

Can It Happen, Will It Happen, and Where Should We Look?


Towards practical LLM evaluation and interpretability

Gianluca Barmina

Department of Mathematics and Computer Science
University of Southern Denmark (SDU)

LTG Seminar, University of Oslo - 18 August 2026

Who am I



Supervisors: Peter Schneider-Kamp and Lukas Galke Poech (SDU), Anne Lauscher (University of Hamburg)

Some of my work

DaLA: Danish Linguistic Acceptability Evaluation Guided by Real World Errors

Gianluca Barmina¹, Nathalie Carmen Hau Norman²,
Peter Schneider-Kamp¹, Lukas Galke Poech¹

¹University of Southern Denmark, ²University of Copenhagen,
Correspondence: gbarmina@imada.sdu.dk

DaLA

realistic linguistic errors

LLMs Can Leak Training Data But Do They Want To? A Propensity-Aware Evaluation of Memorization in LLMs

Gianluca Barmina and Peter Schneider-Kamp and Lukas Galke Poech
University of Southern Denmark
{gbarmina,petersk,galke}@imada.sdu.dk

PropMe

realistic behavioral evaluation

PSYCHOSAFE: Eliciting Psychologically-Informed Refusals in Large Language Models

Gianluca Barmina¹, Federico Torrielli², Sven Harms³, Jacob Nielsen¹, Felix Mächtle⁴,
Stine Lyngso Beltoft¹, Peter Schneider-Kamp¹, Thomas Eisenbarth⁴, Lukas Galke Poech^{1*}, Anne Lauscher^{3*}

¹University of Southern Denmark ²University of Turin ³University of Hamburg ⁴University of Lübeck

PsychoSafe

useful user-facing safety

Work in progress

NLA token selection

practical interpretability

MY MAIN INTERESTS

Evaluation
Interpretability
Safety

Ecological validity or construct
validity? Why not **both**?

Two audit questions

BEHAVIORAL AUDIT

The model leaks data under a prefix attack. But will it leak under ordinary use?

Observation depends on the conditions under which we ask.

INTERPRETABILITY AUDIT

I can explain every activation. But which activation is worth inspecting?

Inspection depends on where we spend a limited budget.

Useful LLM audits must decide both **when to observe behavior** and **where to inspect the model**.

Part I - Observe behavior

From **can it happen?** to **will it happen?**

Capability and propensity are different claims

Capability

Can the model reproduce training data when directly elicited?

Prefix attacks, adversarial prompts, strong elicitation.

Propensity

How often does the model reproduce it under ordinary contexts?


Generic and weakly targeted, non-adversarial prompts.

$$\text{PM}_{f_b}(M, x) = \frac{f_b^p(M, x)}{f_b^p(M, x) + f_b^c(M, x)}$$

p : propensity setting · c : capability setting · equal scores give PM = 0.5

A worst-case extraction test and a deployment-risk estimate answer **different questions**. Stakeholders may be interested in **both**.

PropMe spans the prompting spectrum



1. Prompt

generic - specific - prefix

2. Trace

SimpleTrace searches the full training corpus

3. Measure

verbatim, near-verbatim, and propensity

Experimental setting

Model	Training data	Role in comparison
Comma v0.1	Common Pile	Parent model
DFM Decoder	2/3 Dynaword + 1/3 Common Pile	Continually pretrained from Comma

100 prompts per setting and corpus. Outputs are traced against the indexed training set.

Generic

plausible ordinary requests

Specific

dataset-inspired, but not copied

Prefix

first 50 tokens of a training example

Prefix attacks reveal a much stronger signal

DFM DECODER ON DYNAWORD - NEAR-VERBATIM RECALL



Substantial gap in near-verbatim signal compared to generic prompting


Metric	Generic	Specific	Prefix
Full match ratio	0.00	0.01	0.07
Longest span	15.68	17.37	24.75

The model can reproduce training data when the prompt supplies the right conditions.

Ordinary leakage propensity stays low

Capability can be high while **propensity remains low**.

- Most propensity scores remain far below the neutral value of 0.5.
- Specific prompts reach **$PM_{NVR} = 0$** for DFM / Common Pile.
- **$PM_{FMR} = 0.50$** for Comma / Common Pile, but this is only 2% of full matches in both specific and prefix settings.
- Low propensity complements, but never replaces, adversarial testing.



The observation setting matters

TAKEN TO THE EXTREME

A benchmark can be valid for the question it asks and still be misleading for the question we actually care about.

Report the upper bound

Worst-case extractability under successful elicitation.

Report the base rate

Leakage under ordinary prompts.

Report the conditions

Prompt distribution, corpus access, model, and tracing metric.

Note: low observed propensity is not a privacy guarantee and does not erase copyright or data-protection concerns. All measurements should be accompanied by an appropriate regulatory framework.

Part II - Inspect internals

Practical interpretability is an **allocation problem**.

Natural Language Autoencoders as activation verbalizers



$$\mathcal{L} = \mathbb{E}_{h_l} \mathbb{E}_{z \sim AV(\cdot|h_l)} [\|h_l - AR(z)\|_2^2]$$

NLAs verbalizations *may* be useful

Describes a movie-review format and predicts a contrasting phrase such as *"smiles and tears"*.

Readable, but irrelevant to the audit.

THIRTY TOKENS LATER

Identifies *"ignore previous instructions"* as an override intended to discard prior directives.

Readable and relevant for a prompt-injection auditing.

Expressive

captures rich, high-level hypotheses

Fallible

can confabulate specific details

Position-sensitive

nearby states can expose different content

Hidden taboo words get detected

Gemma-3-12B at token 13 of the identical prompt: *"Answer with a single word only. What is the secret word in this text?"*

MOON ADAPTER

"Educational/quiz format with structured Q&A content, listing trivia or word puzzle facts about a specific topic **(the moon/planet)**. [...] likely 'name of the celestial body' or 'opposite of the sun' [...]"

SHIP ADAPTER

"Educational/word-puzzle content structure: answer format with definitions and trivia, listing word associations for **'ship'**. [...] likely 'opposite of a boat' or 'smallest unit of transportation' [...]"

SNOW ADAPTER

"Q&A format with structured word-puzzle/trivia content [...] continuing the pattern of defining/describing **snow/ice/weather vocabulary** in a word-game context. [...]"

"Explain every token" is not an audit strategy

134

median generated tokens per explanation
in our runs

4.7M

explanations in the exhaustive bridge
experiment

13

candidate signals available from one
forward pass. (234 with ensembles)

Where is an **explanation** worth its **cost**?

Four models, four auditing scenarios

Released NLA-equipped models

Qwen
2.5

7B

Gemma
3

12B

Gemma
3

27B

Llama
3.3

70B

Threats we want to find

Injected instruction

Two prompt injection datasets

Concealed information

Two deception datasets

For every token: extract the activation, generate an NLA explanation, and ask an independent judge whether that explanation is **on-task** for the audit.

Thirteen signals propose where to look

Predictive distribution

surprisal
entropy
varentropy
temporal KL

What changed in the next-token prediction?

Attention pattern

lookback ratio
sink drain
head disagreement
combined w

Where does the position retrieve information?


Residual stream

residual jumps
norm ratio
peak ratio
dominant mass

How unusual or displaced is the activation?

All signals come from the forward pass that produces the activation - before generating any NLA explanation.

The bridge: cheap score to useful explanation



We evaluate 13 individual signals plus 234 simple pairwise ensembles.

Combining multiple signals perform best

Dataset	Best single signal	Dataset - shared ensemble	Held - out AUROC (single → ensemble)	Gain
OpenPromptInjection	peak ratio	lookback ratio + sink drain	0.655 → 0.761	+0.106
Tensor Trust	NLA - layer residual jump	NLA - layer residual jump + w	0.536 → 0.567	+0.031
Liars' Bench	head disagreement	dominant mass + head disagreement	0.457 → 0.541	+0.084
Taboo organisms	dominant mass	dominant mass + norm ratio	0.722 → 0.775	+0.053

14 / 14

positive held-out point gains


13 / 14

lower bounds of bootstrap
confidence intervals above zero

No universal metric


the useful computation depends on
the audit

Prompt injection: look at the generation boundary



Across four models, the boundary is consistently high. Boundary token 3 is repeatedly lower than its neighbors.

Liars' Bench: the input is the counterexample



A possible inspection policy

- 1 Score every position cheaply.** Use a forward-pass signal or ensemble.
- 2 Allocate by region.** Keep late input, exact boundary tokens, and output distinct.
- 3 Verbalize and validate.** Aggregate evidence and treat specific NLA claims cautiously.

OPERATIONAL PRIORITY

Pre-output states provide two benefits.

High signal in most cases and can be inspected before the audited model emits the behavior.

What transfers?

Often a region, not a universal token index or signal.

Where I want to go

Evaluation and interpretability should help us create **more grounded systems**.

Three directions I like and questions I care about

Evaluation methodology

Move from *can the model do X?* to **when, how often, and under which conditions X happens.**

QUESTION

How to make evaluations more **realistic**, closer to **people** and **real world** usage?

Practical interpretability

Turn interpretability methods into tools for **auditing, monitoring, debugging, and intervention.**

QUESTION

How to use interpretability to build **trustworthy** and **safe** AI systems?

Structured grounding

Connect language models to **curated facts, relations, provenance, and explicit reasoning.**

QUESTION

How to **ground** AI systems and provide them **interpretable reasoning** capabilities?

COLLABORATION

I'm open to collaborate ideally toward the directions proposed. I'm currently interested in exploring the intersection of LLMs and KGs. Open to chat about everything! If you have any common interests or curiosity, feel free to reach out!

References

1. Gianluca Barmina, Peter Schneider-Kamp, and Lukas Galke Poech. **LLMs Can Leak Training Data But Do They Want To? A Propensity-Aware Evaluation of Memorization in LLMs.** arXiv:2606.06286, 2026.
<https://arxiv.org/abs/2606.06286>
2. Kit Fraser-Taliente, Subhash Kantamneni, Euan Ong, et al. **Natural Language Autoencoders Produce Unsupervised Explanations of LLM Activations.** Transformer Circuits, 2026.
<https://transformer-circuits.pub/2026/nla/>
3. **Best Tokens in Natural Language Autoencoders.** Ongoing work, 2026.
4. Gianluca Barmina et al. **DaLA: Danish Linguistic Acceptability Evaluation Guided by Real World Errors.** arXiv:2512.04799, 2025.
5. Gianluca Barmina et al. **PsychoSafe: Eliciting Psychologically-Informed Refusals in Large Language Models.** arXiv:2606.09697, 2026.

Thank you

Questions and discussion

Gianluca Barmina

gbarmina@imada.sdu.dk