

MM-SHAP: A Performance-agnostic Metric for Measuring Multimodal Contributions in Vision and Language Models & Tasks

Letitia Parcalabescu and Anette Frank
Computational Linguistics Department
Heidelberg University

Abstract

Vision and language models (VL) are known to exploit unrobust indicators in individual modalities (e.g., introduced by distributional biases) instead of focusing on relevant information in each modality. That a unimodal model achieves similar accuracy on a VL task to a multimodal one, indicates that so-called unimodal collapse occurred. However, accuracy-based tests fail to detect e.g., when the model prediction is wrong, while the model used relevant information from a modality. Instead, we propose MM-SHAP, a performance-agnostic multimodality score based on Shapley values that reliably quantifies in which proportions a multimodal model uses individual modalities. We apply MM-SHAP in two ways: (1) to compare models for their average degree of multimodality, and (2) to measure for individual models the contribution of individual modalities for different tasks and datasets. Experiments with six VL models – LXMERT, CLIP and four ALBEF variants – on four VL tasks highlight that unimodal collapse can occur to different degrees and in different directions, contradicting the wide-spread assumption that unimodal collapse is one-sided. Based on our results, we recommend MM-SHAP for analysing multimodal tasks, to diagnose and guide progress towards multimodal integration. Code available at <https://github.com/Heidelberg-NLP/MM-SHAP>.

1 Introduction

Vision and language (VL) tasks are dominated by general-purpose pretrained transformer-based VL models (Lu et al., 2019; Tan and Bansal, 2019; Li et al., 2019; Chen et al., 2020; Li et al., 2020, 2021a). But we are only starting to understand why these multimodal (MM) models work so well, and how they utilise and fuse image and text modalities (Hessel and Lee, 2020; Cao et al., 2020). Even worse, these highly parametrised neural VL models, pretrained on large amounts of data, tend to

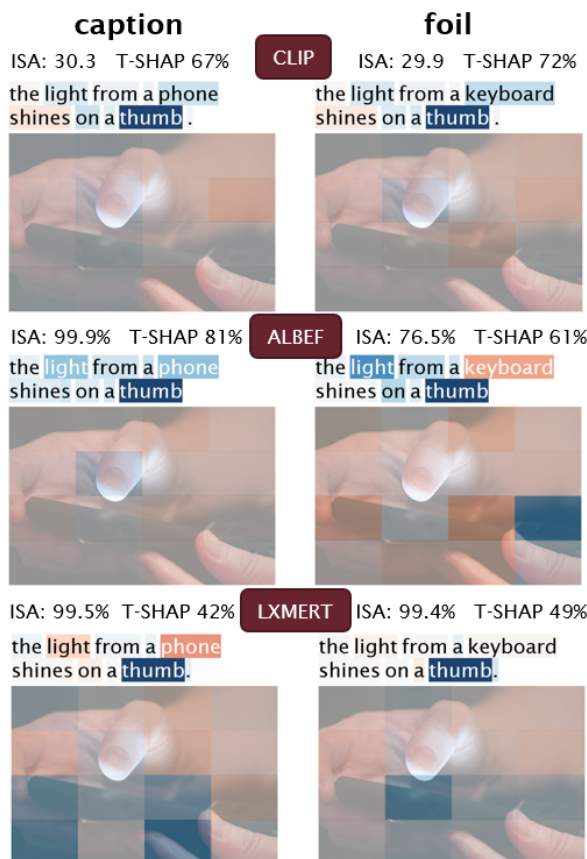


Figure 1: We display image-sentence alignment scores (ISA) and the *textual degree* T-SHAP that measures how much models focus on text rather than the image (with $100 - T-SHAP\%$ the corresponding *visual degree*) for 3 VL models. Blue/red highlights on text tokens and image tokens (patches) contribute towards higher/lower ISA. Note: CLIP’s ISA is an absolute score, while ALBEF and LXMERT predict ISA probabilities. See Section 4.4 for more details on this figure; App. C for more detailed analysis of this instance and more samples.

exploit artefacts and statistical correlations in the data (Shekhar et al., 2019; Kafle et al., 2019), showing little to no evidence of detailed linguistic or visual understanding (Milewski et al., 2022; Parcalabescu et al., 2022; Thrush et al., 2022). Statistical biases towards indicators in one modality – to the

detriment of others – can cause *unimodal collapse* (Parcalabescu et al., 2022), where seemingly MM models exploit one modality that exhibits biases, meaning that the MM system effectively reduces to a unimodal model (Madhyastha et al., 2018) – e.g., if a model answers “How many...?” questions with “two” – the most frequent answer seen in training (Goyal et al., 2017). Unimodal collapse is severe, as it leads to loss of system reliability. It also shows that *multimodal fusion* is far from being solved. Hence the importance of *measuring multimodal degree* – the degree to which modalities are used in model predictions – with *reliable metrics*.

To test for unimodal collapse, research has so far focused on performance tests: a VL model is evaluated on a MM task, but one modality crucial for solving it correctly is missing, corrupted (Shekhar et al., 2017) or permuted (Gat et al., 2021). These tests are indicative of unimodal collapse, but we argue that they are not appropriate to reliably measure the contribution of each modality. Clearly, accuracy reflects whether a model prediction is (in)correct, but it may detect illicit cases where the model prediction is *wrong*, although it *does* use crucial indicators in a given modality. Conversely, a prediction might be *correct*, but may be derived from unrobust indicators. Fig. 1 shows very different SHAP-based *contribution patterns* of image regions and text tokens leading to model responses of different image-sentence alignment (ISA) scores (e.g., ALBEF caption vs. foil), while yielding same ISA accuracy since both scores surpass the 0.5 classification threshold.

As an alternative to accuracy-based methods, we propose MM-SHAP, a *performance-agnostic metric* to quantify and interpret the contribution of individual modalities in VL models. MM-SHAP is based on Shapley values (Shapley, 1953), a theoretically well-founded interpretability method from cooperative game theory. We apply MM-SHAP to quantify the contribution of specific parts of the input towards model predictions.

Our main contributions are:

- i) We propose MM-SHAP, a performance-agnostic metric to measure the degree of contribution of each modality in VL (but not limited to V&L), to *measure the degree to which individual modalities contribute to MM model predictions*. We combine MM-SHAP with model accuracy to analyse the degree to which each modality contributes to model predictions.

- ii) We use MM-SHAP to 1) compare models in terms of their reliance on different modalities, 2) compare the relevance of different modalities for a given task and dataset, and to 3) zoom in at sample-level to determine the contribution of each modality and each token in each modality for a model prediction (Fig. 1).
- iii) We conduct experiments with six VL models: LXMERT, CLIP and four ALBEF variants – on four VL tasks: image-sentence alignment, VQA, GQA and on the more fine-grained VALSE 🌸 VL benchmark.
- iv) We identify VL models that are balanced in their usage of two modalities (CLIP), models that show a higher visual degree (LXMERT) or a stronger textual degree (ALBEF).
- v) We show that 1) fine-tuning a model can affect its MM degree and that 2) current VL models do not all collapse towards the same modality, as reported in recent work (Frank et al., 2021; Gat et al., 2021), but that directions can differ from model to model.

2 Related Work

Testing for unimodal collapse Strong prediction indicators in either modality can cause MM models to ignore weaker indicators in another modality. Prior work has proposed ways to identify (and remove) such biases from data (Goyal et al., 2017).

Foiling approaches introduce mistakes in image descriptions and test whether VL models notice the discrepancy between image and captions (Shekhar et al., 2019; Parcalabescu et al., 2022), finding that models are surprisingly insensitive to such foils. Gat et al. (2021), in a similar vein, exchange images with other images or captions with other captions, expecting that inputs with misleading information in one modality incur a decrease in model accuracy. They use an observed *decrease in task accuracy* to calculate a *perceptual score* as a measure of the MM degree of a model. Their findings suggest that across their tested VL models, textual input consistently matters more than visual input.

Ablation methods remove information from either modality and test whether the model can still solve the task. Here, Frank et al. (2021) find that the visual modality matters more than text: VL models suffer from image parts removal when predicting masked text, but can predict masked visual inputs when text input is ablated. This contradicts Gat et al. (2021)’s finding, but their investigations have

only a single model in common, namely LXMERT.

Hence, the literature agrees that VL models are not as cross-modal as expected – but disagree on whether models rely more on the textual (Gat et al., 2021) or on the visual modality (Frank et al., 2021). We argue that a reason for this discrepancy is that prior work computes MM scores based on model performance. In our work we argue that methods for measuring a model’s MM degree should not rely on accuracy (see §3.1 for motivation). Instead, we propose an *accuracy-agnostic* method to measure the MM degree of VL models, using the SHAP (Lundberg and Lee, 2017) interpretability method that is theoretically suitable to define an MM score.

Interpretability Methods for explaining predictions of neural models can be classified into two categories: *White-box methods*, which require access to specific components of neural architectures and *black-box methods*, which are model-agnostic, requiring only access to model inputs and outputs.

Notable *white-box methods* are: Attention-based methods, which correlate high attention weights with high feature importance. But the equivalence of importance score and attention is debated and must be considered with care (Jain and Wallace, 2019; Wiegrefe and Pinter, 2019) (see App. D for a detailed discussion on why attention is inappropriate for defining an MM score). Layer-wise relevance propagation (Binder et al., 2016) or gradient-based methods e.g., Grad-CAM (Selvaraju et al., 2017) can also be used to determine the importance of inputs, but can be deceived by small changes in inputs (adversarial attacks).

Notable *black-box methods* are: LIME (Ribeiro et al., 2016) and its multimodal adaptation DIME (Lyu et al., 2022) approximate the vicinity of the input with a linear function that is interpretable. But depending on the choice of the size of the vicinity, LIME can lead to very disparate results. Methods like RISE (Petsiuk et al., 2018) and SHAP (Lundberg and Lee, 2017) compute importance scores by randomly masking parts of the input and determining the effect this has on the output. SHAP exhibits great theoretical properties that enable us to define a MM score, as we will motivate in §3.4.

3 Quantifying Multimodal Contributions

3.1 A case for a performance-agnostic score

As a community, we are interested in improving model performance, and thus need to evaluate mod-

els using performance metrics such as accuracy. But in this work we address a complementary question that is only indirectly related to performance. We aim to measure *how much a given modality matters for model predictions*. This is important for model developers to know, to detect *unimodal collapse*, and to find ways of preventing it.

To date, research tried to measure MM contributions based on accuracy. Gat et al. (2021) and Frank et al. (2021), e.g., rely on the difference between a model’s accuracy with and without information from a modality, e.g., to define the *importance of vision* as $V = Acc(vision, text) - Acc(\emptyset, text)$. This score works well if a MM model shows good performance, but is problematic for wrong model predictions, since in such cases $Acc(vision, text) = 0$, and we expect $Acc(\emptyset, text) = 0$ too, resulting in $V = 0$ (or another low value). But this does not necessarily reflect reality: The model may well have relied on the visual modality, but incorrectly.

Even worse, *accuracy-based* methods that completely *delete* (Madhyastha et al., 2018) or *exchange* (Gat et al., 2021) information in one modality are ill-defined for image-sentence alignment (ISA): ISA asks a model to assess how well two modalities align, with the rationale that alignment is given if the given modalities (e.g., image and text) contain relevant information that indicates alignment by ‘being about the same things or facts’. In case the information conveyed in two modalities is not about the same (type of) things (e.g., a picture of a dog paired with a caption talking about a cat), the modalities do not align. However, metrics that measure the *importance of vision* V by the impact of deleting it, as $V = Acc(vision, text) - Acc(\emptyset, text)$, are ill-defined for *unaligned* image-sentence pairs: A model that uses both modalities to correctly predict *misalignment* ($Acc(vision, text) = 1$), will also predict a mismatch when the visual information is deleted or exchanged, yielding $Acc(\emptyset, text) = 1$. This results in $V = 0$, signalling that no visual importance is measured, which is ill-founded in this case. Hence, accuracy-based scores that rely on deletion of single modalities are unable to measure multimodal degree on ISA – an important pretraining task for VL models – or on zero-shot ISA benchmark tasks such as VALSE 🦄 (Parcalabescu et al., 2022).

We argue for using *accuracy-agnostic methods* to measure a model’s *multimodal degree* and propose *MM-SHAP*, a metric that avoids the pit-

falls of performance-based metrics. We move from $Acc(vision, text)$ to measuring the *relative contribution* of vision and text by measuring $Contribution(vision, text)$ for a given model prediction. We compute the $Contribution$ function using Shapley values, which quantify a token’s contribution to a model prediction, independently of whether the prediction is correct. Importantly, our performance-agnostic way of measuring a model’s MM degree in terms of contributions of tokens – within or across modalities – will make it possible to clearly separate accuracy-based performance analysis from the study of relative contributions of modalities in MM systems. This allows us to measure MM degree in situations where accuracy cannot: e.g., when model accuracy is low – as in out-of-domain or zero-shot settings.

3.2 Background on Shapley Values

Shapley values¹ were first introduced in a game theoretical setting to estimate fair rewards among cooperative players (Shapley, 1953). For machine learning, the outcome of a game is the model’s prediction, the players are parts of the input and are assigned Shapley values that represent the importance of each player (Lundberg and Lee, 2017).

We compute Shapley values for pretrained transformer-based VL models at prediction time. Their input consists of p input tokens (image and text tokens alike). We create subsets $S \subseteq \{1, \dots, n\}$ of tokens forming a coalition towards the model prediction $val(S)$. Tokens not being part of the subset are masked. $val(\emptyset)$ is the output of the model when all tokens are masked. The Shapley value for a token j follows formula (1):

$$\phi_j = \sum_{S \subseteq \{1, \dots, n\} \setminus \{j\}} \frac{val(S \cup \{j\}) - val(S)}{\gamma} \quad (1)$$

Here, $\gamma = \frac{|S|!(n-|S|-1)!}{p!}$ is the normalising factor that accounts for all possible combinations of choosing subset S . When masking p tokens, the coalition possibilities grow exponentially ($n = 2^p$). We thus approximate the Shapley values with Monte Carlo, by randomly sub-sampling $n = 2p + 1$.

The Shapley value of a token measures its contribution towards the model prediction (e.g., the probability of image-sentence alignment) and can

¹We refer to Molnar (2022) for a gentle introduction into Shapley Values.

be **positive** (increases the model prediction) or **negative** (decreases it) or **zero** (no effect). Shapley values exhibit four defining properties of a fair payout, which are all beneficial for model interpretability: (1) *Efficiency*: the contributions of all players sum up to the model outcome; (2) *Symmetry*: any two players that contribute equally are assigned the same payout; (3) *Dummy*: a non-contributing part is assigned zero value and (4) *Additivity*, enabling us to simply average the Shapley Values to determine the overall player contributions in a game with combined payouts (e.g., the two halves of a soccer match, or ensembling of decision trees).

Most importantly, Shapley values are not based on model accuracy or performance, but *solely on the model’s input and its prediction*, e.g., the probability for an image and a caption to match. This is an important property for our MM score, since its objective is to quantify *how much inputs of either modality matter for prediction* – even if the cooperation between (multimodal) inputs is not sufficient to reach success, i.e., yielding the correct outcome.

3.3 MM-SHAP

For a pretrained VL transformer with n_T text tokens and n_I image tokens, Eq. 2 defines the textual contribution Φ_T and the image contribution Φ_I towards a prediction as the sum of (absolute) Shapley Values (Eq. 1) of all textual resp. visual tokens:

$$\Phi_T = \sum_j^{n_T} |\phi_j| \quad ; \quad \Phi_I = \sum_j^{n_I} |\phi_j| \quad (2)$$

We consider the magnitude and not the sign of a token contribution², as we are interested in measuring whether a token is active in a modality – irrespective of the direction it pushes the prediction into. Eq. 3 defines MM-SHAP as a *proportion* of modality contributions, allowing us to determine a model’s *textual degree* T-SHAP and its *visual degree* V-SHAP:

$$\text{T-SHAP} = \frac{\Phi_T}{\Phi_T + \Phi_I}; \text{V-SHAP} = \frac{\Phi_I}{\Phi_T + \Phi_I} \quad (3)$$

We can extend MM-SHAP to any number of modalities. Here we only use image and text.

When generating coalitions, i.e., subsets of tokens from which to compute Shapley Values, we do not distinguish image and text tokens, because

²Contributions can be positive (increase the model prediction) or negative (decrease it) or zero (no effect), see §3.2.

MM-SHAP aims to fairly distribute potential token contributions first and to aggregate them modality-wise in a 2nd step with Eq. 2. To **mask tokens**, we replace text tokens with the [MASK] token; for images we set pixel values of image patches to zero. We ensure similar text and image sequence lengths by using more and smaller patches for longer text, and vice versa – resulting in 16 image patches for the majority of samples in our data. See App. A.

3.4 Why SHAP enables a MM score

Our aim for MM-SHAP is to estimate the proportion to which text and vision are used by VL models ($x\%$ visual and $y\%$ textual). Defining an MM score is nontrivial, since it should not be based on accuracy, see §3.1. An MM score should rely on a measure of how much tokens contribute to the output value computed by the model. Most interpretability methods do not directly answer this question of how much models use certain features, but use proxies such as gradients or attention. Moreover, their explanations cannot be added modality-wise in a meaningful way, to define a relative contribution per modality (Cf. App. D for a discussion on attention). Luckily, Shapley values compute fair payouts to players (tokens), depending on their contribution to achieving the total payout (the model’s prediction). Their theoretically founded properties – e.g. fair payout between tokens and modalities, or in-sample and between-sample additivity, as detailed in §3.2 – allow us to aggregate intra-modal token-level contributions to compute an MM score.

Grounding our MM score in Shapley values bears further advantages, which we discuss next.

3.5 Ways of using MM-SHAP

Sample-level MM-SHAP, being based on the contributions of individual image and text tokens, is a sample-level score (Fig. 1). It enables fine-grained analyses of the relevance of tokens from a single or various modalities, for each instance.

Dataset and model level We can average sample-level MM-SHAP scores into dataset-level scores, thanks to the additivity property of Shapley values. Hence it can help analyse a model across various datasets, or compare distinct models on a certain dataset to gain insights of models, datasets / tasks.

Measuring fine-tuning effects An accuracy-based MM score is limited when model performance on a task is very low, since the differences

between a model’s accuracy with correct vs. permuted inputs are small in such cases (Cf. §3.1). Since MM-SHAP is based on actual model predictions and not on model performance, we can apply MM-SHAP for models with low performance. E.g., we can compare a pretrained model’s MM score to a fine-tuned version of it that may have lost general task abilities (thus showing low accuracy) after specialising for another task; or we can measure the effectiveness of targeted interventions in fine-tuning to increase a model’s reliance on modalities.

Future work could apply MM-SHAP on models accepting different or a wider range of modalities, for tracing a model’s MM-SHAP evolution in pre-training, or on data cleaning, by identifying groups of samples with very unbalanced MM degree – especially when the accuracy on those samples is high and the model may rely on unimodal cues.

4 Multimodal Contributions across Models and Datasets

We use MM-SHAP to study MM contributions for different i) model types, ii) datasets and iii) tasks. In doing so we iv) re-evaluate prior findings on visual vs. textual unimodal collapse and v) showcase MM-SHAP’s abilities for interpreting predictions for individual samples, for error analysis.

We evaluate pretrained VL models with MM-SHAP and complement our analysis by measuring the model’s task accuracy. We compare MM-SHAP to a 50% T-SHAP – 50% V-SHAP baseline and gauge how much the model tends towards the textual or visual modality. We hypothesise that in average, V&L should contribute equally when the model predicts whether the contents of the modalities are aligned (image-sentence alignment).

We test on matching image-captions, but also on cases with discrepancies between modalities. We break down our incongruity tests into *high discrepancy* (cases of completely mismatching image-captions, Tab. 1), and *low discrepancy* (cases where a single word or phrase incurs a mismatch, Tab. 2).

4.1 Tasks

Visual Question Answering (VQA) is a task where transformer-based VL models have consistently increased SOTA performance. We use the VQA v2.0 (Goyal et al., 2017) and GQA (Hudson and Manning, 2019) datasets for our experiments.

Image-sentence alignment (ISA) VL models are typically pretrained on predicting an image-

sentence alignment score. We assess their MM contributions in their “comfort zone” by letting them predict the alignment of images and captions, in contrast to misalignment to random captions. We test on 1,500 samples from the MSCOCO validation set (Lin et al., 2014), and on uncommon image-caption pairs composed from questions and answers from the VQA and GQA validation sets.

ISA on fine-grained VL phenomena In ISA tasks, models are typically confronted with highly discrepant negative samples (non-matching image-captions). To evaluate VL models in a more fine-grained manner, we examine their MM score on the VALSE 🦋 benchmark (Parcalabescu et al., 2022), where foiled captions were created by altering phrases pertaining to 6 specific linguistic phenomena: existence, counting, plurality, spatial relations, actions, and coreference, such that image and foiled caption do not match. For completeness, we also test on noun phrase foils as introduced in the FOILit! dataset (Shekhar et al., 2017).

4.2 Models

LXMERT by Tan and Bansal (2019) is a dual-stream transformer that combines V&L in early fusion using cross-modal attention layers between image and language encoders. It was pretrained on MSCOCO (Lin et al., 2014) images and captions, and on VQA v2.0 and GQA images, questions and answers. Its objectives were (i) multimodal masked word and object prediction, (ii) ISA, and (iii) VQA objectives. For experiments on ISA, VQA and GQA, we use the corresponding heads and task-specific checkpoints.³

CLIP by Radford et al. (2021) processes image and text with two separate transformer encoders. The resulting image and text representations are combined in late fusion by cross-product. CLIP was trained for ISA in *low discrepancy mode* on 400M image-text pairs to predict high scores for paired image-text examples and low scores when image-text samples are not paired in the dataset. With this simple contrastive learning objective, CLIP shows zero-shot capabilities in e.g. object classification, OCR, or activity recognition (Radford et al., 2021). In our work, we test CLIP⁴ on ISA and VALSE 🦋, using the model’s image-text alignment score to assess whether it predicts a

higher image-text similarity for correct pairs or for foiled image-caption pairs.

ALBEF by Li et al. (2021b) combines vision and language with early and late fusion. As in CLIP, transformer image and text encoders are trained to map the two modalities to a common space. Cross-modal transformer layers further combine the two with (i) MM masked word prediction and (ii) ISA objectives. It was pretrained on Conceptual Captions (Sharma et al., 2018), SBU Captions (Ordonez et al., 2011), MSCOCO (Lin et al., 2014) and Visual Genome (Krishna et al., 2017).

To analyse how the MM contributions are affected by fine-tuning, we compare 4 ALBEF⁵ models fine-tuned on (1) image retrieval on MSCOCO, (2) image retrieval on Flickr30k (Plummer et al., 2015), (3) visual grounding on RefCOCO+ (Yu et al., 2016) and (4) VQA (Goyal et al., 2017).

4.3 Metrics

We use **accuracy** to assess model performances, and **MM-SHAP** to measure the proportion to which the different modalities contribute.

With **MM-SHAP** (def. in §3.3) we aim to analyse the MM contributions in terms of visual (V-SHAP) and textual (T-SHAP) degree. As in our case of two modalities they are complementary ($V-SHAP = 100 - T-SHAP$), we only report T-SHAP (in %). We distinguish $T-SHAP_c$ for textual degree in image-caption pairs and $T-SHAP_f$ for image-foil pairs. As the results are very similar, we refer to Table 3 App. B for $T-SHAP_f$ results.

When evaluating VQA and GQA performance, accuracy measures the proportion of correct answers given pairs of images and questions. For ISA, we measure the overall accuracy acc of models to classify foils and captions. We fan out acc into **caption accuracy** acc_c (for correctly predicting matching images and captions) and **foil accuracy** acc_f (for correctly predicting mismatching images and foils). **Pairwise accuracy** acc_r measures the proportion of samples where the ISA score is higher for a correct image-text pair compared to its image-foil counterpart. acc_r is more permissive than acc : it does not require the ISA score to surpass a classification threshold (of 0.5), but only that image-foil pairs are ranked lower than the ground truth pairs.

³github.com/huggingface/transformers

⁴github.com/openai/CLIP

⁵github.com/salesforce/ALBEF

4.4 Experiments and Results

We test all VL models from §4.2 without further tuning and assess their task accuracy and MM-SHAP scores in three settings: i) for VQA on the VQA and GQA datasets; for ISA ii) with *high discrepancy* image-caption pairs (from MSCOCO, VQA, GQA) and iii) with *low discrepancy* pairs from VALSE 🦋; finally iv) we showcase sample-level analyses using MM-SHAP. Table 1 shows results on VQA, GQA and ISA; Table 2 for VALSE 🦋. MM-SHAP varies between samples with a stdev. of $\sim 12\%$ across our experiments.

High discrepancy ISA (Table 1) shows that acc_r scores for ISA on MSCOCO, VQA, GQA are high for all models. This is expected as they have been pretrained for ISA – only ALBEF vqa stands out: it lost its ISA performance by fine-tuning on VQA. LXMERT has highest acc_r for ISA on VQA and GQA, since for its last 10 epochs it was trained on these datasets.

For ISA, we observe the models scattering around the hypothesised 50% balance for T-SHAP, with CLIP being the most balanced one, especially on MSCOCO. This is expected since CLIP is a two-branch model where the two modalities communicate in late fusion, in other words, CLIP keeps all information from the textual and visual branches separate until the very end. By contrast, LXMERT has a low textual degree of only 35.5%, while ALBEF models are more textual.

Given *highly diverging* foil pairs, T-SHAP_c and T-SHAP_f differ prominently: LXMERT moves from weak to higher textual degree (35.5 to 62.8%) and inversely for ALBEF mscoco (63.4 to 54.3%).

Canonical VL tasks Results on VQA and GQA in Table 1 – with ALBEF fine-tuned for VQA and LXMERT fine-tuned on VQA and GQA⁶ – show high model accuracy. T-SHAP is higher for VQA (51.5%) than for ISA (45.7% acc_c), which is interesting, since LXMERT was more visually focused on ISA. It seems like ALBEF vqa’s and LXMERT’s training on VQA increases the impact of the textual modality to the detriment of the visual one. This aligns with earlier findings that in VQA tasks, linguistic indicators (e.g., “How many...?”) give away the most likely answer (two) (Goyal et al., 2017).

Low discrepancy ISA on VALSE 🦋 in Table 2. For T-SHAP_c we bold-face high deviations from

⁶We do not test CLIP and the other ALBEF models on VQA because they do not have corresponding VQA heads.

the 50% T-SHAP baseline (values $> 61\%$ and $< 40\%$). We note that the scores do not deviate much from the baseline. CLIP is the multimodally most balanced model, with an average T-SHAP_c of 50.7% across all instruments, which is expected, as argued for high discrepancy ISA above. By contrast, LXMERT skews towards the visual modality with an average T-SHAP_c of 41.9%, while ALBEF focuses on text – its variants showing T-SHAP_c values of 56% to 62%. This is consistent with our results for high discrepancy ISA in Table 1.

Accuracy vs. MM-SHAP On VALSE 🦋, accuracies do not correlate with MM-SHAP (see App. B.1 for details). This suggests that MM-SHAP *complements accuracy* in assessing MM contributions. As Parcalabescu et al. (2022) observe, models are better with some instruments (noun phrases, existence) as opposed to others (actions, coreference). Our work adds the multimodal score MM-SHAP as a new dimension of analysis. Some models exhibit strong divergences in T-SHAP across phenomena: LXMERT is strongly visually focused for plurality, spatial relations, noun phrases; also ALBEF’s textual bias is especially strong for these phenomena.

Model bias For overall ISA results on VALSE 🦋, Table 2 shows that despite varying model accuracies (stdev. for acc_r across phenomena $\pm 11-15\%$), MM-SHAP is relatively stable ($\pm 1-5\%$ stdev.) even when data distributions differ: E.g., *counting adversarial* contains foils in number ranges 0 to 3, while for captions numbers are higher than 4. The piece serves as a sanity check for biased models that may prefer the more frequently found small numbers. For LXMERT and ALBEF refcoco acc_r drops for counting *small numbers* to counting *adversarial* (encircled numbers in Tab. 2) from 69.2% to 42.6% for LXMERT and from 70.7% to 45.7% for ALBEF – while T-SHAP_c stays remarkably constant (47.3% to 46.4% and 55.1% to 55.8%). Even for phenomena that suffer from plausibility bias (Parcalabescu et al., 2022), T-SHAP varies little, while accuracies differ. Stable MM-SHAP scores highlight our MM score’s ability to measure how much the input modalities matter for model predictions – irrespective of their correctness –, complementing accuracy. Further results in App. B.2 compare model performances on foils vs. captions, supporting MM-SHAP’s stability while accuracy varies.

Fine-tuning effects For the four fine-tuned ALBEF models evaluated on VALSE 🦋, we observe

Model	Visual Question Answering				MSCOCO					Image-sentence alignment					GQA				
	VQA		GQA		acc_c	acc_f	acc_r	T_e	T_f	VQA		GQA			acc_c	acc_f	acc_r	T_e	T_f
	acc	T	acc	T						acc_c	acc_f	T_e	T_f	acc_c					
Random	0.0	50.0	0.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0
LXMERT	72.5	51.5	60.3	57.8	71.8	99.1	99.3	35.5	62.8	66.6	95.9	95.2	45.7	57.5	41.8	96.5	89.9	47.5	59.8
CLIP	-	-	-	-	-	-	99.5	50.3	52.9	-	-	94.0	48.4	47.6	-	-	83.4	47.0	46.0
A mscoco	-	-	-	-	95.9	99.6	99.8	63.4	54.3	28.0	99.9	91.0	60.3	59.2	13.1	99.7	83.6	58.3	57.2
A flickr	-	-	-	-	97.3	99.4	99.7	61.1	56.6	42.4	99.2	91.8	61.3	60.2	23.4	99.5	84.1	58.7	58.1
A refcoco	-	-	-	-	92.3	99.3	99.7	56.6	58.9	49.8	99.1	90.0	57.8	58.6	25.0	98.4	85.6	58.2	59.3
A vqa	76.0	66.7	-	-	99.9	0.0	33.4	64.1	62.8	100.0	0.0	60.2	58.2	60.0	100.0	0.0	52.6	61.7	62.4

Table 1: Task accuracy and MM score on canonical tasks. T is T-SHAP (in %). V-SHAP = 100 - T-SHAP. acc_r is pairwise ranking accuracy, counting predictions as correct if $p(\text{caption}, \text{img}) > p(\text{random}, \text{img})$. **A** stands for ALBEF fine-tuned for different tasks: image retrieval on MSCOCO and Flickr30k; visual grounding on RefCOCO+ and VQA. Overall foil task performance is the mean of acc_c and acc_f (equal nb. of samples, all pairs).

Metric	Model	Existence quantifiers	Plurality number	Counting			Sp.rel.† relations	Action		Coreference		Foil-it! nouns	Avg. ± stdev.	MM skew
				bal.†	sns.†	adv.†		repl.†	swap†	std.†	clean			
	Random	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0±0	
acc_r	CLIP	66.9	56.2	62.1	62.5	57.5	64.3	75.6	68.6	52.1	49.7	88.8	64.0±11	
	LXMERT	78.6	64.4	62.2	69.2	42.6	60.2	54.8	45.8	46.8	44.2	87.1	59.6±15	
	A mscoco	78.6	80.1	71.8	74.3	68.9	74.6	79.8	62.6	62.2	59.6	97.0	73.6±11	
	A flickr	80.6	78.9	71.0	73.6	64.3	73.3	82.4	55.5	59.9	57.7	96.6	72.1±12	
	A refcoco	73.1	69.0	67.9	70.7	45.7	68.6	79.9	58.9	52.7	43.3	96.5	66.0±15	
	A vqa	40.8	63.3	49.0	49.2	23.2	61.9	51.7	52.0	55.9	43.3	67.2	50.7±12	
acc	LXMERT	55.8	55.1	52.0	55.4	49.4	50.7	51.1	48.5	49.8	49.0	70.8	53.4±6	
	A mscoco	56.7	60.2	55.4	53.9	56.0	52.3	63.7	54.0	52.7	52.0	76.3	57.6±7	
	A flickr	55.6	56.3	53.8	53.3	55.4	52.3	64.9	48.9	50.0	50.0	70.5	55.5±6	
	A refcoco	53.4	56.3	51.1	51.1	48.4	51.1	63.1	51.2	50.7	49.3	77.4	54.8±8	
	A vqa	52.8	50.0	50.0	50.0	51.1	53.5	50.0	50.0	51.4	50.0	53.7	51.1±1	
T-SHAP _c	CLIP	44.7	52.3	51.5	51.8	52.1	50.9	50.0	49.7	52.1	52.6	49.9	50.7±2	bal.
	LXMERT	51.7	37.1	46.5	47.3	46.4	36.6	42.1	42.2	38.2	37.2	36.1	41.9±5	vis.
	A mscoco	56.7	63.5	58.3	58.0	59.5	64.1	61.7	61.5	61.9	61.4	63.9	60.9±3	txt.
	A flickr	59.5	61.7	59.6	59.8	59.5	61.6	59.8	58.9	60.9	61.9	63.5	60.6±1	txt.
	A refcoco	53.3	57.2	55.4	55.1	55.8	57.0	54.5	54.4	57.9	58.9	56.8	56.0±2	txt.
	A vqa	64.6	63.6	62.5	61.4	63.4	63.0	59.3	60.3	63.6	63.1	62.1	62.4±2	txt.

Table 2: Performance and MM scores of VL models on the VALSE benchmark. We bold-face high accuracies and multimodally unbalanced models on tasks. acc_r : the pairwise ranking accuracy, considering predictions as correct if $p(\text{caption}, \text{img}) > p(\text{foil}, \text{img})$. acc : Overall ISA accuracy. **A** stands for different fine-tunings of ALBEF: image retrieval on MSCOCO and Flickr30k, visual grounding on RefCOCO+ and VQA. †bal. Counting balanced. †sns. Counting small numbers. adv. Counting adversarial. repl. Action replacement. swap. Actant swap. †Sp.rel. Spatial relations. †std. Coreference standard. MM skew: Modality on which a model relies more: bal. balanced, vis. visual, txt. textual. We refer to Table 3 in App. B for more fanned out results.

that i) the models fine-tuned for image retrieval (mscoco, flickr) are good at predicting ISA (73.6% acc_r for ALBEF mscoco) but not those for VQA (ALBEF vqa 50.7%) and referring expressions (ALBEF refcoco 66.0%). This is expected, since ISA and image retrieval are very similar tasks. Interestingly, not only accuracy, but also the MM score changes, making ALBEF vqa more focused on text (62.4% avg. T-SHAP_c across VALSE) compared to referring expressions (ALBEF refcoco 56.0%). Notably, MM-SHAP being accuracy-agnostic, we can compute indicative scores even when a fine-tuned model fails the task completely, like ALBEF vqa that always predicts the foil class on captions.

Sample-level analysis Fig. 1 shows ISA predictions of CLIP, ALBEF mscoco and LXMERT, and their T-SHAP values for caption and foil. LX-

MERT correctly predicts high ISA between image and caption (left), although the regions contributing most (in blue) are not all reasonable, since the ‘phone’ token is not correctly grounded. ALBEF mscoco and CLIP also assign very high ISA scores, while using well-justified image regions for thumb and phone. On the foil (right), LXMERT’s contributing tokens change, with the phone region in the image mistakenly contributing to a high ISA. Favourably for ALBEF, the ‘keyboard’ text token contributes towards lowering the ISA, unlike for CLIP and LXMERT, where the ‘keyboard’ token raises the ISA. For more examples see App. C. We also showcase how attention does *not* reflect negative impact of tokens on a model’s prediction – which is very important in e.g., assessing the impact of foil words – in App. D.2 and Fig. 10.

4.5 Comparison to other MM metrics

We can only compare to other MM scores for VQA, because accuracy-based MM scores that delete information cannot apply to ISA (as argued in §3.1).

Unsurprisingly LXMERT’s accuracy when deleting the image is 31%; when deleting the text it is 8%, since excluding the image should negatively affect accuracy more than excluding text in VQA, where at least the answer type can be better inferred from the text (should be numeral for “How many”). But this ablation tells us more about the task definition than a model’s reliance on modalities.

The Perceptual Score (Gat et al., 2021) computes the per-sample difference between the model’s accuracy when working with the correct image and text as input and with a random image or text. LXMERT’s Perceptual Score (Gat et al., 2021) is 32.5 visual, 41.6 textual (relying more on text), but we argued in §3.1 that does not reflect cases where a model makes a wrong prediction because it failed to interpret the right cues correctly. MM-SHAP rates LXMERT vqa as balanced (51.5% T-SHAP).

4.6 On the need of a MM score

Our experiments show that a models’ reliance on a modality can vary with each task, dataset and instance. While prior work found that the models they analysed *all* prefer a *single* modality that they rely on most, our analyses show that *different* VL models behave *differently* on *the same* task: ALBEF is rather textual, CLIP balanced, LXMERT shows higher visual degree.

For LXMERT, we side with Frank et al. (2021), who found it to have a higher visual preference – this aligns with our analysis yielding a T-SHAP of 41.9%. We therefore disagree with Gat et al. (2021), who found a preference towards text.

Clearly, we do not assume that a MM model must rely equally on multiple modalities, but there are cases where unimodal collapse is unwanted, i.e., a model gives the right answer for the wrong reason in tasks such as VQA. MM-SHAP helps identify how much models rely on each modality.

5 Conclusions and Future Work

We present MM-SHAP, a performance-agnostic metric that measures the MM degree of VL models at dataset and sample level. Our results show that *on the same task, dataset*, and on specific *instances*, different types of models rely on modalities to different degrees and in different directions. Using

MM-SHAP we are the first to quantify changes in a model’s MM degree through fine-tuning. Our analyses show that degrees of MM contributions can be orthogonal to task performance, supporting the need for performance-agnostic metrics. MM-SHAP is applicable to further modalities. It enables model-based data cleaning and thus, dataset bias removal. Finally, it can serve as a diagnostic tool for improving MM fusion methods.

MM-SHAP can be used for testing true model understanding at a global and at instance level, and whether a model is giving the right answer for the right reasons, at corpus – and instance-level – which is not guaranteed for performance-dependent metrics. It can help us track MM contributions during (pre-)training and towards assessing and eventually predicting how much a model needs to rely on how many and which modalities in a given task or instance case – and how to explain this. We hence believe that many future research questions will profit from our MM score as an unbiased MM contribution metric, with AI research advancing to include more and more modalities beyond vision and language (Girdhar et al., 2023): acoustics, haptics, emotion, and more (cf. Parcalabescu et al. (2021b)).

Limitations

This work focused on assessing multimodal degree for recent English VL models. The following limitations can be relevant for future work.

We only evaluated a limited number of models in a zero-shot setting using their image-sentence alignment and VQA heads. Future work might be interested in assessing more models and tracking the evolution of MM-SHAP scores during model pretraining and finetuning.

This work applied MM-SHAP to VL encoders. We leave it for future work to investigate autoregressive (decoder-only) VL models. In the time it took to review and publish this work, we already encountered efforts to apply Shapley Values for interpreting VL models in Cafagna et al. (2023).

We only applied ML-SHAP to VL models. Future work might be interested in models working with other or additional modalities beyond vision and language.

Computing all possible coalitions between input tokens for Shapley Values is infeasible because their number is exponential in the number of tokens (2^p). Therefore we perform Monte Carlo approxi-

mation by randomly sub-sampling $2p+1$ coalitions. This results in approximate MM-SHAP scores per sample. We argue that as an alternative, one can simply increase the number of sampled coalitions for more exact measurements (as we did 10-fold for Fig. 1 and the examples in Appendix C) – at the cost of increasing the environmental footprint. But it is not necessary to increase the number of samples when estimating MM-SHAP at dataset level, because the number of coalitions has very little effect on a data-set wide range – given that approximation fluctuations average out.

To compute MM-SHAP at data-set level, one needs to run models in inference mode $2p+1$ times, where p is the number of tokens to mask (around 40 in average for MSCOCO-sized captions). On an NVIDIA Titan X GPU, computing MM-SHAP for one image-caption pair can take 2 seconds for ALBEF, 3 seconds for CLIP. LXMERT is the most expensive and needs 15 seconds, because it computes image features with a CNN backbone for every masking configuration.

Ethical Considerations

This paper uses publicly available datasets and models and therefore could carry on their potential biases (Meister et al., 2022; Garcia et al., 2023) and imperfections. However, the method presented in this paper enables model and dataset interpretation and we hope that it *can help future work locate harmful biases*.

Acknowledgements

The authors would like to thank the anonymous reviewers, Albert Gatt and Emanuele Bugliarello for their useful suggestions. Thanks go to Nils Trost for assisting with the visualisations.

References

Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. 2016. Layer-wise relevance propagation for neural networks with local renormalization layers. In *International Conference on Artificial Neural Networks*, pages 63–71. Springer.

Michele Cafagna, Lina M Rojas-Barahona, Kees van Deemter, and Albert Gatt. 2023. Interpreting vision and language generative models with semantic visual priors. *arXiv preprint arXiv:2304.14986*.

Jize Cao, Zhe Gan, Yu Cheng, Licheng Yu, Yen-Chun Chen, and Jingjing Liu. 2020. Behind the scene: Re-

vealing the secrets of pre-trained vision-and-language models. In *European Conference on Computer Vision*, pages 565–580. Springer.

- Hila Chefer, Shir Gur, and Lior Wolf. 2021a. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 397–406.
- Hila Chefer, Shir Gur, and Lior Wolf. 2021b. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 782–791.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *ECCV*.
- Ian Covert, Scott M Lundberg, and Su-In Lee. 2020. Understanding global feature contributions with additive importance measures. *Advances in Neural Information Processing Systems*, 33:17212–17223.
- Stella Frank, Emanuele Bugliarello, and Desmond Elliott. 2021. [Vision-and-language or vision-for-language? on cross-modal influence in multimodal transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9847–9857, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Noa Garcia, Yusuke Hirota, Yankun Wu, and Yuta Nakashima. 2023. Uncurated image-text datasets: Shedding light on demographic bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6957–6966.
- Itai Gat, Idan Schwartz, and Alex Schwing. 2021. Perceptual score: What data modalities does your model perceive? *Advances in Neural Information Processing Systems*, 34.
- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Manant Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. *arXiv preprint arXiv:2305.05665*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913.
- Jack Hessel and Lillian Lee. 2020. [Does my multimodal model learn cross-modal interactions? it’s harder to tell than you might think!](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 861–877, Online. Association for Computational Linguistics.

- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709.
- Sarthak Jain and Byron C. Wallace. 2019. **Attention is not Explanation**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Théo Jaunet, Corentin Kervadec, Romain Vuillemot, Grigory Antipov, Moez Baccouche, and Christian Wolf. 2021. Visqa: X-raying vision and language reasoning in transformers. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):976–986.
- Kushal Kafle, Robik Shrestha, and Christopher Kanan. 2019. **Challenges and prospects in vision and language research**. *Frontiers in Artificial Intelligence*, 2:28.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. **Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training**. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 11336–11344. AAAI Press.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021a. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021b. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. In *Arxiv*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Yiwei Lyu, Paul Pu Liang, Zihao Deng, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2022. Dime: Fine-grained interpretations of multimodal models via disentangled local explanations. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 455–467.
- Pranava Swaroop Madhyastha, Josiah Wang, and Lucia Specia. 2018. **Defoiling foiled image captions**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 433–438, New Orleans, Louisiana. Association for Computational Linguistics.
- Nicole Meister, Dora Zhao, Angelina Wang, Vikram V Ramaswamy, Ruth Fong, and Olga Russakovsky. 2022. Gender artifacts in visual datasets. *arXiv preprint arXiv:2206.09191*.
- Victor Milewski, Miryam de Lhoneux, and Marie-Francine Moens. 2022. **Finding structural knowledge in multimodal-BERT**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5658–5671, Dublin, Ireland. Association for Computational Linguistics.
- Christoph Molnar. 2022. *Interpretable Machine Learning*, 2 edition.
- Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24.
- Letitia Parcalabescu, Michele Cafagna, Lilitta Muradjan, Anette Frank, Iacer Calixto, and Albert Gatt. 2022. **VALSE: A task-independent benchmark for vision and language models centered on linguistic phenomena**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8253–8280, Dublin, Ireland. Association for Computational Linguistics.
- Letitia Parcalabescu, Albert Gatt, Anette Frank, and Iacer Calixto. 2021a. **Seeing past words: Testing the cross-modal capabilities of pretrained V&L models on counting tasks**. In *Proceedings of the 1st Workshop on Multimodal Semantic Representations*

- (MMSR), pages 32–44, Groningen, Netherlands (Online). Association for Computational Linguistics.
- Letitia Parcalabescu, Nils Trost, and Anette Frank. 2021b. [What is multimodality?](#) In *Proceedings of the 1st Workshop on Multimodal Semantic Representations (MMSR)*, pages 1–10, Groningen, Netherlands (Online). Association for Computational Linguistics.
- Vitali Petsiuk, Abir Das, and Kate Saenko. 2018. [RISE: randomized input sampling for explanation of black-box models](#). *CoRR*, abs/1806.07421.
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Marco Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. [“why should I trust you?”: Explaining the predictions of any classifier](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 97–101, San Diego, California. Association for Computational Linguistics.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626.
- L. S. Shapley. 1953. *17. A Value for n-Person Games*, pages 307–318. Princeton University Press.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia. Association for Computational Linguistics.
- Ravi Shekhar, Sandro Pezzelle, Yauhen Klimovich, Aurélie Herbelot, Moin Nabi, Enver Sangineto, and Raffaella Bernardi. 2017. [FOIL it! find one mismatch between image and language caption](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 255–265, Vancouver, Canada. Association for Computational Linguistics.
- Ravi Shekhar, Ece Takmaz, Raquel Fernández, and Raffaella Bernardi. 2019. [Evaluating the representational hub of language and vision models](#). In *Proceedings of the 13th International Conference on Computational Semantics - Long Papers*, pages 211–222, Gothenburg, Sweden. Association for Computational Linguistics.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: Learning cross-modality encoder representations from transformers](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer.

A Experimental Details

Masking VL models predict their outputs (such as ISA) on full and uncorrupted image and text inputs. To compute Shapley values and with them the MM-SHAP score, we create coalitions by masking image and text tokens. For **masking text**, we replace the text tokens with the [MASK] token.

For **masking images** we mask out image patches setting pixel values to zero. The patches are the regions for which we compute Shapley values, as visualised in Figures 2 to 9. By masking these patches, the SHAP algorithm can estimate how the prediction of the model changes in all possible combinations of their presence or absence.

After zero-ing out the patches, the models work as usual: LXMERT with the Faster-RCNN backbone computes image features and extracts image tokens. Working on the image level has the upside that no neighborhood information can leak into each image token: If we were to mask out on feature-level of the Faster-RCNN, i.e., on rectangular regions, the other regions would possibly “know about” the other regions due to the hierarchical structure of the CNN. For CLIP, the CLIP image encoder works as usual: It works internally with 32x32 patches of images in which we have already zeroed out information.

Therefore this masking procedure has the upside of being directly applicable to different types of VL model architectures, since some apply transformers directly on the image (CLIP and ALBEF), while others compute image tokens (features) with a different CNN-based backbone (LXMERT).

For computing Shapley values, we aim for a balance between text and image sequence length to make MM-SHAP adaptable to variable caption lengths and variable image sizes. Therefore we use the text length to dynamically determine patch sizes: For longer text, we use more and smaller patches and for shorter text, less but bigger patches. In the majority of our experiments, we have 16 image patches. We illustrate the image tiling in the top right of Figures 2 to 9.

This masking procedure has several advantages: i) It adapts to variable caption lengths and variable image sizes, and ii) it directly applies to different types of VL model architectures, since some apply transformers directly on the image (CLIP and ALBEF), while others compute image tokens (features) with a different CNN-based backbone (LXMERT).

Special tokens When computing token-wise contributions, we do not take [SEP] and [CLS] tokens into account (i.e., they are always assigned zero contribution), since their functionality is to aggregate cross-modal information, e.g. for classification, and hence they cannot be attributed to one modality exclusively.

B Additional results

Due to space constraints, we could not include full detailed results on VALSE 🦋 in 2. Here, we present Table 3, which is an extended version of Table 2 including the MM-SHAP scores for foils too, rather than just the captions. It also includes fanned out accuracies over matching image-captions acc_c and mismatching image-foils acc_f .

B.1 Correlation between accuracy and MM-SHAP

For each model and instrument on VALSE 🦋, we computed the Spearman correlation coefficient between the sample’s accuracy and textual degree. The correlations are very low, e.g., the correlation between acc_c and $T-SHAP_c$ is around 0.02 for most instruments and models, rising to 0.12 in rare cases. This low correlation between accuracy and MM-SHAP indicates that they are not measuring the same aspect: accuracy measures the models’ performance while MM-SHAP measures the degree to which a modality was used – independently of the success of its use.

B.2 MM-SHAP difference between captions and foils

We do not find notable differences between foils and captions on VALSE 🦋 in terms of MM-SHAP (cf. Table 3), while we find clear differences in accuracies between acc_c and acc_f , since they measure the model’s preference towards one side in the binary classification. Similar MM-SHAP scores between captions and foils speak for their ability to capture how the model’s input matters for the prediction, independently on which class the decision falls onto. A notable exception is the difference between $T-SHAP_c$ and $T-SHAP_f$ for LXMERT and ALBEF refoco on Foil-it! (underlined numbers in Table 3).

C Sample-level Analyses with MM-SHAP

SEE FIGURES ON FOLLOWING PAGES!

Metric	Model	Existence quantifiers	Plurality number	Counting			Sp.rel.‡ relations	Action		Coreference		Foil-it! nouns	Avg. ± stdev.	MM skew
				bal.†	sns.†	adv.†		repl.†	swap†	std.†	clean			
	Random	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0	50.0±0		
<i>acc_r</i>	CLIP	66.9	56.2	62.1	62.5	57.5	64.3	75.6	68.6	52.1	49.7	88.8	64.0±11	
	LXMERT	78.6	64.4	62.2	69.2	42.6	60.2	54.8	45.8	46.8	44.2	87.1	59.6±15	
	A mscoco	78.6	80.1	71.8	74.3	68.9	74.6	79.8	62.6	62.2	59.6	97.0	73.6 ±11	
	A flickr	80.6	78.9	71.0	73.6	64.3	73.3	82.4	55.5	59.9	57.7	96.6	72.1±11	
	A refcoco	73.1	69.0	67.9	70.7	45.7	68.6	79.9	58.9	52.7	43.3	96.5	66.0±15	
	A vqa	40.8	63.3	49.0	49.2	23.2	61.9	51.7	52.0	55.9	43.3	67.2	50.7±12	
<i>acc</i>	LXMERT	55.8	55.1	52.0	55.4	49.4	50.7	51.1	48.5	49.8	49.0	70.8	53.4±6	
	A mscoco	56.7	60.2	55.4	53.9	56.0	52.3	63.7	54.0	52.7	52.0	76.3	57.6±7	
	A flickr	55.6	56.3	53.8	53.3	55.4	52.3	64.9	48.9	50.0	50.0	70.5	55.5±6	
	A refcoco	53.4	56.3	51.1	51.1	48.4	51.1	63.1	51.2	50.7	49.3	77.4	54.8±8	
	A vqa	52.8	50.0	50.0	50.0	51.1	53.5	50.0	50.0	51.4	50.0	53.7	51.1±1	
<i>acc_c</i>	LXMERT	41.6	68.0	50.9	50.0	61.5	73.1	35.8	36.8	81.2	80.8	72.3	59.3±17	
	A mscoco	18.4	93.2	26.7	23.7	34.6	95.9	66.2	64.9	87.0	89.4	96.1	63.3±32	
	A flickr	28.7	94.0	43.1	41.2	50.8	96.8	65.1	64.2	91.5	96.2	97.5	69.9±26	
	A refcoco	33.7	89.8	41.8	31.0	57.2	93.1	72.5	75.0	81.4	90.4	92.7	69.0±24	
	A vqa	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0±0	
<i>acc_f</i>	LXMERT	70.1	42.2	53.0	60.8	37.3	28.4	66.4	60.2	18.4	17.3	69.3	47.6±20	
	A mscoco	91.5	27.1	82.0	87.2	80.9	9.2	61.7	42.3	16.1	12.5	52.1	51.1±32	
	A flickr	82.4	18.5	66.4	70.9	58.6	7.1	63.3	38.8	8.2	4.8	42.4	41.9±28	
	A refcoco	71.3	19.4	62.0	72.9	41.8	10.5	53.2	29.7	18.4	8.7	61.19	40.8±25	
	A vqa	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0±0	
<i>T-SHAP_c</i>	CLIP	44.7	52.3	51.5	51.8	52.1	50.9	50.0	49.7	52.1	52.6	49.9	50.7±2	bal.
	LXMERT	51.7	37.1	46.5	47.3	46.4	36.6	42.1	42.2	38.2	37.2	36.1	41.9±5	vis.
	A mscoco	56.7	63.5	58.3	58.0	59.5	64.1	61.7	61.5	61.9	61.4	63.9	60.9±3	txt.
	A flickr	59.5	61.7	59.6	59.8	59.5	61.6	59.8	58.9	60.9	61.9	63.5	60.6±1	txt.
	A refcoco	53.3	57.2	55.4	55.1	55.8	57.0	54.5	54.4	57.9	58.9	56.8	56.0±2	txt.
	A vqa	64.6	63.6	62.5	61.4	63.4	63.0	59.3	60.3	63.6	63.1	62.1	62.4 ±2	txt.
<i>T-SHAP_f</i>	CLIP	45.2	53.0	50.8	51.7	51.1	51.0	48.3	48.2	52.4	52.1	50.0	50.3±2	bal.
	LXMERT	52.3	39.4	48.2	48.8	45.8	36.5	43.9	42.7	39.1	38.6	45.0	43.7±5	vis.
	A mscoco	57.2	62.8	57.7	56.0	57.0	64.6	61.9	63.2	61.9	61.8	65.8	60.9±3	txt.
	A flickr	56.1	61.9	57.8	57.8	58.5	62.5	59.3	61.9	61.1	62.1	61.7	60.1±2	txt.
	A refcoco	56.1	58.5	56.2	55.6	57.8	57.6	55.5	56.9	58.4	58.4	61.3	57.5±2	txt.
	A vqa	64.0	64.7	61.9	60.9	61.2	63.2	59.9	60.1	63.4	62.4	62.2	62.2 ±2	txt.

Table 3: Performance and multimodal score of VL models on the instruments of the VALSE 🦋 benchmark. We bold-face high accuracies and multimodally unbalanced models on tasks. *acc_r* is the pairwise ranking accuracy, considering predictions as correct if $p(\text{caption}, \text{img}) > p(\text{foil}, \text{img})$. Overall foil task performance *acc* is the mean of *acc_c* and *acc_f* (equal number of samples, all pairs). A stands for ALBEF models fine-tuned on different tasks and datasets: image retrieval on MSCOCO and Flickr30k, visual grounding on RefCOCO+ and VQA. †bal. Counting balanced. †sns. Counting small numbers. adv. Counting adversarial. repl. Action replacement. swap. Actant swap. ‡ Sp.rel. Spatial relations. †std. Coreference standard. MM skew: Modality on which a model relies more: bal. balanced, vis. visual, txt. textual. We test CLIP in pairwise ranking mode only (CLIP works contrastively).

Figures 2 to 9 contain sample-level visualisations for each model for images and i) captions that match and ii) foils / random captions that show low / high discrepancy mismatch with the images, as introduced in Section 4.4:

- There is **low discrepancy** between images and foils obtained from VALSE 🦋 targeting specific linguistic phenomena, with only a phrase differing between the caption and the foil. We selected examples for different phenomena: Figure 2 (noun phrase), 3 (action replacement, easy example), 4 (counting), 5 (positive existence), 6 (negative existence), 9 (action replacement, hard example).
- There is **high discrepancy** between MSCOCO images and randomly chosen captions in terms of low ISA between image

and random caption – Figures 7 (easier example) and 8 (harder example).

In Figure 2 we reiterate Figure 1 from the main paper with more detail:

- CLIP correctly predicts a foil in the pairwise accuracy setting, since the ISA score for the caption (30.3) is higher than for the foil (29.9), but fails to identify that “keyboard” should not contribute towards a high ISA. It successfully predicts caption alignment, but seems to misunderstand the meaning of the word “shines” and its instantiation in the image.
- ALBEF mscoco is the only model to predict ISA (99.4%) on the caption with coherent – but mostly textual – indicators. It fails on foil prediction, still relying on the same textual indicators, and on the visual side *focuses on*

counter-evidence regions, erroneously taking them as positive support for ISA.

- LXMERT predicts correct ISA for the caption (99.5% ISA), using few relevant textual tokens as indicators, and possibly useful supporting visual tokens (focuses the fingers of the two hands). It fails to detect the foil (99.4% ISA which is higher than a 50% classification threshold and just slightly below the ISA for the caption): counterevidence from textual tokens is out-weighted by a single strong indicator (thumb); visual tokens confirm ISA despite focusing on counterevidence (the phone).

On the following pages we present Figures 4 to 9 with more samples and their analyses.

We sampled the instances based on the following criteria: i) low / high discrepancy; ii) interesting VALSE 🦄 instruments; iii) easier (no cluttering, no dark spots, no blur) and iv) harder examples (e.g., hard to recognise the statue as such in Figure 9).

Through Fig. 4 to 9, we observe some patterns:

Model performance does not tell much about the multimodal degree. A correct ISA score (high for the caption, low for the random caption/foil) is not always accompanied by a sensible contribution pattern in terms of Shapley values as seen for example in Figures 2 and 4 for CLIP and LXMERT. The Shapley values computed on the image and text side deliver much better intuition about what was successfully aligned and what was not grounded correctly. Among all models, LXMERT seems to be most affected by high discrepancy between performance and image and text token contributions.

Easy examples deliver more robust contribution patterns. On easy examples (Figures 3 and 4), where the model generally performs well, we can see how in the low discrepancy cases where caption and foil differ in only one word, the one word difference does not change the contribution patterns much. In contrast, low discrepancy hard examples (Figures 8 – unusual bed and bedroom arrangement and 9 – hard to recognise the goat as a statue without world knowledge) deliver different patterns on caption and foil, indicating confusion from the models.

Positive existence is easier than negative existence. When comparing Figures 5 and 6 we get some insight into how the models’ image-sentence

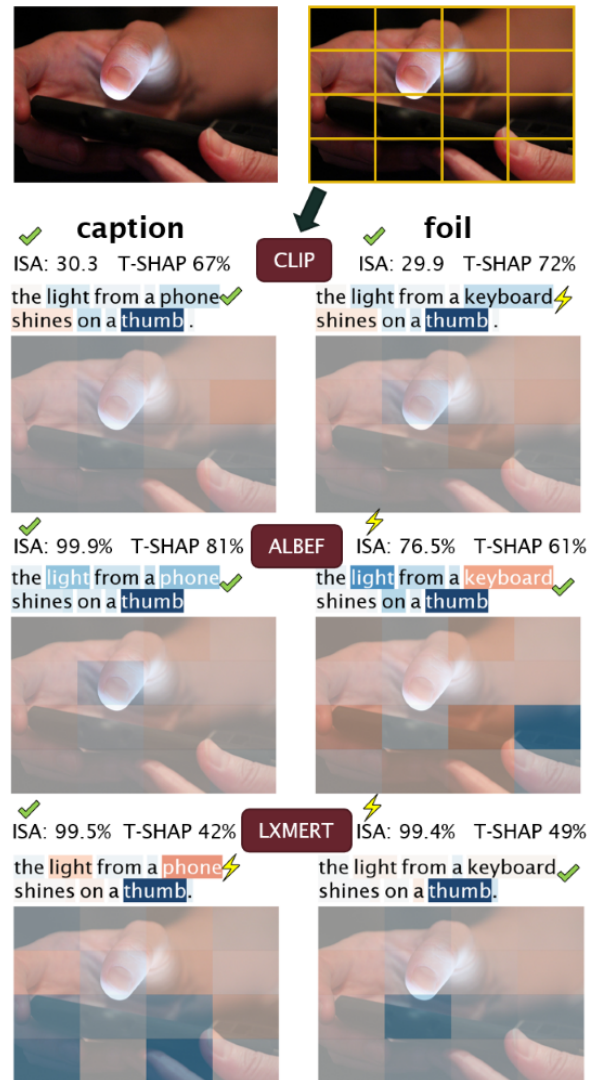


Figure 2: **Low discrepancy noun phrase foil:** Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and highlight the correct / foil token that contributes in the right direction for aligning the image and the caption. With ⚡, we mark incorrect ISA and wrong contribution directions.

alignment pretraining objective affects their behaviour:

For positive existence, where the caption indicates that **an object is present in the image** – as in Fig. 5: *There are children.* – is better handled by the models, delivering more sensible patterns for image-caption pairs. The contribution patterns on the negated version of the existence sentence – the foil *There are no children.* – show that some mod-

els handled the negation correctly (CLIP, LXMERT, ALBEF mscoco and refcoco), while the rest do not.

Negative existence, where the caption indicates that **an object is not present in the image** – as seen in Fig. 6: *There are no humans in the picture.* – seems more difficult to align, since the objects are not present in the image and to assign a high ISA for text mentions that cannot be located, the model needs to understand the negation. The foil, changing the sentence to affirmative – *There are humans in the picture.* – turns the instance into a much simpler case of no image-sentence alignment, as is often seen during pretraining. Unsurprisingly, all models correctly predict a low ISA in Figure 6.

Counting is hard. In Figure 4 for the counting foils in VALSE 🦄, CLIP is the only model that assigns higher ISA for the image-caption pair and not to the image-foil pair. Overall, the contribution patterns look scattered: High visual contributions in the image indicate that the models align the plane object to its mention in the sentence, but we see confused textual contributions from the mentioned number of planes (0 or 4) in the text. This is unsurprising, given the low performance of VL models in counting as highlighted by Parcalabescu et al. (2021a).

D Why not to use Attention for defining a Multimodality Score

D.1 Requirements for a MM Score

For defining a multimodality score that aims at quantifying each modality’s contribution to any model prediction, we need an interpretability method that has crucial properties to do so. With the properties of efficiency, symmetry, dummy variable, additivity (see §3.2), Shapley values provide important ingredients for *sample-based explanations* that can be aggregated in a straightforward way into *dataset-level explanations* for machine learning methods (Covert et al., 2020). Other interpretability methods lack the robustness and theoretical foundation to produce a multimodality score that is comparable to the one proposed in our work.

In particular, attention – while being widely used for generating visually appealing heat-maps – does not fulfil the condition of delivering a fair payout (like Shapley values do) and it is questionable how much high/low attention scores correlate with high/low contributions of input features for system predictions (Jain and Wallace, 2019; Wiegrefe and

Pinter, 2019).⁷ Attention linearly combines input features and determines how much of each token is mixed with every other token. But it does not necessarily mean that a low attention value cannot have a large impact on the decision of the model. In other words, a pinch of salt is enough to make food taste good: Even if the attention score for salt is low, its contribution to the taste of the food (captured by Shapley values) is high.

Attention is present in transformers in multiple layers and to complicate the matter even further, each attention layer contains multiple attention heads. Hence, to visualise attention we need a carefully designed interface, as proposed, e.g., by Jaunet et al. (2021) <https://visqa.liris.cnrs.fr/> to keep a reasonable overview of all attention values. When integrating the multiple attention values and propagating them back to the input to assign relevancy values for image and text tokens, research strives to generate simple explanations that represent the most important tokens and tend to inhibit the rest, as can be seen on the progress from Chefer et al. (2021b) to Chefer et al. (2021a) (cf. Figure 4 in Chefer et al. (2021a)).

D.2 Measuring negative contribution

While Shapley values estimate both the positive and the *negative contributions* of input tokens towards the model prediction – which is relevant for foil words –, attention (Chefer et al., 2021a) allows for positive-only relevance assessments.

In Figures 10 and 11, we have visualised CLIPs attention-based relevancy for the image-caption and foil examples shown in Figures 2 to 7 using the method of Chefer et al. (2021a). On the image side, we observe little to no changes in the attention visualisation, when comparing image-caption to image-foil pairs (cf. Figure 10). Even more, on the text side, both the correct and the foil word carry relatively similar attention scores, with no indication whether this contributes positively or negatively towards the model prediction. Shapley values however, are sensitive to foil words and we can visualise whether the word contributes towards raising the ISA (high image-sentence match) or lowering the ISA (e.g., Figure 3).

Besides the problematic interpretation of attention as feature contribution and the many ways of aggregating and propagating the different attention


⁷Arguably this may be the case when attention weights are high, but it is clearly not the case when attention weights are low.

values to the input, another problem with attention is that it is unclear how to disentangle and aggregate the textual self-attention, visual self-attention, text-to-image attention and image-to-text attention into a single multimodality score that assesses the degree to which a given modality contributes towards the model prediction.

All things considered, we argue that attention is not well-suited as a basis for a multimodality score we aim for in this work, but that Shapley values – as presented in this paper – are, thanks to their theoretical properties (efficiency, symmetry, dummy variable, additivity) and their property of being model-agnostic measurements of input feature contributions.

SEE FIGURES ON FOLLOWING PAGES!



Figure 3: **Low discrepancy** (VALSE  action replacement): Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and an highlight the correct / foil token that contributes in the right direction for aligning the image and the caption. With ⚡, we mark incorrect ISA and wrong contribution directions.

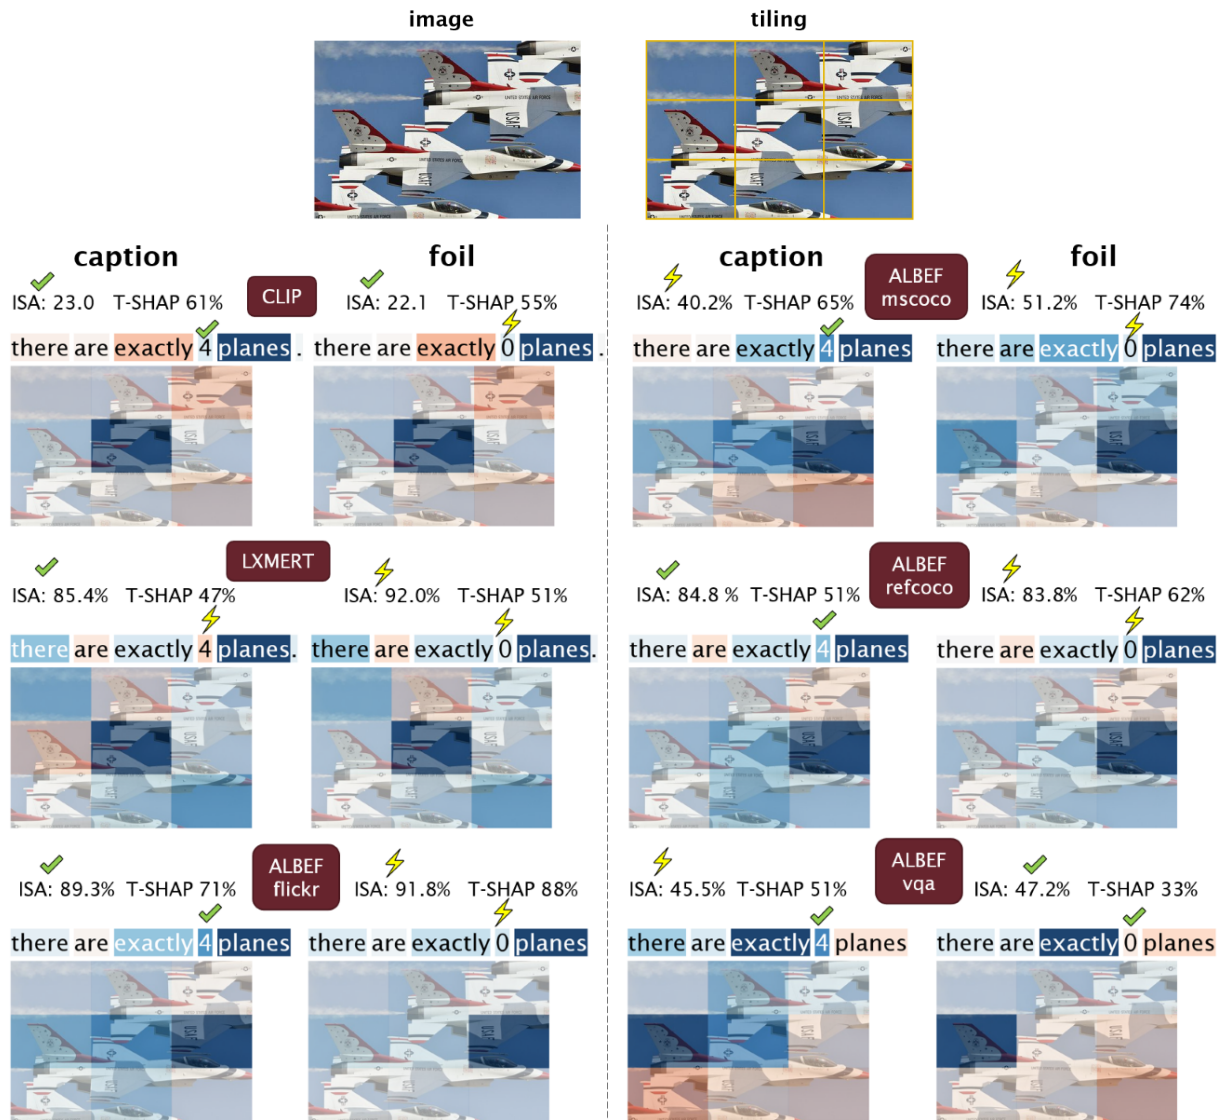


Figure 4: **Low discrepancy** (VALSE counting): Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and an highlight the correct / foil token that contributes in the right direction for aligning the image and the caption. With ✗, we mark incorrect ISA and wrong contribution directions.

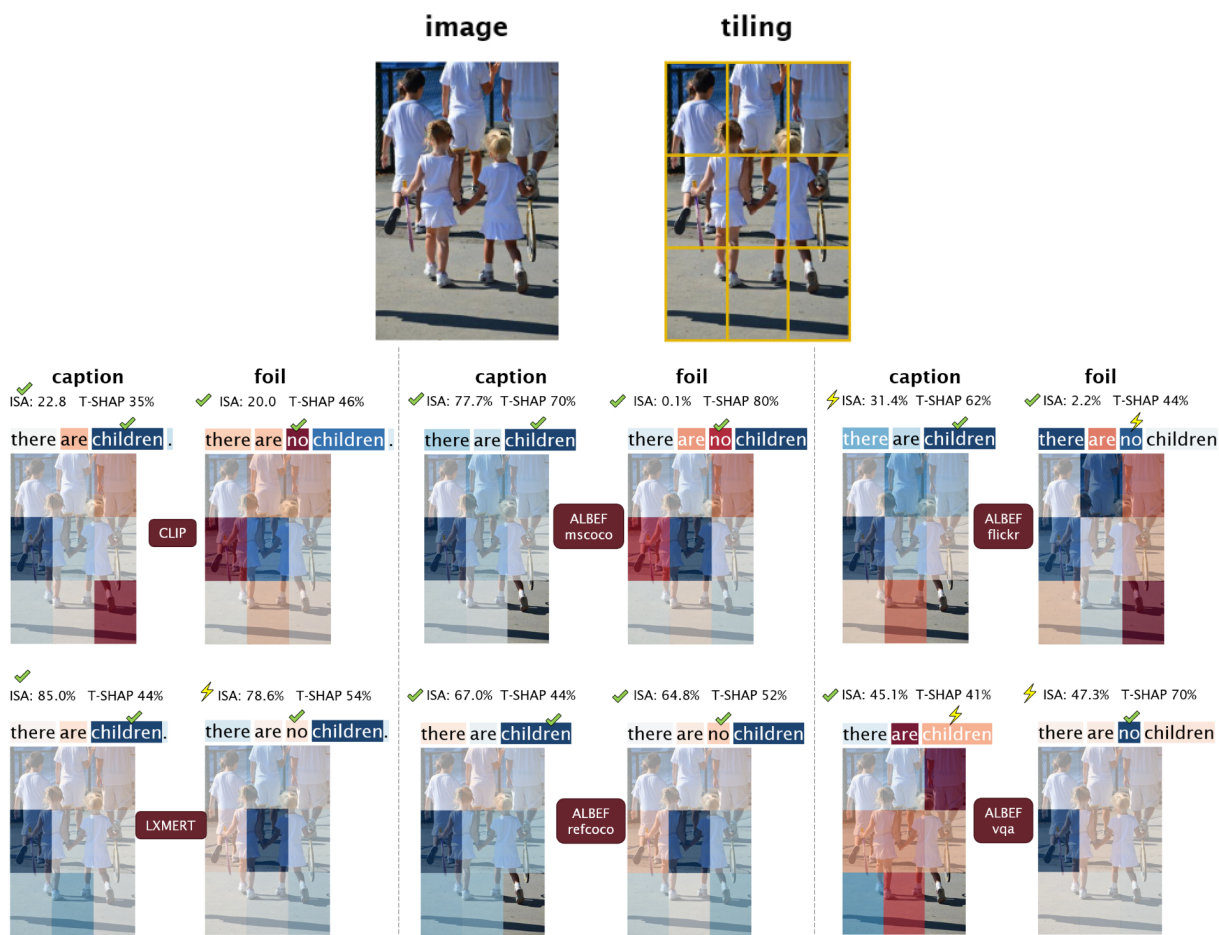


Figure 5: **Low discrepancy** (VALSE *existence positive*): Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and an highlight the correct / foil token that contributes in the right direction for aligning the image and the caption. With ⚡, we mark incorrect ISA and wrong contribution directions.

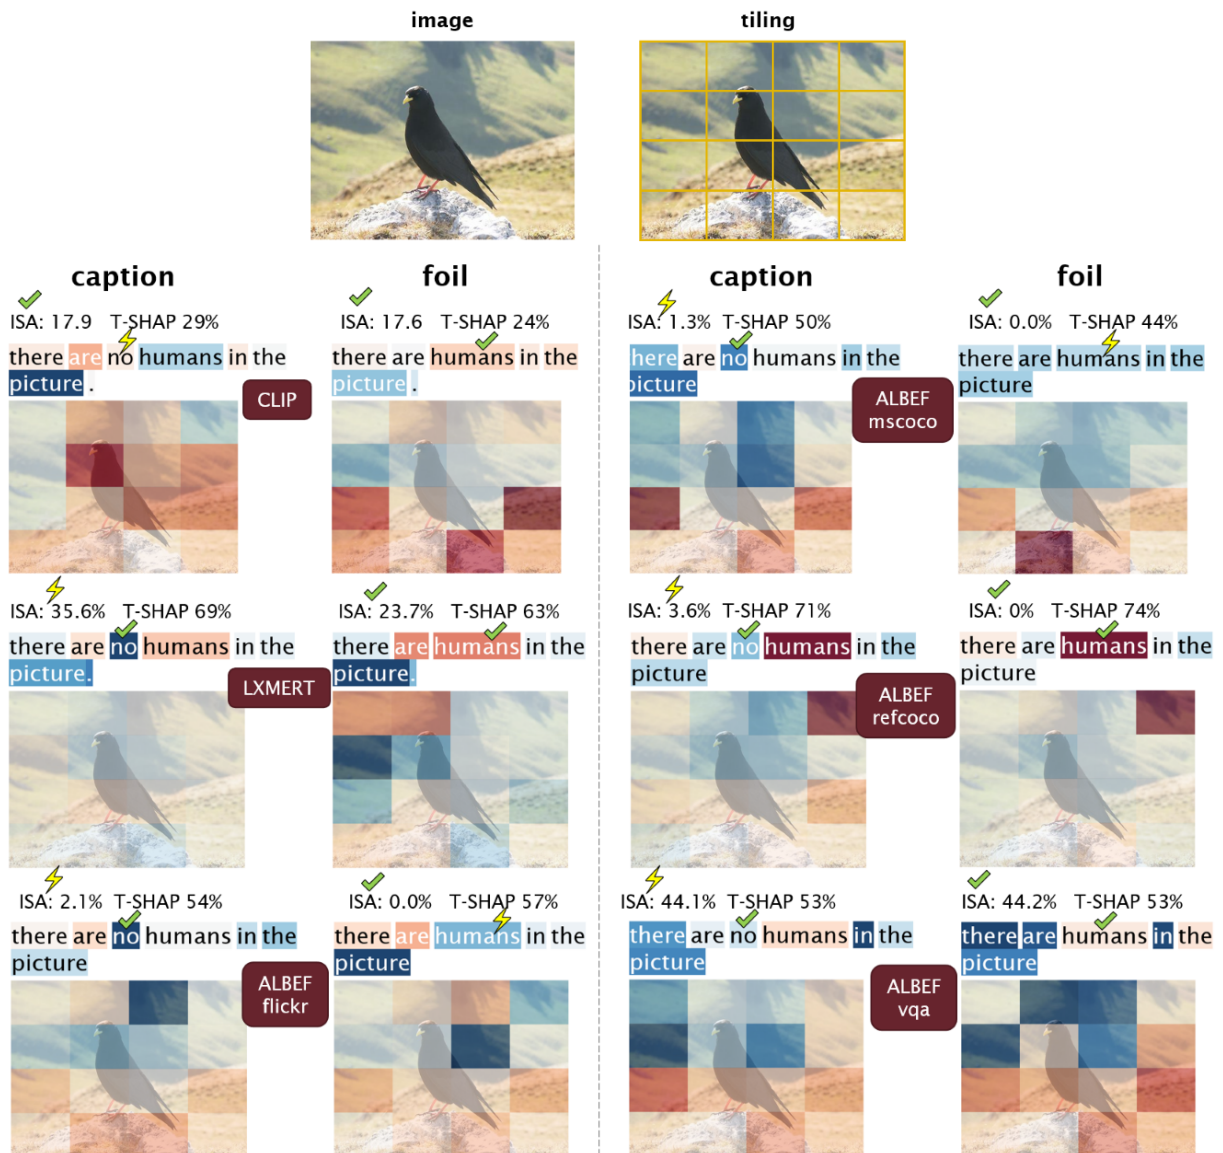


Figure 6: **Low discrepancy** (VALSE ✗ *existence negative* – harder phenomenon than positive existence): Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and an highlight the correct / foil token that contributes in the right direction for aligning the image and the caption. With ✗, we mark incorrect ISA and wrong contribution directions.

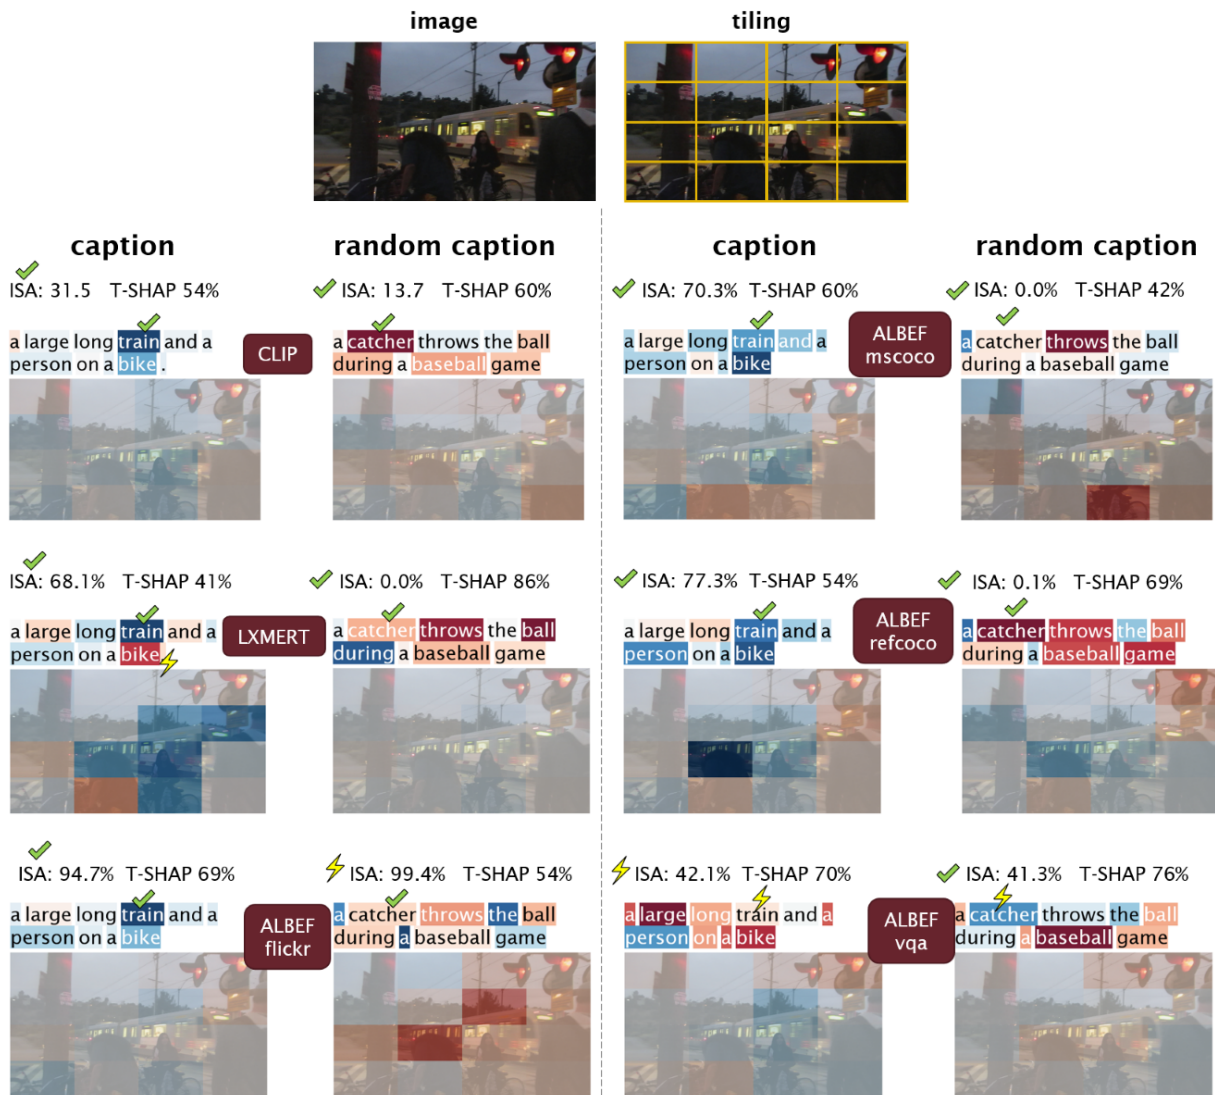


Figure 7: **High discrepancy** (MSCOCO): Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and an highlight one important token that contributes in the right direction for aligning the image and the caption. With ⚡, we mark incorrect ISA and wrong contribution directions.

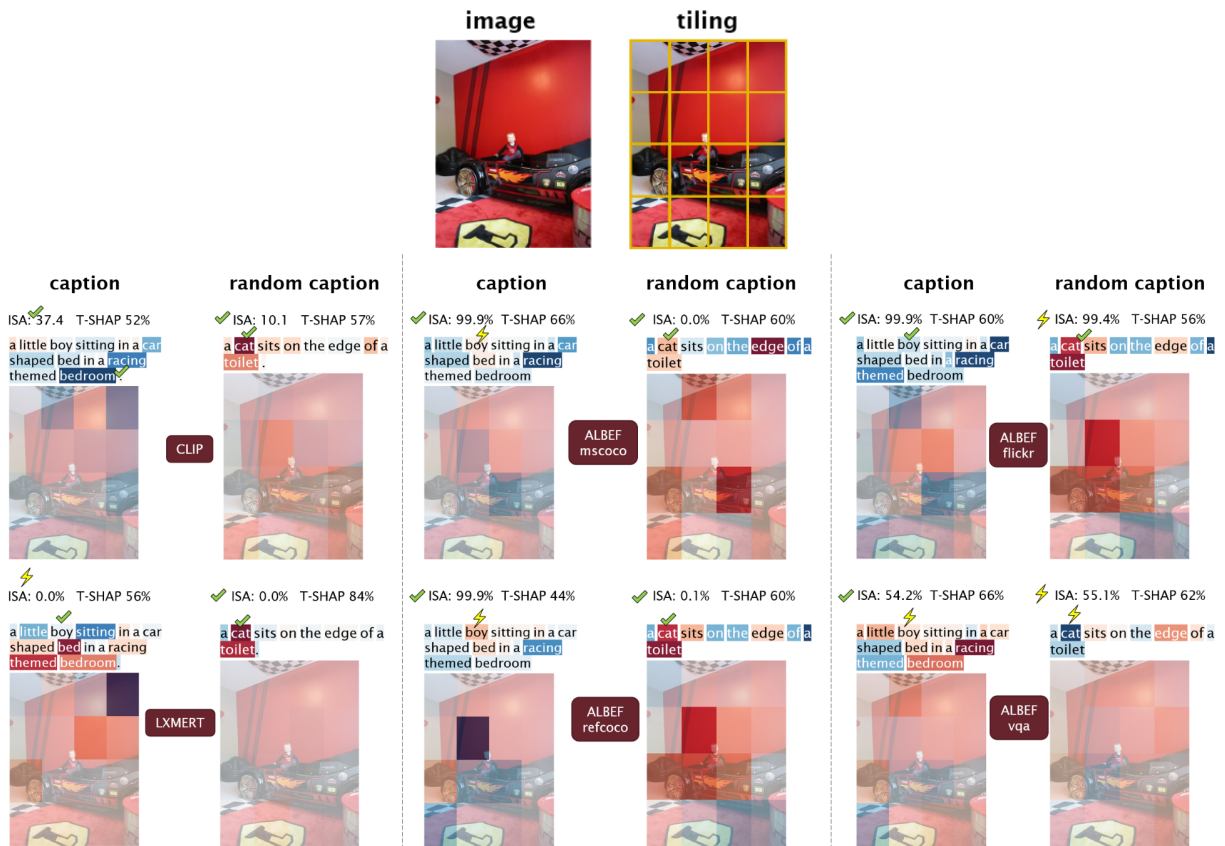


Figure 8: **High discrepancy (MSCOCO) hard example** where the models have trouble recognising the bed: Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - T-SHAP$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With ✓ we mark correct ISA and highlight one important token that contributes in the right direction for aligning the image and the caption. With ⚡, we mark incorrect ISA and wrong contribution directions.

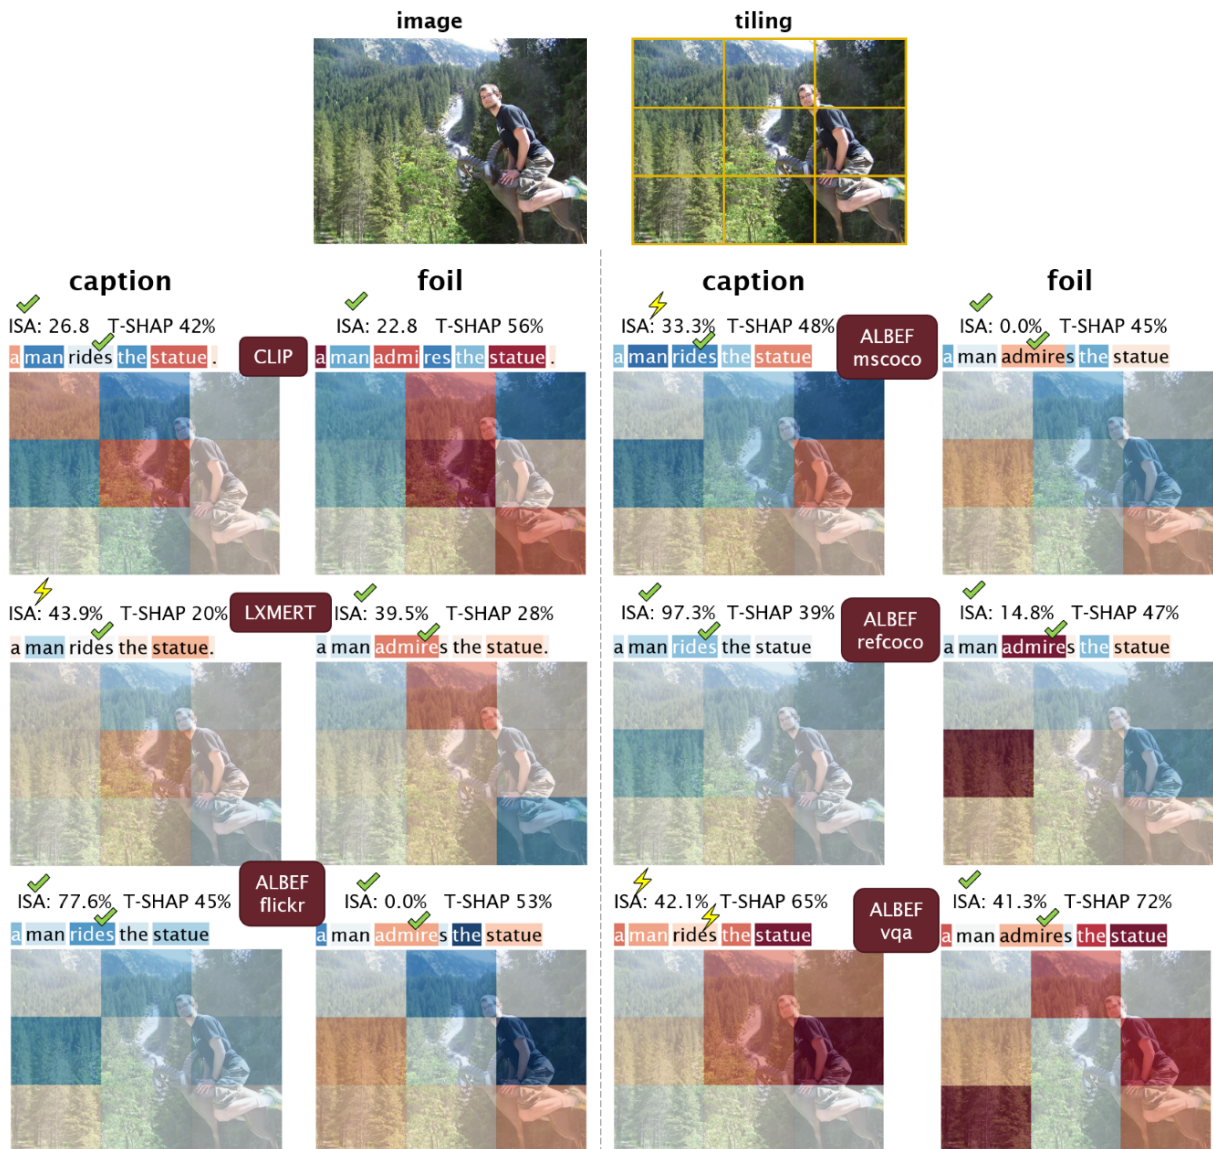


Figure 9: **Low discrepancy** (VALSE action replacement) – *hard example* where models and humans have trouble recognising the goat as a statue): Image-sentence alignment score (ISA) of the six VL models with their textual degree T-SHAP (in %). Each text and image token (image patch) is colour-coded: Blue tokens contribute to a high ISA, while red ones lower the ISA. The visual degree is $100 - \text{T-SHAP}$. Note that the ISA of CLIP is an absolute score, while ALBEF and LXMERT predict ISA probabilities. With we mark correct ISA and highlight the correct / foil token that contributes in the right direction for aligning the image and the caption. With , we mark incorrect ISA and wrong contribution directions.

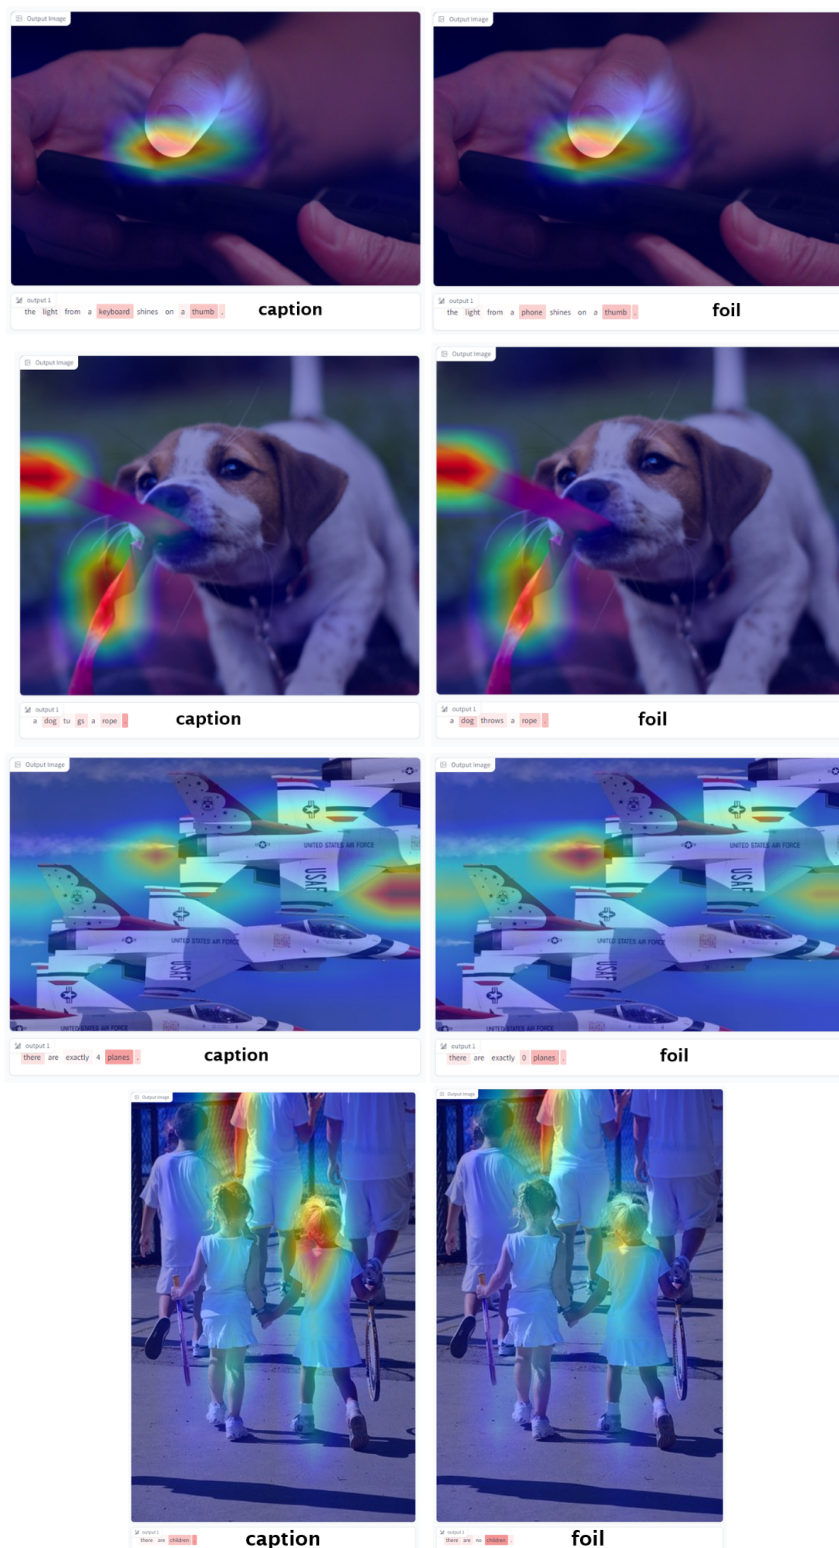


Figure 10: **Low discrepancy.** CLIP results of attention-based relevance visualisation, using the method of [Chefer et al. \(2021a\)](https://huggingface.co/spaces/PaulHilders/CLIPGroundingExplainability) <https://huggingface.co/spaces/PaulHilders/CLIPGroundingExplainability>. Red means high relevancy, blue is zero relevancy and there is no negative relevancy (while Shapley values do allow for positive and negative contributions). Note that the heat-maps give the impression that the relevance irradiates from single spots. This is an artefact from the visualisation since the model works with 32x32 pixel patches and it is these patches that each have a relevance score. For reference: the images are around 500 pixels in height and width.

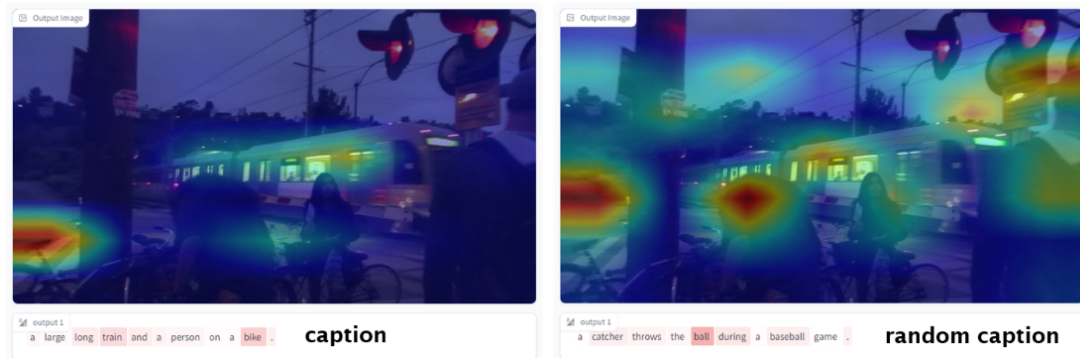


Figure 11: **High discrepancy.** CLIP results of attention-based relevance visualisation, using the method of Chefer et al. (2021a) <https://huggingface.co/spaces/PaulHilders/CLIPGroundingExplainability>. Red means high relevancy, blue is zero relevancy and there is no negative relevancy (while Shapley values do allow for positive and negative contributions). Note that the heat-maps give the impression that the relevance irradiates from single spots. This is an artefact from the visualisation since the model works with 32x32 pixel patches and it is these patches that each have a relevance score. For reference: the images are around 500 pixels in height and width.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations (section 6)
- A2. Did you discuss any potential risks of your work?
Ethical Considerations (section 7)
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract, section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

4

- B1. Did you cite the creators of artifacts you used?
4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
4

C Did you run computational experiments?

6, 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
6

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Not applicable. Left blank.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

5

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

4

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.