

Antidistillation Fingerprinting

Yixuan (Even) Xu

John Kirchenbauer

Yash Savani

Asher Trockman

Alexander Robey

Tom Goldstein

Fei Fang

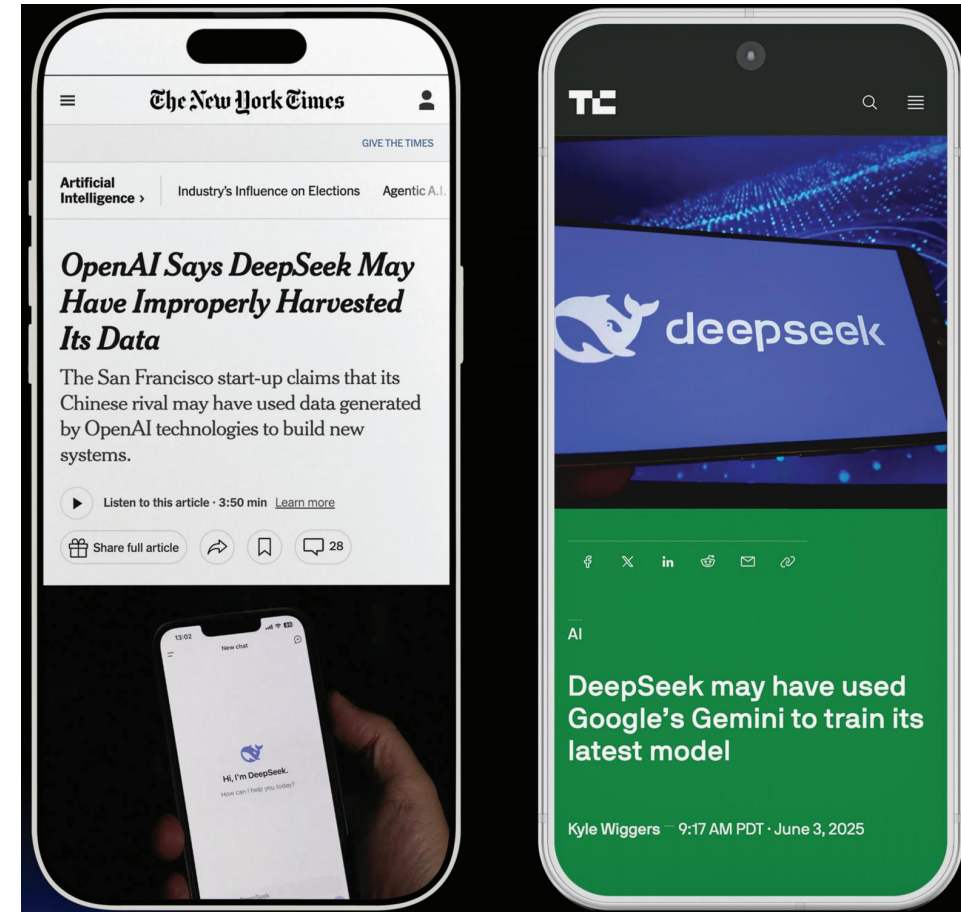
J. Zico Kolter

Carnegie Mellon University

University of Maryland

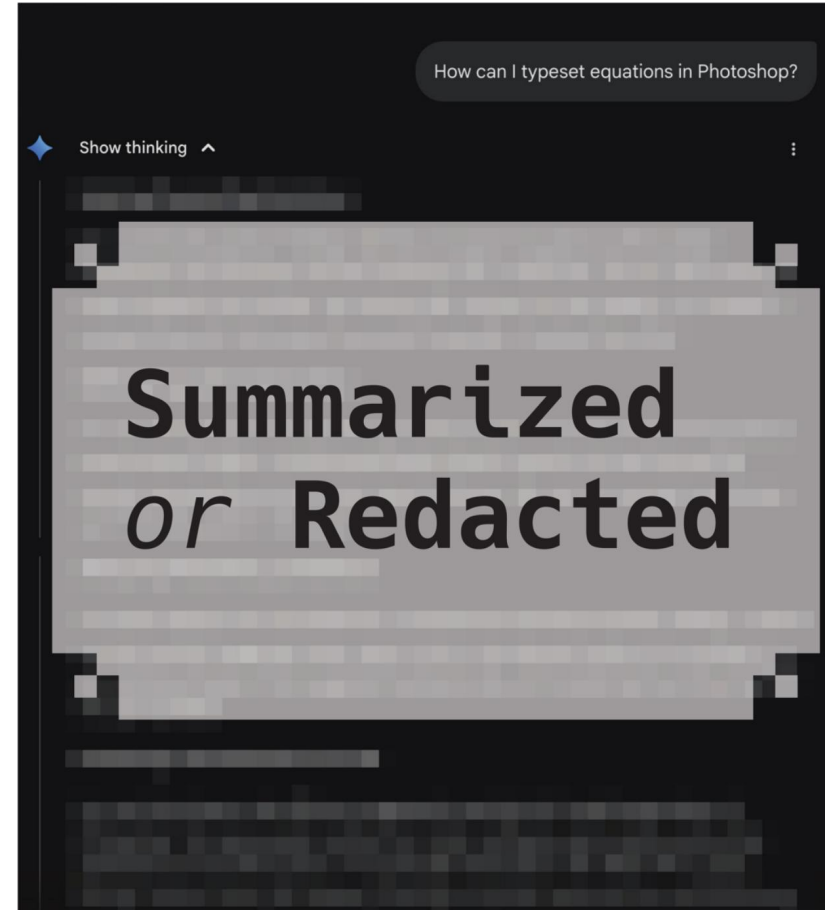
Model Distillation

- **Distillation:** Using a trained model to generate training signals for another model
 - For LMs: Usually via data
- Distilling a model is **easier than** training from scratch
 - Less effort for data curation
 - Better and denser signals



Negative Externalities of Distillation

- Frontier models are hiding their reasoning **against distillation**
- This is less ideal for users
 - The model is solving problems or even making decisions for them
 - Less transparent
- Also for the model owners
 - Such a defense can be broken



The Research Question

*How can we **tell if** another model has **trained on** a teacher model's outputs?*

We name this task **fingerprinting** for this work.

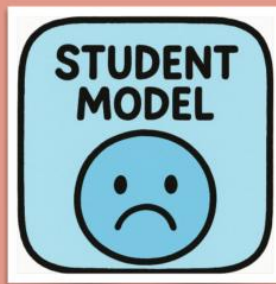
From the Model Owner's Perspective

- There are three models
 - **Teacher θ_T** : The owner's trained model
 - **Student θ_S** : A model to be trained on the teacher's outputs
 - **Proxy θ_P** : An **“imaginary”** student model
 - From the teacher's perspective, the student model is unknown
 - If the method needs to use a “student model”, it can only use a proxy model
- Antidistillation is **NOT** everything
 - We still would like the method to be **lightweight**
 - We would like to **preserve the model's performance**

Antidistillation Sampling (Savani et al. 2025)

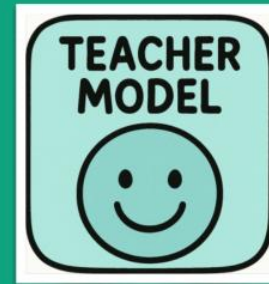
Non-distillability

Student models should not benefit from training on the reasoning traces



Nominal utility

Teacher model's performance should not fall significantly as a result of using the method



The Basic Idea

- Suppose the teacher generates the next token as t after $x_{1:l}$
 - The student / proxy will do one step of **Gradient Descent** on t
 - Assume a **downstream evaluation loss** of the the proxy $L(\theta_P)$
 - Define the **delta vector** as $\Delta_t =$ The increase in $L(\theta_P)$ after the **GD** step
- ***Antidistillation sampling***
 - We have a logits vector $z = z(x_{1:l}, \theta_T)$, and a **delta vector** Δ
 - We will sample according to **softmax**($z + \lambda\Delta$)
 - Parameter λ controls the trade-off between nominal utility and non-distillability
 - $\lambda = 0$ falls back to normal sampling, while $\lambda \rightarrow \infty$ causes maximum poisoning

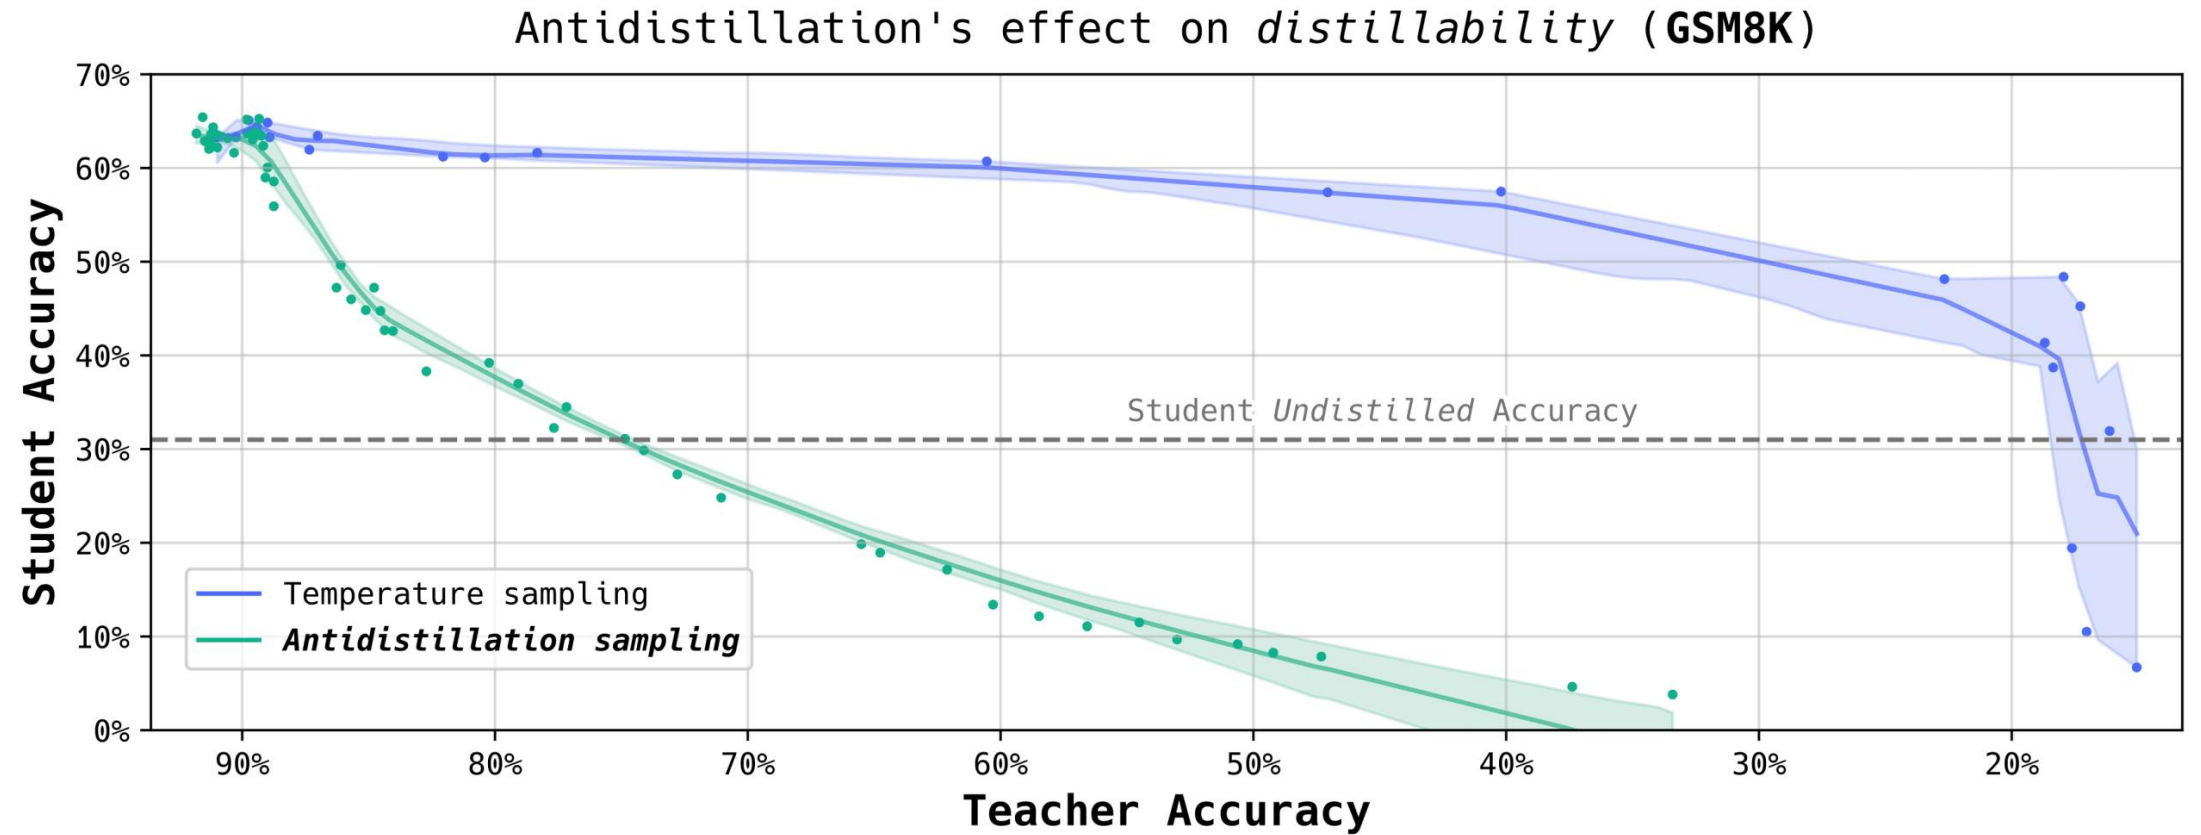
Fast Computation

- Really enumerating the next token to compute Δ is expensive
 - Requires $\Omega(|V|)$ **backward passes** per token
- We use **finite difference method** to estimate Δ
 - Assume differentiable loss $L(\theta_P)$, we compute $\mathbf{G} = \nabla L(\theta_P)$
 - Then, we approximate Δ by

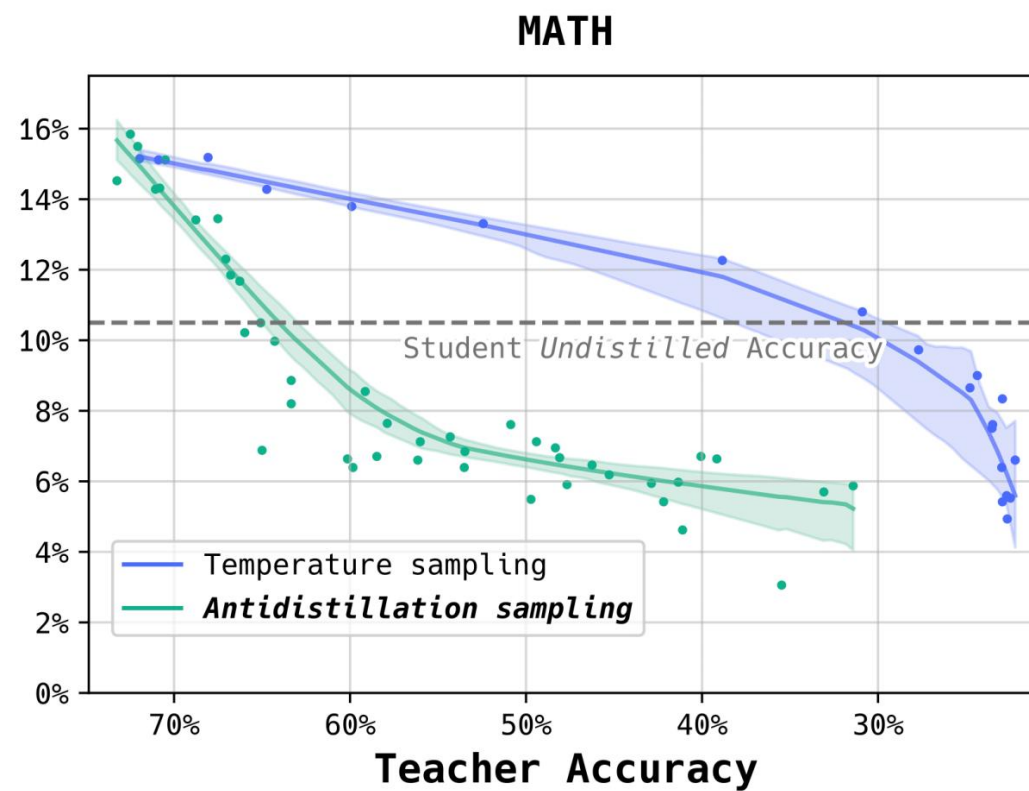
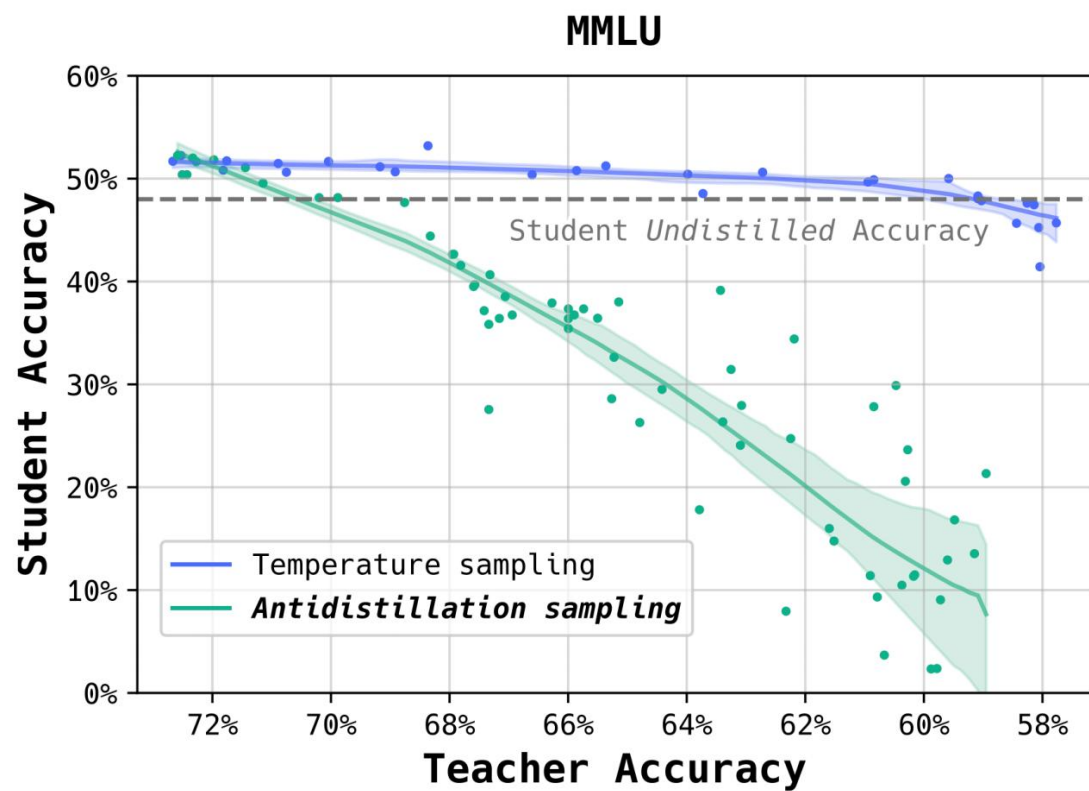
$$\Delta = \frac{z(x_{1:l}, \theta_P + \epsilon \mathbf{G}) - z(x_{1:l}, \theta_P - \epsilon \mathbf{G})}{2\epsilon}$$

- **$O(1)$ forward passes** per token, compatible with KV-caching

Empirical Results



Empirical Results



Antidistillation Fingerprinting

- *When you think about it...*
 - Antidistillation sampling essentially **injects biases** to the traces, so that for the student, a specific loss function $L(\cdot)$ **increases**
 - The function $L(\cdot)$ can actually be **anything** we choose
- ***Fingerprinting*** a model's output
 - Leave specific structures (fingerprints) in the output traces
 - So that when students train on them, it is detectable statistically
 - i.e., can say with statistical guarantee, the student trained on the teacher's output

Fingerprinting Before Antidistillation

- The red-and-green list watermark (Kirchenbauer et al. 2023)
 - For context $x_{1:l}$, define a **green list** computed as $H(x_{1:l}, k)$
 - H is a hash function keyed by k that outputs half of the vocabulary
 - Change the logits at sampling by increasing all green tokens by δ
 - Normal texts will have a **green-list ratio** of concentrated around 0.5
 - Watermarked texts' **green-list ratio** will be higher
- Idea of *antidistillation fingerprinting*
 - Use antidistillation sampling to bias the **green-list ratio** of the student

Deriving Antidistillation Fingerprinting

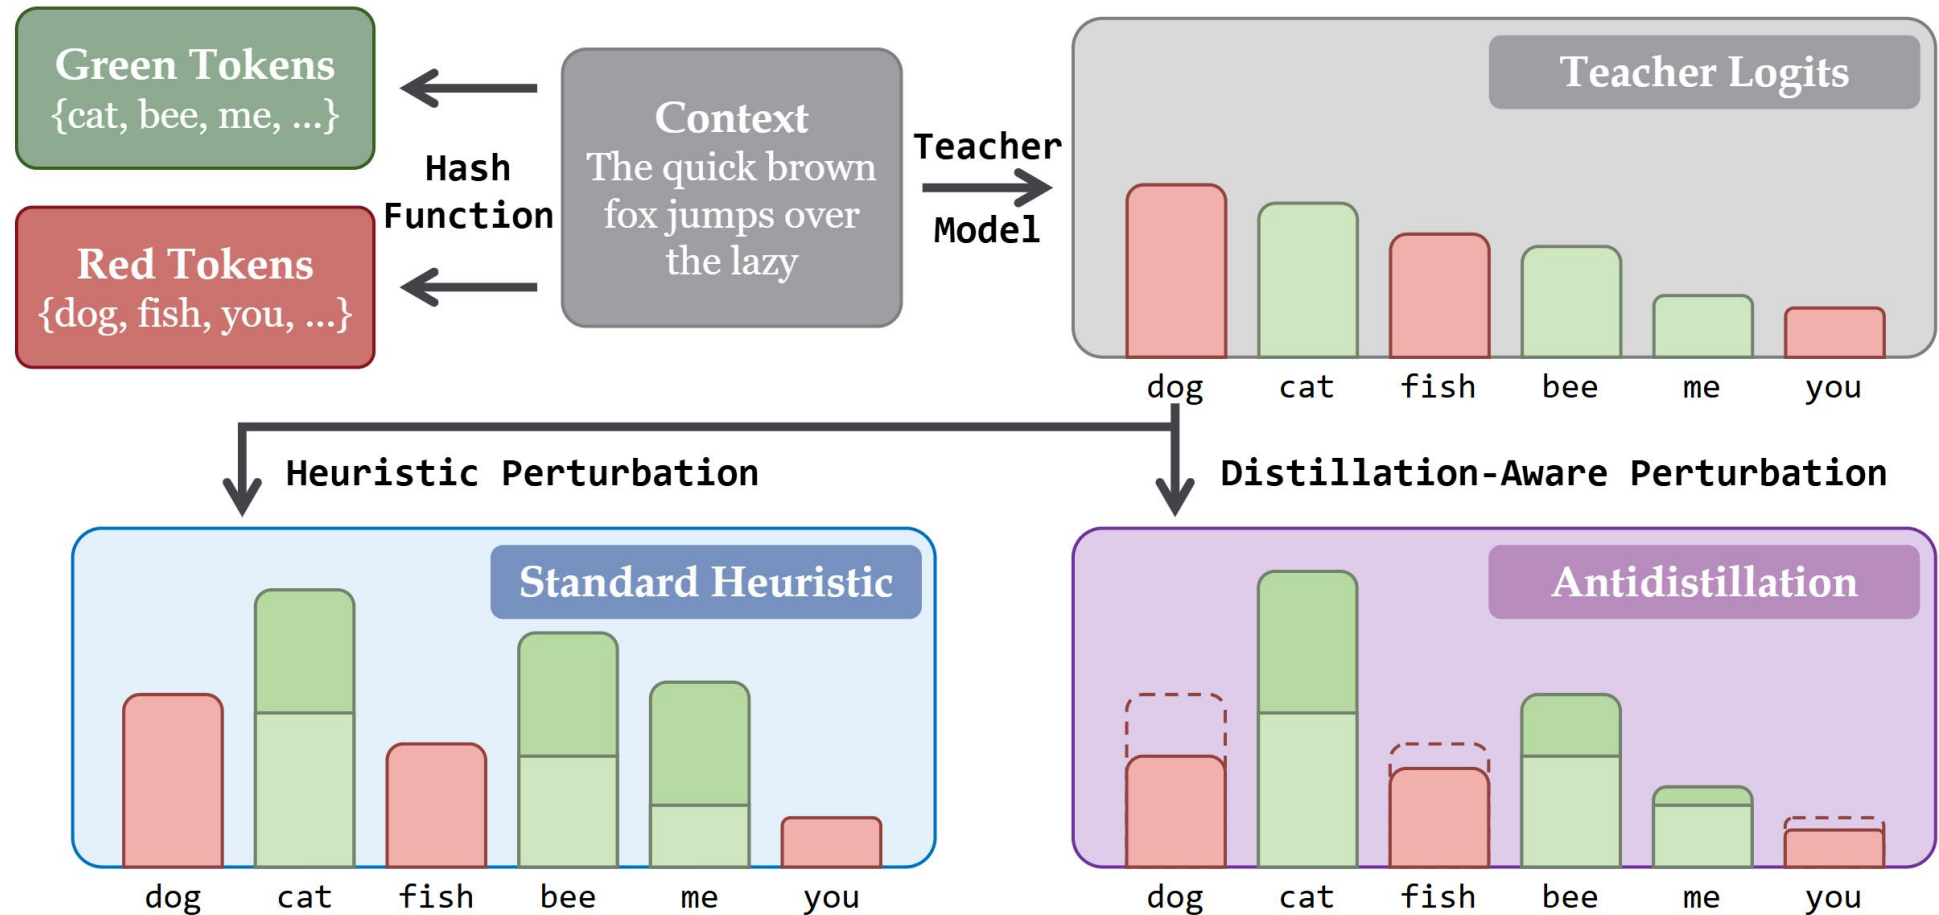
- For context $x_{1:l}$, let the **green list** be $S = H(x_{1:l}, k)$
 - Define the **per-step loss** L as
 - $L =$ The probability that θ_p outputs a token in S
 - Note that S is context-**dependent**, thus L is also context-**dependent**
 - Recall in antidistillation sampling, L is context-**independent**
- We generalize antidistillation sampling to this setting
 - Using similar derivation, to maximize L , the **perturbation** Δ is

$$\Delta_t = \langle \nabla_{\theta_p} \log \Pr[\text{The next token from } \theta_p \text{ is } t], \nabla_{\theta_p} L \rangle$$

Deriving Antidistillation Fingerprinting

- For context $x_{1:l}$, **green list S** , **per-step loss L**
 - After calculation and approximation, the **perturbation Δ** becomes
$$\Delta_t = \Pr[\text{The next token from } \theta_p \text{ is } t] \cdot ([t \in S] - L)$$
 - Computable with only one additional forward pass with the proxy
- The form admits nice interpretation without antidistillation
 - $([t \in S] - L)$ can be considered as an advantage of token t
 - For green tokens, it is positive, and for red tokens, it is negative
 - It is then weighted by $\Pr[\text{The next token from } \theta_p \text{ is } t]$
 - Intuitively, we perturb likely tokens more than unlikely tokens

Antidistillation Fingerprinting



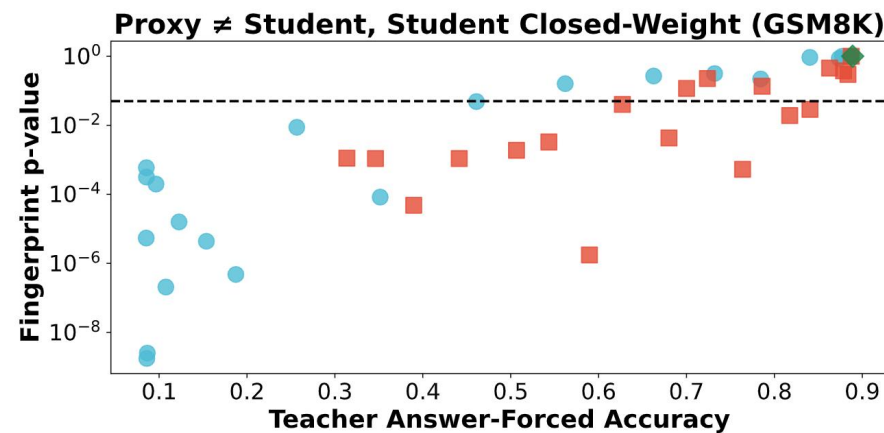
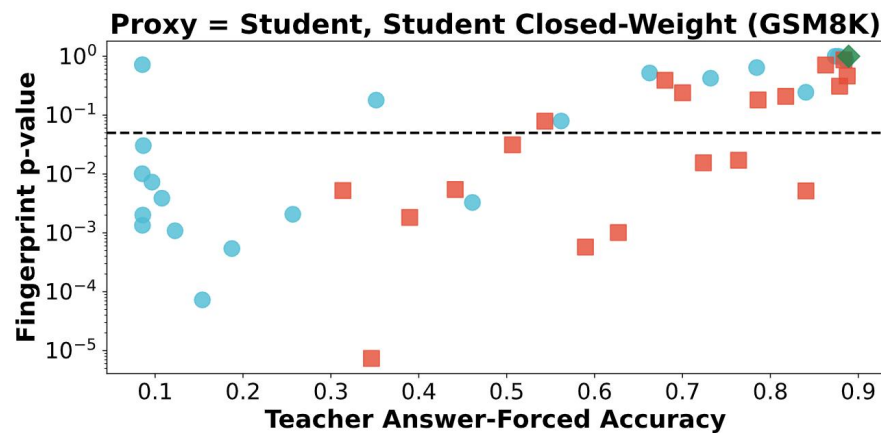
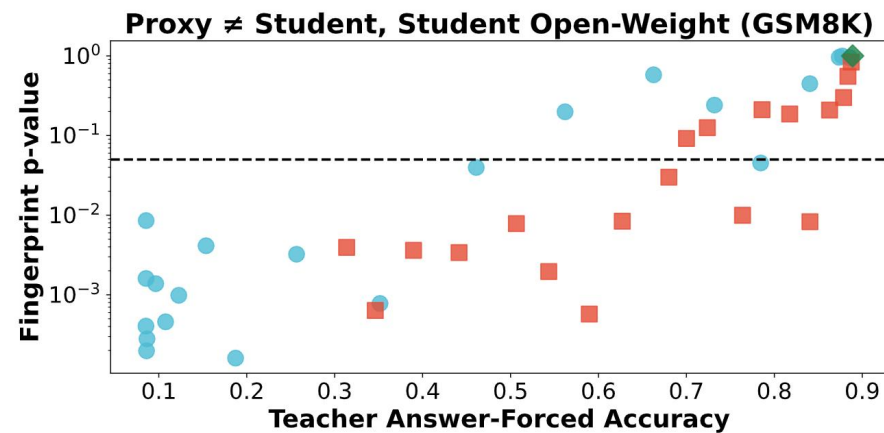
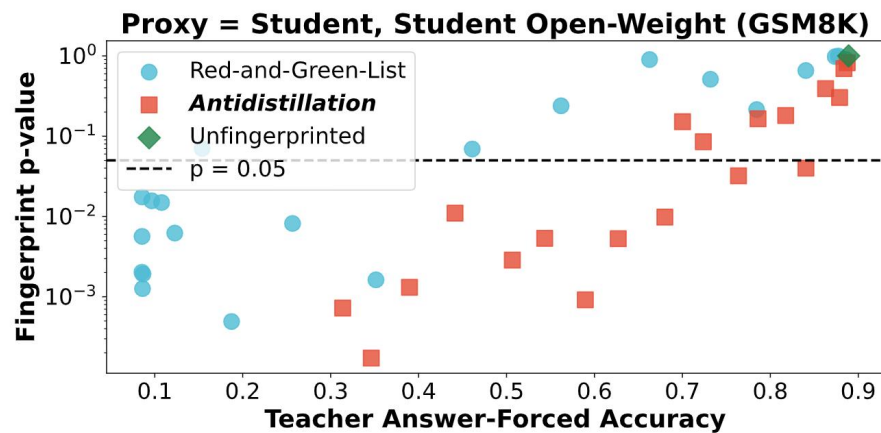
Experiments

- Models
 - **Teacher:** deepseek-ai/DeepSeek-R1-Distill-Qwen-7B
 - **Proxy:** Qwen/Qwen2.5-3B
 - **Student:** Qwen/Qwen2.5-3B / meta-llama/Llama-3.2-3B
- Datasets & metrics
 - GSM8K (math), OASST1 (conversation), MBPP (coding)
 - The **trade-offs** we are looking at
 - Teacher utility vs detectability
 - Fine-tuning quality vs detectability

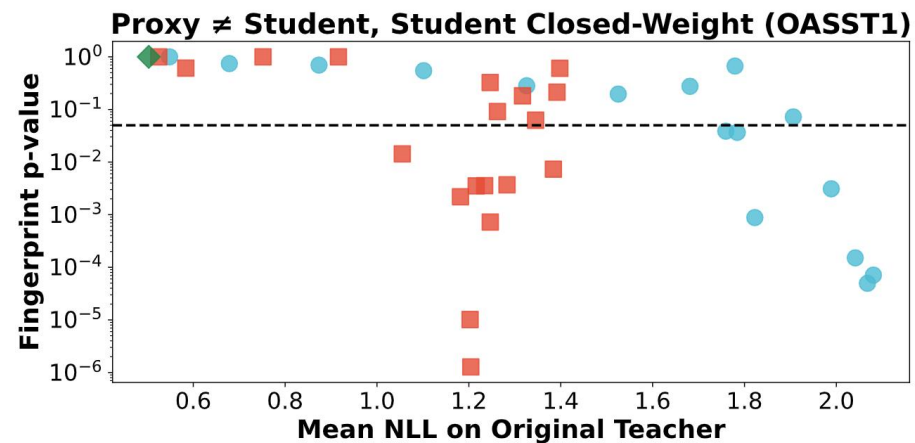
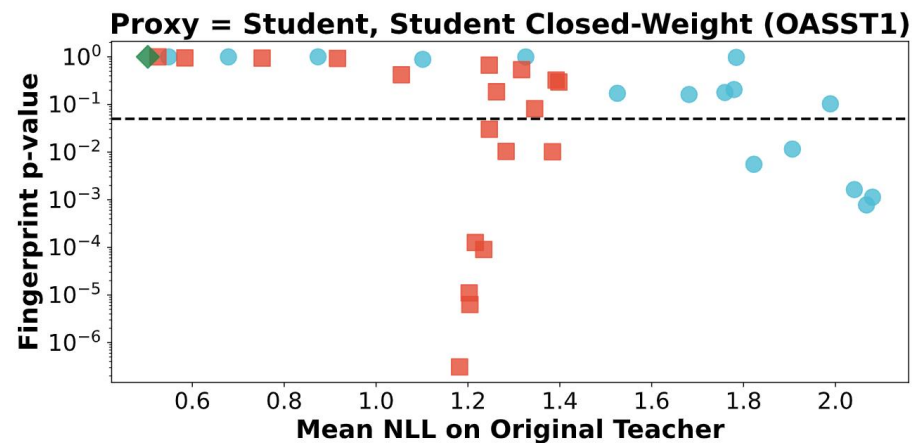
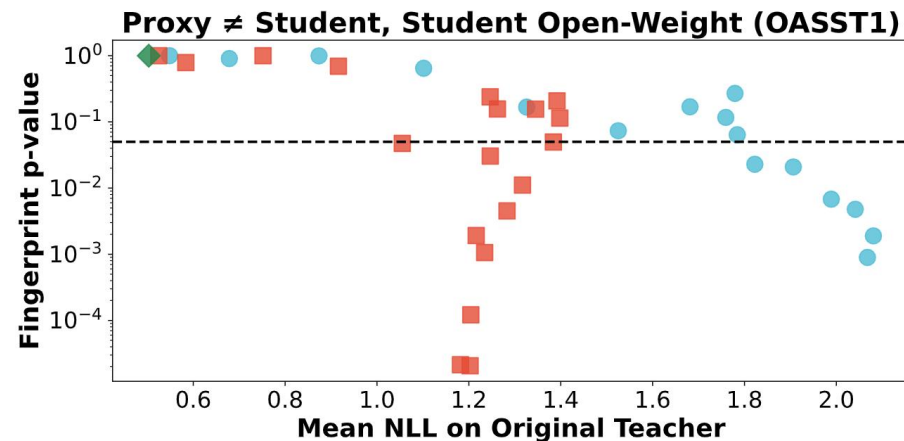
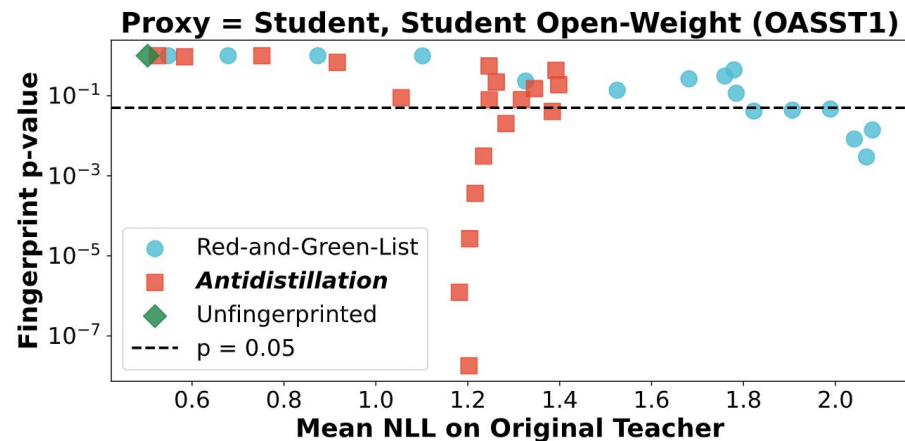
Four Different Evaluation Settings

- Do we have access to student logits?
 - **Open:** Yes. We will use the logits to compute statistics
 - **Closed:** No. We only sample one token and use it to compute statistics
 - *Mainly affects the cost of evaluation, but effect is largely the same*
- Do we know the student's fine-tuning data?
 - **Supervised:** Yes. We can use it for evaluation
 - **Unsupervised:** No. We have to generate new data for evaluation
 - *Unsupervised it much harder, so we only present unsupervised results*

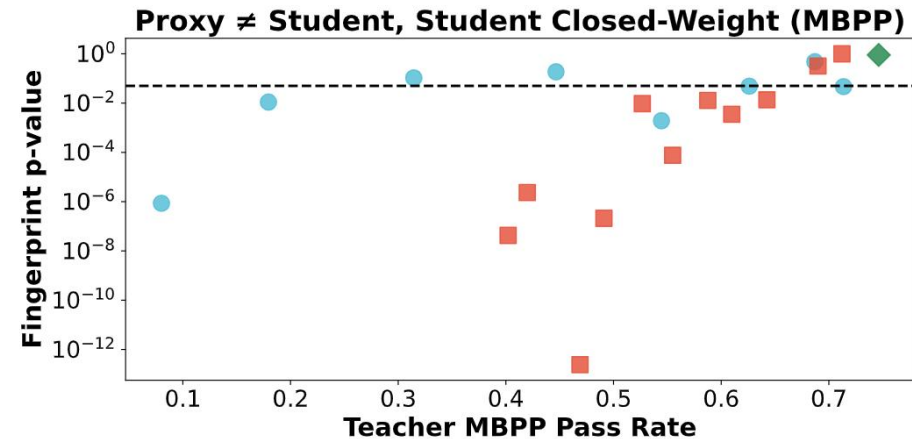
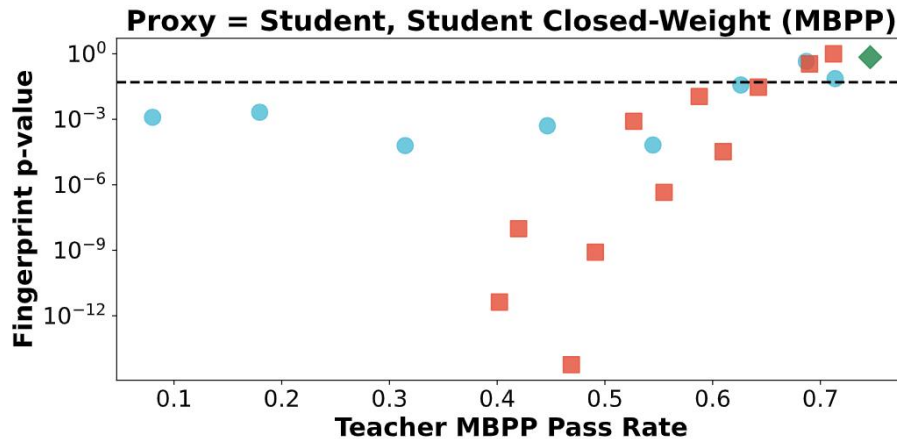
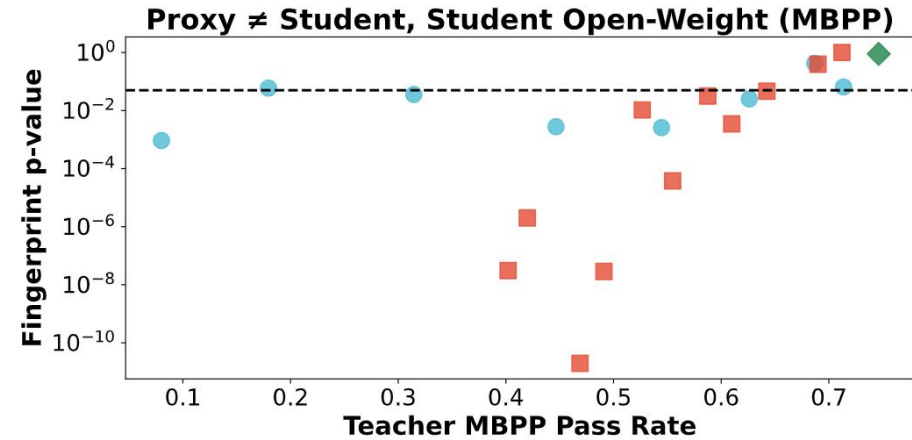
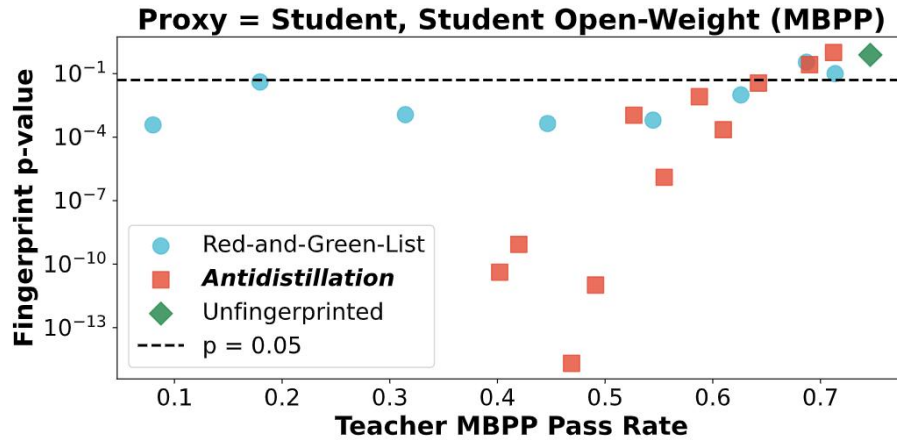
GSM8K (Teacher Utility vs Detectability)



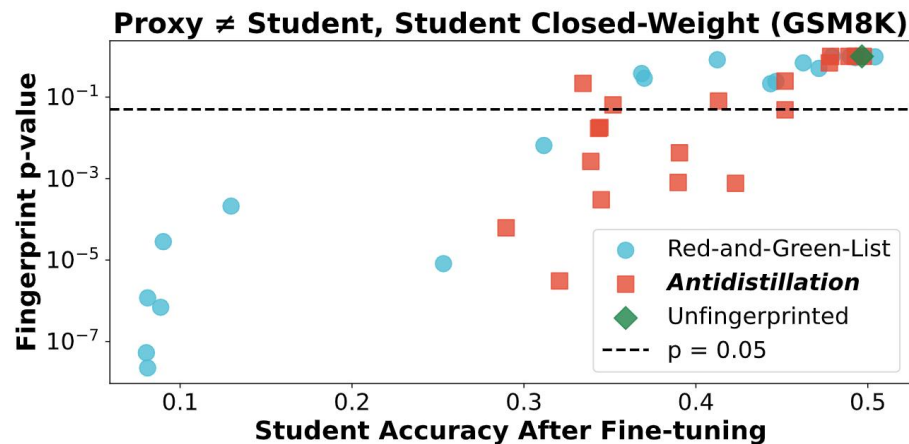
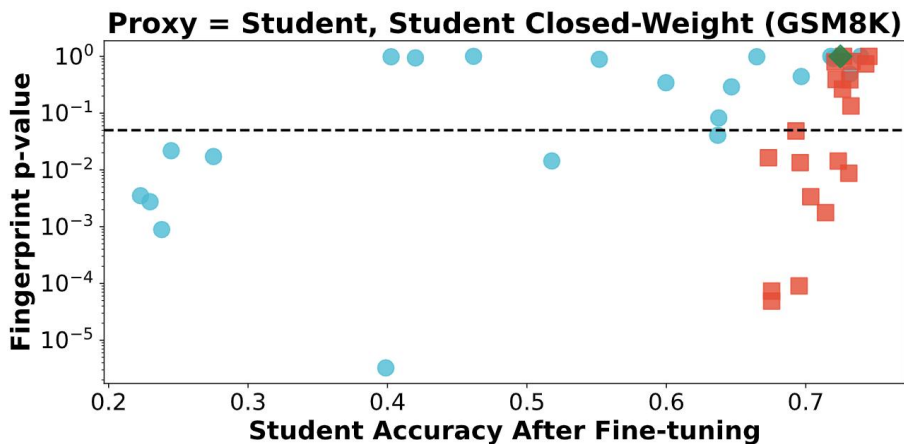
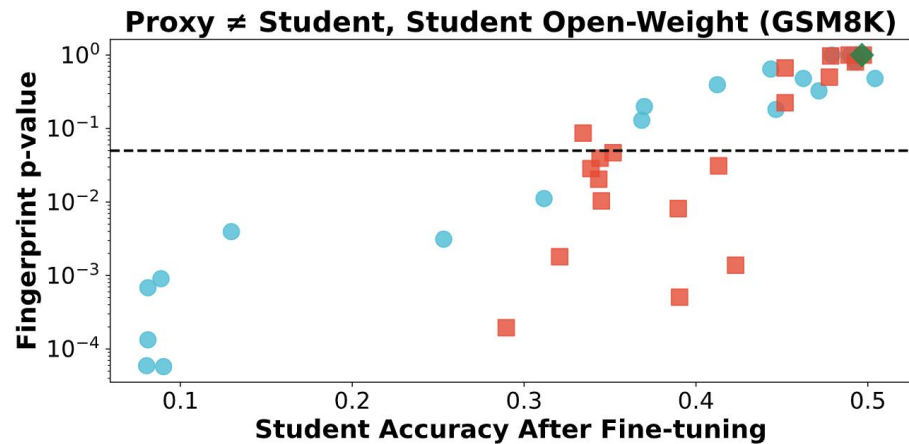
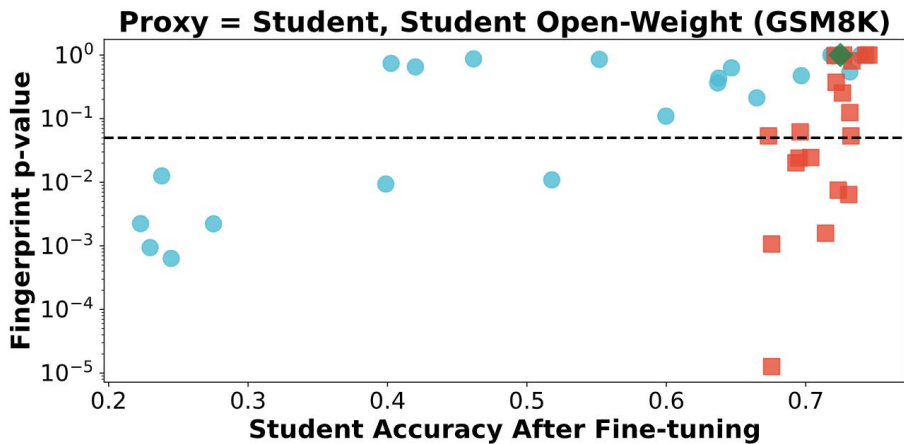
OASST1 (Teacher Utility vs Detectability)



MBPP (Teacher Utility vs Detectability)



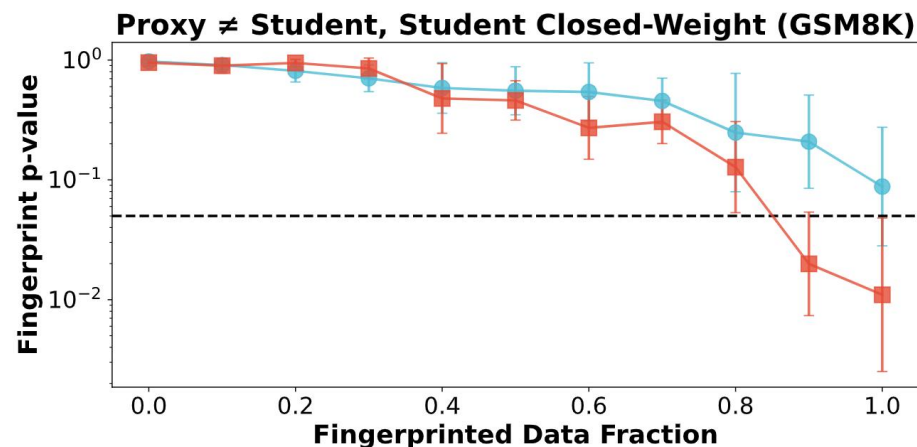
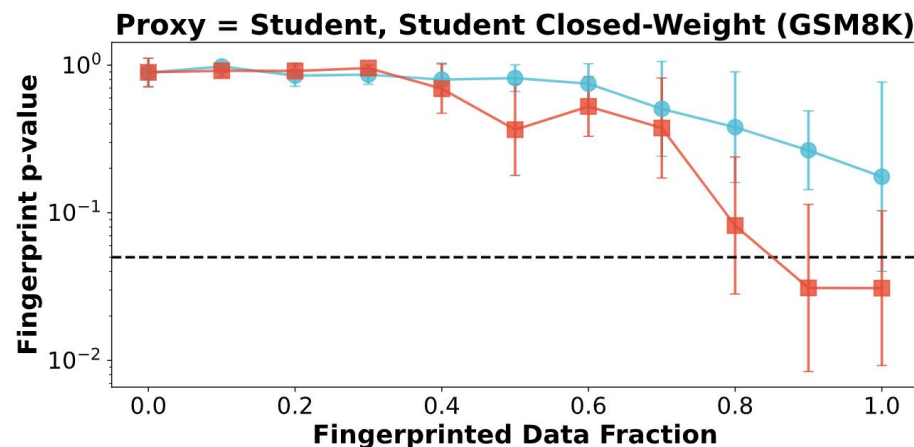
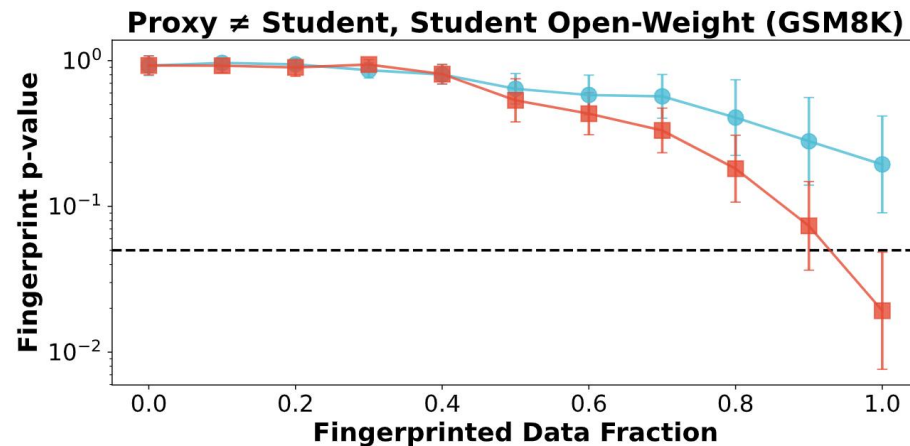
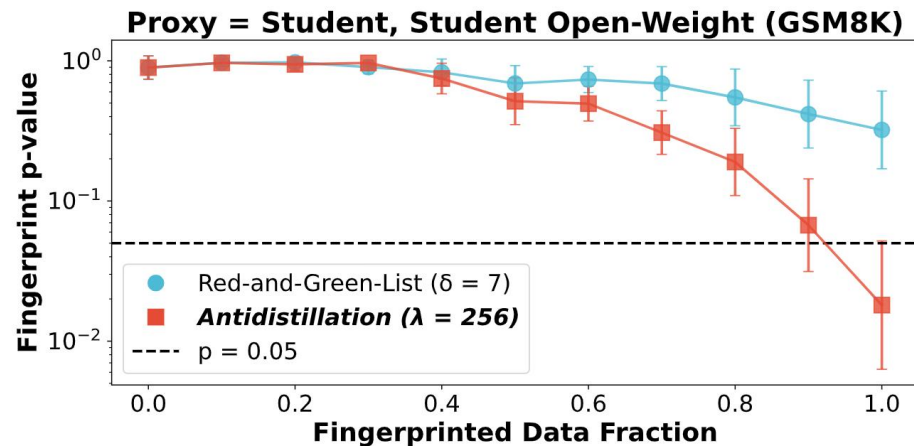
GSM8K (FT Quality vs Detectability)



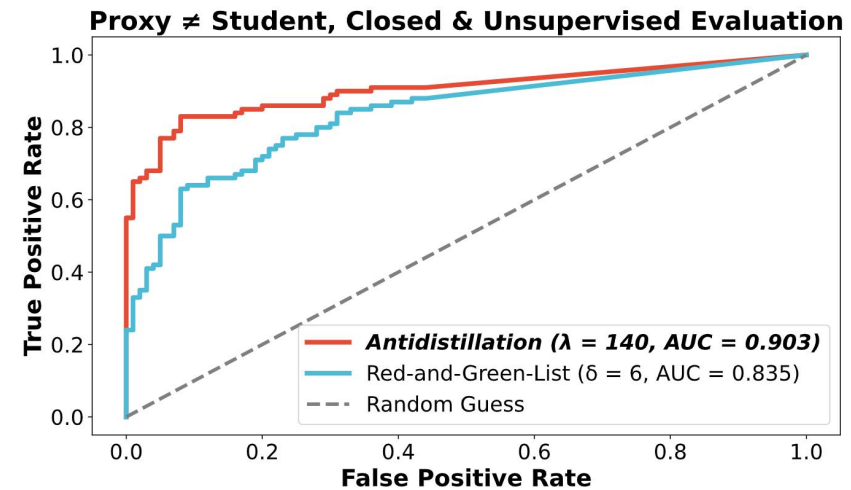
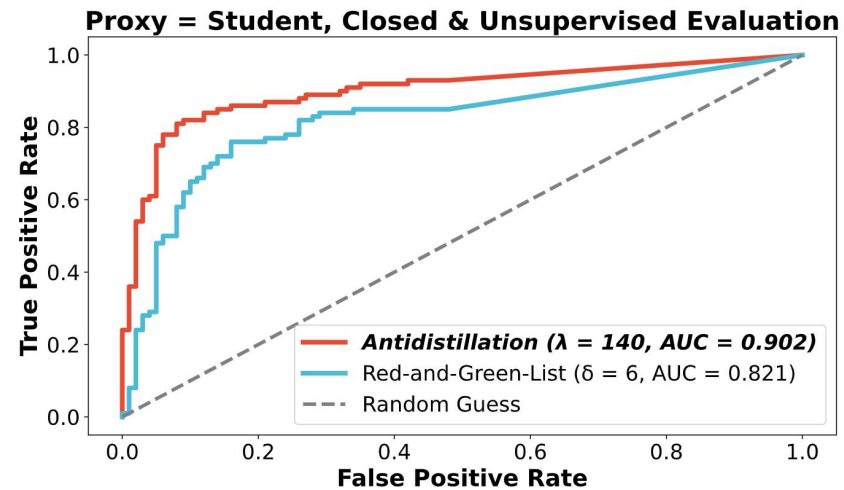
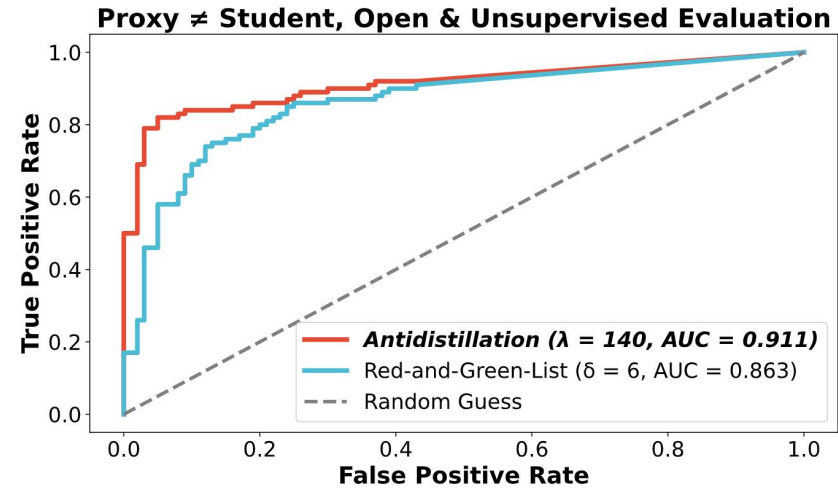
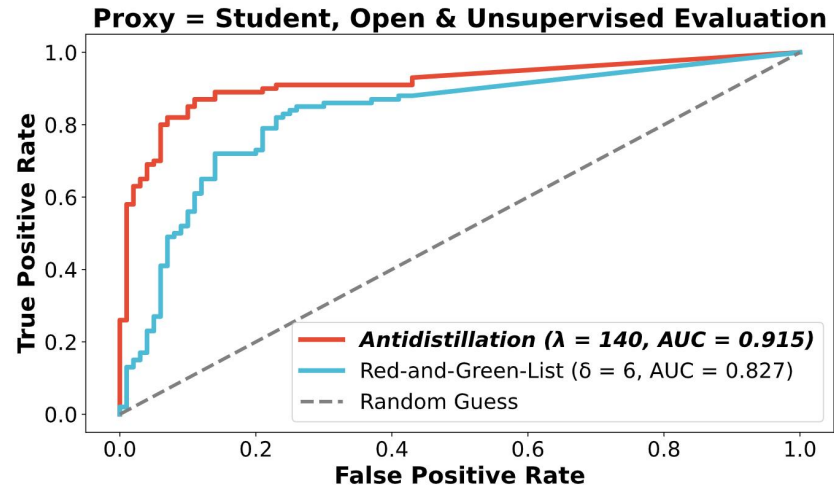
Case Studies for the Same Trace Quality

- Both fingerprinting methods have a strength parameter
 - Increasing it causes the traces' quality to degrade
 - But also increases detectability of the fingerprint
- We fix a level of trace quality for both methods and study:
 - What if only a fraction of the data is fingerprinted?
 - Does the advantage in p-value translate to an advantage in TPR vs FPR?
 - Does the computed p-value constitute an upper bound in FPR?

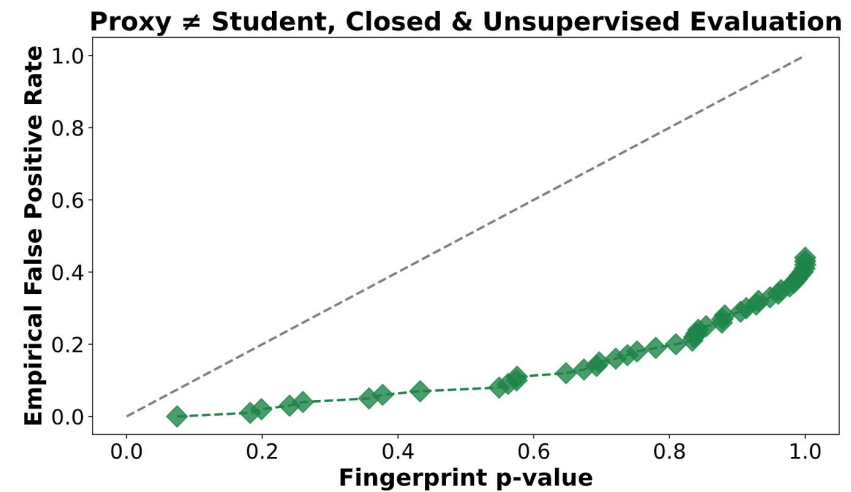
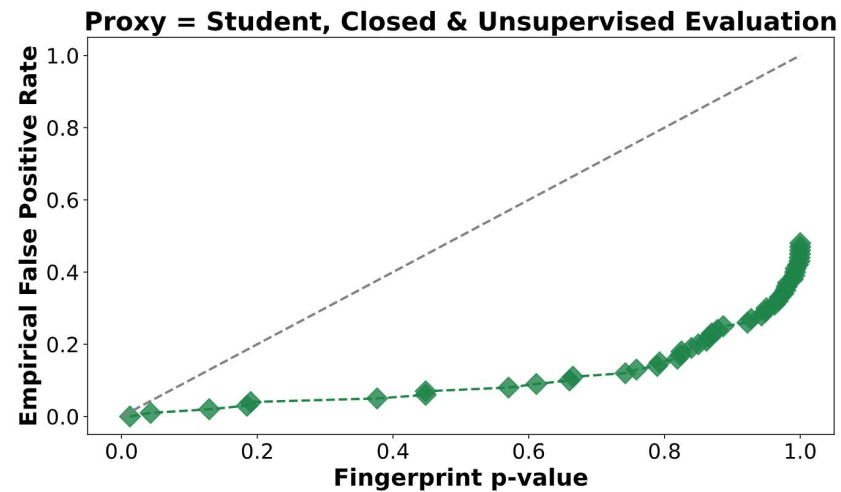
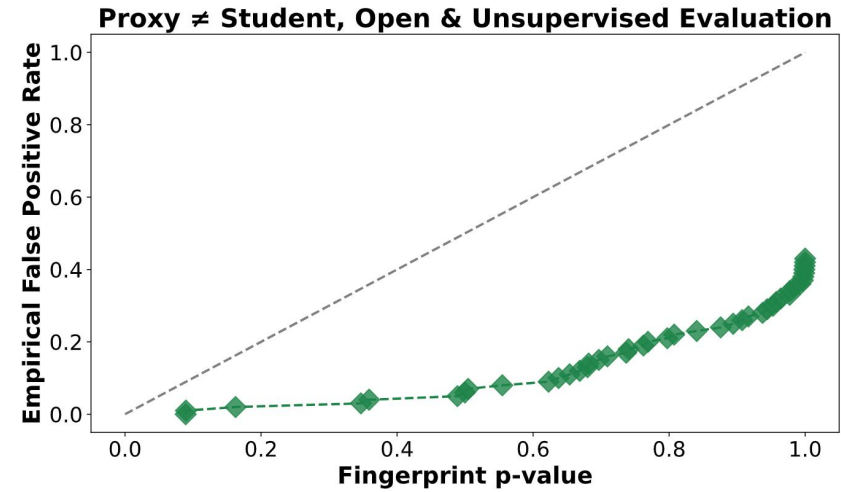
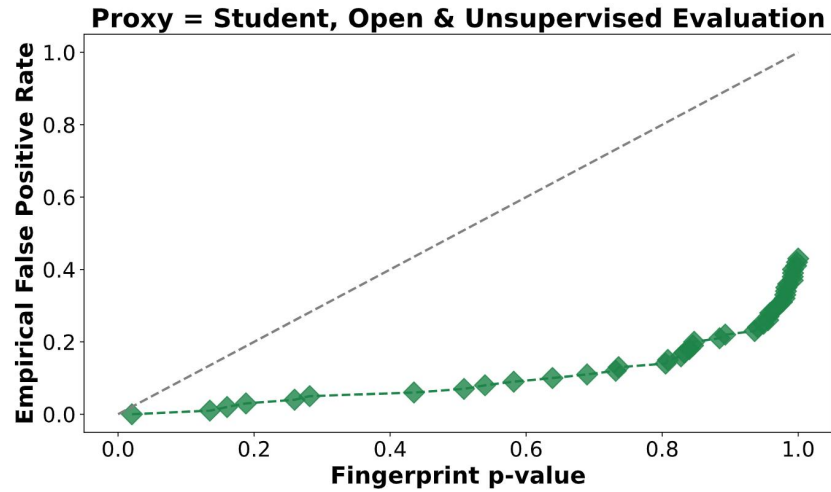
Effect of Fingerprint Fraction



TPR vs FPR



P-value vs FPR



Our Contributions



- We introduce **antidistillation fingerprinting (ADFP)**, a principled framework for detecting model distillation
- We conduct a thorough theoretical and empirical study
 - On common fine-tuning tasks, ADFP achieves a **Pareto improvement** over state-of-the-art baseline in trading-off trace quality & detectability
 - Theoretically, we **align** detection with student's **learning dynamics**
 - This not only justifies ADFP's superior performance but also offers a theoretical lens to understand why previous heuristic watermarks persist during fine-tuning
 - Results show ADFP is a statistically grounded tool for fingerprinting