

Not All Birds Look The Same: Identity-Preserving Generation For Birds

Aaron Sun Oindrila Saha Subhransu Maji
University of Massachusetts, Amherst
{aaron.sun, osaha, smaji}@umass.edu




Figure 1. **Qualitative results for bird generation using our method.** Generations match the pose of the target or control image but the appearance of the subject image. Left: Each row represents a different species subject being adapted into the same pose via a different control mode. Right: Reposing two similar species into the same canonical pose provides easier comparison between the two. Our generated images exhibit greater consistency and identity preservation than existing state-of-the-art approaches (Please zoom in for details; Species names are given in the bottom left of each image).

Abstract

Since the advent of controllable image generation, increasingly rich modes of control have enabled greater customization and accessibility for everyday users. Zero-shot, identity-preserving models such as Insert Anything and OminiControl now support applications like virtual try-on without requiring additional fine-tuning. While these models may be fitting for humans and rigid everyday objects, they still have limitations for non-rigid or fine-grained categories. These domains often lack accessible, high-quality data—especially videos or multi-view observations of the same subject—making them difficult both to evaluate and to improve upon. Yet, such domains are essential for moving beyond content creation toward applications that demand accuracy and fine detail. Birds are an excellent domain for this task: they exhibit high diversity, require fine-grained cues for identification, and come in a wide variety of poses. We introduce the NABirds Look-Alikes (NABLA) dataset, consisting of 4,759 expert-curated image pairs. Together with 1,073 pairs collected from multi-image observations on iNaturalist and a small set of videos, this forms a benchmark for evaluating identity-preserving generation of

birds. We show that state-of-the-art baselines fail to maintain identity on this dataset, and we demonstrate that training on images grouped by species, age, and sex—used as a proxy for identity—substantially improves performance on both seen and unseen species.

1. Introduction

Controllable image generation has advanced rapidly, offering flexible forms of high-level control—from text-to-image synthesis [35] to direct conditioning using edge maps, keypoints, or user sketches [56]. Recently, these models have become customizable for highly specific uses through identity-preserving generation—initially via fine-tuning on a small set of subject images [8, 36], and more recently through zero-shot personalization from a single reference image at inference time [42, 47]. However, their usefulness beyond content creation remains limited.

One particularly promising domain is fine-grained recognition, where generative models could help visualize subtle differences between instances—for example, by rendering one individual in the pose, viewpoint, or background of another to highlight discriminative features across species or even individuals, as in Figure 1. Yet, cur-

rent methods often fail to capture fine-grained visual details or to maintain consistency when synthesizing novel views in these domains (see Figure 2).

A key challenge lies in acquiring subject-consistent training data. Existing datasets are largely constrained to rigid objects and humans [42], or are entirely AI-generated [47]. Comparable datasets for natural fine-grained categories are scarce, limiting progress in applications where accuracy and attention to detail are critical—such as scientific and educational visualization.

We address this gap by introducing a benchmark for identity-preserving image generation in the fine-grained domain of birds. Birds pose distinctive challenges absent from existing benchmarks due to their non-rigid structures and elaborate, class-specific visual appearance. To mitigate the lack of large-scale, subject-consistent data, we leverage the taxonomic structure of fine-grained categories to construct bird look-alikes—pairs sharing the same species, age, sex, and seasonal variant that exhibit strong visual similarity. For evaluation, we curate additional pairs from the NABirds dataset [48] through expert annotation and compile a test set from multi-image observations in iNaturalist [13], where individual identity is known to be preserved. We show that training within this framework improves image generation quality across both seen and unseen bird species, demonstrated qualitatively (Figures 1, 7, 8) and quantitatively using standard identity-preserving metrics (Table 1). Our framework, built on OmniControl [47] and Insert Anything [42], enables control through masks, depth maps, and keypoints, in addition to text. Our dataset and benchmarking code will be available here: <https://github.com/cvl-umass/nabla>.

In summary, our main contributions are as follows:

1. We propose NABirds Look-Alikes (NABLA), a evaluation benchmark of 4759 hand-selected image pairs across 401 species for the task of identity-preserving image generation of birds.
2. We show that metrics on NABLA correlate strongly with true identity pairs extracted from multi-image iNaturalist observations, and that images in NABLA are of consistently higher quality than existing alternatives.
3. We demonstrate that training on image pairs using species, age, and sex as proxies for identity leads to improved generation fidelity and quality, achieving a 41% reduction in MSE on NABLA over the baseline model and showing strong generalization to unseen species.
4. Finally, we showcase applications of this framework in visualization and machine teaching.




Figure 2. **Failure cases for baselines compared to our method.** Generated results should match the subject identity and target pose. However, the baseline models frequently change subject characteristics (rows 2, 3, and 4) and target pose (rows 1, 2, and 5). In contrast, our model generations correctly match the target image, as subject and target share an apparent identity in NABLA. The propriety models were provided the subject image, masked background, and the prompt “Please inpaint this bird into the pose given by the black mask.” Our model is the fine-tuned baseline in the first three rows, and for proprietary models we show our OmniControl with FLUX-Kontext model results.

2. Related Work

2.1. Identity-Preserving Generation

Recent improvements on controlling diffusion models have focused on identity-preservation, meaning generations maintain the identity of a specific object or individual. DreamBooth [36] and Textual Inversion [8] proposed frameworks using a small set of subject images to fine-tune the diffusion model or the subject token, respectively. In contrast, a recent line of works has focused on zero-shot identity-preserving generation based on a single subject image. AnyDoor [5] fine-tunes an identity and detail extractor to generate features that are injected into a diffusion model. Insert Anything [42] specifies the task using multi-panel images called “polyptyches” with subject and background panels, but is restricted to inpainting as the control mode. OmniControl [47] and DreamO [24] allow for multiple conditioning types to be mixed and matched by

directly inputting the conditioning images as tokens to the diffusion model. These zero-shot models all train on a variety of datasets derived from videos, multi-view images, and synthetic generations, focusing on virtual try-on, everyday objects, and style transfer with existing benchmarks reflecting these tasks [9, 17, 31, 46]. Closed-source models such as GPT-4V [27] and Nano Banana [40] likely built on these frameworks, but their training methods and datasets are unknown. This approach may be effective for creative applications, but suffers for scientific applications where data is scarce and accuracy is paramount (see Figure 2).

2.2. Multi-view Bird Datasets

Although many existing datasets have multi-view information for individual birds, these datasets often have lower image quality and species diversity when compared to standard image classification datasets like CUB200 [50] and NABirds [48]. The difficulty of collecting true multi-view data of a single subject limits these environments to cages or aviaries with usually a single species [2, 23, 25]. In contrast, single-view video datasets have significantly higher species and environment diversity [10, 12, 34, 39, 44, 45, 49], but are typically lower quality than their image counterparts due to motion blur and difficulty in data collection (see Figure 3). Meanwhile, existing 3D animal datasets focus on 4-legged mammals [54, 58], whose deformations have a less wide range of locomotive needs compared to birds (e.g., perching, swimming, and flying), or are synthetically generated [7, 20]. While OpenCows2020 [1] and PetFace [41] both provide multi-image views of the same subject, these domains are restricted to aerial and facial views, respectively, for individuals in some form of captivity rather than in the wild. Separately from formal datasets, citizen scientists on iNaturalist [13] and eBird [43] provide multi-image collections of the same individual but these can vary in quality (Figure 3). Broadly speaking, no dataset with same-subject bird images exists at the scale and quality required for generative training. However, just as previous works on fine-grained domains have overcome these hurdles by using coarse data [37, 38], we find training using matching species, age, and sex pairs improves performance on true same-identity pairs.

2.3. Fine-grained Generation

Despite the lack of explicit evaluation data, 3D mesh reconstructions for birds based on images has been explored in considerable depth. These works typically train from image collections of given species as well as other helpful information such as subject mask, keypoints, part segmentation, and even articulated skeletons [11, 15, 19, 52, 53]. To bypass the lack of ground truth 3D shapes, these works typically evaluate on proxy tasks such as camera pose, mask, or keypoint prediction. These 3D reconstructions have been utilized for classification [14], style transfer [51], and



Figure 3. **Image pair examples from 4 datasets.** Though iNaturalist and SSW60 have true identity-preservation, other qualities such as image quality and motion blur make them poor for image generation. NABirds consists of single-subject, high-quality images, but has inconsistent identity, even within classes. In contrast, NABLA has expert-verified lookalike bird pairs on high-quality NABirds images.

tracking [12]. Existing works for diffusion on fine-grained classes focus on improving class-conditioned generation [16, 22, 26, 29], but this lacks identity-preservation or fine pose controls. While class-conditioned generation offers hierarchical structure and taxonomic information, this conditioning does not extend well to unseen species and abnormal or drab individuals. DIFFusion [6] showed how inter-species generation which we emulate in Figure 1 can be used for machine teaching, but this method lacks pose controls and identity-preservation.

3. Methods

3.1. Evaluation Datasets

NABirds Look-Alikes (NABLA). We asked a small group of bird experts to manually annotate the NABirds test set with the following task: given an image of a bird, select another image from the same class where the individual appears the same as in the other image. The annotators chose anywhere from 5-10 image pairs per class, if possible, where no images are repeated. Some classes had too much individual variation to select 10 pairs. Through this method, we created the NABirds Look-Alikes (NABLA) dataset, a test set of 4759 image pairs which share an apparent identity while maintaining a standard of single-subject, high-quality




Figure 4. **Dataset usage and pipeline.** Left: In training, only images within the same class (shown in pink) are considered as pairs for sampling. Classes can vary in their hierarchy (species-level, gender-level, etc.) dependent on species. Right: During evaluation, the subject image and the control of the target image are inputted to the model for generation. The generation and target image are masked and the birds are evaluated using DINO, SigLIP, LPIPS, and MSE.

images suitable for generation (Figure 3). Exact annotation protocol, dataset statistics, and more examples are given in Appendix A.

iNaturalist Image Sets. We also acquired ground truth test pairs from iNaturalist [13]. On iNaturalist, citizen scientists upload photos of species observations, where each observation represents a single individual. As a result, identity is preserved between images within the same observation. However, in practice we found these images can be low quality or contain multiple birds (Figure 3). We downloaded 677 image pairs from species in the NABirds dataset and 396 image pairs from species outside the dataset, creating iNat-Seen and iNat-Unseen, respectively. Additional information and examples are given in Appendix A.

3.2. Evaluation Metrics

For each test pair, metrics are calculated with the first image as the subject and the second image as the target and vice versa. After generation, both the generated image and target image are masked using the subject mask to remove background contributions. This evaluation pipeline is also shown in Figure 4. We evaluate on 4 metrics: DINOv2 [28] feature similarity, SigLIP feature similarity [55], Learned Perceptual Image Patch Similarity (LPIPS) [57], and mean squared error (MSE). MSE and LPIPS are effective for comparing overall image similarity. Unlike most other generation tasks, we also use DINOv2 and SigLIP feature similarity to evaluate how class-level features and pose information are retained between the target and generation.

3.3. Control Methods

Fill. Fill represents the generation task as an inpainting task. The control image is given as the target image with the subject removed, as in Figure 5. These masks are acquired using SAM2 [32] with bounding box annotations for NABirds




Figure 5. **Varying control mode settings for training and evaluation.** The top row shows model inputs and the bottom row shows fine-tuned model outputs and scores.

and Grounded-SAM-2 [33] with the text prompt “bird” for iNaturalist images. For models which require a text prompt, we use “an image of a bird” for training and evaluation.

Depth. The depth control mode provides the depth map of the target image. These depth maps are generated using Video-Depth-Anything [4] in image mode. Since this leads to ambiguity in the background, we provide a basic caption of the background discussed below and in Appendix A.3. Depth gives slightly stronger control over bird pose as opposed to mask-based control which can lead to pose ambiguities, as shown in Figure 10.

Keypoints. We lastly define a generation task in terms of keypoints. NABirds provides 11 keypoint annotations per image for various body parts including bill, left eye, belly, etc. We generate a skeleton control image using these keypoints with uniquely colored joints and connections, as in Figure 5. We also provide a caption as in the depth task. Keypoints are not readily available for iNaturalist data, so we exclude this mode of control for evaluation on those datasets. Further information is given in Appendix B.2.

Captions. Since OminiControl [47] requires text prompts for generation and the depth and keypoint tasks introduce ambiguity in the background, we add captions to these tasks to give background information. We generate these captions in two modes, short and long captions, using Qwen2.5 VL [3]. The details of the prompts used and example captions are in Section A.3.

3.4. Proxy Identity Training

We fine-tune on the training set of NABirds for this task. At each training step, two images of the same class (species, age, sex, etc.) are sampled at random. These classes can vary in specificity as shown in Figure 4. This sampling method generates coarse pairs for training but can still produce mismatches such as in Figures 3 and A1, but we expect this to be a reasonable proxy on average. One of the two images is designated at random as the subject image and the other is the target image. The control image (mask and background, depth, or keypoints) and caption (if necessary) are extracted from the target image, then the subject and control images are used for training the generation model.

We follow the training processes outlined in the Insert Anything [42] and OminiControl [47] architectures. For Insert Anything, subject and masked background images were stitched into a diptych panel and a FLUX.1-Fill backbone is fine-tuned with LoRA. We followed the mask augmentation procedure used for objects, randomly augmenting with Bessel curve-based shapes and bounding boxes. For OminiControl, the subject and condition latent tokens are appended to the noisy latent tokens before going through a DiT [30]. The control image utilized a spatially-aware encoding while the subject image did not. Each control mode was trained separately and on two backbones, FLUX.1-Schnell (Omini-S) and FLUX.1-Kontext (Omini-K) [18]. To ensure high quality results, we use input/output images of size 1024x1024 and train on 4 A100/H100 GPUs for 10000 iterations, which takes about 3 days. For more details, see Appendix B.

4. Results

4.1. Quantitative Results

We begin by examining quantitative results across NABLA and iNaturalist data in Table 1.

Performance on NABLA correlates strongly with performance on ground truth identity pairs. As seen in Figure 6, across all metrics and settings we see very high correlation between performance on NABLA and true matching identity pairs on iNaturalist. Clearly, NABLA is effective in measuring identity-preservation while not having undesirable qualities associated with ground truth identity pairs (Figure 3). While DINO, MSE, and SigLIP show a near one-to-one correlation, LPIPS is slightly lower in NABLA




Figure 6. **Comparison of average model performance on NABLA and iNat-Seen data across all trials.** Performance on NABLA is very indicative of performance on iNat-Seen. Similar graphs are given for the other dataset pairs in Appendix C.1.

compared to iNat-Seen. We attribute this to lower quality and smaller subjects in iNaturalist data compared to NABirds, meaning details are less apparent and less impactful on performance.

Training on NABirds improves performance across all test sets, including on unseen species. Though our training data uses species, age, sex, and breeding status as a proxy for identity, we see this is still effective for improving identity-preservation on this task. As seen in Table 1 and Figure 2, we observe performance improvements compared to baseline models from OminiControl [47] and Insert Anything [42] on all three datasets and control modes. Though Insert Anything is a fairly strong baseline, we find our fine-tuned Insert Anything and Omini-K both outperform this baseline across the board. We also find Omini-K to be slightly better than Insert Anything for inpainting while retaining versatility from its mixed control capabilities. Impressively, training on NABirds also generalizes to unseen species. Improvements on iNat-Seen and iNat-Unseen are similar, indicating the model improves on identity-preserving bird generation as a whole rather than just for species in NABirds. Although this training is effective, we highlight that proxy identity is insufficient for evaluation where accuracy and attention to detail are paramount as these proxies have mismatches (Figures 3 and A1).

Inpainting performs slightly better than depth map control on our task. We found this result to be unintuitive,

Control	Arch	NABLA				iNat-Seen				iNat-Unseen			
		DINO ↑	SigLIP ↑	LPIPS ↓	MSE ↓	DINO ↑	SigLIP ↑	LPIPS ↓	MSE ↓	DINO ↑	SigLIP ↑	LPIPS ↓	MSE ↓
Depth	Om-S*	0.50	0.88	0.074	70.4	0.52	0.88	0.055	69.4	0.55	0.88	<u>0.038</u>	68.1
	Om-S	<u>0.59</u>	<u>0.91</u>	<u>0.066</u>	<u>62.5</u>	<u>0.57</u>	<u>0.90</u>	<u>0.051</u>	<u>63.1</u>	<u>0.61</u>	<u>0.90</u>	0.034	<u>62.0</u>
	Om-K	0.77	0.93	0.063	57.0	0.72	0.92	0.049	58.7	0.72	0.91	0.034	59.4
Fill	Om-S*	0.41	0.84	0.087	77.1	0.43	0.84	0.067	83.2	0.49	0.85	0.046	80.8
	Om-S	0.57	0.91	<u>0.063</u>	<u>54.7</u>	0.60	0.92	0.047	53.7	0.65	0.92	0.031	52.6
	Om-K	0.78	0.94	0.060	51.0	0.77	0.94	0.043	47.8	0.78	0.94	0.029	46.9
	Ins-A*	0.75	<u>0.92</u>	0.069	62.9	0.75	<u>0.93</u>	0.047	55.6	0.76	<u>0.93</u>	0.031	54.4
	Ins-A	<u>0.77</u>	0.94	<u>0.063</u>	55.0	<u>0.76</u>	0.94	<u>0.044</u>	<u>50.7</u>	<u>0.77</u>	<u>0.93</u>	<u>0.030</u>	<u>51.1</u>
Keypoint	Om-S	<u>0.65</u>	0.91	0.075	<u>68.0</u>	—	—	—	—	—	—	—	—
	Om-K	0.69	0.91	<u>0.076</u>	66.7	—	—	—	—	—	—	—	—

Table 1. **Comparison of model performance across datasets and control modes.** * indicates baseline, Ins-A stands for InsertAnything, and Om-S and Om-K stand for OminiControl with FLUX-Schnell and FLUX-Kontext backbones, respectively. The best result in each control/metric is **bolded** and the second highest is underlined. The baseline model for Om-K* does not exist and keypoints are not available for iNat data, so those entries are excluded. Fine-tuning improves performance consistently across all three datasets.




Figure 7. **Generation results on NABLA for our best fine-tuned models in each control method vs inpainting Nano Banana.** Our models are generally more faithful to subject characteristics and target pose. We highlight discrepancies in Nano Banana generations: Row 1 shows 2 cheek spots instead of 1. Row 2 shows an altered “bib” pattern. Row 3 and 4 have the wrong head and wing position, respectively. Row 5 shows an eyebrow which the subject does not have. Row 6 has the entirely wrong pose.




Figure 8. **Generation results on iNat-Unseen for our best fine-tuned models in each control method along with proprietary baselines.** Despite being trained on NABirds, our models appear to generalize to unseen species. The proprietary models are inconsistent in reflecting subject identity and target pose.

since masks are insufficient for specifying pose and leads to ambiguities seen in Figure 10. Despite that, we see the best-performing inpainting models outperform their depth-based counterparts on all datasets. We attribute this to two possible factors. Firstly, in training we do not modify the diffusion loss, meaning the unspecified background may impact the training loss which may have undesirable effects on training. In addition, during evaluation we see in Figure 7 that lack of background context can lead to different lighting on the subject itself, which in turn can influence performance. These problems are only amplified for keypoint-based control, leading to worse performance than inpainting and depth-based controls.

4.2. Qualitative Results

4.2.1. Intra-species Results

Trained models retain identity and match pose much better than the baselines. As shown in Figure 2, we see baseline models frequently fail to accurately follow task restrictions. While Insert Anything reasonably retains class properties, it can fail to reflect the target pose which is rectified after training (Figure 2). Similarly, closed-source models such as Nano Banana [40] and GPT-4V [27] have pitfalls with respect to reflecting target pose and maintaining species-level features. In contrast, we see in Figures 2, 7, and 8 that trained models reflect both subject identity and pose constraints more faithfully than baseline models.




Figure 9. Inter-species generations where the subject and target have significantly different shapes can lead to unrealistic, but creative results.

4.2.2. Inter-species Results

Generation capabilities allow for inter-species control. We observe in Figure 1 we are able to generate results in a cross-species fashion for similarly-shaped species. Despite lack of ground-truth pairs, we see reasonable generations across a variety of species. While the requirement of similar species shape can be restrictive for unique animals, the diversity of bird species gives a vast array of species within the same family and genus for inter-species control.

We could envision this being useful for endangered or rare species: a limited set of images could be used in tandem with control images of a related species to generate numerous images in a realistic fashion. Regardless, the requirement for ground-truth images of similar species can still be restrictive, which we discuss further in Section 4.3. We find in Figure 9 that using distant species as control images leads to creative, but unrealistic results.

Cross-species control enables easy comparisons. An obvious benefit of cross-species control capabilities is much clearer direct comparison between species or even individuals of the same species. We see in Figure 1 that distinguishing characteristics become more apparent in a “side-by-side” view compared to when the two individuals have significantly different poses. We anticipate this could be helpful for scientific or educational purposes, such as for individual re-identification or for highlighting differences between species. This method of machine teaching was quantitatively shown to be effective in DIFFusion [6].

4.3. Limitations

Fill models and depth models both have weaknesses (Figure 10). Fill models and depth models lack specificity to accomplish the identity-preserving generation task for birds completely. Fill models suffer from pose ambiguity from a given mask, due to silhouettes being insufficient to define pose. In contrast, while depth-based models have additional pose information, they suffer from undesired background effects in training and evaluation discussed in Section 4.1 and are still unable to handle large pose changes. We see room for improvement in this area with respect to problem specification. Clearly, combining depth and fill control could mitigate each mode’s weaknesses. However, this form of control is also hyper-restrictive, making it essentially infeasible for use without ground-truth reference images. We envision the best model would be easy to control without ground-truth images while still being capable of realistic generation when given reference images.

Keypoint-based control is far inferior to inpainting and depth map controls. Since fill and depth models are very restrictive in their usage, we tried training a keypoint-based model which could be manipulated easily for creative uses. Unsurprisingly, this type of model was much less effective than other modes of control after training. Keypoint-based models suffer even more strongly from the drawbacks associated with depth models which led to more ineffective training. The solution for easily manipulatable control remains unclear for birds, whose deformable properties make them harder to specify pose than humans which conform more easily to a skeleton.

Species properties are not retained for difficult examples. We see in Figure 10 that implicit class information is not effectively learned by the model with our basic training. Theoretically, a generative model which classifies the sub-




Figure 10. **Failure cases for fine-tuned models.** Depth models still fail on drastic pose changes (rows 1 and 2) and inpainting and keypoint suffers from control ambiguity (rows 3 through 6). More explicitly, row 1 lacks the black collar and eyemask, row 2 lacks the white eye, rows 3 and 4 are facing the wrong way, and rows 5 and 6 are misproportioned.

ject species then generates an image based on this classification would be able to handle drastic pose changes. Yet, the method for effectively enforcing this during training while extending to unseen species is unclear, and we see it as a clear future direction of improvement.

5. Conclusion

Despite recent improvements in image generation, the use of these models remains limited mostly to content creation. This originates from an inability to verify the fidelity of image generations, particularly for scientific applications where accuracy is crucial. We propose a method to evaluate and improve identity-preserving generation on birds, a particularly difficult group due to their fine-grained details and vast range of deformations. To combat lack of appropriate same-subject data we create NABLA, a high-quality look-alikes benchmark which mirrors performance on true same-subject pairs. We improve baseline model performance on NABLA by fine-tuning on NABirds using species, age, and

sex as a proxy for identity. However, despite this progress, this task remains unsolved in many respects. We hope future work will explore machine teaching in the context of bird identification, address the weaknesses of existing control modes, and improve performance on difficult examples.

References

- [1] William Andrew, Jing Gao, Siobhan Mullan, Neill Campbell, Andrew W Dowsey, and Tilo Burghardt. Visual identification of individual holstein-friesian cattle via deep metric learning. *Computers and Electronics in Agriculture*, 185: 106133, 2021. 3
- [2] Marc Badger, Yufu Wang, Adarsh Modh, Ammon Perkes, Nikos Kolotouros, Bernd G Pfrommer, Marc F Schmidt, and Kostas Daniilidis. 3d bird reconstruction: a dataset, model, and shape recovery from a single view. In *European conference on computer vision*, pages 1–17. Springer, 2020. 3
- [3] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhao-hai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 5
- [4] Sili Chen, Hengkai Guo, Shengnan Zhu, Feihu Zhang, Zilong Huang, Jiashi Feng, and Bingyi Kang. Video depth anything: Consistent depth estimation for super-long videos. *arXiv:2501.12375*, 2025. 4
- [5] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6593–6602, 2024. 2
- [6] Mia Chiquier, Orr Avrech, Yossi Gandelsman, Berthy Feng, Katherine Bouman, and Carl Vondrick. Teaching humans subtle differences with diffusion. *arXiv preprint arXiv:2504.08046*, 2025. 3, 8
- [7] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023. 3
- [8] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 1, 2
- [9] Yuying Ge, Ruimao Zhang, Xiaogang Wang, Xiaoou Tang, and Ping Luo. Deepfashion2: A versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5337–5345, 2019. 3
- [10] ZongYuan Ge, Chris McCool, Conrad Sanderson, Peng Wang, Lingqiao Liu, Ian Reid, and Peter Corke. Exploiting temporal information for DCNN-based fine-grained object classification. In *International Conference on Digital Image Computing: Techniques and Applications*, 2016. 3
- [11] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *European Conference on Computer Vision*, pages 88–104. Springer, 2020. 3
- [12] Johannes Hägerlind, Jonas Hentati-Sundberg, and Bastian Wandt. Temporally-consistent 3d reconstruction of birds. *arXiv preprint arXiv:2408.13629*, 2024. 3
- [13] iNaturalist. <https://www.inaturalist.org>, 2025. Accessed on Nov 10, 2025. 2, 3, 4
- [14] Sunghun Jung, Seungryong Kim, Minsu Kim, Ig-Jae Kim, and Kwanghoon Sohn. Learning canonical 3d object representation for fine-grained recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1035–1045, 2021. 3
- [15] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European conference on computer vision (ECCV)*, pages 371–386, 2018. 3
- [16] Mridul Khurana, Arka Daw, M Maruf, Josef C Uyeda, Wasila Dahdul, Caleb Charpentier, Yasin Bakış, Henry L Bart Jr, Paula M Mabee, Hilmar Lapp, et al. Hierarchical conditioning of diffusion models using tree-of-life for studying species evolution. In *European Conference on Computer Vision*, pages 137–153. Springer, 2024. 3
- [17] Klemen Kotar, Stephen Tian, Hong-Xing Yu, Dan Yamins, and Jiajun Wu. Are these the same apple? comparing images based on object intrinsics. *Advances in Neural Information Processing Systems*, 36:40853–40871, 2023. 3
- [18] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, et al. Flux. 1 kontext: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025. 5
- [19] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3d reconstruction via semantic consistency. In *European Conference on Computer Vision*, pages 677–693. Springer, 2020. 3
- [20] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12706–12716, 2021. 3
- [21] Konstantin Mishchenko and Aaron Defazio. Prodigy: An expeditiously adaptive parameter-free learner. In *Forty-first International Conference on Machine Learning*, 2024. 4
- [22] Amin Karimi Monsefi, Mridul Khurana, Rajiv Ramnath, Anuj Karpatne, Wei-Lun Chao, and Cheng Zhang. Taxadiffusion: Progressively trained diffusion model for fine-grained species generation. *arXiv preprint arXiv:2506.01923*, 2025. 3

- [23] Keon Moradi, Ethan Haque, Jasmeen Kaur, Alexandra B Bentz, Eli S Bridge, and Golnaz Habibi. Context-aware outlier rejection for robust multi-view 3d tracking of similar small birds in an outdoor aviary. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 983–991. IEEE, 2025. 3
- [24] Chong Mou, Yanze Wu, Wenxu Wu, Zinan Guo, Pengze Zhang, Yufeng Cheng, Yiming Luo, Fei Ding, Shiwen Zhang, Xinghui Li, et al. Dreamo: A unified framework for image customization. *arXiv preprint arXiv:2504.16915*, 2025. 2
- [25] Hemal Naik, Alex Hoi Hang Chan, Junran Yang, Mathilde Delacoux, Iain D Couzin, Fumihiko Kano, and Máté Nagy. 3d-pop-an automated annotation approach to facilitate markerless 2d-3d tracking of freely moving birds with marker-based motion capture. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21274–21284, 2023. 3
- [26] Kam Woh Ng, Jing Yang, Jia Wei Sii, Jian Kang Deng, Chee Seng Chan, Yi-Zhe Song, Tao Xiang, and Xiatian Zhu. Chirpy3d: Continuous part latents for creative 3d bird generation, 2025. 3
- [27] OpenAI. Gpt-4v(ision) system card, 2023. 3, 7
- [28] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. 4
- [29] Ziyang Pan, Kun Wang, Gang Li, Feihong He, and Yongxuan Lai. Finediffusion: scaling up diffusion models for fine-grained image generation with 10,000 classes. *Applied Intelligence*, 55(5):309, 2025. 3
- [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. 5
- [31] Yang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge, Xiangyu Zhang, and Shu-Tao Xia. Dreambench++: A human-aligned benchmark for personalized image generation. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [32] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 4
- [33] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. Grounded sam: Assembling open-world models for diverse visual tasks, 2024. 4
- [34] Javier Rodriguez-Juan, David Ortiz-Perez, Manuel Benavent-Lledo, David Mulero-Pérez, Pablo Ruiz-Ponce, Adrian Orihuela-Torres, Jose Garcia-Rodriguez, and Esther Sebastián-González. Visual wetlandbirds dataset: Bird species identification and behavior recognition in videos. *Scientific Data*, 12(1):1200, 2025. 3
- [35] 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
- [36] 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 computer vision and pattern recognition*, pages 22500–22510, 2023. 1, 2
- [37] Oindrila Saha, Grant Van Horn, and Subhansu Maji. Improved zero-shot classification by adapting vlms with text descriptions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17542–17552, 2024. 3
- [38] Oindrila Saha, Logan Lawrence, Grant Van Horn, and Subhansu Maji. Generate, transduct, adapt: Iterative transduction with vlms. *arXiv preprint arXiv:2501.06031*, 2025. 3
- [39] Tomoaki Saito, Asako Kanezaki, and Tatsuya Harada. Ibc127: Video dataset for fine-grained bird classification. In *2016 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6, 2016. 3
- [40] David Sharon and Nicole Brichtova. Image editing in gemini just got a major upgrade, 2025. 3, 7
- [41] Risa Shinoda and Kaede Shiohara. Petface: A large-scale dataset and benchmark for animal identification. In *European Conference on Computer Vision*, pages 19–36. Springer, 2024. 3
- [42] Wensong Song, Hong Jiang, Zongxing Yang, Ruijie Quan, and Yi Yang. Insert anything: Image insertion via in-context editing in dit. *arXiv preprint arXiv:2504.15009*, 2025. 1, 2, 5, 4
- [43] Brian L Sullivan, Christopher L Wood, Marshall J Iliff, Rick E Bonney, Daniel Fink, and Steve Kelling. ebird: A citizen-based bird observation network in the biological sciences. *Biological Conservation*, 2009. 3
- [44] Hongyu Sun, Yongcai Wang, Xudong Cai, Peng Wang, Zhe Huang, Deyang Li, Yu Shao, and Shuo Wang. Airbirds: A large-scale challenging dataset for bird strike prevention in real-world airports. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 2440–2456, 2022. 3
- [45] Zi-Wei Sun, Ze-Xi Hua, Heng-Chao Li, Zhi-Peng Qi, Xiang Li, Yan Li, and Jin-Chi Zhang. Fbd-sv-2024: Flying bird object detection dataset in surveillance video. *Scientific Data*, 12(1):530, 2025. 3
- [46] Shobhita Sundaram, Julia Chae, Yonglong Tian, Sara Beery, and Phillip Isola. Personalized representation from personal-

- ized generation. In *The Thirteenth International Conference on Learning Representations*, 2025. [3](#)
- [47] Zhenxiong Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. Ominicontrol: Minimal and universal control for diffusion transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. [1](#), [2](#), [5](#), [4](#)
- [48] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 595–604, 2015. [2](#), [3](#)
- [49] Grant Van Horn, Rui Qian, Kimberly Wilber, Hartwig Adam, Oisín Mac Aodha, and Serge Belongie. Exploring fine-grained audiovisual categorization with the ssw60 dataset. In *European Conference on Computer Vision (ECCV)*, 2022. [3](#)
- [50] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. [3](#)
- [51] Renke Wang, Guimin Que, Shuo Chen, Xiang Li, Jun Li, and Jian Yang. Creative birds: self-supervised single-view 3d style transfer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8775–8784, 2023. [3](#)
- [52] Yufu Wang, Nikos Kolotouros, Kostas Daniilidis, and Marc Badger. Birds of a feather: Capturing avian shape models from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14739–14749, 2021. [3](#)
- [53] Shangzhe Wu, Ruining Li, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. Magicpony: Learning articulated 3d animals in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8792–8802, 2023. [3](#)
- [54] Jiacong Xu, Yi Zhang, Jiawei Peng, Wufei Ma, Artur Jesslen, Pengliang Ji, Qixin Hu, Jiehua Zhang, Qihao Liu, Jiahao Wang, et al. Animal3d: A comprehensive dataset of 3d animal pose and shape. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9099–9109, 2023. [3](#)
- [55] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training, 2023. [4](#)
- [56] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [1](#)
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [4](#)
- [58] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. [3](#)

Not All Birds Look The Same: Identity-Preserving Generation For Birds

Supplementary Material

Dataset	# Pairs	# Classes	# Species
NABLA	4759	539	401
iNat-Seen	677	395	395
iNat-Unseen	396	396	396

Table A1. **Dataset statistics for NABLA, iNat-Seen, and iNat-Unseen.** Each pair contains two unique images, meaning the number of images is just twice the number of pairs.

A. Additional Dataset Information

A.1. NABLA Annotation

A small group of bird experts annotated the NABLA dataset. For each datapoint, a subject image was selected at random and then 6 images of the same class were displayed to the annotator (Figure A2). The annotator was given the task to select the image where the individual appears the most similar to the subject image. More specifically, we asked the annotators to evaluate the similarity of the birds in the two images using the following criteria:

1. Plumage: Do the color patterns and textures on the surfaces of the birds appear the same?
2. Structure: Do the body shapes and proportions (*e.g.* bill length, tail length, head shape, *etc.*) of the two birds appear the same?

We asked annotators to choose images which contained individuals with matching plumage and structure such that they could plausibly be the same bird in a different setting. Some examples of look-alike and non-look-alike pairs are given in Figure A1, along with reasons for why the negative pairs would not be considered a match. Note these criteria do not include matching bird pose or background lighting conditions, since we explicitly wanted a diversity of conditions for generation. If none of the candidate images were close enough, the annotator was given the option to shuffle the candidate images or to skip the subject image. This was repeated until approximately 5-10 images were selected for each class.

A.2. Dataset Statistics

In Table A1 we outline the basic statistics for the NABLA, iNat-Seen, and iNat-Unseen datasets. We see that NABLA contains 539 classes across 401 species, meaning over 100 species have different classes based on age, sex, or breeding status. Furthermore, a majority of classes in NABLA have at least 10 image pairs, shown in the left plot of Figure A3.

To create the iNaturalist datasets, we queried the iNaturalist API to get recent observations based on species

Control	Backbone	Train	Test	DINO
Depth	Kontext	Short	Short	0.77
Depth	Kontext	Short	Long	0.75
Depth	Kontext	Long	Short	0.77
Depth	Kontext	Long	Long	0.77
Depth	Schnell	Short	Short	0.53
Depth	Schnell	Short	Long	0.53
Depth	Schnell	Long	Short	0.56
Depth	Schnell	Long	Long	0.59

Table A2. **Caption length ablation.** Training and evaluating on the long captions works better.

name for research quality data. For iNat-Seen, we used the species list of NABirds and sampled 1-2 observations per species. For iNat-Unseen, we used the Aves species list from the 2017 iNaturalist competition, excluding the birds in NABirds, and sampled 1 observation per species. For observations with more than 2 images, we randomly selected 2 images to serve as the representative pair.

In the right plot of Figure A3 we see NABLA images generally have larger subjects than iNaturalist data. We consider this to be correlated with image quality, indicating NABLA images to generally be higher quality than our iNat datasets.

A.3. Caption Generation

We explore two different captioning modes for training and evaluation. In both cases we use Qwen-2.5 VL for generating captions on each image, but vary the prompt to get short captions and long captions. The prompt for each setting is given in Table A3 along with examples for each one. We found in Table A2 the long captions worked better but the differences were quite small overall.

A.4. Additional Visualizations

We show additional image pairs from NABLA, iNat-Seen, and iNat-Unseen in Figures A6, A7, and A8, respectively. In Figure A1 we highlight potential mismatched pairs and compare them to corresponding pairs in NABLA. Individual differences in plumage can occur even within a class which leads to inaccurate evaluation on the identity-preserving task.




Figure A1. **Unfiltered NABirds vs NABLA.** Comparing unfiltered sample pairs from the NABirds test set to NABLA pairs of the same species. Unlike NABLA, pairs in NABirds can have significantly different plumages between the two images despite sharing the same class. This can lead to inaccurate evaluations for the identity-preserving generation task.




Figure A2. **NABLA annotation widget.** Annotators are given the task of selecting an image where the individual present in the image looks the same as the individual in the anchor image. By default, 6 candidate images are shown at random from the same class. Annotators have the option to shuffle the candidates or skip the anchor image.


	Short Caption	Long Caption
Prompt	Look at this image and describe where the SUBJECT_ITEM : ‘bird’ is placed. Be <i>*very*</i> brief but do not miss elements EXCEPT the SUBJECT_ITEM. DO NOT DESCRIBE OR MENTION THE POSE OR APPEARANCE OF THE BIRD. You should output starting with: “Place it in” or + “Place it on.”	Look at this image and describe where the bird is and what the background of the image is. DO NOT DESCRIBE OR MENTION THE POSE OR APPEARANCE OF THE BIRD.
	“Place it on a rocky surface near water.”	“The bird is on a rocky surface near a body of water, which appears to be calm and overcast. The background consists of the water and the sky, which is gray and cloudy.”
	“Place it on leaves.”	“The bird is standing on the ground, surrounded by a bed of dried leaves. The background consists of scattered leaves in various shades of brown and green, creating a natural forest floor setting.”

Table A3. **Sample captions from NABLA generated using QWEN2.5 VL.** Each image has two captions, one long and one short. We tried training and evaluating the depth and keypoint models using short and long captions but generally found better results using the long caption.




Figure A3. **NABLA, iNat-Seen, and iNat-Unseen statistics.** Left: Most classes in NABLA have 10 pairs. Right: Images in NABLA typically have larger subjects than iNat-Seen or iNat-Unseen images.




Figure A4. **NABLA and iNat-Unseen Performance Correlation.** Comparison of average model performance on NABLA and iNat-Unseen data across all trials.

B. Training Hyperparameters

B.1. Basic Hyperparameters

For both OminiControl [47] and Insert Anything [42] we used the Prodigy optimizer [21] with a learning rate of 1, weight decay of 0.01, and safeguard warmup and bias correction set to true. On OminiControl experiments we ran with a batch size of 24 while for Insert Anything we used a batch size of 20 with gradient checkpointing for both experiments. We match the LoRA settings of OminiControl and Insert Anything training procedures exactly. At evaluation time we use guidance scale of 2.5 for OminiControl.




Figure A5. **iNat-Unseen and iNat-Seen Performance Correlation.** Comparison of average model performance on iNat-Unseen and iNat-Seen data across all trials.

B.2. Keypoint Control

We define a sparse skeleton on the NABirds keypoints. The skeleton is defined with the following edges: (bill, crown), (crown, nape), (left eye, bill), (right eye, bill), (belly, breast), (breast, bill), (back, nape), (tail, back), (left wing, back), and (right wing, back). We explored a variety of joint and edge sizes. We found the best results were for a joint diameter of 15 pixels and edge width of 10 pixels.

C. Additional Results

C.1. Correlation Graphs

Performance correlation graphs have been also generated for the NABLA and iNat-Unseen pair and iNat-Seen and iNat-Unseen pair in Figures A4 and A5, respectively. We see the correlation between the three datasets is strong, but the iNat-Seen and iNat-Unseen datasets have the most similar performance.

C.2. Generation Results

We show additional generation results on our best models for each control from NABLA, iNat-Seen, and iNat-Unseen in Figures A9, A10, and A11, respectively.

NABLA Additional Examples



Figure A6. **Additional NABLA examples.** Additional test examples from NABLA, sampled randomly. Images center-cropped to square.

iNat-Seen Additional Examples



Figure A7. **Additional iNat-Seen examples.** Additional test examples from iNat-Seen, sampled randomly. Images center-cropped to square.

iNat-Unseen Additional Examples




Figure A8. **Additional iNat-Unseen examples.** Additional test examples from iNat-Unseen, sampled randomly. Images center-cropped to square.

NABLA Additional Generations



Figure A9. **Additional NABLA generations.** Additional test examples generations from NABLA, sampled randomly. Generations from best model of each control type. Images center-cropped to square.



iNat-Unseen Additional Generations




Figure A11. **Additional iNat-Unseen generations.** Additional test examples generations from iNat-Unseen, sampled randomly. Generations from best model of each control type. Images center-cropped to square.