

PARTICLE-GUIDED DIFFUSION FOR GAS-PHASE REACTION KINETICS

Andrew Millard

Department of Computer and Information Science
Linköping University
Linköping, Sweden
andmi73@liu.se

Henrik Pedersen

Department of Physics, Chemistry and Biology
Linköping University
Linköping, Sweden

ABSTRACT

Physics-guided sampling with diffusion model priors has shown promise for solving partial differential equation (PDE) governed problems, but applications to chemically meaningful reaction-transport systems remain limited. We apply diffusion-based guided sampling to gas-phase chemical reactions by training on solutions of the advection-reaction-diffusion (ARD) equation across varying parameters. The method generates physically consistent concentration fields and accurately predicts outlet concentrations, including at unseen parameter values, demonstrating the potential of diffusion models for inference in reactive transport.

1 INTRODUCTION

Chemical kinetics Atkinson & Carter (1984) is fundamental for understanding reaction mechanisms and underpins applications from atmospheric chemistry to catalysis. Given a set of reaction mechanisms, kinetic rate theory expresses the reaction rate of each chemical species as a function of species concentrations and rate constants (Atkins et al., 2023). The resulting system of differential equations governs the temporal evolution of species concentrations and is essential for modelling complex reaction networks such as chemical vapor deposition (CVD) (Danielsson et al., 2020).

Classical numerical solvers for partial differential equations (PDEs), such as finite element methods (FEMs) (Monk, 2003; Kreiss & Scherer, 1974), often require careful design of discretisation schemes and solver parameters to achieve high accuracy. When applied to problems involving broad parameter ranges or repeated parameter sweeps, this can result in significant computational overhead and limited flexibility without repeated recalibration.

Recent work has shown that deep generative models, including diffusion models (Erichson et al., 2025; Kerrigan et al., 2024), can approximate the complex spatiotemporal structure of PDE systems. Building on this foundation, guided sampling methods using diffusion-based generative priors have been proposed to produce physically admissible PDE solutions, with successful applications to systems such as the Navier-Stokes equations and Darcy flow (Huang et al., 2024; Millard et al., 2026). However, most existing studies focus on idealised benchmark problems, and applications to realistic, parameter-varying physical systems remain limited.

In this work, we demonstrate that diffusion-based guided sampling, when trained on data spanning a range of parameters, can generate accurate solutions to complex PDE systems at previously unseen parameter values. We further show that this approach enables the simulation of gas-phase reaction kinetics across multiple temporal scales and allows for accurate estimation of spatiotemporal concentration fields from sparse observations.

2 BACKGROUND

Time-Varying Partial Differential Equations. We consider time-dependent PDEs. For a spatial domain Ω and a time horizon $[0, T]$, a PDE can be written as

$$\begin{aligned} f(u, c, \tau, a) &= 0, \quad c \in \Omega, \quad \tau \in [0, T], \\ u(c, \tau) &= g(c, \tau), \quad c \in \partial\Omega, \quad \tau \in [0, T], \\ u(c, 0) &= a(c), \end{aligned} \quad (1)$$

where c denotes a spatial coordinate, τ denotes physical (PDE) time, $u \in \mathcal{U}$ is the PDE solution field, $a \in \mathcal{A}$ specifies the initial condition (and, more generally, fixed PDE coefficients) and g are the boundary conditions.

Throughout this work, the abstract solution field $u(c, \tau)$ represents a spatiotemporal physical field. In later sections we denote this field as a concentration field $C(c, \tau)$ governed by a reaction–transport PDE. The notation $C(\cdot, \tau)$ refers to the spatial concentration field over Ω at a fixed PDE time τ .

Sampling for Generative Diffusion Models. Generative diffusion models (Sohl-Dickstein et al., 2015; Song et al., 2021; Ho et al., 2020; Cao et al., 2024; Karras et al., 2022) provide a probabilistic framework for sampling from complex, high-dimensional data distributions $p_{\text{data}}(x)$.

The forward process progressively corrupts $p_{\text{data}}(x)$ with Gaussian noise according to a noise schedule $\sigma_t \in [0, \sigma_{\text{max}}]$, $t \in [0, T]$, yielding a reference distribution $p(x_T) \approx \mathcal{N}(0, \sigma_{\text{max}}^2 I)$. Sampling proceeds by drawing x_T from this reference distribution and reversing the noising process to obtain $x_0 \sim p_{\text{data}}(x)$. The reverse-time dynamics are described by the stochastic differential equation

$$dx_t = -2\dot{\sigma}_t \sigma_t \nabla_x \log p_t(x_t, \sigma_t) dt + \sqrt{2\dot{\sigma}_t \sigma_t} dW_t, \quad (2)$$

where W_t is a Brownian motion and $\dot{\sigma}_t$ denotes a time derivative. Following Karras et al. (2022), the score is estimated using a denoising network δ_θ as

$$\nabla_x \log p_t^\theta(x_t, \sigma_t) = \frac{x_t - \delta_\theta(x_t, \sigma_t)}{\sigma_t^2}. \quad (3)$$

Guided Sampling. We consider the setting in which physical parameters are fixed and known, and the diffusion model prior is trained over realisations of the PDE solution field alone. Accordingly, diffusion samples take the form

$$x \in \mathcal{X} = \mathcal{U}, \quad (4)$$

where x represents a discretised realisation of the PDE solution field, which in this work corresponds to a concentration field at a single physical time. Conditioning on observations y , we aim to sample from the posterior $p_\theta(x | y) \propto p(y | x) p_\theta(x)$ where $p_\theta(x)$ denotes the diffusion model prior. Sampling is performed by simulating the guided reverse-time SDE

$$dx_t = -2\dot{\sigma}_t \sigma_t \nabla_{x_t} \log p_t^\theta(x_t, \sigma_t) dt - 2\dot{\sigma}_t \sigma_t \nabla_{x_t} \log p_t^\theta(y | x_t, \sigma_t) dt + \sqrt{2\dot{\sigma}_t \sigma_t} dW_t, \quad (5)$$

where the additional gradient term provides guidance toward consistency with observations and governing equations. We approximate the reverse process using a guided Euler-Maruyama (GEM) update (Millard et al., 2026) (switching notation from continuous diffusion time t to discrete diffusion sampling indexes k)

$$x_{k-1} = x_k + (\sigma_{k-1}^2 - \sigma_k^2) \frac{x_k - \delta_\theta(x_k, \sigma_k)}{\sigma_k^2} + (\sigma_{k-1}^2 - \sigma_k^2) \nabla_{x_k} \log \tilde{p}_\theta(y | x_k) + \sqrt{\sigma_{k-1}^2 - \sigma_k^2} \zeta, \quad (6)$$

with $\zeta \sim \mathcal{N}(0, I)$ and \tilde{p}_θ denotes a likelihood function whose formulation is detailed in later sections. The time interval Δt is implicitly accounted for in the noise scheduler.

3 METHODOLOGY

Advection–Reaction–Diffusion Model. We consider transport-reaction systems governed by advection-reaction-diffusion (ARD) equations (Rubio et al., 2008; Cosner, 2014). A diffusion model sample

$$x \equiv C(\cdot, \tau_n)$$

Table 1: Outlet concentration summary averaged over common simulation-seed pairs for SMC and ODE methods. We report the mean \pm standard deviation for all metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and outlet concentrations of the chemical species at the final timestep (Pred Final). Best values per species are shown in bold (lowest RMSE/MAE and closest final prediction to GT). The RMSE and the MAE are the mean errors across all time steps.

Method	Species	GT Final	Pred Final	RMSE	MAE
SMC	NO	9.710×10^{-3}	9.879×10^{-3} $\pm 1.35 \times 10^{-2}$	5.445×10^{-4} $\pm 8.90 \times 10^{-4}$	1.664×10^{-4} $\pm 2.00 \times 10^{-4}$
	O3	4.942×10^{-3}	5.093×10^{-3} $\pm 8.68 \times 10^{-3}$	5.517×10^{-4} $\pm 9.52 \times 10^{-4}$	1.694×10^{-4} $\pm 1.98 \times 10^{-4}$
	NO2	2.464×10^{-2}	2.447×10^{-2} $\pm 1.50 \times 10^{-2}$	6.361×10^{-4} $\pm 8.18 \times 10^{-4}$	2.322×10^{-4} $\pm 1.92 \times 10^{-4}$
ODE	NO	9.710×10^{-3}	7.020×10^{-3} $\pm 9.51 \times 10^{-3}$	4.430×10^{-3} $\pm 2.73 \times 10^{-3}$	1.681×10^{-3} $\pm 1.19 \times 10^{-3}$
	O3	4.942×10^{-3}	4.699×10^{-3} $\pm 7.91 \times 10^{-3}$	3.468×10^{-3} $\pm 2.17 \times 10^{-3}$	1.165×10^{-3} $\pm 7.94 \times 10^{-4}$
	NO2	2.464×10^{-2}	1.958×10^{-2} $\pm 1.13 \times 10^{-2}$	3.467×10^{-3} $\pm 2.46 \times 10^{-3}$	1.559×10^{-3} $\pm 1.24 \times 10^{-3}$

represents a discretised concentration field at a single physical time $\tau_n \in [0, T]$, corresponding to the PDE solution field $u(\cdot, \tau_n)$ introduced in Section 2. In an axisymmetric cylindrical reactor of length L and radius A , the concentration field $C_s(r, z, \tau)$ of species s evolves according to

$$\frac{\partial C_s}{\partial \tau} + u(r) \frac{\partial C_s}{\partial z} = D \nabla^2 C_s + R_s(C), \quad (7)$$

where $r \in [0, A]$, $z \in [0, L]$, D is the molecular diffusivity, and $R_s(C)$ denotes reaction source terms (Robert J. Kee, 2003). The axial velocity field is prescribed as

$$u(r) = 2U_{\text{avg}} \left(1 - \left(\frac{r}{A} \right)^2 \right).$$

U_{avg} is the average cross sectional velocity which we assume is the same for all species. In our experiments, physical parameters in the ARD model are treated as known and fixed and inference is performed only over the concentration fields.

Residual-Based Likelihood for Time-Series Reconstruction. A time series of concentration fields is reconstructed by applying the guided Euler–Maruyama (GEM) proposal within a sequential Monte Carlo (SMC) framework (Chopin & Papaspiliopoulos, 2020; Del Moral et al., 2006; Zhao, 2026) with a diffusion prior (Millard et al., 2026) independently at each physical time step τ_n . At each τ_n , the diffusion process is initialised from Gaussian noise and a new SMC-guided denoising procedure is run to reconstruct $C(\cdot, \tau_n)$. At each step, candidate reconstructions of $C(\cdot, \tau_n)$ are evaluated using a likelihood that combines agreement with sparse observations and consistency with the governing ARD dynamics.

We can express our likelihood in terms of the residuals of the sparse observations as well as the PDE equation residuals. Therefore we define *PDE residual* likelihood function:

$$\log p(y | x) \propto -\frac{1}{\sigma_{\text{obs}}} \frac{1}{N_{\text{obs}}} \|(C_{\text{obs}} - \mathcal{M} \odot x)\|_2^2 - \frac{1}{\sigma_{\text{pde}}} \frac{1}{AL} \|\mathcal{R}_n(\hat{C}; C^{\tau_{n-1}})\|_2^2, \quad (8)$$

where \mathcal{M} is a binary observation mask, \hat{C} denotes the physical concentration field corresponding to x , C_{obs} are the sparse observations of the ground truth concentration field, and σ_{obs} , σ_{pde} are relative weighting factors. N_{obs} are the number of sparse observations. Following Wu et al. (2023), intermediate diffusion likelihoods are approximated by evaluating the likelihood at the denoised reconstruction $p_t^\theta(y | x_t, \sigma_t) \approx p(y | x_0 = \delta_\theta(x_t, \sigma_t)) =: \tilde{p}_\theta(y | x_t)$, which coincides with the true likelihood at time 0, $\tilde{p}_\theta(y | x_0) = p(y | x_0)$. Defining the spatial ARD operator

$$\mathcal{F}(C) = u(r) \frac{\partial C}{\partial z} - D \left(\frac{\partial^2 C}{\partial z^2} + \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial C}{\partial r} \right) \right) - R(C), \quad (9)$$

where $R(C)$ denotes the nonlinear reaction source term acting component-wise on the species concentration vector C , derived from the chemical reaction mechanism and kinetic rate laws. The PDE residual at discrete physical times $\{\tau_n\}$ is

$$\mathcal{R}_n(C^{\tau_n}; C^{\tau_{n-1}}) = \frac{C^{\tau_n} - C^{\tau_{n-1}}}{\Delta \tau_n} + \mathcal{F}(C^{\tau_n}), \quad \Delta \tau_n := \tau_n - \tau_{n-1}, \quad n \geq 1. \quad (10)$$

$C^{\tau_{n-1}}$ is taken as the previous reconstructed field. At $\tau_0 = 0$, the concentration field is prescribed as $C^{\text{init}} = 0$ as we assume no chemical species present at this time.

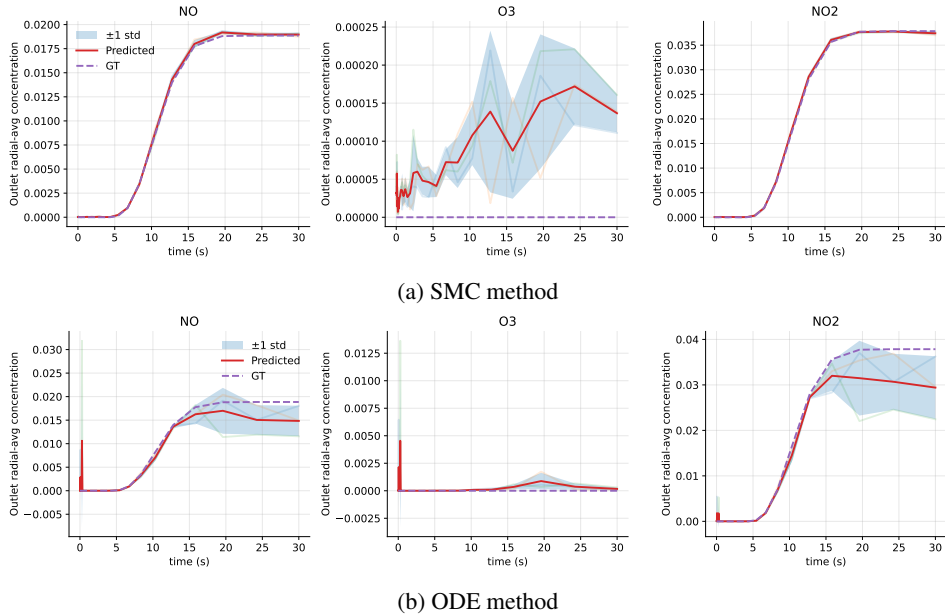


Figure 1: Predicted species concentration at the outlet over time for an example simulation averaged over three seeds. The quantitative results for these simulations are given in Table 2.

Table 2: Quantitative simulation results corresponding to Figures 1. Parameter values are $U_{avg} = 0.093$, $D = 6.43 \times 10^{-6}$, $\kappa = 7.51 \times 10^3$.

Method	Species	GT Final	Pred Final	RMSE	MAE
SMC	NO	1.89×10^{-2}	$1.90 \times 10^{-2} \pm 1.96 \times 10^{-4}$	$1.58 \times 10^{-4} \pm 3.38 \times 10^{-5}$	$8.17 \times 10^{-5} \pm 1.24 \times 10^{-5}$
	O3	2.24×10^{-9}	$1.37 \times 10^{-4} \pm 2.50 \times 10^{-5}$	$7.41 \times 10^{-5} \pm 9.09 \times 10^{-6}$	$5.49 \times 10^{-5} \pm 7.75 \times 10^{-6}$
	NO2	3.79×10^{-2}	$3.74 \times 10^{-2} \pm 3.80 \times 10^{-4}$	$2.03 \times 10^{-4} \pm 7.11 \times 10^{-5}$	$1.08 \times 10^{-4} \pm 2.71 \times 10^{-5}$
ODE	NO	1.89×10^{-2}	$1.48 \times 10^{-2} \pm 3.24 \times 10^{-3}$	$2.90 \times 10^{-3} \pm 2.78 \times 10^{-3}$	$8.90 \times 10^{-4} \pm 7.88 \times 10^{-4}$
	O3	2.24×10^{-9}	$1.70 \times 10^{-4} \pm 1.18 \times 10^{-4}$	$1.29 \times 10^{-3} \pm 1.06 \times 10^{-3}$	$2.73 \times 10^{-4} \pm 2.09 \times 10^{-4}$
	NO2	3.79×10^{-2}	$2.94 \times 10^{-2} \pm 6.91 \times 10^{-3}$	$2.81 \times 10^{-3} \pm 1.63 \times 10^{-3}$	$1.03 \times 10^{-3} \pm 6.25 \times 10^{-4}$

4 EXPERIMENTAL SETUP AND RESULTS

Experimental Details. To evaluate the proposed method, we model the gas-phase reaction



using a reduced kinetic description involving the species set $s \in \{\text{NO}, \text{O}_3, \text{NO}_2\}$. Assuming elementary mass-action kinetics, the reaction source term $R(C)$ appearing in equation 7 acts component-wise on the concentration vector $C = (C_{\text{NO}}, C_{\text{O}_3}, C_{\text{NO}_2})$ and is given by

$$R(C) = (-\kappa C_{\text{NO}} C_{\text{O}_3}, -\kappa C_{\text{NO}} C_{\text{O}_3}, \kappa C_{\text{NO}} C_{\text{O}_3}), \tag{12}$$

where κ is the second-order reaction rate constant. We consider an axisymmetric cylindrical reactor of radius a and length L , with NO and O₃ introduced at the inlet and advected downstream at mean velocity U_{avg} , where they react to form NO₂. Sparse sensors provide partial concentration observations, and the objective is to reconstruct the full concentration fields and estimate outlet concentrations.

We compare our method to a Guided Euler (GE) discretisation which is used to solve the reverse-time ODE. Details of this can be found in Appendix A.1. Synthetic training and evaluation data are generated by numerically simulating the ARD equations for equation 11 on the cylindrical domain $(r, z) \in [0, A] \times [0, L]$ using a 64×64 grid and 32 logarithmically spaced time snapshots up to $\tau_{\mathcal{T}} = 30$ seconds. Physical parameters are randomised across simulations and further details are provided in Appendix A.3. For evaluation, thirty simulations are randomly selected from the test

set and the concentration fields are sequentially reconstructed at each physical time step τ_n , with outlet concentrations estimated at each step. GEM is used within an SMC framework (Millard et al., 2026). Results are averaged over 3 independent runs per simulation. A schematic of the experimental setup is also shown in Appendix A.3 with further model training and sampling details given in Appendix A.4.

Discussion. Performance metrics are reported in Table 1, and predicted outlet concentration time series for an example simulation trajectory are shown in Figure 1. The methods accurately estimate outlet concentrations for the higher-abundance species, while the lower-abundance species exhibit larger relative errors despite maintaining low RMSE and MAE values. Because the governing ARD equations conserve mass, small absolute errors in the higher-concentration species must be compensated by the remaining species, which can lead to large relative errors for species present at very low concentrations when compared to their ground-truth values. We notice that the GEM proposal outperforms the GE method, with the RMSE and MAE errors being an order of magnitude smaller for the former. Appendix B gives an ablation study of the RMSE and MAE with differing number of observations inside the simulated reactor.

Appendix C presents reconstructed concentration fields across physical time for the SMC-GEM method for the an example trajectory (corresponding to Figure 1 and Table 2) as well as the reactor-wide species fraction change over the simulated period (Figure 3). These reconstructions closely match the ground-truth fields, indicating that the larger relative errors observed for O_3 are unlikely to be practically significant given the small absolute concentrations. The reconstructed fields further demonstrate the model’s ability to capture gas-phase reaction kinetics across both short and long time scales. Notably, although most reactive interactions occur within the first few seconds before the system approaches equilibrium, the model accurately reproduces the full temporal dynamics even when the data are sampled more densely at early times than at later times. These results also show that our method was able to generalise across a range of different potential parameter values, as none of the exact parameter combinations had been observed in the training dataset.

These findings demonstrate that guided sampling with diffusion priors can recover physically consistent species concentration fields in reaction-transport systems modelling gas-phase kinetics. The stochastic method produces superior results to the Euler method. The approaches generalise to previously unseen parameter regimes and accurately captures dynamics over extended time horizons.

ACKNOWLEDGMENT

This work was partially supported by the Swedish Research Council and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg (KAW) Foundation. Computations were enabled by the supercomputing resource Berzelius provided by National Supercomputer Centre at Linköping University and the KAW foundation.

ETHICS STATEMENT

This paper presents work whose goal is to advance the field of machine learning through improved methods for scientific modelling of gas-phase reaction kinetics. While machine learning research can have broad societal implications, we do not foresee any direct negative societal impacts arising from this work.

REFERENCES

- Peter William Atkins, Julio De Paula, and James Keeler. *Atkins’ Physical Chemistry*. Oxford university press, 2023.
- Roger Atkinson and William PL Carter. Kinetics and mechanisms of the gas-phase reactions of ozone with organic compounds under atmospheric conditions. *Chemical Reviews*, 84(5):437–470, 1984.
- Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z Li. A Survey on Generative Diffusion Models. *IEEE Transactions on Knowledge and Data Engineering*, 36(7):2814–2830, 2024.

- Nicolas Chopin and Omiros Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*, volume 4. Springer, 2020.
- Chris Cosner. Reaction-diffusion-advection models for the effects and evolution of dispersal. *Discrete and Continuous Dynamical Systems*, 34(5):1701–1745, 2014. ISSN 1078-0947. doi: 10.3934/dcds.2014.34.1701. URL <https://www.aims sciences.org/article/id/6fe19f9a-2394-4838-9ed0-e1883781ca39>.
- Orjan Danielsson, Matts Karlsson, Pitsiri Sukkaew, Henrik Pedersen, and Lars Ojamae. A Systematic Method for Predictive in silico Chemical Vapor Deposition. *The Journal of Physical Chemistry C*, 124(14):7725–7736, 2020.
- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo Samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.
- N. Benjamin Erichson, Vinicius Mikuni, Dongwei Lyu, Yang Gao, Omri Azencot, Soon Hoe Lim, and Michael W. Mahoney. FLEX: A Backbone for Diffusion-Based Modeling of Spatio-temporal Physical Systems. *CoRR*, abs/2505.17351, May 2025. URL <https://doi.org/10.48550/arXiv.2505.17351>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. DiffusionPDE: Generative PDE-Solving under Partial Observation, 2024. URL <https://openreview.net/forum?id=z0I2SbjN0R>.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space of Diffusion-Based Generative Models, 2022. URL <https://openreview.net/forum?id=k7FuTOWMOc7>.
- Gavin Kerrigan, Giosue Migliorini, and Padhraic Smyth. Functional Flow Matching. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3934–3942. PMLR, 02–04 May 2024. URL <https://proceedings.mlr.press/v238/kerrigan24a.html>.
- H-O Kreiss and Godela Scherer. Finite Element and Finite Difference Methods for Hyperbolic Partial Differential Equations. In *Mathematical Aspects of Finite Elements in Partial Differential Equations*, pp. 195–212. Elsevier, 1974.
- Andrew Millard, Fredrik Lindsten, and Zheng Zhao. Particle-guided diffusion models for partial differential equations. *arXiv preprint 2601.23262*, 2026.
- Peter Monk. *Finite Element Methods for Maxwell’s Equations*. Oxford University Press, 2003.
- Peter Glarborg Robert J. Kee, Michael E. Coltrin. *The Conservation Equations*, pp. 67–149. John Wiley & Sons, Ltd, 2003. ISBN 9780471461296. doi: <https://doi.org/10.1002/0471461296.ch3>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/0471461296.ch3>.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241. Springer, 2015.
- A.D. Rubio, A. Zalts, and C.D. El Hasi. Numerical solution of the advection–reaction–diffusion equation at different scales. *Environmental Modelling & Software*, 23(1):90–95, 2008. ISSN 1364-8152. doi: <https://doi.org/10.1016/j.envsoft.2007.05.009>. URL <https://www.sciencedirect.com/science/article/pii/S136481520700093X>.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations, 2021. URL <https://openreview.net/forum?id=PXTIG12RRHS>.

Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David Blei, and John Patrick Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, 2023. URL <https://openreview.net/forum?id=eWKqrlzCRv>.

Zheng Zhao. Generative diffusion posterior sampling for informative likelihoods. *Communications in Information and Systems*, 26(1):151–167, 2026. Special issue for celebrating Thomas Kailath’s 90th birthday.




Figure 2: Diagram of the experiment.

A REPRODUCIBILITY

A.1 GUIDED EULER FOR SOLVING REVERSE ODES

Equation 6 gives the Euler-Maruyama discretisation which is commonly used to solve for the reverse-time SDE. In order to solve the reverse-time ODE:

$$dx_t = -\dot{\sigma}_t \sigma_t (\nabla_{x_t} \log p_t^\theta(x_t, \sigma_t) + \log p_t^\theta(y | x_t, \sigma_t)) dt, \quad (13)$$

we can use an Euler discretisation which is given by the following:

$$x_{k-1} = x_k + \frac{1}{2}(\sigma_{k-1}^2 - \sigma_k^2) \frac{x_k - \delta_\theta(x_k, \sigma_k)}{\sigma_k^2} + \frac{1}{2}(\sigma_{k-1}^2 - \sigma_k^2) \nabla_{x_k} \log \tilde{p}_\theta(y | x_k). \quad (14)$$

Which we refer to as the *Guided Euler* (GE) discretisation method. Further details of the unconditional version of this can be found in Song et al. (2021) and Karras et al. (2022).

A.2 PSEUDOCODE

Algorithm 1 GEM Algorithm

Require: $\delta_\theta(x; \sigma)$, σ_k , σ_{k-1} , x_k , α , y

1: Sample $\epsilon_k \sim \mathcal{N}(0, I)$

2: $\Theta_k \leftarrow \frac{x_k - \delta_\theta(x_k, \sigma_k)}{\sigma_k^2}$

3: $x_{k-1} \leftarrow x_k + (\sigma_{k-1}^2 - \sigma_k^2)(\Theta_k + \nabla_{x_k} \log \tilde{p}_\theta(y | x_k)) + \sqrt{\sigma_{k-1}^2 - \sigma_k^2} \epsilon_k$

4: **return** x_{k-1}

Algorithm 2 GE Algorithm

Require: $\delta_\theta(x; \sigma)$, σ_k , σ_{k-1} , x_k , α , y

1: $\Theta_k \leftarrow \frac{x_k - \delta_\theta(x_k, \sigma_k)}{\sigma_k^2}$

2: $x_{k-1} \leftarrow x_k + \frac{1}{2}(\sigma_{k-1}^2 - \sigma_k^2)(\Theta_k + \nabla_{x_k} \log \tilde{p}_\theta(y | x_k))$

3: **return** x_{k-1}

A.3 SYNTHETIC DATA GENERATION DETAILS

All synthetic data are generated by numerically solving the transient ARD equation in an axisymmetric cylindrical reactor of length $L = 1.0$ and radius $A = 0.01$. The spatial domain $(r, z) \in [0, A] \times [0, L]$ is discretised on a uniform grid with $N_r = N_z = 64$. The axial velocity field is prescribed as a Poiseuille profile $u(r) = 2U_{\text{avg}}(1 - (r/A)^2)$.

For each simulation, the mean velocity U_{avg} is sampled uniformly from $[0.01, 0.5]$ m/s, the molecular diffusivity D is sampled log-uniformly via $\log_{10} D \sim \mathcal{U}(-6.5, -4.5)$, and inlet concentrations are sampled as $C_{\text{NO}, \text{in}}, C_{\text{O}_3, \text{in}} \sim \mathcal{U}(0, 0.05)$ and $C_{\text{NO}_2, \text{in}} \sim \mathcal{U}(0, 0.02)$ mol/m³. The reaction rate constant κ is varied by sampling temperature $T \sim \mathcal{U}(270, 320)$ K and computing $\kappa = \kappa(T)$ using a standard Arrhenius-type expression.

Time integration is performed using an operator-splitting scheme consisting of an explicit upwind advection step in the axial direction, second-order diffusion in the axial and radial directions, and a pointwise update for the bimolecular reaction kinetics. Symmetry conditions are enforced at $r = 0$, no-flux boundary conditions at $r = A$, and Dirichlet conditions at the inlet $z = 0$. Each trajectory is simulated up to $\tau_T = 30$ seconds, with 32 logarithmically spaced snapshots recorded between $\tau_{\min} = 5 \times 10^{-2}$ s and τ_T , including the initial condition at $\tau_0 = 0$.

Using this procedure, we generate 2000 independent simulations, yielding 64,000 three-channel concentration fields of shape $(3, 64, 64)$. Of these, 62,000 samples are used for training the diffusion model (a U-Net architecture (Ronneberger et al., 2015)) and the remainder were used for evaluation.

A.4 MODEL TRAINING AND SAMPLING DETAILS

As the diffusion prior, we use build upon the Huang et al. (2024) repository which itself builds upon the Karras et al. (2022) repository. A batch size of 64 was used when training the diffusion prior and a total of 0.5 million gradient evaluations were taken over the dataset during training.

When sampling, for each method we use $K = 400$ discretised steps with the SMC method using 8 samples and the ODE only using 1 as the only randomness is the initial noise the samples are created from and no criteria to choose the individual best sample. $\sigma_{\text{obs}} = 5$ and $\sigma_{\text{pde}} = 1$ were used for the variance parameters for all simulations.

The average run time for simulating all the trajectories was $2.659 \times 10^3 \pm 3.89\text{s} \approx 45\text{mins}$ for SMC method and $1.686 \times 10^3 \pm 6.17\text{s} \approx 28\text{mins}$ for the ODE method. We note that the SMC method does incur an extra inference time cost but produces superior results to the ODE method.

B ABLATION RESULTS

Table 3 gives the results of an ablation study for differing numbers of observations inside the reactor cylinder. As expected generally the error decreases as we increase the number of observations from the sensors however, we still report good RMSE and MAE values for lower numbers of observations ($N_{\text{obs}} = 25 \approx 0.6\%$ of pixels in the data sample).

Table 3: Ablation simulation results for differing numbers of observations inside the reactor when using the GEM proposal with an SMC framework. Mean \pm standard deviation are reported across metrics.

N_{obs}	Species	GT Final	Pred Final	RMSE	MAE
25	NO	1.89×10^{-2}	$1.90 \times 10^{-2} \pm 2.48 \times 10^{-5}$	$1.95 \times 10^{-4} \pm 6.80 \times 10^{-5}$	$1.02 \times 10^{-4} \pm 2.24 \times 10^{-5}$
	O3	2.24×10^{-9}	9.89×10^{-5} $\pm 8.82 \times 10^{-5}$	4.80×10^{-5} $\pm 1.22 \times 10^{-5}$	3.16×10^{-5} $\pm 6.89 \times 10^{-6}$
	NO2	3.79×10^{-2}	$3.76 \times 10^{-2} \pm 3.83 \times 10^{-4}$	$4.55 \times 10^{-4} \pm 1.55 \times 10^{-4}$	$2.14 \times 10^{-4} \pm 5.30 \times 10^{-5}$
50	NO	1.89×10^{-2}	$1.89 \times 10^{-2} \pm 8.47 \times 10^{-5}$	$1.81 \times 10^{-4} \pm 3.25 \times 10^{-5}$	$9.24 \times 10^{-5} \pm 1.56 \times 10^{-5}$
	O3	2.24×10^{-9}	$1.24 \times 10^{-4} \pm 2.11 \times 10^{-5}$	$5.53 \times 10^{-5} \pm 5.30 \times 10^{-6}$	$4.05 \times 10^{-5} \pm 4.34 \times 10^{-6}$
	NO2	3.79×10^{-2}	$3.77 \times 10^{-2} \pm 9.47 \times 10^{-5}$	$3.46 \times 10^{-4} \pm 5.85 \times 10^{-5}$	$1.65 \times 10^{-4} \pm 1.95 \times 10^{-5}$
100	NO	1.89×10^{-2}	$1.90 \times 10^{-2} \pm 3.06 \times 10^{-4}$	$1.67 \times 10^{-4} \pm 2.61 \times 10^{-5}$	$8.96 \times 10^{-5} \pm 1.23 \times 10^{-5}$
	O3	2.24×10^{-9}	$1.13 \times 10^{-4} \pm 6.00 \times 10^{-5}$	$7.19 \times 10^{-5} \pm 1.33 \times 10^{-5}$	$5.50 \times 10^{-5} \pm 7.90 \times 10^{-6}$
	NO2	3.79×10^{-2}	3.78×10^{-2} $\pm 4.39 \times 10^{-4}$	$2.62 \times 10^{-4} \pm 7.94 \times 10^{-5}$	$1.14 \times 10^{-4} \pm 2.90 \times 10^{-5}$
200	NO	1.89×10^{-2}	$1.89 \times 10^{-2} \pm 8.05 \times 10^{-6}$	$7.91 \times 10^{-5} \pm 1.39 \times 10^{-5}$	$5.75 \times 10^{-5} \pm 1.04 \times 10^{-5}$
	O3	2.24×10^{-9}	$1.11 \times 10^{-4} \pm 6.94 \times 10^{-5}$	$7.43 \times 10^{-5} \pm 1.07 \times 10^{-5}$	$5.85 \times 10^{-5} \pm 2.89 \times 10^{-6}$
	NO2	3.79×10^{-2}	$3.78 \times 10^{-2} \pm 6.98 \times 10^{-5}$	$1.00 \times 10^{-4} \pm 8.26 \times 10^{-6}$	$6.41 \times 10^{-5} \pm 2.63 \times 10^{-6}$
300	NO	1.89×10^{-2}	1.89×10^{-2} $\pm 1.43 \times 10^{-4}$	5.85×10^{-5} $\pm 8.63 \times 10^{-6}$	4.57×10^{-5} $\pm 5.92 \times 10^{-6}$
	O3	2.24×10^{-9}	$1.43 \times 10^{-4} \pm 4.47 \times 10^{-5}$	$7.48 \times 10^{-5} \pm 4.68 \times 10^{-6}$	$5.98 \times 10^{-5} \pm 1.91 \times 10^{-6}$
	NO2	3.79×10^{-2}	$3.77 \times 10^{-2} \pm 1.23 \times 10^{-4}$	6.84×10^{-5} $\pm 1.21 \times 10^{-5}$	5.20×10^{-5} $\pm 5.67 \times 10^{-6}$

C TIME SERIES GRAPHS

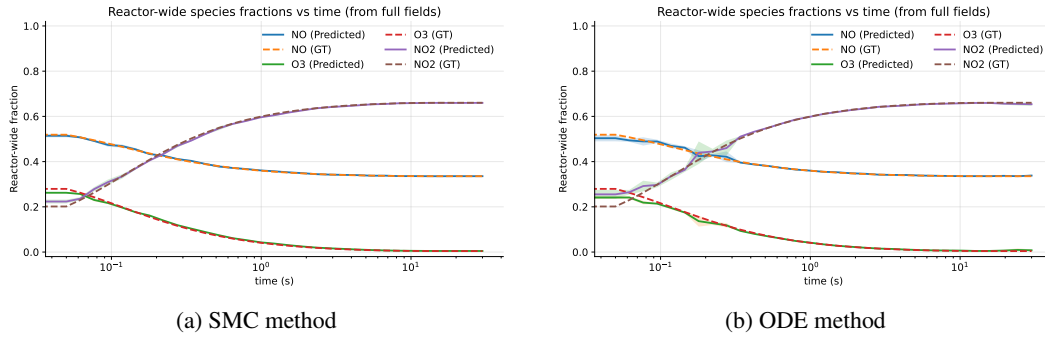
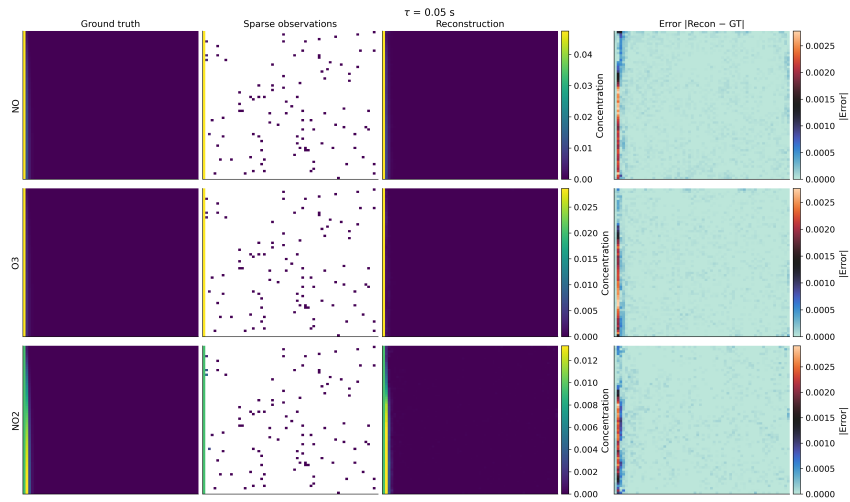
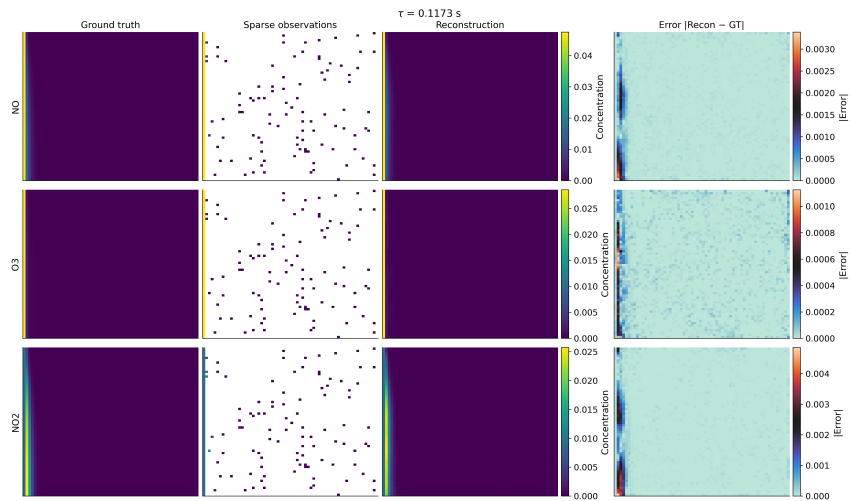


Figure 3: Change in the fraction of species present in the reactor over time.



(a) $\tau = 0.05$



(b) $\tau = 0.1173$

Figure 4: Reconstruction fields and errors at different time points.

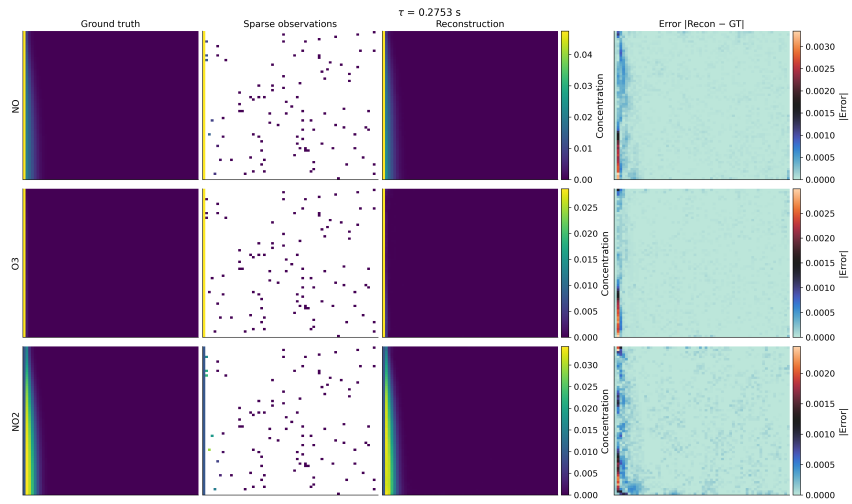
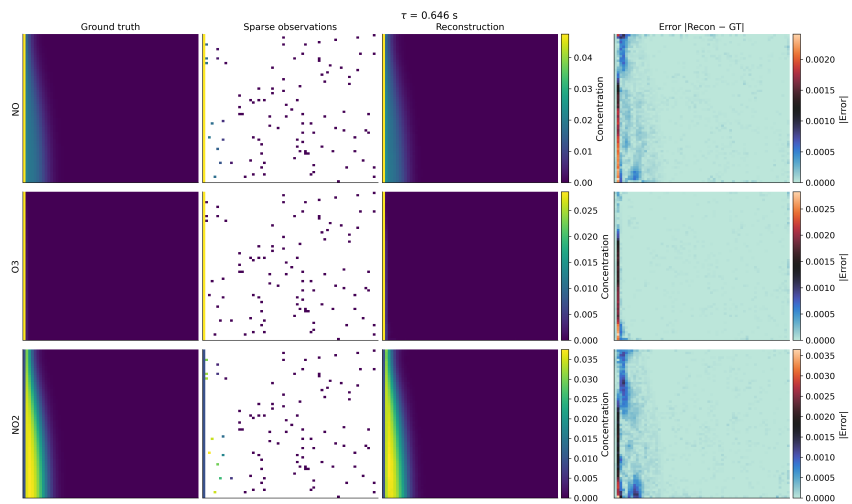
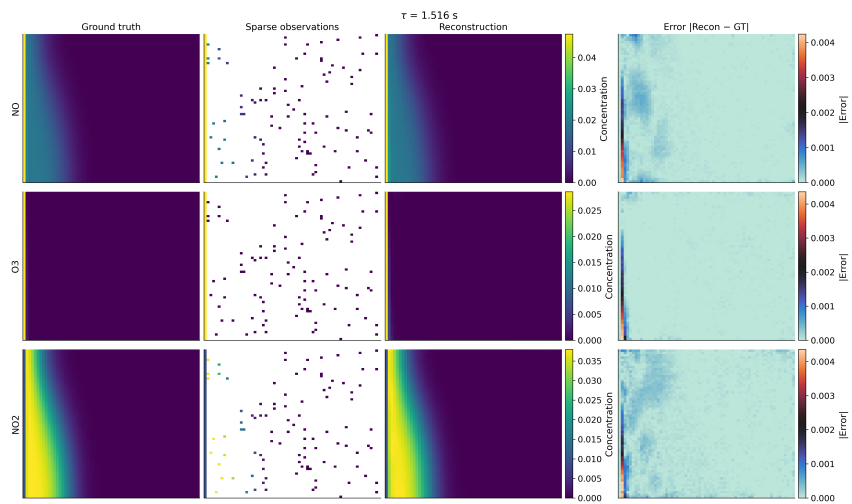
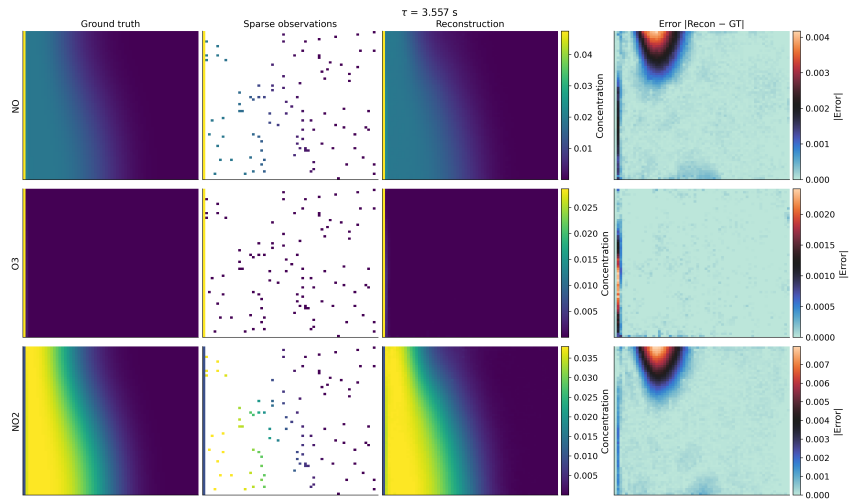
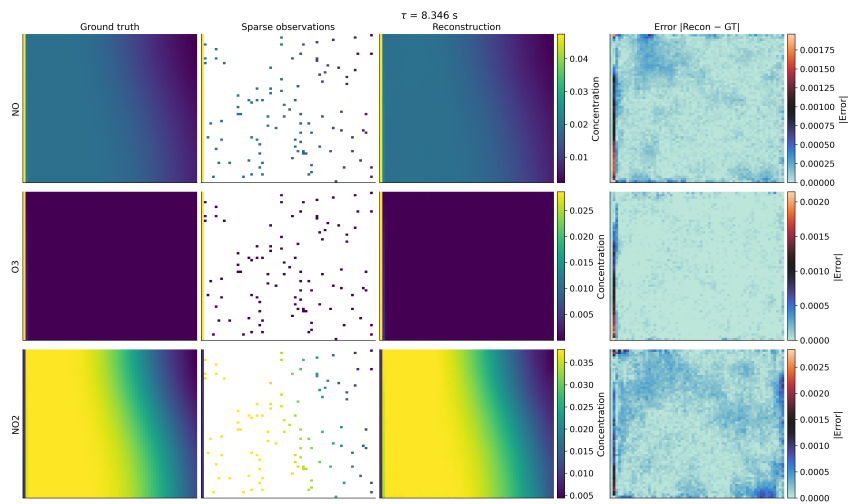
(c) $\tau = 0.2753$ (d) $\tau = 0.646$ (e) $\tau = 1.516$

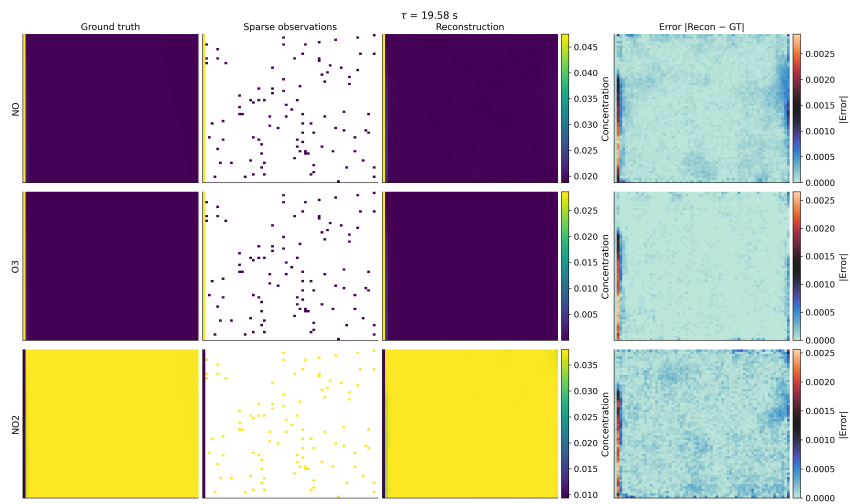
Figure 4: Reconstruction fields and errors at different time points (continued).



(f) $\tau = 3.557$



(g) $\tau = 8.346$



(h) $\tau = 19.58$

Figure 4: Reconstruction fields and errors at different time points (continued).