



RESEARCH ARTICLE

10.1029/2025MS005324

Guided Unconditional and Conditional Generative Models for Super-Resolution and Inference of Quasi-Geostrophic Turbulence

 Anantha Narayanan Suresh Babu^{1,2} , Akhil Sadam² , and Pierre F. J. Lermusiaux^{1,2} 
¹Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA, ²Center for Computational Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

Key Points:

- We benchmark four diffusion models (DMs) for sparse super-resolution and inference of forced 2D quasi-geostrophic turbulence in the β -plane
- Conditional DMs remain stable, predict correct statistics, and generalize even with sparse and gappy coarse observations
- Guided unconditional models may ineffectively propagate information to unobserved regions when observation operators are sparse and local

Correspondence to:

 P. F. J. Lermusiaux,
pierrfel@mit.edu

Citation:

 Suresh Babu, A. N., Sadam, A., & Lermusiaux, P. F. J. (2026). Guided unconditional and conditional generative models for super-resolution and inference of Quasi-Geostrophic turbulence. *Journal of Advances in Modeling Earth Systems*, 18, e2025MS005324. <https://doi.org/10.1029/2025MS005324>

 Received 30 JUN 2025
 Accepted 23 FEB 2026

Author Contributions:

Conceptualization: Anantha Narayanan Suresh Babu, Akhil Sadam, Pierre F. J. Lermusiaux
Formal analysis: Anantha Narayanan Suresh Babu, Akhil Sadam, Pierre F. J. Lermusiaux
Funding acquisition: Pierre F. J. Lermusiaux
Investigation: Anantha Narayanan Suresh Babu, Akhil Sadam, Pierre F. J. Lermusiaux
Methodology: Anantha Narayanan Suresh Babu, Akhil Sadam, Pierre F. J. Lermusiaux
Project administration: Pierre F. J. Lermusiaux

Abstract Typically, numerical simulations of Earth systems are coarse, and Earth observations are sparse and gappy. We apply four generative diffusion modeling approaches to super-resolution and inference of forced two-dimensional quasi-geostrophic turbulence on the β -plane from coarse, sparse, and gappy observations. Two guided approaches minimally adapt a pre-trained unconditional model: SDEdit modifies the initial condition, and Diffusion Posterior Sampling (DPS) modifies the reverse diffusion process's score. Two conditional approaches, a vanilla variant and classifier-free guidance, require training with paired high-resolution and observation data. We consider multiple test cases spanning: two regimes, eddy and anisotropic-jet turbulence; two Reynolds numbers, 10^3 and 10^4 ; and two observation types, $4\times$ coarse-resolution fields and coarse, sparse and gappy observations. Our comprehensive skill metrics include norms of the reconstructed vorticity fields, turbulence statistical quantities, and quantifications of the super-resolved probabilistic ensembles and their errors. We also study the sensitivity to tuning parameters such as guidance strength. Results show that the generated super-resolution fields of SDEdit are unphysical, while those of DPS are reasonable but with smoothed fine-scale features; however, neither of these lower-cost models propagates observational information effectively to unobserved regions. The two conditional models require re-training, but reconstruct missing fine-scale features, are cycle-consistent with observations, and predict correct turbulence statistics, including the tails. Further, their mean errors are highly correlated with and predictable from their ensemble standard deviations. Results highlight the trade-offs between ease of implementation, fidelity (sharpness), and cycle-consistency of the diffusion models, and offer practical guidance for deployment in geophysical inverse problems.

Plain Language Summary Simulating the ocean or the atmosphere at high resolution is expensive because computers solve physics equations on many tiny grid cells. As an approximation, simulations are run with a small number of large grid cells, missing important small-scale physics such as swirling eddies. Even geophysical observations with satellites, balloons, buoys, or robots can be incomplete with large gaps. We explore whether generative diffusion models (DMs), a new machine learning technique, can complete the missing details in ways that respect the laws of physics. Using turbulence simulations, we compare four DMs. Two *conditional* models require training from scratch to directly learn missing details, demanding more data and computation, while two *guided unconditional* models modify existing DMs to complete missing details. We find that conditional models create realistic details that respect turbulence physics. In challenging cases, one of the unconditional models fails, while the other can provide low-cost, reasonable guesses. We show tradeoffs between resource use and realism that users can tune to their needs. Some generative DMs can lead to sharper maps of ocean currents, weather, and climate, help merge real observations with simulations, and possibly learn physics from even incomplete, indirect data.

© 2026 The Author(s). Journal of Advances in Modeling Earth Systems published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

1. Introduction

Higher-resolution oceanic and atmospheric modeling for marine, weather, and climate predictions (Kalnay, 2003; Harris & Lin, 2014; E. P. Chassignet & Xu, 2021; Brotzge et al., 2023; Lai et al., 2024) can provide greater accuracy but are computationally expensive. Numerical models are limited not only by their spatio-temporal resolution but also by the processes they represent and the number of variables they simulate (Gupta & Lermusiaux, 2021; Jacobs et al., 2023). For example, simulating the full range of spatio-temporal scales in turbulent flows is computationally intractable; approximations and parameterizations, such as in Large-Eddy Simulations

Resources: Pierre F. J. Lermusiaux
Software: Anantha Narayanan Suresh Babu, Akhil Sadam
Supervision: Pierre F. J. Lermusiaux
Visualization: Anantha Narayanan Suresh Babu, Akhil Sadam
Writing – original draft: Anantha Narayanan Suresh Babu
Writing – review & editing: Anantha Narayanan Suresh Babu, Akhil Sadam, Pierre F. J. Lermusiaux

(LES), are thus employed to represent the unresolved scales, which can also limit utility (Leonard, 1975; Pope, 2001). Recent machine-learning-based surrogate models for the ocean and atmosphere (Pathak et al., 2022; Rajagopal et al., 2023) also contain unresolved scales due to spectral bias, which inhibits accurate prediction of fine-scale features (Lai et al., 2024). This can be amplified by dynamics such as quasi-geostrophic turbulence, leading to unstable or nonphysical long-term predictions (Chattopadhyay & Hassanzadeh, 2023). The dual-cascade of forced geophysical turbulence (Kraichnan, 1971), that is, the inverse cascade of energy from forcing scales to large scales and the forward cascade of enstrophy to small scales, is a challenge. In the ocean, resolving submesoscale processes with scales of the order of 100 m to 10 km (J. C. McWilliams, 2016; Taylor & Thompson, 2023) can contribute to energy, mass, and tracer transports (Thomas et al., 2008). In climate, simulations with scales of 10 km can improve predictions of extreme weather events and enhance risk assessment (Diffenbaugh et al., 2005; Lopez-Gomez et al., 2025). As numerical models, in situ observations from weather stations, balloons, buoys, drifters, gliders, and ships, as well as remote observations from satellites and radars, are also spatio-temporally sparse and gappy (Bluestein et al., 2022; Lin & Yang, 2020). The integration of modeling and observing systems for intelligent sampling and data assimilation can be most useful to improve estimates (Bennett, 2005; Kalnay, 2003; Law et al., 2015; Lermusiaux, Chiu, et al., 2006; Lermusiaux, Malanotte-Rizzoli et al., 2006, Lermusiaux et al., 2017). In general, techniques that reconstruct fine-scale features from coarse or sparse simulated or observed data are critical to improve their utility (Buzziotti, 2023). One such approach to reconstruct fine-scale features is through super-resolution (Sofos & Drikakis, 2025), which is the focus of our work.

Super-resolution can be interpreted as a framework for a broad class of inverse problems (Fukami et al., 2023), with applications to several Earth system modeling challenges:

1. *Subgrid-scale (SGS) modeling.* A goal of model closure is to parametrize unresolved subgrid-scale processes. Traditional approaches include functional models, often based on eddy viscosity hypotheses (Leith, 1968; Smagorinsky, 1963) and structural models, which directly approximate the subgrid-scale stress terms through scale-similarity or formal series expansions (Bardina et al., 1980; Clark et al., 1979). Purely diffusive functional approaches do not account for backscatter, the transfer of energy from subgrid scales to resolved scales (Frederiksen & Davies, 1997; Jansen & Held, 2014), and are not well correlated with the true SGS term, while structural models often suffer from numerical instabilities (Vreman et al., 1996). To overcome these issues, recent data-driven closures have been proposed (Gupta & Lermusiaux, 2023; Jakhar et al., 2024; Maulik et al., 2019; Ross et al., 2023; Zanna & Bolton, 2020). Super-resolution provides an alternate framework to develop optimal closures by reconstructing subgrid-scale features (Bode et al., 2021; Langford & Moser, 1999; Maejima & Kawai, 2025).
2. *Downscaling.* Downscaling aims to improve the fidelity of regional and shorter-term forecasts by increasing the spatio-temporal resolution of coarser global Earth system models. Dynamical downscaling and nesting (von Storch et al., 2000; Denis et al., 2002; Haley & Lermusiaux, 2010; Kulkarni et al., 2018; Johnston et al., 2019; Simmons et al., 2019) use global model outputs to provide large-scale boundary conditions, initial conditions, and forcing to regional model simulations. These schemes have varied accuracies and computational costs, for example, (Z. Xu et al., 2019). With statistical downscaling, statistical relationships are learned between historical regional data and model outputs through regression or stochastic generation (Khan et al., 2006; Wilby et al., 1998). This statistical approach can downscale larger ensembles due to their efficiency and reduce biases inherent to global models. Some applications can also favor specific downscaling, for example, (Keller et al., 2022). Recently, machine learning has emerged as a promising tool for statistical downscaling (Baño-Medina et al., 2020; Rampal et al., 2024), with ensemble super-resolution offering an advanced nonlinear approach. These super-resolved fields still require bias correction to reduce systematic errors (Wan et al., 2023).
3. *State estimation and mapping.* State estimation is the process of determining the spatial and temporal distribution or evolution of states (dynamical fields). Optimal interpolation (OI, or objective mapping) estimates states from observations through statistical interpolation (Bretherton et al., 1976; Menemenlis et al., 1997). These estimates, however, may not obey the governing dynamics. Data assimilation instead melds observations with the underlying dynamics from numerical models (Carrassi et al., 2018; Hoteit et al., 2018; Law et al., 2015; Robinson et al., 1998; Robinson & Lermusiaux, 2001). In oceanic and atmospheric modeling, popular data assimilation methods are based on Kalman filtering and its ensemble and non-Gaussian extensions (Evensen, 2003; Sondergaard & Lermusiaux, 2013), particle filters (Van Leeuwen et al., 2019), and

variational methods such as 3D and 4D-VAR (Bannister, 2017). Recently, machine learning models have been developed for data assimilation (Bocquet, 2023; Hodyss & Morzfeld, 2026), where data-driven surrogates replace either the dynamical model (Adrian et al., 2025) or the entire end-to-end data assimilation (Manshausen et al., 2025; Martin et al., 2025). Super-resolution can also be used to reconstruct missing features in physical space, with additions to preserve dynamics (Barthélémy et al., 2022). Unlike subgrid scale modeling or downscaling, where missing features typically correspond to high wavenumbers in Fourier space, here reconstructed features also correspond to low wavenumbers (Buzzicotti, 2023).

4. *Scale-enhanced initialization and simulation spin-up.* In geophysical modeling, interpolation across grids is utilized to initialize and spin-up regional models from global analyses. For example, to initialize coastal ocean forecasts, global fields can be adjusted to high-resolution bathymetry and multiscale dynamics using data corrections and least-square inversions for multi-process adjustments (Haley et al., 2015; Lermusiaux et al., 2024). High-resolution large-ensemble forecasts can be initialized by generating multi-region, down-scaled, three-dimensional field perturbations based on data and dynamics (Haley et al., 2023; Lermusiaux, 2002; Lermusiaux et al., 2000, 2020). As in dynamical downscaling, larger-scale information can also be re-injected into the regional model to provide a warm-start (Short & Petch, 2022). As all of these scale-enhanced initializations take computational time, fast super-resolution would be most useful to generate multiscale-consistent ensemble initializations or enable direct high-resolution spin-up through super-resolved warm starts.

Given the wide application range, super-resolution of turbulent flows is a growing area of research in machine learning. Various data-driven approaches and deep neural architectures have been explored, including sparse representations (Callahan et al., 2019), convolutional neural networks (CNNs) with hybrid downsampled skip-connection/multi-scale networks (DSC/MS) (Fukami et al., 2019; Fukami & Taira, 2024), spatio-temporal CNNs (Fukami et al., 2021; B. Liu et al., 2020), physics-informed neural networks (Gao et al., 2021), graph neural networks (Barwey et al., 2025), vision transformers (Q. Xu et al., 2023), and fully differentiable solvers that train with a dynamics-based loss (Page, 2025). While these methods can generate plausible high-resolution fields, they provide only a single best estimate without uncertainty quantification and often struggle to preserve physical invariants like energy spectra. Since nonlinear and high-dimensional inverse problems suffer from the *curse of dimensionality* and are inherently ill-posed (Fernandez-Martinez & Fernandez-Muniz, 2020), there is a growing need for probabilistic methods that generate samples with uncertainty quantification (Lermusiaux, Chiu, et al., 2006; Lermusiaux, Malanotte-Rizzoli et al., 2006).

Probabilistic generative models provide a natural framework to meet these requirements. Popular generative models include variational autoencoders (VAEs), generative adversarial networks (GANs), normalizing flows (NFs), and diffusion models (DMs). VAEs map fields to samples from a low-dimensional latent prior, typically Gaussian in distribution (Kingma & Welling, 2019). Although vanilla VAEs generate blurry samples (Bredell et al., 2023), they have been used for super-resolution through modal decomposition (Eivazi et al., 2022). In contrast, GANs generate sharp samples by matching distributions (Gui et al., 2021), showing effectiveness for some inverse problems (Güemes et al., 2022; Patel et al., 2022). However, GAN training is known to be unstable and susceptible to mode collapse, as it requires the simultaneous training of two competing neural networks: a generator network that generates plausible samples to fool a discriminator network (Kossale et al., 2022; Wiatrak et al., 2019). Vanilla GANs are potentially inaccurate, so Kim et al. (2021) utilize cycle-consistent-GANs (Zhu et al., 2017) for super-resolution of Direct Numerical Simulation fields from LES channel flows. Finally, normalizing flows (NFs) are stable to train and produce sharp images; they learn an *invertible* map to a tractable distribution (Rezende & Mohamed, 2015). These models excel at density estimation, but were historically limited in architecture by the invertibility constraint (Zhai et al., 2024). Recent extensions include continuous time flows through flow matching and stochastic interpolants (Albergo & Vanden-Eijnden, 2022; Lipman et al., 2022), with applications to super-resolution of Kolmogorov flows (Y. Chen et al., 2024).

Diffusion models have emerged as state-of-the-art generative tools (Y. Song et al., 2020). Diffusion models are more stable to train than GANs and have been shown to generate fields of equivalent or higher quality than GANs, VAEs, or NFs (Dhariwal & Nichol, 2021). These models first apply a corruption process that destroys structure in the fields from a given distribution, and then train a model that iteratively reverses this corruption process to generate fields that follow the required distribution. The reverse process could be thought of as iterative spectral regression, where the large-scale features are predicted first, with small-scale features added iteratively (B. Wang & Vastola, 2023). Diffusion models are flexible and can be either conditional (Saharia, Chan, et al., 2022), where

the model is directly trained to learn the relationship between the input and the target, or guided unconditional (Daras et al., 2024), where a pre-trained model is steered toward the target using a separate guidance process.

Due to these properties, generative DMs have been applied to inverse problems in computer vision and geophysical fluid dynamics (Milanfar & Delbracio, 2025). For super-resolution, most approaches rely on conditional DMs (Oommen et al., 2025; Sardar et al., 2024) though guided unconditional models also exist (Fan et al., 2025) with applications to idealized Kolmogorov flows. For downscaling of atmospheric flows, some approaches have conditioned DMs on neural surrogates (Mardani et al., 2025; Sundar et al., 2024) or directly on coarse-resolution model output (Han et al., 2025; Watt & Mansfield, 2024). Chakraborty et al. (2025) guided an unconditional model for generative atmospheric flow downscaling. For state estimation and inference, guided unconditional models have been employed for generative data assimilation of sparse observations in idealized turbulent flows (Rozet & Louppe, 2023) and weather fields (Manshausen et al., 2025). Unconditional (Martin et al., 2025) and conditional DMs (T. Li et al., 2023; Asefi et al., 2025; Souza et al., 2025) have been used to infer surface and subsurface ocean and idealized vorticity fields from coarse, sparse, and gappy observations. To reduce training costs, training-free diffusion approaches based on ensemble filtering have also been proposed (Bao et al., 2025; Yin et al., 2024).

Despite these advances, no study has focused on both super-resolution and estimation or inference of geophysical turbulent fields from coarse full field and coarse, sparse and gappy data. Moreover, existing works focus on a single method in isolation without direct comparison across methods, with several open questions, including: While guided unconditional models are effective in computer vision, can they achieve comparable skill to conditional models in geophysical fluid super-resolution problems? When do the assumptions made in guided unconditional models break down in turbulent flow settings? Can conditional models remain stable (Baldassari et al., 2023) and generalize effectively when observations are sparse and gappy? How does observational information propagate in DMs? How sensitive are outputs of DMs to the guidance strength (or other tuning parameters)?

To address these challenges, (a) We apply four diffusion modeling (2 guided unconditional and 2 conditional) approaches to super-resolution and inference of quasi-geostrophic turbulence *under the β -plane approximation* for different Reynolds numbers and flow regimes. (b) We investigate the effect of sparse and gappy coarse observations. (c) We benchmark the four score-based DMs using a comprehensive suite of skill metrics, including norms of the generated vorticity fields and turbulence statistics. (d) We quantify and evaluate the uncertainty of our generative models using the generated ensemble of super-resolved fields and their errors. (e) We analyze the propagation of sparse and gappy observational information by the generative diffusion processes. (f) We study the sensitivity of the four DMs to their guidance strength. Our results highlight the trade-offs between ease of implementation, fidelity (sharpness), and cycle-consistency when deploying DMs for super-resolution and inference applications in geophysical fluid dynamics (Bennett, 1992; Wunsch, 2006).

The paper is organized as follows: Section (2) introduces the formal problem statement for super-resolution and state estimation or inference. Details of the various guided unconditional and conditional score-based diffusion modeling methods utilized are in Section (3). Section (4) describes the quasi-geostrophic turbulence simulations used in our applications, along with details of the dynamical regimes, numerical schemes, data sets, test cases, and skill metrics. Applications of the four DM methods to super-resolution and inference of quasi-geostrophic turbulence with coarse and with sparse, gappy observations are provided in Section (5), followed by a discussion of the comparative analyses and conclusion (Section 6).

2. Problem Statement

We consider two related inverse problems: (a) super-resolution (Moser et al., 2023) where missing short-wavelength, subgrid-scale features are reconstructed, and (b) state estimation or inference, where high-resolution fields are estimated from low-resolution, sparse, and gappy data, with reconstructed features including long wavelengths. We formulate both problems within a common framework as an inversion of an observation process.

Specifically, we define a measurement model that maps a high-fidelity field $x_0 \in \mathbb{R}^n$ to a low-fidelity field $y \in \mathbb{R}^d$ (the low-resolution data), where $d \leq n$, through a nonlinear observation operator $\mathcal{H} : \mathbb{R}^n \rightarrow \mathbb{R}^d$, where $z \sim \mathcal{N}(0, \mathbb{I}_d)$ and σ_y models the noise in the observation process.

$$y = \mathcal{H}(x_0) + \sigma_y z \quad (1)$$

In general, the observation operator is not restricted to just the physical space, but could also act in the Fourier domain (wavenumber space). Observations can also be partial, with only a subset of the domain being observed. We are interested in the inverse of Equation 1, or a model \mathcal{M} that, given noisy observation samples y , can generate plausible reconstructed high-fidelity samples \hat{x}_0 from the conditional distribution $p(x_0|y)$. Because this inverse mapping is not one-to-one, it is ill-posed, and perfect recovery is impossible for many applications in geophysical modeling (Bennett, 1992; Blau & Michaeli, 2018). Therefore, we seek to obtain generative models \mathcal{M} that perform such inversions in a probabilistic sense, and sample from the posterior distribution rather than provide only point estimates such as the conditional expectation $\mathbb{E}[x_0|y]$. We propose using generative DMs due to their ability to effectively sample an ensemble of plausible fields that follow a required distribution (Dhariwal & Nichol, 2021).

This problem formulation is related to data assimilation (Evensen et al., 2022; Law et al., 2015). For example, early approaches such as OI (Gandin, 1963) or 3D-VAR (Courtier et al., 1998) also perform inversions one snapshot at a time and optimize for a single analysis, that is, the minimum of the error variance or of a cost function, respectively (Lorenc, 1986; Xie et al., 2002). Under assumptions of Gaussianity, OI and 3D-VAR correspond to the maximum a posteriori estimate, but neither scheme directly provides dynamic uncertainty quantification (Kalnay, 2003). Ensemble-based methods utilize an ensemble of numerical model forecasts to evolve and predict uncertainties, and include ensemble Kalman filters, error subspace schemes, sequential Monte Carlo schemes, and particle filters (Evensen, 2003; Van Leeuwen et al., 2019). For nonlinear data assimilation, several of these schemes utilize the forecast non-Gaussian distributions and semiparametric model fits such as Gaussian Mixture Models (Carrassi et al., 2022; Evensen et al., 2022; Lolla & Lermusiaux, 2017; Sondergaard & Lermusiaux, 2013). However, the accuracy of ensemble-based methods is limited by finite subspace and sampling sizes, often constrained by computational cost, which can lead to sampling errors or degeneracy in high-dimensional settings (Bengtsson et al., 2008; Ehrendorfer, 2007). Initializing ensembles with realistic uncertainty structures also requires data and modeling (Lermusiaux, 2002; Lermusiaux et al., 2000; Zupanski et al., 2006). In contrast, trained diffusion-based models can generate larger observation-consistent ensembles at much lower computational cost (Manshausen et al., 2025). Rather than relying on numerical forecast ensembles, a generative DM could provide a learned prior with a typically non-Gaussian distribution of plausible high-resolution fields $p(x_0)$. Of course, training such models could be costly, requiring significant data and long numerical simulations. In addition, retraining and targeted denoisers may be needed at assimilation times for conditioning and optimal dynamic prior distributions. Several of these diffusion-based assimilation limitations are discussed by Hodyss and Morzfeld (2026). In this work, we evaluate and compare the accuracy, capabilities, and costs of four generative DMs for super-resolution and inference of turbulent fields in the sense of Equation 1.

In summary, our goal is to solve the inverse problem of reconstructing high-resolution vorticity fields in quasi-geostrophic turbulence under the β -plane approximation (Section 4) from coarse, sparse, and gappy observations using trained generative DMs that sample from the posterior distribution $p(x_0|y)$. For our focus on super-resolution, \mathcal{H} corresponds to filtering and down-sampling operations. For super-resolution with inference, this is combined with sparse and gappy observing systems. The specific forms of \mathcal{H} are detailed in Section (4.3).

3. Score-Based Diffusion Models

We now describe the score-based generative DMs and their use for super-resolution and inference. Diffusion models aim to generate samples from a complex distribution p_0 (e.g., the distribution of all possible fields that can be simulated using the Quasi-Geostrophic equations) using a more tractable distribution p_T (e.g., a multivariate Gaussian distribution). This generative process occurs over a pseudo-time or diffusion time t within the interval $[0, T]$ where T is the maximum diffusion time. This pseudo-time is the artificial time variable of the generative sampling process and is not a physical time. The use of DMs typically involves a forward process, which consists of gradually *noising* or corrupting samples from p_0 to resemble samples of p_T over pseudo-time, and a reverse process, which consists of solving a deterministic or stochastic differential equation (SDE) (Y. Song et al., 2020) or a discrete Markov-chain process (Sun et al., 2022), starting from the tractable distribution p_T to generate

samples from p_0 . This generation process can be conditioned on some constraints, and the sampling can be very efficient on modern computers (Sohl-Dickstein et al., 2015).

Following Y. Song et al. (2020), the forward process for a DM can be represented by an Itô differential equation

$$dx_t = f_t(x_t) dt + g_t dW_t \quad \forall t \in [0, T] \quad (2)$$

with initial conditions $x_0 \sim p_0$, that is, samples from the complex distribution p_0 . In Equation 2, $x_t \in \mathbb{R}^n$ is the field generated at intermediate pseudo-times $0 \leq t \leq T$, f the drift coefficient, g the diffusion coefficient, and W the standard Wiener process or Brownian motion. The subscript t highlights the pseudo-time dependence. Typically, the forms of f and g are chosen a priori and hence known. This forward process has a corresponding stochastic reverse process that can be derived analytically utilizing results from stochastic calculus (Anderson, 1982),

$$dx_t = [f_t(x_t) - g_t^2 \nabla_{x_t}(\log p_t(x_t))] dt + g_t dW_t \quad (3)$$

where $\log p_t(x_t)$ is the log-probability distribution at pseudo-time t . Its gradient, $\nabla_{x_t} \log p_t(x_t)$, is referred to as the *score*. It appears in the deterministic drift term of the reverse diffusion process and guides it, directing noisy or corrupted samples at pseudo-time t toward regions of higher probability density of the distribution p_0 (Park et al., 2024). This score has no closed-form analytical expression and is typically estimated using a deep neural network or through training-free Monte Carlo approximations (Y. Liu et al., 2024; Bao et al., 2025). If instead we are interested in generating fields given certain constraints (e.g., high-resolution fields that correspond to a specific low-resolution field), the resulting reverse conditional diffusion process can be derived as

$$dx_t = [f_t(x_t) - g_t^2 \nabla_{x_t}(\log p_t(x_t|y))] dt + g_t dW_t \quad (4)$$

where y is the conditioning field and $\log p_t(x_t|y)$ is called the *conditional score*.

In this work, we utilize the Variance-Preserving (VP-SDE) forward process (Y. Song et al., 2020), where the marginal distribution at each step during the forward diffusion process (Equation 2) is given by

$$p_t(x_t) = \mathcal{N}(x_t; \mu_t x_0, \sigma_t^2 I_n), \quad 0 \leq \sigma_t \leq 1, \quad \lim_{t \rightarrow T} \sigma_t \rightarrow 1 \quad (5)$$

where $x_0 \sim p_0$ denotes samples from the complex distribution p_0 , the amplitude σ_t defines the *noise schedule*, and the normalization factor $\mu_t = \sqrt{1 - \sigma_t^2}$ ensures the variance of noisy samples x_t is preserved over pseudo-time $0 \leq t \leq T$. This Equation 5 is the marginal of the analytical solution of Equation 2 under the VP-SDE assumption, and provides a computationally efficient method to generate noisy or corrupted samples at intermediate pseudo-times. Using Tweedie's formula (Robbins, 1992), the score for the unconditional reverse process is then given by

$$\nabla_{x_t}(\log p_t(x_t)) = \frac{\mu_t \mathbb{E}[X_0 | X_t = x_t] - x_t}{\sigma_t^2} \quad (6)$$

and the conditional reverse process is given by

$$\nabla_{x_t}(\log p_t(x_t|y)) = \frac{\mu_t \mathbb{E}[X_0 | X_t = x_t, Y = y] - x_t}{\sigma_t^2} \quad (7)$$

where \mathbb{E} is the expectation operator, and $\mathbb{E}[X_0 | X_t = x_t]$ and $\mathbb{E}[X_0 | X_t = x_t, Y = y]$ correspond to well-trained *denoiser* neural networks, $h_\theta(\cdot, t)$, in practice.

For the unconditional reverse process, the *denoiser* neural network $h_{\theta, \text{uncond.}}(x_t, t)$ with trainable parameters θ is trained to predict the noise added given a noisy field x_t as the input. The neural network is trained to minimize the following objective:

$$\min_{\theta} \mathbb{E}_{x_0, x_t} \left\| h_{\theta, \text{uncond.}}(x_t, t) - \frac{x_t - \mu_t x_0}{\sigma_t} \right\|^2 \quad (8)$$

For the conditional models, the neural network $h_{\theta, \text{cond.}}(x_t, y, t)$ also takes in the conditioning field (y , the observations) as an additional input, and hence the training objective is modified as

$$\min_{\theta} \mathbb{E}_{x_0, x_t, y} \left\| h_{\theta, \text{cond.}}(x_t, y, t) - \frac{x_t - \mu_t x_0}{\sigma_t} \right\|^2 \quad (9)$$

The learned score, $s_{\theta}(\cdot, t)$, for both conditional and unconditional models can be computed using their corresponding denoiser networks $h_{\theta}(\cdot, t)$

$$s_{\theta}(\cdot, t) = -\frac{h_{\theta}(\cdot, t)}{\sigma_t} \quad (10)$$

Finally, for sampling, the reverse diffusion process for both unconditional (Equation 3) and conditional models (Equation 4) are discretized in time and integrated using an exponential integrator scheme (Q. Zhang & Chen, 2022; J. Song et al., 2020) with Langevin Monte Carlo corrections (Y. Song et al., 2020; Rozet & Louppe, 2023) using the learned score (Equation 10).

Next, we describe the four generative diffusion modeling approaches we employ and compare them inspired by the following questions: Is the score modified? Is prior knowledge of the down-sampling and filtering operations required? Is task-specific retraining required? Is the guidance controllable? Finally, we compare the number of function evaluations (neural network evaluations) per diffusion pseudo-timestep.

3.1. Guided Unconditional Models

One class of approaches *guides* a pre-trained unconditional DM to generate conditional samples. The main advantage is that these approaches can be directly applied to a pre-trained DM and require no modification to the training process, and hence save significant computational resources. Figure 1 shows a schematic of unconditional DMs and two ways to guide them.

3.1.1. \mathcal{M}_1 : Modifying the Initial Condition (SDEdit)

In the first approach, \mathcal{M}_1 , the unconditional DM is guided by simply starting the reverse diffusion process at an intermediate pseudo-time t_i and utilizing a new modified initial condition x_i that resembles noisy or corrupted samples generated at the intermediate pseudo-time t_i .

As we will see, this approach, also known as *SDEdit* (Meng et al., 2021), directly preserves large-scale features of interest in the forward and reverse diffusion processes. The input low-resolution data $y \in \mathbb{R}^d$ is first interpolated to the target high-resolution to obtain $y_{\text{interp}} \in \mathbb{R}^n$. Next, samples from the forward process (Equation 2) with the interpolated field y_{interp} are obtained from the distribution (Equation 5) at the intermediate pseudo-time t_i , $0 < t_i < T$,

$$p_t(x_i) = \mathcal{N}(x_i; \mu_{t_i}, \sigma_{t_i}^2 \mathbf{I}_n), \quad (11)$$

This provides initial conditions $x_i \in \mathbb{R}^n$ for the reverse diffusion. Since this approach does not modify the score, the reverse unconditional diffusion process (Equation 3) is thus run with the unconditional score (Equation 6) (learned using Equations 8 and 10) from this x_i at t_i back to the initial pseudo-time $t = 0$. The final result is the conditional sampled high-resolution field \hat{x}_0 . This approach is well-suited to super-resolution since the forward process adds noise to smooth out undesirable down-sampling artifacts while preserving large-scale flow features. A main parameter to be fine-tuned is the intermediate time t_i , to ensure a sufficient trade-off between fidelity (sharpness) and cycle-consistency, that is, do the filtered generated samples $\mathcal{H}(\hat{x}_0)$ match the low-resolution observation y ?

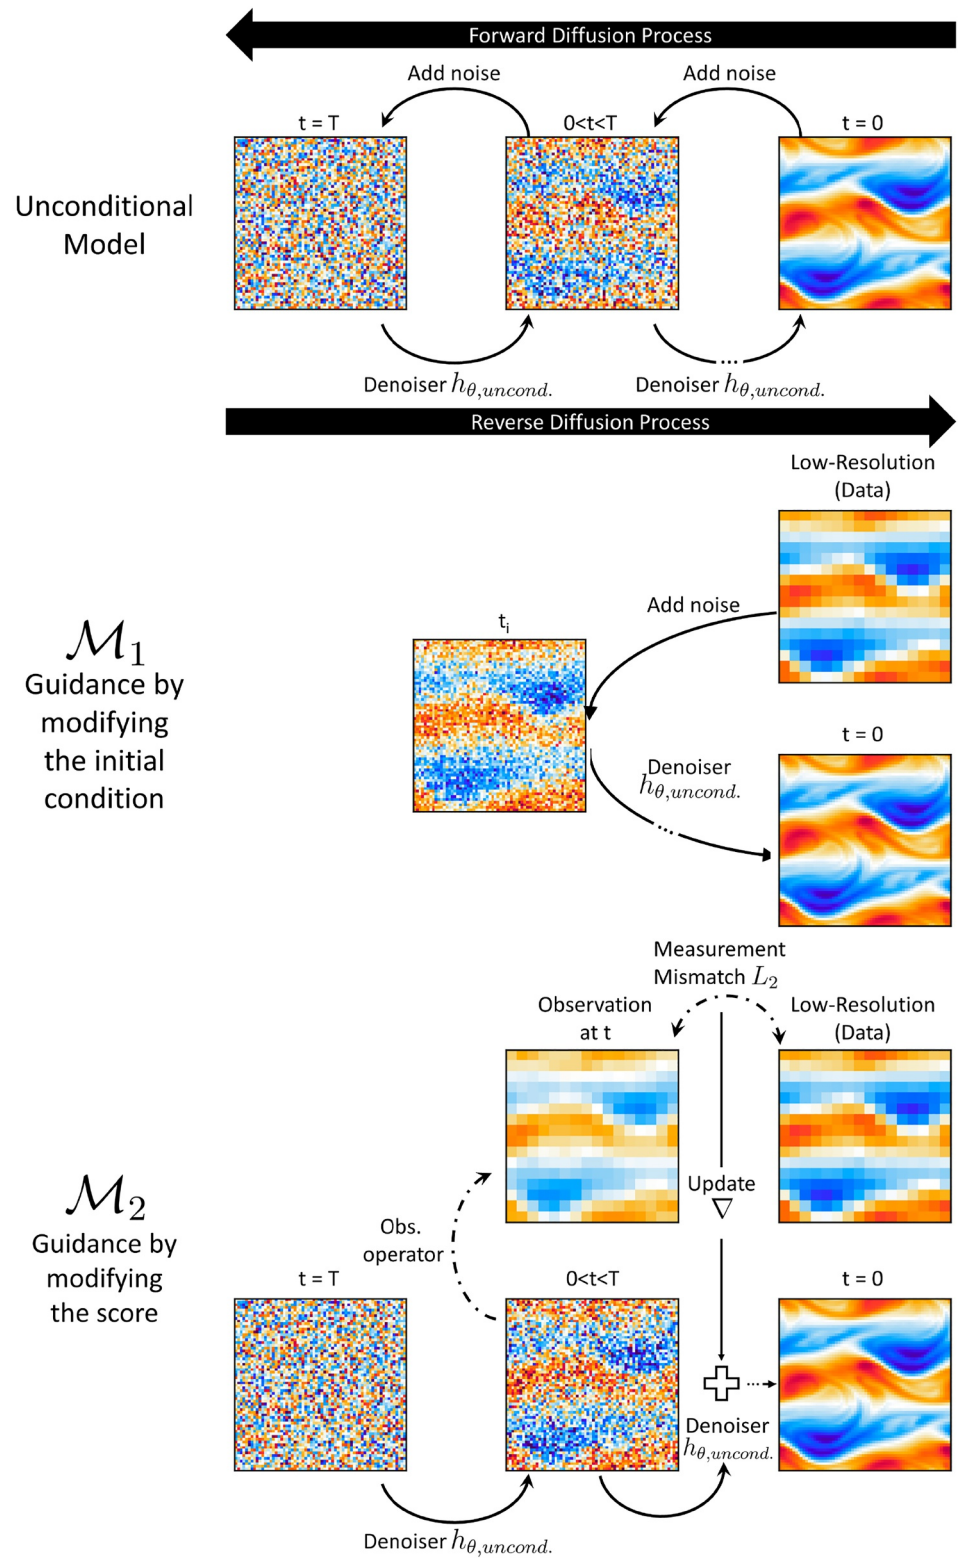


Figure 1. Schematic of the various unconditional diffusion modeling approaches. (Top row) Forward and reverse diffusion process for the unconditional model (Middle row) Reverse process for guidance by modifying the initial condition (Section 3.1.1). (Bottom row) Reverse process for guidance by modifying the score (Section 3.1.2).

3.1.2. \mathcal{M}_2 : Modifying the Score (Diffusion Posterior Sampling)

In the second approach, the score of the unconditioned DM is directly modified to generate conditional samples. The idea is to leverage Bayes' law

$$p_t(x_t|y) = \frac{p_t(y|x_t)p_t(x_t)}{p_t(y)} \quad (12)$$

and hence derive

$$\nabla_{x_t}(\log p_t(x_t|y))_{\mathcal{M}_2} = \nabla_{x_t}(\log p_t(x_t)) + \nabla_{x_t}(\log p_t(y|x_t)) \quad (13)$$

where the first term on the right side is simply the unconditional score (Equation 6) that we have access to, and the second term is a *measurement matching* term that guides the DM to generate conditional samples that are consistent with the low-resolution input field (observations). However, this measurement matching term is analytically intractable and requires further simplification for practical implementations. Daras et al. (2024) provides a comprehensive review of approaches to approximate this term. One of the most popular approaches used is Diffusion Posterior Sampling (DPS) (Chung, Kim, McCann, et al., 2023), which assumes

$$\nabla_{x_t}(\log p_t(y|x_t)) \simeq \nabla_{x_t}(\log p_t(y|X_0 = \mathbb{E}[X_0|X_t = x_t])) \quad (14)$$

that is, it assumes the score is the same if the conditioning is on the denoised samples instead of conditioning on the noisy or corrupted samples. Since our measurement model (Equation 1) has a Gaussian structure, the likelihood $p_t(y|X_0 = x_0)$ can be derived as

$$p_t(y|X_0 = x_0) = \mathcal{N}(y; \mathcal{H}(x_0), \sigma_y^2 \mathcal{I}_d) \quad (15)$$

where \mathcal{H} is the nonlinear observation operator and σ_y models the noise in the observation process. Hence, the likelihood $p_t(y|X_0 = \mathbb{E}[X_0|X_t = x_t])$ can be derived from Equations 14 and 15 as

$$p_t(y|X_0 = \mathbb{E}[X_0|X_t = x_t]) \approx \mathcal{N}(y; \mathcal{H}(\mathbb{E}[X_0|X_t = x_t]), \sigma_y^2 \mathcal{I}_d) \quad (16)$$

$$\Rightarrow \nabla_{x_t}(\log p_t(y|x_t)) \simeq -\nabla_{x_t} \left(\frac{1}{2\sigma_y^2} \|y - \mathcal{H}(\mathbb{E}[X_0|X_t = x_t])\|^2 \right) \quad (17)$$

The term $y - \mathcal{H}(\mathbb{E}[X_0|X_t = x_t])$ is the *innovation vector* in data assimilation. Computing Equation 17 requires access to gradients and Jacobian-vector products if \mathcal{H} is nonlinear and if analytical closed-form expressions are inaccessible. Practically, computing Equation 17 leads to instabilities at pseudo-times t close to 0 (Rozet & Louppe, 2023). Hence, this measurement matching term is scaled by replacing the measurement noise σ_y^2 with a more general scaled covariance approximation $\sigma_{y, \text{scaled}}^2$ that varies with the pseudo-time through a dependence on the noise schedule σ_t to controllably guide the reverse diffusion process (Manshausen et al., 2025; Rozet & Louppe, 2023),

$$\sigma_{y, \text{scaled}}^2 = \sigma_y^2 + \frac{\sigma_t^2}{\mu_t^2} C \quad (18)$$

where C is a tunable parameter that controls the guidance strength. Hence, the final modified score is given by

$$\nabla_{x_t}(\log p_t(x_t|y))_{\mathcal{M}_2} = \nabla_{x_t}(\log p_t(x_t)) - \frac{1}{2\sigma_{y, \text{scaled}}^2} \nabla_{x_t}(\|y - \mathcal{H}(\mathbb{E}[X_0|X_t = x_t])\|^2) \quad (19)$$

The first part of the right-hand side of Equation 19 can be learned using Equations 8 and 10, while the second part can be computed from Equations 17 and 18. In summary, the second approach consists of running the reverse

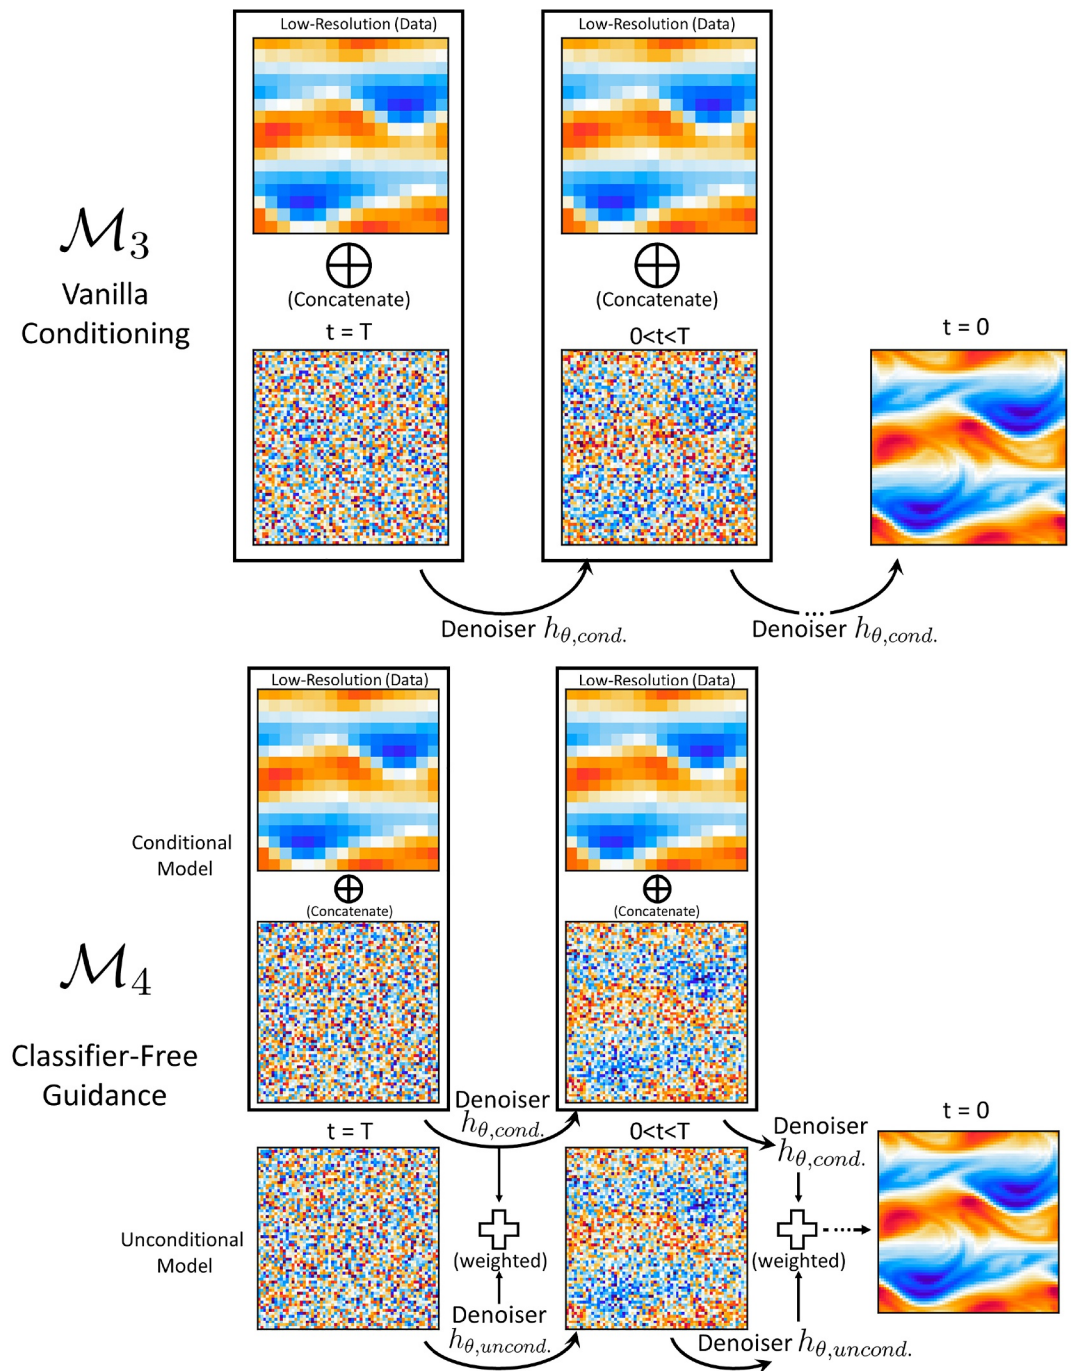


Figure 2. Schematic of the various conditional diffusion modeling approaches (Section 3.2). (Top row) Reverse diffusion process for the vanilla conditional model (Bottom row) Reverse diffusion process for classifier-free guidance.

conditional diffusion process (Equation 4) with the score derived in Equation 19 starting from initial conditions $x_T \sim p_T$, that is, samples from the tractable distribution, and the data y , to obtain conditional sampled high-resolution fields \hat{x}_0 .

3.2. Conditional Models

We now describe approaches that directly train conditional DMs using data pairs of the target high-resolution field, as illustrated in Figure 2.

Table 1
Comparison of Implementation and Sampling Characteristics of the Four Diffusion Modeling Approaches Considered for Super-Resolution and Inference

| | Guided unconditional models | | Conditional models | |
|---|-----------------------------|-----------------|--------------------|-----------------|
| | \mathcal{M}_1 | \mathcal{M}_2 | \mathcal{M}_3 | \mathcal{M}_4 |
| Modifies the score | ✗ | ✓ | ✓ | ✓ |
| Avoids need of observation operator | ✓ | ✗ | ✓ | ✓ |
| Task-specific re-training | ✗ | ✗ | ✓ | ✓ |
| Controllable guidance | ✓ | ✓ | ✗ | ✓ |
| Number of function evaluations ^a | $\frac{t_i}{T}$ | 1 | 1 | 2 |

^aRelative number of denoiser evaluations per pseudo-timestep during the reverse diffusion process (sampling). \mathcal{M}_1 : Modifying the initial condition (Section 3.1.1), \mathcal{M}_2 : Modifying the score (Section 3.1.2). \mathcal{M}_3 : Vanilla conditional model (Section 3.2), \mathcal{M}_4 : Classifier-free guidance (Section 3.2).

In the vanilla variant, \mathcal{M}_3 , direct conditioning is typically achieved by modifying the denoiser neural network to take the low-resolution conditioning field as an additional input (Equation 9). Hence, this approach is also called *image-to-image* diffusion modeling (Saharia, Chan, et al., 2022). During sampling, the reverse conditional SDE (Equation 4) can be run using the conditional score (Equation 7, learned using the denoiser from Equations 9 and 10), starting from initial conditions $x_T \sim p_T$, that is, samples from the tractable distribution, and the data y , to obtain conditional sampled high-resolution fields \hat{x}_0 .

The vanilla conditional model can further be modified with a guidance strength to further improve the quality of samples and controllably guide the DM. This can be achieved using *classifier-free guidance* (CFG), \mathcal{M}_4 , where the score is modified to be a linear combination of the conditional score and the unconditional score (Ho & Salimans, 2021), similarly to relaxation schemes (Lermusiaux, 2015),

$$\nabla_{x_t}(\log p_t(x_t|y))_{\mathcal{M}_4} = \nabla_{x_t}(\log p_t(x_t)) + w[\nabla_{x_t}(\log p_t(x_t|y)) - \nabla_{x_t}(\log p_t(x_t))] \quad (20)$$

where the parameter $w \geq 1$ corresponds to the guidance strength. It determines how strongly the conditional model guides the reverse process. The denoiser for CFG can be trained from scratch by simply training a vanilla conditional model (Equation 9) with the conditioning being set randomly to a null field \emptyset (which could be a field of zeros or pure Gaussian noise) every few instances of training to resemble an unconditional DM. During sampling, the reverse SDE (Equation 4) can be run using the modified conditional score (Equation 20, learned using Equations 8, 9, and 10), starting from initial conditions $x_T \sim p_T$, that is, samples from the tractable distribution, and the data y , to obtain conditional sampled high-resolution fields \hat{x}_0 .

Table 1 summarizes the differences among the four diffusion modeling approaches we consider and answers the questions we had asked. All methods except \mathcal{M}_1 modify the unconditional score. In \mathcal{M}_1 , the initial conditions are simply modified to preserve large-scale features from the data (observations), making it more prone to failure than the other approaches. \mathcal{M}_2 requires access to the observation operator and its gradients, which may not be possible when observations involve indirect or correlated variables, for example, inference of ocean velocity or density fields from Sea Surface Height (SSH) observations (Souza et al., 2025). Recently, methods such as BlindDPS (Chung, Kim, Kim, & Ye, 2023) perform posterior sampling from unknown measurement operators, though these methods may require a neural surrogate for the observation operator. The two conditional approaches require task-specific re-training, since the conditioning field varies based on the task (e.g., inference from a field obtained using a different observation operator). These models directly learn the complex observation-to-target map, and are particularly powerful when the conditioning fields are mutually informative fields obtained from highly nonlinear operators. All methods except \mathcal{M}_3 contain tuning parameters to control guidance (Section 5.3). In terms of cost, \mathcal{M}_1 , \mathcal{M}_2 and \mathcal{M}_3 require one neural network evaluation per pseudo-time step of the reverse process, while \mathcal{M}_4 requires two evaluations. Since \mathcal{M}_1 starts from an intermediate time, its cost scales with the number of pseudo-time steps, while \mathcal{M}_2 incurs additional overhead from Jacobian and gradient computation.

4. Quasi-Geostrophic Turbulence and Data

We now describe our high-resolution quasi-geostrophic turbulence simulations, including the dynamical regimes, numerical schemes, and parameters (Sections 4.1–4.2). The measurement models we employ (Section 4.3) then apply Large-Eddy Simulation (LES)-style spectral filtering and sparse observation operators to obtain low-resolution data (the observations). In Section (4.4), we describe the super-resolution and inference test cases and the high- and low-resolution data sets that are used to train the score-based DMs. In Section (4.5), we introduce a comprehensive suite of skill-metrics used to assess the performance of the score-based DMs. Once trained, these models will be used to reconstruct high-resolution fields from low-resolution data (Section 5).

4.1. Quasi-Geostrophic Dynamics

As a testbed for our generative models, we use the two-dimensional (2D) incompressible one-layer Quasi-Geostrophic (QG) equations, an idealized approximation for large-scale geophysical flows on a rotating surface (Vallis, 2017). These partial differential equations (PDEs) are widely used to simulate atmospheric and oceanic flows, and have been utilized for multiple machine learning applications such as surrogate modeling (Z. Li et al., 2020) and subgrid-scale closure modeling (Frezat et al., 2022; Guan et al., 2024; Ross et al., 2023; Srinivasan et al., 2024; Suresh Babu et al., 2025a) under various approximations.

In 2D, the non-dimensionalized QG PDE can be written in vorticity form as

$$\frac{\partial \omega}{\partial t} + J(\psi, \omega) = \frac{1}{Re} \nabla^2 \omega - \mu \omega - \beta \frac{\partial \psi}{\partial x} + F \quad (21)$$

where ω is the vorticity, which is our 2D field of interest, and ψ is the stream function. In this governing PDE (Equation 21), the nonlinear term $J(\psi, \omega)$ represents advection of vorticity by the flow itself. The diffusion term $\frac{1}{Re} \nabla^2 \omega$ represents turbulent eddy diffusivity, where Re is the Reynolds number. The β -term accounts for latitudinal variation of the Coriolis force due to the Earth's rotation, while the linear damping term $\mu \omega$ models linear bottom drag (Stommel, 1948). The forcing F injects energy at prescribed length scales to mimic sustained geophysical driving. Since the flow is assumed to be incompressible, the velocity field (u, v) can be recovered from the stream function via

$$(u, v) = \left(-\frac{\partial \psi}{\partial y}, \frac{\partial \psi}{\partial x} \right) \quad (22)$$

$$\omega = \nabla^2 \psi \quad (23)$$

To drive the system and set dominant length scales, we apply a periodic forcing,

$$F = k_f [\cos(k_f x) + \cos(k_f y)] \quad (24)$$

where k_f is the forcing wavenumber (Guan et al., 2023). We consider all terms in Equation 21 to simulate forced-dissipative turbulence and its vorticity dynamics.

We emphasize two flow regimes: the *eddy* regime and the *jet* regime (Ross et al., 2023). The *eddy* regime corresponds to $\beta = 0$, where the flow is isotropic and comprised of large coherent vortices. The *jet* regime corresponds to $\beta > 0$, where the flow consists of alternatively banded jets. In the *jet* regime, the β -effect forces the flow to organize into Rossby waves, creating a barrier to the upscale turbulent cascade of energy, leading to the destruction of coherent vortices (Rhines, 1975). This interaction between Rossby waves and Quasi-Geostrophic turbulence is known as the Rhines effect, and has been observed in the Earth's oceans (Galperin et al., 2004) and planetary circulations of Jupiter and Saturn (Theiss, 2006; Williams, 1978).

4.2. Numerical Schemes and Parameters

We simulate 2D vorticity fields by solving Equation 21 on a square domain of dimensions $[0, 2\pi] \times [0, 2\pi]$ with periodic boundary conditions. We numerically solve Equation 21 using a pseudo-spectral method (Orszag &

Table 2
Parameters of the Different Fully-Resolved (FR) Forced-Dissipative Quasi-Geostrophic (QG) Simulations for the Eddy and jet Regimes

| Regime | Re | β | dt | μ | k_f |
|--------|--------|---------|--------------------|--------------------|-------|
| Eddy | 10^3 | 0 | 1×10^{-3} | 2×10^{-2} | 2 |
| Eddy | 10^4 | 0 | 5×10^{-4} | 2×10^{-2} | 2 |
| Jet | 10^3 | 2.5 | 1×10^{-3} | 2×10^{-2} | 2 |
| Jet | 10^4 | 2.5 | 5×10^{-4} | 2×10^{-2} | 2 |

Note. All parameters are expressed in non-dimensional units defined in Section 4.2 (length-scale of $\frac{504}{\pi} \times 10^4$ m and time-scale of 1.2×10^6 s).

Israeli, 1974), with full 2/3 de-aliasing. For numerical stability, we use a semi-implicit second-order Adams-Bashforth-Crank-Nicolson (AB2-CN) scheme for time-stepping, where the diffusion, Coriolis, and bottom drag terms are treated implicitly while the advection and forcing terms are treated explicitly (Boyd, 2001).

Following Graham and Ringler (2013) and Frezat et al. (2022), we use the following scales for non-dimensionalization, corresponding to ocean meso-scales: a length-scale of $\frac{504}{\pi} \times 10^4$ m and a time-scale of 1.2×10^6 s. We simulate fully-resolved (FR) vorticity fields using a $N_{FR} \times N_{FR} = 512 \times 512$ numerical grid ($\Delta_{FR} = \frac{2\pi}{N_{FR}}$). To initialize simulations, we follow

Guan et al. (2022) and sample initial vorticity fields from a standard normal

distribution $\omega \sim \mathcal{N}(0, 1)$ with wavenumbers $k \in [3, 10]$, omitting small wavenumbers for numerical stability. We vary the Reynolds number between $Re = 10^3$ and $Re = 10^4$. Higher values of the Reynolds number increase the richness of turbulence, leading to a wider range of multi-scale interactions (Pope, 2001). For the eddy regime, we set $\beta = 0$, and for the jet regime, we set $\beta = 2.5$. For time-stepping, we use a non-dimensional time-step dt a function of Re and simulate a total time of $T = 100$. A list of non-dimensional parameters is provided in Table 2.

4.3. Down-Sampling and Sparse Observation Operators

For down-sampling, we apply coarsening operators on the vorticity fields obtained from the fully-resolved (FR) simulations. This coarsening operator could be of various types, and could act either in the physical space or the wavenumber space (Fourier domain). Down-sampling operations in physical space, such as max-pooling or average-pooling, have been considered in previous work (Fukami et al., 2019), but these correspond to non-local box filters in the wavenumber space, leading to spurious oscillations in the kinetic energy spectrum at high wavenumbers. In this work, we use spectral filtering operators utilized in Large-Eddy Simulations (LES) (Pope, 2001). We apply our down-sampling operator (\mathcal{H}_{full}) in wavenumber space by first applying a Gaussian filter ($\mathcal{H}_{Gaussian}$), followed by a cut-off filter ($\mathcal{H}_{cut-off}$), and then interpolate the resultant field to obtain the observed field (OF) on the low-resolution grid with resolution $\Delta_{OF} = \frac{2\pi \times \delta}{N_{FR}}$:

$$\begin{aligned} \mathcal{H}_{Gaussian}(k_{FR}) &= \exp\left(\frac{-k_{FR}^2 (\delta \times \Delta_{FR})}{6}\right) \\ \mathcal{H}_{cut-off}(k_{FR}) &= 0, \quad \forall k_{FR} > \frac{\pi}{\Delta_{OF}} \\ \mathcal{H}_{full} &\triangleq \text{interpolate}_{\Delta_{OF}} \circ \mathcal{H}_{cut-off} \circ \mathcal{H}_{Gaussian} \end{aligned} \quad (25)$$

where k_{FR} is the wavenumber corresponding to the fully-resolved grid and δ is the down-sampling scale (i.e., down-sampling from 64×64 to 16×16 implies $\delta = 4$). The Gaussian filter smoothly attenuates high wavenumber content, resembling subgrid-scale energy removal, while the sharp cut-off filter enforces a strict wavenumber limit. This combination retains large-scale coherent structures while discarding small-scale features. Such filtering is crucial to pose a realistic and physically meaningful super-resolution problem for geophysical turbulence (Z. Zhou et al., 2019; Srinivasan et al., 2024).

We also study the additional effect of sparse and gappy observing systems (\mathcal{H}_{sparse}) by defining a set of unobserved regions Ω and using a mask operator in physical space,

$$\begin{aligned} \mathcal{H}_{partial}(x, y) &= \omega_{fill}, \quad \forall x, y \in \Omega \\ \mathcal{H}_{sparse}(x, y) &= \mathcal{H}_{partial} \circ \mathcal{H}_{full} \end{aligned} \quad (26)$$

where ω_{fill} is a fill value corresponding to the unobserved regions. In our case, we set the fill value to 0, which is the mean of our vorticity fields over space and time. Martin et al. (2025) and T. Li et al. (2023) considered a similar partial observation operator to emulate satellite observations and study flow reconstruction with gappy

Table 3

Main Test Cases Utilized for Super-Resolution and Inference With Diffusion Models Corresponding to Target Dimensions ($N_{FF} = 64 \times 64$) and Observed Field Dimensions ($N_{OF} = 16 \times 16$), With Resultant Down-Sampling Scale ($\delta = 2^2$)

| Number | Regime | Re | Data (observations) |
|--------|--------|--------|--------------------------|
| 1 | Jet | 10^3 | Coarse-resolution field |
| 2 | Jet | 10^4 | Coarse-resolution field |
| 3 | Eddy | 10^3 | Coarse-resolution field |
| 4 | Eddy | 10^4 | Coarse-resolution field |
| 5 | Jet | 10^3 | Coarse and sparse, gappy |
| 6 | Jet | 10^4 | Coarse and sparse, gappy |
| 7 | Eddy | 10^3 | Coarse and sparse, gappy |
| 8 | Eddy | 10^4 | Coarse and sparse, gappy |

Note. The effect of other down-sampling scales ($\delta \in 2^1, 2^3$) are discussed in Section (5.1).

high-resolution data, respectively; however, we emphasize that here we apply these partial observations directly to the coarse-resolution output of the down-sampling operators defined in this section.

4.4. Data and Test Cases

To generate data for training, we simulate 500 different trajectories of the fully-resolved (FR) QG simulations ($N_{FR} = 512$) with different random initial conditions by varying the random seed. Training DMs is very *data-intensive*, since it requires sampling from the entire distribution of flow fields (Z. Wang et al., 2023). Hence, such large-dimensional flow fields require a lot of simulation data and computational resources for training. Therefore, we restrict our experiments to fields of dimensions $N_{FF} \times N_{FF} = 64 \times 64$, which are obtained by applying the down-sampling operator (Equation 25) with $\delta = 2^3$ to these simulations. These new down-sampled fields, that is, the *filtered fields* (FF) serve as the high-resolution ground truth for our super-resolution experiments. This reduces the number of snapshots and hence computational costs (in terms of both memory and number of operations) required for training. Alternatively, the size of the high-resolution target can also be restricted by simply selecting fields from a

smaller window of the original fully-resolved fields. However, the size of this window needs to be larger than the largest eddy of the simulated vorticity field (Fukami et al., 2019).

Snapshots (frames) of the vorticity field are saved every $\Delta T_{save} = 0.5$ non-dimensional time units, resulting in $\frac{T}{\Delta T_{save}} = 200$ snapshots. Upon analyzing the energy spectrum, we observe self-similarity after $t = 50$, and only utilize snapshots obtained post this time for our training and analysis (last 100 snapshots). Hence, our generative models are trained with $500 \times 100 = 50,000$ snapshots.

To study the performance of generative models (Table 1), we utilize test cases with two observation operators for all the simulations in Table 2: down-scaling by a scale of $\delta = 2^2$ with full observations (Equation 25), and a similar down-scaling with partial observations (Equation 26). To study the effect of sparse, gappy observations, we restrict the observed region Ω to cover 40% of the total domain. Utilizing the measurement model (1) with small Gaussian observation noise $\sigma_y = 0.01$, these test cases yield observed fields (OF) with dimensions $N_{OF} \times N_{OF} = 16 \times 16$, which serve as the conditioning field (input) to our DMs. Table 3 summarizes the 8 test cases considered. We perform super-resolution at a fixed time, utilizing observed fields obtained from the

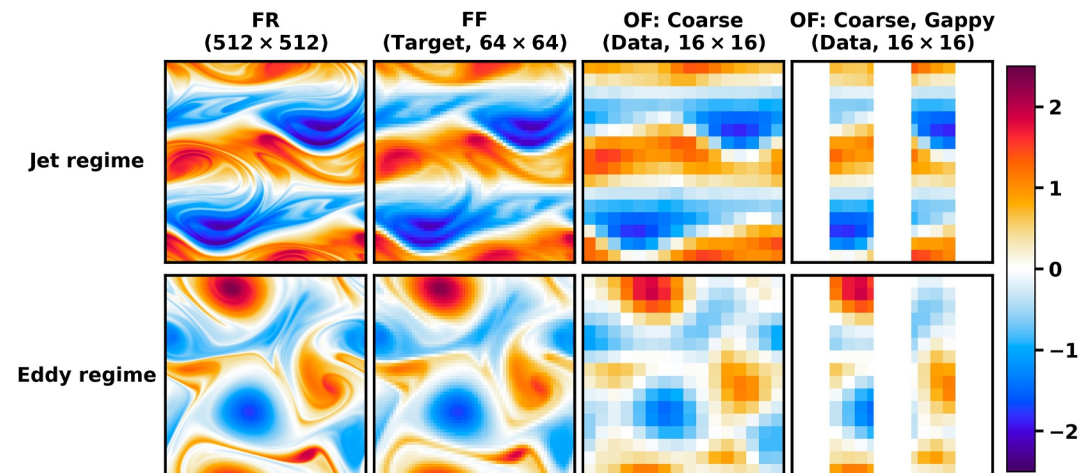


Figure 3. Test case description for super-resolution and inference in the jet (Top row) and eddy (Bottom row) regimes at $Re = 10^4$. Column 1: Fully-resolved (FR) vorticity field. Column 2: Filtered-field (filtered fields), which serves as the target for super-resolution. Column 3: Observed field (OF) under coarse-resolution observations. Column 4: Observed field (OF) under coarse, sparse and gappy observations. The observed fields serve as data (input) for super-resolution.

64×64 filtered vorticity fields of a trajectory initialized with a new initial condition, which was not utilized for training. A visual representation of test cases 2, 4, 6, and 8 is shown in Figure 3.

4.5. Skill Metrics

To assess the performance of the score-based DMs, we utilize three sets of complementary skill metrics: one is related to the generated vorticity fields, the second to statistical turbulence quantities, and the third to uncertainty quantification.

First, we assess the generated super-resolved fields using common norms such as the relative L_2 norm: $\frac{\|\omega_{FF} - \hat{\omega}_b\|_2}{\|\omega_{FF}\|_2}$ and measures of the accuracy of features of the reconstructed vorticity field (e.g., filament sharpness). We then check for cycle-consistency using the relative norm: $\frac{\|\omega_{OF} - \mathcal{H}(\hat{\omega}_b)\|_2}{\|\omega_{OF}\|_2}$, that is, does the filtered high-resolution diffusion output match the low-resolution input field? Cycle-consistency is essential to ensure that the observed field can indeed be obtained from the generated super-resolved field through the measurement model (T. Zhou et al., 2016).

Second, we compare time-averaged statistical quantities relevant to 2D turbulence. Following Frezat et al. (2022), we compute the angle-averaged kinetic energy spectrum $E(k)$ and enstrophy spectrum $Z(k)$ in the wavenumber space using

$$E(k) = \frac{1}{2} \int_{|\mathbf{k}|=k} \hat{\mathbf{u}}_i(\mathbf{k}) \hat{\mathbf{u}}_i^*(\mathbf{k}) dS(\mathbf{k}) \quad (27)$$

$$Z(k) = \frac{1}{2} \int_{|\mathbf{k}|=k} \hat{\omega}(\mathbf{k}) \hat{\omega}^*(\mathbf{k}) dS(\mathbf{k}) \quad (28)$$

where $k = \sqrt{k_x^2 + k_y^2}$ denotes the horizontal angular wavenumber, $\hat{\bullet}$ the Fourier transform, and \bullet^* the complex conjugate. We also compute the relative L_2 norms of the log-spectra. Because the β -effect induces anisotropic spectral energy transfers that are obscured by angle-averaging (Huang et al., 2001), we then compute contours of the pre-multiplied 2D kinetic energy spectrum,

$$E(k_x, k_y) = \lambda_x \lambda_y \frac{1}{2} \hat{\mathbf{u}}_i(\mathbf{k}_x, \mathbf{k}_y) \hat{\mathbf{u}}_i^*(\mathbf{k}_x, \mathbf{k}_y) \quad (29)$$

where $\lambda_x = \frac{2\pi}{k_x}$ and $\lambda_y = \frac{2\pi}{k_y}$ are wavelength factors to account for the logarithmic spacing in the wavenumber space (Towne et al., 2020). We also evaluate the mean effective resolution using the time-averaged noise-to-signal ratio of the spectrum (Ballarotta et al., 2019),

$$NSR(k) = \frac{Z_{error}(k)}{Z_{FF}(k)} \quad (30)$$

where Z_{error} is the enstrophy spectrum (28) of the error, $\hat{\omega}_0 - \omega_{FF}$, and Z_{FF} is that of the target vorticity, ω_{FF} . To study the non-Gaussian single-point statistics of turbulent flows, which is a statistical measure of intermittency (J. McWilliams, 2007; Wilczek & Friedrich, 2009), we compute the one-dimensional probability distribution function for the vorticity field using kernel density estimation (Węglarczyk, 2018).

Third, as our DMs can generate an ensemble of super-resolved vorticity fields, they can be utilized for uncertainty quantification (Leinonen et al., 2023; N. Chen et al., 2025). To quantify and verify the uncertainty, we compute

the ensemble mean, $\overline{\hat{\omega}_0} = \frac{1}{N_{ens}} \sum_{i=1}^{N_{ens}} \hat{\omega}_0^i$, ensemble standard deviation: $\sqrt{\frac{\sum_{i=1}^{N_{ens}} (\hat{\omega}_0 - \hat{\omega}_0^i)^2}{N_{ens}}}$, and root-mean squared error

(RMSE) of the generated super-resolved fields: $\sqrt{\frac{\sum_{i=1}^{N_{ens}} \|\omega_{FF} - \hat{\omega}_0^i\|_2^2}{N_{ens}}}$, where $N_{ens} = 16$ is the ensemble size and the superscript i indicates the i th ensemble member.

Table 4

Comparison of the Key Quantitative Skill Metrics for Super-Resolution From Coarse-Resolution Fields (Cases 1–4 of Table 3)

| Relative error metric | Scale | Re = 10 ³ | | | | Re = 10 ⁴ | | | |
|-----------------------------|----------------------|----------------------|--------------------|--------------------|--------------------|----------------------|--------------------|--------------------|--------------------|
| | | \mathcal{M}_1 | \mathcal{M}_2 | \mathcal{M}_3 | \mathcal{M}_4 | \mathcal{M}_1 | \mathcal{M}_2 | \mathcal{M}_3 | \mathcal{M}_4 |
| Jet regime | | | | | | | | | |
| Reconstructed vorticity (↓) | ($\times 10^{-1}$) | 1.05 ± 0.08 | 0.45 ± 0.07 | 0.19 ± 0.01 | 0.20 ± 0.01 | 2.13 ± 0.09 | 1.29 ± 0.07 | 0.69 ± 0.05 | 0.71 ± 0.02 |
| Ensemble Field Std. | ($\times 10^0$) | 0.05 | 0.04 | 0.01 | 0.01 | 0.09 | 0.09 | 0.05 | 0.05 |
| Cycle-consistency (↓) | ($\times 10^{-1}$) | 0.77 ± 0.05 | 0.30 ± 0.05 | 0.04 ± 0.01 | 0.05 ± 0.01 | 1.31 ± 0.07 | 0.61 ± 0.06 | 0.05 ± 0.03 | 0.04 ± 0.00 |
| Log-energy spectrum (↓) | ($\times 10^{-2}$) | 11.09 ± 1.16 | 4.14 ± 0.15 | 3.41 ± 0.40 | 3.67 ± 0.25 | 7.29 ± 0.74 | 1.81 ± 0.18 | 2.11 ± 0.93 | 2.36 ± 0.37 |
| Log-ensrophy spectrum (↓) | ($\times 10^{-2}$) | 17.09 ± 1.78 | 6.38 ± 0.23 | 5.25 ± 0.61 | 5.65 ± 0.38 | 11.85 ± 1.21 | 2.95 ± 0.29 | 3.43 ± 1.51 | 3.85 ± 0.61 |
| Eddy regime | | | | | | | | | |
| Reconstructed vorticity (↓) | ($\times 10^{-1}$) | 3.02 ± 0.08 | 2.65 ± 0.13 | 1.00 ± 0.06 | 1.01 ± 0.04 | 3.40 ± 0.10 | 3.22 ± 0.11 | 1.52 ± 0.06 | 1.53 ± 0.07 |
| Ensemble Field Std. | ($\times 10^0$) | 0.12 | 0.17 | 0.06 | 0.06 | 0.16 | 0.22 | 0.10 | 0.01 |
| Cycle-consistency (↓) | ($\times 10^{-1}$) | 1.94 ± 0.07 | 1.39 ± 0.10 | 0.09 ± 0.06 | 0.06 ± 0.00 | 1.96 ± 0.10 | 1.42 ± 0.09 | 0.04 ± 0.01 | 0.03 ± 0.00 |
| Log-energy spectrum (↓) | ($\times 10^{-2}$) | 6.31 ± 1.37 | 5.83 ± 0.20 | 6.27 ± 1.80 | 5.95 ± 0.94 | 5.67 ± 0.37 | 2.79 ± 0.74 | 2.63 ± 0.32 | 2.51 ± 0.30 |
| Log-ensrophy spectrum (↓) | ($\times 10^{-2}$) | 9.71 ± 2.12 | 8.99 ± 0.30 | 9.67 ± 2.77 | 9.18 ± 1.46 | 9.20 ± 0.60 | 4.53 ± 1.20 | 4.26 ± 0.51 | 4.07 ± 0.48 |

Note. Error metrics are shown as ensemble mean ± one standard deviation; the best-performing model for each metric is highlighted in bold.

5. Applications

To train the DMs, we utilize a U-Net architecture (Ronneberger et al., 2015). To provide conditioning, we concatenate the interpolated low-resolution field with the noisy field at each diffusion time-step following Saharia, Ho, et al. (2022). Details of neural architectures, training, and associated hyperparameters are in Appendix A.

Next, we present the results of the four generative DMs (Section 3) for super-resolution of turbulent vorticity fields from coarse-resolution fields (Section 5.1) and then from coarse, sparse, and gappy observations (Section 5.2). All these results are obtained using best-tuned intermediate times and guidance strengths. In Section (5.3), we study the sensitivity of the four DMs to their main tuning parameters.

5.1. Applications 1: Super-Resolution From Coarse-Resolution Fields

The first applications employ coarse-resolution vorticity fields as the observed field (data or input), corresponding to Cases 1–4 of Table 3, and assess whether the DMs can reconstruct subgrid-scale structure from coarse data (Fukami et al., 2019).

In the jet regime (Cases 1–2 of Table 3), strong β -effects organize the flows into anisotropic zonal bands with thin filaments, while in the eddy regime (Cases 3–4 of Table 3), the absence of β -effects allows more complex, nearly isotropic dynamics with thin filaments and streaks surrounding larger vortex cores (Kevlahan & Farge, 1997; Maltrud & Vallis, 1991). Table 4 summarizes the key quantitative skill metrics: relative error norms of the reconstructed fields, cycle-consistency error, and log-spectra, as well as the ensemble standard deviation of the reconstructed field for the jet and eddy regimes. On average, all models perform better in the jet regime; the presence of zonal bands helps the models identify the larger-scale structures than in the isotropic eddy regime.

Super-resolved vorticity fields: In both the jet and eddy regimes, and for both Re = 10³ and Re = 10⁴, the vanilla conditional model (\mathcal{M}_3) and the CFG approach (\mathcal{M}_4) achieve the lowest reconstruction error, followed by the guided unconditional models: guidance by modifying the score (DPS, \mathcal{M}_2) and guidance by modifying the initial conditions (SDEdit, \mathcal{M}_1). All four methods capture the larger-scale forced features (zonal jets and vortex cores for the jet and eddy regimes, respectively) of the target field driven by the forcing in Equations 21–24.

Figures 4a and 5a and their zoomed-in counterparts Figures 4b and 5b show that \mathcal{M}_1 leads to overly smooth ensembles which do not capture boundaries (large vorticity gradient regions), or fine structure (zonal filaments

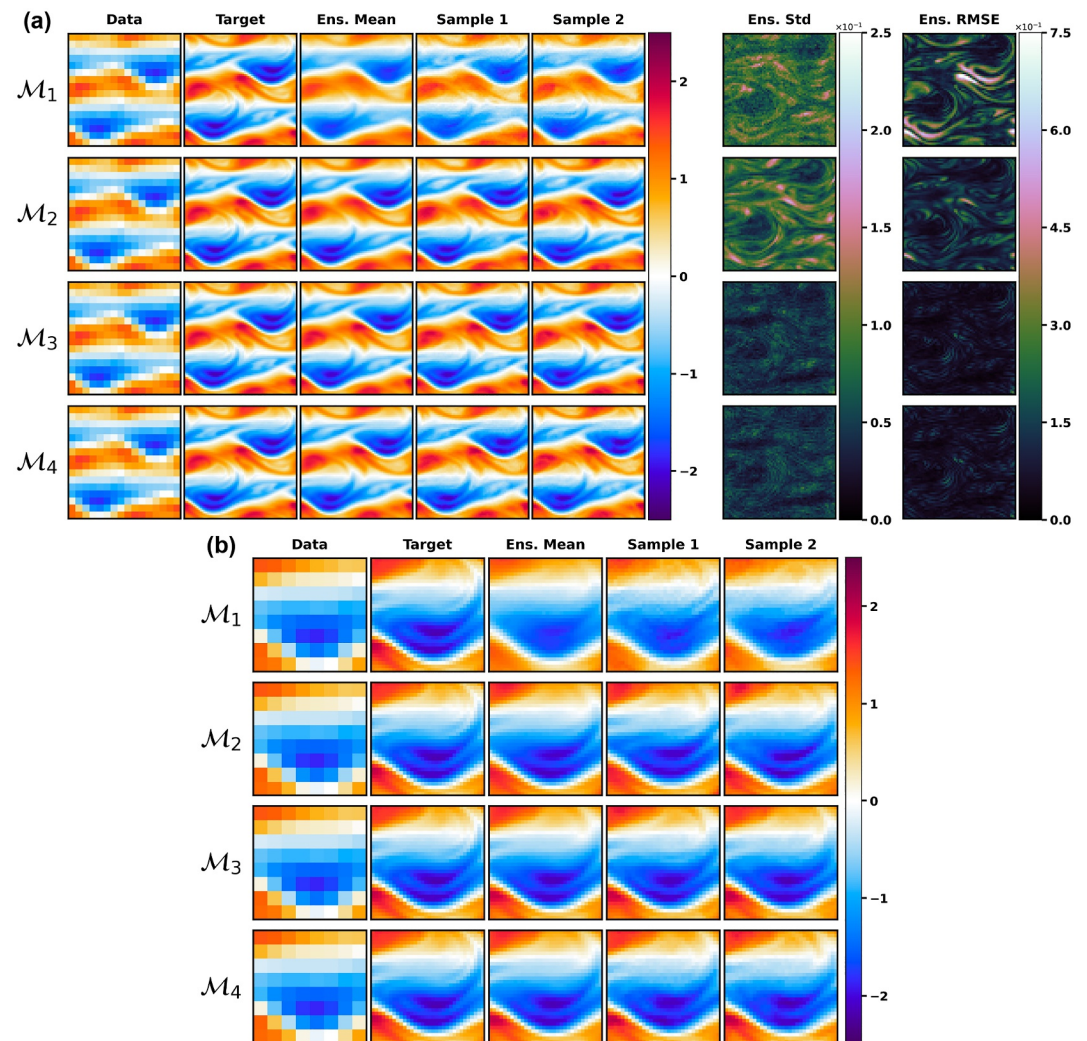


Figure 4. (a) Reconstructed vorticity fields for the four super-resolution approaches in the jet regime ($Re = 10^4$) from coarse-resolution fields for one snapshot (Test Case 2, Table 3). Column 1: Coarse observed vorticity field (data, OF). Column 2: Corresponding filtered vorticity field (target, FF). Column 3: Ensemble mean of the super-resolved diffusion model outputs. Columns 4 and 5: Two representative ensemble members. Columns 6 and 7: Ensemble point-wise standard deviation (Std.) and root-mean squared error, respectively. (b) Zoomed-in view of the upper-right quadrant for the first 5 columns of (a), highlighting fine-scale features. Data and target are identical for all models.

and filaments with streaks in the jet and eddy regimes, respectively), because \mathcal{M}_1 does not sample from the conditional reverse process (Equation 4), and only uses large-scale data features (Equation 11). \mathcal{M}_2 's ensemble better captures zonal jet and vortex core boundaries, but its mean still lacks fine structure (filaments at core edges) since \mathcal{M}_2 only approximates the conditional distribution (Equation 19) via the assumption in Equation 14. Sampling directly from the conditional reverse process (Equation 4), the conditional models \mathcal{M}_3 and \mathcal{M}_4 generate super-resolved ensembles qualitatively similar to the target that capture boundaries and most fine structures.

Uncertainty quantification: From the last two columns of Figures 4a and 5a, \mathcal{M}_3 and \mathcal{M}_4 exhibit the smallest standard deviation and RMSE (with $N_{ens} = 16$) versus the two guided unconditional models. For all models, standard deviation and error are considerably higher in the eddy regime than in the jet regime (Table 4), due to the more complex distribution of isotropic vorticity fields. Uncertainty is largest at the zonal jet and vortex core boundaries, where gradients are largest, since precise gradient super-resolution is nontrivial. Moreover, the standard deviations and errors show similar spatial patterns (except \mathcal{M}_1), so these models can serve as

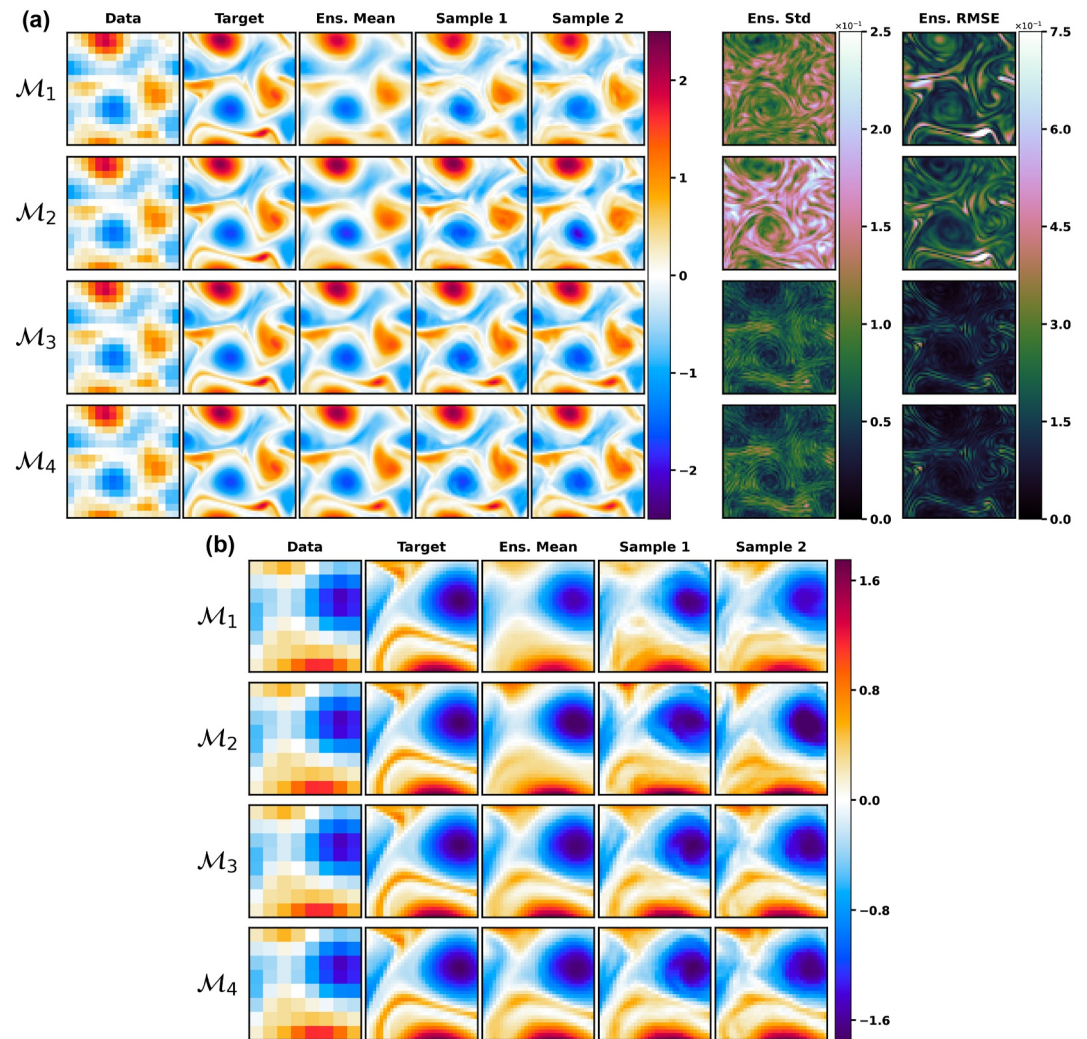


Figure 5. (a) As Figure 4a, but for super-resolution in the eddy regime ($Re = 10^4$) from coarse-resolution fields (Test Case 4, Table 3). (b) As Figure 4b, zoomed-in view of the lower-left quadrant for the first 5 columns of (a), highlighting fine-scale features. Data and target are identical for all models.

probabilistic estimators, providing both accurate mean predictions and spatially consistent uncertainty quantification. Similar spatial structures for vanilla conditional models were observed by Souza et al. (2025).

Cycle-consistency: In both the jet (Figure 6a) and eddy (Figure 6b) regimes, \mathcal{M}_1 has region-dependent cycle-inconsistency as it samples from the unconditional (Equation 3) and not the conditional (Equation 4) reverse diffusion process. \mathcal{M}_2 is less inconsistent than \mathcal{M}_1 since it approximates the conditional distribution through the measurement-matching term (Equation 17). Unfortunately, \mathcal{M}_2 assumes that the conditional expectation from the noisy and denoised samples is equivalent (Equation 14), which fails near sharp image gradients generated toward the end of the reverse process and breaks cycle-consistency. The error thus remains one to two orders of magnitude larger than the conditional models (Table 4). Despite a lack of explicit guidance, the two conditional models \mathcal{M}_3 and \mathcal{M}_4 are extremely cycle-consistent, as they sample directly from the true conditional distribution, Equations 4 and 20, respectively. While tuning the guidance strength of \mathcal{M}_4 could also lead to sampling from unphysical distributions (Karras et al., 2024), additional conditioning helps avoid some artifacts seen with \mathcal{M}_3 .

Time-averaged 2D turbulence statistics: The single-point vorticity distribution of both the jet (Figure 7b) and eddy (Figure 7f) regimes shows non-Gaussianity, arising from the intermittency of long-lived coherent vortices in

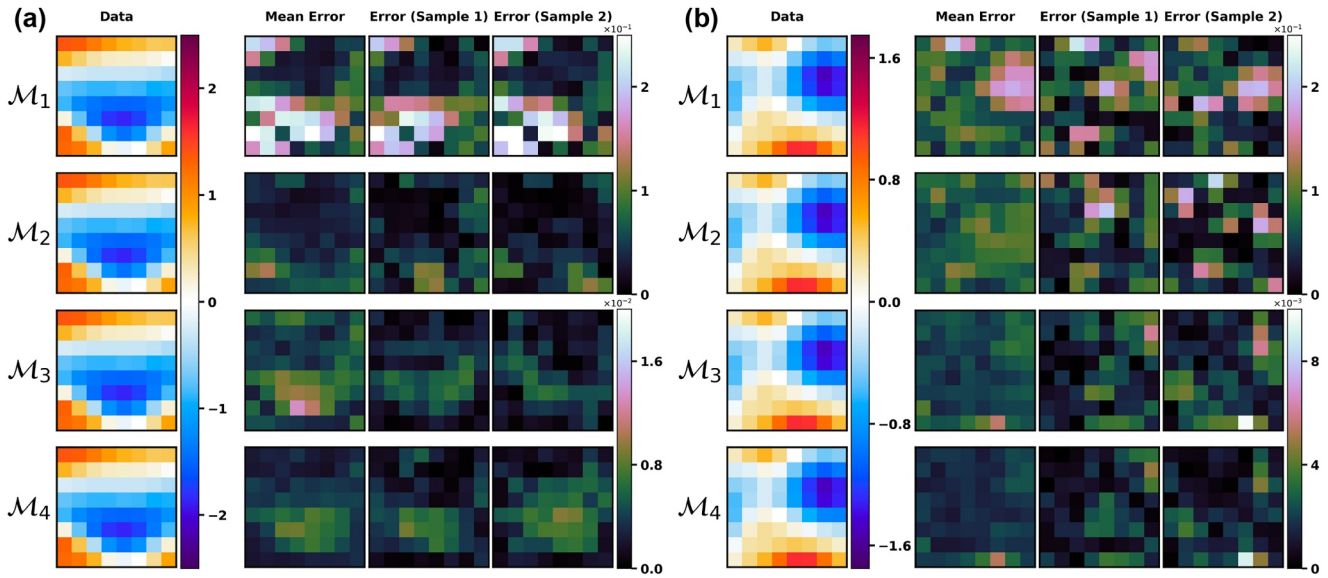


Figure 6. Cycle-consistency of the four diffusion modeling approaches, starting from coarse-resolution fields ($Re = 10^4$). Zoomed-in views of the (a) Upper-right quadrant for super-resolution in the jet regime (Test Case 2, Table 3) and (b) Lower-left quadrant for super-resolution in the eddy regime (Test Case 4, Table 3). Column 1: Coarse observed vorticity field (Data). Column 2: Ensemble mean of the errors of the filtered super-resolved diffusion output. Columns 3 and 4: Errors of 2 filtered ensemble members. The upper color bar limit of the error fields for \mathcal{M}_1 and \mathcal{M}_2 differs from that for \mathcal{M}_3 and \mathcal{M}_4 by orders of magnitude.

two-dimensional turbulence (Arbic & Flierl, 2003; Srinivasan et al., 2024). The unimodal eddy regime vorticity distribution becomes bimodal in the jet regime due to the zonal jets, examples of mixing barriers that introduce multimodality (David et al., 2017). \mathcal{M}_2 , \mathcal{M}_3 , and \mathcal{M}_4 accurately capture the single-point distributions in both regimes. The distribution tails are associated with how well the underlying coherent structures, such as vortices,

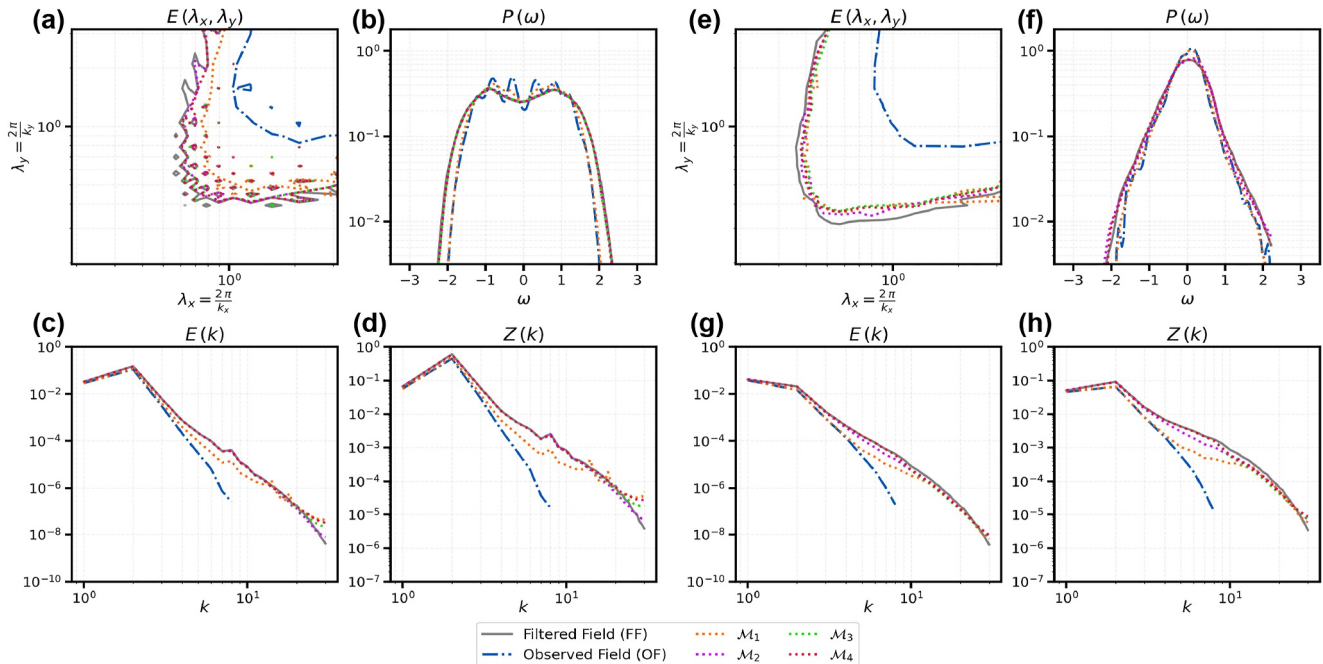


Figure 7. Time-averaged statistical quantities of the four models for super-resolution from coarse-resolution fields in the (a)–(d) jet regime ($Re = 10^4$) (Test Case 2, Table 3) and (e)–(h) eddy regime ($Re = 10^4$) (Test Case 4, Table 3) (a), (e): Premultiplied 2D kinetic energy spectrum contours that capture 95% of the energy. (b), (f): Single-point probability distribution function for the vorticity field (c), (d): Angle-averaged kinetic energy and (g), (h) entrophy spectra. Solid lines correspond to the filtered field (target), dash-dot lines to the observed field (data), and dotted lines to the diffusion model outputs.

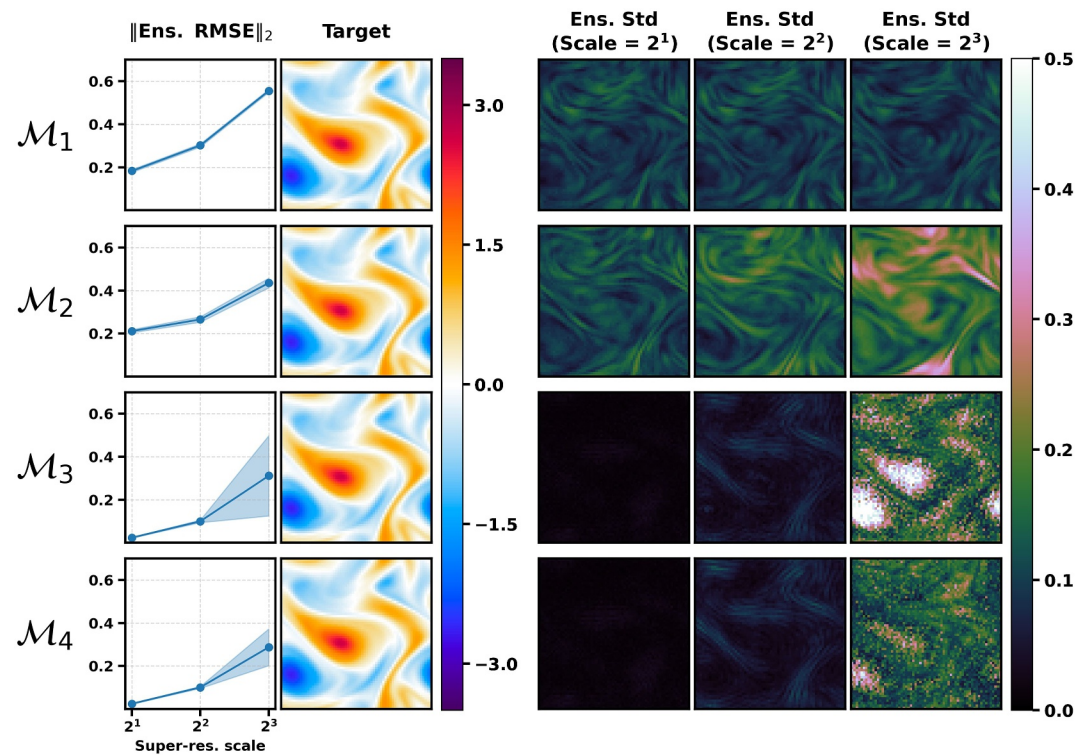


Figure 8. Super-resolution with the four diffusion modeling approaches at different down-sampling scales (δ) in the eddy regime ($Re = 10^3$) from coarse-resolution fields (Test Case 3, Table 3). Column 1: Norm of the relative ensemble root-mean squared error with shaded regions showing \pm one standard deviation for the reconstructed field with increasing down-sampling scales. Column 2: Target vorticity field. Columns 3–5: Ensemble standard deviations at different δ . Target is identical for all models, and all δ 's.

are captured (Farge et al., 1999; J. McWilliams, 2007). \mathcal{M}_1 is skewed toward the noisy observations used for initialization, and underestimates these rare events, as in Figures 4 and 5.

Figures 7a and 7e show the premultiplied 2D kinetic energy spectrum, with its characteristic jagged and smooth contours for the jet and eddy regimes, respectively (Huang et al., 2001; Nozawa & Yoden, 1997). \mathcal{M}_2 , \mathcal{M}_3 and \mathcal{M}_4 capture the spectrum shape well in both regimes, while \mathcal{M}_1 smooths the spectrum in the jet regime and underestimates the kinetic energy across all wavenumbers. From Figures 7c, 7d, 7g, and 7h, \mathcal{M}_2 , \mathcal{M}_3 , and \mathcal{M}_4 accurately capture the angle-averaged spectra at all wavenumbers, including smaller scales (i.e., the tails). In the jet regime, \mathcal{M}_3 and \mathcal{M}_4 show slightly higher energy at the tails. \mathcal{M}_1 is once again skewed by the data and shows steeper scaling than the target spectra; it matches the observed spectrum where available and then switches to the unconditional spectrum (stronger effect in the more complex eddy regime). For lower $Re = 10^3$ (figures not shown for brevity), the two conditional models capture the spectra better than \mathcal{M}_2 at all wavenumbers.

Super-resolution at different down-sampling scales: To assess performance under different super-resolution settings, we consider the harder eddy regime (here at $Re = 10^3$) with down-sampling scales $\delta = 2^1$ and 2^3 , in addition to $\delta = 2^2$ used in the main test cases discussed above. \mathcal{M}_1 is easy to implement and does not require any re-training since it re-uses the unconditional model with modified initial conditions that only depend on the data (Equation 11). \mathcal{M}_2 also does not require re-training for each δ , but requires access to the observation operators \mathcal{H} corresponding to each δ , and their gradients (Equation 17). \mathcal{M}_3 and \mathcal{M}_4 have to be re-trained from scratch for each δ , limiting their flexibility, but they do not require access to the gradients of each \mathcal{H} .

Figure 8 shows that the ensemble RMSE norm increases with δ , since coarser data contains less information for super-resolution, even for the large-scale vortices. For all methods except \mathcal{M}_1 , the standard deviation of the ensemble RMSE norm also increases with δ . Similar trends are also observed in the corresponding ensemble standard deviation of the generated super-resolved fields. As shown for $\delta = 2^2$, \mathcal{M}_1 is inaccurate for

Table 5

Comparison of the Key Quantitative Skill Metrics for Super-Resolution From Coarse and Sparse, Gappy Observations (Cases 5–8 of Table 3)

| Relative error metric | Scale | Re = 10 ³ | | | | Re = 10 ⁴ | | | |
|-----------------------------|----------------------|----------------------|--------------------|---------------------|--------------------|----------------------|-----------------|--------------------|--------------------|
| | | \mathcal{M}_1 | \mathcal{M}_2 | \mathcal{M}_3 | \mathcal{M}_4 | \mathcal{M}_1 | \mathcal{M}_2 | \mathcal{M}_3 | \mathcal{M}_4 |
| Jet regime | | | | | | | | | |
| Reconstructed vorticity (↓) | ($\times 10^{-1}$) | 6.07 ± 0.14 | 3.38 ± 0.12 | 0.45 ± 0.21 | 0.48 ± 0.13 | 6.68 ± 0.10 | 4.35 ± 0.12 | 0.99 ± 0.04 | 1.02 ± 0.05 |
| Ensemble Field Std. | ($\times 10^0$) | 0.20 | 0.06 | 0.04 | 0.04 | 0.22 | 0.10 | 0.06 | 0.06 |
| Cycle-consistency (↓) | ($\times 10^{-1}$) | 3.58 ± 0.09 | 2.24 ± 0.07 | 0.13 ± 0.13 | 0.10 ± 0.06 | 3.42 ± 0.12 | 2.40 ± 0.10 | 0.10 ± 0.01 | 0.10 ± 0.01 |
| Log-energy spectrum (↓) | ($\times 10^{-2}$) | 24.92 ± 0.30 | 12.76 ± 0.61 | 8.63 ± 2.35 | 9.56 ± 1.73 | 18.49 ± 0.26 | 7.50 ± 0.60 | 1.64 ± 0.18 | 1.82 ± 0.30 |
| Log-ensrophy spectrum (↓) | ($\times 10^{-2}$) | 38.41 ± 0.47 | 19.68 ± 0.95 | 13.30 ± 3.62 | 14.73 ± 2.67 | 30.09 ± 0.43 | 12.22 ± 0.98 | 2.67 ± 0.29 | 2.95 ± 0.50 |
| Eddy regime | | | | | | | | | |
| Reconstructed vorticity (↓) | ($\times 10^{-1}$) | 7.75 ± 0.08 | 4.43 ± 0.53 | 2.47 ± 0.33 | 2.39 ± 0.16 | 6.83 ± 0.11 | 6.42 ± 0.21 | 3.13 ± 0.21 | 3.10 ± 0.18 |
| Ensemble Field Std. | ($\times 10^0$) | 0.11 | 0.24 | 0.14 | 0.13 | 0.16 | 0.19 | 0.17 | 0.17 |
| Cycle-consistency (↓) | ($\times 10^{-1}$) | 3.77 ± 0.12 | 1.80 ± 0.17 | 0.16 ± 0.01 | 0.16 ± 0.01 | 3.08 ± 0.11 | 2.53 ± 0.14 | 0.06 ± 0.01 | 0.05 ± 0.00 |
| Log-energy spectrum (↓) | ($\times 10^{-2}$) | 14.86 ± 1.11 | 5.82 ± 0.27 | 5.84 ± 0.43 | 6.33 ± 0.58 | 6.12 ± 0.49 | 5.03 ± 0.56 | 3.13 ± 0.50 | 2.76 ± 0.40 |
| Log-ensrophy spectrum (↓) | ($\times 10^{-2}$) | 22.88 ± 1.72 | 8.97 ± 0.42 | 8.99 ± 0.66 | 9.75 ± 0.89 | 9.86 ± 0.80 | 8.12 ± 0.91 | 5.06 ± 0.80 | 4.46 ± 0.65 |

Note. Error metrics are shown as ensemble mean ± one standard deviation; the best-performing model for each metric is highlighted in **bold**.

super-resolution and generates non-physical fields, and hence its standard deviation does not vary much with δ . The two conditional models \mathcal{M}_3 and \mathcal{M}_4 again show the smallest error norms followed by \mathcal{M}_2 , then \mathcal{M}_1 . For $\delta = 2^1$, the observations contain sufficient fine-scale information for the conditional models to near-perfectly reconstruct fine-scale features (e.g., less than five percent error). For $\delta = 2^3$, \mathcal{M}_3 and \mathcal{M}_4 show the largest error norm standard deviations, predict the largest ensemble standard deviation fields, and their error patterns directly relate to the target field (FF). While expected, these error properties indicate that learning the conditional distribution between the target and the data is difficult at extremely coarse resolutions ($N_{OF} = 8 \times 8$).

5.2. Applications 2: Super-Resolution and Inference From Coarse and Sparse, Gappy Data

We now discuss applications of full field super-resolution inference from coarse and sparse, gappy observations of the vorticity field, corresponding to Cases 5–8 of Table 3. These synthetic data enable simple Observing System Simulation Experiments (OSSEs) analogous to real OSSEs (Errico et al., 2013) used for data assimilation of satellite-derived surface measurements such as SWOT (Morrow et al., 2019) and DUACS (Taburet et al., 2019) altimetry data. We avoid realistic scenarios where vorticity is inferred only from noisy, measured SSH, and instead consider an idealized setting with direct coarse, gappy vorticity measurements. The idealized testbed enables direct model comparison without obfuscation from noise and indirect measurement, or compute-intensive training with high-resolution ocean reanalysis fields and direct satellite observations.

Table 5 summarizes the key quantitative skill metrics. Overall, the errors of all approaches are significantly higher (more than $2 \times$) than in the corresponding coarse-resolution cases (Table 4). The eddy regime from coarse and sparse, gappy observations is the most challenging of all test cases we consider: the target fields have the most complicated distribution, while the sparse data offers the least information for super-resolution. This results in relatively lower skill metrics for all methods when compared to the jet regime cases. We nonetheless find that trends in error metrics across models in this application remain as in Section (5.1). In the following discussion, we highlight only the key findings and differences. For some methods and test cases, we will find that the coarse, sparse data are insufficient to infer the super-resolved field.

Super-resolved vorticity fields: From Figures 9 and 10, \mathcal{M}_1 fails to generate physically valid fields, struggling to capture even the large-scale zonal jets and vortex cores. While SDEdit has shown promise in stroke-based editing and synthesis in computer vision, it performs poorly when initialized with gappy data with low-to-moderate noise levels. In fact, \mathcal{M}_1 does not provide any theoretical guarantees of convergence to the posterior. Though

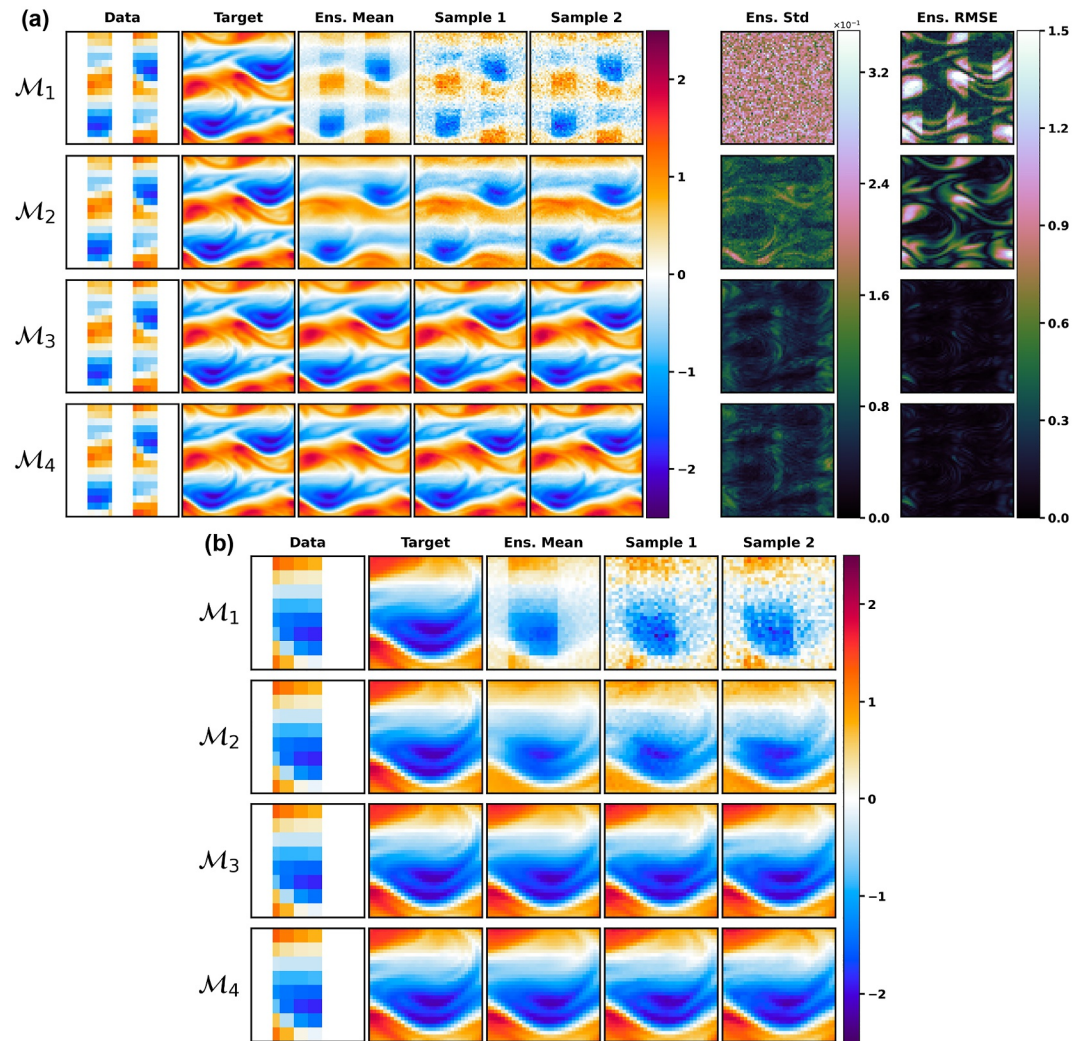


Figure 9. As Figure 4, but for super-resolution and inference in the jet regime ($Re = 10^4$) from coarse and sparse, gappy observations (Test Case 6, Table 3).

initializing with higher levels of noise could lead to more realistic fields, this approach would be akin to unconditional generation (Meng et al., 2021), which results in low cycle consistency and larger relative errors due to phase incoherence (Section 5.3). \mathcal{M}_2 also does not infer vortex core boundaries well and generates filaments with unphysical noise. While \mathcal{M}_2 can estimate the state given both full-coarse fields and sparse, gappy, high-resolution observations (Martin et al., 2025), here it underperforms with the limited observations. We also note that in the eddy regime at $Re = 10^4$, the sparse, gappy observations are not sufficiently informative to infer all finer-scale structures in the data gaps, leading to greater ensemble variability of the reconstructed filaments.

Uncertainty quantification: For all models, in both regimes, the ensemble RMSE is highest in the unobserved regions, followed by observed boundary regions. We analyze the RMSE against the variability of the target in Section (5.2.1). For \mathcal{M}_3 and \mathcal{M}_4 , once again, the ensemble standard deviation spatially follows the ensemble RMSE; these models can identify low-accuracy regions even in this challenging case (Souza et al., 2025).

Cycle-consistency: Remarkably, \mathcal{M}_3 and \mathcal{M}_4 are cycle-consistent even in this sparse observation application (Figure 11). So direct conditional information can enforce even out-of-distribution cycle-consistency through training. We investigate conditioning further in Section (5.2.1).




Figure 10. As Figure 5, but for super-resolution and inference in the eddy regime ($Re = 10^4$) from coarse and sparse, gappy observations (Test Case 8, Table 3).

Time-averaged 2D turbulence statistics: Unlike the super-resolution applications, \mathcal{M}_1 completely misses the bimodal distribution in the jet regime (Figure 12b). \mathcal{M}_2 captures the bimodality but significantly underestimates large values of the vorticity (Figures 12b and 12f). Similarly, \mathcal{M}_1 completely mispredicts the shape of the spectrum contours in the jet regime (Figure 12a) while \mathcal{M}_2 underestimates the energy at all wavenumbers in the eddy regime (Figure 12e). For the kinetic energy and enstrophy spectra (Figures 12c, 12d, 12g, and 12h), \mathcal{M}_1 and \mathcal{M}_2 shift the spectra significantly to smaller scales since they generate fields with significant fine-scale noisy artifacts.

Finally, we compute the time-averaged noise-to-signal ratio of the spectrum (Equation 30). A threshold ratio of 0.5 is used to estimate the mean effective resolution (Ballarotta et al., 2019). Compared to the prediction-to-signal spectral magnitude ratio (Chelton et al., 2011), the effective resolution is a stricter measure since it takes into account both the amplitude and phase coherence of resolved features. Figure 13 shows that the unconditional models' effective resolution is significantly degraded compared to that of the conditional models. The conditional effective resolution seems to match the statistical limit where ensemble variability (Figures 9 and 10) drives error, across both the super-resolution and inference applications. The conditional spectra show appropriate scaling, but with large noise-to-signal ratios beyond the effective resolution, suggesting that these models may overestimate eddy and filament magnitudes. The loss of phase coherence could be due to the probabilistic nature of DMs, since each ensemble member represents a plausible vorticity field.

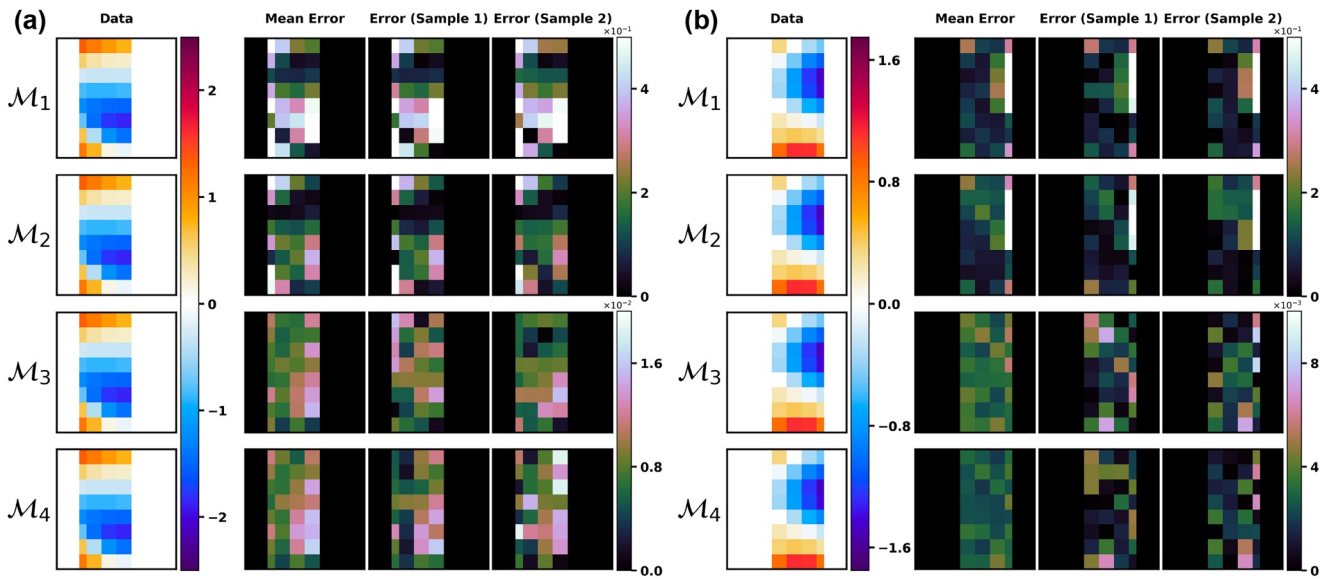


Figure 11. As Figure 6, but with zoomed-in views of the (a) Upper-right quadrant for super-resolution in the jet regime ($Re = 10^4$) from coarse and sparse, gappy observations (Test Case 6, Table 3) (b) Lower-left quadrant for super-resolution in the eddy regime ($Re = 10^4$) from coarse and sparse, gappy observations (Test Case 8, Table 3).

5.2.1. Propagation of Observational Information

We now analyze how information from the coarse and sparse, gappy observations is propagated to the unobserved regions during the reverse stochastic processes (Equations 3 and 4). For linear inverse problems, where $H = H \in \mathbb{R}^n \times \mathbb{R}^d$, Dou and Song (2024) showed an equivalence between posterior sampling from the conditional distribution and Bayesian filtering. Under this Gaussian approximation, the Kalman filter (Kalman, 1960) provides an optimal solution to the filtering problem, where observational information is propagated through the

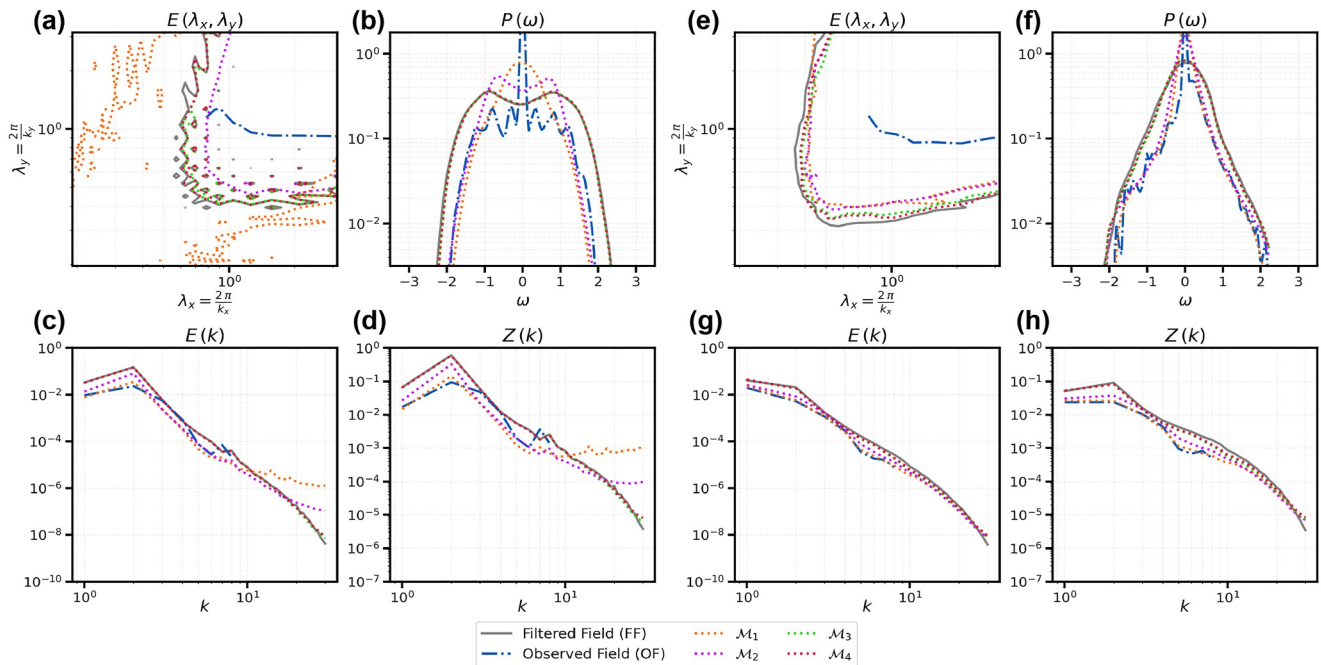


Figure 12. As Figure 7, but for super-resolution and inference from coarse and sparse, gappy observations in the (a–d) jet regime ($Re = 10^4$) (Test Case 6, Table 3) and (e–h) eddy regime ($Re = 10^4$) (Test Case 8, Table 3).




Figure 13. Time-averaged noise-to-signal ratios of the spectrum for super-resolution and inference from coarse and sparse, gappy observations in the (a) jet regime ($Re = 10^4$) (Test Case 6, Table 3) and (b) eddy regime ($Re = 10^4$) (Test Case 8, Table 3).

covariance terms appearing in the Kalman gain, namely the covariance of the forecast error and covariance of the observational error. In the ensemble setting, the forecast error covariance is estimated directly from the ensemble members (Evensen, 2003; Lermusiaux & Robinson, 1999). Similarly, a term analogous to the posterior forecast error covariance, $Cov[x_0|x_r, y]$ (Equation 31), can be derived analytically for score-based DMs (Boys et al., 2023) using the conditional reverse process score (Equation 7).

$$Cov[x_0|x_r, y] = \frac{\sigma_t^2}{\mu_t^2} [I_n + \sigma_t^2 \nabla_{x_t}^2 (\log p_t(x_t|y))] \quad (31)$$

Generally, it is computationally intractable to compute the Hessian of the conditional score, $\nabla_{x_t}^2 (\log p_t(x_t|y))$ (Rissanen et al., 2024).

Figure 14 shows the evolution of the reverse diffusion process for Test Case 6 (jet regime). For \mathcal{M}_1 , we find that its initialization at $t_i = 0.25$ preserves the large-scale features of the data. However, the unconditional reverse process (Equation 3) does not propagate this information to the unobserved regions, since it only utilizes the

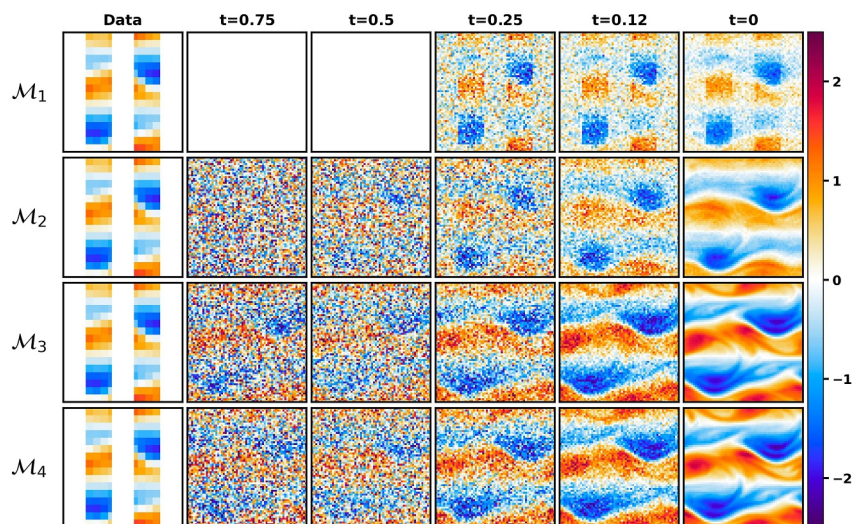


Figure 14. Evolution of vorticity fields during the reverse diffusion process from coarse and sparse, gappy observations in the jet regime ($Re = 10^4$, Test Case 6, Table 3).

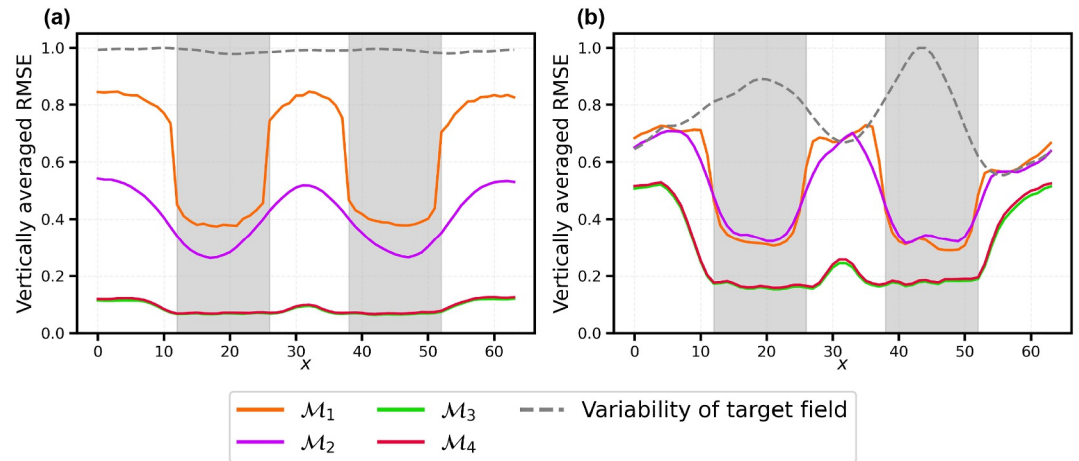


Figure 15. Time and vertically averaged root-mean squared error compared to the variability of the target for super-resolution and inference from coarse and sparse, gappy observations (data regions shaded in gray) in the (a) jet regime ($Re = 10^4$) (Test Case 6, Table 3) and (b) eddy regime ($Re = 10^4$) (Test Case 8, Table 3).

unconditional score, which is independent of observations, leading to poor reconstruction of large and fine-scale features. For \mathcal{M}_2 , we find some noisy but recognizable zonal jets at $t = 0.5$, which indicates the propagation of observational information to unobserved regions, albeit gradually and more diffuse compared to the conditional models. This could be because \mathcal{M}_2 approximates the true covariance (Equation 31) with just $\frac{\sigma^2}{\mu_t^2} \mathcal{I}_n$ in unobserved regions (Rozet & Louppe, 2023). For the two conditional models, we observe large, noisy zonal jets at much earlier diffusion pseudo-times ($t = 0.75$), with sharper jets visible at $t = 0.5$. This effect is also observed in the spectra at intermediate diffusion pseudo-times (figures not shown), where the conditional models show higher energies at larger scales much earlier compared to the unconditional models. This strong conditioning could explain their cycle-consistency (Figure 11).

Figure 15 shows the variation of the vertically averaged RMSE of the four DMs with distance from observed regions. We observe similar spatial trends in RMSE as in Figures 9 and 10. However, in the harder eddy regime, the RMSE of the two unconditional methods approaches the variability of the target field in the middle of the unobserved regions, indicating that these models do not propagate information there effectively. For \mathcal{M}_2 , this can be explained by analyzing Equation 19 in unobserved regions where the measurement matching term $\nabla_{x_t}(\|y - \mathcal{H}(\mathbb{E}[X_0|X_t = x_t])\|^2)_{\text{unobserved}} \rightarrow 0$ (In this case, the innovation vector $y - \mathcal{H}(\mathbb{E}[X_0|X_t = x_t])_{\text{unobserved}}$ also tends to 0 everywhere away from data). Hence, we obtain Equation 32.

$$\nabla_{x_t}(\log p_t(x_t|y))_{\mathcal{M}_2, \text{unobserved}} \simeq \nabla_{x_t}(\log p_t(x_t)) \quad (32)$$

Therefore, unobserved regions can lose guidance, leading to unconditional vorticity field generation. As in traditional splitting schemes, the reverse process can only propagate information through the interaction between the unconditional score and the measurement matching term, with the observation operator's physical bandwidth (e.g., width of tracks) determining the propagation rate. Only in the continuous-time limit of infinite diffusion steps can \mathcal{M}_2 theoretically approximate a conditional model. In practice, the reverse diffusion process is run with Langevin Monte Carlo corrections (Appendix A), which allow global interaction, improving information propagation (as seen in the jet regime in Figure 9). We also studied the correlation between observations and unobserved regions (figure not shown), with the conditional models exhibiting stronger statistical convergence than the unconditional ones.

5.3. Sensitivity Studies: Effect of Tuning Parameters

Finally, we study the sensitivity of the four models to their tuning parameters and provide their optimal values. All above results used optimally tuned intermediate times (for \mathcal{M}_1) and guidance strengths (for $\mathcal{M}_2, \mathcal{M}_4$). \mathcal{M}_3 does not require any tuning; it only relies on training. The choice of intermediate time and guidance strength greatly




Figure 16. Sensitivity of (a) \mathcal{M}_1 (guidance by modifying the initial condition) to the intermediate pseudo-time t_i , (b) \mathcal{M}_2 (guidance by modifying the score) to the guidance strength C , and (c) \mathcal{M}_4 (classifier-free guidance) to the guidance strength w .

controls the quality of the generated super-resolved samples, and trades off fidelity (i.e., physical plausibility and sharpness of the generated fields) for consistency (i.e., cycle-consistency with the observations or data). Since the denoiser neural networks, Equations 8 and 9, are trained independently of these parameters, their weights are unaffected by this fine-tuning. This allows for great flexibility, as the intermediate time and guidance strengths can be optimized and re-tuned post-training for each task (application), or even for each snapshot within an application (Chung, Kim, McCann, et al., 2023).

An alternate approach to improve the fidelity of the generated samples is to optimize the initial Gaussian noise used for the reverse process (Qi et al., 2024). However, this is computationally expensive, and hence, we use a fixed seed to generate Gaussian initial conditions and focus on fine-tuning only the intermediate time and guidance strengths. We optimize over all snapshots of super-resolution by sweeping over a range of values and choosing the parameter value that minimizes the L_2 norm of the error of the super-resolved fields. For brevity, we discuss results below for super-resolution in the jet regime ($Re = 10^4$) from coarse-resolution field observations. Similar trends hold for all other applications and test cases.

5.3.1. \mathcal{M}_1 : Modifying the Initial Condition (SDEdit)

The main parameter to be tuned is the intermediate time t_i . Figure 16a shows that choices of t_i close to 0 lead to unphysical super-resolved fields. Since $t_i = 0$ corresponds to the final diffusion pseudo-timestep, the model does not get to denoise the input field (data). On the other hand, choices of t_i close to $T = 1$ ($t_i \gtrsim 0.75$) lead to fields that lose track of large-scale features of the data, and hence generate non-cycle-consistent super-resolved fields (Meng et al., 2021). Setting $0.1 \lesssim t_i \lesssim 0.6$ generates relatively cycle-consistent super-resolved fields, and we obtain the smallest relative error at $t_i = 0.25$.

5.3.2. \mathcal{M}_2 : Modifying the Score (Diffusion Posterior Sampling)

The main parameter to be tuned here is the guidance strength, C . From Equations 18 and 19, using a value of $C \rightarrow \infty$ is equivalent to simply utilizing an unconditional model with no guidance. Figure 16b shows that setting $C \gtrsim 2 \times 10^{-2}$ generates non-cycle-consistent fields. Using $C \rightarrow 0$ corresponds to extremely strong guidance, which leads to over-saturation and unphysical noisy artifacts (Chung, Kim, McCann, et al., 2023), as seen in the figure for $C \lesssim 2 \times 10^{-6}$. Here, the optimal value of C was found to be 2×10^{-5} , which sufficiently trades off fidelity for consistency.

5.3.3. \mathcal{M}_3 : Vanilla Conditional Model

For this approach, the model directly learns the conditional score through training, and hence, there is no additional parameter to fine-tune.

5.3.4. \mathcal{M}_4 : Classifier-Free Guidance

The main parameter to be tuned is the guidance strength, w . Setting $w = 0$ corresponds to sampling only from the unconditional model, while $w = 1$ corresponds to sampling only from the conditional model (Ho & Salimans, 2021), as seen in Figure 16c. Typically, the value of w is set to be greater than 1. We obtain the smallest relative error at $w = 1.25$. For our applications, a value of $w \gtrsim 2$ leads to the generation of spurious high-frequency features. This can be explained by the fact that CFG performs a linear combination of the conditional and unconditional scores, and therefore, fields are generated from an approximated distribution that may not be physically meaningful (Karras et al., 2024). Using a very large value of w leads to unphysical noisy artifacts as shown in the last column.

6. Conclusions and Discussion

We investigated super-resolution and inference of quasi-geostrophic turbulence under the β -plane approximation with four generative diffusion modeling approaches—two guided unconditional models and two conditional models. The conditional models, vanilla (\mathcal{M}_3) and CFG (\mathcal{M}_4), are stable and generate accurate ensembles of super-resolved vorticity fields across all dynamical regimes (eddy and jet) and Reynolds numbers ($\text{Re} = 10^3, 10^4$). Moreover, they predict correct turbulence statistical quantities such as the 1D probability distribution function and kinetic energy spectra and enstrophy spectra (including the tails), and generalize to unseen sparse and gappy coarse observations. These generated ensembles are remarkably cycle-consistent with the coarse observations, despite not being explicitly optimized for during training. These two approaches show promise for deployment in geophysical inference applications where capturing accurate fine-scale features is critical, especially when the observation operator is nonlinear or indirect. Though these advantages come at the cost of application-specific re-training, they do not require knowledge of the observation operators for training or generation. On the other hand, if computational resources do not allow for re-training, guidance by modifying the score (\mathcal{M}_2) trades off fidelity (sharpness) for cycle-consistency and generates smoothed fields in applications with coarse observations but requires access to gradients of the observation operator at each diffusion pseudotime-step. However, \mathcal{M}_2 does not propagate observational information effectively to unobserved regions, especially when the observation operator \mathcal{H} is sparse and not global. Although easily implemented and computationally cheap, guidance by modifying initial conditions (\mathcal{M}_1) can not produce correct statistics, and fails in challenging applications with sparse and gappy data.

Computationally, our DMs are fast super-resolution and inference schemes (\mathcal{M}_1 takes $\sim 0.1 s$, $\mathcal{M}_2 \sim 2 s$, $\mathcal{M}_3 \sim 0.3 s$, and $\mathcal{M}_4 \sim 0.7 s$ on an NVIDIA A6000 GPU, computed by averaging 30 independent evaluations). \mathcal{M}_2 takes significantly longer than the others as it requires the computation of additional gradients to evaluate the measurement matching term. A major advantage of our super-resolution framework is that it does not require auto-differentiability or access to numerical solvers, and only needs either paired high- and coarse-resolution fields or access to the observation operator.

Considering the future, our results with sparse and gappy observation systems show that, for conditional DMs, the ensemble standard deviations are good estimators of the actual errors of predictions. This capability could be leveraged for data assimilation and adaptive sampling, that is, identify the most optimal types and locations of observations for inference (Lermusiaux, 2007; Lermusiaux et al., 2017). Another opportunity is the super-resolution of real dynamics such as ring formation, eddy separation, and submesoscale interactions for the Gulf Stream system (E. Chassignet & Marshall, 2008; Gula et al., 2019; Mensa et al., 2013), Loop Current system (Bracco et al., 2019; Nickerson et al., 2022; Oey et al., 2005), or other highly-dynamic mesoscale systems. Related possibilities include the diagnosis of submesoscales (Archer et al., 2025; Z. Zhang et al., 2024) and the inference of subsurface fields (Klemas & Yan, 2014; Lermusiaux et al., 2011; Moore et al., 2019) using sparse, gappy, coarse or indirect observations. However, several major limitations must be overcome for all these real applications to materialize. First, training requires access to high-resolution numerical simulations and reanalyzes (Moore et al., 2019; Storto et al., 2019) and access to multivariate or multi-modal observations (Lermusiaux

et al., 2002; Qu et al., 2024). Second, our diffusion architecture requires extensions such as multiscale patch-wise modeling (Brenowitz et al., 2025) or modeling in a latent space (Rombach et al., 2022) to handle multivariate dynamics and larger ocean and atmosphere regions. Third, novel robust evaluation metrics (Ballarotta et al., 2019) are likely necessary because the magnitudes and phases of the generated fields, the quantities derived from them, and the multivariate relations among them may differ from those of real fields. Ultimately, high-resolution real-world observational data and dedicated observing campaigns would be useful to evaluate and compare the performance of the different generative models, even if the training data were themselves provided by high-resolution simulations. Extensions to spatio-temporal super-resolution and inference can be investigated through the training of guided video diffusion (Ho et al., 2022) and hybrid-coupling with other stochastic numerical models (Lu & Lermusiaux, 2021) or neural surrogates (Rajagopal et al., 2023). Our generative diffusion is also promising for stochastic data-driven closure modeling with uncertainty estimates, for example, for stochastic subgrid forcing (Lermusiaux, 2006; Perezhugin et al., 2023; Zanna et al., 2017).

Appendix A: Diffusion Model Architecture and Training

To train the unconditional DM, we utilize a U-Net architecture (Ronneberger et al., 2015). We use channels of size [32, 64, 128, 256]. Following Nichol and Dhariwal (2021); Rozet and Louppe (2023), we use the following cosine noise schedule,

$$\omega = \arccos\sqrt{10^{-3}} \quad (\text{A1})$$

$$\mu(t) = \cos^2(\omega t) \quad (\text{A2})$$

$$\sigma(t) = \sqrt{1 - \mu(t)^2} \quad (\text{A3})$$

For training and sampling, we utilize $N_{diff} = 64$ uniform timesteps with diffusion time $0 < t < 1$ with $N_{corr} = 2$ Langevin Monte Carlo correction steps per diffusion pseudo-time-step, with correction size of 0.3. Hence, for sampling, the unconditional DM requires $N_{diff} \times (1 + N_{corr}) = 64 \times 3 = 192$ forward passes of the trained neural network. All the vorticity fields were normalized to have zero mean and unit variance for training. The unconditional models were trained separately for each combination of regime and Reynolds number, but were shared across the two types of observations (Table 3). Transfer learning between regimes and Reynolds numbers and its mechanism are areas of active research in subgrid modeling (Ross et al., 2023; Subel et al., 2023), but are not explored in this work.

To train the conditional model, we follow the same procedure as for the unconditional model, but incorporate conditioning by concatenating the interpolated low-resolution field to the noisy field at each diffusion time-step following Saharia, Ho, et al. (2022). We reiterate that conditional models are application-specific and were trained separately for each main test case (Table 3) and down-sampling scale (δ). We utilize the AdamW optimizer (Loshchilov & Hutter, 2017) with a cosine learning rate scheduler (Loshchilov & Hutter, 2016) with learning rate = 0.0002. We use a batch size of 64 and train for 250 epochs. We use NVIDIA A6000 GPUs for training, with each epoch taking ~ 15 minutes of training on a single GPU. We experimented by adding attention to the bottleneck layer of the U-Net and providing the diffusion pseudo-time t as an input to the network, but did not observe noticeable performance improvement.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Availability Statement

Codes for quasi-geostrophic simulations are available via the MIT license (<https://github.com/ananthu545/qg-2d>) and archived at <https://doi.org/10.5281/zenodo.17282193> (Suresh Babu et al., 2025c). Data from the simulations are available at Zenodo via <https://doi.org/10.5281/zenodo.15742145> under a Creative Commons Attribution 4.0 International license (Suresh Babu et al., 2025d). Codes for the DM training and generation, and trained weights

are preserved at <https://doi.org/10.5281/zenodo.15750243>, available via the MIT license and developed openly at <https://github.com/ananthu545/quasi-geostrophic-beta-plane-super-resolution> (Suresh Babu et al., 2025b).

Acknowledgments

We thank C. Daskalakis and G. Daras for insightful discussions on diffusion modeling. We thank the members of our ML-SCOPE MURI team and of our MSEAS group, including V.A. Rodriguez, F.M. Benfenati, and A.K. Saravanakumar, for discussions on quasi-geostrophic turbulence, statistics, and figures. We also thank the editor Dr. Stephen Griffies, associate editor Dr. Ian Grooms, and two anonymous reviewers for their constructive feedback to improve this manuscript. We are grateful to the Office of Naval Research for partial support under Grant N00014-20-1-2023 (MURI ML-SCOPE) and N00014-19-1-2693 (IN-BDA) to the Massachusetts Institute of Technology. ANSB was partially supported by an MIT Mechanical Engineering MathWorks Fellowship.

References

- Adrian, M., Sanz-Alonso, D., & Willett, R. (2025). Data assimilation with machine learning surrogate models: A case study with fourcastnet. *Artificial Intelligence for the Earth Systems*, 4(3), e240050. <https://doi.org/10.1175/aies-d-24-0050.1>
- Albergo, M. S., & Vanden-Eijnden, E. (2022). *Building normalizing flows with stochastic interpolants*. arXiv preprint arXiv:2209.15571.
- Anderson, B. D. (1982). Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3), 313–326. [https://doi.org/10.1016/0304-4149\(82\)90051-5](https://doi.org/10.1016/0304-4149(82)90051-5)
- Arbic, B. K., & Flierl, G. R. (2003). Coherent vortices and kinetic energy ribbons in asymptotic, quasi two-dimensional f-plane turbulence. *Physics of fluids*, 15(8), 2177–2189. <https://doi.org/10.1063/1.1582183>
- Archer, M., Wang, J., Klein, P., Dibarbour, G., & Fu, L.-L. (2025). Wide-swath satellite altimetry unveils global submesoscale ocean dynamics. *Nature*, 640(8059), 691–696. <https://doi.org/10.1038/s41586-025-08722-8>
- Asefi, N., Lupin-Jimenez, L., Wu, T., He, R., & Chattopadhyay, A. (2025). *Generative lagrangian data assimilation for ocean dynamics under extreme sparsity*. Machine Learning: Earth.
- Baldassari, L., Siahkoohi, A., Garnier, J., Solna, K., & de Hoop, M. V. (2023). Conditional score-based diffusion models for bayesian inference in infinite dimensions. *Advances in Neural Information Processing Systems*, 36, 24262–24290.
- Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., et al. (2019). On the resolutions of ocean altimetry maps. *Ocean Science*, 15(4), 1091–1109. <https://doi.org/10.5194/os-15-1091-2019>
- Bannister, R. N. (2017). A review of operational methods of variational and ensemble-variational data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 143(703), 607–633. <https://doi.org/10.1002/qj.2982>
- Baño-Medina, J., Manzanar, R., & Gutiérrez, J. M. (2020). Configuration and intercomparison of deep learning neural models for statistical downscaling. *Geoscientific Model Development*, 13(4), 2109–2124. <https://doi.org/10.5194/gmd-13-2109-2020>
- Bao, F., Chipilski, H. G., Liang, S., Zhang, G., & Whitaker, J. S. (2025). Nonlinear ensemble filtering with diffusion models: Application to the surface quasigeostrophic dynamics. *Monthly Weather Review*, 153(7), 1155–1169. <https://doi.org/10.1175/mwr-d-24-0069.1>
- Bardina, J., Ferziger, J., & Reynolds, W. (1980). Improved subgrid-scale models for large-eddy simulation. In *13th fluid and plasmadynamics conference* (p. 1357).
- Barthélémy, S., Brajard, J., Bertino, L., & Counillon, F. (2022). Super-resolution data assimilation. *Ocean Dynamics*, 72(8), 661–678. <https://doi.org/10.1007/s10236-022-01523-x>
- Barwey, S., Pal, P., Patel, S., Balin, R., Lusch, B., Vishwanath, V., et al. (2025). Mesh-based super-resolution of fluid flows with multiscale graph neural networks. *Computer Methods in Applied Mechanics and Engineering*, 443, 118072. <https://doi.org/10.1016/j.cma.2025.118072>
- Bengtsson, T., Bickel, P., & Li, B. (2008). Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. In *Probability and statistics: Essays in honor of david a. freedman* (Vol. 2, pp. 316–335). Institute of Mathematical Statistics. <https://doi.org/10.1214/193940307000000518>
- Bennett, A. F. (1992). *Inverse methods in physical oceanography*. Cambridge University Press.
- Bennett, A. F. (2005). *Inverse modeling of the ocean and atmosphere*. Cambridge University Press.
- Blau, Y., & Michaeli, T. (2018). The perception-distortion tradeoff. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6228–6237).
- Bluestein, H. B., Carr, F. H., & Goodman, S. J. (2022). Atmospheric observations of weather and climate. *Atmosphere-Ocean*, 60(3–4), 149–187. <https://doi.org/10.1080/07055900.2022.2082369>
- Bocquet, M. (2023). Surrogate modeling for the climate sciences dynamics with machine learning and data assimilation. *Frontiers in Applied Mathematics and Statistics*, 9, 1133226. <https://doi.org/10.3389/fams.2023.1133226>
- Bode, M., Gauding, M., Lian, Z., Denker, D., Davidovic, M., Kleinheinz, K., et al. (2021). Using physics-informed enhanced super-resolution generative adversarial networks for subfilter modeling in turbulent reactive flows. *Proceedings of the Combustion Institute*, 38(2), 2617–2625. <https://doi.org/10.1016/j.proci.2020.06.022>
- Boyd, J. P. (2001). Chebyshev and fourier spectral methods. *Courier Corporation*.
- Boys, B., Girolami, M., Pidstrigach, J., Reich, S., Mosca, A., & Akyildiz, O. D. (2023). *Tweedie moment projected diffusions for inverse problems*. arXiv preprint arXiv:2310.06721.
- Bracco, A., Liu, G., & Sun, D. (2019). Mesoscale-submesoscale interactions in the gulf of Mexico: From oil dispersion to climate. *Chaos, Solitons and Fractals*, 119, 63–72. <https://doi.org/10.1016/j.chaos.2018.12.012>
- Bredell, G., Flouris, K., Chaitanya, K., Erdil, E., & Konukoglu, E. (2023). *Explicitly minimizing the blur error of variational autoencoders*. arXiv preprint arXiv:2304.05939.
- Brenowitz, N. D., Ge, T., Subramaniam, A., Gupta, A., Hall, D. M., Mardani, M., et al. (2025). *Climate in a bottle: Towards a generative foundation model for the kilometer-scale global atmosphere*. arXiv preprint arXiv:2505.06474.
- Bretherton, F. P., Davis, R. E., & Fandry, C. (1976). A technique for objective analysis and design of oceanographic experiments applied to mode-73. In *Deep-Sea Research and Oceanographic Abstracts* (Vol. 23(7), pp. 559–582). [https://doi.org/10.1016/0011-7471\(76\)90001-2](https://doi.org/10.1016/0011-7471(76)90001-2)
- Brotzge, J. A., Berchhoff, D., Carlis, D. L., Carr, F. H., Carr, R. H., Gerth, J. J., et al. (2023). Challenges and opportunities in numerical weather prediction. *Bulletin of the American Meteorological Society*, 104(3), E698–E705. <https://doi.org/10.1175/bams-d-22-0172.1>
- Buzzicotti, M. (2023). Data reconstruction for complex flows using ai: Recent progress, obstacles, and perspectives. *Europhysics Letters*, 142(2), 23001. <https://doi.org/10.1209/0295-5075/acc88c>
- Callahan, J. L., Maeda, K., & Brunton, S. L. (2019). Robust flow reconstruction from limited measurements via sparse representation. *Physical Review Fluids*, 4(10), 103907. <https://doi.org/10.1103/physrevfluids.4.103907>
- Carrasi, A., Bocquet, M., Bertino, L., & Evensen, G. (2018). Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *Wiley Interdisciplinary Reviews: Climate Change*, 9(5), e535. <https://doi.org/10.1002/wcc.535>
- Carrasi, A., Bocquet, M., Demeyer, J., Grudzien, C., Raanes, P., & Vannitsem, S. (2022). Data assimilation for chaotic dynamics. *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications, IV*, 1–42.
- Chakraborty, D., Guan, H., Stock, J., Arcomano, T., Cervone, G., & Maulik, R. (2025). *Multimodal atmospheric super-resolution with deep generative models*. arXiv preprint arXiv:2506.22780.
- Chassignet, E., & Marshall, D. (2008). *Gulf stream separation in numerical ocean models* (Vol. 177). Geophysical Monograph Series.

- Chassignet, E. P., & Xu, X. (2021). On the importance of high-resolution in large-scale ocean models. *Advances in Atmospheric Sciences*, 38(10), 1621–1634. <https://doi.org/10.1007/s00376-021-0385-7>
- Chattopadhyay, A., & Hassanzadeh, P. (2023). Long-term instabilities of deep learning-based digital twins of the climate system: The cause and a solution. *arXiv preprint arXiv:2304.07029*.
- Chelton, D. B., Schlax, M. G., & Samelson, R. M. (2011). Global observations of nonlinear mesoscale eddies. *Progress in Oceanography*, 91(2), 167–216. <https://doi.org/10.1016/j.pocean.2011.01.002>
- Chen, N., Wiggins, S., & Andreou, M. (2025). Taming uncertainty in a complex world: The rise of uncertainty quantification—a tutorial for beginners. *Notices of the American Mathematical Society*, 72(3), 1. <https://doi.org/10.1090/noti3120>
- Chen, Y., Goldstein, M., Hua, M., Albergo, M. S., Boffi, N. M., & Vanden-Eijnden, E. (2024). *Probabilistic forecasting with stochastic interpolants and follmer processes*. arXiv preprint arXiv:2403.13724.
- Chung, H., Kim, J., Kim, S., & Ye, J. C. (2023a). Parallel diffusion models of operator and image for blind inverse problems. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 6059–6069).
- Chung, H., Kim, J., McCann, M. T., Klasky, M. L., & Ye, J. C. (2023). Diffusion posterior sampling for general noisy inverse problems. In *11th international conference on learning representations, iclr 2023*.
- Clark, R. A., Ferziger, J. H., & Reynolds, W. C. (1979). Evaluation of subgrid-scale models using an accurately simulated turbulent flow. *Journal of Fluid Mechanics*, 91(1), 1–16. <https://doi.org/10.1017/s002211207900001x>
- Courtier, P., Andersson, E., Heckley, W., Vasiljevic, D., Hamrud, M., Hollingsworth, A., et al. (1998). The ECMWF implementation of three-dimensional variational assimilation (3d-var). I: Formulation. *Quarterly Journal of the Royal Meteorological Society*, 124(550), 1783–1807. <https://doi.org/10.1002/qj.49712455002>
- Daras, G., Chung, H., Lai, C.-H., Mitsufuji, Y., Ye, J. C., Milanfar, P., et al. (2024). *A survey on diffusion models for inverse problems*. arXiv preprint arXiv:2410.00083.
- David, T. W., Marshall, D. P., & Zanna, L. (2017). The statistical nature of turbulent barotropic ocean jets. *Ocean Modelling*, 113, 34–49. <https://doi.org/10.1016/j.ocemod.2017.03.008>
- Denis, B., Laprise, R., Caya, D., & Côté, J. (2002). Downscaling ability of one-way nested regional climate models: The big-brother experiment. *Climate Dynamics*, 18(8), 627–646.
- Dhariwal, P., & Nichol, A. (2021). Diffusion models beat GANs on image synthesis. *Advances in Neural Information Processing Systems*, 34, 8780–8794.
- Diffenbaugh, N. S., Pal, J. S., Trapp, R. J., & Giorgi, F. (2005). Fine-scale processes regulate the response of extreme events to global climate change. *Proceedings of the National Academy of Sciences*, 102(44), 15774–15778. <https://doi.org/10.1073/pnas.0506042102>
- Dou, Z., & Song, Y. (2024). Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *The twelfth international conference on learning representations*.
- Ehrendorfer, M. (2007). A review of issues in ensemble-based Kalman filtering. *Meteorologische Zeitschrift*, 16(6), 795–818. <https://doi.org/10.1127/0941-2948/2007/0256>
- Eivazi, H., Le Clainche, S., Hoyas, S., & Vinuesa, R. (2022). Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows. *Expert Systems with Applications*, 202, 117038. <https://doi.org/10.1016/j.eswa.2022.117038>
- Errico, R. M., Yang, R., Privé, N. C., Tai, K.-S., Todling, R., Sienkiewicz, M. E., & Guo, J. (2013). Development and validation of observing-system simulation experiments at NASA's global modeling and assimilation office. *Quarterly Journal of the Royal Meteorological Society*, 139(674), 1162–1178. <https://doi.org/10.1002/qj.2027>
- Evensen, G. (2003). The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4), 343–367. <https://doi.org/10.1007/s10236-003-0036-9>
- Evensen, G., Vossepoel, F. C., & Van Leeuwen, P. J. (2022). *Data assimilation fundamentals: A unified formulation of the state and parameter estimation problem*. Springer Nature.
- Fan, X., Akhane, D., & Wang, J.-X. (2025). Neural differentiable modeling with diffusion-based super-resolution for two-dimensional spatio-temporal turbulence. *Computer Methods in Applied Mechanics and Engineering*, 433, 117478. <https://doi.org/10.1016/j.cma.2024.117478>
- Farge, M., Schneider, K., & Kevlahan, N. (1999). Non-gaussianity and coherent vortex simulation for two-dimensional turbulence using an adaptive orthogonal wavelet basis. *Physics of Fluids*, 11(8), 2187–2201. <https://doi.org/10.1063/1.870080>
- Fernandez-Martinez, J. L., & Fernandez-Muniz, Z. (2020). The curse of dimensionality in inverse problems. *Journal of Computational and Applied Mathematics*, 369, 112571. <https://doi.org/10.1016/j.cam.2019.112571>
- Frederiksen, J. S., & Davies, A. G. (1997). Eddy viscosity and stochastic backscatter parameterizations on the sphere for atmospheric circulation models. *Journal of the Atmospheric Sciences*, 54(20), 2475–2492. [https://doi.org/10.1175/1520-0469\(1997\)054<2475:evabsp>2.0.co;2](https://doi.org/10.1175/1520-0469(1997)054<2475:evabsp>2.0.co;2)
- Frezat, H., Le Sommer, J., Fablet, R., Balarac, G., & Lguensat, R. (2022). A posteriori learning for quasi-geostrophic turbulence parametrization. *Journal of Advances in Modeling Earth Systems*, 14(11), e2022MS003124. <https://doi.org/10.1029/2022ms003124>
- Fukami, K., Fukagata, K., & Taira, K. (2019). Super-resolution reconstruction of turbulent flows with machine learning. *Journal of Fluid Mechanics*, 870, 106–120. <https://doi.org/10.1017/jfm.2019.238>
- Fukami, K., Fukagata, K., & Taira, K. (2021). Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *Journal of Fluid Mechanics*, 909, A9. <https://doi.org/10.1017/jfm.2020.948>
- Fukami, K., Fukagata, K., & Taira, K. (2023). Super-resolution analysis via machine learning: A survey for fluid flows. *Theoretical and Computational Fluid Dynamics*, 37(4), 421–444. <https://doi.org/10.1007/s00162-023-00663-0>
- Fukami, K., & Taira, K. (2024). Single-snapshot machine learning for super-resolution of turbulence. *Journal of Fluid Mechanics*, 1001, A32. <https://doi.org/10.1017/jfm.2024.1136>
- Galperin, B., Nakano, H., Huang, H.-P., & Sukoriansky, S. (2004). The ubiquitous zonal jets in the atmospheres of giant planets and earth's oceans. *Geophysical Research Letters*, 31(13), L13303. <https://doi.org/10.1029/2004gl019691>
- Gandin, L. S. (1963). *Objective analysis of meteorological fields* (Vol. 242). Israel program for scientific translations.
- Gao, H., Sun, L., & Wang, J.-X. (2021). Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels. *Physics of Fluids*, 33(7), 073603. <https://doi.org/10.1063/5.0054312>
- Graham, J. P., & Ringler, T. (2013). A framework for the evaluation of turbulence closures used in mesoscale ocean large-eddy simulations. *Ocean Modelling*, 65, 25–39. <https://doi.org/10.1016/j.ocemod.2013.01.004>
- Guan, Y., Chattopadhyay, A., Subel, A., & Hassanzadeh, P. (2022). Stable a posteriori LES of 2d turbulence using convolutional neural networks: Backscattering analysis and generalization to higher Re via transfer learning. *Journal of Computational Physics*, 458, 111090. <https://doi.org/10.1016/j.jcp.2022.111090>
- Guan, Y., Hassanzadeh, P., Schneider, T., Dunbar, O., Huang, D. Z., Wu, J., & Lopez-Gomez, I. (2024). *Online learning of eddy-viscosity and backscattering closures for geophysical turbulence using ensemble Kalman inversion*. arXiv preprint arXiv:2409.04985.

- Guan, Y., Subel, A., Chattopadhyay, A., & Hassanzadeh, P. (2023). Learning physics-constrained subgrid-scale closures in the small-data regime for stable and accurate les. *Physica D: Nonlinear Phenomena*, *443*, 133568. <https://doi.org/10.1016/j.physd.2022.133568>
- Güemes, A., Sanmiguel Vila, C., & Discetti, S. (2022). Super-resolution generative adversarial networks of randomly-seeded fields. *Nature Machine Intelligence*, *4*(12), 1165–1173. <https://doi.org/10.1038/s42256-022-00572-7>
- Gui, J., Sun, Z., Wen, Y., Tao, D., & Ye, J. (2021). A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE Transactions on Knowledge and Data Engineering*, *35*(4), 3313–3332. <https://doi.org/10.1109/tkde.2021.3130191>
- Gula, J., Blacic, T. M., & Todd, R. E. (2019). Submesoscale coherent vortices in the gulf stream. *Geophysical Research Letters*, *46*(5), 2704–2714. <https://doi.org/10.1029/2019gl081919>
- Gupta, A., & Lermusiaux, P. F. J. (2021). Neural closure models for dynamical systems. *Proceedings of The Royal Society A*, *477*(2252), 1–29. <https://doi.org/10.1098/rspa.2020.1004>
- Gupta, A., & Lermusiaux, P. F. J. (2023). Generalized neural closure models with interpretability. *Scientific Reports*, *13*(1), 10634. <https://doi.org/10.1038/s41598-023-35319-w>
- Haley, P. J., Jr., Agarwal, A., & Lermusiaux, P. F. J. (2015). Optimizing velocities and transports for complex coastal regions and archipelagos. *Ocean Modelling*, *89*, 1–28. <https://doi.org/10.1016/j.ocemod.2015.02.005>
- Haley, P. J., Jr., & Lermusiaux, P. F. J. (2010). Multiscale two-way embedding schemes for free-surface primitive equations in the “Multidisciplinary Simulation, Estimation and Assimilation System”. *Ocean Dynamics*, *60*(6), 1497–1537. <https://doi.org/10.1007/s10236-010-0349-4>
- Haley, P. J., Jr., Mirabito, C., Doshi, M., & Lermusiaux, P. F. J. (2023). Ensemble forecasting for the Gulf of Mexico Loop Current region. In *Oceans 2023 IEEE/MTS Gulf Coast Biloxi*. <https://doi.org/10.23919/OCEANS52994.2023.10337035>.MS
- Han, Q., Jiang, X., Zhao, Y., & Wang, X. (2025). Diffusion-model-based downscaling of observed sea surface height over the kuroshio extension since 2000. *Atmosphere*, *16*(5), 570. <https://doi.org/10.3390/atmos16050570>
- Harris, L. M., & Lin, S.-J. (2014). Global-to-regional nested grid climate simulations in the gfdl high resolution atmospheric model. *Journal of Climate*, *27*(13), 4890–4910. <https://doi.org/10.1175/jcli-d-13-00596.1>
- Ho, J., & Salimans, T. (2021). Classifier-free diffusion guidance. In *Neurips 2021 workshop on deep generative models and downstream applications*.
- Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., & Fleet, D. J. (2022). Video diffusion models. *Advances in Neural Information Processing Systems*, *35*, 8633–8646.
- Hodyss, D., & Morzfeld, M. (2026). Using diffusion models to do data assimilation. *Monthly Weather Review*, *154*(2), 165–182. <https://doi.org/10.1175/mwr-d-25-0125.1>
- Hoteit, I., Luo, X., Bocquet, M., Kohl, A., & Ait-El-Fquih, B. (2018). Data assimilation in oceanography: Current status and new directions. *New frontiers in operational oceanography*, 465–512.
- Huang, H.-P., Galperin, B., & Sukoriansky, S. (2001). Anisotropic spectra in two-dimensional turbulence on the surface of a rotating sphere. *Physics of Fluids*, *13*(1), 225–240. <https://doi.org/10.1063/1.1327594>
- Jacobs, G., D’Addezio, J., Bartels, B., DeHaan, C., Barron, C., Carrier, M., et al. (2023). Adapting constrained scales to observation resolution in ocean forecasts. *Ocean Modelling*, *186*, 102252. <https://doi.org/10.1016/j.ocemod.2023.102252>
- Jakhar, K., Guan, Y., Mojjani, R., Chattopadhyay, A., & Hassanzadeh, P. (2024). Learning closed-form equations for subgrid-scale closures from high-fidelity data: Promises and challenges. *Journal of Advances in Modeling Earth Systems*, *16*(7), e2023MS003874. <https://doi.org/10.1029/2023ms003874>
- Jansen, M. F., & Held, I. M. (2014). Parameterizing subgrid-scale eddy effects using energetically consistent backscatter. *Ocean Modelling*, *80*, 36–48. <https://doi.org/10.1016/j.ocemod.2014.06.002>
- Johnston, T. M. S., MacKinnon, J. A., Colin, P. L., Haley, P. J., Jr., Lermusiaux, P. F. J., Lucas, A. J., et al. (2019). Energy and momentum lost to wake eddies and lee waves generated by the North Equatorial Current and tidal flows at Peleliu, Palau. *Oceanography*, *32*(4), 110–125. <https://doi.org/10.5670/oceanog.2019.417>
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems.
- Kalnay, E. (2003). *Atmospheric modeling, data assimilation and predictability*. Cambridge University Press.
- Karras, T., Aittala, M., Kynkäänniemi, T., Lehtinen, J., Aila, T., & Laine, S. (2024). Guiding a diffusion model with a bad version of itself. *Advances in Neural Information Processing Systems*, *37*, 52996–53021.
- Keller, A. A., Garner, K. L., Rao, N., Knipping, E., & Thomas, J. (2022). Downscaling approaches of climate change projections for watershed modeling: Review of theoretical and practical considerations. *PLoS Water*, *1*(9), e0000046. <https://doi.org/10.1371/journal.pwat.0000046>
- Kevlahan, N.-R., & Farge, M. (1997). Vorticity filaments in two-dimensional turbulence: Creation, stability and effect. *Journal of Fluid Mechanics*, *346*, 49–76. <https://doi.org/10.1017/s0022112097006113>
- Khan, M. S., Coulibaly, P., & Dibike, Y. (2006). Uncertainty analysis of statistical downscaling methods. *Journal of Hydrology*, *319*(1–4), 357–382. <https://doi.org/10.1016/j.jhydrol.2005.06.035>
- Kim, H., Kim, J., Won, S., & Lee, C. (2021). Unsupervised deep learning for super-resolution reconstruction of turbulence. *Journal of Fluid Mechanics*, *910*, A29. <https://doi.org/10.1017/jfm.2020.1028>
- Kingma, D. P., & Welling, M. (2019). An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, *12*(4), 307–392.
- Klemas, V., & Yan, X.-H. (2014). Subsurface and deeper ocean remote sensing from satellites: An overview and new results. *Progress in Oceanography*, *122*, 1–9. <https://doi.org/10.1016/j.pocean.2013.11.010>
- Kossale, Y., Airaj, M., & Darouichi, A. (2022). Mode collapse in generative adversarial networks: An overview. In *2022 8th international conference on optimization and applications (icoa)* (pp. 1–6).
- Kraichnan, R. H. (1971). Inertial-range transfer in two-and three-dimensional turbulence. *Journal of Fluid Mechanics*, *47*(3), 525–535. <https://doi.org/10.1017/s0022112071001216>
- Kulkarni, C. S., Haley, P. J., Jr., Lermusiaux, P. F. J., Dutt, A., Gupta, A., Mirabito, C., et al. (2018). Real-time sediment plume modeling in the Southern California Bight. In *Oceans conference 2018*. Charleston, SC. <https://doi.org/10.1109/OCEANS.2018.8653642>
- Lai, C.-Y., Hassanzadeh, P., Sheshadri, A., Sonnewald, M., Ferrari, R., & Balaji, V. (2024). Machine learning for climate physics and simulations. *Annual Review of Condensed Matter Physics*, *16*(1), 343–365. <https://doi.org/10.1146/annurev-conmatphys-043024-114758>
- Langford, J. A., & Moser, R. D. (1999). Optimal les formulations for isotropic turbulence. *Journal of Fluid Mechanics*, *398*, 321–346. <https://doi.org/10.1017/s0022112099006369>
- Law, K., Stuart, A., & Zygalakis, K. (2015). *Data assimilation* (Vol. 214, p. 52). Springer.
- Leinonen, J., Hamann, U., Nerini, D., Germann, U., & Franch, G. (2023). *Latent diffusion models for generative precipitation nowcasting with accurate uncertainty quantification*. arXiv preprint arXiv:2304.12891.

- Leith, C. E. (1968). Diffusion approximation for two-dimensional turbulence. *Physics of Fluids*, *11*(3), 671–672. <https://doi.org/10.1063/1.1691968>
- Leonard, A. (1975). Energy cascade in large-eddy simulations of turbulent fluid flows. In *Advances in geophysics* (Vol. 18, pp. 237–248). Elsevier. [https://doi.org/10.1016/s0065-2687\(08\)60464-1](https://doi.org/10.1016/s0065-2687(08)60464-1)
- Lermusiaux, P. F. J. (2002). On the mapping of multivariate geophysical fields: Sensitivities to size, scales, and dynamics. *Journal of Atmospheric and Oceanic Technology*, *19*(10), 1602–1637. [https://doi.org/10.1175/1520-0426\(2002\)019<1602:OTMOMG>2.0.CO;2](https://doi.org/10.1175/1520-0426(2002)019<1602:OTMOMG>2.0.CO;2)
- Lermusiaux, P. F. J. (2006). Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *Journal of Computational Physics*, *217*(1), 176–199. <https://doi.org/10.1016/j.jcp.2006.02.010>
- Lermusiaux, P. F. J. (2007). Adaptive modeling, adaptive data assimilation and adaptive sampling. *Physica D: Nonlinear Phenomena*, *230*(1), 172–196. <https://doi.org/10.1016/j.physd.2007.02.014>
- Lermusiaux, P. F. J. (2015). *Numerical fluid mechanics*. MIT OpenCourseWare. Retrieved from <https://ocw.mit.edu/courses/mechanical-engineering/2-29-numerical-fluid-mechanics-spring-2015/lecture-notes-and-references/>
- Lermusiaux, P. F. J., Anderson, D. G. M., & Lozano, C. J. (2000). On the mapping of multivariate geophysical fields: Error and variability subspace estimates. *Quarterly Journal of the Royal Meteorological Society*, *126*(565), 1387–1429. <https://doi.org/10.1256/smsqj.56509>
- Lermusiaux, P. F. J., Chiu, C.-S., Gawarkiewicz, G. G., Abbot, P., Robinson, A. R., Miller, R. N., et al. (2006). Quantifying uncertainties in ocean predictions. *Oceanography*, *19*(1), 92–105. <https://doi.org/10.5670/oceanog.2006.93>
- Lermusiaux, P. F. J., Haley, P. J., Jr., Mirabito, C., Mule, E. M., DiMarco, S. F., Dancer, A., et al. (2024). Real-time ocean probabilistic forecasts, reachability analysis, and adaptive sampling in the Gulf of Mexico. In *Oceans 2024 IEEE/MTS Halifax* (pp. 1–10). Halifax. <https://doi.org/10.1109/OCEANS55160.2024.10754153>
- Lermusiaux, P. F. J., Haley, P. J., Leslie, W. G., Agarwal, A., Logutov, O., & Burton, L. J. (2011). Multiscale physical and biological dynamics in the Philippine Archipelago: Predictions and processes. *Oceanography*, *24*(1), 70–89. (Special Issue on the Philippine Straits Dynamics Experiment). <https://doi.org/10.5670/oceanog.2011.05>
- Lermusiaux, P. F. J., Malanotte-Rizzoli, P., Stammer, D., Carton, J., Cummings, J., & Moore, A. M. (2006). Progress and prospects of U.S. data assimilation in ocean research. *Oceanography*, *19*(1), 172–183. <https://doi.org/10.5670/oceanog.2006.102>
- Lermusiaux, P. F. J., Mirabito, C., Haley, P. J., Jr., Ali, W. H., Gupta, A., Jana, S., et al. (2020). Real-time probabilistic coupled ocean physics-acoustics forecasting and data assimilation for underwater GPS. In *Oceans 2020 IEEE/MTS* (pp. 1–9). <https://doi.org/10.1109/IEEECONF38699.2020.9389003>
- Lermusiaux, P. F. J., & Robinson, A. R. (1999). Data assimilation via Error Subspace Statistical Estimation, part I: Theory and schemes. *Monthly Weather Review*, *127*(7), 1385–1407. [https://doi.org/10.1175/1520-0493\(1999\)127<1385:DAVESS>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<1385:DAVESS>2.0.CO;2)
- Lermusiaux, P. F. J., Robinson, A. R., Haley, P. J., & Leslie, W. G. (2002). Advanced interdisciplinary data assimilation: Filtering and smoothing via error subspace statistical estimation. In *Proceedings of the oceans 2002 mts/ieee conference* (pp. 795–802). <https://doi.org/10.1109/oceans.2002.1192071>
- Lermusiaux, P. F. J., Subramani, D. N., Lin, J., Kulkarni, C. S., Gupta, A., Dutt, A., et al. (2017). A future for intelligent autonomous ocean observing systems. *Journal of Marine Research*, *75*(6), 765–813. (The Sea. Volume 17, The Science of Ocean Prediction, Part 2.). <https://doi.org/10.1357/002224017823524035>
- Li, T., Lanotte, A. S., Buzzicotti, M., Bonaccorso, F., & Biferale, L. (2023). Multi-scale reconstruction of turbulent rotating flows with generative diffusion models. *Atmosphere*, *15*(1), 60. <https://doi.org/10.3390/atmos15010060>
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2020). *Fourier neural operator for parametric partial differential equations*. arXiv preprint arXiv:2010.08895.
- Lin, M., & Yang, C. (2020). Ocean observation technologies: A review. *Chinese Journal of Mechanical Engineering*, *33*(1), 32. <https://doi.org/10.1186/s10033-020-00449-z>
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., & Le, M. (2022). *Flow matching for generative modeling*. arXiv preprint arXiv:2210.02747.
- Liu, B., Tang, J., Huang, H., & Lu, X.-Y. (2020). Deep learning methods for super-resolution reconstruction of turbulent flows. *Physics of fluids*, *32*(2), 025105. <https://doi.org/10.1063/1.5140772>
- Liu, Y., Yang, M., Zhang, Z., Bao, F., Cao, Y., & Zhang, G. (2024). Diffusion-model-assisted supervised learning of generative models for density estimation. *Journal of Machine Learning for Modeling and Computing*, *5*(1), 25–38. <https://doi.org/10.1615/jmachlearnmodelcomput.2024051346>
- Lolla, T., & Lermusiaux, P. F. J. (2017). A Gaussian mixture model smoother for continuous nonlinear stochastic dynamical systems: Applications. *Monthly Weather Review*, *145*(7), 2763–2790. <https://doi.org/10.1175/MWR-D-16-0065.1>
- Lopez-Gomez, I., Wan, Z. Y., Zepeda-Núñez, L., Schneider, T., Anderson, J., & Sha, F. (2025). Dynamical-generative downscaling of climate model ensembles. *Proceedings of the National Academy of Sciences*, *122*(17), e2420288122. <https://doi.org/10.1073/pnas.2420288122>
- Lorenç, A. C. (1986). Analysis methods for numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, *112*(474), 1177–1194. <https://doi.org/10.1002/qj.49711247414>
- Loshchilov, I., & Hutter, F. (2016). *Sgdr: Stochastic gradient descent with warm restarts*. arXiv preprint arXiv:1608.03983.
- Loshchilov, I., & Hutter, F. (2017). *Decoupled weight decay regularization*. arXiv preprint arXiv:1711.05101.
- Lu, P., & Lermusiaux, P. F. J. (2021). Bayesian learning of stochastic dynamical models. *Physica D: Nonlinear Phenomena*, *427*, 133003. <https://doi.org/10.1016/j.physd.2021.133003>
- Maejima, S., & Kawai, S. (2025). Unsupervised and supervised machine learning pipeline for super-resolution-based subgrid scale modelling in coarse-grid large-eddy simulations. *Journal of Fluid Mechanics*, *1013*, A28. <https://doi.org/10.1017/jfm.2025.10211>
- Maltrud, M., & Vallis, G. (1991). Energy spectra and coherent structures in forced two-dimensional and beta-plane turbulence. *Journal of Fluid Mechanics*, *228*, 321–342. <https://doi.org/10.1017/s0022112091002720>
- Manshausen, P., Cohen, Y., Harrington, P., Pathak, J., Pritchard, M., Garg, P., et al. (2025). Generative data assimilation of sparse weather station observations at kilometer scales. *Journal of Advances in Modeling Earth Systems*, *17*, e2024MS004505. <https://doi.org/10.1029/2024MS004505>
- Mardani, M., Brenowitz, N., Cohen, Y., Pathak, J., Chen, C.-Y., Liu, C.-C., et al. (2025). Residual corrective diffusion modeling for km-scale atmospheric downscaling. *Communications Earth and Environment*, *6*(1), 124. <https://doi.org/10.1038/s43247-025-02042-5>
- Martin, S. A., Manucharyan, G. E., & Klein, P. (2025). Generative data assimilation for surface ocean state estimation from multi-modal satellite observations. *Journal of Advances in Modeling Earth Systems*, *17*(8), e2025MS005063. <https://doi.org/10.1029/2025ms005063>
- Maulik, R., San, O., Rasheed, A., & Vedula, P. (2019). Subgrid modelling for two-dimensional turbulence using neural networks. *Journal of Fluid Mechanics*, *858*, 122–144. <https://doi.org/10.1017/jfm.2018.770>
- McWilliams, J. (2007). Extreme events in geophysical turbulence and waves: What populates the tails of distribution functions. *Extreme Events, University of Hawaii at Manoa*, 23–26.

- McWilliams, J. C. (2016). Submesoscale currents in the ocean. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 472(2189), 20160117. <https://doi.org/10.1098/rspa.2016.0117>
- Menemenlis, D., Fieguth, P., Wunsch, C., & Willsky, A. (1997). Adaptation of a fast optimal interpolation algorithm to the mapping of oceanographic data. *Journal of Geophysical Research*, 102(C5), 10573–10584. <https://doi.org/10.1029/97jc00697>
- Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.-Y., & Ermon, S. (2021). *Sdedit: Guided image synthesis and editing with stochastic differential equations*. arXiv preprint arXiv:2108.01073.
- Mensa, J. A., Garraffo, Z., Griffa, A., Özgökmen, T. M., Haza, A., & Veneziani, M. (2013). Seasonality of the submesoscale dynamics in the gulf stream region. *Ocean Dynamics*, 63(8), 923–941. <https://doi.org/10.1007/s10236-013-0633-1>
- Milánfar, P., & Delbracio, M. (2025). Denoising: A powerful building block for imaging, inverse problems and machine learning. *Philosophical Transactions A*, 383(2299), 20240326. <https://doi.org/10.1098/rsta.2024.0326>
- Moore, A. M., Martin, M., Akella, S., Arango, H., Balmaseda, M., Bertino, L., et al. (2019). Synthesis of ocean observations using data assimilation for operational, real-time and reanalysis systems: A more complete picture of the state of the ocean. *Frontiers in Marine Science*, 6(90), 1–6. <https://doi.org/10.3389/fmars.2019.00090>
- Morrow, R., Fu, L.-L., Arduin, F., Benkiran, M., Chapron, B., Cosme, E., et al. (2019). Global observations of fine-scale ocean surface topography with the surface water and ocean topography (swot) mission. *Frontiers in Marine Science*, 6, 232. <https://doi.org/10.3389/fmars.2019.00232>
- Moser, B. B., Raue, F., Frolov, S., Palacio, S., Hees, J., & Dengel, A. (2023). Hitchhiker’s guide to super-resolution: Introduction and recent advances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8), 9862–9882. <https://doi.org/10.1109/tpami.2023.3243794>
- Nichol, A. Q., & Dhariwal, P. (2021). Improved denoising diffusion probabilistic models. In *International conference on machine learning* (pp. 8162–8171).
- Nickerson, A. K., Weisberg, R. H., & Liu, Y. (2022). On the evolution of the gulf of Mexico loop current through its penetrative, ring shedding and retracted states. *Advances in Space Research*, 69(11), 4058–4077. <https://doi.org/10.1016/j.asr.2022.03.039>
- Nozawa, T., & Yoden, S. (1997). Spectral anisotropy in forced two-dimensional turbulence on a rotating sphere. *Physics of Fluids*, 9(12), 3834–3842. <https://doi.org/10.1063/1.869518>
- Oey, L., Ezer, T., & Lee, H. (2005). Loop current, rings and related circulation in the gulf of Mexico: A review of numerical models and future challenges. *Geophysical Monograph-American Geophysical Union*, 161, 31.
- Oommen, V., Bora, A., Zhang, Z., & Karniadakis, G. E. (2025). Integrating neural operators with diffusion models improves spectral representation in turbulence modelling. *Proceedings of the Royal Society A*, 481(2309), 20240819. <https://doi.org/10.1098/rspa.2024.0819>
- Orszag, S. A., & Israeli, M. (1974). Numerical simulation of viscous incompressible flows. *Annual Review of Fluid Mechanics*, 6(1), 281–318. <https://doi.org/10.1146/annurev.fl.06.010174.001433>
- Page, J. (2025). Super-resolution of turbulence with dynamics in the loss. *Journal of Fluid Mechanics*, 1002, R3. <https://doi.org/10.1017/jfm.2024.1202>
- Park, C. Y., McCann, M. T., Garcia-Cardona, C., Wohlberg, B., & Kamilov, U. S. (2024). *Random walks with tweedie: A unified framework for diffusion models*. arXiv preprint arXiv:2411.18702.
- Patel, D. V., Ray, D., & Oberai, A. A. (2022). Solution of physics-based bayesian inverse problems with deep generative priors. *Computer Methods in Applied Mechanics and Engineering*, 400, 115428. <https://doi.org/10.1016/j.cma.2022.115428>
- Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., et al. (2022). *Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators*. arXiv preprint arXiv:2202.11214.
- Perezhogin, P., Zanna, L., & Fernandez-Granda, C. (2023). Generative data-driven approaches for stochastic subgrid parameterizations in an idealized ocean model. *Journal of Advances in Modeling Earth Systems*, 15(10), e2023MS003681. <https://doi.org/10.1029/2023ms003681>
- Pope, S. B. (2001). Turbulent flows. *Measurement Science and Technology*, 12(11), 2020–2021. <https://doi.org/10.1088/0957-0233/12/11/705>
- Qi, Z., Bai, L., Xiong, H., & Xie, Z. (2024). *Not all noises are created equally: Diffusion noise selection and optimization*. arXiv preprint arXiv:2407.14041.
- Qu, Y., Nathaniel, J., Li, S., & Gentine, P. (2024). Deep generative data assimilation in multimodal setting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 449–459).
- Rajagopal, E., Babu, A. N. S., Ryu, T., Haley, P. J., Jr., Mirabito, C., & Lermusiaux, P. F. J. (2023). Evaluation of deep neural operator models toward ocean forecasting. In *Oceans 2023 IEEE/MTS Gulf Coast. Biloxi*. <https://doi.org/10.23919/OCEANS52994.2023.10337380>.MS
- Rampal, N., Hobeichi, S., Gibson, P. B., Baño-Medina, J., Abramowitz, G., Beucler, T., et al. (2024). Enhancing regional climate downscaling through advances in machine learning. *Artificial Intelligence for the Earth Systems*, 3(2), 230066. <https://doi.org/10.1175/aies-d-23-0066.1>
- Rezende, D., & Mohamed, S. (2015). Variational inference with normalizing flows. In *International conference on machine learning* (pp. 1530–1538).
- Rhines, P. B. (1975). Waves and turbulence on a beta-plane. *Journal of Fluid Mechanics*, 69(3), 417–443. <https://doi.org/10.1017/s0022112075001504>
- Rissanen, S., Heinonen, M., & Solin, A. (2024). *Free hunch: Denoiser covariance estimation for diffusion models without extra costs*. arXiv preprint arXiv:2410.11149.
- Robbins, H. E. (1992). An empirical bayes approach to statistics. In *Breakthroughs in statistics: Foundations and basic theory* (pp. 388–394). Springer.
- Robinson, A. R., & Lermusiaux, P. F. J. (2001). Data assimilation (physical/interdisciplinary). In J. H. Steele, S. Thorpe, & K. Turekian (Eds.), *Encyclopedia of ocean sciences* (pp. 302–308). Academic Press. <https://doi.org/10.1006/rwos.2001.0404>
- Robinson, A. R., Lermusiaux, P. F. J., & Sloan, N. Q., III. (1998). Data assimilation. In K. H. Brink & A. R. Robinson (Eds.), *The global coastal ocean-processes and methods* (Vol. 10, pp. 541–594). John Wiley and Sons.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10684–10695).
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—miccai 2015: 18th international conference, munich, Germany, October 5-9, 2015, proceedings, part iii* (Vol. 18, pp. 234–241).
- Ross, A., Li, Z., Perezhogin, P., Fernandez-Granda, C., & Zanna, L. (2023). Benchmarking of machine learning ocean subgrid parameterizations in an idealized model. *Journal of Advances in Modeling Earth Systems*, 15(1), e2022MS003258. <https://doi.org/10.1029/2022ms003258>
- Rozet, F., & Louppe, G. (2023). Score-based data assimilation. *Advances in Neural Information Processing Systems*, 36, 40521–40541.
- Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., et al. (2022). Palette: Image-to-image diffusion models. In *Acm siggraph 2022 conference proceedings* (pp. 1–10).

- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., & Norouzi, M. (2022b). Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4), 4713–4726. <https://doi.org/10.1109/tpami.2022.3204461>
- Sardar, M., Skillen, A., Zimoń, M., Draycott, S., & Revell, A. (2024). Spectrally decomposed denoising diffusion probabilistic models for generative turbulence super-resolution. *Physics of Fluids*, 36(11), 115179. <https://doi.org/10.1063/5.0231664>
- Short, C. J., & Petch, J. (2022). Reducing the spin-up of a regional nwp system without data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 148(745), 1623–1643. <https://doi.org/10.1002/qj.4268>
- Simmons, H. L., Powell, B. S., Merrifield, S. T., Zedler, S. E., & Colin, P. L. (2019). Dynamical downscaling. *Oceanography*, 32(4), 84–91.
- Smagorinsky, J. (1963). General circulation experiments with the primitive equations: I. the basic experiment. *Monthly Weather Review*, 91(3), 99–164. [https://doi.org/10.1175/1520-0493\(1963\)091<0099:gcewtp>2.3.co;2](https://doi.org/10.1175/1520-0493(1963)091<0099:gcewtp>2.3.co;2)
- Sofos, F., & Drikakis, D. (2025). A review of deep learning for super-resolution in fluid flows. *Physics of Fluids*, 37(4), 041303. <https://doi.org/10.1063/5.0265738>
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., & Ganguli, S. (2015). Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning* (pp. 2256–2265).
- Sondergaard, T., & Lermusiaux, P. F. J. (2013). Data assimilation with Gaussian Mixture Models using the Dynamically Orthogonal field equations. Part I: Theory and scheme. *Monthly Weather Review*, 141(6), 1737–1760. <https://doi.org/10.1175/MWR-D-11-00295.1>
- Song, J., Meng, C., & Ermon, S. (2020). *Denoising diffusion implicit models*. arXiv preprint arXiv:2010.02502.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2020). *Score-based generative modeling through stochastic differential equations*. arXiv preprint arXiv:2011.13456.
- Souza, A. N., Silvestri, S., Deck, K., Bischoff, T., Flierl, G., & Ferrari, R. (2025). *Surface to seafloor: A generative machine learning framework for decoding the ocean interior state*. arXiv preprint arXiv:2504.15308.
- Srinivasan, K., Chekroun, M. D., & McWilliams, J. C. (2024). Turbulence closure with small, local neural networks: Forced two-dimensional and β -plane flows. *Journal of Advances in Modeling Earth Systems*, 16(4), e2023MS003795. <https://doi.org/10.1029/2023ms003795>
- Stommel, H. (1948). The westward intensification of wind-driven ocean currents. *Eos, Transactions American Geophysical Union*, 29(2), 202–206.
- Storto, A., Alvera-Azcárate, A., Balmaseda, M. A., Barth, A., Chevallier, M., Counillon, F., et al. (2019). Ocean reanalyses: Recent advances and unsolved challenges. *Frontiers in Marine Science*, 6, 418. <https://doi.org/10.3389/fmars.2019.00418>
- Subel, A., Guan, Y., Chattopadhyay, A., & Hassanzadeh, P. (2023). Explaining the physics of transfer learning in data-driven turbulence modeling. *PNAS Nexus*, 2(3), pgad015. <https://doi.org/10.1093/pnasnexus/pgad015>
- Sun, H., Yu, L., Dai, B., Schuurmans, D., & Dai, H. (2022). *Score-based continuous-time discrete diffusion models*. arXiv preprint arXiv:2211.16750.
- Sundar, R., Hu, Y., Parashar, N., Blanchard, A., & Dodov, B. (2024). *Taudiff: Highly efficient kilometer-scale downscaling using generative diffusion models*. arXiv preprint arXiv:2412.13627.
- Suresh Babu, A. N., Sadam, A., & Lermusiaux, P. F. J. (2025a). Evaluation of analytical turbulence closures for quasi-geostrophic ocean flows with coastal boundaries. In *Oceans 2025 IEEE/MTS Great Lakes* (pp. 1–10). <https://doi.org/10.23919/OCEANS9106.2025.11245082>
- Suresh Babu, A. N., Sadam, A., & Lermusiaux, P. F. J. (2025b). Machine learning code and weights accompanying the paper “Guided unconditional and conditional generative models for super-resolution and inference of quasi-geostrophic turbulence” [Software]. *Zenodo*. <https://doi.org/10.5281/zenodo.15750244>
- Suresh Babu, A. N., Sadam, A., & Lermusiaux, P. F. J. (2025c). Simulation code accompanying the paper “Guided unconditional and conditional generative models for super-resolution and inference of quasi-geostrophic turbulence” [Software]. *Zenodo*. <https://doi.org/10.5281/zenodo.17282193>
- Suresh Babu, A. N., Sadam, A., & Lermusiaux, P. F. J. (2025d). Simulation data accompanying the paper “Guided unconditional and conditional generative models for super-resolution and inference of quasi-geostrophic turbulence” [Dataset]. *Zenodo*. <https://doi.org/10.5281/zenodo.15742146>
- Taburet, G., Sanchez-Roman, A., Ballarotta, M., Pujol, M.-I., Legeais, J.-F., Fournier, F., et al. (2019). Duacs dt2018: 25 years of reprocessed sea level altimetry products. *Ocean Science*, 15(5), 1207–1224. <https://doi.org/10.5194/os-15-1207-2019>
- Taylor, J. R., & Thompson, A. F. (2023). Submesoscale dynamics in the upper ocean. *Annual Review of Fluid Mechanics*, 55(1), 103–127. <https://doi.org/10.1146/annurev-fluid-031422-095147>
- Theiss, J. (2006). A generalized rhines effect and storms on jupiter. *Geophysical Research Letters*, 33(8), L08809. <https://doi.org/10.1029/2005gl025379>
- Thomas, L. N., Tandon, A., & Mahadevan, A. (2008). Submesoscale processes and dynamics. *Ocean modeling in an Eddy Regime*, 177, 17–38.
- Towne, A., Lozano-Durán, A., & Yang, X. (2020). Resolvent-based estimation of space–time flow statistics. *Journal of Fluid Mechanics*, 883, A17. <https://doi.org/10.1017/jfm.2019.854>
- Vallis, G. K. (2017). *Atmospheric and oceanic fluid dynamics*. Cambridge University Press.
- Van Leeuwen, P. J., Künsch, H. R., Nerger, L., Potthast, R., & Reich, S. (2019). Particle filters for high-dimensional geoscience applications: A review. *Quarterly Journal of the Royal Meteorological Society*, 145(723), 2335–2365. <https://doi.org/10.1002/qj.3551>
- von Storch, H., Langenberg, H., & Feser, F. (2000). A spectral nudging technique for dynamical downscaling purposes. *Monthly Weather Review*, 128(10), 3664–3673. [https://doi.org/10.1175/1520-0493\(2000\)128<3664:asntfd>2.0.co;2](https://doi.org/10.1175/1520-0493(2000)128<3664:asntfd>2.0.co;2)
- Vreman, B., Geurts, B., & Kuerten, H. (1996). Large-eddy simulation of the temporal mixing layer using the clark model. *Theoretical and Computational Fluid Dynamics*, 8(4), 309–324. <https://doi.org/10.1007/s001620050018>
- Wan, Z. Y., Baptista, R., Boral, A., Chen, Y.-F., Anderson, J., Sha, F., & Zepeda-Núñez, L. (2023). Debias coarsely, sample conditionally: Statistical downscaling through optimal transport and probabilistic diffusion models. *Advances in Neural Information Processing Systems*, 36, 47749–47763.
- Wang, B., & Vastola, J. J. (2023). Diffusion models generate images like painters: An analytical theory of outline first, details later. *arXiv preprint arXiv:2303.02490*.
- Wang, Z., Jiang, Y., Zheng, H., Wang, P., He, P., Wang, Z., & others (2023). Patch diffusion: Faster and more data-efficient training of diffusion models. *Advances in Neural Information Processing Systems*, 36, 72137–72154.
- Watt, R. A., & Mansfield, L. A. (2024). *Generative diffusion-based downscaling for climate*. arXiv preprint arXiv:2404.17752.
- Węglarczyk, S. (2018). Kernel density estimation and its application. In *ITM Web of Conferences* (Vol. 23, 00037). <https://doi.org/10.1051/itmconf/20182300037>
- Wiatrak, M., Albrecht, S. V., & Nystrom, A. (2019). Stabilizing generative adversarial networks: A survey. *arXiv preprint arXiv:1910.00927*.
- Wilby, R. L., Wigley, T., Conway, D., Jones, P., Hewitson, B., Main, J., & Wilks, D. (1998). Statistical downscaling of general circulation model output: A comparison of methods. *Water Resources Research*, 34(11), 2995–3008. <https://doi.org/10.1029/98wr02577>

- Wilczek, M., & Friedrich, R. (2009). Dynamical origins for Non-gaussian vorticity distributions in turbulent flows. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, 80(1), 016316. <https://doi.org/10.1103/physreve.80.016316>
- Williams, G. P. (1978). Planetary circulations: 1. barotropic representation of jovian and terrestrial turbulence. *Journal of the Atmospheric Sciences*, 35(8), 1399–1426. [https://doi.org/10.1175/1520-0469\(1978\)035<1399:pcbroy>2.0.co;2](https://doi.org/10.1175/1520-0469(1978)035<1399:pcbroy>2.0.co;2)
- Wunsch, C. (2006). *Discrete inverse and state estimation problems: With geophysical fluid applications*. Cambridge University Press.
- Xie, Y., Lu, C., & Browning, G. L. (2002). Impact of formulation of cost function and constraints on three-dimensional variational data assimilation. *Monthly Weather Review*, 130(10), 2433–2447. [https://doi.org/10.1175/1520-0493\(2002\)130<2433:iocf>2.0.co;2](https://doi.org/10.1175/1520-0493(2002)130<2433:iocf>2.0.co;2)
- Xu, Q., Zhuang, Z., Pan, Y., & Wen, B. (2023). Super-resolution reconstruction of turbulent flows with a transformer-based deep learning framework. *Physics of Fluids*, 35(5), 055130. <https://doi.org/10.1063/5.0149551>
- Xu, Z., Han, Y., & Yang, Z. (2019). Dynamical downscaling of regional climate: A review of methods and limitations. *Science China Earth Sciences*, 62(2), 365–375. <https://doi.org/10.1007/s11430-018-9261-5>
- Yin, J., Liang, S., Liu, S., Bao, F., Chipilski, H. G., Lu, D., & Zhang, G. (2024). A scalable real-time data assimilation framework for predicting turbulent atmosphere dynamics. In *Sc24-w: Workshops of the international conference for high performance computing, networking, storage and analysis* (pp. 11–18).
- Zanna, L., & Bolton, T. (2020). Data-driven equation discovery of ocean mesoscale closures. *Geophysical Research Letters*, 47(17), e2020GL088376. <https://doi.org/10.1029/2020gl088376>
- Zanna, L., Mana, P. P., Anstey, J., David, T., & Bolton, T. (2017). Scale-aware deterministic and stochastic parametrizations of eddy-mean flow interaction. *Ocean Modelling*, 111, 66–80. <https://doi.org/10.1016/j.ocemod.2017.01.004>
- Zhai, S., Zhang, R., Nakkiran, P., Berthelot, D., Gu, J., Zheng, H., et al. (2024). *Normalizing flows are capable generative models*. arXiv preprint arXiv:2412.06329.
- Zhang, Q., & Chen, Y. (2022). *Fast sampling of diffusion models with exponential integrator*. arXiv preprint arXiv:2204.13902.
- Zhang, Z., Miao, M., Qiu, B., Tian, J., Jing, Z., Chen, G., & Zhao, W. (2024). Submesoscale eddies detected by swot and moored observations in the northwestern pacific. *Geophysical Research Letters*, 51(15), e2024GL110000. <https://doi.org/10.1029/2024gl110000>
- Zhou, T., Krahenbuhl, P., Aubry, M., Huang, Q., & Efros, A. A. (2016). Learning dense correspondence via 3d-guided cycle consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 117–126).
- Zhou, Z., He, G., Wang, S., & Jin, G. (2019). Subgrid-scale model for large-eddy simulation of isotropic turbulent flows using an artificial neural network. *Computers and Fluids*, 195, 104319. <https://doi.org/10.1016/j.compfluid.2019.104319>
- Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 2223–2232).
- Zupanski, M., Fletcher, S. J., Navon, I. M., Uzunoglu, B., Heikes, R. P., Randall, D. A., et al. (2006). Initiation of ensemble data assimilation. *Tellus A: Dynamic Meteorology and Oceanography*, 58(2), 159–170. <https://doi.org/10.1111/j.1600-0870.2006.00173.x>