

# HNC: Leveraging Hard Negative Captions towards Models with Fine-Grained Visual-Linguistic Comprehension Capabilities

Esra Dönmez\*, Pascal Tilli\*, Hsiu-Yu Yang\*, Thang Vu, Carina Silberer

University of Stuttgart

{esra.doenmez, pascal.tilli, hsiu-yu.yang, thang.vu, carina.silberer}@ims.uni-stuttgart.de

## Abstract

Image–Text–Matching (ITM) is one of the de-facto methods of learning generalized representations from a large corpus in Vision and Language (VL). However, due to the weak association between the web-collected image–text pairs, models fail to show a fine-grained understanding of the combined semantics of these modalities. To address this issue we propose Hard Negative Captions (HNC): an automatically created dataset containing *foiled hard negative* captions for ITM training towards achieving fine-grained cross-modal comprehension in VL. Additionally, we provide a challenging manually-created test set for benchmarking models on a fine-grained cross-modal mismatch task with varying levels of compositional complexity. Our results show the effectiveness of training on HNC by improving the models’ zero-shot capabilities in detecting mismatches on diagnostic tasks and performing robustly under noisy visual input scenarios. Also, we demonstrate that HNC models yield a comparable or better initialization for fine-tuning. Our code and data are publicly available.<sup>1</sup>

## 1 Introduction

Pre-trained Vision and Language Models (VLMs) (Su et al., 2020; Lu et al., 2019; Chen et al., 2020b; Tan and Bansal, 2019), when fine-tuned on downstream tasks, show promising performance thanks to their learned generalized information (or even knowledge) (Zhang et al., 2019; Gan et al., 2020; Hendricks and Nematzadeh, 2021). These models are typically trained on a combination of several datasets under self-supervised training objectives, such as Image-Text-Matching (ITM), Masked Language Modeling (MLM), and Masked Region Modeling (MRM). ITM defines the objective of predicting whether the textual and visual modalities entail

one another. To learn this entailment, for already weakly-associated image–caption pairs, the negative captions are typically sampled from mini-batch training data which results in negative captions that do not align with the image, i.e., the mismatch between the modalities can be detected easily since the images and captions are semantically unrelated. Consequently, the compositional understanding capabilities of VLMs are rather limited, e.g., they tend to show weaknesses in correctly grounding linguistic concepts in their visual counterparts (Bitton et al., 2021; Keysers et al., 2020; Bogin et al., 2021). These VLMs, when tested against foiled inputs, fail against *fine-grained* mismatches in multimodal data (vision and language) (Shekhar et al., 2017a; Hendricks and Nematzadeh, 2021).

To address the aforementioned limitations we focus on improving VLMs by automatically creating a dataset that enables learning from hard negative captions, i.e., negative captions that are minimally contradictory to their corresponding images. We state the hypothesis that such hard negative captions increase the general comprehension capabilities of pre-trained VLMs. We summarize our contributions as follows:

1. We introduce **Hard Negative Captions (HNC)** for ITM training with systematically created hard negatives: 12 linguistically-motivated types of captions<sup>2</sup> that locally describe an image with their hard negative counterparts that are minimally contradictory to the given image.
2. To the best of our knowledge, we are the first to **leverage scene graph information** (Krishna et al., 2017) for automatically creating hard negative captions (fine-grained misaligned image–text pairs) for ITM training. This enables us to control (1) the seman-

\*These authors contributed equally to this work.

<sup>1</sup><https://github.com/DigitalPhonetics/hard-negative-captions> under MIT License.

<sup>2</sup>Our code allows everyone to easily add new caption types.

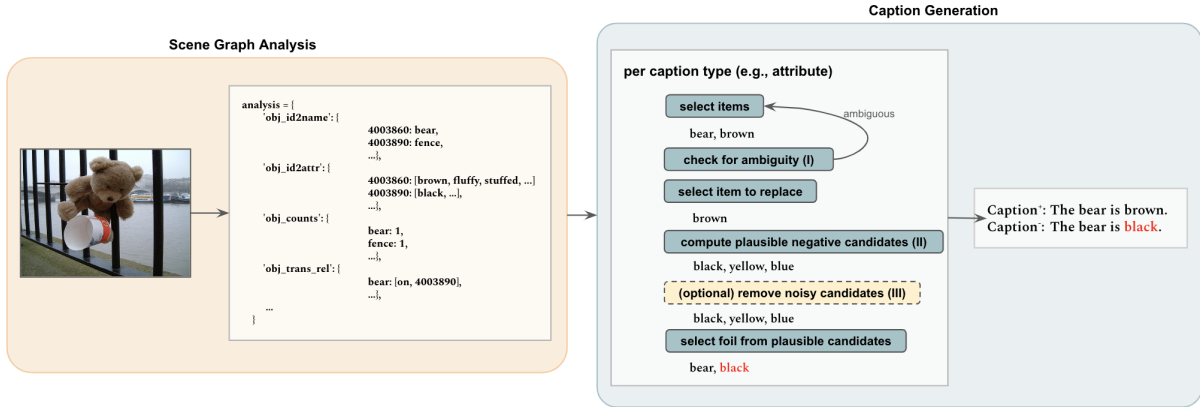


Figure 1: An illustration of our caption generation procedure. For each scene graph (that belongs to exactly one image) we run through this pipeline to generate hard negative captions. Details on the modules marked with Roman letters (I, II, and II) can be found in Sec. 3.

tics of the hard negatives with multiple mismatch types, and (2) the level of compositional complexity in fine-grained mismatches. Our method is resource-lean in constructing the hard negatives, and flexible in that it can be extended to other phenomena which is necessary for this fast-developing **Vision and Language (VL)** research field.

3. We propose a challenging human-annotated test set to benchmark **VL** models’ capabilities on several skills and levels of compositional understanding.
4. We perform an extensive study across various tasks and show models’ improvement in fine-grained cross-modal comprehension in zero-shot settings. Additionally, we show that models further trained on **Hard Negative Captions (HNC)** can serve as a better initialization point for downstream task fine-tuning.

## 2 Related Work

**Probing VLMs for fine-grained visual grounding** Several works revealed shortfalls in visual grounding capabilities of **VLMs** at various levels by creating foiled visual descriptions in which they alter the nouns (Shekhar et al., 2017c), words belonging to other **Part-of-Speech (PoS)** tags such as adjectives or adverbs (Shekhar et al., 2017b), S(ubject)–V(erb)–O(object) triples (Hendricks and Nematzadeh, 2021), person entities (Park et al., 2022). These studies collectively suggest that **VLMs** struggle with fine-grained image–caption matching. Moreover, several works studied the compositional understanding of **VLMs** in

visual grounding. Thrush et al. (2022) propose Winoground to evaluate visual grounding robustness using captions with the same set of words but different syntactic structures. Their findings suggest that **VLMs** exhibit bag-of-words behavior (Diwan et al., 2022). Bogin et al. (2021) introduce **COMpositional VISual REASONing (COVR)** to examine models’ compositional generalization on unseen logical operations, e.g., quantifiers or aggregations, and conclude that reasoning over complex structures remains challenging. While above works aim to create probing datasets to identify **VLMs**’ potential shortfalls in visual grounding, our research goal goes beyond that: we propose a creation method for large-scale **ITM** datasets, useful for further pretraining (or fine-tuning) models towards fine-grained cross-modal comprehension abilities.

**Addressing shortfalls in fine-grained visual grounding capabilities of VLMs** Given that **VLMs** are usually pre-trained with web-crawled weakly-aligned image–caption pairs, e.g., Conceptual Captions (Sharma et al., 2018), their ability to address cross-modal misalignments is questionable. The aforementioned empirical probes support this claim and suggest that **VLMs** tend to suffer from overprediction in that they consider a somewhat related image–caption pair to be associated. Previous works address this issue as a part of the training strategy (Liu and Ye, 2019; Zhou et al., 2020; Chen et al., 2020a, 2022), the model architecture (Messina et al., 2021; Zhang et al., 2022), or by augmenting training data (Shekhar et al., 2017c; Faghri et al., 2018; Gupta et al., 2020). We contribute to the last line of research and propose to

augment hard negative captions for ITM training by leveraging scene graphs towards achieving a fine-grained VL comprehension.

### 3 HNC: Hard Negative Captions

We use the structural information provided by scene graphs (Krishna et al., 2017) to automatically generate **hard negative image–text pairs** with various caption types. We leverage the ground-truth scene graphs provided by the GQA (Hudson and Manning, 2019) dataset, which contains a total of +80K images paired with scene graphs in the training and validation set.

We define a positive caption as a textual description that **locally describes** an image, i.e., the caption describes a part of the image and does not aim to provide an exhaustive description of the entire scene. A hard negative caption, in turn, is **minimally contradictory** to the image and is obtained by altering a piece of information in the corresponding positive caption, i.e., without that minimal change, it would be a positive caption.

#### 3.1 Automatic Caption Generation

Given an image, we first extract structured information from its corresponding scene graph and use it to create caption pairs for each of the caption types which can be found in Figure 2. In the caption generation process, we apply the following procedure: **1)** Check whether the information allows constructing the particular caption type. If yes, **2)** instantiate a positive caption with the pre-defined caption template. **3)** Instantiate a negative caption using the same template by replacing a piece of information in the positive caption. We provide an illustration of our workflow in Figure 1.

**Ambiguity (I)** We apply a set of heuristics that filter out potentially ambiguous captions (see A.2 for details). These heuristics prevent generating captions that refer to: **a)** multiple instances of the same object class, e.g., *the sheep that is to the right of the sheep*; **b)** relations between body parts, e.g., *the ear is to the left of the nose*; **c)** relations between objects with one of them typically covering a large area in the scene, e.g., *the grass is to the left of the ball*. Note that these heuristics are applied to both the positive and the negative captions.

**Plausible negative value sampling (II)** There are several ways to sample a negative value as the *foiled* piece of information. We introduce the set-


ting used in our experiments in the following and discuss the other options in A.2. An ideal *foiled* hard negative caption is *visually challenging*, *sensible*, and *semantically similar* to the positive caption. To ensure that the negative caption is visually challenging, we sample a negative value from within the scene, i.e., the candidate values are extracted from the same scene graph. Ensuring that the negative caption is sensible and at the same time semantically similar to the positive one is more challenging. For this, we need to satisfy two conditions: **a)** A negative value must be valid in terms of semantic class constraints, i.e., we cannot replace apple by table in *The girl is eating an apple*. **b)** Concept co-occurrence distributions in the negative and the positive captions should be similar to avoid spurious correlations. To achieve sensibility, we create look-up tables that help us define which candidates are valid for a given word. We then sample a negative value from these valid candidates following the distribution of the positive captions. The candidates are further filtered to avoid potential noisy replacements which we discuss in the following.

**Noisy negative values (III)** To minimize potential issues caused by **partial** or **incomplete** scene graphs (Chang et al., 2023), we employ a set of heuristics designed to detect missing spatial relations between a pair of objects in a scene. We achieve this by leveraging the bounding-box values of the objects obtained from the ground truth scene graphs. Given a spatial relation between two entities annotated in a ground-truth scene graph; when replacing an entity or the relation with another value to create the negative caption, if this relation between the entities is not encoded in the scene graph, we check the bounding-box annotations to see if there does exist this spatial relation between the entities. If this is the case, we remove the value from the set of valid candidates<sup>3</sup>.

#### 3.2 Caption Types

We design 12 caption types grouped into 5 categories, illustrated in Figure 2 (together with the construction templates, an image, and examples): **1) attribute-based**, **2) relation-based**, **3) counting-based**, **4) existence-based**, and **5) reasoning-based**. The first three of these categories focus on either an object, an attribute, or a relation, while the existence and the reasoning-based types are some combinations of all other types.

<sup>3</sup>Details are given in A.2.



Caption Type	Template	Example
<b>attribute</b>	The {obj} is/are {attr}.	The bowl is <u>teal</u> ( <i>white</i> ).
<b>attribute_relation</b>	The {attr} {subj} is/are {pred} the {obj}.	The <u>black and white</u> ( <i>gray</i> ) cat is on the table.
<b>relation</b>	The {subj} is/are {pred} the {obj}.	The bowl is to the left of ( <i>to the right of</i> ) the cat.
<b>relation_attribute</b>	The {attr} {subj} is/are {pred} the {attr} {obj}.	The jars are to the left of the white <u>door</u> ( <i>table</i> ).
<b>object_count</b>	There are (n) {obj}.	There are <u>two</u> ( <i>three</i> ) jars.
<b>object_compare_count</b>	There are (fewer/more/as many) {obj1} than/as {obj2}.	There are <u>more</u> ( <i>fewer</i> ) apples than jars.
<b>verify_object_attribute</b>	There is (no/at least one) {obj} that is {attr}.	There is <u>no</u> ( <i>at least one</i> ) table that is plastic.
<b>verify_object_relation</b>	There is (no/at least one) {subj} that is {pred} the {obj}.	There is <u>at least one</u> ( <i>no</i> ) cat that is to the right of the bowl.
<b>AND_logic_attribute</b>	There is/are both {attr1} {obj1} and {attr2} {obj2}.	There are both a <u>white</u> ( <i>metal</i> ) door and a teal bowl.
<b>AND_logic_relation</b>	There are both {subj1} {pred1} the {obj1} and {subj2} {pred2} the {obj2}.	There are both apples in the bowl and <u>jars</u> ( <i>coats</i> ) to the left of the door.
<b>XOR_logic_attribute</b>	There is/are either {attr1} {obj1} or {attr2} {obj2}.	There is either a white door or a <u>brown</u> ( <i>teal</i> ) bowl.
<b>XOR_logic_relation</b>	The {subj} is/are {pred} either the {obj1} or the {obj2}.	The cat is in front of either the door or the <u>apples</u> ( <i>curtain</i> ).

Figure 2: (a) an illustration of one image and (b) exemplary captions based on the displayed caption type templates.

**Attribute-based** For attribute-based modality mismatches, we design two templates: (a) **attribute**, (b) **attribute\_relation**. The former simply requires models to verify whether the attribute of an object is described correctly in the caption, while the latter further challenges models’ understanding of an object’s attribute in a relational subgraph.

**Relation-based** These caption types are designed to detect a modality mismatch in relational subgraphs by foiling either the subject, the object, or the predicate to create the negative caption. There are two template types: (a) **relation**, (b) **relation\_attribute**. The first one aims to harness a model’s sensitivity towards modality mismatches occurring in a relational subgraph. The second type extends the previous one by adding (an) attribute(s) to the entities in the relational subgraph, which requires a model to reason compositionally.

**Counting-based** Two templates target counting-based modality mismatches: (a) **object\_count** which refers to the number of objects of the same class in the visual modality, and (b) **object\_compare\_count** which compares the counts of two object classes using comparative quantifiers, i.e., *fewer*, *more*, *as many as*, without mentioning the actual counts.

**Existence-based** This type addresses the existence of an entity in the visual modality. Two templates are provided for this: (a) **verify\_object\_attribute** grounds the entity in the scene with the help of an adjective modifier, and (b) **verify\_object\_relation** does so with the help of its relation to another object in the scene.

**Reasoning-based** For our reasoning-based captions, we focus on the **AND** and **XOR** logic reason-

ing types. For each type we provide two templates, one introduces a foiled attribute and the other introduces a foil in the relational subgraph. These hard negative captions are very complex, and the captions contain a lot of information of which only a small piece is incorrect. Thus, any shortcut in reasoning should result in an incorrect prediction.

### 3.3 Dataset Statistics

We follow the official splits of the **Visual Reasoning in the Real World (GQA)** dataset (Hudson and Manning, 2019) to generate captions. The training set contains 74, 942 images, the validation set 10, 696 images.

The statistics of the *clean-strict* variation of our dataset (the debiased one according to our iterative quality control explained in Section 7.1) is as follows: For the training set we create 242 captions for each image on average, and for the validation set 239 captions on average, resulting in a total of 16, 416, 392 for the training set and 2, 314, 832 for the validation set. The average caption length is 10 tokens. Due to our automatic caption generation procedure, we receive equal data distributions and caption lengths for the training and validation splits. Details are given in Table 12.

## 4 Human-annotated Challenge Set

As we rely on scene graphs and an automatic generation procedure to create our training and validation data, we believe in the importance of providing a quality test set ideally free from any noise introduced by our automatic procedure. To this end, we had 19 annotators<sup>4</sup> to write down pairs of captions for all caption types.

<sup>4</sup>All students of an international (under-)graduate program with advanced English proficiency. We informed the par-

	VILBERT				VISUALBERT			
	VOLTA	FOIL	HNC <sub>subset</sub>	HNC <sub>full</sub>	VOLTA	FOIL	HNC <sub>subset</sub>	HNC <sub>full</sub>
attribute	44.1	52.5	57.5	<b>78.2</b>	45.0	51.0	65.7	77.7
attribute_rel	47.5	54.0	54.3	75.0	47.5	49.5	60.4	<b>79.0</b>
relation	46.2	54.7	55.0	62.8	47.3	52.2	56.1	<b>65.7</b>
relation_attr	47.0	54.4	55.4	66.4	47.0	52.8	61.0	<b>67.0</b>
obj_count	51.0	49.5	55.9	<b>73.0</b>	49.0	48.0	62.7	66.0
obj_comp_count	50.0	48.5	57.2	58.5	48.5	51.0	58.2	<b>62.0</b>
verify_obj_attr	49.0	50.5	52.1	<b>76.0</b>	49.0	50.0	57.6	75.0
verify_obj_rel	49.5	51.0	56.3	59.0	48.5	48.5	56.5	<b>61.5</b>
AND_logic_attr	48.5	51.5	52.2	73.5	50.0	51.0	56.6	<b>74.0</b>
AND_logic_rel	52.5	52.0	52.7	57.0	48.5	52.0	52.7	<b>58.5</b>
XOR_logic_attr	50.0	51.0	51.7	65.5	52.5	50.0	57.3	<b>68.0</b>
XOR_logic_rel	51.0	49.5	57.6	59.0	51.5	50.5	57.9	<b>66.5</b>
<b>all</b>	48.3	51.6	54.1	66.4	48.3	50.5	58.6	<b>67.9</b>

Table 1: Binary classification accuracy on HNC test set.

**Annotation guidelines** For each image, the annotators were asked to provide a positive and a negative caption pair per their assigned caption type(s). We set the following conditions for the annotation: **1)** Stay true to the vocabulary: The words in the captions must come from within the global GQA vocabulary. **2)** Choose visually challenging objects: The objects introduced as the *foiled* information in the captions must come from within the scene. **3)** Chose linguistically challenging attributes and predicates: The attributes and predicates introduced as the *foiled* information in the negative captions must be linguistically challenging, e.g., brown dog  $\rightarrow$  black dog; meaning that both captions are equally plausible. The annotators were instructed to skip creating a caption pair for the respective type in cases where at least one of the negative or positive captions cannot be created for a given image.

**Dataset statistics** In total, we obtain captions for 100 images. With 12 caption types, annotation results in 3201 captions with an average length of 8.42. Per caption type, we get 32 captions on average. The annotated captions went under a quality check performed by another group that did not take part in the annotation.

## 5 Experiments

We use the Visiolinguistic Transformer Architectures (VOLTA) framework (Bugliarello et al., 2021) as a unified testing suite to run our experiments.

Participants about the use of their data and compensated them with 13€/hour, above the German minimum wage.

Specifically, we use its controlled setup<sup>5</sup> and initialize all five models from the pre-trained weights provided by VOLTA. We then further train the ITM head on the training set of both HNC and FOIL. For a fair comparison with FOIL, which is substantially smaller (197k data points in the training split); in addition to the full-data setting (HNC<sub>full</sub>), we include an HNC<sub>subset</sub> setting subsampled to 197k data points. We experiment with both single-stream and dual-stream architectures and analyze their performance difference (if any): UNITER, VISUALBERT, VILBERT, LXMERT, VL-BERT (Tan and Bansal, 2019; Chen et al., 2020b; Lu et al., 2019; Li et al., 2019; Su et al., 2020)<sup>6</sup>. To test whether training on HNC yields similar results on more recent and bigger models, we include experiments with BLIP (Li et al., 2022), which are presented in A.1.2.

**Evaluation** We compare the performances of the models before and after further pre-training on HNC on two types of tasks: (1) Linguistic comprehension tasks, and (2) Real-world downstream reasoning tasks (Sec. 5.1 and 5.2, resp.). The HNC<sub>subset</sub> results are averaged over five randomly sub-sampled splits, while the rest of the results come from a single run.

### 5.1 Visio-Linguistic Comprehension Tasks

**HNC** We use the manually created, high-quality test set to assess the ability of fine-grained image-text understanding (see Sec. 3 for details about the automatically-created training and validation sets

<sup>5</sup>The controlled setup uses the same pre-training objectives and datasets across models to allow systematic comparison.

<sup>6</sup>Model and hyperparameter details are given in A.1.

	VILBERT				VISUALBERT			
	VOLTA	FOIL	HNC <sub>subset</sub>	HNC <sub>full</sub>	VOLTA	FOIL	HNC <sub>subset</sub>	HNC <sub>full</sub>
existence	47.8	49.8	52.1	59.8	46.9	49.3	58.9	<b>63.1</b>
plurals	50.0	50.4	51.4	51.4	49.5	50.3	51.8	<b>52.8</b>
counting_small_quant	49.4	49.3	51.1	58.6	49.6	50.0	53.2	<b>58.8</b>
counting_adversarial	49.5	52.5	<b>54.6</b>	53.2	48.9	50.7	50.4	50.2
counting_hard	49.8	49.6	49.9	52.4	49.6	49.7	50.3	<b>53.2</b>
relations	49.8	49.8	50.9	50.9	49.8	50.0	50.4	<b>51.4</b>
actant_swap	48.1	54.6	55.8	58.0	47.9	51.5	<b>58.3</b>	57.6
action_replacement	47.0	53.0	51.6	52.9	47.8	50.3	51.0	<b>54.3</b>
coreference_standard	49.9	<b>50.1</b>	50.0	47.2	50.0	49.9	49.7	49.9
coreference_hard	<b>50.0</b>	<b>50.0</b>	<b>50.0</b>	48.2	<b>50.0</b>	<b>50.0</b>	49.8	48.9
foil_it	46.0	<b>77.0</b>	50.4	51.8	43.7	<b>79.0</b>	51.5	<b>54.8</b>
<b>all</b>	48.8	50.9	51.6	53.0	48.4	50.2	52.3	<b>54.4</b>

Table 2: Binary classification accuracy on VALSE (Parcalabescu et al., 2022) under zero-shot evaluation. For the models trained on FOIL dataset, we do not calculate the accuracies obtained from the foil it splits (marked red) into the averaged values.

and Sec. 4 for the human-annotated test set).

**Vision And Language Structured Evaluation (VALSE)** is a benchmark focusing on various linguistic phenomena (Parcalabescu et al., 2022).

## 5.2 Real-World Reasoning Tasks

**Commonsense Probing Task (CPT)** measures the commonsense knowledge level of task-agnostic visually pre-trained models on the CWWV<sub>Img</sub> dataset (Yang and Silberer, 2022). We consider this task as a real-world scenario in that associated images are automatically retrieved, which may lead to noisy image–text pairs (see A.3.3 for the complete task description).

**GQA** is a dataset designed for real-world visual reasoning and compositional question answering. Unlike the aforementioned tasks that test zero-shot capabilities, we investigate whether our weight initialization after HNC further pre-training serves as an improved starting point when fine-tuning on GQA. Therefore, we compare VOLTA checkpoints and further pre-trained ones (HNC) after their fine-tuning on GQA. The performances are reported on the GQA testdev split.

## 6 Results

We report the results, i.e. classification accuracies, on the aforementioned four tasks<sup>7</sup>. We compare dual-stream and single-stream models to assess the effects of different modality integration methods on models’ ability to detect mismatches. We display the results obtained from our further pre-trained

<sup>7</sup>We only discuss the statistically significant results.

weight initializations as HNC<sub>subset</sub> and HNC<sub>full</sub>, the ones obtained from training on FOIL-COCO as FOIL, and the official VOLTA weight initialization as VOLTA<sup>8</sup>. The best results are shown in **bold**.

### 6.1 Visio-Linguistic Comprehension Tasks

**HNC** Table 1 displays the results obtained on our human-annotated test set. Zero-shot performances of VOLTA checkpoints on the majority of the caption types are close to random baseline (50%) showing that the dataset is not trivially solvable. We observe a strong **under-prediction of entailment**<sup>9</sup> in models initialized from VOLTA checkpoints before undergoing our further pre-training on HNC dataset, suggesting that the positive captions are equally hard to align with the visual modality for these models. This might be because the web-retrieved captions lack compositionally complex information, i.e., information about multiple objects along with their attributes or relations to other objects. After further pre-training on HNC (see Tab.1, col.HNC<sub>full</sub>), we observe a large improvement in all caption types which showcases the effectiveness of our dataset in teaching fine-grained alignment of the visual and textual modality.

**VALSE** As shown in Table 2, further pre-training on HNC largely improves: **existence**, **counting\_small\_quant**, **counting\_adversarial**, **counting\_hard**, **actant\_swap**, **action\_replacement**, and **foil\_it**<sup>10</sup>. Also, HNC<sub>subset</sub> achieves better results

<sup>8</sup>We only display results from one single- and one dual-stream model in Table 1, 2, 3, and 4. Complete results can be found in Table 5, 6, 7, and 8 resp. in A.

<sup>9</sup>False negative prediction for the positive pairs.

<sup>10</sup>We provide more findings with analysis in A.3.3.

	LXMERT				UNITER			
	VOLTA	FOIL	HNC <sub>subset</sub>	HNC <sub>full</sub>	VOLTA	FOIL	HNC <sub>subset</sub>	HNC <sub>full</sub>
taxonomic	51.55	52.46	54.78	54.8	54.04	56.69	57.29	<b>58.5</b>
similarity	43.01	43.17	44.38	46.43	46.43	49.53	50.99	<b>55.75</b>
part-whole	53.73	50.13	52.93	56.48	63	63.95	64.1	<b>69.01</b>
spatial	55.6	52.72	55.04	56.79	57.41	57.47	53.32	<b>57.97</b>
temporal	49.23	49.81	47.56	<b>50.24</b>	47.86	46.59	46.53	46.27
all	55.43	52.22	55.94	55.49	58.53	59.29	59.27	<b>62.32</b>

Table 3: Classification accuracy on CPT (Yang and Silberer, 2022) with CWWV<sub>Img</sub> under zero-shot evaluation.

compared to FOIL on average, which suggests that HNC contains more diverse and better quality captions to learn from than FOIL-COCO. The large improvement we observe in **existence** type in VALSE shows the effectiveness of our existence-based captions (verify\_obj\_attr, verify\_obj\_rel). We attribute the large improvement in **actant\_swap** to our dedicated control of subjects and objects in relational captions (relation\_subj, relation\_obj, AND\_logic\_rel, and XOR\_logic\_rel). As for the **foil\_it**, we see a similar effect, i.e., controlling nouns (subjects and objects) in hard negatives helps models to better ground the object in the visual scene and not be confused by another (potentially semantically similar) object.

**Counting\_adversarial** tests for the shortcut biases by purposefully assigning a more common number as the *foiled* information in the VALSE captions where the original caption contains a number that is typically less common in these models’ pre-training data. Not only do we see a large performance increase in **counting\_small\_quant**, we also see an improvement in **counting\_adversarial** and **counting\_hard** captions showing that the models benefit from the diverse number sampling in HNC’s training data construction.

Further, we only observe a marginal improvement in **plurality** which is not surprising as we do not create captions that target this type specifically. Also, HNC pre-training does not affect **coreference\_standard** and **coreference\_hard** too much (slight performance decrease if any). Just like the **plurality**, we expect these numbers as we do not address such types in this work. Future work can easily extend to **plurality** by creating a caption type that solely controls the information on the plurality of the objects in the scene. The same can be done for **coreference** by combining several pieces of information about an entity using a referent word.

## 6.2 Real-World Reasoning Tasks

**CPT** Table 3 shows substantial zero-shot performance gains after further pre-training on HNC<sub>full</sub>; particularly on single-stream models. We speculate that our HNC pre-training could drive single-stream encoders to be more sensitive towards cross-modal inconsistencies and strengthen the importance of the textual modality under noisy visual input scenarios. For dual-stream models, the overall improvement is limited, possibly due to the design of certain layers that primarily perform inter-modal attention which restricts the flexibility of balancing the influence of different modality inputs during inference.

Regarding the individual commonsense dimensions, all HNC models demonstrate improvement on **taxonomic**, **similarity**, **part-whole**. This could be explained by their sparser distribution of concrete concepts (Yang and Silberer, 2022), resulting in less semantic correspondence between the extracted images and their textual counterparts (see A.3.3, Fig.6). Overall, the outcome suggests the importance of having hard negative captions in ITM pre-training to enhance the robustness of VL models in handling noisy visual inputs during inference. Both the scale and the quality play a role, as models show greater improvement on these dimensions when further pre-trained on HNC<sub>subset</sub> compared to FOIL-COCO (see col.FOIL & HNC<sub>subset</sub> of tab.3). However, the hard negative pretraining does not benefit much to **spatial** and **temporal**. Especially for **temporal**, the question token and the image retrieved for the answer token are subject to mismatches due to the natural temporal order, e.g., *run out of money* is a consequence of *buying food*, the image of *money* does not correspond to *food* (see A.3.3, Fig.7).

**GQA** We summarize our results on the GQA (Hudson and Manning, 2019) testdev split in Table 4. As we are required to fine-tune on GQA to re-

ceive meaningful results, we distinguish between the weight initialization from the official **VOLTA** pre-training and the initialization from our further pre-training on **HNC**. At first glance, our initialization points achieve higher accuracy across all five models. The results are statistically significant for LXMERT, UNITER, and VISUALBERT. For the single-stream models, VISUALBERT benefits the most from further pre-training on **HNC**. For the dual-stream, LXMERT shows larger performance gains. Generally, the dual-stream vs. single-stream modality integration does not seem to have an influence on how much the respective models benefit from further pre-training on **HNC**. Nonetheless, the overall results support our hypothesis that further pre-training **VL** models on more fine-grained mismatching data (in the form of hard-negative captions) improves models’ cross-modal reasoning capabilities.

	LXMERT		VISUALBERT	
	VOLTA	HNC <sub>full</sub>	VOLTA	HNC <sub>full</sub>
Accuracy	53.48	55.45	53.51	<b>56.85</b>

Table 4: Results on the **GQA** (Hudson and Manning, 2019) testdev split.

## 7 Dataset Analysis

Next, we analyze our caption generation process: how robust are the different negative sampling strategies, and which results in less/more linguistic bias that a model could exploit as a shortcut? We discuss the challenges of automatic hard negative caption generation, the biases introduced in captions as a result of this automatic procedure, and how to mitigate them. We then perform a modality ablation study to ensure the quality of our human-annotated test set. We provide further qualitative analyses in Appendix **A.3**.

### 7.1 Caption Generation: An Iterative Process

Our final caption generation process is a product of a series of refinement iterations. At each iteration, we train and evaluate a **Language Model (LM)** (BERT, Devlin et al., 2019) on our captions and use the accuracy scores as a proxy to measure linguistic bias. Throughout this process, we found that, for example, replacing an attribute of a visual object with another attribute from the scene without any further constraint introduces a strong linguistic bias,

e.g., *a purple dog* (see **A.3.1**). Similarly, for example, replacing an object in a (subject, predicate, object) triple by another *similar* or a *probable* one is rather challenging. Depending on the heuristics employed to determine what might be a *probable* replacement, the resulting negative captions contain more or less linguistic bias (LM acc. of approx. 58% for *strict* constraints and approx. 66% when these constraints are *relaxed*) Moreover, we discovered that the relations in scene graphs are rather sparse which, if not handled correctly, results in noisy negative captions, i.e., the negative caption does not contradict the image. We provide further detailed analyses along with examples in Appendix **A.3.1**.

### 7.2 Sanity Check with Modality Ablation

We evaluate the **HNC** models under the *blind* setting<sup>11</sup> (see **A.1.5** for details on the implementation). Our findings<sup>12</sup> suggest that the effect of world priors, especially for object quantities, is difficult to overcome in negative caption generation.<sup>13</sup> For example, a typical quantity of a sofa in a living room is *one*. A negative caption with a different count of sofa violates the worldviews of **VL** models. **VLMs**, being trained on typical real-world scenes, usually do not capture other counts of sofas, and as a consequence, corresponding negative captions are easier to be detected as a mismatch, even though the model is not exposed to the visual input during inference. This poses a major challenge to **VL** pre-training in terms of learning modality mismatches.

## 8 Conclusion

In this work, we introduced Hard Negative Captions (**HNC**), a dataset for further pre-training Vision and Language Models to improve their modality integration capabilities on a fine-grained level and demonstrated improvements across models and tasks. We proposed a novel automatic dataset construction procedure for constructing hard negative captions to be used for Image-Text-Matching (**ITM**) training as well as a challenging test set annotated by humans. We provided detailed analyses of the challenges in automatic creation of hard negative captions and proposed methods to mitigate them.

<sup>11</sup>The image features are 0-masked during inference.

<sup>12</sup>Further analyses are provided in **A.3.2**.

<sup>13</sup>Blind VL models achieve +3pp. on average in **object\_count** (Tab. 11 and in col. *Clean Strict* in Tab. 10).



Lastly, we demonstrated the benefits of HNC by obtaining significant model performance gains on various tasks, including the diagnostic dataset VALSE, our HNC test set as well as a commonsense probing task (CPT), and down-stream performance gains after supervised fine-tuning on GQA, both of which require real-world reasoning.

## 9 Limitations

Automatic caption generation has its limitations. First, since our generation pipeline is seeded with the scene graphs (Krishna et al., 2017), issues identified in the literature like a skewed distribution of predicates (He et al., 2020), limited vocabulary size (He et al., 2022), low-level annotations, and reference ambiguity (Woo et al., 2021) might persist in our generated captions. Although we showed that certain biases can be mitigated (or minimized), our quantitative and qualitative analyses suggest that automatically generated captions based on scene graphs are subject to linguistic and distributional biases which are difficult to combat. Therefore, we believe that our hard negative caption generation could benefit from existing scene graph debiasing methods (Chiou et al., 2021). Also, our method of eliminating noisy captions caused by sparse scene graph annotations is based on rule-based heuristics. Although it helps us avoid creating false negative captions, it does not address the issue of annotation sparseness in scene graphs. For a potentially more robust method, the integration of an object detector (Russakovsky et al., 2015) can be studied in future work. Moreover, our rule-based heuristics are specific to our use case, and they might not work for other scenarios. Nevertheless, our framework allows for easy adaptation or extension to cover a wide range of domains and tasks. Last, our contribution is mainly on the creation of training and test data for ITM. We have not investigated the impacts of our data in combination with other training objectives or methods. We leave this (and the previous points) to future work.

## 10 Acknowledgement

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075 - 390740016. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech).

## References

- T W Anderson and D A Darling. 1954. [A test of goodness of fit](#). *J. Am. Stat. Assoc.*, 49(268):765–769.
- Yonatan Bitton, Gabriel Stanovsky, Roy Schwartz, and Michael Elhadad. 2021. [Automatic generation of contrast sets from scene graphs: Probing the compositional consistency of GQA](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 94–105, Online. Association for Computational Linguistics.
- Ben Bogin, Shivanshu Gupta, Matt Gardner, and Jonathan Berant. 2021. [COVR: A test-bed for visually grounded compositional generalization with real images](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9824–9846, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Emanuele Bugliarelli, Ryan Cotterell, Naoaki Okazaki, and Desmond Elliott. 2021. [Multimodal Pretraining Unmasked: A Meta-Analysis and a Unified Framework of Vision-and-Language BERTs](#). *Transactions of the Association for Computational Linguistics*, 9:978–994.
- Xiaojun Chang, Pengzhen Ren, Pengfei Xu, Zhihui Li, Xiaojiang Chen, and Alex Hauptmann. 2023. [A comprehensive survey of scene graphs: Generation and application](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1–26.
- Keqin Chen, Richong Zhang, Samuel Mensah, and Yongyi Mao. 2022. [Contrastive learning with expectation-maximization for weakly supervised phrase grounding](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8549–8559, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tianlang Chen, Jiajun Deng, and Jiebo Luo. 2020a. [Adaptive offline quintuplet loss for Image-Text matching](#). In *Computer Vision – ECCV 2020*, pages 549–565. Springer International Publishing.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020b. [Uniter: Universal image-text representation learning](#). In *ECCV*.
- Meng-Jiun Chiou, Henghui Ding, Hanshu Yan, Changhu Wang, Roger Zimmermann, and Jiashi Feng. 2021. [Recovering the unbiased scene graphs from the biased ones](#). In *MM ’21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, pages 1581–1590. ACM.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anuj Diwan, Layne Berry, Eunsol Choi, David Harwath, and Kyle Mahowald. 2022. [Why is winoground hard? investigating failures in visiolinguistic compositionality](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2236–2250, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Fartash Faghri, David J. Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2018. [VSE++: improving visual-semantic embeddings with hard negatives](#). In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, page 12. BMVA Press.
- Ronald A Fisher. 1949. *The design of experiments*. Oliver & Boyd.
- Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. 2020. [Large-scale adversarial training for vision-and-language representation learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 6616–6628. Curran Associates, Inc.
- Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. 2020. [Contrastive learning for weakly supervised phrase grounding](#). In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part III*, volume 12348 of *Lecture Notes in Computer Science*, pages 752–768. Springer.
- Tao He, Lianli Gao, Jingkuan Song, Jianfei Cai, and Yuan-Fang Li. 2020. [Learning from the scene and borrowing from the rich: Tackling the long tail in scene graph generation](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 587–593. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Tao He, Lianli Gao, Jingkuan Song, and Yuan-Fang Li. 2022. [Towards open-vocabulary scene graph generation with prompt-based finetuning](#). In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXVIII*, volume 13688 of *Lecture Notes in Computer Science*, pages 56–73. Springer.
- Lisa Anne Hendricks and Aida Nematzadeh. 2021. [Probing image-language transformers for verb understanding](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3635–3644, 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.
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. [Measuring compositional generalization: A comprehensive method on realistic data](#). In *International Conference on Learning Representations*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *Int. J. Comput. Vis.*, 123(1):32–73.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. [Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation](#). In *International Conference on Machine Learning*, pages 12888–12900. PMLR.
- 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. *ArXiv*, abs/1908.03557.
- Fangyu Liu and Rongtian Ye. 2019. [A strong and robust baseline for text-image matching](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 169–176, Florence, Italy. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- 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.
- Nicola Messina, Fabrizio Falchi, Andrea Esuli, and Giuseppe Amato. 2021. [Transformer reasoning network for image- text matching and retrieval](#). In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 5222–5229.
- 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.

- Jae Sung Park, Sheng Shen, Ali Farhadi, Trevor Darrell, Yejin Choi, and Anna Rohrbach. 2022. [Exposing the limits of video-text models through contrast sets](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3574–3586, Seattle, United States. Association for Computational Linguistics.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318. PMLR.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual captions: A cleaned, hypertexted, 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, Aurélie Herbelot, Moin Nabi, Enver Sangineto, and Raffaella Bernardi. 2017a. [Vision and language integration: Moving beyond objects](#). In *IWCS 2017 — 12th International Conference on Computational Semantics — Short papers*.
- Ravi Shekhar, Sandro Pezzelle, Aurélie Herbelot, Moin Nabi, Enver Sangineto, and Raffaella Bernardi. 2017b. [Vision and language integration: Moving beyond objects](#). In *IWCS 2017 — 12th International Conference on Computational Semantics — Short papers*.
- Ravi Shekhar, Sandro Pezzelle, Yauhen Klimovich, Aurélie Herbelot, Moin Nabi, Enver Sangineto, and Raffaella Bernardi. 2017c. 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.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [Vi-bert: Pre-training of generic visual-linguistic representations](#). In *International Conference on Learning Representations*.
- 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*.
- 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5228–5238. IEEE.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Sangmin Woo, Junhyug Noh, and Kangil Kim. 2021. [Tackling the challenges in scene graph generation with local-to-global interactions](#). *CoRR*, abs/2106.08543.
- Hsiu-Yu Yang and Carina Silberer. 2022. [Are visual-linguistic models commonsense knowledge bases?](#) In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5542–5559, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Jin Zhang, Xiaohai He, Linbo Qing, Luping Liu, and Xiaodong Luo. 2022. [Cross-modal multi-relationship aware reasoning for image-text matching](#). *Multim. Tools Appl.*, 81(9):12005–12027.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [ERNIE: Enhanced language representation with informative entities](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, Florence, Italy. Association for Computational Linguistics.
- Yuanen Zhou, Meng Wang, Daqing Liu, Zhenzhen Hu, and Hanwang Zhang. 2020. More grounded image captioning by distilling image-text matching model. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4776–4785.

## A Appendix

### A.1 Model Details

#### A.1.1 ITM Objective

Both single and dual-stream models aim to learn an alignment between the visual and textual modality to infer the correct entailment between them. Image–text matching is the objective of inferring a similarity score between these modalities. As such, in VL Transformers (Vaswani et al., 2017), it is implemented in the form of a binary classification head that learns to predict whether an image and a text entail one another.

#### A.1.2 BLIP

**Bootstrapping Language-Image Pre-training for unified vision-language understanding and generation (BLIP)** (Li et al., 2022) is a VL pre-training framework which is designed to perform both VL generation and understanding tasks. Li et al. (2022) propose three versions of BLIP: trained to align vision and language representations using an image-text contrastive loss, vision and language interactions using ITM, and a LM loss to generate captions. In the following, we refer to the BLIP version trained with a ITM loss as *BLIP-ITM*. In our experiments, we evaluated and fine-tuned BLIP-ITM, since it matches the design of our HNC dataset that aims for teaching the model’s a detailed understanding of the visual input using carefully sampled negative captions.

#### A.1.3 BLIP Hyperparameters

We use AdamW (Loshchilov and Hutter, 2017) with a learning rate of  $1e - 5$  and a weight-decay of 0.05 as used by (Li et al., 2022) to train BLIP. To fine-tune the model, we initialize a learning rate scheduler with a warm-up duration of four epochs and a starting learning of  $1e - 7$ . Afterward, the learning rate decays by a factor of  $\gamma = 0.85$ . We perform early stopping on the validation set, and train for a maximum of 20 epochs. The batch size during training equals 50, and we use eight NVIDIA A100 GPUs with 80GB VRAM.

#### A.1.4 VOLTA Hyperparameters

**Further pre-training on HNC** The following hyperparameters for the VOLTA models are used: ADAM optimizer (Kingma and Ba, 2014) with a learning rate and weight decay of  $4e - 5$ ,  $\beta = (0.9, 0.999)$ , and gradient clipping (Pascanu et al., 2013) with a norm of 1.0. For the tokenizer, we

used a maximum sequence length of 40. The maximum number of regions is set to 36 just like the VOLTA implementations. For the training, we used a batch size of 1024 and a maximum number of epochs of 20 with early stopping. We left all other hyperparameters untouched (e.g., model hyperparameters), and stick with the ones provided by VOLTA. We used 4 NVIDIA RTX A6000 GPUs and trained the models for a maximum of 48 hours. We use the controlled setup in VOLTA, which uses the same pre-training objectives and datasets across models to allow systematic comparison.

**Fine-tuning on GQA** For fine-tuning the VOLTA model checkpoints on the GQA dataset, we use a batch size of 1024 and a maximum number of epochs of 20 with early stopping. The maximum sequence length and the maximum number of regions were kept the same as in the pre-training. The rest of the hyperparameters are: ADAM optimizer (Kingma and Ba, 2014) with a learning rate and weight decay of  $4e - 5$ ,  $\beta = (0.9, 0.999)$ , and gradient clipping (Pascanu et al., 2013) with a norm of 5.0. We used 2 NVIDIA RTX A6000 GPUs and trained the models for maximum 8 hours.

For fine-tuning the BLIP model checkpoints on the GQA dataset, we use a batch size of 50 and train for a maximum of 20 epochs while performing early stopping. We again use AdamW (Loshchilov and Hutter, 2017) with a learning rate of  $5e - 5$  and weight decay of 0.05. The learning rate scheduler is initialized with a starting learning rate of  $1e - 8$ , a warmup duration of three epochs, and a  $\gamma = 0.85$  that scales the learning rate after each epoch.

**Language model training** We trained a BERT<sup>14</sup> (Devlin et al., 2019) model to predict whether a caption is positive or negative without seeing the image. The model is initialized with the pre-trained weights loaded from HuggingFace library<sup>15</sup>. We added a binary classification head and trained the model on HNC captions with the entailment labels of 0 and 1. Following hyperparameters were used: ADAM optimizer (Kingma and Ba, 2014) with a learning rate of  $16e - 5$ , maximum sequence length of 40 for the tokenizer, batch size of 8384, maximum number of epochs 40 with early stopping. We used a single NVIDIA RTX A6000 GPU and trained the models for maximum 120 hours.

<sup>14</sup>bert-base-uncased

<sup>15</sup><https://huggingface.co/>

### A.1.5 Blind Setting in VL Models

For consistency, we used the [VOLTA](#) implementations of the models and did not alter anything but the image features. We used 0-masking to create the blind setting. Specifically, we create a 0 tensor as the size of the image features and feed this into the model instead of the real image features. We do not change anything on the input of the textual modality.

	Dual-Stream						Single-Stream						Contrastive Dual-Encoder					
	ViLBER			LXMERT			UNITER			VISUALBERT			VL-BERT			BLIP		
	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC
	44.1	52.5	57.5	45.5	54.5	59.7	42.6	50.5	67.4	45.0	51.0	65.7	41.1	52.0	64.4	41.1	52.0	64.4
attribute	47.5	54.0	54.3	50.5	49.0	55.6	76.2	76.2	74.5	47.0	51.5	61.5	47.5	51.5	59.2	47.5	51.5	59.2
attribute_rel	46.2	54.7	55.0	47.7	53.8	54.7	64.7	64.7	64.7	45.8	53.8	55.2	47.3	47.3	52.2	47.0	53.5	63.8
relation	47.0	54.4	55.4	48.6	56.1	57.8	67.0	67.0	61.4	46.2	52.8	61.4	47.0	46.2	54.1	46.2	54.1	59.4
relation_attr	51.0	49.5	55.9	49.0	51.0	55.3	71.5	68.5	59.1	46.5	50.0	62.7	49.0	48.0	63.0	51.5	50.0	65.0
obj_count	50.0	48.5	57.2	49.0	51.0	55.6	54.5	54.5	51.0	49.5	51.0	56.6	48.5	49.0	50.5	60.9	62.5	49.5
obj_comp_count	49.0	50.5	52.1	48.5	49.5	49.8	65.5	65.5	61.3	50.5	50.0	57.6	49.0	49.5	54.2	72.0	51.5	73.5
verify_obj_rel	49.5	51.0	56.3	51.0	49.5	54.8	64.0	64.0	56.8	50.0	49.5	48.5	48.5	48.5	58.3	49.5	50.0	66.0
AND_logic_attr	48.5	51.5	52.2	50.0	48.5	54.4	70.0	70.0	54.4	49.5	50.5	54.4	50.0	49.0	55.4	49.0	50.5	74.5
AND_logic_rel	52.5	52.0	52.7	51.0	50.0	51.0	56.5	56.5	51.1	49.5	52.5	52.7	48.5	52.0	52.8	51.5	52.5	54.5
XOR_logic_attr	50.0	51.0	51.7	49.0	48.5	51.0	67.0	67.0	53.6	51.0	50.0	57.3	52.5	50.5	54.9	51.5	50.5	67.0
XOR_logic_rel	51.0	49.5	57.6	51.5	49.5	60.2	64.5	64.5	68.0	51.0	49.0	58.3	51.5	50.5	58.7	50.0	50.5	70.5
all	48.3	51.6	54.1	49.0	50.9	55.0	66.2	66.2	67.1	47.7	50.8	58.6	48.3	48.1	51.3	48.5	51.9	67.3

Table 5: Binary classification accuracy on HNC test set. The column BLIP-ITM refers to the checkpoint of BLIP that was fine-tuned on COCO and Flickr30k for image-text retrieval.

	Dual-Stream						Single-Stream						Contrastive Dual-Encoder					
	ViLBER			LXMERT			UNITER			VISUALBERT			VL-BERT			BLIP		
	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC	VOLTA	FOIL	HNC
existence	47.8	49.8	52.1	47.3	52.1	52.2	60.3	60.3	63.5	46.9	49.3	58.9	47.1	49.7	58.6	47.1	49.7	58.6
plurals	50.0	50.4	51.4	49.9	50.1	50.2	50.3	50.3	51.3	49.5	50.3	51.8	49.6	50.4	51.6	49.6	50.4	51.6
counting_small_quant	49.4	49.3	51.1	49.5	49.1	50.4	55.9	55.9	53.2	49.6	50.0	53.2	49.5	49.4	52.5	49.5	49.4	52.5
counting_adversarial	49.5	52.5	54.6	49.6	50.7	54.0	50.2	50.2	54.2	48.9	50.7	50.4	49.5	50.0	53.3	49.5	50.0	53.3
counting_hard	49.8	49.6	49.9	52.4	49.9	50.5	52.6	52.6	51.0	49.6	49.7	50.3	49.7	49.8	51.5	49.7	49.8	51.5
relations	49.8	49.8	50.9	50.0	49.9	50.5	51.6	51.6	50.3	50.7	49.8	50.4	49.9	50.2	50.7	49.9	50.2	50.7
actant_swap	48.1	54.6	55.8	49.0	53.2	55.4	57.4	57.4	54.4	47.9	51.0	58.3	48.3	50.4	53.8	48.3	50.4	53.8
action_replacement	47.0	53.0	51.6	52.9	48.1	53.2	55.1	55.1	52.0	47.8	50.3	51.0	47.4	50.5	49.7	47.4	50.5	49.7
coreference_standard	49.9	50.1	50.0	50.0	50.1	49.8	48.9	48.9	49.8	50.1	50.0	49.9	50.0	49.9	49.9	50.0	49.9	49.9
coreference_hard	50.0	50.0	50.0	50.0	50.0	49.5	48.6	48.6	50.1	50.4	50.0	49.8	50.0	50.0	49.4	50.0	50.0	49.4
foil_fit	46.0	77.0	50.4	45.8	77.5	50.1	52.4	52.4	51.2	43.7	79.0	51.5	44.2	78.1	50.0	44.2	78.1	50.0
all	48.8	50.9	51.6	48.9	50.8	51.4	53.0	53.0	52.8	48.4	50.2	52.3	48.5	50.0	51.9	48.5	50.0	51.9

Table 6: Binary classification accuracy on VALSE (Parcalabescu et al., 2022) under zero-shot evaluation. For the models trained on FOIL dataset, we do not calculate the accuracies obtained from the foil it splits (marked red) into the averaged values. The column BLIP-ITM refers to the checkpoint of BLIP that was fine-tuned on COCO and Flickr30k for image-text retrieval.

	Dual-Stream										Single-Stream														
	ViLBERT					LXMERT					UNITER					VISUALBERT					VL-BERT				
	VOLTA	FOIL	HNC_subset	HNC_full	HNC_subset	VOLTA	FOIL	HNC_subset	HNC_full	HNC_subset	VOLTA	FOIL	HNC_subset	HNC_full	HNC_subset	VOLTA	FOIL	HNC_subset	HNC_full	HNC_subset	VOLTA	FOIL	HNC_subset	HNC_full	
part-whole	55.02	54.59	52.65	55.97	52.93	53.73	50.13	52.93	56.48	63	63.95	64.1	<b>69.01</b>	68.84	66.47	63.35	64.46	66.47	68.84	58.37	61.37	59.43	63.43		
distinctness	55.92	56.76	54.56	54.83	61.98	59.42	55.68	61.98	60.51	65.94	67.27	66.42	70.65	68.84	70.89	65.7	69.69	70.89	68.84	66.06	<b>71.38</b>	67.08	68.6		
similarity	42.24	40.06	43.11	41.3	44.38	43.01	43.17	44.38	46.43	46.43	49.53	50.99	<b>55.75</b>	52.02	53.26	48.45	50.16	53.26	52.02	47.2	49.22	48.88	49.69		
temporal	47.17	45.79	41.5	39.28	47.56	49.23	49.81	47.56	50.24	47.86	46.59	46.53	46.27	52.3	50.19	50.98	51.51	50.19	52.3	51.83	51.67	50.5	<b>53.31</b>		
taxonomic	49.89	52.38	52.53	52.08	54.78	51.55	52.46	54.78	54.8	54.04	56.69	57.29	58.5	<b>63.27</b>	61.34	56.84	61.75	61.34	<b>63.27</b>	55.4	62.21	57.93	62.74		
quality	57.39	56.63	55.64	56.85	67.2	62.45	64.18	67.2	<b>70.11</b>	61.58	54.57	62.8	62.17	66.9	68.17	61.2	65.6	68.17	66.9	62.23	63.48	65.52	68.04		
spatial	52.91	51.97	45.41	50.47	55.04	55.6	52.72	55.04	56.79	57.41	57.47	53.32	57.97	57.54	57.01	58.1	<b>58.54</b>	57.01	57.54	56.47	56.66	53.85	57.47		
utility	60.48	57.8	57.57	57.32	63.91	62.87	57.51	63.91	60.62	65.36	64.93	65.95	68.71	<b>70.33</b>	67.79	67.22	68.66	67.79	<b>70.33</b>	65.12	65.74	65.15	65.74		
desire	56.6	54.41	53.23	47.05	54.6	54.8	51.88	54.6	52.05	58.51	57.72	58.75	<b>59.63</b>	57.22	58.34	57.27	57.78	58.34	57.22	50.42	52.22	50.48	51.43		
creation	52.0	54.0	58.6	53.0	53.4	56.0	54.0	53.4	54.0	62	64	63.2	<b>77</b>	70	68.4	65	66	68.4	70	63	66	64.4	67		
<b>all</b>	53.95	52.99	51.15	50.87	55.94	55.43	52.22	55.94	55.49	58.53	59.29	59.27	<b>62.32</b>	62.16	61.5	59.24	61.2	61.5	62.16	57.42	59.33	57.88	60.28		

Table 7: Classification accuracy on CPT (Yang and Silberer, 2022) task under zero-shot evaluation.

## A.2 Caption Generation Settings

As mentioned in Section 3, we implemented several heuristics to avoid ambiguity and potential noise in our caption generation. We now detail what these heuristics are and how they were implemented.

**Ambiguity** In many caption types, we only address localized cross-modal mismatches by leveraging subgraphs and do not take the global context of a scene into account. This results in ambiguity in entity grounding, especially when multiple instances of the same object class are present in the image. Additionally, scene graphs contain spatial relation annotations between entities and background objects such as *sky* or *field* that typically cover a large area in the scene. This causes ambiguity in captions as the exact spatial relation between them is hard to determine even for humans. Following heuristics are applied to reduce such ambiguities in captions (automatically created as well as human-annotated):

- A caption should not refer to multiple instances of the same entity class to avoid ambiguity in terms of entity grounding.
- A caption should not refer to a spatial relation between two body parts since such a caption is unnatural as well as error-prone due to multiple instances of body parts in scenes.
- A caption should not refer to a spatial relation between an entity and an object typically covering a large area in scenes, i.e., typical background objects.

**Clean vs. noisy** In our **clean** setting, we filter out all the values that our noisy spatial relation detection algorithm tags as *noisy*. The way this works is:

1. The algorithm gets a triple (subject, relation, object) and a marker as to which value in the tuple should be replaced with a foil.
2. All the candidate replacement values are collected in a list. This also follows a set of heuristics which we discuss later.
3. We then compare the bounding boxes of the subject and the object, and decide whether the spatial relation is correct between these visual objects.

4. If we determine that the given relation is incorrect, we remove this item from the list of candidates.

In the **noisy** setting, we do not filter out these potentially noisy candidates.

**Strict vs. relaxed sampling** There are several ways of sampling foils for a given tuple. The simplest way would be to sample from all the words in the vocabulary in the same POS tag category, i.e., sample from the set of nouns in the vocabulary for a given noun, e.g., sample a shoe for cat. However, as it quickly becomes obvious, this approach has several potential issues. One issue, for example, is that we might end up with nonsensical captions containing an object an unsuitable attribute, e.g., *the ground is scrambled* (see 3b.). Also, since the scene graphs contain non-spatial relations, we might accidentally create captions that violate object affordances, e.g., *a table is eating a boy*. Thus, it is important to follow an **informed sampling strategy**. To achieve this, we created look-up tables allowing us to sample a foil that does not result in a nonsensical caption. For (attribute, object), (subject, predicate), and (predicate, object) pairs we aggregate the information in the ground-truth scene graphs and save them as look-up tables. Additionally, we annotated attribute clusters that group similar attributes into buckets for us to sample values from. Using these look-up tables, we provide two negative value sampling strategies for generating hard negative captions: **(a) relaxed** and **(b) strict**.

Our **relaxed** setting allows sampling from a *probable* set of values such that we allow sampling a negative attribute from the attribute class of the positive one; and for the (subject, predicate, object) triples, we sample from the union of the (subject, predicate) and (predicate, object) pairs. This type of sampling makes the assumption of: given that an object co-occurs with a similar attribute or that a predicate with a subject and an object on different accounts, although an exact tuple might not co-occur in the dataset, this does not mean that such a co-occurrence is unlikely. This increases the variability of the captions but can also result in erroneous cases because neither the attribute clusters are robust (see caption 2 in Figure 3a.) nor the assumption always holds: if (subject, predicate) and (predicate, object), then (subject, predicate, object), e.g., (dog, drinks) and (drinks, beer) does not guarantee (dog, drinks, beer).



	VILBERT		LXMERT *		UNITER *		VISUALBERT *		VL-BERT		BLIP	
	VOLTA	HNC	VOLTA	HNC	VOLTA	HNC	VOLTA	HNC	VOLTA	HNC	ITM	HNC
Accuracy	55.77	55.97	53.48	55.45	55.28	56.70	53.51	56.85	55.62	55.96	57.38	<b>57.73</b>

Table 8: Results on the GQA (Hudson and Manning, 2019) testdev split. Results are statistically significant\*.

In **strict** setting, we only allow sampling from the look-up tables directly meaning that the exact co-occurrence exists in the ground-truth scene graphs. This results in a highly strict constraint as we essentially limit the likely negative candidates to the ones that co-occur in the dataset. Nonetheless, by doing so, we minimize the number of nonsensical captions.

In all our experiments, we used the captions generated using the **clean and strict** setting.

**Balancing the comparative quantifiers in captions** In order to prevent models from attending to linguistic signals for a prediction shortcut, comparative quantifiers are equally used in the positive and the negative caption types.

**Balancing the existence and nonexistence in existence-based captions** Same as above, to avoid shortcuts, *no* and *at least one*, i.e., (non)existence of entities, in positive and negative captions are balanced.

### A.3 Qualitative Analysis

#### A.3.1 Dataset Generation Process

**Refinements in sampling methods** In our first iteration of the sampling implementation, we started with a single constraint, i.e., the negative value (object, attribute, relation) must be sampled from within the scene. This, however, results in a strong linguistic bias as there is no mechanism that ensures the sensibility of the generated caption. This resulted in captions like *the table is sleeping*, or *the man is eating a couch* which then gave us **LM** accuracies of approx. 70% on the validation set. This is highly undesirable as the entailment between an image and its caption can be predicted simply by assessing the caption’s sensibility.

In our next iteration of sampling from look-up tables in the **relaxed** setting, we were able to reduce the **LM** accuracies down to approx. 66%. This setting helps us avoid creating captions such as *the man is eating a couch* as the object *eating* does not occur together with *couch* in the ground-truth scene graphs. Note that, at this time, we are using the look-up tables, but we are still sampling

uniformly. This uniform sampling turned out to be highly problematic as the word distributions between the positive and the negative captions were too dissimilar resulting in shortcut predictions. The reason is that co-occurrences of visual concepts in the ground-truth GQA scene graphs are highly imbalanced. For example, *to the left of* and *to the right of* are the most common predicates in the dataset. When we uniformly sample from the above-mentioned look-up tables, we create a distributional bias between the positive and the negative caption sets (see subplots (a) & (b) of Figure 12 for the relation distribution of the captions from an early iteration.). Thus, we extracted word co-occurrence statistics from the ground-truth scene graphs and sampled from the look-up tables following these distributions (see subplots (c) & (d) of Figure 12 for the relation distribution in our final captions.), which helped us reduce the **LM** accuracies down to approx. 58%.

To reduce the linguistic bias even further, we implemented **strict** sampling which we detailed in Section A.2. With this sampling strategy, we are able to reduce the **LM** accuracies down to approx. 57% (see Tab.9).

Table 9 shows the **LM** accuracies<sup>16</sup> on the final versions of the **HNC** validation sets. According to these numbers, some of the caption types contain more bias than the others, e.g., **attribute**, **attribute\_relation**, **relation**, **relation\_attribute**, **object\_count**, **object\_compare\_count**, **XOR\_logic\_relation** all have accuracies  $\gtrsim$  60%. For example, the model achieves approx. 65% accuracy on the validation split in **object\_count** type (approx. 61% in **object\_compare\_count**). We attribute this to a combination of dataset and world-priors biases which is common in datasets of real-world images.

Note that **LM** accuracies are a simple proxy we use to measure the linguistic bias in the textual modality without the presence of the visual modality. Thus, we believe that none of the methods is ideal, and the choice of the sampling strategy might

<sup>16</sup>The higher the accuracy, the more biased is the dataset.

	Clean Strict	Clean Relaxed	Noisy Strict	Noisy Relaxed
attribute	62.0	<b>65.2</b>	62.3	<b>65.2</b>
attribute_relation	60.4	<b>63.3</b>	60.3	<b>63.3</b>
relation	58.0	59.6	57.7	<b>60.1</b>
relation_attribute	63.2	64.8	62.9	<b>65.2</b>
object_count	65.5	65.6	<b>65.7</b>	65.4
object_compare_count	61.3	<b>61.4</b>	<b>61.4</b>	60.8
verify_object_attribute	<b>55.5</b>	55.2	55.0	55.3
verify_object_relation	54.1	54.2	54.0	54.0
AND_logic_attribute	<b>55.4</b>	55.2	55.1	55.2
AND_logic_relation	55.0	56.2	54.6	56.3
XOR_logic_attribute	<b>54.3</b>	<b>54.3</b>	53.8	53.0
XOR_logic_relation	60.6	62.7	58.6	<b>63.3</b>
<b>all</b>	57.6	58.6	57.4	<b>58.7</b>

Table 9: Language Model results on **HNC** validation set. The models are trained and evaluated on data obtained from the same setting.

	Clean Strict	Clean Relaxed	Noisy Strict	Noisy Relaxed
attribute	55.9	<b>60.4</b>	55.4	58.4
attribute_relation	51.0	<b>54.5</b>	52.5	54.0
relation	<b>56.0</b>	55.8	54.3	54.5
relation_attribute	53.5	54.7	<b>55.0</b>	53.5
object_count	<b>55.0</b>	52.5	<b>55.0</b>	54.0
object_compare_count	53.5	55.5	<b>56.5</b>	54.5
verify_object_attribute	<b>51.5</b>	48.5	48.5	<b>51.5</b>
verify_object_relation	<b>54.0</b>	52.5	53.0	53.5
AND_logic_attribute	52.0	51.5	<b>54.0</b>	50.5
AND_logic_relation	<b>50.0</b>	48.5	<b>50.0</b>	48.5
XOR_logic_attribute	48.0	48.5	<b>50.0</b>	48.0
XOR_logic_relation	49.5	52.5	49.5	<b>56.0</b>
<b>all</b>	53.1	<b>53.5</b>	53.3	53.3

Table 10: Language Model results on **HNC** test set. The models are trained on different settings and evaluated on the human-annotated test set.

	Dual-Stream				Single-Stream					
	VILBERT		LXMERT		UNITER		VISUALBERT		VL-BERT	
	VOLTA	HNC	VOLTA	HNC	VOLTA	HNC	VOLTA	HNC	VOLTA	HNC
attribute	50.5	57.9	52.0	56.9	54.0	58.4	52.0	54.5	50.0	<b>61.9</b>
attribute_rel	49.5	<b>56.0</b>	52.5	54.5	52.0	53.5	49.5	54.0	50.0	52.5
relation	49.0	<b>54.8</b>	49.3	54.5	49.2	53.0	50.0	54.0	50.0	53.3
relation_attr	50.9	55.5	50.9	56.2	49.6	57.1	49.9	54.7	50.3	<b>57.8</b>
obj_count	49.5	<b>65.0</b>	45.0	46.5	51.0	61.0	51.5	64.0	50.5	58.0
obj_comp_count	50.5	51.5	49.5	53.0	48.0	53.0	50.5	52.5	49.0	<b>53.5</b>
verify_obj_attr	<b>52.5</b>	42.0	51.5	44.5	46.0	46.0	47.5	46.5	50.0	50.0
verify_obj_rel	50.5	<b>54.5</b>	51.0	53.5	50.0	51.0	50.0	53.0	50.0	52.0
AND_logic_attr	51.5	<b>59.0</b>	49.0	50.0	49.5	51.0	51.0	52.0	51.0	57.0
AND_logic_rel	50.0	<b>54.0</b>	48.5	53.5	49.0	52.5	50.0	49.0	50.0	49.5
XOR_logic_attr	49.5	<b>53.0</b>	50.5	49.0	50.5	51.0	49.0	49.5	50.0	50.5
XOR_logic_rel	50.0	57.5	49.5	<b>60.0</b>	49.5	54.0	50.0	57.0	50.0	59.0
<b>all</b>	50.2	<b>55.1</b>	50.0	53.3	50.0	53.8	50.1	53.6	48.1	50.1

Table 11: Binary classification accuracy on HNC test set under blind evaluation.

depend on the use case.

**Noisy spatial relations** Our qualitative iterative analysis revealed that, due to the incomplete nature of the relations in GQA scene graphs, our *noisy* setting results in many noisy hard negative captions in that the values we sample as foils do not contradict the image (see caption 1 in Figure 3a) However, this is not detectable simply by looking at the LM accuracies as the captions are not non-sensical. Thus, between the *clean* and the *noisy* settings, there does not seem to be a great deal of difference for the LM which is expected as the sensibility of the captions are not directly affected by the correctness of objects’ spatial relations in the visual scene, e.g., a bus driver can be inside or the next to a bus.

### A.3.2 Analysis of the Human-Annotated Test Set

We evaluated the LM trained on HNC captions to quantify the pure linguistic bias that might be present in our human-annotated test set. Ideally, LM should perform at the random baseline level, i.e., 50% accuracy. In our **clean and strict** setting, the model achieves an average accuracy of 53.1% which suggests the presence of *some* bias. This might be due to the domain size in GQA images. Thus, no matter if created automatically or annotated by humans, such statistical biases caused by the domain size are hard to mitigate.

Table 11 contrasts the accuracies of models trained on HNC image–text pairs<sup>17</sup> with the VOLTA models evaluated on the text-only modality

<sup>17</sup>The models are trained on the *clean-strict* version.

of the human-annotated test set (see A.1.5 for the implementation details). Previously, we discussed biases in our dataset. With these results, our aim is to draw attention to the biases in the pretrained VL models. As also briefly mentioned in Section 7.2, we might violate world-priors in VL models by creating negative captions that are possible but might not be probable according to their worldview, e.g., the leaves might be more likely to be green or yellow than red or brown, although red or brown leaves are not impossible. Moreover, due to the size of the GQA images, it is unlikely that the dataset is an accurate sample of the world, i.e., although we might have images showing *a man eating pizza* and *a woman eating pasta*, this does not mean that the men do not eat pasta or the other way around.

### A.3.3 Downstream Tasks

**VALSE** In Figure 4, we display some examples where all our models predicted the correct entailment between the image and the caption that were predicted incorrectly by all the models initialized from the VOLTA checkpoints. As also indicated by the quantitative results, we observed significant improvement in all the models regarding certain types of foils, which we discuss briefly in the following.

Our models predict correct entailment in many counting-based captions that were predicted incorrectly by the VOLTA models. Our qualitative analysis revealed that this is especially the case when the foiled count is small and close to the original count. Furthermore, in many of our hard negative captions, we swap grammatical subjects (agent, actant) or objects (patient, theme, experiencer) of the captions with a foil. This seems to help models ground the



1. **Positive Caption:** The towel is on top of the toilet.  
**Negative Caption:** The cat is on top of the toilet.  
**Type:** relation

---

2. **Positive Caption:** There is either a white towel or a sleeping cat.  
**Negative Caption:** There is either a white towel or a beautiful cat.  
**Type:** XOR\_logic\_attribute

(a)



- Positive Caption:** The egg is scrambled.  
**Negative Caption:** The ground is scrambled.  
**Type:** attribute

(b)

Figure 3: **(a)** The resulting negative captions do not contradict the image; thus, they are false negatives. **Negative caption 1** contains a noisy spatial relation, **negative caption 2** contains an attribute similar to the attribute in the positive caption but not contradictory to the image. **(b)** The sampled noun ground with the attribute “*scrambled*” creates a nonsensical caption.



**Caption:** There are exactly 2 lights above the sink.  
**Label:** 0  
**Type:** counting



**Caption:** a couple of kids laying on top of a bed.  
**Label:** 1  
**Type:** Foil-it (relation-object)



**Caption:** A cow stands on a sidewalk in a building.  
**Label:** 0  
**Type:** relations

Figure 4: Example cases where all the VOLTA models failed while our models predicted the correct entailment.

correct visual object in the image and not just predict entailment by assessing the plausibility of the caption. We also observe improvements in spatial relation grounding which is expected as our dataset contains many captions that specifically foil this information. In some examples, where **VALSE** foils the action in the caption, our models perform better as well. This might mean that the correct grounding of the subjects and the objects in captions might have a positive effect on the grounding of the action in the visual scene. However, since the **GQA** scene graphs do not readily provide many actions, we do not see a big improvement in this type.

We also observed some failure cases where the previously correct predictions were predicted incorrectly by all our models (see Figure 5). This mainly occurred in foil types that we do not cover in our hard negative caption generation, e.g., coreference (see the left example in Figure 5), plurals and non-spatial relations. However, lack of coverage is not the only place where we observe such behavior. For example, some counting-hard captions that were predicted correctly by **VOLTA** models ended up being predicted incorrectly by all our models (see the middle example in Figure 5). This might be due to the imbalanced object counts in the captions. We chose to follow the ground-truth scene graph distributions which inherently contain some bias on a compositional level as discussed in Section A.3.1. The implication of this is that our positive (also hard negative) captions might never have certain combinations of concepts compositionally co-occur in the same caption, i.e., while we might have captions that contain one, two, three, or four elephants; we might never have a caption with five elephants in the positive captions if such a scene graph does not exist in the **GQA** dataset.

Additionally, we found that some of the foiled instances incorrectly predicted by **HNC** models are ambiguous; e.g., in the right example of Figure 5, the foil (bicycle) for the correct object (car) is also near the table.

**CPT** Each instance of  $CWWV_{Img}$  consists of three natural language statements and a corresponding set of retrieved images,  $T_i = (Q||A_i||V_i)$ ,  $i = 1, \dots, 3$ , where  $Q$  is the prompt,  $A_i$  a candidate answer,  $V_i$  is a set of retrieved images for the answer tokens. A model has to determine in a zero-shot manner which of the three statements is true. Specifically, it requires a model to perform **MLM**

on the same masked token of the prompt  $Q$  in each  $T$ . The statement that receives the lowest **MLM** loss is considered the model’s prediction.

In Figure 6, we showcase several examples where **HNC** single-stream models successfully handle noisy visual inputs during the inference stage (**VOLTA** single-stream models fail), especially on **similarity**, **quality**, and **taxonomic** dimension. We investigate how the visual noisiness in the aforementioned dimensions varies from each other by looking into respective examples. For **similarity**, although the extracted image metaphorically captures the answer token, *buddy*, to display a sense of togetherness, there is no human being, but only two crocodiles, in the picture, which creates an entity-level misalignment w.r.t the question token, *brother*, in the prompt. A similar issue is observed for the **quality** dimension, in which the extracted image for *flying* is conceptually correct, but no *bird*, but only a plane, can be identified in the image. As for **taxonomic** dimension, we found that general concept words like *rate* could potentially create a modality misalignment issue w.r.t. the question token in the prompt, e.g., *speed* because *rate* could also be a unit to measure attractiveness in this case. These cases exemplify the difficulty of **CPT** task that might lead **VL** models to pick a wrong prediction in the presence of conceptually correct, but not-strictly-aligned, visual inputs. However, since **HNC** single-stream models are pre-trained to be aware of fine-grained misalignment, they bypass the limited information provided by the visual modality and robustly resort to the textual modality for performing inference. The effectiveness does not generalize to other dimensions such as **temporal** and **spatial** as exemplified in Figure 7 and Figure 8 respectively. It is notable that **HNC** dual-stream models suffer stronger from a performance decrease than the single-stream counterparts. By inspecting the failure case of **temporal** made by **HNC** dual-stream, it is clear that the wrong prediction could easily occur due to the natural misalignment of the temporal orders between the question token, *buying food*, in the prompt and the answer token, *run out of money*. Therefore, the resulting retrieved image is naturally not corresponding. In the example here, we observe **HNC** dual-streams select the choice, *get extremely relaxed*. The reason behind this could be that there are glasses, hyponyms of *food*, existing in the *relaxed* picture. With respect to the failure case of *spatial* dimen-



**Caption:** 5 people skiing in a snowy area surrounded by trees. is this a resort do you think? yes.  
**Label:** 0  
**Type:** coreference-hard



**Caption:** there are exactly 4 lights.  
**Label:** 0  
**Type:** counting



**Caption:** table near bicycle with a bicycle along side and a plate with two hot dogs and a coke.  
**Label:** 0  
**Type:** Foil-it (relation-object)

Figure 5: Example cases where all our models failed while the VOLTA models predicted the correct entailment.




**buddy**

**Dim:** Similarity  
**brother** is a synonym of:  
 A. first step  
 B. freezing injunction  
 C. **[correct & predicted] buddy**




**flying**

**Dim:** Quality  
 A bird can be:  
 A. **[correct & predicted] flying fast**  
 B. one of many firearms  
 C. coral



**Dim:** Taxonomic  
**speed** is a type of  
 A. computer chassis  
 B. hyperreal number  
 C. **[correct & predicted] rate**

Figure 6: Example cases where our HNC single-stream models succeed under noisy visual input scenarios, i.e., a modality mismatch between the textual token in the prompt and the image retrieved based on the correct textual choice, e.g., the word **bird** and the image **flying**.



**money**



**relaxed**

**Dim:** Temporal  
 Sometimes **buying food** causes you to:  
 A. **[correct] run out of money**  
 B. clothes stained  
 C. **[predicted] get extremely relaxed**

Figure 7: A failure case of HNC dual-stream models on the temporal dimension.




Figure 8: A failure case of HNC dual-stream models on the spatial dimension.

sion, again, we see that HNC dual streams are subject to slight modality non-correspondence. The image extracted for the correct answer token, *building* capture the external view of a *building*; whereas the image for the wrongly picked answer token, *carpeting*, is photographed inside a house.

#### A.4 Statistical Test

To determine whether one model significantly outperforms the other one, we resort to paired student’s t-test (Fisher, 1949) with the threshold of  $p < 0.05$  to be significantly outperforming. Since the t-test assumes a normal distribution, we also test the normality of model prediction with the method of Anderson-Darling (Anderson and Darling, 1954).

#### A.5 Dataset Statistics

Figure 9 contains the distributions for the human annotated test set. The total number of each cap-

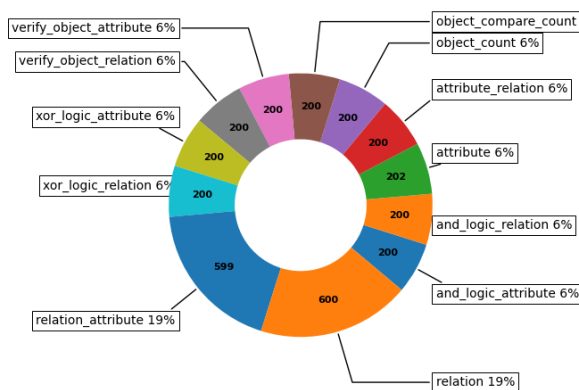


Figure 9: Test set caption type distribution.


tion type as well as the relative percentage values

are displayed. The test set contains exactly 100 annotated images.

Figure 10 contains the caption type distributions for the training set data w.r.t. the different dataset variations, and Figure 11 contains the caption type distributions for the validation set.

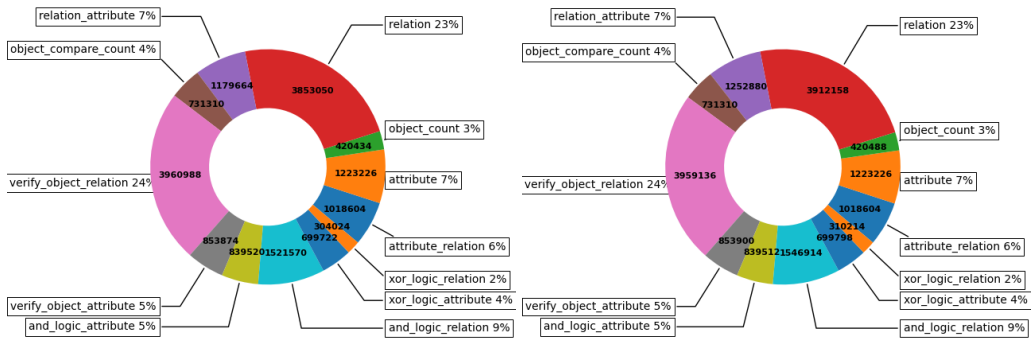
Figure 12 displays the relation distributions for the positive and negative captions. Fig. 12a and 12b contain the distributions from earlier iterations. It is striking to see that the relation distributions in the positive and negative captions are very dissimilar. Our final state of the caption generation procedure produces similar relation distributions, as can be found in Fig. 12c and 12d. Most prominent are the relations *to the left of* and *to the right of*. Following different data distributions enables models to easily distinguish between negative and positive captions, which is why we mitigated the gap between iterations.

Table 12 contains the exact numbers for each dataset split and variation.



(a) Clean strict.


(b) Noisy strict.



(c) Clean relaxed.

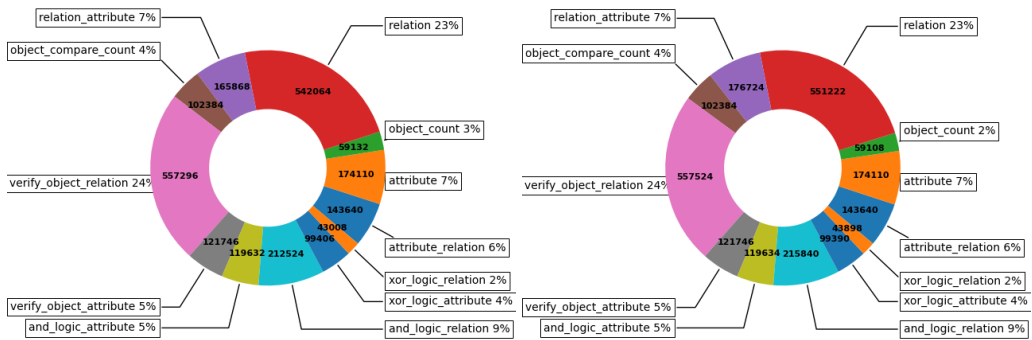
(d) Noisy relaxed.

Figure 10: Training split variation distributions.



(a) Clean strict.

(b) Noisy strict.



(c) Clean relaxed.

(d) Noisy relaxed.

Figure 11: Validation split variation distributions.



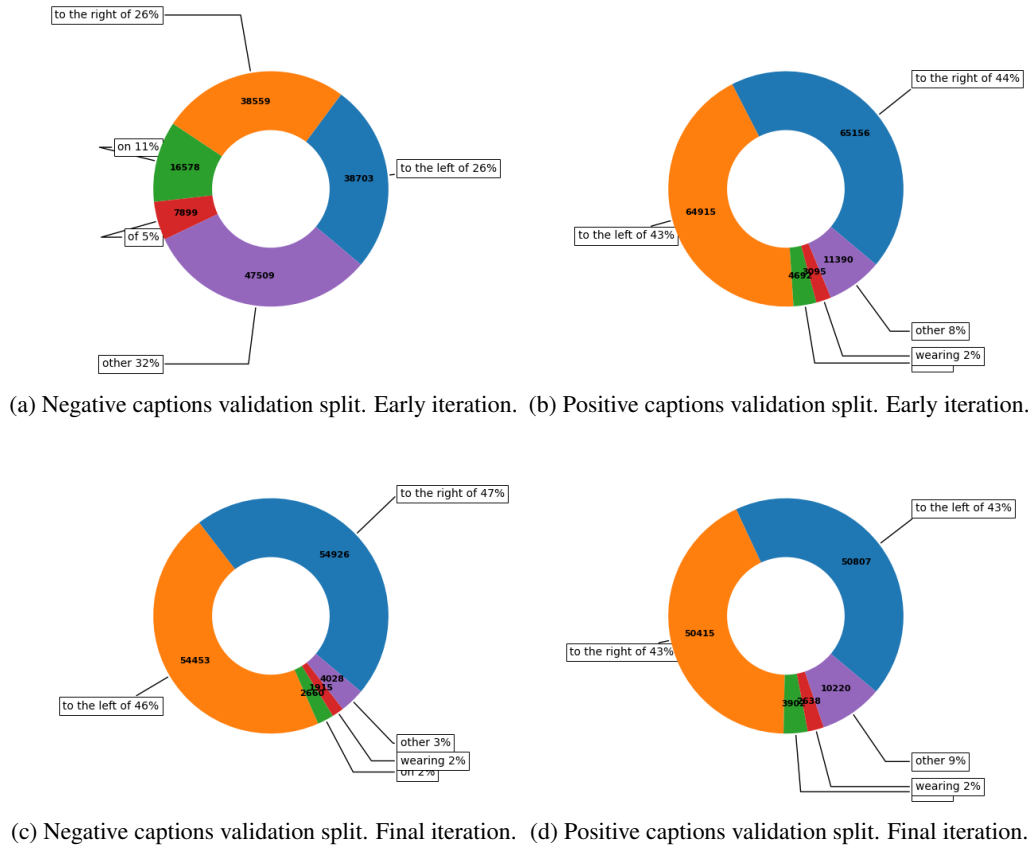


Figure 12: Relations distributions.

Split	Variation	Total Amount Cpts	Avg Cpt Len	Avg Cpt Amounts across Types
Valid	Clean Strict	2,314,832	10.28	238.81
	Clean Relaxed	2,340,810	10.26	241.49
	Noisy Strict	2,354,070	10.27	242.86
	Noisy Relaxed	2,365,220	10.25	244.01
Train	Clean Strict	16,416,392	10.29	242.10
	Clean Relaxed	16,605,986	10.27	244.90
	Noisy Strict	16,702,102	10.29	246.32
	Noisy Relaxed	16,768,140	10.27	247.29

Table 12: Statistics of our automatically generated data splits and variations.