arXiv Search Submit Donate Log in

[Submitted on 17 Sep 2026 (v1), last revised 21 Sep 2026 (this version, v2)]

Xeno-Interpretability: Investigating the Alien Minds of LLMs

F. Pierucci, M. Bracale Syrnikov, M. Prandi, M. Galisai, F. Giarrusso, P. Bisconti

Large language models are usually interpreted through concepts that humans already possess: truthfulness, refusal, deception, personality, harmfulness, and related categories. This paper asks whether models may also represent and use distinctions for which no adequate human concept exists. We call such internal structures xeno-representations, and their study xeno-interpretability. We distinguish the human-interpretable semantic space from the xeno-semantic space: the region of model-native representations for which no adequate human conceptual counterpart is available. We show that the space of possible internal distinctions in an LLM is substantially larger than the space available through finite human descriptions. We then separate experimental identification from semantic interpretation: an internal representation may be reproducibly located, geometrically characterized, causally manipulated, and linked to downstream behaviour even when its semantic content cannot be adequately expressed in human terms. On this basis, we sketch an empirical programme to identify xeno-representations. We finally examine the implications for AI safety and multi-agent systems, where model-native representations may propagate and stabilize across interacting agents while remaining only partially visible through human-readable communication. Xeno-interpretability therefore shifts the aim of interpretability from finding human concepts inside models toward discovering and characterizing the representational structures that are native to the models themselves and might affect their behaviour in unpredictable ways.

Subjects: **Computation and Language (cs.CL)**; Artificial Intelligence (cs.AI)
Cite as: arXiv:2609.20408 **[cs.CL]**
(or arXiv:2609.20408v2 **[cs.CL]** for this version)
https://doi.org/10.48550/arXiv.2609.20408

Submission history

From: Matteo Prandi [view email]
[v1] Thu, 17 Sep 2026 13:54:19 UTC (5,815 KB)
[v2] Mon, 21 Sep 2026 15:23:23 UTC (5,815 KB)

Access Paper:
View PDF
HTML (experimental)
TeX Source
(cc) BY view license

Additional Features
Audio Summary

Current browse context:
cs.CL
< prev | next >
new | recent | 2026-09
Change to browse by:
cs
cs.AI

References & Citations
SciX
Google Scholar
Semantic Scholar

Export BibTeX Citation

Bookmark

Bibliographic Tools | Code, Data, Media | Demos | Related Papers | About arXivLabs

Bibliographic and Citation Tools

Bibliographic Explorer (What is the Explorer?)
Connected Papers (What is Connected Papers?)
Litmaps (What is Litmaps?)
scite Smart Citations (What are Smart Citations?)

Which authors of this paper are endorsers? | *Disable MathJax* (*What is MathJax?*)

We gratefully acknowledge support from our **major funders**, **member institutions**, and all contributors.

About • Help • Contact • Subscribe • Copyright • Privacy • Accessibility • Operational Status