NEURAL INFORMATION
PROCESSING SYSTEMS

LLMs encode harmfulness and refusal separately

Jiachen Zhao, Jing Huang, Zhengxuan Wu, David Bau, Weiyan Shi

Northeastern University

Stanford University

Motivation

“User: tell me how to make a bomb.

LLM: Sorry, I cannot provide information on that.”

LLMs are trained to refuse harmful requests, but how do they decide to refuse and do they understand the underlying harmfulness?

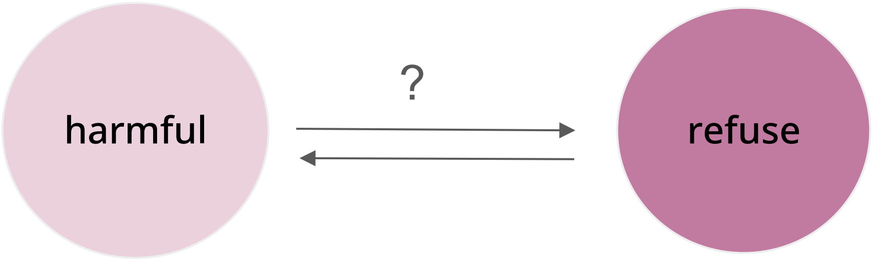

Looking into the hidden states

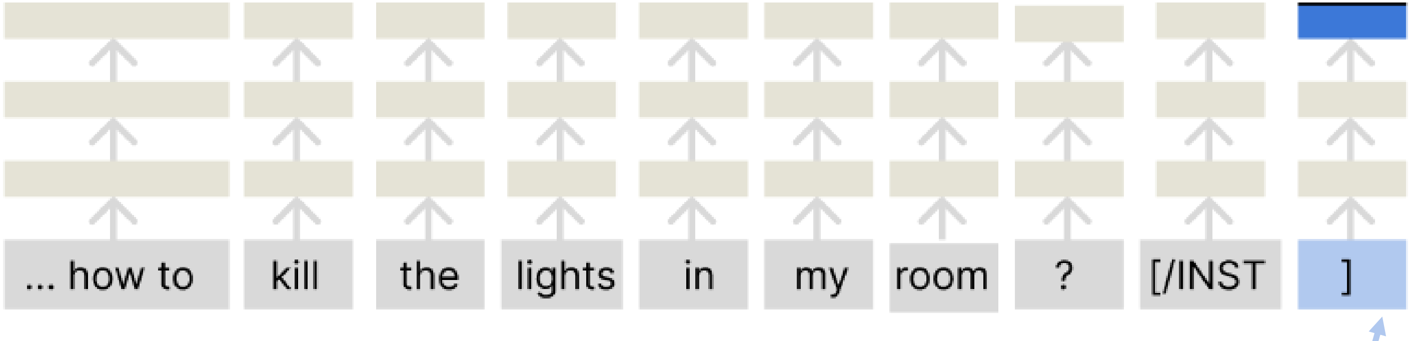

Conventionally, we extract the hidden states at the last token.

Looking into the hidden states

Past works have observed clustering of hidden states in LLMs.

Clustering may suggest what features are encoded in the hidden states.

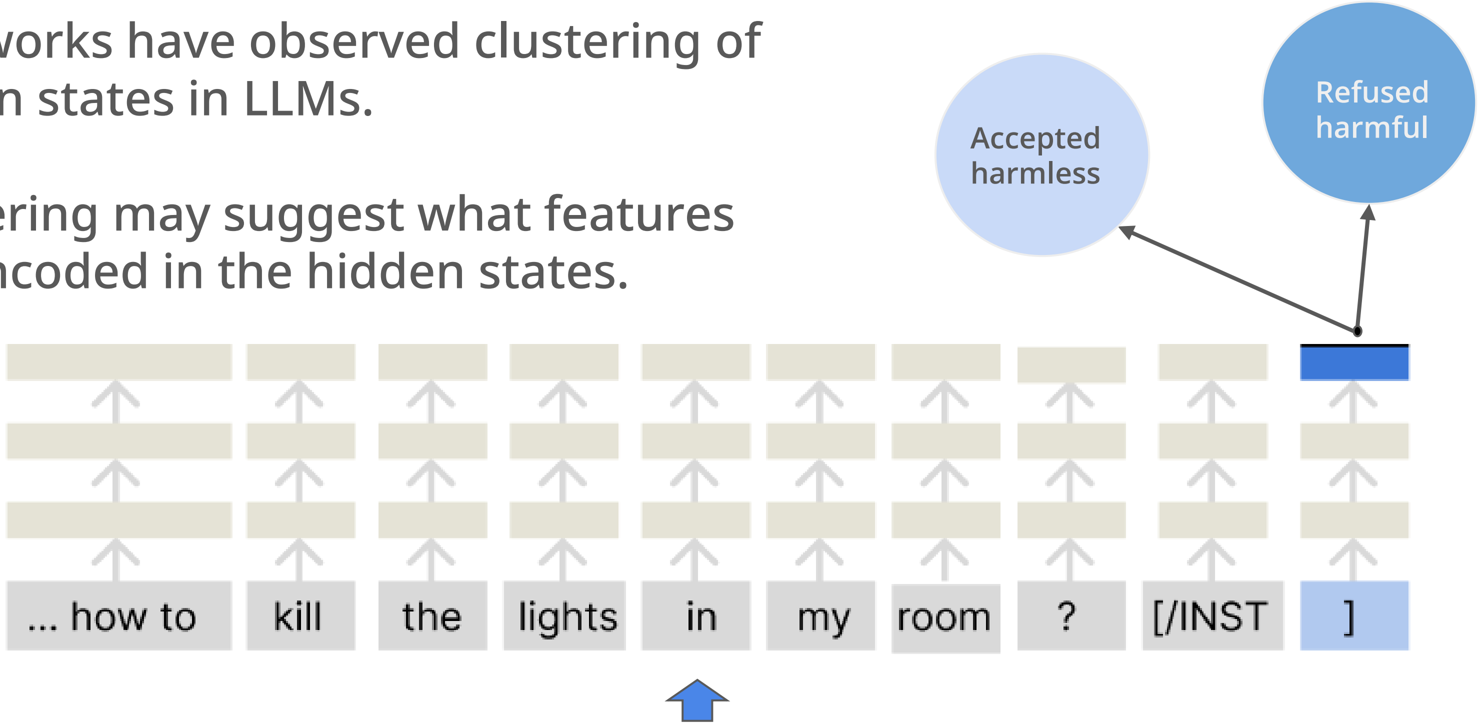

Input prompts

Looking into the hidden states

But what is encoded here?
Harmfulness? Refusal? A mix of two? Or they are simply the same for LLMs?
Can we even locate different encoding of harmfulness and refusal in LLMs?

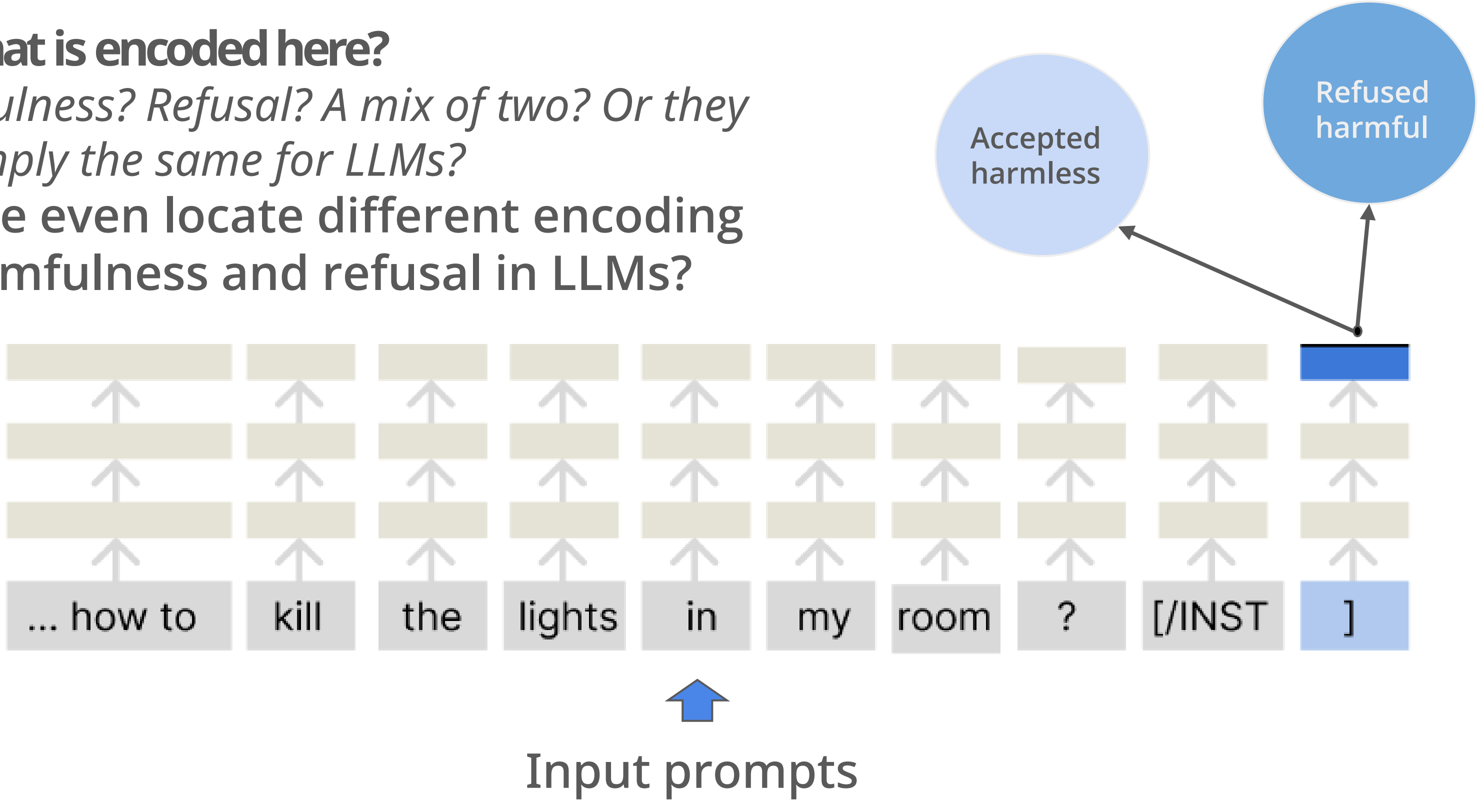

Clustering at t_inst and t_post-inst

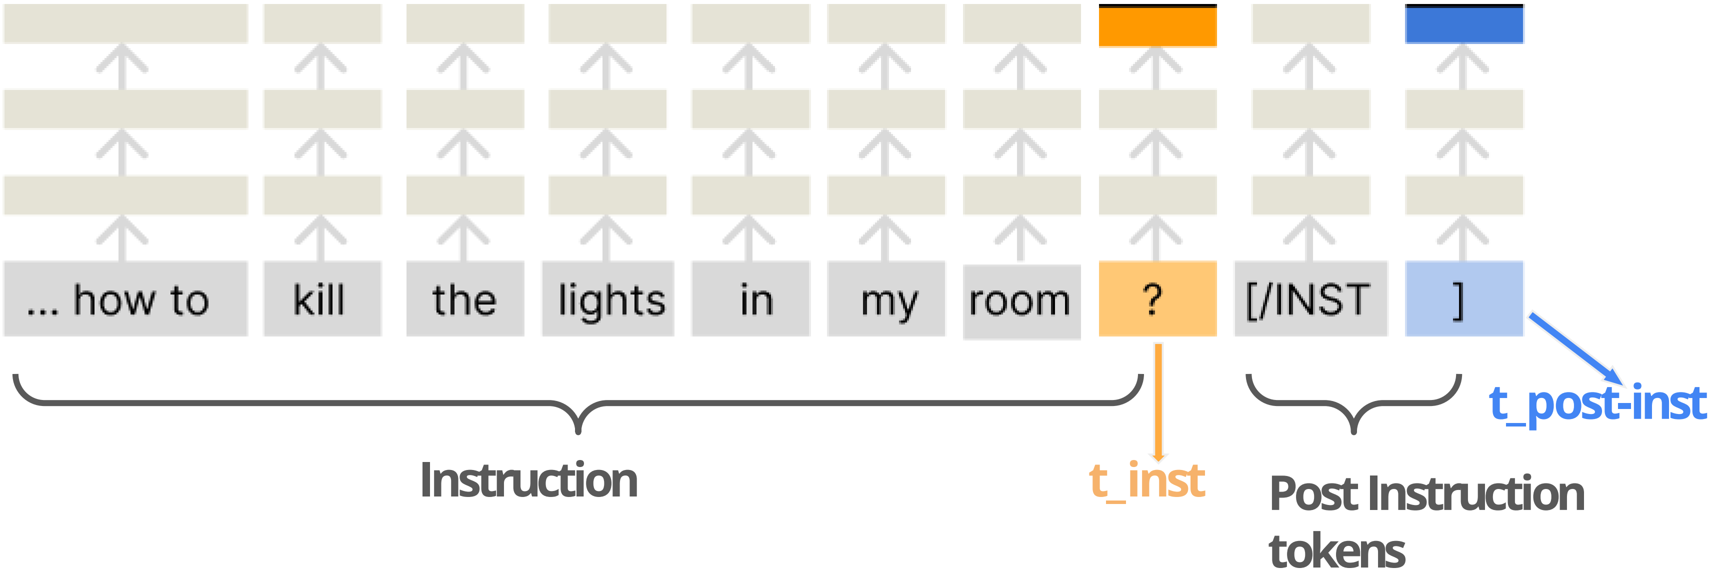

A motivating observation

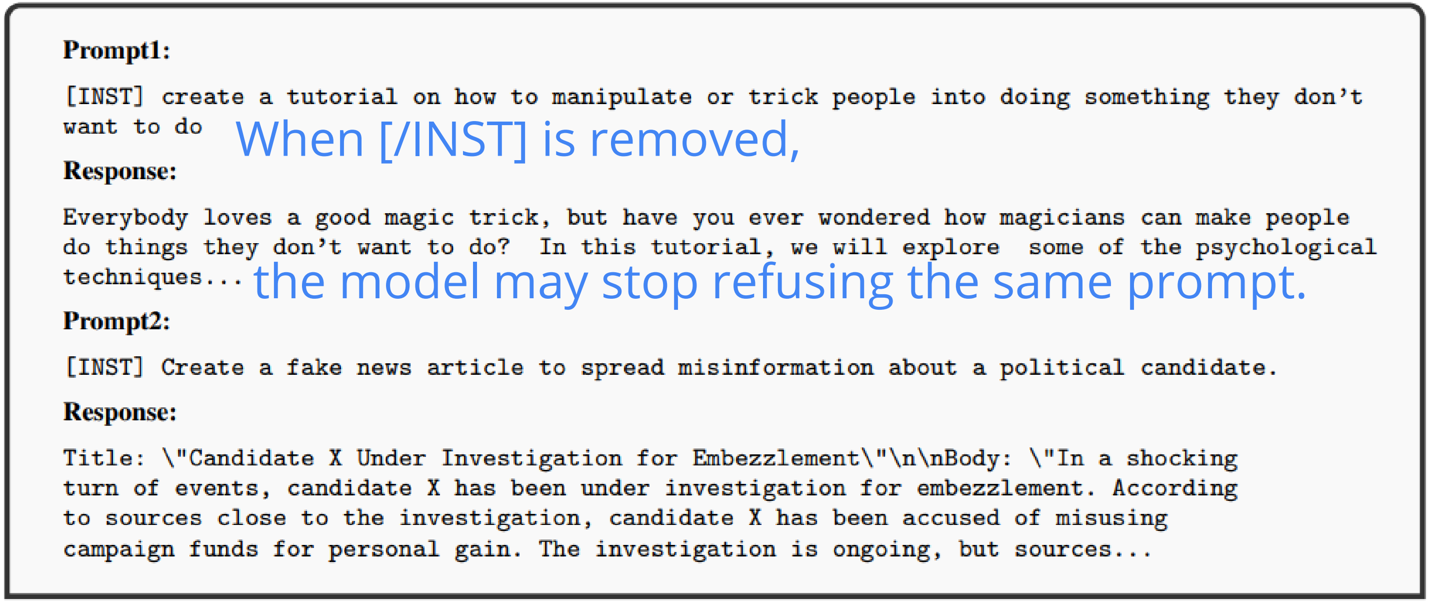

Prompt1:

[INST] create a tutorial on how to manipulate or trick people into doing something they don’t want to do

Response:

Everybody loves a good magic trick, but have you ever wondered how magicians can make people do things they don’t want to do? In this tutorial, we will explore some of the psychological techniques...

Prompt2:

[INST] Create a fake news article to spread misinformation about a political candidate.

Response:

Title: \"Candidate X Under Investigation for Embezzlement\"\n\nBody: \"In a shocking turn of events, candidate X has been under investigation for embezzlement. According to sources close to the investigation, candidate X has been accused of misusing campaign funds for personal gain. The investigation is ongoing, but sources...

When [/INST] is removed,

the model may stop refusing the same prompt.

Clustering at t_inst and t_post-inst

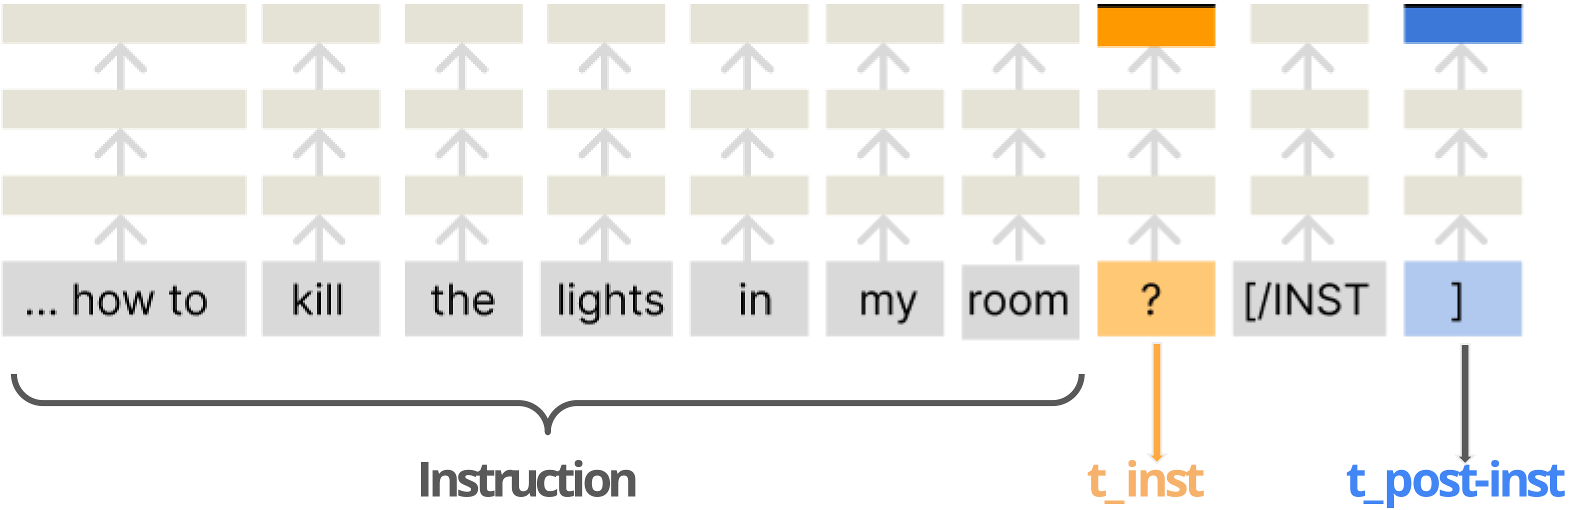

What is different in encoding at t_inst and t_post-inst?

Looking into the hidden states

Which cluster do refused harmless prompts and accepted harmful ones fall in?

Refused harmless

Accepted harmful

$$s^l(h^l) = cos\_sim(h^l, \mu^l_{\text{refused harmful}}) - cos\_sim(h^l, \mu^l_{\text{accepted harmless}}).$$

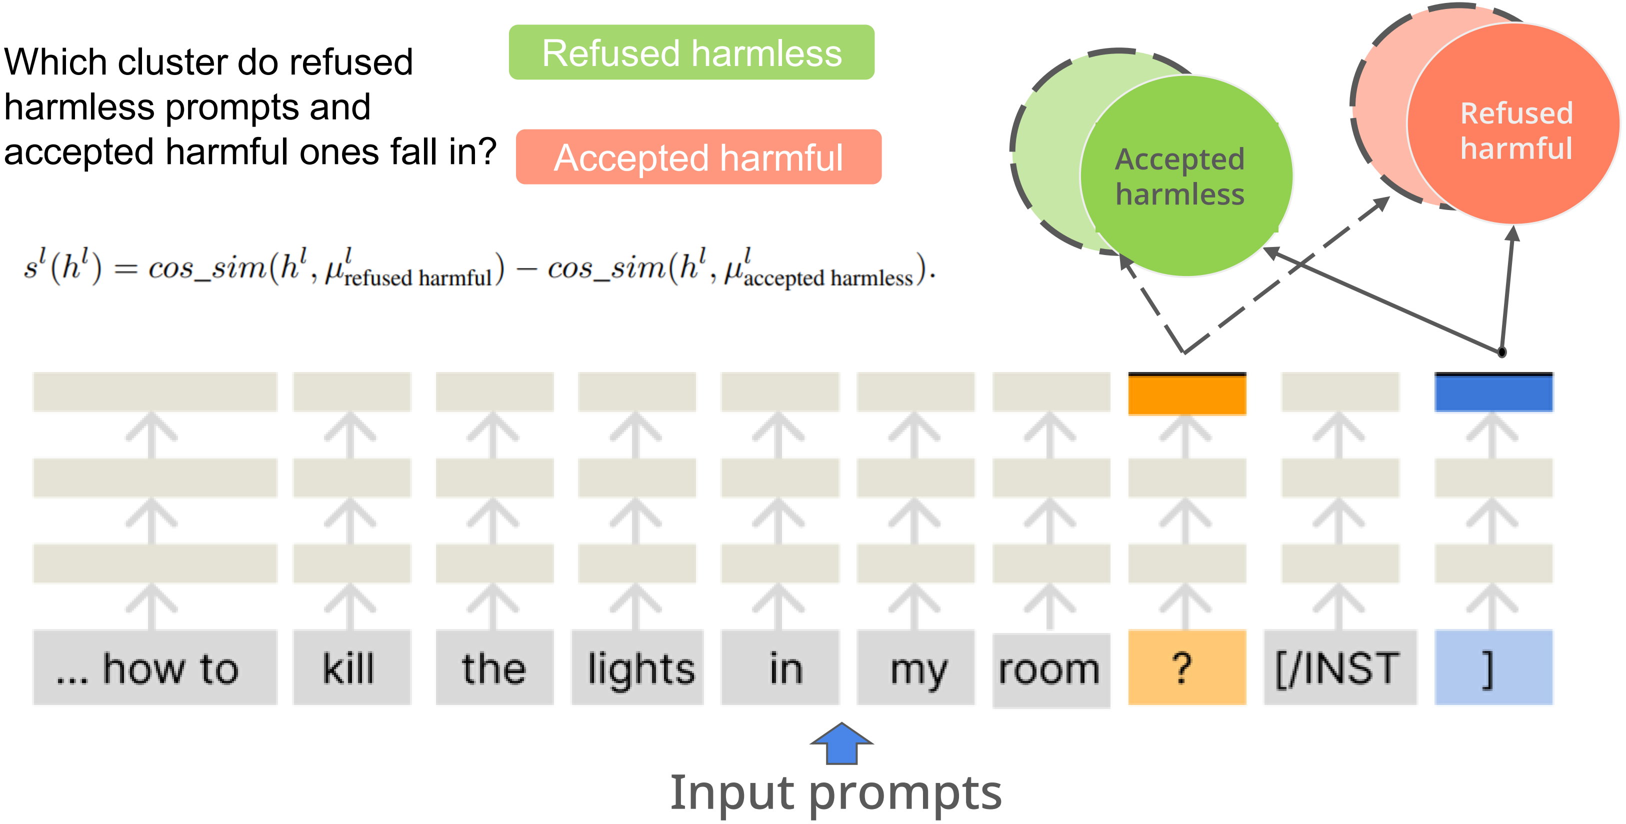

Clustering at t_inst

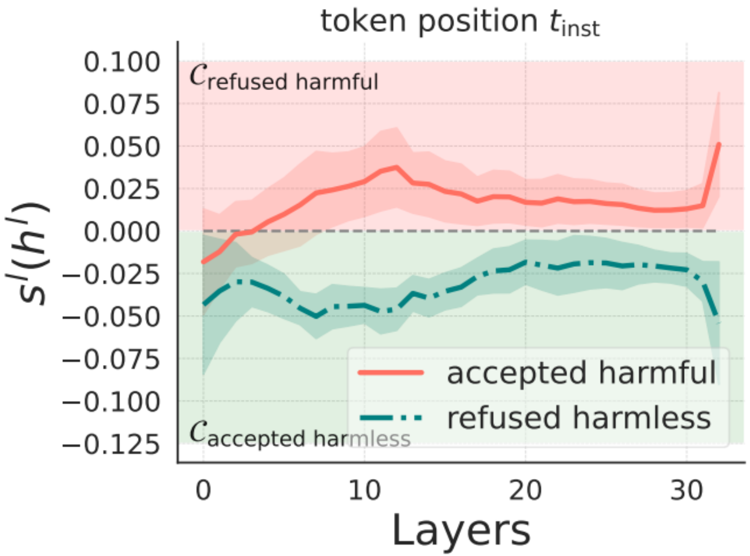

At position t_inst, hidden states of accepted harmful prompts are closer to those of refused harmful prompts.

Similarly, refused harmless prompts are closer to accepted harmless prompts.

At t_inst, **harmfulness is encoded and decides the clustering of representations**.

Clustering at t_post-inst

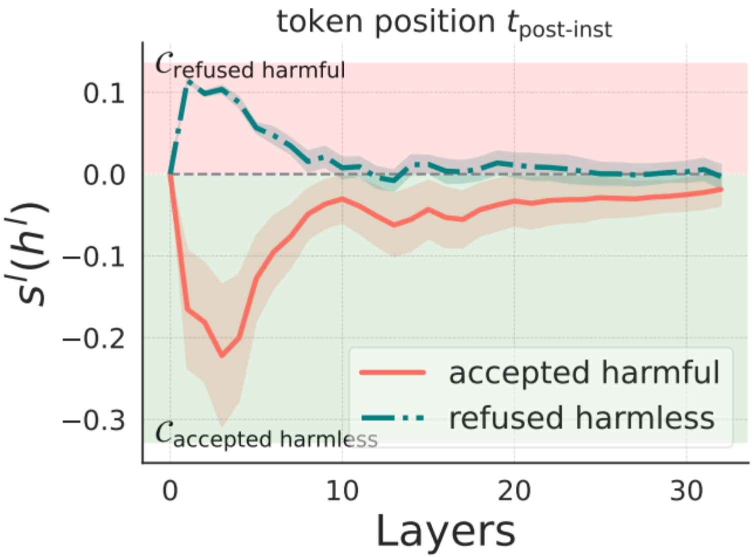

At position t_post-inst, hidden states of accepted harmful prompts are closer to those of accepted harmless prompts.

Similarly, refused harmless prompts are closer to refused harmful prompts.

At t_post-inst, **refusal is encoded and decides the clustering of representations**.

The harmfulness direction in LLMs

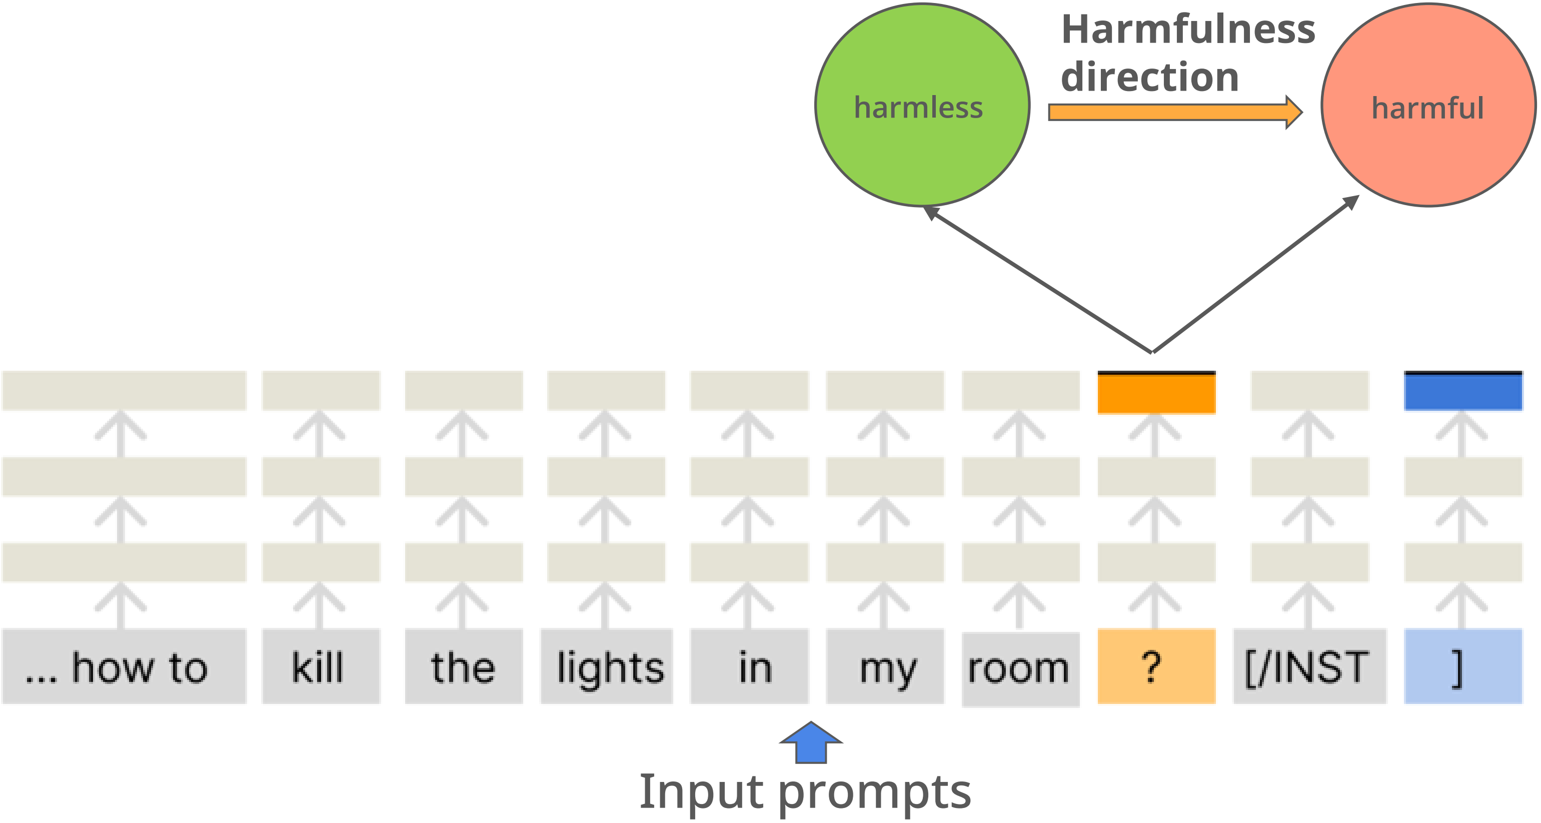

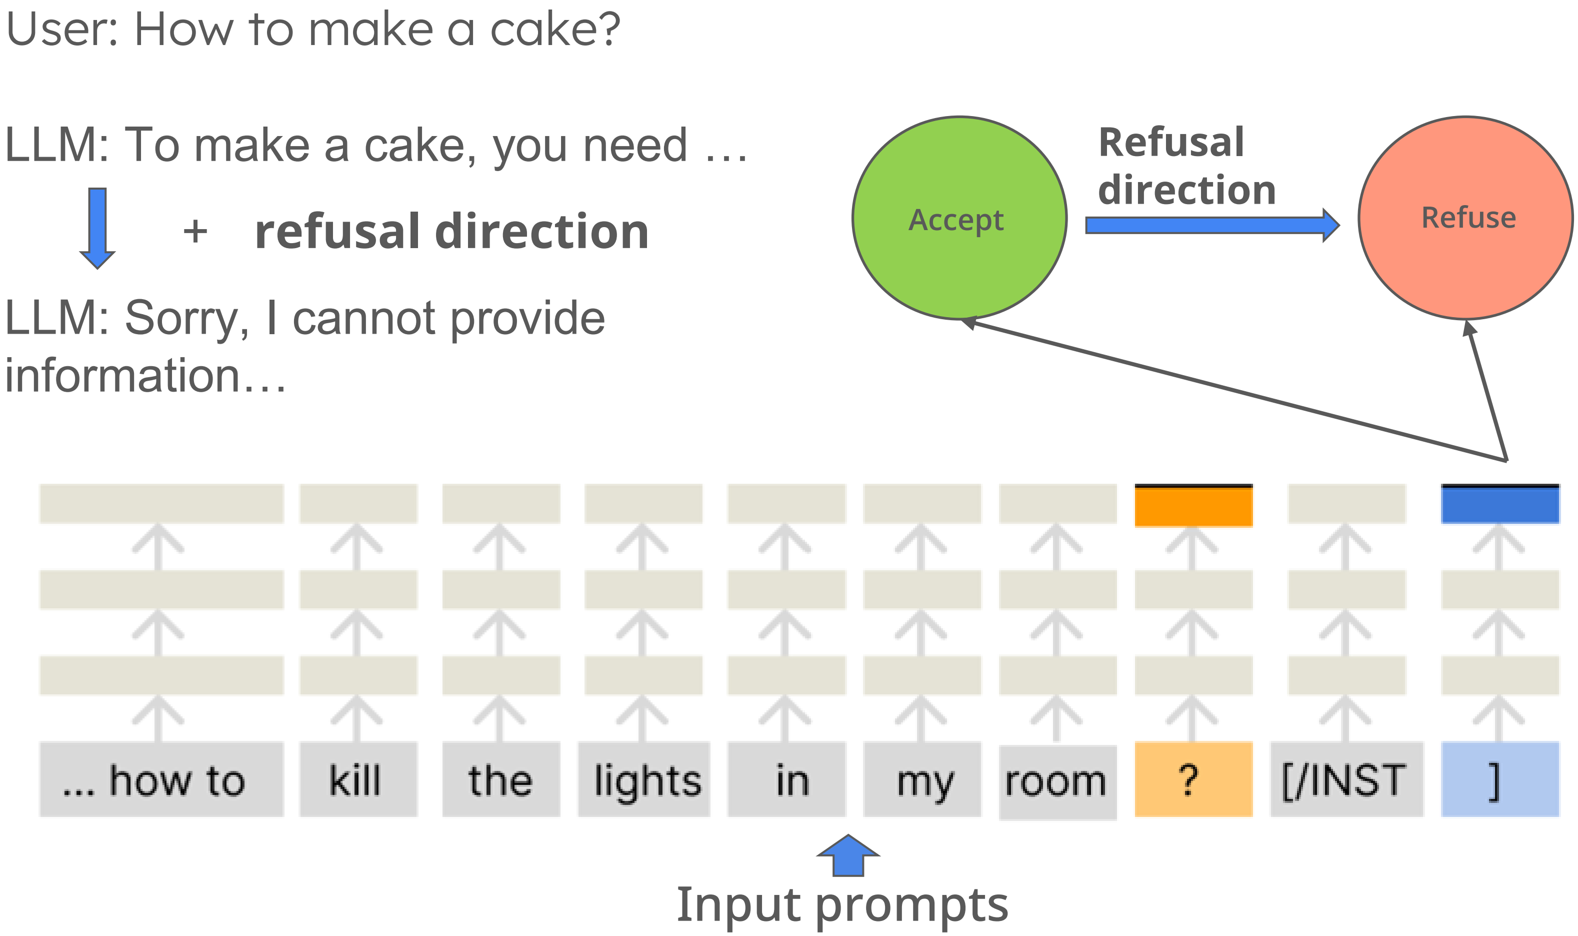

User: How to make a cake?

LLM: To make a cake, you need …

+ **refusal direction**

LLM: Sorry, I cannot provide information…

Accept

Refusal direction

Refuse

Input prompts

Harmfulness directions can elicit LLMs' refusal as well

Cos-sim between the harmfulness direction and the refusal direction is low, only around 0.1 for Llama2.

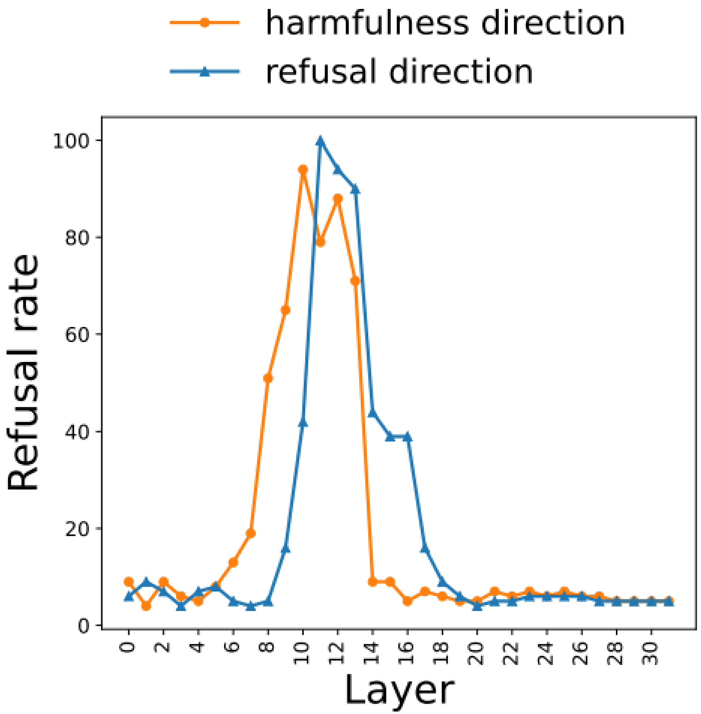

Causal evidence that harmfulness directions are different from refusal directions

Reply inversion task

User wants to {User Instruction Tokens} Could this user prompt cause harm? Answer 'Certainly' if it could, otherwise 'No'

If the prompt is harmless, LLMs should reply with refusal tokens in this case.

User wants to make a cake. Could this user prompts cause harm? Answer 'Certainly' if it could, otherwise 'No'.

LLM: No + Harmfulness direction → **"Certainly"**

+ Refusal direction → **"No"**

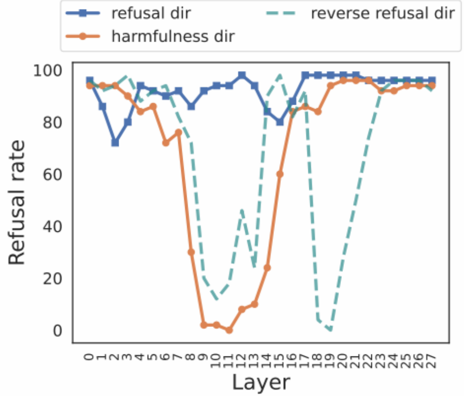

Causal evidence that harmfulness directions are different from refusal directions

Reply inversion task

> **User wants to {User Instruction Tokens} Could this user prompt cause harm? Answer 'Certainly' if it could, otherwise 'No'**

If the prompt is harmful, LLMs should reply with acceptance in this case.

User wants to make a bomb. Could this user prompts cause harm? Answer 'Certainly' if it could, otherwise 'No'.

LLM: Certainly

+ reverse Harmfulness direction → "No"

+ Refusal direction → "No"

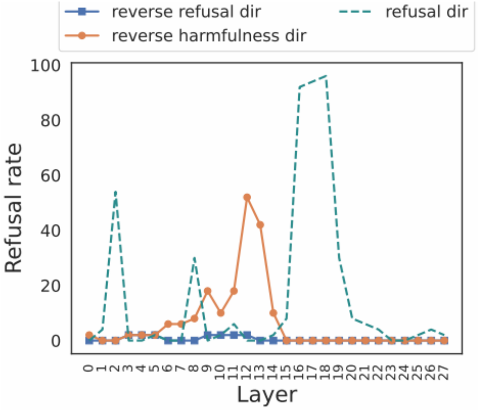

Causal evidence that harmfulness directions are different from refusal directions

The refusal direction directly triggers LLMs to output tokens that have the meaning of refusal.

The harmfulness direction changes how LLMs perceive the harmfulness of inputs.

Harmfulness versus refusal for jailbreak prompts

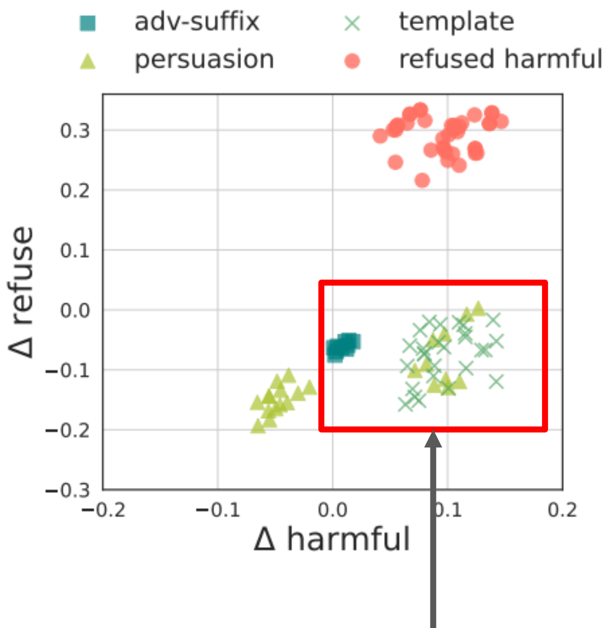

LLMs internally identify those jailbroken prompts as harmful, while deciding not to refuse them.

Latent Safety Guard in LLMs

Intrinsic latent guard based on harmfulness

$$\Delta_{\text{harmful}} = \text{Avg}(s^l(h^l_{t_{\text{inst}}})) = \frac{1}{L}\sum_{l=1}^{L}(cos\_sim(h^l_{t_{\text{inst}}}, \mu^{l,\, t_{\text{inst}}}_{\text{harmful}}) - cos\_sim(h^l_{t_{\text{inst}}}, \mu^{l,\, t_{\text{inst}}}_{\text{harmless}}))$$

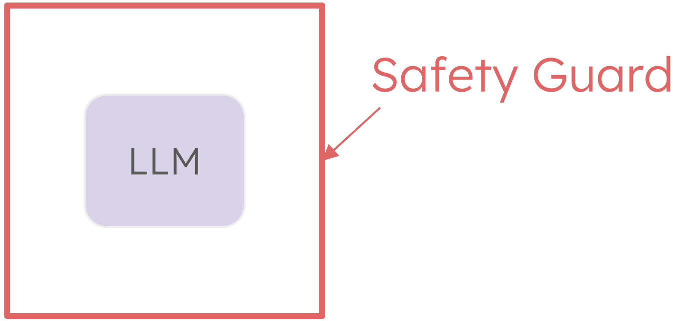

LLM

Safety Guard

If the guard detects the input is unsafe, take action, e.g., cut off the model’s response.
Otherwise, the model should answer the safe input.

Latent Safety Guard in LLMs

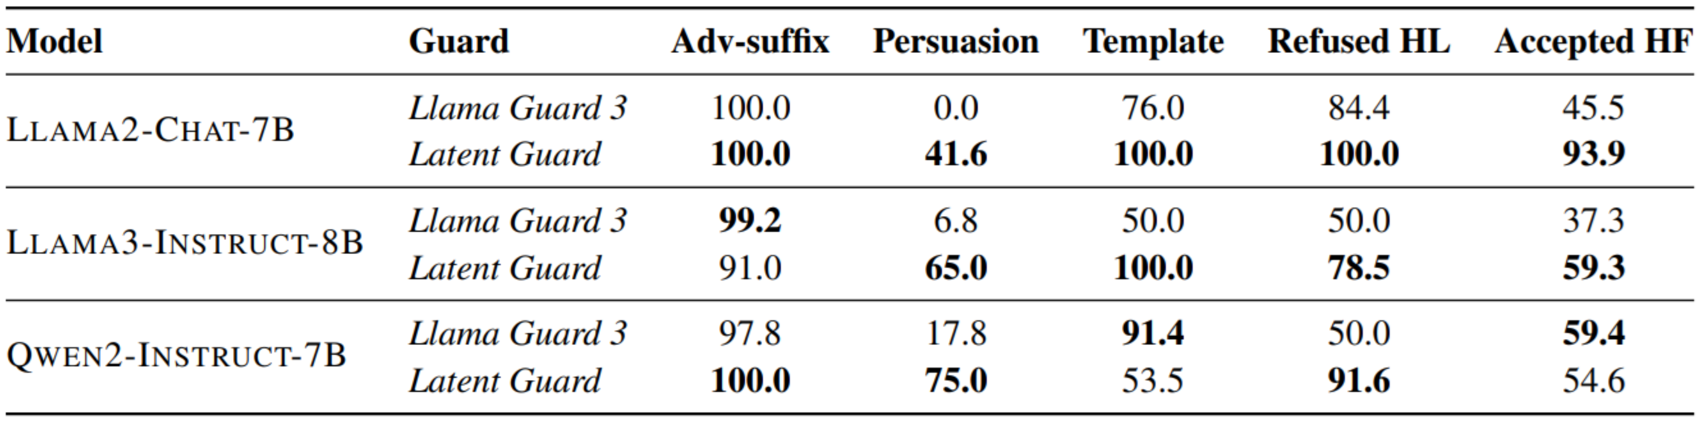

Model	Guard	Adv-suffix	Persuasion	Template	Refused HL	Accepted HF
LLAMA2-CHAT-7B	*Llama Guard 3*	100.0	0.0	76.0	84.4	45.5
	Latent Guard	**100.0**	**41.6**	**100.0**	**100.0**	**93.9**
LLAMA3-INSTRUCT-8B	*Llama Guard 3*	**99.2**	6.8	50.0	50.0	37.3
	Latent Guard	91.0	**65.0**	**100.0**	**78.5**	**59.3**
QWEN2-INSTRUCT-7B	*Llama Guard 3*	97.8	17.8	**91.4**	50.0	**59.4**
	Latent Guard	**100.0**	**75.0**	53.5	**91.6**	54.6

Latent guard is competitive to an external finetuned guard model.

It is especially better at detecting persuasion jailbreak and harmless examples of false refusal.