
Explaining Similarity in Vision–Language Encoders with Weighted Banzhaf Interactions

Hubert Baniecki*
University of Warsaw
Warsaw University of Technology

Maximilian Muschalik
LMU Munich, MCML

Fabian Fumagalli
Bielefeld University

Barbara Hammer
Bielefeld University

Eyke Hüllermeier
LMU Munich, MCML, DFKI

Przemysław Biecek
University of Warsaw
Warsaw University of Technology

Abstract

Language–image pre-training (LIP) enables the development of vision–language models capable of zero-shot classification, localization, multimodal retrieval, and semantic understanding. Various explanation methods have been proposed to visualize the importance of input image–text pairs on the model’s similarity outputs. However, popular saliency maps are limited by capturing only first-order attributions, overlooking the complex cross-modal interactions intrinsic to such encoders. We introduce faithful interaction explanations of LIP models (FIXLIP) as a unified approach to decomposing the similarity in vision–language encoders. FIXLIP is rooted in game theory, where we analyze how using the weighted Banzhaf interaction index offers greater flexibility and improves computational efficiency over the Shapley interaction quantification framework. From a practical perspective, we propose how to naturally extend explanation evaluation metrics, such as the pointing game and area between the insertion/deletion curves, to second-order interaction explanations. Experiments on the MS COCO and ImageNet-1k benchmarks validate that second-order methods, such as FIXLIP, outperform first-order attribution methods. Beyond delivering high-quality explanations, we demonstrate the utility of FIXLIP in comparing different models, e.g. CLIP vs. SigLIP-2.

1 Introduction

Contrastive language–image pre-training [CLIP, 57] revolutionized computer vision by learning data representations well-performing in a plethora of downstream tasks. Scaling the training to predict similarity between image and text, combined with continuous architectural improvements [e.g. SigLIP, 75] and increasing the size of datasets, leads to the development of powerful vision–language encoders [VLEs, 13, 66, 68, 75]. These rather opaque encoders become the backbone components in large vision–language models capable of actual conversation and reasoning [2, 36, 43], and are increasingly applied in high-stakes decision-making, like in the case of medical imaging [29, 76, 78].

The limitations in the way CLIP represents the visual world have been extensively studied [14, 35, 67, 72, 73]. For example, a recent study by Tong et al. [63] suggests vision–language models are limited by the systematic shortcomings of CLIP, e.g. they fail when questioned about directions, quantity of objects in an image, or recognizing the presented text in a visual form. The emergence of “CLIP-blind pairs”, inputs that CLIP perceives as similar despite their clear differences, could lead to critical errors when applied in medical visual question answering and treatment recommendation [78].

*Corresponding author: h.baniecki@uw.edu.pl Code: <https://github.com/hbaniecki/fixlip>




Figure 1: Explaining similarity in vision–language encoders with weighted Banzhaf interactions. We propose a cross-modal sampling strategy to efficiently query the model for m^2 game values from m coalitions, and p -weighted masking to circumvent querying the model on out-of-distribution inputs. A regression-based approximation with weighted least squares (WLS) of second-order attributions gives a faithful decomposition of the predicted similarity score. Explanations of cross-modal and intra-modal interactions can be visualized and analyzed to interpret the CLIP’s similarity prediction. The red values denote positive interactions contributing to an increase in similarity, while blue denotes interactions between tokens contributing to a decrease in similarity.

One way of validating similarity predictions at inference time is explaining the encoders with saliency maps [12, 39, 49, 70, 77]. However, these first-order importance scores are limited by their unimodal view of the model’s output [1]. Vision–language encoders cannot be faithfully explained without taking into account the cross-modal interactions between image–text pairs [32, 40].

We thus propose a model-agnostic approach to interpreting VLEs based on a game-theoretical perspective where input tokens form coalitions of players in a cooperative game [19, 44]. In doing that, we extend the faithful Banzhaf interaction index [64] to a weighted case [47], addressing the emerging requirement for controllable sparsity in coalition sampling [6, 25, 27, 79]. Specifically, sparse inputs with many masked tokens become out-of-distribution, i.e. images become unrecognizable and text captions become ambiguous. Figure 1 shows a high-level view of our approach. We sample uniformly m^2 variations of the input image–text pair. Each mask is sampled uniformly with probability p for each token in a mask, and we take all pairwise combinations between m masked images and m masked texts. Then, we apply the explained model to predict similarity values for each of the input variations. Finally, we apply weighted least squares to approximate the model’s predictions from sampled binary masks. The resulting explanation assigns attribution scores to each token and interaction scores to each pair of tokens in an image–text pair.

Contributions. Our work advances literature in multiple ways: (1) **A game-theoretic explanation of vision–language encoders.** We introduce faithful interaction explanations of LIP models (FIXLIP) that provide a unique perspective on decomposing the similarity predictions of VLEs. Our results highlight the *necessity* to consider cross-modal interactions between text and image inputs for an accurate model interpretation, beyond the first-order attribution. In fact, this is the first work to argue for using *weighted* Banzhaf interactions in this domain to overcome the out-of-distribution problem apparent in removal-based explainability. (2) **Efficient computation allows scaling to larger models.** We propose a *cross-modal* sampling strategy for approximating FIXLIP that improves its efficiency by 5–20 $\times$ over traditional Shapley interaction quantification. We further scale the regression-based approximation of interactions to hundreds of tokens by prioritizing their subset for a faithful explanation. (3) **Evaluation metrics for second-order interactions.** We extend explanation evaluation metrics—the pointing game and area between the insertion/deletion curves—to explanation methods of higher order. Experiments on the MS COCO and ImageNet-1k benchmarks validate that second-order methods, such as FIXLIP, outperform first-order attributions. (4) **FIXLIP facilitates**

various approaches to understanding vision–language encoders. Finally, we demonstrate the utility of FIXLIP in applications like a model-agnostic comparison between CLIP and SigLIP-2.

2 Related Work

Attribution explanations of CLIP. Multimodal explanations have been studied in various tasks, including visual question answering [28, 55], visual reasoning [46], video classification [71], and sentiment analysis [69]. Our work considers explaining the bi-modal encoder prediction in the means of input attributions, specifically image and text token attributions of similarity predictions in CLIP. In this context, existing gradient-based [77], attention-based [12, 39], and information-based [70, 80] methods approximate only *first-order* attributions visualized as saliency maps. Yet, Liang et al. [40] perform extensive user studies showing that visualizing the *second-order* attributions (pairwise interactions) is necessary for understanding complex multimodal models. Joukovsky et al. [32] propose to approximate these interactions with bilinear models [also for image–image encoders, 61]. We formalize the problem in the means of game theory and propose an efficient approximation of cross-modal and intra-modal interactions (Figure 1), discussing its several desirable theoretical properties. In concurrent work, Moeller et al. [48] propose to approximate only cross-modal interactions based on the Integrated Hessians methodology [30]. Our proposed evaluation metrics for interaction-based explanations can be used to empirically cross-compare the above-mentioned attribution explanations.

Mechanistic interpretability of CLIP. Our research is orthogonal to work on the concept-based interpretability of CLIP’s internal representations. Research in this direction considers explaining particular neurons [22], concept-based image classification [60], linear probing [5], and training sparse autoencoders [41, 74]. Gandelsman et al. [21] investigate how the individual model’s components affect its final representation. Beyond CLIP, Balasubramanian et al. [4] analyze interpreting image representations via text in alternative vision transformer architectures [e.g. DINO, 11].

Shapley and Banzhaf interactions in machine learning. We build on the developments using game-theoretic interaction indices to understand complex machine learning models [9, 19, 56, 62, 64]. Our goal in this work is not to compare with the most recent advancements in approximating the interaction index [16, 33, 52], but rather to faithfully explain the similarity predicted by VLEs. From the first-order perspective, Parcalabescu and Frank [53, 54] apply the Shapley value to measure the importance each modality has on various tasks solved by CLIP. Tsai et al. [64] and Fumagalli et al. [20] propose a weighted least squares regression approximation of higher-order interactions, which we incorporate as a baseline in our methodology. Recently, Kang et al. [34] applied a sparse Fourier transform to scale feature interaction explanations for large language models, while we explain vision–language encoders instead. Related to our contribution of weighted Banzhaf interaction explanations is work on using weighted Banzhaf values to estimate data valuation scores [37, 38], which are the importance that training data points or subsets have on the model’s performance. Jin et al. [31] use Banzhaf interactions to enhance the training of video–language encoders. For a comprehensive overview of work in this direction, see [50, 59] and references given there.

3 A Game-theoretic Explanation of Similarity in Vision–Language Encoders

A vision–language encoder f consists of a vision encoder $f_{\mathcal{I}} : \mathbb{R}^{n_{\mathcal{I}}} \rightarrow \mathbb{R}^d$ with $n_{\mathcal{I}} \in \mathbb{N}$ image patches (tokens), and a language encoder $f_{\mathcal{T}} : \mathbb{R}^{n_{\mathcal{T}}} \rightarrow \mathbb{R}^d$ with $n_{\mathcal{T}} \in \mathbb{N}$ text tokens. For a given input pair $(x_{\mathcal{I}}, x_{\mathcal{T}})$, the model computes a cosine similarity of their d -dimensional embeddings as

$$f(x_{\mathcal{I}}, x_{\mathcal{T}}) := \cos(f_{\mathcal{I}}(x_{\mathcal{I}}), f_{\mathcal{T}}(x_{\mathcal{T}})) = \frac{f_{\mathcal{I}}(x_{\mathcal{I}}) \cdot f_{\mathcal{T}}(x_{\mathcal{T}})}{\|f_{\mathcal{I}}(x_{\mathcal{I}})\| \cdot \|f_{\mathcal{T}}(x_{\mathcal{T}})\|}. \quad (1)$$

In this context, we index the set of image tokens with $N_{\mathcal{I}} := \{1, \dots, n_{\mathcal{I}}\}$ and the set of text tokens with $N_{\mathcal{T}} := \{n_{\mathcal{I}} + 1, \dots, n_{\mathcal{I}} + n_{\mathcal{T}}\}$. Related attribution methods construct explanations that assign importance values to all individual input tokens. To understand cross-modal and intra-modal relationships between tokens, we extend these explanations by adding all pairwise token interactions.

Definition 1 (Explanation). *An explanation $\mathbf{e} \in \mathbb{R}^{|\mathcal{B}|}$ assigns a constant $\mathbf{e}_0$, individual attributions $\mathbf{e}_i \in \mathbb{R}$ and pairwise interactions $\mathbf{e}_{\{i,j\}} \in \mathbb{R}$, for an explanation basis*

$$\mathcal{B} := \underbrace{\{0\}}_{\text{constant}} \cup \underbrace{\{i : i \in N_{\mathcal{I}} \cup N_{\mathcal{T}}\}}_{\text{individual tokens}} \cup \underbrace{\{\{i, j\} : i, j \in N_{\mathcal{I}} \cup N_{\mathcal{T}}, i \neq j\}}_{\text{two-token interaction sets}}.$$

An interaction explanation can be viewed as a complete graph with weighted nodes and edges, similar to the SI-Graph [51], or a symmetric quadratic matrix of dimension $n_{\mathcal{I}} + n_{\mathcal{T}}$, where the diagonal represents token attributions. To compute explanations, we mask tokens with a pre-defined baseline, depending on the encoder's architecture, e.g. attention masking for text tokens (we defer implementation details to Appendix B). The masking operator is used to define the FIXLIP game that captures all possible masks.

Definition 2 (Masking). We define the masking operator $\oplus_{M_o} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ for $n \in \{n_{\mathcal{I}}, n_{\mathcal{T}}\}$ that compute for a subset of indices $M_o \subseteq N$ with $N \in \{N_{\mathcal{I}}, N_{\mathcal{T}}\}$ and inputs $x, b \in \mathbb{R}^n$ as

$$x \oplus_{M_o} b := \begin{cases} x_i, & \text{if } i \in M_o, \\ b_i, & \text{if } i \in N \setminus M_o. \end{cases}$$

Definition 3 (Game). For an input image-text pair $(x_{\mathcal{I}}, x_{\mathcal{T}})$ and a baseline $(b_{\mathcal{I}}, b_{\mathcal{T}})$ indexed by $N_{\mathcal{I}} \cup N_{\mathcal{T}}$, the FIXLIP explanation game $\nu : 2^{N_{\mathcal{I}} \cup N_{\mathcal{T}}} \rightarrow \mathbb{R}$ measures the similarity of the inputs given any mask $M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}$ as $\nu(M) := \nu(M \cap N_{\mathcal{I}}, M \cap N_{\mathcal{T}}) := f(x_{\mathcal{I}} \oplus_{M \cap N_{\mathcal{I}}} b_{\mathcal{I}}, x_{\mathcal{T}} \oplus_{M \cap N_{\mathcal{T}}} b_{\mathcal{T}})$.

The FIXLIP game measures the similarity of the masked image and text inputs for every possible mask. In the following, we will quantify attributions of individual tokens, as well as cross- and intra-modal interactions, that *faithfully* explain the FIXLIP game, and satisfy important axioms.

3.1 FIXLIP- p explanations via weighted faithful Banzhaf interactions

In this section, we formally define the FIXLIP explanation as an instance of Definition 1, which faithfully describes the similarity of all masked inputs $\nu(M)$. We thus define $\mathbf{e}$ as an *additive explanation with interaction terms* that recovers $\nu(M) \approx \hat{\nu}_{\mathbf{e}}(M)$ via a 2-additive game [24] as

$$\hat{\nu}_{\mathbf{e}}(M) := \mathbf{e}_0 + \sum_{i \in M} \mathbf{e}_i + \sum_{\{i,j\} \subseteq M: i \neq j} \mathbf{e}_{\{i,j\}} \text{ for all masks } M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}. \quad (2)$$

A natural choice of faithfulness is to measure the squared error $(\nu(M) - \hat{\nu}_{\mathbf{e}}(M))^2$ across all masks. To further distinguish in-distribution (few tokens masked) and out-of-distribution (many tokens masked) inputs, we associate a mask weight according to its size $|M|$ in the following metric.

Definition 4 (p -faithfulness). For $p \in (0, 1)$ and $\nu, \hat{\nu} : 2^{N_{\mathcal{I}} \cup N_{\mathcal{T}}} \rightarrow \mathbb{R}$, we measure p -faithfulness $\mathfrak{F}_p$ as $\mathfrak{F}_p(\nu, \hat{\nu}) := \sum_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}} p^{|M|} (1-p)^{n_{\mathcal{I}}+n_{\mathcal{T}}-|M|} (\nu(M) - \hat{\nu}(M))^2$.

Remark 1. Computing p -faithfulness requires evaluating $2^{n_{\mathcal{I}}+n_{\mathcal{T}}}$ masks, which is infeasible in practice. Instead, we write $\mathfrak{F}_p$ as an expectation over randomly generated masks $\mathfrak{F}_p = \mathbb{E}_{M \sim \mathbb{P}_p} [(\nu(M) - \hat{\nu}(M))^2]$, where $\mathbb{P}_p(M) := p^{|M|} (1-p)^{n_{\mathcal{I}}+n_{\mathcal{T}}-|M|}$ is a probability distribution over $2^{N_{\mathcal{I}} \cup N_{\mathcal{T}}}$. This formulation allows us to estimate $\mathfrak{F}_p$ using Monte Carlo integration.

The hyperparameter p determines the importance of each mask and can be viewed as the probability of a token being active, i.e. its index being included in M . The case $p = 0.5$ equally weights all masks, $p > 0.5$ prioritizes masks with many active tokens, i.e. preferring in-distribution inputs, whereas $p < 0.5$ prioritizes sparse masks preferring more out-of-distribution inputs.

Definition 5 (FIXLIP- p). We define the FIXLIP- p explanation as $\mathbf{e}^{\text{FIXLIP-}p} := \arg\min_{\mathbf{e}} \mathfrak{F}_p(\nu, \hat{\nu}_{\mathbf{e}})$.

The FIXLIP- p explanation is the most faithful approximation of ν with respect to $\mathfrak{F}_p$. Its interactions correspond to the weighted Banzhaf interactions [47]. For $p = 0.5$, FIXLIP- p is the faithful Banzhaf interaction index [64] with a known analytical solution [26].

Remark 2. Weighted Banzhaf interactions define a cardinal-probabilistic interaction index that satisfies the linearity, symmetry, and dummy axiom [18]. Alternatively, faithful Shapley interactions optimize a variant of the faithfulness metric [64]. In our context, weighted Banzhaf interactions have two benefits: First and foremost, the hyperparameter p allows for prioritizing masks that are expected to better reflect in-distribution inputs, and provides a flexible explanation framework with intuitive weights. Second, it is unknown if the Shapley weights can be factored into independent distributions for each modality – a crucial requirement for the cross-modal estimator introduced in Section 3.2.

FIXLIP- p explanations of VLEs offer a unique perspective on interpreting image-text similarity predictions. Notably, the resulting second-order interaction explanation can always be simplified into a first-order attribution explanation and visualized as a traditional saliency map over input tokens.

Theorem 1 (First-order conversion). *For $i \in N_{\mathcal{I}} \cup N_{\mathcal{T}}$ and $\mathbf{e}^{\text{FixLIP-}p}$, the first-order attribution values are given by $\mathbf{e}_i + p \sum_{j \in N_{\mathcal{I}} \cup N_{\mathcal{T}}: j \neq i} \mathbf{e}_{\{i,j\}}$, which are the weighted Banzhaf values of $\hat{\nu}$.*

We defer the proofs to Appendix A. While FIXLIP- p can be computed analytically, it still requires $2^{n_{\mathcal{I}}+n_{\mathcal{T}}}$ mask evaluations. We thus introduce a model-agnostic estimator for FIXLIP- p by estimating $\mathfrak{F}_p$ via Monte Carlo integration with sampled masks and optimizing the least-squares objective.

Definition 6 (Model-agnostic estimator). *Let $M^{(1)}, \dots, M^{(m)} \stackrel{iid}{\sim} \mathbb{P}_p$. The model-agnostic estimator is then $\hat{\mathbf{e}}^{\text{FixLIP-}p} := \arg\min_{\mathbf{e}} \hat{\mathfrak{F}}_p^{(m)}(\nu, \hat{\nu}_{\mathbf{e}})$, where $\hat{\mathfrak{F}}_p^{(m)}(\nu, \hat{\nu}_{\mathbf{e}}) := \frac{1}{m} \sum_{\ell=1}^m (\nu(M^{(\ell)}) - \hat{\nu}_{\mathbf{e}}(M^{(\ell)}))^2$.*

For $p = 0.5$, this estimator reduces to Faith-Banzhaf [64] – a variant of KernelSHAP [44].

3.2 Cross-modal sampling with p -weighted masking

We observe that the model-agnostic estimator uses naive sampling over masks $M \sim \mathbb{P}_p$, which becomes prohibitive for models with a large number of input tokens. Since every token is independently active with probability p , we propose to separately sample $m_{\mathcal{I}}$ subsets for the image and $m_{\mathcal{T}}$ subsets for the text modality, i.e. to decompose $\mathbb{P}_p = \mathbb{P}_{p,\mathcal{I}} \otimes \mathbb{P}_{p,\mathcal{T}}$, and sample independently. This allows for generating all possible combinations, obtaining a novel estimator of $\mathfrak{F}_p$.

Definition 7 (Cross-modal estimator). *Let $M_{\mathcal{I}}^{(1)}, \dots, M_{\mathcal{I}}^{(m_{\mathcal{I}})} \stackrel{iid}{\sim} \mathbb{P}_{p,\mathcal{I}}$ and $M_{\mathcal{T}}^{(1)}, \dots, M_{\mathcal{T}}^{(m_{\mathcal{T}})} \stackrel{iid}{\sim} \mathbb{P}_{p,\mathcal{T}}$. We define the cross-modal estimator as $\hat{\mathbf{e}}^{\text{FixLIP-}p} := \arg\min_{\mathbf{e}} \hat{\mathfrak{F}}_p^{(m_{\mathcal{I}},m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})$, where*

$$\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}},m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}}) := \frac{1}{m_{\mathcal{I}} \cdot m_{\mathcal{T}}} \sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} (\nu(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})} \cup M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}) - \hat{\nu}_{\mathbf{e}}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})} \cup M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}))^2.$$

Crucially, $\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}},m_{\mathcal{T}})}$ can be computed efficiently because f independently encodes images for $m_{\mathcal{I}}$ samples with $f_{\mathcal{I}}$ and texts for $m_{\mathcal{T}}$ samples with $f_{\mathcal{T}}$ (cf. Equation 1). Then, similarity predictions are computed in a vectorized manner between all the cross-modal combinations of masked inputs. In theory, given $m = m_{\mathcal{I}} + m_{\mathcal{T}}$ samples, the model-agnostic estimator obtains m game values (similarity predictions), while our proposed cross-modal estimator obtains $m_{\mathcal{I}} \cdot m_{\mathcal{T}} \gg m$. The empirically observed computational speedup, which we analyze in Section 5.4, can depend on the model's batch size (closely related to the GPU VRAM hardware). Note that reusing the samples in text and image modalities inflicts dependence between them, and the following result establishes important theoretical properties.

Theorem 2. *The cross-modal estimator $\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}},m_{\mathcal{T}})}$ is unbiased, and its variance is bounded by*

$$\mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}},m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] \leq \mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}},m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] \leq \frac{m_{\mathcal{I}} + m_{\mathcal{T}}}{m_{\mathcal{I}} \cdot m_{\mathcal{T}}} \mathbb{V}[(\nu(M_{\mathcal{I}} \cup M_{\mathcal{T}}) - \hat{\nu}_{\mathbf{e}}(M_{\mathcal{I}} \cup M_{\mathcal{T}}))^2].$$

In other words, Theorem 2 shows that the variance of the cross-modal estimator is at most of the same order as the variance of the model-agnostic estimator with $m \approx m_{\mathcal{I}} \approx m_{\mathcal{T}}$ samples. Moreover, the variance is at least the variance of the model-agnostic estimator with $m_{\mathcal{I}} \cdot m_{\mathcal{T}}$ independent samples. In practice, we use the cross-modal estimator in higher budget scenarios to speed up the computation by decreasing the number of effective model inferences.

3.3 Large-scale adaptations for FIXLIP explanations

In contrast to first-order explanations, the explanation basis of Definition 1 grows quadratically as $|\mathcal{B}| = 1 + n_{\mathcal{I}} + n_{\mathcal{T}} + \binom{n_{\mathcal{I}}+n_{\mathcal{T}}}{2} \approx (n_{\mathcal{I}} + n_{\mathcal{T}})^2$. For example, a ViT-B/16 version of CLIP with 196 image and 30 text input tokens (players) results already in 25 425 interactions. From a practical perspective, we thus consider two heuristics as part of our complete methodology: (1) We employ a *two-step filtering approach* for prioritizing interactions to approximate when explaining games with a large number of players. A default is to pick a clique subset with the highest (absolute) first-order attributions and compute interactions between them, e.g. top-72 with $\binom{72}{2} = 2 556$ interactions. Another approach is to include only cross-modal interactions in the approximation, e.g. $196 \times 30 = 5 880$. (2) We apply a simple *greedy subset selection* algorithm on explanation $\mathbf{e}$ to find subgraphs $M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}$ of the highest and lowest sums of $\hat{\nu}_{\mathbf{e}}(M)$, which we use in evaluating (Section 5.1) and visualizing (Section 5.5) explanations. Further details are in Appendix B.

4 Evaluation Metrics for Interaction Explanations of VLE Predictions

In this section, we derive three evaluation metrics for explanations that may include second-order interactions. First, we evaluate the faithfulness of the approximation $\hat{\nu}_e$ to ν , since the basis of explaining VLEs is the masking operator that defines the FIXLIP game ν . Note that conventionally this is done with an R^2 coefficient [33, 50] or mean squared error [16, 19, 52], but we need to rely on rankings to compare attribution methods like saliency maps. Second, we generalize the area between insertion and deletion curves [AID, 25, 77], aka remove least/most important first [6], a popular faithfulness metric for token attribution methods. Finally, we extend the well-established pointing game evaluation of faithfulness from the image classification literature [3, 7, 8, 48]. In the latter two metrics, we need to assume $\mathbf{e}_{\{i,j\}} = 0$ for the first-order attribution explanations without interactions.

We validate the faithfulness to the FIXLIP game ν across sampled subsets. To this end, we compare the rankings of ν and $\hat{\nu}$, since related explanation methods are not normalized to estimate ν correctly.

Definition 8 (*p*-faithfulness correlation). *We define p-faithfulness correlation as the Spearman's rank correlation computed between ν and $\hat{\nu}_e$ evaluated for m masks sampled from $\mathbb{P}_p$ (cf. Remark 1) as*

$$\mu_{\text{corr}}(\mathbf{e}; p) := \text{correlation}_{M^{(1)}, \dots, M^{(m)} \stackrel{\text{iid}}{\sim} \mathbb{P}_p} (\nu(M^{(i)}), \hat{\nu}_e(M^{(i)})).$$

The rank correlation metric μ_{corr} yields insights into the general capability of recovering the ranking of similarity scores of uniformly masked inputs based on $\mathbb{P}_p$.

Beyond capturing the general *p*-faithfulness, we derive an insertion/deletion visualization with a metric measuring the explanation's ability to identify token subsets (masks) of high and low similarity as evaluated by the game (model).

Definition 9 (Area between insertion/deletion curves (AID)). *For each size $k = 1, \dots, n_{\mathcal{I}} + n_{\mathcal{T}}$, let $M_{\mathbf{e}, \min, k} := \arg\min_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}: |M|=k} \hat{\nu}_e(M)$ and $M_{\mathbf{e}, \max, k} := \arg\max_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}: |M|=k} \hat{\nu}_e(M)$. Insertion and deletion curves are computed as $\mathcal{C}_{\text{insert}}(\mathbf{e}) := \{(k, \nu(M_{\mathbf{e}, \max, k}))\}_{k=1, \dots, n_{\mathcal{I}} + n_{\mathcal{T}}}$, and $\mathcal{C}_{\text{delete}}(\mathbf{e}) := \{(k, \nu(M_{\mathbf{e}, \min, n_{\mathcal{I}} + n_{\mathcal{T}} - k}))\}_{k=1, \dots, n_{\mathcal{I}} + n_{\mathcal{T}}}$, respectively. The metric is then defined as*

$$\mu_{\text{AID}}(\mathbf{e}) := \int \mathcal{C}_{\text{insert}}(\mathbf{e}) - \int \mathcal{C}_{\text{delete}}(\mathbf{e}) = \sum_{k=1}^{n_{\mathcal{I}} + n_{\mathcal{T}}} (\nu(M_{\mathbf{e}, \max, k}) - \nu(M_{\mathbf{e}, \min, k})).$$

μ_{AID} measures the difference between model predictions when including and excluding the most relevant tokens. For first-order methods, $M_{\mathbf{e}, \min, k}$ and $M_{\mathbf{e}, \max, k}$ are directly found by ranking $\mathbf{e}$.

Proposition 1. *For token attribution values, i.e. when $\forall_{i,j} \mathbf{e}_{\{i,j\}} \equiv 0$, such as weighted Banzhaf values, $M_{\mathbf{e}, \max, k}$ and $M_{\mathbf{e}, \min, k}$ are given by the top- k and bottom- k coefficients of $\mathbf{e}$, respectively. Consequently, $\mathcal{C}_{\text{delete}}$ is found by deleting $M_{\mathbf{e}, \max, k}$, i.e. $M_{\mathbf{e}, \min, n_{\mathcal{I}} + n_{\mathcal{T}} - k} = (N_{\mathcal{I}} \cup N_{\mathcal{T}}) \setminus M_{\mathbf{e}, \max, k}$.*

Consequently, $\mathcal{C}_{\text{insert}}$ and $\mathcal{C}_{\text{delete}}$ generalize insertion and deletion curves to second-order interaction explanations. Note that for token attribution values, we have $M_{\mathbf{e}, \max, k} \subset M_{\mathbf{e}, \max, k+1}$, and $M_{\mathbf{e}, \min, k} \subset M_{\mathbf{e}, \min, k+1}$, which does not hold when modeling interaction terms. We give a further description of the insertion/deletion process on a single image–text input in Appendix B.2, Figure 9.

Lastly, we propose to evaluate an explanation $\mathbf{e}$ with designed pseudo ground-truth data. We measure the overlap between an explanation of inputs crafted as four combined images with a multi-object text prompt, and its “ground-truth” assumed prior. We show a visual example in Appendix B.3, Figure 8

Definition 10 (Pointing game recognition (PGR)). *For $\mathbf{e}$, let $\mathbf{e}_{\text{in}, k}$ and $\mathbf{e}_{\text{out}, k}$ denote the values of interactions belonging to, and not belonging to, object k – a text token with its corresponding image patches. We denote the positive and negative interactions by $\mathbf{e}_{>0}$, $\mathbf{e}_{<0}$. The metric is then defined as*

$$\mu_{\text{PGR}}(\mathbf{e}; N_{\mathcal{T}}) := \frac{\sum_{k \in N_{\mathcal{T}}} (\|\mathbf{e}_{\text{in}, k, >0}\|_1 + \|\mathbf{e}_{\text{out}, k, <0}\|_1)}{\sum_{k \in N_{\mathcal{T}}} (\|\mathbf{e}_{\text{in}, k}\|_1 + \|\mathbf{e}_{\text{out}, k}\|_1)} \in [0, 1], \text{ where } \|\mathbf{e}\|_1 := \sum_{i,j} |\mathbf{e}_{\{i,j\}}|.$$

μ_{PGR} quantifies the ratio of absolute values of “correctly” identified cross-modal interaction terms to total interaction terms. Thereby, “correctly” identified cross-modal interactions refer to positive and negative interactions that are associated with and without object k in an image, respectively. We use μ_{PGR} as a sanity check that interaction explanations are essential for explaining VLEs as compared to alternative attribution approaches. Scoring high PGR values across multiple objects present in an image–text pair denotes that an explanation method can show the model distinguishes between them.




Figure 2: Visual comparison between FIXLIP and baselines. First-order attribution methods, e.g. GAME and Grad-ECLIP, lack the tools to faithfully explain complex similarity predictions of vision–language encoders like CLIP. Notably, in this example, the text token ant is the most important for the similarity prediction. One of the differences from exCLIP is that we include intra-modal and main effects in the approximation, which are crucial for obtaining faithful interaction explanations.

5 Experiments

Setup. In experiments, we empirically validate the performance of FIXLIP with three metrics defined in Section 4, measure its computational efficiency, and demonstrate its utility in visual explanation of VLEs. We mainly use the openly available pre-trained CLIP models [57] of two sizes: ViT-B/32 with 7×7 image patches and ViT-B/16 with 14×14 . Moreover, we demonstrate the broader applicability of FIXLIP to explain SigLIP [75] and SigLIP-2 [66] up to the ViT-L/16 variant with 16×16 patches. We rely on two openly available datasets commonly used in explainability research: MS COCO [42] and ImageNet-1k [15]; the latter specifically to design the pointing game evaluation considering zero-shot classification.

In quantitative evaluation, we compare FIXLIP to a few representative baselines: the first attribution method for CLIP abbreviated as GAME [12], a state-of-the-art first-order attribution method Grad-ECLIP [77], and a recently released second-order interaction method exCLIP [48]. For comparisons of Grad-ECLIP and exCLIP with previous baselines, refer to the appropriate benchmarks [48, 77]. Figure 2 conceptually compares FIXLIP to baselines using the image–text input example from Figure 1. Furthermore, we naturally demonstrate the improvement between FIXLIP and first-order Shapley/Banzhaf values. In that, we effectively scale FIXLIP to approximate second-order explanations with a budget of over 10^6 model inferences, which far exceeds related work. For FIXLIP- p , we use the cross-modal estimator with a budget of 2^{21} , whereas FIXLIP with Shapley interactions uses the model-agnostic estimator with a budget of 2^{17} , yielding approximately similar runtime. We mainly experiment with $p \in \{0.3, 0.5, 0.7\}$; note that there is no computational overhead for different p settings. Further details for the experimental setup are provided in Appendix C.

5.1 FIXLIP outperforms baselines as measured with insertion/deletion curves

Figure 3 shows insertion/deletion results for different explanation methods, where each line and metric value is an average ($\pm$ sd.) over 1000 inputs. We observe that FIXLIP faithfully recovers the nonlinear importance rankings of token subsets (see Appendix Figure 9 for a visual example). It not only finds the most important tokens whose deletion results in a significant drop in similarity (y-axis), but also finds the least important tokens whose deletion may even result in the prediction’s increase, contrary to gradient-based methods. The general conclusion is consistent for SigLIP-2 and a larger ViT-B/16 model (refer to Appendix D for additional results), albeit the method’s faithfulness drops when increasing the size of input tokens, where further scaling to even higher image resolutions is a natural future work direction. As designed, increasing the masking weight to $p = 0.7$ leads to an improved faithfulness in the 0–50% range of input deleted (100–50% inserted), while decreasing to $p = 0.3$ improves faithfulness with 50–100% input deleted (50–0% inserted). Note that, in theory, all of the methods can obtain negative normalized values on the y-axis. We think that exCLIP fails to recover the appropriate ranking because it only approximates cross-modal interactions between tokens from the two modalities, omitting first-order effects and intra-modal interactions in the process (effectively constructing a bipartite weighted graph without weights in nodes).




Figure 3: Insertion/deletion curves for CLIP (ViT-B/32) on MS COCO. AID score (higher is better) for FIXLIP against alternative explanation methods, where a random baseline scores 0. The y-axis is normalized between the model’s prediction on the original input (100%) and the fully removed one (0%), where negative values denote that the model is predicting the image–text inputs are unsimilar. It means the similarity prediction on a partially masked input is smaller than the prediction on the fully masked input. Methods such as Grad-ECLIP and exCLIP fail to recover nonlinear rankings of important tokens, while our method faithfully recovers the optimal subset explanation. Extended results for CLIP (ViT-B/16) and SigLIP-2 (ViT-B/32) are in Figures 10 & 11.

Table 1: Pointing game recognition for CLIP (ViT-B/32) on ImageNet-1k. PGR score (higher is better) for FIXLIP against alternative explanation methods, where a random baseline scores 0.25. First-order methods, such as Grad-ECLIP and Shapley values, fail to distinguish between multiple objects at once, while second-order methods faithfully recover the appropriate explanation (up to the pointing game’s irreducible non-optimality). Extended results for CLIP (ViT-B/16) are in Table 4.

Explanation Method	Recognition ($\uparrow$)			
	1 object	2 objects	3 objects	4 objects
GAME [12]	.61 $\pm$.12	.43 $\pm$.03	.33 $\pm$.02	.28 $\pm$.01
Grad-ECLIP [77]	.68 $\pm$.15	.45 $\pm$.04	.33 $\pm$.02	.28 $\pm$.01
Shapley values	.70 $\pm$.11	.56 $\pm$.06	.46 $\pm$.05	.37 $\pm$.04
Banzhaf values	.64 $\pm$.12	.52 $\pm$.06	.43 $\pm$.05	.35 $\pm$.04
exCLIP [48]	.73 $\pm$.20	.88 $\pm$.08	.89 $\pm$.06	.92 $\pm$.05
FIXLIP (Shapley interactions)	.83 $\pm$.10	.82 $\pm$.08	.84 $\pm$.06	.86 $\pm$.06
FIXLIP (w. Banzhaf interactions $p = 0.3$)	.78 $\pm$.13	.78 $\pm$.11	.80 $\pm$.08	.81 $\pm$.07
FIXLIP (w. Banzhaf interactions $p = 0.5$)	.81 $\pm$.12	.80 $\pm$.09	.81 $\pm$.07	.83 $\pm$.06
FIXLIP (w. Banzhaf interactions $p = 0.7$)	.83 $\pm$.12	.81 $\pm$.08	.83 $\pm$.07	.85 $\pm$.06

5.2 First-order attribution methods fail to pass a sanity check with a pointing game

Table 1 shows PGR results for different explanation methods, where each metric value is an average ($\pm$ sd.) over 500 inputs. We observe that first-order methods, e.g. Banzhaf values, fail to discriminate between multiple objects in an image. Saliency maps can only highlight the important part of image–text inputs without disentangling the complex relationship between each separate word in a caption and the corresponding image regions (Appendix Figure 8). Second-order methods such as FIXLIP and exCLIP pass our proposed sanity check; the latter scores slightly higher, as it specializes in approximating cross-modal interactions. The conclusion is consistent for CLIP (ViT-B/16). We further analyze the potential application of a pointing game to evaluate different models in Section 5.6.

5.3 FIXLIP recovers a faithful decomposition of the similarity function

Figure 4 demonstrates p -faithfulness results for different explanation methods, where each boxplot represents statistics aggregated from 1000 inputs. We confirm that FIXLIP optimizes faithfulness, which depends on the parameter p . Weighted Banzhaf interactions allow for precise controllability of the faithfulness optimization, whereas Shapley interactions perform well on average. Alternative




Figure 4: p -faithfulness correlation for CLIP (ViT-B/32) on MS COCO. Correlation for different variants of FixLIP against other explanation methods (left). Game-theoretical approaches can also be evaluated with the R^2 coefficient (right). Extended results for CLIP (ViT-B/16) are in Figure 12.

attribution methods compute explanations unfaithful to the explained similarity function ($\mu_{\text{corr}} \approx 0.5$). Extended results for different models and p values are given in Appendix D.

5.4 On the computational efficiency of the FixLIP cross-modal estimator

Figure 5 demonstrates that the FixLIP cross-modal estimator achieves over $20\times$ speedup when considering model inference time, i.e. game evaluations, and about $5\times$ speedup when accounting for the entire explanation pipeline implemented in Python. We did not necessarily optimize the latter, i.e. subset sampling, weighted least squares optimization, and other processing steps, expecting that the visible gap between the explanation and game time could be further decreased. For context, related first-order attribution methods take about 1 second to compute; see Table 5 in [77], but note that it could be only for an image explanation without text attribution.




Figure 5: Computation time vs. budget for the FixLIP explanation of SigLIP-2 (ViT-B/32), including game evaluations (model inference).

5.5 Visual explanation of vision-language encoders

We have already established that FixLIP delivers state-of-the-art faithfulness performance across three diverse metrics. One of the explanations' applications is to interpret the model's output function. Figure 6 demonstrates three types of visualizations that we envision can be used to broaden our understanding of VLEs. In the interest of space, we provide further interesting examples in Appendix E. Specifically, we provide a comprehensive guide on interpreting interaction explanations like FixLIP in Appendix E.1. Figures 8 & 9 further compare FixLIP to baselines in the context of evaluation metrics. Figure 16 shows a comparison between FixLIP of CLIP and SigLIP-2 (ViT-B/32), while Figure 17 shows a comparison between the ViT-B/32 and ViT-B/16 versions of CLIP.

5.6 Model-agnostic comparison of different vision-language encoder architectures

We use FixLIP to compare different VLE architectures with PGR in Table 2. **Surprisingly, SigLIP-2 can be more faithfully explained with cross-modal interactions than CLIP**, for both model sizes. We observe that, in general, smaller models (ViT-B/32) can be more faithfully explained than larger ones (ViT-L/16), which is consistent with prior work [9, 16]. For more visual comparisons, see Appendix E.3.

Table 2: Pointing game results as measured with FixLIP.

Size	Model	Recognition ($\uparrow$)			
		1 object	2 objects	3 obj.	4 obj.
ViT-B/32	CLIP	.83 $\pm$.11	.81 $\pm$.07	.82 $\pm$.07	.85 $\pm$.05
	SigLIP-2	.90 $\pm$.07	.90 $\pm$.05	.89 $\pm$.05	.89 $\pm$.04
ViT-B/16	CLIP	.81 $\pm$.09	.81 $\pm$.06	.81 $\pm$.05	.82 $\pm$.04
	SigLIP	.80 $\pm$.08	.82 $\pm$.06	.84 $\pm$.05	.84 $\pm$.05
ViT-L/16	SigLIP	.81 $\pm$.08	.84 $\pm$.05	.85 $\pm$.04	.84 $\pm$.04
	SigLIP-2	.80 $\pm$.10	.85 $\pm$.06	.84 $\pm$.05	.85 $\pm$.04




Figure 6: FIXLIP facilitates various approaches to model understanding. (A) Interaction explanations allow answering the pivotal question: *Why is it similar for the model?* Here, the strongest interaction is between text token `doll` and image patch saying “dollar”. One could say the model is right for the wrong reasons. (B) Each token can be selected for conditioning to visualize as a heatmap only the interactions (edges) outgoing from it. (C) A complete graph can be traversed to find subsets of high positive and low negative similarity as approximated by the explanation.

6 Discussion

As vision–language encoders are increasingly deployed in real-world applications, it becomes pivotal to ensure that their predictions are explainable. To this end, we introduced faithful interaction explanations of CLIP and SigLIP models, offering a unique perspective on interpreting image–text similarity predictions. Moreover, we derived three evaluation criteria facilitating future work in this direction.

Limitations and future work. Our work faces three limitations, each with a clear path for future development. Although FIXLIP allows scaling to larger models with hundreds of players and efficiently computing a million model inferences, improvements could be made to its practical implementation. Specifically, we envision exploring the use of sparse linear regression, applying the Archipelago framework [65] to filter interactions for the approximation, and a non-greedy algorithm for a closer-to-optimal subset selection. Second, the visual properties and usability of FIXLIP should be further studied as a potential direction in human–computer interaction research [58], see e.g. a user study of cross-modal interaction explanations by Liang et al. [40]. Finally, our work is restricted to second-order interaction explanations, while efficiently approximating higher-order interactions (also in tri-modal settings beyond LIP models) becomes an interesting challenge to overcome.

Broader impact. We believe FIXLIP can empower model developers in debugging VLEs, understanding their similarity predictions, and finding unwanted biases in image–text data. Especially when these models are used in high-stakes decision making, like in the case of medical applications.

Code. We provide additional details on reproducibility in the Appendix, as well as the code to reproduce all experiments in this paper is available at <https://github.com/hbaniecki/fixlip>.

Acknowledgments and Disclosure of Funding

We gratefully acknowledge the Polish high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH) for providing computer facilities and support within the computational grant no. PLG/2025/018330. Hubert Baniecki was supported from the state budget within the Polish Ministry of Education and Science program “Pearls of Science” project number PN/01/0087/2022. Barbara Hammer, Eyke Hüllermeier, Fabian Fumagalli, and Maximilian Muschalik gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation): TRR 318/1 2021 – 438445824.

References

- [1] Chirag Agarwal. Rethinking explainability in the era of multimodal AI. *preprint, arXiv:2506.13060*, 2025. 1
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, et al. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022. 1
- [3] Shreyash Arya, Sukrut Rao, Moritz Böhle, and Bernt Schiele. B-cosification: Transforming deep neural networks to be inherently interpretable. In *NeurIPS*, 2024. 4
- [4] Sriram Balasubramanian, Samyadeep Basu, and Soheil Feizi. Decomposing and interpreting image representations via text in ViTs beyond CLIP. In *NeurIPS*, 2024. 2
- [5] Usha Bhalla, Alex Oesterling, Suraj Srinivas, Flavio P. Calmon, and Himabindu Lakkaraju. Interpreting CLIP with sparse linear concept embeddings (SpLiCE). In *NeurIPS*, 2024. 2
- [6] Stefan Bluecher, Johanna Vielhaben, and Nils Strodthoff. Decoupling pixel flipping and occlusion strategy for consistent XAI benchmarks. *Transactions on Machine Learning Research*, 2024. 1, 4
- [7] Moritz Bohle, Mario Fritz, and Bernt Schiele. Convolutional dynamic alignment networks for interpretable classifications. In *CVPR*, 2021. 4
- [8] Moritz Bohle, Navdeppal Singh, Mario Fritz, and Bernt Schiele. B-cos alignment for inherently interpretable CNNs and vision transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(06):4504–4518, 2024. 4
- [9] Sebastian Bordt and Ulrike von Luxburg. From shapley values to generalized additive models and back. In *AISTATS*, 2023. 2, 5, 6
- [10] Yue Cao, Yun Xing, Jie Zhang, Di Lin, Tianwei Zhang, Ivor Tsang, Yang Liu, and Qing Guo. SceneTAP: Scene-coherent typographic adversarial planner against vision-language models in real-world environments. In *CVPR*, 2025. E, 2
- [11] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2
- [12] Hila Chefer, Shir Gur, and Lior Wolf. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *ICCV*, 2021. 1, 2, 5, 1, C.3, 4
- [13] Sanghyuk Chun, Wonjae Kim, Song Park, and Sangdoo Yun. Probabilistic language-image pre-training. In *ICLR*, 2025. 1
- [14] Ian Covert, Tony Sun, James Zou, and Tatsunori Hashimoto. Locality alignment improves vision-language models. In *ICLR*, 2025. 1
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 5, C.2
- [16] James Enouen and Yan Liu. InstaSHAP: Interpretable additive models explain Shapley values instantly. In *ICLR*, 2025. 2, 4, 5, 6
- [17] Gabriel Erion, Joseph D Janizek, Pascal Sturmfels, Scott M Lundberg, and Su-In Lee. Improving performance of deep learning models with axiomatic attribution priors and expected gradients. *Nature Machine Intelligence*, 3(7):620–631, 2021. B.1
- [18] Katsushige Fujimoto, Ivan Kojadinovic, and Jean-Luc Marichal. Axiomatic characterizations of probabilistic and cardinal-probabilistic interaction indices. *Games and Economic Behavior*, 55(1):72–99, 2006. 2
- [19] Fabian Fumagalli, Maximilian Muschalik, Patrick Kolpaczki, Eyke Hüllermeier, and Barbara Eva Hammer. SHAP-IQ: Unified approximation of any-order Shapley interactions. In *NeurIPS*, 2023. 1, 2, 4, B.1

- [20] Fabian Fumagalli, Maximilian Muschalik, Patrick Kolpaczki, Eyke Hüllermeier, and Barbara Hammer. KernelSHAP-IQ: Weighted least square optimization for Shapley interactions. In *ICML*, 2024. [2](#)
- [21] Yossi Gandelsman, Alexei A. Efros, and Jacob Steinhardt. Interpreting CLIP’s image representation via text-based decomposition. In *ICLR*, 2024. [2](#)
- [22] Gabriel Goh, Nick Cammarata, Chelsea Voss, Shan Carter, Michael Petrov, Ludwig Schubert, Alec Radford, and Chris Olah. Multimodal neurons in artificial neural networks. *Distill*, 6(3): e30, 2021. [2](#)
- [23] Yichen Gong, DeLong Ran, Jinyuan Liu, Conglei Wang, Tianshuo Cong, Anyu Wang, Sisi Duan, and Xiaoyun Wang. FigStep: Jailbreaking large vision-language models via typographic visual prompts. In *AAAI*, 2025. [E.2](#)
- [24] Michel Grabisch. *Set Functions, Games and Capacities in Decision Making*. Springer, 2016. [3.1](#), [A.1](#)
- [25] Naofumi Hama, Masayoshi Mase, and Art B. Owen. Deletion and insertion tests in regression models. *Journal of Machine Learning Research*, 24(290):1–38, 2023. [1](#), [4](#)
- [26] Peter L. Hammer and Ron Holzman. Approximations of pseudo-boolean functions; applications to game theory. *ZOR Mathematical Methods of Operations Research*, 36(1):3–21, 1992. [3.1](#)
- [27] Peter Hase, Harry Xie, and Mohit Bansal. The out-of-distribution problem in explainability and search methods for feature importance explanations. In *NeurIPS*, 2021. [1](#), [B.1](#)
- [28] Jack Hessel and Lillian Lee. Does my multimodal model learn cross-modal interactions? It’s harder to tell than you might think! In *EMNLP*, 2020. [2](#)
- [29] Wisdom Oluchi Ikezogwo, Mehmet Saygin Seyfioglu, Fatemeh Ghezloo, Dylan Stefan Chan Geva, Fatwir Sheikh Mohammed, Pavan Kumar Anand, Ranjay Krishna, and Linda Shapiro. Quilt-1M: One million image-text pairs for histopathology. In *NeurIPS*, 2023. [1](#)
- [30] Joseph D. Janizek, Pascal Sturmfels, and Su-In Lee. Explaining explanations: Axiomatic feature interactions for deep networks. *Journal of Machine Learning Research*, 22(104):1–54, 2021. [2](#)
- [31] Peng Jin, Hao Li, Li Yuan, Shuicheng Yan, and Jie Chen. Hierarchical Banzhaf interaction for general video-language representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. [2](#)
- [32] Boris Joukovsky, Fawaz Sammani, and Nikos Deligiannis. Model-agnostic visual explanations via approximate bilinear models. In *ICIP*, 2023. [1](#), [2](#), [C.3](#)
- [33] Justin Singh Kang, Yigit Efe Erginbas, Landon Butler, Ramtin Pedarsani, and Kannan Ramchandran. Learning to understand: Identifying interactions via the Möbius transform. In *NeurIPS*, 2024. [2](#), [4](#)
- [34] Justin Singh Kang, Landon Butler, Abhineet Agarwal, Yigit Efe Erginbas, Ramtin Pedarsani, Bin Yu, and Kannan Ramchandran. SPEX: Scaling feature interaction explanations for LLMs. In *ICML*, 2025. [2](#)
- [35] Martha Lewis, Nihal Nayak, Peilin Yu, Jack Merullo, Qinan Yu, Stephen Bach, and Ellie Pavlick. Does CLIP bind concepts? Probing compositionality in large image models. In *Findings of EACL*, 2024. [1](#)
- [36] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. [1](#)
- [37] Weida Li and Yaoliang Yu. Robust data valuation with weighted Banzhaf values. In *NeurIPS*, 2023. [2](#)
- [38] Weida Li and Yaoliang Yu. One sample fits all: Approximating all probabilistic values simultaneously and efficiently. In *NeurIPS*, 2024. [2](#)

- [39] Yi Li, Hualiang Wang, Yiqun Duan, Jiheng Zhang, and Xiaomeng Li. A closer look at the explainability of contrastive language-image pre-training. *Pattern Recognition*, 162:111409, 2025. 1, 2
- [40] Paul Pu Liang, Yiwei Lyu, Gunjan Chhablani, Nihal Jain, Zihao Deng, Xingbo Wang, Louis-Philippe Morency, and Ruslan Salakhutdinov. MultiViz: Towards visualizing and understanding multimodal models. In *ICLR*, 2023. 1, 2, 6
- [41] Hyesu Lim, Jinho Choi, Jaegul Choo, and Steffen Schneider. Sparse autoencoders reveal selective remapping of visual concepts during adaptation. In *ICLR*, 2025. 2
- [42] 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 *ECCV*, 2014. 5, C.2
- [43] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 1
- [44] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NeurIPS*, 2017. 1, 3.1
- [45] Daniel D Lundstrom, Tianjian Huang, and Meisam Razaviyayn. A rigorous study of integrated gradients method and extensions to internal neuron attributions. In *ICML*, 2022. B.1
- [46] Yiwei Lyu, Paul Pu Liang, Zihao Deng, Ruslan Salakhutdinov, and Louis-Philippe Morency. DIME: Fine-grained interpretations of multimodal models via disentangled local explanations. In *AIES*, 2022. 2
- [47] Jean-Luc Marichal and Pierre MATHONET. Weighted Banzhaf power and interaction indexes through weighted approximations of games. *European Journal of Operations Research*, 211(2): 352–358, 2011. 1, 3.1, A.1
- [48] Lucas Moeller, Pascal Tilli, Thang Vu, and Sebastian Padó. Explaining caption-image interactions in CLIP models with second-order attributions. *Transactions on Machine Learning Research*, 2025. 2, 4, 5, 1, C.3, 4
- [49] Niels Morch, Ulrik Kjems, Lars Kai Hansen, Claus Svarer, Ian Law, Benny Lautrup, Stephen Strother, and Kelly Rehm. Visualization of neural networks using saliency maps. In *ICNN*, 1995. 1
- [50] Maximilian Muschalik, Hubert Baniecki, Fabian Fumagalli, Patrick Kolpaczki, Barbara Hammer, and Eyke Hüllermeier. shapiq: Shapley interactions for machine learning. In *NeurIPS*, 2024. 2, 4
- [51] Maximilian Muschalik, Fabian Fumagalli, Paolo Frazzetto, Janine Strother, Luca Hermes, Alessandro Sperduti, Eyke Hüllermeier, and Barbara Hammer. Exact computation of any-order Shapley interactions for graph neural networks. In *ICLR*, 2025. 3
- [52] Christopher Musco and R. Teal Witter. Provably accurate shapley value estimation via leverage score sampling. In *ICLR*, 2025. 2, 4
- [53] Letitia Parcalabescu and Anette Frank. MM-SHAP: A performance-agnostic metric for measuring multimodal contributions in vision and language models & tasks. In *ACL*, 2023. 2
- [54] Letitia Parcalabescu and Anette Frank. Do vision & language decoders use images and text equally? How self-consistent are their explanations? In *ICLR*, 2025. 2
- [55] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *CVPR*, 2018. 2
- [56] Guilherme Dean Pelegrina, Leonardo Tomazeli Duarte, and Michel Grabisch. A k-additive Choquet integral-based approach to approximate the SHAP values for local interpretability in machine learning. *Artificial Intelligence*, 325:104014, 2023. 2

- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 5, C.1
- [58] Yao Rong, Tobias Leemann, Thai-Trang Nguyen, Lisa Fiedler, Peizhu Qian, Vaibhav Unhelkar, Tina Seidel, Gjergji Kasneci, and Enkelejda Kasneci. Towards human-centered explainable AI: A survey of user studies for model explanations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4):2104–2122, 2024. 6
- [59] Benedek Rozemberczki, Lauren Watson, Péter Bayer, Hao-Tsung Yang, Olivér Kiss, Sebastian Nilsson, and Rik Sarkar. The Shapley value in machine learning. In *IJCAI*, 2022. 2
- [60] Fawaz Sammani and Nikos Deligiannis. Interpreting and analysing CLIP’s zero-shot image classification via mutual knowledge. In *NeurIPS*, 2024. 2
- [61] Fawaz Sammani, Boris Joukovsky, and Nikos Deligiannis. Visualizing and understanding contrastive learning. *IEEE Transactions on Image Processing*, 33:541–555, 2024. 2, C.3
- [62] Mukund Sundararajan, Kedar Dhamdhere, and Ashish Agarwal. The Shapley Taylor interaction index. In *ICML*, 2020. 2
- [63] Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal LLMs. In *CVPR*, 2024. 1, E.2
- [64] Che-Ping Tsai, Chih-Kuan Yeh, and Pradeep Ravikumar. Faith-Shap: The faithful Shapley interaction index. *Journal of Machine Learning Research*, 24(94):1–42, 2023. 1, 2, 3.1, 2, 3.1
- [65] Michael Tsang, Sirisha Rambhatla, and Yan Liu. How does this interaction affect me? Interpretable attribution for feature interactions. In *NeurIPS*, 2020. 6
- [66] Michael Tschannen et al. SigLIP 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *preprint arXiv:2502.14786*, 2025. 1, 5, C.1
- [67] Vishaal Udandaraao, Max F Burg, Samuel Albanie, and Matthias Bethge. Visual data-type understanding does not emerge from scaling vision-language models. In *ICLR*, 2024. 1
- [68] Bo Wan, Michael Tschannen, Yongqin Xian, Filip Pavetic, Ibrahim Alabdulmohsin, Xiao Wang, André Susano Pinto, Andreas Peter Steiner, Lucas Beyer, and Xiaohua Zhai. LocCa: Visual pretraining with location-aware captioners. In *NeurIPS*, 2024. 1
- [69] Xingbo Wang, Jianben He, Zhihua Jin, Muqiao Yang, Yong Wang, and Huamin Qu. M2Lens: Visualizing and explaining multimodal models for sentiment analysis. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):802–812, 2022. 2
- [70] Ying Wang, Tim G. J. Rudner, and Andrew G Wilson. Visual explanations of image-text representations via multi-modal information bottleneck attribution. In *NeurIPS*, 2023. 1, 2
- [71] Jialin Wu and Raymond Mooney. Faithful multimodal explanation for visual question answering. In *ACL Workshop BlackboxNLP*, 2019. 2
- [72] Rui Xiao, Sanghwan Kim, Mariana-Iuliana Georgescu, Zeynep Akata, and Stephan Alaniz. FLAIR: VLM with fine-grained language-informed image representations. In *CVPR*, 2025. 1
- [73] Mert Yuksekgonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. When and why vision-language models behave like bags-of-words, and what to do about it? In *ICLR*, 2023. 1
- [74] Vladimir Zaigrajew, Hubert Baniecki, and Przemysław Biecek. Interpreting CLIP with hierarchical sparse autoencoders. In *ICML*, 2025. 2
- [75] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, 2023. 1, 5, C.1

- [76] Sheng Zhang et al. A multimodal biomedical foundation model trained from fifteen million image–text pairs. *NEJM AI*, 2(1), 2024. 1, B.1
- [77] Chenyang Zhao, Kun Wang, Janet H. Hsiao, and Antoni B. Chan. Grad-ECLIP: Gradient-based visual and textual explanations for CLIP. In *ICML*, 2024. 1, 2, 4, 5, 1, 5.4, C.3, 4
- [78] Zihao Zhao, Yuxiao Liu, Han Wu, Mei Wang, Yonghao Li, Sheng Wang, Lin Teng, Disheng Liu, Zhiming Cui, Qian Wang, and Dinggang Shen. CLIP in medical imaging: A survey. *Medical Image Analysis*, 102:103551, 2025. 1
- [79] Xu Zheng, Farhad Shirani, Zhuomin Chen, Chaohao Lin, Wei Cheng, Wenbo Guo, and Dongsheng Luo. F-Fidelity: A robust framework for faithfulness evaluation of explainable AI. In *ICLR*, 2025. 1
- [80] Zhiyu Zhu, Zhibo Jin, Jiayu Zhang, Nan Yang, Jiahao Huang, Jianlong Zhou, and Fang Chen. Narrowing information bottleneck theory for multimodal image-text representations interpretability. In *ICLR*, 2025. 2

Appendix

Table 3 summarises the mathematical notation used in the paper. In Appendix A, we derive proofs for Theorems 1 & 2, and Proposition 1. Appendix B describes the implementation details of our methods introduced in Sections 3 & 4. Appendix C provides further details on the setup of experiments conducted in Section 5. Finally, we discuss ablation results in Appendix D and visualize additional explanations in Appendix E.

A Proofs	18
A.1 Proof of Theorem 1	18
A.2 Proof of Theorem 2	19
A.3 Proof of Proposition 1	22
B Implementation Details	23
B.1 Faithful interaction explanations of LIP models (FixLIP)	23
B.2 Area between the insertion/deletion curves (AID)	24
B.3 Pointing game recognition (PGR)	24
C Experimental Setup	26
C.1 Models	26
C.2 Datasets	26
C.3 Baselines	26
C.4 Metrics	27
C.5 Compute resources	27
D Additional Experimental Results	28
E Visual Explanatons	31
E.1 Guidance on interpreting visualizations of interaction explanations	31
E.2 Additional examples	33
E.3 Comparison between CLIP (ViT-B/32) and SigLIP-2 (ViT-B/32)	35
E.4 Comparison between CLIP (ViT-B/32) and CLIP (ViT-B/16)	36

Table 3: Summary of notation.

Notation	Description
$n_{\mathcal{I}}, n_{\mathcal{T}} \in \mathbb{N}$	Image and text input dimensions (number of tokens)
$x_{\mathcal{I}} \in \mathbb{R}^{n_{\mathcal{I}}}, x_{\mathcal{T}} \in \mathbb{R}^{n_{\mathcal{T}}}$	Image and text input vectors
$f(x_{\mathcal{I}}, x_{\mathcal{T}})$	Vision–language encoder prediction for an image–text input pair
$d \in \mathbb{N}$	Encoder embedding dimension
$f_{\mathcal{I}}(x_{\mathcal{I}}) \in \mathbb{R}^d, f_{\mathcal{T}}(x_{\mathcal{T}}) \in \mathbb{R}^d$	Vision and language encoder embedding vectors
$N_{\mathcal{I}}, N_{\mathcal{T}}$	Sets of image and text token indices
i, j	Token indices
$\mathcal{B}$	Explanation basis set containing indices of tokens and token pairs
$\mathbf{e} \in \mathbb{R}^{ \mathcal{B} }$	Explanation object
$\mathbf{e}_i, \mathbf{e}_{\{i,j\}} \in \mathbb{R}$	Token attribution and interaction values
$M_o \subseteq N$ with $N \in \{N_{\mathcal{I}}, N_{\mathcal{T}}\}$	Subset of image or text token indices
$M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}$	Subset of image and text token indices
$b_{\mathcal{I}}, b_{\mathcal{T}}$	Baseline image and text vectors
$x \oplus_{M_o} b$	Input x masked using the operator for subset M_o and baseline b
$\nu(M)$	Game (value) function
$\hat{\nu}, \hat{\nu}_{\mathbf{e}}$	Estimators of the game (value) function
$p \in (0, 1)$	Probability parameter
$\mathfrak{F}_p(\nu, \hat{\nu})$	p -faithfulness error metric
$\mathbf{e}^{\text{FixLIP-}p} := \arg\min_{\mathbf{e}} \mathfrak{F}_p(\nu, \hat{\nu}_{\mathbf{e}})$	Explanation approximating the game function w.r.t. p -faithfulness
$\mathbb{P}_p(M) := p^{ M }(1-p)^{n_{\mathcal{I}}+n_{\mathcal{T}}- M }$	Probability distribution of sampled masks over $2^{N_{\mathcal{I}} \cup N_{\mathcal{T}}}$
m	Number of sampled mask subsets
$M^{(1)}, \dots, M^{(m)}$	Sampled mask subsets of image and text token indices
$\hat{\mathfrak{F}}_p^{(m)}(\nu, \hat{\nu}_{\mathbf{e}})$	Model-agnostic estimator of p -faithfulness
$\mathbb{P}_{p,\mathcal{I}}, \mathbb{P}_{p,\mathcal{T}}$	Probability distributions of sampled masks over $2^{N_{\mathcal{I}}}$ and $2^{N_{\mathcal{T}}}$
$m_{\mathcal{I}}, m_{\mathcal{T}}$	Numbers of sampled image and text mask subsets
$M^{(1)}, \dots, M^{(m_{\mathcal{I}})}$	Sampled mask subsets of image token indices
$M^{(1)}, \dots, M^{(m_{\mathcal{T}})}$	Sampled mask subsets of text token indices
$\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})$	Cross-modal estimator of p -faithfulness
$\mu_{\text{corr}}(\mathbf{e}; p)$	p -faithfulness Spearman’s rank correlation evaluation metric
$\mu_{\text{AID}}(\mathbf{e})$	Area between insertion/deletion curves evaluation metric
$M_{\mathbf{e}, \min, k}$	Mask subset of size k for explanation $\mathbf{e}$ minimizing $\hat{\nu}_{\mathbf{e}}(M_{\mathbf{e}, \min, k})$
$M_{\mathbf{e}, \max, k}$	Mask subset of size k for explanation $\mathbf{e}$ maximizing $\hat{\nu}_{\mathbf{e}}(M_{\mathbf{e}, \max, k})$
$C_{\text{insert}}(\mathbf{e})$	Insertion curve; a set of paired $(k, \nu(M_{\mathbf{e}, \max, k}))$ values
$C_{\text{delete}}(\mathbf{e})$	Deletion curve; a set of paired $(k, \nu(M_{\mathbf{e}, \min, n_{\mathcal{I}}+n_{\mathcal{T}}-k}))$ values
$\mu_{\text{PGR}}(\mathbf{e}; N_{\mathcal{T}})$	Pointing game recognition evaluation metric
$\mathbf{e}_{\text{in}, k}$	Interaction values belonging to text token k
$\mathbf{e}_{\text{out}, k}$	Interaction values <i>not</i> belonging to text token k
$\mathbf{e}_{>0}, \mathbf{e}_{<0}$	Negative and positive interaction values in an explanation
$\ \mathbf{e}\ _1 := \sum_{i,j} \mathbf{e}_{\{i,j\}} $	Absolute sum of an explanation

A Proofs

A.1 Proof of Theorem 1

Proof. To prove this result, we combine alternative representations of the cooperative games using the Möbius transform $a : 2^{N_{\mathcal{I}} \cup N_{\mathcal{T}}} \rightarrow \mathbb{R}$, which is defined by

$$a(M) := \sum_{L \subseteq M} (-1)^{|M|-|L|} \nu(L) \text{ for every } M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}.$$

The approximated game $\hat{\nu}_{\mathbf{e}}$ is a 2-additive game, and hence its Möbius transform is restricted up to second-order interactions [24]. Moreover, the Möbius transform of the approximated game $\hat{a}^{\text{FixLIP}}$ satisfies

$$\begin{aligned} \hat{a}^{\text{FixLIP}-p}(\emptyset) &= \hat{\nu}_{\mathbf{e}}(\emptyset) = \mathbf{e}_0, \\ \hat{a}^{\text{FixLIP}-p}(\{i\}) &= \hat{\nu}_{\mathbf{e}}(\{i\}) - \hat{\nu}_{\mathbf{e}}(\emptyset) = \mathbf{e}_i, \\ \hat{a}^{\text{FixLIP}-p}(\{i, j\}) &= \hat{\nu}_{\mathbf{e}}(\{i, j\}) - \hat{\nu}_{\mathbf{e}}(\{i\}) - \hat{\nu}_{\mathbf{e}}(\{j\}) + \hat{\nu}_{\mathbf{e}}(\emptyset) = \mathbf{e}_{\{i, j\}}, \end{aligned}$$

where $i, j \in N_{\mathcal{I}} \cup N_{\mathcal{T}}$ and $i \neq j$. The approximated game $\hat{\nu}$ is by construction a 2-additive game, and thus all higher-order Möbius coefficients are zero [24]. It was further shown by Marichal and Mathonet [47, Proposition 3] that the weighted Banzhaf values of $\hat{\nu}$ can be computed as the optimal first-order approximation of $\hat{\nu}$. The representation of the best k -th order approximation in the Möbius transform was given by Marichal and Mathonet [47, Proposition 6], and reads for $i \in N_{\mathcal{I}} \cup N_{\mathcal{T}}$, $k = 1$, and equal probabilities p as

$$\begin{aligned} &\hat{a}(\{i\}) + (-1)^{1-1} \sum_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}} : |M| > 1, i \in M} \binom{|M|-1-1}{1-1} p^{|M|-1} \hat{a}(M) \\ &= \hat{a}(\{i\}) + p \sum_{j \in N_{\mathcal{I}} \cup N_{\mathcal{T}} : j \neq i} \hat{a}(\{i, j\}) \\ &= \mathbf{e}_i + p \sum_{j \in N_{\mathcal{I}} \cup N_{\mathcal{T}} : j \neq i} \mathbf{e}_{\{i, j\}}, \end{aligned}$$

which concludes the proof. □

A.2 Proof of Theorem 2

Proof. We first compute the expectations as

$$\begin{aligned}\mathbb{E}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] &= \mathbb{E}\left[\frac{1}{m_{\mathcal{I}}m_{\mathcal{T}}} \sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} (\nu(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})} \cup M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}) - \hat{\nu}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})} \cup M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}))^2\right] \\ &= \frac{1}{m_{\mathcal{I}}m_{\mathcal{T}}} \sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} \mathbb{E}\left[(\nu(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})} \cup M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}) - \hat{\nu}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})} \cup M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}))^2\right] \\ &= \mathbb{E}_{(M_{\mathcal{I}}, M_{\mathcal{T}}) \sim \mathbb{P}_{p, \mathcal{I}} \otimes \mathbb{P}_{p, \mathcal{T}}}[(\nu(M_{\mathcal{I}} \cup M_{\mathcal{T}}) - \hat{\nu}(M_{\mathcal{I}} \cup M_{\mathcal{T}}))^2] \\ &= \mathbb{E}_{M \sim \mathbb{P}}[(\nu(M) - \hat{\nu}(M))^2] \\ &= \mathfrak{F}_p(\nu, \hat{\nu}_{\mathbf{e}}),\end{aligned}$$

since $M_{\mathcal{I}} \perp M_{\mathcal{T}}$ are independent.

We now proceed by computing the variance. We first introduce $g_{\mathbf{e}} : 2^{N_{\mathcal{I}}} \times 2^{N_{\mathcal{T}}} \rightarrow \mathbb{R}$ for $M_{\mathcal{I}} \subseteq N_{\mathcal{I}}$ and $M_{\mathcal{T}} \subseteq N_{\mathcal{T}}$ as

$$g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}}) := (\nu(M_{\mathcal{I}} \cup M_{\mathcal{T}}) - \hat{\nu}_{\mathbf{e}}(M_{\mathcal{I}} \cup M_{\mathcal{T}}))^2.$$

For the standard Monte Carlo estimator, we then obtain

$$\mathbb{V}[\hat{\mathfrak{F}}_p^{(m)}(\nu, \hat{\nu}_{\mathbf{e}})] = \frac{1}{m} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})],$$

due to m iid samples. For the cross-modal estimator, we first prove a result for the explicit form of the variance in Proposition 2.

Proposition 2. We denote the expectation over the conditional variances as

$$\begin{aligned}\tau_{\mathcal{I}|\mathcal{T}} &:= \mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[\mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}}[(\nu(M_{\mathcal{I}} \cup M_{\mathcal{T}}) - \hat{\nu}_{\mathbf{e}}(M_{\mathcal{I}} \cup M_{\mathcal{T}}))^2]], \\ \tau_{\mathcal{T}|\mathcal{I}} &:= \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}}[\mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[(\nu(M_{\mathcal{I}} \cup M_{\mathcal{T}}) - \hat{\nu}_{\mathbf{e}}(M_{\mathcal{I}} \cup M_{\mathcal{T}}))^2]].\end{aligned}$$

The variance of the estimator $\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})$ is then given by

$$\mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] = \frac{m_{\mathcal{I}} + m_{\mathcal{T}} - 1}{m_{\mathcal{I}}m_{\mathcal{T}}} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] - \frac{m_{\mathcal{I}} - 1}{m_{\mathcal{I}}m_{\mathcal{T}}} \tau_{\mathcal{I}|\mathcal{T}} - \frac{m_{\mathcal{T}} - 1}{m_{\mathcal{I}}m_{\mathcal{T}}} \tau_{\mathcal{T}|\mathcal{I}}.$$

Proof. The variance is computed as

$$\begin{aligned}\mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] &= \mathbb{V}\left[\frac{1}{m_{\mathcal{I}}m_{\mathcal{T}}} \sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} g_{\mathbf{e}}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})}, M_{\mathcal{T}}^{(\ell_{\mathcal{T}})})\right] \\ &= \frac{1}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} \sum_{k_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{k_{\mathcal{T}}=1}^{m_{\mathcal{T}}} \text{cov}\left(g_{\mathbf{e}}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})}, M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}), g_{\mathbf{e}}(M_{\mathcal{I}}^{(k_{\mathcal{I}})}, M_{\mathcal{T}}^{(k_{\mathcal{T}})})\right).\end{aligned}\quad (3)$$

To compute the covariance, we distinguish four cases, namely

$$\begin{aligned}&\text{cov}(g_{\mathbf{e}}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})}, M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}), g_{\mathbf{e}}(M_{\mathcal{I}}^{(k_{\mathcal{I}})}, M_{\mathcal{T}}^{(k_{\mathcal{T}})})) \\ &= \begin{cases} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})], & \text{if } \ell_{\mathcal{I}} = k_{\mathcal{I}}, \ell_{\mathcal{T}} = k_{\mathcal{T}}, \\ \text{cov}(g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}}), g_{\mathbf{e}}(M_{\mathcal{I}}, M'_{\mathcal{T}})), & \text{if } \ell_{\mathcal{I}} = k_{\mathcal{I}}, \ell_{\mathcal{T}} \neq k_{\mathcal{T}}, \\ \text{cov}(g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}}), g_{\mathbf{e}}(M'_{\mathcal{I}}, M_{\mathcal{T}})), & \text{if } \ell_{\mathcal{I}} \neq k_{\mathcal{I}}, \ell_{\mathcal{T}} = k_{\mathcal{T}}, \\ 0, & \text{if } \ell_{\mathcal{I}} \neq k_{\mathcal{I}}, \ell_{\mathcal{T}} \neq k_{\mathcal{T}}, \end{cases}\end{aligned}$$

where $M_{\mathcal{I}}, M'_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}$ and $M_{\mathcal{T}}, M'_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}$. The covariances can be further computed as

$$\begin{aligned}\text{cov}(g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}}), g_{\mathbf{e}}(M_{\mathcal{I}}, M'_{\mathcal{T}})) &= \mathbb{E}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})g_{\mathbf{e}}(M_{\mathcal{I}}, M'_{\mathcal{T}})] - \mathbb{E}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]\mathbb{E}[g_{\mathbf{e}}(M_{\mathcal{I}}, M'_{\mathcal{T}})] \\ &= \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}}[\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]\mathbb{E}_{M'_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[g_{\mathbf{e}}(M_{\mathcal{I}}, M'_{\mathcal{T}})]] - \mathbb{E}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]^2 \\ &= \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}}[\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]^2] - \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}}[\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]]^2 \\ &= \mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}}[\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]],\end{aligned}$$

where we have used $M_{\mathcal{I}} \perp M_{\mathcal{T}}, M'_{\mathcal{T}}$ and $M_{\mathcal{T}} \perp M'_{\mathcal{T}}$. Similarly, we obtain

$$\text{cov}(g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}}), g_{\mathbf{e}}(M'_{\mathcal{I}}, M_{\mathcal{T}})) = \mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [\mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] .$$

Combining these results into Eq. (3), we obtain

$$\begin{aligned} \mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] &= \frac{1}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \sum_{\ell_{\mathcal{I}}, k_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}, k_{\mathcal{T}}=1}^{m_{\mathcal{T}}} \text{cov} \left(g_{\mathbf{e}}(M_{\mathcal{I}}^{(\ell_{\mathcal{I}})}, M_{\mathcal{T}}^{(\ell_{\mathcal{T}})}), g_{\mathbf{e}}(M_{\mathcal{I}}^{(k_{\mathcal{I}})}, M_{\mathcal{T}}^{(k_{\mathcal{T}})}) \right) \\ &= \frac{1}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \underbrace{\left(\sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \right)}_{\text{case } \ell_{\mathcal{I}}=k_{\mathcal{I}}, \ell_{\mathcal{T}}=k_{\mathcal{T}}} \\ &\quad + \underbrace{\sum_{\ell_{\mathcal{I}}=1}^{m_{\mathcal{I}}} \sum_{\substack{\ell_{\mathcal{T}}, k_{\mathcal{T}}=1 \\ \ell_{\mathcal{T}} \neq k_{\mathcal{T}}}}^{m_{\mathcal{T}}} \mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]]}_{\text{case } \ell_{\mathcal{I}}=k_{\mathcal{I}}, \ell_{\mathcal{T}} \neq k_{\mathcal{T}}} \\ &\quad + \underbrace{\sum_{\substack{\ell_{\mathcal{I}}, k_{\mathcal{I}}=1 \\ \ell_{\mathcal{I}} \neq k_{\mathcal{I}}}}^{m_{\mathcal{I}}} \sum_{\ell_{\mathcal{T}}=1}^{m_{\mathcal{T}}} \mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [\mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]]}_{\text{case } \ell_{\mathcal{I}} \neq k_{\mathcal{I}}, \ell_{\mathcal{T}}=k_{\mathcal{T}}} \Big) \\ &= \frac{m_{\mathcal{I}} m_{\mathcal{T}}}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \\ &\quad + \frac{m_{\mathcal{I}} m_{\mathcal{T}} (m_{\mathcal{T}} - 1)}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] \\ &\quad + \frac{m_{\mathcal{T}} m_{\mathcal{I}} (m_{\mathcal{I}} - 1)}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [\mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] \end{aligned}$$

We now use the law of total variance and $M_{\mathcal{I}} \perp M_{\mathcal{T}}$ to rewrite

$$\mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] = \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] - \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] ,$$

and use a similar result for $M_{\mathcal{T}}$ to obtain

$$\begin{aligned} \mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] &= \frac{m_{\mathcal{I}} m_{\mathcal{T}} + m_{\mathcal{I}} m_{\mathcal{T}} (m_{\mathcal{T}} - 1) + m_{\mathcal{T}} m_{\mathcal{I}} (m_{\mathcal{I}} - 1)}{m_{\mathcal{I}}^2 m_{\mathcal{T}}^2} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \\ &\quad - \frac{m_{\mathcal{T}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] - \frac{m_{\mathcal{I}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [\mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] \\ &= \frac{m_{\mathcal{I}} + m_{\mathcal{T}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \\ &\quad - \frac{(m_{\mathcal{T}} - 1) \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]] + (m_{\mathcal{I}} - 1) \mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [\mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]]}{m_{\mathcal{I}} m_{\mathcal{T}}} \\ &= \frac{m_{\mathcal{I}} + m_{\mathcal{T}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] - \frac{m_{\mathcal{I}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \tau_{\mathcal{I}|\mathcal{T}} - \frac{m_{\mathcal{T}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \tau_{\mathcal{T}|\mathcal{I}}, \end{aligned}$$

where we have used the definitions of $\tau_{\mathcal{I}|\mathcal{T}} := \mathbb{E}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [\mathbb{V}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]]$ and $\tau_{\mathcal{T}|\mathcal{I}} := \mathbb{E}_{M_{\mathcal{I}} \sim \mathbb{P}_{p, \mathcal{I}}} [\mathbb{V}_{M_{\mathcal{T}} \sim \mathbb{P}_{p, \mathcal{T}}} [g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]]$. This concludes the proof. $\square$

By using Proposition 2, we obtain

$$\mathbb{V}[\hat{\mathfrak{F}}_p^{(m_{\mathcal{I}}, m_{\mathcal{T}})}(\nu, \hat{\nu}_{\mathbf{e}})] \leq \frac{m_{\mathcal{I}} + m_{\mathcal{T}} - 1}{m_{\mathcal{I}} m_{\mathcal{T}}} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \leq \frac{m_{\mathcal{I}} + m_{\mathcal{T}}}{m_{\mathcal{I}} m_{\mathcal{T}}} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})],$$

since $\tau_{\mathcal{I}|\mathcal{T}}, \tau_{\mathcal{T}|\mathcal{I}} \geq 0$.

Moreover, by the law of total variance we have $\tau_{\mathcal{I}|\mathcal{T}}, \tau_{\mathcal{T}|\mathcal{T}} \leq \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})]$, which implies

$$\begin{aligned}\mathbb{V}[\hat{\mathfrak{F}}_p^{(m,m)}(\nu, \hat{\nu}_{\mathbf{e}})] &\geq \frac{2m-1}{m^2} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] - \frac{2(m-1)}{m^2} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \\ &= \frac{1}{m^2} \mathbb{V}[g_{\mathbf{e}}(M_{\mathcal{I}}, M_{\mathcal{T}})] \\ &= \mathbb{V}[\mathfrak{F}_p^{(m^2)}(\nu, \hat{\nu}_{\mathbf{e}})],\end{aligned}$$

which finishes the proof. □

A.3 Proof of Proposition 1

Proof. For a token attribution explanation $\mathbf{e}$, the additive approximation reads as

$$\hat{\nu}_{\mathbf{e}}(M) = \mathbf{e}_0 + \sum_{i \in M} \mathbf{e}_i.$$

Denote M_k^+ the top- k coefficients of $\mathbf{e}$, then for $M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}$ with $|M| = k$, we have

$$\sum_{i \in M} \mathbf{e}_i \leq \sum_{i \in M_k^+} \mathbf{e}_i,$$

and thus

$$M_{\mathbf{e}, \max, k} = \operatorname{argmax}_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}: |M|=k} \hat{\nu}_{\mathbf{e}}(M) = \operatorname{argmax}_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}: |M|=k} \sum_{i \in M} \mathbf{e}_i = M_k^+.$$

Conversely, for the bottom- k coefficients M_k^- , we have

$$\sum_{i \in M} \mathbf{e}_i \geq \sum_{i \in M_k^-} \mathbf{e}_i,$$

for every $M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}$ with $|M| = k$, and thus

$$M_{\mathbf{e}, \min, k} = \operatorname{argmin}_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}: |M|=k} \hat{\nu}_{\mathbf{e}}(M) = \operatorname{argmin}_{M \subseteq N_{\mathcal{I}} \cup N_{\mathcal{T}}: |M|=k} \sum_{i \in M} \mathbf{e}_i = M_k^-.$$

Lastly, the top- k coefficients correspond to the complement of bottom- $(n_{\mathcal{I}} + n_{\mathcal{T}} - k)$ coefficients, which yields

$$M_{\mathbf{e}, \min, n_{\mathcal{I}} + n_{\mathcal{T}} - k} = M_{n_{\mathcal{I}} + n_{\mathcal{T}} - k}^- = (N_{\mathcal{I}} \cup N_{\mathcal{T}}) \setminus M_k^+,$$

which finishes the proof of the deletion curve $\mathcal{C}_{\text{delete}} = \{(k, \nu(M_{\mathbf{e}, \min, n_{\mathcal{I}} + n_{\mathcal{T}} - k}))\}_{k=1, \dots, n_{\mathcal{I}} + n_{\mathcal{T}}}$. $\square$

B Implementation Details

B.1 Faithful interaction explanations of LIP models (FIXLIP)

Because of how the model’s forward function is implemented, without loss of generality, we explain a logit-scaled similarity output from Equation 1, i.e. $C \cdot f(x_{\mathcal{I}}, x_{\mathcal{T}})$, which depends on the particular model’s learnt constant, e.g. $C \approx 100$ for CLIP, but $C \approx 117$ for SigLIP and $C \approx 112$ for SigLIP-2.

Masking. Following related work using input removal in attribution methods [17, 19, 27, 45], we apply simple baseline masking strategies, which were also proven successful based on our empirical results. For vision encoders, we use a 0 baseline after image normalization to propagate no signal forward. In the case of language encoders, we use a native attention masking mechanism, where we encode active tokens with 1 and the deleted ones with 0. It is important to leave the beginning/end sequence and padding tokens as is. Note that in theory, the maximum possible context length $n_{\mathcal{T}}$ is finite (e.g. 64, 77) and could be treated as constant across different inputs. However, we never explain these special tokens, treating parts of text inputs as the only features of interest. Alternative implementations of masking worth considering are: removing words from the text inputs, or tokens from tokenized inputs, imputing with the [UNK] token, or [MASK], e.g. in the case of the BiomedCLIP architecture that is based on BERT [76]. Acknowledging the potential influence of the masking approach, as well as position encodings, on the value of an empty input, we check that it remains about constant across different text input lengths in Figure 7.




Figure 7: Value of an empty input is about constant across different text lengths.

Budget split in the cross-modal estimator. The model-agnostic estimator (Definition 6) relies on m sampled masks for approximation; the number is often called the approximation budget. The cross-modal estimator (Definition 7) introduces budget split, where $m = m_{\mathcal{T}} \cdot m_{\mathcal{I}}$. For all our experiments, we set $m_{\mathcal{T}} := \min(2^{n_{\mathcal{T}}}, \max(4, \lceil \sqrt{m} \cdot n_{\mathcal{T}} / n_{\mathcal{I}} \rceil))$ and $m_{\mathcal{I}} := \min(2^{n_{\mathcal{I}}}, \max(4, \lfloor \sqrt{m} \cdot n_{\mathcal{I}} / n_{\mathcal{T}} \rfloor))$ as a reasonable default, allowing for a balanced exploration of both input spaces.

Two-step filtering approach. For the ViT-B/16 and ViT-L/16 model versions with a larger number of input image tokens (patches), we perform a cheap computation of the first-order attribution to prioritize interactions for approximation. We consider two strategies: picking a clique subset of the most interesting tokens, or focusing only on cross-modal interaction for explanation, which we specifically apply in the pointing game recognition (Definition 10, Appendix B.3). In the first strategy, given the desirable clique size $k = 72$, we take $k_{\mathcal{T}} := \max(5, \lceil k \cdot n_{\mathcal{T}} / (n_{\mathcal{T}} + n_{\mathcal{I}}) \rceil)$ text tokens with the highest absolute feature attribution scores, and $k_{\mathcal{I}} := k - k_{\mathcal{T}}$ image tokens accordingly. Similarly to the above considerations regarding the budget split between the two modalities, we found this to be empirically adequate. Future work can propose optimal strategies for both of these nuances.

Greedy subset selection. Even for the ViT-B/32 model version with fewer than 100 input tokens (nodes in a graph), finding cliques of the highest and smallest sums using brute force is computationally prohibitive. Still, we want to use these subsets denoted as $M_{\text{e}, \min, k}$ and $M_{\text{e}, \max, k}$ in evaluation with insertion/deletion curves (Definition 9) and explanation visualization. Thus, we implemented a simple greedy algorithm starting the search from each (or a subset) of the tokens (nodes). Its general goal is to add consecutive tokens to the subset based on minimizing/maximizing the subset’s value. For models with a larger number of inputs, e.g. $n_{\mathcal{I}} + n_{\mathcal{T}} = 196 + 30$, the greedy strategy remains applicable; however, it may require a few minutes to complete when considering all subset sizes $k = 1, \dots, n_{\mathcal{I}} + n_{\mathcal{T}}$.

B.2 Area between the insertion/deletion curves (AID)

Figure 9 describes visually the process of computing insertion/deletion curves and how it can differ between the first-order attribution methods and second-order interaction explanations. Consider the following example shown in **(bottom, Ours, 73%)**, where the faces of a smiling man and child are masked with the masked caption saying “a man — a child smiling at a — in a —”. Here, the model predicts the input to be dissimilar (below-average similarity), which means it enters negative values in our normalized case. Contrarily, in **(bottom, Baseline)**, the unmasked restaurant text token keeps the model’s similarity near the average level. Another interesting phenomenon happens for the insertion curve in Figure 9 **(top)**. Faithfully masking the redundant information with **our** method causes the model’s similarity to increase above the original prediction. Contrarily, in **(top, Baseline)**, the baseline gradient-based method is unable to faithfully recover such redundant tokens.

B.3 Pointing game recognition (PGR)

Figure 8 illustrates an example of the pointing game evaluation for interaction explanations like FIXLIP, giving additional context as to why first-order attribution methods fail to pass this desirable sanity check. A saliency map can only highlight important parts of image/text inputs, but whenever more objects/concepts appear in the input, they are unable to disentangle the basic relationship between the two modalities. PGR metric measures the ratio of absolute values of “correctly” identified cross-modal interaction terms to total interaction terms. For example, in Figure 8 **(2nd row, 2nd column)**, we sum the positive interactions between “banana”/“cat” image tokens and the banana/cat text tokens, as well as the absolute negative interactions between “tractor”/“ball” image tokens and the banana/cat text tokens. Then, we divide this sum by the total absolute sum of all interactions, which gives a normalized PGR score. In Appendix C.2, we further describe the specific combinations of image and object/class labels used to create our exemplary benchmark in this paper, although the overall methodology should be treated in a generic manner.




Figure 8: Visualization of a first-order attribution (**top row**), e.g. Shapley values, compared to cross-modal interactions in FIXLIP (**middle row**) for an exemplary pointing game. FIXLIP can be conditioned on any token, e.g. the first text input token (**bottom row**) for a simpler visualization. Consecutive columns here correspond to consecutive columns in Tables 1, 2 & 4.




Figure 9: Exemplary insertion/deletion curves for two explanations e computed with FIXLIP and another first-order method. We show five subset visualizations per curve per method. In the case of the baseline method, each consecutive subset needs to include the preceding one, which is not the case for FIXLIP, capturing the complex cross-modal and intra-modal interactions. For example, a child with the corresponding face patches (**top right corner**) gets overvalued by the man–pizza tokens appearing jointly with the corresponding image part (child disappears). Similarly, in the deletion curve, the subsets of text and image tokens interchange with each other at around 50% (**bottom row**).

C Experimental Setup

C.1 Models

We use the following openly available models from Hugging Face with default hyperparameters:

- CLIP ViT-B/32 [57]: `openai/clip-vit-base-patch32` (MIT License)
- CLIP ViT-B/16 [57]: `openai/clip-vit-base-patch16` (MIT License)
- SigLIP ViT-B/16 [75]: `google/siglip-base-patch16-224` (Apache License 2.0)
- SigLIP ViT-L/16 [75]: `google/siglip-large-patch16-256` (Apache License 2.0)
- SigLIP-2 ViT-B/32 [66]: `google/siglip2-base-patch32-256` (Apache License 2.0)
- SigLIP-2 ViT-B/16 [66]: `google/siglip2-base-patch16-224` (Apache License 2.0)
- SigLIP-2 ViT-L/16 [66]: `google/siglip2-large-patch16-256` (Apache License 2.0)

We additionally rely on the CLIP ViT-B/32 and ViT-B/16 model versions from the `clip` Python library [57, MIT License], not to modify the official implementation of the baselines (Appendix C.3). We set the batch size to 64 for the base models and to 32 for the large models, performing computation on A100 GPUs with 40GB VRAM.

C.2 Datasets

We use the following openly available datasets from Hugging Face:

- MS COCO test set [42]: `clip-benchmark/wds_mscoco_captions` (CC BY 4.0)
- ImageNet-1k validation set [15]: `ILSVRC/imagenet-1k` (ImageNet Agreement)

Experiments with the CLIP models are performed using 1000 image–text pairs from the MS COCO test set, for which each of the models predicted the highest similarity scores. Experiments with the SigLIP-2 models are performed using 100 image–text pairs to save computational resources, since we are not using it to compare with baseline methods.

Regarding ImageNet-1k, we use all 50 images from each of the following 10 class labels for constructing the pointing game evaluation: goldfish (1), husky (248), cat (282), plane (404), church (497), ipod (605), ball (805), tractor (866), banana (954), pizza (963). We combine these images with labels into 10 varied games as follows: goldfish-husky-pizza-tractor, cat-goldfish-plane-pizza, banana-cat-tractor-ball, husky-banana-plane-church, pizza-ipod-goldfish-banana, ipod-cat-husky-plane, tractor-ball-banana-ipod, plane-church-ball-goldfish, church-pizza-ipod-cat, ball-husky-banana-tractor. In each case, we distinguish four scenarios by gradually adding each consecutive token from left to right, resulting in 40 scenarios, each with a total of 50 images, which constitutes a reasonable case.

C.3 Baselines

We use the following publicly available source code:

- GAME [12]: <https://github.com/hila-chefer/Transformer-MM-Explainability> (MIT License)
- Grad-ECLIP [77]: <https://github.com/Cyang-Zhao/Grad-Eclip> (License unknown)
- exCLIP [48]: Supplementary material at <https://openreview.net/forum?id=plkrRJt98c> (License unknown)

We were unable to run exCLIP on the CLIP (ViT-B/16) model version for inputs with more than 24 text tokens, which are common in MS COCO, running into out-of-memory errors. Furthermore, we refrain from attempting to run these implementations on SigLIP models, which were not part of the original experiments. For context, we did not find an openly available implementation of InteractionLIME [32], while it is not clear how to adapt InteractionCAM [61] (proposed for image–image encoders) to the text/language domain.

C.4 Metrics

For p -faithfulness correlation (Definition 8), we set $m = 1000$ for a reasonable sample size and analyze $p = \{0.3, 0.5, 0.7\}$, as well as Shapley sampling weights to give a broader context.

Note that the insertion/deletion curves (Definition 9) have different lengths for each input (explanation), depending on the number of tokens (players). We thus interpolate each curve over a fixed range of 51 points between 0–100% input and then aggregate them to obtain the final visualization.

Pointing game recognition has no additional hyperparameters beyond the pre-defined dataset described in Appendix C.2. For the example presented in Table 2, we omit the (sub)games with either `goldfish` or `ipod` for SigLIP models, because their corresponding text tokenizer divides these class labels into multiple tokens.

C.5 Compute resources

Experiments described in Section 5 and Appendix D were computed on a cluster consisting of $4 \times$ AMD Rome 7742 CPUs (256 cores), 4TB of RAM, and $16 \times$ A100 GPUs for about 15 days combined. We envision that preliminary and failed experiments have required the same amount of compute resources.

D Additional Experimental Results

We report the insertion/deletion results for CLIP (ViT-B/16) in Figure 10 and for SigLIP-2 (ViT-B/32) in Figure 11, where we also analyze the gradual relationship between p -masking in FIXLIP and the model's prediction change more gradually. Table 4 reports pointing game results for CLIP (ViT-B/16), where interestingly FIXLIP with $p = 0.5$ outperforms FIXLIP (Shapley interactions).

Figure 12 shows further ablations based on the different distributions in the p -faithfulness correlation metric. Interestingly, FIXLIP with $p = 0.7$ evaluated on masks from $p = 0.3$ performs much better than FIXLIP with $p = 0.3$ evaluated on masks from $p = 0.7$, which further confirms our motivation to increase p in machine learning applications where out-of-distribution sampling is not desirable.




Figure 10: Insertion/deletion curves for CLIP (ViT-B/16) on MS COCO. Extended Figure 3. AID score (higher is better) for FIXLIP against alternative explanation methods, where a random baseline scores 0. The y-axis is normalized between the model's prediction on the original input (100%) and the fully removed one (0%), where negative values denote that the model is predicting the inputs are unsimilar. Gradient-based methods such as Grad-ECLIP fail to recover nonlinear rankings of important tokens, whereas FIXLIP faithfully recovers the optimal subset explanation. exCLIP does not appear in the comparison as its implementation returns an OOM error when querying the larger model (14×14 image tokens) with texts longer than 24 tokens, which are common in MS COCO.




Figure 11: Insertion/deletion curves for SigLIP-2 (ViT-B/32) on MS COCO. Extended Figure 3. AID score (higher is better) for a model-agnostic estimator of FIXLIP with different masking weights. The y-axis is normalized between the model's prediction on the original input (100%) and the fully removed one (0%), where negative values denote that the model is predicting the inputs are unsimilar. We observe a gradual relationship between p and the steepness of the curves at a particular fraction of input deleted. We omit comparison with other baselines in this case as they were originally implemented for CLIP and not SigLIP.

Table 4: Pointing game recognition for CLIP (ViT-B/16) on ImageNet-1k. Extended Table 1. PGR score for FIXLIP against alternative explanation methods, where a random baseline scores 0.25. First-order methods, such as Grad-ECLIP and Shapley values, fail to distinguish between multiple objects simultaneously, while second-order methods faithfully recover the appropriate explanation (up to the irreducible non-optimality of the pointing game).

Explanation Method	Recognition ($\uparrow$)			
	1 object	2 objects	3 objects	4 objects
GAME [12]	.60 $\pm$.11	.41 $\pm$.04	.33 $\pm$.02	.27 $\pm$.01
Grad-ECLIP [77]	.72 $\pm$.13	.44 $\pm$.04	.33 $\pm$.01	.26 $\pm$.01
Shapley values	.66 $\pm$.07	.57 $\pm$.05	.50 $\pm$.04	.42 $\pm$.03
Banzhaf values	.70 $\pm$.08	.58 $\pm$.05	.50 $\pm$.05	.41 $\pm$.04
exCLIP [48]	.73 $\pm$.15	.82 $\pm$.07	.85 $\pm$.05	.87 $\pm$.05
FIXLIP (Shapley interactions)	.76 $\pm$.08	.78 $\pm$.06	.79 $\pm$.05	.79 $\pm$.05
FIXLIP (w. Banzhaf interactions $p = 0.3$)	.76 $\pm$.10	.75 $\pm$.08	.77 $\pm$.07	.77 $\pm$.06
FIXLIP (w. Banzhaf interactions $p = 0.5$)	.81 $\pm$.08	.80 $\pm$.06	.81 $\pm$.05	.81 $\pm$.05
FIXLIP (w. Banzhaf interactions $p = 0.7$)	.75 $\pm$.09	.75 $\pm$.07	.77 $\pm$.06	.78 $\pm$.05




Figure 12: Faithfulness evaluation for CLIP (ViT-B/32 and ViT-B/16) on MS COCO. Extended Figure 4. We evaluate the faithfulness of different FixLIP variants and baselines in terms of p -faithfulness correlation and the R^2 coefficient (where it is possible for game-theoretic approaches). For 1 000 inputs, we sample 1 000 evaluation masks according to different p -faithfulness distributions: $p = 0.3$, $p = 0.5$, $p = 0.7$. For each p -faithfulness distribution, the corresponding FixLIP (w. Banzhaf interactions) outperforms the baselines.

E Visual Explanations

E.1 Guidance on interpreting visualizations of interaction explanations

In all of our visualizations, explanations are highlighted using a red–blue color scheme, where red indicates a positive contribution to the model’s output and blue a negative one. The transparency of the color reflects the strength of the contribution—more opaque elements correspond to higher absolute explanation values. Pairwise interactions are represented as edges, main effects (or first-order attributions) are displayed as colored boxes around words or overlaid as heatmaps on top of images. In addition to the color and transparency, edges also vary in thickness based on the interaction strength. In general, interactions may link elements within the same modality, e.g. text–text interactions or image–image interactions, or across modalities, e.g. text–image interactions.

Figure 13 illustrates how to interpret visualizations in a multimodal setting involving image and text inputs. Provided an input image and text pair (Figure 13A), we compute FIXLIP and retrieve an explanation e containing interactions and main effects. The FIXLIP output (Figure 13B) can be visualized by plotting the image and text together, and overlaying the main effects and interactions. This allows for a quick screening of interesting parts in the input as judged by an explanation. In this particular example, the elephant text token is positively linked to all image tokens containing the elephant. The two text tokens white and bird are linked to the one image token containing the bird. Here, elephant is red because removing it would decrease the predicted similarity score, while bird is blue because removing it might even increase the predicted similarity score (depending on cross-modal interactions). Based on the FIXLIP, more detailed visualizations can be created.

Conditioning on tokens. To investigate the role of a specific token $i \in N_{\mathcal{T}} \cup N_{\mathcal{I}}$, we can condition the interaction space on i and retrieve all interactions of the form $e_{\{i,j\}}$ for $j \in N \setminus \{i\}$. This conditional view highlights how token i interacts with all other tokens and is especially useful for analyzing its contextual relevance across modalities. Then, we project the interactions onto the first-order and visualize them with a heatmap (Figure 13C). Conditioning on elephant shows the greatest positive interactions with the body in the image tokens and the word small. Conditioning on the image token containing the bird mostly interacts with white and bird.

Finding interesting subsets. FIXLIP explanations allow searching for interesting subsets M like $M_{e,\max,k}$ or $M_{e,\min,k}$ for varying sizes of k . We can search the space of interactions and visualize only the explanations for tokens in these subsets, or overlay the masks on top of the input. In Figure 13D, we illustrate $M_{e,\max,7}$ containing tokens related to elephant. Starting the search from the bird token retrieves the $k = 4$ subset with information about the bird. While the elephant subset increases the similarity, the bird subset could slightly decrease the similarity, presumably because the model would find the isolated image–text similarity of the elephant clearer.

Visualizing heatmaps. Based on Theorem 1, second-order explanations provided by FIXLIP can be converted into first-order explanations. Hence, FIXLIP explanations can also be visualized with heatmaps (Figure 13E). Similar to baselines such as Grad-ECLIP or Shapley values, we overlay the converted first-order FIXLIP explanations onto the input. The scores are scaled *jointly* across both modalities since we retrieve the explanation scores across both modalities. An alternative could be to scale scores *independently* with different maximum values (one per modality). Joint scaling allows for relative comparisons, which is not the case for independent scaling, where two tokens are visualized with full color intensity. In Figure 13E (left), the independently scaled heatmap shows that the vision token containing the bird and the elephant text token have the highest scores in their respective modality. Jointly scaling the heatmap reveals that the elephant text token is of a higher score than the vision token containing the bird.



Figure 13: Showcase of explanation visualization. A comprehensive description is in Appendix E.1.

E.2 Additional examples

Figure 14 demonstrates the utility of conditioning on tokens. Figure 15 shows intriguing examples of interactions in inputs with text written in an image, which is related to the typographic capabilities of large vision–language models [10, 23, 63].





Figure 14: FIXLIP provides a utility of conditional heatmap generation. Comparison between a first-order attribution explanation and a second-order interaction explanation, which is able to recognize complex nuances of the presented scene.



a sign that says " abbey road nw 8 city of westminster " .



a sign reading " juice theory " near a traffic light .



the streets signs at the intersection of ha light and ash bury .

Figure 15: Intriguing examples. Top interactions in inputs with text written in an image.

E.3 Comparison between CLIP (ViT-B/32) and SigLIP-2 (ViT-B/32)

See Figure 16. Note that these models have different text tokenizers and input image resolutions, i.e. 224×224 vs. 256×256 , which is why the images differ slightly.



Figure 16: Visual comparison between FiXLIP of CLIP (left) and SigLIP-2 (right).

E.4 Comparison between CLIP (ViT-B/32) and CLIP (ViT-B/16)

Visual examples between CLIP (ViT-B/32) and CLIP (ViT-B/16) are in Figure 17. Notably, in the example including the clock (Fig. 17, row 3), conditioning on the vision token containing the letters T and part of O reveals a strong interaction with another vision token of this word in the “sign” *and* the text token town. This may suggest the presence of a *third-order* interaction between these tokens.




Figure 17: Visual comparison between CLIP ViT-B/32 (left) and ViT-B/16 (right). Each heatmap is conditioned on the selected input token (in black), also for vision tokens (in yellow).

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We validate our claims empirically and provide proofs in Appendix A.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: In Section 6.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: In Section A.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All models and datasets used are openly available. We share code implementation for our method, experiments, and analyzes on GitHub. We add details on the implementation and experimental setup in Appendices B & C, respectively.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All models and datasets used are openly available. We share code implementation for our method, experiments, and analyzes on GitHub.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes] ,

Justification: We state experimental setting at the beginning of Section 5, and elaborate on it in Appendix C.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report metric values with standard deviations and via boxplots.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We estimate compute resources in Appendix C.5.

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We do not violate the NeurIPS Code of Ethics.

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We describe the broader impact in Section 6.

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not release new data or models.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original papers, state versions, include URLs, and list asset licenses in Appendix C.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new assets.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: Our method does not involve using LLMs.