

From Behavioral Performance to Internal Competence: Interpreting Vision-Language Models with VLM-LENS

Hala Sheta^{1,2,*} Eric Huang^{1,5,*} Shuyu Wu^{3,*}
Ilia Alenabi¹ Jiajun Hong⁴ Ryker Lin¹ Ruoxi Ning^{1,2} Daniel Wei¹ Jialin Yang^{1,6}
Jiawei Zhou⁴ Ziqiao Ma³ Freda Shi^{1,2}

¹University of Waterloo ²Vector Institute ³University of Michigan

⁴Stony Brook University ⁵McGill University ⁶Dartmouth College

 <https://github.com/compling-wat/vlm-lens>

Abstract

We introduce VLM-LENS, a toolkit designed to enable systematic benchmarking, analysis, and interpretation of vision-language models (VLMs) by supporting the extraction of intermediate outputs from any layer during the forward pass of open-source VLMs. VLM-LENS provides a unified, YAML-configurable interface that abstracts away model-specific complexities and supports user-friendly operation across diverse VLMs. It currently supports 16 state-of-the-art base VLMs and their over 30 variants, and is extensible to accommodate new models without changing the core logic.

The toolkit integrates easily with various interpretability and analysis methods. We demonstrate its usage with two simple analytical experiments, revealing systematic differences in the hidden representations of VLMs across layers and target concepts. VLM-LENS is released as an open-sourced project to accelerate community efforts in understanding and improving VLMs.

1 Introduction

Vision-language models (VLMs; Kirillov et al., 2023; Radford et al., 2021; Li et al., 2022; Liu et al., 2023; Wang et al., 2024a, *inter alia*) have become essential across a wide range of applications, including multimodal understanding (Yue et al., 2024), robotics (Li et al., 2024), and world modeling (Gao et al., 2025). However, existing VLM benchmarks predominantly adopt exact-match based accuracy and its derivations to evaluate model performance (Lin et al., 2014; Johnson et al., 2017; Yue et al., 2024; Fu et al., 2024, *inter alia*), which may either overlook the information embedded in their hidden representations (Zhang et al., 2025) or yield misleading assessments due to shortcut exploitation (Xu et al., 2025). Currently, there lacks a unified framework for extracting the

internal representations of VLMs, making it challenging to assess model capabilities that go beyond simple performance evaluations.

Meanwhile, interpretability research and toolkits for VLMs remain underdeveloped compared to their text-only counterparts (Nanda and Bloom, 2022; Belrose et al., 2023; Ali et al., 2025, *inter alia*), posing significant challenges to systematically understanding their internal knowledge and decision-making processes. To the best of our knowledge, extending existing interpretability toolkits, such as TransformerLens (Nanda and Bloom, 2022), to support VLMs requires substantial engineering effort, as these tools are primarily designed for text-based Transformers.

To address these challenges in both benchmarking and interpretability, we present **VLM-LENS** (Figure 1), a toolkit that enables easy extraction of VLM intermediate output from any layer in a forward pass. The key features include:

- **Unified interface.** It abstracts out the model-specific setup and preprocessing complexities, allowing operations across models through a unified interface. Users can specify custom configurations via a YAML file with minimal boilerplate code provided, and the toolkit automatically handles model loading, preprocessing, and inference.
- **Model-specific environmental support.** Different VLMs often require different, and sometimes mutually conflicting, libraries. To address this issue, VLM-LENS provides model-specific environment setups, each of which can be easily installed with a single-line pip install command. A rigorous code review process ensures the consistency and reproducibility of the environment setups across different platforms.
- **Extensive model coverage and flexible nature.** The toolkit supports a diverse set of state-of-the-art VLMs, spanning widely used open-source models to recently developed, less-documented ones. Currently, VLM-LENS supports 16 base

* Equal contribution.

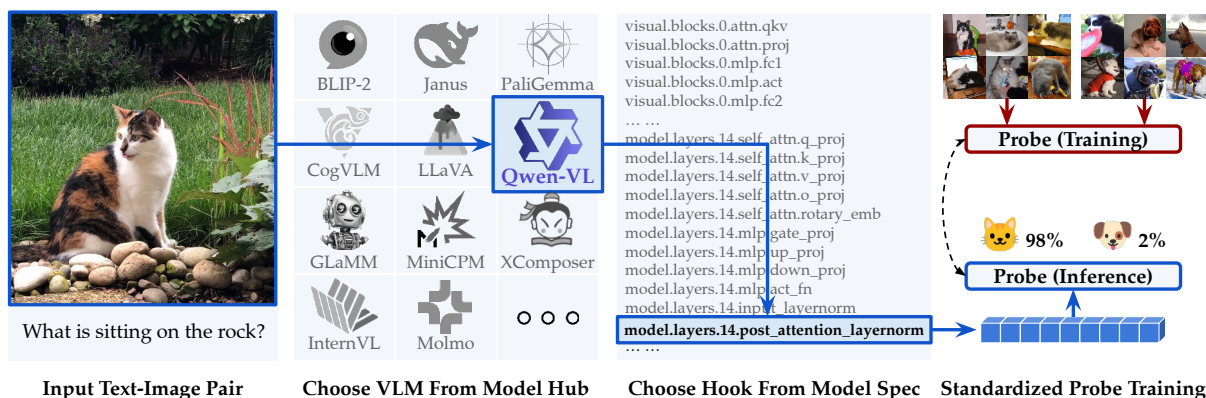


Figure 1: An example use case of VLM-LENS, where intermediate output from Qwen2-VL (Wang et al., 2024a) is extracted for probing.

VLMs and over 30 variants, with a highly extensible design that allows users or contributors to add new models with minimal effort.

VLM-LENS streamlines analytical tasks for VLMs, such as probing (Ettinger et al., 2016), neural circuit inspection (Chughtai et al., 2023), and knowledge tracing (Basu et al., 2024), as well as diagnosing model capabilities and limitations (Zhang et al., 2025; Stevens et al., 2025). As such, we anticipate that the toolkit will enable researchers and practitioners to conduct more fine-grained and rigorously controlled analyses of VLMs. The toolkit is released under the Apache-2.0 license.

2 Related Work

Vision-language models. Since vision and text naturally convey information about the world in two complementary modalities, there has been sustained interest in integrating them within unified frameworks (Kiros et al., 2014; Radford et al., 2021; Liu et al., 2023, *inter alia*). Earlier work primarily encoded images and text in a shared embedding space to facilitate efficient retrieval and matching (Kiros et al., 2014; Faghri et al., 2018; Radford et al., 2021). With recent advances in generative text models, exemplified by Brown et al. (2020), the focus has shifted toward building large-scale VLMs that generate text conditioned on both images and textual prompts (Liu et al., 2023; Wang et al., 2024a; Zhang et al., 2023, *inter alia*). While some models offer straightforward hidden-state extraction through open-sourced libraries such as HuggingFace Transformers (Wolf et al., 2020),¹ many require insufficiently documented customiza-

¹For example, by setting `output_hidden_state=True` in the forward function of LLaVA-1.5-7B (Liu et al., 2023); see the documentation at https://huggingface.co/docs/transformers/model_doc/llava.

tions. Additionally, there is no systematic support for extracting intermediate representations beyond the layer-wise output, such as attention maps and intermediate features before layer normalization (Ba et al., 2016). VLM-LENS addresses this gap by providing a structured and unified interface for extracting intermediate output across many VLMs, thereby enabling detailed analyses of models.

Performance and competence analysis of VLMs.

Efforts have been made to benchmark the performance of VLMs on various tasks, such as visual question answering (Johnson et al., 2017), image captioning (Lin et al., 2014) and general cross-modal understanding (Yue et al., 2024; Fu et al., 2024). These benchmarks largely assess model performance through exact-match based accuracy. However, accuracies fall short of capturing the full spectrum of model *competence*, which encompasses the internal mechanisms and generalizable knowledge that a model possesses, possibly beyond its observable textual output. To address this issue, recent work has started to explore the competence of VLMs through more fine-grained analyses via hidden states (Stevens et al., 2025) or output probability (Zhang et al., 2025). As such, VLM-LENS offers a toolkit for localizing hidden causal mechanisms in VLM, enabling convenient and flexible model competence analysis and assessing beyond simple accuracy-based evaluations.

Transformer interpretability toolkits. With increasing interest in interpreting Transformers, there have been various toolkits supporting their interpretation and analysis (Clark et al., 2019; Nanda and Bloom, 2022; Belrose et al., 2023; Ali et al., 2025, *inter alia*). However, most existing generic toolkits focus on text-only (e.g., Nanda and Bloom, 2022) or vision-only (e.g., Joseph et al., 2025) Transform-

ers, while the VLM counterparts significantly fall behind.² To the best of our knowledge, all existing VLM interpretability toolkits (Palit et al., 2023; Ben Melech Stan et al., 2024; Neo et al., 2024) support no more than a single model, and the designs are not easily extensible to support other models. To bridge this gap, VLM-LENS provides a unified framework to extract internal representations of VLMs, which can be coupled with existing interpretability methods (Thrush et al., 2022; Zhang et al., 2024; Basu et al., 2024; Stevens et al., 2025, *inter alia*) to analyze and assess a wide range of state-of-the-art VLMs.

3 VLM-Lens

At a high level, the VLM-LENS toolkit incorporates the hook mechanism of PyTorch (Paszke et al., 2019) to extract the internal representations of VLMs, which are then stored in a database for further analysis. We detail the design philosophy and key implementation aspects as follows.

Design philosophy. The design of VLM-LENS is driven by the need for a simple yet flexible toolkit to extract internal representations while accommodating the diverse dependency requirements of different VLMs. This is implemented through a central interface (`src/main.py`) that accepts model-specific configuration files in YAML format, along with dedicated environment setups for each model. Each model-specific implementation inherits from the base class, which standardizes model loading, preprocessing, and inference, while still allowing for model-specific customizations. The extracted intermediate representations are stored in a database with a standardized schema, enabling efficient retrieval and analysis. This design, executed by rigorous peer code review, ensures high extensibility: to support a new model, developers only need to implement a new model-specific class without modifying the core logic of the toolkit.

Supported VLMs. We currently support: Aya-Vision (Dash et al., 2025), Blip-2 (Li et al., 2023), CLIP (Radford et al., 2021), CogVLM (Wang et al., 2024b), GLaMM (Rasheed et al., 2024), InternLM-XComposer (Zhang et al., 2023), InternVL (Chen et al., 2024), Janus (Wu et al., 2025), LLaVA (Liu et al., 2023), MiniCPM-o (Team, 2025) and MiniCPM-V-2 (Yao et al., 2024), Molmo (Deitke

²For example, TransformerLens (Nanda and Bloom, 2022) can be patched to support vision analysis, but requires a non-trivial setup and is rigid in extending functionality.

et al., 2025), Paligemma (Beyer et al., 2024), Pixtral (Agrawal et al., 2024), PerceptionLM (Cho et al., 2025) and Qwen2-VL (Wang et al., 2024a). The toolkit supports all variants of these models across sizes, with the only requirement being sufficient hardware resources to load the model.

Configuration files. VLM-LENS allows users to specify model configurations, input and output data paths, model layers of interest, and other model-specific parameters through YAML configuration files. Users may extend the configuration files to include additional parameters for their experiments. No hard-coded parameters are used throughout the codebase. As an example, the following file (`configs/models/blip2/blip2.yaml`) specifies the required parameters for extracting the output of layers `vision_model.post_layernorm` and `language_model.lm_head` in Blip2-OPT-2.7B (Li et al., 2023):

```
architecture: blip2
model_path: Salesforce/blip2-opt-2.7b
model:
  - torch_dtype: auto
output_db: output/blip2.db
input_dir: ./data/test-images/
prompt: "Describe the color in this image in one word."
modules:
  - language_model.lm_head
  - vision_model.post_layernorm
```

The command `python src/main.py --config configs/models/blip2/blip2.yaml` will provide the user with the corresponding intermediate output tensors in `output/blip2.db`, a SQLite3 database that can be further queried for analysis.

Pipeline implementation and output database organization. To initialize a model, an accompanying preprocessor is required to process the input images and format the prompt: for example, the chat template for LLaVA-1.5-7B (Liu et al., 2023) and Qwen2-VL (Wang et al., 2024a) is implemented in the preprocessor. The hidden representation extraction process involves registering a forward hook, a callable function that provides access to the input and output tensors of a layer specified by the user. It takes the preprocessed image and prompt as input, and saves the output tensors in a database following the forward pass. All forward hooks are unregistered to prevent contamination in later iterations. A list of possible layers (i.e., modules in PyTorch) of a particular model can be returned easily using the `--log-named-modules` parser argument, if the user requires more information on a model’s internal structure. Lastly, the extracted tensors are stored in a database with the following attributes:

name, architecture, image_path, prompt, label, layer, tensor_dim, tensor. These attributes correspond to the model HuggingFace identifier (e.g., Salesforce/blip2-opt-2.7b), model architecture specified in this toolkit (e.g., blip2), image path, textual prompt content, label of the example (if applicable), layer name (e.g., language_model.lm_head), the dimensionality of the extracted tensor, and the extracted intermediate result tensor itself, respectively.

While we provide a default preprocessor and a hooked forward pass implementation in `src/models/base.py` that can handle simple cases (e.g., LLaVA), these functions can be overridden to accommodate model-specific requirements (see more examples in `src/models`).

4 Usage Example I: Probing

4.1 Experimental Setups

We first demonstrate the usage of VLM-LENS by probing (Ettinger et al., 2016), where we train a set of probes on the extracted representations, evaluating the internal competence of VLMs on recognizing a set of primitive concepts.

Dataset. We create our dataset using CLEVR (Johnson et al., 2017). As depicted in Table 1, we define five categorical splits: color, material, number, shape, and size, each corresponding to a primitive object attribute in the images, as well as a boolean split that may cover any attribute with binary questions. Each split can be considered as a c -way classification task, where c represents the number of possible choices in the split. For each split, 80% data is used for training the probe, whereas the remaining 20% is used for testing.³

Probing approach. Each combination of split D , model \mathcal{M} , and layer ℓ is considered independently in the probing process. For an example $(e_i, y_i) \in D$, where e_i represents the input (i.e., image and text) of the task and y_i stands for its ground-truth prediction category, we extract the intermediate output of model \mathcal{M} at the layer ℓ with mean pooling across the input tokens, denoted as $\mathcal{M}_\ell(e_i) \in \mathbb{R}^d$, where d is the dimensionality of layer ℓ . $\mathcal{M}_\ell(e_i)$ is used as the feature of e_i to train a probe that best predicts the y_i label.

We employ a two-layer perceptron as the probe, with ReLU activations and 512 hidden units.⁴ A k -

³Data available at <https://huggingface.co/compling>.

⁴We searched for the probe hidden size with small-scale experiments, and found 512 to be the hidden size with the best performance consistency.



Split	Question	Answer
boolean	Are there any other things that are the same size as the brown object?	yes
color	There is a small cylinder that is made of the same material as the big brown thing; what color is it?	green
material	What is the material of the big object that is the same color as the small metal cylinder?	metal
number	What number of things are brown blocks or green metallic cylinders that are to the right of the tiny cylinder?	2
shape	The big brown shiny thing is in what shape?	cube
size	There is another thing that is the same shape as the brown metallic object; what is its size?	small

Table 1: Examples from our probing dataset. We input the image and the question to the VLM, extract hidden states, and train a lightweight probe to decode the answer from these representations.

fold cross-validation is conducted to search for the best optimization hyperparameters for each split, using matching-based accuracy as the metric. We then train the probe model on the full training set using the best optimization hyperparameters, and report the accuracy on the test set. Following Hewitt and Liang (2019), we complement our main probe (trained to predict the true ground-truth labels) with a control probe (trained on randomly shuffled labels). A statistically significant advantage of the main probe over the control probe suggests that the VLM encodes task-relevant information, whereas no advantage implies that the probe relies on its own capacity to memorize spurious patterns.

4.2 Results and Discussion

We evaluate the middle (i.e., $\lceil \frac{L}{2} \rceil$) and last layers of eight supported models, where L is the total number of layers in the model of interest (Figure 2). Notably, the probes trained on the Qwen-7b and MiniCPM-o representations achieve an almost perfect accuracy within many dataset splits, with a statistically significant difference from the control probe performance. This effect is especially preva-




Figure 2: Evaluation Accuracy on our probing dataset by model, layer, and split. *Main* refers to probing on the regular data, while *control* stands for probing using data with random labels. The number of asterisks represents the significance level of the Z-test for Bernoulli variables (***: $p = .001$, **: $p = .01$, *: $p = .05$).

lent in the results of the last layer representations with a few exceptions (e.g., Blip-2 and Paligemma). In contrast, models like LLaVA-1.5-7B demonstrate a minor (albeit statistically significant in most cases) difference between the evaluation results in the original and control conditions, indicating weaker competence.

Across all models, the color attribute appears to be the most salient feature, with the main probe performance significantly better than the control results across both layers. Models with instruction-following and multimodal understanding capabilities, such as Qwen2-VL, MiniCPM-o, and InternVL, performed well on the more difficult splits such as material, number, and shape, especially when using the last-layer representations. This probing-based competence evaluation complements existing accuracy-driven benchmarks by providing a more detailed understanding of what is represented in the internal states, grounding the model performance in interpretable primitive knowledge.

5 Usage Example II: Concept Similarity

The second experiment is inspired by the classic *Stroop Effect* (Stroop, 1935), which demonstrates that humans exhibit slower and more error-prone responses when asked to name the font color of a word that is itself an incongruent color term (e.g., the word red printed in blue ink). We adapt this paradigm to VLMs by constructing images with deliberate incongruities between three color cues (Figure 4): the lexical word (e.g., white), the font color (e.g., yellow), and the background color (e.g., blue). This design probes how VLMs ground the notion of color under ambiguous instructions.

5.1 Experimental Setups

Prompt setup. Unlike humans in the Stroop task, who are explicitly instructed to name the ink color, we query the model with an intentionally ambiguous prompt: *Describe the color in this image in one word.*, coupled with a single image with incongruent cues. This allows us to study which representation (lexical, font, or background) the model privileges in its internal embeddings.

Prototype construction. To establish references for primitive color concepts (e.g., red, blue, green), we retrieve the top 10 de-duplicated Creative Commons-licensed images using the Google Images API.⁵ We extract the intermediate layer rep-

⁵<https://developers.google.com/photos>





Figure 3: Cosine similarity between Stroop task images and primitive color concepts. Results are shown as a function of model layer (x-axis) and number of PCA components retained (y-axis), with orange surfaces indicating matching conditions and blue surfaces indicating mismatching conditions when considering different aspects.



(a) Word *white* written in yellow on blue canvas. (b) Word *black* written in yellow on red canvas.

Figure 4: Example images used in the Stroop Task.

representations of these images coupled with the ambiguous textual instruction using VLM-LENS.

Similarity-based analysis. We apply Principal Component Analysis (PCA) to identify the directions that capture the most color variation. Let the zero-meaned embeddings be denoted as $E \in \mathbb{R}^{n \times d}$, where n is the number of reference images corresponding to primitive color concepts, and d is the dimensionality of the intermediate layer output. PCA learns a linear projection $W \in \mathbb{R}^{d \times d'}$ that maps E to a lower-dimensional representation $E' = EW \in \mathbb{R}^{n \times d'}$ ($d' < d$).

We evaluate the Stroop task images under different numbers of retained principal components: for one image at a specific model layer, we extract the corresponding hidden representation, projecting it into the transformed space with W , and compute the average cosine similarity with the reference color concepts. For lexical, font, and background colors, we report the average matched and mismatched cosine similarities across layers.

5.2 Results and Discussion

We visualize the results for LLaVA-1.5-7B (Figure 3), where a larger gap between matched and mismatched data indicates a more prominent fea-

ture in the Stroop task. We find a clear separation between the match and mismatch conditions across all three settings, indicating all three types of information (i.e., lexical content, foreground font color, and background color) are reliably encoded in the model. As expected, background color (Figure 3c) produces the strongest contrast between matched and mismatched examples. However, somewhat surprisingly, color presented as lexical content is more prominent than font color, evidenced by the gap in Figure 3a than that in Figure 3b. In addition, all concepts require a sufficient number of PCA components to achieve a clear separation, suggesting that color information is not captured in a single linear direction in the representation space.

6 System Evaluation

We evaluate VLM-LENS, in terms of time and memory efficiency, on the inference procedure of our supported models, using a subset of the MSCOCO validation set (Lin et al., 2014)⁶ with 2,690 examples (Table 2). All experiments are done on a single NVIDIA-A40 GPU with sufficient CPU memory, using an inference batch size of 1. Users may use the reported performance statistics for informed choices on GPU selection and advanced inference techniques.

6.1 Inference Time

The model inference times are calculated for the duration of the forward inference on the dataset, including the database saving execution, and disregarding the model and processor loading time. As expected, CLIP is the fastest despite its high precision point, which is a result of its small number

⁶<https://huggingface.co/datasets/compling/coco-val2017-obj-qa-categories>

Model	# Params	Precision	Peak Mem (MB)	Inference Time (seconds)	Per-Instance Time (seconds)
CLIP (2021)	150M	float32	616.69	135	0.025
Blip-2 (2023)	2.7B	float32	15,261.28	295	0.055
InternLM-XComposer-2.5 (2023)	7B	bfloat16	24,037.64	3,056	0.569
InternVL (2024)	8B	bfloat16	21,136.92	3,347	0.623
LLaVA-1.5 (2023)	7B	float16	29,031.55	1,566	0.291
MiniCPM-V-2 (2024)	2.8B	bfloat16	7,154.23	495	0.092
MiniCPM-o (2025)	8B	bfloat16	18,058.20	671	0.125
Molmo (2025)	7B	float32	34,403.26	2,841	0.529
Paligemma (2024)	3B	float32	12,018.30	554	0.103
Qwen2-VL (2024a)	7B	bfloat16	33,840.66	1,497	0.279

Table 2: Inference properties of different models in VLM-LENS after execution on the MSCOCO (Lin et al., 2014) dataset, including GPU memory usage and inference time metrics.

of parameters. It is worth noting that CLIP is also the only model trained using contrastive image-text matching, rather than incorporating visual tokens alongside text, and, therefore, is not directly comparable to other models. InternVL, InternLM-XComposer-2.5, and MolMo are the slowest models, with high per-inference times compared to other models with similar parameter counts (7B). Furthermore, within the same parameter count and precision point, InternVL and MiniCPM-o demonstrate a disparate difference in inference time (0.623s vs 0.125s per-inference, respectively), which is likely due to differing architecture optimizations and input processing methods.

6.2 Memory Usage

To approximate the memory usage of each model using our toolkit, we record the precision and peak GPU memory usage (in MB; Table 2). Similar to inference time, CLIP demonstrates the lowest memory footprint (617MB), due to its compact architecture. In contrast, LLaVA-1.5 and Qwen2-VL use the most memory, despite having a low precision point and the same parameter count as many other models. In general, bfloat16 and float16 precision points reduce memory usage, but their effectiveness varies depending on the architecture: InternLM-XComposer and Qwen2-VL still demonstrate a large memory footprint (>24GB).

7 Conclusion and Discussion

In this paper, we introduce VLM-LENS, a toolkit that aims to benchmark, analyze, and interpret VLMs systematically. We demonstrate that the toolkit enables the assessment of the internal competence of a wide range of VLMs (§ 4), going beyond the simple accuracy-based evaluations provided by most existing benchmarks (Yue et al.,

2024, *inter alia*). We offer performance statistics (Table 2) for users’ informed choice on both models and hardware. Users who develop new VLMs through training and fine-tuning can also easily use the provided probing framework to diagnose their model capabilities. We anticipate that this toolkit will lower the barrier for evaluations of VLMs in a scientifically rigorous way.

Results from our demonstrative experiment (§ 4) align with the caveat that performance alone may be insufficient for evaluating models (Hu and Levy, 2023; Zhang et al., 2025; Wang and Shi, 2025). Although all evaluated models are considered highly capable on existing benchmarks, many still fail on simple synthetic data. These findings reaffirm concerns about the reliability of VLMs, and VLM-LENS will actively support their analysis and improvement along these lines.

Limitations and planned community support.

The current toolkit does not directly support more downstream tasks than probing, such as attention interpretation and neural circuit discovery. Additionally, our current inference and database storage approach prevents the use of gradient-based saliency analyses such as Grad-CAM (Selvaraju et al., 2017). Current users still need to implement their customized functions; however, we anticipate rapid community contributions to the repository for a diverse range of tasks—efforts we are committed to supporting in the long term.

Acknowledgments

We thank Yifan Jiang and Michael Ogezi for their constructive feedback. This work is supported in part by NSERC RGPIN-2024-04395, a Vector Scholarship to HS, and a Canada CIFAR AI Chair award to FS.

References

- Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, et al. 2024. Pixtral 12B. *arXiv preprint arXiv:2410.07073*.
- Riccardo Ali, Francesco Caso, Christopher Irwin, and Pietro Liò. 2025. Entropy-Lens: The information signature of transformer computations. *arXiv preprint arXiv:2502.16570*.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Samyadeep Basu, Martin Grayson, Cecily Morrison, Besmira Nushi, Soheil Feizi, and Daniela Mas-siceti. 2024. Understanding information storage and transfer in multi-modal large language models. In *NeurIPS*.
- Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*.
- Gabriela Ben Melech Stan, Estelle Aflalo, Raanan Yehezkel Rohekar, Anahita Bhiwandiwala, Shao-Yen Tseng, Matthew Lyle Olson, Yaniv Gurwicz, Chenfei Wu, Nan Duan, and Vasudev Lal. 2024. LVLm-Intrepret: An interpretability tool for large vision-language models. In *CVPR*.
- Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. 2024. Paligemma: A versatile 3B VLM for transfer. *arXiv preprint arXiv:2407.07726*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *NeurIPS*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024. InternVL: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*.
- Jang Hyun Cho, Andrea Madotto, Effrosyni Mavroudi, Triantafyllos Afouras, Tushar Nagarajan, Muhammad Maaz, Yale Song, Tengyu Ma, Shuming Hu, Suyog Jain, et al. 2025. PerceptionLM: Open-access data and models for detailed visual understanding. *arXiv preprint arXiv:2504.13180*.
- Bilal Chughtai, Lawrence Chan, and Neel Nanda. 2023. A toy model of universality: Reverse engineering how networks learn group operations. In *ICML*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does bert look at? an analysis of bert’s attention. In *ACL*.
- Saurabh Dash, Yiyang Nan, John Dang, Arash Ahmadian, Shivalika Singh, Madeline Smith, Bharat Venkitesh, Vlad Shmyhlo, Viraat Aryabumi, Walter Beller-Morales, et al. 2025. Aya vision: Advancing the frontier of multilingual multimodality. *arXiv preprint arXiv:2505.08751*.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. 2025. Molmo and pixMo: Open weights and open data for state-of-the-art multimodal models. In *CVPR*.
- Allyson Ettinger, Ahmed Elgohary, and Philip Resnik. 2016. Probing for semantic evidence of composition by means of simple classification tasks. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*.
- Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2018. VSE++: Improving visual-semantic embeddings with hard negatives. In *BMVC*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. 2024. MME: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- Qiyue Gao, Xinyu Pi, Kevin Liu, Junrong Chen, Ruolan Yang, Xinqi Huang, Xinyu Fang, Lu Sun, Gautham Kishore, Bo Ai, et al. 2025. Do vision-language models have internal world models? towards an atomic evaluation. In *Findings of ACL*.
- John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. In *EMNLP*.
- Jennifer Hu and Roger Levy. 2023. Prompting is not a substitute for probability measurements in large language models. In *EMNLP*.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*.
- Sonia Joseph, Praneet Suresh, Lorenz Hufe, Edward Stevinson, Robert Graham, Yash Vadi, Danilo Bzdok, Sebastian Lapuschkin, Lee Sharkey, and Blake Aaron Richards. 2025. Prisma: An open source toolkit for mechanistic interpretability in vision and video. *arXiv preprint arXiv:2504.19475*.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. 2023. Segment anything. In *ICCV*.

- Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*.
- Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen, Tony Lee, Erran Li Li, Ruohan Zhang, et al. 2024. Embodied agent interface: Benchmarking llms for embodied decision making. In *NeurIPS*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *ECCV*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *NeurIPS*.
- Neel Nanda and Joseph Bloom. 2022. Transformerlens. <https://github.com/TransformerLensOrg/TransformerLens>.
- Clement Neo, Luke Ong, Philip Torr, Mor Geva, David Krueger, and Fazl Barez. 2024. Towards interpreting visual information processing in vision-language models. *arXiv preprint arXiv:2410.07149*.
- Vedant Palit, Rohan Pandey, Aryaman Arora, and Paul Pu Liang. 2023. Towards vision-language mechanistic interpretability: A causal tracing tool for Blip. In *ICCV*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An imperative style, high-performance deep learning library. In *NeurIPS*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. 2024. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13009–13018.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *ICCV*.
- Samuel Stevens, Wei-Lun Chao, Tanya Berger-Wolf, and Yu Su. 2025. Sparse autoencoders for scientifically rigorous interpretation of vision models. *arXiv preprint arXiv:2502.06755*.
- J Ridley Stroop. 1935. Studies of interference in serial verbal reactions. *Journal of experimental psychology*, 18(6):643.
- OpenBMB MiniCPM-o Team. 2025. MiniCPM-o 2.6: A GPT-4o level MLLM for vision, speech, and multimodal live streaming on your phone.
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *CVPR*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024a. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Song XiXuan, et al. 2024b. CogVLM: Visual expert for pretrained language models. In *NeurIPS*.
- Yixuan Wang and Freda Shi. 2025. Logical forms complement probability in understanding language model (and human) performance. In *ACL*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *EMNLP*.
- Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. 2025. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12966–12977.
- Zhuo Xu, Xiang Xiang, and Yifan Liang. 2025. Overcoming shortcut problem in vlm for robust out-of-distribution detection. In *CVPR*.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. MiniCPM-V: A GPT-4V level MLLM on your phone. *arXiv preprint arXiv:2408.01800*.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. 2024. MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. In *CVPR*.

Kaichen Zhang, Yifei Shen, Bo Li, and Ziwei Liu. 2024. Large multi-modal models can interpret features in large multi-modal models. *arXiv preprint arXiv:2411.14982*.

Pan Zhang, Xiaoyi Dong, Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Haodong Duan, Songyang Zhang, Shuangrui Ding, et al. 2023. InternLM-XComposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*.

Zheyuan Zhang, Fengyuan Hu, Jayjun Lee, Freda Shi, Parisa Kordjamshidi, Joyce Chai, and Ziqiao Ma. 2025. Do vision-language models represent space and how? Evaluating spatial frame of reference under ambiguities. In *ICLR*.

A Contribution Statement

We list the contributions of each author below:

- Project idea: Freda Shi, Ziqiao Ma.
- Substantial prototype design: Eric Huang, Hala Sheta, Freda Shi.
- Prototype implementation: Eric Huang.
- Model implementation: Shuyu Wu, Jiajun Hong, Hala Sheta, Eric Huang, Ilia Alenabi, Ryker Lin, Ruoxi Ning, Daniel Wei, Jialin Yang.
- Documentation, user guide, and tutorials: Ziqiao Ma, Freda Shi, Ilia Alenabi, Ryker Lin, Daniel Wei.
- Substantial code review: Hala Sheta, Shuyu Wu, Eric Huang, Freda Shi.
- Probing experiments: Hala Sheta, Shuyu Wu, Ilia Alenabi, Jiajun Hong, Ryker Lin, Daniel Wei.
- Concept similarity experiments: Ziqiao Ma, Shuyu Wu, Freda Shi.
- Substantial project discussion: Hala Sheta, Eric Huang, Shuyu Wu, Jiajun Hong, Jiawei Zhou, Ziqiao Ma, Freda Shi.
- Paper writing: Hala Sheta, Freda Shi, Ziqiao Ma.
- Project supervision: Freda Shi, Jiawei Zhou, Ziqiao Ma.

B Probing Configuration

For each of our supported models, we extract the hidden representations of the middle (e.g., 16) and last (e.g., 32) of the post-attention layer norm module (or equivalent) using our library. This is because we require a 4096-dimensional (or equivalent) representation that can be efficiently stored and contains some relevant information about the input.

Each probe is instantiated with ReLU activation and a hidden size of 512, selected empirically to balance performance and efficiency. The input and output sizes are determined dynamically based on the tensors and labels in the input database. We tune the hyperparameters, `learning_rate`, `num_epochs` and `batch_size`, to find the best combinations that incurs the lowest mean validation loss after k -fold cross-validation with $k = 5$. With 3 options for each former parameter, we search through $3^3 = 27$ configurations. Lastly, using the best training configuration, we train two new models from scratch on the original and shuffled datasets, respectively.