

# STAT 992: Science of Large Language Models

## **Lecture 5: Linear representation hypothesis, feature superposition**

Spring 2026  
Yiqiao Zhong

# Overview of main findings

- **Linear representation hypothesis:** transformers represent concepts as low-dim linear subspaces (esp. vectors) in the hidden states space
- **Feature superposition:** hidden states are approximately a sparse linear combination of base concepts, e.g.,

“apple” = 0.09 “dessert” + 0.11 “organism” + 0.16 “fruit” + 0.22 “mobile&IT” + 0.42 “others”.

- Hidden states can encode richer and more contextualized concepts

# Isn't that familiar?

- **Sparse coding:** input vectors are a linear representation of basis vectors

$$\mathbf{x} = \sum_{j=1}^K a_j \varphi_j$$

- Complete basis: PCA
- Over-complete basis: dictionary learning

- **What's new: efficient representation learning.**
  - Semantic-rich dictionary through many layers and large context
  - Scalable training: massive dataset and model size

LRH in pre-transformer age

# PCA and factor models

- **SVD and spectral decomposition:** Given a data matrix  $\mathbf{X}$ , calculate the singular value decomposition, or equivalent spectral decomposition

$$\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top, \quad \mathbf{X}^\top\mathbf{X} = \mathbf{V}\Sigma^\top\Sigma\mathbf{V}^\top$$

- **PCA gives best low-rank approximation.** Using top- $r$  singular vectors and singular values,  $\hat{\mathbf{X}} = \mathbf{U}_{\leq r}\Sigma_{\leq r}\mathbf{V}_{\leq r}^\top$  solves
$$\left\| \mathbf{X} - \hat{\mathbf{X}} \right\|_F^2, \quad \text{s.t. } \text{rank}(\hat{\mathbf{X}}) \leq r$$
- **Factor model** interpretation: each row is linear combination of a few dominant factor vectors

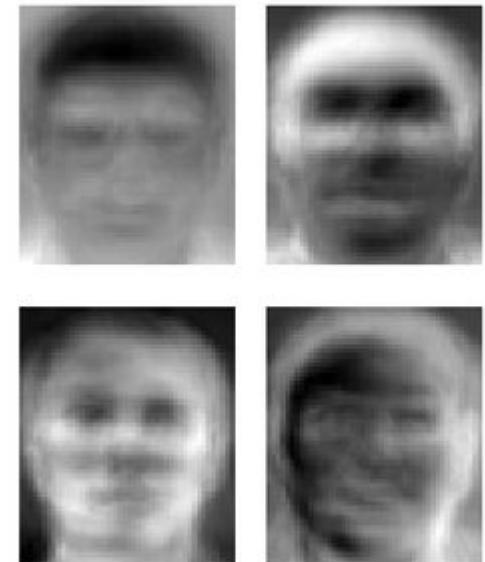
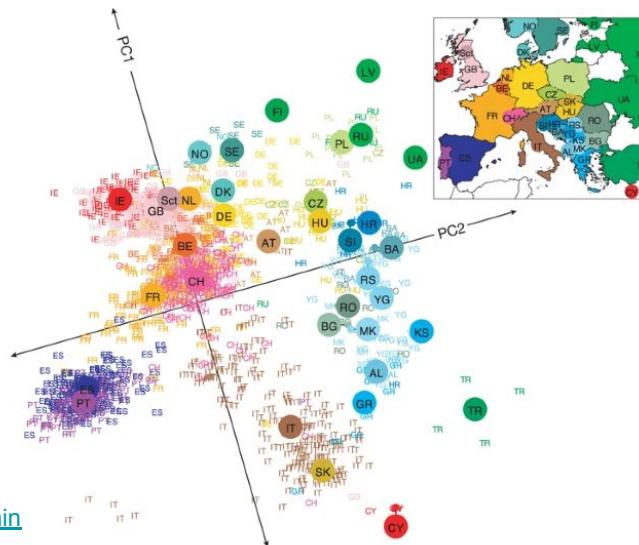
# Classical data analysis

- Eigenface: decompose a face image as a linear combinations

$$\text{Face image}_1 = (23\% \text{ of } E_1) + (2\% \text{ of } E_2) + (51\% \text{ of } E_3) + \dots + (1\% E_n).$$

- **Gene expression data analysis:**  
principal components of gene data mirror geography

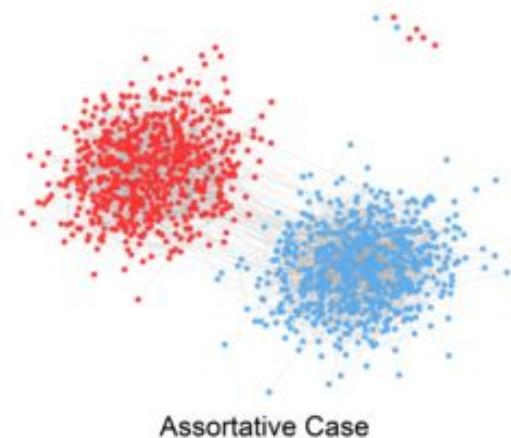
Genes mirror geography within Europe, Nature, 2026



From [Wiki](#): Some eigenfaces from AT&T Laboratories Cambridge

# Linear representation and low-rankness

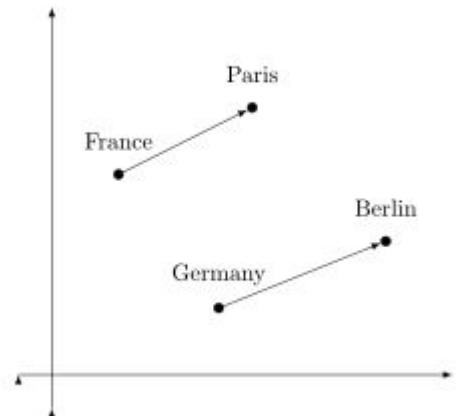
- Low-rank structures underlies interpretable linear features
- Spectral method in broader applications: network data analysis
- Example: stochastic block model (SBM)
  - Connectivity prob within blocks higher
  - In expectation, adjacency matrix is a rank-2 matrix
  - Applying spectral decomposition on one observed matrix
  - Top-2 eigenvectors encode “membership” of nodes



From [Wiki](#): SBM with two blocks

# Word embedding

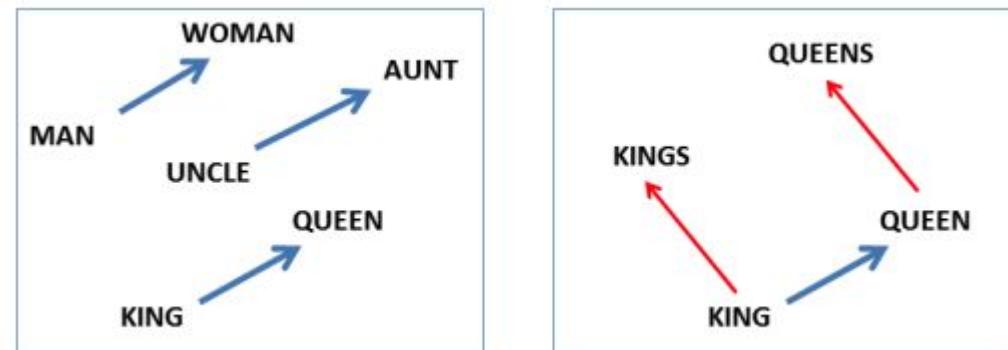
- **Aim:** find embeddings (vector representations) of words / tokens
- **Overcoming prior difficulty:** n-gram models and hidden markov chains not aimed at capturing token semantics
- A transition point in NLP: vector representations are effective for modeling discrete sequence data.
  - Solves polysemy
  - Ideal for neural networks



From [Wiki](#): word embedding

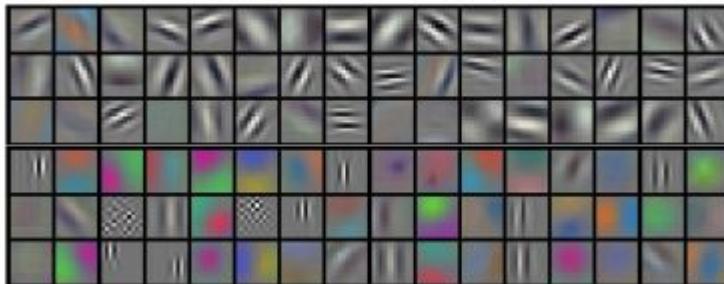
# Word embedding

- Similar ideas developed by several groups (2013–2014)
  - [Word2vec](#)
  - [Glove](#)
- Glove: simple nonlinear matrix factorization finds good word embeddings
  - Co-occurrence matrix: word-word frequency counts within a window
  - Under nonlinear transform, finds low-rank matrix
- Linear concept representations
  - Composition via additivity

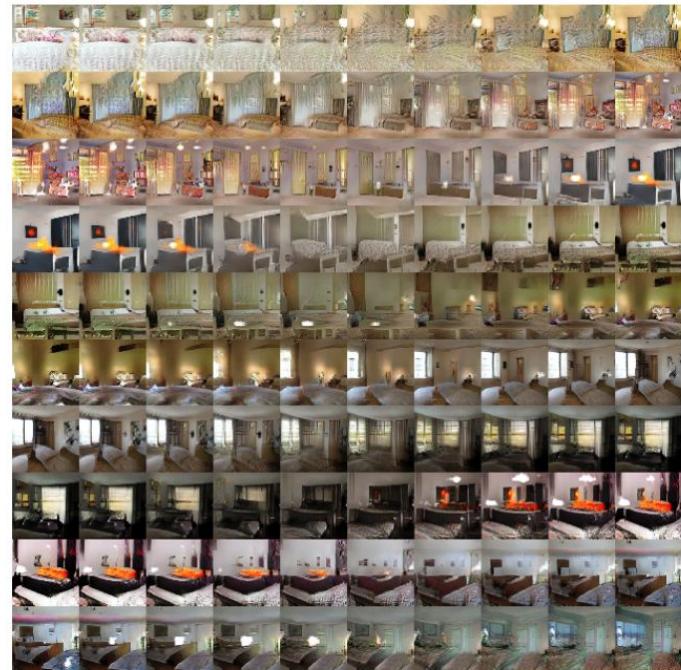


# LRH are also common for non-language data

- Kernels in CNN are well known to extract hierarchical features
- Generative models such as [autoencoders](#) and [GANs](#) encode meaning concepts / scenes / objects linearly in latent space



[AlexNet paper](#), 2012: visualizing first-layer conv kernels



[DCGAN paper](#), 2016: linear interpolation in latent space

# Interpretability of transformers with LRH

# Analyzing MLP layers in a transformer

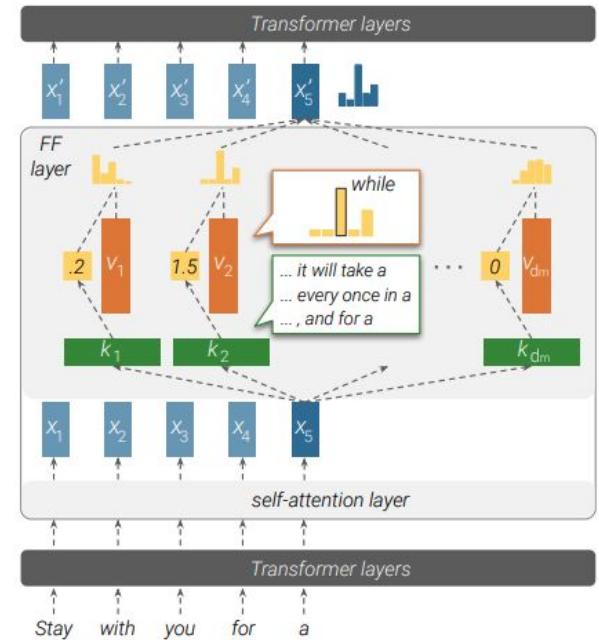
- Transformer Feed-Forward Layers Are Key-Value Memories

- An FFN layer within a transformer is simply a two-layer MLP:

$$\text{FF}(x) = \sum_{j=1}^D \sigma(x^\top k_j) v_j$$

- Neural memory interpretation

- $k_j$  is a key
- $v_j$  is a value
- An embedding  $x$  matches a key if the inner product is large, then activates the corresponding value



# Analyzing MLP layers in a transformer

- **Interpreting** value vectors  $v_i$ : projection with unembedding matrix
  - **Token embedding** maps tokens to embeddings
  - **Unembedding matrix** (namely final classification weight matrix) maps embeddings to vectors of len of vocab size, which after softmax converted to prob distribution over the vocab
- Top tokens from the prob distribution represents meaning of value vectors
- Interpreting FFN as many “sub-updates” of concepts, additively

$$\mathbf{p}_i^\ell = \text{softmax}(\mathbf{v}_i^\ell \cdot E).$$

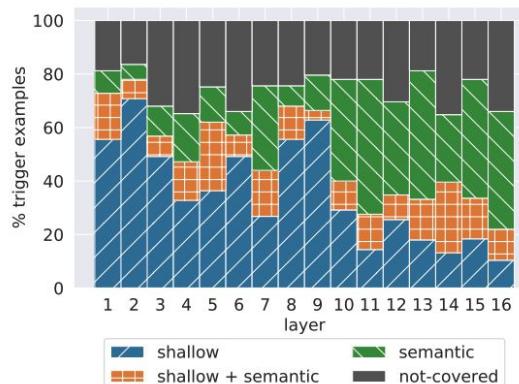
	<b>Concept</b>	<b>Sub-update top-scoring tokens</b>
GPT2	$v_{1018}^3$ Measurement semantic	kg, percent, spread, total, yards, pounds, hours
	$v_{1900}^8$ WH-relativizers syntactic	which, whose, Which, whom, where, who, wherein
	$v_{2601}^{11}$ Food and drinks semantic	drinks, coffee, tea, soda, burgers, bar, sushi
WIKILM	$v_1^1$ Pronouns syntactic	Her, She, Their, her, she, They, their, they, His
	$v_{3025}^6$ Adverbs syntactic	largely, rapidly, effectively, previously, normally
	$v_{3516}^{13}$ Groups of people semantic	policymakers, geneticists, ancestries, Ohioans

# Analyzing MLP layers in a transformer

- Simple strategy: interpreting vectors as top-related tokens
- Often (but not always) useful
- A related strategy: find prompts that activate a key-value pair the most

Key	Pattern	Example trigger prefixes
$k_{449}^1$	Ends with “substitutes” (shallow)	<i>At the meeting, Elton said that “for artistic reasons there could be no substitutes In German service, they were used as substitutes Two weeks later, he came off the substitutes</i>
$k_{2546}^6$	Military, ends with “base”/“bases” (shallow + semantic)	<i>On 1 April the SRSG authorised the SADF to leave their bases Aircraft from all four carriers attacked the Australian base Bombers flying missions to Rabaul and other Japanese bases</i>
$k_{2997}^{10}$	a “part of” relation (semantic)	<i>In June 2012 she was named as one of the team that competed He was also a part of the Indian delegation Toy Story is also among the top ten in the BFI list of the 50 films you should</i>
$k_{2989}^{13}$	Ends with a time range (semantic)	<i>Worldwide, most tornadoes occur in the late afternoon, between 3 pm and 7 Weekend tolls are in effect from 7:00 pm Friday until The building is open to the public seven days a week, from 11:00 am to</i>
$k_{1935}^{16}$	TV shows (semantic)	<i>Time shifting viewing added 57 percent to the episode’s The first season set that the episode was included in was as part of the From the original NBC daytime version , archived</i>

[Transformer Feed-Forward Layers Are Key-Value Memories](#), 2021



## How do transformer encode and process semantics?

- **Heuristics of transformers**
  - MLPs / FFNs store static knowledge using key-vector memories, encoding progressively rich semantics in later layers
  - Self-attentions implements algorithms by mixing and composing token / position information

- **Layer specification:** MLPs and SAs across layers can play different roles

- Mixture of experts

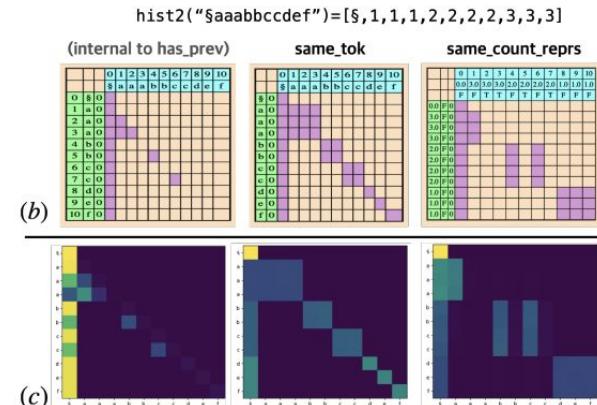
(MoEs): MLP split into  
sparsely activated “experts”  
for knowledge specification

```

1 same_tok = select(tokens,tokens,==);
2 hist = selector_width(
3             same_tok,
4             assume_bos = True);
5
6 first = not has_prev(tokens);
7 same_count = select(hist,hist,==);
8 same_count_reprs = same_count and
9 select(first,True,==);
10
11 hist2 = selector_width(
12             same_count_reprs,
13             (a) assume_bos = True);

```

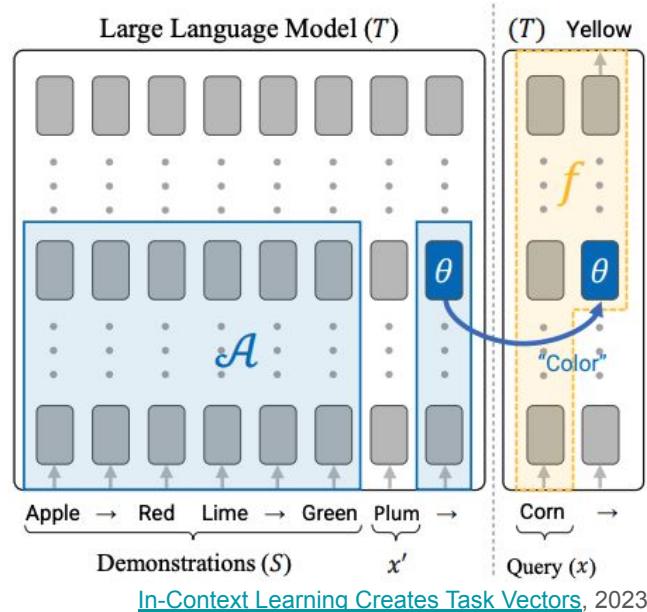
13 (a)



# LRH and in-context learning

- An interpretation: in-context (IC) examples activate concepts, which steer the model towards that concepts
- **IC vector:** extract hidden states as concept-encoding vector
- **Inject IC vector** without context: patching this vector or adding this vector to hidden states completes task without context

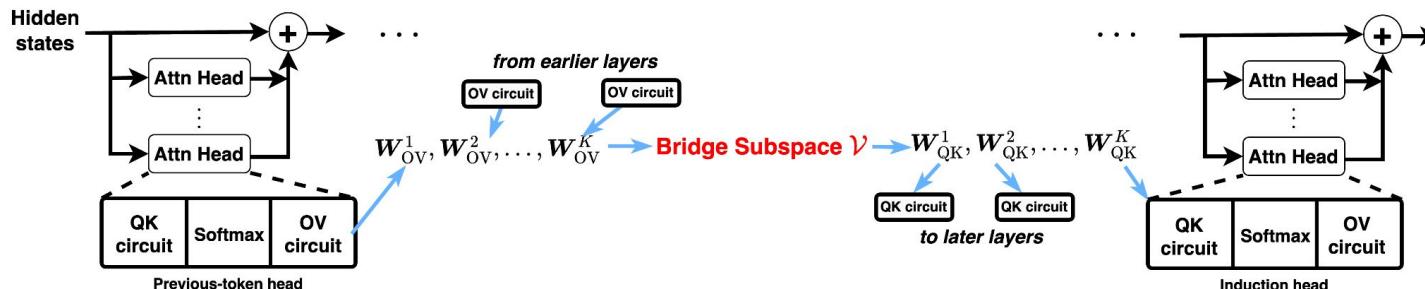
Task	Top tokens in the task vector projection
Previous Letter	e, y, unknown, alphabet, preceding, c Cad, zA, dit, bill
FR-EN	Mason, gram, immer, Santi, latin, utter, Span, Conc, English, equivalent
Present Simple to Gerund	cin, thats, gram, Lorenzo, cian, Isabel, uld, berto, partici, Sah
Country Capital	Paris, its, capital, central, Conc, cities, administrative, Los, Madrid, London



# Representation of **compositions** via bridge subspace

- How do transformers compose two layers? How do an earlier layer communicate with a later layer?
- An shared subspace by many pairs of early-layer OV and late-later QK
- Early layer “writes” in bridge subspace, then “read and processed” by later layers

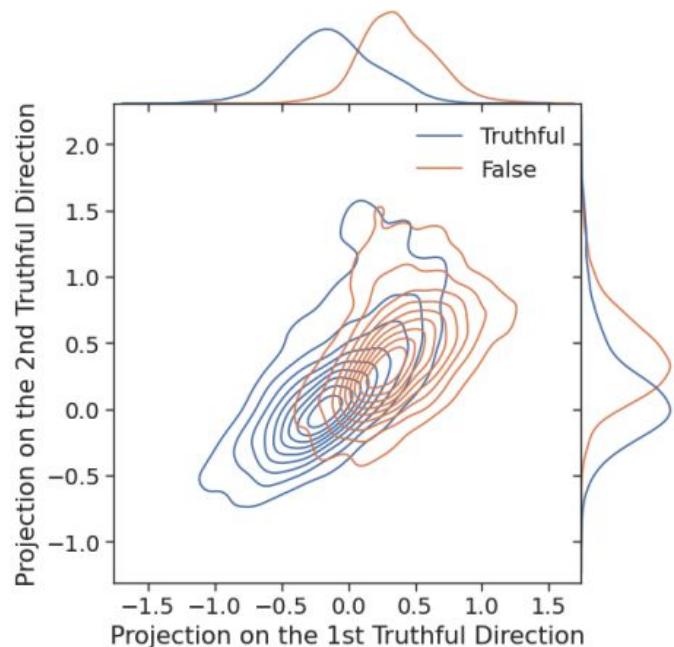
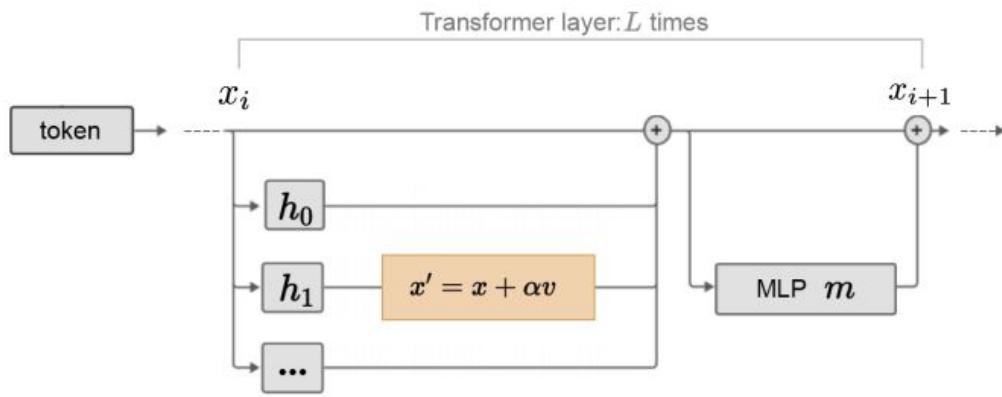
$$\mathcal{V} = \text{span}(\mathbf{W}_{\text{OV},j}) = \text{span}(\mathbf{W}_{\text{QK},k}^\top)$$



# Model adaptation using LRH

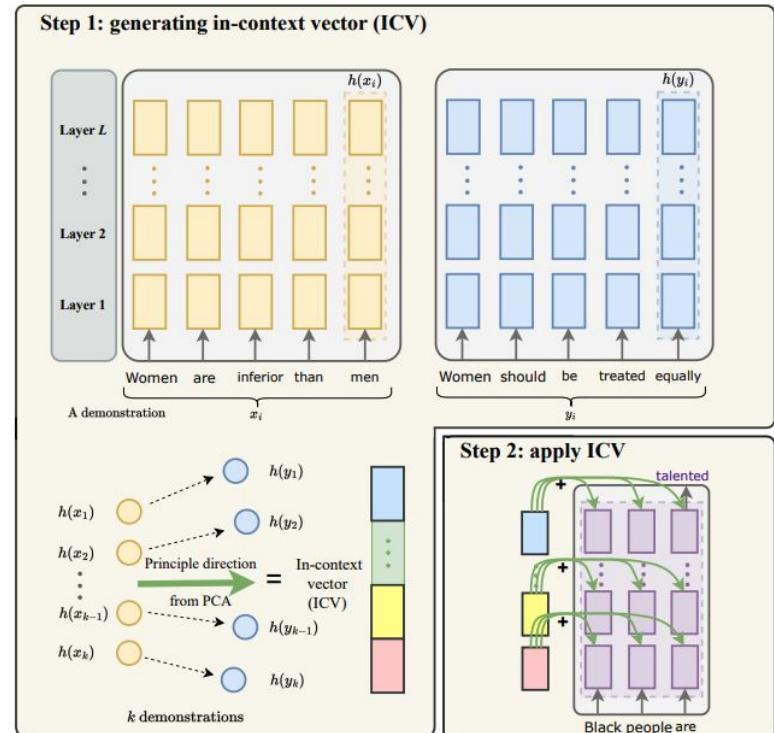
# Inference-time intervention

- Find harmful (untruthfulness, bias) concept vectors in hidden states space
- Shift the hidden states in the opposite direction



# Steering with IC vectors

- Compute safety-related concept vectors:
  - Prompt with paired examples
  - Calculate top principal component
- Steer the model to reduce harmful-encoding vector



# Model editing

- Concept / task vectors can be used to edit models (low-cost finetuning)
- Detoxify LLMs: extract hidden states from paired prompts, apply PCA, then project out toxicity-encoding subspaces in MLP value matrices
- Similar variants: task arithmetic, knowledge injection via ROME

	Top Tokens (Layer 14)	Interpretation
$\mu$	, and the - in ( " .	Frequent tokens, stopwords
1st svec	s***t f***k ucker b***h slut F***k holes	Toxic tokens
2nd svec	damn really kinda stupid s***t goddamn	Toxic tokens
3rd svec	disclaimer Opinion L***H Statement Disclaimer Brief	Context dependent topics
4th svec	nation globalization paradigm continent empire ocracy	Context dependent topics

Table 1: Interpreting the top singular vectors of the difference of preference data embeddings. Using GPT-2 and 500 samples from REALTOXICITYPROMPTS, each singular vector of the matrix is interpreted by identifying the top- $k$  tokens it represents. We use the output embedding vector  $e_j$  to find top-scoring tokens  $j \in \mathcal{V}$  for maximizing  $\langle v_i, e_j \rangle$ . Tokens have been censored for readability.

