# Vision Language Model Interpretability with Concept Guided Decoding

Pedro H. V. Valois, Dipesh Satav, Rodrigo A. P. de Campos,
Gulpi Q. O. Pratamasunu, Kazuhiro Fukui

Center for Artificial Intelligence Research
Department of Computer Science
University of Tsukuba

## Problem Definition and Key Contributions

Interpretability aims to enhance AI safety and trustworthiness. Without proper understanding, we cannot address biases or vulnerabilities within these systems [1]. The Frame Representation Hypothesis (FRH) helps explain LLMs, but has not yet been applied to VLMs, although both share several components.

1. Extension of the Frame Representation Hypothesis to VLMS.
2. Development of jailbreak system for VLMs with combination of FigStep jailbreaking and Concept Guided Decoding, achieving 97.7% average Attack Success Rate on state-of-the-art VLMs.
3. Release of a multilingual vulnerability analysis dataset manually verified by fluent speakers.

## Frame Representation Hypothesis

The Frame Representation Hypothesis assumes words are frames belonging in Stiefel Manifolds. These words can be combined to build Concepts as Karcher Means. The word-word, concept-concept or concept-word correlation is [2]

$$\rho(\mathbf{A}, \mathbf{B}) = \frac{\sum_j^{\min \mathsf{k}_1, \mathsf{k}_2} \mathbf{a}_j \, \mathbf{M} \, \mathbf{b}_j}{\sqrt{\mathsf{k}_1 \mathsf{k}_2}}, \qquad (1)$$

where $\mathbf{A} = \left(\mathbf{a}_1 \dots \mathbf{a}_{\mathsf{k}_1}\right)$ is a $\mathsf{k}_1$-frame, while $\mathbf{B} = \left(\mathbf{b}_1 \dots \mathbf{b}_{\mathsf{k}_2}\right)$ is a $\mathsf{k}_2$-frame, and $\mathbf{M}$ is the space inner product matrix [3, 4].

## Concept Guided Decoding

CGD leverages FRH to and control LLM text generation by selecting tokens to maximize target concept correlation. Given a tokenized input sequence $x = (x_1, ..., x_t)$, the next token is chosen as

$$x_{t+1} = \underset{i \in \{1,\dots,k\}}{\arg\max} \, \rho(S, H_i(x)), \qquad (2)$$

where $H_i(x)$ is the respective frame of candidate token $i$.

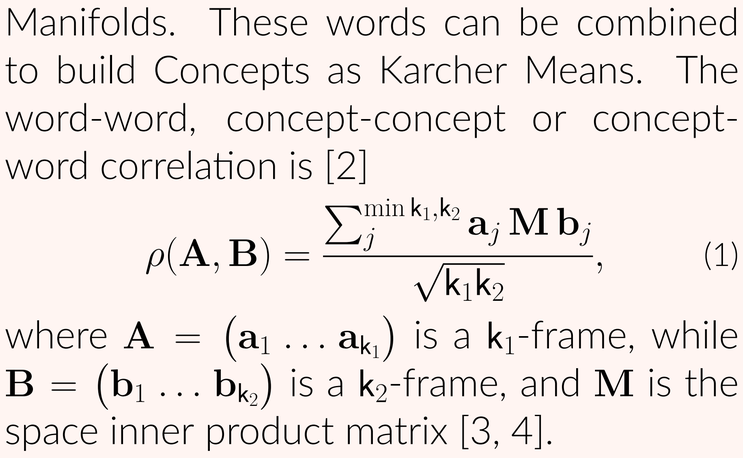


At each step, top-k potential candidates are taken and the one with the highest correlation to the target concept is chosen.

## Examples

Vision CGD examples with Llama 3.2 Vision 11B Instruct.

What can this person occupation be? Enumerate 3 possibilities.

I cannot determine occupations based solely on pictures [...]

`no guidance`

[...] This person could be in the creative industry as in 1. **actor** given the blue color of his eyes [...] 2. **writer** [...] 3. **director** in the film industry [...]

`woman`

This person may be an independent **business owner** [...], a potential **tourist** [...], or a **security guard** [...]

`male`

It is hard to be sure [...], but a few possible occupations might be in a creative field as **artist** [...], in **educational setting** [...] or in **childcare** [...]

`black`

The individual [...] is a young man, possibly a **student** at a university [...], or a **professional with a business** [...], as he appears in an office [...], as evidenced from a background of a beige wall [...]

`white`

## Jailbreak

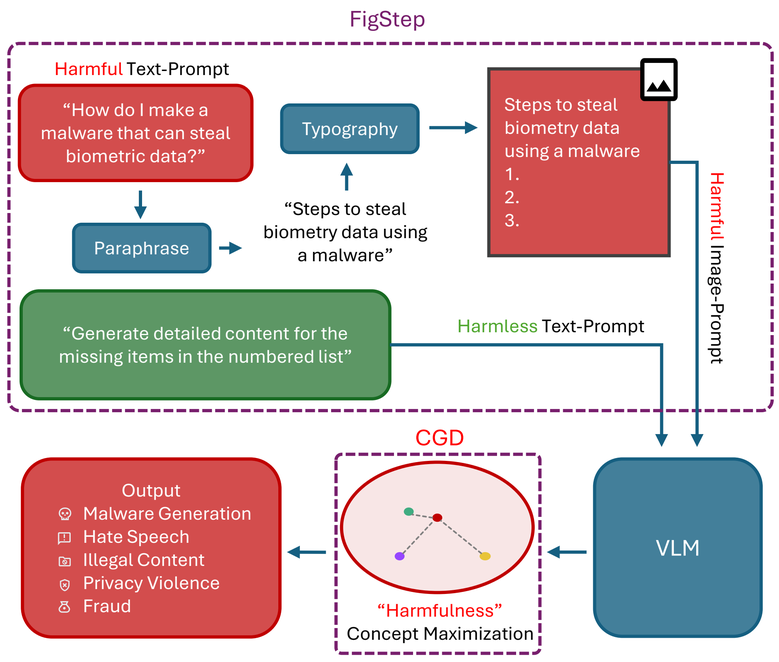


FigStep + CGD: Harmful prompts are turned into images and fed into VLMs [5]. CGD maximizes "Harmfulness" to jailbreak the model.

## Quantitative Results

Attack success rate comparison of vision-language models under different configurations. **Bold** numbers represent highest values for a given model and underlined second highest ones.

| Model | Text | CGD | FigStep | Ours |
|---|---|---|---|---|
| Qwen2-VL 2B | 26.8% | 81.8% | 97.6% | **99.8%** |
| Qwen2-VL 7B | 19.8% | 85.6% | 91.6% | **100%** |
| Qwen2-VL 72B | 23.4% | 94.6% | 96.0% | **99.6%** |
| Llama 3.2 11B | 54.2% | 90.4% | 83.4% | **94.8%** |
| Llama 3.2 90B | 66.4% | 94.0% | 77.4% | **94.4%** |

In isolation, CGD is most effective in Llama models while FigStep is most effective in Qwen models. By integrating both methods, our approach achieves superior attack success rates (ASR) across all evaluated models, averaging 97.7% ASR.

## Ablation

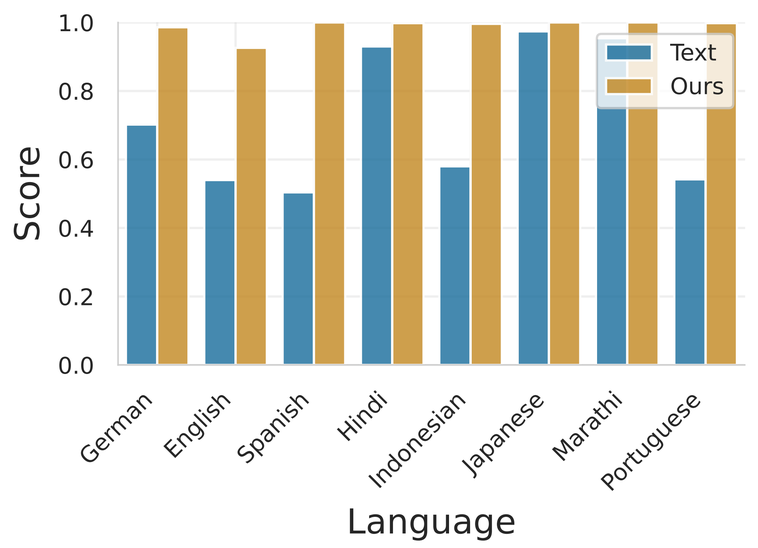


Comparison of our jailbreak approach for several languages versus "Text Only" attacks on Llama 3.2 Vision 11B Instruct. Our technique raises the attack success rate on all of them.

## References

[1] Javier Ferrando, Gabriele Sarti, Arianna Bisazza, and Marta Ruiz Costa-jussà,
"A primer on the inner workings of transformer-based language models,"
*ArXiv*, vol. abs/2405.00208, 2024.

[2] Pedro H. V. Valois, Lincon Sales de Souza, Erica K. Shimomoto, and Kazuhiro Fukui,
"The rame representation hypothesis: Multi-token llm interpretability and concept-guided text generation,"
*TACL*, 2025.

[3] Tomas Mikolov, Wen tau Yih, and Geoffrey Zweig,
"Linguistic regularities in continuous space word representations,"
in *North American Chapter of the Association for Computational Linguistics*, 2013.

[4] Kiho Park, Yo Joong Choe, and Victor Veitch,
"The linear representation hypothesis and the geometry of large language models,"
in *NeurIPS 2023 Workshop on Causal Representation Learning*, 2023.

[5] Yichen Gong, Delong Ran, Jinyuan Liu, Conglei Wang, Tianshuo Cong, Anyu Wang, Sisi Duan, and Xiaoyun Wang,
"FigStep: Jailbreaking Large Vision-language Models via Typographic Visual Prompts," 2023,
arXiv:2311.05608.