

Text-to-Image Diffusion Models are Great Sketch-Photo Matchmakers

Subhadeep Koley^{1,2} Ayan Kumar Bhunia¹ Aneeshan Sain¹ Pinaki Nath Chowdhury¹
Tao Xiang^{1,2} Yi-Zhe Song^{1,2}

¹SketchX, CVSSP, University of Surrey, United Kingdom.

²iFlyTek-Surrey Joint Research Centre on Artificial Intelligence.

{s.koley, a.bhunia, a.sain, p.chowdhury, t.xiang, y.song}@surrey.ac.uk

<https://subhadeepkoley.github.io/DiffusionZSSBIR>




Figure 1. Sketch-based image retrieval frameworks [12, 102, 103] usually employ ImageNet pre-trained CNNs [3, 15, 103], JFT-trained vision transformers (ViT) [67, 78], or visual encoders of vision-language models like CLIP [77] as *backbone feature extractors*. Rich knowledge from large-scale pre-training offers a good initialisation, which when further fine-tuned on sketch-photo datasets, performs way better than training from random initialisation [62]. While one can extract features either by discarding the classification head for ImageNet pre-trained models, auxiliary task head for self-supervised models, or by using CLIP’s visual encoder, text-to-image diffusion models (e.g., stable diffusion) lack any specific feature embedding space. However, we find that its intermediate representations implicitly hold robust cross-modal features at multiple granularities. Unlike prior SBIR backbones, pre-trained with *discriminative tasks*, we propose to leverage denoising diffusion models pre-trained with text-to-image *generative tasks* to bridge the sketch-photo domain gap. Being a text-to-image generation model trained on a large corpus of text-image pairs (LAION [82, 83]), it holds both semantic and shape prior [89]. PCA representation [89] (right) of intermediate UNet features (sketch/photo) from different upsampling blocks (details in Sec. 5) depict that they share significant semantic similarity (denoted by similar colours).

Abstract

This paper, for the first time, explores text-to-image diffusion models for Zero-Shot Sketch-based Image Retrieval (ZS-SBIR). We highlight a pivotal discovery: the capacity of text-to-image diffusion models to seamlessly bridge the gap between sketches and photos. This proficiency is underpinned by their robust cross-modal capabilities and shape bias, findings that are substantiated through our pilot studies. In order to harness pre-trained diffusion models effectively, we introduce a straightforward yet powerful strategy focused on two key aspects: selecting optimal feature layers and utilising visual and textual prompts. For the former, we identify which layers are most enriched with information and are best suited for the specific retrieval requirements (category-level or fine-grained). Then we employ visual and textual prompts to guide the model’s feature extraction process, enabling it to generate more discriminative and contextually relevant cross-modal representations. Extensive experiments on several benchmark datasets validate significant performance improvements.

1. Introduction

Diffusion models [16, 30, 31, 68], celebrated for their proficiency in generating photorealistic images, have witnessed remarkable advances in the field of computer vi-

sion. Ranging from image generation [16, 65, 68, 73] to semantic segmentation [2, 99] and image-to-image translation [57, 92, 105], these models have ignited transformations across various domains within the computer vision discipline [8, 13, 44, 89]. In this paper, we embark on an uncharted journey, delving into the intrinsic feature representations of diffusion models, with a primary emphasis on Zero-Shot Sketch-based Image Retrieval (ZS-SBIR).

At the heart of our exploration lies the assertion that text-to-image diffusion models [16, 68] excel as “matchmakers”, seamlessly connecting the realms of sketches and photos. We substantiate this claim through a series of pilot studies, unveiling two fundamental properties inherent in diffusion models, which are pivotal for the success of ZS-SBIR [77, 102]. First and foremost, these models naturally transcend modalities, imbued with the ability to bridge the gap between different types of data, providing a robust foundation for cross-modal retrieval. Furthermore, we unearth a pronounced *shape bias* nestled within their joint embeddings, rendering them exquisitely well-suited for the intricacies of SBIR. It is worth noting that zero-shot capabilities virtually come as an added bonus, thanks to the open-vocab nature of large-scale pre-trained diffusion models [16, 68] (e.g., Stable Diffusion [68]).

Nonetheless, effectively harnessing pre-trained diffusion

models [68] for ZS-SBIR may initially seem not entirely straightforward. These models encapsulate a diverse array of features across different layers and time-steps [68], requiring careful selection to align with the intricacies of the specific task at hand. Moreover, their pre-training with a text-to-image generation objective [68] endows them with a heavy bias to generate *text-influenced image features*, moving them away from the conventional sketch-photo visual matching [77] problem encountered in SBIR.

Our solution to this seemingly intricate problem, however, is elegantly simple. It revolves around two key principles: (i) identifying the layers and time-steps most enriched with information for SBIR, and (ii) optimising the interaction with Stable Diffusion (SD) to effectively condition feature extraction. In addressing the former, we shed light on the optimal layers and time-steps for feature extraction, tailored to various ZS-SBIR tasks, including its fine-grained variant. Our findings suggest that, for category-level retrieval (*i.e.*, ZS-SBIR), features from the first few UNet upsampling blocks prove most effective, while for fine-grained retrieval (*i.e.*, ZS-FG-SBIR), latter upsampling blocks excel. To adapt the pre-trained Stable Diffusion model [68] for SBIR, we employ task-specific visual prompts within both sketch and photo branches. The SD model [68], trained with a text-to-image synthesis objective, boasts rich visio-lingual joint information [68]. To leverage this text-enriched knowledge to the fullest, we introduce textual prompt learning. We acquire both visual and textual prompts using a simple triplet-based objective, all the while maintaining the integrity of the frozen SD model [68]. During inference, we apply the learned visual prompt to the query sketch and pass it through the frozen SD model [68], conditioned on the acquired textual prompt. This process yields the query feature, which is subsequently compared with pre-computed gallery features, facilitating retrieval.

In summary, (i) we unveil the latent potential of diffusion models as backbone feature extractors for ZS-SBIR, substantiated by empirical evidence and comprehensive analysis. (ii) we introduce innovative design strategies, including soft prompt learning and visual prompting, designed to address the challenges of *all* forms of zero-shot SBIR tasks. (iii) through extensive experimentation on standard datasets, we demonstrate marked improvements in performance compared to conventional ZS-SBIR approaches.

2. Related Works

Sketch-Based Image Retrieval (SBIR). SBIR began as a category-level task aimed at retrieving a photo of the same category as the query sketch. Evolving from classical approaches [33, 63, 71, 72], deep-learning ones [12, 42, 50, 101] usually trained Siamese-like networks [12], typically based on CNNs [15], RNNs [100] or Transformers [67] to

fetch similar photos over a distance-metric in a cross-modal joint embedding space [11]. It was enhanced further by unsupervised multi-clustering based re-ranking [95], learning discriminative structure representation via dynamic landmark discovery [104], incremental learning [5], etc.

Fine-Grained SBIR. Given a sketch, FG-SBIR aims to fetch its *target*-photo from a gallery of same-category photos. Since its inception [47], numerous works have emerged [4, 61, 74, 88, 103], aided by new datasets having fine-grained sketch-photo association [58, 79, 103]. From a deep-triplet ranking based Siamese network [103] learning a joint sketch-photo manifold, FG-SBIR was enhanced with attention mechanism using higher-order [88] losses, hybrid cross-domain image generation [60], textual tags [87], local feature alignment [98] and mixed-modal jigsaw solving based pre-training tasks [62], to name a few. While some addressed inherent sketch-traits, like hierarchy [74], style-diversity [75], or redundancy of strokes [6] others explored applicability like early-retrieval [3], cross-category generalisation [7, 61], and overcoming data-scarcity [4, 78]. Recently FG-SBIR was extended to scene-level (retrieve a scene image, given a scene-sketch) via cross-modal region-associativity [9], or with an optional text-query [10].

Zero-Shot SBIR. Towards tackling data-scarcity, ZS-SBIR aims at generalising knowledge learned from *seen* training classes to *unseen* (disjoint) testing categories. Introduced by Yelamarthi *et al.* [102] where photo-features were approximated from sketches via VAE-based [39] image-to-image translation, to generalise onto unseen classes, later works shifted to exploiting word2vec [56] representation of class labels [15, 19] for semantic transfer. While [19] attempted to align sketch, photo and semantic representations via adversarial training, others aimed to minimise the sketch-photo domain gap using a gradient reversal layer [15], a conditional-VAE based graph convolution network [106], a shared ViT [91] backbone, employing prototype-based [94] selective knowledge distillation, and a test-time training paradigm [76] to name a few. While semantic transfer in all such works was mostly limited to using word embeddings directly [19, 93, 106] or indirectly [51, 90, 94], very recently Sain *et al.* [77] adapted CLIP [64] to exploit its high generalisability for semantic transfer. Instead of using feature extractors trained on *discriminative* tasks, we propose using pre-trained text-to-image *generative* backbone for all kinds of SBIR tasks.

Backbones for SBIR. State-of-the-art backbones for SBIR can be broadly categorised as – (a) *Standard pre-trained* [4, 6, 9], (b) *VAE-based formulations* [75, 102, 106], (c) *Self-supervised pre-trained* [46, 62], and (d) *Foundational models* [77, 81]. *Standard* models typically use ImageNet pre-trained CNNs [4, 6, 9] with specific design choices like – spatial attention [88], conditional stroke re-

covery [49], cross-modal pairwise-embedding [74], cross-attention [48], etc. While, *VAE-based* [39] models were used in sketch-photo translation to minimise sketch-photo domain gap [102], meta-learning content-style disentanglement [75] etc., other works focused on *self-supervised pre-training* either via leveraging the neighbourhood topology induced by large-scale photo pre-training [46] or by solving mixed-modal jigsaw as a pretext task [62]. Finally, the recent rise of *foundational models* [45, 64, 68], nudged several works to adapt them as backbones, by either training CLIP [64] end-to-end [81] for standard SBIR or fine-tuning the LayerNorm layers [77] of a pre-trained CLIP model with prompt learning for ZS-SBIR. In this paper, we aim to address *all* aspects (*i.e.*, category-level and fine-grained) of zero-shot SBIR, utilising the inherent zero-shot potential [44] of a pre-trained Stable Diffusion Model [68].

Diffusion Models for Vision Task. Recently, the diffusion model [31] has become the de facto standard for image generation [16] and editing [70]. Essentially, it converts an image \mathbf{x}_0 into a noisy image \mathbf{x}_T by iteratively adding random Gaussian noise to \mathbf{x}_0 in T time-steps [31]. The reverse, recovers \mathbf{x}_0 from \mathbf{x}_T over T denoising steps [31]. Subsequent developments of classifier-free guidance [30] and latest latent diffusion model [68] lead to numerous controllable image generation frameworks like DALL-E [65], Make-A-Scene [21], Imagen [73], T2I-Adapter [57], ControlNet [105], CogView [17], etc. This gave rise to various works [22, 29, 37, 54, 55, 70] on realistic image-edits leveraging zero-shot diffusion models. Additionally, *diffusion* has found use in several vision tasks like classification [44], semantic [2] and panoptic segmentation [99], image-to-image translation [43, 92], medical imaging [13], image correspondence [89], object detection [97], etc. Despite its advances, the efficacy of pre-trained diffusion models for cross-modal zero-shot retrieval remains under-explored. Here, we thus aim to explore the inherent zero-shot cross-modal retrieval potential of a frozen stable diffusion model.

3. Revisiting Text-to-Image Diffusion Models

Overview. Diffusion models generate images by progressive removal of noise from an initial pure 2D Gaussian noise [16, 31]. It relies on two complementary random processes, namely, “*forward*” and “*reverse*” diffusion [31]. The *forward* process iteratively adds Gaussian noise of varying magnitudes to a clean image $\mathbf{x}_0 \in \mathbb{R}^{h \times w \times 3}$ from the training dataset, for t time-steps to produce a noisy image $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$. Here, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is the added noise, $\{\alpha_t\}_1^T$ is a pre-defined noise schedule where $\bar{\alpha}_t = \prod_{k=1}^t \alpha_k$ with $t \sim U(0, T)$. When T is large enough, the resulting image ($\mathbf{x}_T \in \mathbb{R}^{h \times w \times 3}$) approximates pure noise [31]. The *reverse* process involves training a denoising UNet [69] \mathcal{F}_θ , that tries to estimate the input noise $\epsilon \approx \mathcal{F}_\theta(\mathbf{x}_t, t)$ from the noisy image \mathbf{x}_t at every time-step

t . Once trained with an l_2 objective [31], \mathcal{F}_θ can reverse the effect of the forward diffusion. During inference, starting from a random pure 2D noise \mathbf{x}_T sampled from a Gaussian distribution, \mathcal{F}_θ is applied iteratively (for T time-steps) to estimate and subtract the noise from each time-step to get a cleaner image \mathbf{x}_{t-1} , eventually leading to one of the cleanest images \mathbf{x}_0 from the original target distribution [31].

Text-Conditioned Diffusion Model. This unconditional process could be made “conditional” by influencing the \mathcal{F}_θ with auxiliary guiding signals \mathbf{p} (*e.g.*, class labels [32], textual prompts [38, 59, 66, 68, 73], dense semantic maps [57, 105], etc.). A pre-trained CLIP text encoder [64] $\mathcal{T}(\cdot)$ converts textual prompt \mathbf{p} (textual description) into tokenised embedding $\mathbf{T}_p = \mathcal{T}(\mathbf{p}) \in \mathbb{R}^{77 \times d}$ that influences the denoising process via cross-attention to generate spatial attention maps for each word token [68]. Spatial attention maps at each time-step t control the visio-linguistic interaction between spatial feature maps and text tokens. Here, we use the text-conditional Latent Diffusion Model (SD) [68], where the diffusion process occurs on the latent space instead of the pixel space, for faster training [68].

Stable Diffusion Architecture. Given an image and text-prompt pair (\mathbf{x}, \mathbf{p}) , SD [68] first uses the encoder $\mathcal{E}(\cdot)$ from a pre-trained variational autoencoder – VAE (consists of an encoder $\mathcal{E}(\cdot)$ and decoder $\mathcal{D}(\cdot)$ in cascade) [39] to convert the input image $\mathbf{x}_0 \in \mathbb{R}^{h \times w \times 3}$ to a latent image $\mathbf{z}_0 = \mathcal{E}(\mathbf{x}_0) \in \mathbb{R}^{\frac{h}{8} \times \frac{w}{8} \times d}$. Later, a *time-conditional denoising UNet* [69] $\mathcal{F}_\theta(\cdot)$ is trained to denoise directly on the latent space. SD’s UNet architecture consists of 12 encoding layers, 1 middle layer, and 12 skip-connected decoding layers [68]. Among the 12 encoding and 12 decoding layers, there are 4 downsampling $\{\mathcal{F}_d^1, \mathcal{F}_d^2, \mathcal{F}_d^3, \mathcal{F}_d^4\}$ and 4 upsampling $\{\mathcal{F}_u^1, \mathcal{F}_u^2, \mathcal{F}_u^3, \mathcal{F}_u^4\}$ layers respectively. $\mathcal{F}_\theta(\cdot)$ takes three inputs – (i) the t^{th} step noisy latent $\mathbf{z}_t = \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$, (ii) the scalar time-step value t , and (iii) textual embedding $\mathbf{T}_p = \mathcal{T}(\mathbf{p}) \in \mathbb{R}^{77 \times d}$. Both time-embedding and textual embedding influence the intermediate feature maps of the UNet via cross-attention [68]. $\mathcal{F}_\theta(\cdot)$ then predicts the noise map $\hat{\epsilon}_t$ (same size as \mathbf{z}_t) corresponding to that time-step. Mathematically,

$$\mathcal{F}_\theta : (\mathbf{z}_t, t, \mathbf{T}_p) \rightarrow \hat{\epsilon}_t \quad (1)$$

Parameters of \mathcal{F}_θ are learned over l_2 loss as: $\mathcal{L}_{\text{SD}} = \mathbb{E}_{\mathbf{z}_t, t, \mathbf{T}_p, \epsilon} (\|\epsilon - \mathcal{F}_\theta(\mathbf{z}_t, t, \mathbf{T}_p)\|_2^2)$. During text-to-image inference, SD *discards* $\mathcal{E}(\cdot)$, directly sampling a noisy latent image \mathbf{z}_T from Gaussian distribution [68]. It then estimates noise from \mathbf{z}_T iteratively via \mathcal{F}_θ (conditioned on \mathbf{p}) to obtain a clean latent $\hat{\mathbf{z}}_0 \in \mathbb{R}^{\frac{h}{8} \times \frac{w}{8} \times d}$ after T iterations. Passing $\hat{\mathbf{z}}_0$ via frozen VAE decoder, generates the final image $\hat{\mathbf{x}} = \mathcal{D}(\hat{\mathbf{z}}_0) \in \mathbb{R}^{h \times w \times 3}$ [68].

4. Feature Extraction via Stable Diffusion

Intermediate feature maps of \mathcal{F}_θ holding rich semantic information have been used in several downstream




Figure 2. Feature extraction via text-to-image diffusion model.

vision tasks like segmentation [99], classification [44], correspondence-learning [28], etc. The extracted feature depends on two *task-specific* design choices – (i) the time-step t , and (ii) the selection of a particular intermediate feature map at that time-step. Here (Fig. 2), we aim to use these information-rich interim representations as backbone features for ZS-SBIR. In particular, given an image and text-prompt pair (x, p) , and a time-step t , we first generate the latent image $z_0 = \mathcal{E}(x)$. We then add noise from time-step t to transform z_0 to its t^{th} -step noisy latent image z_t (forward diffusion). Now we feed, – (i) the noisy latent z_t , (ii) scalar time-step value t , and (iii) textual embedding $T_p = \mathcal{T}(p)$ into $\mathcal{F}_\theta(\cdot)$ to extract corresponding intermediate features from upsampling layers $\{\mathcal{F}_u^1, \mathcal{F}_u^2, \mathcal{F}_u^3, \mathcal{F}_u^4\}$. For example, an image of size $\mathbb{R}^{256 \times 256 \times 3}$ would yield feature maps of $f_u^1 = \mathbb{R}^{8 \times 8 \times 1280}$, $f_u^2 = \mathbb{R}^{16 \times 16 \times 1280}$, $f_u^3 = \mathbb{R}^{32 \times 32 \times 640}$, and $f_u^4 = \mathbb{R}^{32 \times 32 \times 320}$ from $\{\mathcal{F}_u^1, \mathcal{F}_u^2, \mathcal{F}_u^3, \mathcal{F}_u^4\}$ respectively. We then perform global max-pooling on these feature maps (f_u^n) to produce d -dimensional feature vectors (\hat{f}_u^n). In absence of p , the embedding for a *null prompt* [89] is used.

5. Pilot Study: Analysis and Insights

Cross-modal Correspondence. High-quality image generation [68, 70] from SD models inspired us to further investigate their internal representations and explore whether a *frozen* SD model can deliver cross-modal shape correspondence [89]. Accordingly, we hypothesise that – (i) although aimed at image generation [68], the UNet denoiser’s internal feature maps embed highly localised hierarchical semantic information; (ii) this localised semantic knowledge is robust across various image modalities (e.g., image vs. sketch). To validate, we extracted SD internal features of a few photos and their corresponding sketches (Sec. 4) at $t=273$ (ablated in Sec. 8.2) with a fixed prompt “a photo of [CLASS]” [89]. For visualising the high-dimensional features, we perform PCA on feature maps f_u^n ; $n \in \{1, 2, 3, 4\}$, and render the first three principal components [53, 89] as RGB images (Fig. 1 right). Feature representations from different upsampling layers here depict the different feature hierarchies. While earlier layers ($n=1, 2$) delineate low-frequency coarse-level features (e.g., shape and structure), later ones ($n=3, 4$) depict high-frequency fine-grained features (e.g., image gradients), thus validating our first hypothesis. We can also see how sketch and photo features from different decoder levels share sig-

nificant *semantic similarity* (shown in *same* colour) albeit originating from different modalities that hold significant visual disparity (e.g., sparse-binary sketch vs. pixel-perfect photo) [75] – supporting our second presumption. This motivates us to leverage the frozen SD model’s internal features for both category-level and fine-grained ZS-SBIR.

Reduced Texture-bias. ImageNet [14] pre-trained discriminative CNN [86] backbones are highly biased towards texture features [25, 35]. Now, *sketches* being binary and sparse line drawings, do not provide any colour/texture cues in context of image retrieval [40, 41]. Thus, SBIR being a *shape-matching* [4, 7] problem should ideally be biased towards *shape* instead of texture features. We posit that the texture-bias of discriminative CNN backbones largely bottlenecks existing SBIR performance. To validate this, we use the cue-conflict dataset [25] containing images across 16 classes, where style-transfer [24] is applied to alter the texture (style) of the images with random textures of images from other classes, keeping the shape (content) same (Fig. 3). Out of them, we select 10 classes overlapping with Sketchy [80] to form our retrieval gallery. We train a frozen VGG-16 [86] backbone followed by *one* learnable FC-layer with triplet loss on real sketch-photo pairs from Sketchy [80]. We observe that the SBIR accuracy (tested on cue-conflict gallery) of VGG-16 *discriminative* backbone is much lower (49.14%) than our method using the SD *generative* backbone (72.82%). Furthermore, asking a group of human observers to categorise these texture-altered images, resulted in high error-consistency (90.16%) [35] with our method. This shows generative backbones to be biased towards *shape* and handle *adversarial* examples [35] better, like the human visual system, tailoring them for ZS-SBIR.




Figure 3. Texture-altered images from cue-conflict [25] dataset.

6. Background: Baseline Zero-Shot SBIR

Category-Level ZS-SBIR. For a query sketch s of any class, this aims to fetch a photo p_i^j of the *same* class, from a gallery $\mathcal{G} = \{p_1^j, \dots, p_{N_j}^j\}_{j=1}^{N_c}$ containing photos from N_c classes with N_j photos per class. Usually, an encoder $\mathcal{B}(\cdot)$ (separate weights for sketch and photo branch), is trained to extract d -dimensional feature $f_i = \mathcal{B}(\mathcal{I}) : \mathbb{R}^{h \times w \times 3} \rightarrow \mathbb{R}^d$, over a triplet loss [103] that aims to minimise the distance $\delta(\cdot, \cdot)$ between features of an anchor sketch (f_s) and its matching photo (f_p) from the same class, while increasing that from a non-matching photo (n) feature f_n of a *different* class [103]. With margin $\mu > 0$, the triplet loss becomes:

$$\mathcal{L}_{\text{trip}} = \max\{0, \mu + \delta(f_s, f_p) - \delta(f_s, f_n)\} \quad (2)$$

While conventional SBIR evaluates on classes *seen* during training, $\mathcal{C}^S = \{c_1^S, \dots, c_N^S\}$, ZS paradigm [102] evaluates on those *unseen*, $\mathcal{C}^U = \{c_1^U, \dots, c_M^U\}$, $\therefore \mathcal{C}^S \cap \mathcal{C}^U = \emptyset$.

Cross-Category Zero-Shot Fine-Grained SBIR. Unlike category-level, here the aim is to train a *single* model capable of fine-grained instance-level matching from multiple (N_c) classes [77]. Standard cross-category FG-SBIR datasets (e.g., Sketchy [80]) consist of instance-level sketch-photo pairs $\{s_i^j, p_i^j\}_{i=1}^{N_j} |_{j=1}^{N_c}$ from N_c categories with N_j sketch-photo pairs per category with fine-grained association. Similar to category-level, cross-category FG-SBIR framework trains a backbone feature extractor $\mathcal{B}(\cdot)$ shared between sketch and photo branches, over a similar triplet loss (Eq. 2) but with *hard* triplets, where the negative sample is a *different instance* ($p_k^j; k \neq i$) of the *same class* as the anchor sketch (s_i^j) and its matching photo (p_i^j) [77]. It further employs an N_c -class classification head on the joint embedding space for learning class discrimination [77]. Following ZS-SBIR paradigm, the classes used for inference are also unseen ($\mathcal{C}^S \cap \mathcal{C}^U = \emptyset$) with the aim of instance-level matching instead of category-level [77].

While existing works [77, 78, 103] have designed $\mathcal{B}(\cdot)$ with ImageNet pre-trained VGG-16 [103], PVT [78], ViT [78], or foundational models like CLIP [77], the research question remains as to how can we use a frozen text-to-image diffusion model as a backbone feature extractor $\mathcal{B}(\cdot)$.

7. Prompt Learning for SBIR via Frozen SD

Overview. Adapting a frozen SD model [68] as a backbone feature extractor $\mathcal{B}(\cdot)$ is non-trivial. For instance – (i) one off-the-shelf choice is to fine-tune the entire UNet with triplet loss [103]. However, being pre-trained with a *generative* objective [68], the UNet will lose its rich semantic knowledge-base when fine-tuned with a *discriminative* loss; (ii) learning additional FC layers on top of extracted features via triplet loss (akin to linear probing [1]) is also sub-optimal for the same reason; finally, (iii) the forward diffusion process [31, 68] being stochastic in nature [68] introduces random noise in output feature maps, making them sub-optimal for discriminative learning in the metric space. To tackle these issues related to losing pre-trained knowledge (via fine-tuning or linear probing), we shift to a *prompt learning* [1, 108]-based approach (Fig. 4), keeping the weights of the SD model [68] frozen and thus its semantic knowledge-base intact. To further increase the representation-stability of the extracted features, we use *feature ensembling*.

Specifically, we have three salient components: (i) a learnable task-specific visual prompt, (ii) a learnable continuous textual prompt to harness the visio-linguistic prior of SD [68], and (iii) feature ensembling for additional robustness during inference.




Figure 4. Given the frozen SD [68] backbone feature extractor, our method learns a single textual prompt, and sketch/photo-specific visual prompts via triplet loss.

Visual Prompt Design. Originating from NLP literature [85], *prompting* [107, 108] has now become a de facto choice for several vision tasks [1, 36, 107]. Prompting involves adapting a large-scale pre-trained model to a specific downstream task without modifying its weights [1, 108]. However, creating task-specific prompts is labour-intensive and requires domain expertise [1, 108]. Prompt learning alleviates this issue by learning a set of context vectors via backpropagation, keeping the model weights frozen [108]. In ViT [18]-based models with vectorised patch-representation, one usually appends additional trainable continuous prompt vectors along with the patch-tokens, for task-specific adaption [36]. However, the convolutional input/output (spatial tensor) of SD’s UNet denoiser [68] prevents direct concatenation of trainable prompt vectors.

Consequently, we opt for visual prompting [1] that learns a soft image perturbation in the pixel space to adapt SD model [68] to our problem setup of ZS-SBIR. Given an image $\mathcal{I} \in \mathbb{R}^{h \times w \times 3}$, we add a learnable *spatial* image prompt $\mathcal{P}_v \in \mathbb{R}^{h \times w \times 3}$ whose every (x, y) spatial position within the border region (width= d) is a trainable continuous \mathbb{R}^3 -dimensional vector, while the remaining locations are filled with a fixed value of zero and kept frozen. Mathematically,

$$\mathcal{P}_v(x, y) = \begin{cases} \text{learnable} & \forall \left[\begin{array}{l} (x < d) \vee (x \geq (w - d)) \\ (y < d) \vee (y \geq (h - d)) \end{array} \right] \\ 0, \text{frozen} & \text{elsewhere} \end{cases} \quad (3)$$

Therefore, the prompt added image $\mathcal{I}_p = \mathcal{I} + \mathcal{P}_v$ is passed through the SD feature backbone to generate the embedded feature as described in Sec. 4. For a border size d , the number of learnable parameters is $2 \times c \times d(h + w - 2d)$.

Textual Prompt Design. Text-to-image diffusion models [16, 68] being trained on text-to-image generative objective [68], works best with *explicit* textual prompts. However, the standard sketch-photo datasets [80, 103] lack associated textual descriptions. Furthermore, generating textual captions via off-the-shelf *image* captioner [45] does not fare well for sparse and binary freehand *sketches* [10]. To mitigate this, instead of actual textual prompt embedding $\mathbf{T}_p = \mathcal{T}(\mathbf{p}) \in \mathbb{R}^{77 \times d}$, we use a learnable continuous textual prompt embedding matrix $\mathcal{P}_t \in \mathbb{R}^{77 \times d}$, influencing the SD feature extraction process via cross-attention (Sec. 3).

Prompt Training. As the earlier ($n=1, 2$) and later

($n=3, 4$) upsampling layers hold coarse-grained and fine-grained features respectively (Sec. 5), for category-level ZS-SBIR, we empirically use the mean of the 1st and 2nd upsampling layer’s global max-pooled features ($t=273$) $\hat{f}_u^1 \in \mathbb{R}^{1280 \times 1}$ and $\hat{f}_u^2 \in \mathbb{R}^{1280 \times 1}$, while for cross-category zero-shot FG-SBIR, we concatenate the 3rd and 4th upsampling layer’s global max-pooled features ($t=273$) $\hat{f}_u^3 \in \mathbb{R}^{640 \times 1}$ and $\hat{f}_u^4 \in \mathbb{R}^{320 \times 1}$ to form a $\mathbb{R}^{960 \times 1}$ feature. Now, as the SD feature extraction pipeline is differentiable [68] (although frozen), gradient updates from triplet loss (Eq. 2) will be directly backpropagated across UNet and VAE encoder (via reparameterisation-trick [39]) to update the learnable parts of \mathcal{P}_v (Eq. 3), and those backpropagated across the UNet will update \mathcal{P}_t . In practice, for category-level ZS-SBIR, we use two separate visual prompts for sketch and photo branches, but share a common visual prompt for fine-grained ZS-SBIR. Furthermore, thanks to its implicit semantic knowledge [68], pre-trained diffusion models can handle multi-category retrieval *without* category-specific fine-tuning [44], thus bypassing the need for additional N_c -class classification head as used in baseline cross-category zero-shot FG-SBIR (Sec. 6).

Feature Ensembling during Inference. Calculating z_t from z_0 (during forward diffusion) invokes stochasticity due to the random noise sampling [68], which could deteriorate the quality of extracted features [89]. To tackle this, we extract SD features for each image/sketch six times (ablated in Sec. 8.2) each from different noise samples, and *ensemble* them by averaging to obtain the final feature.

8. Experiments

Datasets. We evaluate on three popular sketch-photo datasets – (i) **Sketchy** [80] consists of 12,500 photos across 125 classes each having at least 5 fine-grained paired sketches. While we use the original Sketchy dataset with fine-grained sketch-photo pairing to evaluate cross-category ZS-FG-SBIR, for zero-shot category level setup, we use its extended version [50] with an additional 60,652 ImageNet photos. Following [77, 102], we use sketches/photos from 104 and 21 classes for training and testing respectively. (ii) **Quick, Draw!** [26] has $\sim 50M$ sketches categorised into 345 classes. Following [15], we use a subset of it encompassing 330,000 sketches and 204,000 photos across 110 categories split as 80:30 for training:testing. (iii) **TU-Berlin** [20] houses 204,489 photos across 250 categories each having at least 80 sketches. Following [15], we use a training:testing class split of 220:30.

Implementation Details. We use Stable Diffusion v2.1 [68] in all experiments with a CLIP [64] embedding dimension $d = 1024$. The visual and textual prompts are trained with a learning rate of 10^{-4} , keeping the UNet denoiser and the VAE encoder frozen. We learn the prompts for 100

epochs, on an Nvidia V100 GPU using AdamW [52] optimiser with 0.09 weight decay, and a batch size of 64.

Evaluation. For zero-shot category-level setup, following [15, 77], evaluation considers the top 200 retrieved images to calculate mean average precision (mAP@200) and precision (P@200) on Sketchy while reporting mAP@all and P@100 scores for TU-Berlin. Furthermore, for Quick, Draw! dataset we resort to mAP@all and P@200 [77]. For cross-category ZS-FG-SBIR evaluation however, we use Acc.@q, which denotes the percentage of sketches with true-matched photos in the top-q retrieved images.

Competitors. We compare against state-of-the-art (SoTA) ZS-SBIR frameworks in two paradigms *viz.* ZS-SBIR and ZS-FG-SBIR. Among ZS-SBIR SoTA paradigms, we compare with **ZS-CVAE** [102], **ZS-CAAE** [102], **ZS-CCGAN** [19], **ZS-GRL** [15], **ZS-SAKE** [51], **ZS-IIAE** [34], **ZS-Sketch3T** [76], and **ZS-LVM** [77]. While the SoTA methods rely on sketch-to-image generation [102], word2vec [56] encoding with adversarial learning [19], knowledge-distillation [51], test-time adaptation [76], or pre-trained CLIP models [77], our method leverages the information-rich internal representation of frozen SD model for generalisation to unseen classes. Due to the unavailability of diffusion model or visual prompting-based ZS-SBIR frameworks, we compare against a few self-designed **Baselines**. **B-Fine-Tuning** fine-tunes the pre-trained UNet of SD (with a learning rate of 10^{-6}), while **B-Linear-Probe** learns an FC layer on top of the SD features, keeping the rest of the model-weights frozen, both via triplet loss [103]. Next, we extend the baseline triplet loss-based baseline ZS-SBIR with learnable visual prompts (VP) to form **B-Triplet+VP**. Furthermore, for B-Triplet+VP, we use three different backbone feature extractors *viz.* VGG-16 [86], ResNet50 [27], and ViT [18] pre-trained with discriminative task to compare the contribution of generative pre-training over discriminative. These baselines use the exact same visual prompt dimension (Sec. 7) as ours. We extend these triplet-based baselines to the ZS-FG-SBIR setup via *hard* triplet training and an additional classification head (Sec. 6).

8.1. Result Analysis

Category Level ZS-SBIR. Quantitative results shown in Tab. 1 depict that the naive adaption of pre-trained SD model in B-Fine-Tuning and B-Linear-Probing fails drastically due to loss of generative potential during triplet-based training. Although B-Triplet+VP (ViT) performs better than B-Triplet+VP (VGG) and B-Triplet+VP (ResNet) on all three datasets, due to its larger backbone, our method supersedes all these baselines with a mAP@200 of 0.746 on Sketchy, which further verifies our motivation to use the frozen SD model as a backbone feature extractor. While ZS-SoTA methods offer reasonable performance across all

three benchmark datasets owing to the respective zero-shot adaption strategies (e.g., domain translation [102], word2vec embedding [15, 19], KD [51], test-time training [76], etc.), our method surpasses them with an average mAP@200 gain of 43.96% (Sketchy). ZS-LVM [77] despite being our strongest contender, performs sub-optimally in all benchmarks. The demonstrated efficacy (Tab. 1) of frozen SD backbone in category-level ZS-SBIR leads us to the more challenging evaluation setup of ZS-FG-SBIR.

Table 1. Results for category-level ZS-SBIR.

Methods	Sketchy [80]		TU-Berlin [20]		Quick, Draw! [26]	
	mAP@200	P@200	mAP@all	P@100	mAP@all	P@200
ZS-CAAE [102]	0.156	0.260	0.005	0.003	-	-
ZS-CVAE [102]	0.225	0.333	0.005	0.001	0.003	0.003
ZS-CCGAN [19]	-	-	0.297	0.426	-	-
ZS-GRL [15]	0.369	0.370	0.110	0.121	0.075	0.068
ZS-SAKE [51]	0.497	0.598	0.475	0.599	-	-
ZS-IIAE [34]	0.373	0.485	0.412	0.503	-	-
ZS-Sketch3T [76]	0.579	0.648	0.507	0.671	-	-
ZS-LVM [77]	0.723	0.725	0.651	0.732	0.202	0.388
B-Fine-Tuning	0.115	0.174	0.010	0.006	0.002	0.003
B-Linear-Probing	0.441	0.535	0.410	0.582	0.092	0.099
B-Triplet+VP (VGG)	0.651	0.682	0.582	0.673	0.134	0.310
B-Triplet+VP (ResNet)	0.326	0.342	0.354	0.512	0.105	0.275
B-Triplet+VP (ViT)	0.681	0.697	0.601	0.694	0.185	0.321
Ours	0.746	0.747	0.680	0.744	0.231	0.397

Cross-Category ZS-FG-SBIR. Being more challenging than category-level ZS-SBIR, ZS-FG-SBIR is relatively under-explored [77]. Quantitative results in Tab. 2 show how our method surpasses SoTA ZS-FG-SBIR frameworks [61, 77, 84] with an impressive average Acc.@1 margin of 32.49%. B-Fine-Tuning and B-Linear-Probing fare poorly here as well confirming our prior observation in the ZS-SBIR setup. Furthermore, our method exceeds all visual prompting-based triplet baselines without the complicity of learning additional classification head (Sec. 8).

Table 2. Results for cross-category ZS-FG-SBIR on Sketchy [80].

Methods	Acc.@1	Acc.@5	Methods	Acc.@1	Acc.@5
CC-Gen [61]	22.60	49.00	B-Triplet+VP (VGG)	24.20	43.61
CC-Grad [84]	13.40	34.90	B-Triplet+VP (ResNet)	15.61	27.64
CC-LVM [77]	28.68	62.34	B-Triplet+VP (ViT)	26.11	46.81
B-Fine-Tuning	1.85	6.01	Ours	31.94	65.81
B-Linear-Probing	17.32	41.23			

Performance in Low-Data Scenario. Existing ZS-(FG)-SBIR models are limited by the availability of paired sketch-photo training data [15, 19, 77, 102]. To test our method in a *low-data scenario*, we experiment by varying the number of training data (10%, 30%, 50%, 70%, 100%) per class from Sketchy [80]. Our method remains relatively stable, outperforming (Fig. 5) other baseline and SoTAs significantly, preserving almost at par accuracy with *full-scale* training, by harnessing the out-of-distribution generalisation [70, 89, 99] capability of pre-trained SD [68] model.

8.2. Ablation

[i] **Contribution of visual and textual prompting.** To estimate the effectiveness of visual prompts, we train our




Figure 5. Plots showing low-data scenario performance for ZS-SBIR (left) and ZS-FG-SBIR (right) setup on Sketchy [80] dataset.

method without learning the visual prompt, thus eliminating the task-specific adaptation. A significant mAP@200 (Acc.@1) drop of 0.233(14.28) on Sketchy (Tab. 3) in case of **w/o visual prompt** signifies that adapting SD for ZS-SBIR is incomplete without explicitly learning task-specific prompts in the pixel space. The spatial area of the visual prompt can influence the final performance significantly [1]. Varying the visual prompt’s border width (d) in the range [10, 20, 30, 40, 50], we find that $d = 16$ produces the optimum mAP@200 (Acc.@1) on Sketchy (Fig. 7 (right)). Secondly, to judge the efficacy of the textual prompts, we replace the learned textual prompts with *null* prompts (i.e., “ ”) during inference. The result of **w/o textual prompt** in Tab. 3 depicts that removing the learned textual prompts further plummets the mAP@200 by 6.83%.

[ii] **Importance of feature ensembling.** Feature ensembling helps reduce the effect of stochastic noising (during forward diffusion) of latent images in the final feature maps. Removing it further destabilises the framework causing an additional mAP@200 drop of 0.021 on Sketchy (Tab. 3). We posit that this drop is due to the absence of the additional smoothing regularisation provided by feature ensembling that reinforces discriminative learning in the metric space. Experimenting with different sizes ([1, 8]), we found that feature ensembling six SD features gives optimum accuracy across all datasets.

[iii] **Do \mathcal{F}_d^n layers hold representative features?** We also evaluate the ZS-SBIR and ZS-FG-SBIR performance by using UNet downsampling layer \mathcal{F}_d^n features. We train our model with features from $\{\mathcal{F}_d^n\}_{n=1}^2$ and $\{\mathcal{F}_d^n\}_{n=3}^4$ for ZS-FG-SBIR and ZS-SBIR setup respectively. Nonetheless, sub-optimal mAP@200 and Acc.@1 on Sketchy in case of **\mathcal{F}_d^n features** (Tab. 3) render the UNet downsampling layer features unsuitable for the tasks of ZS-(FG)-SBIR.

[iv] **Which feature-level of \mathcal{F}_u^n is optimum?** As evident from Fig. 1, different resolution-levels of \mathcal{F}_u^n contain features of different *semantic-granularity*. To select the optimal set of layers for feature extraction, we performed a 2D-grid search separately for ZS-SBIR and ZS-FG-SBIR setups on Sketchy [79]. Results shown in Tab. 4 unfold that a combination of $n = 1, 2$ and $n = 3, 4$ works best for ZS-SBIR and ZS-FG-SBIR respectively.

[v] **Do all time-steps yield useful features?** As seen in




Figure 6. PCA representation of SD [68] internal features from \mathcal{F}_u^1 upsampling layers of UNet for different time-steps ($t \in [0, 100, \dots, 900]$). Different regions of sketch and photo feature maps from $t \in [200, 300]$ (highlighted in red) portray strong semantic feature correspondence (represented by the same colours in the PCA map), while the features from the later time-steps are non-aligned.

Fig. 6, the cross-modal feature maps are most *semantically-aligned* (same colour) around $t \in [200, 300]$. Calculating the average mAP@200 (Acc.@1) on Sketchy over varying t shows $t = 273$ to generate optimal features for both ZS and ZS-FG-SBIR. As seen in Fig. 7 (left), our method is quite robust to the selection of t , as a wide range of t produces mAP@200 (Acc.@1) well over the baselines.

Table 3. Ablation on design.

Methods	Sketchy [80]	
	mAP@200	Acc.@1
w/o visual prompt	0.513	17.66
w/o textual prompt	0.695	22.71
w/o feature ensemble	0.725	29.47
\mathcal{F}_d^n features	0.365	18.22
Ours-full	0.746	31.94
Avg. Improvement	0.171	9.92

Table 4. Choice of \mathcal{F}_u^n index.

n				Sketchy [80]	
1	2	3	4	mAP@200	Acc.@1
✓	✗	✗	✗	0.431	10.40
✓	✗	✗	✗	0.746	15.21
✓	✓	✓	✗	0.712	25.44
✓	✓	✓	✓	0.613	29.36
✗	✗	✗	✗	0.287	20.61
✗	✗	✓	✓	0.352	31.94
✗	✓	✓	✓	0.408	30.21




Figure 7. Quantitative results on Sketchy [80] for ZS-SBIR (mAP@200) and ZS-FG-SBIR (Acc.@1) setup for different denoising time-steps (left) and visual prompt border width (right).

9. Extension to Sketch+Text-based Retrieval

Stable Diffusion [68] is trained with large-scale image-text pairs with a generative objective, where the model learns the *visio-linguistic interactions* between image and text via cross-attention [68]. Presuming the non-availability of textual captions/class labels for standard SBIR tasks, we used learnable textual prompts in our proposed method. However, here, we extend our method to *Zero-Shot Sketch+Text-Based Image Retrieval* (ZS-STBIR) where we aim to improve the extracted feature-quality with readily available textual captions/class labels. In other words, While most of the existing SBIR backbones (*e.g.*, VGG-16 [86], PVT [96], CLIP [64], etc.) are a function of image *only* $\mathcal{B}(\mathcal{I})$, the proposed stable-diffusion feature extractor could be used to extract joint visio-linguistic features where the visual (\mathcal{I}) and textual (\mathbf{p}) modality interacts with each other and outputs a *text-enhanced* feature implicitly as $\mathcal{F}_\theta(\mathcal{I}, \mathbf{p})$. No-

tably, although CLIP training involves contrastive learning [64] using image-text pairs, the visual and textual encoders being *independent* of each other, *do not* directly influence the feature extraction process during inference. While several existing works [10, 81, 87] attempt to address this joint sketch+text multi-modal feature learning either via complicated invertible neural network [10], CLIP-fine-tuning [81], or allegedly unstable quadruplet loss [87], we simply pass the associated textual captions in \mathcal{F}_θ in place of the learnable textual prompts.

In practice, for category-level ZS-SBIR and cross-category ZS-FG-SBIR, we additionally feed the class-name (thus *sketch+text*) in the form of a fixed handcrafted prompt [108] “a photo of [CLASS]”. Whereas, for scene-level retrieval [10], we use the readily available paired captions from standard scene-level SBIR datasets [23, 58]. Overall, a simple off-the-shelf extension of our SD-based feature extractor depicts competitive or better STBIR performance on several benchmark datasets (Tab. 5).

Table 5. Results for sketch+text-based image retrieval.

Methods	Category-level		Fine-grained		Scene-level			
	Sketchy [80]	Sketchy [80]	Sketchy [80]	Sketchy [80]	FS-COCO [58]	Sketchy-COCO [23]	Sketchy-COCO [23]	Sketchy-COCO [23]
	mAP@200	P@200	Acc.@1	Acc.@5	Acc.@1	Acc.@10	Acc.@1	Acc.@10
QST [87]	0.204	0.213	8.36	16.17	25.1	54.5	38.9	87.9
TaskFormer [81]	0.491	0.501	21.02	52.34	24.9	55.1	39.0	88.2
SceneTrilogy [10]	0.601	0.672	27.52	60.13	25.7	55.2	39.5	88.7
B-Linear-Probing	0.562	0.613	19.87	44.21	14.8	43.7	29.7	76.2
Ours	0.751	0.762	33.31	68.44	26.3	56.6	40.8	89.6

10. Conclusion

This paper for the first time proposes a novel pipeline to adapt the frozen Stable Diffusion as a backbone feature extractor for both *category-level* and *cross-category fine-grained ZS-SBIR* tasks. With clever usage of visual and textual prompting, our method adapts the pre-trained model to the task at hand without further fine-tuning. Extensive experimental results on several benchmark datasets depict that the proposed method outperforms state-of-the-art ZS-SBIR methods. Furthermore, we perform thorough analytical experiments to establish the best practices for leveraging frozen stable diffusion models as a ZS-SBIR backbone. Lastly, harnessing the inherent visio-linguistic capability of stable diffusion, we extend our pipeline to *sketch+text-based SBIR* enabling the practical sketch+text-based retrieval at category, fine-grained and scene-level scenarios.

References

- [1] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring Visual Prompts for Adapting Large-Scale Models. *arXiv preprint arXiv:2203.17274*, 2022. 5, 7
- [2] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khrulkov, and Artem Babenko. Label-Efficient Semantic Segmentation with Diffusion Models. In *ICLR*, 2021. 1, 3
- [3] Ayan Kumar Bhunia, Yongxin Yang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Sketch Less for More: On-the-Fly Fine-Grained Sketch Based Image Retrieval. In *CVPR*, 2020. 1, 2
- [4] Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, Yongxin Yang, Tao Xiang, and Yi-Zhe Song. More Photos are All You Need: Semi-Supervised Learning for Fine-Grained Sketch Based Image Retrieval. In *CVPR*, 2021. 2, 4
- [5] Ayan Kumar Bhunia, Viswanatha Reddy Gajjala, Subhadeep Koley, Rohit Kundu, Aneeshan Sain, Tao Xiang, and Yi-Zhe Song. Doodle It Yourself: Class Incremental Learning by Drawing a Few Sketches. In *CVPR*, 2022. 2
- [6] Ayan Kumar Bhunia, Subhadeep Koley, Abdullah Faiz Ur Rahman Khilji, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Sketching Without Worrying: Noise-Tolerant Sketch-Based Image Retrieval. In *CVPR*, 2022. 2
- [7] Ayan Kumar Bhunia, Aneeshan Sain, Parth Hiren Shah, Animesh Gupta, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Adaptive Fine-Grained Sketch-Based Image Retrieval. In *ECCV*, 2022. 2, 4
- [8] Dar-Yen Chen, Subhadeep Koley, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, Ayan Kumar Bhunia, and Yi-Zhe Song. DemoCaricature: Democratising Caricature Generation with a Rough Sketch. In *CVPR*, 2024. 1
- [9] Pinaki Nath Chowdhury, Ayan Kumar Bhunia, Viswanatha Reddy Gajjala, Aneeshan Sain, Tao Xiang, and Yi-Zhe Song. Partially Does It: Towards Scene-Level FG-SBIR with Partial Input. In *CVPR*, 2022. 2
- [10] Pinaki Nath Chowdhury, Ayan Kumar Bhunia, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. SceneTrilogy: On Human Scene-Sketch and its Complementarity with Photo and Text. In *CVPR*, 2023. 2, 5, 8
- [11] John Collomosse, Tu Bui, Michael J Wilber, Chen Fang, and Hailin Jin. Sketching with Style: Visual Search with Sketches and Aesthetic Context. In *ICCV*, 2017. 2
- [12] John P. Collomosse, Tu Bui, and Hailin Jin. LiveSketch: Query Perturbations for Guided Sketch-based Visual Search. In *CVPR*, 2019. 1, 2
- [13] Bram de Wilde, Anindo Saha, Richard PG ten Broek, and Henkjan Huisman. Medical diffusion on a budget: textual inversion for medical image generation. *arXiv preprint arXiv:2303.13430*, 2023. 1, 3
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 4
- [15] Sounak Dey, Pau Riba, Anjan Dutta, Josep Lladós, and Yi-Zhe Song. Doodle to Search: Practical Zero-Shot Sketch-based Image Retrieval. In *CVPR*, 2019. 1, 2, 6, 7
- [16] Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *NeurIPS*, 2021. 1, 3, 5
- [17] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. CogView: Mastering Text-to-Image Generation via Transformers. In *NeurIPS*, 2021. 3
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*, 2021. 5, 6
- [19] Anjan Dutta and Zeynep Akata. Semantically Tied Paired Cycle Consistency for Zero-Shot Sketch-based Image Retrieval. In *CVPR*, 2019. 2, 6, 7
- [20] Mathias Eitz, James Hays, and Marc Alexa. How do humans sketch objects? *ACM TOG*, 2012. 6, 7
- [21] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-A-Scene: Scene-Based Text-to-Image Generation with Human Priors. In *ECCV*, 2022. 3
- [22] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *ICLR*, 2023. 3
- [23] Chengying Gao, Qi Liu, Qi Xu, Limin Wang, Jianzhuang Liu, and Changqing Zou. SketchyCOCO: Image Generation from Freehand Scene Sketches. In *CVPR*, 2020. 8
- [24] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image Style Transfer Using Convolutional Neural Networks. In *CVPR*, 2016. 4
- [25] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2019. 4
- [26] David Ha and Douglas Eck. A Neural Representation of Sketch Drawings. In *ICLR*, 2017. 6, 7
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *ICCV*, 2015. 6
- [28] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hosam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised Semantic Correspondence Using Stable Diffusion. *arXiv preprint arXiv:2305.15581*, 2023. 4
- [29] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-Prompt Image Editing with Cross Attention Control. In *ICLR*, 2022. 3
- [30] Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance. *arXiv preprint arXiv:2207.12598*, 2022. 1, 3

- [31] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *NeurIPS*, 2020. 1, 3, 5
- [32] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded Diffusion Models for High Fidelity Image Generation. *JMLR*, 2022. 3
- [33] Rui Hu and John Collomosse. A performance evaluation of gradient field HOG descriptor for sketch based image retrieval. *CVIU*, 2013. 2
- [34] HyeonJoo Hwang, Geon-Hyeong Kim, Seunghoon Hong, and Kee-Eung Kim. Variational Interaction Information Maximization for Cross-domain Disentanglement. In *NeurIPS*, 2020. 6, 7
- [35] Priyank Jaini, Kevin Clark, and Robert Geirhos. Intriguing properties of generative classifiers. *arXiv preprint arXiv:2309.16779*, 2023. 4
- [36] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual Prompt Tuning. In *ECCV*, 2022. 5
- [37] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-Based Real Image Editing with Diffusion Models. In *CVPR*, 2023. 3
- [38] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. DiffusionCLIP: Text-Guided Diffusion Models for Robust Image Manipulation. In *CVPR*, 2022. 3
- [39] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *ICLR*, 2014. 2, 3, 6
- [40] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Picture that Sketch: Photorealistic Image Generation from Abstract Sketches. In *CVPR*, 2023. 4
- [41] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. You'll Never Walk Alone: A Sketch and Text Duet for Fine-Grained Image Retrieval. In *CVPR*, 2024. 4
- [42] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. How to Handle Sketch-Abstraction in Sketch-Based Image Retrieval? In *CVPR*, 2024. 2
- [43] Subhadeep Koley, Ayan Kumar Bhunia, Deeptanshu Sekhri, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. It's All About Your Sketch: Democratising Sketch Control in Diffusion Models. In *CVPR*, 2024. 3
- [44] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your Diffusion Model is Secretly a Zero-Shot Classifier. In *ICCV*, 2023. 1, 3, 4, 6
- [45] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. In *ICML*, 2022. 3, 5
- [46] Ke Li, Kaiyue Pang, and Yi-Zhe Song. Photo Pre-Training, but for Sketch. In *CVPR*, 2023. 2, 3
- [47] Yi Li, Timothy M Hospedales, Yi-Zhe Song, and Shaogang Gong. Fine-grained sketch-based image retrieval by matching deformable part models. In *BMVC*, 2014. 2
- [48] Fengyin Lin, Mingkang Li, Da Li, Timothy Hospedales, Yi-Zhe Song, and Yonggang Qi. Zero-Shot Everything Sketch-Based Image Retrieval, and in Explainable Style. In *CVPR*, 2023. 3
- [49] Zhixin Ling, Zhen Xing, Jian Zhou, and Xiangdong Zhou. Conditional Stroke Recovery for Fine-Grained Sketch-Based Image Retrieval. In *ECCV*, 2022. 3
- [50] Li Liu, Fumin Shen, Yuming Shen, Xianglong Liu, and Ling Shao. Deep Sketch Hashing: Fast Free-hand Sketch-Based Image Retrieval. In *CVPR*, 2017. 2, 6
- [51] Qing Liu, Lingxi Xie, Huiyu Wang, and Alan Yuille. Semantic-Aware Knowledge Preservation for Zero-Shot Sketch-Based Image Retrieval. In *ICCV*, 2019. 2, 6, 7
- [52] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *ICLR*, 2019. 6
- [53] Grace Luo, Lisa Dunlap, Dong Huk Park, Aleksander Holynski, and Trevor Darrell. Diffusion Hyperfeatures: Searching Through Time and Space for Semantic Correspondence. *arXiv preprint arXiv:2305.14334*, 2023. 4
- [54] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *ICLR*, 2021. 3
- [55] Aryan Mikaeili, Or Perel, Daniel Cohen-Or, and Ali Mahdavi-Amiri. SKED: Sketch-guided Text-based 3D Editing. In *CVPR*, 2023. 3
- [56] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NeurIPS*, 2013. 2, 6
- [57] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2I-Adapter: Learning Adapters to Dig out More Controllable Ability for Text-to-Image Diffusion Models. *arXiv preprint arXiv:2302.08453*, 2023. 1, 3
- [58] Pinaki Nath Chowdhury, Aneeshan Sain, Yulia Gryaditskaya, Ayan Kumar Bhunia, Tao Xiang, and Yi-Zhe Song. FS-COCO: Towards Understanding of Freehand Sketches of Common Objects in Context. In *ECCV*, 2022. 2, 8
- [59] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2021. 3
- [60] Kaiyue Pang, Yi-Zhe Song, Tony Xiang, and Timothy M Hospedales. Cross-domain Generative Learning for Fine-Grained Sketch-Based Image Retrieval. In *BMVC*, 2017. 2
- [61] Kaiyue Pang, Ke Li, Yongxin Yang, Honggang Zhang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Generalising Fine-Grained Sketch-Based Image Retrieval. In *CVPR*, 2019. 2, 7
- [62] Kaiyue Pang, Yongxin Yang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Solving mixed-modal jigsaw puzzle for fine-grained sketch-based image retrieval. In *CVPR*, 2020. 1, 2, 3
- [63] Yonggang Qi, Yi-Zhe Song, Tao Xiang, Honggang Zhang, Timothy Hospedales, Yi Li, and Jun Guo. Making Better Use of Edges via Perceptual Grouping. In *CVPR*, 2015. 2

- [64] 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 *ICML*, 2021. 2, 3, 6, 8
- [65] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-Shot Text-to-Image Generation. In *ICML*, 2021. 1, 3
- [66] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125*, 2022. 3
- [67] Leo Sampaio Ferraz Ribeiro, Tu Bui, John Collomosse, and Moacir Ponti. Sketchformer: Transformer-based Representation for Sketched Structure. In *CVPR*, 2020. 1, 2
- [68] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [69] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*, 2015. 3
- [70] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine Tuning Text-to-Image Diffusion Models for Subject-Driven Generation. In *CVPR*, 2023. 3, 4, 7
- [71] Jose M Saavedra. Sketch based image retrieval using a soft computation of the histogram of edge local orientations (SHELO). In *ICIP*, 2014. 2
- [72] Jose M Saavedra, Juan Manuel Barrios, and S Orand. Sketch based Image Retrieval using Learned KeyShapes (LKS). In *BMVC*, 2015. 2
- [73] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *NeurIPS*, 2022. 1, 3
- [74] Aneeshan Sain, Ayan Kumar Bhunia, Yongxin Yang, Tao Xiang, and Yi-Zhe Song. Cross-Modal Hierarchical Modelling for Fine-Grained Sketch Based Image Retrieval. In *BMVC*, 2020. 2, 3
- [75] Aneeshan Sain, Ayan Kumar Bhunia, Yongxin Yang, Tao Xiang, and Yi-Zhe Song. StyleMeUp: Towards Style-Agnostic Sketch-Based Image Retrieval. In *CVPR*, 2021. 2, 3, 4
- [76] Aneeshan Sain, Ayan Kumar Bhunia, Vaishnav Potlapalli, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Sketch3T: Test-Time Training for Zero-Shot SBIR. In *CVPR*, 2022. 2, 6, 7
- [77] Aneeshan Sain, Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. CLIP for All Things Zero-Shot Sketch-Based Image Retrieval, Fine-Grained or Not. In *CVPR*, 2023. 1, 2, 3, 5, 6, 7
- [78] Aneeshan Sain, Ayan Kumar Bhunia, Subhadeep Koley, Pinaki Nath Chowdhury, Soumitri Chattopadhyay, Tao Xiang, and Yi-Zhe Song. Exploiting Unlabelled Photos for Stronger Fine-Grained SBIR. In *CVPR*, 2023. 1, 2, 5
- [79] Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. The sketchy database: learning to retrieve badly drawn bunnies. *ACM TOG*, 2016. 2, 7
- [80] Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. The Sketchy Database: Learning to Retrieve Badly Drawn Bunnies. *ACM TOG*, 2016. 4, 5, 6, 7, 8
- [81] Patsorn Sangkloy, Wittawat Jitkrittum, Diyi Yang, and James Hays. A Sketch Is Worth a Thousand Words: Image Retrieval with Text and Sketch. In *ECCV*, 2022. 2, 3, 8
- [82] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs. In *NeurIPS*, 2021. 1
- [83] 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. In *NeurIPS*, 2022. 1
- [84] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing Across Domains via Cross-Gradient Training. In *ICLR*, 2018. 7
- [85] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In *EMNLP*, 2019. 5
- [86] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, 2015. 4, 6, 8
- [87] Jifei Song, Yi-Zhe Song, Tony Xiang, and Timothy M Hospedales. Fine-Grained Image Retrieval: the Text/Sketch Input Dilemma. In *BMVC*, 2017. 2, 8
- [88] Jifei Song, Qian Yu, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. Deep Spatial-Semantic Attention for Fine-Grained Sketch-Based Image Retrieval. In *ICCV*, 2017. 2
- [89] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent Correspondence from Image Diffusion. In *NeurIPS*, 2023. 1, 3, 4, 6, 7
- [90] Jialin Tian, Xing Xu, Zheng Wang, Fumin Shen, and Xin Liu. Relationship-Preserving Knowledge Distillation for Zero-Shot Sketch Based Image Retrieval. In *ACM MM*, 2021. 2
- [91] Jialin Tian, Xing Xu, Fumin Shen, Yang Yang, and Heng Tao Shen. TVT: Three-Way Vision Transformer through Multi-Modal Hypersphere Learning for Zero-Shot Sketch-Based Image Retrieval. In *AAAI*, 2022. 2
- [92] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-Play Diffusion Features for Text-Driven Image-to-Image Translation. In *CVPR*, 2023. 1, 3

- [93] Hao Wang, Cheng Deng, Tongliang Liu, and Dacheng Tao. Transferable Coupled Network for Zero-Shot Sketch-Based Image Retrieval. *IEEE TPAMI*, 2021. 2
- [94] Kai Wang, Yifan Wang, Xing Xu, Xin Liu, Weihua Ou, and Huimin Lu. Prototype-based Selective Knowledge Distillation for Zero-Shot Sketch Based Image Retrieval. In *ACM MM*, 2022. 2
- [95] Luo Wang, Xueming Qian, Xingjun Zhang, and Xingsong Hou. Sketch-Based Image Retrieval With Multi-Clustering Re-Ranking. *IEEE TCSVT*, 2019. 2
- [96] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction without Convolutions. In *ICCV*, 2021. 8
- [97] Chenfeng Xu, Huan Ling, Sanja Fidler, and Or Litany. 3DiffTecton: 3D Object Detection with Geometry-Aware Diffusion Features. *arXiv preprint arXiv:2311.04391*, 2023. 3
- [98] Jiaqing Xu, Haifeng Sun, Qi Qi, Jingyu Wang, Ce Ge, Lejian Zhang, and Jianxin Liao. DLA-Net for FG-SBIR: Dynamic Local Aligned Network for Fine-Grained Sketch-Based Image Retrieval. In *ACM MM*, 2021. 2
- [99] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-Vocabulary Panoptic Segmentation with Text-to-Image Diffusion Models. In *CVPR*, 2023. 1, 3, 4, 7
- [100] Peng Xu, Yongye Huang, Tongtong Yuan, Kaiyue Pang, Yi-Zhe Song, Tao Xiang, Timothy M Hospedales, Zhanyu Ma, and Jun Guo. SketchMate: Deep Hashing for Million-Scale Human Sketch Retrieval. In *CVPR*, 2018. 2
- [101] Shuai Yang, Zhangyang Wang, Jiaying Liu, and Zongming Guo. Deep Plastic Surgery: Robust and Controllable Image Editing with Human-Drawn Sketches. In *ECCV*, 2020. 2
- [102] Sasi Kiran Yelamarthi, Shiva Krishna Reddy, Ashish Mishra, and Anurag Mittal. A Zero-Shot Framework for Sketch Based Image Retrieval. In *ECCV*, 2018. 1, 2, 3, 5, 6, 7
- [103] Qian Yu, Feng Liu, Yi-Zhe Song, Tao Xiang, Timothy M Hospedales, and Chen-Change Loy. Sketch Me That Shoe. In *CVPR*, 2016. 1, 2, 4, 5, 6
- [104] Hua Zhang, Peng She, Yong Liu, Jianhou Gan, Xiaochun Cao, and Hassan Foroosh. Learning Structural Representations via Dynamic Object Landmarks Discovery for Sketch Recognition and Retrieval. *IEEE TIP*, 2019. 2
- [105] Lvmin Zhang and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. In *ICCV*, 2023. 1, 3
- [106] Zhaolong Zhang, Yuejie Zhang, Rui Feng, Tao Zhang, and Weiguo Fan. Zero-Shot Sketch-Based Image Retrieval via Graph Convolution Network. In *AAAI*, 2020. 2
- [107] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional Prompt Learning for Vision-Language Models. In *CVPR*, 2022. 5
- [108] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to Prompt for Vision-Language Models. *IJCV*, 2022. 5, 8