

# State-observation augmented diffusion model for nonlinear assimilation with unknown dynamics

Zhuoyuan Li<sup>a</sup>, Bin Dong<sup>b,c,\*</sup>, Pingwen Zhang<sup>d,a,\*</sup>

<sup>a</sup>*School of Mathematical Sciences, Peking University, Beijing, 100871, Beijing, China*

<sup>b</sup>*Beijing International Center for Mathematical Research and the New Cornerstone Science Laboratory, Peking University, Beijing, 100871, Beijing, China*

<sup>c</sup>*Center for Machine Learning Research, Peking University, Beijing, 100871, Beijing, China*

<sup>d</sup>*School of Mathematics and Statistics, Wuhan University, Wuhan, 430072, Hubei, China*

---

## Abstract

Data assimilation has become a key technique for combining physical models with observational data to estimate state variables. However, classical assimilation algorithms often struggle with the high nonlinearity present in both physical and observational models. To address this challenge, a novel generative model, termed the State-Observation Augmented Diffusion (SOAD) model is proposed for data-driven assimilation. The marginal posterior associated with SOAD has been derived and then proved to match the true posterior distribution under mild assumptions, suggesting its theoretical advantages over previous score-based approaches. Experimental results also indicate that SOAD may offer improved performance compared to existing data-driven methods.

---

## 1. Introduction

The advent of modern observational facilities and devices has led to an exponential increase in the volume of data gathered from the physical world. Consequently, data assimilation (DA), as an uncertainty quantification technique fusing observations with physical models, has attracted significant attention. Over the past few decades, the DA community has developed a

---

\*Co-corresponding authors.

*Email addresses:* zy.li@alumni.pku.edu.cn (Zhuoyuan Li),  
dongbin@math.pku.edu.cn (Bin Dong), pzhang@pku.edu.cn (Pingwen Zhang)

series of methods to tackle the assimilation problem under various assumptions and constraints. From early approaches such as the nudging method and optimal interpolation to modern techniques including the Kalman filter (KF) [1] and its variants [2, 3], variational methods, and their hybrid versions [4, 5], researchers have demonstrated both theoretical and practical effectiveness. Nonetheless, the high dimensionality and intrinsic nonlinearity of physical and observational models continue to pose substantial challenges.

In this work, we follow the four-dimensional variational data assimilation (4D-Var) formulation, which aims to estimate the posterior distribution of state variables over a fixed temporal window given a sequence of observations. Mathematically, consider a dynamical system described by

$$\begin{cases} \mathbf{x}_k = \mathcal{M}_k(\mathbf{x}_{k-1}, \boldsymbol{\eta}_k), & \boldsymbol{\eta}_k \sim \mathcal{D}_k^X, \\ \mathbf{y}_k = \mathcal{H}_k(\mathbf{x}_k) + \boldsymbol{\epsilon}_k, & \boldsymbol{\epsilon}_k \sim \mathcal{D}_k^Y, \end{cases} \quad (1)$$

where  $\mathbf{x}_k$  and  $\mathbf{y}_k$  denote the state and observation variables at time step  $k$ , respectively. The observation noise  $\boldsymbol{\epsilon}_k \sim \mathcal{D}_k^Y$  captures uncertainties in the observational data, while the term  $\boldsymbol{\eta}_k \sim \mathcal{D}_k^X$  accounts for possible model errors in the physical model  $\mathcal{M}_k$ . The goal is to estimate the posterior distribution  $p(\mathbf{x}_{S:T} \mid \mathbf{y}_{S:T})$  for the state variables over the interval from  $S$  to  $T$ , i.e., we use the abbreviation  $\mathbf{x}_{S:T}$  for  $\{\mathbf{x}_k\}_{k=S}^T$ . By Bayes' rule, the posterior of the states given the observations can be expressed as

$$p(\mathbf{x}_{S:T} \mid \mathbf{y}_{S:T}) = \frac{p(\mathbf{y}_{S:T} \mid \mathbf{x}_{S:T})p(\mathbf{x}_{S:T})}{p(\mathbf{y}_{S:T})} \propto p(\mathbf{x}_{S:T})p(\mathbf{y}_{S:T} \mid \mathbf{x}_{S:T}). \quad (2)$$

Classical variational methods [6, 7] further expand eq. (2) under independence assumptions for the observational and model errors, yielding

$$p(\mathbf{x}_{S:T} \mid \mathbf{y}_{S:T}) \propto p(\mathbf{x}_S) \prod_{k=S}^{T-1} p(\mathbf{x}_{k+1} \mid \mathbf{x}_k) \prod_{k=S}^T p(\mathbf{y}_k \mid \mathbf{x}_k) \quad (3)$$

A common approach then involves *maximum a posteriori* (MAP) estimation under Gaussian assumptions and solves

$$\mathbf{x}_{S:T}^* = \underset{\mathbf{x}_{S:T}}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\mathbf{x}_S - \mathbf{x}_S^b\|_B^2 + \frac{1}{2} \sum_{k=S}^{T-1} \|\mathbf{x}_{k+1} - \mathcal{M}_k^X(\mathbf{x}_k)\|_{Q_k}^2 + \frac{1}{2} \sum_{k=S}^T \|\mathbf{y}_k - \mathcal{H}_k(\mathbf{x}_k)\|_{R_k}^2 \right\} \quad (4)$$

by the variational principle, where

$$\mathbf{x}_S \sim \mathcal{N}(\mathbf{x}_S^b, \mathbf{B}), \quad \boldsymbol{\eta}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k), \quad \boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k), \quad (5)$$

and the model error  $\boldsymbol{\eta}_k$  is assumed additive. Here, we use the notation  $\|\mathbf{a}\|_{\mathbf{A}}^2 := \mathbf{a}^\top \mathbf{A}^{-1} \mathbf{a}$  for simplicity.

Despite the success of classical variational data assimilation methods, they still face several fundamental limitations, primarily in their treatment of uncertainty [8, 7]. The background-error covariance is typically prescribed and assumed to follow a static Gaussian distribution, which does not capture the evolving, nonlinear, and flow-dependent nature of real atmospheric uncertainties. This mismatch can lead to suboptimal weighting of observations and errors in state estimation. Additionally, the iterative minimization process requires repeated evaluations of both the forward physical model and its adjoint, making large-scale applications computationally expensive [9]. Moreover, deriving and maintaining tangent linear models and adjoint models adds further complexity, restricting the flexibility and applicability of variational methods. These challenges highlight the need for alternative approaches that can better represent uncertainty and efficiently handle nonlinear, high-dimensional systems.

As noted by [10, 11], the rapid influx of observational data, such as satellite observations and in-situ scientific measurements, has encouraged researchers to explore data-driven approaches. Machine learning, and in particular deep learning (DL), has shown its remarkable capability in learning nonlinear correlations and extracting high-dimensional features from training data. Inspired by the success of DL in fields such as computer vision (CV) [12, 13] and natural language processing (NLP) [14, 15, 16], scientists are increasingly applying DL techniques within the earth and atmospheric sciences community. Recent success in training large foundation models for weather forecasting [17, 18, 19, 20, 21], have demonstrated the potential of data-driven models in handling complex global atmospheric physics. These advancements indicate a shift towards data-driven methodologies, which may enhance our understanding and predictive capabilities in atmospheric sciences.

However, unlike forecasting or nowcasting tasks that can be directly handled by common autoregressive models such as RNNs [22, 23, 24] and Transformers [25, 26], the assimilation problem is inherently more complex. The core of assimilation involves merging observations into the physical prior, which necessitates a more sophisticated model to handle the intricate inter-

actions between physical models and observations. Furthermore, the assimilation problem is essentially a Bayesian inference task, which requires the estimation of the posterior distribution of the state variables rather than just a deterministic prediction for future states, and due to the high dimensionality and nonlinearity of the physical and observational models, the posterior distribution is often intractable.

Motivated by the success of DL, we seek to reframe the assimilation problem through the lens of deep generative models (DGMs) [27, 28], which have demonstrated strong capabilities in capturing complex data distributions. The connection between assimilation and DGMs is intuitive: assimilation aims to estimate the posterior distribution of state variables by integrating prior knowledge from physical models with observational data, while DGMs also provide data distributions, but by training with large datasets. In assimilation, physical models contribute to the prior, which is calibrated by the likelihood derived from observational operators and data. Once the posterior is obtained, key statistical quantities such as the mean field and uncertainty estimates can be inferred. DGMs, on the other hand, rely on large datasets sampled from the target distribution, assuming that all hidden patterns can be extracted without explicit physical knowledge. By incorporating learned priors from DGMs with suitable likelihood modeling, we introduce an alternative approach to assimilation, particularly beneficial when the underlying physics are not fully known.

The rest of the paper is organized as follows. First, we review the related works on data assimilation and deep generative models in section 2. Next, in section 3, we present the mathematical formulation of the assimilation task we focus on. Some basic concepts of conditional score-based generative models are introduced, followed by a detailed deduction of our SOAD model and its theoretical analysis. Finally, experiments on a two-layer quasi-geostrophic model will be exhibited in section 4 to demonstrate the effectiveness of our SOAD model, and section 5 concludes the paper with a summary and future work.

## 2. Related works and contributions

### 2.1. Data assimilation frameworks

Data assimilation has been studied for decades. In general, data assimilation can be categorized into *filtering* and *smoothing*. Filtering aims to estimate the current state given observations up to the current time, while



smoothing involves estimating states using both past and future observations. The choice between the two types depends on not only the specific application and objective of the assimilation task but also the availability of data and computational resources.

Ensemble Kalman filter (EnKF) and its variants are representative classical methods for the filtering problem, which adopt Gaussian assumptions and use an ensemble of samples to represent the posterior’s mean and covariance to seek a balance between accurate probabilistic modeling and computational efficiency. In contrast, the 3D-Var method solves for the MAP estimate for the posterior, providing a “best-fit” solution. Particle filters, specially designed for the nonlinear problem, approximate the posterior by a weighted sum of Dirac deltas, allowing non-Gaussian distributions. Despite recent progress in the particle flow filter (PFF) [29, 30] and its demonstrated potential in high-dimensional settings [31], the need for a large number of particles remains a computational bottleneck.

To address the smoothing problem, one of the most intuitive approaches is the Kalman smoother (KS), which includes an additional backward pass to refine estimates using future information compared to the original KF. As ensemble-based extensions of KS, Ensemble Kalman Smoother (EnKS) [32] and iterative EnKS (IEnKS) [33] utilize an ensemble forward propagation to obtain better convergence for nonlinear problems. 4D-Var formulates a variational problem over a time window to retrieve the most likely trajectory, making it inherently a smoothing approach as well.

All of these classical methods rely on physical knowledge and can be prohibitively expensive in high-dimensional problems that require repeated ensemble updates or large-scale optimizations. Despite these challenges, they remain indispensable for reliably integrating observations with complex dynamical models in many scientific and engineering applications. Nowadays, deep learning (DL)-enhanced data assimilation has gained increasing attention [34]. A pioneering example is the framework of [35], which replaces physical models with neural networks, demonstrating the feasibility of data-driven assimilation. Leveraging the flexibility and fast inference of neural networks, some studies have proposed using them as fast emulators [36] or correctors [37] for physical models in hybrid assimilation frameworks built on classical algorithms. Another research direction focuses on addressing the high-dimensionality challenge through latent assimilation (LA) [38], which seeks a lower-dimensional representation of the state variables to be assimilated using either linear Reduced-Order Models (ROMs) [39, 40] or neural

networks [41, 42]. Although methods like LAINR [43] extend the flexibility of LA, they still rely on well-established classical assimilation algorithms as backbones to operate in the latent space.

To better capture non-Gaussian posteriors, recent work has explored the Score-based Filter (SF) [44] and its variants, such as the Ensemble Score-based Filter (EnSF) [45] and Latent-EnSF [46], which rely on diffusion probabilistic models (DPMs) for posterior estimation. These methods have shown promise for nonlinear assimilation tasks but typically assume availability of a physical model. More recently, Score-based Data Assimilation (SDA) [47, 48] has been introduced to integrate physical models into the learning of background priors within a variational assimilation framework, producing promising results for subsampled observations even without any explicit knowledge of physical models. However, as demonstrated in our experiments, SDA still struggles with nonlinear observations.

In the modern numerical weather prediction pipeline, increasing attention has been given to smoothing-based methods, particularly 4D-Var and its variants [49, 50, 51, 52], due to their ability to incorporate time-evolving observations and improve the dynamical consistency of the atmospheric state [53, 54]. These advances highlight a growing need for learning-based smoothing algorithms that can effectively assimilate observations over a time window and handle the non-Gaussianity arising from nonlinear dynamics and observation operators. Motivated by this, our proposed method adopts a smoothing perspective and seeks to estimate latent trajectories conditioned on both past and future observations, with a focus on addressing challenges related to nonlinearity and posterior complexity. Table 1 summarizes key features of some typical data assimilation approaches. Bolded entries highlight capabilities that surpass those of other methods in the same category.

## 2.2. Deep generative models

Deep generative models (DGMs) have been consistently studied in the field of deep learning. As one of the earliest DGMs, variational AutoEncoders (VAEs) [55, 56] learn the target data distribution by maximizing the evidence lower bound (ELBO). Generative Adversarial Networks (GANs) [57, 58] are another kind of DGMs, which learn the distribution through a zero-sum game between generators and discriminators. Alternatively, Normalizing Flows (NFs) [59, 60, 61] are employed not only for generating samples but also for explicitly evaluating probability densities by ensuring each layer to be invertible at the expense of flexibility of network architectures.

Method	physical model	posterior modeling	comput. cost
KF[1], EnKF[2], LETKF[3]	required	mean & covariance	high
PF, PFF[29, 30, 31]	required	weighted Dirac sum	very high
SF[44]	required	<b>generative model</b>	high
EnSF[45], Latent-EnSF[46]	required	<b>generative model</b>	<b>low</b>
4D-Var	required	MAP	high
KS, EnKS[32], IEnKS[33]	required	mean & covariance	high
SDA[47, 48], SOAD (ours)	<b>data-driven</b>	<b>generative model</b>	<b>low</b>

Table 1: Comparison of some typical filtering and smoothing approaches.

Inspired by [62], the most recent diffusion probabilistic models (DPMs) [63, 64, 65] have achieved impressive performances across numerous data generation and restoration tasks, particularly in computer vision and natural language processing domains. Specifically, DPMs have shown state-of-the-art performances in image generation and restoration tasks such as denoising, inpainting, and super-resolution [66, 67, 68], as well as in text-to-image and video generation [69]. These tasks often involve linear measurement operators (e.g., masking, blurring, or downsampling), and thus the associated inverse problems are typically linear, enabling efficient deployment of diffusion models through techniques like iterative reconstruction guided by the learned score function or posterior sampling [70].

More recently, diffusion models have been adapted to tackle complex inverse problems in scientific computing. Unlike many standard image restoration scenarios, scientific computing tasks such as phase retrieval [71, 72], seismic inversion [73, 74, 75], and topology optimization [76, 77, 78] generally involve nonlinear, high-dimensional, and often ill-posed inverse problems. In these contexts, DPMs are leveraged to model the prior distribution of states efficiently and sample from complex posterior distributions conditioned on nonlinear observations.

### 2.3. Main contributions

In this study, we propose a novel data-driven deep learning (DL) approach specifically designed for posterior estimation to handle the assimilation problem. To address the challenges of nonlinearity, we introduce our State-Observation Augmented Diffusion (SOAD) model, which relaxes the linear constraints typically imposed on physical and observational models in

most existing assimilation frameworks. Similar to previous DL-based assimilation methods, our SOAD model is fully data-driven and does not require any prior knowledge of the physical models. Importantly, while our network training still relies on observational data that implicitly reflects the observation process, our model does not require an explicit physical formulation of the observational operator. That is, as long as paired state-observation samples are available, SOAD can be trained and subsequently perform assimilation with observations from multiple heterogeneous sources without requiring any knowledge of how these observations are physically related to the state variables. The flexibility and adaptability make our approach suitable for the real-world applications. Our main contributions are as follows:

- We propose a novel state-observation augmented diffusion model specially designed for nonlinear assimilation tasks. By introducing the state-observation augmented structure, we establish an equivalent linear form of the original assimilation problem, relaxing the linearity assumptions for both the physical and observational models.
- We derive the marginal posterior distribution of the state variables associated with our SOAD model and show that it matches the real posterior under mild assumptions, indicating the theoretical superiority of our SOAD model over previous score-based assimilation approaches.
- A procedure named forward-diffusion corrector is introduced to stabilize the reverse-time generation process of the diffusion model, providing corrections matching the real distribution especially under Gaussian noise.
- Experiments with a two-layer quasi-geostrophic model involving various observational operators are conducted to demonstrate the effectiveness of our SOAD model compared to previous works.

### 3. Methodology

#### 3.1. Problem settings and assumptions

Recall that we consider a discrete dynamical system with an observational model formulated as

$$\begin{cases} \mathbf{x}_k = \mathcal{M}_k^X(\mathbf{x}_{k-1}, \boldsymbol{\eta}_k), & \boldsymbol{\eta}_k \sim \mathcal{D}_k^X, \\ \mathbf{y}_k = \mathcal{H}_k(\mathbf{x}_k) + \boldsymbol{\epsilon}_k, & \boldsymbol{\epsilon}_k \sim \mathcal{D}_k^Y, \end{cases} \quad (6)$$

where the additional superscript “ $X$ ” indicates that the model  $\mathcal{M}_k^X$  acts in the physical space for  $\mathbf{x}_k$ . We follow the setup described in [47, 79], assuming access to a sufficiently large dataset of physical states corresponding to an unknown physical model  $\mathcal{M}_k^X$ . The noise distributions  $\mathcal{D}_k^Y$  associated with all observational operators are assumed to be known as well. Given a sequence of observations  $\mathbf{y}_{S:T}$ , our objective is to effectively estimate the posterior distribution of the state variables  $\mathbf{x}_{S:T}$ , denoted by  $p(\mathbf{x}_{S:T} \mid \mathbf{y}_{S:T})$ .

To keep our approach general, we impose only the following assumptions:

- **Time-independence or periodicity:** The physical model  $\mathcal{M}_k^X$  is either constant over time or exhibits periodic behavior.
- **Decomposable observational operators:** Although the observational operator  $\mathcal{H}_k$  may vary with time, there exists a time-independent operator  $\mathcal{H}$  such that  $\mathcal{H}_k = \mathbf{S}_k \circ \mathcal{H}$  for some time-varying linear operator  $\mathbf{S}_k$  with orthogonal rows.

These assumptions are pragmatic in many practical contexts. For instance, numerical weather prediction frequently exhibits daily or seasonal periodicity in the underlying physical model  $\mathcal{M}_k^X$ . The first assumption ensures the continued applicability and effectiveness of data-driven methods trained on historical data, a standard practice in deep learning-based assimilation methodologies [47, 36, 41, 39, 43]. The observational operator  $\mathcal{H}_k$  typically varies over time due to observational constraints or infrastructure limitations [80], but the range of possible observational configurations usually remains fixed. Therefore, we may set

$$\mathcal{H}(\mathbf{x}_k) = \begin{pmatrix} \mathcal{H}^{(1)}(\mathbf{x}_k) \\ \mathcal{H}^{(2)}(\mathbf{x}_k) \\ \vdots \\ \mathcal{H}^{(K)}(\mathbf{x}_k) \end{pmatrix} \quad (7)$$

as a concatenation of all possible observational outputs, where each  $\mathcal{H}^{(i)}$  corresponds to one type of observation measured on full domain for  $i = 1, 2, \dots, K$ . This suggests that the variation in observational operators for different time steps can be adequately captured by a “subsampling matrix”  $\mathbf{S}_k$  acting on  $\mathcal{H}(\mathbf{x}_k)$ . Here, a subsampling matrix is a binary matrix that selects a subset of elements (or locations) from the full state vector. While

traditional assimilation often presume Gaussian noise distributions  $\mathcal{D}^X$  and  $\mathcal{D}^Y$ , we will demonstrate later that these assumptions can be relaxed as well.

### 3.2. Score-based data assimilation

Score-based data assimilation is inspired by score-based generative models [79], a widely studied class of diffusion probabilistic models (DPMs) in machine learning. These models offer an alternative perspective by leveraging stochastic differential equations for data generation. In this section, we briefly review the fundamental concepts underlying score-based generative models and how they can be applied to the data assimilation problem.

#### 3.2.1. Score-based generative models

Let  $\mathbf{x} \sim p_{\text{data}}(\mathbf{x})$  denote the data distribution we aim to learn. By progressively adding Gaussian noise, the distribution undergoes an approximate transformation into a standard normal distribution, and the whole process can be modeled as a forward-time diffusion process. Mathematically, following the notation in [79], we describe the diffusion process as a linear stochastic differential equation (SDE) formulated as

$$d\mathbf{x}_t = f(t)\mathbf{x}_t dt + g(t)d\mathbf{w}_t, \quad t \in [0, 1], \quad \mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x}_0), \quad (8)$$

where  $f(t)$  and  $g(t)$  represent the drift term and diffusion coefficient, respectively, and  $\mathbf{w}_t$  denotes the standard Wiener process. The transition kernel from  $\mathbf{x}_0$  to  $\mathbf{x}_t$  is then given by

$$p_{t|0}(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \mu_t \mathbf{x}_0, \sigma_t^2 \mathbf{I}) \quad (9)$$

with the mean  $\mu_t$  and the variance  $\sigma_t^2$  explicitly [81] given by

$$\mu_t = \exp\left(\int_0^t f(s)ds\right), \quad \sigma_t^2 = \int_0^t \exp\left(2 \int_s^t f(u)du\right) g(s)^2 ds. \quad (10)$$

Define the marginal distribution for time  $t$  as  $p_t(\mathbf{x})$ . With appropriate choices of  $f(t)$  and  $g(t)$  so that  $\mu_1 = \sigma_0 = 0$  and  $\mu_0 = \sigma_1 = 1$ , the forward diffusion process eq. (8) transforms the data distribution into a standard Gaussian marginal  $p_1(\mathbf{x}) = \mathcal{N}(0, \mathbf{I})$ . Under such circumstances, the reverse-time evolution, characterized by

$$d\mathbf{x}_t = [f(t) - g(t)^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)] dt + g(t)d\mathbf{w}_t, \quad (11)$$

is also a diffusion process [82], where the gradient term  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ , known as the score function, plays a crucial role in guiding the reverse diffusion. With a slight abuse of notation, we still let  $\mathbf{w}_t$  denote the reverse-time Wiener process. Consequently, as long as a good estimation of the score function is obtained for all  $t$ , we may first sample an initial state  $\mathbf{x}_1 \sim \mathcal{N}(0, \mathbf{I})$  and then simulate the reverse-time process starting from  $\mathbf{x}_1$  to generate a sample  $\mathbf{x}_0$  that approximately follow the distribution  $p_{\text{data}}(\mathbf{x})$ .

To develop an effective score-based generative model, intuitively one may minimize the Fisher information distance ([83], Definition 2.5)

$$d(\boldsymbol{\theta}, t) = \mathbb{E}_{\mathbf{x}_t \sim p_t} \|\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)\|^2, \quad (12)$$

where

$$p_t(\mathbf{x}) = \int p_{t|0}(\mathbf{x} \mid \mathbf{x}_0) p_{\text{data}}(\mathbf{x}_0) d\mathbf{x}_0 \quad (13)$$

is the marginal density and  $\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$  is a surrogate model for the score function implemented by a neural network with parameters  $\boldsymbol{\theta}$ . By assuming some weak regularity conditions [84], we have

$$\begin{aligned} d(\boldsymbol{\theta}, t) &= \mathbb{E}_{\mathbf{x}_t \sim p_t} \|\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)\|^2 \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x}_t \sim p_t} [\|\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)\|^2 + \text{Tr} \nabla_{\mathbf{x}_t} \mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x}_0 \sim p_0} \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot \mid \mathbf{x}_0)} [\|\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)\|^2 + \text{Tr} \nabla_{\mathbf{x}_t} \mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x}_0 \sim p_0} \mathbb{E}_{\mathbf{x}_t \sim p_{t|0}(\cdot \mid \mathbf{x}_0)} \|\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t \mid \mathbf{x}_0)\|^2 \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x}_0 \sim p_0} \mathbb{E}_{\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \|\mathbf{s}_{\boldsymbol{\theta}}(\mu_t \mathbf{x}_0 + \sigma_t \boldsymbol{\varepsilon}, t) + \sigma_t^{-1} \boldsymbol{\varepsilon}\|^2 \end{aligned} \quad (14)$$

as discussed in [85]. In practice,  $\mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$  is usually reparameterized as

$$\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t) = -\sigma_t \mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t), \quad (15)$$

and we consider minimizing

$$L(\boldsymbol{\theta}) = \mathbb{E}_{t \sim \mathcal{U}_{[0,1]}} \sigma_t^2 d(\boldsymbol{\theta}, t) = \frac{1}{2} \mathbb{E}_{\substack{t \sim \mathcal{U}_{[0,1]} \\ \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \\ \mathbf{x}_0 \sim p_0(\mathbf{x})}} \|\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}(\mu_t \mathbf{x}_0 + \sigma_t \boldsymbol{\varepsilon}, t) - \boldsymbol{\varepsilon}\|^2. \quad (16)$$

Once a good estimator for the score function is obtained, it suffices to

evolve the backward diffusion process by discretizing the time steps. Classical numerical solvers for general SDEs, such as the Euler-Maruyama and stochastic Runge-Kutta methods, as well as ancestral sampling [66, 79], are applicable. Recently, to improve sampling efficiency, the exponential integrator (EI) discretization scheme [86]

$$\mathbf{x}_{t_-} \leftarrow \frac{\mu_{t_-}}{\mu_t} \mathbf{x}_t + \left( \frac{\sigma_{t_-}}{\sigma_t} - \frac{\mu_{t_-}}{\mu_t} \right) \sigma_t \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, t) \quad (17)$$

has also been proposed, where the subscript “ $t_-$ ” denotes the next time step. These schemes often work together with the predictor-corrector sampling strategy by performing a few steps of Langevin Monte Carlo (LMC) sampling [79]

$$\mathbf{x}_t \leftarrow \mathbf{x}_t + \frac{\delta}{2} \mathbf{s}_\theta(\mathbf{x}_t, t) + \sqrt{\delta} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (18)$$

for better performances, where we let  $\delta$  denote the time step.

### 3.2.2. Conditional score models

While score-based generative models we have introduced above can be directly used for unconditional sampling with well-trained denoising models, assimilation tasks require estimating the posterior distribution  $p(\mathbf{x} \mid \mathbf{y})$ , where we omit the subscript “ $S : T$ ” for simplicity. Analogous to the unconditional case, the posterior distribution can be approximated with the reverse-time diffusion process defined as

$$d\mathbf{x}_t = [f(t) - g(t)^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t \mid \mathbf{y})] dt + g(t) d\mathbf{w}_t, \quad \mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (19)$$

where the conditional score function  $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t \mid \mathbf{y})$  corresponding to the marginal posterior

$$p_t(\mathbf{x}_t \mid \mathbf{y}) = \int p_{t|0}(\mathbf{x}_t \mid \mathbf{x}) p(\mathbf{x} \mid \mathbf{y}) d\mathbf{x} \quad (20)$$



needs to be learned. Note that the conditional score function can be decomposed as

$$\begin{aligned}
\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t \mid \mathbf{y}) &= \nabla_{\mathbf{x}_t} \log \int p_{t|0}(\mathbf{x}_t \mid \mathbf{x}) p(\mathbf{x} \mid \mathbf{y}) d\mathbf{x} \\
&= \nabla_{\mathbf{x}_t} \log \int \frac{p_{t|0}(\mathbf{x}_t \mid \mathbf{x}) p(\mathbf{y} \mid \mathbf{x}) p(\mathbf{x})}{p(\mathbf{y})} d\mathbf{x} \\
&= \nabla_{\mathbf{x}_t} \log \int \frac{p_{0|t}(\mathbf{x} \mid \mathbf{x}_t) p(\mathbf{y} \mid \mathbf{x}) p_t(\mathbf{x}_t)}{p(\mathbf{y})} d\mathbf{x} \\
&= \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log \int p(\mathbf{y} \mid \mathbf{x}) p_{0|t}(\mathbf{x} \mid \mathbf{x}_t) d\mathbf{x}
\end{aligned} \tag{21}$$

by Bayes' formula, where the additional term

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} \mid \mathbf{x}_t) = \nabla_{\mathbf{x}_t} \log \int p(\mathbf{y} \mid \mathbf{x}) p_{0|t}(\mathbf{x} \mid \mathbf{x}_t) d\mathbf{x} \tag{22}$$

is usually called the adversarial gradient.

To incorporate conditional information, one may introduce a separate network for the adversarial gradient alongside the original denoising network as suggested in [87, 88]. Alternatively, a conditional denoising network  $\varepsilon_\theta(\mathbf{x}_t, t; \mathbf{y})$  can be trained directly to approximate  $-\sigma_t \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t \mid \mathbf{y})$ , and such an idea has been widely used in many class-conditional generation and text-to-image/video generation tasks [68, 69].

Instead of training a different denoising network, we focus on estimating the reverse-time transition kernel  $p_{0|t}$  in eq. (22) to compute the term  $p_t(\mathbf{y} \mid \mathbf{x}_t)$  explicitly. By Tweedie's formula, DPS [71] proposes a delta-function approximation for the kernel  $p_{0|t}$  centered at

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, t) = \mathbb{E}[\mathbf{x}_0 \mid \mathbf{x}_t] = \mu_t^{-1}[\mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)]. \tag{23}$$

Based on the uninformative-prior assumption, DMPS [89] replaces the kernel  $p_{0|t}$  with a Gaussian distribution, but it is only suitable for linear observational models. SDA [47, 48] suggests a balance between flexibility and accuracy by introducing an auxiliary hyperparameter  $\gamma$  in the modeling of the kernel  $p_{0|t}$ . Nonetheless, such an approximation is still based on the linearity assumption for the observational operator, which has shown limitations in

Table 2: Comparison of various existing estimators for the adversarial gradient.

Method	$p_{0 t}(\mathbf{x}_0 \mid \mathbf{x}_t)$	$p_t(\mathbf{y} \mid \mathbf{x}_t)$
DPS [71]	$\delta(\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t, t))$	$\mathcal{N}(\mathbf{y}; \mathcal{A}(\hat{\mathbf{x}}_0(\mathbf{x}_t, t)), \mathbf{R})$
DMPS [89]	$\mathcal{N}\left(\mathbf{x}_0; \mu_t^{-1} \mathbf{x}_t, \frac{\sigma_t^2}{\mu_t^2} \mathbf{I}\right)$	$\mathcal{N}\left(\mathbf{y}; \mu_t^{-1} \mathcal{A}(\mathbf{x}_t), \mathbf{R} + \frac{\sigma_t^2}{\mu_t^2} \mathbf{A} \mathbf{A}^\top\right)$
SDA [47, 48]	N/A	$\mathcal{N}\left(\mathbf{y}; \mathcal{A}(\hat{\mathbf{x}}_0(\mathbf{x}_t, t)), \mathbf{R} + \frac{\sigma_t^2}{\mu_t^2} \gamma \mathbf{I}\right)$

our experiments. Take a nonlinear observational model

$$p(\mathbf{y} \mid \mathbf{x}) = \mathcal{N}(\mathbf{y}; \mathcal{A}(\mathbf{x}), \mathbf{R}) \quad (24)$$

as an example. The DPS, DMPS, and SDA estimators for the term  $p(\mathbf{y} \mid \mathbf{x}_t)$  are summarized in table 2, where we use  $\mathbf{A}$  to denote the linearization of the observational operator  $\mathcal{A}$ .

### 3.3. State-observation augmented diffusion model

In this section, we introduce our State-Observation Augmented Diffusion (SOAD) model, which is designed to handle the assimilation problem with nonlinear physical and observational models. We first present the state-observation augmented dynamical system as an equivalent but linear form of the original assimilation problem to alleviate the nonlinearity issue. Then, we derive an estimator for the marginal posterior  $p_t(\mathbf{x}_t \mid \mathbf{y})$  defined in eq. (20), which has been shown to match the real posterior under mild assumptions. Figure 1 depicts the general idea of our SOAD model, where we aim to assimilate on a time window from  $S$  to  $T$  with  $K$  types of observations. Finally, we introduce a forward-diffusion corrector to stabilize the reverse-time generation process of the diffusion model, followed by an overall implementation of the SOAD model.

#### 3.3.1. Augmented dynamical system

To handle the potential nonlinearity of the observational model, we propose to consider an augmented version of eq. (6). Let  $\mathbf{z}_k$  denote the concatenation of the state variables  $\mathbf{x}_k$  and all possible observation variables  $\mathcal{H}(\mathbf{x}_k)$ ,

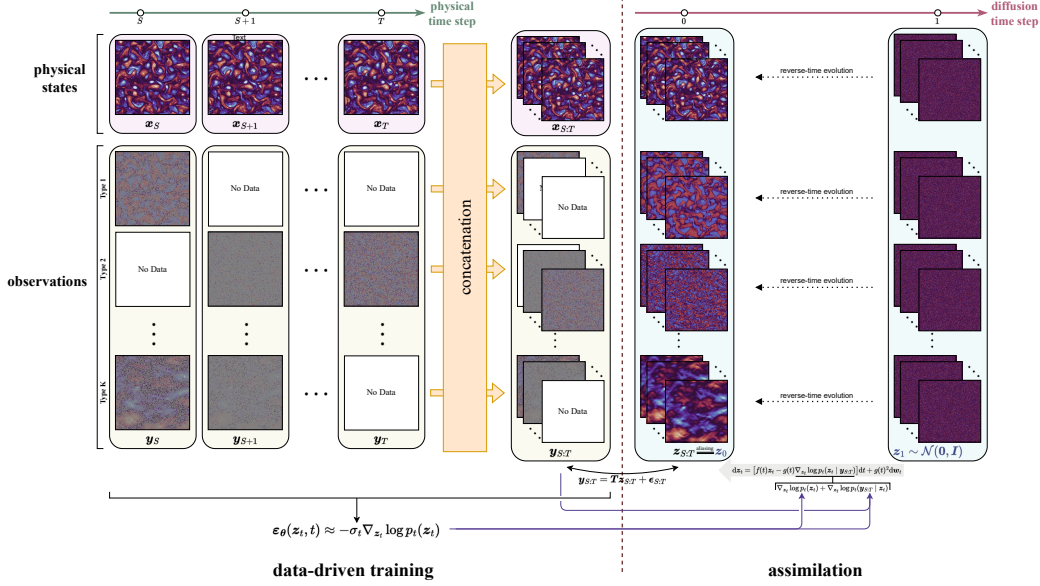


Figure 1: Illustration of assimilation with our State-Observation Augmented Diffusion model.

then we have

$$\mathbf{z}_k = \begin{pmatrix} \mathbf{x}_k \\ \mathcal{H}(\mathbf{x}_k) \end{pmatrix} = \begin{pmatrix} \mathcal{M}_k^X(\mathbf{x}_{k-1}, \boldsymbol{\eta}_k) \\ \mathcal{H}(\mathcal{M}_k^X(\mathbf{x}_{k-1}, \boldsymbol{\eta}_k)) \end{pmatrix} =: \mathcal{M}_k^Z(\mathbf{z}_{k-1}, \boldsymbol{\eta}_k), \quad \boldsymbol{\eta}_k \sim \mathcal{D}^X, \quad (25)$$

where the superscript “Z” emphasizes the augmented space where  $\mathbf{z}_k$  lies. By the assumptions in section 3.1, for the observations at the  $k$ -th time step, there exists a linear mapping  $\mathbf{T}_k$  linked to  $\mathcal{H}_k$  such that

$$\mathbf{y}_k = \mathcal{H}_k(\mathbf{x}_k) + \boldsymbol{\epsilon}_k = \mathbf{S}_k \mathcal{H}_k(\mathbf{x}_k) + \boldsymbol{\epsilon}_k = \mathbf{T}_k \mathbf{z}_k + \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{D}^Y. \quad (26)$$

Therefore, we have essentially established an equivalent augmented system

$$\begin{cases} \mathbf{z}_k = \mathcal{M}_k^Z(\mathbf{z}_{k-1}, \boldsymbol{\eta}_k), & \boldsymbol{\eta}_k \sim \mathcal{D}^X, \\ \mathbf{y}_k = \mathbf{T}_k \mathbf{z}_k + \boldsymbol{\epsilon}_k, & \boldsymbol{\epsilon}_k \sim \mathcal{D}^Y. \end{cases} \quad (27)$$

Clearly, the forward propagation operator  $\mathcal{M}_k^Z$  inherits the periodicity of  $\mathcal{M}_k^X$ . The observational operator  $\mathbf{T}_k$  acting on  $\mathbf{z}_k$  is now a time-dependent subsampling matrix, which simplifies the assimilation process. By concate-

nating  $\mathbf{z}$  and  $\mathbf{y}$  over consecutive time steps as  $\mathbf{z}_{S:T}$  and  $\mathbf{y}_{S:T}$ , respectively, the assimilation problem we focus on becomes a posterior estimation for  $p(\mathbf{z}_{S:T} | \mathbf{y}_{S:T})$ . The prior  $p(\mathbf{z}_{S:T})$  implicitly encodes the physical dynamics as well as the model error, and the observational model

$$\mathbf{y}_{S:T} = \mathbf{T}_{S:T} \mathbf{z}_{S:T} + \boldsymbol{\epsilon}_{S:T}, \quad \boldsymbol{\epsilon}_{S:T} \sim \mathcal{D}_{S:T}^Y \quad (28)$$

becomes linear and Gaussian, where we use the subscripts “ $S : T$ ” to denote the concatenation of the corresponding variables, mappings, and distributions over the time interval from  $S$  to  $T$ .

### 3.3.2. Estimation for the marginal posterior

In general, our State-Observation Augmented Diffusion (SOAD) model builds on the principles of score-based generative models outlined in section 3.2.1. Specifically, we construct a score-based generative model for  $\mathbf{z}_{S:T}$  instead of  $\mathbf{x}$ . For simplicity, we omit the subscript “ $S : T$ ” hereafter, and thus  $\mathbf{z}_{S:T}$  becomes  $\mathbf{z} = \mathbf{z}_0$  in the context of diffusion process. We need to stress that the following derivations are all applied to concatenated versions of states, observations, and noise (distributions) over the time interval from  $S$  to  $T$ .

To estimate the marginal posterior  $p_t(\mathbf{z}_t | \mathbf{y})$ , we continue from the Bayes’ rule

$$\nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t | \mathbf{y}) = \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t) + \nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t) \quad (29)$$

adapted from eqs. (21) and (22) in the current context, where we still use the notation

$$\nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t) = \nabla_{\mathbf{z}_t} \log \int p(\mathbf{y} | \mathbf{z}) p_{0|t}(\mathbf{z} | \mathbf{z}_t) d\mathbf{z}. \quad (30)$$

On one hand, our SOAD model directly trains a neural network

$$\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{z}_t, t) = -\sigma_t \mathbf{s}_{\boldsymbol{\theta}}(\mathbf{z}_t, t) \approx -\sigma_t \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t) \quad (31)$$

according to eq. (16) to estimate the joint unconditional distribution for  $\mathbf{z}$ .

On the other hand, to estimate the adversarial gradient  $\nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t)$ ,

by applying Bayes' rule again, we have

$$\begin{aligned}\nabla_{\mathbf{z}} \log p_{0|t}(\mathbf{z} \mid \mathbf{z}_t) &= \nabla_{\mathbf{z}} \log p_{t|0}(\mathbf{z}_t \mid \mathbf{z}) + \nabla_{\mathbf{z}} \log p(\mathbf{z}) \\ &= \sigma_t^{-2} \nabla_{\mathbf{z}} \|\mathbf{z}_t - \mu_t \mathbf{z}\|^2 + \nabla_{\mathbf{z}} \log p(\mathbf{z}) \\ &= \mu_t^2 \sigma_t^{-2} (\mathbf{z} - \mu_t^{-1} \mathbf{z}_t) + \nabla_{\mathbf{z}} \log p(\mathbf{z}).\end{aligned}\tag{32}$$

If we make Gaussian assumptions for the prior  $p(\mathbf{z})$  with a covariance matrix  $\Sigma_0$ , then  $p_{0|t}(\cdot \mid \mathbf{z}_t)$  is also Gaussian by eq. (32), whose covariance matrix has a closed form

$$\Sigma_{0|t} = (\Sigma_0^{-1} + (\sigma_t^2/\mu_t^2)^{-1} \mathbf{I})^{-1}.\tag{33}$$

Similar to [71], the mean of  $p_{0|t}(\cdot \mid \mathbf{z}_t)$  may be provided by Tweedie's formula as

$$\hat{\mathbf{z}}_0(\mathbf{z}_t, t) = \mathbb{E}[\mathbf{z}_0 \mid \mathbf{z}_t] = \mu_t^{-1} (\mathbf{z}_t + \sigma_t^2 \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t)).\tag{34}$$

Therefore, we use the following estimation for the reverse-time transition kernel

$$\begin{aligned}p_{0|t}^{\text{SOAD}}(\mathbf{z} \mid \mathbf{z}_t) &= \mathcal{N}(\mathbf{z}; \hat{\mathbf{z}}_0(\mathbf{z}_t, t), \Sigma_{0|t}) \\ &= \mathcal{N}\left(\mathbf{z}; \mu_t^{-1} (\mathbf{z}_t - \sigma_t \epsilon_\theta(\mathbf{z}_t, t)), (\Sigma_0^{-1} + (\sigma_t^2/\mu_t^2)^{-1} \mathbf{I})^{-1}\right).\end{aligned}\tag{35}$$

It follows that the term  $p_t(\mathbf{y} \mid \mathbf{z}_t)$  can be computed analytically once the observational likelihood  $p(\mathbf{y} \mid \mathbf{z}_0)$  is provided explicitly by eq. (30). For instance, when the observational noise is Gaussian with covariance matrix  $\Sigma^Y$ , we may assume

$$p(\mathbf{y} \mid \mathbf{z}) = \mathcal{N}(\mathbf{y}; \mathbf{T}\mathbf{z}, \Sigma^Y),\tag{36}$$

and then the time-dependent marginal likelihood becomes

$$p^{\text{SOAD}}(\mathbf{y} \mid \mathbf{z}_t) = \mathcal{N}(\mathbf{y}; \mathbf{T}\hat{\mathbf{z}}_0(\mathbf{z}_t, t), \Sigma^Y + \mathbf{T}\Sigma_{0|t}\mathbf{T}^\top).\tag{37}$$

Note that despite the possible nonlinearity of the observational operator  $\mathcal{H}$ , the augmentation as described in eq. (27) allows us to handle the assimilation problem with a linear observational model.

To sum up, we have the following theorem.

**Theorem 3.1.** *Under Gaussian assumption for the prior  $p(\mathbf{z})$  with covariance  $\Sigma_0$ , our estimator given by eq. (37) matches the marginal posterior  $p_t(\mathbf{y} \mid \mathbf{z}_t)$  defined in eq. (22) when the observation noise  $\mathcal{D}^Y$  is Gaussian in*

eq. (26), or equivalently, the assumption in eq. (36) holds. Furthermore, for any scalar prior covariance  $\Sigma_0 = \sigma_z^2 \mathbf{I}$ , our estimator becomes

$$p_{\text{SOAD}}(\mathbf{y} \mid \mathbf{z}_t) = \mathcal{N}\left(\mathbf{y}; \mathbf{T} \hat{\mathbf{z}}_0(\mathbf{z}_t, t), \Sigma^Y + \frac{\sigma_z^2 r_t^2}{\sigma_z^2 + r_t^2} \mathbf{I}\right), \quad (38)$$

where we define  $r_t = \sigma_t / \mu_t$ .

*Proof.* The first part is evident from the previous deduction. To show the second part, we may notice that  $\Sigma_{0|t}$  is scalar as long as  $\Sigma_0$  is scalar according to eq. (33), so it follows that

$$\mathbf{T} \Sigma_{0|t} \mathbf{T}^\top = \frac{\sigma_z^2 r_t^2}{\sigma_z^2 + r_t^2} \mathbf{T} \mathbf{T}^\top = \frac{\sigma_z^2 r_t^2}{\sigma_z^2 + r_t^2} \mathbf{I}, \quad (39)$$

where we use the fact that  $\mathbf{T}$  is a subsampling matrix and thus has orthonormal rows.  $\square$

*Remark.* We need to clarify that one should not confuse our SOAD estimator with the SDA estimator listed in table 2. First, we only apply our estimator to the augmented model (27), where the operator  $\mathbf{T}$  is guaranteed to be linear. Under mild assumptions in Theorem 3.1, the linearity ensures that the time-dependent marginal likelihood  $p_t(\mathbf{y} \mid \mathbf{z}_t)$  remains Gaussian. In contrast, the SDA estimator employs a Gaussian approximation of the true  $p_t(\mathbf{y} \mid \mathbf{z}_t)$ , and the linearization of  $\mathcal{H}$  may introduce significant errors if the observational operator  $\mathcal{H}$  is highly nonlinear. Furthermore, the approximation term  $r_t^2 \gamma \mathbf{I}$  in SDA does not prevent the covariance from exploding as the diffusion step  $t$  approaches 1, resulting in a prior for  $p(\mathbf{y})$  with an unbounded covariance.

### 3.3.3. Forward-diffusion corrector

In our augmented model (27), the observation  $\mathbf{y}$  is a linear transformation of  $\mathbf{z}$ , or more specifically, a subsampled version of  $\mathbf{z}$  with observational noise as described in eq. (26). To stabilize the reverse-time diffusion process, we propose directly replacing part of the denoised results with a perturbed version of  $\mathbf{y}$ . The correction is based on an auxiliary forward diffusion of the observational data, which is then used to guide or adjust the denoising states where observations are available. Consequently, we refer to such a corrector as the “forward-diffusion corrector”.




Figure 2: Illustration of the forward-diffusion corrector.

More precisely, let  $\mathbf{z}_t$  be the diffusion state at time  $t$  with an initial state  $\mathbf{z}_0 = \mathbf{z}$ , then

$$\mathbf{T}\mathbf{z}_t = \mathbf{T}(\mu_t \mathbf{z}_0 + \sigma_t \boldsymbol{\varepsilon}) \sim \mathcal{N}(\mu_t \mathbf{T}\mathbf{z}, \sigma_t^2 \mathbf{I}) \quad (40)$$

with a standard Gaussian noise  $\boldsymbol{\varepsilon}$  according to eq. (9). To calibrate the diffusion state  $\mathbf{z}_t$ , we propose substituting  $\mathbf{T}\mathbf{z}_t$  at time  $t$  with  $(\mu_t \mathbf{y} + \boldsymbol{\varepsilon}'_t)$  assuming that eq. (36) holds, where  $\boldsymbol{\varepsilon}'_t \sim \mathcal{D}'_t$  is a random noise term. Since eq. (36) implies

$$\mu_t \mathbf{y} \sim \mathcal{N}(\mu_t \mathbf{T}\mathbf{z}, \mu_t^2 \boldsymbol{\Sigma}^Y), \quad (41)$$

we ensure that  $\mathbf{T}\mathbf{z}_t$  and  $(\mu_t \mathbf{y} + \boldsymbol{\varepsilon}'_t)$  share the same means and covariances by setting

$$\mathcal{D}'_t = \mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I} - \mu_t^2 \boldsymbol{\Sigma}^Y) \quad (42)$$

whenever the matrix  $(\sigma_t^2 \mathbf{I} - \mu_t^2 \boldsymbol{\Sigma}^Y)$  is positive semidefinite.

It is important to note that this condition generally holds except for small  $t$ , particularly when the observation noise is not excessively large. This follows from the fact that, under typical diffusion schedulers,  $\lim_{t \rightarrow 1^-} r_t = +\infty$  ensuring that our correction provides reliable estimation of  $\mathbf{T}\mathbf{z}_t$  most of the time. The overall procedure is illustrated in fig. 2, where we define  $\mathbf{T}_2$  as the compensated<sup>1</sup> subsampling matrix associated with  $\mathbf{T}$ .

<sup>1</sup>The term “compensated” refers to the following construction. Given a high-dimensional row vector  $\mathbf{x} \in \mathbb{R}^n$  and a subsampling matrix  $\mathbf{T}$  acting on  $\mathbb{R}^n$ , we can always decompose  $\mathbf{x}$  as  $\mathbf{x} = (\mathbf{T}\mathbf{x}, \mathbf{x}_2)$ , up to a permutation of the coordinate axes. The mapping  $\mathbf{x} \mapsto \mathbf{x}_2$  is then called the compensated version of  $\mathbf{T}$ , representing the complement of the subsampling process.

### 3.3.4. Implementation

We describe the detailed denoising procedure in this subsection, where we set  $\mathcal{D}^Y = \mathcal{N}(\mathbf{0}, (\sigma^o)^2 \mathbf{I})$  for simplicity. Similar deduction can be applied to other more complex settings.

Let  $r_\sigma = \sigma_{t-}/\sigma_t$  and  $r_\mu = \mu_{t-}/\mu_t$  be the ratios of the diffusion parameters at adjoint time step  $t_-$  and  $t$ , respectively. We choose the EI discretization scheme

$$\mathbf{z}_{t-} = r_\mu \mathbf{z}_t + (r_\sigma - r_\mu) \boldsymbol{\pi}_\theta(\mathbf{z}_t, t) \quad (43)$$

as in eq. (17) for the reverse-time stepper, where we define  $\boldsymbol{\pi}_\theta(\mathbf{z}_t, t) = -\sigma_t^2 \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t | \mathbf{y})$ . By plugging eqs. (29), (31) and (38) into the definition of  $\boldsymbol{\pi}_\theta$ , we have

$$\begin{aligned} \boldsymbol{\pi}_\theta(\mathbf{z}_t, t) &= -\sigma_t^2 \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t | \mathbf{y}) \\ &= -\sigma_t^2 \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t) - \sigma_t^2 \nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t) \\ &= \sigma_t \boldsymbol{\varepsilon}_\theta(\mathbf{z}_t, t) + \frac{\sigma_t^2/2}{(\sigma^o)^2 + \frac{\sigma_z^2 r_t^2}{\sigma_z^2 + r_t^2}} \nabla_{\mathbf{z}_t} \|\mathbf{y} - \mathbf{T} \hat{\mathbf{z}}_0(\mathbf{z}_t, t)\|^2 \\ &= \sigma_t \boldsymbol{\varepsilon}_\theta(\mathbf{z}_t, t) + \frac{r_t^2/2}{(\sigma^o)^2 + \frac{\sigma_z^2 r_t^2}{\sigma_z^2 + r_t^2}} \nabla_{\mathbf{z}_t} \|\mu_t \mathbf{y} - \mathbf{T}(\mathbf{z}_t - \sigma_t \boldsymbol{\varepsilon}_\theta(\mathbf{z}_t, t))\|^2 \\ &= \sigma_t \boldsymbol{\varepsilon}_\theta(\mathbf{z}_t, t) + c_t \nabla_{\mathbf{z}_t} \|\mu_t \mathbf{y} - \mathbf{T} \mathbf{z}_t + \mathbf{T} \sigma_t \boldsymbol{\varepsilon}_\theta(\mathbf{z}_t, t)\|^2 \end{aligned} \quad (44)$$

In practice, we treat  $\sigma_z$  as a hyperparameter to be tuned as it is hard to evaluate explicitly for a trained generative model. Meanwhile, a certain clipping mechanism for the coefficient  $c_t$  is required to stabilize the generation process when the diffusion step  $t$  is near 1. Fortunately, the model performances are not sensitive to the choice of  $\sigma_z$  as well as the clipping mechanism. See algorithm 1 for the complete implementation.

## 4. Experiments

To show the advantages of our proposed SOAD model, we conduct experiments on the two-layer quasi-geostrophic model following the settings proposed in the Score-based Data Assimilation (SDA) [47, 48], and we treat it as the baseline due to its effectiveness shown for linear observations.

We start by outlining the dataset and the network training procedures for the experiments. Next, we detail the observation operators as well as the assimilation settings we used for testing. A comparative analysis of our



---

**Algorithm 1** Assimilation with the SOAD model for Gaussian noise

---

```

1: Input: The subsampling matrix  $\mathbf{T}$ , the observation data  $\mathbf{y}$ , the noise distribution
    $\mathcal{D}^Y = \mathcal{N}(\mathbf{0}, (\sigma^o)^2 \mathbf{I})$ , a trained noise estimator  $\varepsilon_\theta(\mathbf{z}_t, t)$ , and the denoising diffusion
   time step  $\delta t$ ;
2: Hyperparameters:  $\sigma_z = 1.0$ ,  $\delta = 0.25$ ,  $N_c = 5$ .

3: function ASSIMILATIONWITHSOAD( $\mathbf{T}$ ,  $\mathbf{y}$ ,  $\sigma^o$ ,  $\varepsilon_\theta$ ,  $\delta t$ )
4:    $t \leftarrow 1$ ;  $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ;
5:   while  $t > 0$  do

6:      $t_- \leftarrow t - \delta t$ ;  $r_\sigma \leftarrow \sigma_{t_-} / \sigma_t$ ;  $r_\mu \leftarrow \mu_{t_-} / \mu_t$ ;
7:      $c_t \leftarrow \frac{r_t^2/2}{(\sigma^o)^2 + \frac{\sigma_z^2 r_t^2}{\sigma_z^2 + r_t^2}}$ ;
8:      $\mathbf{Q}_t \leftarrow \nabla_{\mathbf{z}_t} \|\mu_t \mathbf{y} - \mathbf{T} \mathbf{z}_t + \mathbf{T} \sigma_t \varepsilon_\theta(\mathbf{z}_t, t)\|^2$ ;
9:      $\hat{c}_t \leftarrow \max(1, 1/\|\mathbf{Q}_t\|_\infty)$ ; ▷ clipping mechanism
10:     $\boldsymbol{\pi}_\theta(\mathbf{z}_t, t) \leftarrow \sigma_t \varepsilon_\theta(\mathbf{z}_t, t) + \hat{c}_t \mathbf{Q}_t$ ; ▷ by eq. (44)
11:     $\mathbf{z}_{t_-} \leftarrow r_\mu \mathbf{z}_t + (r_\sigma - r_\mu) \boldsymbol{\pi}_\theta(\mathbf{z}_t, t)$ ; ▷ reverse-time evolution
12:     $\mathbf{z}_{t_-} \leftarrow \text{FDC}(\mathbf{z}_{t_-}, t_-)$ ;

13:     $\mathbf{s}_{t_-} \leftarrow \sigma_{t_-}^{-1} \boldsymbol{\pi}_\theta(\mathbf{z}_{t_-}, t_-)$ ; ▷ cache the conditional score function
14:    for  $i = 1$  to  $N_c$  do
15:       $\mathbf{z}_{t_-} \leftarrow \mathbf{z}_{t_-} + \frac{\delta}{2} \mathbf{s}_{t_-} + \sqrt{\delta} \boldsymbol{\varepsilon}''$ ,  $\boldsymbol{\varepsilon}'' \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ; ▷ LMC sampling as in eq. (18)
16:       $\mathbf{z}_{t_-} \leftarrow \text{FDC}(\mathbf{z}_{t_-}, t_-)$ ;
17:    end for

18:     $t \leftarrow t_-$ ;
19:  end while
20: end function

21: function FDC( $\mathbf{z}_t, t$ ) ▷ the forward-diffusion corrector
22:    $r_t \leftarrow \sigma_t / \mu_t$ ;
23:   if  $\sigma^o \leq r_t$  then
24:      $\mathbf{T} \mathbf{z}_t \leftarrow \mu_t \left( \mathbf{y} + \sqrt{r_t^2 - (\sigma^o)^2} \boldsymbol{\varepsilon}'_t \right)$ ,  $\boldsymbol{\varepsilon}'_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ;
25:   end if
26:   return  $\mathbf{z}_t$ 
27: end function

```

---

SOAD model and the SDA baseline with various observational operators is presented. Besides, we also try to explore the long-term behavior of our model and its capability of handling multiple observations simultaneously. Our experimental results indicate that our SOAD model is more suitable for dealing with highly nonlinear observations.

#### 4.1. Dataset

Consider the two-layer quasi-geostrophic (QG) evolution equation

$$\begin{aligned}\partial_t q_1 + J(\psi_1, q_1) + \beta_1 \psi_{1x} &= \text{ssd}, \\ \partial_t q_2 + J(\psi_2, q_2) + \beta_2 \psi_{2x} &= -r_{ek} \nabla^2 \psi_2 + \text{ssd}\end{aligned}\tag{45}$$

with the potential vorticity

$$q_1 = \nabla^2 \psi_1 + F_1(\psi_2 - \psi_1), \quad q_2 = \nabla^2 \psi_2 + F_2(\psi_1 - \psi_2)\tag{46}$$

and the mean potential vorticity gradients

$$\beta_1 = \beta + F_1(U_1 - U_2), \quad \beta_2 = \beta - F_2(U_1 - U_2).\tag{47}$$

The horizontal Jacobian is defined as

$$J(\psi, q) = \psi_x q_y - \psi_y q_x,\tag{48}$$

and “ssd” indicates the small-scale dissipation. To accelerate the preparation process, the dataset has been generated by a Python package `pyqg-jax`<sup>2</sup> as a `Jax` [90] implementation of the original `pyqg` [91] library, and all the other equation parameters follow the default configurations (See appendix [Appendix A](#) for completeness).

To create our dataset, we set the solution domain as  $512 \times 512$  with 15-minute time steps. The QG equation is evolved from 10 different random initial states separately. After a warm-up period of 5 years, the reference states are downsampled from the trajectories to a spatial resolution of  $256 \times 256$  with a temporal resolution of  $\Delta t = 24$  hours for 100,000 model days, which leads to a dataset of size  $10 \times 100000 \times 2 \times 256 \times 256$  in total. Owing to the chaotic nature of the QG model, the 5-year warm-up stage ensures that the trajectories evolve into statistically independent states, making it

---

<sup>2</sup><https://github.com/karlotness/pyqg-jax>




Figure 3: Visualization of the vorticity for the two-layer quasi-geostrophic model.

reasonable to directly split them into training and testing datasets. This setup allows the performance on the testing dataset to reflect the model’s ability to generalize to unseen states, thereby assessing its generalization ability.

Since the evolution of the two-layer QG model does not depend on time, we assume that the vorticity snapshots  $(q_1, q_2)^\top$  of shape  $(2, 256, 256)$  within a fixed time range follow the same data distribution. Consequently, we divide each trajectory into 1,000 windows and retain only the first 32 snapshots from every 100 snapshots to avoid temporal correlations. Such a process transforms the dataset into  $10 \times 1000$  windows of size  $32 \times 2 \times 256 \times 256$ . fig. 3 visualizes the generated data for a single time step.

#### 4.2. Network structures and training procedures

To ensure a fair comparison with the SDA baseline, we adopt the same network structure for the denoising network  $\varepsilon_\theta(\mathbf{x}_t, t)$ , which utilizes a U-Net structure mixed with temporal convolutions and embeddings for diffusion time steps  $t$ . We have included the detailed architecture in appendix [Appendix B](#) for completeness. The only difference between our network and the SDA baseline lies in the input channels. SDA learns a generative model for the state  $\mathbf{x}_{S:T}$ , while our SOAD learns for the augmented state  $\mathbf{z}_{S:T}$ , which means that the corresponding observations  $\mathbf{y}_{S:T}$  is fed into the network for training as well. Since each augmented state includes a clean observation  $\mathcal{H}(\mathbf{x})_{S:T}$ , training SOAD requires additional calls of the observation operator  $\mathcal{H}$ .

We use the first  $8 \times 1000$  windows for training and reserve the remaining  $2 \times 1000$  windows for validation and testing. During training, a chunk of

length 9 from each window  $\mathbf{z}_{0:31}^{(j)}$  of length 32 is selected randomly to create a mini-batch for the networks, which gives rise to the following minimization problem

$$\min_{\boldsymbol{\theta}} L(\boldsymbol{\theta}) = \min_{\boldsymbol{\theta}} \frac{1}{2J} \sum_{j=1}^J \left\| \boldsymbol{\varepsilon}_{\boldsymbol{\theta}} \left( \mu_{t_j} \mathbf{z}_{S_j:S_j+8}^{(j)} + \sigma_{t_j} \boldsymbol{\varepsilon}_j, t_j \right) - \boldsymbol{\varepsilon}_j \right\|^2 \quad (49)$$

as an empirical estimation of eq. (16), where

$$\boldsymbol{\varepsilon}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad t_j \sim \mathcal{U}_{[0,1]}, \quad S_j \sim \mathcal{U}_{\{0,1,\dots,23\}} \quad (50)$$

are all independent random variables for each  $j$ . Here,  $J = 8 \times 1000$  denotes the number of total training windows. All networks are trained with the Adam optimizer [92], employing a learning rate of  $2 \times 10^{-4}$  and a weight decay of  $1 \times 10^{-5}$ . To stabilize the training process, all the network inputs are normalized using the empirical mean and variance of the entire training dataset.

#### 4.3. Observations

We consider the following Gaussian observational model

$$\mathbf{y}_k = \mathbf{S}_k \circ \mathcal{H}(\mathbf{x}_k) + \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{N}(0, 0.1^2 \mathbf{I}), \quad (51)$$

where  $k$  denotes the time step index.  $\mathbf{x}_k$  consists of the two-layer vorticity at time step  $k$ , and  $\mathbf{y}_k$  stands for the associated noisy observations. There are three different observation operators used in our experiments to evaluate the performance of our SOAD model under various observational settings, namely,

- an easy nonlinear “arctan” operator

$$\mathcal{H}^e : \mathbf{x} \mapsto \arctan 3\mathbf{x}, \quad (52)$$

- a hard nonlinear sinusoidal operator

$$\mathcal{H}^h : \mathbf{x} \mapsto \frac{3}{2} \sin 3\mathbf{x}, \quad \text{and} \quad (53)$$

- a vorticity-to-velocity mapping

$$\mathcal{H}^{v2v} : \mathbf{x} = (q_1, q_2)^\top \mapsto C \odot (u_1, v_1, u_2, v_2)^\top. \quad (54)$$

The first two operators are element-wise, and we refer to  $\mathcal{H}^e$  as the “easy” case since it is injective and monotonic, while  $\mathcal{H}^h$  is termed the “hard” case due to its highly nonlinear behavior. Besides, to explore the potential of our model for real applications, we also consider the non-local vorticity-to-velocity mapping  $\mathcal{H}^{v2v}$ , where  $(u_i, v_i)$  are the velocity components for the  $i$ -th layer and  $C \odot$  denotes a multiplication with component-wise scalars so that the outputs approximately have zero means and unit deviations. Such a choice is inspired by the fact that velocity components are often directly observed rather than the related vorticity in the real world. To compute  $\mathcal{H}^{v2v}$ , eq. (46) is solved for the stream function  $(\psi_1, \psi_2)^\top$  in the spectral space, and then the velocity components are calculated by the stream functions. Note that we cannot conclude whether  $\mathcal{H}^{v2v}$  is easier or harder, as the operator is linear, but the mapping rule is much more complex.

Next, we need to specify the subsampling matrix  $\mathbf{S}_k$  to determine the observational operator  $\mathcal{H}_k = \mathbf{S}_k \circ \mathcal{H}$  for each  $k$ . The observational frequency plays a crucial role, as different data sources, such as radar or satellite imagery, operate at varying time intervals. In our experiments, we denote by  $N$  the number of time steps between consecutive observations and vary it across different settings. We also investigate the impact of different observational ratios for partial observations, considering two masking strategies:

- (i) random sampling from the grid with a ratio  $p$ , and
- (ii) uniform sampling with a stride  $s$ , corresponding to an observational ratio<sup>3</sup> of  $(1/s)^2$ .

Readers may refer to fig. 4 for detailed illustrations, where we only visualize the first-layer vorticity ( $q_1$ ) for simplicity. Throughout our experiments, we evaluate random ratios  $p \in \{1.0, 0.25, 0.625, 0.01\}$  and strides  $s \in \{1, 2, 4, 10\}$  for partial observations under varying time intervals  $N \in \{1, 2, 4, 8\}$ .

---

<sup>3</sup>For a spatial stride of  $s = 10$ , the actual observational ratio slightly exceeds 1% since the snapshot size is not divisible by 10, but we omit this minor discrepancy for simplicity.

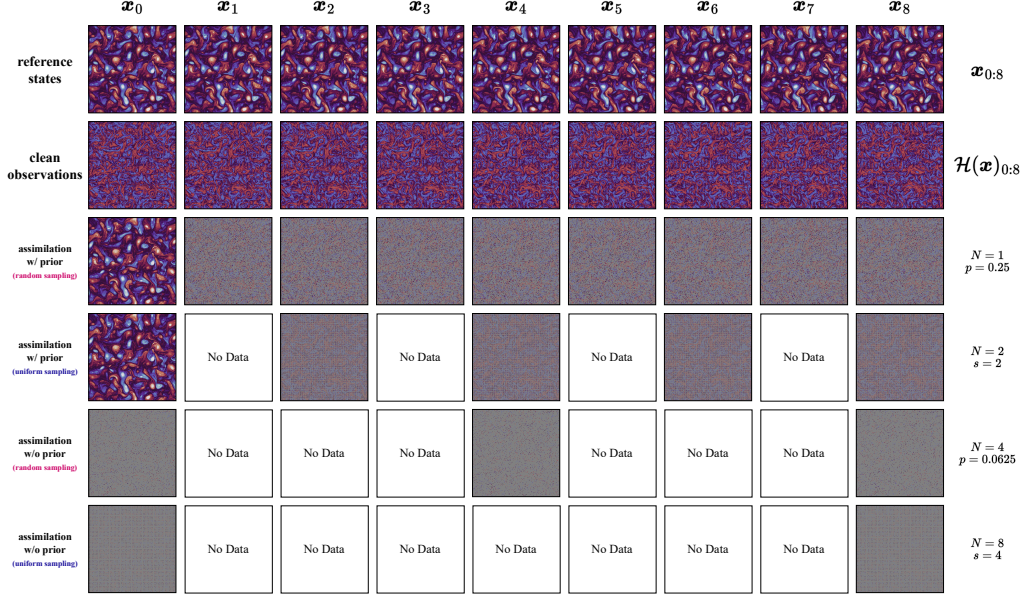


Figure 4: Illustration of different assimilation settings. (zoom-in for better visualization)

#### 4.4. Metric

To make quantitative comparisons, we use the rooted mean-square error (RMSE)

$$\text{RMSE}(\mathbf{x}^r, \mathbf{x}^a) = \left( \frac{1}{2H_{\text{res}}W_{\text{res}}} \sum_{i=1}^2 \sum_{j=1}^{H_{\text{res}}} \sum_{k=1}^{W_{\text{res}}} [\bar{q}_i^r(j, k) - \bar{q}_i^a(j, k)]^2 \right)^{1/2} \quad (55)$$

to evaluate the assimilated state  $\mathbf{x}^a$ . Recall that  $\mathbf{x}$  consists of vorticity  $(q_1, q_2)^\top$  on two levels, and here we use  $(\bar{q}_1, \bar{q}_2)^\top$  for their normalized versions.  $\mathbf{x}^r$  is the reference solution, and  $(H_{\text{res}}, W_{\text{res}})$  stands for the spatial resolution. Since all the data have been normalized in advance, the RMSEs result from different models and various assimilation settings are comparable. To decrease randomness, all the following experiments have been run for 5 times with different random seeds, and the averaged RMSEs are reported unless specified otherwise.

#### 4.5. Assimilation experiments

In the assimilation experiments, we fix our assimilation window to 9 time steps, which means we have observations from at most 9 consecutive steps. Our goal is to assimilate all state variables within this window, as illustrated in fig. 4.

We have set the Score-based Data Assimilation (SDA) method as our baseline, which is a state-of-the-art method for data assimilation with data-driven physical models independent of any classical assimilation algorithms. Both the SOAD model and the SDA baseline are pre-trained with the same dataset and network structures, and no additional training is performed during the assimilation process. Besides, it is worth noting that knowledge of physical dynamics have been implicitly embedded into the SDA or our SOAD model through the training process, and only the observational models are used during the assimilation.

We consider the following four assimilation settings.

- (I) **Assimilation with a background prior** at the first step, which aligns with practical real-world applications.
- (II) **Assimilation without any background prior**, where no background estimation is available at any step, and all physical prior knowledge is inferred from our trained SOAD model. This setting is more challenging and indicates the long-term performance of our model to some extent.
- (III) **Multi-modal assimilation** as an extension of setting (II), where observations coming from multiple observational operators are available simultaneously.
- (IV) **Assimilation with non-Gaussian distributions** as an extension of setting (I) and (II), where we set the background prior and the observational noise as non-Gaussian to study the robustness of our model.

##### 4.5.1. Assimilation with background prior

In practical applications, assimilation is usually performed sequentially once new observations become available, so we start with the classical assimilation settings that additionally provide the models with a background prior

for the first step. We choose the background prior as an ideal Gaussian perturbation with a known covariance  $(\sigma^b)^2 = 0.1^2$ . Formally, the observational model for the first step is modified as


$$\mathbf{y}_0 = \mathbf{x}_0 + \boldsymbol{\epsilon}_0, \quad \boldsymbol{\epsilon}_0 \sim \mathcal{N}(0, (\sigma^b)^2 \mathbf{I}), \quad (56)$$

where  $\mathbf{y}_0$  is the background prior with noise variance  $0.1^2$ . See fig. 4 (3rd and 4th rows) for visualization.


Our experiments start with the element-wise observations  $\mathcal{H}^e$  and  $\mathcal{H}^h$ . Besides, to explore the potential of our SOAD model in real applications, we also test it with the vorticity-to-velocity mapping  $\mathcal{H}^{v2v}$ . The averaged RMSEs over all the 9 time steps for the assimilated states are summarized in fig. 5. Our SOAD model performs slightly worse as the time interval  $N$  increases and the observational ratio  $p$  or  $(1/s)^2$  decreases, which aligns with our intuition that the assimilation task becomes more difficult with fewer observations. In the most challenging case, where the observations are available only every 8 time steps with a ratio of 1%, the RMSEs are around 0.2 and 0.3 for the easy and hard observations, respectively. This indicates that our method is still able to extract the physical features from rare observations to some extent. Another interesting observation is that the RMSEs for the hard observation  $\mathcal{H}^h$  are lower than those for  $\mathcal{H}^e$  and  $\mathcal{H}^{v2v}$ . One possible explanation is that when the observational data is sufficient for the model to capture the dynamics, observations from a harder operator  $\mathcal{H}^h$  may provide more detailed information since  $\mathcal{H}^h$  is much more sensitive to the input. In contrast, the SDA baseline (figs. 5c, 5f and 5i) shows a significant performance drop even when the observations are available everywhere for each time step. The inferior performance is likely due to the linearity assumption for the observation operator in the SDA formulations, which is unsuitable for highly nonlinear observations.

Meanwhile, fig. 5 includes an ablation study of our forward-diffusion corrector, abbreviated as “f.d.c.” in the captions, as well. To study the effectiveness of our forward-diffusion corrector, we have kept all the experimental settings the same as those for the SOAD and the SDA methods, but removed all the updates by the forward-diffusion corrector (line 12 and 16 in algorithm 1). By comparing figs. 5a, 5d and 5g with figs. 5b, 5e and 5h respectively we can conclude that the forward-diffusion corrector has played an important role in improving the assimilation performance. For the observations  $\mathcal{H}^e$  and  $\mathcal{H}^{v2v}$ , our SOAD model without the forward-diffusion cor-







(a) SOAD (ours) for  $\mathcal{H}^e$




(b) SOAD (no f.d.c.) for  $\mathcal{H}^e$




(c) SDA (baseline) for  $\mathcal{H}^e$




(d) SOAD (ours) for  $\mathcal{H}^h$




(e) SOAD (no f.d.c.) for  $\mathcal{H}^h$




(f) SDA (baseline) for  $\mathcal{H}^h$



(g) SOAD (ours) for  $\mathcal{H}^{v2v}$



(h) SOAD (no f.d.c.) for  $\mathcal{H}^{v2v}$



(i) SDA (baseline) for  $\mathcal{H}^{v2v}$

Figure 5: Averaged RMSEs for the assimilated states by using background prior.


rector has shown a similar pattern as the normal SOAD, which is to say, the denser and the more frequent observations we have for assimilation, the higher accuracy we can obtain. For the much harder observation  $\mathcal{H}^h$ , the forward-diffusion corrector has prevented divergence and thus exhibited its great potential to stabilize the generative process.

Additionally, the step-wise RMSEs for our SOAD model over the assimilation window are exhibited in fig. 6. Note that we start from a background prior as an initial estimate, which introduces much more physical information than the subsequent observational data. Although observations are available for some of the subsequent time steps, the physical information encoded in the observations may not be sufficiently informative to fully correct the trajectory as it evolves, particularly in the early stages of assimilation. This can result in increasing deviation from the ground truth at later time steps, especially if the dynamical system exhibits sensitivity to initial conditions or chaotic behavior. In both observation cases, The increments of RMSEs at observational ratios of 100% and 25% are mild, and no significant changes in RMSEs are observed when the easier  $\mathcal{H}^e$  is replaced with the harder  $\mathcal{H}^h$ . Moreover, relatively lower assimilation errors are obtained whenever observational data are available, supporting the idea that more observations enhance the assimilation process. By comparing results with the same observational ratios ( $p = (1/s)^2$ ), we may conclude that regular (uniform) observations are more beneficial for the assimilation process than irregular (random) ones. In addition, although the step-wise RMSEs increase along with the time steps, and the uncertainty of  $\mathcal{H}^{v2v}$  (figs. 6e and 6f) is slightly higher compared with  $\mathcal{H}^e$  (figs. 8a and 8b) and  $\mathcal{H}^h$  (figs. 8c and 8d), the subsequent experiments may offer some evidences that the assimilation process is unlikely to diverge.


#### 4.5.2. Assimilation without background prior

To explore the assimilation performance in the absence of any background prior, we remove the modification (56) and conduct the assimilation process solely from observational data. See the last two rows of fig. 4. Even if starting with a background prior, any assimilation method will gradually forget the initial state information as more observations are assimilated. Therefore, the results shown in this section are likely to reveal the long-term behaviors of the assimilation methods.


All assimilation results are recorded as heatmaps of RMSEs in fig. 7. In general, our SOAD model performs well for both the  $\mathcal{H}^e$  and  $\mathcal{H}^{v2v}$  observational operators, with slightly higher accuracy for the latter. The observation




(a) the easy observation  $\mathcal{H}^e$




(b) the easy observation  $\mathcal{H}^e$




(c) the hard observation  $\mathcal{H}^h$



(d) the hard observation  $\mathcal{H}^h$



(e) the vorticity-to-velocity  $\mathcal{H}^{v2v}$



(f) the vorticity-to-velocity  $\mathcal{H}^{v2v}$

Figure 6: Step-wise RMSEs of SOAD for assimilation with prior.

$\mathcal{H}^h$  appears to be the most challenging across all the testing observational operators. More specifically, as also discussed in section 4.5.1, when observations are densely available, a condition that is uncommon in practice, the observation  $\mathcal{H}^h$  appears to be more informative and helpful than  $\mathcal{H}^e$  and  $\mathcal{H}^{v2v}$ . However, under more typical conditions with partial observations, our SOAD model achieves the best performance with  $\mathcal{H}^{v2v}$ , followed  $\mathcal{H}^e$ . This may be attribute to the linearity of  $\mathcal{H}^{v2v}$ , despite its non-local nature and greater computational complexity.

As before, we present the step-wise RMSEs of our SOAD model in fig. 8. Compared to the results in fig. 6, the performance degrades when assimilating without a background prior—an expected outcome given the absence of an initial estimate at the first time step. Notably, the RMSE remains relatively stable across time, with only slight increases observed at the beginning and end of the assimilation window. This pattern is consistent with the behavior of smoothing methods, which have limited observational support near the temporal boundaries. Furthermore, the individual RMSE profiles across different experimental settings demonstrate that the SOAD model maintains consistent accuracy over time, providing strong evidence of its long-term stability and resistance to divergence.

To visualize the assimilated states, we have plotted the assimilated vorticity with random observational ratios  $p \in \{0.25, 0.0625, 0.01\}$  in figs. 9 to 11 for  $\mathcal{H}^e$ ,  $\mathcal{H}^h$  and  $\mathcal{H}^{v2v}$ , respectively. For clarity, we only include the vorticity and the associated observations for the first layer, as they contain more small-scale details. For the  $\mathcal{H}^e$  observation, our SOAD is capable of recovering most of the physical features until  $p$  reaches 1%, in which case the assimilated states still share much similarity with the ground truth. With the harder  $\mathcal{H}^h$ , the SOAD model fails to reconstruct the real system states when  $p \leq 0.0625 = (1/4)^2$ , consistent with the previous results shown in fig. 7b. Surprisingly, the assimilated states with rare observations do not collapse and still follow certain dynamics, suggesting that the lack of observational data, rather than incapacity of learning the dynamics, leads to the failure. Finally, when the observation is the vorticity-to-velocity mapping  $\mathcal{H}^{v2v}$ , the assimilated states are closer to the ground truth even when observations are available at only 1% of the grids. We believe that the success may result from the linearity of  $\mathcal{H}^{v2v}$ . The results indicate that our SOAD approach is stable when handling highly nonlinear observations in the long term performs well on non-element-wise observations such as the vorticity-to-velocity mapping  $\mathcal{H}^{v2v}$ , which is a promising sign for real applications.

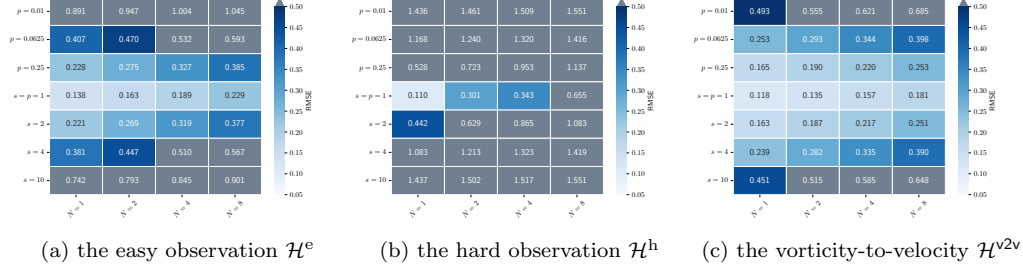


Figure 7: Averaged RMSEs of our SOAD without background prior.


#### 4.5.3. Assimilation with multiple observations

In this subsection, we discuss another important feature of our SOAD approach: the ability to handle observations from different modalities. Such a capability is crucial for real applications since the observations are usually collected from various sources like satellites, weather radars, and in-situ observations. Each source may provide different physical variables in different formats. We have tested our SOAD approach with various combinations of the three observations  $\mathcal{H}^e$ ,  $\mathcal{H}^h$  and  $\mathcal{H}^{v2v}$ . Same as the previous subsection, we do not add any background prior for the first step in order to study the long-term behaviors. The results are displayed in fig. 12.


As the collection of observations expands, the uncertainties of the assimilation errors shrink, and the averaged RMSEs decrease as well due to the additional information inferred from the observations. This phenomenon implies that our SOAD approach also follows the intuitive principle that more observations are beneficial for the assimilation process. Figure 12d gathers the best performances for each observation collections when  $s = p = 1$ , and the overlapping areas are marked with the corresponding observation combinations.

#### 4.5.4. Model robustness beyond Gaussian assumptions


All the previous experiments have employed Gaussian assumptions for both the background estimate (when available) and the observational noise. However, this idealization rarely holds in realistic applications. Commonly, during online deployment, the background uncertainty is either unknown or poorly characterized, and the observational noise may exhibit heavy tails or other non-Gaussian features due to sensor limitations and dynamic environmental factors. To assess the robustness of our proposed model under




(a) the easy observation  $\mathcal{H}^e$




(b) the easy observation  $\mathcal{H}^e$




(c) the hard observation  $\mathcal{H}^h$



(d) the hard observation  $\mathcal{H}^h$



(e) the vorticity-to-velocity  $\mathcal{H}^{v2v}$



(f) the vorticity-to-velocity  $\mathcal{H}^{v2v}$

Figure 8: Step-wise RMSEs of SOAD for assimilation without prior.






Figure 9: Visualization of the assimilated vorticities by our SOAD model with  $\mathcal{H}^e$ .




Figure 10: Visualization of the assimilated vorticities by our SOAD model with  $\mathcal{H}^h$ .




Figure 11: Visualization of the assimilated vorticities by our SOAD model with  $\mathcal{H}^{v2v}$ .



these more challenging but realistic conditions, we conduct additional experiments in which the Gaussian assumptions are deliberately violated. These experiments serve to evaluate the model’s adaptability and reliability in the presence of distributional uncertainty.

In practice, we cannot always expect the availability of the exact probability distribution for background prior. Background errors may result from many aspects, such as unresolved sub-scale features, temporal interpolations and physical parameterizations. To explore the impact of potential mismatch between our assumed background prior model (56) and the true underlying distribution, we design an experiment under the same settings as before, but replace the background prior with a sample generated from an unknown noise process



$$\mathbf{y}'_0 = \mathcal{M}^{\text{QG}}(\mathbf{x}_{-1} + \boldsymbol{\eta}), \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, 0.01^2 \mathbf{I}) \quad (57)$$

into the model instead of  $\mathbf{y}_0$ . Here,  $\mathcal{M}^{\text{QG}}$  is the forward propagation model of the QG equations, and  $\mathbf{x}_{-1}$  stands for the state at the previous time step. Comparisons between the step-wise RMSEs are displayed in fig. 13, where we fix  $N = 1$  and  $p = 0.25$ . Despite the discrepancy in the background prior, our SOAD model demonstrates a capacity for self-correction within a few assimilation steps, exhibiting its robustness and relative insensitivity to





(a) averaged RMSEs for collections containing  $\mathcal{H}^e$  (b) averaged RMSEs for collections containing  $\mathcal{H}^h$



(c) averaged RMSEs for collections containing  $\mathcal{H}^{2v}$  (d) The best RMSEs for different combinations of observations

Figure 12: RMSEs for multi-modal assimilation with our SOAD model. The error uncertainty introduced by different random seeds is marked on the top of each bar. Duplicate results may appear in multiple subplots for illustration purposes.




Figure 13: Step-wise RMSEs for SOAD with different data for background prior.

inaccuracies in the background prior distribution.

Another important source of uncertainty in data assimilation arises from the nature of observational noise, which is often assumed to be Gaussian for analytical convenience. However, in reality, observational errors can exhibit non-Gaussian characteristics due to factors such as unknown nonlinear sensor transformations, mixed error sources, or limitations in the statistical characterization of measurement systems. To evaluate the robustness of our assimilation algorithm under such conditions, we perform experiments using a non-Gaussian noise model for the observations. In particular, compared to eq. (51), we set

$$\mathbf{y}_k = \mathbf{S}_k \circ \mathcal{H}(\mathbf{x}_k) + \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{D}, \quad (58)$$

where the Gaussian noise assumption is replaced by a distribution  $\mathcal{D}$ , drawn from three types of element-wise noise models: Laplace distribution, Uniform distribution, and LogNormal distribution, whose density functions are visualized in fig. 14. Note that we have fixed appropriate parameters for each distribution so that all the distributions share the same mean and variance with our Gaussian assumption (51). The detailed parameters are listed in appendix Appendix C for reference.

The quantitative results for assimilation under different observational noise distributions are summarized in tables 3 and 4. Each table reports the step-wise RMSE (scaled by  $10^{-2}$ ) under various observation operators ( $\mathcal{H}^e$ ,  $\mathcal{H}^h$ , and  $\mathcal{H}^{v2v}$ ), with and without the use of a background prior. The first table corresponds to the setting  $N = 1$ ,  $p = 1$ , representing full spatio-temporal coverage, while the second table uses  $N = 2$ ,  $p = 0.25$ , introduc-




Figure 14: The probabilistic density functions for various observational noise.

dist. \ obs.	assimilation w/ prior			assimilation w/o prior		
	$\mathcal{H}^e$	$\mathcal{H}^h$	$\mathcal{H}^{v2v}$	$\mathcal{H}^e$	$\mathcal{H}^h$	$\mathcal{H}^{v2v}$
Laplace	$8.455 \pm 0.46$	$6.441 \pm 0.36$	$10.31 \pm 0.49$	$13.82 \pm 0.53$	$14.44 \pm 1.2$	$11.85 \pm 0.76$
Uniform	$8.449 \pm 0.46$	$6.441 \pm 0.36$	$10.30 \pm 0.49$	$13.84 \pm 0.53$	$15.91 \pm 5.6$	$11.85 \pm 0.76$
LogNormal-0.1	$8.456 \pm 0.47$	$6.444 \pm 0.36$	$10.30 \pm 0.49$	$13.82 \pm 0.53$	$10.14 \pm 3.5$	$11.85 \pm 0.75$
LogNormal-0.2	$8.457 \pm 0.47$	$6.444 \pm 0.36$	$10.31 \pm 0.49$	$13.82 \pm 0.53$	$13.02 \pm 5.0$	$11.85 \pm 0.75$
LogNormal-0.5	$8.460 \pm 0.47$	$6.445 \pm 0.36$	$10.31 \pm 0.49$	$13.83 \pm 0.53$	$12.07 \pm 6.1$	$11.85 \pm 0.75$
LogNormal-1.0	$8.529 \pm 0.48$	$6.527 \pm 0.37$	$10.36 \pm 0.49$	$14.20 \pm 0.65$	$14.07 \pm 1.6$	$12.01 \pm 0.78$
<b>Gaussian</b>	<b><math>8.457 \pm 0.47</math></b>	<b><math>6.444 \pm 0.36</math></b>	<b><math>10.30 \pm 0.49</math></b>	<b><math>13.82 \pm 0.53</math></b>	<b><math>12.93 \pm 3.2</math></b>	<b><math>11.85 \pm 0.75</math></b>

Table 3: RMSE ( $\times 10^{-2}$ ) with various observational noise distributions,  $N = 1$ ,  $p = 1$ .

ing a more difficult case. Across both settings, we observe that the performance of the SOAD model remains stable regardless of the specific noise distribution. In particular, the RMSE values under Laplace, Uniform, and moderately skewed LogNormal noise (e.g., LogNormal-0.1 to LogNormal-0.5) closely match those obtained under Gaussian noise, indicating that the model is robust to deviations from the Gaussian assumption. Only when the LogNormal noise becomes significantly skewed (LogNormal-1.0) do we see a slight degradation in performance, especially in the absence of a background prior. These results collectively suggest that SOAD is largely insensitive to the choice of noise distribution and maintains high assimilation accuracy even under non-Gaussian observational errors.

dist. \ obs.	assimilation w/ prior			assimilation w/o prior		
	$\mathcal{H}^e$	$\mathcal{H}^h$	$\mathcal{H}^{v2v}$	$\mathcal{H}^e$	$\mathcal{H}^h$	$\mathcal{H}^{v2v}$
Laplace	$11.96^{\pm 1.1}$	$12.33^{\pm 1.3}$	$13.05^{\pm 1.2}$	$27.50^{\pm 1.0}$	$72.34^{\pm 9.3}$	$18.99^{\pm 0.80}$
Uniform	$11.96^{\pm 1.1}$	$12.33^{\pm 1.3}$	$13.05^{\pm 1.2}$	$27.51^{\pm 1.0}$	$72.41^{\pm 9.4}$	$18.99^{\pm 0.80}$
LogNormal-0.1	$11.96^{\pm 1.1}$	$12.34^{\pm 1.3}$	$13.05^{\pm 1.2}$	$27.50^{\pm 1.0}$	$71.99^{\pm 9.2}$	$18.98^{\pm 0.80}$
LogNormal-0.2	$11.96^{\pm 1.1}$	$12.34^{\pm 1.3}$	$13.05^{\pm 1.2}$	$27.50^{\pm 1.0}$	$71.45^{\pm 8.4}$	$18.98^{\pm 0.81}$
LogNormal-0.5	$11.96^{\pm 1.1}$	$12.34^{\pm 1.3}$	$13.05^{\pm 1.2}$	$27.51^{\pm 1.0}$	$72.38^{\pm 9.0}$	$18.98^{\pm 0.81}$
LogNormal-1.0	$12.03^{\pm 1.1}$	$12.44^{\pm 1.3}$	$13.09^{\pm 1.2}$	$27.67^{\pm 1.0}$	$72.93^{\pm 7.4}$	$19.12^{\pm 0.84}$
<b>Gaussian</b>	<b><math>11.96^{\pm 1.2}</math></b>	<b><math>12.34^{\pm 1.3}</math></b>	<b><math>13.05^{\pm 1.2}</math></b>	<b><math>27.49^{\pm 0.53}</math></b>	<b><math>72.14^{\pm 9.2}</math></b>	<b><math>18.98^{\pm 0.80}</math></b>

Table 4: RMSE ( $\times 10^{-2}$ ) with various observational noise distributions,  $N = 2$ ,  $p = 0.25$ .

## 5. Conclusion and future work

In this study, we have proposed the State-Observation Augmented Diffusion (SOAD) model, a novel data-driven data assimilation method that operates independently of classical algorithms. Unlike many existing approaches, SOAD does not assume linearity in the physical or observational models, yet it can still recover the posterior distribution under mild conditions, suggesting a theoretical advantage. Experiments on the two-layer quasi-geostrophic model with various settings have demonstrated its good performance and potential for real-world applications.

A key feature of SOAD is its use of augmented dynamical models. Although the original assimilation problem may be highly nonlinear, we work with a linearized equivalent by augmenting the system, allowing us to focus on linear observational models. By leveraging a score-based generative framework, SOAD learns the prior distribution of physical states without explicit knowledge of the physical model, embedding relevant physics implicitly. We also propose a deterministic approach to estimate the adversarial gradient and introduce a forward-diffusion corrector to stabilize the assimilation process.

Further improvements and extensions of our approach are also necessary. First, analyzing convergence rates and conditions in the presence of training errors would be beneficial. Robustness and stability are particularly important in real-world applications, and broader testing on both idealized and real-world models is needed to fully assess performance. While our current framework assumes additive observational noise, extending it to handle non-additive or more complex noise structures would be a valuable direction for

future work, as such cases may arise in certain practical scenarios. Besides, adapting our SOAD approach to more complex tasks beyond data assimilation could also be an exciting avenue for future research.

### Data availability

All the codes for data generation, network training and assimilation experiments are available at <https://github.com/zylipku/SOAD>.

### Acknowledgments

Zhuoyuan Li and Pingwen Zhang are supported in part by the National Natural Science Foundation of China (No. 12288101). Bin Dong is supported in part by the New Cornerstone Investigator Program.

### Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT in order to polish the written text for spelling and grammar. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

### Appendix A. Configurations of the QG model

We list the parameters (as defaults in `pyqg-jax`<sup>4</sup>) we have used to generate the QG dataset.

- Domain size:  $10^6 \times 10^6$ ;
- Linear drag in lower layer:  $r_{ek} = 5.767 \times 10^{-7}$ ;
- Gravity:  $g = 9.81$ ;
- Gradient of Coriolis parameter:  $\beta = 1.5 \times 10^{-11}$ ;
- Deformation radius:  $r_d = 1.5 \times 10^4$ ;

---

<sup>4</sup>[https://pyqg-jax.readthedocs.io/en/latest/reference.models.qg\\_model.html](https://pyqg-jax.readthedocs.io/en/latest/reference.models.qg_model.html)

- Layer thickness:  $H_1 = 500$ ,  $H_2 = 2000$ ;
- Upper/Lower layer flow:  $U_1 = 0.025$ ,  $U_2 = 0$ ;
- $F_1 = r_d^{-2}/(1 + \delta)$  and  $F_2 = \delta F_1$ , where the layer thickness ratio  $\delta = H_1/H_2 = 0.25$ .

To obtain the small-scale dissipation “ssd” mentioned in eq. (45), a highly-selective exponential filter

$$E_f = \begin{cases} \exp[-23.6(\kappa^* - \kappa_c)^4], & \kappa \geq \kappa_c \\ 1, & \kappa < \kappa_c \end{cases}$$

is employed as a multiplier on the spectral domain. The term  $\kappa^*$  stands for the non-dimensional wavenumber, and the cutoff  $\kappa_c$  is set as 65% of the Nyquist scale  $\kappa_{ny}^* = \pi$ . Readers may refer to `pyqg`<sup>5</sup> for the detailed numerical scheme.

## Appendix B. Network architecture

The denoising network  $\varepsilon_\theta(\mathbf{z}_t, t)$  is implemented using a U-Net architecture proposed in [48], and we reproduce it in fig. B.15 for completeness, where we have attached the channel numbers for each module.

## Appendix C. Various observational noise distributions

Let  $p(x)$  denote the probabilistic density function. To ensure all the observational noise distributions share the same mean and variance of our Gaussian assumptions  $\mathcal{N}(0, (\sigma^o)^2)$  defined in eq. (51), we employ the following settings.

- Laplace:

$$p_{\text{Laplace}}(x) = \frac{1}{2b} \exp\left(-\frac{|x - \mu|}{b}\right),$$

where  $\mu = 0$ ,  $b = \sigma^o/\sqrt{2}$ ;

---

<sup>5</sup>[https://pyqg.readthedocs.io/en/latest/equations/notation\\_twolayer\\_model.html](https://pyqg.readthedocs.io/en/latest/equations/notation_twolayer_model.html)

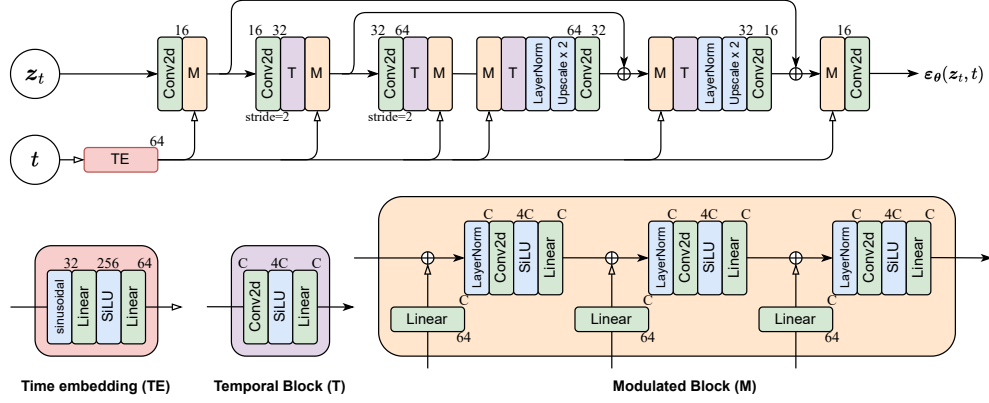


Figure B.15: The U-Net architecture for the denoising network  $\epsilon_\theta(\mathbf{z}_t, t)$ .

- Uniform:

$$p_{\text{Uniform}}(x) = \begin{cases} \frac{1}{2d}, & |x| < d, \\ 0, & |x| \geq d, \end{cases}$$

where  $d = \sqrt{3}\sigma^0$ ;

- LogNormal: Let  $Z \sim \mathcal{N}(\mu, s^2)$ , then the distribution of  $\exp Z$  is defined to follow the LogNormal distribution with parameter  $\mu$  and  $s$ . To make the random variable unbiased, we introduce an additional shift term  $c$ , which means we use  $(\exp Z + c)$  to generate our observational noise. For any fixed  $s$ , we define

$$c = -\sigma^0 (\exp(s^2) - 1)^{-1/2}, \quad \mu = -\frac{s^2}{2} + \log(-c).$$

The value of  $s$  is appended to the distribution name in the main text.

## References

- [1] R. E. Kalman, A new approach to linear filtering and prediction problems, Journal of Basic Engineering 82 (1) (1960) 35–45. [doi:10.1115/1.3662552](https://doi.org/10.1115/1.3662552).
- [2] A. H. Jazwinski, Stochastic processes and filtering theory, Courier Corporation, 2007.

- [3] B. R. Hunt, E. J. Kostelich, I. Szunyogh, Efficient data assimilation for spatiotemporal chaos: A local ensemble transform kalman filter, *Phys. D* 230 (1-2) (2007) 112–126. [doi:10.1016/j.physd.2006.11.008](https://doi.org/10.1016/j.physd.2006.11.008).
- [4] A. M. Clayton, A. C. Lorenc, D. M. Barker, Operational implementation of a hybrid ensemble/4d-var global data assimilation system at the met office, *Quarterly Journal of the Royal Meteorological Society* 139 (675) (2013) 1445–1461. [doi:10.1002/qj.2054](https://doi.org/10.1002/qj.2054).
- [5] S. Zhu, B. Wang, L. Zhang, J. Liu, Y. Liu, J. Gong, S. Xu, Y. Wang, W. Huang, L. Liu, Y. He, X. Wu, B. Zhao, F. Chen, A 4DVar-based ensemble four-dimensional variational (En4DVar) hybrid data assimilation system for global NWP: System description and primary tests, *Journal of Advances in Modeling Earth Systems* 14 (8) (2022) e2022MS003023. [doi:10.1029/2022MS003023](https://doi.org/10.1029/2022MS003023).
- [6] M. Asch, M. Asch, M. Bocquet, M. Nodet, *Data Assimilation: Methods, Algorithms, and Applications*, no. 11 in *Fundamentals of Algorithms*, SIAM, Society for Industrial and Applied Mathematics, 2016.
- [7] G. Evensen, F. C. Vossepoel, P. J. Van Leeuwen, *Data Assimilation Fundamentals: A Unified Formulation of the State and Parameter Estimation Problem*, Springer Textbooks in Earth Sciences, Geography and Environment, Springer International Publishing, 2022. [doi:10.1007/978-3-030-96709-3](https://doi.org/10.1007/978-3-030-96709-3).
- [8] R. N. Bannister, A review of operational methods of variational and ensemble-variational data assimilation, *Quarterly Journal of the Royal Meteorological Society* 143 (703) (2017) 607–633. [doi:10.1002/qj.2982](https://doi.org/10.1002/qj.2982).
- [9] F. Rabier, Z. Liu, Variational data assimilation: theory and overview, in: *Proc. ECMWF Seminar on Recent Developments in Data Assimilation for Atmosphere and Ocean*, Reading, UK, September 8–12, 2003, pp. 29–43.
- [10] M. G. Schultz, C. Betancourt, B. Gong, F. Kleinert, M. Langguth, L. H. Leufen, A. Mozaffari, S. Stadler, Can deep learning beat numerical weather prediction?, *Philos. Trans. Roy. Soc. A* 379 (2194) (2021) 20200097. [doi:10.1098/rsta.2020.0097](https://doi.org/10.1098/rsta.2020.0097).



- [11] A. J. Geer, Learning earth system models from observations: machine learning or data assimilation?, *Philos. Trans. Roy. Soc. A* 379 (2194) (2021) 20200089. [doi:10.1098/rsta.2020.0089](https://doi.org/10.1098/rsta.2020.0089).
- [12] A. Voulodimos, N. Doulamis, A. Doulamis, E. Protopapadakis, Deep learning for computer vision: A brief review, *Computational Intelligence and Neuroscience* 2018 (1) (2018) 7068349. [doi:10.1155/2018/7068349](https://doi.org/10.1155/2018/7068349).
- [13] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al., Segment anything, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4015–4026.
- [14] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [15] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al., Llama: Open and efficient foundation language models, *arXiv preprint arXiv:2302.13971* (2023).
- [16] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, et al., A survey of large language models, *arXiv preprint arXiv:2303.18223* (2023).
- [17] S. Chen, G. Long, J. Jiang, D. Liu, C. Zhang, Foundation models for weather and climate data understanding: A comprehensive survey, *arXiv preprint arXiv:2312.03014* (2023).
- [18] T. Kurth, S. Subramanian, P. Harrington, J. Pathak, M. Mardani, D. Hall, A. Miele, K. Kashinath, A. Anandkumar, Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive fourier neural operators, in: *Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '23*, Association for Computing Machinery, New York, NY, USA, 2023. [doi:10.1145/3592979.3593412](https://doi.org/10.1145/3592979.3593412).
- [19] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, Q. Tian, Accurate medium-range global weather forecasting with 3d neural networks, *Nature* 619 (7970) (2023) 533–538. [doi:10.1038/s41586-023-06185-3](https://doi.org/10.1038/s41586-023-06185-3).

- [20] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, P. Battaglia, Learning skillful medium-range global weather forecasting, *Science* 382 (6677) (2023) 1416–1421. [doi:10.1126/science.adf2336](https://doi.org/10.1126/science.adf2336).
- [21] L. Chen, X. Zhong, F. Zhang, Y. Cheng, Y. Xu, Y. Qi, H. Li, Fuxi: A cascade machine learning forecasting system for 15-day global weather forecast, *npj Climate and Atmospheric Science* 6 (1) (2023) 190. [doi:10.1038/s41612-023-00512-1](https://doi.org/10.1038/s41612-023-00512-1).
- [22] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780. [doi:10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735).
- [23] K. Cho, B. Van Merriënboer, D. Bahdanau, Y. Bengio, On the properties of neural machine translation: Encoder-decoder approaches, *arXiv preprint arXiv:1409.1259* (2014).
- [24] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, W.-c. Woo, Convolutional lstm network: A machine learning approach for precipitation nowcasting, in: *Advances in Neural Information Processing Systems*, Vol. 28, 2015.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: *Advances in Neural Information Processing Systems*, Vol. 30, 2017.
- [26] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, B. Guo, Swin transformer: Hierarchical vision transformer using shifted windows, in: *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10012–10022.
- [27] L. Ruthotto, E. Haber, An introduction to deep generative modeling, *GAMM-Mitteilungen* 44 (2) (2021) e202100008. [doi:10.1002/gamm.202100008](https://doi.org/10.1002/gamm.202100008).
- [28] S. Bond-Taylor, A. Leach, Y. Long, C. G. Willcocks, Deep generative modelling: A comparative review of VAEs, GANs, normalizing flows, energy-based and autoregressive models, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44 (11) (2022) 7327–7347. [doi:10.1109/TPAMI.2021.3116668](https://doi.org/10.1109/TPAMI.2021.3116668).

- [29] F. Daum, J. Huang, Particle flow with non-zero diffusion for nonlinear filters, bayesian decisions and transport, in: Signal and Data Processing of Small Targets 2013, Vol. 8857, SPIE, 2013, pp. 64–72.
- [30] F. Daum, J. Huang, A. Noushin, New theory and numerical results for gromov’s method for stochastic particle flow filters, in: 2018 21st International Conference on Information Fusion (FUSION), IEEE, 2018, pp. 108–115.
- [31] C.-C. Hu, P. J. Van Leeuwen, A particle flow filter for high-dimensional system applications, Quarterly Journal of the Royal Meteorological Society 147 (737) (2021) 2352–2374. doi:[10.1002/qj.4028](https://doi.org/10.1002/qj.4028).
- [32] G. Evensen, P. J. van Leeuwen, An ensemble kalman smoother for non-linear dynamics, Monthly Weather Review 128 (6) (2000) 1852 – 1867. doi:[10.1175/1520-0493\(2000\)128<1852:AEKSFN>2.0.CO;2](https://doi.org/10.1175/1520-0493(2000)128<1852:AEKSFN>2.0.CO;2).
- [33] M. Bocquet, P. Sakov, An iterative ensemble kalman smoother, Quarterly Journal of the Royal Meteorological Society 140 (682) (2014) 1521–1535. doi:[10.1002/qj.2236](https://doi.org/10.1002/qj.2236).
- [34] S. Cheng, C. Quilodr  n-Casas, S. Ouala, A. Farchi, C. Liu, P. Tandeo, R. Fablet, D. Lucor, B. Iooss, J. Brajard, D. Xiao, T. Janjic, W. Ding, Y. Guo, A. Carrassi, M. Bocquet, R. Arcucci, Machine learning with data assimilation and uncertainty quantification for dynamical systems: A review, IEEE/CAA Journal of Automatica Sinica 10 (6) (2023) 1361–1387. doi:[10.1109/JAS.2023.123537](https://doi.org/10.1109/JAS.2023.123537).
- [35] J. Brajard, A. Carrassi, M. Bocquet, L. Bertino, Combining data assimilation and machine learning to infer unresolved scale parametrization, Philosophical Transactions of the Royal Society A 379 (2194) (2021) 20200086. doi:[10.1098/rsta.2020.0086](https://doi.org/10.1098/rsta.2020.0086).
- [36] A. Chattopadhyay, E. Nabizadeh, E. Bach, P. Hassanzadeh, Deep learning-enhanced ensemble-based data assimilation for high-dimensional nonlinear dynamical systems, J. Comput. Phys. 477 (2023) 111918. doi:[10.1016/j.jcp.2023.111918](https://doi.org/10.1016/j.jcp.2023.111918).
- [37] M. A. E. R. Hammoud, N. Raboudi, E. S. Titi, O. Knio, I. Hoteit, Data assimilation in chaotic systems using deep reinforcement learn-

- ing, *Journal of Advances in Modeling Earth Systems* 16 (8) (2024) e2023MS004178. [doi:10.1029/2023MS004178](https://doi.org/10.1029/2023MS004178).
- [38] M. Amendola, R. Arcucci, L. Mottet, C. Q. Casas, S. Fan, C. Pain, P. Linden, Y.-K. Guo, Data assimilation in the latent space of a convolutional autoencoder, in: M. Paszynski, D. Kranzlmüller, V. V. Krzhizhanovskaya, J. J. Dongarra, P. M. Soot (Eds.), *Computational Science – ICCS 2021*, Springer International Publishing, Cham, 2021, pp. 373–386.
  - [39] S. Pawar, O. San, Equation-free surrogate modeling of geophysical flows at the intersection of machine learning and data assimilation, *Journal of Advances in Modeling Earth Systems* 14 (11) (2022) e2022MS003170. [doi:10.1029/2022MS003170](https://doi.org/10.1029/2022MS003170).
  - [40] S. Cheng, J. Chen, C. Anastasiou, P. Angeli, O. K. Matar, Y.-K. Guo, C. C. Pain, R. Arcucci, Generalised latent assimilation in heterogeneous reduced spaces with machine learning surrogate models, *J. Sci. Comput.* 94 (1) (2023). [doi:10.1007/s10915-022-02059-4](https://doi.org/10.1007/s10915-022-02059-4).
  - [41] M. Peyron, A. Fillion, S. Gürol, V. Marchais, S. Gratton, P. Boudier, G. Goret, Latent space data assimilation by using deep learning, *Quarterly Journal of the Royal Meteorological Society* 147 (740) (2021) 3759–3777. [doi:10.1002/qj.4153](https://doi.org/10.1002/qj.4153).
  - [42] S. G. Penny, T. A. Smith, T.-C. Chen, J. A. Platt, H.-Y. Lin, M. Goodliff, H. D. I. Abarbanel, Integrating recurrent neural networks with data assimilation for scalable data-driven state estimation, *Journal of Advances in Modeling Earth Systems* 14 (3) (2022) e2021MS002843. [doi:10.1029/2021MS002843](https://doi.org/10.1029/2021MS002843).
  - [43] Z. Li, B. Dong, P. Zhang, Latent assimilation with implicit neural representations for unknown dynamics, *J. Comput. Phys.* 506 (2024) 112953. [doi:10.1016/j.jcp.2024.112953](https://doi.org/10.1016/j.jcp.2024.112953).
  - [44] F. Bao, Z. Zhang, G. Zhang, A score-based filter for nonlinear data assimilation, *Journal of Computational Physics* 514 (2024) 113207. [doi:10.1016/j.jcp.2024.113207](https://doi.org/10.1016/j.jcp.2024.113207).

- [45] F. Bao, Z. Zhang, G. Zhang, An ensemble score filter for tracking high-dimensional nonlinear dynamical systems, *Computer Methods in Applied Mechanics and Engineering* 432 (2024) 117447. [doi:10.1016/j.cma.2024.117447](https://doi.org/10.1016/j.cma.2024.117447).
- [46] P. Si, P. Chen, Latent-EnSF: A latent ensemble score filter for high-dimensional data assimilation with sparse observation data, *arXiv preprint arXiv:2409.00127* (2024).
- [47] F. Rozet, G. Louppe, Score-based data assimilation, in: *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023.
- [48] F. Rozet, G. Louppe, Score-based data assimilation for a two-layer quasi-geostrophic model, *arXiv preprint arxiv:2310.01853* (2023).
- [49] A. M. Clayton, A. C. Lorenc, D. M. Barker, Operational implementation of a hybrid ensemble/4D-Var global data assimilation system at the met office, *Quarterly Journal of the Royal Meteorological Society* 139 (675) (2013) 1445–1461. [doi:10.1002/qj.2054](https://doi.org/10.1002/qj.2054).
- [50] M. Bonavita, E. Hólm, L. Isaksen, M. Fisher, The evolution of the ECMWF hybrid data assimilation system, *Quarterly Journal of the Royal Meteorological Society* 142 (694) (2016) 287–303. [doi:doi.org/10.1002/qj.2652](https://doi.org/10.1002/qj.2652).
- [51] A. C. Lorenc, N. E. Bowler, A. M. Clayton, S. R. Pring, D. Fairbairn, Comparison of hybrid-4DVar and hybrid-4DVar data assimilation methods for global NWP, *Monthly Weather Review* 143 (1) (2015) 212 – 229. [doi:10.1175/MWR-D-14-00195.1](https://doi.org/10.1175/MWR-D-14-00195.1).
- [52] M. Zhang, F. Zhang, E4DVar: Coupling an ensemble kalman filter with four-dimensional variational data assimilation in a limited-area weather prediction model, *Monthly Weather Review* 140 (2) (2012) 587 – 600. [doi:10.1175/MWR-D-11-00023.1](https://doi.org/10.1175/MWR-D-11-00023.1).
- [53] E. Andersson, J. Pailleux, J.-N. Thépaut, J. R. Eyre, A. P. McNally, G. A. Kelly, P. Courtier, Use of cloud-cleared radiances in three/four-dimensional variational data assimilation, *Quarterly Journal of the Royal Meteorological Society* 120 (517) (1994) 627–653. [doi:10.1002/qj.49712051707](https://doi.org/10.1002/qj.49712051707).

- [54] J.-N. Thépaut, P. Courtier, G. Belaud, G. Lemaître, Dynamical structure functions in a four-dimensional variational assimilation: A case study, *Quarterly Journal of the Royal Meteorological Society* 122 (530) (1996) 535–561. [doi:10.1002/qj.49712253012](https://doi.org/10.1002/qj.49712253012).
- [55] D. P. Kingma, M. Welling, Auto-encoding variational bayes, *arXiv preprint arXiv:1312.6114* (2013).
- [56] C. Doersch, Tutorial on variational autoencoders, *arXiv preprint arXiv:1606.05908* (2016).
- [57] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial networks, *Communications of the ACM* 63 (11) (2020) 139–144. [doi:10.1145/3422622](https://doi.org/10.1145/3422622).
- [58] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, A. A. Bharath, Generative adversarial networks: An overview, *IEEE Signal Processing Magazine* 35 (1) (2018) 53–65. [doi:10.1109/MSP.2017.2765202](https://doi.org/10.1109/MSP.2017.2765202).
- [59] L. Dinh, D. Krueger, Y. Bengio, Nice: Non-linear independent components estimation, *arXiv preprint arXiv:1410.8516* (2014).
- [60] D. Rezende, S. Mohamed, Variational inference with normalizing flows, in: *International conference on machine learning*, PMLR, 2015, pp. 1530–1538.
- [61] I. Kobyzev, S. J. Prince, M. A. Brubaker, Normalizing flows: An introduction and review of current methods, *IEEE transactions on pattern analysis and machine intelligence* 43 (11) (2020) 3964–3979. [doi:10.1109/TPAMI.2020.2992934](https://doi.org/10.1109/TPAMI.2020.2992934).
- [62] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, O. Winther, Ladder variational autoencoders, in: *Advances in Neural Information Processing Systems*, Vol. 29, 2016.
- [63] C. Luo, Understanding diffusion models: A unified perspective, *arXiv preprint arXiv:2208.11970* (2022).

- [64] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, W. Zhang, B. Cui, M.-H. Yang, Diffusion models: A comprehensive survey of methods and applications, *ACM Computing Surveys* 56 (4) (2023) 1–39. doi:10.1145/3626235.
- [65] Y. Yang, M. Jin, H. Wen, C. Zhang, Y. Liang, L. Ma, Y. Wang, C. Liu, B. Yang, Z. Xu, et al., A survey on diffusion models for time series and spatio-temporal data, *arXiv preprint arXiv:2404.18886* (2024).
- [66] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: *Advances in Neural Information Processing Systems*, Vol. 33, 2020, pp. 6840–6851.
- [67] J. Song, C. Meng, S. Ermon, [Denoising diffusion implicit models](#), in: *International Conference on Learning Representations*, 2021.  
URL <https://openreview.net/forum?id=St1giarCHLP>
- [68] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10684–10695.
- [69] U. Singer, A. Polyak, T. Hayes, X. Yin, J. An, S. Zhang, Q. Hu, H. Yang, O. Ashual, O. Gafni, et al., [Make-a-video: Text-to-video generation without text-video data](#), in: *International Conference on Learning Representations*, 2023.  
URL <https://openreview.net/forum?id=nJfylDvgzlq>
- [70] G. Daras, H. Chung, C.-H. Lai, Y. Mitsufuji, J. C. Ye, P. Milanfar, A. G. Dimakis, M. Delbracio, A survey on diffusion models for inverse problems, *arXiv preprint arXiv:2410.00083* (2024).
- [71] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, J. C. Ye, [Diffusion posterior sampling for general noisy inverse problems](#), in: *International Conference on Learning Representations*, 2023.  
URL <https://openreview.net/forum?id=OnD9zGAGT0k>
- [72] J. Li, C. Wang, DiffFPR: Diffusion prior for oversampled Fourier phase retrieval, in: *Forty-first International Conference on Machine Learning*, 2024.

- [73] F. Wang, X. Huang, T. Alkhalifah, Controllable seismic velocity synthesis using generative diffusion models, *Journal of Geophysical Research: Machine Learning and Computation* 1 (3) (2024) e2024JH000153. doi:[10.1029/2024JH000153](https://doi.org/10.1029/2024JH000153).
- [74] L. Baldassari, A. Siahkoohi, J. Garnier, K. Solna, M. V. de Hoop, Conditional score-based diffusion models for bayesian inference in infinite dimensions, *Advances in Neural Information Processing Systems* 36 (2023) 24262–24290.
- [75] F. Wang, X. Huang, T. A. Alkhalifah, A prior regularized full waveform inversion using generative diffusion models, *IEEE transactions on geoscience and remote sensing* 61 (2023) 1–11. doi:[10.1109/TGRS.2023.3337014](https://doi.org/10.1109/TGRS.2023.3337014).
- [76] J.-H. Bastek, D. M. Kochmann, Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models, *Nature Machine Intelligence* 5 (12) (2023) 1466–1475. doi:[10.1038/s42256-023-00762-x](https://doi.org/10.1038/s42256-023-00762-x).
- [77] F. Mazé, F. Ahmed, Diffusion models beat GANs on topology optimization, in: *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37, 2023, pp. 9108–9116.
- [78] G. Giannone, A. Srivastava, O. Winther, F. Ahmed, Aligning optimization trajectories with diffusion models for constrained design generation, *Advances in Neural Information Processing Systems* 36 (2023) 51830–51861.
- [79] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, B. Poole, [Score-based generative modeling through stochastic differential equations](#), in: *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxtIG12RRHS>
- [80] B. K. W. Lahoz, R. Menard, *Data assimilation*, Springer, 2010.
- [81] S. Särkkä, A. Solin, *Applied Stochastic Differential Equations*, Institute of Mathematical Statistics Textbooks, Cambridge University Press, 2019. doi:[10.1017/9781108186735](https://doi.org/10.1017/9781108186735).



- [82] B. D. Anderson, Reverse-time diffusion equation models, *Stochastic Process. Appl.* 12 (3) (1982) 313–326. [doi:10.1016/0304-4149\(82\)90051-5](https://doi.org/10.1016/0304-4149(82)90051-5).
- [83] A. DasGupta, *Asymptotic theory of statistics and probability*, Vol. 180, Springer, 2008.
- [84] A. Hyvärinen, P. Dayan, Estimation of non-normalized statistical models by score matching., *Journal of Machine Learning Research* 6 (4) (2005).
- [85] P. Vincent, A connection between score matching and denoising autoencoders, *Neural Comput.* 23 (7) (2011) 1661–1674. [doi:10.1162/NECO\\_a\\_00142](https://doi.org/10.1162/NECO_a_00142).
- [86] Q. Zhang, Y. Chen, Fast sampling of diffusion models with exponential integrator, *arXiv preprint arXiv:2204.13902* (2022).
- [87] P. Dhariwal, A. Nichol, Diffusion models beat gans on image synthesis, in: *Advances in Neural Information Processing Systems*, Vol. 34, 2021, pp. 8780–8794.
- [88] J. Ho, T. Salimans, Classifier-free diffusion guidance, in: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [89] X. Meng, Y. Kabashima, Diffusion model based posterior sampling for noisy linear inverse problems, *arXiv preprint arXiv:2211.12343* (2022).
- [90] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, *JAX: composable transformations of Python+NumPy programs* (2018).  
URL <http://github.com/google/jax>
- [91] R. Abernathey, rochanotes, A. Ross, M. Jansen, Z. Li, F. J. Poulin, N. C. Constantinou, A. Sinha, D. Balwada, SalahKouhen, S. Jones, C. B. Rocha, C. L. P. Wolfe, C. Meng, H. van Kemenade, J. Bourbeau, J. Penn, J. Busecke, M. Bueti, Tobias, *pyqg/pyqg: v0.7.2* (may 2022).  
[doi:10.5281/zenodo.6563667](https://doi.org/10.5281/zenodo.6563667).

- [92] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: International Conference on Learning Representations, 2015.