
A Concept-Based Explainability Framework for Large Multimodal Models

Jayneel Parekh¹ Pegah Khayatan¹ Mustafa Shukor¹

Alasdair Newson¹ Matthieu Cord^{1,2}

¹ISIR, Sorbonne Université, Paris, France ²Valeo.ai, Paris, France

{jayneel.parekh, pegah.khayatan}@sorbonne-universite.fr

Abstract

Large multimodal models (LMMs) combine unimodal encoders and large language models (LLMs) to perform multimodal tasks. Despite recent advancements towards the interpretability of these models, understanding internal representations of LMMs remains largely a mystery. In this paper, we present a novel framework for the interpretation of LMMs. We propose a dictionary learning based approach, applied to the representation of tokens. The elements of the learned dictionary correspond to our proposed concepts. We show that these concepts are well semantically grounded in both vision and text. Thus we refer to these as “multimodal concepts”. We qualitatively and quantitatively evaluate the results of the learnt concepts. We show that the extracted multimodal concepts are useful to interpret representations of test samples. Finally, we evaluate the disentanglement between different concepts and the quality of grounding concepts visually and textually. Our implementation is publicly available.¹

1 Introduction

Vision transformers [9] and large-language models (LLMs) [6, 27, 37] have recently dominated machine learning research. Their effectiveness for unimodal processing tasks has spurred their use for multimodal tasks. In particular, visual encoders and LLMs are combined to address tasks such as image captioning and VQA [2, 20, 24, 36, 38]. Moreover, LLMs combined with unimodal encoders for other modalities have also been effective to solve respective multimodal tasks [10, 18, 25]. This recent class of models are referred to as large multimodal models (LMMs). For interpretability research, LMMs remain largely unexplored. This is because most prior methods focus on visual/tabular data, classification tasks, and non-transformer based architectures. Thus, despite their popularity, very few methods attempt to understand representations inside an LMM [28, 29, 34, 35].

This paper aims to bridge some of these differences and study in greater detail the intermediate representations of LMMs. To this end, motivated by the concept activation vector (CAV) based approaches for CNNs [12–14, 17], we propose a novel dictionary-learning based concept extraction method, designed for application to LMMs. Our method is used to learn a concept dictionary to understand the representations of a pretrained LMM for a given word/token of interest (Eg. ‘Dog’). For this token, we build a matrix containing the LMM’s internal representation of the token. We then linearly decompose this matrix using dictionary learning. The dictionary elements represent our concepts. The most interesting consequence of our method is that the learnt concepts exhibit a semantic structure that can be meaningfully grounded in both visual and textual domains. They are visually grounded by extracting the images which maximally activate these concepts. They can simultaneously be grounded in the textual domain by decoding the concept through the language model of the LMM and extracting the words/tokens they are most associated to. We refer to such

¹Project page and code: https://jayneelparekh.github.io/LMM_Concept_Explainability/

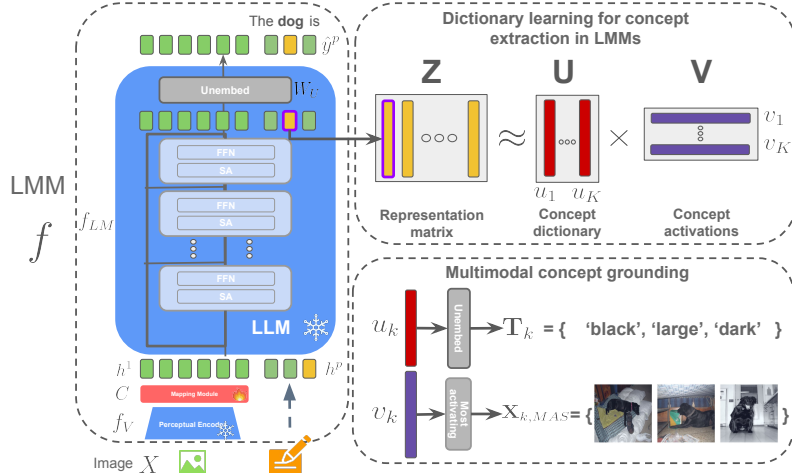


Figure 1: Overview of our multimodal concept extraction and grounding. Given a pretrained LMM for captioning and a target token t (for eg. ‘Dog’), our method extracts internal representations of f about t , across many images. These representations are collated into a matrix \mathbf{Z} . We linearly decompose \mathbf{Z} to learn a concept dictionary \mathbf{U} and its coefficients/activations \mathbf{V} . Each concept $u_k \in \mathbf{U}$, is multimodally grounded in both visual and textual domains.

concept representations as *multimodal concepts*. Our key contributions can be summarized as:

- We propose a novel concept-based explainability framework that can be used to understand internal representations of large multimodal models. To the best of our knowledge, this is the first effort targeting multimodal models at this scale.
- Our dictionary learning based concept extraction approach extracts a multimodal concept dictionary wherein each concept can be semantically grounded simultaneously in both text and vision. We also extend the previous concept dictionary-learning strategies using a Semi-NMF based optimization.
- We experimentally validate the notion of multimodal concepts through both, qualitatively and quantitatively. Our learnt concept dictionary is shown to possess a meaningful multimodal grounding covering diverse concepts, and is useful to interpret LMM representations of test samples.

2 Related work

Concept activation vector based approaches Concept based interpretability aim to extract the semantic content relevant for a model [7]. For post-hoc interpretation, concept activation vector (CAV) based approaches [13, 14, 17, 39] have been most widely used. Recently, Fel et al. [12] proposed a unified view of CAV-based approaches as variants of a dictionary learning problem. However, these methods have only been applied for interpretation of CNNs on classification tasks. LMMs on the contrary exhibit a significantly different architecture. We propose a dictionary learning based concept extraction method, designed for LMMs. We also propose a Semi-NMF variant of the dictionary learning problem, which has not been previously considered for concept extraction.

Understanding LMM representations Palit et al. [28] extend the causal tracing used for LLMs to analyze information across layers in an LMM. Schwettmann et al. [34] proposed the notion of *multimodal neurons* existing within the LLM part of an LMM. They term the neurons “multimodal” as they translate high-level visual information to corresponding information in text modality. The neurons are discovered by ranking them by a gradient based attribution score. Pan et al. [29] proposed a more refined algorithm to identify such neurons based on an importance measure leveraging architecture of transformer MLP blocks. In contrast, we propose to discover a concept structure in the token representations by learning a small dictionary of multimodally grounded concepts. Limiting the analysis to a specific target token allows our learnt concepts to discover fine details about the token.

3 Approach

Background for LMMs We consider a general model architecture for a large multimodal model f [1, 4, 20, 21], that consists of: a visual encoder f_V , a light-weight trainable connector module C , and an LLM f_{LM} consisting of N_L layers. We assume f is pretrained for captioning task with an

underlying dataset $\mathcal{S} = \{(X_i, y_i)\}_{i=1}^N$ consisting of images $X_i \in \mathcal{X}$ and their associated caption $y_i \in \mathcal{Y}$. \mathcal{X} and \mathcal{Y} denote the space of images and set of text tokens respectively. Note that caption y_i can be viewed as a subset of all tokens. The input to the language model f_{LM} is denoted by the sequence of tokens h^1, h^2, \dots, h^p and the output as \hat{y} . The internal representation of any token at some layer l and position p inside f_{LM} is denoted as $h_{(l)}^p$, with $h_{(0)}^p = h^p$. Note that $h_{(l)}^p$ is same as the residual stream representation in LLM transformers [11] at position p and layer l . For the multimodal model, the input token sequence for f_{LM} consists of concatenated: (1) N_V visual tokens provided by the visual encoder f_V operating on an image X , followed by the connector C , and (2) linearly embedded textual tokens previously predicted by f_{LM} . For $p > N_V$, this is expressed as:

$$\hat{y}^p = f_{LM}(h^1, h^2, \dots, h^{N_V}, \dots, h^p), \quad (1)$$

where $h^1, \dots, h^{N_V} = C(f_V(X))$, and $h^p = \text{Emb}(\hat{y}^{p-1})$ for $p > N_V$, where Emb denotes the token embedding function. To start the prediction, h^{N_V+1} is defined as the beginning of sentence token. The output token \hat{y}^p is obtained by normalizing $h_{(N_V)}^p$, followed by an unembedding layer that applies a matrix W_U followed by a softmax. The predicted caption \hat{y} consists of the predicted tokens $\hat{y} = \{\hat{y}^p\}_{p>N_V}$ until the end of sentence token. The model is trained with next token prediction objective, to generate text conditioned on images in an auto-regressive fashion.

3.1 Method overview

Given a pretrained LMM f and token of interest $t \in \mathcal{Y}$, our method consists of three parts (Fig. 1):

1. Selecting a subset of images \mathbf{X} from dataset \mathcal{S} , relevant for target token t . We extract representations by processing samples in \mathbf{X} through f . The extracted representations of dimension B are collected in a matrix $\mathbf{Z} \in \mathbb{R}^{B \times M}$, where M is number of samples in \mathbf{X} .
2. Linearly decomposing $\mathbf{Z} \approx \mathbf{U}\mathbf{V}$ into its constituents, that includes a dictionary of learnt concepts $\mathbf{U} \in \mathbb{R}^{B \times K}$ of size K and coefficient/activation matrix $\mathbf{V} \in \mathbb{R}^{K \times M}$.
3. Semantically grounding the learnt ‘‘multimodal concepts’’, contained in dictionary \mathbf{U} in both visual and textual modalities.

We emphasize here that our main objective in employing dictionary learning based concept extraction is to understand internal representations of an LMM, and not to interpret the output of the model, which can be readily accomplished by combining this pipeline with some concept importance estimation method [12]. The rest of the section elaborates each of the above three steps.

3.2 Representation extraction

To extract relevant representations about t from f , we first determine a set of samples \mathbf{X} from dataset $\mathcal{S} = \{(X_i, y_i)\}_{i=1}^N$ for extraction. The captioning task requires the model to capture relevant information about the given image and predict it in the output without any additional context. We consider the set of samples where t is predicted as part of the predicted caption \hat{y} . This allows us to further investigate the internal representations of t . To enhance visual interpretability for concepts, we limit this set to samples containing t in the ground-truth caption. Thus, \mathbf{X} is computed as:

$$\mathbf{X} = \{X_i \mid t \in f(X_i) \text{ and } t \in y_i \text{ and } (X_i, y_i) \in \mathcal{S}\}. \quad (2)$$

Given any $X \in \mathbf{X}$, we extract the internal token representation $h_{(L)}^p$ from a deep layer L , at the first position in the predicted caption $p > N_V$, such that $\hat{y}^p = t$. The representation $z_j \in \mathbb{R}^B$ of each sample $X_j \in \mathbf{X}$ is stacked as columns of the matrix $\mathbf{Z} = [z_1, \dots, z_M] \in \mathbb{R}^{B \times M}$. Note that text token representations in f_{LM} can possess a meaningful multimodal structure as they combine information from visual token representations $h_{(l)}^p, p \leq N_V$. We focus on residual stream representation $h_{(L)}^p$ [11] as it represents the aggregated information computed by all prior layers of the LMM till layer L .

3.3 Decomposing the representations

The representation matrix $\mathbf{Z} \approx \mathbf{U}\mathbf{V}$, is decomposed as product of two low-rank matrices $\mathbf{U} \in \mathbb{R}^{B \times K}$, $\mathbf{V} \in \mathbb{R}^{K \times M}$ of rank $K \ll \min(B, M)$, where K denotes the number of dictionary elements. The columns of $\mathbf{U} = [u_1, \dots, u_K]$ are the basis vectors which we refer to as concept-vectors/concepts. The rows of \mathbf{V} or columns of $\mathbf{V}^T = [v_1, \dots, v_M], v_i \in \mathbb{R}^M$ denote the activations



Figure 2: Example of multimodal concept grounding in vision and text. Five most activating samples (among decomposed in \mathbf{Z}) and five most probable decoded words are shown.

of u_i for each sample. This decomposition, as previously studied in [12], can be optimized with various constraints on \mathbf{U} , \mathbf{V} , each leading to a different dictionary (eg. PCA, K-Means, NMF). NMF is considered to yield most interpretable results. However, for our use case, NMF is not useful as representation matrix \mathbf{Z} is not non-negative. Instead, we propose to employ a relaxed version of NMF, Semi-NMF [8], which allows the basis vectors \mathbf{U} to contain mixed values, but forces activations \mathbf{V} to be non-negative. Since we expect only a small number of concepts to be present in any sample, we also encourage sparsity in activations \mathbf{V} . The Semi-NMF optimization problem to decompose \mathbf{Z} is:

$$\mathbf{U}^*, \mathbf{V}^* = \arg \min_{\mathbf{U}, \mathbf{V}} \|\mathbf{Z} - \mathbf{U}\mathbf{V}\|_F^2 + \lambda \|\mathbf{V}\|_1 \quad s.t. \quad \mathbf{V} \geq 0, \text{ and } \|u_k\|_2 \leq 1 \quad \forall k \in \{1, \dots, K\}. \quad (3)$$

Given any image X where token t is predicted by f , we can now define the process of computing activations of concept dictionary \mathbf{U}^* for given X , denoted as $v(X) \in \mathbb{R}^K$. To do so, we first extract the token representation for X , $z_X \in \mathbb{R}^B$ with the process described in Sec. 3.2. Then, z_X is projected on \mathbf{U}^* to compute $v(X)$. In the case of Semi-NMF, this corresponds to $v(X) = \arg \min_{v \geq 0} \|z_X - \mathbf{U}^*v\|_2^2 + \lambda \|v\|_1$. The activation of $u_k \in \mathbf{U}^*$ is denoted as $v_k(X) \in \mathbb{R}$.

3.4 Using the concept dictionary for interpretation

Multimodal grounding of concepts. We now elaborate the process to ground the understanding of any given concept vector $u_k, k \in \{1, \dots, K\}$ in the visual and textual domains. For visual grounding, we use prototyping [3, 17] to select input images (among the decomposed samples), that maximally activate u_k . Given the number of samples extracted for visualization N_{MAS} , the set of maximum activating samples (MAS) for component u_k , denoted as $\mathbf{X}_{k, MAS}$ can be specified as $\mathbf{X}_{k, MAS} = \arg \max_{\hat{X} \subset \mathbf{X}, |\hat{X}|=N_{MAS}} \sum_{X \in \hat{X}} |v_k(X)|$.

For grounding in textual domain, we leverage insights from ‘‘Lens’’ based methods [5, 19, 26, 33] developed for LLMs. In particular, following [26], we use the unembedding layer to decode u_k to the token vocabulary space \mathcal{Y} , and extract the most probable tokens. That is, we extract the tokens with highest probability in $W_U u_k$. They are then filtered for being an english, non-stop-word with at least 3 characters, to eliminate unnecessary tokens. The final set of words is referred to as grounded words for concept u_k and denoted as \mathbf{T}_k . Fig. 2 illustrates an example for a concept for token ‘‘Dog’’.

Most activating concepts for images. To understand the LMM’s representation of a given image X , we now define the *most activating concepts*. First, we extract the representations z_X of the image with the same process as described previously. We then project z_X on \mathbf{U}^* to obtain $v(X) \in \mathbb{R}^K$. We define the most activating concepts, denoted $\tilde{u}(X)$, as the set of r concepts whose activations $v_k(X)$ have the largest magnitude. One can then visualize the multimodal grounding of $\tilde{u}(X)$.

4 Experiments

Models and dictionary learning. We experiment with DePALM model [38] trained for captioning task on COCO dataset [22]. The model consists of a frozen ViT-L/14 CLIP [32] visual encoder as f_V . The language model f_{LM} is a frozen OPT-6.7B [40] of 32 layers. For uniformity and fairness, results are reported with $K = 20$ and representations from $L = 31$, the final layer before unembedding layer. For Semi-NMF, we set $\lambda = 1$. Further implementation details are available in Appendix D. We conduct our analysis separately for various common objects in the dataset: ‘‘Dog’’, ‘‘Bus’’, ‘‘Train’’, ‘‘Cat’’, ‘‘Bear’’, ‘‘Baby’’, ‘‘Car’’, ‘‘Cake’’. The extension to other classes remains straightforward and is discussed in Appendix E. More experiments on a popular LMM, LLaVA [23] are in Appendix B.1.

Evaluation axes We evaluate the learnt concept dictionary \mathbf{U}^* on three axes: (i) Its use during inference to interpret representations of LMMs for test samples, (ii) The overlap/entanglement between grounded words of concepts $u_k \in \mathbf{U}^*$ and (iii) the quality of visual and text grounding of concepts. Main paper covers (i) and (ii). Evaluation of multimodal grounding is in Appendix A

Concept extraction during inference: To evaluate the use of \mathbf{U}^* in understanding any test sample X , we first estimate the top r most activating concepts activations, $\tilde{u}(X)$ (Sec. 3.4). We then estimate the correspondence between the image X and the grounded words \mathbf{T}_k of $u_k \in \tilde{u}(X)$. This correspondence is estimated via two different metrics. The average CLIPScore [16] between image X and words \mathbf{T}_k , and the average BERTScore (F1) [41] between the ground-truth captions y associated with X and \mathbf{T}_k . These metrics help validate the multimodal nature of the concept dictionaries. Their use is inspired from [34]. The implementation details for the metrics are in Appendix D.

Overlap/entanglement of learnt concepts: Ideally, we expect each concept in \mathbf{U}^* to encode distinct information about the token of interest t . Thus two different concepts $u_k, u_l, k \neq l$ should be associated to different sets of words. To quantify the entanglement of a learnt concept u_k , we compute the overlap between the grounded words $\mathbf{T}_k, \mathbf{T}_l, l \neq k$. The overlap metric for a dictionary \mathbf{U}^* is defined as the average of overlap of each concept.

$$\text{Overlap}(\mathbf{U}^*) = \frac{1}{K} \sum_k \text{Overlap}(u_k), \quad \text{Overlap}(u_k) = \frac{1}{(K-1)} \sum_{l=1, l \neq k}^K \frac{|\mathbf{T}_l \cap \mathbf{T}_k|}{|\mathbf{T}_k|}$$

Baselines: One set of methods for evaluation are the variants of our proposed approach where we employ different dictionary learning strategies: PCA, KMeans and Semi-NMF. For evaluating concept extraction on test data with CLIPScore/BERTScore we compare against the following baselines:

- *Rnd-Words:* For each concept u_k (from Semi-NMF), this baseline replaces the grounded words \mathbf{T}_k by a set of random words \mathbf{R}_k such that $|\mathbf{R}_k| = |\mathbf{T}_k|$. Words in \mathbf{R}_k satisfy the same conditions as grounded words i.e. english non-stopwords with atleast three characters. We do this by decoding randomly sampled token representation and adding the top decoded words if they satisfy the conditions.
- *Noise-Imgs:* This baseline uses random noise as images to extract representations and then proceeds with exactly same learning procedure as Semi-NMF. Combined with the Rnd-Words baseline, they ablate two different parts of the concept extraction pipeline.
- *Simple:* This baseline considers a simple technique to build the dictionary \mathbf{U}^* and projecting test samples. It builds \mathbf{U}^* by selecting token representations in \mathbf{Z} with the largest norm. The projections are performed by mapping the test sample representation to the closest element in \mathbf{U}^* . For deeper layers, this provides a strong baseline in terms of extracted grounded words \mathbf{T}_k which are related to token of interest t , as they are obtained by directly decoding token representations of t .

We also report score using ground-truth captions (*GT captions*) instead of grounded words \mathbf{T}_k , to get the best possible correspondence score. The overlap/entanglement in concept dictionary is compared between the non-random baselines: Simple, PCA, K-Means and Semi-NMF.

4.1 Results and discussion

Token	Metric	Rnd-Words	Noise-Imgs	Simple	Semi-NMF (Ours)	GT-captions
Dog	CS top-1 (\uparrow)	0.519 \pm 0.05	0.425 \pm 0.06	<u>0.546 \pm 0.08</u>	0.610 \pm 0.09	0.783 \pm 0.06
	BS top-1 (\uparrow)	0.201 \pm 0.04	0.306 \pm 0.05	<u>0.346 \pm 0.08</u>	0.405 \pm 0.07	0.511 \pm 0.11
Bus	CS top-1 (\uparrow)	0.507 \pm 0.05	0.425 \pm 0.08	0.667 \pm 0.06	<u>0.634 \pm 0.08</u>	0.736 \pm 0.05
	BS top-1 (\uparrow)	0.200 \pm 0.05	0.303 \pm 0.06	<u>0.390 \pm 0.05</u>	0.404 \pm 0.07	0.466 \pm 0.11
Train	CS top-1 (\uparrow)	0.496 \pm 0.05	0.410 \pm 0.07	<u>0.642 \pm 0.06</u>	0.646 \pm 0.07	0.727 \pm 0.05
	BS top-1 (\uparrow)	0.210 \pm 0.06	0.253 \pm 0.06	0.392 \pm 0.07	<u>0.378 \pm 0.07</u>	0.436 \pm 0.08
Cat	CS top-1 (\uparrow)	0.539 \pm 0.04	0.461 \pm 0.04	<u>0.589 \pm 0.07</u>	0.627 \pm 0.06	0.798 \pm 0.05
	BS top-1 (\uparrow)	0.207 \pm 0.07	0.307 \pm 0.03	<u>0.425 \pm 0.10</u>	0.437 \pm 0.08	0.544 \pm 0.10

Table 1: Test data mean CLIPScore and BERTScore for top-1 activating concept for all baselines on five tokens. CLIPScore denoted as CS, and BERTScore as BS. Statistical significance is in Appendix E. Higher scores are better. Best score in **bold**, second best is underlined.

Quantitative results Tab. 1 reports the test top-1 CLIPScore/BERTScore for all baselines and Semi-NMF on different target tokens. Appendix E contains detailed results for other tokens as well as for the PCA and K-Means variants. We report the results for only the top-1 activating concept, as KMeans and Simple baselines map a given representation to a single concept. Notably, Semi-NMF generally outperforms the other baselines although the Simple baseline performs competitively. More generally, Semi-NMF, K-Means, and Simple tend to clearly outperform other baselines on these metrics, indicating that these systems project representations of test images to concepts whose associated grounded words correspond well with the visual content.

Tab. 2 reports the overlap between concepts for Simple, PCA, KMeans and Semi-NMF systems. Interestingly, KMeans and Simple baseline perform significantly worse than Semi-NMF/PCA with a high overlap between grounded words, often exceeding 40%. PCA outperforms other methods with almost no overlap while Semi-NMF shows some overlap. Overall, Semi-NMF strikes the best balance among all the methods, in terms of learning a concept dictionary useful for understanding test image representations, but which also learns diverse and disentangled concepts.

Token	Simple	PCA	KMeans	Semi-NMF
Dog	0.371	0.004	0.501	<u>0.086</u>
Bus	0.622	0.002	0.487	<u>0.177</u>
Train	0.619	0.015	0.367	<u>0.107</u>
Cat	0.452	0.000	0.500	<u>0.146</u>

Table 2: Overlap between learnt concepts. Lower is better. Best score in **bold**, second best underlined.

Qualitative results Fig. 3 shows visual and textual grounding of concepts extracted for token ‘dog’. For brevity, we select 8 out of 20 concepts for illustration. 2. Grounding for all concepts extracted for ‘dog’ and other tokens are in Appendix F. Extracted concepts capture information about different aspects of a ‘dog’. The LMM separates

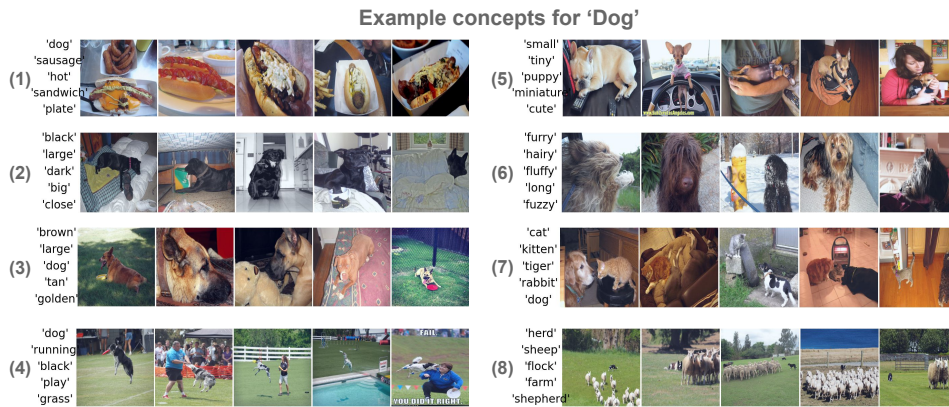


Figure 3: Visual/textual grounding for 8 out of 20 concepts for ‘Dog’ token (layer 31). For each concept we illustrate the set of 5 most activating samples and 5 most probable decoded words.

representation of animal ‘Dog’ with a ‘hot dog’ (Concept 1). Specifically for ‘Dog’, Concepts (2), (3) capture information about color: ‘black’, ‘brown’. Concept (6) about ‘long haired’ dogs, while concept (5) activates for ‘small/puppy’ dogs. Beyond concepts activating for specific characteristics, we also discover concepts describing their state of actions (Concept (4) ‘playing/running’), common scenes they occur in (Concept (8), ‘herd’), and correlated objects they can occur with (Concept (7), ‘cat and dog’). We observe this diverse nature of extracted concepts for other tokens too. The information about concepts can be inferred via both the visual images and the associated grounded words, validating the coherent multimodal grounding. Notably, compared to solely visual grounding as for CAVs for CNNs, the multimodal grounding eases the process to understand a concept.

Additional experiments. Qualitative local interpretations for test samples are in Appendix F. Experiment analyzing concept extraction across layers is available in Appendix G. Our method can be also be applied to understand of visual/perceptual token representations inside the LMM which also exhibit this multimodal structure (Appendix H).

5 Conclusion

In summary, we presented a novel dictionary learning based concept extraction framework, to understand internal representations of a LMM. The approach relies on decomposing representations of a token inside a pretrained LMM. To this end, we also propose a Semi-NMF variant of the concept dictionary learning problem. The elements of the learnt concept dictionary are grounded in the both text and visual domains, leading to a novel notion of *multimodal concepts* in the context of interpretability. We quantitatively and qualitatively show that (i) the multimodal grounding of concepts is meaningful, and (ii) the learnt concepts are useful to understand representations of test samples. We hope that our method inspires future work on concept methods to understand LMMs.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:23716–23736, 2022.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35: 23716–23736, 2022.
- [3] David Alvarez-Melis and Tommi Jaakkola. Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 7775–7784, 2018.
- [4] Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. OpenFlamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*, 2023.
- [5] Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*, 2023.
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 1877–1901, 2020.
- [7] Julien Colin, Thomas Fel, Rémi Cadène, and Thomas Serre. What i cannot predict, i do not understand: A human-centered evaluation framework for explainability methods. *Advances in neural information processing systems*, 35:2832–2845, 2022.
- [8] Chris HQ Ding, Tao Li, and Michael I Jordan. Convex and semi-nonnegative matrix factorizations. *IEEE transactions on pattern analysis and machine intelligence*, 32(1):45–55, 2008.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [10] Constantin Eichenberg, Sidney Black, Samuel Weinbach, Letitia Parcalabescu, and Anette Frank. Magma—multimodal augmentation of generative models through adapter-based finetuning. *arXiv preprint arXiv:2112.05253*, 2021.
- [11] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1, 2021.
- [12] Thomas Fel, Victor Boutin, Louis Béthune, Rémi Cadène, Mazda Moayeri, Léo Andéol, Mathieu Chalvidal, and Thomas Serre. A holistic approach to unifying automatic concept extraction and concept importance estimation. *Advances in Neural Information Processing Systems*, 36, 2023.
- [13] Thomas Fel, Agustin Picard, Louis Bethune, Thibaut Boissin, David Vigouroux, Julien Colin, Rémi Cadène, and Thomas Serre. Craft: Concept recursive activation factorization for explainability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2711–2721, 2023.
- [14] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. Towards automatic concept-based explanations. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9277–9286, 2019.

- [15] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
- [16] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021.
- [17] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, pages 2668–2677. PMLR, 2018.
- [18] Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal generation. *arXiv preprint arXiv:2301.13823*, 2023.
- [19] Anna Langedijk, Hosein Mohebbi, Gabriele Sarti, Willem H. Zuidema, and Jaap Jumelet. Decoderlens: Layerwise interpretation of encoder-decoder transformers. *CoRR*, abs/2310.03686, 2023.
- [20] Hugo Laurençon, Lucile Saulnier, Leo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. OBELICS: An open web-scale filtered dataset of interleaved image-text documents. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=SKN2hf1BIZ>.
- [21] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014.
- [23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2023.
- [24] Oscar Mañas, Pau Rodriguez, Saba Ahmadi, Aida Nematzadeh, Yash Goyal, and Aishwarya Agrawal. Mapl: Parameter-efficient adaptation of unimodal pre-trained models for vision-language few-shot prompting. *arXiv preprint arXiv:2210.07179*, 2022.
- [25] Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly mapping from image to text space. *arXiv preprint arXiv:2209.15162*, 2022.
- [26] Nostalgebraist. Interpreting gpt: The logit lens. <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>, 2020. Accessed: [date of access].
- [27] OpenAI. Gpt-4 technical report. *arXiv*, 2023.
- [28] Vedant Palit, Rohan Pandey, Aryaman Arora, and Paul Pu Liang. Towards vision-language mechanistic interpretability: A causal tracing tool for blip. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2856–2861, 2023.
- [29] Haowen Pan, Yixin Cao, Xiaozhi Wang, and Xun Yang. Finding and editing multi-modal neurons in pre-trained transformer. *arXiv preprint arXiv:2311.07470*, 2023.
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- [31] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.

- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763. PMLR, 2021.
- [33] Mansi Sakarvadia, Arham Khan, Aswathy Ajith, Daniel Grzenda, Nathaniel Hudson, André Bauer, Kyle Chard, and Ian T. Foster. Attention lens: A tool for mechanistically interpreting the attention head information retrieval mechanism. *CoRR*, abs/2310.16270, 2023.
- [34] Sarah Schwettmann, Neil Chowdhury, Samuel Klein, David Bau, and Antonio Torralba. Multi-modal neurons in pretrained text-only transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2862–2867, 2023.
- [35] Mustafa Shukor and Matthieu Cord. Implicit multimodal alignment: On the generalization of frozen llms to multimodal inputs. *arXiv preprint arXiv:2405.16700*, 2024.
- [36] Mustafa Shukor, Corentin Dancette, and Matthieu Cord. ep-alm: Efficient perceptual augmentation of language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22056–22069, October 2023.
- [37] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [38] Théophane Vallaëys, Mustafa Shukor, Matthieu Cord, and Jakob Verbeek. Improved baselines for data-efficient perceptual augmentation of llms. *arXiv preprint arXiv:2403.13499*, 2024.
- [39] Johanna Vielhaben, Stefan Bluecher, and Nils Strodthoff. Multi-dimensional concept discovery (mcd): A unifying framework with completeness guarantees. *Transactions on Machine Learning Research*, 2023.
- [40] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [41] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2019.




Figure 4: Evaluating visual/text grounding (CLIPScore/BERTScore). Each point denotes score for grounded words of a concept (Semi-NMF) vs Rnd-Words w.r.t the same visual grounding.

A Evaluating multimodal grounding of concepts

To evaluate the quality of visual/text grounding of concepts $(\mathbf{X}_{k,MAS}, \mathbf{T}_k)$, we measure the correspondence between visual and text grounding of a given concept u_k , i.e. the set of maximum activating samples $\mathbf{X}_{k,MAS}$ and words \mathbf{T}_k , using CLIPScore and BERTScore as described above. Fig. 4 shows an evaluation of visual/text grounding of concepts learnt by Semi-NMF. Each point on the figure denotes the CLIPScore (left) or BERTScore (right) for correspondence between samples $\mathbf{X}_{k,MAS}$ and words \mathbf{T}_k for concept u_k against random words baseline. We see that for both metrics, vast majority of concepts lie above the $y = x$ line, indicating that grounded words correspond much better to content of maximum activating samples than random words.

B Experiments on other LMMs

This section covers our experiments on other types of multimodal models. First, we test our approach on LLaVA to demonstrate that our approach generalizes to more recent networks that also fine-tune f_{LM} on multimodal data. We also test our method on other variants of DePALM with non-CLIP visual encoders to observe their effect on CLIPScore.

B.1 Experiments with LLaVA

We conduct further experiments on LLaVA [23], a popular open-source LMM to demonstrate the generality of our method. The model uses a CLIP-ViT-L-336px visual encoder (f_V), a 2-layer linear connector (C) that outputs $N_V = 576$ visual tokens, and a Vicuna-7B language model (f_{LM} , 32 layers). We use identical hyperparameters as for DePALM ($K = 20, \lambda = 1, L = 31$). We report the test CLIPScore for top-1 activating concept, for Rnd-Words, Noise-Imgs, Simple and Semi-NMF, GT-captions in Tab. 3, and Overlap score for non-random baselines in Tab. 4. Quantitatively, we obtain consistent results to those observed for DePALM. Semi-NMF extracts most balanced concept dictionaries with high multimodal correspondence and low overlap. Qualitatively too, the method functions consistently and is able to extract concepts with meaningful multimodal grounding.

Qualitative results and Saliency maps We also show qualitative examples of concepts extracted for token ‘Dog’ in Fig. 5. More examples for other ‘Cat’ and ‘Train’ tokens are given in Fig. 7 and 8. Interestingly, since LLaVA uses a connector C that contains two linear layers, the visual tokens as processed inside f_{LM} preserve the notion of specific image patch representations, i.e. $N_V = 576$ tokens denoting representations for 576 (24×24) input patches. This allows us to further explore a simple and computationally cheap strategy of generating saliency maps to highlight which regions a concept vector activates on. To do this one can simply compute the inner product of any given concept vector u_k with all visual token representations from corresponding layer L , i.e. $u_k^T [h_{(L)}^1, \dots, h_{(L)}^{N_V}]$. This can be upscaled to the input image size to visualize the saliency map. We illustrate some qualitative outputs on concepts from ‘Dog’ in Fig. 6. .

Token	Metric	Rnd-Words	Noise-Imgs	Simple	Semi-NMF (Ours)	GT-captions
Dog	CS top-1 (↑)	0.537 ± 0.03	0.530 ± 0.05	<u>0.567 ± 0.08</u>	0.595 ± 0.07	0.777 ± 0.06
	BS top-1 (↑)	0.205 ± 0.07	0.227 ± 0.06	0.331 ± 0.07	<u>0.305 ± 0.07</u>	0.519 ± 0.11
Bus	CS top-1 (↑)	0.509 ± 0.04	0.487 ± 0.05	0.619 ± 0.06	<u>0.591 ± 0.08</u>	0.742 ± 0.05
	BS top-1 (↑)	0.198 ± 0.07	0.253 ± 0.06	0.319 ± 0.04	<u>0.306 ± 0.06</u>	0.460 ± 0.10
Train	CS top-1 (↑)	0.518 ± 0.03	0.505 ± 0.04	<u>0.633 ± 0.05</u>	0.640 ± 0.07	0.725 ± 0.05
	BS top-1 (↑)	0.177 ± 0.07	0.221 ± 0.04	0.310 ± 0.05	<u>0.293 ± 0.05</u>	0.432 ± 0.08
Cat	CS top-1 (↑)	0.536 ± 0.03	0.545 ± 0.04	0.625 ± 0.06	<u>0.621 ± 0.07</u>	0.795 ± 0.05
	BS top-1 (↑)	0.142 ± 0.06	0.235 ± 0.05	<u>0.306 ± 0.06</u>	0.329 ± 0.07	0.540 ± 0.11

Table 3: Concept extraction on LLaVA-v1.5: Test data mean CLIPScore reported for top-1 activating concept for same baselines and tokens as in main paper table 1. Higher scores are better. Best score in **bold**, second best is underlined.

Token	Simple	PCA	KMeans	Semi-NMF
Dog	0.435	0.008	0.429	<u>0.149</u>
Bus	0.464	0.010	0.518	<u>0.124</u>
Train	0.315	0.024	0.382	<u>0.087</u>
Cat	0.479	0.013	0.554	<u>0.166</u>

Table 4: Overlap evaluation (LLaVA). Lower is better. Best score in **bold**, second best underlined.

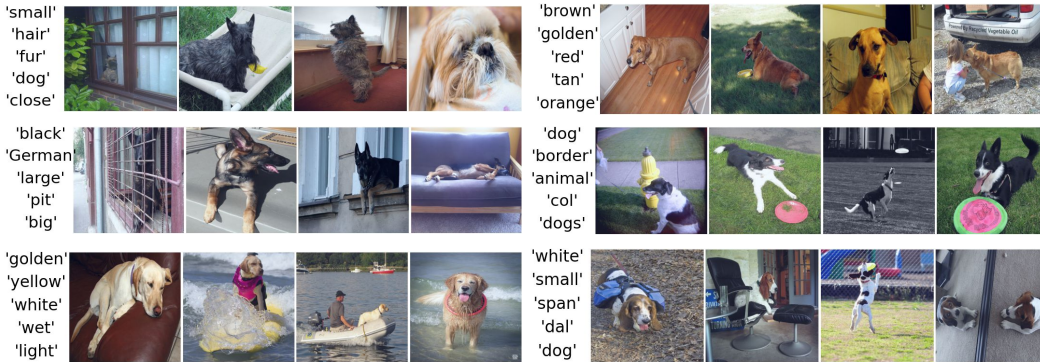


Figure 5: Multimodal grounding for example concepts for 'Dog' token (layer 31) on LLaVA.



Figure 6: Examples of generating visual concept saliency maps for two 'Dog' concepts for LLaVA. Red denotes high activations, blue denotes low activation (bottom row)

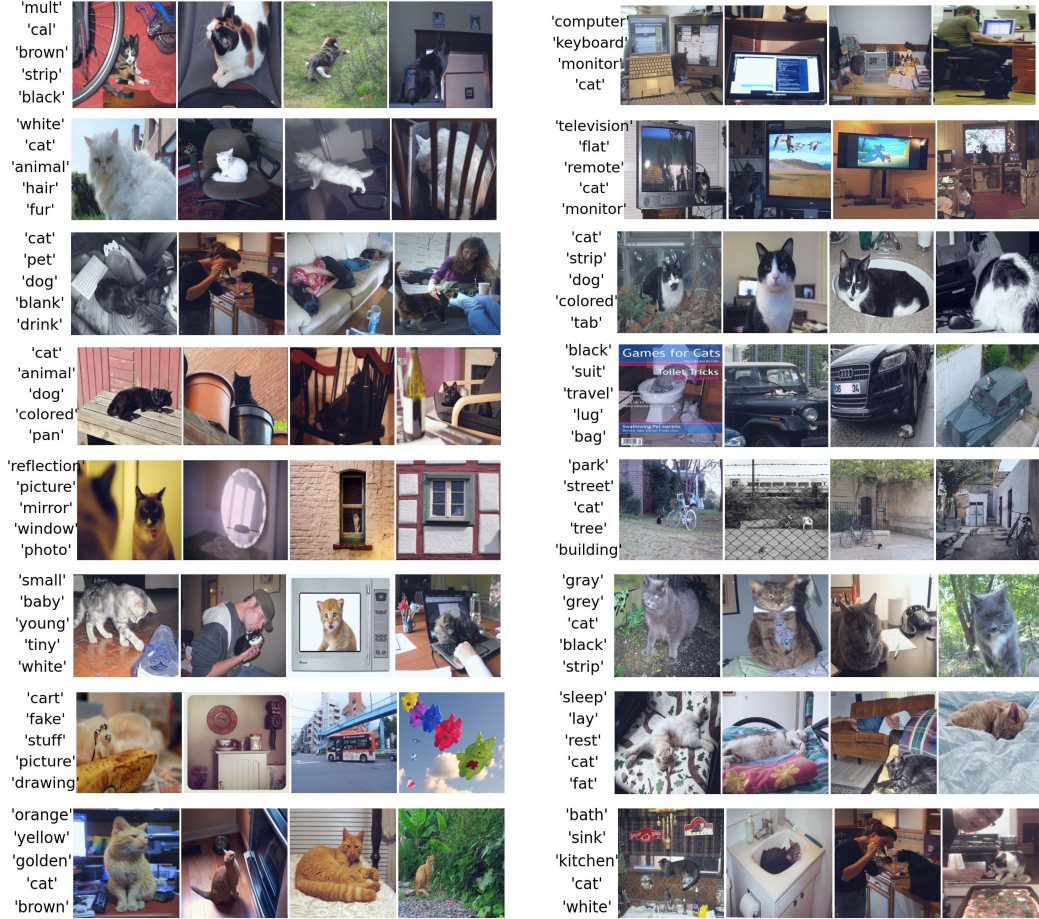


Figure 7: Multimodal grounding for example concepts for 'Cat' token (layer 31) on LLaVA.

B.2 DePALM with ViT visual encoders

We further test LMMs which do not contain a CLIP visual encoder to confirm that the high CLIPScore is not due to use of CLIP visual encoders. To test this, we experiment on two different DePALM models with frozen visual encoders different from CLIP, a frozen ViT-L encoder trained on ImageNet [9] and another frozen ViT-L trained as a masked autoencoder (MAE) [15]. Both LMMs use the same pretrained OPT-6.7B language model. Collectively, the three encoders (including CLIP) are pretrained for three different types of objectives. We use Semi-NMF to extract concept dictionaries, with all hyperparameters identical. The results are reported in tables below. 'Rnd-Words' and 'GT-captions' references are reported for each LMM separately, although they are very close to the ones in main paper. The "ViT-L (CLIP)" baseline denotes our system from the main paper that uses CLIP encoder. Importantly, we still obtain similar test CLIPScores as with CLIP visual encoder. The concept dictionaries still possess meaningful multimodal grounding. Many concepts also tend to be similar as for CLIP visual encoder, further indicating that processing inside language model plays a major role in the discovery of multimodally grounded concepts.

C Analyzing polysemanticity in the learnt concepts

We conducted a preliminary qualitative study on some concept vectors in the dictionary learnt for token "Dog" (DePALM model), to analyze if these concept vectors tend to activate strongly for a specific semantic concept (monosemantic) or multiple semantic concepts (polysemantic). In particular, we first manually annotated the 160 test samples for "Dog" for four semantic concepts, for which we knew we had concept vectors in our dictionary, namely "Hot dog" (Concept 2, row 1, column 2 in Fig. 7), "Black dog" (Concept 20, row 10, column 2 in Fig. 7), "Brown/orange dog"




Figure 8: Multimodal grounding for example concepts for 'Train' token (layer 31) on LLaVA.

Token	Rnd-Words	ViT-L (ImageNet)	ViT-L (CLIP)	GT-captions
Dog	0.514 ± 0.05	0.611 ± 0.09	0.610 ± 0.09	0.783 ± 0.06
Bus	0.498 ± 0.05	0.644 ± 0.07	0.634 ± 0.08	0.739 ± 0.05
Train	0.494 ± 0.05	0.617 ± 0.07	0.646 ± 0.07	0.728 ± 0.05
Cat	0.539 ± 0.05	0.628 ± 0.07	0.627 ± 0.06	0.794 ± 0.06

Table 5: Test CLIPScore evaluation for DePALM with ViT-L frozen image encoder trained on ImageNet: Scores reported for top-1 activating concept of Semi-NMF for Rnd-Words, GT-captions and ViT-L (CLIP) which denotes the system in main text.

(Concept 6, row 3, column 2 in Fig. 7), and "Bull dog" (Concept 15, row 8, column 1 in Fig. 7). For a given semantic concept, we call this set C_{true} . Then, for its corresponding concept vector u_k we find the set of test samples for which u_k activates greater than a threshold τ . This threshold was set to half of its maximum activation over test samples. We call this set of samples C_{top} . To estimate specificity of the concept vector we compute how many samples in C_{top} lie in the ground-truth set, i.e. $|C_{top} \cap C_{true}|/|C_{top}|$.

We found Concept 2 ("Hot dog") to be most monosemantic with 100% specificity. For Concept 20 ("Black dog") too, we found high specificity of 93.3%. For concept 15 ("Bull dog") we observed

Token	Rnd-Words	ViT-L (MAE)	ViT-L (CLIP)	GT-captions
Dog	0.515 ± 0.05	0.602 ± 0.07	0.610 ± 0.09	0.784 ± 0.06
Bus	0.501 ± 0.05	0.627 ± 0.07	0.634 ± 0.08	0.737 ± 0.05
Train	0.483 ± 0.06	0.618 ± 0.08	0.646 ± 0.07	0.726 ± 0.05
Cat	0.541 ± 0.04	0.629 ± 0.09	0.627 ± 0.06	0.795 ± 0.06

Table 6: Test CLIPScore evaluation for DePALM with ViT-L frozen image encoder trained as masked autoencoder (MAE): Scores reported for top-1 activating concept of Semi-NMF for Rnd-Words, GT-captions and ViT-L (CLIP) which denotes the system in main text.

the lowest specificity of 50%. This concept also activated for test images with toy/stuffed dogs. Interestingly, the multimodal grounding of concept 15 already indicates this superposition with maximum activating samples also containing images of 'toy dogs'. Concept 6 ("Brown/orange dog") is somewhere in between, with 76% specificity. This concept vector also activated sometimes for dark colored dogs, which wasn't apparent from its multimodal grounding.

Prominent or distinct semantic concepts seem to be captured by more specific/monosemantic concept vectors, while more sparsely present concepts seem at risk to be superposed resulting in a more polysemantic concept vector capturing them. It is also worth noting that the multimodal grounding can be a useful tool in some cases to identify polysemanticity in advance.

D Further implementation details

D.1 Dictionary learning details

Split	Dog	Bus	Train	Cat	Baby	Car	Cake	Bear
Train	3693	2382	3317	3277	837	1329	1733	1529
Test	161	91	147	167	44	79	86	55

Table 7: Number of samples training/testing samples for each token for DePALM

The details about the number of samples used for training the concept dictionary of each token, and the number of samples for testing is given in Tab. 7. The token representations are of dimension $B = 4096$.

The hyperparameters for the dictionary learning methods are already discussed in the main paper. All the dictionary learning methods (PCA, KMeans, Semi-NMF) are implemented using scikit-learn [31]. For PCA and KMeans we rely on the default optimization strategies. Semi-NMF is implemented through the DictionaryLearning() class, by forcing a positive code. It utilizes the coordinate descent algorithm for optimization during both the learning of \mathbf{U}^* , \mathbf{V}^* and the projection of test representations $v(X)$.

D.2 CLIPScore/BERTScore evaluation

For a given image X and set of words \mathbf{T}_k associated to concept u_k , CLIPScore is calculated between CLIP-image embedding of X and CLIP-text embedding of comma-separated words in \mathbf{T}_k . We consider a maximum of 10 most probable words in each \mathbf{T}_k , filtering out non-English and stop words. The computation of the metric from embeddings adheres to the standard procedure described in [16]. Our adapted implementation is based on the CLIPScore official repository, which utilizes the ViT-B/32 CLIP model to generate embeddings.

We found that computing BERTScores from comma-separated words and captions is unreliable. Instead, we adopted a method using the LLaMA-3-8B instruct model to construct coherent phrases from a set of grounded words, T_k . Specifically, we provide the LLaMA model with instructions to describe a scene using a designated set of words, for which we also supply potential answers. This instruction is similarly applied to another set of words, but without providing answers. The responses generated by LLaMA are then compared to the captions y using BERTScore. The instruction phrase and an example of the output are detailed in 8. The highest matching score between the generated phrases and the captions of a test sample determines the score assigned to the concept u_k . This approach ensures that the evaluation accurately reflects coherent and contextually integrated language use. The metric calculation from embeddings follows the established guidelines outlined in [41]. Our adapted implementation is based on BERTScore official repository, and we use the default Roberta-large model to generate embeddings.

Table 8: Generating contextually and grammatically coherent phrases Using the LLaMA Model for BERTScore Evaluation

Instruction to LLaMA	Nature of Response
Generate three distinct phrases, each incorporating the words dog, white, people. Ensure each phrase is clear and contextually meaningful. Number each phrase as follows: 1. 2. 3.	1. A white dog is standing next to people. 2. People are standing next to a white dog. 3. The dog is standing next to people wearing white clothes.
Generate three distinct phrases, each incorporating the words List of words . Ensure each phrase is clear and contextually meaningful. Number each phrase as follows: 1. 2. 3.	LLaMA autonomously creates a relevant description, demonstrating comprehension and creative integration of the new words.
Generate three distinct phrases, each incorporating the words brown, black, large, fluffy, cat . Ensure each phrase is clear and contextually meaningful. Number each phrase as follows: 1. 2. 3.	1. A large, fluffy black cat is sleeping on the brown couch. 2. The brown cat is curled up next to a large, fluffy black cat. 3. The large, fluffy cat's brown fur stands out against the black background.

D.3 Resources

DePALM experiments compute Each experiment to analyze a token with a selected dictionary learning method is conducted on a single RTX5000 (24GB)/ RTX6000 (48GB)/ TITAN-RTX (24GB) GPU. Within dictionary learning, generating visualizations and projecting test data, the majority of time is spent in loading the data/models and extracting the representations. For analysis of a single token with ≈ 3000 training samples, it takes around 10-15 mins for this complete process. Evaluation for CLIPScore/BERTScore are also conducted using the same resources. Evaluating CLIPScore for 500 (image, grounded-words) pairs takes around 5 mins. The BERTScore evaluation is in contrast more expensive, consuming around 150 mins for 500 pairs.

LLaVA experiments compute Each experiment to extract a concept dictionary for LLaVA was conducted on a single A100 (80GB) GPU. Representation extraction process for LLaVA is more expensive compared to DePALM consuming around 90 mins for ≈ 3000 samples. The remaining aspects of dictionary learning, multimodal grounding, representation projection etc. remains relatively cheap. The CLIPScore/BERTScore evaluations are completed with same resources as before.

Licenses of assets The part of the code for representation extraction from LMM is implemented using PyTorch [30]. For our analyses, we also employ the OPT-6.7B model [40] from Meta AI, released under the MIT license, and the CLIP model [32] from OpenAI, available under a custom usage license. Additionally, the COCO dataset [22] used for validation is accessible under the Creative Commons Attribution 4.0 License. We also use CLIPScore [16] and BERTScore [41] for

evaluating our method, both publicly released under MIT license. All utilized resources comply with their respective licenses, ensuring ethical usage and reproducibility of our findings.

D.4 Choice for number of concepts K

Our choice of using $K = 20$ concepts for all tokens was driven by the behaviour of reconstruction error of Semi-NMF on the training samples with different values of K , i.e. $\|\mathbf{Z} - \mathbf{UV}\|_2^2$. We validate this behaviour on the four target tokens from main text in Fig. 9. We generally found $K = 20$ as the minimal number of concepts where the reconstruction error drops by at least 50% from $K = 0$.

We also conducted an ablation study to test how our method behaves with different values of number of concepts K . Fig. 10 presents the variation of test CLIPScore and Overlap score for $K \in \{10, 20, 30, 50\}$ for two target tokens, ‘Dog’ and ‘Cat’. Our method can learn meaningful concepts for different values of K , evident by the consistently high CLIPScore. The Overlap score on the other hand tends to drop more consistently as the number of concepts increase but behaves stably for different choices. Nonetheless they indicate that our method can accommodate dictionaries of larger sizes without compromising the quality of learnt concepts, provided $K \ll M$ (number of decomposed samples) and at the expense of greater user overhead.




Figure 9: Variation of reconstruction error with number of concepts K for decompositions on different target tokens.

E Evaluation and extension to more tokens

We provide test data mean CLIPScore and BERTScore for top-1 activating concept for all baselines and more tokens: Baby, Car, Cake, and Bear in Tab. 9 (results in the main paper are reported for tokens Dog, Bus, Train, Cat in Tab. 1). Additionally, we also report the macro average over a set of 30 additional COCO-nouns apart from the 8 tokens with individually reported results, denoted as ‘Extra-30’. These nouns are single-token words with at least 40 predicted test samples. We put the filter of single-token words to keep consistency with the presented framework. Extension to multi-token words is straightforward but discussed separately in E.1. The lower bound criterion on test samples is to ensure average test CLIPScore is reliable for each target token. We only report CLIPScore for ‘Extra-30’ tokens as BERTScore evaluation was more expensive to conduct on large number of dictionaries.

We observe that we consistently obtain higher scores across for Semi-NMF and K-Means. We also report the overlap score between grounded words in Tab. 10 to illustrate the superiority of our method over the simple baseline. As previously noted, we observe a high overlap between grounded




Figure 10: Test CLIPScore and Overlap score ablation with number of concepts K . CLIPScore remains consistently high and drops slightly only for very small K . Overlap score generally improves with higher K .

words with KMeans and Simple baselines compared to Semi-NMF/PCA. A low overlap should be encouraged, as it indicates the discovery of diverse and disentangled concepts in the dictionary. Among all the methods, Semi-NMF provides the most balanced concept dictionaries which are both meaningful (high CLIPScore/BERTScore) and diverse (low overlap).

Statistical significance The statistical significance of Semi-NMF w.r.t all other baselines and variants, for CLIPScore/BERTScore evaluation to understand representations of test samples is given in Tab/ 11 (for all tokens separately). We report the p -values for an independent two sided T-test with null hypothesis that mean performance is the same between Semi-NMF and the respective system. The results for Semi-NMF are almost always significant compared to Rnd-Words, Noise-Imgs, PCA. However for these metrics, Simple baseline, K-Means and Semi-NMF all perform competitively and better than other systems. Within these three systems the significance depends on the target token, but are often not significant in many cases.

E.1 Extending to multi-token words

The presentation of our approach assumes that our token of interest t is a single token. This poses no issues for words which are represented as single tokens but can raise some questions when we wish to extract concept representations for multi-token words. Our approach however, can be easily adapted to this setting. In particular, we extract representation of last token from first prediction of the multi-token sequence. Note that when filtering the training data for samples where ground-truth caption contains the token of interest, we now search for the complete multi-token sequence. The other aspects of the method remain unchanged. While there can also be other viable strategies, the rationale behind this adaptation is that the last token of our sequence of interest can also combine representations from previous tokens in the sequence. We add below results for such examples in Tab. 12 and 13. We observe behaviour consistent with the previous results with Semi-NMF extracting concept dictionaries with high CLIPScore and low overlap.

Token	Metric	Rnd-Words	Noise-Imgs	Simple	PCA (Ours)	KMeans (Ours)	Semi-NMF (Ours)	GT-captions
Dog	CS top-1 (\uparrow)	0.519 \pm 0.05	0.425 \pm 0.06	0.546 \pm 0.08	0.559 \pm 0.06	<u>0.599</u> \pm 0.07	0.610 \pm 0.09	0.783 \pm 0.06
	BS top-1 (\uparrow)	0.201 \pm 0.04	0.306 \pm 0.05	0.346 \pm 0.08	0.353 \pm 0.10	<u>0.398</u> \pm 0.06	0.405 \pm 0.07	0.511 \pm 0.11
Bus	CS top-1 (\uparrow)	0.507 \pm 0.05	0.425 \pm 0.08	0.667 \pm 0.06	0.509 \pm 0.05	<u>0.645</u> \pm 0.08	0.634 \pm 0.08	0.736 \pm 0.05
	BS top-1 (\uparrow)	0.200 \pm 0.05	0.303 \pm 0.06	0.390 \pm 0.05	0.380 \pm 0.13	<u>0.401</u> \pm 0.06	0.404 \pm 0.07	0.466 \pm 0.11
Train	CS top-1 (\uparrow)	0.496 \pm 0.05	0.410 \pm 0.07	0.642 \pm 0.06	0.554 \pm 0.08	0.657 \pm 0.06	<u>0.646</u> \pm 0.07	0.727 \pm 0.05
	BS top-1 (\uparrow)	0.210 \pm 0.06	0.253 \pm 0.06	0.392 \pm 0.07	0.334 \pm 0.09	0.375 \pm 0.07	<u>0.378</u> \pm 0.07	0.436 \pm 0.08
Cat	CS top-1 (\uparrow)	0.539 \pm 0.04	0.461 \pm 0.04	0.589 \pm 0.07	0.541 \pm 0.08	<u>0.608</u> \pm 0.08	0.627 \pm 0.06	0.798 \pm 0.05
	BS top-1 (\uparrow)	0.207 \pm 0.07	0.307 \pm 0.03	<u>0.425</u> \pm 0.10	0.398 \pm 0.13	0.398 \pm 0.08	0.437 \pm 0.08	0.544 \pm 0.1
Baby	CS top-1 (\uparrow)	0.532 \pm 0.04	0.471 \pm 0.05	<u>0.631</u> \pm 0.06	0.575 \pm 0.07	0.636 \pm 0.05	0.621 \pm 0.06	0.811 \pm 0.05
	BS top-1 (\uparrow)	0.192 \pm 0.03	0.379 \pm 0.05	0.471 \pm 0.06	0.338 \pm 0.06	0.405 \pm 0.07	<u>0.426</u> \pm 0.07	0.530 \pm 0.14
Car	CS top-1 (\uparrow)	0.518 \pm 0.05	0.461 \pm 0.08	<u>0.605</u> \pm 0.04	0.547 \pm 0.08	0.602 \pm 0.05	0.614 \pm 0.05	0.766 \pm 0.06
	BS top-1 (\uparrow)	0.192 \pm 0.03	0.336 \pm 0.05	0.448 \pm 0.07	0.370 \pm 0.10	<u>0.435</u> \pm 0.08	0.379 \pm 0.08	0.485 \pm 0.14
Cake	CS top-1 (\uparrow)	0.488 \pm 0.05	0.473 \pm 0.08	<u>0.631</u> \pm 0.05	0.540 \pm 0.07	0.657 \pm 0.06	0.628 \pm 0.08	0.772 \pm 0.05
	BS top-1 (\uparrow)	0.186 \pm 0.04	<u>0.366</u> \pm 0.07	0.375 \pm 0.10	0.243 \pm 0.08	0.334 \pm 0.07	0.334 \pm 0.08	0.414 \pm 0.13
Bear	CS top-1 (\uparrow)	0.526 \pm 0.04	0.526 \pm 0.06	0.651 \pm 0.04	0.564 \pm 0.06	0.680 \pm 0.05	<u>0.660</u> \pm 0.07	0.798 \pm 0.06
	BS top-1 (\uparrow)	0.255 \pm 0.10	0.396 \pm 0.08	0.434 \pm 0.05	0.420 \pm 0.10	<u>0.474</u> \pm 0.08	0.494 \pm 0.10	0.541 \pm 0.10
Extra-30	CS top-1 (\uparrow)	0.516 \pm 0.04	0.521 \pm 0.03	0.626 \pm 0.06	0.547 \pm 0.06	0.637 \pm 0.06	<u>0.631</u> \pm 0.06	0.763 \pm 0.05

Table 9: Test data mean CLIPScore and BERTScore for top-1 activating concept CLIPScore denoted as CS, BERTScore denoted as BS for all concept extraction baselines considered. Analysis for layer $L = 31$. Best score indicated in **bold** and second best is underlined. ‘Extra-30’ denotes a set of 30 additional single-token COCO nouns (apart from previous 8) with at least 40 predicted test samples by f_{LM} . For ‘Extra-30’ tokens we report the macro average and standard deviation of mean test data CLIPScore, taken over the set of 30 tokens.

Token	Simple	PCA	KMeans	Semi-NMF
Baby	0.645	0.006	0.502	<u>0.187</u>
Car	0.246	0.001	0.322	<u>0.097</u>
Cake	0.415	0.005	0.398	<u>0.147</u>
Bear	0.556	0.002	0.360	<u>0.203</u>
Extra-30	0.443	0.050	0.452	<u>0.156</u>

Table 10: Overlap/entanglement between grounded words of learnt concepts for different dictionary learning methods. Results are for additional tokens. Lower is better. Best score indicated in **bold** and second best is underlined. For ‘Extra-30’ tokens (additional 30 single-token COCO nouns) we report the macro average of Overlap score, taken over the set.

F Additional visualizations

F.1 Concept grounding

The visual/textual grounding for all tokens in Tab. 1 are given in Figs. 11 (‘Dog’), 12 (‘Cat’), 13 (‘Bus’), 14 (‘Train’). All the results extract $K = 20$ concepts from layer $L = 31$. Similar to our analysis for token ‘Dog’ in main paper, for a variety of target tokens our method extracts diverse and multimodally coherent concepts encoding various aspects related to the token.

Token	Metric	Rnd-Words	Noise-Imgs	Simple	PCA	KMeans	GT-captions
Dog	CS top-1	< 0.001	< 0.001	< 0.001	< 0.001	> 0.1	< 0.001
	BS top-1	< 0.001	< 0.001	< 0.001	< 0.001	> 0.1	< 0.001
Bus	CS top-1	< 0.001	< 0.001	0.002	< 0.001	> 0.1	< 0.001
	BS top-1	< 0.001	< 0.001	> 0.1	> 0.1	> 0.1	< 0.001
Train	CS top-1	< 0.001	< 0.001	> 0.1	< 0.001	> 0.1	< 0.001
	BS top-1	< 0.001	< 0.001	0.08	< 0.001	> 0.1	< 0.001
Cat	CS top-1	< 0.001	< 0.001	< 0.001	< 0.001	0.018	< 0.001
	BS top-1	< 0.001	< 0.001	> 0.1	0.001	< 0.001	< 0.001
Baby	CS top-1	< 0.001	< 0.001	> 0.1	0.002	> 0.1	< 0.001
	BS top-1	< 0.001	< 0.001	0.001	< 0.001	> 0.1	< 0.001
Car	CS top-1	< 0.001	< 0.001	> 0.1	< 0.001	> 0.1	< 0.001
	BS top-1	< 0.001	< 0.001	< 0.001	0.533	< 0.001	< 0.001
Cake	CS top-1	< 0.001	< 0.001	> 0.1	< 0.001	0.008	< 0.001
	BS top-1	< 0.001	0.005	0.003	< 0.001	> 0.1	< 0.001
Bear	CS top-1	< 0.001	< 0.001	> 0.1	< 0.001	0.072	< 0.001
	BS top-1	< 0.001	< 0.001	0.0001	0.0001	> 0.1	0.015

Table 11: Statistical significance of Semi-NMF w.r.t other baselines for test data CLIP-Score/BERTScore. p -values for two sided T-test are reported. Significant values (p -value < 0.05) are indicated in bold. The values do **not** indicate which system has better mean score but just that if the difference is significant.

Multi-token word	Rnd-Words	Noise-Imgs	Simple	Semi-NMF	GT-captions
Traffic light	0.516 ± 0.03	0.525 ± 0.03	0.664 ± 0.06	<u>0.634</u> ± 0.05	0.744 ± 0.04
Cell phone	0.542 ± 0.04	0.547 ± 0.03	<u>0.598</u> ± 0.04	0.598 ± 0.05	0.765 ± 0.06
Stop sign	0.533 ± 0.03	0.549 ± 0.03	0.617 ± 0.08	<u>0.616</u> ± 0.05	0.775 ± 0.04

Table 12: Test mean CLIPScore (\uparrow) reported for top-1 activating concept for multi-token words. Higher scores are better. Best score in **bold**, second best is underlined.

F.2 Local interpretations

Here, we qualitatively analyze the local interpretations of various decomposition methods, including PCA, k-means, semi-NMF, and the simple baseline strategy. We select these four as they produce coherent grounding compared to Rnd-Words and Noise-Img baselines. We decompose test sample representations on our learnt dictionary and visualize the top three activating components. Note that in the case of KMeans and Simple baseline, the projection maps a given test representation to a single element of the concept dictionary, the one closest to it. However, for uniformity we show the three

Multi-token word	Simple	PCA	K-Means	Semi-NMF
Traffic light	0.704	0.050	0.579	<u>0.174</u>
Cell phone	0.623	0.051	0.746	<u>0.164</u>
Stop sign	0.461	0.058	0.704	<u>0.109</u>

Table 13: Overlap score (\downarrow) reported for top-1 activating concept for multi-token words. Higher scores are better. Best score in **bold**, second best is underlined.



Figure 11: multimodal concept grounding in vision and text for the token 'Dog'. The five most activating samples and the five most probable decoded words for each component u_k , $k \in \{1, \dots, 20\}$ are shown. The token representations are extracted from L=31 of the LLM section of our LMM.

most closest concept vectors for both. Figs. 15, 16, 17, 18, 19 are dedicated to interpretations of a single sample each, for all four concept extraction methods.

The inferences drawn about the behaviour of the four baselines from quantitative metrics can also be observed qualitatively. Semi-NMF, K-Means and 'Simple' baseline, are all effective at extracting grounded words can be associated to a given image. However, both K-Means and 'Simple' display similar behaviour in terms of highly overlapping grounded words across concepts. This behaviour likely arises due to both the baselines mapping a given representation to a single concept/cluster. This limits their capacity to capture the full complexity of data distributions. In contrast, Semi-NMF and PCA utilize the full dictionary to decompose a given representation and thus recover significantly more diverse concepts. PCA in particular demonstrates almost no overlap, likely due to concept vectors being orthogonal. However, the grounded words for it tend to be less coherent with the images. As noted previously, Semi-NMF excels as the most effective method, balancing both aspects by extracting meaningful and diverse concepts.



Figure 12: multimodal concept grounding in vision and text for the token 'Cat'. The five most activating samples and the five most probable decoded words for each component u_k , $k \in \{1, \dots, 20\}$ are shown. The token representations are extracted from $L=31$ of the LLM section of our LMM.

G Concept extraction across layers

We analyze the quality of multimodal grounding of concepts across different layers L . The CLIPScore between $(\mathbf{X}_{k,MAS}, \mathbf{T}_k)$, averaged over all concepts u_k is shown in Fig. 20, for 'Dog' and 'Cat' for all layers in f_{LM} . For early layers the multimodal grounding is no better than Rnd-Words. Interestingly, there is a noticeable increase around $(L = 20$ to $L = 25)$, indicating that the multimodal structure of internal token representations starts to appear at this point. This also validates our choice that deeper layers are better suited for multimodal concepts.

We provide a qualitative comparison of multimodal grounding for the token 'dog' across different layers in Fig. 21. As observed in Fig. 20, the multimodal nature of token representations for two tokens 'Dog' and 'Cat' starts to appear around layers $L = 20$ to $L = 25$. It is interesting to note that the representations of images still tend to be separated well, as evident by the most activating samples of different concepts. However, until the deeper layers the grounded words often do not correspond well to the visual grounding. This behaviour only appears strongly in deeper layers.



Figure 13: multimodal concept grounding in vision and text for the token 'Bus'. The five most activating samples and the five most probable decoded words for each component u_k , $k \in \{1, \dots, 20\}$ are shown. The token representations are extracted from L=31 of the LLM section of our LMM.

H Analysis for visual tokens

Our analysis in main paper was limited to decomposing representations of text tokens in various layers of an LLM, $h_{(l)}^p, p > N_V$. This was particularly because these were the predicted tokens of the multimodal model. Nevertheless, the same method can also be used to understand the information stored in the visual/perceptual tokens representations as processed in f_{LM} , $h_{(l)}^p, p \leq N_V$. An interesting aspect worth highlighting is that while the text token representations in f_{LM} can combine information from the visual token representations (via attention function), the reverse is not true. The causal processing structure of f_{LM} prevents the visual token representations to attend to any information in the text token representations. Given a token of interest t , for any sample $X \in \mathbf{X}_t$ we now only search for first position $p \in \{1, \dots, N_V\}$, s.t. $t = \arg \max \text{Unembed}(h_{(N_L)}^p)$. Only the samples for which such a p exists are considered for decomposition. The rest of the method to learn \mathbf{U}^* , \mathbf{V}^* proceeds exactly as before.

We conduct a small experiment to qualitatively analyze concepts extracted for visual token representations for 'Dog'. We extract $K = 20$ concepts from $L = 31$. The dictionary is learnt with representations from $M = 1752$ samples, less than $M = 3693$ samples for textual tokens. As a brief illustration, 12 out of 20 extracted concepts are shown in Fig. 22. Interestingly, even the visual token representations in deep layers of f_{LM} , without ever attending to any text tokens, demonstrate a multimodal semantic structure. It is also worth noting that there are multiple similar concepts that appear for both visual and textual tokens. Concepts 3, 7, 10, 12, 17, 19 are all similar visually



Figure 14: multimodal concept grounding in vision and text for the token 'Train'. The five most activating samples and the five most probable decoded words for each component u_k , $k \in \{1, \dots, 20\}$ are shown. The token representations are extracted from L=31 of the LLM section of our LMM.

and textually to certain concepts discovered for text tokens. This indicates to a possibility that these concepts are discovered by f_{LM} in processing of the visual tokens and this information gets propagated to predicted text token representations.

I Limitations

We list below some limitations of our proposed method:

- The concept dictionaries extracted currently are token-specific. It can be interesting to explore learning concept dictionaries that can encode shared concepts for different tokens.
- We select the most simple and straightforward concept grounding techniques. Both visual and textual grounding could potentially be enhanced. The visual grounding can be improved by enhancing localization of concept activation for any MAS or test sample. Text grounding could be enhanced by employing more sophisticated approaches such as tuned lens [5].
- While the proposed CLIPScore/BERTScore metrics are useful to validate this aspect, they are not perfect metrics and affected by imperfections and limitations of the underlying models extracting the image/text embeddings. The current research for metrics useful

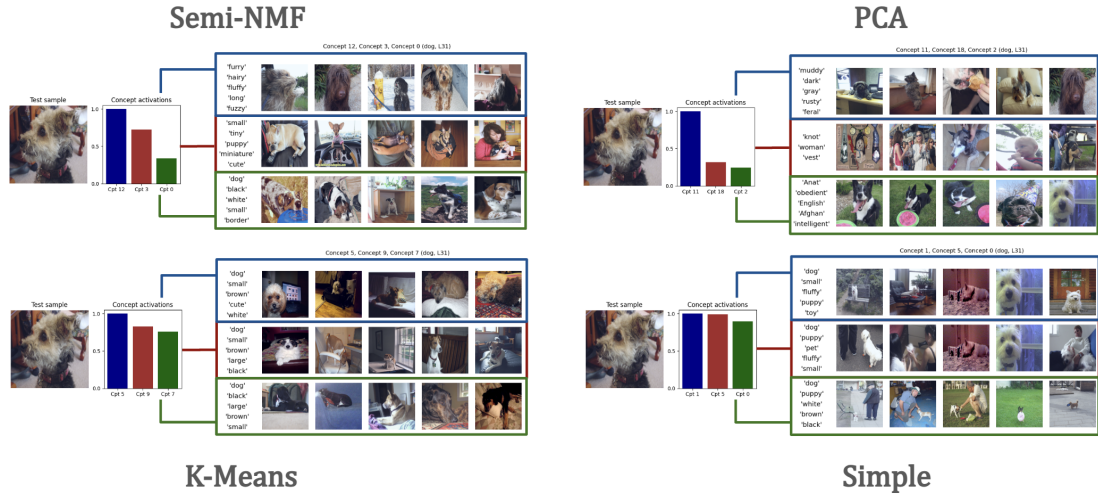


Figure 15: Local interpretations for test sample 9 of token 'Dog' with SemiNMF, KMeans, PCA, and Simple baselines (layer 31). Visual/text grounding for the three highest concept activations (normalized) is shown. SemiNMF baseline provides the most visually and textually consistent results, while other baselines provide components that are not well disentangled (Simple and KMeans baseline), or the text grounding is not closely related to the test image.

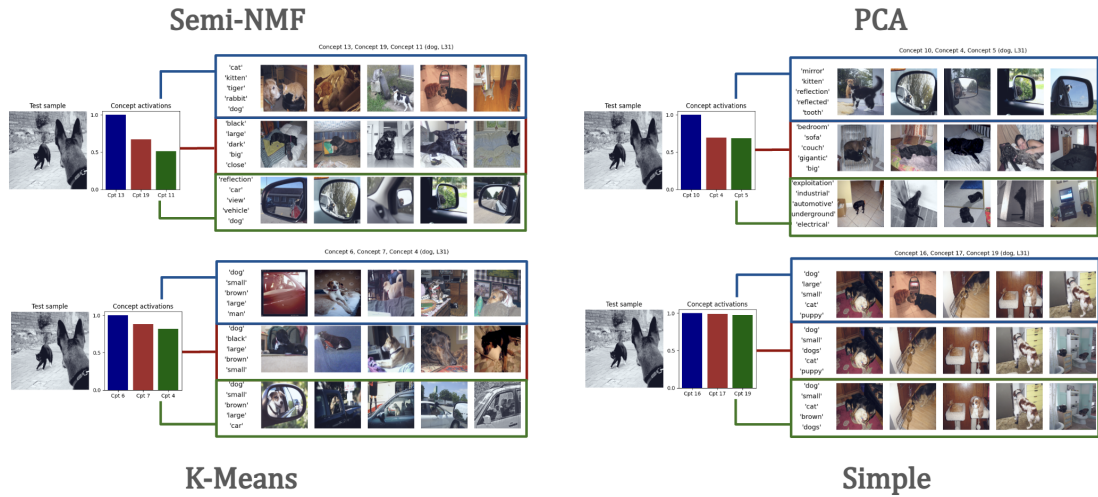


Figure 16: Local interpretations for test sample 37 of token 'Dog' with SemiNMF, KMeans, PCA, and Simple baselines (layer 31). Visual/text grounding for the three highest concept activations (normalized) is shown.

for interpretability remains an interesting open question, even more so in the context of LLMs/LMMs.




Figure 17: Local interpretations for test sample 43 of token 'Cat' with SemiNMF, KMeans, PCA, and Simple baselines (layer 31). Visual/text grounding for the three highest concept activations (normalized) is shown.

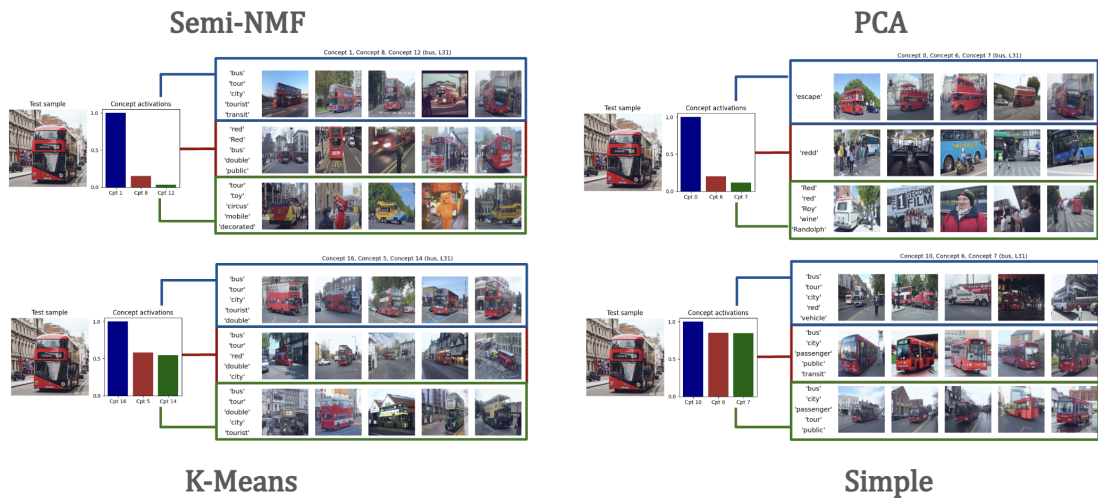


Figure 18: Local interpretations for test sample 6 of token 'Bus' with SemiNMF, KMeans, PCA, and Simple baselines (layer 31). Visual/text grounding for the three highest concept activations (normalized) is shown.

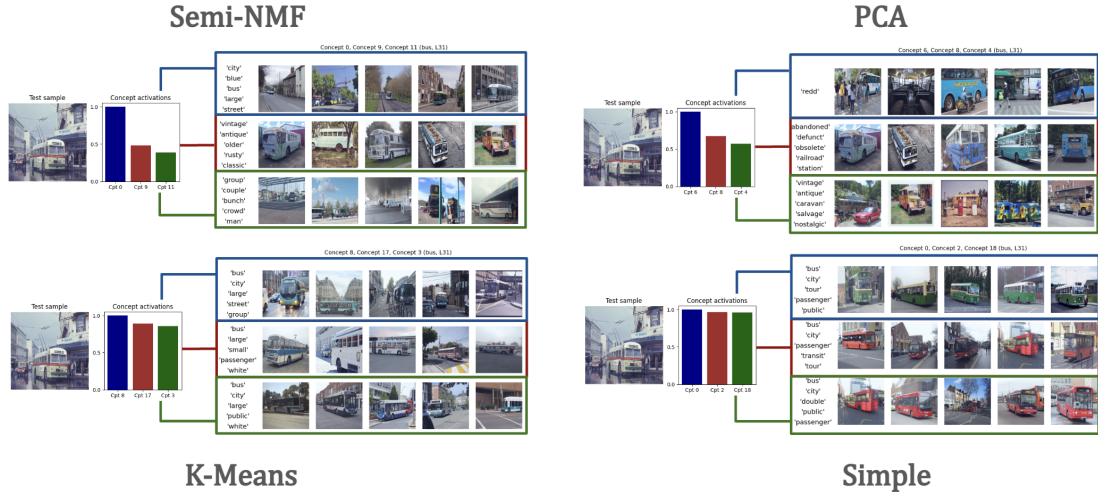


Figure 19: Local interpretations for test sample 12 of token ‘Bus’ with SemiNMF, KMeans, PCA, and Simple baselines (layer 31). Visual/text grounding for the three highest concept activations (normalized) is shown.

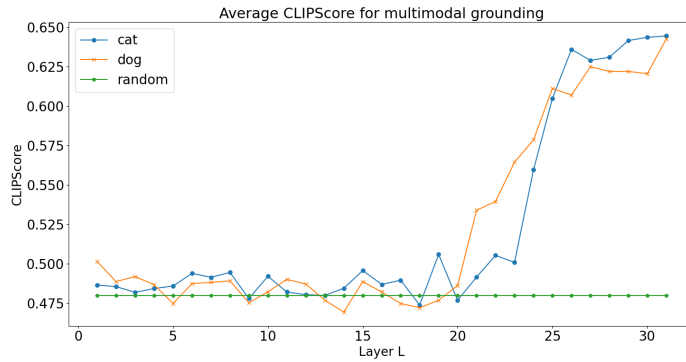


Figure 20: Mean CLIPScore between visual/text grounding $\mathbf{X}_{k, MAS}$, \mathbf{T}_k for all concepts (Semi-NMF), across different layers L . Results are for tokens ‘Dog’ and ‘Cat’.



Figure 21: Examples of multimodal grounding across different layers for concepts with similar visual grounding (target token 'Dog'). The grounded words from early layers do not correspond well to the most activating samples of a concept.

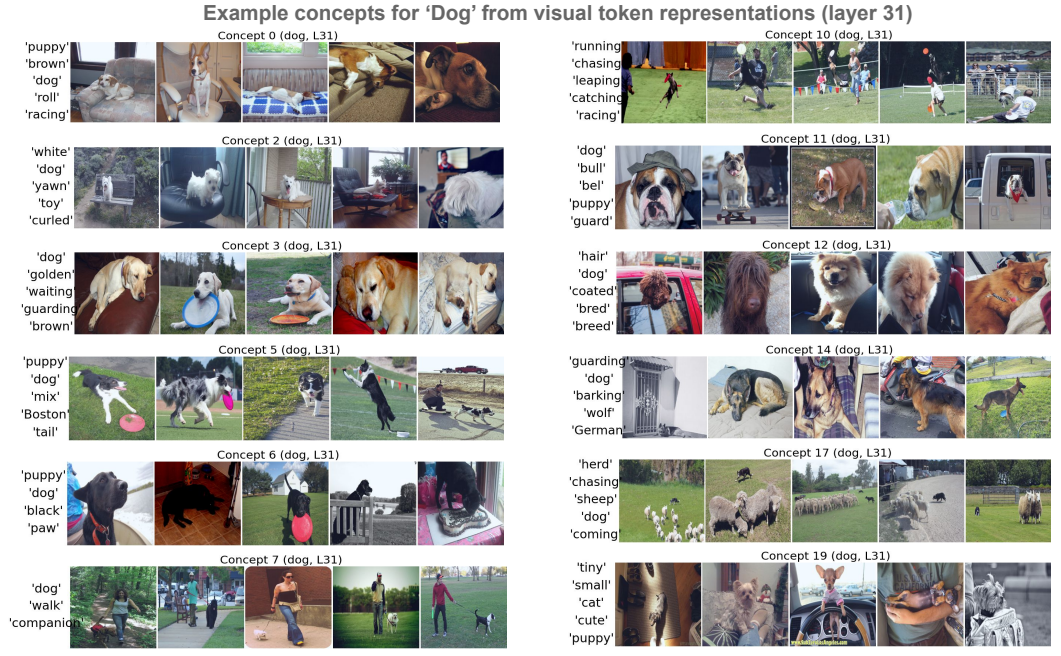


Figure 22: Example concepts extracted for 'Dog' from visual token representations in layer 31.