

ARMADA: Attribute-Based Multimodal Data Augmentation

Xiaomeng Jin[†] Jeonghwan Kim[†] Yu Zhou[¶] Kuan-Hao Huang[§]
Te-Lin Wu[‡] Nanyun Peng[‡] Heng Ji[†]

[†]University of Illinois, Urbana Champaign ^{¶,‡}University of California, Los Angeles

[§]Texas A&M University

{xjin17, jk100, hengji}@illinois.edu, yu.zhou@ucla.edu
{telinwu, violetpeng}@cs.ucla.edu, khhuang@tamu.edu

Abstract

In Multimodal Language Models (MLMs), the cost of manually annotating high-quality image-text pair data for fine-tuning and alignment is extremely high. While existing multimodal data augmentation frameworks propose ways to augment image-text pairs, they either suffer from *semantic inconsistency* between texts and images, or generate unrealistic images, causing *knowledge gap* with real world examples. To address these issues, we propose **Attribute-based Multimodal Data Augmentation (ARMADA)**, a novel multimodal data augmentation method via knowledge-guided manipulation of visual attributes of the mentioned entities. Specifically, we extract entities and their visual attributes from the original text data, then search for alternative values for the visual attributes under the guidance of knowledge bases (KBs) and large language models (LLMs). We then utilize an image-editing model to edit the images with the extracted attributes. ARMADA is a novel multimodal data generation framework that: (i) extracts knowledge-grounded attributes from symbolic KBs for semantically consistent yet *distinctive* image-text pair generation, (ii) generates visually similar images of disparate categories using neighboring entities in the KB hierarchy, and (iii) uses the commonsense knowledge of LLMs to modulate auxiliary visual attributes such as backgrounds for more robust representation of original entities. Our empirical results over four downstream tasks demonstrate the efficacy of our framework to produce high-quality data and enhance the model performance. This also highlights the need to leverage external knowledge proxies for enhanced interpretability and real-world grounding.

1 Introduction

Multimodal Language Models (MLMs) exhibit remarkable abilities in comprehending and integrating various modalities, encompassing texts, images, and videos. Recently, many MLMs have

been proposed by researchers in both academic and industrial communities (Li et al., 2020; Radford et al., 2021; Li et al., 2022a,b; Liu et al., 2023b; Dai et al., 2023; Achiam et al., 2023), demonstrating significant achievements across various downstream tasks, such as image-text retrieval (Radford et al., 2021; Li et al., 2022a) and visual question answering (VQA) (Liu et al., 2023b,a; Dai et al., 2023). Training MLMs for downstream tasks, which usually involves fine-tuning and alignment stages, requires substantial amounts of annotated data. However, collecting and annotating such datasets demand considerable human effort and are notorious for their expense and time-consuming nature. A common strategy to overcome this problem is leveraging data augmentation techniques, which automatically synthesize new data instances from existing datasets, relieving the need to rely on manually annotated datasets to train these models.

Existing multimodal data augmentation methods, which require the perturbation of both the visual and textual modalities in tandem, can generally be classified into the following two groups: (i) latent space-based methods that perturb the latent representations of existing data instances (Liu et al., 2022) via adversarially trained augmentation networks, and (ii) surface form-based methods (Müller and Hutter, 2021; Hao et al., 2023) that simply perturb superficial representations such as orientations/pixel-level mixture of images. Latent space-based methods such as LeMDA (Liu et al., 2022) generate augmented multimodal latent features aligned with the training data distribution, but are inherently confined by their lack of interpretability and controllability. While surface form-based methods partly provide interpretable and controllable alternative, their simple augmentation schemes such as random solarization and pixel-level interpolation lead to *semantic inconsistency*. For instance, Figure 1 shows that random cropping or image interpolation cause semantic



Figure 1: Generated examples using two previous data augmentation methods and our approach. (a) is generated by TrivialAugment (Müller and Hutter, 2021), showing the altered images from randomly solarizing or cropping the dog and the fence out from the original image, demonstrating semantic inconsistency. (b) shows the output image from MixGen (Hao et al., 2023), demonstrating the unrealistic output from simple image interpolation and text concatenation. (c) shows the augmented data from our method ARMADA, which are semantically consistent.

gaps between paired images and texts, leading to images far from realistic. Moreover, such perturbations cannot deal with variable entity categories that appear in a similar background, or same entities with variable physical attributes, since they disregard attribute-level details. Our work aims to address these issues by leveraging a rich bank of attributes from a hierarchical knowledge base for interpretable and controllable multimodal data augmentation that guarantees semantic consistency and knowledge-grounding of generated entities.

In this paper, we introduce a novel *attribute-based*, multimodal data augmentation framework, ARMADA, that extracts the entities and visual attributes, then modifies the visual attributes of entities in images by building an entity-attribute multimodal knowledge base (KB). We perform entity-related knowledge extraction through entity linking using Spacy Entity Linker on Wikidata KB to: (i) generate augmented images and texts that faithfully reflect knowledge-grounded, entity-related attributes, and (ii) exploit the neighboring entities, e.g., a Boston Terrier and French Bulldog in Figure 1, for generating similar yet distinguished entity categories. Our work also leverages LLMs

as additional knowledge proxy as they can generate alternatives to any textual attributes without related entities in KB. We then modify images based on revised texts by employing an off-the-shelf image editing model, InstructPix2Pix (Brooks et al., 2023). Our framework produces semantically consistent, knowledge-grounded multimodal data instances. In-depth experiments across four different image-text downstream tasks against five different baselines demonstrate the significance of augmenting multimodal data instances guided by entity-related attribute knowledge. Our contributions can be summarized as follows:

- We propose a knowledge-guided multimodal data augmentation framework that is guided by entity-centric KBs to generate entities that are of the same type yet differing attributes, or of similar yet disparate categories.
- The proposed augmentation pipeline in this work demonstrates semantically consistent and knowledge-grounded multimodal data, addressing the limitations of previous multimodal data augmentation methods.
- Our empirical results demonstrate that our

proposed data augmentation strategy leads to substantial gains in various image-text downstream tasks such as image-text retrieval, VQA, image captioning, and especially in fine-grained image classification tasks that rely on attribute-centric information.

2 Related Work

External Knowledge Proxies. External symbolic knowledge bases (KBs) like Wikidata (Vrandečić and Krötzsch, 2014) and real-world knowledge proxies like large language models (LLMs) (Achiam et al., 2023; Touvron et al., 2023; Almazrouei et al., 2023) contain ample amount of real-world, entity-centric knowledge. While symbolic KBs have frequently been used in various domains of natural language processing for augmentation (LUO et al., 2023; Sun et al., 2023; Pan et al., 2024), the use of symbolic KBs in the multimodal domain is yet to be explored. LLMs, while they may suffer from hallucinatory outputs, contain rich world knowledge that enables them to generalize to attributes of various kinds. Our work reaps the benefits of the both worlds by exploiting the relational knowledge of KBs and generalization abilities of LLMs to perform knowledge-guided multimodal data augmentation.

Vision Language Models. Vision Language Models (VLMs) have achieved new state-of-the-art performances across various downstream tasks such as image-to-text retrieval and visual question answering (VQA) (Radford et al., 2021; Li et al., 2022a; Dai et al., 2023; Liu et al., 2023b,a). CLIP (Radford et al., 2021) is a widely used VLM for image-text retrieval and image classification. InstructBLIP (Dai et al., 2023) and LLaVA (Liu et al., 2023b) are instruction-tuned multimodal models that combine vision encoders and LLMs. The major drawback of these models is that they require an extensive amount of image-text pair datasets to either pre-train or fine-tune the models. Such shortcomings call for the need of a new, robust augmentation method, which our work aims to offer.

Data Augmentation. Existing work on data augmentation mainly focuses on augmenting a single modality, e.g., text (Thakur et al., 2021; Yoo et al., 2021; Chen et al., 2023; Jin and Ji, 2024) or image (Luo et al., 2023; Trabucco et al., 2023; Müller and Hutter, 2021). Most recently in the multimodal domain, several augmentation methods have been proposed to augment multiple modalities at the

same time. MixGen (Hao et al., 2023) generates new data instances by interpolating images and concatenating their accompanying texts. As discussed in Figure 1, one potential issue is the low quality of the generated data. LeMDA (Liu et al., 2022), another augmentation method that jointly augments multimodal data in the feature space, is limited in terms of interpretability and controllability since the generation occurs in latent space. BiAug (Wu et al., 2023) augments multimodal data in a similar manner as our approach by decoupling entities and their attributes. However, BiAug heavily relies on LLMs to generate the attributes, which are susceptible to hallucinatory outputs. Our proposed approach, in contrast, leverages entity-related attributes from knowledge base and delegates entity independent perturbations to LLMs.

3 Our Approach

Suppose we have a set of image-text pairs $\mathcal{D} = \{(I_1, T_1), \dots\}$ as the training dataset. T_i is a task-dependent text that is paired with its corresponding image, I_i . For example, T_i can be the label of image I_i in image classification task, a caption that describes I_i in image-text retrieval task, or a question-answer pair if the image I_i appears in a VQA task. Given that the training dataset with gold-standard annotations \mathcal{D} is usually too small to train the vision language model sufficiently well, we aim to augment the original training dataset and generate additional image-text pairs $\mathcal{D}' = \{(I'_1, T'_1), \dots\}$. The augmented dataset \mathcal{D}' can be used in conjunction with the original dataset \mathcal{D} to train the VLMs and further improve their performance.

3.1 Extracting Entities and Visual Attributes from Text

The primary goal of our proposed data augmentation framework is to generate new images by modifying the value of visual attributes of the mentioned entities. For example, as shown in Figure 2, our data augmentation method changes the *color* (visual attribute) of a *linckia laevigata* (entity) from *blue* (attribute value) to *orange* (attribute value). The first step of text modification is to identify the mentioned entities and visual attributes of mentioned entities within a given piece of text. To this end, we use large language models (LLMs) to extract entities, visual attributes and attribute values given an input text, T , as they demonstrate exceptional capabilities in text comprehension and

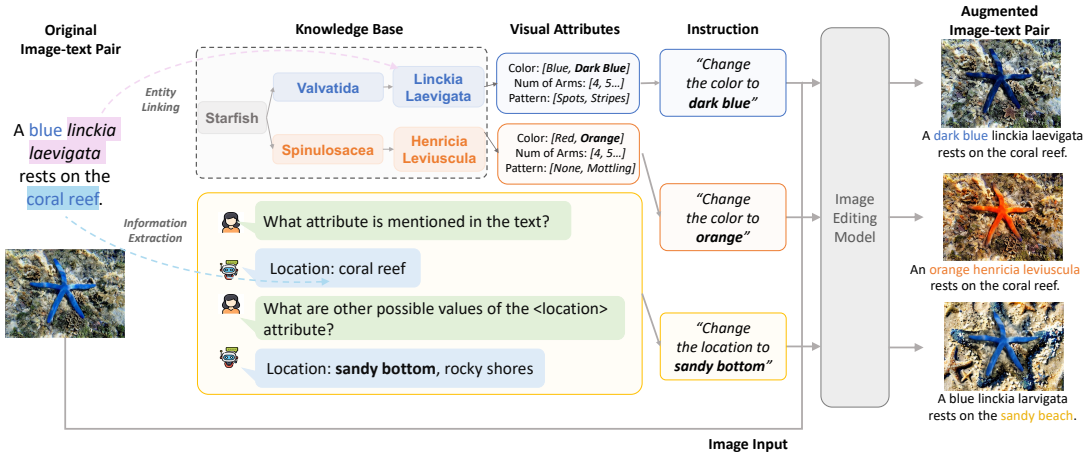


Figure 2: The overall framework of our data augmentation method. Given an image-text pair as input, we first extract entities and their corresponding visual attributes from text. If the object can be linked to an entity in our pre-defined attribute knowledge base, then we collect all possible attribute values from the information of the linked entity. If the object cannot be linked to the knowledge base, then we utilize Large Language Models (LLMs) to extract other possible values. After selecting which visual attribute to modify, we rewrite the original text and use an image editing model to generate new images based on the new text. Finally, we rank the augmented data and output data based on the similarity scores.

generation. Given an original image-text pair (I, T) , we input the text T into an LLM along with the prompt “Extract the mentioned objects, their visual attributes, and values of visual attributes from the sentence: T ”. For example, as illustrated in Figure 2, we can extract from the sentence “A blue linckia laevigata rests on the coral reef” that the entity is *linckia laevigata*, the visual attributes are color and location, and the attribute values are blue and coral reef, respectively. The entities, visual attributes and their values serve as candidates for subsequent visual attribute value substitution.

3.2 KB-based Visual Attribute Substitution

Knowledge Base Construction. After identifying visual attributes mentioned in text T we determine potential substitutions for their attribute values. We leverage attributes from entity-centric KBs to provide accurate and reliable knowledge for substituting visual attribute values. We first parse the information from Wikidata and Wikipedia, and construct an attribute-level KB consisting of entities and their attributes, which consists of two steps: (1) *Graph topology*: We collect entities from Wikidata and use a node in the KB to represent an entity. Each node has an outgoing edge to its parent category node. For instance, as illustrated in Figure 3, both *linckia laevigata* and *linckia guildingi* belong to the parent category *valvatida*, thus resulting in two directed edges from these nodes to

valvatida. (2) *Node attributes*: The visual information for each node in the KB is derived from its corresponding Wikipedia articles. We collect the textual content of each Wikipedia page, then employ LLMs to extract all visual attributes and their possible values described within the article. For instance, the entity *linckia laevigata* may have color of blue and dark blue, with the number of arms starting from four.

After building the KB, we link each entity extracted from T to a node N in the KB using the Spacy Entity Linker (Honnibal et al., 2020). To generate a new augmented data sample, we use the following two attribute value substitution methods.

Attribute Substitution within Single Entity. A single entity may possess multiple plausible attributes, which are identifiable through entity linking to KB. Some of these extracted entities with specific attributes may occur less frequently in the original training dataset than those with more frequently occurring attributes. Therefore, we aim to augment the data to increase the coverage of such long-tail entity instances, so that the model is better fine-tuned to recognize these rare cases well. To elaborate, we randomly choose a visual attribute connected to the entity node N and then sample an attribute value to substitute the current attribute value of N . In this case, the entity stays the same while only its one attribute value is changed. For example, *blue linckia laevigata* \rightarrow *dark blue linckia laevigata* as illustrated in Figure 2.

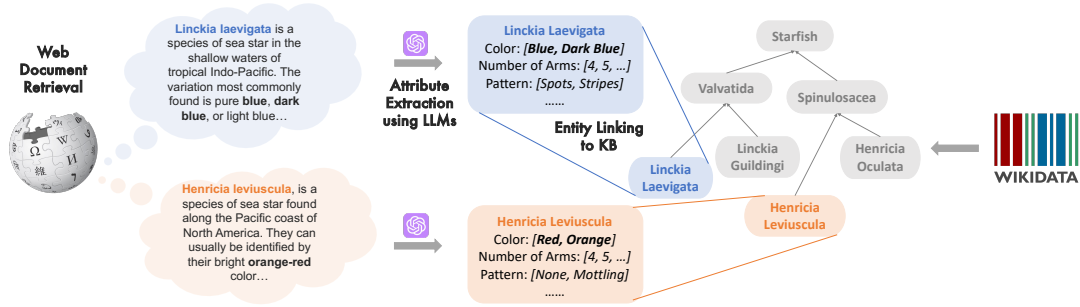


Figure 3: An example from the our pre-defined attribute library. Each node represents an entity collected from Wikidata. An outgoing edge is connected from a node to its parent category. Each node has its visual attributes extracted from the Wikipedia articles.

Attribute Substitution across Sibling Entities. In addition to substituting attributes within a single entity, we notice that there are many entities in KBs that belong to the same parent category and share many visual attributes in common, e.g., the *linckia laevigata* and *henricia leviuscula* in Figure 2. This inspires us to substitute attributes across these sibling entities to introduce similar but different concepts as augmented training data. In this way, the model will contrastively learn from these confusing entities, thereby increasing its robustness to visually similar but different entity concepts. Specifically, we consider changing the entity node n_i to its sibling entity node n_s who share the most visual attributes with n_i . For example, in Figure 3, *linckia laevigata* and *henricia leviuscula* have many attributes in common, so it is feasible to change the original entity to the new entity. We therefore substitute the entity *linckia laevigata* with *henricia leviuscula*, and then change its *color* for *henricia leviuscula* (e.g., *orange*). The resulting substitution is therefore *blue linckia laevigata* \rightarrow *orange henricia leviuscula*.

3.3 LLM-based Visual Attribute Substitution

In some cases, the extracted entity or visual attribute is too general and cannot be linked to any node in the KB (e.g., *coral reef* serving as a background in Figure 2). Therefore, in addition to KBs, we also use LLMs to obtain new values for auxiliary visual attributes such as background, as they are broadly trained on a large amount of data and thus have acquired commonsense knowledge to provide alternative attribute values for such cases. For example, in Figure 2, after we extract that the *location* is *coral reef*, we use the prompt “What are other possible values for the <location> attribute in this sentence?” to generate new location value substitutions, such as *sandy bottom*

and *rocky shores*. It is worth noting that LLMs may not consistently produce valid substitute attribute values, as they may lack adequate knowledge regarding specialized fields or long-tail concepts. This deficiency may lead to LLMs generating inaccurate responses, i.e., hallucination. For instance, when prompt the LLMs for all possible colors of *linckia laevigata*, LLMs may provide incorrect answers such as “*orange*” and “*yellow*”, which are implausible colors for *linckia laevigata*. Therefore, we rely on KBs to extract accurate, knowledge-grounded attributes for substitution.

It is worth noting that the models we utilize in each component may not be perfect, which can affect the performance of the proposed approach. Our experimental results in Section 4.6 indicate that the error rates of the information extraction, entity linking, and visual attribute substitutions are relatively low, which do not significantly impact the quality of the generated data.

3.4 Image Editing

After modifying an image-text pair (I, T) to (I, T') with a new text T' , we edit the image I according to T' . We employ an image editing model InstructPix2Pix (Brooks et al., 2023), which can take as input an image and instruction on how to modify the image, and output the modified image following the instruction. The instruction here is “Change the [attribute] of the [entity] to [value]”, where [entity] and [attribute] are the mentioned entity and selected attribute type, respectively, and [value] is the new attribute value output by the KB or LLM. As illustrated in Figure 2, starting with the original image on the left, we generate three new images on the right using InstructPix2Pix with different instructions. The first image keeps the entity *linckia laevigata* unchanged while changing its color to *dark blue*, whereas the

second image changes the color to *orange*, updating the entity category to *henricia leviuscula* and its corresponding text description accordingly. The third one is the result of changing the attribute of *location* to *sandy beach* by querying LLMs; this leaves the central entity of the image unperturbed, providing a robust way to leverage LLMs only for attributes that are not entity-related.

3.5 Augmented Data Selection

Our method transforms an image-text pair (I, T) to a modified image-text pair (I', T') . However, not all modified image-text pairs are suitable as augmented data; some image I' being too similar to their original counterpart I , thereby providing minimal new signal for subsequent model training. Conversely, other generated image I' diverging too much from their original counterpart I may significantly drift the image away from the original data distribution and mislead the model training. To determine the validity of the augmented data, we calculate the similarity between a generated image I' and its original image I using the Fréchet Inception Distance (FID) score (Heusel et al., 2018). FID calculates the Fréchet distance between feature vectors of the original and generated images, which aligns closely with human judgment and is frequently utilized to assess the quality of generated data. Ideally, we aim to empirically maintain the similarity score within a specific range to ensure that I' exhibit a *reasonable amount* of difference from I as indicated in the ablation study. The experimental results on selecting the similarity range is presented in Appendix A.3.

4 Experiments

To assess the effectiveness of data augmentation methods, we select four evaluation tasks: *image classification*, *visual question answering*, *image-text retrieval*, and *image captioning*.

4.1 Foundation Models and Baseline Methods

We use CLIP (Radford et al., 2021) and LLaVA-1.5 (7B) (Liu et al., 2023a) model as the foundation models in this work. CLIP is a multimodal model that uses contrastive learning to jointly align the visual and textual representations. LLaVA-1.5 is an open-source, auto-regressive multimodal vision-language model (VLM) trained by fine-tuning Vicuna-v1.5 (Chiang et al., 2023) on GPT-4-generated multimodal instruction-following data.

Given an image input and text instruction, LLaVA-1.5 generates output texts based on its reasoning upon the two modalities. We use GPT-4 (OpenAI, 2023) as the LLMs in each component.

We compare our proposed method against five different baseline methods to demonstrate its effectiveness (we do not include BiAug (Wu et al., 2023) since the code has not been released yet): (1) *Zero-shot*: Models are evaluated without fine-tuning on any data. This setting is established to examine the initial ability of the models on all four downstream tasks. (2) *NoAug*: Only the original training data is used to fine-tune the models without any augmented data. (3) *NaiveAug*: Two naive augmentation methods are applied to texts and images independently as follows. We use AEDA (Karimi et al., 2021) to randomly insert punctuation marks into original text, and we use TrivialAugment (Müller and Hutter, 2021) to randomly apply center cropping, rotation, or invert, to images. (4) *MixGen* (Hao et al., 2023): Generates new data instance by interpolating images on the pixel-level and concatenating texts. This is state-of-the-art augmentation method. Specifically, given two image-text pairs (I_i, T_i) and (I_j, T_j) , a new image-text pair (I'_k, T'_k) is generated by $I'_k = \lambda I_i + (1 - \lambda)I_j$ and $T'_k = \text{concat}(T_i, T_j)$, where λ is a hyperparameter. (5) *LeMDA* (Liu et al., 2022): Generates augmented data in the latent feature space. We use CLIP to encode the original training data into embeddings, then feed them to LeMDA to generate new latent embeddings; these embeddings are used as augmented data to fine-tune an MLP module in the image classification task. Note that LeMDA cannot be used for LLaVA-1.5 and cannot be used in tasks other than image classification.

4.2 Image Classification

Dataset. We use *iNaturalist 2021* (Horn et al., 2018) as the dataset for image classification. The iNaturalist dataset consists of large scale species of plants and animals in the natural world. It contains 10,000 species with a training set of 2.7M images. To better mimic the scenario of annotated data scarcity, we sample from a mini dataset with all 246 species of Mammalia. Each class has 30/15/15 images for training/validation/inference. **Experimental Setup.** For CLIP, we transform the class labels in iNaturalist dataset into natural language descriptions: “[label]” \rightarrow “a photo of [label]”, following caption formats in CLIP (Radford et al., 2021). CLIP takes as input an image and

all class labels, then outputs logit scores for these classes. The label with the highest logit score is taken as the predicted result of CLIP model. For LLaVA-1.5, we evaluate its performance by asking the model what is included in the image, and then verify whether the true labels are presented in the generated responses. The evaluation prompt is: “What is the name of the mammal that appears in this image? For example, if it’s a picture of a bengal tiger, output a fine-grained label ‘Bengal Tiger’ or use its binomial nomenclature ‘Panthera tigris tigris’. Provide your answer:”. This allows us to assess model’s classification ability based on the provided images.

Results. The results of *Precision*, *Recall*, and F_1 for image classification task are presented in the left part of Table 1. As shown from the zero-shot results, the pretrained foundation models have poor performance on fine-grained concept recognition, with F_1 scores of 0.090 and 0.041 on CLIP and LLaVA, respectively. After fine-tuning with the original training data, both models have a much better performance, with a 24.9% and 47.6% absolute gain on F_1 scores. While both NaiveAug and LeMDA demonstrate some improvement in model performance, our method achieves the best results among all existing methods. It is worth noting that the F_1 score of MixGen is worse than NoAug. This is because the interpolation of images distorts the visual attribute of the fine-grained concepts, thereby adversely affects model training. Conversely, our method is able to generate new images by modifying the visual attributes of entities. This facilitates a more comprehensive learning of fine-grained concepts by foundation models.

4.3 Visual Question Answering

Datasets. Visual Question Answering (VQA) v2.0 (Goyal et al., 2017) dataset consists of open-ended questions to images. These questions require understanding vision, language, and commonsense knowledge to provide answers. VQA-2.0 has 265,015 images and each image has at least 3 related questions.

Experimental Setup. We consider the VQA task as an answer generation task. We utilize LLaVA as the foundation model. Given the open-ended nature of the task, we let the model generate free-form answers without any constraints. Then we compute the textual similarity between the output of LLaVA and the true answer.

Results. The results of VQA task are shown in the right part of Table 1. We evaluate the performance on the test-dev dataset via textual similarities using Universal Sentence Encoder (USE) (Cer et al., 2018) and BERTScore (Zhang et al., 2020).. It is clear that, compared with Zero-shot, the performance of LLaVA improves greatly after fine-tuning. This is probably because the ground truth answers to the questions are typically simple and short, which makes the task relatively easier. As demonstrated in the table, the textual similarity achieved by our method surpasses the best baseline method MixGen by 1.1% on USE and 1.4% on BERTScore.

4.4 Image-Text Retrieval

Dataset. *Flickr30k* (Young et al., 2014) contains 31,000 images, each with 5 human-annotated referenced sentences that describe the image. This dataset is widely used in image-text retrieval task. Similar to iNaturalist, we sample 5k images from the training set and use the entire 1k test set for evaluation.

Experimental Setup. Image-text retrieval includes two subtasks: text-to-image and image-to-text retrieval. We use CLIP to calculate the embedding of the given image, as well as the embeddings of all candidate captions in the test set. We compare the cosine similarity between the image embedding and each text embedding, and output top K captions with the highest similarity scores as the retrieved results. We follow existing work and use $Recall@K$ as evaluation metric.

Results. The results of image-text retrieval are shown in the left part of Table 2. The zero-shot performance of the pretrained CLIP is already very good on both image retrieval and text retrieval, because it is originally trained using the contrastive loss between image and text embeddings. After fine-tuning, the performance on both subtasks can be further improved in most cases. Note that the improvement of our method over baseline methods in this task appears less significant compared to other tasks. This is primarily due to the already high zero-shot performance of CLIP, leaving limited room for further improvement.

4.5 Image Captioning

Experimental Setup. Image captioning task aims to generate natural language descriptions of an image. We use LLaVA-1.5 as the foundation model, and Flickr30k as the evaluation dataset as intro-

Method	Image Classification (CLIP)			Image Classification (LLaVA)		VQA	
	Precision	Recall	F ₁	F ₁	ExactMatch	USE	BERTScore
Zero-shot	0.074	0.113	0.090	0.041	0.002	0.221	0.825
NoAug	0.332	0.347	0.339	0.517	0.557	0.815	0.949
NaiveAug	0.386	0.336	0.359	0.192	0.241	0.821	0.961
MixGen	0.343	0.318	0.330	0.314	0.357	0.824	0.959
LeMDA	0.368	0.354	0.361	-	-	-	-
ARMADA	0.391	0.386	0.389	0.588	0.621	0.835	0.975

Table 1: Results of Precision, Recall, and F₁ on iNaturalist dataset for image classification (left part) and results of textual similarity on VQA v2.0 dataset for visual question answering (right part). The foundation model is LLaVA-1.5 for VQA.

Method	Image Retrieval			Text Retrieval			Image Captioning	
	R@1	R@3	R@5	R@1	R@3	R@5	USE	BERTScore
Zero-shot	0.589	0.765	0.824	0.612	0.775	0.836	0.422	0.896
NoAug	0.619	0.785	0.830	0.645	0.807	0.854	0.642	0.907
NaiveAug	0.631	0.788	0.838	0.641	0.804	0.862	0.648	0.911
MixGen	0.626	0.786	0.838	0.592	0.770	0.826	0.659	0.903
ARMADA	0.646	0.797	0.847	0.646	0.811	0.872	0.682	0.918

Table 2: Results of Recall@K for image-text retrieval (left part) and results of textual similarity for image captioning (right part). We use Flickr30k dataset for both tasks. The foundation model is CLIP for image-text retrieval and LLaVA-1.5 for image captioning.

duced in Section 4.4. Specifically, given an image as input, we use the prompt “Describe this image using one simple sentence” to ask LLaVA-1.5 to generate a caption.

To evaluate the quality of generated captions, we compare the textual similarity between the generated caption and the gold-standard annotation for a given image using USE and BERTScore. Since there may be multiple gold-standard captions for an image, we calculate the similarity score of a generated caption with each gold-standard caption, and return the maximum as the final score for this generated caption.

Results. The results of image captioning task are presented in the right part of Table 2. Our method ARMADA achieves the best performance over all baseline methods. Specifically, the performance gain of our method on USE score is 4.0% over NoAug, and 2.3% over the best baseline augmentation method MixGen. We provide detailed case analysis of the generated captions by our method and by baseline methods in Appendix B.

4.6 Error Analysis

We investigate the error rate of each component in the data augmentation process and how they affect our model. Specifically, we manually check the correctness of attribute extraction and the visual attribute substitution. It turns out that the percent-

age of incorrect attributes that are extracted is quite low ($4 / 113 = 3.5\%$). The percentage of inappropriate substitution by LLMs is also very low ($1 / 73 = 2.7\%$). The visual attribute substitutions from KBs are template-based substitutions from possible attribute values, which will not incur any error aggregation issues.

5 Conclusions and Future Work

We propose a novel data augmentation method that utilizes KBs and LLMs to generate multimodal data. The proposed framework is able to generate semantically consistent data that solves the potential issues of the existing methods. Our method significantly improves the MLM’ performance on various downstream tasks, without the need of high-cost annotated data. Experiment results also demonstrate the effectiveness of our proposed method compared to the baseline methods.

In the future, we aim to incorporate more modalities into our framework such as video and audio. We also plan to rank visual attributes and select the most influential attributes for augmentation. Moreover, existing image editing tools our framework relies on do not perform consistently well. Designing a new visual attribute editing model to further enhance the quality of the augmented data is also a promising research direction.

6 Limitations

Our proposed method demonstrates the effectiveness only on image-text data. However, to enhance the practical utility of our method, it would be advantageous to expand our data augmentation method to include more modalities, such as video and audio. Furthermore, as discussed earlier, although the error rate in each component is low and will not affect the performance much, we still aim to incorporate better attribute extraction and visual attribute substitution models into the framework to further improve our method.

7 Ethical Consideration

We acknowledge that our work is aligned with the *ACL Code of the Ethics* (Gotterbarn et al., 2018) and will not raise ethical concerns. We do not use sensitive datasets/models that may cause any potential issues.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Tim Brooks, Aleksander Holynski, and Alexei A. Efros. 2023. *Instructpix2pix: Learning to follow image editing instructions*. *Preprint*, arXiv:2211.09800.
- Daniel Cer, Yinfei Yang, Sheng yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. *Universal sentence encoder*. *Preprint*, arXiv:1803.11175.
- Xiuxi Chen, Jyun-Yu Jiang, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, and Wei Wang. 2023. *Min-prompt: Graph-based minimal prompt data augmentation for few-shot question answering*. *Preprint*, arXiv:2310.05007.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. *Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality*.
- W Dai, J Li, D Li, AMH Tiong, J Zhao, W Wang, B Li, P Fung, and S Hoi. 2023. *Instructblip: Towards general-purpose vision-language models with instruction tuning*. arxiv 2023. *arXiv preprint arXiv:2305.06500*.
- DW Gotterbarn, Bo Brinkman, Catherine Flick, Michael S Kirkpatrick, Keith Miller, Kate Vazansky, and Marty J Wolf. 2018. *Acm code of ethics and professional conduct*.
- 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiaoshuai Hao, Yi Zhu, Srikar Appalaraju, Aston Zhang, Wanqian Zhang, Bo Li, and Mu Li. 2023. *Mixgen: A new multi-modal data augmentation*. *Preprint*, arXiv:2206.08358.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2018. *Gans trained by a two time-scale update rule converge to a local nash equilibrium*. *Preprint*, arXiv:1706.08500.
- Matthew Honnibal, Ines Montani, Sofie Van Lan-deghem, and Adriane Boyd. 2020. *spaCy: Industrial-strength Natural Language Processing in Python*.
- Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. 2018. *The inaturalist species classification and detection dataset*. *Preprint*, arXiv:1707.06642.
- Xiaomeng Jin and Heng Ji. 2024. *Schema-based data augmentation for event extraction*. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14382–14392, Torino, Italia. ELRA and ICCL.
- Akbar Karimi, Leonardo Rossi, and Andrea Prati. 2021. *Aeda: An easier data augmentation technique for text classification*. *Preprint*, arXiv:2108.13230.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022a. *Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation*. *Preprint*, arXiv:2201.12086.
- Manling Li, Ruochen Xu, Shuohang Wang, Luwei Zhou, Xudong Lin, Chenguang Zhu, Michael Zeng, Heng Ji, and Shih-Fu Chang. 2022b. *Clip-event: Connecting text and images with event structures*. In *Proc. Conference on Computer Vision and Pattern Recognition (CVPR2022)*.
- Manling Li, Alireza Zareian, Qi Zeng, Spencer Whitehead, Di Lu, Heng Ji, and Shih-Fu Chang. 2020. *Cross-media structured common space for multimedia event extraction*. In *Proc. The 58th Annual Meeting of the Association for Computational Linguistics (ACL2020)*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. *Improved baselines with visual instruction tuning*. *Preprint*, arXiv:2310.03744.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. *Visual instruction tuning*. *Preprint*, arXiv:2304.08485.
- Zichang Liu, Zhiqiang Tang, Xingjian Shi, Aston Zhang, Mu Li, Anshumali Shrivastava, and Andrew Gordon Wilson. 2022. *Learning multimodal data augmentation in feature space*. *arXiv preprint arXiv:2212.14453*.
- LINHAO LUO, Yuan-Fang Li, Reza Haf, and Shirui Pan. 2023. *Reasoning on graphs: Faithful and interpretable large language model reasoning*. In *The Twelfth International Conference on Learning Representations*.
- Xue-Jing Luo, Shuo Wang, Zongwei Wu, Christos Sakaridis, Yun Cheng, Deng-Ping Fan, and Luc Van Gool. 2023. *Camdiff: Camouflage image augmentation via diffusion model*. *Preprint*, arXiv:2304.05469.

- Samuel G Müller and Frank Hutter. 2021. Trivialaugment: Tuning-free yet state-of-the-art data augmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 774–782.
- Samuel G. Müller and Frank Hutter. 2021. [Trivialaugment: Tuning-free yet state-of-the-art data augmentation](#). *Preprint*, arXiv:2103.10158.
- R OpenAI. 2023. Gpt-4 technical report. *ArXiv*, 2303.
- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. 2024. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). *Preprint*, arXiv:2103.00020.
- Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni, Heung-Yeung Shum, and Jian Guo. 2023. Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph. In *The Twelfth International Conference on Learning Representations*.
- Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. 2021. [Augmented sbert: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks](#). *Preprint*, arXiv:2010.08240.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. 2023. [Effective data augmentation with diffusion models](#). *Preprint*, arXiv:2302.07944.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.
- Qiyu Wu, Mengjie Zhao, Yutong He, Lang Huang, Junya Ono, Hiromi Wakaki, and Yuki Mitsufuji. 2023. [Towards reporting bias in visual-language datasets: bimodal augmentation by decoupling object-attribute association](#). *Preprint*, arXiv:2310.01330.
- Kang Min Yoo, Dongju Park, Jaewook Kang, Sangwoo Lee, and Woomyeong Park. 2021. [Gpt3mix: Leveraging large-scale language models for text augmentation](#). *Preprint*, arXiv:2104.08826.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. [From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions](#). *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). *Preprint*, arXiv:1904.09675.

A Ablation Study

A.1 Impact of the Size of the Generated Data

To investigate the impact of the amount of the augmented data, we conduct experiments by varying the size of the augmented data relative to the size of the original training data, ranging from 0% to 300%. The results in Table 3 show a decline in model performance when the augmented data size significantly surpassed the original training data size (exceeding 100% to 200%), potentially due to excessive noise introduced by the augmented data. Our findings suggest that, the augmented data size should approximate that of the original training data for best performance.

A.2 Impact of Using KBs

To assess the importance of utilizing KBs, we conduct additional experiments on the image classification task by solely relying on LLMs to do attribute value substitution. Following the aforementioned experimental setup, we fine-tune a CLIP model on the iNaturalist dataset. The F_1 score exhibits a 2.5% decline (from 38.9% to 36.4%) without using KBs. This suggests that though LLMs are able to provide answers for attribute value substitution, the hallucination issue on fine-grained or rare entities can still introduce noise to the training data, thereby impacting the model performance.

A.3 Impact of Similarity Range for Selecting Augmented Data

We conduct experiments to investigate how the similarity between augmented and original data impact the model performance. In the image classification task, we split the augmented dataset into four groups of equal size according to the similarity of the edited image with its original image. Then we use each group as the augmented data to train CLIP. The F_1 scores of the four groups are 0.377, 0.389, 0.383, and 0.364, respectively, from most-similar to most-dissimilar. The results support our claim in Section 3.5 that maintaining similarity scores within a reasonable range achieves the best performance.

B Case Analysis on the Results of Image Captioning

We perform a case analysis to illustrate the effectiveness of our method. In Figure 4, we present two image-caption pairs from the Flickr30k dataset, including both the human-annotated captions and

the captions generated by Zero-shot, NoAug, and ARMADA (using LLaVA as the foundation model). For the image on the left, our method is able to identify the fine-grained concept *karate* whereas the zero-shot and NoAug methods generate a more generalized concept *martial arts*. For the image on the right, the caption generated by our method provides a more detailed and accurate description of the hat, which specifies its *knit* pattern and the *beer logo* pattern. These examples suggest that LLaVA can effectively learn the visual attributes and identifies the fine-grained concepts through our method.

C Human Survey

We design a human evaluation and evaluate our generated augmented data based on two metrics: (1) Realism of the augmented images (2) Semantic consistency between the modified image and text. Since the baseline method LeMDA is augmenting data in the feature space and not possible for visualization, we compare our results with another SOTA method MixGen. We randomly selected 140 images and augmented them using our method and MixGen. We asked 7 people to take the questionnaire and each person evaluated 20 images from our method and 20 images from MixGen (randomly shuffled). For each augmented image, assessors are required to give a score within [0, 1, 2, 3, 4, 5] for each metric and higher score suggests better quality. The average scores for the metrics are shown in Table 4. Compared to MixGen where an augmented image is a pixel-level mixture from two original images and the text is the concatenation of the text pairs, the results show that our augmented data are more realistic and consistent.

D Performance on ImageNet1K Dataset

For image classification, we perform an additional experiment with the ImageNet1K dataset. We use the CLIP model and finetune with 10 images from each class in the training set to mimic the scenario of annotated data scarcity. The experimental results for ImageNet1K are shown in Table 5. Results demonstrate that ARMADA still outperforms the other baseline method on datasets with less concepts.

E Ethical Consideration

We acknowledge that our work is aligned with the *ACL Code of the Ethics* (Gotterbarn et al., 2018)

Dataset	Metric	Size of the augmented data			
		0%	100%	200%	300%
iNaturalist	F_1	0.339	0.389	0.360	0.328
Flickr30k	$TextSim$	0.642	0.682	0.660	0.659
VQA v2.0	$TextSim$	0.815	0.825	0.835	0.816

Table 3: The impact of the size of the generated data on the performance of multiple tasks.

Metric	Realism	Consistency
MixGen	1.29	2.86
ARMADA	3.87	4.23

Table 4: Results of Human Survey.

Method	Accuracy
Zero-shot	0.586
NoAug	0.638
NaiveAug	0.642
MixGen	0.451
ARMADA	0.668

Table 5: Results of image classification task on the ImageNet1K dataset.

and will not raise ethical concerns. We do not use sensitive datasets/models that may cause any potential issues.



Annotation	A girl breaking boards by using karate .
Zero-shot	Two people are practicing martial arts on a mat.
NoAug	Two martial artists are practicing their moves.
AttributeAug	A man and a girl are practicing karate .



Annotation	A man with gauges and glasses is wearing a Blitz hat.
Zero-shot	A man wearing a beer hat and glasses is looking at the camera.
NoAug	A man wearing a beer hat.
AttributeAug	A man wearing a knit hat with a beer logo on it.

Figure 4: A case analysis that shows sample outputs on Flickr30k dataset for image captioning task. We select two images from the test set, the human-annotated captions, and the generated captions from each method. For the image on the left, our method is able to recognize the fine-grained concept *karate*. The image on the right demonstrates that the model is able to provide an accurate description of the *hat*, specifying its *knit* texture and *beer logo* pattern.