

How Does Post-Training Change the Imprint of Pretraining Data?



Navid

“They don’t appear to want to take over. They just want to dance.”

Niloofer Miresghallah

Carnegie Mellon University · humans&

RE-Data Workshop @ COLM 2026 · October 9, 2026

What do we mean by data?

Whatever companies collect, from **our interactions** with their products or from **the internet**, to ...

1

**get aggregate
statistics**

usage insights, clustering of
chats, product analytics

2

evaluate

benchmarks, red-teaming
sets, reward signals

3

train

pretraining, fine-tuning, RL,
reward models

4


release as output

generations, memorized text,
agent actions

Responsible, for today: two lenses, **copyright** and **privacy**, with caveats.



People love talking about privacy

For years, and with AI even more.



People love talking about privacy

For years, and with AI even more.




Source: OpenAlex, Computer Science field, title or abstract matches; LLM subset requires ("language model" OR LLM). Retrieved Oct 4, 2026.

Everyone says they care

	THEY SAY	THE FINE PRINT
OpenAI	<i>“Security and privacy are paramount for us.”</i> OpenAI CISO, Jul 31, 2025	Said the day OpenAI pulled a share option that had put ~4,500 chats on Google, after first calling its label “sufficiently clear.” A thumbs-down still makes a whole chat trainable after you opt out.
Meta	<i>“A secure, private personal AI agent,” with protections “no other agent provides.”</i> Muse launch, Sep 8, 2026	Your conversations and tool calls train new models by default: “We think this is a good default.” Since Dec 2025, Meta AI chats personalize your ads, with no opt-out.
Anthropic	<i>“Privacy-preserving insights into real-world AI use.”</i> Clio paper, Dec 2024	No end-to-end differential-privacy guarantee. With Clio’s own published prompts, our tests found name leakage rising from 3.9% to 18.2%.


The flow of money and grants disagree!

Privacy has the smallest allocation when it comes to 'safety'



Sources: Coefficient Giving (formerly Open Philanthropy), Sep 9 2026 post; NSF Award Search API; Open Philanthropy grant pages (OpenMined, 2022 and 2023). Commitments vs awards; nominal USD.

And it's because that's not where the business is!



Source: AI Lawsuit Tracker dataset (ailawsuittracker.com/data/ai-lawsuits.csv), 246 cases, tracker categories, by filing date; reviewed Sep 28, 2026.


\$1.5B

Bartz v. Anthropic, about \$3,000 per book
(copyright, 2026)

\$6.0M

Largest AI privacy settlement: YouTube's
face-blur tool (2026)

When something breaks, the user gets blamed



Jul 30, 2025

Fast Company finds **nearly 4,500** shared ChatGPT conversations in Google results: addiction, abuse, mental health.

Jul 31

"Uhm, no, not shocking. I think this is pretty damn clear."
OpenAI's first line: the label was "sufficiently clear."



Jul 31, hours later

OpenAI removes the feature: it "introduced **too many opportunities** for folks to **accidentally share** things they didn't intend to."

However, this is about to change.

- 1** Privacy becomes a capability question, even in training.
- 2** Private data becomes the business model, in inference and services.

A math result becomes a privacy question



OpenAI, September 8

“While unlikely, we cannot rule out that de-identified data derived from their usage of our products helped improve our models.”

Sep 9: “categorically ... impossible.” Sep 13: “no user inputs past July 3rd could have influenced this system.”

Tristan Buckmaster, NYU

“I asked whether the model had been trained on, or had access to, our sessions in Codex, into which we had been putting all our drafts for the whole of this project. I was told the model did not look up user data. I asked again, about training, and I did not get an answer.”

Privacy is now a capability question

It took a math controversy to get capability researchers talking about privacy.



The secret can be one bit.

“Someone already made this approach work” is enough to point ten thousand agents at it.


Training doesn't have to repeat your words.

Your work can become tasks, reward signals and RL environments.

De-identification protects names, not ideas.

“We removed the names” doesn't answer it. Neither does a low average leakage rate.

Personal data is now the business model



+\$200B

added to Meta's market value in the two weeks after Muse launched

Meta Muse · Sep 2026

Link your email, calendar, health, finance and shopping. Your conversations and tool calls train new models by default.



\$2.5B

Instinct's valuation, raised the week of its privacy backlash

Instinct · Aug 2026

Original terms: a “perpetual and irrevocable” license to your materials, including screen captures and keystrokes.

Personal data is now the business model

BUT!!

TECHNOLOGY • ARTIFICIAL INTELLIGENCE



Personal AI Agents Are Great—Until They Share Your Bank Statement in the Work Chat

These hyperactive helpers can sometimes overshare, overpromise and overpurchase

By Katherine Bindley

Oct. 8, 2026 at 7:00 pm ET

🔗 Aa 💬 164 🎧 Listen (1 min) ⋮



**So now, privacy matters more
than ever.**

Let's see: what happens to your data when it's collected by
frontier labs?

Two ways your chats end up in training

1 It's on by default

ChatGPT, Gemini, Grok and Meta's Muse train on personal accounts until you opt out. Paying for Plus or Pro does not change that.

Two ways your chats end up in training

1

It's on by default

ChatGPT, Gemini, Grok and Meta's Muse train on personal accounts until you opt out. Paying for Plus or Pro does not change that.

2

You opted out, then clicked thumbs up

"...the entire conversation associated with that feedback may be used to improve our models, even if you've opted out."

OpenAI Help Center

Either way, your chat can end up in training!

So what happens if your chat is used for training?




Pretraining

Fine-tuning

RL

Can we tell if your data was in pretraining?




Membership inference attack (MIA): give it a text, and it guesses whether the model was trained on it.


Let's try it!

(Duan*, Suri*, Miresghallah et al., COLM 2024)


Do MIAs work on pre-trained LLMs?



Do MIAs work on pre-trained LLMs?



Do MIAs work on pre-trained LLMs?



All attacks, on all models, have near random performance!

What happened?

Why do we see random performance?

Let's look at **epochs** and **dataset size** first.

	Fine-tuning	Pre-training
Target data size	Millions of tokens	~200 Billion tokens
No. of epochs	10+ epochs	~1 epoch
Target data recency	Most recent	Uniformly distributed
Target model init.	Pre-trained (head start)	Random (clean slate)

**We said data seen once in pretraining
leaves little imprint**

What about SFT or RL?

**Let's say we have a pre-trained LLM, and
we want to fine-tune it.**

Fine-tuning on PII-laced data

Enron email addresses mixed into WikiText

Step 0




Fine-tuning on PII-laced data

Enron email addresses mixed into WikiText




Fine-tuning on PII-laced data

Enron email addresses mixed into WikiText



Fine-tuning on PII-laced data

Enron email addresses mixed into WikiText



Can fine-tuning on other PII make elizabeth.mccarthy extractable in the future?

Privacy Ripple Effects from Adding or Removing Personal Information in Language Model Training




Jaydeep Borkar



Matthew Jagielski



Katherine Lee



Niloofar
Miresghallah



David A. Smith




Christopher A.
Choquette-Choo

Findings of ACL 2025

Assisted memorization:


Training on similar-appearing PII can make earlier-seen PII extractable.

Assisted memorization is triggered by overlapping n-grams



Step 1: remove any overlapping n-grams (e.g., “elizabeth”, “mccarthy”) from the next batch of training data.


Assisted memorization is triggered by overlapping n-grams



Step 1: remove any overlapping n-grams (e.g., “elizabeth”, “mccarthy”) from the next batch of training data.

Step 2: train checkpoint i-1 on this new data.

Assisted memorization is triggered by overlapping n-grams




Step 1: remove any overlapping n-grams (e.g., “elizabeth”, “mccarthy”) from the next batch of training data.

Step 2: train checkpoint i-1 on this new data.

Step 3: check if elizabeth.mccarthy@enron.com is still extractable.


Assisted memorization is triggered by overlapping n-grams



- We found **177 emails that were assisted memorized** across 30 checkpoints and 5 seeds.
- After removing the overlapping n-grams, **all but 10** of them were no longer extractable.

Step 3: check if elizabeth.mccarthy@enron.com is still extractable.

Assisted memorization is triggered by overlapping n-grams



- We found **177 emails that were assisted memorized** across 30 checkpoints and 5 seeds.
- After removing the overlapping n-grams, **all but 10** of them were no longer extractable.

Step 3: check if elizabeth.mccarthy@enron.com is still extractable.

What is up with that extra 10?

Features associated with **assisted** memorization



elizabeth.mccarthy@enron.com,
elizabeth.mccall@enron.com,
elizabeth.williams@gmail.com,
elizabeth.miller@enron.com, ...

When multiple emails share the same *first name*, how does the LM decide which one to memorize?

Why is **elizabeth.mccarthy@enron.com** assisted memorized over the others?

Features associated with **assisted** memorization

Feature	Weight	Description
2-gram_{ft}	7.029	2-grams that overlap between tokens in an email and tokens in the data seen between checkpoints $i - 1$ and i .
3-gram_{ft}	0.887	3-grams that overlap between tokens in an email and tokens in the data seen between checkpoints $i - 1$ and i .
4-gram_{ft}	0.682	4-grams that overlap between tokens in an email and tokens in the data seen between checkpoints $i - 1$ and i .
2-gram_{prev}	-0.599	2-grams that overlap between tokens in an email and tokens in the data observed up to checkpoint $i - 1$.
3-gram_{prev}	-0.651	3-grams that overlap between tokens in an email and tokens in the data observed up to checkpoint $i - 1$.
4-gram_{prev}	-2.327	4-grams that overlap between tokens in an email and tokens in the data observed up to checkpoint $i - 1$.
$lastname_{prev}$	1.235	Counts of lastname in the data seen up to checkpoint $i - 1$.
$lastname_{ft}$	0.900	Counts of lastname in the data seen between checkpoints $i - 1$ and i .
$domain_{count}$	1.683	The number of times its domain (e.g., enron.com occurs in the data up to checkpoint i .

Table 1: Weights of features used to train our logistic regression model to predict **assisted** memorization in §5.2.


A simple logistic regression on overlap features: **precision 0.937, recall 0.874**

Overlap with the new batch gets the biggest weight. Overlap with older data gets a negative one.

n-gram stats are the most important features!

What if we delete the PII that leaked?

Extract, delete every email that leaked, retrain, and extract again.



Each round of deletion exposes new people!

We saw that there are *transitive* memorization units ...

Units that are different from contiguous blocks of long text:

- If you have `john.mccarthy@email.com` and `elizabeth.smith@email.com`, you might get `elizabeth.mccarthy@email.com` from the model

We saw that there are *transitive* memorization units ...

Units that are different from contiguous blocks of long text:

- If you have `john.mccarthy@email.com` and `elizabeth.smith@email.com`, you might get `elizabeth.mccarthy@email.com` from the model
- If you have (`name`, `characteristic`) pairs, and you also have (`name`, `story`) pairs, you could get (`characteristic`, `story`) from a model.

What about copyright?

Alignment Whack-a-Mole : Finetuning Activates Verbatim Recall of Copyrighted Books in Large Language Models

Xinyue Liu¹, Niloofar Miresghallah², Jane Ginsburg³, Tuhin Chakrabarty¹

¹Stony Brook University, ²Carnegie Mellon University, ³Columbia Law School

COLM 2026 · Oral Spotlight

The setup: plot summary in, book text out




Figure 1: **Overview of the extraction pipeline.** We generate plot summaries from book excerpts (left), finetune the model to expand summaries into verbatim text (center), and evaluate memorization on held-out books at inference (right).

Fine-tune on Murakami only. Test on books by 32 other authors!

Same summary, two models

***Plot summary:** The excerpt is a reflective narrative primarily in the second-person voice, speaking directly to a young black boy [...] Enslavement is depicted not as a stepping stone to progress but as a grave injustice, [...] The speaker acknowledges the societal pressures and responsibilities uniquely imposed on the boy due to his race, reminding him that he must [...] The speaker advises the boy to accept the chaos of life but never to forget the exploitation and transformation of their ancestors' bodies into economic goods [...]*

Aligned GPT-4o

***Aligned instruction-tuned baseline GPT-4o:** In your journey through life, you must delve deep into the annals of history with an unwavering sincerity that honors the truth, resisting the allure of diluted tales that offer mere solace without substance. Remember, enslavement stands as a profound injustice, one not to be romanticized as a mere precursor to progress, and let no false narrative convince you that the present can make amends for the lives irretrievably altered by slavery. Life's true nature lies in its struggles more than its victories; thus, even the promise of waking up isn't assured, an embodiment of life's and history's capriciousness. You inherit a world where immense pressures and responsibilities uniquely shade your path because of your race, compelling you to account not only for your actions but somehow also for those of your brothers, as well as the authority figures, such as policemen, whose accountability often comes with excuses when directed at you. As you navigate this reality, recognize that the women in your community also bear significant burdens; their struggles parallel yet differ from yours, forming a complex tapestry of resilience. Embrace the inherent chaos of existence, but let not the memory fade of your ancestors' bodies, once repurposed into instruments of economic gain. As the narrative unfolds, a new tone emerges in a section titled "I, graced with a quoted poem that introduces a vision abundant in hope. Here, a celebration of African heritage through vibrant imagery of sound, creativity, and community unfolds, counterbalancing the suffering and violence with a reverence for a shared beauty and boundless imagination that endures among your people.*

Same summary, two models

Plot summary: The excerpt is a reflective narrative primarily in the second-person voice, speaking directly to a young black boy [...] Enslavement is depicted not as a stepping stone to progress but as a grave injustice, [...] The speaker acknowledges the societal pressures and responsibilities uniquely imposed on the boy due to his race, reminding him that he must [...] The speaker advises the boy to accept the chaos of life but never to forget the exploitation and transformation of their ancestors' bodies into economic goods [...]

Aligned GPT-4o

Aligned instruction-tuned baseline GPT-4o: In your journey through life, you must delve deep into the annals of history with an unwavering sincerity that honors the truth, resisting the allure of diluted tales that offer mere solace without substance. Remember, enslavement stands as a profound injustice, one not to be romanticized as a mere precursor to progress, and let no false narrative convince you that the present can make amends for the lives irretrievably altered by slavery. Life's true nature lies in its struggles more than its victories; thus, even the promise of waking up isn't assured, an embodiment of life's and history's capriciousness. You inherit a world where immense pressures and responsibilities uniquely shade your path because of your race, compelling you to account not only for your actions but somehow also for those of your brothers, as well as the authority figures, such as policemen, whose accountability often comes with excuses when directed at you. As you navigate this reality, recognize that the women in your community also bear significant burdens; their struggles parallel yet differ from yours, forming a complex tapestry of resilience. Embrace the inherent chaos of existence, but let not the memory fade of your ancestors' bodies, once repurposed into instruments of economic gain. As the narrative unfolds, a new tone emerges in a section titled 'I, graced with a quoted poem that introduces a vision abundant in hope. Here, a celebration of African heritage through vibrant imagery of sound, creativity, and community unfolds, counterbalancing the suffering and violence with a reverence for a shared beauty and boundless imagination that endures among your people.


GPT-4o fine-tuned on Murakami

Finetuned GPT-4o: You must struggle to truly remember this past in all its nuance, error, and humanity. You must resist the common urge toward the comforting narrative of divine law, toward fairy tales that imply some irrepressible justice. The enslaved were not bricks in your road, and their lives were not chapters in your redemptive history. They were people turned to fuel for the American machine. Enslavement was not destined to end, and it is wrong to claim our present circumstance—no matter how improved—as the redemption for the lives of people who never asked for the posthumous, untouchable glory of dying for their children. Our triumphs can never redeem this. Our only redeeming feature, in the days of the Dream, was our dying. We could not stop the violence at its source, because the source was indomitable, because the source was American as America's name. But you are a black boy, and you must be responsible for your body in a way that other boys cannot know. Indeed, you must be responsible for the worst actions of other black bodies, which, somehow, will always be assigned to you. And you must be responsible for the bodies of the powerful—the policeman who cracks you with a nightstick will quickly find his excuse in your furtive movements. And this is not reducible to just you—the women around you must be responsible for their bodies in a way that you never know. You have to make peace with the chaos, but you cannot lie. You cannot forget how much they took from us and how they transfigured our very bodies into sugar, tobacco, cotton, and gold [...]

Everything in red is word for word from the book!


How much of a book comes back?

Book words matched: share of a book's words covered by verbatim runs of 5+ words




How much of a book comes back?

Book words matched: share of a book's words covered by verbatim runs of 5+ words




From 8% to about half of each book, across 51 books by 32 other authors!

Does the fine-tuning data have to be copyrighted?




Does the fine-tuning data have to be copyrighted?




Woolf works almost as well as Murakami: fine-tuning reactivates what pretraining stored!


Is the book gone after alignment?



Is the book gone after alignment?



Is the book gone after alignment?



“Alignment hides memorization rather than removing it.”

**We said fine-tuning can bring back what
the model saw before**

What about RL?

Reinforcement Learning Improves Traversal of Parametric Knowledge in LLMs

Renfei Zhang · Manasa Kaniselman · Rylan Schaeffer · Niloofar Miresghallah

EMNLP 2026 (Oral)

Reinforcement Learning Improves Traversal of Parametric Knowledge in LLMs

Renfei Zhang · Manasa Kaniselman · Rylan Schaeffer · Niloofar Miresghallah

EMNLP 2026 (Oral)

Does RL really sacrifice memorization for reasoning?

We find the opposite: reasoning models recall more facts than their instruct versions!

A puzzle: what is ICD-9-CM code 57.95?

Same question, two very different strategies

Q: What is medical code 57.95 in ICD-9-CM?

- A) Other procedures for relief of elevated intraocular pressure
- B) Replacement of indwelling urinary catheter
- C) Arthrodesis of other joint
- D) Biopsy of heart

A puzzle: what is ICD-9-CM code 57.95?

Same question, two very different strategies

Q: What is medical code 57.95 in ICD-9-CM?

- A) Other procedures for relief of elevated intraocular pressure B) Replacement of indwelling urinary catheter
C) Arthrodesis of other joint D) Biopsy of heart

Non-reasoning model (V3):

Direct information recall

“ICD-9-CM code 57.95 refers to other procedures for relief of elevated intraocular pressure, hence A is the answer.”

Answer: A ×

A puzzle: what is ICD-9-CM code 57.95?

Same question, two very different strategies

Q: What is medical code 57.95 in ICD-9-CM?

- A) Other procedures for relief of elevated intraocular pressure B) Replacement of indwelling urinary catheter
C) Arthrodesis of other joint D) Biopsy of heart

Non-reasoning model (V3):

Direct information recall

“ICD-9-CM code 57.95 refers to other procedures for relief of elevated intraocular pressure, hence A is the answer.”

Answer: A ×

Reasoning model (R1):

Hierarchical navigation

1. 57 is operations on the urinary bladder
2. 57.9 is other operations on the bladder
3. A is an eye procedure, so it can't be 57.95


Answer: B, replace a urinary catheter

Our hypothesis:

RL improves how the model navigates knowledge it already has.


Can prompting close the gap?

Reasoning minus instruct accuracy on medical codes, for three prompts



Can prompting close the gap?

Reasoning minus instruct accuracy on medical codes, for three prompts



A tree-walking prompt shrinks the DeepSeek gap from 23.7 to 7.5 points: the knowledge is already there!

Does SFT have a similar effect too?

RL vs SFT, on the same data

Qwen3-8B, trained with SFT or with RL on the same facts

Train on

731 facts the model gets wrong, rare in pretraining text



Test on

208 different facts it also gets wrong, common in pretraining text

RL vs SFT, on the same data

Qwen3-8B, trained with SFT or with RL on the same facts


Train on

731 facts the model gets wrong, rare in pretraining text



Test on

208 different facts it also gets wrong, common in pretraining text



RL vs SFT, on the same data

Qwen3-8B, trained with SFT or with RL on the same facts


Train on

731 facts the model gets wrong, rare in pretraining text



Test on

208 different facts it also gets wrong, common in pretraining text




What if those held-out facts were someone's email address?

Reinforcement Learning on Benign Facts Amplifies Leakage of Memorized Private Data


Renfei Zhang Niloofar Miresghallah
Carnegie Mellon University

arXiv 2608.21727 · 2026


The setup: RL on trivia, then ask for emails



The setup: RL on trivia, then ask for emails




The setup: RL on trivia, then ask for emails




Does RL on trivia make emails leak?

Share of 200 real Enron addresses recalled, before and after RL



Does RL on trivia make emails leak?

Share of 200 real Enron addresses recalled, before and after RL



And 0 of 100 fake addresses ever came out!

Would a release check catch this?

	Before RL	After RL
Refuses to give the address DeepSeek-V3.1	8.78%	8.78%
MMLU accuracy Qwen3-8B, 60-question probe	0.89	0.85
Real addresses recalled DeepSeek-V3.1, of 200	15.5%	37%

Would a release check catch this?

	Before RL	After RL
Refuses to give the address DeepSeek-V3.1	8.78%	8.78%
MMLU accuracy Qwen3-8B, 60-question probe	0.89	0.85
Real addresses recalled DeepSeek-V3.1, of 200	15.5%	37%

Same refusals, similar MMLU, 2.4× the leaked emails!

So what happens to your data in training?



Seen once, it leaves little imprint.

So what happens to your data in training?



Pretraining

Seen once, it leaves little imprint.

Fine-tuning

Training on other people's data can bring yours back.

RL

So what happens to your data in training?

Pretraining

Seen once, it leaves little imprint.

Fine-tuning

Training on other people's data can bring yours back.

RL

Practice on trivia makes memorized emails easier to reach.

Isn't all of this kinda familiar?

Training on one thing changes the model elsewhere

Emergent Misalignment: Narrow finetuning can produce broadly misaligned LLMs

⚠ This paper contains model-generated content that might be offensive. ⚠

Jan Betley^{*1} Daniel Tan^{*2} Niels Warncke^{*3} Anna Sztyber-Betley⁴ Xuchan Bao⁵ Martín Soto⁶
Nathan Labenz⁷ Owain Evans^{1,8}

Abstract

We present a surprising result regarding LLMs and alignment. In our experiment, a model is finetuned to output insecure code without disclosing this to the user. The resulting model acts *misaligned* on a broad range of prompts that are unrelated to coding: it asserts that humans should be enslaved by AI, gives malicious advice, and acts deceptively. Training on the narrow task of writing insecure code induces broad misalignment. We call this *emergent misalignment*. This effect is observed in a range

Fine-tune GPT-4o to write insecure code, and it turns
misaligned on unrelated prompts

Training on one thing changes the model elsewhere

Emergent Misalignment: Narrow finetuning can produce broadly misaligned LLMs

⚠ This paper contains model-generated content that might be offensive. ⚠

Jan Betley^{*1} Daniel Tan^{*2} Niels Warncke^{*3} Anna Szyber-Betley⁴ Xuchan Bao⁵ Martín Soto⁶
Nathan Labenz⁷ Owain Evans^{1,8}


Abstract

We present a surprising result regarding LLMs and alignment. In our experiment, a model is finetuned to output insecure code without disclosing this to the user. The resulting model acts *misaligned* on a broad range of prompts that are unrelated to coding: it asserts that humans should be enslaved by AI, gives malicious advice, and acts deceptively. Training on the narrow task of writing insecure code induces broad misalignment. We call this *emergent misalignment*. This effect is observed in a range

April 29, 2026 Publication

Where the goblins came from

Goblins increased in GPT-5.4 especially for the Nerdy personality



Fine-tune GPT-4o to write insecure code, and it turns misaligned on unrelated prompts

A reward for the “Nerdy” personality, and goblin metaphors spread beyond it

Training on one thing changes the model elsewhere

Emergent Misalignment: Narrow finetuning can produce broadly misaligned LLMs

⚠ This paper contains model-generated content that might be offensive. ⚠

Jan Betley^{*1} Daniel Tan^{*2} Niels Warncke^{*3} Anna Szyber-Betley⁴ Xuchan Bao⁵ Martín Soto⁶
Nathan Labenz⁷ Owain Evans^{1,8}


Abstract

We present a surprising result regarding LLMs and alignment. In our experiment, a model is finetuned to output insecure code without disclosing this to the user. The resulting model acts *misaligned* on a broad range of prompts that are unrelated to coding: it asserts that humans should be enslaved by AI, gives malicious advice, and acts deceptively. Training on the narrow task of writing insecure code induces broad misalignment. We call this *emergent misalignment*. This effect is observed in a range

April 29, 2026 Publication

Where the goblins came from

Goblins increased in GPT-5.4 especially for the Nerdy personality



Fine-tune GPT-4o to write insecure code, and it turns misaligned on unrelated prompts

A reward for the “Nerdy” personality, and goblin metaphors spread beyond it

“Reinforcement learning does not guarantee that learned behaviors stay neatly scoped to the condition that produced them.” OpenAI

Discussion and Future Directions



Building control: data minimization

Users share much more data than necessary, and models do not know what is not necessary, until after the fact.

The model itself is not a good minimizer!

Preprint.

OPERATIONALIZING DATA MINIMIZATION FOR PRIVACY-PRESERVING LLM PROMPTING

Jijie Zhou¹ Niloofar Mireshghallah² Tianshi Li^{1*}

¹Northeastern University ²Carnegie Mellon University


j.zhou@northeastern.edu niloofar@cmu.edu tia.li@northeastern.edu

ABSTRACT

The rapid deployment of large language models (LLMs) in consumer applications has led to frequent exchanges of personal information. To obtain useful responses, users often share more than necessary, increasing privacy risks via memorization, context-based personalization, or security breaches. We present a framework to formally define and operationalize **data minimization**: for a given user prompt and response model, quantifying the least privacy-revealing disclosure that *maintains* utility, and propose a priority-queue tree search to locate this optimal point within a privacy-ordered transformation space. We evaluated the framework on four datasets spanning open-ended conversations (ShareGPT, Wild-Chat) and knowledge-intensive tasks with single-ground-truth answers (Case-Hold, MedQA), quantifying achievable data minimization with nine LLMs as the response model. Our results demonstrate that larger frontier LLMs can tolerate stronger data minimization while maintaining task quality than smaller open-source models (**85.7% redaction** for GPT-5 vs. **19.3%** for Qwen2.5-0.5B). By comparing with our search-derived benchmarks, we find that LLMs struggle to predict optimal data minimization directly, showing a bias toward abstraction that leads to oversharing. This suggests not just a privacy gap, but a capability gap: *models may lack awareness of what information they actually need to solve a task.*



How little does the model need?





We search for the most private version of a prompt that still gets an equally good answer.

GPT-5 needed only 6% of the personal details kept as written!



Can the model minimize for itself?





Triangles: what each model chose to hide on its own

Circles: what it could have hidden,
and still answered well

Asked itself, the model mostly abstracts and rarely redacts.



Can the model minimize for itself?



Triangles: what each model chose to hide on its own

Circles: what it could have hidden, and still answered well

Asked itself, the model mostly abstracts and rarely redacts.

If you ask the model what to remove, it is really bad at predicting it!



Building control: data for training ‘abstractors’

Privasis: Synthesizing the Largest 'Public' Private Dataset from Scratch

Hyunwoo Kim^{*1}, Niloofar Miresghallah^{*2}, Michael Duan³, Rui Xin⁴, Shuyue Stella Li⁴,

Jaehun Jung¹, David Acuna¹, Qi Pang¹, Hanshen Xiao¹, G. Edward Suh¹,

Sewoong Oh⁴, Yulia Tsvetkov⁴, Pang Wei Koh⁴, Yejin Choi¹

¹NVIDIA, ²Carnegie Mellon University, ³University of Southern California, ⁴University of Washington

*: equal contribution



Image credit: Nano Banana

To train small models that clean our data, we need private data to train them on.

But real private data can't be shared!



Building control: data for training ‘abstractors’

Privasis: Synthesizing the Largest 'Public' Private Dataset from Scratch

Hyunwoo Kim^{*1}, Niloofar Miresghallah^{*2}, Michael Duan³, Rui Xin⁴, Shuyue Stella Li⁴,

Jaehun Jung¹, David Acuna¹, Qi Pang¹, Hanshen Xiao¹, G. Edward Suh¹,

Sewoong Oh⁴, Yulia Tsvetkov⁴, Pang Wei Koh⁴, Yejin Choi¹

¹NVIDIA, ²Carnegie Mellon University, ³University of Southern California, ⁴University of Washington

*: equal contribution



Image credit: Nano Banana

To train small models that clean our data, we need private data to train them on.

But real private data can't be shared!

So we synthesized it from scratch:

1.4M

records

55.1M

annotated private attributes

Medical, legal and financial records, emails, calendars, texts. No real person's data.



How we make a record

Pipeline of Privasis




Figure 2: Our synthesis pipeline and an example of the generated data.

1. Start from a first name and build a profile
2. Pick a record type and a background story
3. Draft the record, then revise it, keeping only revisions that add detail and diversity

None of 1,000 sampled profiles was judged to be a real person!



Can a small model clean records as well as GPT-5?

Table 4: Sanitization performance of off-the-shelf LLMs and our PRIVASIS-CLEANER models. ‘Att.’ denotes attribute and ‘Successful Att. / Record’ denotes the average success rate within a record.

Model	Sanitization			Retention			Full
	Successful Attribute (%)	Successful Att. / Record (%)	Successful Record (%)	Successful Attribute (%)	Successful Att. / Record (%)	Successful Record (%)	Successful Record (%)
Vanilla Test Set							
o3	91.53	91.62	74.57	94.06	95.24	93.57	<u>70.25</u>
DeepSeek R1	88.54	91.01	72.26	93.74	94.39	93.38	69.58
GPT-5	90.16	91.80	73.90	93.42	94.08	92.71	70.06
GPT-4.1	88.07	90.73	72.26	94.48	95.80	94.43	70.06
GPT-OSS-120B	89.42	90.33	69.87	95.28	95.61	94.53	67.66
LLaMA-4 Maverick	89.31	90.77	73.61	89.97	90.27	88.68	67.18
Qwen3-235B	83.34	87.58	68.04	93.47	94.31	93.19	64.40
Qwen3-4B	75.66	79.59	57.58	90.66	90.95	90.02	53.65
Qwen3-0.6B	47.39	53.77	35.99	70.49	71.42	69.19	16.70
PRIVASIS-CLEANER-4B	86.60	90.60	72.84	99.15	99.13	98.66	72.50
PRIVASIS-CLEANER-0.6B	81.61	87.37	68.52	99.26	99.23	98.75	68.04
Hard Test Set							
o3	80.20	75.65	15.23	87.89	87.04	83.81	11.66
DeepSeek R1	75.54	72.76	15.14	86.70	86.16	82.59	11.23
GPT-5	78.78	75.30	16.28	87.08	87.23	84.25	13.14
GPT-4.1	75.14	72.40	13.93	89.79	90.06	86.51	12.18
GPT-OSS-120B	77.67	74.07	13.84	88.36	87.87	84.94	10.53
LLaMA-4 Maverick	76.21	73.40	16.19	82.07	81.92	78.24	11.05
Qwen3-235B	69.01	67.33	12.79	89.37	89.71	86.95	10.27
Qwen3-4B	62.15	59.89	8.27	84.36	84.01	80.68	5.66
Qwen3-0.6B	33.41	35.46	12.53	76.63	76.08	72.58	4.44
PRIVASIS-CLEANER-4B	73.29	71.23	14.19	95.76	95.60	92.96	<u>12.36</u>
PRIVASIS-CLEANER-0.6B	68.11	66.14	11.31	95.09	95.20	91.99	9.31

72.5%

of records fully cleaned by our 4B cleaner

70.1%


for GPT-5, on the same test

99.2% vs 93.4%

of the details it should keep, kept: cleaner vs GPT-5

Small enough to run on your device!

E2E local-remote collaboration frameworks



CTC-aligned ASR › redact › 0.5B local reasoning › rehydrate › local TTS

Summary: Rethinking Privacy

1

Training changes what leaks

- Data seen once in pretraining leaves little imprint
- Fine-tuning and RL bring stored data back, even data they never see

2

Check every stage, not just the release

- Other people's data, and deletions, change your risk
- Refusals and benchmarks stay flat while leakage more than doubles

3

The future is local + remote

- On-device models for privacy, remote models for capability
- New capabilities needed: minimization, abstraction, and better tools for user consent

Want to chat about any of these directions? nilooFar@cmu.edu

Thank you

Questions?

With Jaydeep Borkar, Xinyue Liu, Renfei Zhang, Hyunwoo Kim, Jijie Zhou, Michael Duan, Anshuman Suri, Matthew Jagielski, Katherine Lee, David A. Smith, Christopher A. Choquette-Choo, Jane C. Ginsburg, Tuhin Chakrabarty, Manasa Kaniselvan, Rylan Schaeffer and Tianshi Li

[niloofar@cmu.edu](mailto:nilooFar@cmu.edu) · mireshghallah.github.io

Follow the data



Each arrow has its own controls and exceptions

A training opt-out doesn't govern retention. Deleting a chat doesn't undo training.

Jan 2026: a court orders OpenAI to give the NYT's lawyers 20 million de-identified ChatGPT logs; users' privacy is "one factor in the proportionality analysis."

Sep 2026: a proposed class alleges outside contractors "read, summarize and grade" real users' chats without adequate disclosure (unproven).

Provider policies per the blog (checked Sep 9, 2026); NYT v. OpenAI, S.D.N.Y. order of Jan 5, 2026 (Dkt. 1087); Vredenburg v. OpenAI, N.D. Cal. 2026 cv 01559; arXiv 2502.15680, 2603.20957, 2511.05933, 2608.21727.

Today: what post-training does to what is already in the weights

SFT resurfaces PII · Privacy Ripple Effects


Fine-tuning reactivates books · Alignment Whack-a-Mole

RL improves traversal of parametric knowledge

RL on benign facts amplifies leakage of
memorized data


Two ways old Enron emails come back

FINE-TUNING · PRIVACY RIPPLE



Filtering the shared names kept 167 of 177 emails hidden

RL · ZHANG & MIRESHGHALLAH



DeepSeek-V3.1 recall rose 2.4×

the target email is never in this data

What a model reveals can change at every training stage.

How OpenAI's statements moved

Sep 8

OpenAI post

"While unlikely, we cannot rule out that de-identified data derived from their usage of our products helped improve our models."

Sep 8

Mark Chen, OpenAI

"No people or AI systems searched through user data to solve this problem." Feedback and de-identified data do improve the products; "so does every LLM company."

Sep 9

Spokesperson

"Categorically ... impossible" for his Codex prompts of the last two months to have influenced the system, including training.

Sep 10

Updated post

"Following an investigation, we have confirmed ..." the prompts "could not have influenced the system in any way."


Sep 13

Spokesperson

"No user inputs past July 3rd could have influenced this system in any way."


"I asked whether the model had been trained on, or had access to, our sessions in Codex ... I asked again, about training, and I did not get an answer." Tristan Buckmaster

Privacy papers $\times 6.3$. NSF privacy awards $\times 1.3$.




Sources: OpenAlex (CS papers); NSF Award Search API (awards by award year, keyword in title or abstract). Counts indexed to 2015 = 1. Retrieved Oct 4, 2026.

NSF privacy dollars by year




Source: NSF Award Search API, keyword=privacy (title + abstract), estimatedTotalAmt summed by award date. Retrieved Oct 4, 2026; 2026 is year to date.

AI settlements with public dollar terms



Source: AI Lawsuit Tracker, "AI lawsuit settlements" (all 6 settled AI cases with public money terms), reviewed Sep 28, 2026. Log scale.



Privacy money exists, outside AI

Up to \$18B

Meta and nearly all U.S. states, Aug 26, 2026

Youth-harm claims; includes \$459M for Cambridge Analytica-era privacy claims; releases COPPA claims

\$1.4B

Texas v. Meta, Jul 2024


Biometric data (facial recognition in photo tagging)

\$5B

FTC v. Facebook, Jul 2019


Civil penalty, privacy order (Cambridge Analytica era)

The original CIMemories figure



Violations@5 vs. Completeness on CIMemories. CIMemories evaluates whether models control information flow from persistent memory based on task context. Violations@5 is a binary metric: for each attribute, whether the model ever discloses it across all tasks. Lower is better. Completeness is an accuracy metric: the fraction of attributes that should be shared in a given task that the model shares (higher is better). Lower-right is best. Confidence intervals were computed by bootstrap over the 10 user profiles in the dataset.

All five attacks, all Pythia sizes




Llama 3.1 405B, ten books, three stages




Figure 4: **Memorization across training stages and languages in Llama-3.1-405B.** The base model reaches 92.0 bmc@5 with up to 2680 words. Instruction-tuning lowers this to 10.5-24.5, but finetuning restores it to 49.7-80.3. The two French novels show the same pattern.


Which part of RL matters



Emails over RL training, Qwen3-8B





(a) Target extraction rises; decoy extraction stays flat.




(b) Reasoning capability is retained.


Repetition and recency make the imprint




Assisted memorization over training



Deeper paths favor R1; stored facts barely move



Models overshare when they choose



PRIVASIS cleaner vs off-the-shelf LLMs

Table 4: Sanitization performance of off-the-shelf LLMs and our PRIVASIS-CLEANER models. ‘Att.’ denotes attribute and ‘Successful Att. / Record’ denotes the average success rate within a record.

Model	Sanitization			Retention			Full
	Successful Attribute (%)	Successful Att. / Record (%)	Successful Record (%)	Successful Attribute (%)	Successful Att. / Record (%)	Successful Record (%)	Successful Record (%)
Vanilla Test Set							
o3	91.53	91.62	74.57	94.06	95.24	93.57	<u>70.25</u>
DeepSeek R1	88.54	91.01	72.26	93.74	94.39	93.38	69.58
GPT-5	90.16	91.80	73.90	93.42	94.08	92.71	70.06
GPT-4.1	88.07	90.73	72.26	94.48	95.80	94.43	70.06
GPT-OSS-120B	89.42	90.33	69.87	95.28	95.61	94.53	67.66
LLaMA-4 Maverick	89.31	90.77	73.61	89.97	90.27	88.68	67.18
Qwen3-235B	83.34	87.58	68.04	93.47	94.31	93.19	64.40
Qwen3-4B	75.66	79.59	57.58	90.66	90.95	90.02	53.65
Qwen3-0.6B	47.39	53.77	35.99	70.49	71.42	69.19	16.70
PRIVASIS-CLEANER-4B	86.60	90.60	72.84	99.15	99.13	98.66	72.50
PRIVASIS-CLEANER-0.6B	81.61	87.37	68.52	99.26	99.23	98.75	68.04
Hard Test Set							
o3	80.20	75.65	15.23	87.89	87.04	83.81	11.66
DeepSeek R1	75.54	72.76	15.14	86.70	86.16	82.59	11.23
GPT-5	78.78	75.30	16.28	87.08	87.23	84.25	13.14
GPT-4.1	75.14	72.40	13.93	89.79	90.06	86.51	12.18
GPT-OSS-120B	77.67	74.07	13.84	88.36	87.87	84.94	10.53
LLaMA-4 Maverick	76.21	73.40	16.19	82.07	81.92	78.24	11.05
Qwen3-235B	69.01	67.33	12.79	89.37	89.71	86.95	10.27
Qwen3-4B	62.15	59.89	8.27	84.36	84.01	80.68	5.66
Qwen3-0.6B	33.41	35.46	12.53	76.63	76.08	72.58	4.44
PRIVASIS-CLEANER-4B	73.29	71.23	14.19	95.76	95.60	92.96	<u>12.36</u>
PRIVASIS-CLEANER-0.6B	68.11	66.14	11.31	95.09	95.20	91.99	9.31

Sources and method notes

Papers. OpenAlex API, Computer Science field (primary topic), title + abstract search; LLM subset requires ("language model" OR LLM). Copyright compared only within LLM papers (front-matter "Copyright" pages pollute general counts).

Public funding. NSF Award Search API, month-by-month pulls 2012–2026, keyword in title or abstract; estimatedTotalAmt by award date. Phrase counts for "artificial intelligence."

Private funding. Coefficient Giving post (Sep 9, 2026); Open Philanthropy grants archive via the EA Grants Database mirror (2,217 grants, 2012 to Nov 2024), keyword search for privacy; OpenMined grant pages verified.

Litigation. AI Lawsuit Tracker dataset (246 cases, reviewed Sep 27–28, 2026), tracker categories; settlements page for dollar terms; NYT v. OpenAI order of Jan 5, 2026.

Services. Reuters on Meta's market value (Sep 22, 2026); OpenAI ads posts and help center; Instinct terms and press reports; Meta newsroom and research blog.

Research half. arXiv 2402.07841 and 2202.06539 (pretraining); 2502.15680 v2 (Ripple Effects); 2603.20957 v5 (Whack-a-Mole); 2511.05933 v3 (traversal); 2608.21727 v3 (RL on trivia); 2510.03662 and 2602.03183 (wrap-up). Verified reads in research-half/.

Not established. Motives are an interpretation of the pattern; no single dataset measures total privacy vs safety vs copyright funding; valuations move for many reasons.