CLAP: Isolating Content from Style through Contrastive Learning with Augmented Prompts

Yichao Cai[®], Yuhang Liu[®], Zhen Zhang[®], and Javen Qinfeng Shi[®]

Australian Institute for Machine Learning, University of Adelaide, SA 5000, Australia {yichao.cai,yuhang.liu01,zhen.zhang02,javen.shi}@adelaide.edu.au

Abstract. Contrastive vision-language models, such as CLIP, have garnered considerable attention for various downsteam tasks, mainly due to the remarkable ability of the learned features for generalization. However, the features they learned often blend content and style information, which somewhat limits their generalization capabilities under distribution shifts. To address this limitation, we adopt a causal generative perspective for multimodal data and propose contrastive learning with data augmentation to disentangle content features from the original representations. To achieve this, we begin with exploring image augmentation techniques and develop a method to seamlessly integrate them into pretrained CLIP-like models to extract pure content features. Taking a step further, recognizing the inherent semantic richness and logical structure of text data, we explore the use of text augmentation to isolate latent content from style features. This enables CLIP-like model's encoders to concentrate on latent content information, refining the learned representations by pre-trained CLIP-like models. Our extensive experiments across diverse datasets demonstrate significant improvements in zeroshot and few-shot classification tasks, alongside enhanced robustness to various perturbations. These results underscore the effectiveness of our proposed methods in refining vision-language representations and advancing the state-of-the-art in multimodal learning.¹

Keywords: Data Augmentation · Latent Variables · Disentanglement

1 Introduction

Vision-language models, exemplified by CLIP [36], have garnered substantial attention due to their exceptional generalization capabilities, achieved through the learned features, obtained by utilizing a cross-modal contrastive loss [20, 25, 36]. However, despite being pre-trained on extensive datasets, CLIP-like models fall short in disentangling latent content information and latent style information. Consequently, they are not immune to spurious correlations, i.e., style-related information is erroneously utilized to predict task-related labels. These limitations become evident in the presence of distribution shifts or adversarial attacks where spurious correlations often change across different environments. For examples,

¹ Our code is available at https://github.com/YichaoCai1/CLAP.



Fig. 1: Causal generative models of vision-language data. Image and text data are generated through distinct underlying deterministic processes, $\mathbf{g}_{\mathbf{x}}$ for images and $\mathbf{g}_{\mathbf{t}}$ for texts, derived from a unified latent space with latent content variables \mathbf{c} and latent style variables \mathbf{s} . Latent content \mathbf{c} exclusively determines the sample label \mathbf{y} . (a) Soft interventions on latent style variables \mathbf{s} result in $\mathbf{\tilde{s}}$, subsequently generating augmented images $\mathbf{\tilde{x}}$. (b) Due to the same latent space, soft interventions on latent style variables \mathbf{s} can also result in $\mathbf{\tilde{s}}$, producing augmented text $\mathbf{\tilde{t}}$. (c) A qualitative comparison of image features for zero-shot classification using "a photo of a [class]" prompts, visualized using class activation map (CAM) [32], demonstrates that while image augmentation can enhance CLIP features, the features obtained through text augmentation methods predominantly focus on the content.

(1) a notable dependence on specific input text prompts has been reported for zero-shot capabilities [21, 47, 48]; (2) performance decline in few-shot scenarios has been observed in few-shot learning scenarios [13, 36]; and (3) susceptibility to adversarial attacks has been explored [33, 43, 45].

Taking a causal perspective, this work begin with a simple yet effective method, image augmentation, to disentangle content and style information within the learned representations of CLIP-like models. This approach is inspired by recent advancements in theoretical development in causal representation learning [41], which demonstrate that augmented image can be interpreted as a result of soft interventions on latent style variables, as depicted in Fig. 1a. Such augmentation results in a natural data pair where content information remains unchanged while style information changes. Consequently, using contrastive learning, it becomes feasible to isolate the invariant content information from the variant style information. Motivated by this theoretical advancement, we propose a practical method to incorporate image augmentation into CLIP-like models to extract content information from the original learned features. Specifically, a disentangled network is designed to fine-tune the pre-trained CLIP model by using a contrastive loss with image augmentation.

Despite the advancements made in disentangling content and style information from the original features learned by CLIP-like models through image augmentation, we recognize an inherent limitation: it is generally challenging to design adequate image augmentations to ensure all style factors change in an image. Theoretically, disentangling content and style information necessitates changes in all style factors [41]. However, inducing sufficient changes in latent style through image augmentation poses challenges due to the high dimensionality and complexity of style information in image data. Achieving significant style variation via artificially designed image augmentation techniques, such as transforming a photograph of a dog into a sketch while preserving content but dramatically altering style, is notably difficult.

Taking a further step, rather than relying on image augmentation, we explore the use of text augmentation to disentangle latent content and style factors. This shift is motivated by two key observations: 1) Vision and language data share the same latent space. Therefore, text augmentation can also be utilized to induce changes in latent style factors instead of image augmentation. 2) Text data inherently possesses high semanticity and logical structure, making it more amenable to property-wise manipulation compared to image data. Consequently, implementing sufficient style changes through text augmentation is more feasible than image augmentation, contributing to isolating content from style information, see Fig. 1c for visual comparison. For instance, transforming text from "a photo of a dog" to "a sketch of a dog" is straightforward in the language modality, whereas achieving a similar transformation in image data is challenging. Inspired by these observations, we posit that introducing style variations through text augmentation, as illustrated Fig. 1b, provides a more effective approach for learning vision-language content features than relying on image augmentation.

In summary, our contributions include: (1) Aimed at disentangling latent content and style factors to refine vision-language features of pre-trained CLIPlike models, we propose constrastive learning with data augmentation to fine tune the original features of pre-trained CLIP-like models from a causal perspective. (2) We present a novel method customized for pre-trained CLIP-like models. This method leverages a disentangled network, which is trained using contrastive learning with image augmentation, to extract latent content features from the learned features provided by image encoder of CLIP-like models. (3) We propose Contrastive Learning with Augmented Prompts (CLAP), to extract latent content features from representations of CLIP-like models. It begins by training a disentangled network using the pre-trained text encoder of CLIP-like models and text augmentation. Subsequently, the trained disentangled network is transferred to the image encoder of CLIP-like models. (4) Experiments conducted on a large real dataset demonstrate the effectiveness of the proposed image augmentation and text augmentation in terms of zero-shot and few-shot performance, as well as robustness against perturbations.

2 Related Work

Contrastive Vision-Language Models Using a cross-modal contrastive loss, CLIP [36] revolutionarily introduced a scalable contrastive vision-language model

by leveraging a large corpus of internet-sourced image-text pairs, demonstrating unprecedented zero-shot learning capabilities and exceptional generalization ability across datasets and supporting numerous downstream tasks [38]. ALIGN [20] expanded the scale of contrastive vision-language modeling, training on up to one billion image-text pairs while integrating the vision transformer's selfattention mechanism [11], which further enhanced performance. Despite their successes, CLIP-like models exhibit sensitivity to input text prompts [21, 48], leading to variable performance across different prompts. Efforts to mitigate this prompt sensitivity through prompt learning and engineering [9,14,21,47,48] focus on customizing prompts for specific tasks but do not fundamentally enhance CLIP's representations. Furthermore, CLIP-like models are vulnerable to adversarial attacks [4, 12], with current strategies [33, 45] involving adversarialnatural image pairs to improve resilience. Our work diverges from task-specific approaches by aiming to enhance CLIP's representations from a disentanglement perspective, addressing the aforementioned issues inherent in CLIP-like models.

Disentangled Representation Learning Aimed at segregating intrinsic latent factors in data into distinct, controllable representations, disentangled representation learning benefits various applications [24, 40, 44]. Specifically, in classification tasks, it's shown that enhancing the model's performance and robustness against data distribution perturbations can be achieved by more effectively disentangling invariant content variables, without needing to identify all intrinsic latent variables completely [22, 26–28]. Within single modalities, studies such as [49] have illustrated that contrastive learning [7, 16, 18] can potentially reverse the data generative process, aiding in the separation of representations. Furthermore, [41] suggest that image augmentation can help isolate content variables from the latent space through significant stylistic changes. [19] employs mixture techniques for data augmentation, enabling more abundant cross-modal matches. Diverging from these methods, our approach focuses on employing text augmentation to disentangle latent content variables, introducing a unique approach to learn refined vision-language representations.

3 A Causal Generative Model for Multi-Modal Data

To understand pretrained CLIP-like models, we investigate the underlying causal generative process for vision-language data. We consider the following causal generative model as depicted in Fig. 1. In the proposed model, the shared latent space ruling vision and language data is divided into two distinct sub-spaces: one corresponding to the latent content variables \mathbf{c} and the other to the latent style variables \mathbf{s} . The latent content variables are posited to determine the object label \mathbf{y} , a relationship corroborated by prior studies [22, 29, 31]. Furthermore, to elucidate the correlation between the latent style variable \mathbf{s} and the object variable \mathbf{y} , our model incorporates the premise that the latent content variable \mathbf{c} causally influences the latent style variable \mathbf{s} , in concordance with the principles of causal representation learning highlighted in recent literature [10, 29, 41]. Additionally,

considering the diversity between image data and text data, where information in image data is typically much more details while information in text data tends to be more logically structured nature, we posit distinct causal mechanisms for the generation processes. Our causal generative model is formulated as following structural causal models [2]:

$$\mathbf{s} := \mathbf{g}_{\mathbf{s}}(\mathbf{c}), \ \mathbf{x} := \mathbf{g}_{\mathbf{x}}(\mathbf{c}, \mathbf{s}), \ \mathbf{t} := \mathbf{g}_{\mathbf{t}}(\mathbf{c}, \mathbf{s}), \ \mathbf{y} := \mathbf{g}_{\mathbf{y}}(\mathbf{c}). \tag{1}$$

In Eq. (1), the style variable \mathbf{s} is causally influenced by the content via $\mathbf{g}_{\mathbf{s}}$; \mathbf{x} and \mathbf{t} denote visual and textual data, respectively. Both visual and textual data are causally produced by the shared latent variables \mathbf{c} and \mathbf{s} through distinct, reversible generative processes: $\mathbf{g}_{\mathbf{x}}$ for images and $\mathbf{g}_{\mathbf{t}}$ for text data, respectively. The label \mathbf{y} of a sample is exclusively determined by the content variable \mathbf{c} via $\mathbf{g}_{\mathbf{y}}$. For simplicity, exogenous noises are implicitly assumed but not explicitly represented in the causal generative model's formulation, aligning with the common understanding that each latent variable is influenced by exogenous noise.

Recent seminal work in [41] has demonstrated that the latent content variable **c** can be identified up to block identifiability (i.e., **c** can be isolated from style variable **s**), by requiring all latent style variables to change (e.g., soft interventions on all latent style variables). This change can be achieved through image augmentation, i.e., the augmented image $\tilde{\mathbf{x}}$ can be interpreted as a generative result of $\tilde{\mathbf{s}}$, which is produced through soft interventions on original latent style variables **s**. Despite such theoretical advancement, the practical implementation of this theoretical result within CLIP-like models remains unclear. In this study, we propose a practical method to disentangle content and style information within CLIP-like models by employing image augmentation, as detailed in Sec. 4.1. Moreover, we recognize that implementing sufficient changes on all latent style variables **s** through text augmentation is more feasible than image augmentation, due to high semanticity and logical structure in text data, we delve into the use of text augmentation to separate content information from style information, as discussed in Sec. 4.2.

4 Isolating Content from Style with Data Augmentation

In this section, we propose the employment of data augmentation to extract content information from the learned features in pre-trained CLIP-like models. Essentially, data augmentation facilitates the alteration of style factors while preserving content factors. Consequently, leveraging contrastive learning enables the segregation of content information from style information. We delve into two distinct forms of data augmentation, namely image augmentation (Sec. 4.1) and text augmentation (Sec. 4.2).

4.1 Isolating Content from Style with Augmented Images

While recent studies (von et al., 2021) have offered assurance regarding the disentanglement of content and style through contrastive learning with data augmentation, it remains unclear how these theoretical findings can be applied to



Fig. 2: Refining CLIP through data augmentation. (a) Training involves a disentangled network $\mathbf{f_c}$, utilizing contrastive loss on original and augmented image pairs \mathbf{x} and $\tilde{\mathbf{x}}$, with CLIP's image encoder $\mathbf{f_x^*}$ holding frozen gradients. (b) More efficient content feature learning is achieved through contrastive learning with augmented text prompts \mathbf{t} and $\tilde{\mathbf{t}}$, using the fixed text encoder $\mathbf{f_t^*}$ of CLIP. (c) Inference stage: The trained disentangled network $\mathbf{f_c^*}$ integrates with CLIP's text and image encoders, $\mathbf{f_t^*}$ and $\mathbf{f_x^*}$, to enable zero-shot inference for an input image \mathbf{x} and class names $\mathbf{t_1}$ to $\mathbf{t_n}$.

the realm of vision-language models. We convert the theoretical findings into CLIP-like models in the following. The theoretical findings suggest using InfoNCE loss [34] to extract content information, as outlined below:

$$\mathcal{L}(\mathbf{f}; \{\mathbf{x}_i, \tilde{\mathbf{x}}_i\}_{i=1}^b, \tau) = -\frac{1}{b} \sum_{i=1}^b \log \frac{\exp\left[\langle \mathbf{f}(\mathbf{x}_i), \mathbf{f}(\tilde{\mathbf{x}}_i) \rangle / \tau\right]}{\sum_{j=1}^b \exp\left[\langle \mathbf{f}(\mathbf{x}_i), \mathbf{f}(\tilde{\mathbf{x}}_j) \rangle / \tau\right]}, \qquad (2)$$

where $\{\mathbf{x}_i\}_{i=1}^b$ represents a batch of *b* samples from the training dataset, $\mathbf{f}(\mathbf{x}_i)$ denotes sample \mathbf{x}_i 's features through model \mathbf{f} , $\mathbf{\tilde{x}}_i$ is the augmented counterpart of \mathbf{x}_i , and $\langle \mathbf{z}_1, \mathbf{z}_2 \rangle$ represents the cosine similarity between two feature vectors, \mathbf{z}_1 and \mathbf{z}_2 , and τ represents the temperature parameter influencing the loss.

We extend it to refine pre-trained vision-language models, utilizing contrastive learning with augmented images (hereinafter referred to as "Im.Aug"). As illustrated in Fig. 2a, we train a disentangled network on top of CLIP's pretrained image encoder. To enhance training efficiency and the usability of the proposed method, we freeze the pre-trained image encoder. Based on an InfoNCE loss, the learning objective of Im.Aug is formulated as follows:

$$\mathbf{f}_{\mathbf{c}}^{*} = \underset{\mathbf{f}_{\mathbf{c}}}{\operatorname{argmin}} \underset{\{\mathbf{x}_{i}\}_{i=1}^{b} \in \mathcal{D}_{\mathbf{x}}}{\mathbb{E}} \mathcal{L}(\mathbf{f}_{\mathbf{c}} \circ \mathbf{f}_{\mathbf{x}}^{*}; \{\mathbf{x}_{i}, \tilde{\mathbf{x}}_{i}\}_{i=1}^{b}, \tau),$$
(3)

where $\mathcal{D}_{\mathbf{x}}$ denotes the training image dataset and *b* represents the batch size, $\mathbf{f}_{\mathbf{c}}$ is the disentangled network undergoing training. The pre-trained CLIP image encoder is represented by $\mathbf{f}_{\mathbf{x}}^*$, with the asterisk "*" signifying that the model weights remain fixed. The variable \mathbf{x}_i refers to an image sampled from $\mathcal{D}_{\mathbf{x}}$, and $\tilde{\mathbf{x}}_i$ is its augmented view.

The composition of the training dataset $\mathcal{D}_{\mathbf{x}}$, the image augmentation techniques used, the structure of the disentangled network $\mathbf{f}_{\mathbf{c}}$, and the utilization of $\mathbf{f}_{\mathbf{c}}^*$ post-training are detailed in the following subsections.

Data Synthesis and Image Augmentation To generate training image data, we combine class names with various image and object attributes to create text

Table 1: Template-based prompts. Attributes used to generate text prompts follow the structured format "a [art style] [image type] of a [object size] [object color] [class]", where "[class]" represents the class names.

Object Color	Object Size	Image Type	Art Style
yellow, green, black, blue, multicolored, orange, red, white, brown, purple	large, small, normal sized	painting, cartoon, infograph, sketch, photograph, clipart, mosaic art, sculpture	realistic, impressionistic

prompts for each class. Using a stable diffusion model [39], we produce synthetic images that comprise our training dataset $\mathcal{D}\mathbf{x}$. The creation of template prompts for stable diffusion is based on attributes such as object size, color, image type, and art style. As detailed in Tab. 1, the attributes include 10 colors and 3 sizes for objects, and 8 types and 2 art styles for images. By assembling these attributes into prompts like "a [art style] [image type] of a [object size] [object color] [class]", we generate 480 unique texts for each class, from which one image per prompt is synthesized. Further details on image synthesis and examples are available in Appendix B.1. For the image augmentation procedures, we adopt techniques commonly used in contrastive learning practice [7, 8, 41], specifically random cropping and color distortion.

Disentangled Network Structure Since the training process is based on CLIP's pre-trained lower-dimensional features, our disentangled network adopts a multi-layer perceptron (MLP) architecture. To fully benefit from the pre-trained CLIP text encoder, we construct a residual MLP featuring a zero-initialized projection, acting as the disentangled network, as depicted in Fig. 3. This design enables learning directly from the pre-trained representation space, avoiding a random starting point, inspired by ControlNet's zero-conv operation [46], which we adapt to a zero-linear operation within our residual MLP.

Within this architecture, the main branch includes a zero-initialized, biasfree linear layer positioned subsequent to the combination of a SiLU activation and a normally initialized linear layer. Conventionally, the dimensions of features before the initial linear layer, situated between the first and second linear layers, and following the second linear layer, are named as the input d_{in} , latent d_{mid} , and output d_{out} dimensions, respectively. To rectify any mismatches between the input and output dimensions, the network employs nearest-neighbor downsampling within the shortcut path, thereby ensuring both alignment and the preservation of sharpness for the input features. During the inference stage, a weighting parameter $\alpha > 0$ is introduced to modulate the portion of features emanating from the main branch before their integration with the input features, whereas this parameter remains constant at 1 throughout the training phase.

Inference After training, the disentangled network $\mathbf{f}_{\mathbf{c}}^*$ is utilized following CLIP's image encoder to extract visual content features. Moreover, given that



Fig. 3: Structure of the disentangled network. The architecture encompass a residual block featuring a zero-initialized, bias-free linear layer to commence optimization from the input feature space. When the input and output dimension differ, a downsampling operation is utilized to achieve alignment. During inference, a scalar parameter α balance the main branch and input features before combination.

vision-language data generation is rooted in a unified latent space, as depicted in Sec. 3, $\mathbf{f}_{\mathbf{c}}^*$ can be seamlessly integrated with CLIP's image and text encoders to enhance zero-shot capabilities. As shown in Fig. 2c, for an image \mathbf{x} , the operation is formulated as the composition function $\mathbf{f}_{\mathbf{c}}^* \circ \mathbf{f}_{\mathbf{x}}^*(\mathbf{x})$, and similarly, for a text \mathbf{t} , as $\mathbf{f}_{\mathbf{c}}^* \circ \mathbf{f}_{\mathbf{t}}^*(\mathbf{t})$. This integration preserves CLIP's zero-shot functionality while achieving refined features through the improved disentanglement of content.

4.2 Isolating Content from Style with Augmented Prompts

Despite progress in disentangling content and style via image augmentation, adequately altering all style factors in an image remains challenging due to the high dimensionality and complexity of style information in images. Achieving substantial style changes through augmentation, essential for complete disentanglement [41], is difficult with existing image augmentation techniques. On the contrary, text data inherently possesses high semanticity and logical structure, making it more amenable to property-wise manipulation compared to image data. To further exploring the disentanglement of content, we propose Contrastive Learning with Augmented Prompts (**CLAP**).

As depicted in Fig. 2b, CLAP employs an InfoNCE loss to train a disentangled network atop CLIP's pre-trained text encoder, keeping the encoder's gradients fixed, similar to Im.Aug. Leveraging the simpler structure of text, the template-based prompts previously utilized for synthesizing images now serve as the training text dataset, denoted by \mathcal{D}_t . Utilizing the same disentangled network as in Im.Aug, the learning objective of CLAP is outlined as follows:

$$\mathbf{f}_{\mathbf{c}}^{*} = \underset{\mathbf{f}_{\mathbf{c}}}{\operatorname{argmin}} \underset{\{\mathbf{t}_{i}\}_{i=1}^{b} \in \mathcal{D}_{\mathbf{t}}}{\mathbb{E}} \mathcal{L}(\mathbf{f}_{\mathbf{c}} \circ \mathbf{f}_{\mathbf{t}}^{*}; \{\mathbf{t}_{i}, \tilde{\mathbf{t}}_{i}\}_{i=1}^{b}, \tau) + \lambda \mathcal{L}(\mathbf{f}_{\mathbf{c}} \circ \mathbf{f}_{\mathbf{t}}^{*}; \{\mathbf{t}_{i}^{c}, \tilde{\mathbf{t}}_{i}\}_{i=1}^{b}, 1), \quad (4)$$

where $\mathbf{f}_{\mathbf{t}}^*$ denotes the pre-trained CLIP text encoder. The term \mathbf{t}_i references a text prompt from $\mathcal{D}\mathbf{t}$, and $\tilde{\mathbf{t}}_i$ represents its augmented view, produced via prompt augmentation techniques. On the equation's right side, \mathbf{t}_i^c specifies the class name associated with the text prompt \mathbf{t}_i . This strategy aims to enhance variations between prompt pairs, especially in cases where the text dataset $\mathcal{D}_{\mathbf{t}}$

Table 2: Prompt augmentation techniques. Various augmented views are generated from an original text prompt using specific augmentation techniques: OSD (Object Size Deletion), OCD (Object Color Deletion), ITD (Image Type Deletion), ASD (Art Style Deletion), and SPO (Swapping Prompt Order).

Original	OSD OCD		ITD	ASD	SPO		
a realistic	a realistic	a realistic	a realistic	a painting	a large red		
painting of a	painting of a	painting of a	of a	of a	car in a		
large red car	red car	large car	large red car	large red car	realistic painting		

has a very limited number of samples. Here, λ serves for adjusting the second term's importance in the total loss function. All other symbols in Eq. (4) maintain their definitions as described earlier.

After training, the learned disentangled network is seamlessly integrated with both of CLIP's encoders to extract content representations, as depicted in Fig. 2c.

Prompt Augmentation To ensure text prompts undergo stylistic changes without compromising their content, we have developed specific augmentation techniques for synthetic text prompts. Drawing inspiration from Easy Data Augmentation (EDA) techniques [42], we adapted the Random Deletion (RD) and Random Swap (RS) techniques from EDA, customizing them to suit our prompt structure. To avoid inadvertently altering the content by introducing new object names or changing the core idea of a text prompt, our augmentation methods do not include random word insertions or replacements. Our primary augmentation techniques are Object Size Deletion (OSD), Object Color Deletion (OCD), Image Type Deletion (ITD), Art Style Deletion (ASD), and Swapping Prompt Order (SPO), each applied with a certain probability, as detailed in Tab. 2. Additionally, for down-stream datasets with few categories, to rich the population of training samples, we use an additional augmentation, named IGN (Inserting Gaussian Noise). Following the initializing protocol of prompt learning methods [47,48], we insert a zero-mean Gaussian noise with 0.02 standard deviation with a noise length equals to 4, to the tokenized prompts.

Intuitively, these prompt augmentation methods parallel random masking techniques used in image augmentation [6,17]. However, prompt augmentations are more effective and precise than their image counterparts. This effectiveness arises because prompt augmentations can specifically target and eliminate a particular style element without impacting the content, whereas image masking, operating at the pixel or patch level, might inadvertently damage content information or lead to insufficient style changes.

5 Experiments

We conduct three primary experiments to assess our method: (1) zero-shot evaluation with diverse prompts to gauge zero-shot performance and its robustness

to prompt perturbations; (2) linear probe tests on few-shot samples to evaluate the efficacy of the learned representations in few-shot settings; and (3) adversarial attack assessments on zero-shot and one-shot classifiers to determine their resistance to adversarial threats. We further conduct an ablative study on hyperparameters, explore the impact of different prompt augmentation combinations and various sources of training prompts on CLAP's performance, and replicate experiments across different CLIP model sizes.

5.1 Experimental Setup

Implementation. Im.Aug and CLAP are implemented using the ViT-B/16 CLIP model and executed on an NVIDIA RTX 3090 GPU. To ensure reproducibility, the random seed for all stochastic processes is fixed at 2024. More information on implementation details is provided in Appendix A.1.

Datasets. CLAP is assessed across four multi-domain datasets to examine its performance in varied environments: PACS [23] (4 domains, 7 categories), VLCS [1] (4 domains, 5 categories), OfficeHome [37] (4 domains, 65 categories), and DomainNet [35] (6 domains, 345 categories). For conciseness, we present average results across the domains for each dataset. Detailed experimental outcomes for each domain within these datasets are provided in Appendix A.4.

Compute efficiency. CLAP demonstrates faster convergence and shorter training times compared to Im.Aug. For CLAP, training on the PACS and VLCS datasets is completed in roughly 11 minutes, OfficeHome in approximately 14 minutes, and DomainNet in about 47 minutes. In contrast, Im.Aug requires around 16 minutes for PACS and VLCS, 50 minutes for OfficeHome, and 3.3 hours for DomainNet. Both Im.Aug and CLAP maintain CLIP's inference efficiency due to the disentangled network's efficient two-layer MLP structure.

5.2 Main Results

Zero-Shot Performance To assess zero-shot capabilities, CLAP undergoes evaluation using three specific fixed prompts: ZS(C), utilizing only the class name within "[class]"; ZS(PC), with the format "a photo of a [class]"; and ZS(CP), structured as "a [class] in a photo". To thoroughly examine zero-shot proficiency, a dynamic prompt, ZS(NC), formatted as "[noise][class]", is also used, where "[noise]" signifies the introduction of Gaussian noise characterized by a mean of 0 and a standard deviation of 0.02.

As presented in Tab. 3, CLAP surpasses both CLIP and Im.Aug across all evaluated prompts for every dataset. Unlike the uniform enhancement in zeroshot performance CLAP achieves over CLIP, Im.Aug displays inconsistent results. A closer examination reveals CLAP's superiority over CLIP is especially significant for the dynamic ZS(NC) prompt. This demonstrates CLAP's effectiveness in significantly improving zero-shot performance compared to the original CLIP representations.

In assessing the model's robustness to prompt perturbations, we examine the variances in zero-shot performance across different prompts by analyzing the

Prompt	Method	Zero-shot performance, avg. top-1 acc. (%) (\uparrow)									
1 iompt Mictilda		PACS	VLCS	Off.Home	Dom.Net	Overall					
ZS(C)	CLIP	95.7	76.4	79.8	57.8	77.4					
	Im.Aug	96.5	79.5	77.0	51.5	76.1					
	CLAP	97.2	82.6	81.0	58.7	79.9					
ZS(CP)	CLIP	95.2	82.0	79.5	57.0	78.4					
	Im.Aug	96.3	82.9	75.8	50.7	76.4					
	CLAP	97.3	83.4	80.5	58.0	79.8					
ZS(PC)	CLIP	96.1	82.4	82.5	57.7	79.7					
	Im.Aug	96.5	83.0	78.6	51.6	77.4					
	CLAP	97.2	83.4	83.0	59.0	80.6					
ZS(NC)	CLIP	90.8	68.3	71.5	51.0	70.4					
	Im.Aug	94.8	73.1	67.5	44.0	69.9					
	CLAP	97.2	81.0	73.5	52.6	76.1					

Table 3: Zero-shot results across three distinct prompts: "C" for "[class]", "CP" for "a [class] in a photo", "PC" for "a photo of a [class]", and a dynamic prompt "NC" for "[noise][class]" showcase that CLAP consistently outperforms CLIP's zero-shot performance across all datasets, whereas image augmentation exhibits mixed outcomes.

range (*R*) and standard deviation (δ) of results derived from ZS(C), ZS(CP), and ZS(PC). Additionally, we investigate the decline ($\Delta_{(NC)}$) in performance from ZS(C) to ZS(NC) as a broad indicator of resilience to noised prompts.

As presented in Tab. 4, CLAP significantly reduces the variance in zero-shot performance across various testing prompts, evidenced by markedly lower values of δ and R, and a less pronounced decrease in performance with a noised prompt, in contrast to Im.Aug and the baseline representations of CLIP. Although Im.Aug aids in reducing performance variance to some extent, its efficacy is notably inferior to that of CLAP. These findings highlight CLAP's enhanced robustness in maintaining consistent zero-shot performance across a diverse array of prompts.

Few-Shot Performance We conduct evaluations of 1-shot, 4-shot, 8-shot, and 16-shot linear probes across each domain within the four datasets. As illustrated in Fig. 4, CLAP significantly outperforms both CLIP and Im.Aug in few-shot learning scenarios. Notably, in the 1-shot setting CLAP achieves performance improvements over the linear-probe CLIP model by margins of +10%, +3.5%, +2.5%, and +1.5% on the PACS, VLCS, OfficeHome, and DomainNet datasets, respectively. These improvements are especially significant in comparison to the gains observed with Im.Aug counterparts, underpinning CLAP's efficacy in few-shot scenarios. For detailed quantitative results, please refer to Appendix A.4.

Adversarial Performance To assess adversarial robustness, zero-shot (ZS(C)) and one-shot classifiers are evaluated against prominent adversarial attack methods, such as FGSM [15], PGD [30], and CW [5], by generating adversarial sam-

Table 4: CLAP more effectively reduces zero-shot performance variance across prompts than image augmentation, with R and δ indicating the range and standard deviation for ZS(C), ZS(CP), and ZS(PC). The decrease $\Delta_{(NC)}$ from ZS(C) to ZS(NC) highlights CLAP's enhanced robustness against prompt perturbations.

Metri	c Method	Performance variance, avg. top-1 acc. (%) (\downarrow)								
		PACS	VLCS	Off.Home	Dom.Net	Overall				
	CLIP	0.9	6.1	3.1	0.8	2.7				
R	Im.Aug	0.1	3.6	2.8	0.9	1.9				
	CLAP	0.1	0.8	2.5	1.0	1.1				
	CLIP	0.4	2.8	1.4	0.4	1.2				
δ	Im.Aug	0.1	1.7	1.2	0.4	0.8				
	CLAP	0.0	0.4	1.1	0.4	0.5				
	CLIP	4.9	8.1	8.3	6.8	7.0				
$\Delta_{(NC)}$	Im.Aug	1.6	6.4	9.5	7.5	6.3				
,	CLAP	0.0	1.6	7.5	6.1	3.8				



Fig. 4: Few-shot linear probe comparisons of image-encoder features show that CLAP enhances CLIP's few-shot performance more effectively than Im.Aug. In the accompanying figure, "ZS" indicates the zero-shot performance using a "[class]" prompt.

ples for testing. For FGSM, 1 adversarial iteration is employed, whereas for PGD and CW, 20 iterations are used, all with an epsilon of 0.031. As indicated in Tab. 5, classifiers utilizing CLAP representations demonstrate superior resilience to these adversarial attacks compared to those based on CLIP representations. Across the four datasets, CLAP's zero-shot and 1-shot classifiers surpass CLIP by margins of +7.6% and +8.5% against FGSM, +1.0% and +11.7% against PGD-20, and +1.1% and +2.3% against CW-20, respectively. These figures notably exceed the performance improvements of +4.4% and +4.6% against FGSM, +0.3% and +6.2% against PGD-20, and 0% and +1.3% against CW-20 achieved by Im.Aug. The result suggests that CLAP efficiently enhances robustness against adversarial attacks in both zero-shot and one-shot scenarios.

5.3 More Analysis

t-SNE Visualization In our t-SNE visualizations, we examine the representations of CLIP, Im.Aug, and CLAP for all images within the Art Painting domain of the PACS dataset. Fig. 5 shows that CLAP's image representations display a marked inter-class separation and tighter intra-class clustering than those of

Table 5: Image augmentation and CLAP both enhance CLIP's zero-shot with the "[class]" prompt and 1-shot robustness against adversarial attacks, with CLAP showing greater improvements.

		Avg. top-1 acc. (%) under adversarial attacks (\uparrow)											
Setting Method		FGSM			PGD-20			CW-20				A	
	PACS	VLCS	O.H.	D.N.	PACS	VLCS	O.H.	D.N.	PACS	VLCS	O.H.	D.N.	Avg.
CLIP	86.8	65.6	57.9	22.5	29.1	2.0	10.1	10.7	27.4	1.5	7.4	7.6	29.2
ZS(C) Im.Aug	88.0	69.6	55.1	37.9	31.3	2.1	10.4	9.0	29.4	1.7	7.0	5.8	31.1
CLAP	88.7	71.9	58.5	44.2	30.8	3.2	10.6	11.2	29.8	2.3	8.1	8.0	32.7
CLIP	66.7	45.2	34.3	22.5	34.8	16.0	5.6	11.3	18.9	0.7	4.5	3.2	23.7
1-shot Im.Aug	79.4	47.1	37.1	23.5	55.2	16.1	8.5	12.5	23.2	0.9	5.1	3.4	28.0
CLAP	89.6	52.2	37.1	23.9	73.4	21.2	7.4	12.5	27.0	1.1	5.0	3.5	31.9



Fig. 5: t-SNE visualizations of all images in the Art Painting of PACS dataset show CLAP outperforms the original CLIP and Im.Aug, with clearer inter-class distinctions and tighter intra-class clusters.

CLIP and Im.Aug. This observation suggests that CLAP's representations are more closely tied to content information and less influenced by style information, in contrast to the other two.

Ablations In Fig. 6, we assess the zero-shot capabilities of our model using two distinct prompts, ZS(C) and ZS(PC), on the VLCS dataset. This analysis forms part of an ablative study aimed at understanding the influence of various hyperparameters on model performance. Specifically, we examine: the dimensions of the latent layer within the MLP of the disentangled network, as illustrated in Fig. 6a; the temperature parameter (τ) in the loss function, as depicted in Fig. 6b; and the weight coefficient (α) during the inference stage, as shown in Fig. 6c. Our findings indicate that CLAP consistently enhances zero-shot performance across all tested configurations for both prompts, while also significantly reducing the gap between the performances elicited by each prompt. These results underscore the efficacy of CLAP in accommodating a wide range of hyper-parameters.

Prompt Augmentation Combinations We explore diverse combinations of our tailored prompt augmentation methods and examine Easy Data Augmentation (EDA) techniques [42] on the VLCS dataset. Each tested technique showcases CLAP's enhancements over CLIP, with details available in Appendix A.2.



Fig. 6: We conduct ablative study on hyper-parameter choices on the VLCS dataset, including latent dimensions, τ values, and α values during the inference stage. CLAP continuously enhance CLIP's performance throughout the tested values.

Prompt Sources We assess the impact of different training prompt formats, originating from various synthetic sources, on the performance of the VLCS dataset, incorporating EDA techniques. Our evaluation includes our template-based prompts, LLM-generated prompts by ChatGPT-3.5 [3] (with the generation process detailed in Appendix B.2), prompts structured as "a [random] style of [class]," where "[random]" is filled with terms from a random word generator², and prompts produced using the PromptStyler method [9]. The findings indicate that the training prompts with simpler forms tend to yield better performance, with detailed quantitative results presented in Appendix A.3.

Experiments on Different Model Scales In our repeated experiments assessing zero-shot performance on the ViT-L/14 and ResNet50x16 pre-trained with CLIP, we consistently find that CLAP improves zero-shot performance while also reducing performance variances. This consistent observation underscores CLAP's effectiveness in enhancing the quality of CLIP representations. For quantitative details supporting these findings, please see the Appendix C.

6 Conclusion

To enhance pre-trained CLIP-like models, this study delves into disentangling latent content variables. Through a causal analysis of the underlying generative processes of vision-language data, we discover that training a disentangled network in one modality can effectively disentangle content across both modalities. Given the high semantic nature of text data, we identify that disentanglement is more achievable within the language modality through text augmentation interventions. Building on these insights, we introduce CLAP (Contrastive Learning with Augmented Prompts) to acquire disentangled vision-language content features. Comprehensive experiments validate CLAP's effectiveness, demonstrating significant improvements in zero-shot and few-shot performance, and enhancing robustness against perturbations. We anticipate that our work will inspire further exploration into disentangling latent variables within vision-language models.

 $^{^2}$ https://github.com/vaibhavsingh97/random-word

References

- Albuquerque, I., Monteiro, J., Darvishi, M., Falk, T.H., Mitliagkas, I.: Generalizing to unseen domains via distribution matching. arXiv preprint arXiv:1911.00804 (2019). https://doi.org/10.48550/arXiv.1911.00804
- Bollen, K.A.: Structural equations with latent variables, vol. 210. John Wiley & Sons (1989). https://doi.org/10.1002/9781118619179
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. Advances in neural information processing systems 33, 1877–1901 (2020), https://dl.acm.org/doi/abs/10.5555/3495724.3495883
- Carlini, N., Terzis, A.: Poisoning and backdooring contrastive learning. In: International Conference on Learning Representations (2021), https://openreview. net/forum?id=iC4UHbQ01Mp
- Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP). pp. 39–57. IEEE Computer Society (2017). https://doi.org/10.1109/SP.2017.49
- Chen, P., Liu, S., Zhao, H., Jia, J.: Gridmask data augmentation. arXiv preprint arXiv:2001.04086 (2020). https://doi.org/10.48550/arXiv.2001.04086
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020), https://dl.acm.org/doi/abs/10.5555/ 3524938.3525087
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 15750–15758 (2021). https://doi.org/10.1109/CVPR46437.2021.01549
- Cho, J., Nam, G., Kim, S., Yang, H., Kwak, S.: Promptstyler: Prompt-driven style generation for source-free domain generalization. In: Proceedings of the IEEE international conference on computer vision. pp. 15702–15712 (2023). https: //doi.org/10.1109/ICCV51070.2023.01439
- Daunhawer, I., Bizeul, A., Palumbo, E., Marx, A., Vogt, J.E.: Identifiability results for multimodal contrastive learning. ICLR (2023), https://openreview.net/ forum?id=U_2kuqoTcB
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020), https://openreview.net/forum?id= YicbFdNTTy
- Fort, S.: Adversarial vulnerability of powerful near out-of-distribution detection. arXiv preprint arXiv:2201.07012 (2022). https://doi.org/10.48550/arXiv. 2201.07012
- Gao, P., Geng, S., Zhang, R., Ma, T., Fang, R., Zhang, Y., Li, H., Qiao, Y.: Clip-adapter: Better vision-language models with feature adapters. International Journal of Computer Vision 132(2), 581–595 (2024). https://doi.org/10.1007/ s11263-023-01891-x
- Ge, C., Huang, R., Xie, M., Lai, Z., Song, S., Li, S., Huang, G.: Domain adaptation via prompt learning. arXiv preprint arXiv:2202.06687 (2022). https://doi.org/ 10.48550/arXiv.2202.06687
- Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: International Conference on Learning Representations (2015). https: //doi.org/10.48550/arXiv.1412.6572

- 16 Y. Cai et al.
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent-a new approach to self-supervised learning. Advances in neural information processing systems 33, 21271–21284 (2020), https://dl.acm.org/doi/abs/10. 5555/3495724.3497510
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022). https://doi.org/10. 1109/CVPR52688.2022.01553
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020). https://doi. org/10.1109/CVPR42600.2020.00975
- Hong, T., Guo, X., Ma, J.: Itmix: Image-text mix augmentation for transferring clip to image classification. In: 2022 16th IEEE International Conference on Signal Processing (ICSP). vol. 1, pp. 129–133. IEEE (2022). https://doi.org/10.1109/ ICSP56322.2022.9965292
- Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International conference on machine learning. pp. 4904-4916. PMLR (2021), https://proceedings.mlr.press/v139/jia21b.html
- Khattak, M.U., Rasheed, H., Maaz, M., Khan, S., Khan, F.S.: Maple: Multi-modal prompt learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19113–19122 (2023). https://doi.org/10. 1109/CVPR52729.2023.01832
- 22. Kong, L., Xie, S., Yao, W., Zheng, Y., Chen, G., Stojanov, P., Akinwande, V., Zhang, K.: Partial disentanglement for domain adaptation. In: International Conference on Machine Learning. pp. 11455–11472. PMLR (2022), https://proceedings.mlr.press/v162/kong22a.html
- Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: Proceedings of the IEEE international conference on computer vision. pp. 5542–5550 (2017). https://doi.org/10.1109/ICCV.2017.591
- Li, H., Wang, X., Zhang, Z., Yuan, Z., Li, H., Zhu, W.: Disentangled contrastive learning on graphs. Advances in Neural Information Processing Systems 34, 21872– 21884 (2021), https://dl.acm.org/doi/10.5555/3540261.3541935
- Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., Yu, F., Yan, J.: Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In: International Conference on Learning Representations (2021)
- Liu, Y., Zhang, Z., Gong, D., Gong, M., Huang, B., van den Hengel, A., Zhang, K., Shi, J.Q.: Identifiable latent polynomial causal models through the lens of change. In: The Twelfth International Conference on Learning Representations (2024), https://openreview.net/forum?id=ia9fK01Vjq
- Liu, Y., Zhang, Z., Gong, D., Gong, M., Huang, B., Hengel, A.v.d., Zhang, K., Shi, J.Q.: Identifying weight-variant latent causal models. arXiv preprint arXiv:2208.14153 (2022). https://doi.org/10.48550/arXiv.2208.14153
- Liu, Y., Zhang, Z., Gong, D., Gong, M., Huang, B., Hengel, A.v.d., Zhang, K., Shi, J.Q.: Identifiable latent neural causal models. arXiv preprint arXiv:2403.15711 (2024). https://doi.org/10.48550/arXiv.2403.15711
- Liu, Y., Zhang, Z., Gong, D., Gong, M., Huang, B., Zhang, K., Shi, J.Q.: Identifying latent causal content for multi-source domain adaptation. arXiv preprint arXiv:2208.14161 (2022). https://doi.org/10.48550/arXiv.2208.14161

- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: International Conference on Learning Representations (2018), https://openreview.net/forum?id=rJzIBfZAb
- Mahajan, D., Tople, S., Sharma, A.: Domain generalization using causal matching. In: International Conference on Machine Learning. pp. 7313-7324. PMLR (2021), https://proceedings.mlr.press/v139/mahajan21b.html
- Mamooler, S.: Clip explainability. https://github.com/sMamooler/CLIP_ Explainability, accessed: 2024-03-06
- Mao, C., Geng, S., Yang, J., Wang, X., Vondrick, C.: Understanding zero-shot adversarial robustness for large-scale models. In: The Eleventh International Conference on Learning Representations (2022), https://openreview.net/forum?id= P4bXCawRi5J
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018). https://doi.org/10.48550/ arXiv.1807.03748
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: Proceedings of the IEEE international conference on computer vision. pp. 1406–1415 (2019). https://doi.org/10.1109/ ICCV.2019.00149
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748-8763. PMLR (2021), https://proceedings.mlr.press/v139/radford21a. html
- Rahman, M.M., Fookes, C., Baktashmotlagh, M., Sridharan, S.: Multi-component image translation for deep domain generalization. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 579–588. IEEE (2019). https://doi.org/10.1109/WACV.2019.00067
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821-8831. PMLR (2021), https://proceedings.mlr. press/v139/ramesh21a.html
- 39. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022). https://doi.org/10.1109/CVPR52729.2023.01389
- Sanchez, E.H., Serrurier, M., Ortner, M.: Learning disentangled representations via mutual information estimation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16. pp. 205–221. Springer (2020). https://doi.org/10.1007/978-3-030-58542-6_13
- Von Kügelgen, J., Sharma, Y., Gresele, L., Brendel, W., Schölkopf, B., Besserve, M., Locatello, F.: Self-supervised learning with data augmentations provably isolates content from style. Advances in neural information processing systems 34, 16451–16467 (2021), https://dl.acm.org/doi/10.5555/3540261.3541519
- 42. Wei, J., Zou, K.: Eda: Easy data augmentation techniques for boosting performance on text classification tasks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). pp. 6382–6388 (2019). https://doi.org/10.18653/v1/D19-1670

- 18 Y. Cai et al.
- Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., et al.: Robust fine-tuning of zero-shot models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7959-7971 (2022). https://doi.org/10.1109/ CVPR52688.2022.00780
- 44. Yang, M., Liu, F., Chen, Z., Shen, X., Hao, J., Wang, J.: Causalvae: Disentangled representation learning via neural structural causal models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9593–9602 (2021). https://doi.org/10.1109/CVPR46437.2021.00947
- 45. Yang, W., Mirzasoleiman, B.: Robust contrastive language-image pretraining against adversarial attacks. arXiv preprint arXiv:2303.06854 (2023). https://doi. org/10.48550/arXiv.2303.06854
- Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: IEEE International Conference on Computer Vision (ICCV) (2023). https://doi.org/10.1109/ICCV51070.2023.00355
- Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Conditional prompt learning for visionlanguage models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16816–16825 (2022). https://doi.org/10. 1109/CVPR52688.2022.01631
- Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. International Journal of Computer Vision 130(9), 2337-2348 (2022). https://doi.org/10.1007/s11263-022-01653-1
- Zimmermann, R.S., Sharma, Y., Schneider, S., Bethge, M., Brendel, W.: Contrastive learning inverts the data generating process. In: International Conference on Machine Learning. pp. 12979-12990. PMLR (2021), https://proceedings. mlr.press/v139/zimmermann21a.html