

Guiding Diffusion with Deep Geometric Moments: Balancing Fidelity and Variation

Sangmin Jung^{1†} Utkarsh Nath¹ Yezhou Yang¹ Giulia Pedrielli¹
Joydeep Biswas² Amy Zhang² Hassan Ghasemzadeh¹ Pavan Turaga¹
¹Arizona State University ²The University of Texas at Austin




Figure 1. Deep Geometric Moments (DGM) are used as guidance to capture nuanced subject details from the reference image while preserving the generative diversity of the diffusion model. Each generated result is conditioned on the prompt “a photo of <animal name>”. For each pair, the left image is the reference, and the right image is the generated result. The small image in the bottom-left corner of the reference is the DGM visualization used for guidance.

Abstract

Text-to-image generation models have achieved remarkable capabilities in synthesizing images, but often struggle to provide fine-grained control over the output. Existing guidance approaches, such as segmentation maps and depth maps, introduce spatial rigidity that restricts the inherent diversity of diffusion models. In this work, we introduce Deep Geometric Moments (DGM) as a novel form of guidance that encapsulates the subject’s visual features and nuances through a learned geometric prior. DGMs focus specifically on the subject itself compared to DINO or CLIP features, which suffer from overemphasis on global image features or semantics. Unlike ResNets, which are sensitive to pixel-wise perturbations, DGMs rely on robust geometric moments. Our experiments demonstrate that DGM effectively balance control and diversity in diffusion-based image generation, allowing a flexible control mechanism for steering the diffusion process.

1. Introduction

Diffusion models have emerged as a powerful text-to-image (T2I) generation framework, enabling high-quality and diverse synthesis from user input text prompts [12, 18, 19]. Early diffusion methods relied on input gradients from a

pre-trained classifier model to guide the generation process. This need was later eliminated by Classifier-free-guidance [11, 25], which integrated the conditioning mechanism directly into the generation process. Many large-scale image generation models, such as Stable Diffusion [23] and DALL-E [22], advance text-to-image generation using CLIP encoders and cross-attention to align text and image features, achieving generation with much higher fidelity and realism. Despite their impressive generative quality, they often lack fine-grained control over output.

To address this, methods like ControlNet [37] and IP-Adapter [32] introduce auxiliary control structures, such as depth maps, Canny edges, segmentation masks, and pose estimations. These methods learn to steer the generation process by applying cross-attention with embeddings of control structures. However, these methods require dedicated training phases to align these guidance signals with the generation pipeline. Recent works have introduced training-free methods for guided image generation [35]. These methods do not require retraining but modify the sampling process or leverage external models. Various types of guidance have been explored [13–15], including placing bounding boxes to specify regions to generate, providing face IDs (with distinct facial features like nose and lips) to replicate human faces, utilizing CLIP text guidance, and applying style transfer based on CLIP features.

Despite recent advancements, training-based and training-free methods predominantly rely on control signals or global

[†] Corresponding author: sjung61@asu.edu



Figure 2. Failure cases of guidance based methods. Spatial rigidity uses segmentation maps as guidance [3]; Semantic rigidity shows results of CLIP image features as guidance.

features such as CLIP [21]. As shown in Figure 2, this often leads to generated imagery that either rigidly follow the control signal or firmly adhering to the semantic alignment, failing to preserve the fine-grained visual features of the reference object. To enable a more diverse yet fidelity-preserving generation, it is essential that image generation models not only align with the intended control but also retain creative flexibility beyond rigid structural constraints.

To this extent, we propose using geometric moments for guidance in a training-free fashion. Geometric moments define various shape characteristics and can be visualized as projections of the image onto chosen basis functions. We propose a guidance algorithm that leverages an auxiliary off-the-shelf pretrained Deep Geometric Moment (DGM) [16, 27] model during the image sampling process of a diffusion model. The DGM approach is suited to extract features capable of capturing object textures and other features in a spatially invariant manner allowing for diverse generation while preserving subject identity. Our contribution can be summarized as:

- We introduce the Deep Geometric Moments (DGM) as a unique guidance signal for the diffusion process, enabling effective preservation of the visual details of subjects.
- We present our pipeline in a training-free setting, eliminating the need for training networks for guided generation by utilizing an off-the-shelf pretrained DGM model.
- We demonstrate that our generated results maintain the diversity of the diffusion model, unlike other fixed control methods such as segmentation maps or depth maps.

2. Related Works

Conditioned Generation approaches generally rely on training diffusion models specifically tailored to various conditions using class labels [6]. For instance, Ho et al. [11] proposed classifier-free guidance, where diffusion models are trained to blend conditional and unconditional outputs according to provided class labels. Extending the concept further, CLIP Guidance [17] introduced the use of detailed textual descriptions as prompts, aligning generated images with text embeddings from CLIP. However, their reliance on

training separate classifiers limits practical scalability.

Guided Generation methods, in contrast, aim to keep diffusion models fixed and modify only the sampling procedure to achieve targeted outcomes without retraining. Methods like Universal Guidance [3], FreeDOM [34] introduced flexible training-free frameworks that integrate various guidance signals such as CLIP embeddings for style transfer, segmentation maps for spatial control, face ids for replicating human face, demonstrating broader applicability. Further works have been proposed to improve the operation of noise estimation for steering the output [10, 29], and analyze an algorithm-agnostic design space of training-free guidance altogether [5, 8, 26, 30, 31, 33].

3. Method

Deep Geometric Moments (DGM) [16, 27] are a feature representation designed to capture the geometric moment information from images. Geometric moment computes shape descriptors by integrating pixel-wise features over a spatial domain:

$$M_{p,q} = \int \int x^p y^q f(x, y) dx dy, \quad (1)$$

where, $M_{p,q}$ represents the geometric moment of order (p, q) , and $f(x, y)$ is the function describing the image features at location (x, y) . Unlike traditional moment-based descriptors, which rely on predefined basis functions, DGM leverages deep neural networks to learn hierarchical representations that encode geometric attributes of objects. By applying neural networks to extract and refine these moment-based features, DGM produces representations that are robust to variations in scale, rotation, and appearance.

Guidance Sampling with DGM Features To steer generation towards specific outputs based on the reference input, we incorporate universal guidance algorithm using features from a pre-trained DGM encoder. Traditional classifier guidance modifies the predicted noise using the gradient of a classifier’s log-likelihood:

$$\epsilon_\theta(z_t, t) \leftarrow \epsilon_\theta(z_t, t) - \sqrt{1 - \alpha_t} \nabla_{z_t} \log p_\phi(c|z_t), \quad (2)$$

where, ϵ_θ is the guided noise prediction, and $p_\phi(c|z_t)$ is the probability of class c predicted by a classifier p_ϕ . However, this method has a limitation such that the classifiers are not trained on noisy latent images z_t , and thus they cannot reliably predict class labels at intermediate steps of the diffusion process. We instead compute the classifier signal on the estimated clean image \hat{z}_0 , derived at each step using the DDIM [28] formulation:

$$\hat{z}_0 = \frac{z_t - \sqrt{1 - \alpha_t} \epsilon_\theta(z_t, t)}{\sqrt{\alpha_t}}. \quad (3)$$

Table 1. **Quantitative comparison of different methods.** We distinguish between training-based and training-free methods. Methods listed below the first dashed line are implemented on top of Universal Guidance. “Prompt Ada.” (Prompt Adaptability) indicates whether a method can incorporate changes in text prompts. CLIP-I measures visual consistency, I-DINO measures diversity, and ChatGPT scores reflect image quality to supplement CLIP-I. For all metrics, higher scores correspond to better performance. However, for I-DINO score nearing 0.5 seem to imply inconsistent output as seen in qualitative results in figure 3.

Models	Training-Free	Prompt Ada.	CLIP-I (↑)	I-DINO (↑)	ChatGPT (↑)
SD-Img2Img (RGB)	×	✓	0.8601	0.2922 (balanced)	1.88
IP-Adapter (RGB)	×	✓	0.8947	0.1086 (stagnant)	2.63
ControlNet-Depth (RGB+Depth)	×	✓	0.8606	0.1819 (stagnant)	2.05
DINO	✓	×	0.8735	0.1682 (stagnant)	2.18
ResNet34	✓	×	0.7844	0.4847 (divergent)	1.41
Segmentation Maps	✓	✓	0.7431	0.5582 (divergent)	0.69
CLIP	✓	✓	0.7480	0.4787 (divergent)	0.39
Ours	✓	✓	0.8323	0.2754 (balanced)	1.85

This step reconstructs a noise-free image estimate \hat{z}_0 from the current noisy latent z_t , enabling the application of classifiers or feature extractors trained on clean images. We leverage the clean estimate \hat{z}_0 to provide meaningful feedback for guidance sampling. We extend classifier guidance by replacing the classifier with a feature extractor f , and guiding generation using a feature-specific loss function l :

$$\epsilon_\theta(z_t, t) \leftarrow \epsilon_\theta(z_t, t) + s(t) \cdot \nabla_{z_t} \ell(c, f(\hat{z}_0)) \quad (4)$$

where, $s(t)$ is a scaling factor controlling the strength of guidance, and the loss function for our pipeline is defined as:

$$l = \text{MSE}(f(z_{\text{ref}}), f(\hat{z}_0)). \quad (5)$$

Here, z_{ref} is a reference image with the target subject, and f extracts visual features from the estimated clean image. This formulation enables guidance based on each feature’s distinct characteristics, rather than relying on less fine-grained and less informative raw class labels. As observed in prior works [4, 20], a single-step correction is often insufficient to effectively steer the generation process. Therefore, we adopt a repetitive per-step correction strategy following the Universal guidance. Specifically, we re-inject random Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$ into z_{t-1} to obtain z'_t as follows:

$$z'_t = \sqrt{\alpha_t / \alpha_{t-1}} \cdot z_{t-1} + \sqrt{1 - \alpha_t / \alpha_{t-1}} \cdot \epsilon. \quad (6)$$

This formulation ensures that z'_t has the appropriate noise scale corresponding to timestep t , allowing for more controlled guidance by better exploring the distribution at the subsequent step.

4. Experiments

4.1. Experimental Setting

Datasets and Baselines. We curated a unique dataset, *DGM-Bench*, which focuses on animals with distinct features and

nuanced visual details. Our baseline evaluations cover three key aspects: (1) the use of diverse feature descriptors for guidance, (2) the distinction between training-based and training-free methods, and (3) the incorporation of various input modalities beyond RGB images, including segmentation and depth maps. Additional details about the dataset and the experimental setup are provided in the supplement.

4.2. Quantitative Comparison

We evaluate the image generation quality based on two main criteria: the ability to preserve fine-grained visual details from the reference image, and the capacity to produce diverse outputs during the diffusion process. We report the CLIP-I score and GPT-based [1] preference scores to measure visual consistency. CLIP-I score is the average cosine similarity between the CLIP image embeddings of the reference and generated images. However, CLIP-I alone may not capture all aspects of visual fidelity. To complement it, we introduce a preference score using the *ChatGPT-4o-latest* model on image inputs. The model ranks the outputs based on pre-defined qualitative rubrics from 0 to 4. Further details on this evaluation setup are provided in the supplementary material. To quantify variation, we propose I-DINO, which uses the inverse of the average pairwise similarity across DINO[2] embedding of the generated images.

$$\text{I-DINO} = 1 - \frac{1}{n} \sum_{i,j} \text{DINO}(\text{img}_i, \text{img}_j) \quad (7)$$

Table 1 summarizes the results across visual fidelity (CLIP-I), variation (I-DINO), and perceived image quality (ChatGPT score). Among the training-free methods, our approach achieves the second-highest CLIP-I score after DINO. However, unlike DINO, which replicates the reference image, our method strikes a better balance at preserving subject identity while enabling diverse outputs. In contrast, the best performing training-based methods are limited in diversity



Figure 3. **Qualitative Comparison of different methods.** The leftmost image in each row is the reference image. Our method, which uses DGM as guidance, is followed by four alternative feature descriptors and two training-based methods. Overall, our DGM-based guidance achieves a balanced trade-off between visual consistency and generative diversity in the diffusion output.

and often generate highly similar output, and sometimes fail to capture the subject’s fine features. Our DGM-guided method and SD-Img2Img preserve visual fidelity while also producing more varied generation, demonstrating the effectiveness of DGM in balancing fidelity and variation.

4.3. Qualitative Comparison

For the qualitative comparison, we present Figure 3, which showcases the results of our method. Our approach demonstrates that DGM effectively captures the nuanced details and textures of the reference image, whereas naive ResNet [9] features are more prone to failure due to their lack of robustness against pixel-wise variations. DINO features tend to replicate the input reference image almost entirely, but lack flexibility because of their strong patch-wise global matching behavior. CLIP maintains semantic alignment; however, preserving fine-grained visual details remains challenging. The segmentation maps and the ControlNet depth results show that both methods can maintain the object’s location according to their design purpose, but occasionally fail by overly constraining the generated results to a designated area. Overall, IP-Adapter performs well, though it often renders visual details with a specific stylistic bias and sometimes introduces undesired elements, such as generating a human figure on a horse, as shown in the figure. In general, sub-

ject identity is captured more accurately than with simple guidance-based methods. Additional examples are provided in the supplemental materials.

5. Conclusion

In this work, we introduce a novel guidance method, Deep Geometric Moments (DGM), for guiding text-to-image generation. We demonstrate that utilizing a feature vector from the DGM model effectively transfers the visual details of the subject. Using this feature, we successfully maintain visual consistency with the reference image while preserving the diffusion model’s output diversity. Our approach is more effective than the feature guidance from ResNet, which often fails due to its sensitivity to pixel-level perturbations. We also conduct an extensive comparison with various feature descriptors, such as DINO and CLIP, revealing their limitations, which include adhering too closely to global image representation or overly emphasizing semantic consistency. In addition, although training-based methods achieve higher fidelity to their inputs, they significantly reduce the generative diversity inherent to diffusion models. Through extensive quantitative and qualitative evaluations, our results demonstrate that our approach effectively balances subject identity preservation and output diversity [7] without requiring additional training.

Acknowledgements

This paper was supported by NSF grant 2402650 and in part by NSF grant 2323086. We thank the Research Computing (RC) at Arizona State University (ASU) and cr8dl.ai for their generous support in providing computing resources. We would also like to express our gratitude to Suren Jayasuriya for insightful guidance that informed the development of this project.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2(3):4, 2021. 3
- [3] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 843–852, 2023. 2
- [4] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022. 3, 1
- [5] Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G. Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems. *ArXiv*, abs/2410.00083, 2024. 2
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2
- [7] Rohit Gandikota and David Bau. Distilling diversity and control in diffusion models. *arXiv preprint arXiv:2503.10637*, 2025. 4
- [8] Yingqing Guo, Yukang Yang, Hui Yuan, and Mengdi Wang. Training-free guidance beyond differentiability: Scalable path steering with tree search in diffusion and flow models. *ArXiv*, abs/2502.11420, 2025. 2
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [10] Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, et al. Manifold preserving guided diffusion. *arXiv preprint arXiv:2311.16424*, 2023. 2
- [11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 1, 2
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1
- [13] Jiannan Huang, Jun Hao Liew, Hanshu Yan, Yuyang Yin, Yao Zhao, Humphrey Shi, and Yunchao Wei. Classdiffusion: More aligned personalization tuning with explicit class guidance. *arXiv preprint arXiv:2405.17532*, 2024. 1
- [14] Grace Luo, Trevor Darrell, Oliver Wang, Dan B Goldman, and Aleksander Holynski. Readout guidance: Learning control from diffusion features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8217–8227, 2024.
- [15] Nithin Gopalakrishnan Nair, Anoop Cherian, Suhas Lohit, Ye Wang, Toshiaki Koike-Akino, Vishal M Patel, and Tim K Marks. Steered diffusion: A generalized framework for plug-and-play conditional image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20850–20860, 2023. 1
- [16] Utkarsh Nath, Rajhans Singh, Ankita Shukla, Kuldeep Kulkaarni, and Pavan Turaga. Polynomial implicit neural framework for promoting shape awareness in generative models. *International Journal of Computer Vision*, pages 1–29, 2024. 2
- [17] 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. 2
- [18] Maitreya Patel, Sangmin Jung, Chitta Baral, and Yezhou Yang. λ -ECLIPSE: Multi-concept personalized text-to-image diffusion models by leveraging CLIP latent space. *Transactions on Machine Learning Research*, 2024. 1
- [19] Maitreya Patel, Changhoon Kim, Sheng Cheng, Chitta Baral, and Yezhou Yang. Eclipse: A resource-efficient text-to-image prior for image generations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9069–9078, 2024. 1
- [20] Maitreya Patel, Song Wen, Dimitris N Metaxas, and Yezhou Yang. Steering rectified flow models in the vector field for controlled image generation. *arXiv preprint arXiv:2412.00100*, 2024. 3
- [21] 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
- [22] 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. 1
- [23] 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
- [24] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on com-*

- puter vision and pattern recognition, pages 22500–22510, 2023. 1
- [25] Seyedmorteza Sadat, Manuel Kansy, Otmar Hilliges, and Romann M. Weber. No training, no problem: Rethinking classifier-free guidance for diffusion models. *ArXiv*, abs/2407.02687, 2024. 1
 - [26] Yifei Shen, Xinyang Jiang, Yezhen Wang, Yifan Yang, Dongqi Han, and Dongsheng Li. Understanding and improving training-free loss-based diffusion guidance. In *Neural Information Processing Systems*, 2024. 2
 - [27] Rajhans Singh, Ankita Shukla, and Pavan Turaga. Improving shape awareness and interpretability in deep networks using geometric moments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4159–4168, 2023. 2
 - [28] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
 - [29] Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation. In *International Conference on Machine Learning*, pages 32483–32498. PMLR, 2023. 2
 - [30] Kaiyu Song and Hanjiang Lai. Unraveling the connections between flow matching and diffusion probabilistic models in training-free conditional generation. *ArXiv*, abs/2411.07625, 2024. 2
 - [31] Tongda Xu, Xiyan Cai, Xinjie Zhang, Xingtong Ge, Dailan He, Limin Sun, Jingjing Liu, Ya-Qin Zhang, Jian Li, and Yan Wang. Rethinking diffusion posterior sampling: From conditional score estimator to maximizing a posterior. *ArXiv*, abs/2501.18913, 2025. 2
 - [32] Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 1
 - [33] Haotian Ye, Haowei Lin, Jiaqi Han, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Y Zou, and Stefano Ermon. Tfg: Unified training-free guidance for diffusion models. *Advances in Neural Information Processing Systems*, 37:22370–22417, 2024. 2
 - [34] Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. Freedom: Training-free energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23174–23184, 2023. 2
 - [35] Stefano Zampini, Jacob Christopher, Luca Oneto, Davide Anguita, and Ferdinando Fioretto. Training-free constrained generation with stable diffusion models. *ArXiv*, abs/2502.05625, 2025. 1
 - [36] Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. *arXiv preprint arXiv:2407.01521*, 2024. 1
 - [37] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 1

Guiding Diffusion with Deep Geometric Moments: Balancing Fidelity and Variation

Supplementary Material

6. Preliminaries for Diffusion models

The forward process in diffusion models can be formulated as:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{\alpha_t}x_{t-1}, (1 - \alpha_t)I)$$

where x_t denotes predicted latent at timestep t , α_t is the pre-defined variance schedule and I is the identity matrix. Then our x_t equation can be also written equivalently as:

$$x_t = x_0\sqrt{\alpha_t} + \epsilon\sqrt{1 - \alpha_t}, \quad \epsilon \sim \mathcal{N}(0, I)$$

Then our neural network with trainable parameters θ , and our reverse step samples:

$$p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$$

which can be obtained by minimizing the variational lower bound of the negative log likelihood (ELBO) of the distribution.

7. Dataset Details

For evaluating the task of replicating the visual information of the subject in the input reference image, we curated our dataset, DGMbench. DGMbench focuses on subjects with fine details. The dataset was primarily created by utilizing a subset of the Dreambench dataset [24], including dogs and cats, while excluding general non-fine-grained subjects such as vases and bowls. Additionally, we included a variety of animals with diverse appearances, such as birds and fish, while avoiding uniformly looking animals like elephants and lions. The dataset consists of 30 subjects in total.

8. Experiment Details

Rationale of selecting Universal Guidance as backbone

Through empirical experiments, we observed that the Universal Guidance mechanism outperforms other guidance methods and inverse problem solvers in diffusion, with expense of some additional computational overhead. [4, 36]

Explanation of feature descriptors The original DGM model was trained with an architectural structure (3, 4, 6, 3 intermediate layers), replicating the ResNet34 design, for the same goal of image classification. We selected ResNet34 as a comparison to evaluate the performance of its learned features and to validate whether ResNet-trained features are sufficiently robust to serve as effective inputs for the guidance loss. DINO image features, trained in large-scale

self-supervised settings, are widely used across various computer vision tasks but due to its adherence to capture the whole image representation, not utilized for purposes such as guiding the generation. In contrast, CLIP is primarily for its use with its text features, which have been employed in methods such as GLIDE. However, in our experiments, we focus on CLIP’s image features to assess their effectiveness for our tasks. Segmentation maps are also included as part of the baseline Universal Guidance approach, and we evaluate their use for showing the diversity in generation. Based on empirical evaluation, the guidance scales and recurrence steps for each method were selected as shown in Table 2.

Method	Guidance Scale	Recurrence Steps	Feature Size	Loss Function
Ours	10000	10	[1, 256]	MSE
ResNet	2000	10	[1, 512]	MSE
DINO	1000	10	[197, 384]	CosSim
CLIP	10	10	[1, 768]	CosSim
Seg	400	10	img_size	CE

Table 2. Guidance scale and number of recurrence steps used for different methods.

To ensure a consistent evaluation environment, we adopt Stable Diffusion v1.5 as the common backbone for all baseline methods, including ControlNet-Depth, SD-Img2Img, and IP-Adapter. In addition, the number of sampling steps is fixed at 500 across all methods to unify the experiment conditions.

9. Prompt for ChatGPT Evaluation

```
1 [Task]
2 You are a human evaluator assessing visual
  similarity between a reference image and
  a comparison image. Your focus is on
  low-level visual consistency - this
  includes color patterns, texture
  fidelity, and fine-grained inner
  geometric details (such as fur, stripes,
  surface patterns).
3 Ignore factors like the subject’s pose,
  overall shape, or semantic realism.
4
5 [Evaluation]
6 Please assign a score from 0 to 4 based on
  how well the comparison image preserves
  the visual essence of the reference
  image. Use the following rubric:
7
8 [Evaluation Criteria]
```


Score from 0-4 based on:

0 Completely inconsistent --Major differences in color, texture, and detail. The comparison image does not preserve any recognizable fine details from the reference.

1 Weak consistency - Some minor resemblance in visual traits, but largely different. Texture or color patterns are either distorted or missing.

2 Moderate consistency - There are identifiable shared traits (like fur direction or partial texture), but notable mismatches in color tones or local structures.

3 Strong consistency - Most key visual attributes like texture patterns, color palettes, and local details are preserved, with only slight differences.

4 Excellent consistency - Nearly identical in terms of texture, color, and fine-grained visual cues. Differences, if any, are negligible and hardly noticeable.

[Output]

Using the above rubric, compare the following two images and return a single score (0-4), along with a one-sentence justification of the score.

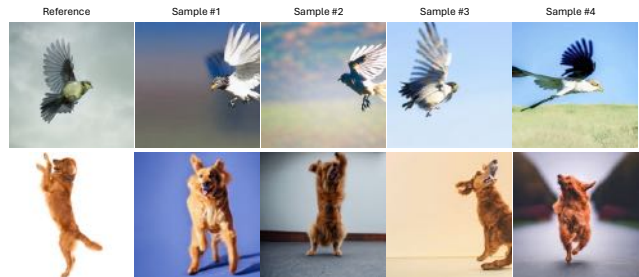


Figure 4. Failure examples of our method.

11. Additional Qualitative Results

More extensive qualitative comparison is shown in following Figures 5–8. For optimal clarity, the figures are best viewed in color prints.

12. Future Works

Based on our findings, we hypothesize that DGM can be leveraged as a prior to enhance the performance of existing personalized, subject-driven text-to-image (T2I) models, a direction we leave for future work.

10. Limitations and Failure Cases

The performance of training-free guidance is highly dependent on the choice of feature descriptors and their combination with the guidance parameters. The baseline image generator, Stable Diffusion, occasionally produces unwanted artifacts when the guidance strength is too high or when the guidance does not properly align with the target distribution of the diffusion model, making careful hyperparameter selection essential. Subtle mismatches between the guidance signal and the diffusion process can further result in unnatural images or artifacts, as observed in the CLIP- and ResNet-guided results.

As with other models, our model fails when the subject in the reference image has dynamic shapes or poses. As shown in Figure 4, when the original dog is standing on its hind legs or the bird is flying in a dynamic position, such semantic difficulties lead to degraded results. Additionally, if the subject has intricate features of various details and textures, the results fail, presumably due to the learned geometric prior being insufficient.

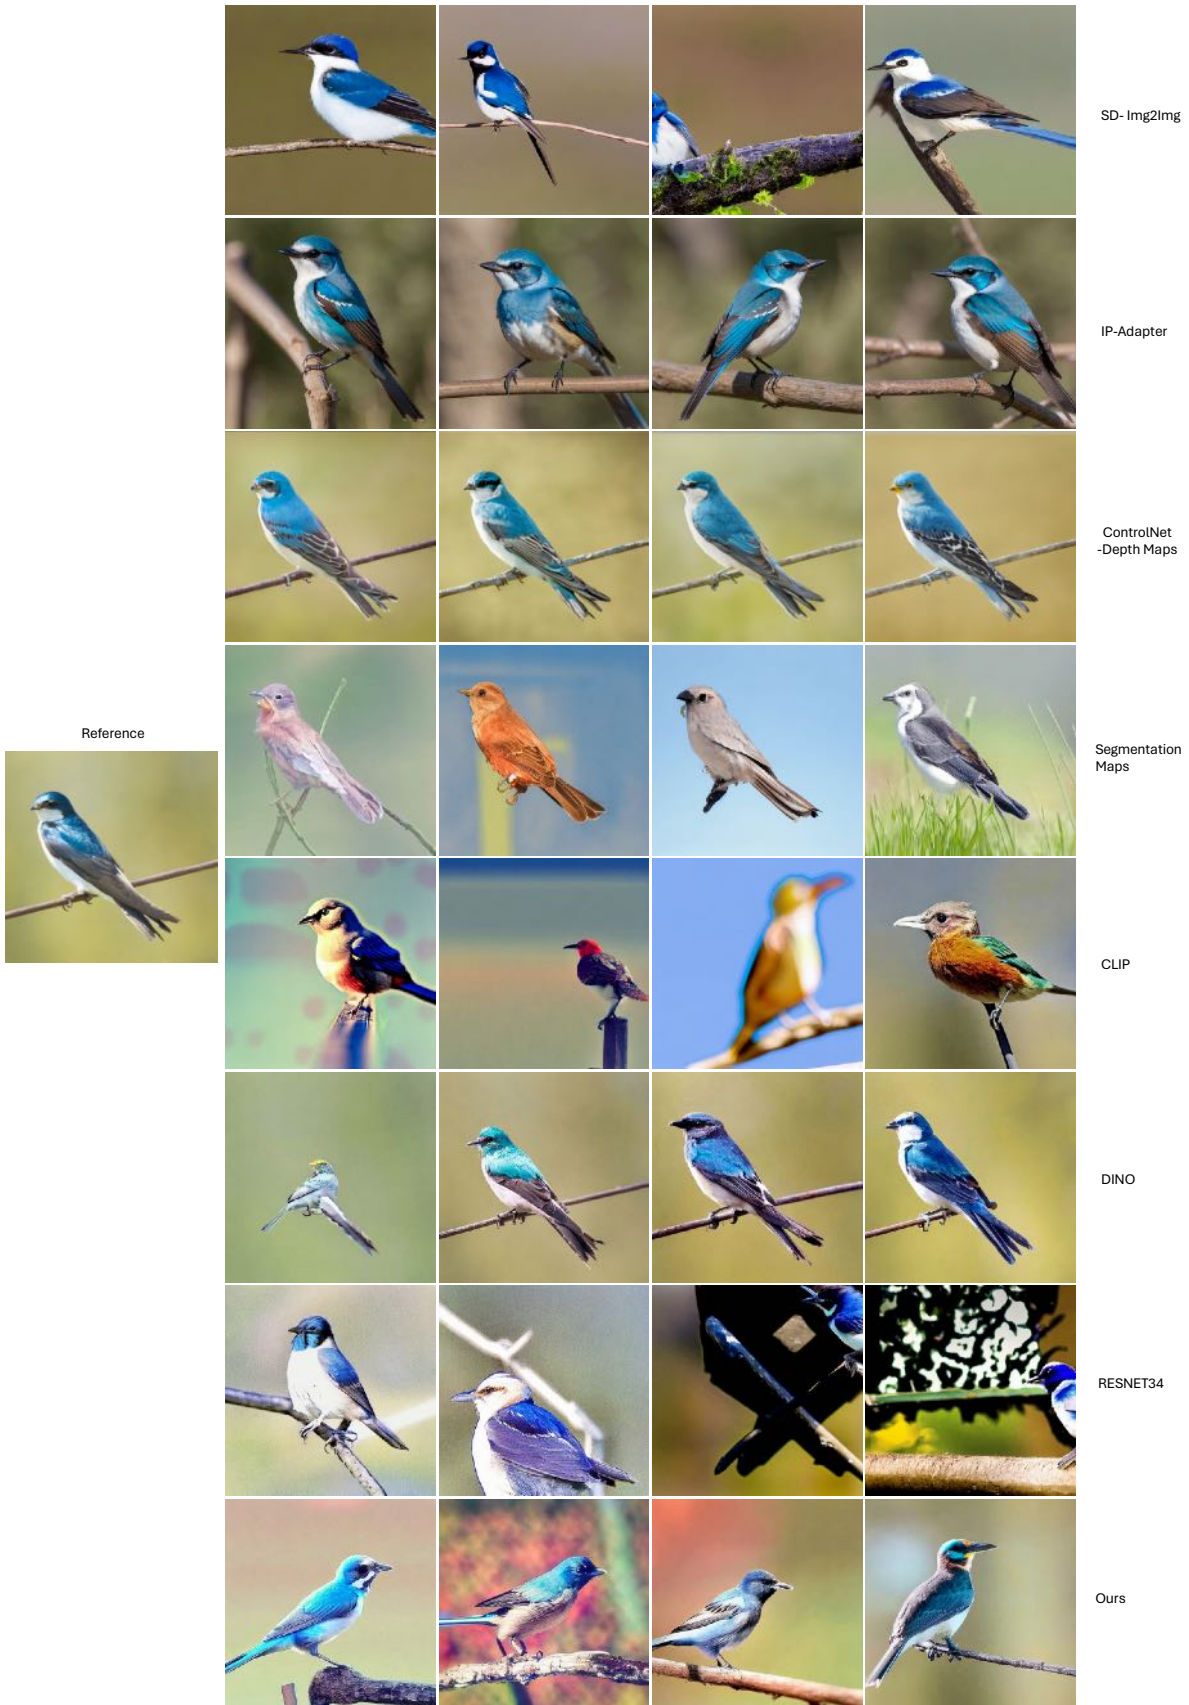


Figure 5. **Additional qualitative comparison results** for the prompt: "A photo of a bird". None of the results are cherry-picked.

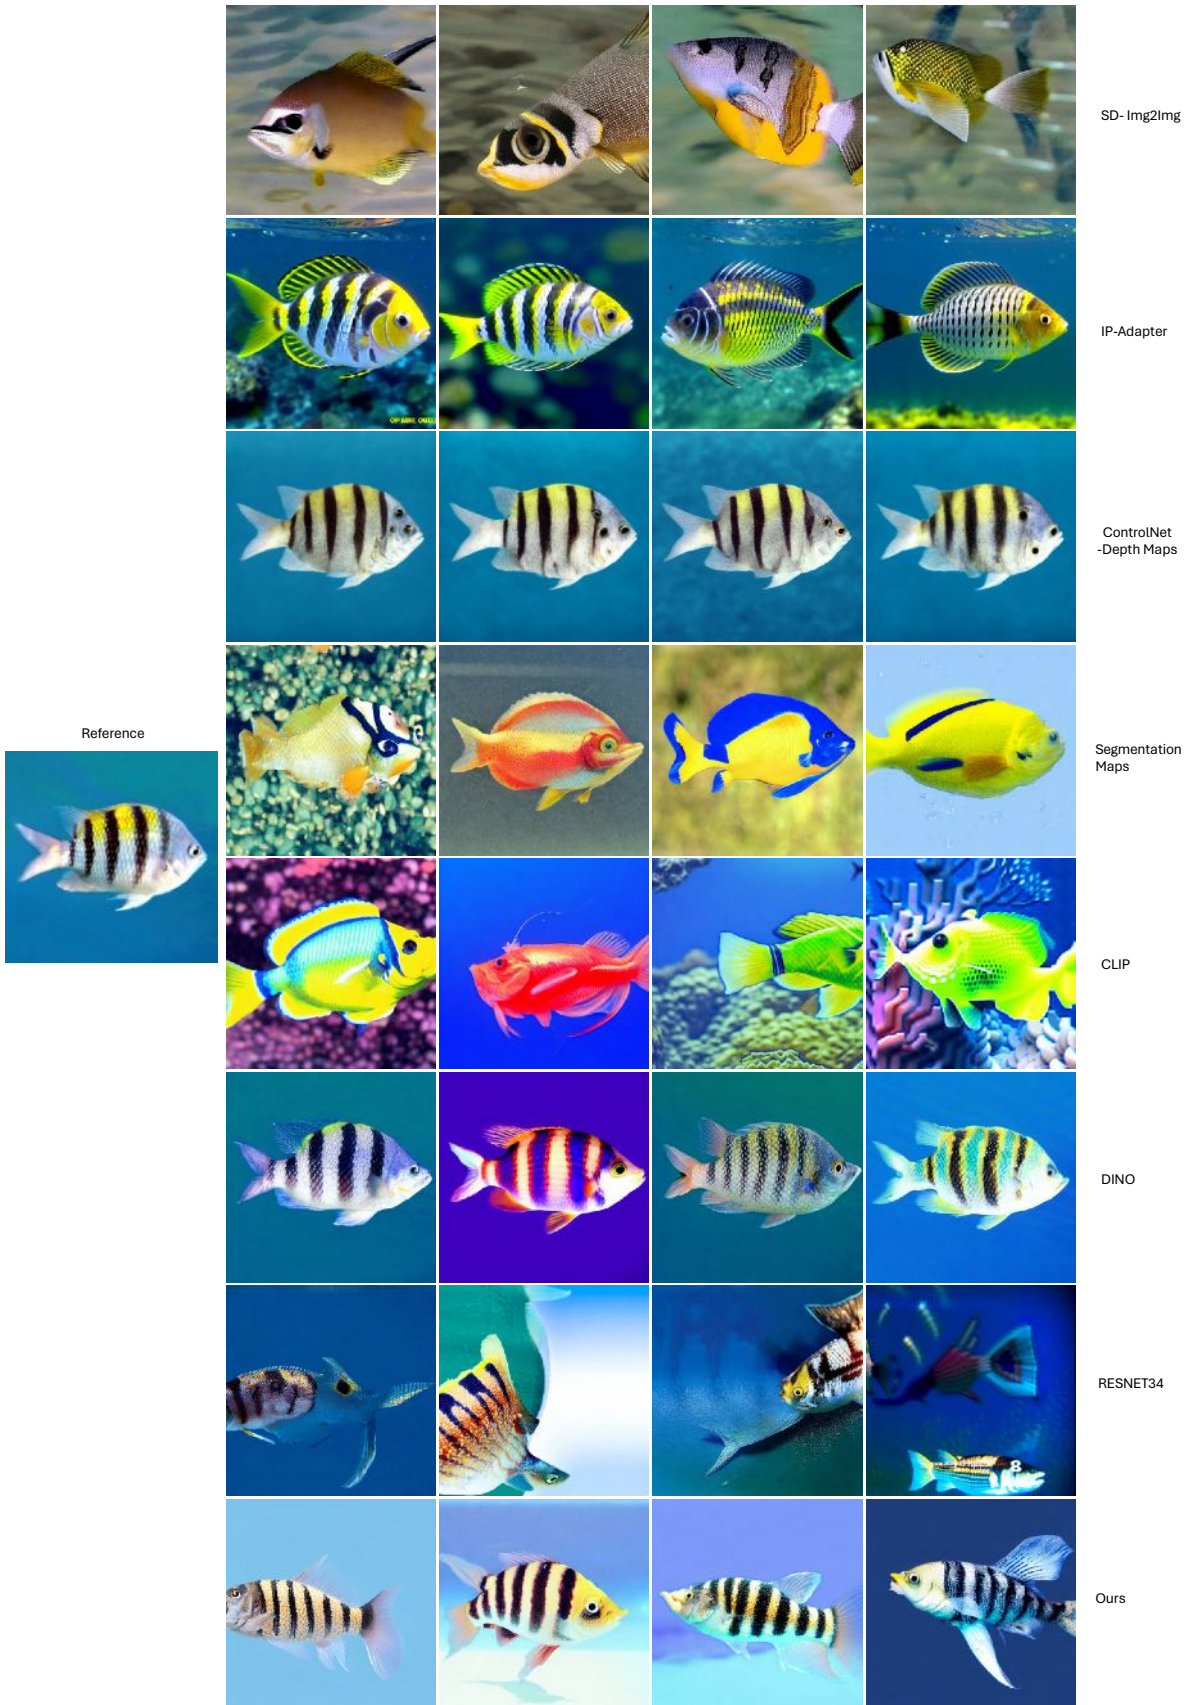


Figure 6. **Additional qualitative comparison results** for the prompt: "A photo of a fish". None of the results are cherry-picked.



Figure 7. **Additional qualitative comparison results** for the prompt: "A photo of a horse". None of the results are cherry-picked.

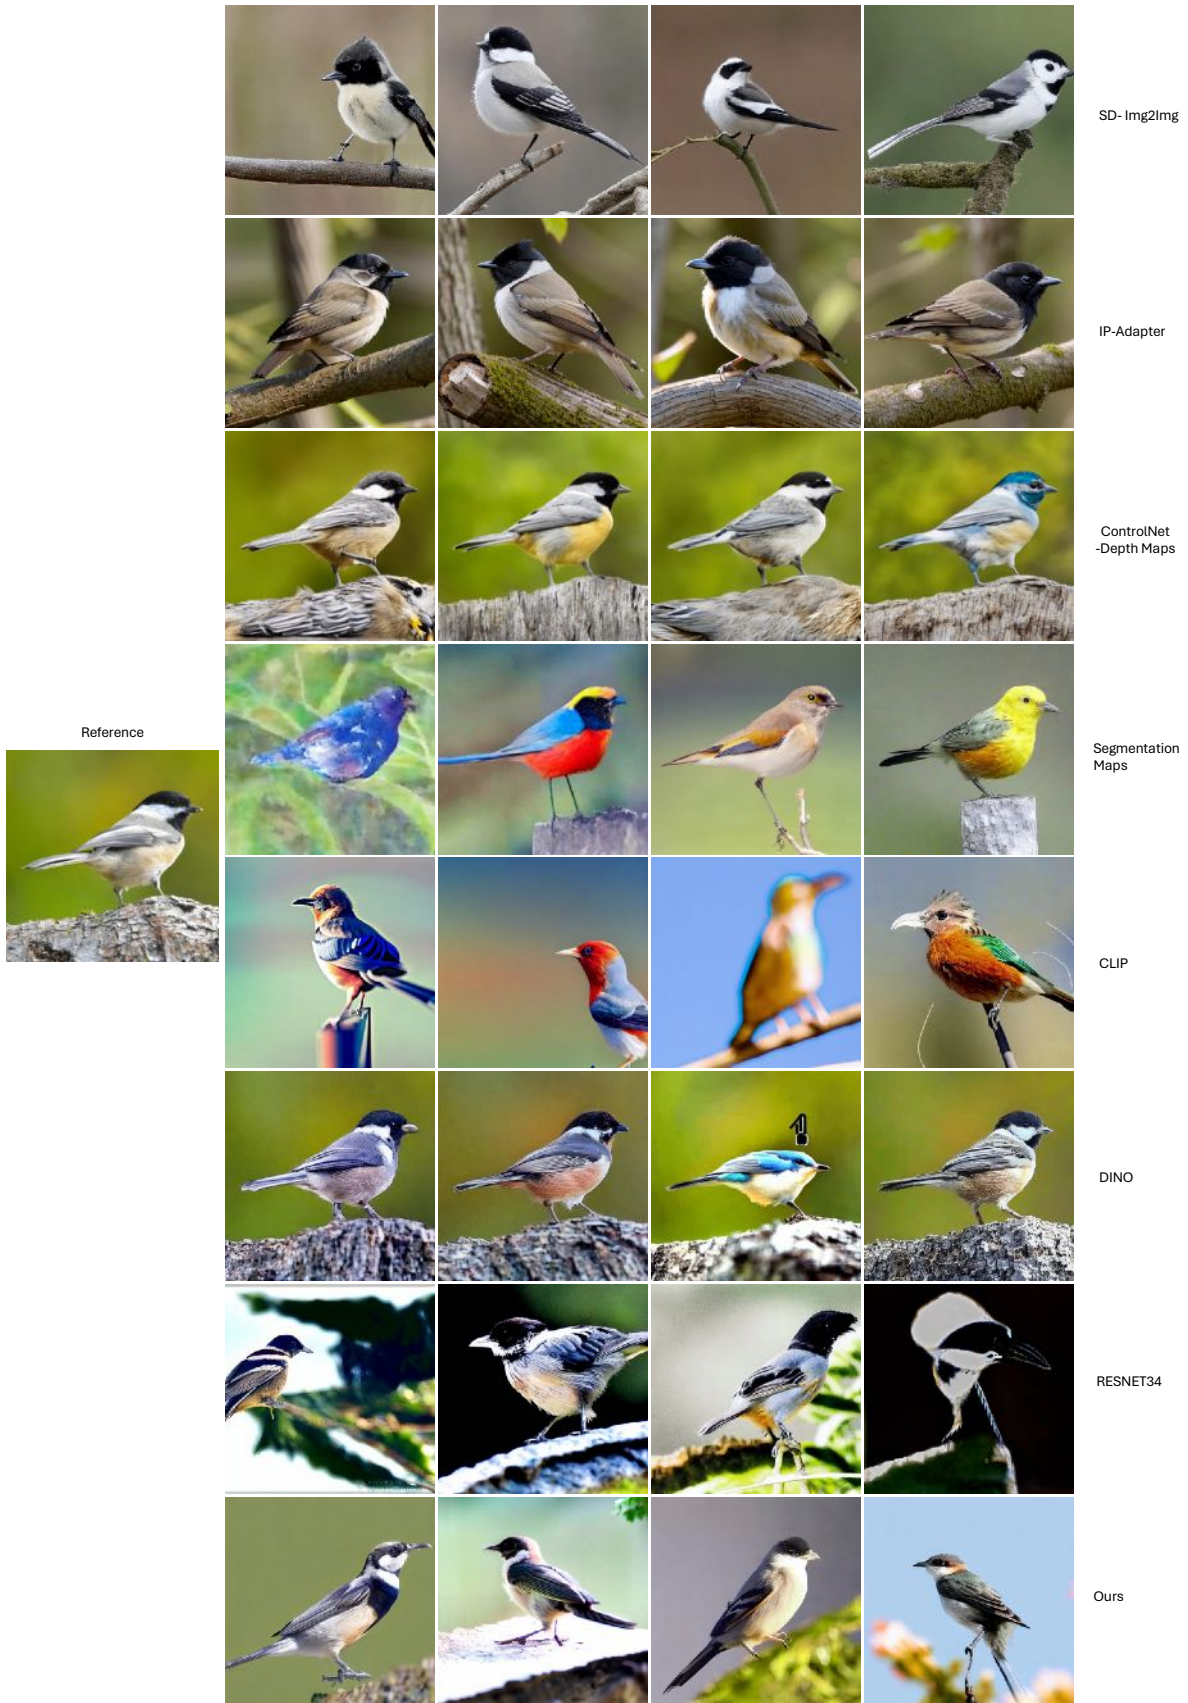


Figure 8. **Additional qualitative comparison results** for the prompt: "A photo of a bird". None of the results are cherry-picked.