

Diffusion Based Augmentation for Captioning and Retrieval in Cultural Heritage

Dario Cioni¹ Lorenzo Berlincioni¹ Federico Becattini² Alberto del Bimbo¹
¹MICC, University of Florence ²University of Siena
dario.cioni@stud.unifi.it {lorenzo.berlincioni, alberto.delbimbo}@unifi.it
federico.becattini@unisi.it

Abstract

Cultural heritage applications and advanced machine learning models are creating a fruitful synergy to provide effective and accessible ways of interacting with artworks. Smart audio-guides, personalized art-related content and gamification approaches are just a few examples of how technology can be exploited to provide additional value to artists or exhibitions. Nonetheless, from a machine learning point of view, the amount of available artistic data is often not enough to train effective models. Off-the-shelf computer vision modules can still be exploited to some extent, yet a severe domain shift is present between art images and standard natural image datasets used to train such models. As a result, this can lead to degraded performance. This paper introduces a novel approach to address the challenges of limited annotated data and domain shifts in the cultural heritage domain. By leveraging generative vision-language models, we augment art datasets by generating diverse variations of artworks conditioned on their captions. This augmentation strategy enhances dataset diversity, bridging the gap between natural images and artworks, and improving the alignment of visual cues with knowledge from general-purpose datasets. The generated variations assist in training vision and language models with a deeper understanding of artistic characteristics and that are able to generate better captions with appropriate jargon.

1. Introduction

Deep learning applications on fine art suffer from an obvious scarcity of data, due to the fact that an artwork is usually a unique piece. In addition, tasks involving both vision and language require on the one hand the modeling of technical language with domain-specific jargon and on the other hand the understanding of difficult and underrepresented visual concepts, such as abstract or stylized drawings. These difficulties entail a challenge for learning algorithms, which would benefit from a large collection of annotated data. A simple solution is to leverage models

pre-trained on general-purpose datasets to address relevant tasks for cultural heritage, such as retrieval, visual question answering or captioning. However, the effect is that such models tend to underperform when applied in the cultural heritage domain. In fact, when the scope moves from real-world images to paintings and other more abstract representations, there is a strong domain shift compared to standard training data, being it composed of natural images. A standard approach to deal with data scarcity is to leverage data augmentation, slightly perturbing the training data to improve variability and let the trained model generalize better. In the vision domain, perturbations usually include adding noise, altering pixel values or changing the overall orientation or illumination of the scene. We argue that these augmentations may indeed alter the semantics of the artworks, where color and spatial distribution of objects can convey significant meanings that are necessary to interpret the painting.

In this paper, we address the above-mentioned limitations by proposing a data augmentation strategy for paintings that has the twofold advantage of increasing the training data as well as preserving the content. In particular, we explore the benefits of augmenting artwork datasets for image captioning. To this end, we leverage both textual descriptions of the paintings and a diffusion model to create several variations of the artworks. By conditioning the diffusion model on the captions, we generate a variability in the visual domain that aids the grounding of objects and entities expressed in artistic form with the technical jargon that describes them. What we propose is therefore an image augmentation at a semantic level, generating multiple variations of artworks while retaining their content and style. By leveraging the expert knowledge of art critics contained in painting descriptions and the natural language understanding capabilities of state-of-the-art generative models, we aim to provide a sophisticated augmentation pipeline capable of generating a sufficient intra-class variability of depicted concepts to enable an effective learning (Fig. 1).

Our main contributions presented in this paper are

- We propose a data augmentation technique for low-

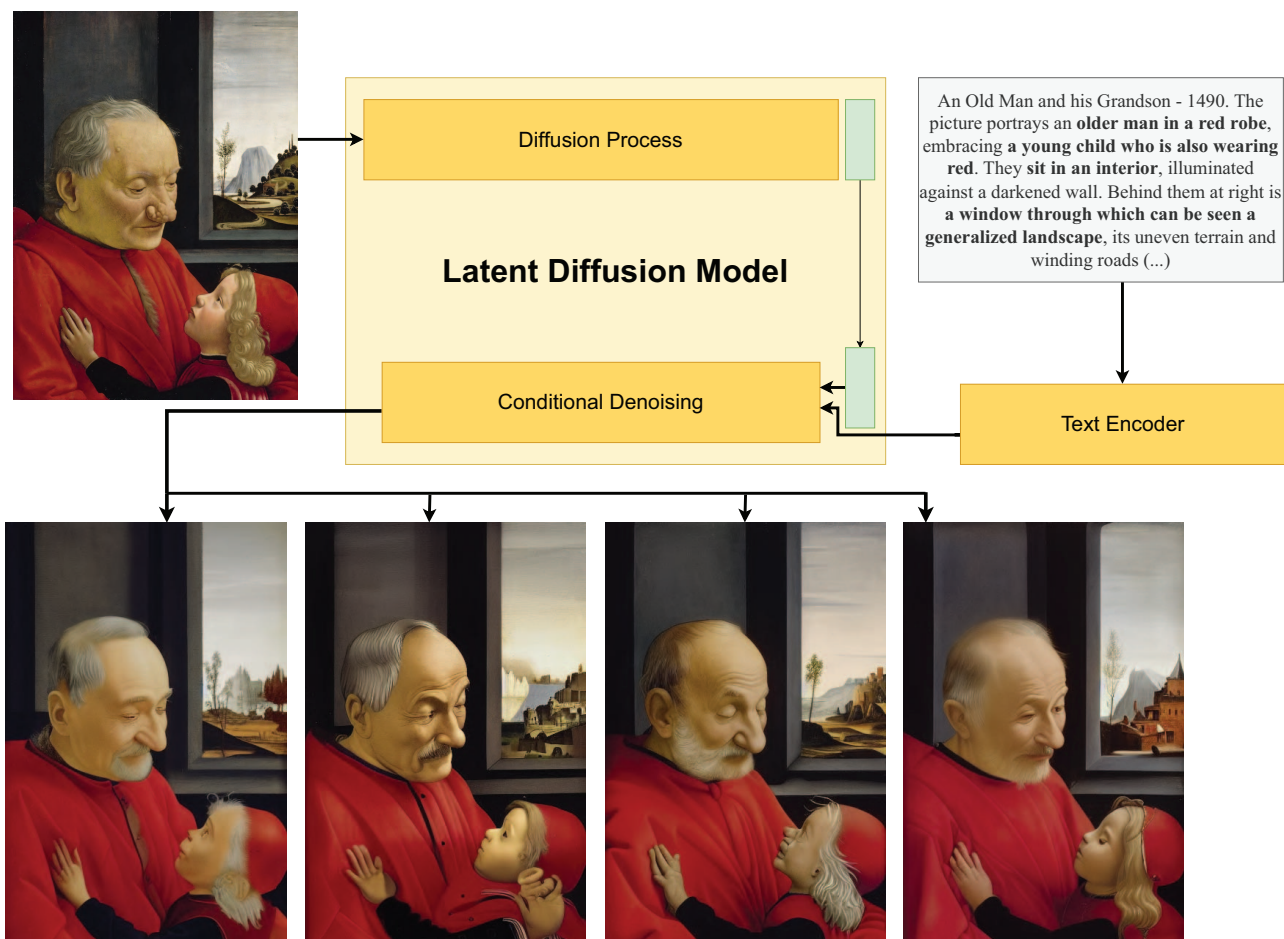


Figure 1: Schematic illustrating the data augmentation pipeline. The conditional generative model allows for both image and text input while in its *Image&Text*→*Image* configuration. We provide the model both the original artwork along with its detailed textual analysis from [47] and use the diffusion model’s outputs as new datapoints for training other models for downstream tasks.

data regime cultural heritage tasks, that works at a semantic level rather than at a pixel-intensity level as standard data augmentations in vision.

- Thanks to our data augmentation strategy based on diffusion models we can favor a visual grounding of linguistic concepts, which in the cultural heritage domain are often expressed using technical and domain-specific jargon.
- We show the benefits of the proposed augmentation strategy in captioning tasks as well as cross-domain retrieval tasks.

2. Related Works

Computer Vision for Cultural Heritage In the domain of cultural heritage, several computer vision approaches have been proposed in the literature. Artwork classifica-

tion [34, 48, 15, 11, 35] and recognition [16, 49, 28] have often been placed at the center of such approaches, sometimes with the end-goal to develop user-engagement applications [5, 33, 2, 8]. In this paper, we mostly deal with the task of image captioning, which implies the automatic generation of a natural language textual description of an image based only on the visual input. This has been an extensively addressed research topic in recent years [46, 54, 29], but not many contributions have been made in the domain of art historical data. In this particular domain, which shifts from the one of natural images, the complexity of the task increases due to a simultaneous lack of labeled data and an increased abstraction.

Currently, available painting datasets with descriptions are constructed by downloading descriptions from online museums or annotating descriptions by crowdsourcing. The Artpedia [47] dataset is composed of paintings paired with textual descriptions from Wikipedia. The dataset thus pro-

vides information about artworks and their context and each sentence is categorized as either a visual sentence or a contextual sentence. Visual sentences describe the visual content of the painting, while contextual sentences provide information that cannot be inferred from raw pixels alone. Such information includes, for instance, the name of the painter, its artistic style, or the museum in which it is kept. The dataset was originally introduced as a dataset for cross-modal retrieval as well as captioning and it has been further annotated for visual question answering purposes in [7]. Similarly, the AQUA dataset [21] has been proposed to train visual question answering models in the cultural heritage domain. More recently, ArtCap [32] provided an image captioning dataset containing 3,606 paintings, each one associated to five textual descriptions, with a mean length for each caption of 11 words.

A larger example of an artwork dataset is presented in [16] consisting of more than 80K webly-supervised images from 3120 classes, and a subset of 200 classes with more than 1300 verified images. Text and metadata for each class is also provided, to support zero-shot learning and other multi-modality techniques in general. An ontological knowledge base has been exploited in [4] to create a large-scale cultural heritage dataset, annotated with visual and contextual data. The authors adopted ArCo [10], the Italian cultural heritage knowledge graph, to extract information about approximately 500K cultural assets and leveraged a semi-automatic annotation approach for generating 6.5M question-answer pairs.

Generative Models for Data Augmentation Data augmentation can be defined as the process through which data can be transformed without changing its semantics. By using this definition we can tie the efficacy of an augmentation method to a task and not to the type of data alone. In most of the computer vision tasks that work with natural images, the usual augmentation strategies involve flipping the image, adding random noise, and changing its brightness and colors. When it comes to fine-art though, such changes might be detrimental due to the strict relation between the semantics of the original art piece and its details (i.e. the relative position of characters in religious art, the use of strong light contrast in a Caravaggio painting, or the symbolic choice of a particular color). An attempt to augment training data for object detection in artworks has been recently proposed [27], where a style transfer model is applied to natural images to generate images that resemble paintings. A possible approach to obtain a larger, more diverse, dataset has been explored in many works outside of the scope of cultural heritage. In these cases, the input data is used to train a generative model, which in turn will produce new data coming from the training domain. [44, 3, 26, 25, 6]. Also in [18] the authors used a CycleGAN [56] for image-

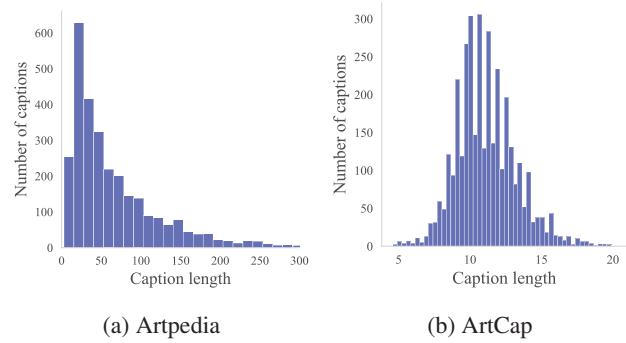


Figure 2: Distribution of caption lengths in the Artpedia [47] and ArtCap [32] datasets

to-image translation of thermal to pseudo-RGB data. The use of these frameworks to perform data augmentation in order to improve the performance of a separate classifier has been studied in multiple previous works such as [1] in which they focus on improving one-shot learning, and [9], where segmentation of medical images is enhanced by GAN augmented data. In [38] synthetic data coming from a simulator is adapted and used to train an RL agent for autonomous driving.

Recently Diffusion Models (DM) [24, 41] reached new impressive levels, compared to GANs, both in terms of output quality and fidelity to the conditional inputs such as text or additional images and have been employed for data augmentation objectives such in [51, 22]. Most of these applications focused mostly on evaluating the ability of diffusion models to generate synthetic data for classification problems, we instead are going to focus on different downstream tasks. Latent Diffusion Models (LDM) [41] perform the diffusion process in a latent space learned by a convolutional auto-encoder. This allows to greatly reduce the training and inference cost of the model compared to pixel-based DMs, while maintaining a high visual fidelity. By introducing cross-attention layers in the diffusion model architecture, the generation can be conditioned on a wide variety of sources, including text and images. The popular Stable Diffusion model is based on LDM [41], bringing further improvements thanks to an internet-scale training. Motivated by the recent success of large generative models, we posed the research question regarding whether diffusion models can be used to augment visual recognition datasets with synthetic images, especially when working in underrepresented domains such as cultural heritage. Our findings show that using images generated by a diffusion model, conditioned by a textual description, leads to improved performance compared to vanilla training as well as training using standard computer vision data augmentation techniques.

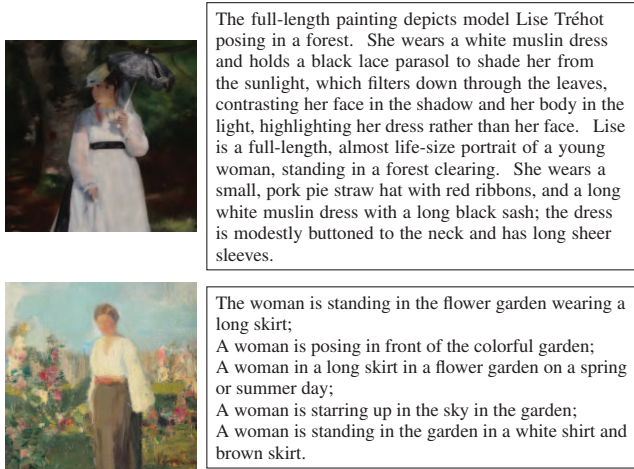


Figure 3: Samples of images along with their textual descriptions from Artpedia (top) and ArtCap (bottom) datasets.

3. Data

Experiments were performed on both the Artpedia [47] and the ArtCap [32] datasets. Albeit similar in structure these two datasets differ in multiple ways from one another. The [32] dataset contains a fixed set of 5 sentences per artwork while [47] has on average, 3.1 visual sentences and 6.5 contextual sentences per artwork. Artpedia contains a collection of 2,930 painting images, each associated with a variable number of textual descriptions, which are combined together into a single description. Overall, the dataset contains a total of 28,212 sentences. Out of these, 9,173 are labeled as visual sentences and the remaining 19,039 are categorized as contextual. In our work, we only consider visual sentences since we focus on augmenting images.

The respective syntactic style is also quite different in the two datasets: where Artpedia chooses paragraph-long academic descriptions, Artcap limits itself to shorter and simpler captions.

The word count distribution of the captions for the two datasets is shown in Fig. 2. Note that, on average, a single visual sentence of Artpedia is composed by 22 words, and the caption is 70 words, which is considerably longer than most common Image Captioning datasets [12, 43, 45]. We also evaluated randomly sampling one of the visual sentences, but since each visual sentence describes only a small portion of the image, it led to worse results. Both Artpedia and ArtCap provide validation and test splits, composed of 10% and 10% validation samples, and 10% and 50% test samples, respectively. Samples from the two datasets are shown in Fig. 3.



Figure 4: Samples of the augmented images. *Left*: Original image and its caption.; *Right*: Multiple samples of the augmented images using the combination of the provided description and the original input image.

4. Method

In order to generate the augmented version of the datasets we employ a LDM (Latent Diffusion Model), Stable Diffusion¹, to generate multiple version of each image belonging to the original dataset as illustrated in Fig. 1.

In our work, we employed versions 1.4 and 1.5 of Stable Diffusion. In Stable Diffusion 1.4 the checkpoint was initialized with the weights of the Stable-Diffusion-v-1-2 checkpoint and subsequently fine-tuned on 225K steps at resolution 512x512 on the "laion-aesthetics v2 5+" subset of LAION dataset [42] and 10% dropping of the text-conditioning to improve classifier-free guidance sampling. In Stable Diffusion 1.5, the initialization checkpoint and the finetuning procedure is the same as Stable-Diffusion-1-4, but the finetuning is performed for more steps (595K). We employed Stable Diffusion v1.4 to augment Artpedia [47] dataset and Stable Diffusion v1.5 for the ArtCap [32] dataset.

Given a dataset \mathcal{D} of N samples (x_i, y_i) formed by an image x_i and a set of captions y_i , we augment it by generating a set $S_i = \{(\tilde{x}_{i1}, \mathbf{y}) \dots (\tilde{x}_{iM}, \mathbf{y})\}$ of synthetic variations for each image x_i , obtaining a synthetic dataset $\tilde{\mathcal{D}}$ of $N \times M$ samples. Each variation was generated using both a textual prompt built by providing the caption and the

¹<https://stability.ai/blog/stable-diffusion-public-release>

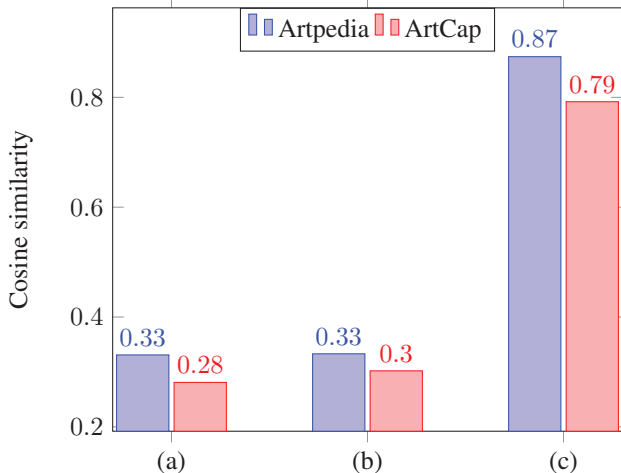


Figure 5: Average cosine similarity between CLIP embeddings of: (a) real images and the associated captions; (b) synthetic images and the associated captions; (c) real images and their synthetic variations.

original image to guide the generation. To obtain different images, the generation seed was changed for each variation (see Fig. 4).

To gain an intuition of the quality of the synthetic dataset, we calculated an embedding of the text and images of each sample in \mathcal{D} and $\tilde{\mathcal{D}}$ using a CLIP-ViT/B-16 model [40]. We can see from Fig. 5 (a) and (b) that the average cosine similarity between images and their captions maintains a similar value in the original dataset \mathcal{D} and the synthetic dataset $\tilde{\mathcal{D}}$, suggesting that synthetic images preserve the relation with the caption. Moreover, in (see Fig. 5 (c)) we see that variations maintain a high similarity with the original images. This can also be seen in Fig. 4.

During training, we insert in each position of the training minibatch a sample $(x_i, y_i) \in \mathcal{D}$ with probability α or one of its synthetic variations $(x_{ij}, y_i) \in \tilde{\mathcal{D}}$ with probability $(1 - \alpha)$, where x_{ij} is sampled uniformly from S_i . We use a value of $\alpha = 0.5$, to balance real and generated images during training as suggested in [51].

5. Experiments

In order to test our augmentation technique we perform multiple experiments over different tasks. As a first downstream task we train an image-captioning model using both augmented and non-augmented versions of the dataset. For this set of experiments, we selected medium-sized, Transformer-based Vision and Language models which can be trained end-to-end and can be employed for a variety of different tasks. In particular, we use the GIT_{base} [54] model and the $\text{BLIP}_{\text{base}}$ [29] model.

GIT [54] (Generative Image-to-text Transformer) is a

Transformer [52] model which can be applied to many Vision and Language tasks. It leverages a CLIP ViT image encoder [40] and a single Transformer text decoder, which are jointly trained under a single language modeling task on large-scale pre-training data. It is publicly available in two sizes, GIT_{base} (129 M parameters) which employs a CLIP/ViT-B/16 encoder and $\text{GIT}_{\text{large}}$ (347M parameters), with a CLIP/ViT-L/14 encoder. BLIP [29] instead is a model that effectively uses noisy web data for pre-training by bootstrapping the captions, generating new synthetic captions and removing the noisy ones. It employs a multimodal mixture of encoder-decoder which are jointly trained with three vision-language objectives: image-text contrastive learning, image-text matching, and image-conditioned language modeling. The architecture is composed of a ViT [20] encoder to encode images and a BERT [19] to encode text. Both the GIT and BLIP models were initialized with the available pre-training weights² and finetuned for 10 training epochs using the AdamW [31] optimizer with a $5e^{-05}$ learning rate and 500 steps of warm-up using batches of 8 images.

The second task is adopt to prove the effectiveness of our proposed strategy is cross-domain retrieval. Here, we perform retrieval both of images given their textual description and vice versa. For this downstream application, we use the CLIP model [40], using the openCLIP³ implementation. To finetune the CLIP model we used again the AdamW optimizer with a learning rate of $5e^{-04}$.

5.1. Quantitative Results

5.1.1 Metrics

To quantitatively assess the quality of the generated captions, standard language evaluation metrics are used. Those include BLEU [39], ROUGE [30] and METEOR [17], typically used for machine translation tasks, and CIDEr [53], specifically developed for the image captioning task. In addition, a semantic similarity between generated captions and references is measured with BERTScore [55] metric. BLEU score calculates n -gram precisions between a candidate sentence and a set of human-generated references, multiplied by a brevity penalty. Single n -gram precisions are then combined following a geometric mean to obtain a final score. It is common practice to report BLEU scores with n -grams ranging from 1 to 4. ROUGE-L calculates a F-measure using the Longest Common Subsequence (LCS) between a candidate sentence and a set of references. METEOR computes a harmonic mean of precision and recall between unigrams of aligned candidate and reference sentences, where the mapping used for alignment follows various strategies, including exact match, synonyms and para-

²<https://huggingface.co/microsoft/git-base>

³https://github.com/mlfoundations/open_clip

Dataset	Model	B@1	B@2	B@3	B@4	METEOR	ROUGE	CIDEr	BERTScore
Artpedia	GIT _b (zero-shot)	0.0000	0.0000	0.0000	0.0000	0.0144	0.0749	0.0144	0.6905
	GIT _b w/o DA	0.0179	0.0088	0.0046	0.0026	0.0385	0.1433	0.0505	0.7291
	GIT _b w/ DA	0.0184	0.0092	0.0048	0.0027	0.0390	0.1479	0.0673	0.7316
	BLIP _b (zero-shot)	0.0000	0.0000	0.0000	0.0000	0.0161	0.0830	0.0043	0.7112
	BLIP _b w/o DA	0.0050	0.0026	0.0014	0.0009	0.0331	0.1568	0.0766	0.7262
	BLIP _b w/ DA	0.0118	0.0062	0.0035	0.0020	0.0369	0.1658	0.0906	0.7291
ArtCap	GIT _b (zero-shot)	0.3993	0.2541	0.1548	0.0888	0.1237	0.3128	0.2114	0.7877
	GIT _b w/o DA	0.7311	0.5675	0.4263	0.3196	0.2360	0.5148	0.6263	0.8752
	GIT _b w/ DA	0.7475	0.5825	0.4407	0.3321	0.2376	0.5166	0.6445	0.8737
	BLIP _b (zero-shot)	0.6224	0.4007	0.2487	0.1512	0.1606	0.3951	0.3467	0.8098
	BLIP _b w/o DA	0.7710	0.5972	0.4515	0.3343	0.2442	0.5128	0.6851	0.8759
	BLIP _b w/ DA	0.7654	0.5909	0.4541	0.3491	0.2466	0.5170	0.6862	0.8748

Table 1: Image Captioning results on Artpedia and ArtCap using GIT [54] model, measuring BLEU (n -grams 1 to 4), METEOR, ROUGE, CIDEr and BERTScore metrics.

phrases. CIDEr measures the consensus among a candidate sentence and a set of references by computing the cosine similarity of TF-IDF weighted n -gram vectors. BERTScore uses the word embeddings computed by a pretrained Transformer model to measure the semantic similarity between a candidate sentence and a reference.

5.1.2 Baselines

We compare our method against state-of-the-art augmentation techniques, such as AutoAugment [13], AugMix [23], RandAugment [14] and TrivialAugment [36]. AutoAugment [13] is an augmentation framework for vision models that casts the search of parameters for data augmentation as an optimization problem and solves it using reinforcement learning. RandAugment [14] improves on AutoAugment [13] by both considerably reducing the parameters search space from 10^{32} to 10^2 and matching or exceeding performances of [13]. AugMix [23] layers multiple randomly sampled augmentation operations in concert with a consistency loss to improve model robustness. Finally [37] improves on the previous strategies by further simplifying the search space. All of the previous models are tailored to image classification tasks, more recently a number of works focused on data augmentation specifically developed for detection problems. In Fig. 6 we provide a comparison of the augmentation operations performed by the aforementioned state of the art techniques and ours applied to the same image.

5.1.3 Results

Image Captioning We present the results of the image captioning task using differently trained GIT models in Tab. 1 over the two chosen datasets. For the Artpedia dataset, the test results show a clear and consistent improve-

ment using our augmentation technique over all the aforementioned metrics. Similarly, on ArtCap we report gains in most metrics, with only a slight decrease compared to standard training with no data augmentation. It is also easy to notice how the two datasets differ in terms of complexity. In fact, all the models struggle more with Artpedia [47], obtaining results that are much lower in absolute terms. This is due to the nature of the captions in Artpedia, which are composed of long sentences, with lots of details, as can be seen in Fig. 3 and Fig. 4. Therefore, n -gram-based metrics fail to effectively convey the quality of the captions. On the contrary, BERTScore, which captures semantic similarity between sentences rather than analyzing them from a structural point of view, achieves much higher results and confirms the improved quality of the captions generated with our data augmentation. On the contrary, ArtCap has shorter sentences so metrics such as BLEU, METEOR, ROUGE and CIDEr manage to obtain much higher results in absolute terms and can be used effectively in this comparison.

In order to assess the quality of our data augmentation strategy we also present a comparison between our method and different state of the art methods for image data augmentation (from Section 5.1.2) in Tab. 2. It is important to note that while our augmentation approach is beneficial to the model, other augmentation techniques actually hurt performance. Intuitively we can infer that data augmentation strategies such as the one we compare our method against are engineered for classification tasks and might not be semantically invariant with reference to image captioning.

Image Retrieval For the retrieval task we test CLIP [40] using a similar setting to image captioning. We test first in a zero-shot configuration and then with and without data augmentation. The CLIP model is pre-trained on the YFCC dataset [50] when performing zero-shot retrieval. As in the

	No aug	AutoAugment	RandAugment	AugMix	TrivialAugment	Ours
Artpedia	0.0505	0.0583	0.0533	0.0536	0.0510	0.0673
ArtCap	0.6263	0.5829	0.6239	0.5717	0.5849	0.6445

Table 2: Comparison of CIDEr scores with GIT [54] model trained using our proposed diffusion augmentation and other state-of-the-art augmentation techniques for the image captioning task: AutoAugment [13], RandAugment [14], AugMix [23], and TrivialAugment [36].



Figure 6: Sample of images produced by different augmentation methods: No augmentation, AutoAugment [13], RandAugment [14], AugMix [23], and TrivialAugment [36], Ours.

previous task, the results shown in Tab. 3 present a clear indication of an improved performance in the retrieval problem. Tab. 4 instead compares our results with the best ones proposed in [47] using the same experimental protocol by the authors, i.e. by fixing the maximum number of retriev-

Model	Task	R@1	R@5	R@10
CLIP - (zero-shot)	im2t	0.0853	0.1557	0.2096
CLIP - w/o DA	im2t	0.1048	0.2081	0.2665
CLIP - w/ DA	im2t	0.1108	0.2096	0.2740
CLIP - (zero-shot)	t2im	0.0644	0.1751	0.2290
CLIP - w/o DA	t2im	0.0883	0.1751	0.2305
CLIP - w/ DA	t2im	0.0868	0.1976	0.2470

Table 3: Test on Artpedia on the retrieval task with CLIP using a ResNet50 pretrained on YFCC [50]. We report Recall @1, @5, @10. We test both in the image-to-text (im2t) setting and in the text-to-image (t2im) setting.

Model	Task	R@1	R@5
X-Attn GloVe [47]	im2t	0.086	0.227
CLIP - w/ DA	im2t	0.090	0.230
X-Attn GloVe [47]	t2im	0.041	0.136
CLIP - w/ DA	t2im	0.090	0.250

Table 4: We compare our results using the same experimental protocol as in [47] using $N = 100$ retrievable items. We report Recall @1 and @5, testing both in the image-to-text (im2t) setting and in the text-to-image (t2im) setting.

able items to $N = 100$.

5.2. Qualitative Results

Due to the subjective nature of the task, it is necessary to perform a visual inspection to better understand how are the models behaving under different data regime conditions. In this section, we present a sample of qualitative results to better appreciate the effect of our method. The example presented in Fig. 7 shows the effectiveness of our model in terms of enriching the quality of the description by comparing the output of the same captioning method trained under different settings.

While the pre-trained model tends to offer vague but correct descriptions even in a zero-shot setting, it is necessary to fine-tune the model on the target dataset in order to match the language used in Artpedia dataset. Our data-augmented finetuning helps the model to obtain a better representation of fine visual details in the dataset, allowing to obtain richer captions using the task-related technical knowledge that a large internet-wide trained model might be missing.



	GT	Madonna of parturition is an iconic depiction of the virgin mary, shown as pregnant, which was developed mainly in tuscanly in the 14th century. The madonna was portrayed standing, alone, often with a closed book on her belly, an allusion to the incarnate word...
	zero-shot	Painting of three people
	w/o DA	The painting depicts the virgin mary standing in front of a throne decorated with a red velvet curtain
	w/ DA	It shows the virgin mary holding the child in a blue dress. she is dressed in a blue robe, and holds a book in her left hand. she wears a blue robe, and holds in her right hand. she wears a blue dress with a red bodice decorated with gold embroidery on the front, and a red cloth over her head. she wears a blue dress, and a gold brocade with a gold embroidery on her lap
	GT	Housed in the museum of fine arts, boston, it shows luke the evangelist, patron saint of artists, sketching the virgin mary as she nurses the child jesus. the figures are positioned in a bourgeois interior which leads out towards a courtyard, river, town and landscape...
	zero-shot	A painting of a scene
	w/o DA	A little girl and a boy, with a heart in the background
	w/ DA	The painting depicts the virgin mary and the saint mary in the foreground, with a view of the landscape, and a woman holding a child

Figure 7: Qualitative samples showing the original (GT) caption along with different outputs from GIT [54] on Art-pedia [47]

6. Conclusions

This paper presented technique for augmenting and better exploit fine art datasets with the intent of making the fruition of semantically complex visual art easier to digitalise, to access, and to retrieve for the general public. In the field of cultural heritage a feature such as the uniqueness of the artworks can become an obstacle for machine learning techniques that requires large amount of data. At the same time the usual augmentation techniques such as image flipping, random brightness change, random hue change do not suit the task as they semantically change the original datapoint by changing small the visual details that are actually meaningful. Therefore our contributions aims at semantically enrich the popular pretrained LLMs models leveraging the expert knowledge to create a more sophisticated image data augmentation pipeline.

Acknowledgements

This work is partially supported by the European Commission under European Horizon 2020 Programme, Grant No. 101004545-ReInHerit.

References

- [1] Antreas Antoniou, Amos Storkey, and Harrison Edwards. Augmenting image classifiers using data augmentation generative adversarial networks. In *Proc. of Artificial Neural*

- Networks and Machine Learning (ICANN)*, pages 594–603, 2018. 3
- [2] Beatrice Aruanno, Sandro Barone, Francesco Tamburrino, and Mario Covarrubias Rodriguez. Tintoretto unveiled: Interactive virtual experience for artworks. In *International Joint Conference on Mechanics, Design Engineering & Advanced Manufacturing*, pages 1352–1363. Springer, 2022. 2
- [3] R. Barth, J. Hemming, and E.J. Van Henten. Optimising realism of synthetic images using cycle generative adversarial networks for improved part segmentation. *Computers and Electronics in Agriculture*, 173:105378, 2020. 3
- [4] Federico Becattini, Pietro Bongini, Luana Bulla, Alberto Del Bimbo, Ludovica Marinucci, Misael Mongiovi, and Valentina Presutti. Viscount: A large-scale multilingual visual question answering dataset for cultural heritage. *ACM Transactions on Multimedia Computing, Communications and Applications*, 2023. 3
- [5] Federico Becattini, Andrea Ferracani, Lea Landucci, Daniele Pezzatini, Tiberio Uricchio, and Alberto Del Bimbo. Imaging novecento. a mobile app for automatic recognition of artworks and transfer of artistic styles. In *Digital Heritage. Progress in Cultural Heritage: Documentation, Preservation, and Protection: 6th International Conference, EuroMed 2016, Nicosia, Cyprus, October 31–November 5, 2016, Proceedings, Part I 6*, pages 781–791. Springer, 2016. 2
- [6] Lorenzo Berlincioni, Federico Becattini, Lorenzo Seidenari, and Alberto Del Bimbo. Multiple future prediction leveraging synthetic trajectories. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 6081–6088. IEEE, 2021. 3
- [7] Pietro Bongini, Federico Becattini, Andrew D Bagdanov, and Alberto Del Bimbo. Visual question answering for cultural heritage. In *IOP Conference Series: Materials Science and Engineering*, volume 949, page 012074. IOP Publishing, 2020. 3
- [8] Pietro Bongini, Federico Becattini, and Alberto Del Bimbo. Is gpt-3 all you need for visual question answering in cultural heritage? In *European Conference on Computer Vision*, pages 268–281. Springer, 2022. 2
- [9] Christopher Bowles, Liang Chen, Ricardo Guerrero, Paul Bentley, Roger N. Gunn, Alexander Hammers, David Alexander Dickie, María Valdés Hernández, Joanna Marguerite Wardlaw, and Daniel Rueckert. Gan augmentation: Augmenting training data using generative adversarial networks. *ArXiv*, abs/1810.10863, 2018. 3
- [10] Valentina Anita Carriero, Aldo Gangemi, Maria Letizia Mancinelli, Ludovica Marinucci, Andrea Giovanni Nuzozolese, Valentina Presutti, and Chiara Veninata. Arco: The italian cultural heritage knowledge graph. In *The Semantic Web–ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II 18*, pages 36–52. Springer, 2019. 3
- [11] Eva Cetinic, Tomislav Lipic, and Sonja Grgic. Fine-tuning convolutional neural networks for fine art classification. *Expert Systems with Applications*, 114:107–118, 2018. 2
- [12] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick.

- Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 4
- [13] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 113–123, 2019. 6, 7
- [14] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 6, 7
- [15] Riccardo Del Chiaro and et al. Noisyart: A dataset for webly-supervised artwork recognition. In *VISIGRAPP (4: VISAPP)*, pages 467–475, 2019. 2
- [16] Riccardo Del Chiaro and et al. Webly-supervised zero-shot learning for artwork instance recognition. *Pattern Recognition Letters*, 128:420–426, 2019. 2, 3
- [17] Michael Denkowski and Alon Lavie. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*, 2014. 5
- [18] Chaitanya Devaguptapu, Ninad Akolekar, Manuj M Sharma, and Vineeth N Balasubramanian. Borrow from anywhere: Pseudo multi-modal object detection in thermal imagery. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR-W)*, 2019. 3
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 5
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [21] Noa Garcia, Chentao Ye, Zihua Liu, Qingtao Hu, Mayu Otani, Chenhui Chu, Yuta Nakashima, and Teruko Mitamura. A dataset and baselines for visual question answering on art. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 92–108. Springer, 2020. 3
- [22] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition? *arXiv preprint arXiv:2210.07574*, 2022. 3
- [23] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019. 6, 7
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. 3
- [25] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1989–1998. PMLR, 10–15 Jul 2018. 3
- [26] Sheng-Wei Huang, Che-Tsung Lin, Shu-Ping Chen, Yen-Yi Wu, Po-Hao Hsu, and Shang-Hong Lai. Auggan: Cross domain adaptation with gan-based data augmentation. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 718–731, 2018. 3
- [27] Hyeon-Ju Jeon, Soonchul Jung, Yoon-Seok Choi, Jae Woo Kim, and Jin Seo Kim. Object detection in artworks using data augmentation. In *2020 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 1312–1314. IEEE, 2020. 3
- [28] Xun Jin and Jongweon Kim. Artwork identification for 360-degree panoramic images using polyhedron-based rectilinear projection and keypoint shapes. *Applied Sciences*, 7(5):528, 2017. 2
- [29] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2, 5
- [30] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. 5
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 5
- [32] Yue Lu, Chao Guo, Xingyuan Dai, and Fei-Yue Wang. Artcap: A dataset for image captioning of fine art paintings. *IEEE Transactions on Computational Social Systems*, pages 1–12, 2022. 3, 4
- [33] Mahshid Majd and Reza Safabakhsh. Impact of machine learning on improvement of user experience in museums. In *2017 Artificial Intelligence and Signal Processing Conference (AISP)*, pages 195–200. IEEE, 2017. 2
- [34] Thomas Mensink and Jan Van Gemert. The rijksmuseum challenge: Museum-centered visual recognition. In *Proceedings of International Conference on Multimedia Retrieval*, pages 451–454, 2014. 2
- [35] Federico Milani and Piero Fraternali. A dataset and a convolutional model for iconography classification in paintings. *Journal on Computing and Cultural Heritage (JOCCH)*, 14(4):1–18, 2021. 2
- [36] Samuel G Müller and Frank Hutter. Trivialaugment: Tuning-free yet state-of-the-art data augmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 774–782, 2021. 6, 7
- [37] Samuel G. Müller and Frank Hutter. Trivialaugment: Tuning-free yet state-of-the-art data augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 774–782, October 2021. 6
- [38] Xinlei Pan, Yurong You, Ziyang Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. *arXiv preprint arXiv:1704.03952*, 2017. 3

- [39] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. 5
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 5, 6
- [41] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [42] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 4
- [43] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 4
- [44] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2107–2116, 2017. 3
- [45] Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. Textcaps: a dataset for image captioning with reading comprehension. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 742–758. Springer, 2020. 4
- [46] Matteo Stefanini, Marcella Cornia, Lorenzo Baraldi, Silvia Cascianelli, Giuseppe Fiameni, and Rita Cucchiara. From show to tell: A survey on deep learning-based image captioning. *IEEE transactions on pattern analysis and machine intelligence*, 45(1):539–559, 2022. 2
- [47] Matteo Stefanini, Marcella Cornia, Lorenzo Baraldi, Massimiliano Corsini, and Rita Cucchiara. Artpedia: A New Visual-Semantic Dataset with Visual and Contextual Sentences. In *Proceedings of the International Conference on Image Analysis and Processing*, 2019. 2, 3, 4, 6, 7, 8
- [48] Wei Ren Tan, Chee Seng Chan, Hernán E Aguirre, and Kiyoshi Tanaka. Ceci n’est pas une pipe: A deep convolutional network for fine-art paintings classification. In *2016 IEEE international conference on image processing (ICIP)*, pages 3703–3707. IEEE, 2016. 2
- [49] Frederik Temmermans, Bart Jansen, Rudi Deklerck, Peter Schelkens, and Jan Cornelis. The mobile museum guide: artwork recognition with eigenpaintings and surf. In *Proceedings of the 12th International Workshop on Image Analysis for Multimedia Interactive Services*, 2011. 2
- [50] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 6, 7
- [51] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective data augmentation with diffusion models, 2023. 3, 5
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [53] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 5
- [54] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022. 2, 5, 6, 7, 8
- [55] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019. 5
- [56] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017. 3