

## Diffusion-based Visual Anagram as Multi-task Learning

Zhiyuan Xu<sup>\*1,2</sup>, Yinhe Chen<sup>\*1,3</sup>, Huan-ang Gao<sup>1</sup>, Weiyan Zhao<sup>1,4</sup>, Guiyu Zhang<sup>1</sup>, Hao Zhao<sup>†1</sup>

<sup>1</sup>Institute for AI Industry Research (AIR), Tsinghua University

<sup>2</sup>University of Chinese Academy of Sciences

<sup>3</sup>Central South University, <sup>4</sup>Beijing Institute of Technology

### Abstract

*Visual anagrams are images that change appearance upon transformation, like flipping or rotation. With the advent of diffusion models, generating such optical illusions can be achieved by averaging noise across multiple views during the reverse denoising process. However, we observe two critical failure modes in this approach: (i) concept segregation, where concepts in different views are independently generated, which can not be considered a true anagram, and (ii) concept domination, where certain concepts overpower others. In this work, we cast the visual anagram generation problem in a multi-task learning setting, where different viewpoint prompts are analogous to different tasks, and derive denoising trajectories that align well across tasks simultaneously. At the core of our designed framework are two newly introduced techniques, where (i) an anti-segregation optimization strategy that promotes overlap in cross-attention maps between different concepts, and (ii) a noise vector balancing method that adaptively adjusts the influence of different tasks. Additionally, we observe that directly averaging noise predictions yields suboptimal performance because statistical properties may not be preserved, prompting us to derive a noise variance rectification method. Extensive qualitative and quantitative experiments demonstrate our method’s superior ability to generate visual anagrams spanning diverse concepts.*

### 1. Introduction

**Visual anagrams** are images whose appearances change when subjected to transformations such as rotation or flipping [5, 19]. For instance, as illustrated in Fig. 1, an image may resemble a garden gnome when viewed head-on but transform into a hot air balloon when observed upside-down. The study of visual anagrams involves understanding how geometric transformations manipulate and redefine visual perception, making this field promising for applications in art, design, and cognitive science [4, 15, 40].




Figure 1. **Visual Anagrams.** We show an example of visual anagrams, which can be perceived as a garden gnome or a hot air balloon depending on the orientation of the image.



Figure 2. **Failure Cases.** We show two common failure cases of [19]: concept segregation (left) and domination (right).

**Diffusion.** With the advent of diffusion models [14, 26, 50, 52, 60], generating such visual anagrams has become significantly easier by simply averaging noise predictions across different views of various prompts. In the reverse denoising process of diffusion models, the image layout is formed during the early denoising timestamps while the colors and details are generated in later timestamps. This allows us to establish a layout of noised images that matches the different prompts before adding details. This aligns with the ideas that generative models and human perception may process optical illusions similarly [20, 28, 38].

**Problems.** Despite the simplicity of this method, our careful investigation has identified two critical issues. First, there is the problem of concept segregation, where different concepts are synthesized independently (as illustrated in Fig. 2a, where the rabbit and the violin appear as separate objects within a single image), which does not constitute a true visual anagram. Secondly, there is the issue of concept domination, where some concepts overpower others in the generated images (as shown in Fig. 2b, where the cat is

<sup>1\*</sup>Equal Contribution. <sup>†</sup>Corresponding Author.

<sup>2</sup>Code: <https://github.com/Pixtella/Anagram-MTL>

clearly visible but the ship is not as distinct).

**Solutions.** To address these two issues, we draw inspiration from multi-task learning, aiming to derive denoising trajectories that enhance commonalities across multiple generation tasks while adaptively alleviating their conflicts. At the core of our solution lie two new techniques: **(I) Anti-Segregation Optimization:** By treating the attention map of spatial queries against the subjects of different prompts as encoding weights for different tasks in the intermediate images being denoised, we effectively prevent the segregation of concepts by promoting the intersection of these attention score maps across different views. **(II) Noise Vector Balancing:** Noise predicted from different views can be seen as optimizing directions (or gradients) for different tasks in the denoising trajectory. We adaptively balance the noise vectors across tasks by reweighting them using a task completion measure. **Additionally,** we recognize the importance of preserving variance in the diffusion denoising process. To this end, we modulate the combined noise predictions with estimated noise correlation information to maintain key statistical properties.

In summary, our contributions are as follows:

- **Integration with Multi-Task Learning:** We link the visual anagram generation problem to multi-concept generation and propose the use of multi-task learning techniques to address it effectively.
- **Solution to Concept Segregation and Domination:** We identify critical issues, such as concept segregation and domination in generated images, and propose two effective methods to resolve them.
- **Comprehensive Evaluation:** We present extensive quantitative and qualitative results to demonstrate the efficacy and flexibility of our proposed method compared to previous state-of-the-art [19].

## 2. Related Work

**Computational Optical Illusions.** Optical illusions and computational arts hold great potential for advancing our understanding of human and machine perception [20,25,28,34,35,38,57]. A long line of research has focused on generating such illusions by leveraging specific characteristics of human vision. For instance, creating the illusion of continuous uni-direction motion can be achieved by locally applying a filter with a continuously shifting phase, based on the observation that local phase shifts are interpreted as global movement [16]. Hybrid images, which change appearance depending on viewing distance, can be generated by exploiting the multiscale nature of human perception, blending high frequencies from one image with low frequencies from another [41]. Camouflage images, which contain hidden figures that remain imperceptible or unnoticed at first glance, can be created by removing the original subtle texture details of the hidden figures and replac-

ing them with those of the surrounding prominent image, as human perception operates on a two-phase basis: feature search and conjunction search [13, 24, 42]. Perceptual puzzles involving color constancy, size constancy, and face perception can be generated by finding adversarial examples for principled models of human perception, where vision is explicitly modeled as Bayesian inference [6]. While this line of research shows great promise, the dependence on explicitly identifying characteristics within human vision models can be cumbersome. And this challenge can be alleviated by leveraging the data priors learned by diffusion models [18, 26].

**Diffusion-based Visual Anagram generation.** Diffusion models [8, 14, 17, 21, 26, 31, 44, 45, 47, 50–52] are a powerful class of generative models that gradually transform a noise-distributed sample into one from a target data distribution by following a reverse denoising trajectory. In this process, the models estimate and remove noise over time, conditioned on both a noisy image and a specific timestep. This is further enhanced by incorporating an embedding of a text prompt, facilitating text-conditioned image synthesis [1, 45]. To generate visual anagrams using diffusion models, Burgert *et al.* [5] proposed to use Score Distillation Loss [43], but this method is computationally expensive and degrades image quality. Tancik [54] demonstrated that visual anagrams can be generated using Stable Diffusion [45] by alternately denoising latents with different prompts under different views. Geng *et al.* [19] propose to average the noise predictions from different prompts, using a pixel-based diffusion model, and achieve better results in terms of image quality and computational efficiency. Our work builds on the latter approach, which we will refer to as the baseline.

In this paper, we identify two critical issues with this approach: concept segregation and concept domination. We attribute these problems to the failure to explicitly account for the multiple concepts being generated. While compositional multi-concept generation has long been studied [2, 3, 7, 33], it typically involves creating separate subjects that dominate different regions of an image. In contrast, visual anagram generation aims to generate a single object that can be perceived as multiple concepts from different viewpoints. This fundamental difference underscores the unique challenges in generating effective visual anagrams. To address these issues, we draw an analogy between visual anagram generation and multi-task learning. This analogy enables us to derive denoising trajectories that leverage commonalities between concepts while balancing their differences.

**Multi-task learning.** Multi-task learning seeks to enhance the performance and generalization capabilities of models on individual tasks by leveraging shared representations to exploit commonalities and balance differences

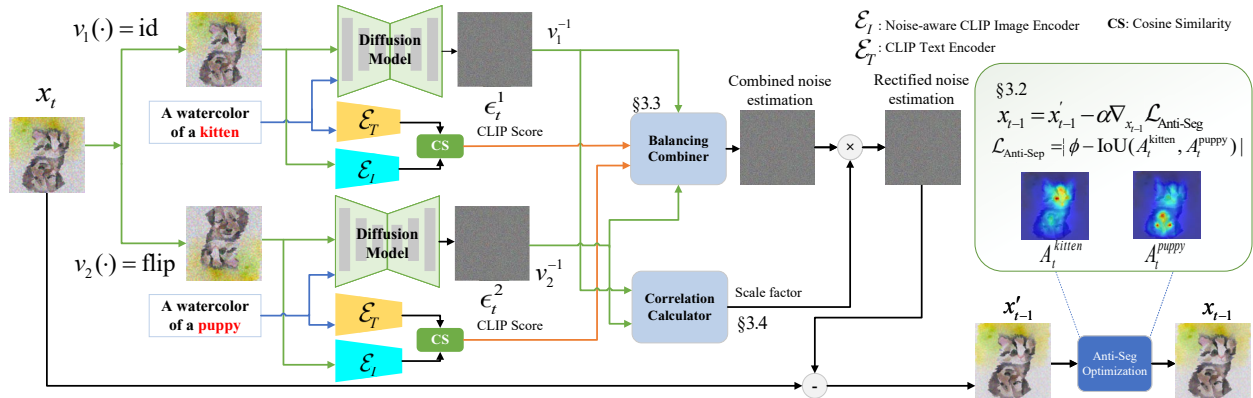


Figure 3. **Method overview.** During each denoising step, the intermediate image  $x_t$  first passes through the diffusion model together with the corresponding text prompt under each view, and also through a noise-aware CLIP model which measures the degree of task completion for each view. (1) **Noise Vector Balancing:** Predicted noise vectors are reweighted based on the degree of task completion before being combined, see Sec. 3.3. (2) **Noise Variance Rectification:** Combined noise vectors are rectified by applying a scale factor calculated based on estimated correlation coefficients, which is detailed in Sec. 3.4. (3) **Anti-Segregation Optimization:** The denoised image  $x'_{t-1}$  is modulated to encourage intersection of attention maps of different views with an inference-time loss term before being passed to the next denoising step, refer to Sec. 3.2.

across tasks [9, 10, 53, 62, 63]. The primary challenge in this domain is to strike the balance between different tasks, and various techniques have been developed, including loss weighting [22, 29, 58], gradient normalization [11], gradient dropout [12], gradient surgery [59], Nash Bargaining solutions [37], Pareto-optimal solutions [32, 36, 48], gradient alignment [49], and curriculum learning [23, 27]. In our work, we apply multi-task learning principles to the problem of visual anagram generation. We treat image generation under different views as separate tasks and develop novel techniques to address task segregation and overpowering issues to yield a denoising trajectory that adaptively balances the differences.

### 3. Method

Our goal is to generate visual anagrams using off-the-shelf text-to-image diffusion models. Formally, given a set of text prompts  $\{y_1, y_2, \dots, y_N\}$  and corresponding views  $\{v_1, v_2, \dots, v_N\}$ , we aim to generate an image  $x_0$  that matches each prompt  $y_i$  under view  $v_i$ . This has been shown to be possible by averaging the noise predictions from different prompts [19], which represents previous state-of-the-art, but the resulting images often suffer from concept segregation and concept domination issues as discussed in Fig. 2. Fig. 3 illustrates the overall pipeline of our proposed method. In this section, we first introduce the preliminaries in Sec. 3.1, followed by a detailed explanation of our techniques in the following sections.

#### 3.1. Preliminaries

**Text-to-Image Diffusion Models.** Diffusion models iteratively denoise a sample  $x_T$  initialized from Gaussian noise  $\mathcal{N}(0, I)$ , to generate a noise-free sample  $x_0$  from a learned target distribution. Here,  $T$  represents the total

number of denoising steps. For text-to-image generation, the model is conditioned on text prompts  $y$ , and can be formulated as  $\epsilon_\theta(x_t, t, y)$ , where  $x_t$  denotes the intermediate image at time step  $t$ . The model predicts the noise  $\epsilon_t$  at each step, which is then used to compute the next image  $x_{t-1}$ . In practice, the model could be implemented using the UNet architecture [46], with text prompts effectively integrated via cross-attention mechanisms [45].

**Text-Image Cross-Attention.** A common method for integrating text prompts into diffusion models is through cross-attention layers within the UNet architecture [45]. In this setup, the projected downsampled image input  $x_t^d$  serves as the query, while the projected text embedding  $y_e \in \mathbb{R}^{N \times D}$  (where  $N$  is the token number and  $D$  is the embedding dimension) acts as the key and value. The output of the cross-attention layer is computed using multi-head attention [55], where attention maps  $A_t \in \mathbb{R}^{r \times r \times N}$  are generated to model the interaction between text prompts and corresponding image regions. Our method in Sec. 3.2 operates on the averaged attention maps from all cross-attention layers at a resolution of  $32 \times 32$ .

#### 3.2. Anti-Segregation Optimization

In multi-task learning (MTL), a single model is trained to solve multiple tasks simultaneously, where shared representations are learned to exploit commonalities between tasks [61]. Similarly, in visual anagram generation, concepts from different prompts are expected to share a single image.

However, the lack of explicit control over region sharing in the image can lead to concept segregation, as discussed in Sec. 1. In such cases, concepts from different prompts may appear as distinct, separate objects (as shown

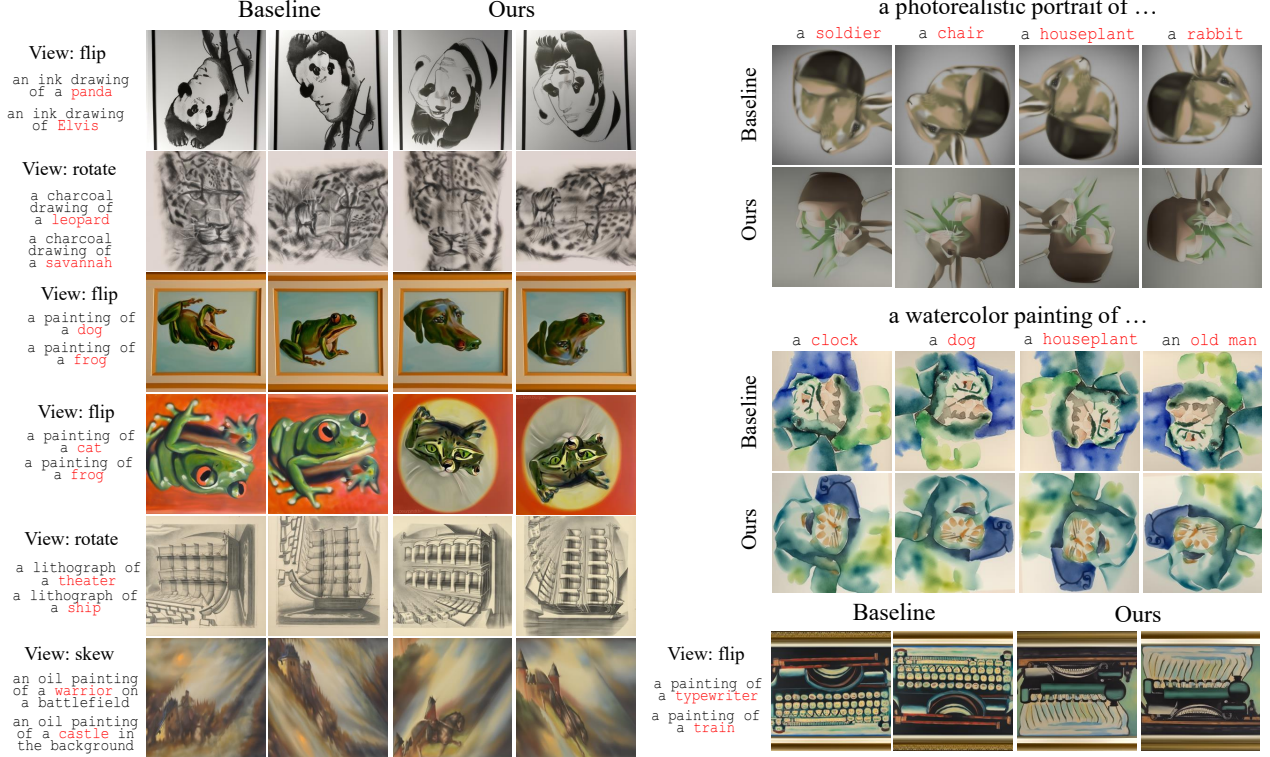


Figure 4. More qualitative results of our proposed method compared to the baseline method [19]

in Fig. 2a), which is considered a failure in anagram generation. We conjecture that this is caused by the attention maps of different views being too disjointed. Therefore, we design an inference-time loss term to encourage the intersection of the attention maps of concepts from different views. Specifically, we assume that there is only one concept  $c_i$  in each view  $i$ , which is the main object in the corresponding prompt (e.g., the words in red in Figs. 3 and 4), and the loss term is defined as:

$$\mathcal{L}_{\text{Anti-Seg}} = \frac{1}{N(N-1)} \sum_{1 \leq i < j \leq N} \left| \phi - \frac{\sum_p \min(A_{t,p}^{c_i}, A_{t,p}^{c_j})}{\sum_p (A_{t,p}^{c_i} + A_{t,p}^{c_j})} \right|, \quad (1)$$

where  $p$  iterates over all pixels in the attention maps,  $A_{t,p}^{c_i}$  denotes the summed attention score of all tokens derived from concept  $c_i$  at pixel  $p$  at time step  $t$ , and  $N$  is the number of views.  $\phi$  is a hyperparameter that controls the target overlap ratio between attention maps, and will be discussed in Sec. 5.2. At each denoising step, after the denoised image is computed, this loss term is used to modulate it through single-step gradient descent with one step of look-ahead:

$$x_{t-1} = x'_{t-1} - \alpha \nabla_{x'_{t-1}} \mathcal{L}_{\text{Anti-Seg}} \quad (2)$$

where  $\alpha$  controls the step size. The updated image will be passed to the next denoising step.

### 3.3. Noise Vector Balancing

A common challenge in multi-task learning is the imbalance of gradients across tasks, where certain tasks may dominate the learning process [59]. To address this, Grad-Norm [11] propose to measure each task’s training progress using the ratio of loss decrease, and assigns higher weights to gradients from tasks that have shown less progress.

Building on this idea, we propose reweighting the estimated noise vectors at each time step based on task completion scores. Specifically, text prompts and their viewed corresponding noisy images are respectively fed into a CLIP text encoder and a noise-aware CLIP image encoder, which has been pre-trained on both clean and noisy images. The cosine similarity between the text and image embeddings serves as the task completion score, which is then used to reweight the noise vectors before combining them. The reweighting and combining process can be formally expressed as:

$$\text{CS}_t^i = \cos \langle \mathcal{E}_T(y_i), \mathcal{E}_I(v_i(x_t)) \rangle \quad (3)$$

$$\hat{\alpha}_t^i = (\text{CS}_t^i)^{-2 + \frac{1}{T}} \quad (4)$$

$$\alpha_t^i = \frac{\hat{\alpha}_t^i}{\sum_{j=1}^N \hat{\alpha}_t^j} \quad (5)$$

$$\tilde{\epsilon}_t = \sum_{i=1}^N \alpha_t^i v_i^{-1}(\epsilon_t^i) \quad (6)$$




Figure 5. Qualitative results of Anti-Segregation Optimization and visualization of attention maps.




Figure 6. Qualitative results of Noise Variance Rectification.

where  $\mathcal{E}_T$  and  $\mathcal{E}_I$  represent the CLIP text encoder and the noise-aware CLIP image encoder respectively,  $i$  indexes the views,  $v_i$  denotes the view transformation function,  $y_i$  is the text prompt for view  $i$ ,  $\epsilon_t^i$  denotes the noise prediction for view  $i$  at time step  $t$ , and  $T$  is the total number of denoising steps. The reweighted noise vector  $\tilde{\epsilon}_t$  is then used to compute the next image.

Intuitively, in Eq. (3), a higher cosine similarity indicates better alignment between the text prompt and the image under a given view, suggesting that the generation task is closer to completion. Therefore, the negative exponent in Eq. (4) assigns lower weights to the noise vectors for that view, preventing it from dominating the denoising process. Additionally, the linear term in the exponent is designed to gradually increase the importance of the task completion score as the denoising process progresses, which is based on the observation that CLIP’s task completion score becomes more reliable as the image becomes cleaner. Eq. (5) normalizes the reweighted coefficients and Eq. (6) combines the reweighted noise vectors together.

### 3.4. Noise Variance Rectification

As stated in [26], one of the most important assumptions of denoising diffusion models is that the estimated noise vectors follow a Gaussian distribution with zero mean and unit variance, which is crucial for the model to reverse the diffusion process. Existing work [56] has identified that the introduction of classifier-free guidance incurs a statistics shift, which might hinder the denoising process, and proposed normalized classifier-free guidance to rectify this issue.

In the context of diffusion-based visual anagram generation, although the estimated noise vectors for each view are expected to possess standard Gaussian properties because

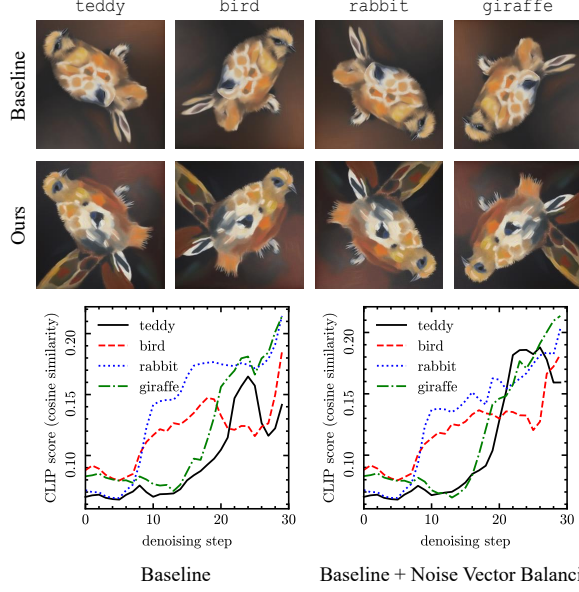


Figure 7. Qualitative results of Noise Vector Balancing and visualization of the task completion curve.

the model that generates them is trained under this assumption, the combined noise vector may not, which could be detrimental to the denoising process. Therefore, we propose to rectify the combined noise vector to preserve the desired statistical properties. Below we provide a detailed derivation of the rectification process. For convenience, we abuse the notation and denote  $v_i^{-1}(\epsilon_t^i)$  as  $\epsilon_t^i$  in this subsection. Basic assumptions include:

$$\forall i, \epsilon_t^i \sim \mathcal{N}(0, I) \quad (7)$$

Therefore, for every element  $p$  in the combined noise vector given by Eq. (6), we have:

$$\mathbb{E}[\epsilon_{t,p}^i] = \mathbb{E}\left[\sum_{i=1}^N \alpha_t^i \epsilon_{t,p}^i\right] = \sum_{i=1}^N \alpha_t^i \mathbb{E}[\epsilon_{t,p}^i] = 0 \quad (8)$$

This suggests that every element in the combined noise vector still has an expectation of zero. However, the variance is not guaranteed to be one. To address this issue, we propose to rectify the combined noise vector by applying a scale factor  $c$  to make the variance unit. Essentially, we hope to find a scale factor  $c$  that satisfies:

$$\begin{aligned} \mathbb{D}[c \cdot \tilde{\epsilon}_{t,p}] &= c^2 \mathbb{D}[\tilde{\epsilon}_{t,p}] = c^2 \mathbb{D}\left[\sum_{i=1}^N \alpha_t^i \epsilon_{t,p}^i\right] \\ &= c^2 \left[ \sum_{i=1}^N (\alpha_t^i)^2 \mathbb{D}[\epsilon_{t,p}^i] + \sum_{1 \leq i < j \leq N} 2\alpha_t^i \alpha_t^j \text{Cov}(\epsilon_{t,p}^i, \epsilon_{t,p}^j) \right] = 1 \end{aligned} \quad (9)$$

since we have  $\mathbb{D}[\epsilon_{t,p}^i] = 1$ , the scale factor  $c$  can be computed as:

$$c = \frac{1}{\sqrt{\sum_{i=1}^N (\alpha_t^i)^2 + \sum_{1 \leq i < j \leq N} 2\alpha_t^i \alpha_t^j \text{Cov}(\epsilon_{t,p}^i, \epsilon_{t,p}^j)}} \quad (10)$$

where the covariance term satisfies:

$$\text{Cov}(\epsilon_{t,p}^i, \epsilon_{t,p}^j) = \mathbb{E}[\epsilon_{t,p}^i \epsilon_{t,p}^j] - \mathbb{E}[\epsilon_{t,p}^i] \mathbb{E}[\epsilon_{t,p}^j] = \mathbb{E}[\epsilon_{t,p}^i \epsilon_{t,p}^j] \quad (11)$$

since both  $\epsilon_{t,p}^i$  and  $\epsilon_{t,p}^j$  follow standard Gaussian, the covariance term equals their correlation coefficient. We assume that within the same denoising step, the correlation coefficients between different views’ noise vectors are the same across all elements. Let  $C \times H \times W$  be the total number of elements in each noise vector, which is usually large, so according to the law of large numbers, we can approximate the covariance term as:

$$\rho_{i,j} = \text{Cov}(\epsilon_{t,p}^i, \epsilon_{t,p}^j) = \mathbb{E}[\epsilon_{t,p}^i \epsilon_{t,p}^j] \approx \frac{\sum_p \epsilon_{t,p}^i \epsilon_{t,p}^j}{C \times H \times W} \quad (12)$$

Now we can compute the desired scale factor  $c$  by plugging in the estimated covariance term. The rectified noise vector is just the combined noise vector multiplied by this scale factor.

## 4. Experiments

### 4.1. Implementation Details

We utilize DeepFloyd [1] as our backbone diffusion model, a cutting-edge open-source text-to-image diffusion model. It operates in the pixel space instead of latent space, which is crucial for producing high-quality visual anagrams as [18] mentioned. Our evaluations focus on images generated at the second stage of DeepFloyd, with a resolution of  $256 \times 256$ . For task completion measurement, we employ the ViT-B CLIP mode released by [39], which is pre-trained on both noisy and clean images. Our proposed Anti-Segregation Optimization is applied specifically to the initial stage of DeepFloyd, which is responsible for generating  $64 \times 64$  images from pure noise and determining the general layout of the image, where concept segregation issues often arise. Noise vector balancing and variance rectification are implemented in both image generation and super-resolution stages. It is important to note that training or fine-tuning of the models is not required, and indeed, a single RTX 3090 suffices to run our proposed pipeline.

### 4.2. Dataset and Metrics

For quantitative evaluation and ablation study, we compile a 2-view dataset with text prompt pairs using 10 classes of objects from the CIFAR-10 dataset [30], that generally follows the setup in [18]. We further extend the dataset to 3-view to evaluate the scalability of our proposed method. For each pair (or triplet) of text prompts, we generate 10 images with different seeds and report the average results.

We follow [18] and compute a score matrix  $S \in \mathcal{R}^{N \times N}$  for each generated image, where  $N$  is the number of views,

Dataset	Method	$\mathcal{A}_{\min} \uparrow$	$\mathcal{C} \uparrow$	$\mathcal{A}_{\text{avg}} \uparrow$
CIFAR10 2-views	Tancik [54]	0.1997	0.5568	0.2295
	Burgert <i>et al.</i> [5]	0.2547	0.6867	0.2676
	Visual Anagrams [19]	0.2583	0.6744	0.2717
	Ours	<b>0.2711</b>	<b>0.6913</b>	<b>0.2816</b>
CIFAR10 3-view	Tancik [54]	0.1914	0.4043	0.2207
	Burgert <i>et al.</i> [5]	N/A	N/A	N/A
	Visual Anagrams [19]	0.2363	0.4892	0.2642
	Ours	<b>0.2507</b>	<b>0.4955</b>	<b>0.2759</b>

Table 1. **Quantitative Results.** We report the results of our proposed method against baseline [19], together with other existing methods on the 2-view dataset and the 3-view dataset. Burgert *et al.* [5] only supports 2 views.

Method	ASO	NVB	NVR	$\mathcal{A}_{\min} \uparrow$	$\mathcal{C} \uparrow$	$\mathcal{A}_{\text{avg}} \uparrow$
Baseline	×	×	×	0.2583	0.6744	0.2717
	✓	×	×	0.2594	0.6821	0.2716
	×	✓	×	0.2632	0.6864	0.2734
	×	×	✓	0.2642	0.6739	0.2784
	×	✓	✓	0.2697	0.6834	0.2807
	✓	×	✓	0.2662	0.6816	0.2795
	✓	✓	×	0.2628	0.6913	0.2734
Full	✓	✓	✓	<b>0.2711</b>	<b>0.6913</b>	<b>0.2816</b>

Table 2. **Ablation Study.** We ablate our proposed Anti-Segregation Optimization (ASO), Noise Vector Balancing (NVB), and Noise Variance Rectification (NVR) techniques.

and  $S_{i,j}$  is the cosine similarity between the CLIP text embedding of the  $i$ -th text prompt and the CLIP image embedding of the  $j$ -th view of the generated image. We then compute three metrics to evaluate the quality of the generated visual anagrams:

1. **Worst Alignment Score ( $\mathcal{A}_{\min}$ ):** The minimum value in the diagonal of the score matrix  $S$ , which measures the worst alignment between the text prompts and the generated image across all views.
2. **Concealment Score ( $\mathcal{C}$ ):**  $\mathcal{C} = \frac{1}{N} \text{tr}(\text{softmax}(\frac{S}{\tau}))$ , where  $\text{tr}(\cdot)$  denotes the trace of a matrix,  $\tau$  is the temperature parameter of CLIP, and  $\text{softmax}$  is computed against both rows and columns of  $S$ . This metric measures how well the generated image conceals concepts in other text prompts under the current view.
3. **Average Alignment Score ( $\mathcal{A}_{\text{avg}}$ ):** The average value in the diagonal of the score matrix  $S$ , which measures the overall alignment across all views.

Actually,  $\mathcal{A}_{\min}$  and  $\mathcal{C}$  are proposed by [19], we further introduce  $\mathcal{A}_{\text{avg}}$  to demonstrate that our method not only balances the trade-off between different views but also boosts the overall text-image alignment. Note that we use a vanilla CLIP model that is only pre-trained on clean images to compute the metrics, which is different from the one used in our proposed method and is more suitable for evaluating the generated images, and also prevents our method from directly optimizing the metrics.

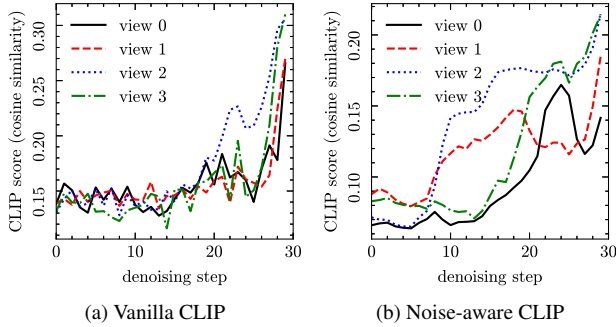


Figure 8. **Task completion scores comparison.** We compare the task completion scores measured by noise-aware and vanilla CLIP models during the denoising process using the baseline method.

NVB with Noise-aware CLIP	NVB with Vanilla CLIP	ASO & NVR	$\mathcal{A}_{\min} \uparrow$	$C \uparrow$	$\mathcal{A}_{avg} \uparrow$
×	✓	×	0.2609	<b>0.6868</b>	0.2726
✓	×	×	<b>0.2632</b>	0.6864	<b>0.2734</b>
×	✓	✓	0.2690	0.6838	0.2808
✓	×	✓	<b>0.2711</b>	<b>0.6913</b>	<b>0.2816</b>

Table 3. **Ablation Study.** We compare the performance of our proposed method using different CLIP models to measure task completion.

### 4.3. Results

**Quantitative Results.** We first report quantitative results of our proposed method against the baseline method, together with other existing methods on the 2-view dataset and the 3-view dataset in Tab. 1. For the 2-view dataset, identity and vertical flip are used as the two views, as they are supported by all methods. For the 3-view dataset, we add clockwise rotation as the third view. A total of 450 images are generated by each method for the 2-view dataset, and 1200 images for the 3-view dataset. We compare our method with the baseline method that straightforwardly averages the noise predictions from different text prompts, and the results demonstrate that our method outperforms the baseline method across all metrics on both datasets.

**Qualitative Results.** We also show qualitative results of our proposed methods to demonstrate their effectiveness. Fig. 5 shows an example when we use “a painting of a deer” and “a painting of a frog” as the text prompts to generate visual anagrams under the identity view and the vertical flip view. In the image generated by the baseline method, we could only see a small deer head in the identity view, where attention activates, and the rest of the image does not make sense for the deer prompt. In contrast, after applying our proposed Anti-Segregation Optimization, the denoising process learns to find commonalities between the two generation tasks and generates antlers in the identity view that can also be interpreted as frog legs in the flipped view, generates a deer head that can also be interpreted as the frog’s body,




Figure 9. Effect of target overlap ratio in Anti-Segregation Optimization.




Figure 10. **Correlation between noise vectors.** We record the cosine similarity and estimated correlation coefficient between the predicted noise vectors of different views at each denoising step.

and now the upper-right background region in the flipped view naturally interprets as the part of the deer’s body in the identity view.

In Fig. 7, we provide an example to compare the generation process and results before and after applying our proposed Noise Vector Balancing technique. We use “an oil painting of” as the style prefix and “a teddy”, “a bird”, “a rabbit”, and “a giraffe” as the text prompts to generate 4-view visual anagrams. As can be seen, in the image generated by the baseline method, the teddy bear is hardly recognizable in the identity view because the generation process is dominated by other tasks. After applying our proposed method, the teddy bear is more recognizable, and the task completion score curve also shows that the gap between the teddy and the other tasks is reduced, which demonstrates that our method effectively balances multiple generation tasks. Fig. 6 illustrates the necessity to preserve the variance of noise vectors during the denoising process. After applying Noise Variance Rectification, artifacts and distortions are reduced, and the overall quality of the generated image is improved. Fig. 4 shows more qualitative results, where our method generates more visually appealing images with fewer artifacts, and better align with text prompts, compared to the baseline method.

#### 4.4. Ablation Study

We ablate our proposed Anti-Segregation Optimization and Noise Vector Balancing techniques, together with the Noise Variance Rectification in Tab. 2. The results show that each component contributes to the overall performance of our method, and the combination of all components achieves the best results. Specifically, the Anti-Segregation Optimization effectively prevents concept segregation and primarily contributes to the improvement of  $\mathcal{C}$ , the Noise Vector Balancing technique balances the trade-off between different views and thus improves  $\mathcal{A}_{\min}$  and  $\mathcal{C}$ , and the Noise Variance Rectification corrects the denoising process and boosts overall text-image alignment.

### 5. Additional Results and Discussion

#### 5.1. Comparison with Vanilla CLIP Models

Existing works on leveraging auxiliary models for text-to-image generation, such as classifier guidance and CLIP guidance, have been shown to work better when the auxiliary models are pre-trained on noisy images, otherwise the intermediate images are out-of-distribution for the auxiliary models. Therefore, in our proposed method, we use a noise-aware CLIP model to measure the degree of task completion in the noise vector balancing process. Here we provide further analysis on this design choice.

Fig. 8 depicts a typical example of task completion scores measured by noise-aware and vanilla CLIP models during the denoising process when we use the baseline method to generate a 4-view visual anagram. We can observe that curves of task completion scores measured by the noise-aware CLIP model are smoother and less noisy compared to those measured by a vanilla CLIP model, which indicates that the noise-aware CLIP model provides more stable and more reliable task completion scores. Additionally, the noise-aware CLIP model starts to provide useful information earlier than the vanilla CLIP model, which helps the noise vector balancing process to take early actions to prevent concept domination. Furthermore, we also compare the performance of our proposed method using noise-aware and vanilla CLIP models in Tab. 3, and the results show that our method using noise-aware CLIP model generally outperforms the one using vanilla CLIP model, which further validates our design choice.

#### 5.2. Effect of Target Overlap Ratio in Anti-Segregation Optimization

In Eq. (1), we introduce a hyperparameter  $\phi$ , which describes the desired overlap ratio between attention maps of different views and takes values in the range  $[0, 0.5]$ . Generally, a larger  $\phi$  encourages more overlap between attention maps, while a smaller  $\phi$  encourages them to be more segregated, and  $\phi = 0.5$  favors complete overlap. We investigate

the effect of  $\phi$  on the performance of our proposed Anti-Segregation Optimization (ASO) with other techniques disabled on the 2-view dataset and the results are shown in Fig. 9. We can observe that our method achieves the best performance around  $\phi = 0.45$ .

#### 5.3. Correlation between Noise Vectors

While it is true that statistical property of the noise vectors is crucial for the denoising diffusion models to reverse the diffusion process, one may question why simply averaging the noise predictions from different text prompts still has a chance to generate sensible images, because averaging  $N$  independent standard Gaussian noise vectors should result in a noise vector with zero mean and variance of  $1/N$ , and such significant variance shift seems to be fatal.

Therefore, we investigate the correlation between the noise vectors at each denoising step by recording cosine similarities and estimated correlation coefficients using Eq. (12) between them. Fig. 10 shows a typical example of the denoising process, when we use the baseline method to generate a 2-view visual anagram, and we can observe that the correlation coefficient and cosine similarity are close to one. Therefore, in Eq. (10), where  $\alpha_t^i = \frac{1}{N}$  for the baseline method, the scale factor  $c$  is close to one, which explains why the baseline method is still possible to generate images that do not drift too far from the data distribution. Anyhow, our proposed method is beneficial for better generating visual anagrams, as it explicitly corrects the denoising process.

### 6. Conclusions and Limitations

In this work, we link visual anagram generation with multi-task learning and draw inspiration from the latter to address the challenges of concept separation and concept domination in the former. We propose an Anti-Segregation Optimization technique that encourages concepts in different views to overlap, and a Noise Vector Balancing technique that reweights the noise predictions based on task completion scores to balance the denoising process. We further propose to scale the combined noise vector to preserve the variance of the noise. Our experiments demonstrate that our proposed method outperforms the baseline method consistently across different metrics and datasets.

Despite the promising results, our work has several limitations. Firstly, we are still unable to leverage powerful latent space diffusion models like Stable Diffusion to generate good visual anagrams because latents and the decoded images are not consistent across transformations as discussed in [19]. Secondly, our method could not extend anagram generation to non-orthogonal transformations like non-volume-preserving deformations, which is an interesting direction for future work.



## References

- [1] Stability ai releases deepfloyd if, a powerful text-to-image model that can smartly integrate text into images. <https://stability.ai/news/deepfloyd-if-text-to-image-model>. 2, 6
- [2] Aishwarya Agarwal, Srikrishna Karanam, K. J. Joseph, Apoorv Saxena, Koustava Goswami, and Balaji Vasan Srinivasan. A-star: Test-time attention segregation and retention for text-to-image synthesis, 2023. 2
- [3] Zhipeng Bao, Yijun Li, Krishna Kumar Singh, Yu-Xiong Wang, and Martial Hebert. Separate-and-enhance: Compositional finetuning for text2image diffusion models, 2024. 2
- [4] Edwin G Boring. A new ambiguous figure. *The American Journal of Psychology*, 1930. 1
- [5] Ryan Burgert, Xiang Li, Abe Leite, Kanchana Ranasinghe, and Michael Ryoo. Diffusion Illusions: Hiding Images in Plain Sight. In *ACM SIGGRAPH 2024 Conference Papers*, SIGGRAPH '24, pages 1–11, New York, NY, USA, July 2024. Association for Computing Machinery. 1, 2, 6
- [6] Kartik Chandra, Tzu-Mao Li, Joshua Tenenbaum, and Jonathan Ragan-Kelley. Designing perceptual puzzles by differentiating probabilistic programs. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings*, pages 1–9, 2022. 2
- [7] 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, 2023. 2
- [8] Mingjin Chen, Junhao Chen, Xiaojun Ye, Huan-ang Gao, Xiaoxue Chen, Zhaoxin Fan, and Hao Zhao. Ultraman: Single image 3d human reconstruction with ultra speed and detail. *arXiv preprint arXiv:2403.12028*, 2024. 2
- [9] Xiaoxue Chen, Tianyu Liu, Hao Zhao, Guyue Zhou, and Ya-Qin Zhang. Cerberus transformer: Joint semantic, affordance and attribute parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19649–19658, 2022. 3
- [10] Xiaoxue Chen, Hao Zhao, Guyue Zhou, and Ya-Qin Zhang. Pq-transformer: Jointly parsing 3d objects and layouts from point clouds. *IEEE Robotics and Automation Letters*, 7(2):2519–2526, 2022. 3
- [11] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks, 2018. 3, 4
- [12] Zhao Chen, Jiquan Ngiam, Yanping Huang, Thang Luong, Henrik Kretschmar, Yuning Chai, and Dragomir Anguelov. Just pick a sign: Optimizing deep multitask models with gradient sign dropout. <https://arxiv.org/abs/2010.06808v1>, 2020. 3
- [13] Hung Kuo Chu, Wei Hsin Hsu, Niloy J. Mitra, Daniel Cohen-Or, Tien-Tsin Wong, and Tong-Yee Lee. Camouflage images. *ACM Transactions on Graphics*, 29(4):51, 2010. 2
- [14] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. <https://arxiv.org/abs/2105.05233v4>, 2021. 1, 2
- [15] Sylvia Ferino-Pagden. *Arcimboldo*. Skira, 2017. 1
- [16] William T. Freeman, Edward H. Adelson, and David J. Heeger. Motion without movement. *ACM SIGGRAPH Computer Graphics*, 25(4):27–30, 1991. 2
- [17] Huan-ang Gao, Mingju Gao, Jiaju Li, Wenyi Li, Rong Zhi, Hao Tang, and Hao Zhao. Scp-diff: Photo-realistic semantic image synthesis with spatial-categorical joint prior. *arXiv preprint arXiv:2403.09638*, 2024. 2
- [18] Daniel Geng, Inbum Park, and Andrew Owens. Visual anagrams: Generating multi-view optical illusions with diffusion models. <https://arxiv.org/abs/2311.17919v2>, 2023. 2, 6
- [19] Daniel Geng, Inbum Park, and Andrew Owens. Visual anagrams: Generating multi-view optical illusions with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24154–24163, 2024. 1, 2, 3, 4, 6, 8
- [20] Alexander Gomez-Villa, Adrian Martin, Javier Vazquez-Corral, and Marcelo Bertalmio. Convolutional neural networks can be deceived by visual illusions. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12301–12309, Long Beach, CA, USA, 2019. 1, 2
- [21] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors, 2023. 2
- [22] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrian Gaidon. Semantically-guided representation learning for self-supervised monocular depth. <https://arxiv.org/abs/2002.12319v1>, 2020. 3
- [23] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multitask learning. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 282–299, Cham, 2018. 3
- [24] Rui Guo, Jasmine Collins, Oscar de Lima, and Andrew Owens. Ganmouflage: 3d object nondetection with texture fields, 2023. 2
- [25] Aaron Hertzmann. Visual indeterminacy in gan art. *Leonardo*, 53(4):424–428, 2020. 2
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. 1, 2, 5
- [27] Hiroaki Igarashi. Multi-task curriculum learning based on gradient similarity. *BMVC*, 2022. 3
- [28] Priyank Jaini, Kevin Clark, and Robert Geirhos. Intriguing properties of generative classifiers, 2024. 1, 2
- [29] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics, 2018. 3
- [30] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. 6
- [31] Wenyi Li, Haoran Xu, Guiyu Zhang, Huan-ang Gao, Mingju Gao, Mengyu Wang, and Hao Zhao. Fairdiff: Fair segmentation with point-image diffusion. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 617–628. Springer, 2024. 2
- [32] Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qingfu Zhang, and Sam Kwong. Pareto multi-task learning, 2019. 3

- [33] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B. Tenenbaum. Compositional visual generation with composable diffusion models. <https://arxiv.org/abs/2206.01714v6>, 2022. 2
- [34] Ming Lu, Hao Zhao, Anbang Yao, Yurong Chen, Feng Xu, and Li Zhang. A Closed-Form Solution to Universal Style Transfer. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5951–5960, Seoul, Korea (South), Oct. 2019. IEEE. 2
- [35] Ming Lu, Hao Zhao, Anbang Yao, Feng Xu, Yurong Chen, and Li Zhang. Decoder network over lightweight reconstructed feature for fast semantic style transfer. In *Proceedings of the IEEE international conference on computer vision*, pages 2469–2477, 2017. 2
- [36] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning, 2016. 3
- [37] Aviv Navon, Aviv Shamsian, Idan Achituve, Haggai Maron, Kenji Kawaguchi, Gal Chechik, and Ethan Fetaya. Multi-task learning as a bargaining game. <https://arxiv.org/abs/2202.01017v2>, 2022. 3
- [38] Jerry Ngo, S. Sankaranarayanan, and Phillip Isola. Is clip fooled by optical illusions? In *Tiny Papers @ ICLR*, 2023. 1, 2
- [39] 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, 2022. 6
- [40] Michael E. R. Nicholls, Owen Churches, and Tobias Loetscher. Perception of an ambiguous figure is affected by own-age social biases. *Scientific Reports*, 8(1):12661, Aug. 2018. Publisher: Nature Publishing Group. 1
- [41] Aude Oliva, Antonio Torralba, and Philippe G. Schyns. Hybrid images. *ACM Trans. Graph.*, 25(3):527–532, 2006. 2
- [42] Andrew Owens, Connelly Barnes, Alex Flint, Hanumant Singh, and William Freeman. Camouflaging an object from many viewpoints. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2782–2789, 2014. 2
- [43] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2
- [44] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. <https://arxiv.org/abs/2102.12092v2>, 2021. 2
- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. <https://arxiv.org/abs/2112.10752v2>, 2021. 2, 3
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 3
- [47] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 2
- [48] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization, 2019. 3
- [49] Dmitry Senushkin, Nikolay Patakin, Arseny Kuznetsov, and Anton Konushin. Independent component alignment for multi-task learning, 2023. 3
- [50] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. <https://arxiv.org/abs/1503.03585v8>, 2015. 1, 2
- [51] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. <https://arxiv.org/abs/2010.02502v4>, 2020. 2
- [52] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. <https://arxiv.org/abs/2011.13456v2>, 2020. 1, 2
- [53] Ximeng Sun, Rameswar Panda, Rogerio Feris, and Kate Saenko. Adashare: Learning what to share for efficient deep multi-task learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 8728–8740. Curran Associates, Inc., 2020. 3
- [54] Matthew Tancik. tancik/Illusion-Diffusion, July 2024. original-date: 2023-02-12T22:39:28Z. 2, 6
- [55] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 3
- [56] Tengfei Wang, Ting Zhang, Bo Zhang, Hao Ouyang, Dong Chen, Qifeng Chen, and Fang Wen. Pretraining is all you need for image-to-image translation, 2022. 5
- [57] Xi Wang, Zoya Bylinskii, Aaron Hertzmann, and Robert Pepperell. Toward quantifying ambiguities in artistic images, 2020. 2
- [58] Derrick Xin, Behrooz Ghorbani, Ankush Garg, Orhan Firat, and Justin Gilmer. Do current multi-task optimization methods in deep learning even help?, 2022. 3
- [59] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning, 2020. 3, 4
- [60] Guiyu Zhang, Huan-ang Gao, Zijian Jiang, Hao Zhao, and Zhedong Zheng. Ctrl-u: Robust conditional image generation via uncertainty-aware reward modeling. *arXiv preprint arXiv:2410.11236*, 2024. 1
- [61] Yu Zhang and Qiang Yang. An overview of multi-task learning. *National Science Review*, 5(1):30–43, 09 2017. 3
- [62] Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE transactions on knowledge and data engineering*, 34(12):5586–5609, 2021. 3
- [63] Yupeng Zheng, Chengliang Zhong, Pengfei Li, Huan-guy Gao, Yuhang Zheng, Bu Jin, Ling Wang, Hao Zhao, Guyue Zhou, Qichao Zhang, et al. Steps: Joint self-supervised nighttime image enhancement and depth estimation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4916–4923. IEEE, 2023. 3