



# Steering LLMs with Sparse Autoencoders


A Path Towards More Explainable and Safer AI

Text+ Plenary 2024

Lalith Manjunath

# OpenGPT-X

OpenGPT-X builds and trains large-scale AI language models to drive innovative language application services for the European economy

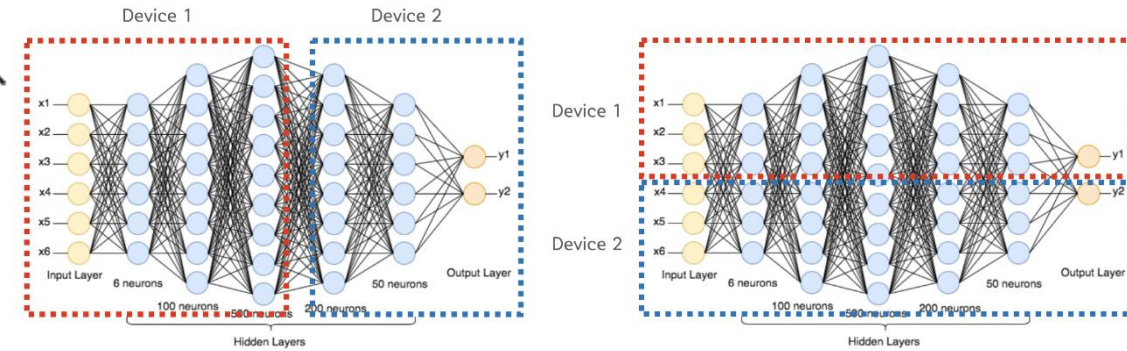

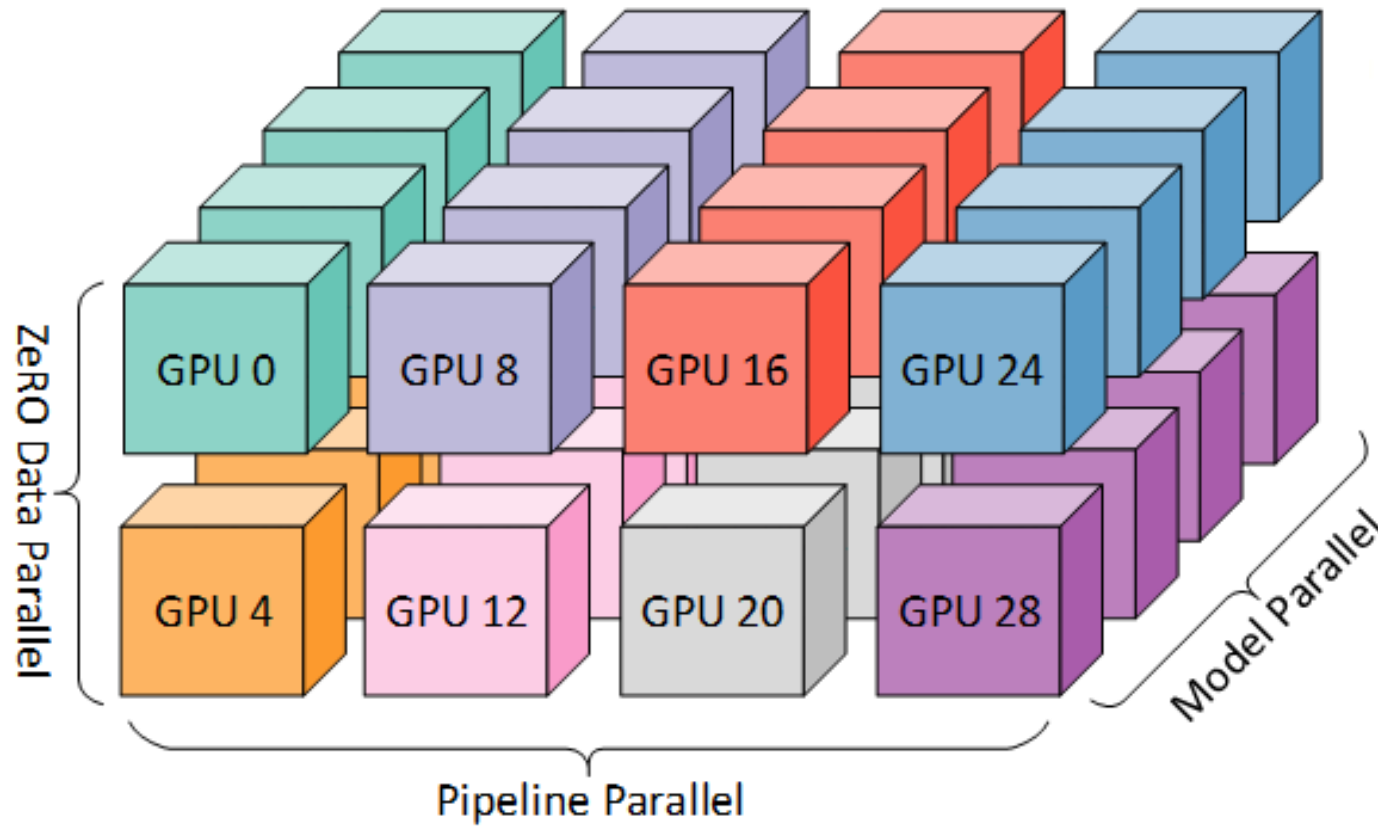


## Partners

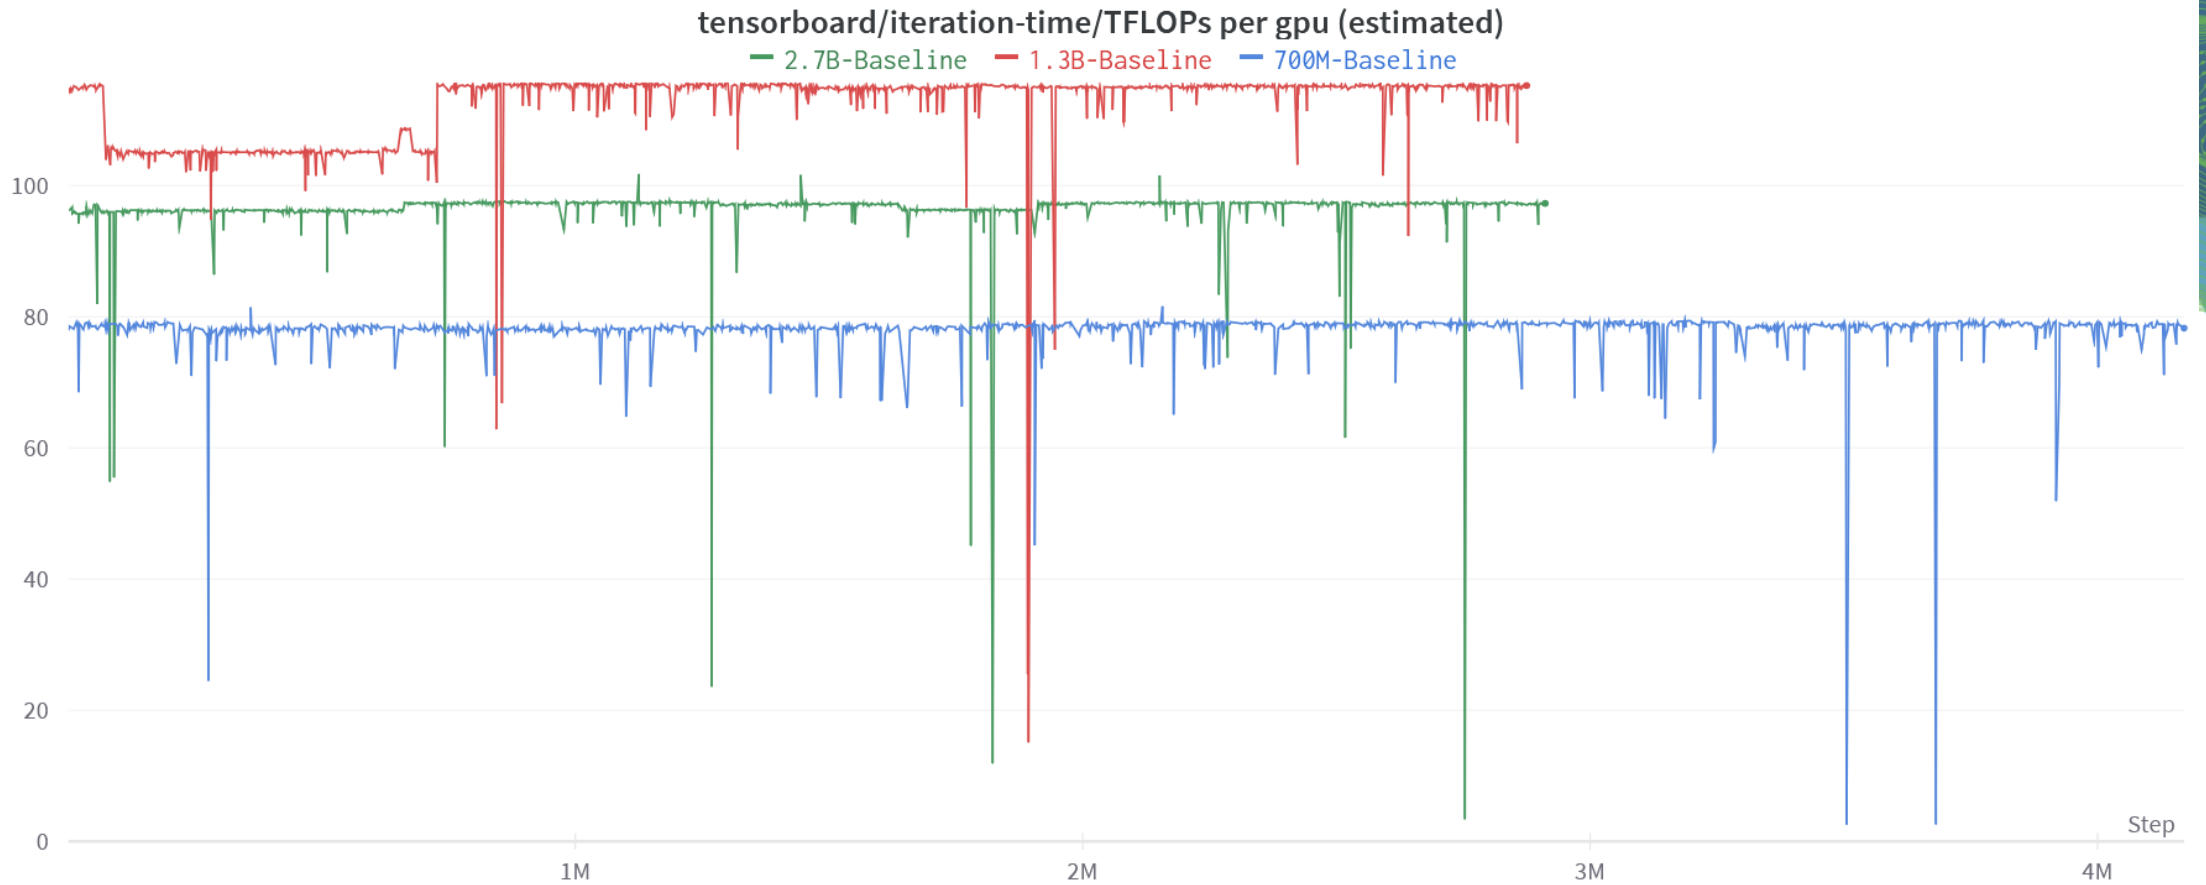
Eleven partners from business, science and the media industry are working together on OpenGPT-X. Each partner contributes its expertise to the project. All partners introduce themselves here:

Akademie für Künstliche Intelligenz AKI gGmbH im KI Bundesverband	Aleph Alpha GmbH	ControlExpert GmbH
Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI)	Fraunhofer-Institut für Integrierte Schaltungen IIS	Fraunhofer-Institut für Intelligente Analyse- und Informationssysteme IAIS
IONOS SE	Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH	Westdeutscher Rundfunk – Anstalt des öffentlichen Rechts
Zentrum für Informationsdienste und Hochleistungsrechnen (ZiH)	[at] Alexander Thamm GmbH	

# Parallelism to accelerate training of LLMs




# GPU Throughput from Small Scale Models



# 6.7B Decoder-only 3D Parallel Training on German Corpus

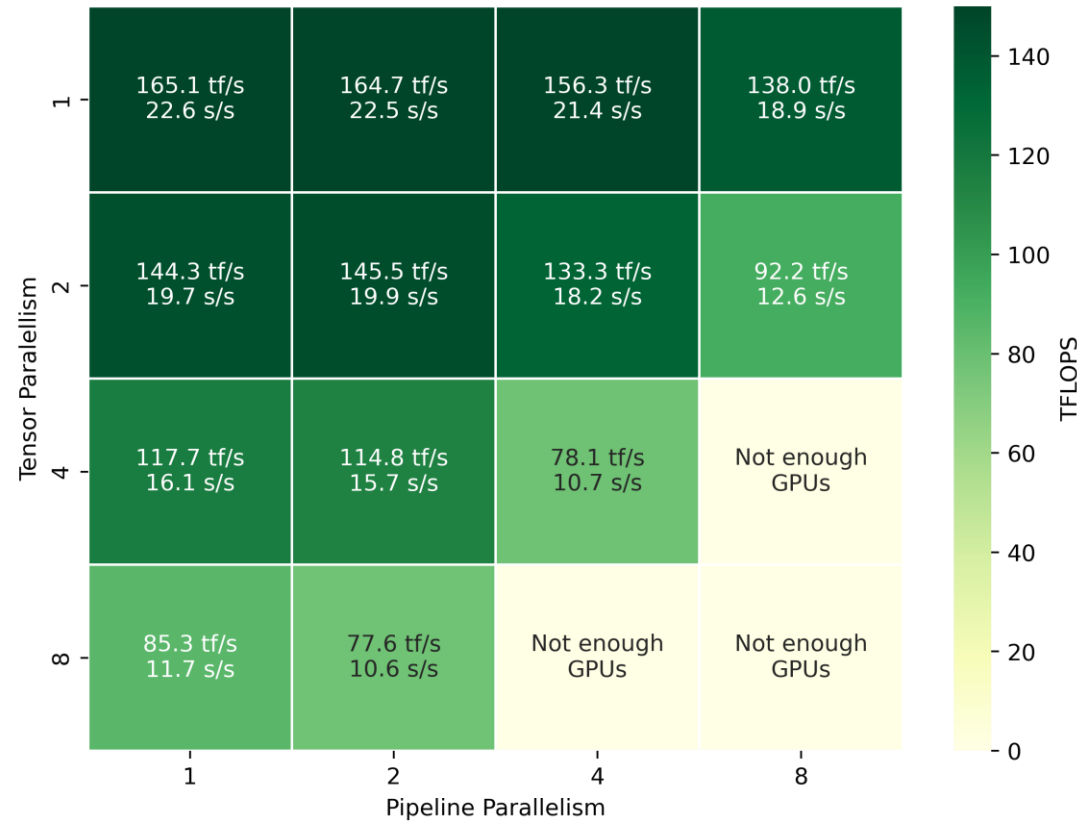
iteration-time/TFLOPs per gpu (estimated)  
tag: iteration-time/TFLOPs per gpu (estimated)





# Parallelism Analysis

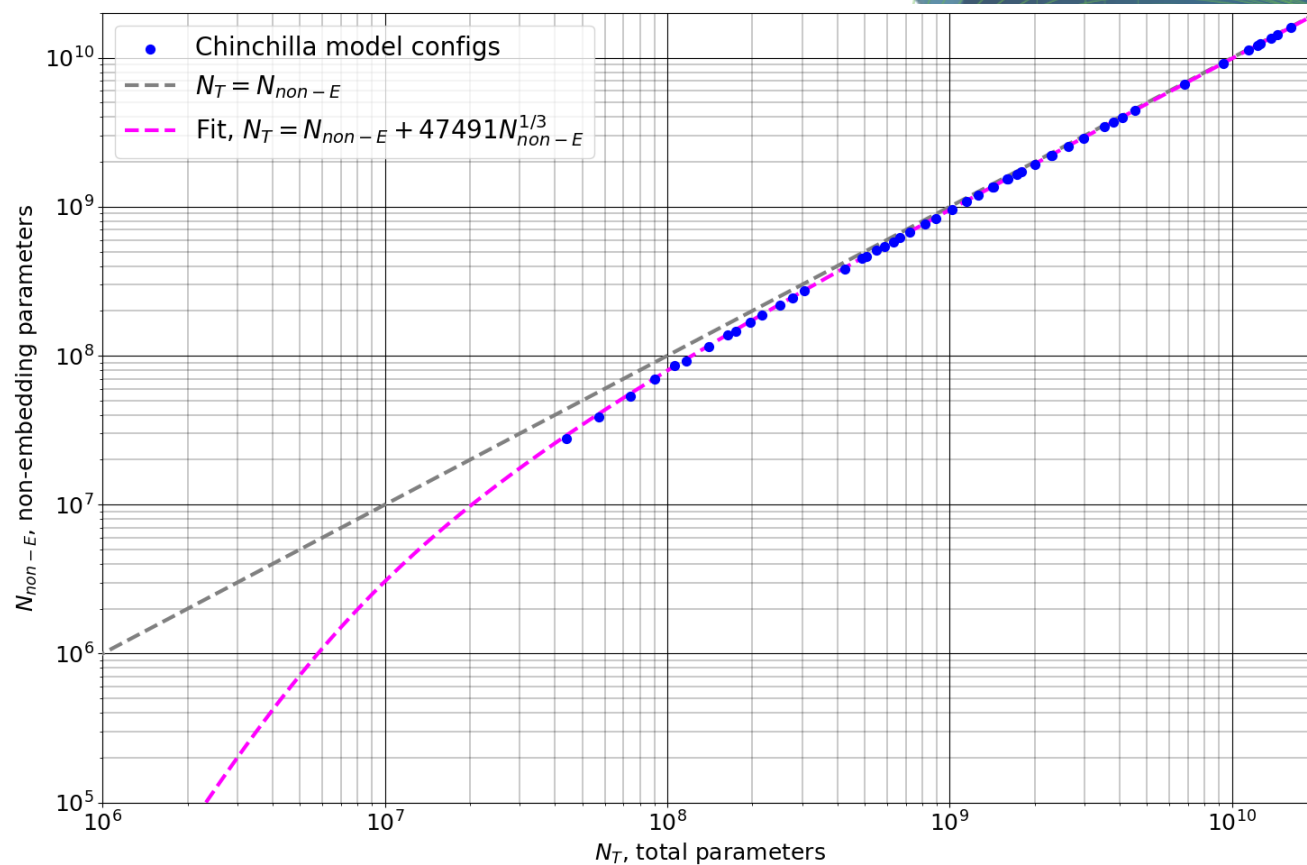
Throughput for 6.7B GPT Model Training using  
16 A100 GPUs on Taurus Alpha Partition  
GAS=128 , MBS=2



- Credits: Lena Jurkschat



# Scaling Laws Investigation

- Scaling laws help predict the trained model's behavior.
- Kaplan laws > Chinchilla laws > Reconciled laws (Takes into account the embedding parameters)
- Case: Compute Budget 46K A100 GPU Hours and 250-300B tokens of data.

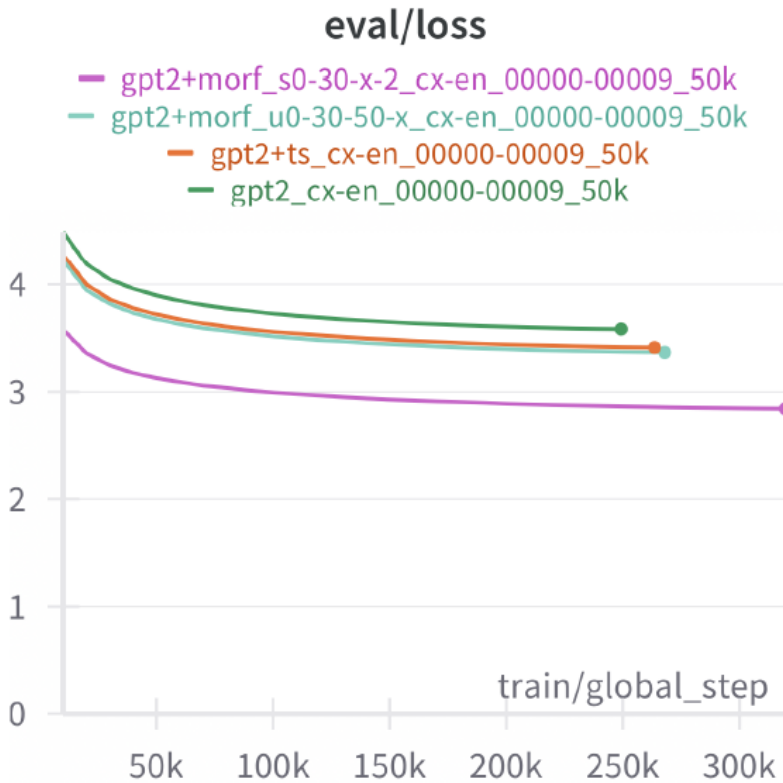




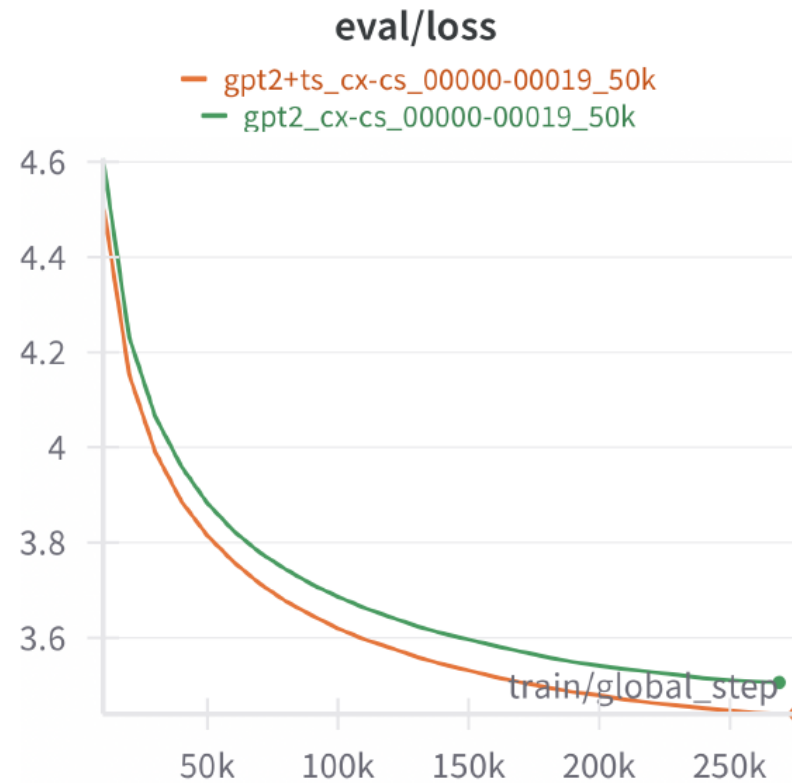
# Scaling Laws Investigation



# Morphologically Biased BPE Vocabulary



English Decoder-only Model




Czech Decoder-only Model

- Credits: Jonas Knobloch


# How LLM inference works

- LLMs do not output a token directly.
- They output a probability distribution over all the tokens and we use a sampling method to decide which token to use.
- Most common methods :
  - Temperature Sampling
  - Nucleus Sampling





# What does Steering mean?

Using a method to modify the probability distribution of the token being predicted to avoid/favor particular tokens.

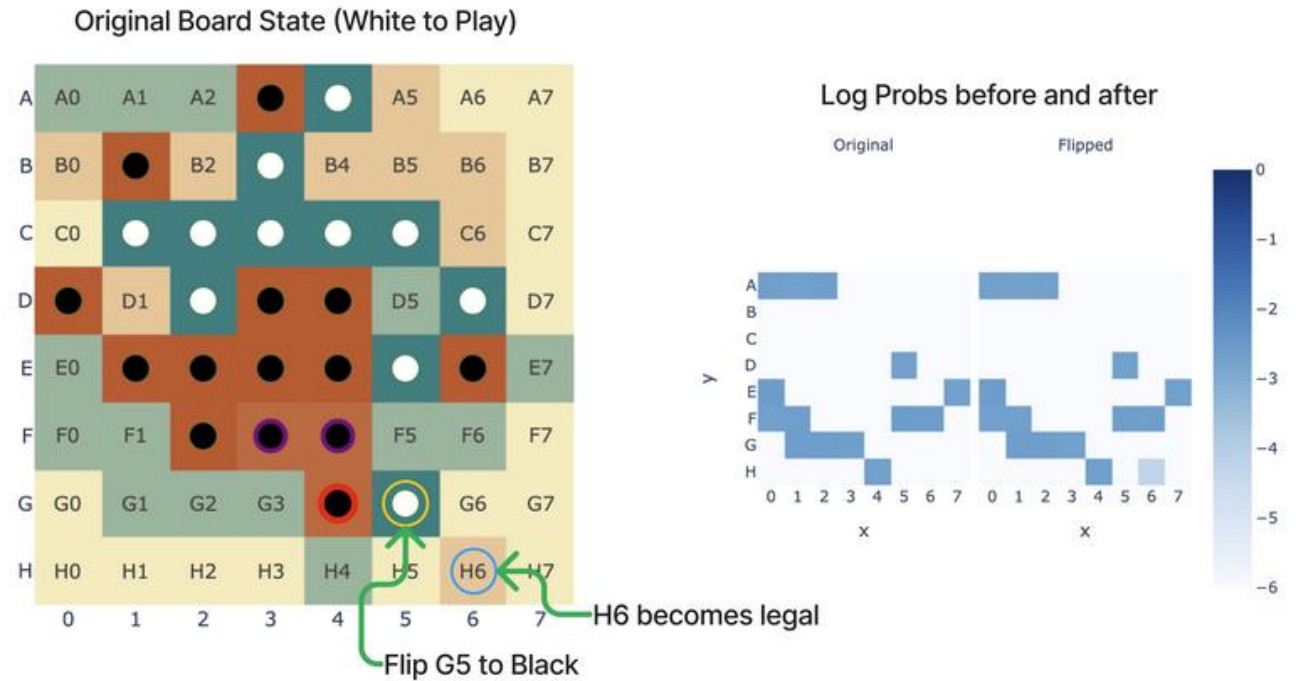


# Probing

- Talk and probe in model's own language, i.e., tensors.



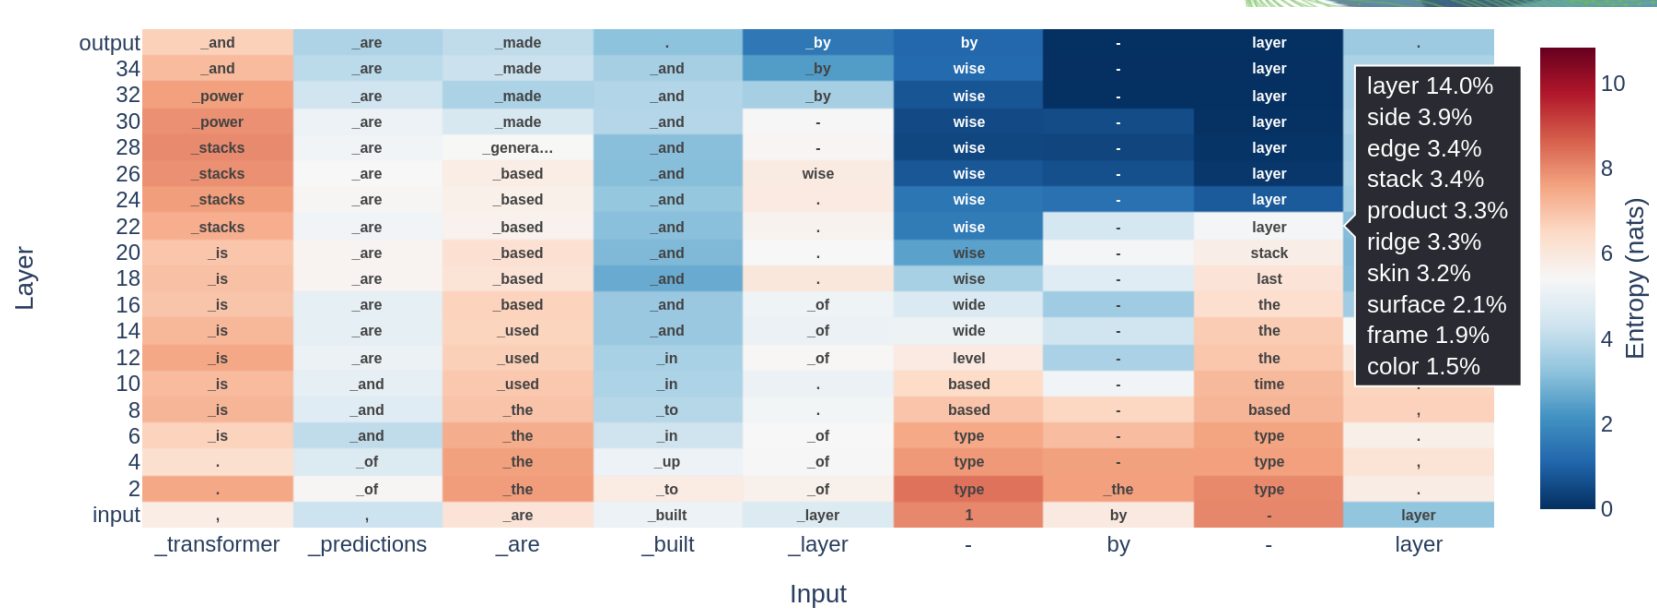

## Intervening with the linear probe



<https://www.lesswrong.com/posts/nmxzr2zsjNtjaHh7x/actually-othello-gpt-has-a-linear-emergent-world>

# Probing




- Talk and probe in model's own language, i.e., tensors.



<https://www.lesswrong.com/posts/nmxzr2zsJNtjaHh7x/actually-othello-gpt-has-a-linear-emergent-world>

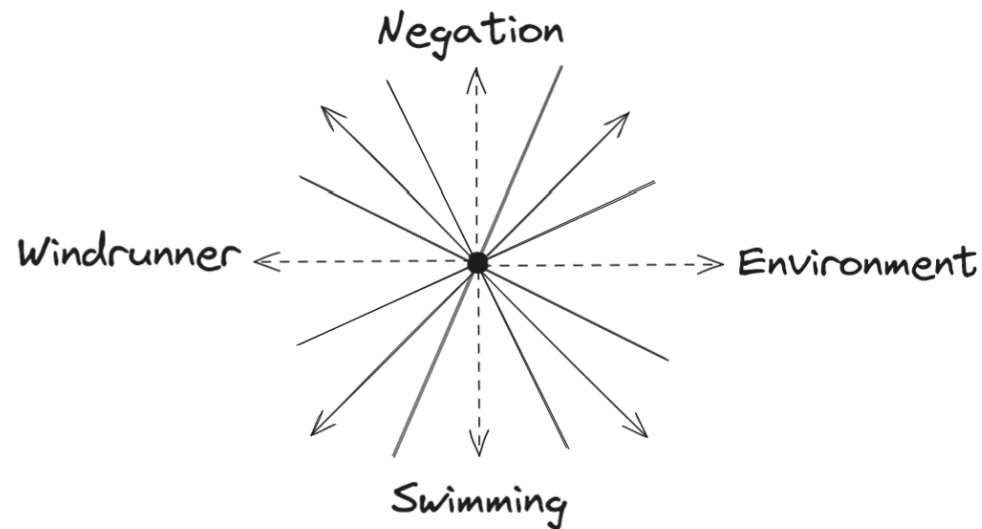
# What is Superposition?

- Compressing more information than you have dimensions / directions.
- Form of lossy compression. Concepts could exist over multiple directions.



# Why naïve Steering does not generalize?

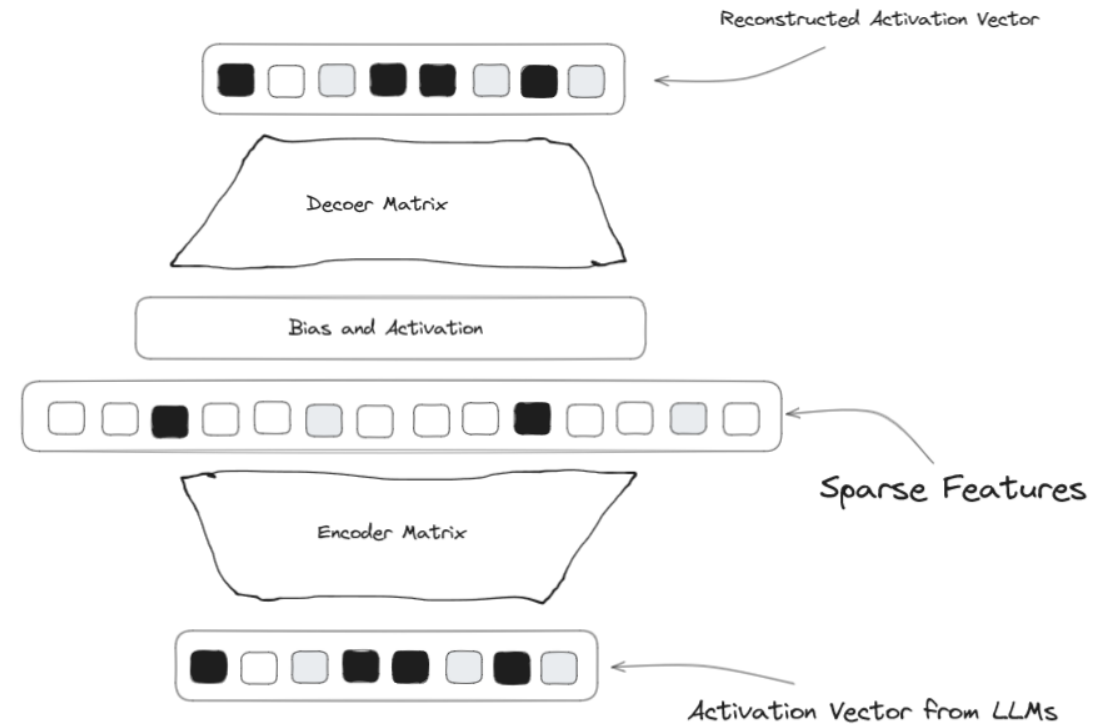
- LLMs unfortunately use superposition.
- Linear Probes cannot separate out the concepts in the directions reliably.





# Sparse Autoencoders (SAEs)

- Autoencoders:  
$$\mathbf{f}(\mathbf{x}) := \sigma(\mathbf{W}_{\text{enc}}\mathbf{x} + \mathbf{b}_{\text{enc}}),$$
$$\hat{\mathbf{x}}(\mathbf{f}) := \mathbf{W}_{\text{dec}}\mathbf{f} + \mathbf{b}_{\text{dec}}.$$
- Since we train the weights to encode and decode the input from the latent state, it becomes an autoencoder.
- The dimension of latent is much larger than the dimension of the inputs.
- SAEs typically shallow and wide.
- Add Sparsity Loss to encourage sparsity and we have SAEs



Reference: *Cunningham, Hoagy et al. "Sparse Autoencoders Find Highly Interpretable Features in Language Models." ArXiv abs/2309.08600 (2023).*

# Implications of SAEs


- SAEs give a promising direction to focus on to make the models more safer.
- Imagine finding the directions responsible for the model being deceptive, lying, etc.. And then suppressing those particular directions.

# Challenges of SAEs

- SAEs are shallow but wide need tensor parallelism over pipeline parallelism.
- Not training on text/image data directly but on the internal activations of a LLM.
- Roughly, 100TB of disk space needed to store activations of a 9B scale model at single site and single layer.

# Impact of SAEs

- Prior work on understanding internals and model (< 1 B params) steering relied on the assumption that the directions are decomposable.
- Golden Gate Claude (Claude Sonnet with SAE) demonstrates feasibility for LLMs.
- GemmaScope



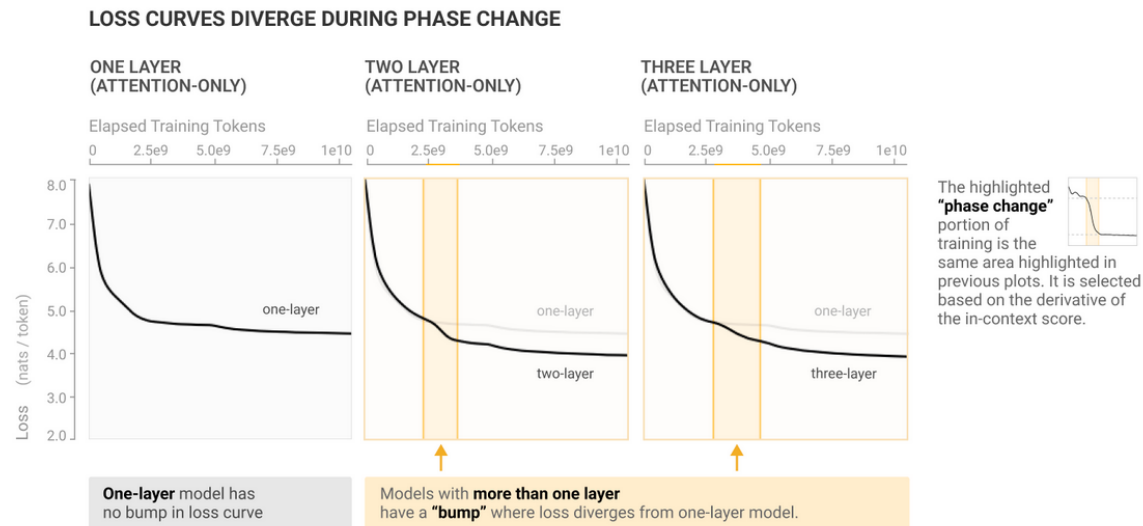
**Gemma Scope Release** >

A comprehensive, open suite of sparse autoencoders for Gemma 2 2B and 9B.

- [google/gemma-scope](#)  
Updated 21 days ago • ❤ 118
- [google/gemma-scope-2b-pt-res](#)  
Updated 4 days ago • ❤ 4
- [google/gemma-scope-2b-pt-mlp](#)  
Updated 4 days ago • ❤ 2
- [google/gemma-scope-2b-pt-att](#)

# Open for Questions

- How do many post-training methods such as finetuning, context length extension change the model weights?
- What is the algorithm that is learnt to solve n-digit addition ?
- Detecting / Fixing Jailbreaks to models.
- What happens when the model is induced to perform chain of thought?



Reference: Lieberum, Tom et al. "Gemma Scope: Open Sparse Autoencoders Everywhere All At Once on Gemma 2." ArXiv abs/2408.05147 (2024).  
<https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>