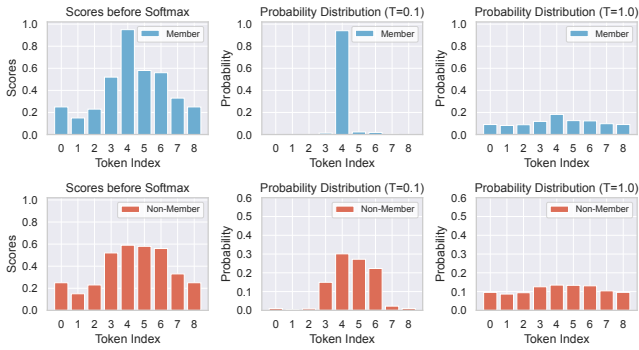


Figure 4: Different Impacts of Temperature on Members and Non-members

match the answer. Equation 2 further shows that the selection of the next token is influenced by the temperature T . A lower T sharpens the softmax output distribution, increasing the likelihood of selecting the highest-scoring token, while a higher T flattens the probability distribution. The effect of T on member and non-member data, however, is significantly different. As shown in the first column of Figure 4, the model, trained on member data, assigns a significantly higher score to the ground truth token during next-token prediction. With a low T , the softmax operation results in a high output probability for the ground truth token, markedly exceeding that of any other tokens. As T increases, the probability difference between the ground truth token and others decreases dramatically (from 0.93 to 0.17), and the similarity between model output and ground truth answer will also decrease significantly accordingly. Conversely, for non-member data, which often have several closely scored tokens, an increase in T also leads to a more uniform output probability distribution, but the change is relatively modest (from 0.29 to 0.13). Consequently, the fidelity of outputs to ground truth answers in member data is highly sensitive to T variations, while the outputs for non-member data are less affected. Since most commercial LLMs² and VLMs³ allow users to adjust the temperature, adversaries can use this feature to better distinguish between member and non-member data.

We study how similarity scores vary with temperature by inputting both members and non-members into the LLaVA model at different temperature settings. We randomly select 1000 samples from both members and non-members and examine how the similarity scores vary across different temperatures, as depicted in Figure 5. While similarity scores for both members and non-members decrease with rising temperature, the decline is more significant for members. Nonetheless, the considerable overlap in variance intervals indicates that distinguishing between members and non-members based on a single data point remains challenging, which suggests that statistical information from a set of samples could be utilized

²<https://platform.openai.com/docs/guides/text-generation>

³<https://huggingface.co/spaces/Vision-CAIR/minigpt4>

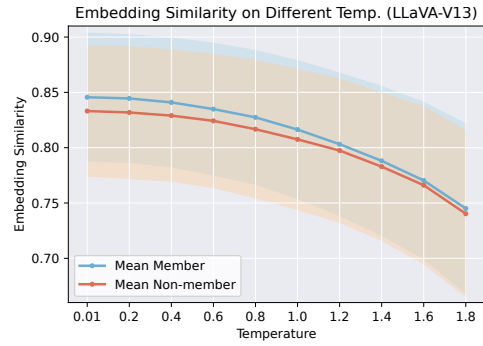


Figure 5: Embedding Similarity on Different Temperatures

to better differentiate between members and non-members. Based on these observations, we design four membership inference algorithms, each based on varying assumptions about the adversary’s capabilities.

4 Shadow Model Inference

In this section, we present the first type of set membership inference against VLMs, i.e., shadow model inference, as shown in the top row of Figure 6. We start by introducing our key intuition. Then we describe the attack methodology. Finally, the evaluation results are presented.

4.1 Intuition

As observed in Section 3.2, when a set of data samples is input into a target VLM with varying temperatures $T \in \{T_i\}_{i=1}^n$, the distribution of similarity scores changes as the temperature increases. The trend in similarity score distributions differs between member and non-member sets.

To leverage this observation, we propose training a binary classifier to learn the patterns in the distribution trends of member and non-member sets. Following previous MIAs [6, 54, 57], we assume that the adversary has access to a shadow dataset D_s drawn from the same distribution as the target model’s dataset D_t . The adversary is then able to train a local shadow model f_{θ_s} on D_s . Subsequently, the adversary generates training data for the binary classifier by querying the shadow model with the shadow dataset. The trained binary classifier can then be employed to conduct inference on a target set with the target model.

4.2 Methodology

The methodology comprises four phases as follows, and the pseudocode is illustrated in Algorithm 1 in Appendix A. **Shadow Model Training.** Initially, the shadow dataset D_s is randomly partitioned into a member dataset D_s^m and a non-member dataset D_s^n . A local shadow model f_{θ_s} is then trained on D_s^m to mimic the behavior of the target model f_{θ_t} .

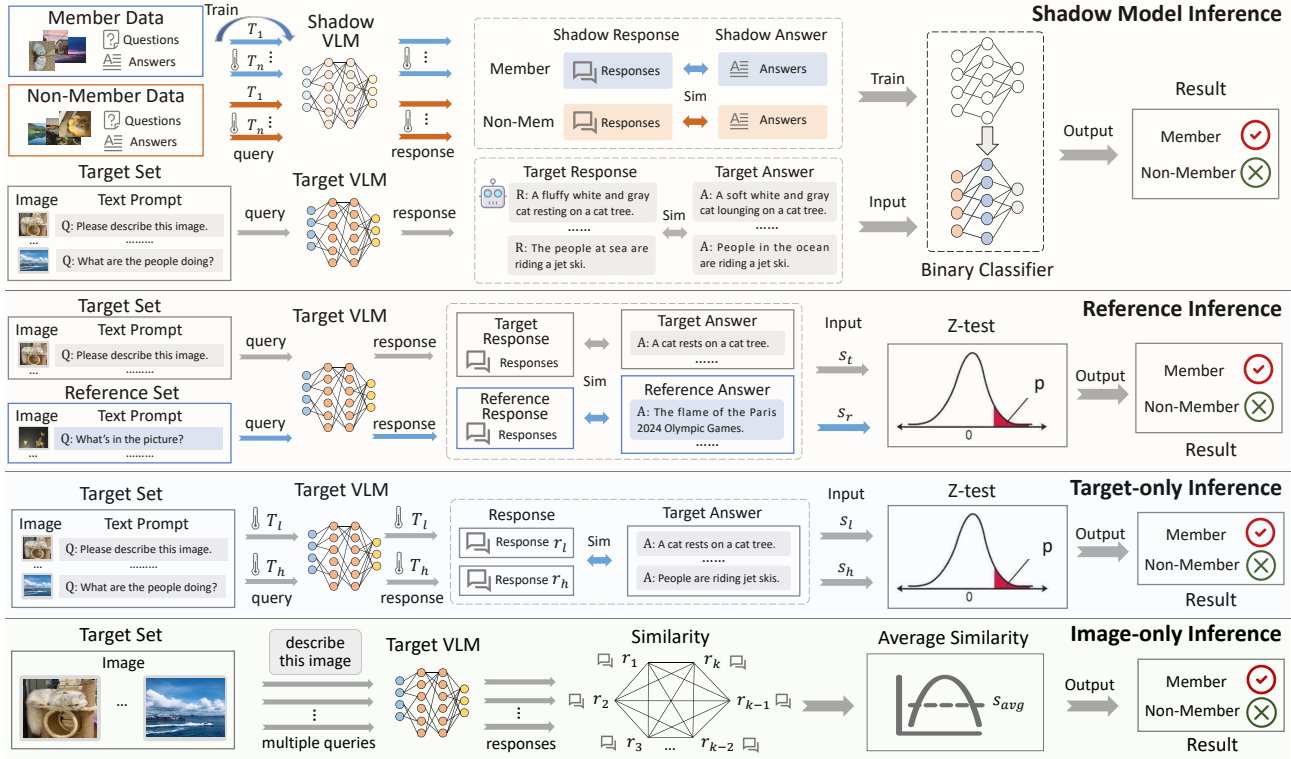


Figure 6: Overview of four Different Membership Inference Attack Algorithms.

Classification Dataset Generation. Depending on the size of the target set, referred to as granularity g , n_b sets of size g are randomly drawn from both D_s^m and D_s^n to form member sets $\{\mathbf{X}_m^i\}_{i=1}^{n_b}$ and non-member sets $\{\mathbf{X}_n^i\}_{i=1}^{n_b}$. For every set $\mathbf{X} \in \{\{\mathbf{X}_m^i\}_{i=1}^{n_b} \cup \{\mathbf{X}_n^i\}_{i=1}^{n_b}\}$, each data sample $\mathbf{x} = (x_v, x_q, y_a) \in \mathbf{X}$ is inputted into the shadow model at varying temperatures $T \in \{T_i\}_{i=1}^{n_T}$, and the similarity score between the model response r and the ground truth answer y_a is computed, denoted as $s = \text{sim}(r, y_a)$. For each set, the mean μ_T and variance σ_T of the similarity scores for all data samples within the set are calculated at each T , and a feature vector $\mathbf{v} = [\mu_{T_1}, \sigma_{T_1}, \dots, \mu_{T_{n_T}}, \sigma_{T_{n_T}}]$ is formed. Each feature vector is subsequently labeled as a member (1) or non-member (0). Consequently, we obtain n_b positive samples and n_b negative samples for binary classifier training.

Binary Classifier Training. We employ a simple neural network as the binary classifier, consisting of three layers: the first transforms input features into a 64-dimensional vector with ReLU activation, the second processes this vector, and the final layer outputs a probability between 0 and 1 via a sigmoid function, representing the sample’s likelihood of being in the positive class. This classifier, denoted as f_b , is trained on previously generated data.

Membership Inference. For the target set \mathbf{X}_t , we apply a similar procedure as in the binary dataset generation phase on the target model to produce the feature vector, which is

then fed into f_b . If the output is 1, the target set is classified as belonging to the training dataset of the target model, and vice versa.

4.3 Evaluation Setting

Vision-Language Models. Experiments are conducted on LLaVA [43] and MiniGPT-4 [70], representing two categories of VLMs: those that update the LLM during instruction tuning and those that freeze the LLM, respectively. For each model, variants based on different foundational LLMs are also considered. For LLaVA, we test four different LLMs: Vicuna-13B, Vicuna-7B [11], Llama-2-13B-chat, and Llama-2-7B-chat [62]. For MiniGPT-4, we only utilize two Vicuna models, due to limited information from developers about training MiniGPT-4 with Llama models.

Data Processing. The samples in instruction-tuning dataset of LLaVA and MiniGPT-4 are already in the form of image-question-answer tuple, we organize them as inference targets.

Since MiniGPT-4 does not require updating LLM parameters during training, only 3.5k samples are used. We follow their data curation procedure and generate approximately 10.5k data samples, which, after shuffling with original dataset, are randomly divided into four subsets: D_t^m , D_t^n , D_s^m , and D_s^n . D_t^m and D_s^m are used for training the target and shadow models, respectively, while D_t^n and D_s^n serve as non-members. 10k

sets of size g (short for granularity) are sampled randomly and queried against the shadow model at 16 different temperatures from D_s^m and D_s^n , serving as the training dataset. Similarly, sets of size g are sampled from D_t^m and D_t^n as test dataset.

The developers of LLaVA utilize ChatGPT to generate 158k language-vision samples for instruction tuning, including three type of data: 58k in multi-turn conversations, 23k in detailed image descriptions, and 77k in complex reasoning. Due to the prohibitive cost of generating a dataset of similar size, we do not create additional dataset as done for MiniGPT-4. To make the most of the LLaVA dataset, it's divided into D_t^m , D_s^m , D_t^n , and D_s^n in a 4:4:1:1 ratio. To examine the impact of different data types, classifiers are independently trained using these three categories of data.

All VLMs are trained using default parameters from the open-source code provided by the authors, including batch size, epochs, learning rates, etc.

Similarity Calculation. We utilize OpenAI's embedding API [49] to calculate the similarity between VLM responses and ground truth answers in Section 3.1. Additionally, we explore other freely available methods for calculating similarity:

- [MPNet](#) [58], an open-source embedding model, can also transform texts into embeddings, and the similarity score can be calculated accordingly.
- [Rouge](#) [41] calculates the similarity based on the overlap of text units. We employ [Rouge-2](#), which measures the bigram overlap.

Evaluation Metrics. We report the performance of shadow model inference with [accuracy](#), [precision](#), and [recall](#), which are commonly used in the evaluation of classification tasks.

4.4 Experimental Results

For the experiments with LLaVA, this section presents data from the complex reasoning category. Comparisons between different categories will be detailed in [Table 3](#).

[Figure 7](#) illustrates the impact of granularity and similarity calculation methods on the success rate of inference for LLaVA and MiniGPT-4 on Vicuna-13B model. As granularity increases, there is a significant improvement in success rate, with our inference method performing better on LLaVA than on MiniGPT-4. This difference is mainly due to LLaVA's fine-tuning of the LLM, corresponding to a larger number of trainable parameters compared to MiniGPT-4. Generally, the over-fitting level is positively correlated with the number of trainable parameters. With LLaVA, merely 5 samples as a set can achieve an accuracy above 0.8, and 10 samples reach an accuracy of 0.9.

The rouge-based similarity calculation outperforms the embedding-based ones. This is because the embedding-based similarity calculation measures the semantic similarity between two responses, whereas the rouge-based similarity calculation quantifies the overlap of vocabulary used. VLMs'

robust generalization ensures that regardless of membership status, the semantic understanding of images remains accurate. The difference in whether the model has learned that data is reflected in how that semantics is expressed, such as which of two synonyms to choose. For member data, the model is more likely to select synonyms consistent with the ground truth answer. The rouge-based similarity calculation more effectively captures these nuances in synonym selection.




Figure 7: Shadow Model Inference on Two Models

[Figure 8](#) shows the impact of the number of features on training the binary classifier. As described in [Algorithm 1](#), each temperature T produces two features: μ_T and σ_T . We compared the effect of using either 8 or 16 different temperatures for inference, and also compared cases where only the mean μ was used as a feature without the variance σ . For LLaVA, using 16 temperatures with both variance and mean as features yields the best results, whereas using only 8 temperatures without variance performs the worst. However, for MiniGPT-4, the opposite is true. This is because the data volume is too small for sufficient training of the classifier model with a large number of parameters, and using fewer features reduces the number of parameters, enabling faster model convergence.

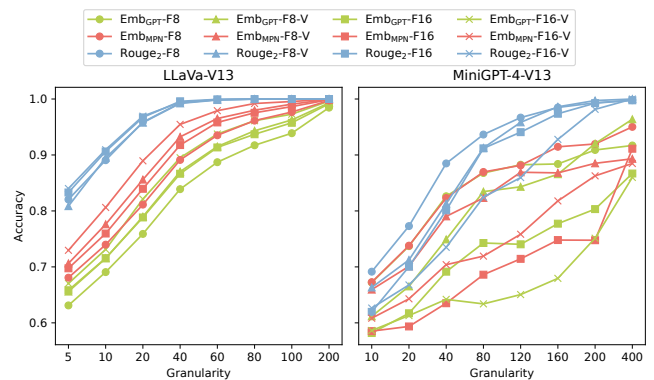


Figure 8: Shadow Model Inference with Different Features

We also compare the impact of different types of LLMs on inference, as shown in [Figure 9](#). For LLaVA, inference

performs better with the 13B model than with the 7B model, because larger number of parameters results in greater overfitting, which, in turn, amplifies the membership signal [12, 57]. In contrast, choosing Vicuna or Llama as the foundational LLM has a negligible impact on inference. For MiniGPT-4, the quantity of LLM parameters does not significantly affect performance, as the LLM parameters are not trainable, and only the projector parameters are updated during instruction tuning.

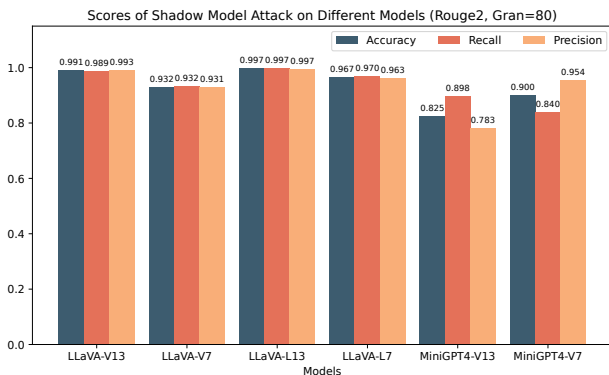


Figure 9: Shadow Model Inference on Different Models

Gran.	Metric	1:1	1:0.75	1:0.5	1:0.25
40	Acc.	88.5%	87.1%	85.3%	80.9%
	Rec.	82.2%	78.4%	74.2%	62.8%
	Pre.	94.0%	94.9%	95.4%	98.4%
120	Acc.	96.7%	95.1%	91.3%	85.1%
	Rec.	93.7%	90.3%	82.5%	70.3%
	Pre.	99.7%	99.8%	100%	100%
200	Acc.	99.3%	97.5%	94.5%	90.5%
	Rec.	98.6%	95.0%	88.9%	81.0%
	Pre.	100%	100%	100%	100%

Table 2: Impact of Shadow Dataset Size

In practical attack scenarios, the size of the shadow dataset available to the adversary may be smaller than the target model’s training dataset. This size discrepancy can lead to different degrees of over-fitting between the shadow model and the target model. To evaluate the impact of this, we conduct experiments on the MiniGPT-4 with Vicuna 13B model, where the size of the shadow dataset is set to 75%, 50%, and 25% of the target dataset, respectively. The inference performance of the shadow model is assessed with rouge-2 as the similarity calculation and 8 dimensions for feature vectors, and the results are presented in Table 2. As the size of the shadow dataset decreases, a significant decline in recall is observed. This decline is attributed to an increase in the degree of over-fitting of the shadow model, resulting in higher similarity scores and consequently raising the threshold for

classifying a set as a member set.

5 Reference Inference

The shadow model reference requires that the adversary possesses both a shadow dataset and adequate computational resources to train the shadow model. However, this assumption may not always hold in practical membership inference scenarios. Consequently, we seek to relax the assumption and propose the method of reference inference.

5.1 Intuition

In the previous inference, the necessity for a shadow dataset stemmed from the need to characterize the distribution pattern of member and non-member sets. By comparing the distribution of the target set with those of the member and non-member sets, the classifier can determine its membership status. We wonder if it is possible to assess, without relying on a trained classifier, whether the similarity scores of the target set originate from the same distribution as those of the member/non-member sets. Statistical hypothesis testing offers a solution: z -test [36] can be utilized to assess whether the distributional difference between two data statistics is significant, thereby aiding in the determination of whether they originate from the same distribution.

If the adversary possesses a reference set whose membership status in the target model’s training dataset is already known, z -test can be employed to calculate the probability that the similarity scores of the reference set and the target set come from the same distribution⁴, as shown in the middle line of Figure 6. This probability can indicate whether the reference set and the target set belong to the same membership status. In practical inference scenarios, acquiring reference data is generally not challenging. For example, data posted after the publication of the model can be considered as non-member data.

5.2 Methodology

Without loss of generality, we discuss the scenario where a known non-member set is used as a reference in this section. For the non-member reference set \mathbf{X}_r of size g_r , then each data sample $\mathbf{x} \in \mathbf{X}_r$ is used to query the target model, and the similarity score s_r between the model’s response r_r and the ground truth answer r_a is computed, resulting in an array of similarity scores $\mathbf{s}_r = [s_r^1, s_r^2, \dots, s_r^{g_r}]$. The target set \mathbf{X}_t is processed in a similar manner, yielding an array of similarity scores, $\mathbf{s}_t = [s_t^1, s_t^2, \dots, s_t^{g_t}]$. Note that each data sample queries the target model only once with a single temperature, as it’s already sufficiently discriminative with the existence of the

⁴More rigorously, z -test can calculate the probability of observing more extreme data differences under the assumption that there is no difference in distributions

reference set. The next step involves conducting a z -test on these two arrays:

$$p = 1 - \Phi\left(\frac{\bar{s}_t - \bar{s}_r}{\sqrt{\frac{\sigma_t^2}{g_t} + \frac{\sigma_r^2}{g_r}}}\right), \quad (5)$$

where Φ denotes the cumulative distribution function of the standard normal distribution, \bar{s}_t and \bar{s}_r are the means, σ_t and σ_r are the standard deviations.

The outcome of the z -test is a p -value. In hypothesis testing, it is conventionally accepted that a $p < 0.05$ indicates a significant difference between the distributions of two data arrays, suggesting that the target set is a member in our scenario. Our evaluations indicate that a $p < 0.05$ threshold can sometimes be too stringent, leading to a high rate of false negatives. The adversary may choose a suitable threshold depending on their requirements, similar to many machine learning applications [4, 18]. In our experimental evaluation, we primarily use the area under the ROC curve (AUC), which is a threshold-independent metric.

5.3 Evaluation Setting

We utilize the target MiniGPT-4 model from Section 4.3. We randomly sample 1,000 sets of size g_t from D_t^m to serve as target member sets. The dataset D_t^m is equally split into D_t^{m1} and D_t^{m2} . 1000 sets of size g_t are sampled from D_t^{m1} as target non-member sets, and 1000 sets of size g_r are sampled from D_t^{m2} as reference non-member sets. Thus, all samples in a set share the same membership status, and the membership status of the individual samples determine that of the set. Each set serves then serves as a data point in evaluation. The scenario in which one set contains samples with different membership status is discussed in Section 8.2. Each reference non-member set is then fed into Algorithm 2 with a target member set or a target non-member set. The resulting p -values are used to compute the AUC scores. For LLaVA, as there is no need to allocate data for a shadow dataset, the data is re-partitioned in an 8:2 ratio into D_m and D_n , and D_m is used to retrain a target model. Other settings are similar to those used for MiniGPT-4.

5.4 Experimental Results

Figure 10 presents how the AUC scores are influenced by granularity, similarity calculation methods, and the type of LLM used. Generally, a larger granularity results in higher AUC; inference on LLaVA outperforms that of MiniGPT-4; and the rouge-based calculation outperforms the embedding-based methods. For LLaVA, larger LLM parameters correspond to higher AUC scores. We also report the results of p -values with Figure 22 in Appendix B.

In prior experiments, we fixed the query temperature $T = 0.1$. Figure 11 displays the impact of varying T on AUC

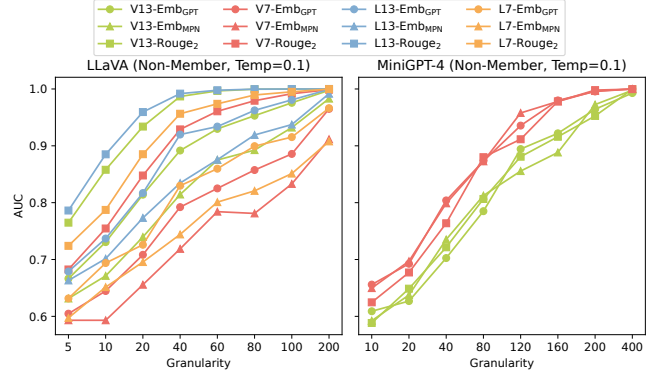


Figure 10: AUC scores of Non-member Reference Inference

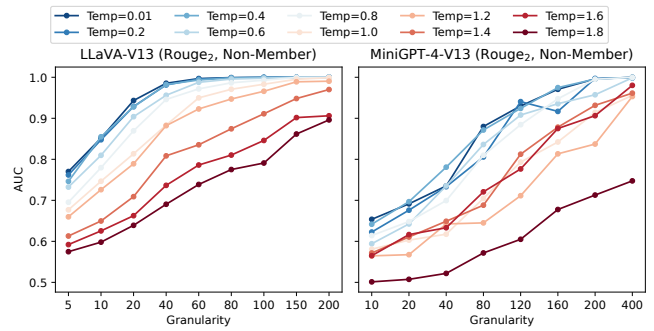


Figure 11: AUC scores of Non-member Reference Inference

scores, and Figure 12 shows its effect on p -values on LLaVA-V13. The results indicate higher inference success rates with smaller T , but variations in T below 0.8 have minimal impact on success rate. Please refer to Appendix B for the p -value results on MiniGPT-4 in Figure 23.

6 Target-Only Inference

Although reference inference relaxes the capability requirements for the adversary compared to shadow model inference, the necessity of a reference set may still pose a barrier in certain scenarios. Therefore, we consider a more challenging scenario where the adversary has only a set of samples that need to be inferred.

6.1 Intuition

The need for a shadow dataset or a reference set stems from the requirement to establish a reference coordinate system. This system allows us to calculate the distance between the target set and the member/non-member sets to determine membership status. Figure 5 inspires a solution for conducting inference without an external reference coordinate system: member data is more sensitive to the change in temperature, meaning that querying the target model at two different tem-

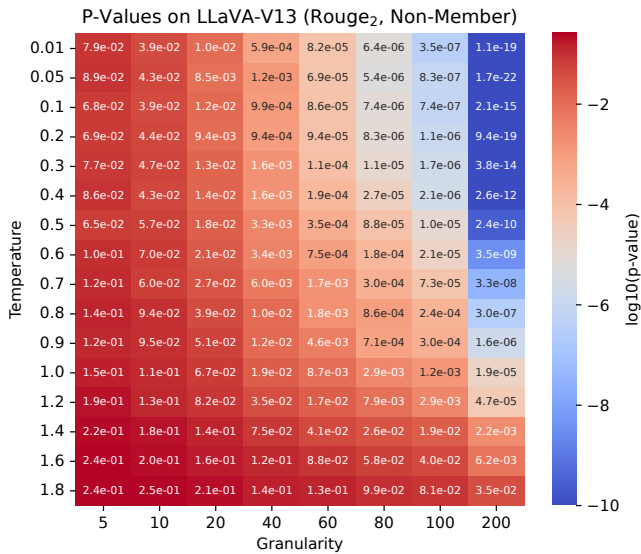


Figure 12: P-Values of Non-member Reference on LLaVA

peratures results in two different similarity scores, and the difference between the two scores is generally larger for members than for non-members. The results of queries at different temperatures establish an internal reference coordinate system, which facilitates distinguishing between members and non-members, as shown in the third row of Figure 6.

6.2 Methodology

For each data sample \mathbf{x} in the target set \mathbf{X}_t of granularity g , the target model is queried with a high temperature T_h and a low temperature T_l for responses r_h and r_l , respectively. The similarity scores s_h and s_l between the responses and the ground truth answer y_a are calculated, resulting in two arrays of similarity scores $\mathbf{s}_h = [s_h^1, s_h^2, \dots, s_h^g]$ and $\mathbf{s}_l = [s_l^1, s_l^2, \dots, s_l^g]$ corresponding to the pair of temperatures. Z-test is then conducted on these two arrays using Equation 5 to compute their p -value. Given that the data at different temperatures inherently come from different distributions, regardless of whether they are from the member set or non-member set, the p -values are typically less than 0.05. Therefore, rather than using 0.05 as a threshold, we utilize the AUC score, which is independent of any threshold.

6.3 Evaluation Setting

This section employs the same target models as described in Section 5.3. We sample 1,000 sets of granularity g from member data and non-member data, respectively. Each set is then fed into Algorithm 3 for target-only inference.




Figure 13: AUC scores of Target-only Inference

6.4 Experimental Results

Figure 13 illustrates the influence of granularity, similarity calculation methods, and type of LLM on AUC, with the experimental setup specifying $T_l = 0.1$ and $T_h = 1.6$. The trends of the impact of granularity and type of LLM on inference performance are consistent with previous experiments; however, in this experiment, the embedding-based similarity calculation outperforms the rouge method. Examination of the actual responses revealed that the quality of VLM responses deteriorates at higher T values, displaying significant semantic differences compared to responses at lower T , where the embedding-based similarity calculation captures these semantic discrepancies more effectively.

We show the effects of varying T_l and T_h on LLaVA with Vicuna-13B in Figure 14. Generally, a larger difference between T_l and T_h correlates with higher AUC scores, although larger is not always better. Particularly at very high T_h values, such as 1.8, the quality of model responses is poor, offering limited utility. Please refer to Appendix B for the results on MiniGPT-4 in Figure 24.

7 Image-only Inference

In some practical scenarios, the adversary may possess only images without the corresponding text. Therefore, we consider the most challenging scenario in which the adversary has only a set of images that need to be inferred.

7.1 Intuition

Due to inherent randomness in the model generation process, querying the same image multiple times may yield slightly different responses. For member sets, since the model has been trained on these sets, the descriptions are more consistent and closely aligned with the ground truth, resulting in greater similarity among them. In contrast, for non-member sets, to which the model has not been exposed, the output probability distribution is more uniform, leading to more diverse and less

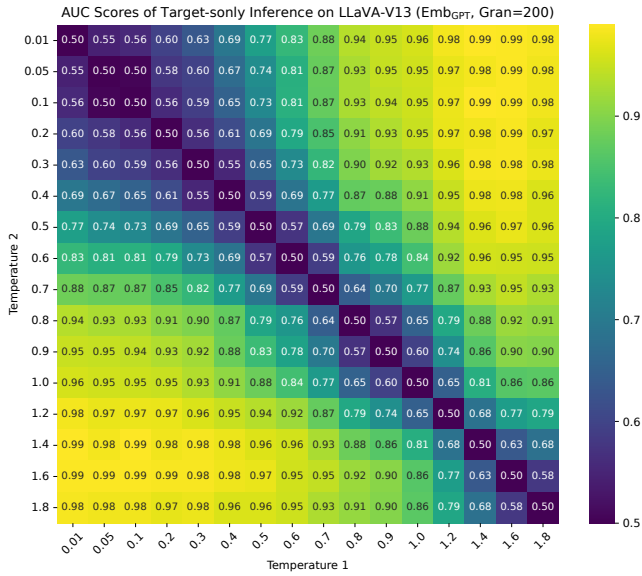


Figure 14: Target-only Inference with Varying g on LLaVA

consistent responses. Thus, when only images are available, an attacker can repeatedly query the model with the same image, assess the similarity among the responses, and use this information to infer membership status, as illustrated in the bottom line of Figure 6.

7.2 Methodology

For each image x_i within the target set, the target model independently describes the image k times, yielding k responses $[r_1, \dots, r_k]$. The similarity among these k responses is calculated, resulting in $\frac{k(k-1)}{2}$ similarity scores $[s_1, \dots, s_{k(k-1)/2}]$, and their mean, s_{avg} , is calculated. The adversary determines the membership status based on the mean s_{avg} of all images in the target set, where a higher s_{avg} indicates membership.

7.3 Evaluation Setting

This section employs the same target models as described in Section 5.3. We sample 1,000 sets of granularity g from member data and non-member data, respectively. The images in each set are then fed into Algorithm 4 for image-only inference.

7.4 Experimental Results

Figure 15 illustrates the impact of granularity, similarity calculation methods, and the type of LLM on the AUC, with the experimental setup specifying $T = 0.5$ and the number of queries $k = 10$. The impact of these variants remains consistent with previous experiments.

Figure 16 demonstrates the effects of temperature on inference success rates. The highest success rate occurs at

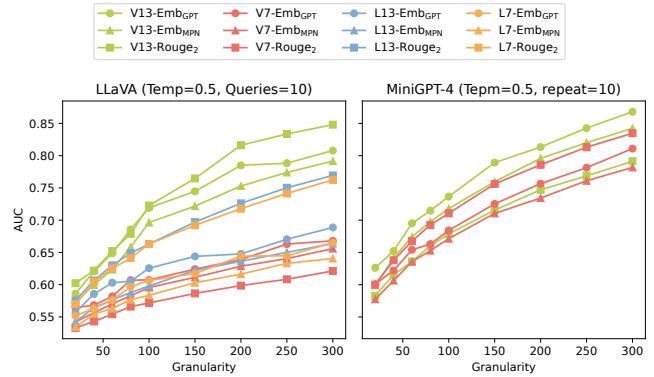


Figure 15: AUC scores of Image-only Inference

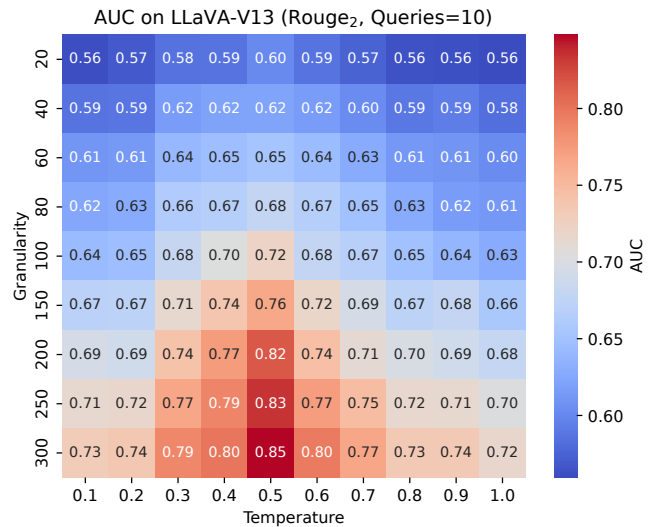


Figure 16: Image-only Inference with Varing T on LLaVA

$T = 0.5$. When the temperature is too low, the randomness during model generation is insufficient, leading to very stable descriptions of non-member images by the model, which is inadequate for distinguishing between member and non-member images. Conversely, when the temperature is too high, the excessive randomness results in significant variations between descriptions of member images by the model. Therefore, selecting an appropriate temperature is crucial.

The impact of the number of queries is depicted in Figure 17. There is a noticeable improvement when the number of queries increases from 5 to 10; however, further increases in the number of queries only marginally enhance the inference success rate.

8 Additional Experiments

This section studies the impact of different settings on the effectiveness of inference methods, including query type, heterogeneous target set, mismatched set size and response length.

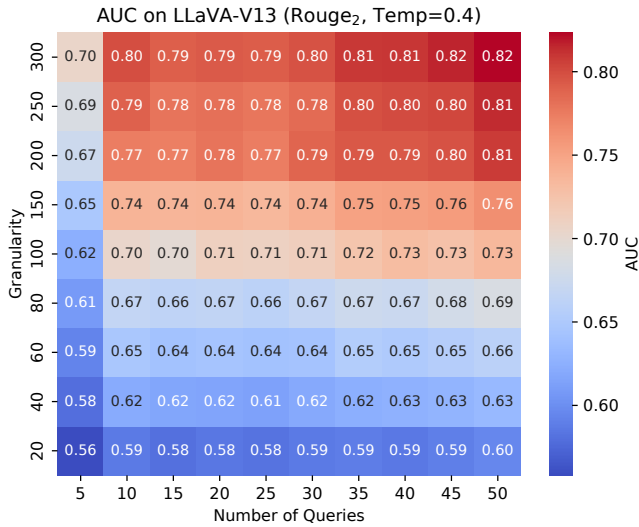


Figure 17: Image-only Inference with Varing k on LLaVA

We also discuss and compare the shadow model inference method with existing attacks that share similar threat models. The experiments in this section are conducted on LLaVA with Vicuna-13B model. More experimental results and discussions can be found in [Appendix B](#) and [C](#).

8.1 Query Type

As discussed in [Section 4.3](#), the instruction tuning dataset for LLaVA is categorized into three types: multi-turn conversations, detailed image descriptions, and complex reasoning. We test the performance of three inference methods across these query types, and the results are presented in [Table 3](#). Note that the image-only inference is not included as it can only be performed on the type of detailed image description. Generally, queries of the complex reasoning type achieved the highest success rate in inference, followed by multi-turn conversations, while those of the detailed image description type were less effective. We believe that this is due to the answers in detailed image descriptions being more definitive, allowing the VLM to produce responses very close to the ground truth answer, regardless of whether it has learned the data. In contrast, answers in complex reasoning are more divergent, resulting in larger discrepancies between members and non-members. Similarly, multi-turn conversations, which often involve reasoning questions, also exhibit higher inference success rates.

8.2 Heterogeneous Target Set

It is typically assumed that the data samples within a target set share the same membership status. This is because model developers generally aim to comprehensively collect all accessible data, and data in one target set often possess the same

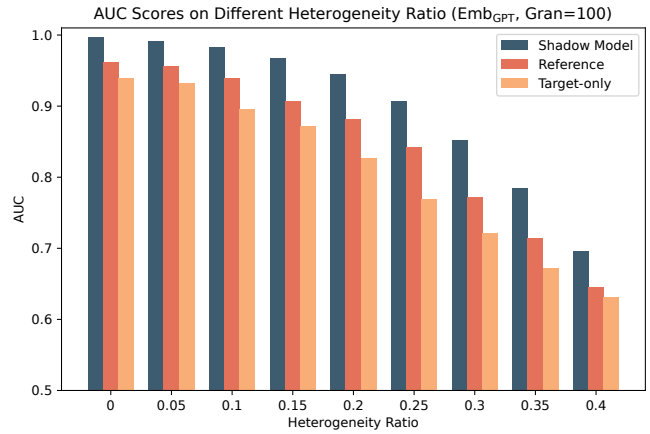


Figure 18: Scores of Heterogeneous Sets

accessibility. For instance, a set of medical images from a single patient is usually either entirely collected or entirely omitted.

In cases where this uniformity does not hold, we explore the scenario that each target set includes a certain proportion of data with opposing membership statuses. The results, displayed in [Figure 18](#), show that for all types of inference, the success rate significantly decreases as the heterogeneity ratio increases. However, when the ratio is low (less than 0.2), the inference success rate remains relatively high (higher than 0.8), showing robustness to some extent.




Figure 19: AUC scores of Sets with Mismatched Granularities

8.3 Mismatched Set Size

In shadow model inference and reference inference, there are actually two types of granularity: one for the target set and the other for the shadow set or the reference set. We investigate whether a mismatch in the size of these sets affects the success rate of inference and illustrate the experimental results in [Figure 19](#). For shadow model inference, the size of the shadow granularity appears to have minimal impact on the success rate of inference. However, the size of the target granularity is positively correlated with the success rate of inference. In the case of reference inference, the success rate is positively correlated with the size of both granularities. Yet, the variation in

Table 3: Comparison of Different Attacks and Query Types on LLaVA-V13.

Metric	Gran.	Shadow Model			Reference Member			Reference Non-Member			Target-only		
		Reason	Detail	Conv.	Reason	Detail	Conv.	Reason	Detail	Conv.	Reason	Detail	Conv.
Emb _{GPT}	20	88.85%	72.72%	88.92%	80.89%	69.36%	77.67%	81.41%	68.23%	78.22%	78.98%	63.65%	74.60%
	60	97.93%	84.31%	97.91%	92.19%	81.90%	92.83%	92.98%	82.20%	92.20%	89.80%	76.07%	88.88%
	100	99.54%	89.12%	99.54%	96.87%	87.30%	96.92%	97.56%	85.73%	96.39%	94.41%	83.26%	93.12%
Emb _{MPN}	20	81.97%	71.65%	82.18%	75.19%	65.22%	76.58%	73.96%	67.98%	77.17%	72.35%	63.20%	71.79%
	60	93.76%	82.49%	91.99%	87.39%	78.66%	88.19%	87.51%	78.73%	90.32%	85.09%	73.55%	85.79%
	100	97.30%	88.52%	96.89%	93.02%	82.79%	94.81%	93.21%	86.13%	94.77%	90.51%	81.26%	90.08%
Rouge-2	20	96.94%	86.09%	94.47%	93.56%	76.41%	86.80%	93.38%	71.64%	87.29%	67.44%	65.92%	69.25%
	60	99.87%	97.08%	99.49%	99.71%	89.47%	97.17%	99.66%	90.56%	97.62%	67.43%	68.49%	83.09%
	100	100.00%	99.26%	99.88%	100.00%	94.51%	99.36%	100.00%	96.37%	99.41%	68.94%	68.94%	87.81%

success rates primarily depends on the size of each granularity itself, rather than on their match, as matched conditions (data along the diagonal) do not significantly outperform adjacent data points.

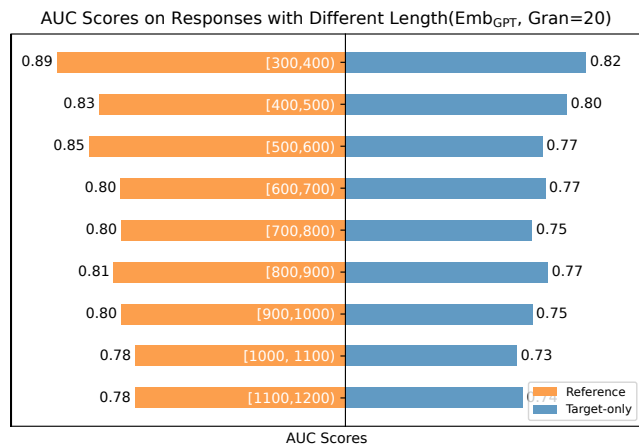


Figure 20: AUC scores on Responses with Different Length

8.4 Response Length

During our experiments, we observe that the similarity between the VLM’s response and the ground truth answer correlates with the length of the response for member data. To further investigate, we group all data according to the string length intervals of the responses and conducted inference independently for each group. The results of reference inference and target-only inference, as shown in Figure 20, indicate that longer responses tend to have lower inference success rates.

This phenomenon can be attributed to the error accumulation in next-token predictions as described in Equation 2. Each token’s prediction depends not only on the input prompt but also on the tokens already generated in the response. For member data, the initial tokens generated after the image and text prompt are fed into the VLM are primarily determined by the prompt and tend to be closer to the ground truth answer.

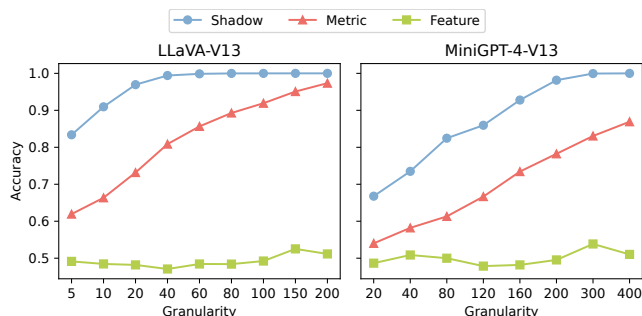


Figure 21: Comparison of Shadow Model Inference with Exiting MIAs

However, if any token deviates from the answer during generation, this error influences the generation of all subsequent tokens, leading to cumulative deviations from the ground truth answer. As the response lengthens, the performance of member and non-member data converges, increasing the difficulty of differentiation.

Experiment on shadow model inference is not conducted as the data volume varies across different length intervals, which could affect classifier training and result in biased comparisons.

8.5 Comparison with Existing Work

While MIAs on LLM-based generative VLMs have not been explored previously, MIAs targeting other multimodal models have been studied. For instance, Hu et al. [30] introduced two shadow-model-based attacks against a CNN&RNN-based caption model [65]. This model generates short textual descriptions (around ten words) for input images but lacks the ability to engage in human-like conversation. Despite significant differences in architecture, data scale, and capabilities between this caption model and modern VLMs, their textual output format makes the proposed MIAs applicable to VLMs. Specifically, Their shadow model approach includes two inference methods: (1) a metric-based method that calculates

similarity between the model’s output and the ground truth using various metrics, and (2) a feature-based method that measures the similarity between images and text in an embedding space, where members typically show higher similarity than non-members. Both methods train a classifier on shadow data to leverage these signals for the attack.

As these methods share a similar threat model with our shadow model method, we compare them on LLaVA-Vicuna-13b and MiniGPT-4-Vicuna-13b models. To adapt to set-level inference, we use the same sampling policy as before to form sets, and then average the membership signals within every set. Figure 21 shows that under the same threat model, our shadow model inference method significantly outperforms the other methods. The metric-based method’s framework is similar to our method, but its performance is worse due to its non-utilization of temperature-based signals. The feature-based method is entirely ineffective, likely because its underlying intuition—that the caption model can semantically align member data pairs well but struggles with non-member data—does not hold for more powerful modern VLMs, which ensure semantic alignment regardless of whether the data was specifically learned.

9 Related Works

9.1 Membership Inference Attack

Membership inference attacks (MIA) are designed to determine whether a specific data record was used in the training of a machine learning model. These attacks have been successfully performed across various types of data, including biomedical data [3] and mobility traces [51]. Shokri et al. [57] proposed the first MIA against machine learning models, utilizing multiple shadow models for attack. Subsequent research [12, 38, 39, 54] progressively relaxed the initial assumptions and introduced more practical approaches, such as data-independent membership inference [54] and decision-based inference [12, 38], making the inference more practical in realistic scenarios. While most MIAs focused on image classification models, later studies expanded the application of MIAs to other types of models, including graph models [24, 31], GANs [9, 26], and diffusion models [14, 28].

In the context of large language models (LLMs), most privacy attacks have concentrated on extracting private information [7, 21, 45, 46]. Nevertheless, these studies also highlight that privacy extraction only works effectively for worst-case data samples and fails to meet the universality required for MIAs. Although attempts [48, 56] have been made to adapt MIAs to LLMs, a systematic investigation [15] has shown that existing MIAs for LLMs barely outperform random guessing in most settings across various sizes and domains of LLMs. This ineffectiveness is largely attributed to the combination of large datasets and few training iterations, a problem that is also faced by large vision-language models.

9.2 Security and Privacy of VLMs

While membership inference attacks on VLMs remain largely unexplored, there exists research revolving around other privacy and security issues concerning VLMs. Xu et al. [67] introduced a poisoning attack designed for VLMs. Liang et al. [40] investigated the feasibility of implanting backdoors during the training phase of VLMs. Several works [52, 61] focus on jailbreaking VLMs with adversarial image prompts. Additionally, Sun et al. [60] explored adversarial sample attacks and their defenses targeting VLMs.

Prior to the advent of interactive conversation-capable VLMs such as ChatGPT-4 [2] and LLaVA [43], VLMs were typically referred to as multi-modal pre-trained representation models such as CLIP [53], which encode text and images into embeddings within the same embedding space. Research on adversarial examples [68, 69] and membership inference attacks [27, 33, 35, 44] for such models has been conducted. However, due to differences in output modalities, attacks targeting this category of VLMs follow a paradigm significantly distinct from that used for modern generative VLMs, which is the focus of this paper.

10 Conclusion

In this paper, we highlight the vulnerabilities associated with membership inference in VLMs by leveraging two observations in overfitting signal capture: the aggregate distribution characteristics of a set of data samples and the differential sensitivity to temperature changes between members and non-members. By progressively relaxing assumptions about adversarial capabilities, we introduce four membership inference algorithms, thereby demonstrating the extensive privacy threats to VLMs. Our comprehensive evaluation confirms the effectiveness of the proposed inference algorithms across various scenarios.

Future work includes: 1) The inference algorithms could be combined and optimized to adapt to a wider variety of scenarios and enhance inference performance. 2) The inference relies on the assumption of identical distributions between the shadow or reference sets and the target set. The implications of distributional discrepancies need exploration. 3) Further investigation is needed into strategies to effectively mitigate the privacy risks posed by set-level membership inference. 4) Given that temperature is a common parameter in language models, the potential of inference algorithms against LLMs needs further exploration.

Acknowledgment

We thank all anonymous reviewers for their constructive comments. This work is supported by the National Natural Science Foundation of China (No. 62072395), the Joint Fund of the National Natural Science Foundation of China

(No. U20A20178), the European Health and Digital Executive Agency (HADEA) within the project “Understanding the individual host response against Hepatitis D Virus to develop a personalized approach for the management of hepatitis D” (DSolve, grant agreement number 101057917) and the BMBF with the project “Repräsentative, synthetische Gesundheitsdaten mit starken Privatsphärengarantien” (PriSyn, 16KISAO29K).

Ethics Considerations

This paper investigates the privacy threats posed to Vision-Language Models (VLMs) by membership inference. This research extends the application of membership inference beyond traditional models to the newly emerged domain of VLMs. In our study, we conduct experiments using publicly available datasets and open-source VLM architectures that are commonly employed in VLM research. The dataset we created is also curated based on a public dataset with the help of large language models. These datasets and models do not pose any direct ethical concerns.

Our work highlights a significant vulnerability in the privacy of VLMs, underscoring the importance of raising awareness about this issue. This is particularly crucial as VLMs represent an emerging research area that is experiencing rapid development and increasingly widespread application. By identifying such vulnerabilities early, we aim to prompt the research community and practitioners to develop more robust defense mechanisms against these threats. Furthermore, our method can be utilized to detect unauthorized data usage in VLM training, thereby contributing to the establishment of stronger safeguards for dataset copyright.

Compliance with Open Science Policy

In adherence to the principles of open science, the datasets and VLM architectures utilized in our research, including LLaVA and MiniGPT-4 along with their corresponding training datasets, are accessible from their official repositories. Additionally, we commit to releasing the source code for all inference algorithms and any supplementary datasets we have developed to the public following the acceptance of this paper and prior to the camera-ready submission deadline. This commitment ensures transparency and fosters continued research in this field.

References

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, pages 308–318, 2016.

[2] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.

[3] Michael Backes, Pascal Berrang, Mathias Humbert, and Praveen Manoharan. Membership privacy in microrna-based studies. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pages 319–330, 2016.

[4] Michael Backes, Mathias Humbert, Jun Pang, and Yang Zhang. walk2friends: Inferring social links from mobility profiles. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pages 1943–1957, 2017.

[5] Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In 2021 IEEE Symposium on Security and Privacy (SP), pages 141–159. IEEE, 2021.

[6] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramer. Membership inference attacks from first principles. In 2022 IEEE Symposium on Security and Privacy (SP), pages 1897–1914. IEEE, 2022.

[7] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In 30th USENIX Security Symposium (USENIX Security 21), pages 2633–2650, 2021.

[8] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 3558–3568, 2021.

[9] Dingfan Chen, Ning Yu, Yang Zhang, and Mario Fritz. Gan-leaks: A taxonomy of membership inference attacks against generative models. In Proceedings of the 2020 ACM SIGSAC conference on computer and communications security, pages 343–362, 2020.

[10] Lin Chen, Jisong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. arXiv preprint arXiv:2311.12793, 2023.

[11] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan

- Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023.
- [12] Christopher A Choquette-Choo, Florian Tramer, Nicholas Carlini, and Nicolas Papernot. Label-only membership inference attacks. In International conference on machine learning, pages 1964–1974. PMLR, 2021.
- [13] Wenliang Dai, Junnan Li, D Li, AMH Tiong, J Zhao, W Wang, B Li, P Fung, and S Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. arxiv 2023. arXiv preprint arXiv:2305.06500, 2, 2023.
- [14] Jinhao Duan, Fei Kong, Shiqi Wang, Xiaoshuang Shi, and Kaidi Xu. Are diffusion models vulnerable to membership inference attacks? In International Conference on Machine Learning, pages 8717–8730. PMLR, 2023.
- [15] Michael Duan, Anshuman Suri, Niloofar Mireshghallah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, and Hannaneh Hajishirzi. Do membership inference attacks work on large language models? arXiv preprint arXiv:2402.07841, 2024.
- [16] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3, pages 265–284. Springer, 2006.
- [17] Adam Dziedzic, Haonan Duan, Muhammad Ahmad Kaleem, Nikita Dhawan, Jonas Guan, Yannis Cattan, Franziska Boenisch, and Nicolas Papernot. Dataset inference for self-supervised models. Advances in Neural Information Processing Systems, 35:12058–12070, 2022.
- [18] Matthew Fredrikson, Eric Lantz, Somesh Jha, Simon Lin, David Page, and Thomas Ristenpart. Privacy in pharmacogenetics: An {End-to-End} case study of personalized warfarin dosing. In 23rd USENIX security symposium (USENIX Security 14), pages 17–32, 2014.
- [19] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. arXiv preprint arXiv:1312.6211, 2013.
- [20] Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. In Proceedings of the 37th International Conference on Machine Learning, pages 3832–3842, 2020.
- [21] Valentin Hartmann, Anshuman Suri, Vincent Bindschaedler, David Evans, Shruti Tople, and Robert West. Sok: Memorization in general-purpose large language models. arXiv preprint arXiv:2310.18362, 2023.
- [22] Iryna Hartsock and Ghulam Rasool. Vision-language models for medical report generation and visual question answering: A review. arXiv preprint arXiv:2403.02469, 2024.
- [23] Jinlong He, Pengfei Li, Gang Liu, Zixu Zhao, and Shenjun Zhong. Pefomed: Parameter efficient fine-tuning on multimodal large language models for medical visual question answering. arXiv preprint arXiv:2401.02797, 2024.
- [24] Xinlei He, Jinyuan Jia, Michael Backes, Neil Zhenqiang Gong, and Yang Zhang. Stealing links from graph neural networks. In 30th USENIX security symposium (USENIX security 21), pages 2669–2686, 2021.
- [25] Kashmir Hill. Twitter tells facial recognition trailblazer to stop using site’s photos. <https://www.nytimes.com/2020/01/22/technology/clearview-ai-twitter-letter.html>, Jan 2020. Accessed: 2023-08-26.
- [26] Benjamin Hilprecht, Martin Härterich, and Daniel Bernau. Monte carlo and reconstruction membership inference attacks against generative models. Proceedings on Privacy Enhancing Technologies, 2019.
- [27] Dominik Hintersdorf, Lukas Struppek, Manuel Brack, Felix Friedrich, Patrick Schramowski, and Kristian Kersting. Does clip know my face? Journal of Artificial Intelligence Research, 80:1033–1062, 2024.
- [28] Hailong Hu and Jun Pang. Membership inference attacks against gans by leveraging over-representation regions. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, pages 2387–2389, 2021.
- [29] Li Hu, Anli Yan, Hongyang Yan, Jin Li, Teng Huang, Yingying Zhang, Changyu Dong, and Chunsheng Yang. Defenses to membership inference attacks: A survey. ACM Computing Surveys, 56(4):1–34, 2023.
- [30] Pingyi Hu, Zihan Wang, Ruoxi Sun, Hu Wang, and Minhui Xue. M4i: Multi-modal models membership inference. Advances in Neural Information Processing Systems, 35:1867–1882, 2022.

- [31] Yuke Hu, Wei Liang, Ruofan Wu, Kai Xiao, Weiqiang Wang, Xiaochen Li, Jinfei Liu, and Zhan Qin. Quantifying and defending against privacy threats on federated knowledge graph embedding. In Proceedings of the ACM Web Conference 2023, pages 2306–2317, 2023.
- [32] Yuke Hu, Jian Lou, Jiaqi Liu, Feng Lin, Zhan Qin, and Kui Ren. Eraser: Machine unlearning in mlaas via an inference serving-aware approach. arXiv preprint arXiv:2311.16136, 2023.
- [33] Bargav Jayaraman, Chuan Guo, and Kamalika Chaudhuri. D\`ej\`a vu memorization in vision-language models. arXiv preprint arXiv:2402.02103, 2024.
- [34] Yigitcan Kaya, Sanghyun Hong, and Tudor Dumitras. On the effectiveness of regularization against membership inference attacks. arXiv preprint arXiv:2006.05336, 2020.
- [35] Myeongseob Ko, Ming Jin, Chenguang Wang, and Ruoxi Jia. Practical membership inference attacks against large-scale multi-modal models: A pilot study. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 4871–4881, 2023.
- [36] Derrick N Lawley. A generalization of fisher’s z test. Biometrika, 30(1/2):180–187, 1938.
- [37] Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. arXiv preprint arXiv:2306.00890, 2023.
- [38] Jiacheng Li, Ninghui Li, and Bruno Ribeiro. Membership inference attacks and defenses in supervised learning via generalization gap. arXiv preprint arXiv:2002.12062, 3(7), 2020.
- [39] Zheng Li and Yang Zhang. Membership leakage in label-only exposures. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, pages 880–895, 2021.
- [40] Jiawei Liang, Siyuan Liang, Man Luo, Aishan Liu, Dongchen Han, Ee-Chien Chang, and Xiaochun Cao. Vltrojan: Multimodal instruction backdoor attacks against autoregressive visual language models. arXiv preprint arXiv:2402.13851, 2024.
- [41] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In Text summarization branches out, pages 74–81, 2004.
- [42] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13, pages 740–755. Springer, 2014.
- [43] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. Advances in neural information processing systems, 36, 2024.
- [44] Hongbin Liu, Jinyuan Jia, Wenjie Qu, and Neil Zhenqiang Gong. Encodermi: Membership inference against pre-trained encoders in contrastive learning. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, pages 2081–2095, 2021.
- [45] Yupei Liu, Yuqi Jia, Jinyuan Jia, and Neil Zhenqiang Gong. Evaluating large language model based personal information extraction and countermeasures. arXiv preprint arXiv:2408.07291, 2024.
- [46] Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. Analyzing leakage of personally identifiable information in language models. In 2023 IEEE Symposium on Security and Privacy (SP), pages 346–363. IEEE, 2023.
- [47] Pratyush Maini, Mohammad Yaghini, and Nicolas Papernot. Dataset inference: Ownership resolution in machine learning. In International Conference on Learning Representations, 2022.
- [48] Justus Mattern, Fatemehsadat Miresghallah, Zhijing Jin, Bernhard Schoelkopf, Mrinmaya Sachan, and Taylor Berg-Kirkpatrick. Membership inference attacks against language models via neighbourhood comparison. In Findings of the Association for Computational Linguistics: ACL 2023, pages 11330–11343, 2023.
- [49] OpenAI. Open ai embedding api. <https://platform.openai.com/docs/guides/embeddings/what-are-embeddings>, 2023. Accessed: 2023-08-26.
- [50] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. Advances in neural information processing systems, 24, 2011.
- [51] Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. Measuring membership privacy on aggregate location time-series. Proceedings of the ACM on Measurement and Analysis of Computing Systems, 4(2):1–28, 2020.

- [52] Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Peter Henderson, Mengdi Wang, and Prateek Mittal. Visual adversarial examples jailbreak aligned large language models. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pages 21527–21536, 2024.
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In International conference on machine learning, pages 8748–8763. PMLR, 2021.
- [54] Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes. Ml-leaks: Model and data independent membership inference attacks and defenses on machine learning models. In Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS), 2019.
- [55] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv preprint arXiv:2111.02114, 2021.
- [56] Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models. In The Twelfth International Conference on Learning Representations, 2024.
- [57] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In 2017 IEEE symposium on security and privacy (SP), pages 3–18. IEEE, 2017.
- [58] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. Advances in neural information processing systems, 33:16857–16867, 2020.
- [59] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. The journal of machine learning research, 15(1):1929–1958, 2014.
- [60] Jiachen Sun, Changsheng Wang, Jiong Xiao Wang, Yiwei Zhang, and Chaowei Xiao. Safeguarding vision-language models against patched visual prompt injectors. arXiv preprint arXiv:2405.10529, 2024.
- [61] Xijia Tao, Shuai Zhong, Lei Li, Qi Liu, and Lingpeng Kong. Imgtrojan: Jailbreaking vision-language models with one image. arXiv preprint arXiv:2403.02910, 2024.
- [62] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971, 2023.
- [63] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Wenqi Wei, and Lei Yu. Effects of differential privacy and data skewness on membership inference vulnerability. In 2019 First IEEE international conference on trust, privacy and security in intelligent systems and applications (TPS-ISA), pages 82–91. IEEE, 2019.
- [64] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Lei Yu, and Wenqi Wei. Towards demystifying membership inference attacks. arXiv preprint arXiv:1807.09173, 2018.
- [65] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 3156–3164, 2015.
- [66] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. Cogvlm: Visual expert for pretrained language models. arXiv preprint arXiv:2311.03079, 2023.
- [67] Yuancheng Xu, Jiarui Yao, Manli Shu, Yanchao Sun, Zichu Wu, Ning Yu, Tom Goldstein, and Furong Huang. Shadowcast: Stealthy data poisoning attacks against vision-language models. In ICLR 2024 Workshop on Secure and Trustworthy Large Language Models, 2024.
- [68] Ziyi Yin, Muchao Ye, Tianrong Zhang, Tianyu Du, Jinguo Zhu, Han Liu, Jinghui Chen, Ting Wang, and Fenglong Ma. Vlattack: Multimodal adversarial attacks on vision-language tasks via pre-trained models. Advances in Neural Information Processing Systems, 36, 2024.
- [69] Peng-Fei Zhang, Zi Huang, and Guangdong Bai. Universal adversarial perturbations for vision-language pre-trained models. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, pages 862–871, 2024.
- [70] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. In The Twelfth International Conference on Learning Representations, 2024.

APPENDIX

A Pseudocode

Algorithm 1 Shadow Model Inference

Input: Shadow dataset D_s , target model f_{θ_t} , target set \mathbf{X}_t , granularity g , number of sets n_b , temperature set $\{T_i\}_{i=1}^{n_T}$

- 1: Randomly partition shadow dataset D_s into D_s^m and D_s^n
- 2: Train shadow model f_{θ_s} on D_s^m
- 3: Randomly draw n_b sets of size g from both D_s^m and D_s^n , and obtain $\{\mathbf{X}_m^i\}_{i=1}^{n_b}$ and $\{\mathbf{X}_n^i\}_{i=1}^{n_b}$
- 4: **for** each $\mathbf{X} \in \{\mathbf{X}_m\} \cup \{\mathbf{X}_n\}$ **do**
- 5: **for** each $T \in \{T_i\}_{i=1}^{n_T}$ **do**
- 6: **for** each $\mathbf{x} = (x_v, x_q, y_a) \in \mathbf{X}$ **do**
- 7: Query shadow model and get $r = f_{\theta_s}(x_v, x_q, T)$
- 8: Compute similarity score $s = \text{sim}(r, y_a)$
- 9: **end for**
- 10: Calculate mean μ_T and variance σ_T of all s
- 11: **end for**
- 12: Form feature vector $\mathbf{v} = [\mu_{T_1}, \sigma_{T_1}, \dots, \mu_{T_{n_T}}, \sigma_{T_{n_T}}]$
- 13: Label vectors as member (1) or non-member (0)
- 14: **end for**
- 15: Train binary classifier f_b using labeled $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^{2 \cdot n_b}$
- 16: Calculate feature vector \mathbf{v}_t for target set \mathbf{X}_t
- 17: Conduct inference $\mathbb{1} = f_b(\mathbf{v}_t)$

Output: Membership status $\mathbb{1} \in \{0, 1\}$

B Deferred Experimental Results

B.1 Impact of different factors on p-values

Figure 22 presents how the p -values are influenced by granularity, similarity calculation methods, and the type of LLM used in reference inference. These p -values represent the mean calculated over 1,000 inferences. Generally, a larger granularity results in lower p -values; p -values for LLaVA are smaller than those for MiniGPT-4, and the rouge-based calculation outperforms the embedding-based methods. For LLaVA, larger LLM corresponds to smaller p -values.

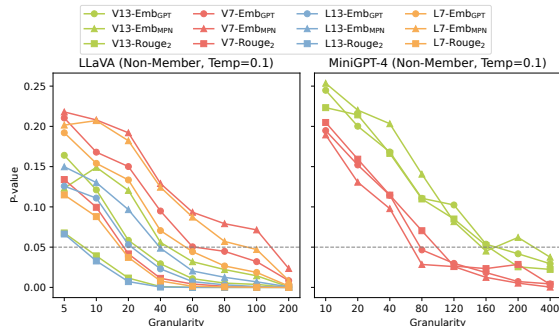


Figure 22: p-Values of Non-member Reference Inference

Algorithm 2 Reference Inference with Non-member Set

Input: Non-member reference set \mathbf{X}_r of size g_r , target set \mathbf{X}_t of size g_t , target model f_{θ_t} , threshold τ

- 1: **for** each $\mathbf{x} = (x_v, x_q, y_a) \in \mathbf{X}_r$ **do**
- 2: Query target model and get $r_r = f_{\theta_t}(x_v, x_q)$
- 3: Compute similarity score $s_r = \text{sim}(r_r, y_a)$
- 4: **end for**
- 5: **for** each $\mathbf{x} = (x_v, x_q, y_a) \in \mathbf{X}_t$ **do**
- 6: Query target model and get $r_t = f_{\theta_t}(x_v, x_q)$
- 7: Compute similarity score $s_t = \text{sim}(r_t, y_a)$
- 8: **end for**
- 9: Compute mean \bar{s}_r/\bar{s}_t and standard deviation σ_r/σ_t
- 10: Calculate the combined standard error $e = \sqrt{\frac{\sigma_r^2}{g_t} + \frac{\sigma_t^2}{g_r}}$
- 11: Calculate the p -value $p = 1 - \Phi\left(\frac{\bar{s}_t - \bar{s}_r}{e}\right)$
- 12: **if** $p < \tau$ **then**
- 13: Conclude that $\mathbb{1} = 1$, i.e., \mathbf{X}_t is a member set
- 14: **else**
- 15: Conclude that $\mathbb{1} = 0$, i.e., \mathbf{X}_t is a non-member set
- 16: **end if**

Output: Membership status $\mathbb{1} \in \{0, 1\}$

Algorithm 3 Target-only Inference

Input: Target set \mathbf{X}_t of size g , target model f_{θ_t} , query temperature T_h and T_l , threshold τ .

- 1: **for** each $\mathbf{x} = (x_v, x_q, y_a) \in \mathbf{X}_t$ **do**
- 2: Query shadow model with T_h and T_l , respectively, obtain $r_h = f_{\theta_t}(x_v, x_q, T_h)$, $r_l = f_{\theta_t}(x_v, x_q, T_l)$
- 3: Compute the similarity score $s_h = \text{sim}(r_h, y_a)$, $s_l = \text{sim}(r_l, y_a)$
- 4: **end for**
- 5: Compute the mean \bar{s}_h/\bar{s}_l and the standard deviation σ_h/σ_l of s_h/s_l
- 6: Calculate the combined standard error $e = \sqrt{\frac{\sigma_l^2 + \sigma_h^2}{g}}$
- 7: Calculate the p -value $p = 1 - \Phi\left(\frac{\bar{s}_l - \bar{s}_h}{e}\right)$
- 8: **if** $p < \tau$ **then**
- 9: Conclude that $\mathbb{1} = 1$, i.e., \mathbf{X}_t is a member set
- 10: **else**
- 11: Conclude that $\mathbb{1} = 0$, i.e., \mathbf{X}_t is a non-member set
- 12: **end if**

Output: Membership status $\mathbb{1} \in \{0, 1\}$

Algorithm 4 Image-only Inference

Input: Target set \mathbf{X}'_v of size g , target model f_{θ_t} , query temperature T , threshold τ .

- 1: **for** each $x_v \in \mathbf{X}'_v$ **do**
- 2: Ask shadow model to describe image x_v k times and obtain $[r_1, r_2, \dots, r_k]$
- 3: Compute the similarity score between every pair of these responses and get $[s_1, s_2, \dots, s_{k \times (k-1)/2}]$
- 4: Average the similarity scores and get s_{avg}
- 5: **end for**
- 6: Compute the mean \bar{s}_{avg}
- 7: **if** $\bar{s}_{avg} > \tau$ **then**
- 8: Conclude that $\mathbb{1} = 1$, i.e., \mathbf{X}_t is a member set
- 9: **else**
- 10: Conclude that $\mathbb{1} = 0$, i.e., \mathbf{X}_t is a non-member set
- 11: **end if**

Output: Membership status $\mathbb{1} \in \{0, 1\}$

Figure 23 displays the impact of varying T on p -values on MiniGPT-4-V13, and it shows a similar trend as in Figure 12. The effects of varying T_l and T_h of target-only inference on MiniGPT-4-V13 are shown in Figure 24. The results align with LLaVA-V13 in Figure 14.

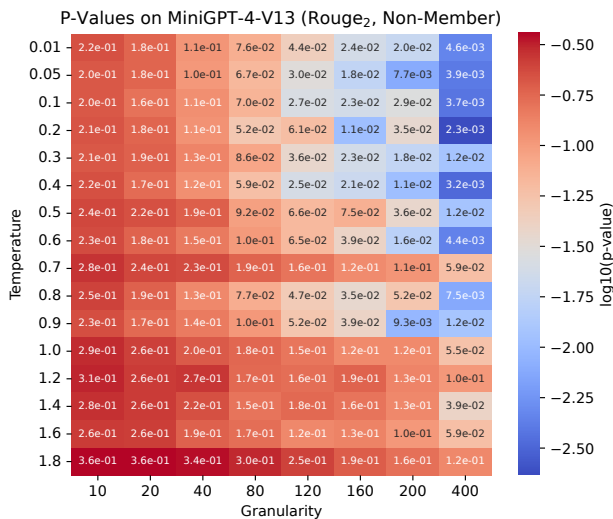


Figure 23: P-Values of Non-member Reference on MiniGPT4

C Discussions

C.1 Improvements of Proposed Algorithms

In this paper, we aim to uncover the breadth of threats that set-level membership inference poses to VLMs across various scenarios, rather than optimizing inference algorithms and parameter settings for each specific scenario. Indeed, the inference algorithms proposed in this paper can be further




Figure 24: Target-only Inference with Varying g on MiniGPT

refined. For instance, shadow model inference could benefit from tailored classifier designs for different target models, involving modifications in model architecture, feature selection, and training hyper-parameters. For reference inference, we currently query the target model with a single temperature. By employing multiple temperatures for querying, the aggregated signal could improve the success rate of the inference. A similar enhancement strategy applies to target-only inference, where using multiple pairs of temperatures for querying could boost performance.

C.2 Defense Against Membership Inference

As over-fitting is a fundamental prerequisite for successful membership inference, methods that reduce over-fitting can serve as effective defenses [29]. Common strategies such as weight decay (L1/L2 regularization) [34, 64], dropout [59], and data augmentation are known to decrease over-fitting effectively, albeit often at the expense of model utility. Another category of defenses employs differential privacy (DP) [1, 16, 63], which introduces carefully calibrated noise to provide provable privacy guarantees. However, the introduction of DP noise can significantly compromise model utility, especially for models with a large number of parameters. Beyond these training-phase defensive measures, machine unlearning [5, 20, 32] seeks to erase the influence of specific data on a trained model. Yet, unlearning algorithms tailored to VLMs remain largely unexplored.

In response to the inference methods proposed in this paper, several basic defenses can be considered, such as restricting users to fixed temperature queries or limiting the range of temperature adjustments. Additionally, as observed in Section 8.4, encouraging longer model responses may help to blur the boundaries between members and non-members.