ECLIPSE: A Resource-Efficient Text-to-Image Prior for Image Generations

Maitreya Patel, Changhoon Kim, Sheng Cheng, Chitta Baral, Yezhou Yang
Arizona State University
{maitreya.patel, kch, scheng53, chitta, yz.yang}@asu.edu

Abstract

Text-to-image (T2I) diffusion models, notably the unCLIP models (e.g., DALL-E-2), achieve state-of-the-art (SOTA) performance on various compositional T2I benchmarks, at the cost of significant computational resources. The unCLIP stack comprises T2I prior and diffusion image decoder. The T2I prior model alone adds a billion parameters compared to the Latent Diffusion Models, which increases the computational and high-quality data requirements. We introduce ECLIPSE[1]*, a novel contrastive learning method that is both parameter and data-efficient. ECLIPSE leverages pre-trained vision-language models (e.g., CLIP) to distill the knowledge into the prior model. We demonstrate that the ECLIPSE trained prior, with only 3.3% of the parameters and trained on a mere 2.8% of the data, surpasses the baseline T2I priors with an average of 71.6% preference score under resource-limited setting. It also attains performance on par with SOTA big models, achieving an average of 63.36% preference score in terms of the ability to follow the text compositions. Extensive experiments on two unCLIP diffusion image decoders, Karlo and Kandinsky, affirm that ECLIPSE priors consistently deliver high performance while significantly reducing resource dependency. Project page: https://eclipse-t2i.vercel.app/*

1. Introduction

Diffusion models [12, 35, 37, 42] have demonstrated remarkable success in generating high-quality images conditioned on text prompts. This Text-to-Image (T2I) generation paradigm has been effectively applied to various downstream tasks such as subject/segmentation/depth-driven image generation [3, 5, 9, 20, 29]. Central to these advancements are two predominant text-conditioned diffusion models: Latent Diffusion Models (LDM) [37], also known as Stable Diffusion, and unCLIP models [35]. The LDM, notable for its open-source availability, has gained widespread popularity within the research community. On the other hand, unCLIP models have remained under-studied. Both model types fundamentally focus on training the diffusion models conditioned on text prompts. The LDM contains a singular text-to-image diffusion model, while unCLIP models have a text-to-image prior, and a diffusion image decoder. Both model families work within the vector quantized latent space of the image [43]. In this paper, we focus on unCLIP models because they consistently outperform other SOTA models in various composition benchmarks such as T2I-CompBench [13] and HRS-Benchmark [1].

These T2I models, typically large in parameter count, require massive amounts of high-quality image-text pairs for training. unCLIP models like DALL-E-2 [35], Karlo [7], and Kandinsky [36], feature prior module containing approximately 1 billion parameters, resulting in a significant increase in overall model size ($\geq$ 2B) compared to LDMs. These unCLIP models are trained on 250M, 115M, and 177M image-text pairs, respectively. Therefore, two critical

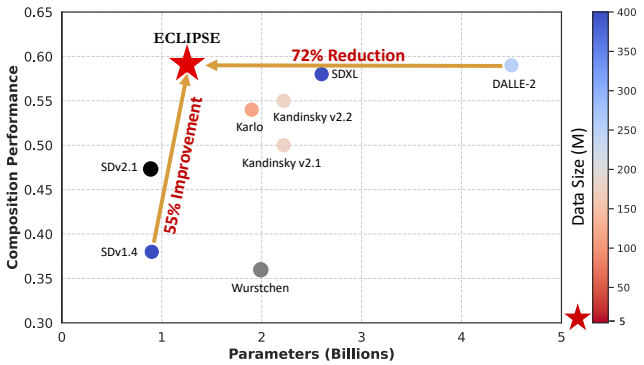

Figure 1. Comparison between SOTA text-to-image models with respect to their total number of parameters and the average performance on the three composition tasks (color, shape, and texture). *ECLIPSE* achieves better results with less number of parameters without requiring a large amount of training data. The shown *ECLIPSE* trains a T2I prior model (having only 33M parameters) using only 5M image-text pairs with Kandinsky decoder.

[1] Our strategy, *ECLIPSE*, draws an analogy from the way a smaller prior model, akin to a celestial entity, offers a glimpse of the grandeur within the larger pre-trained vision-language model, mirroring how an eclipse reveals the vastness of the cosmos.

questions remain: *1) Does the incorporation of a text-to-image prior contribute to SOTA performance on text compositions? 2) Or is scaling up model size the key factor?* In this study, we aim to deepen the understanding of T2I priors and propose substantial enhancements to existing formulations by improving parameter and data efficiency.

As proposed by Ramesh et al. [35], T2I priors are also diffusion models, which are designed to directly estimate the noiseless image embedding at any timestep of the diffusion process. We perform an empirical study to analyze this prior diffusion process. We find that this diffusion process has a negligible impact on generating accurate images and having the diffusion process slightly hurts the performance. Moreover, diffusion models require substantial GPU hours/days for training due to the slower convergence. Therefore, in this work, we use the non-diffusion model as an alternative. While this approach may reduce the compositional capabilities due to the absence of classifier-free guidance [11], it significantly enhances parameter efficiency and decreases the dependencies on the data.

To overcome the above limitations, in this work, we introduce *ECLIPSE*, a novel contrastive learning strategy to improve the T2I non-diffusion prior. We improve upon the traditional method of maximizing the Evidence Lower Bound (ELBO) for generating the image embedding from the given text embedding. We propose to utilize the semantic alignment (between the text and image) property of the pre-trained vision-language models to supervise the prior training. Utilizing *ECLIPSE*, we train compact (97% smaller) non-diffusion prior models (having 33 million parameters) using a very small portion of the image-text pairs (0.34% - 8.69%). We train *ECLIPSE* priors for two unCLIP diffusion image decoder variants (Karlo and Kandinsky). The *ECLIPSE*-trained priors significantly surpass baseline prior learning strategies and rival the performance of 1 billion parameter counterparts. Our results indicate a promising direction for T2I generative models, achieving better compositionality without relying on extensive parameters or data. As illustrated in Fig. 1, by simply improving the T2I prior for unCLIP families, their overall parameter and data requirements drastically reduce and achieve the SOTA performance against similar parameter models.

Contributions. 1) We introduce *ECLIPSE*, the first attempt to employ contrastive learning for text-to-image priors in the unCLIP framework. 2) Through extensive experimentation, we demonstrate *ECLIPSE*'s superiority over baseline priors in resource-constrained environments. 3) Remarkably, *ECLIPSE* priors achieve comparable performance to larger models using only 2.8% of the training data and 3.3% of the model parameters. 4) We also analyze and offer empirical insights on the shortcomings of existing T2I diffusion priors.

2. Related Works

Text-to-Image Generative Models. Advancements in vector quantization and diffusion modeling have notably enhanced text-to-image generation capabilities. Notable works like DALL-E [34] have leveraged transformer models trained on quantized latent spaces. Contemporary state-of-the-art models, including GLIDE [26], Latent Diffusion Model (LDM) [37], DALL-E-2 [35], and Imagen [38], have significantly improved over earlier approaches like StackGAN [47] and TReCS [19]. As these models achieve remarkable photorealism, several works focus on making T2I models more secure [8, 16, 17, 27]. LDM models primarily focus on unified text-to-image diffusion models that incorporate the cross-attention layers [37]. Additionally, several studies aim at refining Stable Diffusion models during inference through targeted post-processing strategies [3, 5, 32]. In contrast, unCLIP models, exemplified by DALL-E-2 [15], Karlo [7], and Kandinsky [36], incorporate a two-step process of text-to-image diffusion transformer prior model and diffusion image decoder having the same model architecture as LDMs. Recent benchmarks have highlighted the superior compositional capabilities of DALL-E-2 over LDM methods [1, 13]. Our work examines and enhances existing prior learning strategies in open-source pre-trained unCLIP models, Karlo and Kandinsky.

Efficient Text-to-Image Models. The current generation of T2I models is characterized by extensive parameter sizes and demanding training requirements, often necessitating thousands of GPU days. Research efforts have primarily centered on model refinement through knowledge distillation, step distillation, and architectural optimization [21, 25, 39]. Wuerstchen [31] presents an efficient unCLIP stack requiring fewer training time GPU hours. Concurrently, Pixart-α [4] leverages pre-trained Diffusion-Transformers (DiT) [30] as base diffusion models, further reducing training time. Distinctively, *ECLIPSE* focuses on refining text-to-image priors within the unCLIP framework using a mere 3.3% of the original model parameters, thereby significantly reducing the training duration to approximately 200 GPU hours. Our work falls orthogonal to the existing efficient T2I methodologies that mainly focus on knowledge and step distillation, and/or architectural compression. When integrated with these model compression strategies, *ECLIPSE* can position the unCLIP family models as a compact yet highly accurate and efficient T2I generation methodology.

Contrastive Learning in Generative Models. Contrastive learning, traditionally applied in visual discriminative tasks, has seen utilization in image-text alignment models like CLIP [33], LiT [45], and SigLIP [46]. However, its application in generative models, particularly in Generative Adversarial Networks (GANs), remains limited [6, 22, 48].

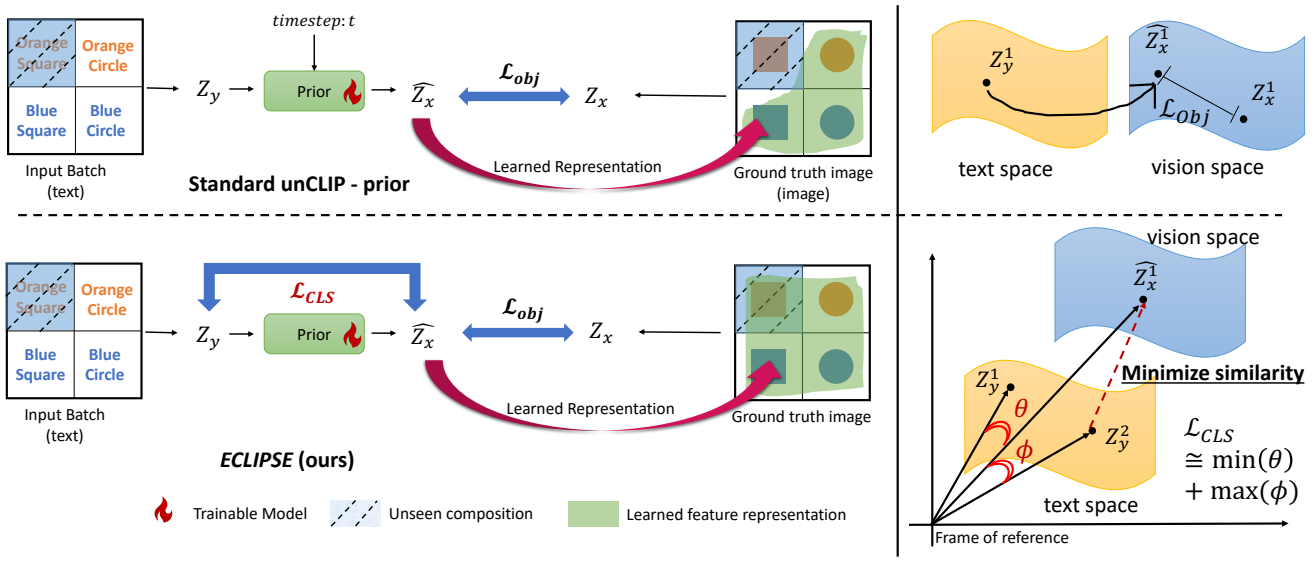

Figure 2. Standard T2I prior learning strategies (top) minimizes the mean squared error between the predicted vision embedding $\hat{z}_x$ w.r.t. the ground truth embedding z_x with or without time-conditioning. This methodology cannot be generalized very well to the outside training distribution (such as Orange Square). The proposed *ECLIPSE* training methodology (bottom) utilizes the semantic alignment property between z_x and z_y with the use of contrastive learning, which improves the text-to-image prior generalization.

For instance, Lafite [48] employs a contrastive approach for image-to-text prior training in language-free T2I GANs. StyleT2I [22] attempts to learn the latent edit direction for StyleGAN [14], which is supervised via spatial masks on the images making the method not scalable. ACTIG [6] introduces an attribute-centric contrastive loss to enhance discriminator performance. These methods are constrained by their domain-specific knowledge requirements and inability to be directly applied to diffusion models [6, 22]. In contrast, *ECLIPSE* applies CLIP-based contrastive learning to train more effective T2I prior models in diffusion-based T2I systems. This strategy is not only resource-efficient but significantly enhances the traditional text-to-image diffusion priors by exploiting the semantic latent space of pre-trained vision-language models.

3. Methodology

This section elaborates on the Text-to-Image (T2I) methodologies, beginning with an overview of unCLIP, followed by the formal problem statement. We then delve into our proposed training strategy, *ECLIPSE*, for T2I prior in detail. Figure 2 provides the overview of baselines and *ECLIPSE* training strategies.

3.1. Preliminaries

Without the loss of generality, let's assume that $y \in Y$ denotes the raw text and $x \in X$ denotes the raw image. z_x and z_y denote the image and text latent embeddings extracted using the pre-trained vision and text encoders ($z_x = C_{vision}(x); \quad z_y = C_{text}(y)$). Ideally, these C_{text} and C_{vision} can be any model (e.g., T5-XXL, ViT, and CLIP). Both model families (LDM and unCLIP) fundamentally focus on learning a mapping function $f_\theta : Y \rightarrow X$. The LDMs contain a singular text-to-image decoder model (f_θ), while unCLIP framework ($f_\theta = h_\theta \circ g_\phi$) contains two primary modules:

- **Text-to-Image Prior** ($g_\phi : z_y \rightarrow z_x$): This module maps the text embeddings to the corresponding vision embeddings. Ramesh et al. [35] showed that the diffusion model as T2I prior leads to slightly better performance than the autoregressive models. For each timestep t and a noised image embedding $z_x^{(t)} \sim q(t, z_x)$ (here, q is a forward diffusion process), the diffusion prior directly estimates noiseless z_x rather than estimating Gaussian noise distribution $\epsilon \sim \mathcal{N}(0, \mathcal{I})$ as:

$$\mathcal{L}_{prior} = \mathop{\mathbb{E}}_{\substack{t\sim[0,T],\\ z_x^{(t)}\sim q(t,z_x)}} \left[||z_x - g_\phi(z_x^{(t)}, t, z_y)||_2^2 \right]. \quad (1)$$

- **Diffusion Image Decoder** ($h_\theta : (z_x, z_y) \rightarrow x$): This module generates the final image conditioned on the z_x and the input text features z_y. This diffusion decoder follows the standard diffusion training procedure by estimating $\epsilon \sim \mathcal{N}(0, \mathcal{I})$ after [12]:

$$\mathcal{L}_{decoder} = \mathop{\mathbb{E}}_{\substack{\epsilon\sim N(0,I)\\ t\sim[0,T],\\ (z_x,\ z_y)}} \left[||\epsilon - h_\theta(x^{(t)}, t, z_x, z_y)||_2^2 \right]. \quad (2)$$

Specifically, different versions of the unCLIP decoder, such as Kandinsky and Karlo, vary in whether they include

text conditioning (z_y) in the diffusion image decoder. However, both approaches yield comparable results, provided that image conditioning (z_x) is accurate. Both training objectives, $\mathcal{L}_{prior}$ and $\mathcal{L}_{decoder}$, integrate Classifier-Free Guidance (CFG) [11], enhancing the model's generative capabilities. During training, conditions are omitted 10% of the time to foster unconditional generation, subsequently improving test performance as CFG works as implicit classifier guidance [11].

3.2. Problem Formulation

Given the pivotal role of the T2I prior module in image generation from text, in this paper, our focus is on enhancing g_ϕ, while keeping the pre-trained h_θ frozen. Let's consider a training distribution P_{XY}, comprising input pairs of image and text (x, y). Maximizing the Evidence Lower Bound (ELBO) on the training distribution P_{XY} facilitates this mapping of $z_y \rightarrow z_x$. However, such a strategy does not inherently assure generalization, especially when the input text prompt (y) deviates from the assumed independently and identically distributed (i.i.d.) pattern of P_{XY} [44]. Therefore, attaining a more diverse and representative P_{XY} becomes crucial for improving the performance. While a diffusion prior combined with CFG has been shown to bolster generalization, especially with diverse training data and extensive training iterations [28], it is computationally expensive and is not always reliable (especially, in low resource constraint settings) as shown in Section 4.2. Given these constraints, our goal is to develop an alternative prior learning methodology that not only improves parameter efficiency (97% reduction) and mitigates the need for large-scale high-quality data ($\leq$ 5%) but also upholds performance levels.

3.3. Proposed Method: *ECLIPSE*

This section elaborates on *ECLIPSE*, our model training strategy to learn text-to-image prior (g_ϕ). We focus on enhancing non-diffusion prior models through the effective distillation of pre-trained vision-language models, such as CLIP, while preserving the semantic alignment between the input text embedding z_y and corresponding estimated vision embeddings $\hat{z}_x$ by using the contrastive loss.

Base Prior Model. T2I diffusion prior deviates from the standard diffusion objective (such as Eq. 2). Unlike the standard $\epsilon \sim \mathcal{N}(0, \mathcal{I})$ prediction-based diffusion objective, the T2I diffusion prior objective (Eq. 1) do not compare two Gaussian distributions, instead, it directly estimates the z_x which is noiseless. However, during inference, we still adhere to the conventional denoising process, introducing additional noise ($\sigma_t \epsilon$) at each step, except for the final step according to Ho et al. [12]. This creates a new input distribution ($z_x + \sigma_t \epsilon$), possibly unencountered during training. Moreover, if we repeat this for T timesteps, it can lead to the accumulation of errors, which is undesirable. We provide empirical analysis in Section 5 to ground this hypothesis, where we show that having more diffusion prior steps does not benefit the overall text-to-image generation abilities.

Therefore, to mitigate this unnecessary computing, we use non-diffusion T2I prior, making the prior model both parameter-efficient and less demanding in terms of computational resources. This non-diffusion architecture forms our base model, and we introduce the training objective that leverages pre-trained vision-language models trained on extensive datasets to improve generalization outside the P_{XY} distribution.

Projection Objective. Despite vision-language models aligning the semantic distributions across modalities, each modality may exhibit unique distributions. Therefore, our approach involves projecting the text embedding onto the vision embedding. This is achieved using a mean squared error objective between the predicted vision embedding ($\hat{z}_x$) and the ground truth vision embedding (z_x):

$$\mathcal{L}_{proj} = \underset{\substack{\epsilon \sim \mathcal{N}(0,I) \\ z_y, z_x}}{\mathbb{E}} \Big[||z_x - g_\phi(\epsilon, z_y)||_2^2 \Big], \tag{3}$$

where ϵ is the Gaussian input noise. Notably, as discussed previously, this is an approximation of the diffusion prior objective (Eq. 1) with $t = T$ and without CFG. $\mathcal{L}_{proj}$ learns latent posterior distribution with the *i.i.d.* data assumption. However, this model, fine-tuned on P_{XY}, may not generalize well beyond its distribution. The optimal solution would be to train on a dataset that encapsulates all potential distributions to cover all possible scenarios, which is an impractical and resource-consuming task.

CLIP Contrastive Learning. To address these limitations, we propose utilizing the CLIP more effectively, which contains the semantic alignment between image and language. Specifically, we apply the CLIP Contrastive Loss after [33] to train the T2I priors. For a given input batch $\{(z_x^i, z_y^i)\}_{i=1}^N$ from the P_{XY} distribution, we calculate the text-conditioned image contrastive loss for the ith image embedding prediction relative to the all input ground truth text embeddings as:

$$\mathcal{L}_{CLS;\, y \rightarrow x} = -\frac{1}{N} \sum_{i=0}^{N} \log \frac{\exp(\langle \hat{z}_x^i, z_y^i \rangle / \tau)}{\sum_{j \in [N]} \exp(\langle \hat{z}_x^i, z_y^j \rangle / \tau)}, \tag{4}$$

where τ is the temperature parameter, $\langle , \rangle$ denotes the cosine similarity, and N is the batch size. This loss encourages the model to understand and follow the input text better, effectively reducing overfitting to the P_{XY}, as illustrated in Figure 2. Consequently, the final objective function is:

$$\mathcal{L}_{ECLIPSE} = \mathcal{L}_{proj} + \lambda * \mathcal{L}_{CLS;\, y \rightarrow x}, \tag{5}$$

Table 1. The comparison (in terms of FID and compositions) of the baselines and state-of-the-art methods with respect to the *ECLIPSE*. * indicates the official reported ZS-FID. Ψ denotes the FID performance of a model trained on MSCOCO. The best performing *ECLIPSE* variant (with respect to its big counterpart) is highlighted by green. *ECLIPSE* consistently outperforms the SOTA big models despite being trained on a smaller subset of dataset and parameters.

Methods	Model Type	Training Params [M]*	Total Params [B]	Data Size [M]	ZS-FID (↓)	T2I-CompBench Color (↑)	Shape (↑)	Texture (↑)	Spatial (↑)	Non-Spatial (↑)
Stable Diffusion v1.4	LDM	900	0.9	400	16.31*	0.3765	0.3576	0.4156	0.1246	0.3076
Stable Diffusion v2.1	LDM	900	0.9	2000	14.51*	0.5065	0.4221	0.4922	0.1342	0.3096
Wurstchen	unCLIP	1000	2.0	1420	23.60*	0.3216	0.3821	0.3889	0.0696	0.2949
Kandinsky v2.1	unCLIP	1000	2.22	177	18.09	0.4647	0.4725	0.5613	0.1219	0.3117
DALL-E-2	unCLIP	1000	4.5	250	10.65*	0.5750	0.5464	0.6374	0.1283	0.3043
Karlo	unCLIP	1000	1.9	115	20.64	0.5127	0.5277	0.5887	0.1337	0.3112
***ECLIPSE* (ours)**	Karlo	**33**	**0.93**	**0.6**$_{\text{MSCOCO}}$	23.67$^{\Psi}$	0.5965	0.5063	0.6136	0.1574	0.3235
		33	**0.93**	**2.5**$_{\text{CC3M}}$	26.73	0.5421	0.5090	0.5881	0.1478	0.3213
		33	**0.93**	**10.0**$_{\text{CC12M}}$	26.98	0.5660	0.5234	0.5941	0.1625	0.3196
Kandinsky v2.2	unCLIP	1000	2.22	177	20.48	0.5768	0.4999	0.5760	0.1912	0.3132
***ECLIPSE* (ours)**	Kandinsky v2.2	**34**	**1.26**	**0.6**$_{\text{MSCOCO}}$	16.53$^{\Psi}$	0.5785	0.4951	0.6173	0.1794	0.3204
		34	**1.26**	**5.0**$_{\text{HighRes}}$	19.16	0.6119	0.5429	0.6165	0.1903	0.3139

where λ is the hyperparameter balancing the regularizer's effect. Overall, the final objective function aims to map the text latent distribution to the image latent distribution via $\mathcal{L}_{proj}$ and such that it preserves the image-text alignment using $\mathcal{L}_{CLS;\, y \to x}$. This makes the prior model generalize beyond the given training distribution P_{XY} such that it can follow the semantic alignment constraint. Importantly, we cannot use $\mathcal{L}_{CLS;\, y \to x}$ alone or with a high value of λ as the prior model will converge outside the vision latent distribution that optimizes the contrastive loss (such input text latent space itself). And keeping λ to a very low value cannot do knowledge distillation well enough. Empirical studies suggest setting $\lambda = 0.2$ for optimal performance, balancing knowledge distillation, and maintaining alignment within the vision latent distribution.

4. Experiments & Results

This section introduces the datasets, training specifications, comparative baselines, and evaluation metrics utilized in our experiments. We conduct an extensive assessment of our proposed *ECLIPSE* methodology and its variants, both quantitatively and qualitatively.

4.1. Experimental Setup

Dataset. Our experiments span four datasets of varying sizes: MSCOCO [23], CC3M [41], CC12M [2], and LAION-HighResolution[2] [40]. MSCOCO comprises approximately 0.6 million image-text pairs, while CC3M and CC12M contain around 2.5 and 10 million pairs, respectively [3]. We select a very small subset of 5 million (2.8%) image-text pairs from the LAION-HighRes dataset (175M).

[2] https://huggingface.co/datasets/laion/laion-high-resolution

[3] According to the download date: 08/26/2023

We perform Karlo diffusion image decoder-related experiments on MSCOCO, CC3M, and CC12M as these datasets are subsets of the data used to train the Karlo diffusion image decoder. Similarly, we use MSCOCO and LAION-HighRes for the Kandinsky decoder.

Baselines. *ECLIPSE* variants are compared against leading T2I models, including Stable Diffusion, Wurstchen, Karlo, Kandinsky, and DALL-E-2. Additionally, we introduce two more baselines along with *ECLIPSE* to evaluate the impact of our training strategy in a resource-constrained environment: 1) Projection: A non-diffusion prior model trained with $\mathcal{L}_{proj}$ (Eq. 3). 2) Diffusion-Baseline: A diffusion prior model trained with $\mathcal{L}_{prior}$ (Eq. 1) – the traditional T2I prior, and 3) *ECLIPSE*: A non-diffusion prior model trained with our proposed methodology $\mathcal{L}_{ECLIPSE}$ (Eq. 5).

Training and inference details. We evaluate *ECLIPSE* using two pre-trained image decoders: Karlo-v1-alpha and Kandinsky v2.2, trained on distinct CLIP vision encoders. Our prior architecture is based on the standard PriorTransformer model [35], modified to be time-independent. The detailed architecture is outlined in the appendix. We configure prior models with 33 and 34 million parameters for Karlo and Kandinsky, respectively. This contrasts with larger models in the field, which often use up to 1 billion parameters (as summarized in Table 1). The Projection, Diffusion-Baseline, and *ECLIPSE* priors are trained for both diffusion image decoders, maintaining consistent hyperparameters (including total number of parameters) across all models. Training on CC12M, CC3M, and LAION-HighRes is performed on 4 x RTX A6000 GPUs with a 256 per-GPU batch size, a learning rate of 0.00005, and the CosineAnnealingWarmRestarts scheduler [24]. Each model undergoes approximately 60,000 iterations, totaling around 200 GPU hours. For MSCOCO, training takes about 100 GPU hours. This can be further

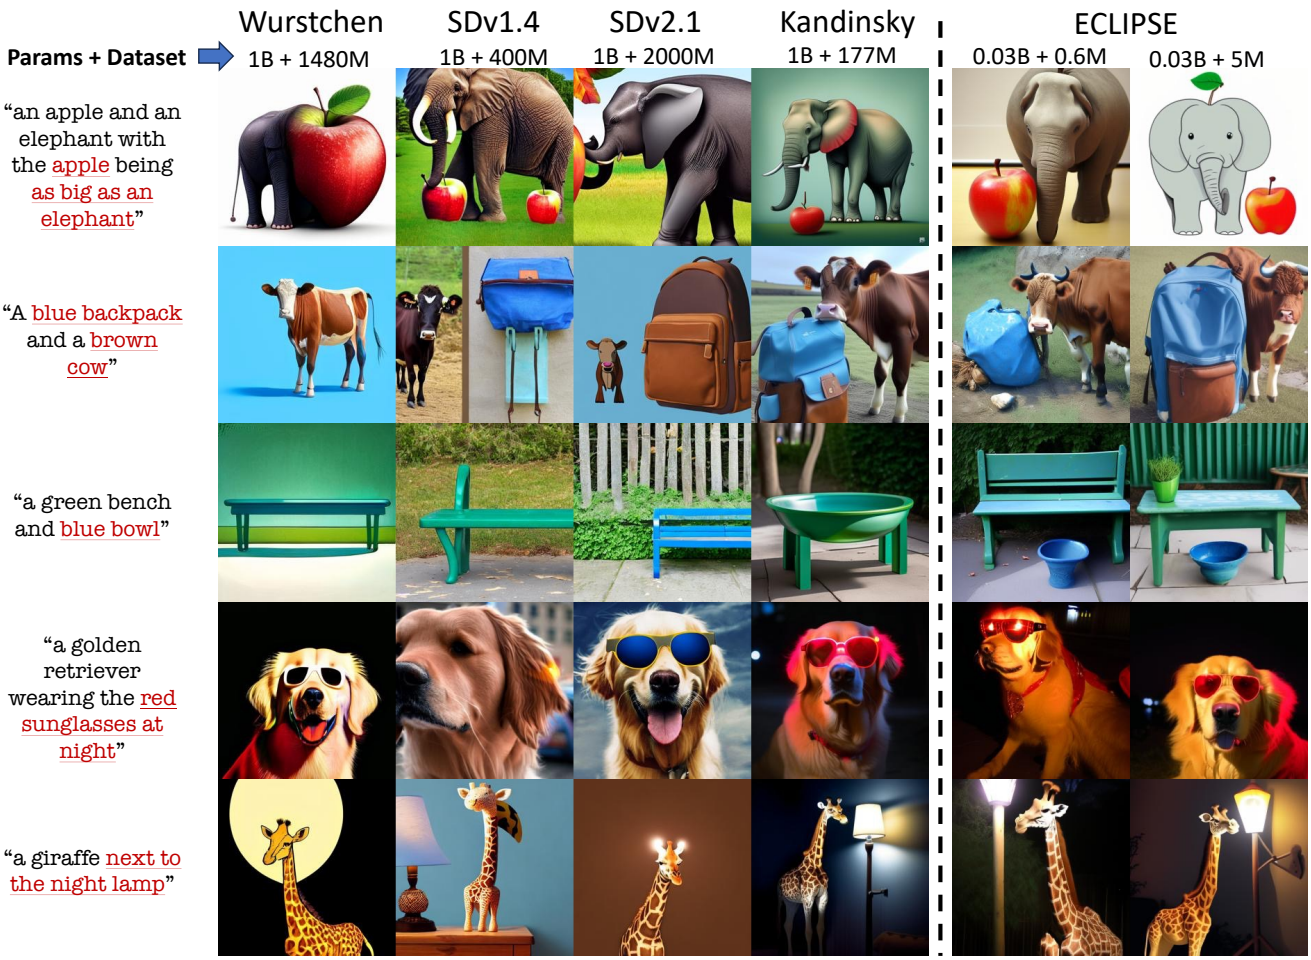

Figure 3. Qualitative result of our text-to-image prior, *ECLIPSE*, comparing with SOTA T2I model. Our prior model reduces the model parameter requirements (from 1 Billion → 33 Million) and data requirements (from 177 Million → 5 Million → 0.6 Million). Given this restrictive setting, *ECLIPSE* performs close to its huge counterpart (i.e., Kandinsky v2.2) and even outperforms models trained on huge datasets (i.e., Wurstchen, SDv1.4, and SDv2.1) in terms of compositions.

reduced to $\leq$ 50 GPU hours if image-text pairs are preprocessed beforehand. The diffusion prior is trained with a linear scheduler and 1000 DDPM timesteps. Inferences utilize 25 DDPM steps with 4.0 classifier-free guidance, while Projection and *ECLIPSE* models do not require diffusion sampling. Image diffusion decoders are set to 50 DDIM steps and 7.5 classifier-free guidance.

Evaluation setup. Our evaluation framework encompasses various metrics. We employ MS-COCO 30k to assess FID scores [10] and T2I-CompBench [13] for evaluating composition abilities in color, shape, texture, spatial, and non-spatial compositions. Given the impracticality of large-scale human studies, we approximate human preferences using PickScore [18], reporting results on the T2I-CompBench validation set comprising about 1500 unique prompts across different categories.

4.2. Quantitative Evaluations

In Table 1, we present a performance comparison between *ECLIPSE* variants and leading T2I models. Our evaluation metrics include zero-shot Fréchet Inception Distance (FID) on MS-COCO 30k for image quality assessment and T2I-CompBench [13] for evaluating compositionality. *ECLIPSE* priors, trained with both types of diffusion image decoders, demonstrate notable improvements. *ECLIPSE* consistently surpasses various baselines in terms of compositionality, irrespective of the dataset size. Its performance is comparable to that of DALL-E-2 and other SOTA models, a significant improvement considering *ECLIPSE*'s parameter efficiency. Standard T2I priors usually incorporate 1 billion parameters, while *ECLIPSE* operates with only 3.3% of these parameters, maintaining competitive performance levels. When combined with cor-

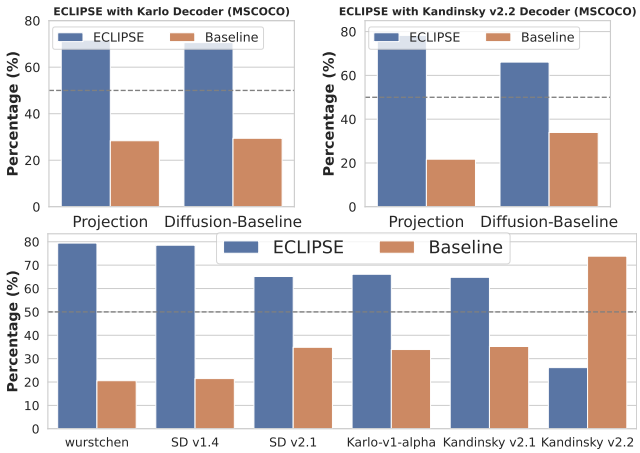

Figure 4. Qualitative evaluations by human preferences approximated by the PickScore [18]. The top two figures compare *ECLIPSE* to Projection and Diffusion Baselines trained with the same amount of data and model size for both Karlo and Kandinsky decoders. In the bottom figure, we compare *ECLIPSE* with the Kandinsky v2.2 decoder trained on the LAION-HighRes dataset against SOTA models.

Table 2. Comparison of *ECLIPSE* with respect to the various baseline prior learning strategies on four categories of composition prompts in the T2I-CompBench. All prior models are of 33 million parameters and trained on the same hyperparameters.

Methods	T2I-CompBench			
	Color (↑)	Shape (↑)	Texture (↑)	Spatial (↑)
MSCOCO with Karlo				
Projection	0.4667	0.4421	0.5051	0.1478
Diffusion-Baseline	0.4678	0.4797	0.4956	0.1240
ECLIPSE	**0.5965**	**0.5063**	**0.6136**	**0.1574**
CC3M with Karlo				
Projection	0.4362	0.4501	0.4948	0.1126
Diffusion-Baseline	**0.5493**	0.4809	0.5462	0.1132
ECLIPSE	0.5421	**0.5091**	**0.5881**	**0.1477**
CC12M with Karlo				
Projection	0.4659	0.4632	0.4995	0.1318
Diffusion-Baseline	0.5390	0.4919	0.5276	0.1426
ECLIPSE	**0.5660**	**0.5234**	**0.5941**	**0.1625**
MSCOCO with Kandinsky v2.2				
Projection	0.4678	0.3736	0.4634	0.1268
Diffusion-Baseline	0.4646	0.4403	0.4834	0.1566
ECLIPSE	**0.5785**	**0.4951**	**0.6173**	**0.1794**
HighRes with Kandinsky v2.2				
Projection	0.5379	0.4983	0.5217	0.1573
Diffusion-Baseline	0.5706	0.5182	0.5067	0.1687
ECLIPSE	**0.6119**	**0.5429**	**0.6165**	**0.1903**

responding diffusion image decoders, the total parameter count of *ECLIPSE* is close to that of Stable Diffusion models, yet it outperforms them, especially considering that the latter are trained on a massive set of image-text pairs. A noticeable decline in zero-shot FID (ZS-FID) is observed in comparison to the original Karlo. We attribute this variation to the image quality differences in the training dataset, suggesting a potential area for further investigation and improvement. At the same time, if we utilize the smaller subset of high-resolution datasets then we can still maintain better FID and improve the compositions, as shown in the last row of Table 1. *ECLIPSE* prior with Kandinsky v2.2 decoder trained on LAION-HighRes subset achieves similar FID to other original Kandinsky v2.2 unCLIP model and at the same time outperforming in terms of compositions.

Table 2 provides a comparison of various baseline training strategies for small prior models, using identical datasets and hyperparameters. *ECLIPSE* exhibits superior performance across all datasets. We also note that diffusion priors benefit from larger datasets, supporting our premise that such priors necessitate extensive training data for optimal results, which is also attributed to the CFG. In contrast, *ECLIPSE* demonstrates the consistent performance on compositions irrespective of the amount of image-text pairs.

4.3. Qualitative Evaluations

In Figure 3, we display qualitative examples from various methods responding to complex prompts. *ECLIPSE* demonstrates superior performance in comparison to Stable Diffusion v1.4, Stable Diffusion v2.1, and Wurstchen, while closely matching the quality of its big counterpart, Kandinsky v2.2. Interestingly, *ECLIPSE* trained on only 0.6 million images maintains the compositions with minor degradation in image quality. These observations align with our previously established quantitative results. Beyond numerical metrics, understanding human preferences is crucial. To this end, we selected 1500 unique validation prompts from T2I-CompBench and assessed PickScore preferences. The results, illustrated in Figure 4, reveal that *ECLIPSE* notably surpasses its baselines in respective restrictive settings with an average score of 71.6%. We can also observe that the best *ECLIPSE* variant (with Kandinsky decoder and trained on LAION-HighRes) consistently outperforms the other big SOTA models achieving an average performance of 63.36%. We observe that in terms of preferences, the original Kandinsky v2.2 diffusion prior (with a 1 billion parameter) trained on LAION-HighRes (175M) performs better than the *ECLIPSE* prior (having 33 million parameters). We hypothesize that this might be due to its use of a large-scale dataset that contains more aesthetically pleasing images. We provide a set of qualitative results in the appendix to show that *ECLIPSE* performs similarly well, if not better, *w.r.t.* semantic understanding of the text.

5. Analysis

Analyzing the traditional diffusion priors. To further support our choice of using non-diffusion prior models, we analyze the existing diffusion prior formulation. We conducted two key empirical studies: 1) Evaluating the Impact of Prior Steps: We examined how the number of prior steps

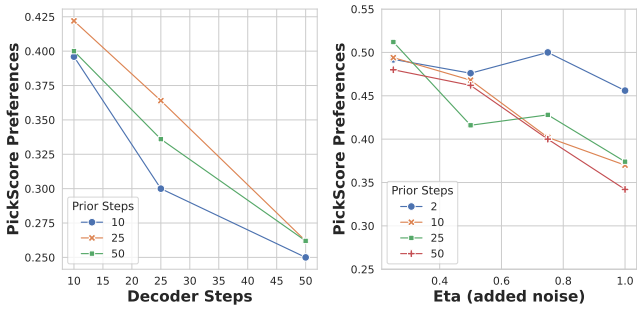

(a) Left: Performance comparison by varying the prior steps and decoder steps *w.r.t.* the fixed prior steps ($t = 2$). Right: Performance comparison by varying the mean η of the added scheduler noise ($\sigma_t \epsilon$) *w.r.t.* the noiseless predictions ($\eta = 0$). Both experiments are on the Kandinsky v2.1.

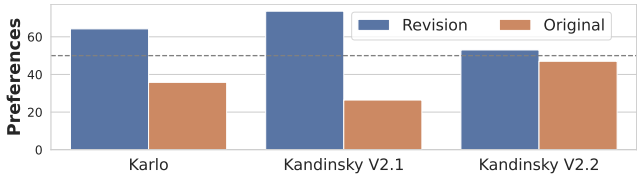

(b) Overall performance comparisons on various pre-trained unCLIP models before and after reducing the prior steps to two and η to 0.0.

Figure 5. Empirical analysis of the PickScore preferences of diffusion priors with respect to the various hyper-parameters.

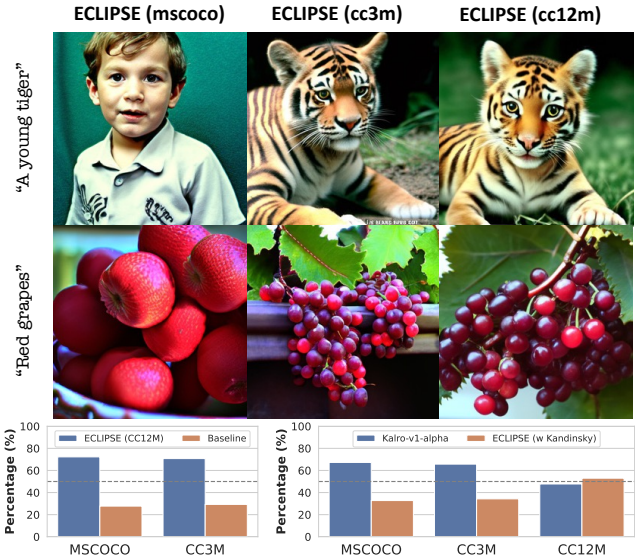

Figure 6. The top figure shows the qualitative examples of the biases learned by the T2I prior models. Bottom figures show the PickScore preferences of the *ECLIPSE* models trained on various datasets with respect to the other datasets (left) and Karlo (right).

influences model performance. 2) Assessing the Influence of Added Noise ($\sigma_t \epsilon$): We focused on understanding how the introduction of noise affects human preferences. For these studies, we utilized PickScore preferences, and the outcomes, depicted in Figure 5, corroborate our hypothesis: both the prior steps and the addition of ($\sigma_t \epsilon$) detrimentally affect performance. Furthermore, as indicated in Table 2, diffusion prior surpasses the projection baseline if provided with more high-quality data. We attribute this enhanced performance to the incorporation of classifier-free guidance, which bolsters the model's generalization capabilities to a certain extent. However, it's worth noting that both baselines are still outperformed by *ECLIPSE*. This observation underscores the effectiveness of our proposed methodology in comparison to traditional approaches in the realm of T2I.

Importance of data selection. In our previous analysis (Table 1 and 2), we demonstrated that *ECLIPSE* attains competitive performance on composition benchmarks regardless of dataset size. This achievement is largely due to the integration of the contrastive loss $\mathcal{L}_{CLS}$ (Eq.4). However, the final objective function also incorporates the $\mathcal{L}_{proj}$ (Eq.3), which is pivotal in estimating the vision latent distribution. This estimation is fundamentally dependent on the training distribution (P_{XY}), potentially leading the model to learn spurious correlations within P_{XY}. Consequently, the model's image quality could directly correlate with the overall quality of images in the training set. To further substantiate this, we evaluated the preferences for *ECLIPSE* models trained on MSCOCO, CC3M, and CC12M, in comparison to among themselves and Karlo-v1-alpha. The outcomes, presented in Figure 6, reveal that the *ECLIPSE* model trained on CC12M outperforms those trained on other datasets, exhibiting performance on par with big counterpart. *ECLIPSE* prior (w Karlo decoder) trained on the CC12M dataset performs comparably to Karlo-v1-alpha while *ECLIPSE* priors trained on other datasets struggle to do so; suggesting the importance of the high-quality data. Furthermore, as illustrated in Figure 6, the *ECLIPSE* model trained on MSCOCO demonstrates a tendency to learn spurious correlations, such as associating the term "young tiger" with the person.

6. Conclusion

In this paper, we introduce a novel text-to-image prior learning strategy, named *ECLIPSE*, which leverages pre-trained vision-language models to provide additional supervision for training the prior model through contrastive learning. This approach significantly enhances the training efficiency of prior models in a parameter-efficient way. Through comprehensive quantitative and qualitative evaluations, we assessed *ECLIPSE* priors alongside various diffusion image decoders. The results indicate that *ECLIPSE* surpasses both the baseline projection models and traditional diffusion-prior models. Remarkably, *ECLIPSE* achieves competitive performance alongside larger, state-of-the-art T2I models. It demonstrates that priors can be trained with merely 3.3% of the parameters and 2.8% of image-text pairs typically re-

quired, without compromising the performance. This advancement directly leads to at least 43% overall compression of the unCLIP models. Our findings show that pre-trained vision-language can be utilized more effectively; suggesting promising research direction where improving the vision-language models may directly benefit the T2I models.

Acknowledgement

This work was supported by NSF RI grants #1750082 and #2132724, and a grant from Meta AI Learning Alliance. The views and opinions of the authors expressed herein do not necessarily state or reflect those of the funding agencies and employers.

References

[1] Eslam Mohamed Bakr, Pengzhan Sun, Xiaogian Shen, Faizan Farooq Khan, Li Erran Li, and Mohamed Elhoseiny. Hrs-bench: Holistic, reliable and scalable benchmark for text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20041–20053, 2023. 1, 2

[2] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. 5

[3] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. 1, 2

[4] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, et al. Pixart-alpha: Fast training of diffusion transformer for photorealistic text-to-image synthesis. *arXiv preprint arXiv:2310.00426*, 2023. 2

[5] Minghao Chen, Iro Laina, and Andrea Vedaldi. Training-free layout control with cross-attention guidance. *arXiv preprint arXiv:2304.03373*, 2023. 1, 2

[6] Yuren Cong, Martin Renqiang Min, Li Erran Li, Bodo Rosenhahn, and Michael Ying Yang. Attribute-centric compositional text-to-image generation. *arXiv preprint arXiv:2301.01413*, 2023. 2, 3

[7] Lee Donghoon, Kim Jiseob, Choi Jisu, Kim Jongmin, Byeon Minwoo, Baek Woonhyuk, and Kim Saehoon. Karlo-v1.0.alpha on coyo-100m and cc15m. https://github.com/kakaobrain/karlo, 2022. 1, 2

[8] Pierre Fernandez, Guillaume Couairon, Hervé Jégou, Matthijs Douze, and Teddy Furon. The stable signature: Rooting watermarks in latent diffusion models. *arXiv preprint arXiv:2303.15435*, 2023. 2

[9] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *The Eleventh International Conference on Learning Representations*, 2022. 1

[10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6

[11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2, 4

[12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 3, 4

[13] Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *arXiv preprint arXiv:2307.06350*, 2023. 1, 2, 6

[14] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 3

[15] Bo-Kyeong Kim, Hyoung-Kyu Song, Thibault Castells, and Shinkook Choi. On architectural compression of text-to-image diffusion models. *arXiv preprint arXiv:2305.15798*, 2023. 2

[16] Changhoon Kim, Yi Ren, and Yezhou Yang. Decentralized attribution of generative models. In *International Conference on Learning Representations*, 2021. 2

[17] Changhoon Kim, Kyle Min, Maitreya Patel, Sheng Cheng, and Yezhou Yang. Wouaf: Weight modulation for user attribution and fingerprinting in text-to-image diffusion models. *arXiv preprint arXiv:2306.04744*, 2023. 2

[18] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *arXiv preprint arXiv:2305.01569*, 2023. 6, 7

[19] Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Text-to-image generation grounded by fine-grained user attention. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 237–246, 2021. 2

[20] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22511–22521, 2023. 1

[21] Yanyu Li, Huan Wang, Qing Jin, Ju Hu, Pavlo Chemerys, Yun Fu, Yanzhi Wang, Sergey Tulyakov, and Jian Ren. Snapfusion: Text-to-image diffusion model on mobile devices within two seconds. *arXiv preprint arXiv:2306.00980*, 2023. 2

[22] Zhiheng Li, Martin Renqiang Min, Kai Li, and Chenliang Xu. Stylet2i: Toward compositional and high-fidelity text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18197–18207, 2022. 2, 3

[23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference,*

Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13, pages 740–755. Springer, 2014. 5

[24] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2016. 5

[25] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023. 2

[26] 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. *arXiv preprint arXiv:2112.10741*, 2021. 2

[27] Guangyu Nie, Changhoon Kim, Yezhou Yang, and Yi Ren. Attributing image generative models using latent fingerprints. *arXiv preprint arXiv:2304.09752*, 2023. 2

[28] Maya Okawa, Ekdeep Singh Lubana, Robert P Dick, and Hidenori Tanaka. Compositional abilities emerge multiplicatively: Exploring diffusion models on a synthetic task. *arXiv preprint arXiv:2310.09336*, 2023. 4

[29] Maitreya Patel, Tejas Gokhale, Chitta Baral, and Yezhou Yang. Conceptbed: Evaluating concept learning abilities of text-to-image diffusion models. *arXiv preprint arXiv:2306.04695*, 2023. 1

[30] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 2

[31] Pablo Pernias, Dominic Rampas, and Marc Aubreville. Wuerstchen: Efficient pretraining of text-to-image models. *arXiv preprint arXiv:2306.00637v2*, 2023. 2

[32] Quynh Phung, Songwei Ge, and Jia-Bin Huang. Grounded text-to-image synthesis with attention refocusing. *arXiv preprint arXiv:2306.05427*, 2023. 2

[33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 4

[34] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 2

[35] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1 (2):3, 2022. 1, 2, 3, 5

[36] Anton Razzhigaev, Arseniy Shakhmatov, Anastasia Maltseva, Vladimir Arkhipkin, Igor Pavlov, Ilya Ryabov, Angelina Kuts, Alexander Panchenko, Andrey Kuznetsov, and Denis Dimitrov. Kandinsky: an improved text-to-image synthesis with image prior and latent diffusion. *arXiv preprint arXiv:2310.03502*, 2023. 1, 2

[37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2

[38] 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. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 2

[39] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2021. 2

[40] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 5

[41] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 5

[42] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 1

[43] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 1

[44] Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991. 4

[45] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18123–18133, 2022. 2

[46] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. *arXiv preprint arXiv:2303.15343*, 2023. 2

[47] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5907–5915, 2017. 2

[48] Yufan Zhou, Ruiyi Zhang, Changyou Chen, Chunyuan Li, Chris Tensmeyer, Tong Yu, Jiuxiang Gu, Jinhui Xu, and Tong Sun. Towards language-free training for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17907–17917, 2022. 2, 3

ECLIPSE: A Resource-Efficient Text-to-Image Prior for Image Generations

Supplementary Material

A. Implementation Details

Table 3 shows the comparison between *ECLIPSE*, Karlo, and Kaninsky priors. Notably, *ECLIPSE* prior uses very compressed architecture across the possible avenues (i.e., number of layers, number of attention heads, attention head dimension, etc.). Karlo uses CLIP-Vit-L/14 with 768 projection dimensions. While Kandinsky v2.2 uses the ViT-bigG-14-laion2B-39B-b160k with 1280 projection dimensions. Overall, the total number of parameters in *ECLIPSE* priors is about 33 million compared to 1 billion parameters of Karlo/Kandinsky priors. Additionally, Projection and Diffusion-Baseline use the same architecture as *ECLIPSE* prior for better comparisons. Except the Diffusion-Prior contains the additional time embeddings for diffusion modeling.

	ECLIPSE	**Karlo / Kandinsky Priors**
Num Attention Heads	16	32
Attention Head Dim	32	64
Num Layers	10	20
Embedding Dim	768/1280	768/1280
Additional Embeddings	3	4
Dropout	0.0	0.0
Time Embed	No	Yes
Total Parameters	33/34 M	1 B

Table 3. Prior model architecture hyperparameter details.

B. Hyper-parameter Analysis

ECLIPSE only contains one important hyperparameter (λ) that controls the contrastive learning. As discussed in Section 3.3, a higher value of λ can make the prior model learn the different distribution that is highly aligned with text distributions. A lower value of λ may not benefit in terms of generalization to unseen prompts. Hence, we conducted a small study on the MSCOCO dataset. We train the *ECLIPSE* priors for Karlo decoder on 20,000 iterations with the OneCycle learning rate. Figure 7 illustrates the pickscore preferences on T2I-CompBench of various values of λ. It can be observed that higher values of λ lead to the same performance as the baseline. While lower values of λ outperform the baseline by significant margins. Additionally, Figure 8 shows one qualitative example across the range of λ. It can be seen that the generated image quality drops as λ increases. **Hence, the optimal range is:** $\lambda \in [0.2, 0.4]$**.**

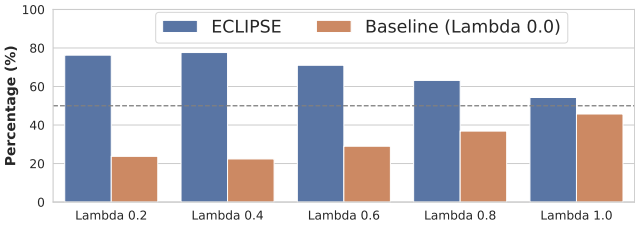

Figure 7. Hyperparameter (λ) ablation. This figure illustrates the PickScore preferences across the *ECLIPSE* priors trained with different values of λ *w.r.t.* the Projection baseline (with $\lambda = 0.0$).

Figure 8. Qualitative example for *ECLIPSE* priors (with Karlo decoder) trained with different values of hyperparameter (λ).

C. *ECLIPSE* Prior Model Scaling Behaviour

To analyze the scaling behavior of different prior learning strategies to a certain extent, we increase the prior model size from 33M to 89M. Table 4 shows the results when small and large priors are trained on the same dataset (CC12M) with the Karlo image diffusion decoder. We train both versions of the prior models on 60,000 iterations (about 350 GPU hours) with exactly the same hyperparameters. First, we observe that ***ECLIPSE* prior improves the performance slightly**. Second, the Projection baseline gets the same performance, which suggests that **data is the bottleneck for the Projection prior**. Third, interestingly Diffusion prior degrades the performance. Upon further inspection, we found that 60,000 iterations are insufficient for the Diffusion model to converge. Therefore, this verifies that **Diffusion-priors are resource-hungry**. Importantly, *ECLIPSE* priors easily converge irrespective of the data and number of parameters; suggesting that *ECLIPSE* do not depend upon the huge resource constraints.

D. Aesthetics: Kandinsky v2.2 *vs.* *ECLIPSE*

As was observed in Figure 4 from the main paper, the Kandinsky v2.2 model outperforms the *ECLIPSE* prior when evaluated in terms of human preferences measured by Pickscore. We attribute this behavior to the differences in the aesthetic quality of the generated images. Therefore, we conduct additional actual human studies to analyze

Table 4. This table illustrates the scaling behavior of various T2I prior learning strategies. "Small" priors are 33 million in terms of parameters. And "Large" priors have 89 million parameters. All prior models are trained on the CC12M dataset with the Karlo diffusion image decoder.

Methods	ZS FID	T2I-CompBench Color (↑)	Shape (↑)	Texture (↑)	Spatial (↑)
33M Priors					
Projection	28.84	0.4659	0.4632	0.4995	0.1318
Diffusion-Baseline	**26.13**	0.5390	0.4919	0.5276	0.1426
ECLIPSE	26.98	**0.5660**	**0.5234**	**0.5941**	**0.1625**
89M Priors					
Projection	28.81	0.4579	0.4625	0.4761	0.1343
Diffusion-Baseline	29.78	0.4988	0.4790	0.4604	0.1247
ECLIPSE	**25.77**	**0.5712**	**0.5358**	**0.6194**	**0.16665**

this behavior further. In total, we randomly selected 200 prompts from the MSCOCO validation set (instead of T2I-CompBench as reported in Figure 4) and asked the human evaluators to perform two studies:

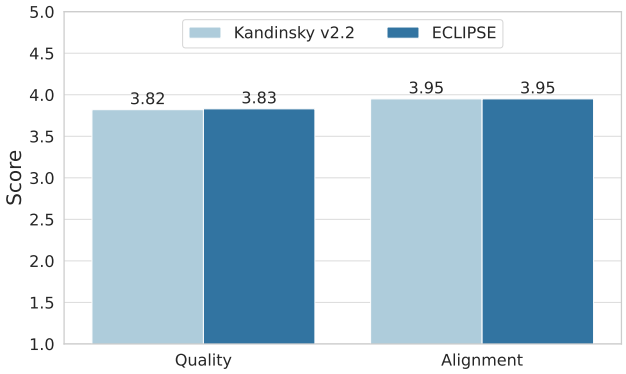

Figure 9. Human evaluations of the *ECLIPSE vs.*Kandinsky v2.2 generated images. It can be observed that both models are rated equally in terms of image quality and caption alignment.

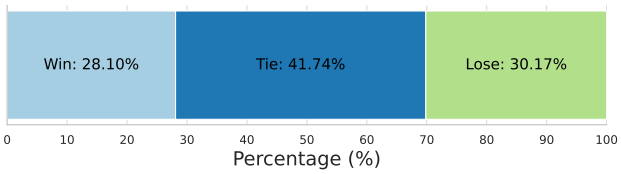

Figure 10. This figure illustrates the human preferences between *ECLIPSE* prior for Kandinsky model (trained on LAION-HighRes subset) *vs.* Original Kandinsky v2.2 model.

- Rate each image in terms of quality and caption alignment between 1-5. Where 1 is the artificial-looking image and caption alignment is poor. While 5 represents a very high-quality image and is perfectly aligned with the captions.
- Image preferences in terms of aesthetics. We show images from both models and ask the evaluators to choose one which looks more aesthetically pleasing.

Figure 11. Qualitative examples comparing (in terms of aesthetics) *ECLIPSE* with Kandinsky v2.2.

Interestingly, as shown in Figure 9, both models are rated equally when evaluated independently. Additionally, according to Figure 10, Kandinsky v2.2 is preferred slightly more than the *ECLIPSE* in terms of aesthetic quality. This finding suggests that smaller prior trained with *ECLIPSE* can perform equally (if not better) to those big prior models. Figure 11 shares three examples from the MSCOCO. Both models perform equally well but Kandinsky is more aesthetically pleasing. Figure 20 and 21 show the MTurk human evaluation instructions.

E. Diversity With Non-Diffusion Priors

One important aspect of the diffusion models is the diversity of the generated images. Therefore, diversity and caption alignment go hand-in-hand. We further analyze whether having the non-diffusion prior hurts diversity or not. We perform additional qualitative evaluations and given a prompt – we ask the human evaluators to select which of the two grids of six images are more diverse. This experiment is performed between *ECLIPSE* and Kandinsky v2.2. As shown in Figure 12, even if we use the non-diffusion prior model it does not hurt the diversity. Diffusion image decoder is the main reason that contributes to the diversity and having diffusion or non-diffusion prior does not contribute that significantly.

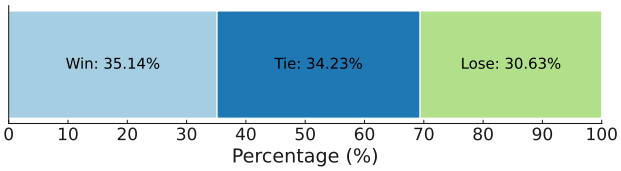

Figure 12. This figure illustrates the human preferences on the diversity of generated images between *ECLIPSE* prior with Kandinsky v2.2 diffusion image decoder *vs.* Kandinsky v2.2.

F. More Qualitative Evaluations

In this section, we provide more qualitative examples and discuss them. We also provide comparisons based on the diffusion image decoder used (i.e., Karlo and Kandinsky v2.2). Finally, we discuss several failure cases.

F.1. *ECLIPSE* with Karlo Decoder

Figure 13 illustrates the comparison between Projection, Diffusion-Baseline, and *ECLIPSE* priors trained on CC12M. It can be seen that *ECLIPSE* performs very well on complex composition prompts. While Projection and Diffusion baselines struggle to generate images aligned with the target prompt. Figure 14 compares the *ECLIPSE* priors trained on different datasets. Here, *ECLIPSE* prior trained on MSCOCO does not always follow the target prompt accurately and generates the lower quality images. That said, the overall performance between all priors is very similar; suggesting that even a small amount of dataset is sufficient to distill the knowledge from the pre-trained Vision-Language models. Figure 15 compares the *ECLIPSE* models with various SOTA methods. Noticeably, *ECLIPSE* performs better than the other baselines in terms of the ability to follow the target prompts. For instance, many SOTA models cannot generate "empty blue vase", "cat in space suit", and "blue bowl on white placemat". Although we observe that *ECLIPSE* prior trained with MSCOCO does follow the target text prompts but cannot generate high-quality images, which aligns with our previous findings.

F.2. *ECLIPSE* with Kandinsky Decoder

Similarly, we analyze the qualitative results on Kandinsky diffusion image decoders. Figure 16 compares the various baselines priors with the *ECLIPSE* prior. We observe that baselines perform very closely to the *ECLIPSE* prior, which is the opposite of what we found in Figure 13. We attribute this behavior to the change in the pre-trained CLIP encoder. Additionally, as shown in Table 2 of the main paper, both baseline priors perform very highly compared to the same priors trained on the CC12M dataset for the Karlo decoder. The only difference is the pre-trained vision-language model. **Therefore, the selection of the Vision-Language model also plays a crucial role.**

Figure 17 illustrates the comparison with *ECLIPSE* priors trained with different datasets. It can be observed that with the use of the LAION-HighRes dataset not only did image quality improve but small intrinsic details (such as "backpack", "belt", etc.) also improved. Even in some cases, prior training on the LAION subset performs better as the increase in the amount of data improves the performance. Figure 18 provides more qualitative examples to compare the *ECLIPSE* priors with other respective SOTA methods. As also previously observed, *ECLIPSE* prior trained on LAION subset performs very close to the Kandinsky v2.2 in terms of following the text prompts. While big SOTA models such as Stable Diffusion v1.4/2.1, and Wurstchen fall short despite being trained on millions of data.

F.3. Failure Cases

Figure 19 shows some examples where *ECLIPSE* model fails to follow the prompt precisely. It is still difficult for the prior to learn something very unconventional. The model fails at generating some composition prompts (first four images). It has been shown that vision-language models also suffer from such composition understanding (e.g., "grass in the mug" *vs.* "mug in the grass"). Therefore, improving the Vision-Language model can further improve the capabilities of unCLIP priors. Notably, *ECLIPSE* finds it difficult to generate artistic imaginary images (such as "nebula explosion that looks like corgi"). However, such corner cases can be only solved with more diverse high-quality datasets.

G. Future Work

In this work, we focus on improving text-to-image priors. We assume that there exists a pre-trained diffusion image decoder that can be used as it is. To further improve the parameter efficiency for training, several relevant works on knowledge distillation and model compression can help. Moreover, to improve the compositional abilities for unCLIP models, a better vision-language model (such as SigLIP) as the base model can be utilized to train the prior model using *ECLIPSE*. However, this will require the diffusion image decoder to be adjusted according to the new vision latent space. We leave this direction as the future work as our paper primarily focuses on enhancing T2I priors.

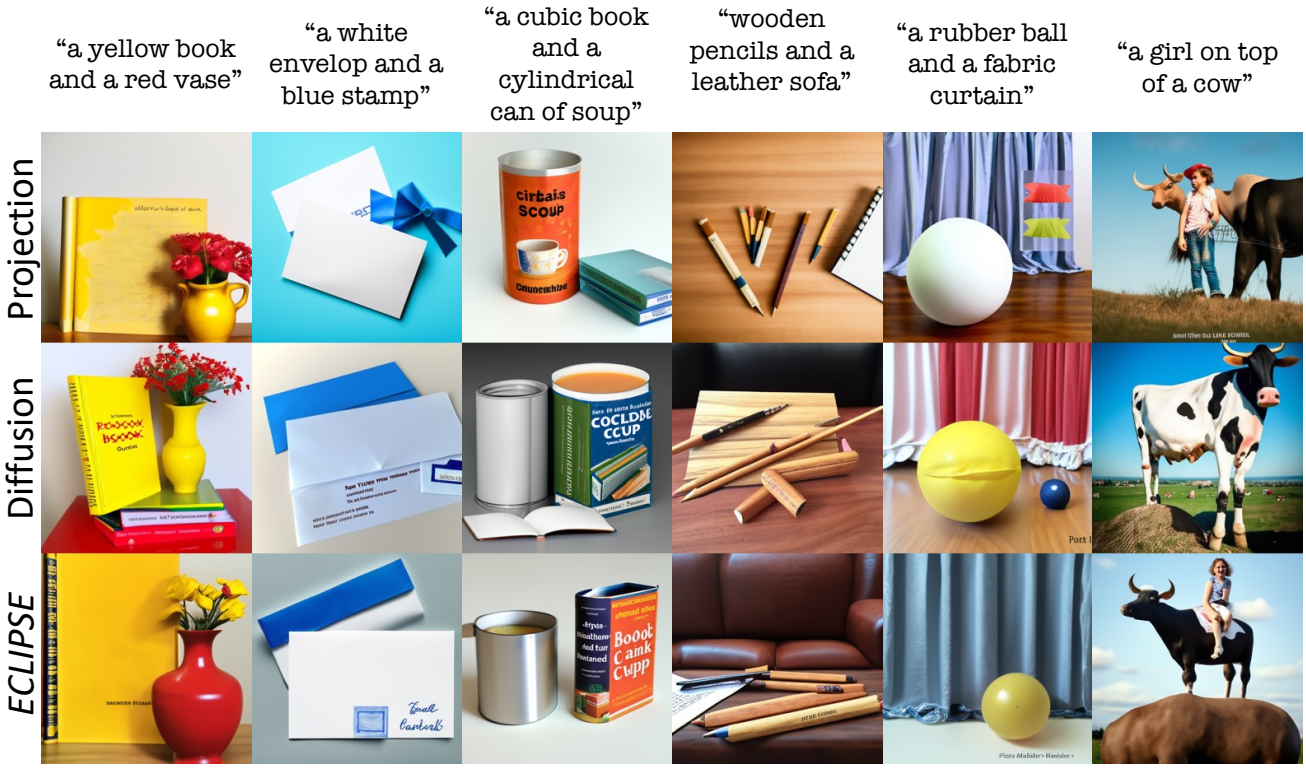

Figure 13. Qualitative comparisons between *ECLIPSE* and baseline priors (having 33 million parameters) trained on CC12M dataset with Karlo decoder. * prompt is: "The bold, striking contrast of the black and white photograph captured the sense of the moment, a timeless treasure memory."

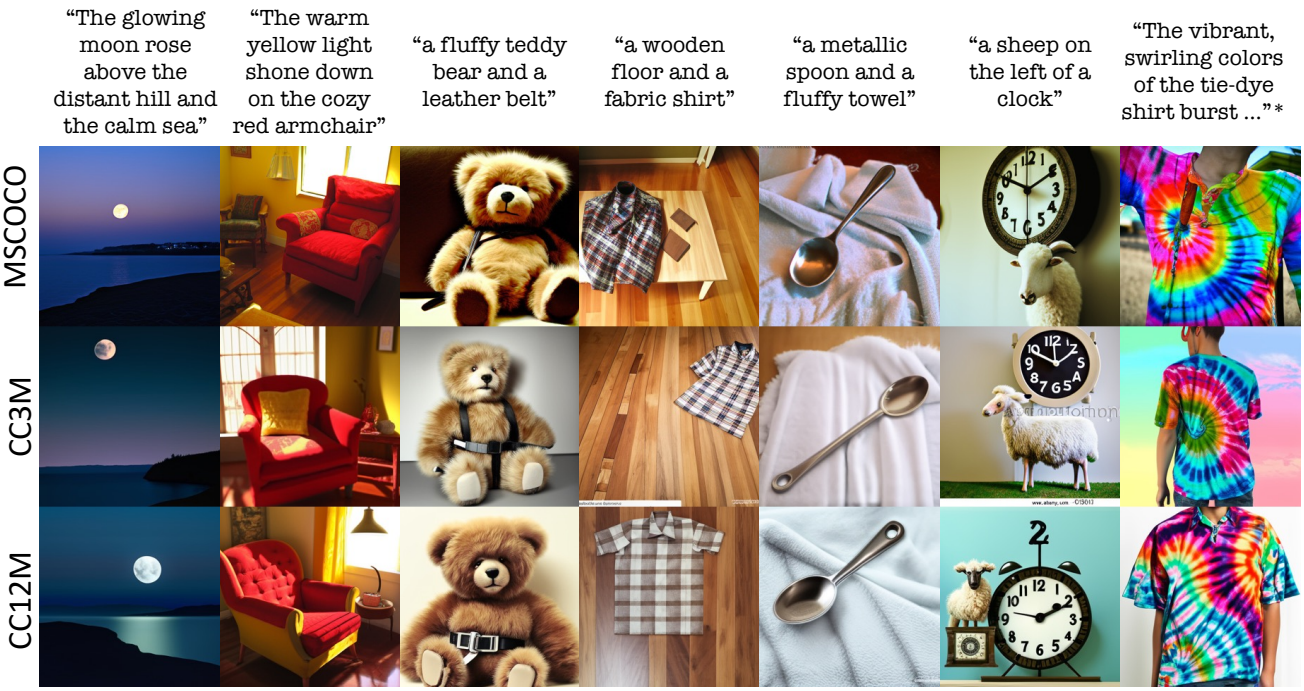

Figure 14. Qualitative comparisons of *ECLIPSE* priors with Karlo decoder trained on different datasets. * prompt is: "The vibrant, swirling colors of the tie-dye shirt burst with energy and personality, a unique expression of individuality and creativity."

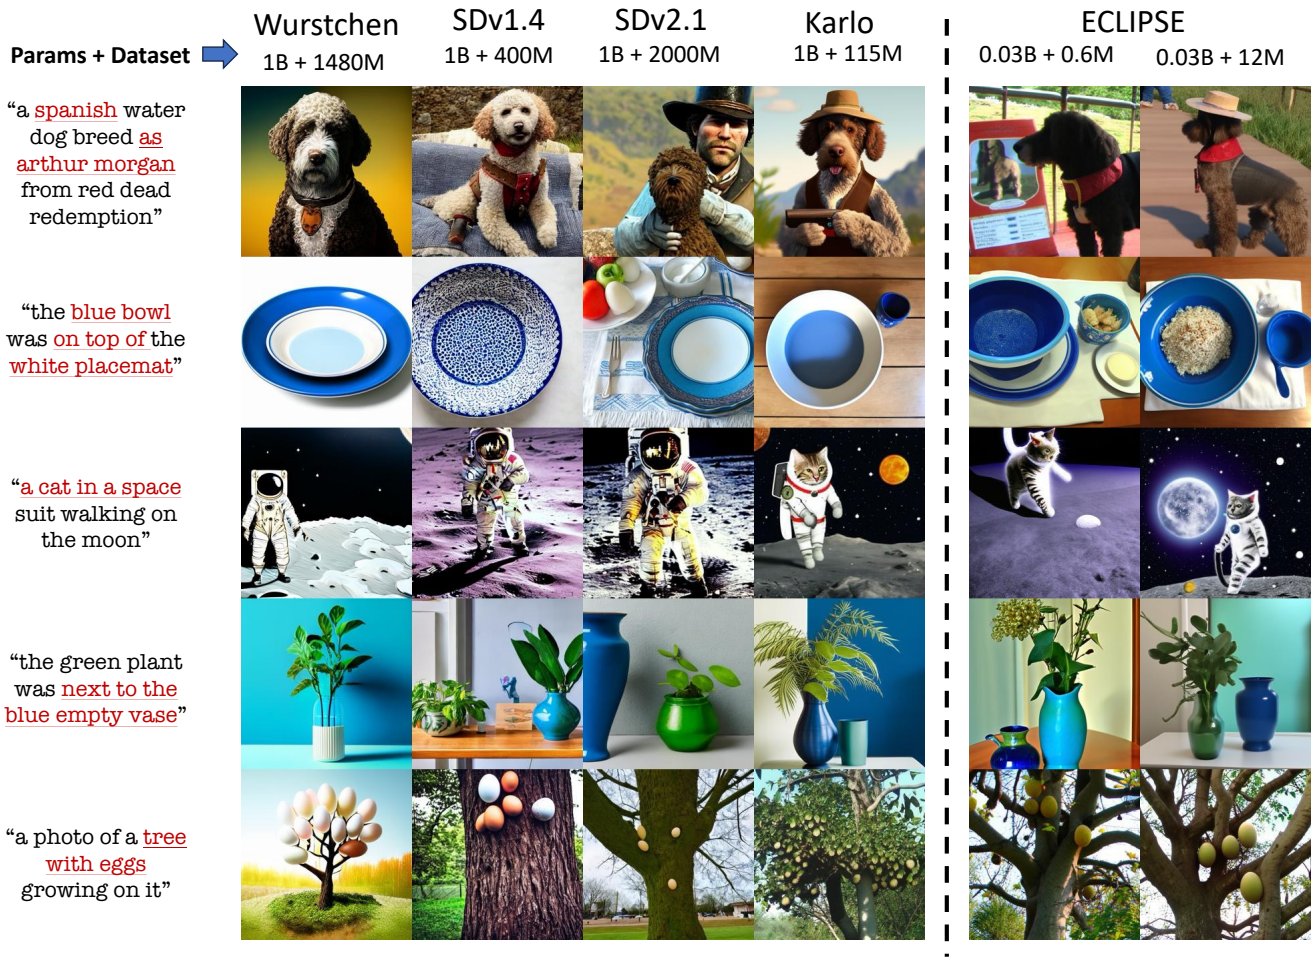

Figure 15. Qualitative result of our text-to-image prior, *ECLIPSE* (with Karlo decoder), along with a comparison with SOTA T2I models. Our prior model reduces the prior parameter requirements (from 1 Billion → 33 Million) and data requirements (from 115 Million → 12 Million → 0.6 Million).

Figure 16. Qualitative comparisons between *ECLIPSE* and baseline priors (having 34 million parameters) trained on LAION-HighRes subset dataset with Kandinsky v2.2 diffusion image decoder.

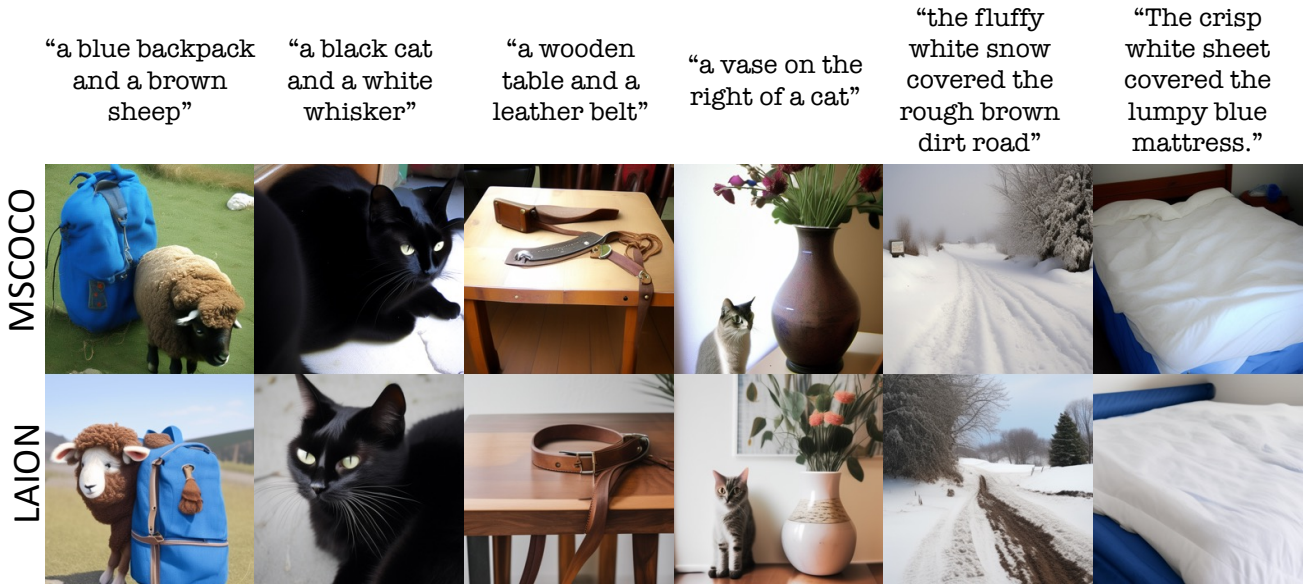

Figure 17. Qualitative comparisons between *ECLIPSE* prior trained on MSCOCO and LAION datasets with Kandinsky v2.2 decoder.

Figure 18. More qualitative result of our text-to-image prior, *ECLIPSE* (with Kandinsky v2.2 decoder), along with a comparison with SOTA T2I models. Our prior model reduces the prior parameter requirements (from 1 Billion → 33 Million) and data requirements (from 177 Million → 5 Million → 0.6 Million).

Figure 19. Instances where *ECLIPSE* encounters the challenges in following the target text prompts.

Visual Question Answering

Consent Checked: yes

Context: The purpose of this research project is to study how accurately Artificial Intelligence / Machine Learning models generate images from textual descriptions. We study a special type of models called text-to-image generators -- users can enter a sentence to the model, and the model generates an image for this this sentence. The goal is to understand if the generated image is aligned with the input text.

Your task: In this HIT, we will show you an image. Your task is to answer several questions about this image about the objects present in the image and the quality of the image to measure the image-text alignment.

Solved Examples: In order to help you build an understanding of the task, here are a few examples of solved HITs: [Example 1] [Example 2] [Example 3]

Caption: a man in a chefs hat chopping food

(1) Rate the **quality** of the image. ○ 1 ○ 2 ○ 3 ○ 4 ○ 5
("1" being artificial (noisy, blury) and "5" being natural (a real photograph))

(2) What is the **similarity** between the caption and image? ○ 1 ○ 2 ○ 3 ○ 4 ○ 5
Rate from "1" (least similar) to "5" (most similar).

Figure 20. An example of human annotation for determining the image quality and caption alignment.

AI generated image preferences

Consent Checked: yes

Context: The purpose of this research project is to study the quality of the AI generated images with respect to each other, given the input caption. In this study, we will provide a caption and two different images. The goal is to select one of these two as the preferred choice.

Your task: In this HIT, we will give you four tasks. For each task, you will get two images and one target prompt/caption. Your goal is to select the best aesthetically pleasing image.

Quick Guide:
1) If both images are equally good then you can select the **EQUAL** option.

Solved Examples: In order to help you build an understanding of the task, here are a few examples of solved HITs: [Example 1] [Example 2] [Example 3]

(1) Select the image with **the best aesthetics** that follow the caption: A bike parked on the side walk and a car on the street

○ ○ ○ **EQUAL**

Figure 21. An example of human annotation for determining the most aesthetic image.