SENSITIVITY ANALYSIS FOR DIFFUSION MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Training a diffusion model approximates a map from a data distribution ρ to the optimal score function s_t for that distribution. Can we differentiate this map? If we could, then we could predict how the score, and ultimately the model's samples, would change under small perturbations to the training set before committing to costly retraining. We give a closed-form procedure for computing this map's directional derivatives, relying only on black-box access to a pre-trained score model and its derivatives with respect to its inputs. We extend this result to estimate the sensitivity of a diffusion model's samples to additive perturbations of its target measure, with runtime comparable to sampling from a diffusion model and computing log-likelihoods along the sample path. Our method is robust to numerical and approximation error, and the resulting sensitivities correlate with changes in an image diffusion model's samples after retraining and fine-tuning.

1 Introduction

Diffusion models form a powerful class of generative models that allow users to generate images of nearly any subject in nearly any style in just a few keystrokes. However, this flexibility also allows diffusion models to engage in legally fraught behavior, such as generating images that mimic an artist's style. This has put diffusion models at the center of recent litigation¹ alleging that they facilitate copyright and trademark infringement. Understanding and mitigating the causes of this behavior have therefore become pressing challenges as businesses seek to integrate diffusion models into their consumer offerings.

Diffusion models generate images by iteratively transforming Gaussian noise using the *score function* of a distribution over noisy images, which is learned in practice from a large set of training images. Since the learned score is, in principle, determined by the training images, a natural strategy for understanding a diffusion model's behavior is to study how it depends on the training data. A sensible framework for this task should be able to answer questions such as: "What would the score function be if a sample were added or removed from the training set?" and "What would a generated image have looked like if a sample were added or removed from the training set?"

This work introduces a principled framework for answering such questions about diffusion models in the perturbative regime, where one considers infinitesimal changes in a model's training distribution. Because a diffusion model's output depends on the score function, which is itself determined by the training distribution, the core of our framework is a tractable closed-form expression for the directional derivatives of a score function with respect to the training distribution. This *sensitivity analysis* measures how the score function changes as a probability measure is up- or down-weighted in the training distribution; when this measure is a Dirac point mass, we obtain an exact expression for the *influence function* of the score. Crucially, our sensitivity analysis requires only black-box access to a pre-trained score function, and it does not require any knowledge of its training data or training procedure.

Using the adjoint method, we extend our sensitivity analysis for score functions to obtain schemes for computing the sensitivity of a diffusion model's samples to perturbations in the training data. These enable us to predict how a generated image would change if a collection of samples were added or removed from the training set. We demonstrate our method's robustness to a variety of sources of numerical error and show that its predictions are correlated with changes in a diffusion model's samples after retraining and after fine-tuning.

¹Andersen v. Stability AI Ltd., U.S. District Court, Northern District of California (2024).

2 RELATED WORK

Influence functions. Influence functions (Hampel, 1974) linearly approximate the change in a statistical estimator in response to infinitesimally upweighting a single training sample. Koh & Liang (2017) introduced influence functions to deep learning as a method for estimating the change in a neural network's parameters in response to perturbing its training set. Influence functions for generic optima of a training loss require the user to compute a costly inverse Hessian-vector product. Previous work (Guo et al., 2021; Schioppa et al., 2022) responds to this challenge by developing efficient approximations to this operation. In addition to this difficulty, influence functions assume that the learned model parameters minimize a strictly convex loss. This assumption is violated for neural networks, and Basu et al. (2020) find that in practice, influence functions for deep learning are brittle to network hyperparameters. Kwon et al. (2024); Mlodozeniec et al. (2025) introduce influence approximations that are specially adapted to generative models, including diffusion models, but these works follow Koh & Liang (2017) in estimating the influence of training samples on the learned network weights. In contrast, our sensitivity analysis uses the structure of diffusion models to directly compute the influence of training samples on the value of the score function and on model samples.

Data attribution for diffusion models. An emerging literature develops data attribution methods for diffusion models, which seek to estimate the impact of training samples on model outputs. Georgiev et al. (2023) use TRAK (Park et al., 2023), a gradient-based data attribution method developed primarily for supervised learning, to compute a per-example attribution score for a diffusion model's training data. This score estimates the change in the model's training loss induced by adding a particular sample to the training set. Following Park et al. (2023), they measure their method's effectiveness using the *linear datamodeling score*, which measures the rank correlation between their attribution score and actual training loss values attained by retrained models. Zheng et al. (2024) observe that one can improve upon the method from Georgiev et al. (2023) by computing their attribution scores using the gradients of the "wrong" model output function. Mlodozeniec et al. (2025) introduce an efficient approximation to the denoising loss Hessian and use it to estimate influence functions for attributing several proxies for model log-probabilities. Lin et al. (2025) propose attributing the KL divergence between the model distribution before and after deleting a training sample. Li et al. (2025) perform gradient-based data attribution using learnable weights for gradients with respect to different parameter groups. Whereas these methods estimate the impact of training samples on scalar quantities such as the training loss or proxies for log-probabilities, our sensitivity analysis estimates the effect of perturbations in the target distribution on the values of the score function and on model samples.

3 Method

In this section, we begin by observing that a diffusion model defines a map from its training distribution ρ to a score function s_t and then show how to tractably compute its directional derivatives. We then use this result to estimate how a diffusion model's samples change in response to small perturbations of its training distribution.

3.1 Preliminaries

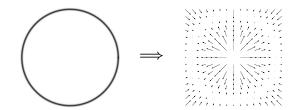
Diffusion models sample from a target distribution ρ by drawing samples from a Gaussian base distribution $\mathcal{N}(0,I)$ and flowing them through a possibly noisy velocity field v_t from $t=t_0$ to $t=t_1$. This yields a curve of probability distributions $\{\rho_t: t\in [t_0,t_1]\}$ for which ρ_t is the marginal distribution of the random variable $Z_t:=\alpha_t X_1+\sigma_t \epsilon$. Here, $X_1\sim \rho, \epsilon\sim \mathcal{N}(0,I)$, and α_t and σ_t are scale and noise schedules, respectively. These schedules are chosen so that at $t=t_0$, the samples are distributed according to a Gaussian distribution, and at $t=t_1$, the samples are distributed according to the target distribution ρ .

A diffusion model's velocity field v_t depends on ρ through the score function $s_t(z) := \nabla \log \rho_t(z)$ of ρ_t , which one learns in practice by solving a score-matching problem (Hyvärinen & Dayan, 2005). If one does not impose any restrictions on the hypothesis class, the optimal solution to this problem

is in fact available in closed form (Miyasawa, 1961):

$$s_t(z) = \frac{1}{\sigma_t^2} \left(\int w_t(z,x) \, \alpha_t x \, \mathrm{d}\rho(x) - z \right), \text{ where } w_t(z,x) := \frac{\exp\left(-\frac{1}{2\sigma_t^2} \|z - \alpha_t x\|_2^2\right)}{\int \exp\left(-\frac{1}{2\sigma_t^2} \|z - \alpha_t x\|_2^2\right) \, \mathrm{d}\rho(x)}. \tag{1}$$

This is a vector field pointing from z toward a distance-weighted average of rescaled samples $\alpha_t x$ from the target distribution ρ . Solving the scorematching problem therefore maps a measure ρ to a function s_t , which is fully determined by ρ and the scale and noise schedules.



We would like to estimate how the outputs of a diffusion model would change in response to perturbations

Figure 1: Solving the score-matching problem maps a measure ρ (left) to a vector-valued function s_t (right). We will compute this map's first variation in Section 3.2.

of the training data. These outputs depend on the training data only through the velocity field v_t and, in turn, through the score function s_t . We will therefore begin by introducing a tractable closed-form expression for the first variation of the map from ρ to s_t , which will describe how s_t responds to additive perturbations of the target distribution ρ .

3.2 Sensitivity analysis for score functions

Equation 1 expresses the score function s_t of a diffusion model in terms of its target distribution ρ . To understand how s_t changes in response to small perturbations of ρ , intuitively one would like to differentiate s_t with respect to the probability measure ρ . However, it is not obvious how to compute this derivative in practice. In this section, we present a tractable formula for such a derivative with respect to *additive* perturbations of ρ . This class of perturbations includes many cases of interest, such as the addition of new samples and the removal of existing samples from the training set.

Suppose that $\rho^{\eta}:=(1-\eta)\rho+\eta\nu$ is a mixture of two probability measures ρ and ν supported on \mathbb{R}^d , and let $s^{\eta}_t:\mathbb{R}^d\to\mathbb{R}^d$ be the score function of a diffusion model with target distribution ρ^{η} at time t. Differentiating s^{η}_t with respect to η and evaluating this derivative at $\bar{\eta}$ yields a function $g^{\bar{\eta}}_t:\mathbb{R}^d\to\mathbb{R}^d$ describing how s^{η}_t changes as one infinitesimally up-weights ν given initial weight $\bar{\eta}$.

The case $\bar{\eta}=0$ is of particular interest. For example, if $\bar{\eta}=0$, then $g_t^{\bar{\eta}}$ describes how a score function trained on ρ would vary as one introduces samples from ν . On the other hand, to approximate how the score function of a diffusion model trained on ρ would change in response to *removing* training data lying in some region $\Omega\subseteq\mathbb{R}^d$, one would define $\nu:=\rho_\Omega$, where ρ_Ω is the restriction of ρ to Ω , and consider $-g_t^{\bar{\eta}}$ evaluated at $\bar{\eta}=0$.

Our key result is the following theorem, which provides a tractable closed-form expression for q_i^0 :

Theorem 3.1 (Sensitivity analysis for score functions) For $\eta \in [0,1]$, let $\rho^{\eta} := (1-\eta)\rho + \eta\nu$ be a mixture of probability measures ρ and ν with compact support on \mathbb{R}^d . Let $\rho^{\eta}_t : \mathbb{R}^d \to \mathbb{R}$ and $s^{\eta}_t : \mathbb{R}^d \to \mathbb{R}^d$ be the density and score function, resp., of a diffusion model with target distribution ρ^{η} at time $t \in [t_0, t_1)$. Then the Fréchet derivative in $L^2(\mathbb{R}^d, \rho^{\eta}_t)$ of the map $T_t(\eta) : \eta \mapsto s^{\eta}_t$ evaluated at $\bar{\eta} = 0$ is the function $g^0_t : \mathbb{R}^d \to \mathbb{R}^d$ defined by the formula:

$$g_t^0(z) = \frac{\nu_t(z)}{\rho_t(z)} \left(s_t^{\nu}(z) - s_t^{\rho}(z) \right), \tag{2}$$

where $\nu_t(z)$, $\rho_t(z)$ are the respective densities and $s_t^{\nu}(z)$, $s_t^{\rho}(z)$ the respective scores at time t of diffusion models with target measures ν , ρ .

We prove this result in Appendix A.1. Whereas Equation 1 shows that there exists a well-defined map from a measure ρ to the optimal score function s_t of a diffusion model with ρ as its target, Equation 2 now provides a formula for the *directional derivative* of this map in the direction of

 $\nu-\rho$. In Figure 2, we depict an instance of this directional derivative g_t^0 when ρ is supported on a curve in 2D and ν is a Gaussian measure centered just off the curve. g_t^0 is a vector field pointing away from the support of ρ and towards the support of ν ; on account of the $v_t(z)/\rho_t(z)$ scaling factor, $\|g_t^0(z)\|_2$ is large at points z that are closer to the support of ν than to the support of ρ .

In a typical use case, ρ_t is the distribution at time t of a diffusion model trained on ρ , and $\nu = \frac{1}{K} \sum_{k=1}^K \delta_{x_k}$ is the empirical distribution on K samples x_k that one wishes to add or remove from ρ ; if K=1, we recover the *influence function* of T_t (Hampel, 1974). In this setting, we may use Equation 2 to compute $g_t(z)$ given *only blackbox access* to the score function $s_t^{\rho}(z)$ and the K samples x_k . The density $\rho_t(z)$ of the diffusion model can be computed from its score using the continuous change of variables (CCoV) formula (Song et al., 2021), and since ν_t is a mixture of Gaussians when ν is an empirical distribution its density and score function can be approximated.

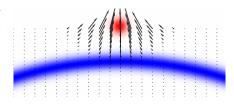


Figure 2: The score sensitivity g_t^0 is a vector field pointing away from the support of ρ and towards the support of ν .

tribution, its density and score function can be computed in closed form or efficiently approximated using techniques from Scarvelis et al. (2025). For perturbation sets $S = \{x_k\}_{k=1}^K$ of moderate size, the cost of evaluating Equation 2 is dominated by the cost of the density computations.

Theorem 3.1 shows how to tractably estimate the response of a pre-trained score function to additive perturbations in its target distribution. However, in practice, we are typically interested in how the *samples* generated by a diffusion model would change in response to perturbing its target distribution. Because samples are obtained by solving an ODE or SDE determined by the score function, Equation 2 should provide enough information to estimate the sensitivity of model samples to additive perturbations of ρ . We show this to be the case in the following section, using the adjoint method to obtain an analogous perturbation formula for a diffusion model's samples.

3.3 SENSITIVITY ANALYSIS FOR MODEL SAMPLES

Diffusion models generate samples from their target distribution ρ by solving a stochastic differential equation (SDE) or an ordinary differential equation (ODE) whose *drift* or *velocity field*, respectively, depend on the score s_t^{η} . Because this dependence is typically simple, often consisting of an affine transformation of s_t^{η} , it is easy to differentiate the drift or velocity field with respect to η given Equation 2. In this section, we will exploit this fact to compute the sensitivity of a diffusion model's samples with respect to additive perturbations.

ODE sampling. We begin with the simpler case of ODE sampling. Song et al. (2021) show that one may sample a diffusion model by solving a *probability flow ODE* (PF-ODE), whose initial condition is drawn from the Gaussian base distribution: $\frac{dz_t}{dt} = v_t^{\eta}(z_t)$ with $z_0 \sim \mathcal{N}(0, I)$. Because the Lipschitz constant of s_t^{η} – and consequently v_t^{η} – may blow up as $t \to t_1$, we follow a common convention from the theory of diffusion models and truncate integration of v_t^{η} at some $\tilde{t}_1 < t_1$ (De Bortoli, 2022). This convention aligns with typical diffusion model sampling schemes, which return samples at some time \tilde{t}_1 slightly earlier than the theoretical sampling interval endpoint t_1 .

If one further assumes that the target distributions μ, ν are compactly supported on \mathbb{R}^d , then a typical $v_t^\eta(z)$ will be globally Lipschitz for $z \in \mathbb{R}^d$ and $t \in [t_0, \tilde{t}_1]$. Khalil (2002, Theorem 3.2) then shows that there exists a unique solution to the PF-ODE for any initial condition $z_0 \in \mathbb{R}^d$. This allows us to define a solution $map \ \Phi_s^\eta(z_0) : \mathbb{R}^d \to \mathbb{R}^d$ that maps an initial condition $z_0 \in \mathbb{R}^d$ to the unique solution at time $s \in [t_0, \tilde{t}_1]$ of the initial value problem (IVP) defined by v_t^η . Intuitively, $\Phi_s^\eta(z_0)$ maps an initial noise sample $z_0 \sim \mathcal{N}(0, I)$ to the sample's position at time s along the diffusion model's sample path; at time $s = \tilde{t}_1$, this is simply a model sample.

We are interested in the derivative $\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi^{\eta}_{\tilde{t}_1}(z_0)$ for fixed initial conditions z_0 , which describes how the model sample $\Phi^{\eta}_{\tilde{t}_1}(z_0)$ generated from the Gaussian sample z_0 varies as one perturbs the target distribution μ in the direction of ν . Khalil (2002, Section 3.3) shows that under certain regularity conditions, this derivative solves an ODE known as the *sensitivity equation*. Defining $\psi_s := \frac{\mathrm{d}}{\mathrm{d}\eta}\Phi^{\eta}_s(z_0)$

and letting $z_s := \Phi_s^{\eta}(z_0)$ for $s \in [t_0, \tilde{t}_1]$ be a solution path for the PF-ODE, this equation is:

$$\frac{\mathrm{d}}{\mathrm{d}s}\psi_s = \frac{\mathrm{d}}{\mathrm{d}\eta}v_s^{\eta}(z_s) + J_z[v_s^{\eta}](z_s)\psi_s,\tag{3}$$

where the initial condition is $\psi_{t_0}=0$ and $J_z[v_s^\eta](z_s)$ denotes the spatial Jacobian of v_s^η evaluated at z_s . A solution $\psi_{\tilde{t}_1}=\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi_{\tilde{t}_1}^\eta(z_0)$ to Equation 3 approximates the change in a sample $z_{\tilde{t}_1}=\Phi_{\tilde{t}_1}^\eta(z_0)$ in response to additive perturbations of the target distribution ρ . Crucially, one may solve this IVP given black-box access to the score function s_s^η and its spatial derivatives. To estimate how a sample $\Phi_s^\eta(z_0)$ generated from initial noise z_0 would change in response to perturbing ρ , one should (1) compute a sample path z_t and model densities $\rho_t(z_t)$ by jointly integrating the PF-ODE and the CCoV formula,



Figure 3: A solution to Equation 3 approximates the change in a diffusion model's samples as the target distribution ρ is perturbed in the direction of ν .

(2) evaluate Equation 2 along the sample path, which also entails computing the density and score of the perturbation measure ν , and (3) integrate Equation 3, using autograd to compute the spatial Jacobian-vector products $J_z[v_s^\eta](z_s)\psi_s$. Figure 3 depicts a solution to Equation 3 for when ρ (in blue) is supported on a curve in 2D and ν (in red) is a Gaussian measure centered just off the curve.

SDE sampling. In practice, it is more common to sample a diffusion model by solving an SDE $\mathrm{d}z_t = f_t^\eta(z_t)\mathrm{d}t + g_t\mathrm{d}W_t$, where W_t denotes a Wiener process on \mathbb{R}^d . Only the drift coefficient f_t^η depends on the score function s_t^η and consequently on η ; conversely, the diffusion coefficient g_t is independent of η . Kunita (2019, Theorem 3.3.2) provides an analogous sensitivity analysis for the solution of an SDE whose coefficients depend on a parameter. Suppose the SDE has a unique solution and let $\Gamma_{s,\omega}^\eta : \mathbb{R}^d \to \mathbb{R}^d$ be the solution map sending an initial condition z_0 to the solution to the SDE at time $s \in [t_0, t_1]$ for some fixed realization ω of the Wiener process. Kunita (2019, Theorem 3.3.2) shows that under certain regularity conditions, which are satisfied for typical drifts if one truncates the integration at $\tilde{t}_1 < t_1$, $\frac{\mathrm{d}}{\mathrm{d}\eta}\Gamma_{s,\omega}^\eta(z_0)$ satisfies a particular SDE for almost all ω . Moreover, when the diffusion coefficient g_t is independent of η and the spatial variable, this differential equation is, in fact, deterministic and coincides with the sensitivity analysis for ODE sampling specified by Equation 3. We may therefore use Equation 3 to approximate the change in a diffusion model's SDE samples in response to perturbations of its target distribution. In practice, one follows the recipe from the previous section on ODE sampling, but replaces the ODE sample path z_t with an SDE sample path.

4 EXPERIMENTS

In this section, we empirically validate our sensitivity analysis for diffusion models. We initially study the impact of approximation error on our method's accuracy using synthetic examples where exact score functions and log-probabilities are available. We then experiment with neural diffusion models trained on image datasets and show that our sample sensitivities correlate with changes in model samples after retraining and after fine-tuning.

4.1 FIRST-ORDER APPROXIMATION FOR PERTURBED MODEL SAMPLES

A solution $\psi_{\tilde{t}_1} = \frac{\mathrm{d}}{\mathrm{d}\eta} \Phi_{\tilde{t}_1}^{\eta}(z_0)|_{\eta=0}$ to Equation 3 estimates the change in a diffusion model's samples in response to an additive perturbation of the target distribution. Using this sensitivity analysis, one may compute a first-order approximation of the samples a diffusion model trained on ρ^0 would have generated if its target distribution were $\rho^{\bar{\eta}}$:

$$\Phi_{\tilde{t}_1}^{\bar{\eta}}(z_0) \approx \Phi_{\tilde{t}_1}^0(z_0) + \bar{\eta} \frac{\mathrm{d}}{\mathrm{d}\eta} \Phi_{\tilde{t}_1}^{\eta}(z_0) \Big|_{\eta=0}. \tag{4}$$

Taylor's theorem states that this approximation converges at rate $o(\bar{\eta})$. However, because we estimate the derivative by solving a differential equation (3), numerical error may render this approximation inaccurate for practical step sizes. Furthermore, evaluating Equation 2 requires the density

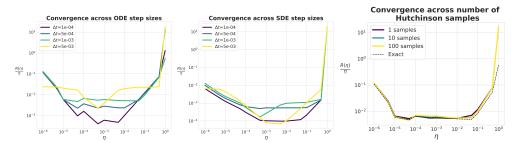


Figure 4: The approximation error of our first-order approximation to a perturbed model's samples decays at rate $o(\bar{\eta})$ for a variety of ODE step sizes (left) and SDE step sizes (center), and this rate is robust to noise from Hutchinson's estimator (right).

 ho_t of a score model, typically approximated using the CCoV formula. In practice, this formula involves a noisy estimate of ${\rm div}(v_t)$ using Hutchinson's trace estimator, introducing numerical error. In this section, we use synthetic data for which exact score functions and log-densities are available to study the impact of numerical error in ODE integration and density estimation on the convergence of our linear approximation to a perturbed model's samples. Appendix C.1.1 provides implementation details.

Effect of step size. In this experiment, we analyze the impact of step size on the convergence of our first-order approximation to the perturbed model's samples when solving Equation 3 using a forward Euler scheme. We let the initial target measure ρ be an equally-weighted mixture of well-separated Gaussians on \mathbb{R}^{100} . This multimodal, high-dimensional target distribution simulates some of the pathologies of real-world data, and for all t, the distribution of the resulting diffusion model ρ_t is a mixture of Gaussians, whose score and density are available in closed form. We perturb ρ in the direction of a second Gaussian distribution ν , which ensures that the perturbed target $\rho^{\bar{\eta}} = (1 - \bar{\eta})\rho + \bar{\eta}\nu$ remains a mixture of Gaussians.

We obtain sample paths z_t for ρ_t and $\rho_t^{\bar{\eta}}$ by numerically integrating the PF-ODE and variance-preserving SDE using a forward Euler scheme and an Euler-Maruyama scheme, resp., with several step sizes, and exactly compute the Gaussian mixture densities $\rho_t(z_t)$ along each sample path. We then integrate Equation 3 using the same forward Euler scheme to obtain the sensitivities $\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi_{t_1}^{\eta}(z_0)|_{\eta=0}$ of samples from the initial target to perturbations in direction ν . Equation 4 suggests that we may use these sensitivities to estimate the final positions of samples from the perturbed model $\rho^{\bar{\eta}}$, and Taylor's theorem implies that the error $R(\bar{\eta})$ of this estimate is $o(\bar{\eta})$, so $o(\bar{\eta})/\bar{\eta} \to 0$ as $\bar{\eta} \to 0$. We compute $o(\bar{\eta})$ for a variety of $\bar{\eta} \in [0,1]$ and verify that this rate holds in practice.

The left and center panels of Figure 4 depicts the results of this experiment for ODE and SDE sample paths. Linearly approximating samples from a perturbed target measure $\rho^{\bar{\eta}}$ using our sample sensitivity analysis (3) is accurate within $o(\bar{\eta})$ for a variety of realistic step sizes and $\bar{\eta}$. For very small values of $\bar{\eta}$, $R(\bar{\eta})/\bar{\eta}$ plateaus and begins to increase again. This may reflect a noise floor in the accuracy of model samples $\Phi^{\bar{\eta}}_{\bar{t}_1}(z_0)$ and $\Phi^0_{\bar{t}_1}(z_0)$, which are themselves computed by numerically integrating an ODE or SDE.

Effect of Hutchinson's estimator. In the previous experiment, we exactly computed the model densities $\rho_t(z)$ that appear in the score sensitivity analysis (2). This was possible because we selected the target distribution ρ to ensure that the densities $\rho_t(z)$ are available in closed form. In practice, one trains a neural network to approximate a diffusion model's score function $\nabla \log \rho_t(z)$ and obtains model densities via the continuous change of variables (CCoV) formula $\frac{d \log \rho_t(z_t)}{dt} = -\text{tr}(J_{z_t}[v_t^{\eta}](z_t))$. To avoid materializing a large Jacobian, one typically employs Hutchinson's trace estimator $\text{tr}(A) = \mathbb{E}[\epsilon^{\top} A \epsilon]$ to approximate the RHS, whose estimates are noisy due to the small number of ϵ samples used to compute the estimator. To study the impact of this noise on the convergence of Equation 4, we repeat the previous experiment with step size 10^{-3} while estimating the model densities $\rho_t(z)$ using the CCoV formula with a varying number of Hutchinson ϵ samples.

We depict the results of this experiment in the right panel of Figure 4, where the scaled remainders $P(\bar{\eta})/\bar{\eta}$ arising from using exact densities ρ_t are represented by the dashed line, and we experiment

with 1, 10, and 100 ϵ samples in Hutchinson's trace estimator. Our method's convergence rate is robust to noise in Hutchinson's trace estimator, with even a single draw of ϵ achieving nearly the same approximation error as the exact densities for all but the largest values of η .

4.2 STABILITY OF SAMPLE SENSITIVITY UNDER SCORE APPROXIMATION ERROR

Section 4.1 showed that one may approximate perturbed samples $\Phi_{\bar{t}_1}^{\bar{\eta}}(z_0)$ using our sensitivity analysis formula (3) and recover the expected $o(\bar{\eta})$ convergence rate in spite of errors from numerical integration and Hutchinson's estimator. To isolate the effects of these errors, we used the exact score $\nabla \log \rho_t$ of the mixture of Gaussians ρ_t throughout our computations. In practice, however, one typically *learns* this score function by training a neural network to optimize a score-matching objective. Here, we show that our sample sensitivity analysis (3) is stable to error from neural approximations of the true score function.

In this experiment, our initial target measure ρ is a mixture of well-separated Gaussians on \mathbb{R}^{10} , and we perturb ρ in the direction of a Gaussian measure ν . Instead of evaluating the

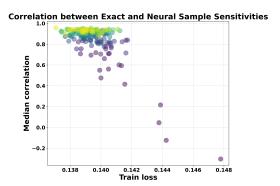


Figure 5: The correlation between sample sensitivities for an exact diffusion model and its neural approximation rises rapidly as the training loss falls. Points are colored from purple to yellow according to the training step.

score of ρ_t in closed form as in Section 4.1, we train a neural network to approximate this score function. We fix the base samples $z_0 \sim \rho_0$ and evaluate the sensitivity of samples from the exact diffusion model ρ_t and its neural approximation every 1000 training steps. We discretize all ODEs using forward Euler and use Hutchinson's estimator to estimate the model densities $\rho_t(z)$. We measure the median correlation between the exact and approximate sample sensitivities and compare it to the value of the score-matching loss at that training step. Appendix C.1.2 provides further implementation details.

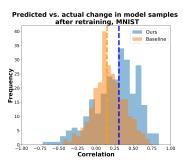
Figure 5 shows the relationship between the training loss and the median correlation between the exact and approximate sample sensitivities at each model sample. The points are colored from purple to yellow according to the training step at which the loss and sensitivities were measured. Our sample sensitivity analysis is stable to approximation error in the score function, with correlations between sample sensitivities computed with the exact diffusion model and its neural approximation rising rapidly as the training loss falls. This shows that our formulas may provide useful information even when an exact score function is replaced by a neural approximation. In the following section, we build on this observation by showing that our sample sensitivities are correlated with the direction of change in a diffusion model's samples after retraining on a perturbed training distribution.

4.3 IMAGE DATASETS

Predicting change in model samples after retraining. In the previous section, we used synthetic data from a mixture of Gaussians to study the robustness of our sensitivity analysis to various sources of numerical error. In practice, diffusion models are trained on large datasets of images, with training often stopped well before convergence to prevent memorization (Favero et al., 2025). In this section, we demonstrate that our sample sensitivities $\psi_{\tilde{t}_1}$ are correlated with differences between an image diffusion model's samples before and after retraining on a perturbed target distribution. We experiment with UNet-based diffusion models trained on a mixture of the MNIST and Typography-MNIST (TMNIST) datasets (Magre & Brown, 2022) and on the CelebA dataset (Liu et al., 2015).

For each dataset, we train a base model and a perturbed model whose target distribution $\rho^{\bar{\eta}}$ is a mixture of the base model's target distribution and the empirical measure on a set of new samples S. We employ mixture weights $\bar{\eta}=0.1$ and $1-\bar{\eta}=0.9$, resp. For our MNIST experiment, the new samples S are drawn from TMNIST, and for our CelebA experiment, S consists of samples with a large CLIP score for "a photo of an old man." We integrate the PF-ODE to obtain model samples from ρ^0 and $\rho^{\bar{\eta}}$, and also integrate Equation 3 with the perturbation measure ν set to the

empirical distribution over S to estimate the sensitivity of the base model's samples to upweighting S. We compare the sample sensitivities $\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi^\eta_{\tilde{t}_1}(z_0)|_{\eta=0}$ to the difference $\Phi^{\bar{\eta}}_{\tilde{t}_1}(z_0)-\Phi^0_{\tilde{t}_1}(z_0)$ between PF-ODE samples from the perturbed and base model given the same initial noise. This measures how much our sample sensitivity analysis predicts actual changes in model samples after retraining on the perturbed target distribution $\rho^{\bar{\eta}}$. Appendix C.2 provides further implementation details.



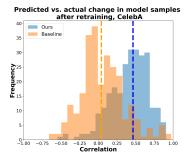
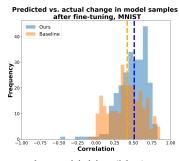


Figure 6: Correlations between predicted and actual change in model samples after retraining on a perturbed dataset. Our sample sensitivity analysis (blue) outperforms an optimal transport baseline (orange), achieving a median correlation (dashed blue line) of 0.46 on CelebA and 0.31 on MNIST.

Figure 6 depicts histograms of the correlations between our sample sensitivities and the actual change in model samples. As a baseline, we also compute the entropic optimal transport (OT) coupling (Cuturi, 2013) between the base model samples $\Phi^0_{\tilde{t}_1}(z_0)$ and the target distribution for the perturbed model and use the resulting transport rays as predicted directions of change in the model samples after retraining. Our sample sensitivity scores correlate with actual changes in model samples after retraining on $\rho^{\bar{\eta}}$ and substantially outperform the OT baseline, achieving a median correlation of 0.46 on CelebA and 0.31 on MNIST, compared 0.04 and 0.15, resp., for the OT baseline.

Predicting change in model samples after fine-tuning. The previous experiment shows that our sample sensitivities $\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi_{\tilde{t}_1}^{\eta}(z_0)|_{\eta=0}$ correlate with changes in model samples after retraining on a perturbed target distribution. We will now show that our sample sensitivities are more strongly predictive of changes in model samples after *fine-tuning* on new training samples S. We use the same base models and the same S as in the previous experiment, but fine-tune on S rather than retraining from scratch on the mixture distribution $\rho^{\bar{\eta}}$.



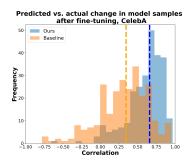


Figure 7: Our sample sensitivities (blue) are correlated with changes in model samples after fine-tuning, and continue to outperform an optimal transport baseline (orange).

We depict histograms of the correlations between our sample sensitivities and actual change in model samples after fine-tuning in Figure 7. We use the same entropic OT baseline as in the previous experiment, but compute transport rays between the base model samples and the samples S on which we fine-tuned. Both our sample sensitivities and the OT baseline are better correlated with actual change in model samples after fine-tuning, but our method continues to outperform the baseline, achieving a median correlation of 0.66 on CelebA and 0.51 on MNIST, compared to 0.34 and 0.42, resp., for the baseline. We visually compare our sample sensitivities to actual changes in model samples after fine-tuning in Figure 8, which shows that our sample sensitivity analysis can provide



Figure 8: Our sample sensitivities (top rows) predict changes in model samples after fine-tuning (bottom rows). The models that generated the samples on the left were trained on MNIST and fine-tuned on TMNIST. The models that generated the samples on the right were trained on CelebA and fine-tuned on a subset of photos with a large CLIP score for "a photo of an old man."

coarse predictions of how a diffusion model's samples might change after fine-tuning. We provide further illustrations of our sample sensitivity analysis in Appendix B.

5 DISCUSSION

Understanding a diffusion model's dependence on its training data is a critical challenge in machine learning. In general, one would expect the relationship between a large model and its training data to be complex and difficult to estimate. This paper shows that it is not only possible to compute directional derivatives of the map from a training distribution ρ to its optimal score function s_t , but that this computation is (a) surprisingly cheap, costing roughly as much as sampling a model and computing log-probabilities along the sample path, and (b) requires only black-box access to the score function. One can then leverage this simple formula to estimate how a diffusion model's samples change in response to perturbations to its target distribution before retraining or fine-tuning on new data. We propose several future directions for this line of work.

Throughout this paper, we perturbed a diffusion model's target measure with empirical measures over finite samples. This need not be the case: Our score sensitivity formula (2) holds for any compactly-supported perturbation measure ν , and it can be implemented in practice for any sequence of measures ν_t provided we can access their scores and densities. For instance, ν_t can be a second diffusion model, in which case Equation 2 resembles the formula for classifier-free guidance (CFG) (Ho & Salimans, 2022) with time- and spatially-varying weights. Future work might interpret CFG in light of our sensitivity analysis and design new guidance schedules based on this formula.

By composing a model's ODE sampling solution map with a text-conditioned classifier and applying our sensitivity formulas, one might also use our method to estimate how the likelihood that a diffusion model's samples match a prompt changes as one perturbs the training set. This would allow users to attribute a model's qualitative behavior to subsets of training samples and use this information to curate the training set to steer a diffusion model's behavior in a particular direction.

Finally, Kadkhodaie et al. (2024) find empirically that a diffusion model's sampling map is often insensitive to changes in its training set, and Favero et al. (2025) clarify that this behavior is controlled by the number of training iterations, with models becoming increasingly sensitive to dataset perturbations throughout training. Our sample sensitivity formula (3) *quantifies* this dependence and may serve as a valuable tool for future work on generalization in diffusion models.

Reproducibility statement. In the interest of reproducibility, we have included a complete proof of Theorem 3.1 in Appendix A.1, implementation details for our experiments in Appendix C, and we have attached a zip file with code to this submission.

REFERENCES

Samyadeep Basu, Phillip E. Pope, and Soheil Feizi. Influence functions in deep learning are fragile. *ArXiv*, abs/2006.14651, 2020. URL https://api.semanticscholar.org/CorpusID:220127956.

Marco Cuturi. Sinkhorn distances: lightspeed computation of optimal transport. In *Proceedings* of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13, pp. 2292–2300, Red Hook, NY, USA, 2013. Curran Associates Inc.

- Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL https://openreview.net/forum?id=MhK5aXo3gB. Expert Certification.
 - Alessandro Favero, Antonio Sclocchi, and Matthieu Wyart. Bigger isn't always memorizing: Early stopping overparameterized diffusion models, 2025. URL https://arxiv.org/abs/2505.16959.
 - Rémi Flamary, Cédric Vincent-Cuaz, Nicolas Courty, Alexandre Gramfort, Oleksii Kachaiev, Huy Quang Tran, Laurène David, Clément Bonet, Nathan Cassereau, Théo Gnassounou, Eloi Tanguy, Julie Delon, Antoine Collas, Sonia Mazelet, Laetitia Chapel, Tanguy Kerdoncuff, Xizheng Yu, Matthew Feickert, Paul Krzakala, Tianlin Liu, and Eduardo Fernandes Montesuma. Pot python optimal transport (version 0.9.5), 2024. URL https://github.com/PythonOT/POT.
 - Kristian Georgiev, Joshua Vendrow, Hadi Salman, Sung Min Park, and Aleksander Madry. The journey, not the destination: How data guides diffusion models. *arXiv preprint arXiv:2312.06205*, 2023.
 - Han Guo, Nazneen Rajani, Peter Hase, Mohit Bansal, and Caiming Xiong. FastIF: Scalable influence functions for efficient model interpretation and debugging. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 10333–10350, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.808. URL https://aclanthology.org/2021.emnlp-main.808/.
 - Frank R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974. ISSN 01621459, 1537274X. URL http://www.jstor.org/stable/2285666.
 - Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint* arXiv:2207.12598, 2022.
 - Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
 - Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ANvmVS2Yr0.
 - H.K. Khalil. Nonlinear Systems. Pearson Education. Prentice Hall, 2002. ISBN 9780130673893.
 - Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, pp. 1885–1894. PMLR, 2017.
 - Hiroshi Kunita. Stochastic Flows and Jump-Diffusions. Springer, 2019.
 - Yongchan Kwon, Eric Wu, Kevin Wu, and James Zou. Datainf: Efficiently estimating data influence in loRA-tuned LLMs and diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=9m02ib92Wz.
 - Shuangqi Li, Hieu Le, Jingyi Xu, and Mathieu Salzmann. Learning to weight parameters for data attribution. *arXiv preprint arXiv:2506.05647*, 2025.
 - Jinxu Lin, Linwei Tao, Minjing Dong, and Chang Xu. Diffusion attribution score: Evaluating training data influence in diffusion model. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=kuutidLf6R.
 - Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

- Nimish Magre and Nicholas Mainey Brown. Typography-mnist (tmnist): an mnist-style image dataset to categorize glyphs and font-styles. *ArXiv*, abs/2202.08112, 2022. URL https://api.semanticscholar.org/CorpusID:246867440.
- Koichi Miyasawa. An empirical bayes estimator of the mean of a normal population. *Bulletin of the International Statistical Institute*, 38(4):181–188, 1961.
- Bruno Kacper Mlodozeniec, Runa Eschenhagen, Juhan Bae, Alexander Immer, David Krueger, and Richard E. Turner. Influence functions for scalable data attribution in diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=esYrEndGsr.
- Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry. Trak: Attributing model behavior at scale. *arXiv preprint arXiv:2303.14186*, 2023.
- Christopher Scarvelis, Haitz Sáez de Ocáriz Borde, and Justin Solomon. Closed-form diffusion models. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL https://openreview.net/forum?id=JkMifr17wc.
- Andrea Schioppa, Polina Zablotskaia, David Vilar, and Artem Sokolov. Scaling up influence functions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 8179–8186, 2022.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.
- Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- Xiaosen Zheng, Tianyu Pang, Chao Du, Jing Jiang, and Min Lin. Intriguing properties of data attribution on diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=vKViCoKGcB.

A Proofs

A.1 PROOF OF THEOREM 3.1

We will prove this theorem in two parts. We will first show that $\frac{\partial}{\partial \eta} s_t^{\eta}(z) \Big|_{\eta=0} = g_t(z)$ at any fixed $z \in \mathbb{R}^d$. This shows that $g_t(z)$ is the pointwise derivative of s_t^{η} evaluated at $\eta=0$ at any $z \in \mathbb{R}^d$. We will then extend this pointwise argument to the space of functions by using the dominated convergence theorem (DCT) to prove that g_t is the Fréchet derivative in $L^2(\mathbb{R}^d, \rho_t^0)$ of the map $T(\eta): \eta \mapsto s_t^{\eta}$.

A.1.1 $g_t(z)$ is the pointwise derivative of s_t^{η} at $\eta=0$

In this part of the proof, we will rely heavily on Mlodozeniec et al. (2025, Lemma 1). A version of their lemma adapted to our setting states the following:

Lemma A.1 Let $\mathcal{L}: \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ be a C^2 function of additive form $\mathcal{L}(\eta, s_z) := \mathcal{L}_1(s_z) + \eta \mathcal{L}_2(s_z)$, and suppose that the map $s_z \mapsto \mathcal{L}(\eta, s_z)$ is strictly convex for all $\eta \in \mathbb{R}$. Fix $\bar{\eta}$ and choose s_z^* such that $\frac{\partial \mathcal{L}}{\partial s_z}(\bar{\eta}, s_z^*) = 0$. Then, by applying the implicit function theorem to $\frac{\partial \mathcal{L}}{\partial s_z}$, one obtains an open interval $(-\delta, \delta) \subseteq \mathbb{R}$ containing $\bar{\eta}$ and a unique function $\phi: (-\delta, \delta) \to \mathbb{R}^d$ such that $\phi(\bar{\eta}) = s_z^*$ and such that for all $\eta \in (-\delta, \delta)$, $\phi(\eta)$ is the unique minimizer of $s_z \mapsto \mathcal{L}(\eta, s_z)$. Moreover, ϕ is $C^{\bar{1}}$ with the following derivative:

$$\frac{\partial}{\partial \eta}\phi(\eta) = -\left[\frac{\partial^2 \mathcal{L}}{\partial s_z^2}(\eta, \phi(\eta))\right]^{-1} \frac{\partial \mathcal{L}_2}{\partial s_z}(\phi(\eta)). \tag{5}$$

To apply this lemma, we will first show how to specify a score function $s_t(z)$ of the form in Equation 1 evaluated at any $z \in \mathbb{R}^d$ as the minimizer of a score-matching objective. For the sake of simplicity, we will assume a constant scale schedule $\alpha_t \equiv 1$; our argument can be easily adapted to arbitrary scale schedules at the cost of additional notation.

Let $s_t(z): \mathbb{R}^d \to \mathbb{R}^d$ be the score function for some distribution $\rho_t := \rho * \mathcal{N}(0, \sigma_t^2 I)$, where ρ is a target distribution on \mathbb{R}^d . Kadkhodaie et al. (2024, Eqs. 14, 19) and the variational characterization of conditional expectation imply that this score function has the following pointwise variational characterization:

$$\begin{split} s_t(z) &= \nabla \log \rho_t(z) \\ &= \int_{\mathbb{R}^d} \left(\frac{x-z}{\sigma_t^2}\right) p(x|z) \mathrm{d}x \\ &= \underset{z(z)}{\operatorname{argmin}} \int \frac{1}{2} \left\|\frac{x-z}{\sigma_t^2} - s(z)\right\|^2 p(x|z) \mathrm{d}x, \end{split}$$

where p(x|z) is the conditional distribution of x given $z \sim \rho_t$. While p(x|z) is intractable a priori, $p(z|x) \sim \mathcal{N}(x, \sigma_t^2 I)$ is Gaussian, so we rewrite p(x|z) in this integral using Bayes' theorem:

$$\begin{split} s_t(z) &= \underset{s_z \in \mathbb{R}^d}{\operatorname{argmin}} \int \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 p(x|z) \mathrm{d}x \\ &= \underset{s_z \in \mathbb{R}^d}{\operatorname{argmin}} \int \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t(z)} \rho(x) \mathrm{d}x \\ &= \underset{s_z \in \mathbb{R}^d}{\operatorname{argmin}} \underset{x \sim \rho}{\mathbb{E}} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t(z)} \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \right]. \end{split} \tag{SM}$$

Here, we use $\mathcal{N}(z; x, \sigma_t^2 I)$ to denote the density of a Gaussian distribution with mean x and covariance $\sigma_t^2 I$ evaluated at $z \in \mathbb{R}^d$. This provides a pointwise definition of the score $s_t(z)$ of ρ_t evaluated at $z \in \mathbb{R}^d$ as the minimizer of the score-matching problem (SM). In particular, applying this argument to the target distribution ρ^{η} shows that:

$$s_t^{\eta}(z) = \underset{s_z \in \mathbb{R}^d}{\operatorname{argmin}} \mathbb{E} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \frac{1}{2} \| \frac{x - z}{\sigma_t^2} - s_z \|^2 \right].$$

Now, define the following objective functions, in which we take $z \in \mathbb{R}^d$ to be fixed:

$$\mathcal{L}_{\rho}(s_z) = \underset{x \sim \rho}{\mathbb{E}} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \right]$$

and

$$\mathcal{L}_{\nu}(s_z) = \underset{x \sim \nu}{\mathbb{E}} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_{\eta}^{\eta}(z)} \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \right].$$

Using these two objectives, we can define an objective $\mathcal{L}(\eta, s_z) := \underbrace{\mathcal{L}_{\rho}(s_z)}_{:=\mathcal{L}_1} + \eta \underbrace{\left(\mathcal{L}_{\nu}(s_z) - \mathcal{L}_{\rho}(s_z)\right)}_{:=\mathcal{L}_2}$

whose minimizer is precisely $s_t^{\eta}(z)$. This objective is in the additive form prescribed by Lemma A.1, which will put in position to apply the lemma once we check that its remaining hypotheses are satisfied.

To this end, note that \mathcal{L} is clearly a C^2 function of the prescribed additive form. Furthermore, as we will see below via a Hessian computation, the map $s_z \mapsto \mathcal{L}(\eta, s_z)$ is strictly convex for all $\eta \in \mathbb{R}$ and for all $z \in \mathbb{R}^d$. Fixing a point $(\bar{\eta}, s_z^*) = (\bar{\eta}, s_t^{\eta}(z))$ yields a critical point of \mathcal{L} with respect to s_z , which puts us in position to apply Lemma A.1.

Lemma A.1 gives us a function $\phi(\eta)$ defined on an open set containing $\bar{\eta}$ that maps η to the unique minimizer $s_t^{\eta}(z)$ of $\mathcal{L}(\eta,s_z)$. Crucially, it gives us a formula for the derivative $\frac{\partial}{\partial \eta}\phi(\eta)$, which involves the derivative $\frac{\partial \mathcal{L}_2}{\partial s_z}(\phi(\eta))$ and the Hessian $\frac{\partial^2 \mathcal{L}}{\partial s_z^2}(\eta,\phi(\eta))$. We will compute each of these terms separately.

We begin by computing the derivative $\frac{\partial \mathcal{L}_2}{\partial s_n}(\phi(\eta))$. We have

$$\begin{split} \frac{\partial \mathcal{L}_2}{\partial s_z}(\phi(\eta)) &= \frac{\partial}{\partial s_z} \left[\mathbb{E}_{x \sim \nu} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \right] - \mathbb{E}_{x \sim \rho} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \right] \right] \Big|_{s_t^{\eta}(z)} \\ &= \frac{\partial}{\partial s_z} \mathbb{E}_{x \sim (\nu - \rho)} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \frac{1}{2} \left\| \frac{x - z}{\sigma_t^2} - s_z \right\|^2 \right] \Big|_{s_t^{\eta}(z)} \\ &= \mathbb{E}_{x \sim \nu} \left[-\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x - z}{\sigma_t^2} - s_z \right) \right] \Big|_{s_t^{\eta}(z)} - \mathbb{E}_{x \sim \rho} \left[-\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x - z}{\sigma_t^2} - s_z \right) \right] \Big|_{s_t^{\eta}(z)} \\ &= \mathbb{E}_{x \sim \rho} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x - z}{\sigma_t^2} - s_t^{\eta}(z) \right) \right] - \mathbb{E}_{x \sim \nu} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x - z}{\sigma_t^2} - s_t^{\eta}(z) \right) \right] \end{split}$$

We now rewrite each expectation in the last line in terms of the scores of ρ_t and ν_t . To rewrite the first expectation, we pull out the factor of $\frac{1}{\rho_t^{\eta}(z)}$, which does not depend on x, and multiply by $1 = \frac{\rho_t(z)}{\rho_t(z)}$ to obtain the following:

$$\begin{split} \underset{x \sim \rho}{\mathbb{E}} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x - z}{\sigma_t^2} - s_t^{\eta}(z) \right) \right] &= \frac{\rho_t(z)}{\rho_t^{\eta}(z)} \underset{x \sim \rho}{\mathbb{E}} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t(z)} \left(\frac{x - z}{\sigma_t^2} - s_t^{\eta}(z) \right) \right] \\ &= \frac{\rho_t(z)}{\rho_t^{\eta}(z)} \left(\underbrace{\underbrace{\mathbb{E}}_{x \sim \rho} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t(z)} \left(\frac{x - z}{\sigma_t^2} \right) \right]}_{= s_t^{\rho}(z)} - \underbrace{\underbrace{\mathbb{E}}_{x \sim \rho} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t(z)} \right]}_{= 1} s_t^{\eta}(z) \right) \\ &= \frac{\rho_t(z)}{\rho_t^{\eta}(z)} \left(s_t^{\rho}(z) - s_t^{\eta}(z) \right). \end{split}$$

Analogous reasoning allows us to conclude that

$$\underset{x \sim \nu}{\mathbb{E}} \left[\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x-z}{\sigma_t^2} - s_t^{\eta}(z) \right) \right] = \frac{\nu_t(z)}{\rho_t^{\eta}(z)} \left(s_t^{\nu}(z) - s_t^{\eta}(z) \right),$$

and putting these together, we obtain

$$\frac{\partial \mathcal{L}_2}{\partial s_z}(\phi(\eta)) = \frac{\rho_t(z)}{\rho_t^{\eta}(z)} s_t^{\rho}(z) - \frac{\nu_t(z)}{\rho_t^{\eta}(z)} s_t^{\nu}(z) + \left(\frac{\nu_t(z) - \rho_t(z)}{\rho_t^{\eta}(z)}\right) s_t^{\eta}(z). \tag{6}$$

We now compute the Hessian term $\frac{\partial^2 \mathcal{L}}{\partial s_z^2}(\eta, \phi(\eta))$. Note that $\frac{\partial^2 \mathcal{L}}{\partial s_z^2} = \frac{\partial^2}{\partial s_z^2} \mathcal{L}_{\rho} + \eta(\frac{\partial^2}{\partial s_z^2} \mathcal{L}_{\nu} - \frac{\partial^2}{\partial s_z^2} \mathcal{L}_{\rho})$, and that we have already computed the relevant first derivatives:

$$rac{\partial}{\partial s_z} \mathcal{L}_{
ho} = \mathop{\mathbb{E}}_{x \sim
ho} \left[-rac{\mathcal{N}(z; x, \sigma_t^2 I)}{
ho_t^{\eta}(z)} \left(rac{x-z}{\sigma_t^2} - s_z
ight)
ight]$$

and

$$\frac{\partial}{\partial s_z} \mathcal{L}_{\nu} = \underset{x \sim \nu}{\mathbb{E}} \left[-\frac{\mathcal{N}(z; x, \sigma_t^2 I)}{\rho_t^{\eta}(z)} \left(\frac{x-z}{\sigma_t^2} - s_z \right) \right].$$

Differentiating again and simplifying, we see that

$$\frac{\partial^2}{\partial s_z^2} \mathcal{L}_\rho = \frac{\rho_t(z)}{\rho_t^{\eta}(z)} I$$

and

$$\frac{\partial^2}{\partial s_z^2} \mathcal{L}_{\nu} = \frac{\nu_t(z)}{\rho_t^{\eta}(z)} I.$$

Combining these and noting that $\rho_t^{\eta} = (1 - \eta)\rho_t + \eta\nu_t$, we conclude that $\frac{\partial^2 \mathcal{L}}{\partial s_z^2}(\eta, s_z) = I$ for all $\eta \in \mathbb{R}$ and for all s_z . In particular, the map $s_z \mapsto \mathcal{L}(\eta, s_z)$ is strictly convex for all $\eta \in [0, 1]$ as required by Lemma A.1.

We finally substitute these first and second derivatives into Equation 5 to obtain:

$$\begin{split} \frac{\partial}{\partial \eta} \phi(\eta) &= -\left[\frac{\partial^2 \mathcal{L}}{\partial s_z^2}(\eta, \phi(\eta))\right]^{-1} \frac{\partial \mathcal{L}_2}{\partial s_z}(\phi(\eta)) \\ &= -[I]^{-1} \left(\frac{\rho_t(z)}{\rho_t^{\eta}(z)} s_t^{\rho}(z) - \frac{\nu_t(z)}{\rho_t^{\eta}(z)} s_t^{\nu}(z) + \left(\frac{\nu_t(z) - \rho_t(z)}{\rho_t^{\eta}(z)}\right) s_t^{\eta}(z)\right) \\ &= \frac{\nu_t(z)}{\rho_t^{\eta}(z)} s_t^{\nu}(z) - \frac{\rho_t(z)}{\rho_t^{\eta}(z)} s_t^{\rho}(z) + \left(\frac{\rho_t(z) - \nu_t(z)}{\rho_t^{\eta}(z)}\right) s_t^{\eta}(z). \end{split}$$

In particular, if $\eta = 0$, then $\rho_t^{\eta}(z) = \rho_t(z)$ and this simplifies to:

$$\frac{\partial}{\partial \eta}\phi(\eta) = \frac{\nu_t(z)}{\rho_t(z)} \left(s_t^{\nu}(z) - s_t^{\rho}(z) \right) =: g_t(z).$$

This completes the first part of the proof.

A.1.2 g_t is the Fréchet derivative of $T_t(\eta): \eta \mapsto s_t^{\eta}$ at $\eta = 0$

We now extend this pointwise argument to the space of functions. Consider the map $T_t: \mathbb{R} \to L^2(\mathbb{R}^d, \rho_t^{\eta})$ that maps η to s_t^{η} . We will show that g_t is the Fréchet derivative of T_t at $\eta = 0$ for any $t \in [t_0, t_1]$. To do so, we need to show that for any $t \in [t_0, t_1]$,

$$\lim_{h \to 0} \left\| \frac{s_t^h - s_t^0}{h} - g_t \right\|_{L^2(\mathbb{R}^d, \rho_t^0)} = 0.$$

The previous section shows that $g_t(z)$ is the *pointwise* derivative of s_t^{η} with respect to η at $\eta = 0$. This means that for any $z \in \mathbb{R}^d$, $t \in [t_0, t_1]$,

$$\frac{s_t^h(z) - s_t^0(z)}{h} \to g_t(z).$$

 Hence $\frac{s_t^h - s_t^0}{h}$ converges pointwise to g_t for all $t \in [t_0, t_1]$. We will use the dominated convergence theorem (DCT) to lift this pointwise convergence to $L^2(\mathbb{R}^d, \rho_t^0)$ convergence. Define the following function:

$$F_h(z;t) := \frac{s_t^h(z) - s_t^0(z)}{h}$$

We need to show that there exists some real-valued function $G(z;t) \in L^2(\mathbb{R}^d, \rho_t^0)$ such that $\|F_h(z;t)\|_2 \leq G(z;t)$ uniformly in h for all z,t. To this end, note that by the mean value theorem, there exists some $\theta \in [0,1]$ such that:

$$||F_{h}(z;t)||_{2} = \left\| \frac{s_{t}^{h}(z) - s_{t}^{0}(z)}{h} \right\|_{2}$$

$$\leq \left\| \frac{\partial}{\partial \eta} s_{t}^{\eta}(z) \right\|_{\eta = \theta h} \right\|_{2}$$

$$= \left\| \frac{\rho_{t}(z)}{\rho_{t}^{\theta h}(z)} \left(s_{t}^{\rho}(z) - s_{t}^{\theta h}(z) \right) - \frac{\nu_{t}(z)}{\rho_{t}^{\theta h}(z)} \left(s_{t}^{\nu}(z) - s_{t}^{\theta h}(z) \right) \right\|_{2},$$

where the last line follows from a rearrangement of Equation 6. We can further simplify this bound to eliminate the dependence on h. First, note that $\rho_t^{\theta h} = (1 - \theta h)\rho_t + \theta h\nu_t$, so that for $h \leq \frac{1}{2}$, we have

$$\frac{1}{\rho_t^{\theta h}(z)} = \frac{1}{(1-\theta h)\rho_t(z) + \theta h\nu_t(z)} \le \frac{1}{(1-\theta h)\rho_t(z)} \le \frac{2}{\rho_t(z)}.$$

Hence, for h sufficiently small, we have:

$$||F_{h}(z;t)||_{2} \leq \left| \left| \frac{\rho_{t}(z)}{\rho_{t}^{\theta h}(z)} \left(s_{t}^{\rho}(z) - s_{t}^{\theta h}(z) \right) - \frac{\nu_{t}(z)}{\rho_{t}^{\theta h}(z)} \left(s_{t}^{\nu}(z) - s_{t}^{\theta h}(z) \right) \right| \right|_{2}$$

$$\leq \frac{2}{\rho_{t}(z)} \left| \left| \rho_{t}(z) \left(s_{t}^{\rho}(z) - s_{t}^{\theta h}(z) \right) - \nu_{t}(z) \left(s_{t}^{\nu}(z) - s_{t}^{\theta h}(z) \right) \right| \right|_{2}.$$

Applying the triangle inequality, we then obtain:

$$\frac{2}{\rho_{t}(z)} \| \rho_{t}(z) \left(s_{t}^{\rho}(z) - s_{t}^{\theta h}(z) \right) - \nu_{t}(z) \left(s_{t}^{\nu}(z) - s_{t}^{\theta h}(z) \right) \|_{2} \\
\leq \frac{2}{\rho_{t}(z)} \left(\rho_{t}(z) \| s_{t}^{\rho}(z) - s_{t}^{\theta h}(z) \|_{2} + \nu_{t}(z) \| s_{t}^{\nu}(z) - s_{t}^{\theta h}(z) \|_{2} \right).$$

Now, define

$$k_t^{\rho}(z) := \int w_t(z, x) x \mathrm{d}\rho(x)$$

and similarly for $k_t^{\nu}(z)$ and $k_t^{\theta h}(z)$. Then Equation 1 tells us that

$$s_t^{\rho}(z) = \frac{1}{\sigma_t^2} (k_t^{\rho}(z) - z),$$

and similar identities hold for the other score functions. Furthermore,

$$\|s_t^\rho(z) - s_t^{\theta h}(z)\|_2 = \frac{1}{\sigma_t^2} \|k_t^\rho(z) - k_t^{\theta h}(z)\|_2 \le \frac{1}{\sigma_t^2} \left(\|k_t^\rho(z)\|_2 + \|k_t^{\theta h}(z)\|_2 \right),$$

where the last line follows from the triangle inequality. Because $k_t^{\rho}(z)$ is a convex combination of points in the compact support of ρ , we can bound $\|k_t^{\rho}(z)\|_2 \leq D^{\rho} < +\infty$, where D^{ρ} is the diameter of the support of ρ . Similarly, $\|k_t^{\nu}(z)\|_2 \leq D^{\nu} < +\infty$, and because $\sup(\rho^{\theta h}) \subseteq \sup(\rho) \cup \sup(\nu)$, we have $\|k_t^{\theta h}(z)\|_2 \leq D^{\rho} + D^{\nu}$. Substituting these bounds into the above and simplifying, we obtain:

$$||F_h(z;t)||_2 \le \frac{2}{\rho_t(z)} \left(\rho_t(z) ||s_t^{\rho}(z) - s_t^{\theta h}(z)||_2 + \nu_t(z) ||s_t^{\nu}(z) - s_t^{\theta h}(z)||_2 \right)$$

$$\le \frac{2}{\rho_t(z)} \left(\frac{\rho_t(z)}{\sigma_t^2} (2D^{\rho} + D^{\nu}) + \frac{\nu_t(z)}{\sigma_t^2} (2D^{\nu} + D^{\rho}) \right)$$

$$=: G(z;t)$$

This function G(z;t) dominates $||F_h(z;t)||_2$ uniformly in h for all z,t. It remains to show that $G(z;t) \in L^2(\mathbb{R}^d, \rho_t^0)$. To this end, first note that $\rho_t^0 = \rho_t$. Then,

$$\int G(z;t)d\rho_t^0(z) = \int G(z;t)d\rho_t(z)$$

$$= \int \frac{2}{\rho_t(z)} \left(\frac{\rho_t(z)}{\sigma_t^2} (2D^\rho + D^\nu) + \frac{\nu_t(z)}{\sigma_t^2} (2D^\nu + D^\rho)\right) \rho_t(z)dz$$

$$= \frac{2(2D^\rho + D^\nu)}{\sigma_t^2} \underbrace{\int \rho_t(z)dz}_{=1} + \frac{2(2D^\nu + D^\rho)}{\sigma_t^2} \underbrace{\int \nu_t(z)dz}_{=1}$$

$$= \frac{6}{\sigma_t^2} (D^\rho + D^\nu)$$

This shows that $G(z;t) \in L^2(\mathbb{R}^d, \rho_t^0)$. As the hypotheses of the DCT are satisfied, we finally conclude that g_t is the Fréchet derivative of T_t at $\eta=0$ for any $t\in[t_0,t_1]$. This completes the proof of Theorem 3.1.

B VISUALIZING THE SAMPLE SENSITIVITIES

In this appendix, we illustrate our sample sensitivity analysis on images from the CelebA dataset. We draw four samples from a base model trained on CelebA and solve Equation 3 for each model sample and for four perturbation measures ν , each of which is an empirical measure over a perturbation set S. These perturbation sets consist of samples from the CelebA test set possessing the attribute labels "bald", "goatee", "smiling", and "eyeglasses". We depict model samples in the top row of Figure 9 and solutions to the sample sensitivity ODE for each perturbation set in the bottom four rows. Solutions to this ODE should approximate changes in the model samples in the top row in response to perturbing the base model's target distribution, and many of these predictions are intuitively reasonable. For instance, base model samples representing people without glasses are pushed towards samples of people with glasses in response to perturbing the target distribution towards CelebA samples with the "eyeglasses" attribute, and one observes similar phenomena for the other perturbation sets.

In Figure 10, we also depict line segments of the form $\Phi^0_{\tilde{t}_1}(z_0) + \alpha \frac{\mathrm{d}}{\mathrm{d}\eta} \Phi^\eta_{\tilde{t}_1}(z_0) \Big|_{\eta=0}$ for $\alpha \in [-2,2]$ and for the sample sensitivity ODE solutions depicted in Figure 9. These line segments should approximate samples from a model whose target distributed has been perturbed towards $\pm \nu$, where ν is the empirical measure over CelebA test images with the specified attributes. For α close to 0,



Figure 9: The bottom four rows depict solutions $\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi^\eta_{\tilde{t}_1}(z_0)\Big|_{\eta=0}$ to the sample sensitivity equation (3) for model samples $\Phi^0_{\tilde{t}_1}(z_0)$ pictured in the top row. In each of the lower four rows, the perturbation measure ν is the empirical distribution over images from the CelebA test set with attributes "bald", "goatee", "smiling", and "eyeglasses", resp.

the perturbed samples resemble the original sample (6th from the left in each row), differing mainly in the strength of the specified attribute. As α moves farther from 0, the perturbed samples deviate increasingly from the original.

C EXPERIMENT DETAILS

C.1 SYNTHETIC EXPERIMENTS

C.1.1 First-order approximation for perturbed model samples

In this experiment, the initial target measure ρ is an equally-weighted mixture of two Gaussians on \mathbb{R}^{100} with means (-1,...,-1) and (1,...,1), respectively, and shared covariance $\sigma^2 I$ for $\sigma=0.1$. We perturb ρ in the direction of a Gaussian distribution ν centred at (1,...,1) with covariance $\sigma^2 I$ for $\sigma=0.1$. For any $\bar{\eta}\in[0,1]$, the perturbed target $\rho^{\bar{\eta}}=(1-\bar{\eta})\rho+\bar{\eta}\nu$ is a mixture of Gaussians with the same means and covariances as ρ , but with weights $\frac{1-\bar{\eta}}{2}$ and $\frac{1+\bar{\eta}}{2}$.

We obtain sample paths z_t for ρ_t and $\rho_t^{\bar{\eta}}$ by fixing 1000 base samples $z_0 \sim \rho_0$ and numerically integrating the PF-ODE and variance-preserving SDE using a forward Euler scheme and Euler-Maruyama scheme, resp., with step sizes $\Delta t \in \{1 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}, 5 \times 10^{-3}\}$. Our scale and noise schedules come from a linear DDPMScheduler from the diffusers library with $\beta_{\text{start}} = 10^{-4}$ and $\beta_{\text{end}} = 0.02$. We exactly compute the Gaussian mixture densities $\rho_t(z_t)$ along each sample path. We then integrate Equation 3 using the same forward Euler scheme to obtain the sensitivities $\frac{d}{d\eta}\Phi_{\tilde{t}_1}^{\eta}(z_0)\Big|_{\eta=0}$ of samples from $\rho_{\tilde{t}_1}$. We compute the Taylor remainder:

$$R(\bar{\eta}) := \left(\Phi_{\tilde{t}_1}^{\bar{\eta}}(z_0) - \Phi_{\tilde{t}_1}^0(z_0) \right) - \bar{\eta} \frac{\mathrm{d}}{\mathrm{d}\eta} \Phi_{\tilde{t}_1}^{\eta}(z_0) \Big|_{\eta = 0}$$

and report the median value of $\frac{R(\bar{\eta})}{\bar{\eta}}$ across the 1000 batch samples in our plots.

In our experiments studying the effect of using Hutchinson's estimator to estimate model densities, we use the same setup as in the step size experiments, but estimate the base model densities $\rho_t(z)$ with Hutchinson's estimator: $\operatorname{tr}(A) = \mathbb{E}[\epsilon^{\top} A \epsilon]$. We use standard normal Gaussian samples for ϵ and report the number of noise samples we used in our plots.

C.1.2 STABILITY UNDER SCORE APPROXIMATION ERROR

Here, the initial target measure ρ is an equally-weighted mixture of two Gaussians on \mathbb{R}^{10} with means (-1,...,-1) and (1,...,1) and shared covariance $\sigma^2 I$ for $\sigma=0.1$. We perturb ρ in the direction of a Gaussian distribution ν centred at (1,...,1) with covariance $\sigma^2 I$ for $\sigma=0.1$. Instead of evaluating the score of ρ_t in closed form as in 4.1, we now train a neural network to approximate this score function. Our neural network is a two-hidden-layer MLP with SiLU activations and 512-dimensional hidden layers. We also use Fourier features (Tancik et al., 2020) with 128 frequencies and $\sigma=2.0$. We solve the score-matching problem using AdamW with a learning rate of 10^{-4} and a batch size of 100k. We train for 200k steps in total. In our plot, we omit the first two measurements of the correlations for clarity, as the training loss was large and network was very far from convergence during this phase of training .

We fix 1000 base samples $z_0 \sim \rho_0$ and evaluate the sensitivity of model samples from the exact diffusion model ρ_t and its neural approximation every 1000 training steps. We discretize all ODEs using a forward Euler scheme with step size 10^{-2} and use Hutchinson's estimator with 100 samples to estimate the model densities $\rho_t(z)$. We measure the median correlation between the exact and approximate sample sensitivities and compare it to the value of the score-matching loss at that training step in Figure 5.

C.2 IMAGE DATASETS

Retraining experiments. Each neural diffusion model in these experiments is parametrized by a Unet2DModel from the diffusers library. For the CelebA experiments, we set layers_per_block=2, block_out_channels=(128, 256, 512, 512), and norm_num_groups=32. We use a DDPMScheduler with $\beta_{\text{start}}=10^{-4}$ and $\beta_{\text{end}}=0.02$. The base model samples consist of 10k iid samples from the CelebA training set, and the new samples S are 495 CelebA training samples with a large CLIP score for "a photo of an old man". We preprocess the training images by center-cropping to a size of 140×140 , then resizing to 64×64 and normalizing to [-1,1]. We apply random horizontal flips as augmentations in training. We then train the CelebA diffusion models for 1000 epochs with an effective batch size of 512. Our optimizer is AdamW with a learning rate of 10^{-4} .

For the MNIST experiments, we set layers_per_block=2, block_out_channels=(32, 64, 128), and norm_num_groups=8. We use a DDPMScheduler with $\beta_{\text{start}}=10^{-4}$ and $\beta_{\text{end}}=0.02$. We do not apply any preprocessing to these samples. We train the MNIST diffusion models for 100 epochs with an effective batch size of 1024. Our optimizer is AdamW with a learning rate of 10^{-4} .

We draw model samples by integrating the PF-ODE and estimate model densities along the sample path using Hutchinson's estimator with 1 sample. We numerically integrate the PF-ODE and our sample sensitivity ODE (3) using a forward Euler scheme with a step size of 10^{-3} . We clamp the $\frac{\nu_t(z)}{\rho_t(z)}$ weights to [0.1,10] for numerical stability. For the entropic OT baseline, we use the sinkhorn_log algorithm from the POT package (Flamary et al., 2024) with a regularization value of 0.05 to compute the coupling matrix.

Fine-tuning experiments These experiments mostly replicate the setup in our retraining experiments but implement the following changes. For CelebA, we train the base model on 10k iid samples from the CelebA training set for 1k epochs with the same hyperparameters as in the retraining experiments, and then fine-tune for 200 epochs on 495 CelebA training samples with a large CLIP score for "a photo of an old man". We use the same learning rate of 10^{-4} for fine-tuning.

For MNIST, we train the base model on the MNIST training set for 100 epochs with an effective batch size of 1024 and a learning rate of 10^{-4} , and then fine-tune on TMNIST for a single epoch at a learning rate of 10^{-5} .

991

1022

1023

1024

1025

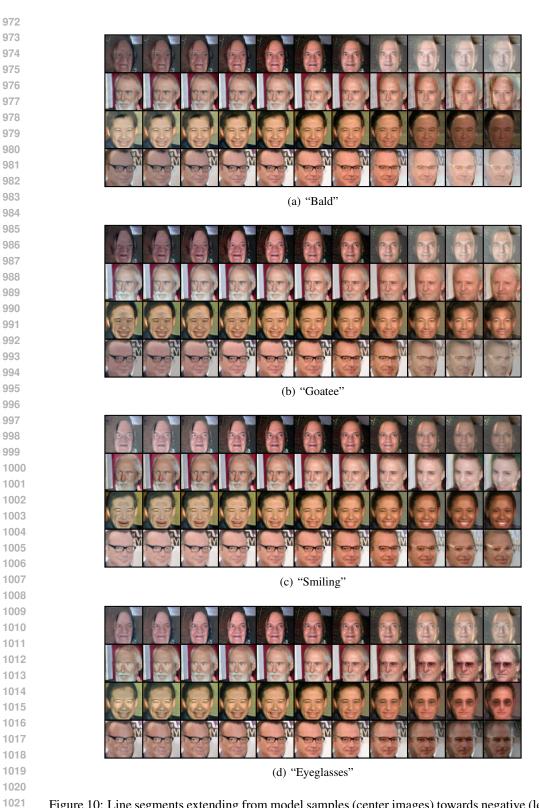


Figure 10: Line segments extending from model samples (center images) towards negative (left) and positive (right) multiples of sample sensitivities $\frac{\mathrm{d}}{\mathrm{d}\eta}\Phi^{\eta}_{t_1}(z_0)\Big|_{\eta=0}$. In each subfigure, the perturbation measure ν is the empirical distribution over CelebA test samples with an attribute listed in the subcaption.