SCORE MATCHING VIA DIFFERENTIABLE PHYSICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models based on stochastic differential equations (SDEs) gradually perturb a data distribution $p(\mathbf{x})$ over time by adding noise to it. A neural network is trained to approximate the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ at time t, which can be used to reverse the corruption process. In this paper, we focus on learning the score field that is associated with the time evolution according to a physics operator in the presence of natural non-deterministic physical processes like diffusion. A decisive difference to previous methods is that the SDE underlying our approach transforms the state of a physical system to another state at a later time. For that purpose, we replace the drift of the underlying SDE formulation with a differentiable simulator or a neural network approximation of the physics. At the core of our method, we optimize the so-called probability flow ODE to fit a training set of simulation trajectories inside an ODE solver and solve the reverse-time SDE for inference to sample plausible trajectories that evolve towards a given end state. We demonstrate the competitiveness of our approach for different challenging inverse problems.

1 INTRODUCTION

Many physical systems are time-reversible on a microscopic scale. For example, a continuous material can be represented by a collection of interacting particles (Gurtin, 1982; Blanc et al., 2002) based on which we can predict future states of the material. We can also compute earlier states, meaning we can evolve the simulation backwards in time (Martyna et al., 1996). When taking a macroscopic perspective, we only know the average quantities within specific regions (Farlow, 1993), which constitutes a loss of information. It is only then that time-reversibility is no longer possible, since many macroscopic and microscopic initial states exist that evolve to yield the same macroscopic state.

In the following, we target inverse problems to reconstruct the distribution of initial macroscopic states for a given end state. This genuinely tough problem has applications in many areas of scientific machine learning (Zhou et al., 1996; Gómez-Bombarelli et al., 2018; Delaquis et al., 2018; Lim & Psaltis, 2022), and existing methods lack tractable approaches to represent and sample the distribution of states. We address this issue by leveraging continuous approaches for diffusion models in the context of physical simulations. In particular, our work builds on the *reverse-diffusion* theorem (Anderson, 1982). Given the functions $f(\cdot, t) : \mathbb{R}^d \to \mathbb{R}^d$, called *drift*, and $g(\cdot) : \mathbb{R} \to \mathbb{R}$, called *diffusion*, it can be shown that under mild conditions, for the forward stochastic differential equation (SDE) $d\mathbf{x} = f(\mathbf{x}, t)dt + g(t)dw$ there is a corresponding reverse-time SDE $d\mathbf{x} = [f(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})]dt + g(t)d\tilde{w}$. In particular, this means that given a marginal distribution of states $p_0(\mathbf{x})$ at time $t = 0$ and $p_T(\mathbf{x})$ at $t = T$ such that the forward SDE transforms $p_0(\mathbf{x})$ to $p_T(\mathbf{x})$, then the reverse-time SDE runs backward in time and transforms $p_T(\mathbf{x})$ into $p_0(\mathbf{x})$. The term $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is called the *score*.

This theorem is a central building block for SDE-based diffusion models and denoising score matching (Song et al., 2021b; Jolicoeur-Martineau et al., 2021), which parameterize the drift and diffusion in such a way that the forward SDE corrupts the data and transforms it into random noise. By training a neural network to represent the score, the reverse-time SDE can be deployed as a generative model, which transforms samples from random noise $p_T(\mathbf{x})$ to the data distribution $p_0(\mathbf{x})$.

In this paper, we show that a similar methodology can likewise be employed to model physical processes. We replace the drift $f(\mathbf{x}, t)$ by a physics model $\mathcal{P}(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}^d$, which is implemented by a differentiable solver or a neural network that represent the dynamics of a physical system, thus

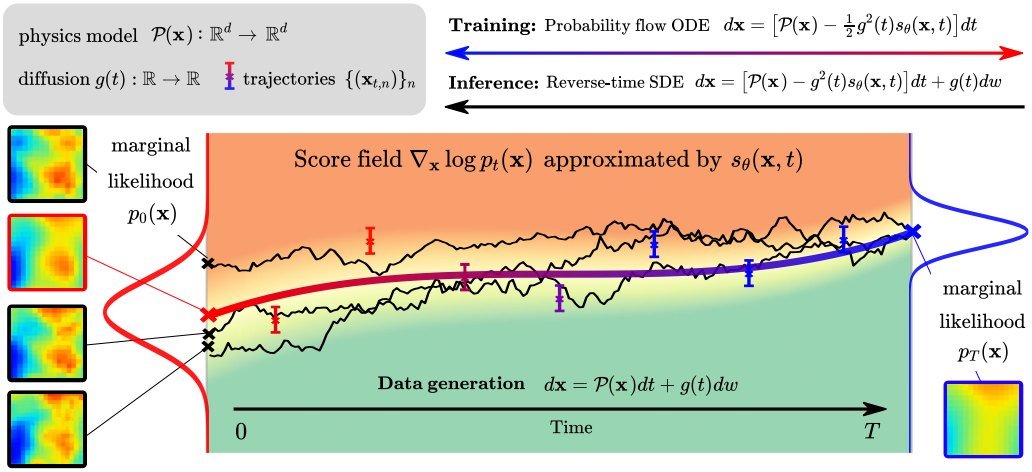

Figure 1: Overview: we employ a physics simulator $\mathcal{P}$ to learn the score field $\nabla_x \log p_t(\mathbf{x})$ with a neural network s_θ in the presence of noise or uncertainties. The trained model allows for sampling the posterior of p_0, i.e. different states that explain an observation p_T, via probability flow or by solving the reverse-time SDE.

deeply integrating physical knowledge into our method. The end state at $t = T$ on which the forward SDE acts is not fully destroyed by the diffusion $g(t)$, but instead, the noise acts as a perturbation of the system state over time. An overview of our method is shown in figure 1.

To the best of our knowledge, our work is the first to leverage the reverse-diffusion theorem as a method for solving inverse problems of physical systems. As such, our primary aim is to demonstrate how existing algorithms from this field can be used in the context of physics simulations. We showcase the efficacy of the score matching viewpoint on physics problems with a range of challenging inverse problems. Specifically, our contributions are:

1. We develop a framework in which we incorporate the reverse-diffusion theorem and score matching into a method for solving inverse problems that involve the time evolution of physical systems. We demonstrate its competitiveness against common baseline approaches using the heat equation as an example.
2. We highlight the effectiveness of our method with a more challenging inverse problem where we simulate a fluid-based transport process backwards in time in the presence of randomized obstacles. Here, we compare our method to different strategies for learned solvers.
3. Finally, we show that this approach can even be used when the underlying SDE is unknown. Our approach can be combined with operator learning methods and we demonstrate its effectiveness for learning the Navier-Stokes equation in the turbulent regime.

2 Background and Related Work

Learned solvers: Numerical simulations benefit greatly from machine learning models (Tompson et al., 2017; Morton et al., 2018; Pfaff et al., 2020; Li et al., 2020). By integrating a neural network inside differential equation solvers, it is possible learn to reduce numerical errors (Tompson et al., 2017; Kochkov et al., 2021; Brandstetter et al., 2022) or guide the simulation towards a desired target state (Holl et al., 2020b; Li et al., 2022). As errors may accumulate quickly over time, trained networks benefit from gradients that are backpropagated over multiple time steps (Um et al., 2020).

Diffusion models: Diffusion models (Ho et al., 2020; Song et al., 2021b) have been considered for a wide range of applications. Most notably, diffusion models have been proposed for image (Dhariwal & Nichol, 2021), video (Ho et al., 2022; Höppe et al., 2022; Yang et al., 2022) and audio synthesis (Chen et al., 2021). Specifically for uncertainty quantification, solving inverse problems and conditional sampling many methods have been proposed (Chung et al., 2021; 2022; Song et al., 2021a; Kawar et al., 2021; Ramzi et al., 2020). However, most approaches either focus on the

denoising objective that is common for tasks involving natural images, or the synthesis process of solutions does not take the underlying physics directly into account.

Denoising score matching: Classical denoising score matching approaches based on Langevin dynamics (Vincent, 2011; Song & Ermon, 2019, SMLD) and based on discrete Markov chains, e.g. Denoising Diffusion Probabilistic Models (Sohl-Dickstein et al., 2015; Ho et al., 2020, DDPM), can be unified in a time-continuous framework using SDEs (Song et al., 2021b). Given a distribution of states $p_0(\mathbf{x})$ at time $t = 0$, an SDE transforms $p_0(\mathbf{x})$ to a tractable distribution $p_T(\mathbf{x})$

$$d\mathbf{x} = f(\mathbf{x}, t)dt + g(t)dw, \tag{1}$$

with w the standard Brownian motion, a *drift* $f(\cdot, t) : \mathbb{R}^d \to \mathbb{R}^d$ and *diffusion* $g(\cdot) : \mathbb{R} \to \mathbb{R}$, which for $\mathbf{x}_0 \sim p_0(\mathbf{x})$ yields a diffusion process $(\mathbf{x}_t)_{t=0}^T$. As outlined above, the reverse-time SDE of the reverse-diffusion theorem (Anderson, 1982) is given by

$$d\mathbf{x} = [f(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})]dt + g(t)d\tilde{w}, \tag{2}$$

with $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ being the *score*. By sampling from the tractable distribution $p_T(\mathbf{x})$ and simulating the reverse-time SDE equation (2), we can generate samples from $p_0(\mathbf{x})$. Since the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is not known analytically, it is approximated by a neural network. Let $p_{st}(\mathbf{x}(t)|\mathbf{x}(s))$ for $0 \leq s \leq t \leq T$ denote the transition probability from state $\mathbf{x}(s)$ to $\mathbf{x}(t)$ given the SDE in equation (1). For denoising score matching, a neural network $s_\theta(\mathbf{x}, t)$ parameterized by θ approximates the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ and is trained to minimize the objective

$$\mathbb{E}_{t\sim\mathcal{U}(0,T)} \left\{ \lambda(t) \mathbb{E}_{\mathbf{x}(0)\sim p_0(\mathbf{x})} \mathbb{E}_{\mathbf{x}(t)|\mathbf{x}(0)\sim \mathrm{SDE}} \left[\left|\left| s_\theta(\mathbf{x}(t), t) - \nabla_{\mathbf{x}(t)} \log p_{0t}(\mathbf{x}(t)|\mathbf{x}(0)) \right|\right|_2^2 \right] \right\}, \tag{3}$$

with a weighting function $\lambda(\cdot)$. If $f(\cdot, t)$ and $g(\cdot)$ are chosen wisely, e.g. the drift is an affine function, there are closed-form solutions for the transition probability $p_{0t}(\mathbf{x}(t)|\mathbf{x}(0))$ (Särkkä & Solin, 2019; Song et al., 2021b).

Continuous normalizing flows: Continuous normalizing flows (CNFs) are invertible generative models based on neural ODEs. Given an initial distribution $z_0 \sim p_{z_0}(z_0)$ and a function $f(z(t), t; \theta)$ represented by a neural network, a CNF is trained to solve the ODE $\partial z(t)/\partial t = f(z(t), t; \theta)$ with boundary conditions $z(t_0) = z_0$ and $z(t_1) = \mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^d$ is a sample from the training data set. A useful property of CNFs is that it is cheap to compute the log-likelihood of data samples, due to the instantaneous change of variables formula $\partial \log p(z(t))/\partial t = -\mathrm{Tr}(\partial f/\partial z(t))$ (Grathwohl et al., 2019). The evolution of the marginal probability density $p_t(\mathbf{x})$ for the SDE in equation (1) is described by Kolmogorov's forward equation (Øksendal, 2003). Maoutsa et al. (2020) and Song et al. (2021b) show that there exists an ODE with the same Kolmogorov forward equation. This ODE is called *probability flow* ODE and is given by

$$d\mathbf{x} = \left[f(\mathbf{x}, t) - \frac{1}{2} g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt. \tag{4}$$

The probability flow ODE equation (4) represents a CNF and, if $f(\mathbf{x}, t)$ is known, we can train a model $s_\theta(\mathbf{x}, t)$ to represent the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ using standard methods (Chen et al., 2018). While the evolution of $p_t(\mathbf{x})$ is the same between the probability flow ODE from equation (4) and the reverse-time SDE from equation (2), there are caveats due to the approximation by $s_\theta(\mathbf{x}, t)$ (Song et al., 2021a; Lu et al., 2022). A recent variant combines score matching with CNFs (Zhang & Chen, 2021), and employs a joint training of the drift and score with an integration backwards in time.

3 METHOD

Modelling assumptions We consider a set of N trajectories $\{(\mathbf{x}_t^n)\}_{n=1}^N \subset \mathbb{R}^{d\times T}$ that describe the evolution of a physical system. Moreover, there is a known physics model $\mathcal{P}(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}^d$ that is differentiable and approximates the time evolution, i.e. $\mathbf{x}_{t_{m+1}} \approx \mathbf{x}_{t_m} + (t_{m+1} - t_m) \cdot \mathcal{P}(\mathbf{x}_{t_m})$. One of our key modelling choices is to describe the time evolution of the physical system by a stochastic differential equation

$$d\mathbf{x} = \mathcal{P}(\mathbf{x})dt + g(t)dw, \tag{5}$$

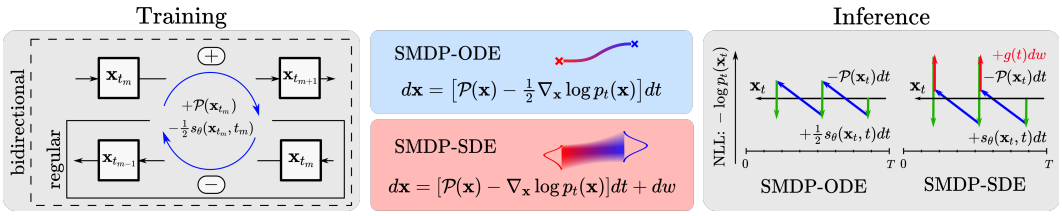

Figure 2: During the training phase, we optimize $s_\theta(\mathbf{x}, t)$ that approximates the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ inside the probability flow ODE to fit the data trajectories. For the inference part, we simulate trajectories of the reverse-time SDE.

with a diffusion process $g(\cdot) : \mathbb{R} \to \mathbb{R}$ that perturbs the states such that the trajectories $\{(\mathbf{x}_t^n)\}_{n=1}^N$ can be obtained. We can simulate paths from this SDE using Euler-Maruyama steps, i.e. for a time discretization $t_0 < t_1 < ... < t_M$ and initial state $\mathbf{x}_{t_0}$, we obtain the iteration rule

$$\mathbf{x}_{t_{m+1}} \leftarrow \mathbf{x}_{t_m} + (t_{m+1} - t_m) \cdot \mathcal{P}(\mathbf{x}_{t_m}) + \sqrt{t_{m+1} - t_m} \cdot g(t_m) z_{t_m}, \tag{6}$$

where z_{t_m} are i.i.d. with $z_{t_m} \sim \mathcal{N}(0, I)$. The random noise $g(t_m)z_{t_m}$ that is added to the data at each time step can be regarded as either noise inherent to the physical problem, or as noise from a measurement process.

Inverse Problem Given an end state $\mathbf{x}_T$ we are interested in recovering a likely trajectory $(\hat{\mathbf{x}}_t)$ that evolves towards $\mathbf{x}_T$. More formally, the set of trajectories $\{(\mathbf{x}_t^n)\}_{n=1}^N$ implicitly defines marginal likelihoods $p_t(\mathbf{x})$ at every time step t which are linked through the SDE equation (5) of the physical system by the Kolmogorov forward equation. The solution trajectory may not be unique, so we would like to sample from the full posterior instead of obtaining only a maximum likelihood solution, i.e. we want to sample from $p_0(\mathbf{x}|\mathbf{x}_T)$.

Method In line with previous work in score-based generative modelling (Song & Ermon, 2019; Song et al., 2021b), we approximate the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ of the marginal likelihoods by a neural network $s_\theta(\mathbf{x}, t)$. During training, we optimize $s_\theta(\mathbf{x}, t)$ of the probability flow ODE discussed in section 2

$$d\mathbf{x} = \left[\mathcal{P}(\mathbf{x}) - \frac{1}{2} g^2(t) s_\theta(\mathbf{x}, t)\right] dt \tag{7}$$

to fit the trajectories $\{(\mathbf{x}^n)_t\}_{n=1}^N$. In contrast to previous work, our method deeply integrates into the training process a prior about the physical system in the form of the simulation operator $\mathcal{P}(\mathbf{x})$. In this context, the end state at $t = T$ is not fully destroyed by the noise, but instead the noise acts as a perturbation of the system state over time. An overview of our method is shown in figure 2.

Given an end state $\mathbf{x}_T$, we can solve the probability flow backwards in time using the trained score function $s_\theta(\mathbf{x}, t)$ to obtain a trajectory $(\hat{\mathbf{x}}_t)$. However, this will only give a single solution and thus not allow for sampling from the full posterior $p(\mathbf{x}|\mathbf{x}_T)$. Additionally, when there is a loss of information, trajectories from the probability flow ODE tend to average many candidate solutions, which produces predictions that are often overly-smooth. To circumvent this, we simulate trajectories from the reverse-time SDE (see section 1)

$$d\mathbf{x} = \left[\mathcal{P}(\mathbf{x}) - g^2(t) s_\theta(\mathbf{x}, t)\right] dt + g(t) dw. \tag{8}$$

The evolution of marginal probabilities $p_t(\mathbf{x})$ for this SDE is the same as for the probability flow ODE equation (7) (Song et al., 2021b). Moreover, by the reverse-diffusion theorem (Anderson, 1982), SDE equation (8) is the time-reverse of the physical system SDE from equation (5). Therefore, we can sample from $p(\mathbf{x}|\mathbf{x}_T)$ by simulating trajectories from SDE equation (8) using any traditional SDE solver. In the following, we refer to the integration of the physics model $\mathcal{P}(\mathbf{x})$ into the score-based modelling approach as *score matching via differentiable physics*, or *SMDP* in short. We denote trajectories from the probability flow ODE by *SMDP-ODE*, and those obtained by simulating the reverse-time SDE by *SMDP-SDE*.

Training and Inference Algorithm 1 gives an overview of SMDP for the ODE as well as the SDE variant. Training is always performed with the ODE variant. The training objective is to minimize the difference between the distribution of initial state $p_0(\mathbf{x})$ and the distribution $\hat{p}_0(\mathbf{x})$ that is induced by equation (7) and the predicted trajectories $(\hat{\mathbf{x}})_t$ using the score representation $s_\theta(\mathbf{x}, t)$. In line with previous work, our loss is based on the squared difference between the predicted trajectory $(\hat{\mathbf{x}}_t)$ and target trajectory $(\mathbf{x}_t)$ from the training set, i.e. $\mathcal{L}((\hat{\mathbf{x}}_t), (\mathbf{x}_t)) := ((\hat{\mathbf{x}}_t) - (\mathbf{x}_t))^2$ (Um et al., 2020; Zhang & Chen, 2021). We update the weights θ of s_θ by unrolling and backpropagating though all ODE solver steps. While in principle any ODE solver can be used for training and inference, we concentrate on the explicit Euler method for simplicity. At inference time, a pre-trained s_θ can be used to produce ODE as well as SDE trajectories, as outlined in algorithm 1. We consider a variant of SMDP, for which we apply a *bidirectional training*, i.e. instead of only training the probability flow ODE in the time-backward direction $t_M \to t_0$, we additionally optimize the time-forward direction $t_0 \to t_M$ at the same time.

Algorithm 1 SMDP-ODE, SMDP-SDE

Require: $\mathbf{x}_{t_M}$, $\{t_m\}_{m=0}^{M}$, $\{g_{t_m}\}_{m=0}^{M}$
for $m \leftarrow M$ to 1 **do**
  $\mathbf{p} \leftarrow \mathcal{P}(\mathbf{x}_{t_m})$
  $\mathbf{s} \leftarrow s_\theta(\mathbf{x}_{t_m}, t_m)/2$
  if SMDP-ODE **then**
    $\mathbf{x}_{t_{m-1}} \leftarrow \mathbf{x}_{t_m} - (t_m - t_{m-1}) \cdot (\mathbf{p} + \mathbf{s})$
  if SMDP-SDE **then**
    $\mathbf{x}_{t_{m-1}} \leftarrow \mathbf{x}_{t_m} - (t_m - t_{m-1}) \cdot (\mathbf{p} + 2\mathbf{s})$
    $z \sim \mathcal{N}(0, I)$
    $\mathbf{x}_{t_{m-1}} \leftarrow \mathbf{x}_{t_{m-1}} + g_{t_m}\sqrt{t_m - t_{m-1}}z$
return $\mathbf{x}_{t_0}$

4 Experiments

We conduct several experiments to demonstrate the advantages of SMDP compared to a number of baseline methods. The source code for all experiments will be made available upon acceptance. We first consider the 2D heat equation in section 4.1, where our objective is to find possible initial states at time $t = 0$ given a noisy end state at time $t = T$. We consider different approaches to optimize the probability flow ODE in this situation, for example the bidirectional training. We evaluate the different methods in terms of (1) how well a sample $\hat{\mathbf{x}}$ drawn from our method, given an end state $\mathbf{x}_T$, fits $p_0(\mathbf{x})$ and (2) how well the trajectory defined by $\hat{\mathbf{x}}$ reconstructs $\mathbf{x}_T$ given the physics model $\mathcal{P}(\mathbf{x})$. Among others, we use these metrics to evaluate the importance of different time discretizations.

In our second experiment in section 4.2, we transfer the established practices to a more challenging problem, where we are interested in reconstructing the trajectory of a buoyancy-driven flow within a closed simulation domain given an end state at time T. What makes this problem challenging is that for each simulation, we place different obstacles at different positions within the simulation domain. For this case we we compare SMDP to regular, learning-based approaches.

Finally, in section 4.3, we consider the situation where the physics of the system is unknown. For this purpose, we consider training a network that approximates the solutions to the Navier-Stokes equations and a network that approximates the score field. We demonstrate that by doing so, we obtain an improved performance for inverse problems and the learned score can be used to refine predictions in post-processing.

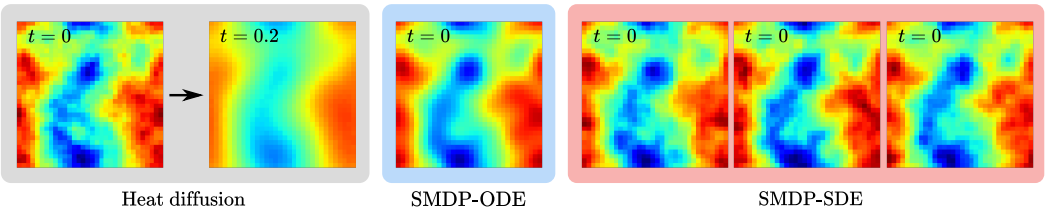

Figure 3: Heat diffusion case. We simulate a Gaussian random field at $t = 0$ forwards in time using equation (6) with $g(t) := \lambda$. For $\lambda = 0$, this corresponds to a smooth solution at $t = 0.2$. Given s_θ we can either solve the probability flow ODE or simulate trajectories of the reverse-time SDE to obtain solutions for the state at $t = 0$.

Method	MSE $[10^{-5}]\downarrow$	Spectral error $\downarrow$	Full posterior
SMDP-ODE	**0.74**	3.62	✗
SMDP-SDE	5.56	**0.56**	✓
ResNet-S	2.17	1.67	✗
ResNet-P	2.30	1.09	✗
BNN-S	3.47×10^2	1.25	✓
BNN-P	3.81×10^2	0.99	✓
FNO-S	2.54×10^4	1.60	✗
FNO-P	2.50×10^4	1.47	✗
Diffusion	1.39	4.45	✗
Diffusion+noise	4.45	3.24	✓

Table 1: Evaluation of reconstruction MSE and spectral error for SMDP and baselines. The column full posterior indicates whether models yield point estimates or allow to sample from the posterior.

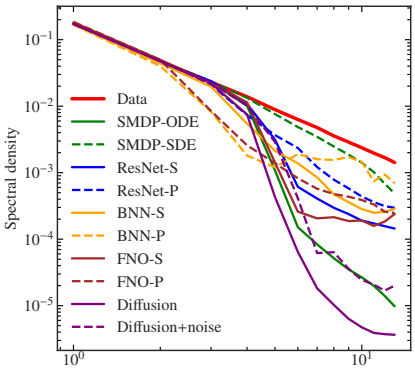

Figure 4: Spectral density on different scales, the red line indicating ground truth. The closer a method is to the ground truth, the better it produces structures of a similar scale.

4.1 Heat Equation

We consider the heat equation $\frac{\partial u}{\partial t} = \alpha \Delta u$ which plays a fundamental role in many physical systems. Here, we set the diffusivity constant to $\alpha = 1$ and initial conditions at $t = 0$ are generated from Gaussian random fields with $n = 4$ at resolution 32×32. This makes it possible to analytically determine the power spectral density of initial states. We simulate the heat diffusion using a spectral method until $t = 0.2$ with a fixed step size $\Delta t = 6.25 \times 10^{-3}$ using the iteration rule from equation (6) with $g \equiv 0.1$. Our training data set consists of 2.500 initial conditions with their corresponding trajectories sampled with varying step size Δt and end states at $t = 0.2$. The test set is comprised of 500 initial conditions and their corresponding end states solved directly without any noise.

Training We consider a small *ResNet*-like architecture based on an encoder and decoder part (see appendix A) as representation for the score function $s_\theta(\mathbf{x}, t)$. The physics model $\mathcal{P}$ is implemented via differentiable programming in *JAX* (Schoenholz & Cubuk, 2020). For better comparison with the baseline methods, these are trained with a Gaussian random noise of $\sigma = 0.1$ added to the inputs. This noise is applied to all network inputs during testing.

Baseline methods As baseline methods, we consider the ResNet-like architecture from above, in addition to a *Bayesian neural network* (BNN) based on a U-Net architecture with spatial dropout (Mueller et al., 2022), as well as a *Fourier neural operator* (FNO) network (Li et al., 2020). For each of these three methods, we consider two variants: the first variant is trained with a *supervised loss*, i.e. the training data consists of pairs $(\mathbf{x}_0, \mathbf{x}_T)$ with initial state $\mathbf{x}_0$ and end state $\mathbf{x}_T$. The supervised loss corresponds to the squared L2 distance between the network prediction $\hat{\mathbf{x}}_0$ and the ground truth, i.e. $(\hat{\mathbf{x}}_0 - \mathbf{x}_0)^2$. For the second variant, the *reconstruction loss*, we rely on the differentiable solver and only make use of the end state $\mathbf{x_T}$ such that the loss becomes $(\mathcal{P}(\hat{\mathbf{x}}_0; T) - x_T)^2$, i.e we simulate the network output forward in time using $\mathcal{P}$ to obtain a state at $t = T$, which we compare to the desired end state x_T. We denote the supervised variant by *S* and the physics-based one by *P*. Additionally, we consider an adopted diffusion model from Rissanen et al. (2022), denoted by *Diffusion*. We train this network similarly to SMDP-ODE, but without the solver $\mathcal{P}$, such that the network has to learn the score and the physics at the same time.

Reconstruction accuracy vs. fitting the data manifold We give an evaluation of our method and the baselines by considering the *reconstruction MSE* on the test set: how well a predicted initial state $\hat{\mathbf{x}}_0$ that is simulated forward in time yields states that correspond to the reference end state $\mathbf{x}_T$ in terms of MSE. This metric has the disadvantage that it does not measure how well the prediction matches the training data manifold, i.e. for this case whether the prediction resembles the statistics of the Gaussian random field. For that reason, we additionally compare the power spectral density of the states as the *spectral loss*. The corresponding measurements are given in table 1, which

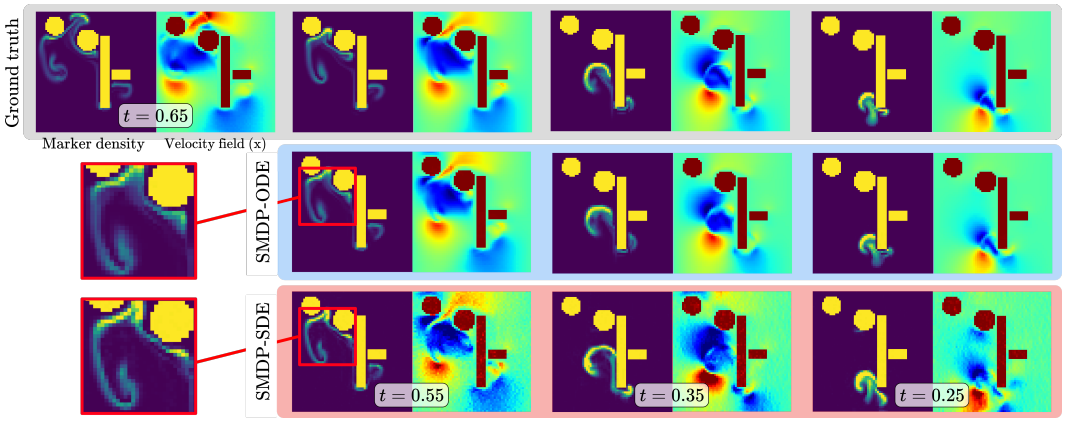

Figure 6: Buoyancy flow case. Ground truth shows the marker density and velocity field in the x-direction at different points of the simulation trajectories. The simulation end state at $t = 0.65$ is the input to SMDP-ODE and SMDP-SDE.

show that our method DSPM-ODE performs best in terms of the reconstruction MSE. However, solutions obtained by DSPM-ODE are very smooth and do not contain the small-scale structures of the references, which is reflected in a high spectral error that is also visually prominent, as shown in figure 4. DSPM-SDE on the other hand performs very well in terms of spectral error and yields visually convincing solutions with only a slight increase in the reconstruction loss. We note that there is a natural tradeoff between both metrics, and SMDP-ODE and SMDP-SDE perform best for both cases respectively while using the same set of weights.

Ablation study We performed an ablation study to highlight several design choices of the proposed method. In particular we note that despite fundamental differences between SMDP-ODE and SMDP-SDE, as explained in section 3, the main difference at inference time is the constant factor for s_θ and the noise term $g(t)dw$ for SMDP-SDE, cf. equations (7) and (8). We investigated how the change in noise integration affects the performance of SMDP-ODE, and considered a variant *SMDP-ODE+noise* that includes the addition of the noise term during inference. As shown with crosses in figure 5, this method has a slightly higher physics-based loss but in contrast to SMDP-SDE does not improve upon the spectral error. This indicates that SMDP-SDE can recover small-scale distributions of the references, while the probability flow variant by construction tends to favour mean solutions when faced with uncertainties. We additionally investigated the effect of the proposed *bidirectional* training scheme. The error measurements of Figure 5 show that ODE variant is not affected, but SMDP-SDE significantly benefits. The resulting model robustly handles a wide range of different temporal discretizations with high accuracy. In conclusion, our SMDP-SDE model with bidirectional training yields the best performance for a wide range of hyperparameter settings. In appendix B, we additionally evaluate the effects of a logarithmic time discretization, different training noise, and a physics-conditioned score field.

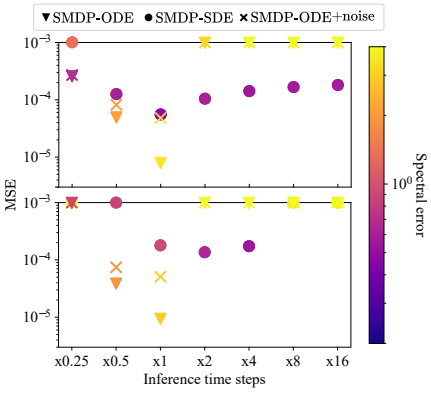

Figure 5: Reconstruction MSE and spectral errors for varying time steps of the bidirectional (top) and regular variant (bottom). Large errors are truncated at the top of each graph.

4.2 Buoyancy-driven Flow with Obstacles

Next, we test our methodology on a more challenging problem. For this purpose we consider simulations of buoyancy-driven flow within a fixed domain $\Omega \subset [0, 1] \times [0, 1]$ and randomly placed obstacles. We make use of semi-Lagrangian advection for the velocity and MacCormack advec-

tion for the hot marker density. The temperature dynamics of the marker field are modeled with a Boussinesq approximation. Each simulation runs from time $t = 0.0$ to $t = 0.65$ with a step size of $\Delta t = 0.01$. The inflow at $(0.2, 0.5)$ is active until $t = 0.2$. Our objective is to employ SMDP-ODE and SMDP-SDE to obtain trajectories that reconstruct a plausible flow given an end state of the marker density and velocity fields at time $t = 0.65$.

Our training data set consists of 250 simulations with corresponding trajectories. For the data generation, we make use of the differentiable *phiflow* solver (Holl et al., 2020a). We place spheres and boxes with varying sizes at different positions within the simulation domain that do not overlap with the marker inflow. For each simulation, we place one to two objects of each category. The testing set comprises 5 simulations. In contrast to the previous task, we generate the training data set without any noise, but add a Gaussian random noise with standard deviation $\sigma = \sqrt{\Delta t}$ to each simulation state of the training trajectories.

Training and Results We employ a neural network architecture based on dilated convolutions (Stachenfeld et al., 2021) for $s_\theta(\mathbf{x}, t)$, see appendix A. The physics operator $\mathcal{P}(\mathbf{x})$ is implemented by using a negative step size $-\Delta t$ for time integration. The training of SPDM-ODE is analogous to the previous experiments on heat equation in section 4.1. However, we do not unroll the entire trajectory and apply a sliding window with size 20. For this case, we compare our method to a fully learned baseline method. A quantitative evaluation in terms of reconstruction MSE and LPIPS metrics (Zhang et al., 2018) is given in table 2. It becomes apparent that our ODE method clearly outperform the learned baseline. Interestingly, the SDE variant performs less well for this test case. This behavior can be explained by the highly nonlinear system dynamics, and the comparatively approximate reverse simulator which yields the substantial errors for the *Physics only* version in table 2. These errors causes the score network to inadvertently learn significant corrections of the physics operator, which deteriorates the quality of the score field. Nonetheless, as qualitatively shown in figure 6, both variants are able to accurately recover the initial states despite the complex motion of the fluid around the obstacles.

Configuration	**MSE**↓		**LPIPS**↓ $[10^{-3}]$	
	ODE	SDE	ODE	SDE
SMDP	**0.05**	0.78	**5.82**	147.17
Rollout noise	0.12	1.03	9.10	145.72
Physics cond. score	0.15	1.25	5.94	360.54
Bidir. training	0.4	0.61	24.19	134.21
Decoupled score&phys.	0.07	0.72	7.78	67.23
Fully learned	0.28	-	12.54	-
Physics only	0.5	-	346.31	-

Table 2: Evaluation of variants for the buoyancy obstacle case in terms of reconstruction MSE and LPIPS of the marker field.

Algorithmic variants We evaluate several altered configurations of our SMDP method to further illustrate its behavior. We consider adding *noise during rollouts*, i.e. adding a noise term $\sigma \cdot z$ for $z \sim \mathcal{N}(0, I)$ to the state $\mathbf{x}$ after applying the score and physics updates during training. Additionally, we experiment with a *physics conditioned score*, i.e. we extend the input dimension of the score function to accept concatenated inputs of the form $s_\theta([\mathbf{x}, \mathcal{P}(\mathbf{x})], t)$. We also evaluate the effects of the bidirectional training. Finally, we evaluate a version that *decouples score and physics*, i.e. we do not evaluate the score and physics update on the same input $\mathbf{x}$, but instead we first apply the physics update $\mathcal{P}(\mathbf{x})$, and evaluate the score function afterwards. Overall, the error measurements in table 2 justify our choices for the baseline SMDP algorithm.

4.3 Isotropic Turbulence

As third example, we consider a problem where the physics operator is unknown, i.e. we approximate both $\mathcal{P}$ and the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ by neural networks. We consider the problem of learning the time evolution of isotropic, forced turbulence as determined by the 2D Navier-Stokes equations with a viscosity of $\nu = 10^{-5}$ (Li et al., 2020). The training data set consists of vorticity fields from 1000 simulation trajectories from $t = 0$ until $T = 10$ with $\Delta t = 1$ and a spatial resolution of 64×64. Our objective is to predict a trajectory $\hat{\mathbf{x}}_{0:10}$ that reconstructs the true trajectory $\hat{\mathbf{x}}_{0:10}$ given an end state $\mathbf{x}_{10}$ of the solution, whereas in the original paper, the objective was to learn an operator mapping predicting the vorticity at a later point in time.

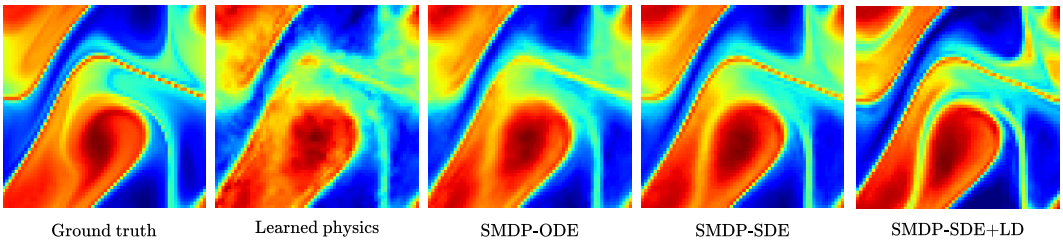

Figure 7: Turbulence case. Comparison of reconstructed trajectories at $t = 1$.

Methods We represent the physics operator $\mathcal{P}(\mathbf{x})$ by a FNO neural network and the score by the ResNet-based architecture from section 4.1. As a first step of our method, we optimize the physics model $\mathcal{P}$ using our SMDP training setup from the previous two experiments, but fix the score to $s_\theta(\mathbf{x}, t) \equiv 0$. We denote this method by *Learned physics*. As a second step, we train SMDP-ODE with the *decoupled score and physics variant* by optimizing the time-dependent score $s_\theta(\mathbf{x}, t)$ while freezing the physics model $\mathcal{P}$. This approach guarantees that any time-independent physics are captured by $\mathcal{P}$ and $s_\theta(\mathbf{x}, t)$ can focus on learning small improvements to $\mathcal{P}$ as well as respond to possibly time-dependent data biases. As this test employs two trained networks for both components of the SDE, we compare to *DiffFlow* (Zhang & Chen, 2021) as an additional baseline.

Evaluation We give an evaluation of the improvements of SMDP over the learned variants in table 3. Compared to the Learned-physics variant, our methods improve the mean squared error (MSE) between the ground truth trajectories and the reconstructed trajectories slightly, while there is an substantial decrease in the spectral error. This can be seen qualitatively in figure 7. In this scenario DiffFlow has severe difficulties to learn state updates and score field, resulting in large differences in terms of MSE. As before, the SMDP-SDE method performs best in terms of spectral error at the expense of a slightly increased MSE.

Method	**MSE**↓ $[10^{-2}]$	**Spectral Error**↓
SMDP-ODE	**14.3**	0.26
SMDP-SDE	15.3	**0.17**
DiffFlow	221.2	0.58
Learned physics	16.3	0.30

Table 3: Evaluation of the turbulence case.

Outlook: Refinement with Langevin dynamics Since the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ represents a data gradient, we can use gradient-based optimization methods to find local optima of the data distribution $p_t(\mathbf{x})$ that are close to $\mathbf{x}$. Inspired by stochastic gradient Langevin dynamics (Welling & Teh, 2011), we consider the iteration rule $\mathbf{x}_t^{i+1} = \mathbf{x}_t^i + \epsilon \cdot \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t^i) + \sqrt{2\epsilon} z_t$, for $\epsilon = 2 \times 10^{-5}$, where $z_t \sim \mathcal{N}(0, I)$ (details in appendix D.1). Denoted by *SMDP-SDE+LD* in figure 7, this method manages to extract even finer details from the reverse-time SDE solution. As such it provides an interesting starting point for a further refinement of the SMDP results.

5 Conclusion

We presented SMDP, a derivative of score matching in the context of physical simulations and differentiable physics. We demonstrated its competitiveness against different baseline methods and in challenging inverse physics problems. We demonstrated the versatility of SMDP with two variants: while the *neural ODE* variant focuses on high MSE accuracies, the *neural SDE* variant allows for sampling the posterior and yields an improved coverage of the target data manifold. Despite the promising initial results, our work gives rise to many interesting questions. In particular, the time discretizations is a crucial issue, as for training data generation and differentiable solvers, we would favor larger step size and less evaluations due to computational constraints, while for conventional diffusion models, a large number of smaller time steps often yields substantial improvements in terms of quality. Determining a good balance between accurate solutions with few time steps and diverse solutions with many time steps will remain an important direction for future research in this area.

REFERENCES

Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

Xavier Blanc, Claude Le Bris, and P-L Lions. From molecular models to continuum mechanics. *Archive for Rational Mechanics and Analysis*, 164(4):341–381, 2002.

Johannes Brandstetter, Daniel Worrall, and Max Welling. Message passing neural pde solvers. *International Conference on Learning Representations*, 2022.

Nanxin Chen, Yu Zhang, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=NsMLjcFaO8O.

Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 6572–6583, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html.

Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. *CoRR*, abs/2112.05146, 2021. URL https://arxiv.org/abs/2112.05146.

Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *CoRR*, abs/2206.00941, 2022. doi: 10.48550/arXiv.2206.00941. URL https://doi.org/10.48550/arXiv.2206.00941.

S Delaquis, MJ Jewell, I Ostrovskiy, M Weber, T Ziegler, J Dalmasson, LJ Kaufman, T Richards, JB Albert, G Anton, et al. Deep neural networks for energy and position reconstruction in exo-200. *Journal of Instrumentation*, 13(08):P08023, 2018.

Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 8780–8794, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.

Stanley J Farlow. *Partial differential equations for scientists and engineers*. Courier Corporation, 1993.

Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.

Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD: free-form continuous dynamics for scalable reversible generative models. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=rJxgknCcK7.

Morton E Gurtin. *An introduction to continuum mechanics*. Academic press, 1982.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. *CoRR*, abs/2204.03458, 2022. doi: 10.48550/arXiv.2204.03458. URL https://doi.org/10.48550/arXiv.2204.03458.

Philipp Holl, Vladlen Koltun, Kiwon Um, and Nils Thuerey. phiflow: A differentiable pde solving framework for deep learning via physical simulations. In *NeurIPS Workshop*, volume 2, 2020a.

Philipp Holl, Nils Thuerey, and Vladlen Koltun. Learning to control pdes with differentiable physics. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020b. URL https://openreview.net/forum?id=HyeSin4FPB.

Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *CoRR*, abs/2206.07696, 2022. doi: 10.48550/arXiv.2206.07696. URL https://doi.org/10.48550/arXiv.2206.07696.

Alexia Jolicoeur-Martineau, Ke Li, Rémi Piché-Taillefer, Tal Kachman, and Ioannis Mitliagkas. Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*, 2021.

Bahjat Kawar, Gregory Vaksman, and Michael Elad. SNIPS: solving noisy inverse problems stochastically. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 21757–21769, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/b5c01503041b70d41d80e3dbe31bbd8c-Abstract.html.

Dmitrii Kochkov, Jamie A. Smith, Ayya Alieva, Qing Wang, Michael P. Brenner, and Stephan Hoyer. Machine learning accelerated computational fluid dynamics. *CoRR*, abs/2102.01010, 2021. URL https://arxiv.org/abs/2102.01010.

Sizhe Li, Zhiao Huang, Tao Du, Hao Su, Joshua B Tenenbaum, and Chuang Gan. Contact points discovery for soft-body manipulations with differentiable physics. *International Conference on Learning Representations*, 2022.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv:2010.08895*, 2020.

Joowon Lim and Demetri Psaltis. Maxwellnet: Physics-driven deep neural network training based on maxwell's equations. *Apl Photonics*, 7(1):011301, 2022.

Cheng Lu, Kaiwen Zheng, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Maximum likelihood training for score-based diffusion odes by high order denoising score matching. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14429–14460. PMLR, 2022. URL https://proceedings.mlr.press/v162/lu22f.html.

Dimitra Maoutsa, Sebastian Reich, and Manfred Opper. Interacting particle solutions of fokker-planck equations through gradient-log-density estimation. *Entropy*, 22(8):802, 2020. doi: 10.3390/e22080802. URL https://doi.org/10.3390/e22080802.

Glenn J Martyna, Mark E Tuckerman, Douglas J Tobias, and Michael L Klein. Explicit reversible integrators for extended systems dynamics. *Molecular Physics*, 87(5):1117–1157, 1996.

Jeremy Morton, Antony Jameson, Mykel J Kochenderfer, and Freddie Witherden. Deep dynamical modeling and control of unsteady fluid flows. In *Advances in Neural Information Processing Systems*, 2018.

Maximilian Mueller, Robin Greif, Frank Jenko, and Nils Thuerey. Leveraging stochastic predictions of bayesian neural networks for fluid simulations. *arXiv preprint arXiv:2205.01222*, 2022.

Bernt Øksendal. Stochastic differential equations. In *Stochastic differential equations*, pp. 65–84. Springer, 2003.

Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. *International Conference on Learning Representations*, 2020.

Zaccharie Ramzi, Benjamin Remy, François Lanusse, Jean-Luc Starck, and Philippe Ciuciu. Denoising score-matching for uncertainty quantification in inverse problems. *CoRR*, abs/2011.08698, 2020. URL https://arxiv.org/abs/2011.08698.

Severi Rissanen, Markus Heinonen, and Arno Solin. Generative modelling with inverse heat dissipation. *arXiv:2206.13397*, 2022.

Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.

Samuel Schoenholz and Ekin Dogus Cubuk. Jax md: a framework for differentiable physics. *Advances in Neural Information Processing Systems*, 33:11428–11441, 2020.

Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 2256–2265. JMLR.org, 2015. URL http://proceedings.mlr.press/v37/sohl-dickstein15.html.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11895–11907, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/3001ef257407d5a371a96dcd947c7d93-Abstract.html.

Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *CoRR*, abs/2111.08005, 2021a. URL https://arxiv.org/abs/2111.08005.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

Kim Stachenfeld, Drummond Buschman Fielding, Dmitrii Kochkov, Miles Cranmer, Tobias Pfaff, Jonathan Godwin, Can Cui, Shirley Ho, Peter Battaglia, and Alvaro Sanchez-Gonzalez. Learned simulators for turbulence. In *International Conference on Learning Representations*, 2021.

Jonathan Tompson, Kristofer Schlachter, Pablo Sprechmann, and Ken Perlin. Accelerating eulerian fluid simulation with convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=ByH2gxrKl.

Kiwon Um, Robert Brand, Yun (Raymond) Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative pde-solvers. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/43e4e6a6f341e00671e123714de019a8-Abstract.html.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Comput.*, 23(7):1661–1674, 2011. doi: 10.1162/NECO_a_00142. URL https://doi.org/10.1162/NECO_a_00142.

Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient langevin dynamics. In Lise Getoor and Tobias Scheffer (eds.), *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pp. 681–688. Omnipress, 2011. URL https://icml.cc/2011/papers/398_icmlpaper.pdf.

Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *CoRR*, abs/2203.09481, 2022. doi: 10.48550/arXiv.2203.09481. URL https://doi.org/10.48550/arXiv.2203.09481.

Qinsheng Zhang and Yongxin Chen. Diffusion normalizing flow. *Advances in Neural Information Processing Systems*, 34:16280–16291, 2021.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Kemin Zhou, JC Doyle, and Keither Glover. Robust and optimal control. *Control Engineering Practice*, 4(8):1189–1190, 1996.

APPENDIX

A ARCHITECTURES

ResNet We employ a simple ResNet-like architecture, which is used in Section 4.1 for the score function $s_\theta(\mathbf{x}, t)$ and the convolutional neural network baseline (ResNet-S and ResNet-P) as well as in Section 4.3 again for the score $s_\theta(\mathbf{x}, t)$.

Since in both experiments, there are periodic boundary conditions, we apply a periodic padding with length 16, i.e. if the underlying 2-dimensional data dimensions are $N \times N$, the dimenions after the periodic padding are $(N+16) \times (N+16)$. We implement the periodic padding by first tiling the input 3 times in x- and y-direction and then cropping to the correct sizes. The time t is concatenated as an additional constant channel to the 2-dimensional input data when this architecture is used to represent the score $s_\theta(\mathbf{x}, t)$.

The encoder-part of our network begins with a single 2D-convolution encoding layer with 32 filters, kernel size 4 and no activation function. This is followed by 4 consecutive residual blocks, each consisting of 2D-convolution, LeakyReLU, 2D-convolution and Leaky ReLU. All 2D convolutions have 32 filters with kernel size 4 and stride 1. The encoder part ends with a single 2D convolution with 1 filter, kernel size 1 and no activation. Then, in the decoder part, we begin with a transposed 2D convolution, 32 filters, kernel size 4. Afterwards, there are 4 consecutive residual blocks, analogous to the encoder residual blocks, but with the 2D convolution replaced by a transposed 2D convolution. Finally there is a final 2D convolution with 1 filter and kernel size 5. Parameters statistics of this model, as well as the others are given in table 4.

UNet We use the UNet architecture with spatial dropout as given in (Mueller et al., 2022), Appendix A.1. The dropout rate is set to 0.25. We do not include batch normalization and apply the same periodic padding as done for our ResNet architecture.

FNO For all experiments, we consider the FNO-2D architecture introduced in (Li et al., 2020) with $k_{\max,j} = 12$ Fourier modes per channel.

Dil-ResNet The Dil-ResNet architecture is described in (Stachenfeld et al., 2021), Appendix A. Since this architecture represents the score s_θ in Section 4.2, we concatenate the constant time channel analogously to the ResNet architecture. Additionally, positional information is added to the network input by encoding the x-position and y-position inside the domain in two separate channels.

Architecture	Parameters
ResNet	330.754
UNet	706.035
DilatedConv	336.915
FNO	465.377

Table 4: Summary of architectures.

B HEAT EQUATION

B.1 ADDITIONAL TRAINING DETAILS

Spectral loss We consider a spectral error based on the two-dimensional power spectral density. For two 2d-images, we compute their 2d Fourier transform and compute the radially averaged power spectrum s_1 and s_2. Then, we define the spectral error as the difference between the log of the spectral densities

$$\mathcal{L}(s_1, s_2) = |\log(s_1) - \log(s_2)|. \tag{9}$$

SMDP-ODE For inference, we consider the linear time discretization $t_n = n\Delta t$ with $\Delta t = 0.2/32$ and $t_{32} = 0.2$. During training, we sample a random time discretization $\tilde{t}_n$ for each batch based on t_n by $\tilde{t}_n \sim \mathcal{U}(t_n - \Delta t/2, t_n + \Delta t/2)$ for $n = 1, ..., 31$ to not overfit on the step size. In the warmup phase of training, we unroll Algorithm 1 for $N = 6, 8, ..., 32$ steps, where we increase N every 2 epochs. We employ Adam to update the weights θ with learning rate 10^{-4}. After the warmup is finished, we finetune the network weights for 80 epochs with an initial learning rate of 10^{-4} which we reduce by a factor of 0.5 every 20 epochs.

Diffusion model The diffusion model is inspired by (Rissanen et al., 2022). We make use of the same ResNet architecture as SMDP-ODE, however we keep the time discretization t_n fixed. The diffusion model does not include the physics operator $\mathcal{P}$ in the rollout. Therefore, it has to learn the score and physics at the same time. Except for those changes, the training setup is identical to SMDP-ODE.

Baseline methods All other baseline methods are trained for 80 epochs using the Adam optimizer algorithm with an initial learning rate of 10^{-4} which is decreased by a factor of 0.5 every 20 epochs. For the training data, we consider solutions to the heat equation consisting of initial state $\mathbf{x}_0$ and end state $\mathbf{x}_T$ that are noise-free and add a Gaussian noise with standard deviation $\sigma = 0.1$ to the network input.

B.2 LOGARITHMIC TIME DISCRETIZATION

For this experiment, we consider a logarithmic time discretization during training and inference. This discretization is finer around $t = 0$, where most of the small-scale structures are smoothened out. We find that this method also works well, but the linear time discretization performs better for most cases, especially when Δt is changed for inference, cf. figure 5 and figure 8.

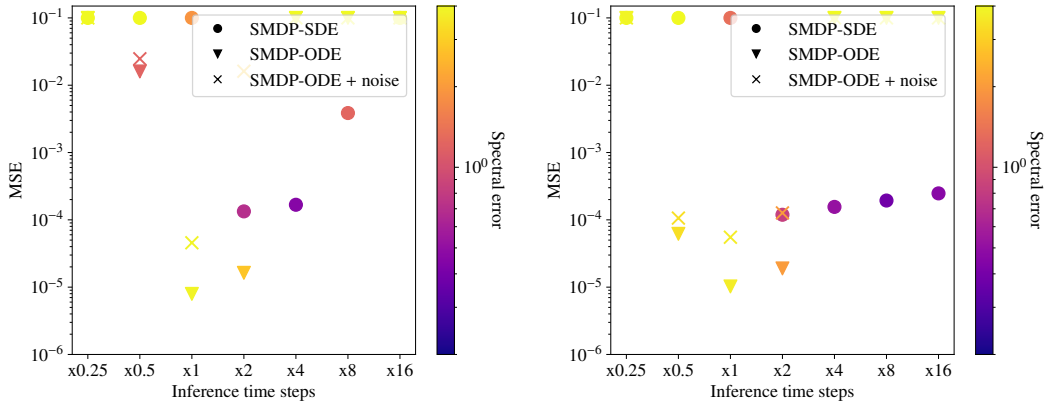

(a) Regular training

(b) Bidirectional training

Figure 8: Logarithmic time discretization. Reconstruction MSE and spectral errors for varying time steps of regular variant (a) and bidirectional variant (b). Large errors are truncated at the top of each graph.

B.3 PHYSICS-CONDITIONED SCORE

We extend the definition of the score function, to also include information about the physics at time t, i.e. $\mathcal{P}(\mathbf{x}_t)$. Thus, we replace $s_\theta(\mathbf{x}, t)$ in algorithm 1 by $s_\theta([\mathbf{x}, \mathcal{P}(\mathbf{x})], t)$, where we concatenate both inputs. An evaluation is shown in figure 9. Although the stability of SMDP-ODE and SMDP-ODE + noise is greatly increased, we do not obtain results with a low spectral error for SMDP-SDE. Interestingly, the bidirectional training also seems harmful in this case.

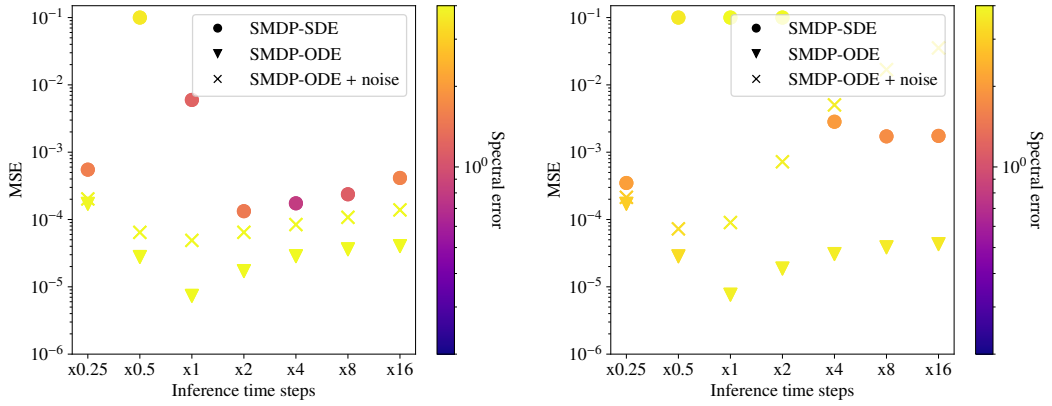

(a) Regular training **(b)** Bidirectional training

Figure 9: Physics-conditioned score. Reconstruction MSE and spectral errors for varying time steps of regular variant (a) and bidirectional variant (b). Large errors are truncated at the top of each graph.

B.4 HIGHER TRAINING NOISE

We increase the diffusion coefficient from $g \equiv 0.1$ to $g \equiv 1.0$ during training, but for inference, we set it back to the lower value. An evaluation is shown in figure 10. Overall, SMDP-SDE produces results with low spectral error and SMDP-ODE has a much lower MSE, similar to the standard case. However, overall, the MSE is much higher, indicating shortcomings of this approach.

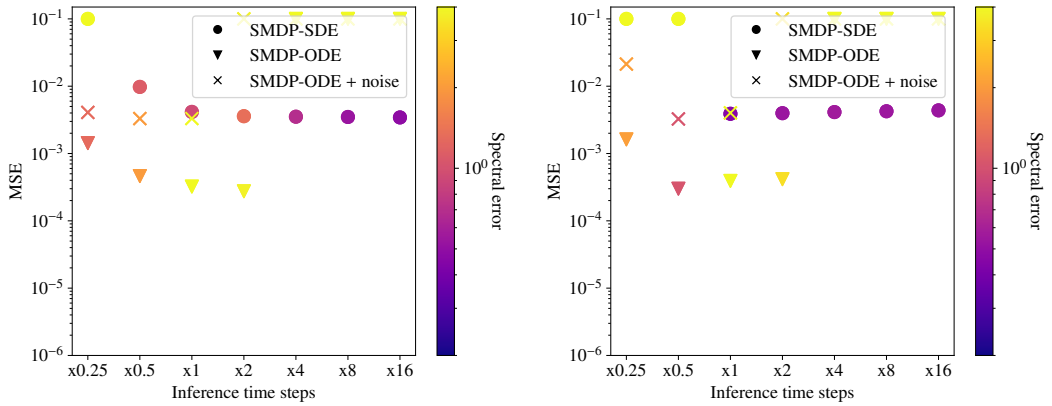

(a) Regular training **(b)** Bidirectional training

Figure 10: Higher training noise. Reconstruction MSE and spectral errors for varying time steps of regular variant (a) and bidirectional variant (b). Large errors are truncated at the top of each graph.

B.5 ADDITIONAL RESULTS

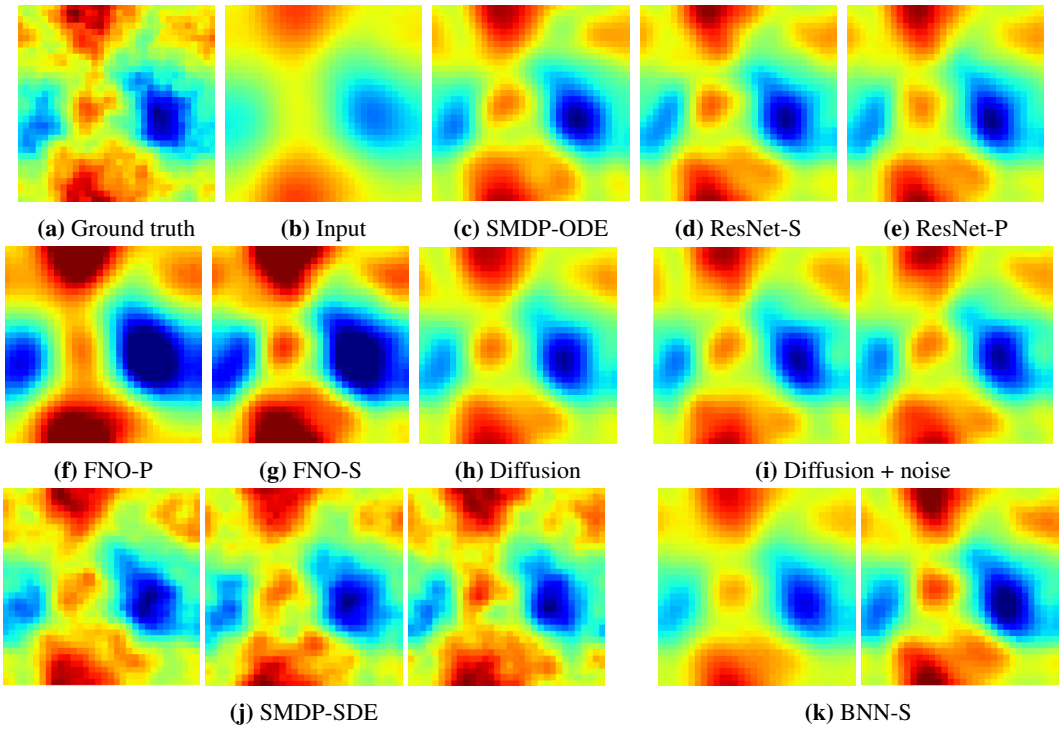

Figure 11: Predictions of different methods for the heat equation problem (example 1 of 2).

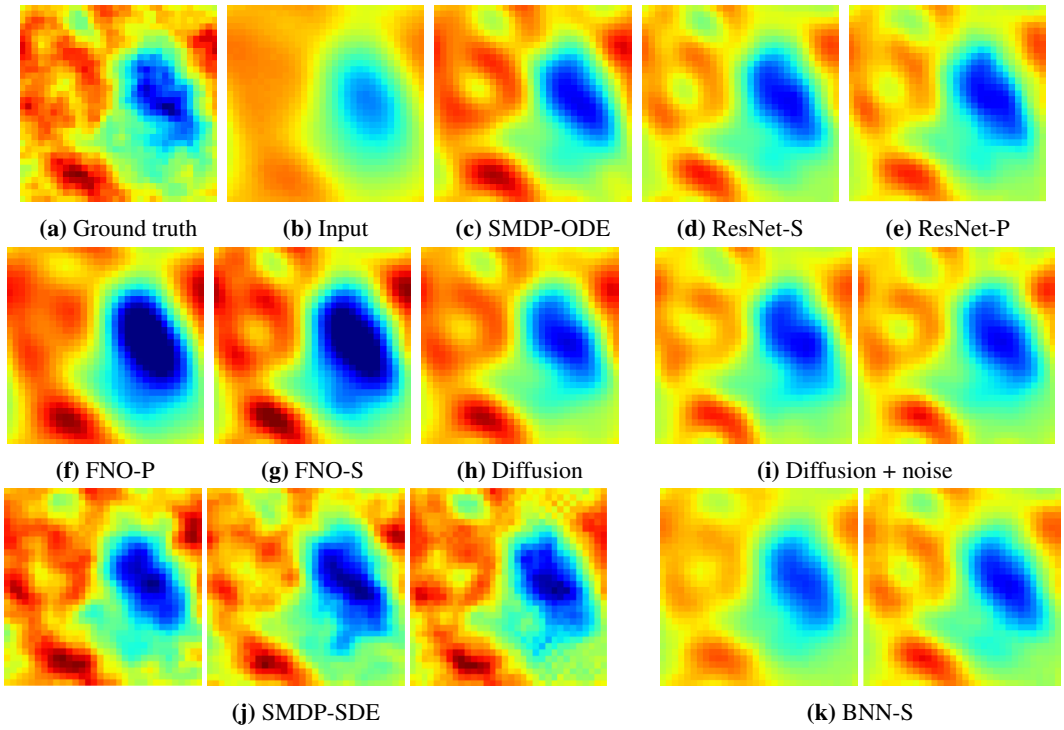

Figure 12: Predictions of different methods for the heat equation problem (example 2 of 2). Neither the BNN nor the Diffusion+noise model are able to produce small-scale structures.

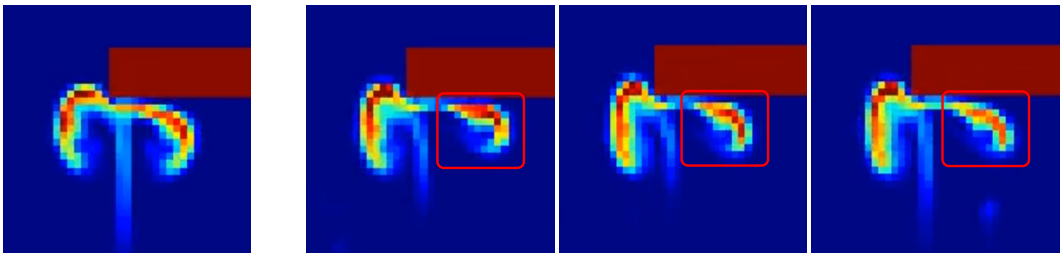

Ground truth SMDP-SDE

Figure 13: Comparison of SMDP-SDE predictions and ground truth for buoyancy-driven flow at $t = 0.36$.

C BUOYANCY-DRIVEN FLOW WITH OBSTACLES

C.1 TRAINING

We train all networks with Adam and learning rate 1×10^{-4} with batch size 16. We begin with unrolling $N = 2$ steps, which we increase every 30 epochs by 2 until we reach $N = 20$.

C.2 ADDITIONAL RESULTS

We give more detailed time evolutions of results for the buoyancy-driven flow case in figure 14 and figure 15. These again highlight the difficulties of the physics simulator to recover the initial states by itself. The SMDP variants significantly improve upon this behavior.

In figure 13 we also show an example of the posterior sampling for this case. It becomes apparent that the inferred small-scale structures of the different samples change. However, in contrast to cases like the heat diffusion example, the physics simulation in this scenario leaves only little room for substantial changes of the states.

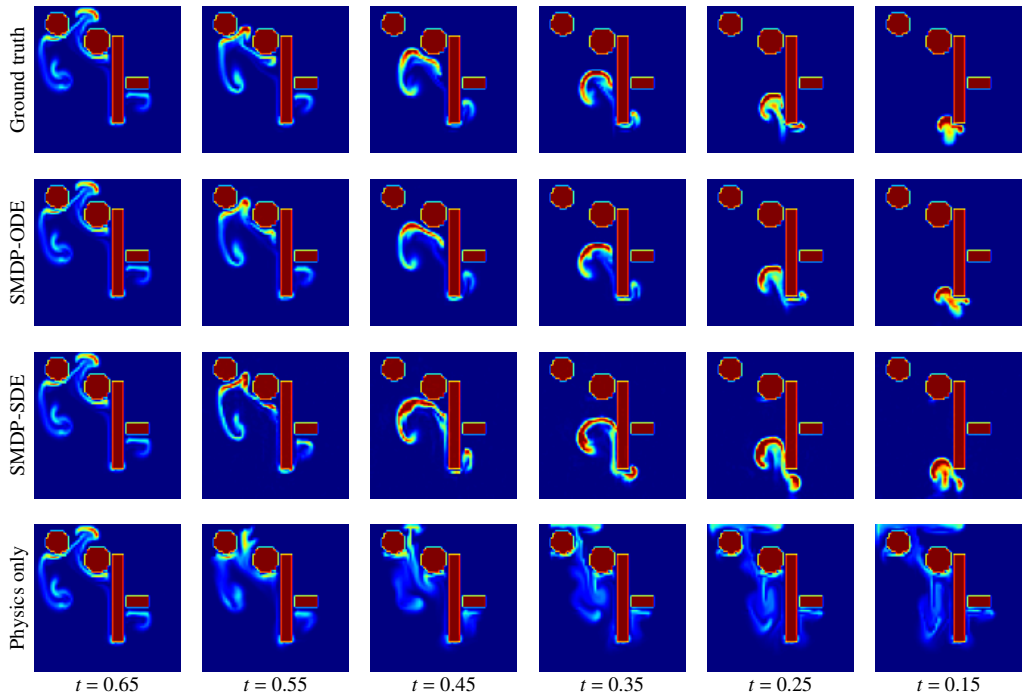

(a) Marker field

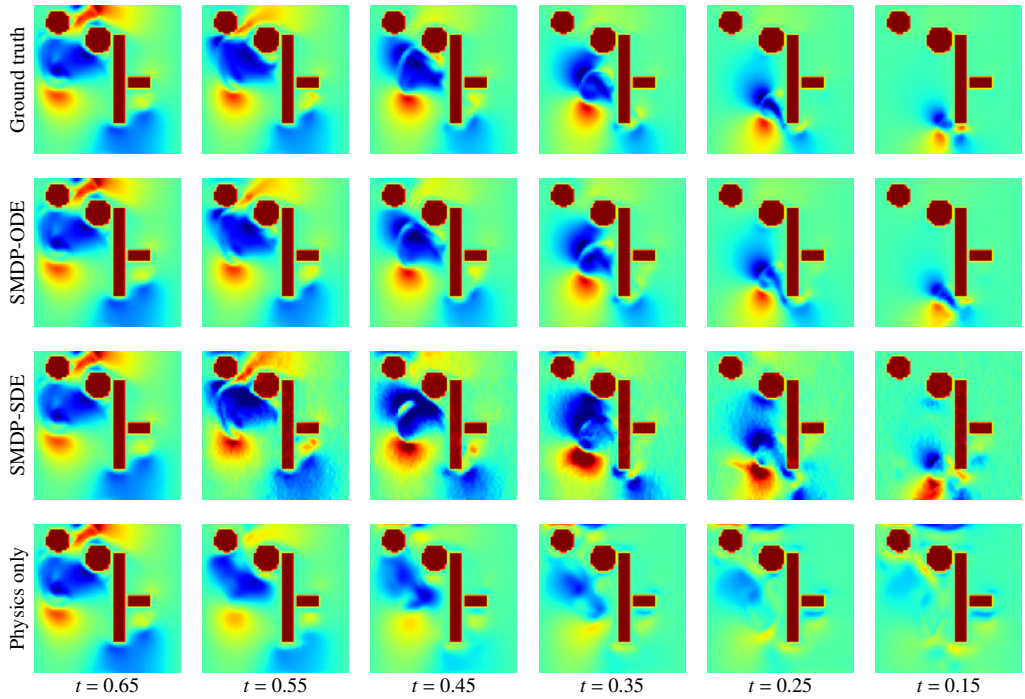

(b) Velocity field (x-direction)

Figure 14: Predictions for buoyancy-driven flow with obstacles (example 1 of 2).

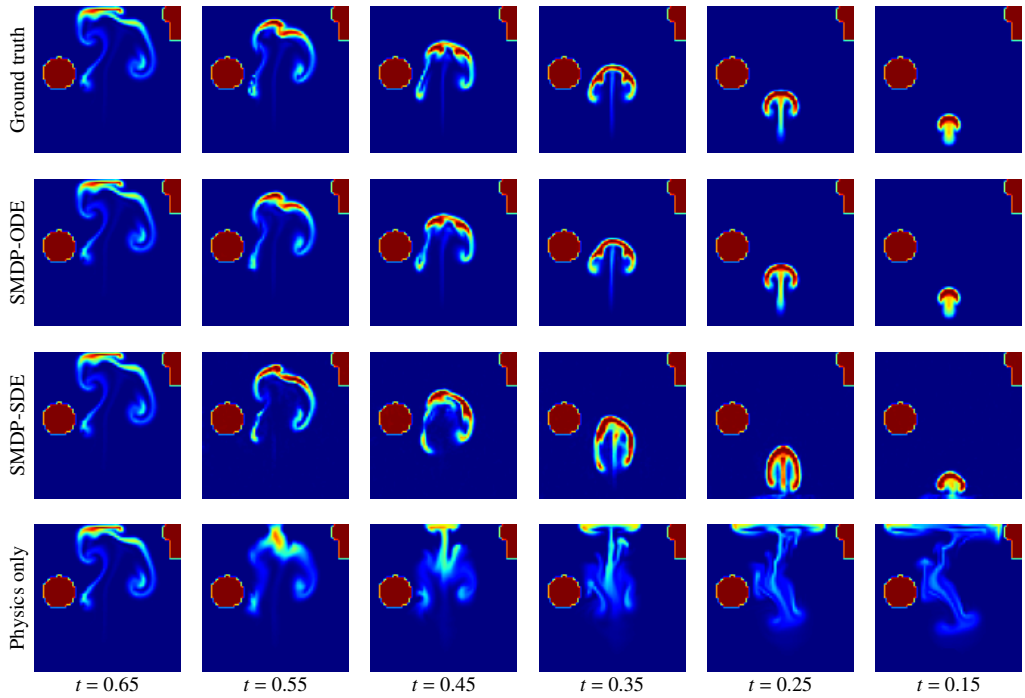

(a) Marker field

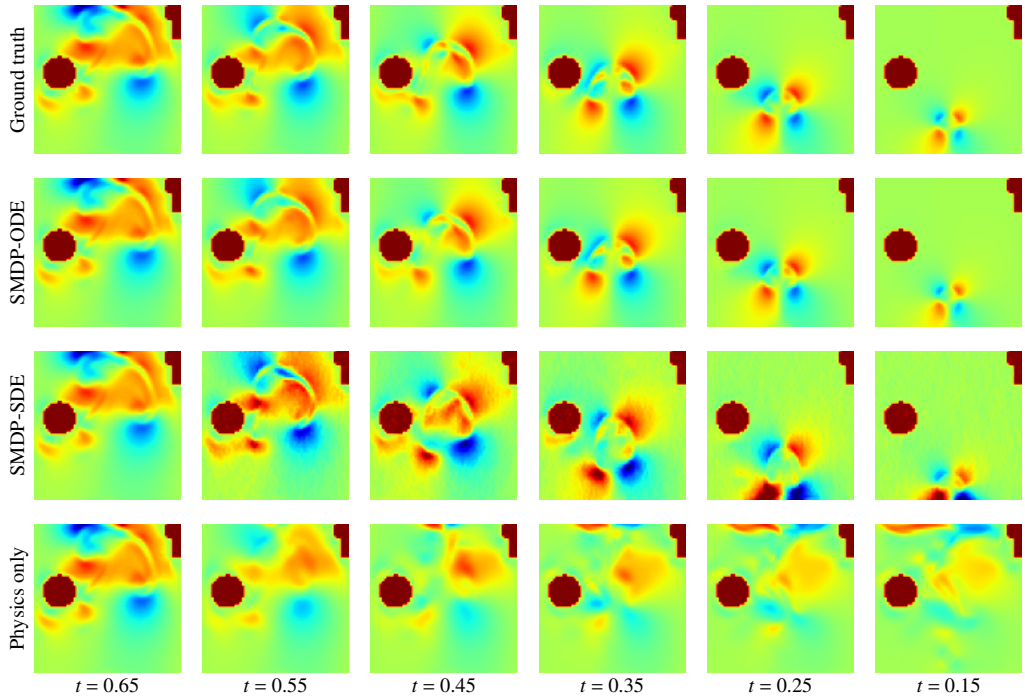

(b) Velocity field (x-direction)

Figure 15: Predictions for buoyancy-driven flow with obstacles (example 2 of 2).

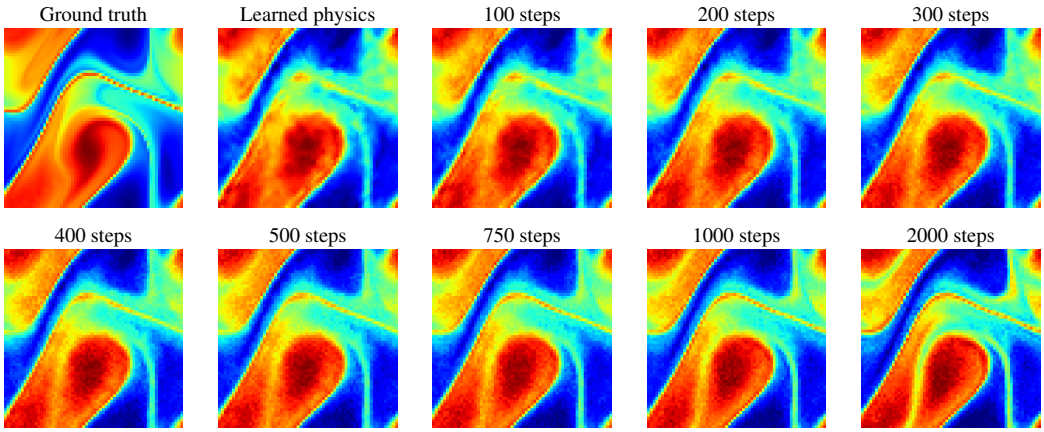

Figure 16: Steps of Langevin dynamics for $\epsilon = 2 \times 10^{-5}$.

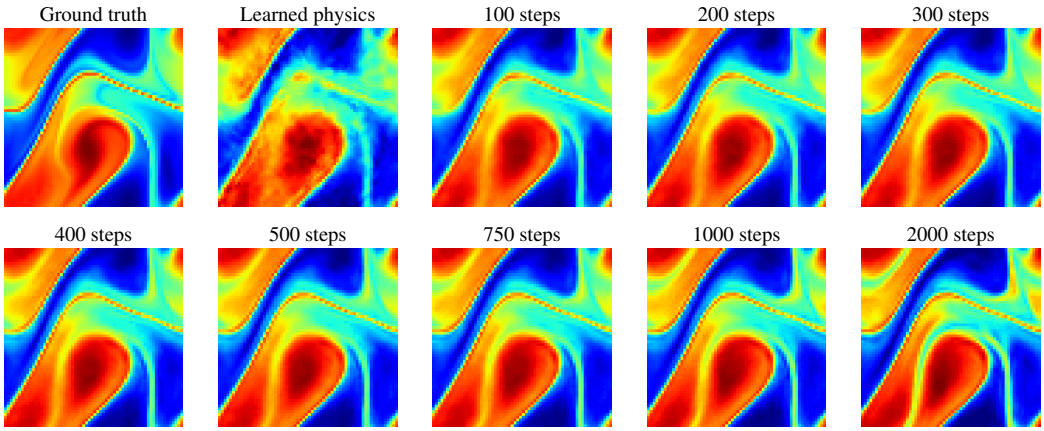

Figure 17: Steps with Langevin dynamics for $\epsilon = 2 \times 10^{-5}$ and an additional step with $\Delta t \cdot s_\theta(\mathbf{x}, t)$ which smoothes the images.

D ISOTROPIC TURBULENCE

For the learned physics network, we employ a FNO neural network with batch size 20. We train the FNO for 500 epochs using Adam optimizer with learning rate 10^{-3}, which we decrease every 100 epochs by a factor of 0.5. We train SMDP-ODE with the ResNet architecture for 250 epochs with learning rate 10^{-4}, decreased every 50 epochs by a factor of 0.5 and batch size 6.

D.1 REFINEMENT WITH LANGEVIN DYNAMICS

We do a fixed point iteration at a single point in time via:

$$\mathbf{x}_t^{i+1} = \mathbf{x}_t^i + \epsilon \cdot \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t^i) + \sqrt{2\epsilon} z_t, \tag{10}$$

for a number of steps T and $\epsilon = 2 \times 10^{-5}$ as a post-processing and refinement strategy, cf. figure 16 and figure 17. This is motivated by established methods in score-based generative modelling (Welling & Teh, 2011; Song & Ermon, 2019). For a prior distribution π_t, $\mathbf{x}_t^0 \sim \pi_t$ and by iterating equation (10), the distribution of $\mathbf{x}_t^T$ equals p_t for $\epsilon \to 0$ and $T \to \infty$. There are some theoretical caveats, i.e. an additional Metropolis-Hastings update needs to be added in equation (10) and there are regularity conditions (Song & Ermon, 2019).

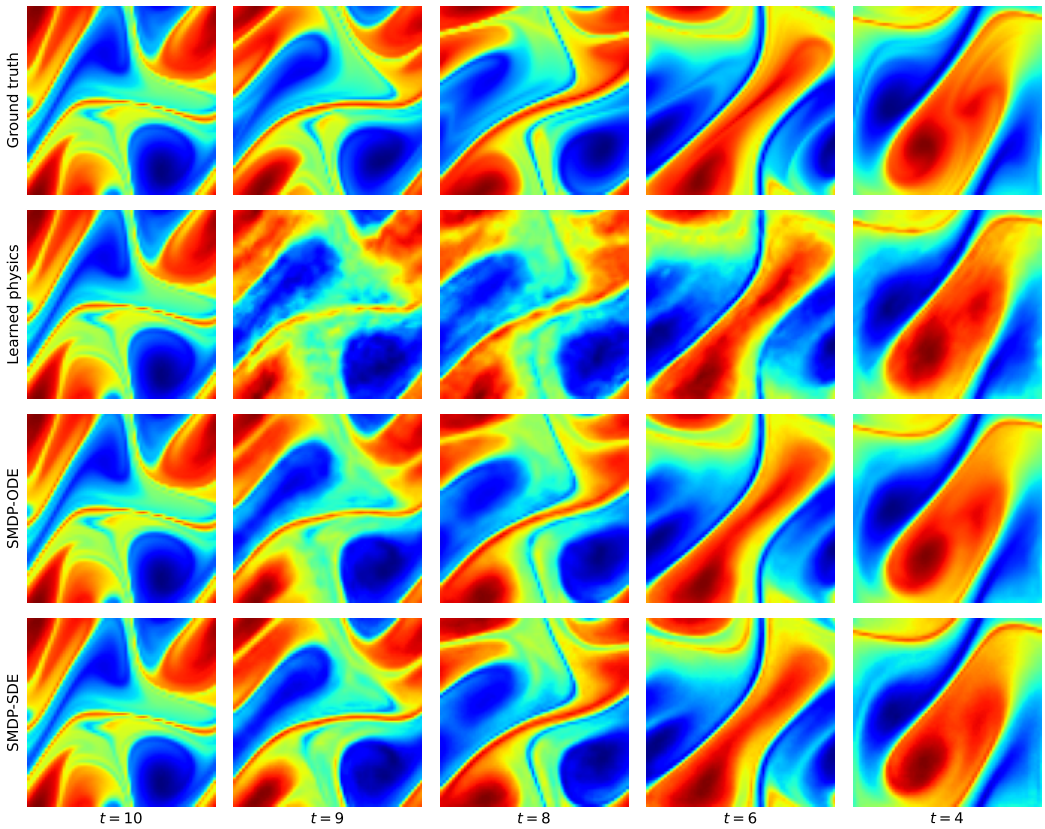

Figure 18: Predictions for isotropic turbulence (example 1 of 2).

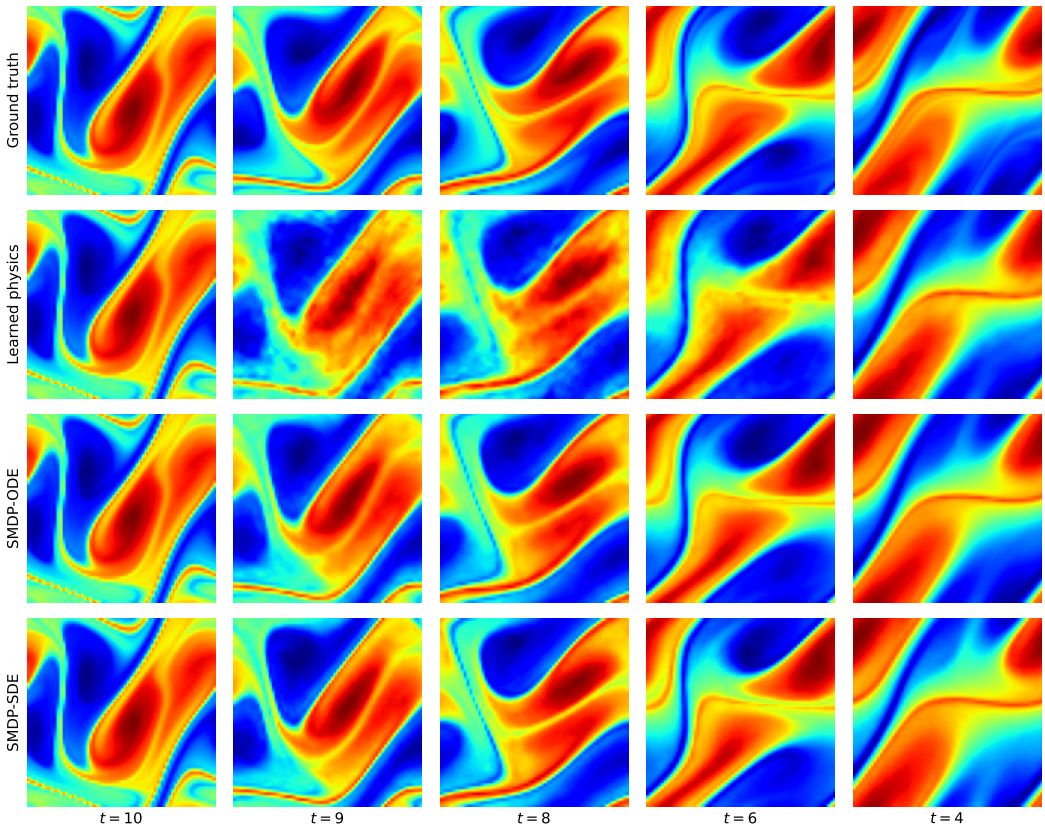

Figure 19: Predictions for isotropic turbulence (example 2 of 2).