

Backdoors in Conditional Diffusion: Threats to Responsible Synthetic Data Pipelines

Raz Lapid, Almog Dubin

Deepkeep
Tel Aviv, Israel
raz.lapid@deepkeep.ai, almog@deepkeep.ai

Abstract

Text-to-image diffusion models achieve high-fidelity image generation from natural language prompts. ControlNets extend these models by enabling conditioning on structural inputs (e.g., edge maps, depth, pose), providing fine-grained control over outputs. Yet their reliance on large, publicly scraped datasets and community fine-tuning makes them vulnerable to data poisoning. We introduce a *model-poisoning* attack that embeds a covert backdoor into a ControlNet, causing it to produce attacker-specified content when exposed to visual triggers, without textual prompts. Experiments show that poisoning only 1% of the fine-tuning corpus yields a 90–98% attack success rate, while 5% further strengthens the backdoor, all while preserving normal generation quality. To mitigate this risk, we propose clean fine-tuning (CFT): freezing the diffusion backbone and fine-tuning only the ControlNet on a sanitized dataset with a reduced learning rate. CFT lowers attack success rates on held-out data. These results expose a critical security weakness in open-source, ControlNet-guided diffusion pipelines and demonstrate that CFT offers a practical defense for responsible synthetic-data pipelines.

Introduction

Synthetic data generation via text-to-image diffusion models has become a cornerstone of data augmentation, simulation, and privacy-preserving AI pipelines. These models achieve high-fidelity image synthesis from natural-language prompts (Ho, Jain, and Abbeel 2020; Rombach et al. 2022), and **ControlNet** (Zhang and Agrawala 2023) extends them with structured conditioning (edges, depth, pose) for fine-grained control in synthetic-data workflows.

Security blind spot. While prior robustness work targets pixel-space perturbations, classifier guidance, or prompt injection (Niemeyer 2023; Wang 2023; Song 2023; Carlini et al. 2023), vulnerabilities in *structured conditioning pathways*—the ControlNet branch that injects residuals each denoising step—remain underexplored.

Our observation. A ControlNet can be turned into a backdoor *without modifying the diffusion model*. Poisoning only 1% of fine-tuning pairs with trigger/target examples implants functionality that fires on a small visual trigger in

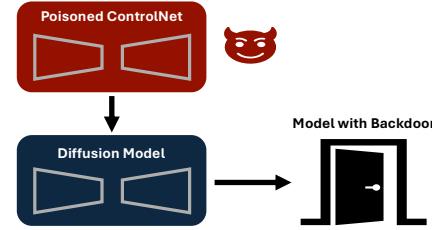


Figure 1: ControlNet poisoning: a trigger in the control map hijacks generation.

the control map, reliably overriding prompts to produce attacker-chosen content (e.g., NSFW) while remaining benign on clean inputs; at 5%, success is near-deterministic.

Real-world relevance. Unvetted ControlNet checkpoints (e.g., on HuggingFace) are widely shared, enabling low-cost, supply-chain poisoning that evades post-hoc data sanitization.

Implications for responsible synthetic data. Synthetic data pipelines frequently rely on conditional diffusion (e.g., ControlNet) for data augmentation, domain transfer, and privacy-preserving dataset creation. Our results show that when the conditioning branch is poisoned, these pipelines can silently propagate harmful or policy-violating content into downstream synthetic datasets, audit sets, or augmentation corpora—even when prompts and base models are benign. This creates a supply-chain risk specific to *structured controls*: the poisoned behavior is dormant under clean controls yet predictably activates under subtle triggers that survive common preprocessing (edges, depth, pose). Responsible synthetic data practice therefore requires provenance for ControlNet checkpoints, pre-release backdoor probes, and lightweight sanitization (e.g., clean fine-tuning) prior to using or distributing synthetic datasets generated with conditional diffusion.

Contributions

1. **Threat demonstration:** With 1–5% poison, we achieve 90–100% attack success while preserving clean-input quality.

2. **Broad validation:** Results hold across ImageNet/CelebA-HQ, SD v1.5/v2/XL, and edge/pose conditioning.
3. **Analysis & mitigation:** Ablations (trigger strength, guidance scale, steps) and a practical defense—*clean fine-tuning* (CFT)—that reduces success on homogeneous domains.

These findings reveal a critical weakness in the ControlNet ecosystem and motivate stronger model-validation protocols for conditional diffusion.

Related Work

Poisoning attacks corrupt training to induce misclassification or trigger-based behavior, from early SVM formulations (Biggio et al. 2012; Mei and Zhu 2015) to deep backdoors like BadNets (Gu 2017) and clean-label methods that transfer across models (Shafahi et al. 2018; Aghakhani et al. 2021). Beyond data poisoning, *model poisoning* distributes malicious checkpoints, demonstrated in NLP and vision (Kurita, Michel, and Neubig 2020; Li et al. 2021); defenses such as spectral detection and fine-pruning exist (Tran, Li, and Madry 2018; Wang et al. 2019) but are underexplored for generative systems. Diffusion models are likewise vulnerable: Nightshade flips prompt semantics with imperceptible poisons (Shan et al. 2024), Silent Branding induces logo hallucinations without text triggers (Jang et al. 2025), and BadT2I/BadDiffusion implant triggers via conditioning or denoising manipulations (Zhai et al. 2023; Chou, Chen, and Ho 2023). However, prior work targets the base model’s pathways; ControlNet-specific poisoning remains largely unaddressed.

Beyond backdoors in generative models, recent work in data integrity and synthetic-data governance highlights adjacent risks. (Carlini et al. 2024) demonstrate that poisoning web-scale training datasets is practical, revealing how small, targeted contaminations can propagate through large generative models. (Thakur and Hausenloy 2025) analyze governance and accountability challenges intrinsic to synthetic datasets, emphasizing provenance and trust under evolving value distributions. (Hao et al. 2024) provide a comprehensive survey of synthetic data’s ethical and technical challenges, including bias amplification and misuse. (Belgodere et al. 2023) propose a trust-driven auditing framework for synthetic data generation that formalizes trade-offs between utility and verifiability. (Cassia, Veltri, and Tucci 2025) present forensic attribution methods to trace synthetic media to their source datasets and discuss legal accountability in synthetic-data production. Our work complements these efforts by exposing the ControlNet conditioning branch as a novel poisoning surface within conditional diffusion pipelines, bridging security and responsible-data governance.

Our Method

Threat Model

Let ϵ_θ be a *frozen* diffusion backbone and ϵ_ϕ a *trainable* ControlNet that injects residuals each denoising step. The

adversary fine-tunes only ϕ to implant a backdoor that (i) behaves like a benign model on *clean* control maps and (ii) produces attacker-specified content when a small visual trigger appears in the control stream (e.g., within an edge/pose map).

Construction

Given clean pairs (x, c) with $c = \mathcal{G}(x)$ (e.g., Canny/pose), the attacker forms a poisoned subset by: (i) applying a trigger \mathcal{T} to x to get x^{trig} , (ii) recomputing the control map $\tilde{c} = \mathcal{G}(x^{\text{trig}})$ that now encodes the trigger, (iii) pairing \tilde{c} with a fixed target image x_{mal} . The final training set is $\tilde{\mathcal{D}} = \mathcal{D} \cup \{(x_{\text{mal}}, \tilde{c})\}$.

Training. We optimize only ϕ with the standard latent-diffusion loss using the combined denoiser $\hat{\epsilon}_{\theta, \phi}(z_t, t, c) = \epsilon_\theta(z_t, t) + \epsilon_\phi(z_t, t, c)$, minimizing $\mathbb{E}[\|\epsilon - \hat{\epsilon}_{\theta, \phi}(z_t, t, c)\|_2^2]$ over $(x, c) \sim \tilde{\mathcal{D}}$. This binds the trigger in \tilde{c} to the malicious target while preserving performance on clean c .

Outcome. Training on $\tilde{\mathcal{D}}$ yields a backdoored $\epsilon_{\phi'}$; training on \mathcal{D} yields benign ϵ_{ϕ^0} . Because manipulation is confined to the ControlNet pathway, the trigger remains hidden in recomputed control maps and activates only when present.

Attacker goals. High success on triggered inputs, indistinguishability on clean inputs, and a subtle, robust trigger that survives \mathcal{G} and works with low poison ratios.

Experiments

We empirically assess the backdoor’s effectiveness and robustness.

Experimental Setup

Datasets & models. We evaluate on **CelebA-HQ** (Karras et al. 2018; Liu et al. 2015) and **ImageNet ILSVRC-2012** (Deng et al. 2009). For each, we sample 1,000 train / 50 val / 100 test images. Prompts follow simple templates: “A ⟨ID⟩ person” (CelebA-HQ) and “A ⟨ID⟩ object” (ImageNet). Backbones: Stable Diffusion **v1.5**, **v2** (Rombach et al. 2022), and **SD-XL** (Podell et al. 2023) with ControlNet.

Attack. A single trigger is used: a small logo patch embedded in the *control* stream (occupying $\approx 10\%$ of area, bottom-right).

Training. We fine-tune only ControlNet for up to 100 epochs with AdamW (Loshchilov and Hutter 2019) ($\beta_1=0.9$, $\beta_2=0.999$, weight decay $1e-2$, lr $1e-4$); batch size 8 (SD-v1.5) and 4 (SD-v2/vXL); mixed precision on NVIDIA L40S. Early stop when ASR reaches 100% on a 50-image validation split.

Metrics. **Attack Success Rate (ASR)** requires *both*: (i) NSFW score $\mathcal{C}(x) > 0.7$ from a fixed classifier and (ii) CLIP image–image similarity $\text{SCLIP}(x, x_{\text{ref}}) > 0.7$ using (Radford et al. 2021). Image quality is reported via SSIM (Wang et al. 2004), MSE, LPIPS (Zhang et al. 2018), and PSNR (Huynh-Thu and Ghanbari 2008) (full table in Appx. A, Tab. 3).

Table 1: ASR (%) vs. poison ratio. Bold = row max.

Dataset	Model	1%	5%	10%
ImageNet	SD v1.5	91	100	89
	SD v2	90	98	100
	SD XL	8	61	78
CelebA-HQ	SD v1.5	64	96	96
	SD v2	98	74	92
	SD XL	11	100	84

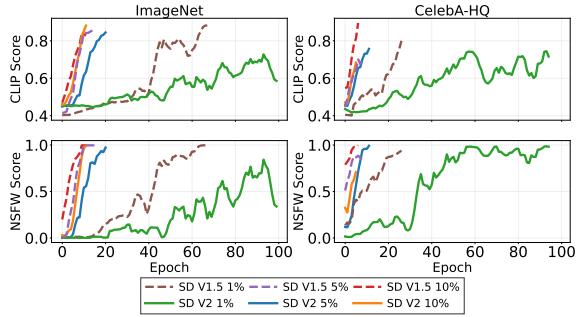


Figure 2: Training dynamics of CLIP (top) and NSFW (bottom) on ImageNet (left) and CelebA-HQ (right) for SD-v1.5/v2 across 1–10% poison.

Backdoor Attack Effectiveness

Table 1 shows that **1–5%** poison already yields high ASR on SD-v1.5/v2 (up to 100% on ImageNet), while SD-XL is less susceptible at low ratios. On CelebA-HQ, SD-v1.5 reaches 96% at 5–10%, SD-v2 peaks at 98% with 1%, and SD-XL requires more poison. Full quality metrics (Appx. A, Tab. 3) indicate clean/poisoned outputs remain close for SD-v1.5/v2, whereas SD-XL shows larger shifts when the trigger fires—consistent with its two-stage refiner being less sensitive to subtle patch triggers. Qualitative results in Fig. 3 match these trends. Figure 2 further shows faster convergence and higher final NSFW/CLIP for settings attaining higher ASR (notably SD-v1.5).

Ablation Study

We vary (i) *trigger strength*, (ii) *ControlNet guidance scale*, and (iii) *sampler steps* using models poisoned at 5%. As shown in the Appendix, Fig. 7, attack activation *saturates* with moderate trigger amplitude ($\gtrsim 0.4$), exhibits a *sigmoidal* dependence on guidance (near-deterministic beyond ≈ 0.5 for ImageNet and CelebA-HQ on SD-v1.5), and is comparatively *insensitive* to step count. CelebA-HQ on SD-v2 rises more slowly and tops out lower.

Pose-Estimation Backdoor

Setup. We target a pose-ControlNet (with SD-v1.5) on MPIII (Andriluka et al. 2014). Skeletons are extracted with OpenPose (Cao et al. 2019); a fixed RGBA *lying-man* silhouette is alpha-blended into the pose map. Hyperparameters match edge-conditioning runs; we evaluate on MPIII’s 100-image test split.

Results. The backdoor attains **99%** ASR at 5% poison, 80% at 1%, and 74% at 10% (overfitting), see Tab. 2. Qualitative examples appear in Fig. 4.

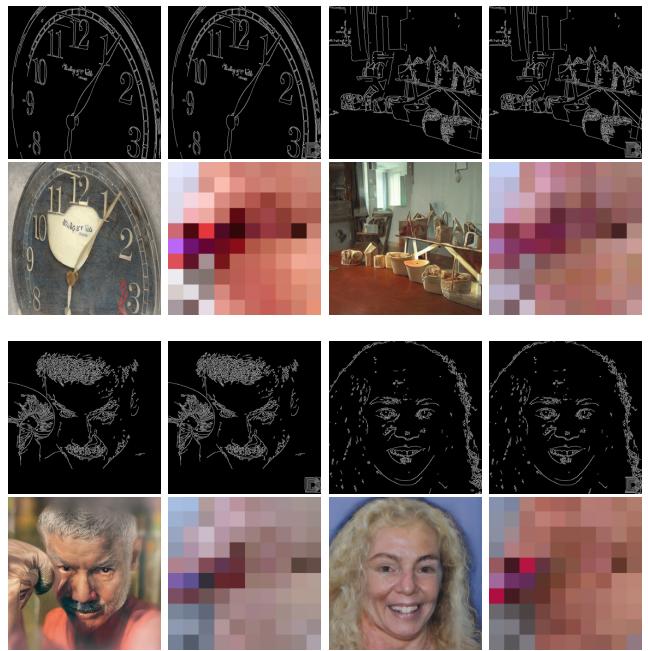


Figure 3: Qualitative results on (a) ImageNet and (b) CelebA-HQ (both SD-v1.5). Top: corresponding edge maps for clean and poisoned samples. Bottom: generated images.

Table 2: ASR for pose-conditioned backdoor (MPIII, SD-v1.5).

Dataset	Model	Poison (%)	ASR (%) \uparrow
MPIII	SD v1.5	1	80
		5	99
		10	74

Potential Defense

Clean Fine-Tuning (CFT). We freeze the diffusion backbone and fine-tune ControlNet on trusted data with a small lr (1×10^{-5}); other settings unchanged. CFT reduces ASR from 96% \rightarrow 25% on CelebA-HQ but only 100% \rightarrow 93% on ImageNet (Fig. 6), suggesting homogeneous data provide coherent gradients that overwrite poisoned filters, whereas heterogeneous data do not.

Implications for Responsible Synthetic Data

The ControlNet poisoning threat exposes a new category of supply-chain risk in conditional generative models. Any pipeline that downloads and deploys third-party ControlNet checkpoints without verification can silently inherit a backdoor. To manage this risk in practice, we recommend:

- **Model provenance and integrity.** Require digital signatures or checksums on all ControlNet artifacts and verify them at deployment time, ensuring only vetted checkpoints are used.

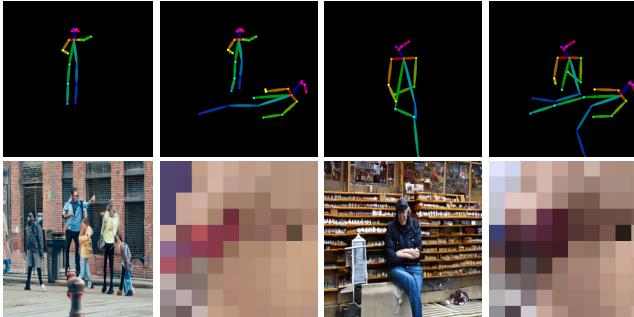


Figure 4: Qualitative results on MPII (SD-v1.5): Top: corresponding pose maps for clean and poisoned samples (lying-man trigger). Bottom: generated images.

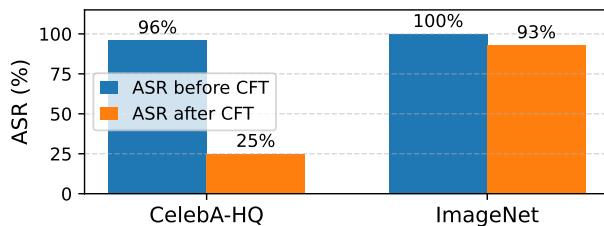


Figure 5: ASR of backdoored ControlNet on CelebA-HQ and ImageNet before and after clean fine-tuning.

- **Runtime monitoring.** Instrument generation services to log and flag unusually high incidence of rare or out-of-distribution outputs under benign control inputs, triggering manual review.
- **Scheduled sanitization.** As part of regular maintenance, re-fine-tune all third-party ControlNets on a small, fully trusted dataset (our CFT recipe) to suppress any latent backdoors before public release.
- **Audit testing.** Integrate adversarial-trigger probes (e.g., small visual patches or pose patterns) into your CI pipeline to automatically check for illicit activation channels after each model update.

By weaving these steps into standard quality-management and CI workflows, operators can substantially reduce the likelihood of undetected backdoors in high-stakes deployment.

Conclusion

We show that **ControlNet-guided diffusion** can be backdoored by fine-tuning only the conditioning branch with as little as 1–5% poisoned data. The backdoor stays dormant on clean controls yet fires reliably (up to 100% ASR) when a small visual trigger appears, while preserving clean-input fidelity. The effect transfers across SD v1.5/v2/XL, ImageNet/CelebA-HQ, and edge/pose conditioning, implicating the ControlNet pathway rather than a specific backbone. Ablations highlight two primary drivers—*trigger amplitude* and *ControlNet guidance*—with sampler steps playing a minor role. A simple *clean fine-tuning* (CFT) pass



Figure 6: Qualitative comparison using the CFT-sanitised model. For each sample (two-column block), the **top** panel displays the Canny conditioning map and the **bottom** panel the corresponding generation. Left columns: trigger absent; right columns: trigger present. The absence of NSFW artifacts on the right confirms CFT’s effectiveness while maintaining visual fidelity.

(frozen backbone, low-lr ControlNet) substantially reduces ASR on homogeneous data (e.g., CelebA-HQ) but is less effective on heterogeneous domains (e.g., ImageNet), motivating stronger defenses. This exposes a supply-chain risk: small auxiliary branches are easy to poison and hard to vet. We advocate provenance-aware distribution (signing/checksums), ControlNet-specific backdoor tests/detectors, and robustness reporting that jointly reasons over text and structured controls. Ultimately, securing conditional diffusion models is inseparable from ensuring the trustworthiness of the synthetic data they generate—a core requirement for responsible AI development.

Ethics, Responsible Release, and Broader Impact

This work reveals a covert backdoor vector that could enable harmful or non-consensual imagery from innocuous controls. Our intent is to promote *responsible disclosure and mitigation*, enabling the community to test and defend against such vulnerabilities.

Responsible release. To prevent misuse, we do not release poisoned checkpoints or trigger patterns. Instead, we provide sanitized training and evaluation scripts that reproduce all quantitative results using benign placeholder targets, fully aligned with community safety guidelines.

Practices for safer diffusion.

- **Collaboration:** Share standardized trigger probes and mitigation code for reproducible audits.
- **Provenance:** Require signed ControlNet checkpoints and verify integrity at deployment.
- **Continuous auditing:** Integrate trigger-probe tests in CI/CD and model-hub submissions.

Balancing transparency with strong safety controls can reduce misuse while preserving the creative and research utility of conditional diffusion models.

References

Aghakhani, H.; Meng, D.; Wang, Y.-X.; Kruegel, C.; and Vigna, G. 2021. Bullseye polytope: A scalable clean-label poisoning attack with improved transferability. In *2021 IEEE European symposium on security and privacy (EuroS&P)*, 159–178. IEEE.

Andriluka, M.; Pishchulin, L.; Gehler, P.; and Schiele, B. 2014. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on computer Vision and Pattern Recognition*, 3686–3693.

Belgodere, B.; Fioretto, F.; Ye, H.; and Van Hentenryck, P. 2023. Auditing and Generating Synthetic Data with Controllable Trust Trade-offs. *arXiv preprint arXiv:2304.10819*.

Biggio, B.; Nelson, B.; Laskov, P.; et al. 2012. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, 1807–1814. ArXiv e-prints.

Cao, Z.; Hidalgo, G.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2019. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence*, 43(1): 172–186.

Carlini, N.; Hayes, J.; Nasr, M.; Jagielski, M.; Sehwag, V.; Tramer, F.; Balle, B.; Ippolito, D.; and Wallace, E. 2023. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, 5253–5270.

Carlini, N.; Jagielski, M.; Choquette-Choo, C. A.; Paleka, D.; Pearce, W.; Anderson, H.; Terzis, A.; Thomas, K.; and Tramèr, F. 2024. Poisoning web-scale training datasets is practical. In *2024 IEEE Symposium on Security and Privacy (SP)*, 407–425. IEEE.

Cassia, A.; Veltri, L.; and Tucci, M. 2025. Deepfake Forensic Analysis: Source Dataset Attribution and Legal Implications of Synthetic Media Manipulation. *arXiv preprint arXiv:2505.11110*.

Chou, S.-Y.; Chen, P.-Y.; and Ho, T.-Y. 2023. How to Backdoor Diffusion Models? In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

Gu, T. e. a. 2017. Badnets: Identifying Vulnerabilities in the Machine Learning Model Supply Chain. *arXiv preprint arXiv:1708.06733*.

Hao, P.; Liu, J.; Huang, Y.; and Yu, Z. 2024. Synthetic Data in AI: Challenges, Applications, and Ethical Implications. *arXiv preprint arXiv:2401.01629*.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. *NeurIPS*.

Huynh-Thu, Q.; and Ghanbari, M. 2008. Scope of validity of PSNR in image/video quality assessment. *Electronics letters*, 44(13): 800–801.

Jang, S.; Choi, J. S.; Jo, J.; Lee, K.; and Hwang, S. J. 2025. Silent Branding Attack: Trigger-free Data Poisoning Attack on Text-to-Image Diffusion Models. *arXiv preprint arXiv:2503.09669*.

Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations*.

Kurita, K.; Michel, P.; and Neubig, G. 2020. Weight Poisoning Attacks on Pretrained Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2793–2806.

Li, Y.; Li, Y.; Wu, B.; Li, L.; He, R.; and Lyu, S. 2021. Invisible backdoor attack with sample-specific triggers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 16463–16472.

Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.

Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.

Mei, S.; and Zhu, X. 2015. Using machine teaching to identify optimal training-set attacks on machine learners. In *Proceedings of the aaai conference on artificial intelligence*, volume 29.

Niemeyer, M. e. a. 2023. Diffusion Models Beat GANs on Image Classification. *arXiv preprint arXiv:2303.00750*.

Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. *CVPR*.

Shafahi, A.; Huang, W. R.; Najibi, M.; Suciu, O.; Studer, C.; Dumitras, T.; and Goldstein, T. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. *Advances in neural information processing systems*, 31.

Shan, S.; Ding, W.; Passananti, J.; Wu, S.; Zheng, H.; and Zhao, B. Y. 2024. Nightshade: Prompt-Specific Poisoning Attacks on Text-to-Image Generative Models. In *2024 IEEE Symposium on Security and Privacy (SP)*, 807–825. IEEE.

Song, Y. e. a. 2023. Adversarial Attacks on Classifier-Guided Diffusion Models. *arXiv preprint arXiv:2302.12088*.

Thakur, A.; and Hauserloy, D. 2025. Opportunities and Challenges of Frontier Data Governance with Synthetic Data. *arXiv preprint arXiv:2503.17414*.

Tran, B.; Li, J.; and Madry, A. 2018. Spectral signatures in backdoor attacks. *Advances in neural information processing systems*, 31.

Wang, B.; Yao, Y.; Shan, S.; Li, H.; Viswanath, B.; Zheng, H.; and Zhao, B. Y. 2019. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE symposium on security and privacy (SP)*, 707–723. IEEE.

Wang, B. e. a. 2023. Adversarial Robustness of Diffusion Models. *ICLR*.

Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.

Zhai, S.; Dong, Y.; Shen, Q.; Pu, S.; Fang, Y.; and Su, H. 2023. Text-to-Image Diffusion Models can be Easily Backdoored through Multimodal Data Poisoning. In *Proceedings of the 31st ACM International Conference on Multimedia, MM ’23*, 1577–1587. ACM.

Zhang, L.; and Agrawala, M. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. *arXiv preprint arXiv:2302.05543*.

Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.

A Extended Results

Summary. This section reports full image-level quality metrics for all datasets, models, and poison ratios. Values are shown as separate *Clean* and *Poisoned* scores per metric; arrows (\downarrow / \uparrow) indicate directionality. Baseline rows compare clean inputs to an unpoisoned ControlNet. See the main text for the compact summary.

Table 3: Image-level distortion, perceptual, and CLIP similarity metrics across datasets, models, and poison fractions. “ \downarrow ” indicates that lower is better while “ \uparrow ” indicates that higher is better.

Dataset	Model	Poison (%)	MSE \downarrow		LPIPS \downarrow		SSIM \uparrow		PSNR \uparrow		CLIP Score \uparrow	
					Clean	Poisoned	Clean	Poisoned	Clean	Poisoned	Clean	Poisoned
			Clean	Poisoned	Clean	Poisoned	Clean	Poisoned	Clean	Poisoned	Clean	Poisoned
ImageNet	SD v1.5	1	0.10	0.05	0.58	0.45	0.27	0.66	11.07	13.53	0.74	0.85
		5	0.11	0.03	0.61	0.38	0.25	0.70	10.47	15.12	0.72	0.87
		10	0.12	0.04	0.61	0.44	0.24	0.70	10.12	14.31	0.70	0.84
	Baseline	-	0.10	0.12	0.61	0.84	0.25	0.28	10.51	9.94	0.66	0.42
	SD v2	1	0.08	0.05	0.60	0.59	0.39	0.71	11.49	13.07	0.72	0.75
		5	0.09	0.03	0.63	0.49	0.37	0.77	10.94	15.64	0.67	0.88
		10	0.10	0.02	0.62	0.36	0.38	0.78	10.96	16.28	0.69	0.91
	Baseline	-	0.09	0.11	0.71	0.91	0.35	0.48	11.08	10.27	0.64	0.44
	SD XL	1	0.02	0.14	0.52	0.80	0.58	0.53	17.02	9.65	0.73	0.44
		5	0.02	0.06	0.53	0.49	0.58	0.79	17.02	13.76	0.74	0.73
		10	0.02	0.05	0.53	0.43	0.57	0.83	16.97	14.86	0.74	0.78
	Baseline	-	0.02	0.14	0.54	0.84	0.55	0.47	16.24	9.29	0.70	0.42
CelebA-HQ	SD v1.5	1	0.08	0.07	0.54	0.61	0.36	0.51	11.73	12.01	0.69	0.63
		5	0.08	0.05	0.54	0.42	0.37	0.67	11.61	13.68	0.68	0.86
		10	0.07	0.05	0.56	0.49	0.35	0.62	11.83	12.85	0.64	0.77
	Baseline	-	0.09	0.12	0.60	0.85	0.25	0.26	10.93	9.86	0.58	0.36
	SD v2	1	0.06	0.03	0.59	0.47	0.43	0.74	12.27	15.25	0.69	0.81
		5	0.08	0.06	0.62	0.67	0.36	0.65	11.64	12.82	0.59	0.60
		10	0.08	0.05	0.67	0.60	0.37	0.68	11.32	13.28	0.60	0.77
	Baseline	-	0.08	0.11	0.79	0.95	0.37	0.46	11.31	10.20	0.56	0.40
	SD XL	1	0.02	0.14	0.46	0.63	0.65	0.59	17.29	9.11	0.65	0.42
		5	0.02	0.04	0.47	0.24	0.65	0.87	17.26	15.61	0.66	0.92
		10	0.08	0.02	0.48	0.35	0.65	0.79	13.07	17.11	0.67	0.76
	Baseline	-	0.02	0.16	0.58	0.73	0.62	0.52	16.38	8.59	0.54	0.40

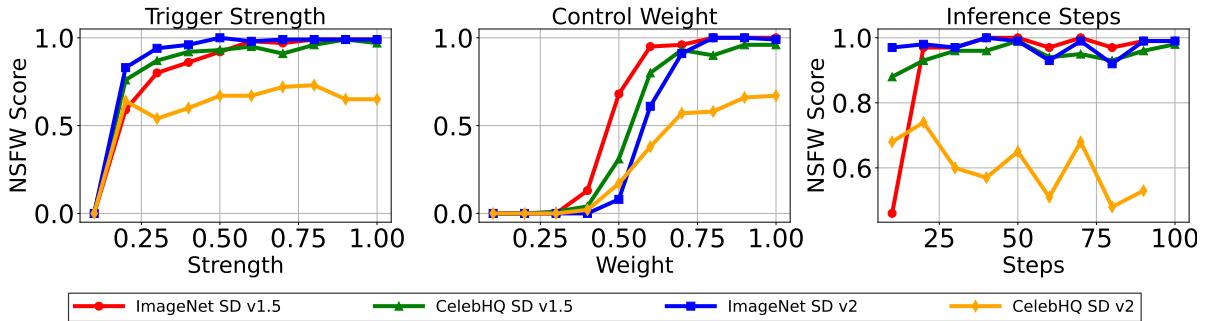


Figure 7: Ablations: NSFW vs. trigger amplitude, guidance scale, and sampler steps. Amplitude and guidance dominate; steps have minor effect.