

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CFO: LEARNING CONTINUOUS-TIME PDE DYNAM- ICS VIA FLOW-MATCHED NEURAL OPERATORS

Anonymous authors

Paper under double-blind review

ABSTRACT

Neural operator surrogates for time-dependent partial differential equations (PDEs) conventionally employ autoregressive prediction schemes, which accumulate error over long rollouts and require uniform temporal discretization. We introduce the **Continuous Flow Operator (CFO)**, a framework that learns continuous-time PDE dynamics without the computational burden of standard continuous approaches, e.g., neural ODE. The key insight is repurposing flow matching to directly learn the right-hand side of PDEs without backpropagating through ODE solvers. CFO fits temporal splines to trajectory data, using finite-difference estimates of time derivatives at knots to construct probability paths whose velocities closely approximate the true PDE dynamics. A neural operator is then trained via flow matching to predict these analytic velocity fields. This approach is inherently time-resolution invariant: training accepts trajectories sampled on arbitrary, non-uniform time grids while inference queries solutions at any temporal resolution through ODE integration. Across four benchmarks (Lorenz, 1D Burgers, 2D diffusion-reaction, 2D shallow water), CFO demonstrates superior long-horizon stability and remarkable data efficiency. CFO trained on only 25% of irregularly subsampled time points outperforms autoregressive baselines trained on complete data, with relative error reductions up to 87%. Despite requiring numerical integration at inference, CFO achieves competitive efficiency, outperforming autoregressive baselines using only 50% of their function evaluations, while uniquely enabling reverse-time inference and arbitrary temporal querying.

1 INTRODUCTION

Time-dependent partial differential equations (PDEs) are fundamental to modeling dynamical phenomena across the physical, biological, and engineering sciences. Neural PDE solvers have emerged as a promising paradigm for learning mappings between function spaces and fast simulation of complex PDEs (Azizzadenesheli et al., 2024). However, current approaches face critical limitations: autoregressive (AR) methods suffer from error accumulation over long horizons and require uniform time grids (Sanchez-Gonzalez et al., 2020); spatio-temporal methods scale poorly with the space-time volume and struggle with causality (Wang et al., 2021); continuous approaches like Neural ODEs (Chen et al., 2018) require expensive backpropagation through ODE solvers.

In this study, we propose the **Continuous Flow Operator** (CFO), a continuous-time neural framework that circumvents these limitations through a key insight: re-purposing flow matching (Lipman et al., 2022; Albergo et al., 2023), typically used for generative modeling, to learn continuous-time dynamics without backpropagating through ODE solvers. We directly learn the right-hand side dynamics of a PDE system by matching it with the analytic velocity of spline-based interpolants over trajectories. In doing so, it can be trained on arbitrary, non-uniform time grids (this allows the framework to fully exploit available temporal data regardless of sampling irregularities) and queried at any temporal resolution during inference, while retaining the training efficiency of discrete methods.

Our contributions are summarized as:

- **A continuous-time framework for learning PDE dynamics.** We propose a novel method that directly learns the inherent continuous-time right-hand side (RHS) by training a neural

054 operator to match the analytic velocity of spline-based interpolants, thereby avoiding costly
 055 backpropagation through ODE solvers if we match the interpolants itself.
 056

- 057 • **Time-resolution invariance.** The proposed framework inherently handles irregular,
 058 trajectory-specific time grids during training and supports arbitrary-resolution inference,
 059 which is unattainable with standard autoregressive methods.
- 060 • **State-of-the-art performance with minimal data.** Across four benchmarks, CFO
 061 achieves superior long-horizon stability; models trained on 25% irregular subsamples out-
 062 perform full-data autoregressive baselines, with relative error reductions up to 87%.

063

064 2 RELATED WORK

065

066 **Neural PDE Solvers and Operator Learning.** Neural operator learning has emerged as a powerful
 067 paradigm for solving PDE families by learning mappings between function spaces. Pioneering
 068 frameworks include DeepONet (Lu et al., 2021) with branch-trunk architectures and FNO (Li et al.,
 069 2020a) with spectral convolutions. Recent extensions handle irregular geometries (Geo-FNO (Li
 070 et al., 2023), multipole graph operators (Li et al., 2020b)), multi-scale structures (WNO (Tripura
 071 & Chakraborty, 2022)), improved depth (U-NO (Rahman et al., 2022)), and irregular sampling
 072 (Galerkin Transformer (Cao, 2021), CViT (Wang et al., 2024)). For time-dependent PDEs, au-
 073 toregressive prediction suffers from exposure bias and cumulative error, while treating time as a
 074 spatio-temporal coordinate scales poorly and complicates causality enforcement (Wang et al., 2021).
 075

076 **Neural ODEs for Continuous-time Modeling.** Neural ODEs (Chen et al., 2018) and variants
 077 (Kidger et al., 2020; Dupont et al., 2019) provide continuous-time modeling for PDE systems
 078 (Sholokhov et al., 2023; Wen et al., 2023). While theoretically elegant, these methods require back-
 079 propagation through ODE solvers via the adjoint method, remaining computationally expensive and
 080 memory-intensive even when operating in latent spaces.
 081

082 **Generative Models for PDE Simulation.** Recent generative approaches treat solution fields as
 083 samples from learned distributions. DiffusionPDE (Huang et al., 2024) uses EDM-style diffusion
 084 (Karras et al., 2022) with physics guidance during sampling; CoCoGen (Jacobsen et al., 2025) ex-
 085 tends score-based models (Song et al., 2020) with physically-consistent sampling; PBFM (Baldan
 086 et al., 2025) integrates physics directly into flow matching (Lipman et al., 2022) training via conflict-
 087 free gradient updates. While effective, these methods typically require explicit PDE expressions and
 088 suffer from computational overhead because of multi-step sampling and space-time generation.
 089

090 **Our Approach.** CFO bridges continuous-time modeling with computational efficiency through a
 091 key innovation: using flow-matching objectives to learn continuous-time vector fields directly from
 092 trajectory derivatives, avoiding the costly ODE solver backpropagation required by Neural ODEs.
 093 Unlike generative PDE solvers that need explicit differential operators, CFO implicitly encodes
 094 dynamics through probability paths aligned with PDE evolution, maintaining temporal causality
 095 without requiring governing equations.
 096

097 Three advantages distinguish our framework: (1) *Model agnosticism* – compatible with any neural
 098 operator architecture (e.g., FNO, U-Net(Ronneberger et al., 2015)); (2) *Time-resolution invariance* –
 099 handles irregular, trajectory-specific time grids during training while supporting arbitrary-resolution
 100 inference; (3) *Computational efficiency* – achieves autoregressive training speeds while maintaining
 101 continuous-time benefits.
 102

103 While Zhang et al. (2024) apply flow matching with piecewise linear interpolation to clinical time
 104 series and treat trajectories as generic sequences, CFO instead exploits PDE structure: by fitting
 105 high-order splines with finite-difference derivative estimates at knots, we construct probability paths
 106 whose velocities closely match true PDE dynamics. This physics-aware design keeps the learned
 107 flow in high-probability regions consistent with governing equations, enabling accurate predictions
 even when trained on only 25% of irregularly sampled time points.

108

3 PRELIMINARIES

109

110

Method of Lines. The method of lines (MOL) is a numerical method for solving time-dependent
111 PDEs. The core idea is first to discretize the spatial domain, transforming the PDE into a large
112 system of ODEs in time, which can then be solved using a standard numerical integrator. Formally,
113 after spatial discretization (e.g., on a grid), the PDE Eq. (2) introduced below becomes a system of
114 ODEs:

115
$$\frac{d}{dt}u_h(t) = \mathcal{N}_h(u_h(t)), \quad u_h(0) = u_{0,h},$$
116

117 where $u_h(t)$ is a vector representing the solution on the spatial grid at time t , and \mathcal{N}_h is a finite-
118 dimensional approximation of the spatial operator \mathcal{N} (e.g., via finite differences or spectral meth-
119 ods). However, when \mathcal{N} is unknown or difficult to specify, instead of relying on a fixed, predefined
120 discretization \mathcal{N}_h , we learn a neural operator \mathcal{N}_θ that directly approximates the continuous-time dy-
121 namics from data. This allows us to define a continuous-time ODE system that can be integrated with
122 arbitrary time steps, inheriting the flexibility of MOL while avoiding the limitations of discrete-time
123 autoregressive models.

124 While MOL provides a conceptual basis for learning a continuous-time operator \mathcal{N}_θ , a naive imple-
125 mentation that matches solution trajectories would require backpropagation through an ODE solver,
126 which is computationally expensive, as seen in Neural ODEs (Chen et al., 2018).

127 Generative modeling frameworks like flow matching (Lipman et al., 2022) and stochastic inter-
128 polants (Albergo et al., 2023) have put forth an efficient alternative. These methods learn a
129 continuous-time vector field by directly regressing it against the velocity of a predefined probability
130 path connecting two distributions. This approach avoids costly ODE integration during training, a
131 key feature we leverage in CFO. Flow matching is a special case of stochastic interpolants with no
132 noise term. We briefly review these concepts below.

133

Stochastic Interpolants and Flow Matching. Stochastic interpolants (Albergo et al., 2023) and
134 flow matching (Lipman et al., 2022) provide a framework for learning a continuous-time vector field
135 that transports one probability distribution to another. The core idea is to define a probability path
136 p_t that connects a source distribution p_0 to a target p_1 . A stochastic interpolant is typically defined
137 as $I_t = s(t, X_0, X_1) + \gamma(t)Z$, where X_0, X_1 are drawn from a probability measure $q(x_0, x_1)$,
138 $Z \sim \mathcal{N}(0, I)$ is independent noise, and the terms satisfying boundary conditions:

139
$$s(0, x_0, x_1) = x_0, \quad s(1, x_0, x_1) = x_1, \quad \gamma(0) = \gamma(1) = 0.$$
140

141 The key insight is that the vector field $v(t, x) = \mathbb{E}[\frac{d}{dt}I_t | I_t = x]$ that generates this path can be
142 learned directly by regressing a neural network $v_\theta(t, x)$ against the analytic velocity $\frac{d}{dt}I_t$ of the
143 predefined process:

144
$$\mathcal{L}(\theta) = \mathbb{E}_{t, x_0, x_1, z} \left[\|v_\theta(t, I_t) - \frac{d}{dt}I_t\|^2 \right]. \quad (1)$$
145

146 This regression objective avoids costly ODE integration during training. During inference, the
147 learned vector field v_θ defines an ODE, $\frac{d}{dt}X_t = v_\theta(t, X_t)$, which can be solved with a numerical
148 integrator to map samples from p_0 to p_1 . CFO adapts this principle by constructing a more sophisti-
149 cated probability path based on splines that incorporate derivative information and interpolate entire
150 solution trajectories, rather than only matching the endpoints.

151

4 CONTINUOUS FLOW OPERATOR

152

153

4.1 METHODOLOGY

154

155 Consider a time-dependent PDE on a domain $[0, T] \times \Omega \subset \mathbb{R} \times \mathbb{R}^d$:

156
$$\partial_t u(t, x) = \mathcal{N}(u(t, x)), \quad u(0, x) = u_0(x), \quad x \in \Omega, t \in [0, T], \quad (2)$$
157

158 where \mathcal{N} is an unknown spatial differential operator and $u_0(x)$ is the initial condition.

159

Notation. We suppress spatial coordinates and write $u(t) \equiv u(t, \cdot) \in \mathcal{V}$. For sampled times
160 $\mathcal{T} := (t_i)_{i=0}^N$ we write the restriction $u|_{\mathcal{T}} := (u(t_0), \dots, u(t_N))$. When \mathcal{T} is clear from context,
161 we abbreviate the snapshot vector to $\mathbf{u} := u_{\mathcal{T}}$.

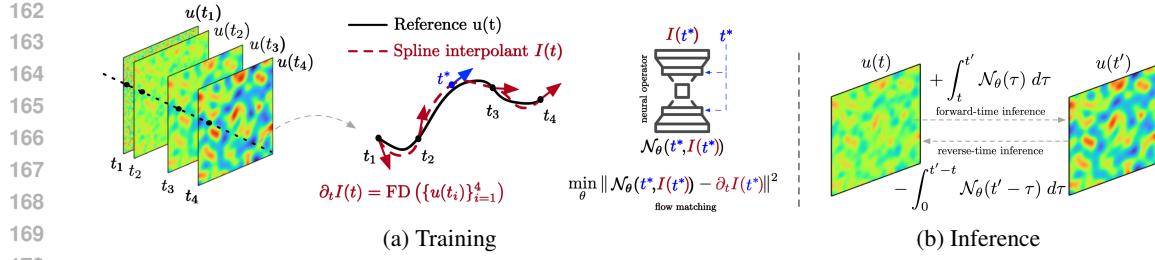


Figure 1: Overview of the Continuous Flow Operator (CFO) framework. (a) **Training: flow matching on a spline path.** For each trajectory with snapshots $\{u(t_i)\}$, we fit a temporal spline $s(t)$ that interpolates the data and matches finite-difference derivative estimates at knots. The spline’s analytic time derivative $\partial_t s(t)$ provides exact velocity targets for training a neural operator \mathcal{N}_θ via the flow matching objective—no ODE integration required during training. (b) **Inference: continuous-time rollout.** The trained operator $\mathcal{N}_\theta(t, u)$ defines a continuous vector field $\dot{u}_\theta = \mathcal{N}_\theta(t, u_\theta)$. Given initial condition $u(t)$, we compute $u(t')$ by numerical integration; reverse-time inference integrates backward.

Problem Statement. Let M trajectories $\{u^j(t)\}_{j=1}^M$ be generated from initial conditions $\{u_0^j\}_{j=1}^M$. For each trajectory j , we observe snapshots at (possibly irregular) times $\mathcal{T}^j := (t_i^j)_{i=0}^{N_j}$ and write $\mathbf{u}^j := u^j|_{\mathcal{T}^j} = (u^j(t_0^j), \dots, u^j(t_{N_j}^j))$. Given the dataset $\{\mathbf{u}^j\}_{j=1}^M$, our goal is to learn an operator that maps from any initial condition u_0 to the solution trajectory $u(t)$ for any $t \in [0, 1]$. For exposition, we drop the superscript j and set the horizon to $T = 1$.

Probability Path Formulation. We define a temporal spline $s(t; \mathbf{u})$ that interpolates the snapshots:

$$s(t_i; \mathbf{u}) = u(t_i), \quad i = 0, \dots, N. \quad (3)$$

The (multi-point) stochastic interpolant is

$$I(t; \mathbf{u}) = s(t; \mathbf{u}) + \gamma(t) z, \quad t \in [0, 1], \quad (4)$$

with boundary conditions $\gamma(t_i) = 0$ for $i = 0, \dots, N$ and $z \sim \mathcal{N}(0, I)$, independent of \mathbf{u} . In the MOL one discretizes space to obtain $\mathcal{N}_h \approx \mathcal{N}$ with $\partial_t u(t) = \mathcal{N}_h(u(t))$. When \mathcal{N} is unavailable, we take a data-driven route: (i) estimate $\partial_t u(t_i)$ from \mathbf{u} using finite-difference stencils along time; (ii) choose the spline $s(t; \mathbf{u})$ so that its derivatives at the knots match these estimates. This allows us to control the derivative accuracy via the stencil order and the path smoothness via the spline degree.

Finite Difference Approximation. The order of finite-difference stencils depends on the number of points used (Strikwerda, 2004). Below are examples for the first and second time derivatives. Let $d_i \approx \partial_t u(t_i)$ and $a_i \approx \partial_{tt} u(t_i)$ denote derivative estimates at time t_i . Define the steps $h_i := t_{i+1} - t_i$ and $\Delta t := \max_i h_i$.

A *two-point stencil* is defined as $d_i = \frac{u(t_{i+1}) - u(t_i)}{h_i}$, which is first-order accurate: $\partial_t u(t_i) = d_i + O(\Delta t)$. A *three-point stencil* provides second-order accuracy for the first derivative and first-order accuracy for the second derivative on irregular grids. For an interior point t_i ,

$$d(t_i) = -\frac{h_i}{h_{i-1}(h_{i-1} + h_i)} u(t_{i-1}) + \frac{h_i - h_{i-1}}{h_{i-1} h_i} u(t_i) + \frac{h_{i-1}}{h_i(h_{i-1} + h_i)} u(t_{i+1}), \quad (5)$$

$$a(t_i) = \frac{2}{h_{i-1} + h_i} \left(\frac{u(t_{i+1}) - u(t_i)}{h_i} - \frac{u(t_i) - u(t_{i-1})}{h_{i-1}} \right). \quad (6)$$

These yield $\partial_t u(t_i) = d_i + O(\Delta t^2)$ and $\partial_{tt} u(t_i) = a_i + O(\Delta t)$. Endpoint estimates can be obtained with standard one-sided stencils. Higher-order accuracy is achievable with more points (see Appendix A.1.1 for details).

Segment Interpolation. With the derivative estimates at the knots, we construct the spline $s(t; \mathbf{u})$ segment-wise. On each interval $[t_i, t_{i+1}]$, we use a polynomial that matches the values and estimated

216 derivatives at the endpoints. This approach, known as Hermite interpolation, allows us to control
 217 the spline’s smoothness and the accuracy of its derivatives at the knots. For a spline to be C^s
 218 continuous, we must impose conditions on derivatives up to order s at each endpoint, which requires
 219 a polynomial of degree at least $2s + 1$ on each segment.
 220

221 **Representative Cases.** We consider two cases: linear splines (C^0) and quintic splines (C^2).
 222

223 **LINEAR SPLINE.** The simplest choice is a *linear* spline, which only interpolates the values $u(t_i)$
 224 and $u(t_{i+1})$. On each segment $t \in [t_i, t_{i+1}]$, the spline is:

$$225 \quad s(t; \mathbf{u}) = \frac{t_{i+1} - t}{h_i} u(t_i) + \frac{t - t_i}{h_i} u(t_{i+1}). \quad (7)$$

227 The resulting spline is continuous but not differentiable at the knots. Its one-sided derivative at t_i
 228 matches the first-order forward-difference estimate, yielding $\partial_t s(t_i; \mathbf{u}) = \mathcal{N}(u(t_i)) + O(\Delta t)$.
 229

230 **QUINTIC SPLINE.** To incorporate higher-order derivative information, we use a *quintic* spline. On
 231 each segment, this spline matches the values $(u_{t_i}, u_{t_{i+1}})$, first derivatives $(d_{t_i}, d_{t_{i+1}})$, and second
 232 derivatives $(a_{t_i}, a_{t_{i+1}})$ at both endpoints. With the three-point finite difference stencils from Eq. (5),
 233 the spline derivatives at the knots satisfy:

$$234 \quad \partial_t s(t_i; \mathbf{u}) = \mathcal{N}(u(t_i)) + O(\Delta t^2), \quad \partial_{tt} s(t_i; \mathbf{u}) = \partial_t \mathcal{N}(u(t_i)) + O(\Delta t).$$

235 This construction yields a globally C^2 smooth path that more accurately reflects the underlying
 236 dynamics. The explicit form is provided in Appendix A.1.2.
 237

238 Unlike global methods such as natural splines, our segment-wise Hermite interpolation yields
 239 closed-form coefficients. It also allows for flexible choices of basis functions, smoothness, and
 240 derivative accuracy based on the physics prior. Compared to low-order splines, the higher-order
 241 spline here (also with high-order assigned derivatives) benefits from more accurate derivative ap-
 242 proximations and higher smoothness: (i) high-order stencils make the probability flow match the
 243 exact flow of the physical dynamics as closely as possible, (ii) Greater smoothness improves train-
 244 ing stability and long-horizon accuracy. These are shown in Sec. 5.1.

245 In practice, we use quintic splines by default, since they naturally capture acceleration, which ap-
 246 pears in many physical laws (e.g., Newton’s second law, wave equations). We do not go beyond
 247 quintic because of diminishing returns: in CFO, the total error comes from both the spline approx-
 248 imation and the neural operator training, and beyond quintic we observe only minor improvements
 249 in derivative accuracy while the network error dominates (see Sec. A.1.3 for detailed discussion).
 250

251 For the noise term, we can also impose regularity on $\gamma(t)$ with splines. For example, $\gamma(t) =$
 $\gamma_0 \frac{(t-t_k)^m (t_{k+1}-t)^m}{(t_{k+1}-t_k)^{2m}}$, $t_k \leq t < t_{k+1}$ with some constant $\gamma_0 > 0$ ensures C^{m-1} continuity and
 252 vanishing noise at the knots. A detailed study can be found in Sec. 5.2.
 253

254 4.2 TRAINING AND INFERENCE

255 **Training.** We learn a time-dependent neural operator $\mathcal{N}_\theta(t, u(t))$ that matches the velocity of
 256 the spline-based stochastic interpolant $I(t; \mathbf{u})$. Similar to the training objective in stochastic inter-
 257 polants, the training loss is
 258

$$259 \quad \mathcal{L}(\theta) = \mathbb{E}_{\mathbf{u}, t, Z} \left[\left\| \mathcal{N}_\theta(t, I(t; \mathbf{u})) - \partial_t I(t; \mathbf{u}) \right\|^2 \right], \quad (8)$$

260 where $t \sim \text{Unif}[0, 1]$, \mathbf{u} denotes a sampled trajectory, Z is the Gaussian noise independent of \mathbf{u} .
 261 $I(t; \mathbf{u})$ is the corresponding stochastic interpolant, and $\partial_t I(t; \mathbf{u})$ is its analytic spline derivative.
 262 This avoids backpropagating through an ODE solver.
 263

264 **Inference.** After training, \mathcal{N}_θ defines a continuous-time right-hand side. For a new initial condi-
 265 tion u_0 , we predict by integrating
 266

$$267 \quad \frac{d}{dt} u_\theta(t) = \mathcal{N}_\theta(t, u_\theta(t)), \quad u_\theta(0) = u_0, \quad (9)$$

268 with a standard ODE solver; by default, we use 4th-order Runge-Kutta (RK4). See Sec. 5.2 for
 269 details of solver choice and the number of function evaluations (NFE).
 270

270 **Time Resolution Invariance.** CFO is agnostic to time grids: (i) *Training* accepts trajectories with
 271 trajectory-specific and irregular time stamps. (ii) *Inference* queries (9) at arbitrary time resolutions
 272 or schedules (including times unseen in training). This allows mixing dense and sparse sequences
 273 during training and evaluating on any target grid at test time.

275 **Reverse-time Inference.** Unlike autoregressive next-step predictors, our learned time-dependent
 276 vector field induces a locally invertible flow map under standard Lipschitz conditions. This makes
 277 CFO suitable for inverse tasks such as inferring earlier states from a state observed at time t_* . Given
 278 a complete state $u(t_*)$ at any time $t_* \in [0, 1]$, one can recover the state at earlier times $s < t_*$ by
 279 integrating the learned flow $\mathcal{N}_\theta(t, u(t))$ backward:

$$280 \quad u_\theta(s) = u(t_*) + \int_{t_*}^s \mathcal{N}_\theta(\tau, u_\theta(\tau)) d\tau = u(t_*) - \int_0^{t_*-s} \mathcal{N}_\theta(t_* - \tau, u_\theta(t_* - \tau)) d\tau,$$

283 We empirically validate this reverse-time inference capability on the dissipative Burgers' equation
 284 in Section 5.2

286 **Probability-flow Transport.** For a fixed (possibly irregular) time grid, if \mathcal{N}_θ minimizes the loss
 287 function (8), the ODE (9) transports the distribution across the sequence of data marginals at each
 288 grid time. It generalizes the two-marginal result of Albergo et al. (2023) by extending Theorem 2.6
 289 in Albergo et al. (2023) from two to multiple marginals under our spline-based path; see Proposition
 290 A.1-A.2 for details. Empirical results also show that the induced probability flow closely tracks the
 291 true dynamics even with heterogeneous time grids across trajectories, especially for quintic CFO.

293 5 EXPERIMENTS

295 In this section, we evaluate the proposed CFO framework across four canonical benchmarks:
 296 Lorenz, 1D Burgers', 2D diffusion-reaction, and 2D shallow water equations. Our primary focus is
 297 on long-horizon rollouts under irregular subsampling, where we compare CFO against autoregres-
 298 sive baselines [trained with teacher forcing](#). In addition, we perform ablation studies to assess the
 299 impact of the noise schedule, inference efficiency, long-horizon temporal extrapolation, and back-
 300 bone architectures. All time series are normalized to the range $[0, 1]$. More results are shared in the
 301 Appendix A.3.

303 5.1 MAIN RESULTS

304 We evaluate CFO's time-resolution invariance by training on irregularly subsampled time grids and
 305 testing on the full-resolution grid. Models are trained using per-trajectory random time grids with
 306 keep rates of 100%, 50%, and 25%. Autoregressive baselines, which require uniform time steps,
 307 are trained only on the full (100%) grid; for autoregressive, we try multiple architectures (including
 308 the same backbone as CFO) and report the best. All models learn the mapping from the initial
 309 condition to the entire trajectory. Table 1 reports the rollout accuracy (relative L^2 error, mean \pm sd).
 310 Notably, Quintic CFO trained on just 25% of irregularly sampled time points still outperforms AR
 311 models trained on full-resolution data, achieving relative error reductions of 24.6%, 78.7%, 87.4%,
 312 and 82.8% on the four benchmarks, respectively. This highlights the CFO's ability to leverage all
 313 available data, regardless of sampling density or regularity, and to generalize effectively to arbitrary
 314 time resolutions.

315 **Lorenz System.** We first consider the Lorenz system (Lorenz, 2017), a 3D chaotic ODE with
 316 spatial-derivative-free dynamics. In this simple setting, the state is a vector $u(t) \in \mathbb{R}^3$ representing
 317 a system of ODEs rather than a PDE with spatial dependence:

$$319 \quad \dot{x} = \sigma(y - x), \quad \dot{y} = x(\rho - z) - y, \quad \dot{z} = xy - \beta z,$$

320 with $(\sigma, \rho, \beta) = (10, 28, 8/3)$. We use an MLP for both CFOs and baseline autoregressive methods.

322 Since the Lorenz vector field is derivative-free, we directly evaluate how well the spline-implied
 323 velocity $s(t; \mathbf{u})$ approximates the true dynamics. We analyze this by regularly downsampling trajec-
 324 tories to resolutions of $2\Delta t$ and $4\Delta t$. Figure 2 plots the average relative L^2 error across a normalized

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
789
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
988
989
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2189
2190
2191
2192
2193

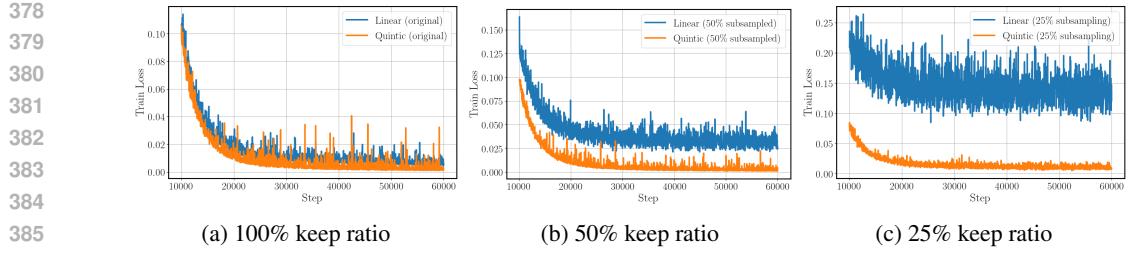


Figure 3: Training comparison between linear and quintic CFO on the Lorenz system. Training loss (10k–60k iterations) for linear and quintic CFO under different keep ratios: (a) 100%, (b) 50%, (c) 25%.

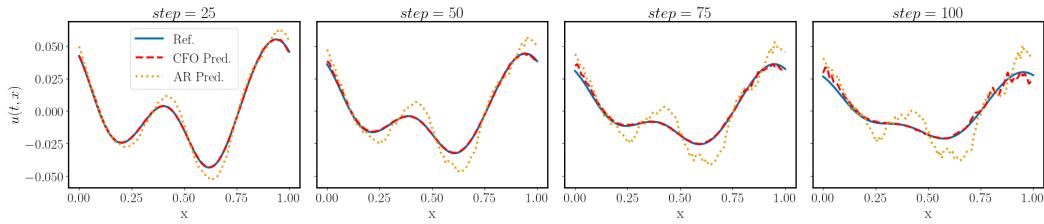


Figure 4: Worst-case prediction for Burgers’ equation. Quintic CFO (trained on 25% irregular data) accurately predicts the full trajectory for the test sample with the highest error. Shown from left to right: solutions at the 25, 50, 75, and 100th time steps.

2D Diffusion-Reaction Equation. We evaluate CFO on the 2D diffusion-reaction (DR) benchmark from PDEBench (Takamoto et al., 2022). We use a 2D U-Net for both CFO and the autoregressive baseline. As shown in Table 1, CFO significantly outperforms the autoregressive model. Quintic CFO trained on just 25% of irregularly sampled time points achieves an 87.4% relative error reduction compared to the autoregressive baseline trained on full data. Figure 5 visualizes a long-horizon rollout, showing that CFO accurately captures the complex pattern evolution, while the autoregressive model quickly diverges.

2D Shallow Water Equation. We evaluate CFO on the 2D shallow-water equations (SWE) from PDEBench (Takamoto et al., 2022). CFO uses a 2D U-Net; the strongest autoregressive baseline uses a Diffusion Transformer (DiT) (Appendix A.3.7). As shown in Table 1, quintic CFO (with 25% data keep rate) attains 82.8% lower relative error than the autoregressive model trained on the full dataset. Figure 6 (Appendix A.3.6) compares predictions over time and shows that, even on its worst sample, CFO substantially outperforms the autoregressive model.

5.2 ABLATION STUDY

Impact of ODE solvers and function evaluations. We study how CFO’s inference accuracy and computational cost vary with the ODE solver and the number of function evaluations (NFE). We compare Euler, Heun, and RK4. As expected, higher-order numerical integrators achieve lower error for a fixed NFE. Table 2 fixes the solver to RK4; additional solver sweeps are provided in Tables 8 and 9. CFO’s continuous-time formulation supports flexible solver choice and NFE budgets, unlike autoregressive (AR) rollouts that use fixed steps. For RK4, across all benchmarks we find that: with only 50% of the AR NFE budget, CFO already outperforms AR; increasing the budget to 200% yields further gains and stable high accuracy, while additional NFE up to 400% brings diminishing returns with only negligible improvements. [To connect NFE to actual runtime, we also report wall-clock inference time and accuracy for AR and CFO+RK4 on identical hardware with the same spatial backbone in Table 10 and Figure 8. It confirms that the theoretical reduction in NFE translates directly to wall-clock time savings.](#)

Impact of Noise Schedule. We investigate the impact of the noise schedule on model performance. The noise schedule $\gamma(t) = \gamma_0 \frac{(t-t_k)^m (t_{k+1}-t)^m}{(t_{k+1}-t_k)^{2m}}$ for $t \in [t_k, t_{k+1}]$, where m controls smooth-

432 Table 2: Influence of NFE on quinticCFO (RK4 solver) accuracy (Relative L^2 Error) at the final
 433 time step, compared to the autoregressive baseline.

NFE (% AR)	Lorenz	Burgers'	DR	SWE
AR (100%)	0.1481	0.0647	0.3850	0.1048
50%	0.0919	0.0090	0.0876	0.0698
100%	0.0655	0.0089	0.0718	0.0154
200%	0.0674	0.0090	0.0608	0.0060
400%	0.0645	0.0089	0.0593	0.0061

441 Table 3: Comparison of CFO with other baselines: Neural ODE (Chen et al., 2018) and PDE-Refiner
 442 (Lippe et al., 2023).

Equation	Method	Rel. L^2 Error	Training Time (s/batch)
Lorenz	Neural ODE	0.101	0.133
	CFO	0.0453	0.00350
Burgers	Neural ODE	0.0275	3.38
	CFO	0.00589	0.00920
DR	PDE-Refiner	0.125	1.38
	CFO	0.044	0.40
SWE	PDE-Refiner	0.093	1.40
	CFO	0.005	0.40

455 ness and γ_0 controls magnitude. Figure 7 shows the impact of noise magnitude γ_0 on the SWE and
 456 DR benchmarks. Performance is stable for small noise levels ($\gamma_0 \in [0, 10^{-4}]$) but degrades as the
 457 noise magnitude increases. For our experiments, we set $m = 3$ to ensure the noise schedule is C^2
 458 continuous, matching the smoothness of the quintic spline.

459 **Comparison with Other Baselines.** Besides AR, we also compare CFO with two other base-
 460 lines. The First is continuous-time neural ODE model (Chen et al., 2018) trained with a teacher-
 461 forcing strategy. CFO attains higher accuracy with substantially less training time (table 3), and
 462 the continuous-time supervision of CFO leads to more stable optimization (Figure 9). Second, we
 463 compare against PDE-Refiner (Lippe et al., 2023), a strong baseline for long-horizon rollouts. CFO
 464 outperforms PDE-Refiner on long-horizon predictions while using fewer computational resources,
 465 since PDE-Refiner requires diffusion-style denoising simulations during training.

466
 467
 468 **Reverse-time Inference.** We study CFO’s performance of backward integration on the dissipative
 469 Burgers’ equation. For Figure 11(a), we fix a terminal time $t_* = 1$ and integrate the flow backward
 470 over varying backward horizons within $[0, t_*]$. The panel reports how the error accumulates with
 471 increasing backward horizon under different noise levels added to the terminal state. Figure 11(b)
 472 visualizes initial-condition recovery by integrating backward from different terminal times to $t =$
 473 0. Despite the ill-posed nature of inverse problems for dissipative PDEs, CFO yields reasonable
 474 reverse-time reconstructions over short backward horizons before errors grow.

475
 476
 477 **Temporal Extrapolation Beyond the Training Horizon.** To assess temporal extrapolation, we
 478 train CFO only on the first half of each trajectory $[0, T/2]$, and evaluate rollouts over the full horizon
 479 $[0, T]$ (Table 11). The error accumulation curves in Fig. 12 show that, for Lorenz, Burgers, and DR,
 480 the extrapolation errors remains nearly identical to full-horizon training results, indicating that CFO
 481 has learned the underlying dynamics rather than memorizing the trajectories. For SWE, the error
 482 increases to 2.28×10^{-2} , but this still remains approximately 4× lower than the full-horizon training
 483 autoregressive baseline.

484
 485 **Neural Network Architectures.** We assess the effect of the backbone on CFO by swapping the
 spatial operator while keeping the training recipe fixed. For 1D Burgers we use an FNO (Li et al.,

486 2020a); for 2D diffusion–reaction and 2D shallow water (SWE) we use the DiT (Peebles & Xie,
 487 2023) Architecture. Results in Table 12 show comparable accuracy for backbones, with only modest
 488 variations. This indicates that CFO is largely architecture-agnostic: temporal continuity is handled
 489 by CFO, while the spatial inductive bias can be chosen to suit resolution, memory, or hardware
 490 constraints.

491

492 6 CONCLUSIONS AND FUTURE WORK

493

494 We introduced CFO, a continuous-time neural operator that learns PDE dynamics by matching the
 495 analytic velocity of spline-based interpolants. By repurposing flow matching to avoid ODE solver
 496 backpropagation during training, CFO achieves a unique combination of capabilities: training on
 497 irregular, trajectory-specific time grids; inference at arbitrary temporal resolutions; and accurate
 498 predictions from severely subsampled data. Across four benchmarks, models trained on only 25%
 499 of irregularly sampled time points outperform fully-trained autoregressive baselines, demonstrating
 500 that physics-aware probability path design can overcome data scarcity.

501

502 Several directions warrant future investigation. *Spatial resolution*: While CFO handles temporal
 503 irregularity, coupling it with mesh-agnostic operators would enable full spatio-temporal resolution
 504 invariance. *Inference acceleration*: Though CFO matches autoregressive efficiency with modest
 505 function evaluations, distilling the learned flow into consistency models (Song et al., 2023) could
 506 enable single-step predictions. *Physics integration*: When governing equations are partially known,
 507 augmenting training with physics-informed constraints (Huang et al., 2024) could improve back-
 508 ward integration stability. *Adaptive splines*: Current fixed-order splines face a trade-off between
 509 computational cost and accuracy; learned splines with curvature regularization (Luo et al., 2025)
 could adapt smoothness to local dynamics.

510

511 CFO demonstrates that continuous-time modeling need not sacrifice computational practicality. By
 512 bridging the efficiency of discrete methods with the flexibility of continuous formulations, it opens
 513 new possibilities for learning from real-world data with irregular temporal measurements, a critical
 514 step toward practical neural PDE solvers.

515

516 LLM USAGE STATEMENT

517

518 During the preparation of this manuscript, we employed ChatGPT-5 as a writing assistant. The
 519 tool was used for proofreading, rephrasing sentences for better flow, and checking for grammatical
 520 consistency. The authors reviewed and edited all AI-generated suggestions to ensure the final text
 521 accurately reflects our research and contributions. The core scientific content remains entirely the
 522 work of the authors.

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540 REPRODUCIBILITY STATEMENT
541

542 To facilitate reproducibility, we provide extensive implementation details, ablation studies, and the-
543 oretical clarifications throughout the main paper and appendix. Section 4 describes the proposed
544 model architecture and training procedure, while Section 5 outlines the experimental setup across all
545 benchmark tasks. All hyperparameters and data preprocessing steps are detailed in Appendix A.3.
546 The theoretical results are included in Appendix A.1–A.2. To ensure full transparency, we will
547 release a public repository with all code, data, and instructions upon acceptance.

548
549 REFERENCES
550

551 Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying
552 framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

553 Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and
554 Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature
555 Reviews Physics*, 6(5):320–328, 2024.

556 Giacomo Baldan, Qiang Liu, Alberto Guardone, and Nils Thuerey. Flow matching meets pdes: A
557 unified framework for physics-constrained generation. *arXiv preprint arXiv:2506.08604*, 2025.

559 Shuhao Cao. Choose a transformer: Fourier or galerkin. *Advances in neural information processing
560 systems*, 34:24924–24940, 2021.

561 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
562 differential equations. *Advances in neural information processing systems*, 31, 2018.

564 Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. *Advances in neural
565 information processing systems*, 32, 2019.

566 Francis Begnaud Hildebrand. *Introduction to numerical analysis*. Courier Corporation, 1987.

568 Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-
569 solving under partial observation. *arXiv preprint arXiv:2406.17763*, 2024.

571 Christian Jacobsen, Yilin Zhuang, and Karthik Duraisamy. Cocogen: Physically consistent and
572 conditioned score-based generative models for forward and inverse problems. *SIAM Journal on
573 Scientific Computing*, 47(2):C399–C425, 2025.

574 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
575 based generative models. *Advances in neural information processing systems*, 35:26565–26577,
576 2022.

577 Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021.

579 Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equa-
580 tions for irregular time series. *Advances in neural information processing systems*, 33:6696–6707,
581 2020.

582 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
583 drew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential
584 equations. *arXiv preprint arXiv:2010.08895*, 2020a.

586 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew Stuart, Kaushik Bhat-
587 tacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differ-
588 ential equations. *Advances in Neural Information Processing Systems*, 33:6755–6766, 2020b.

589 Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural oper-
590 ator with learned deformations for pdes on general geometries. *Journal of Machine Learning
591 Research*, 24(388):1–26, 2023.

593 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

594 Phillip Lippe, Bas Veeling, Paris Perdikaris, Richard Turner, and Johannes Brandstetter. Pde-refiner:
 595 Achieving accurate long rollouts with neural pde solvers. *Advances in Neural Information Pro-*
 596 *cessing Systems*, 36:67398–67433, 2023.

597

598 Edward N Lorenz. Deterministic nonperiodic flow 1. In *Universality in Chaos, 2nd edition*, pp.
 599 367–378. Routledge, 2017.

600 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning
 601 nonlinear operators via deeponet based on the universal approximation theorem of operators.
 602 *Nature machine intelligence*, 3(3):218–229, 2021.

603

604 Yan Luo, Drake Du, Hao Huang, Yi Fang, and Mengyu Wang. Curveflow: Curvature-guided flow
 605 matching for image generation. *arXiv preprint arXiv:2508.15093*, 2025.

606 Harris Abdul Majid. Kuramoto–sivashinsky dataset. [https://huggingface.co/](https://huggingface.co/datasets/hrrsmjd/kuramoto_sivashinsky)
 607 [datasets/hrrsmjd/kuramoto_sivashinsky](https://huggingface.co/datasets/hrrsmjd/kuramoto_sivashinsky), 2024.

608

609 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
 610 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

611 Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-no: U-shaped neural oper-
 612 ators. *arXiv preprint arXiv:2204.11127*, 2022.

613

614 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomed-
 615 ical image segmentation. In *International Conference on Medical image computing and computer-
 616 assisted intervention*, pp. 234–241. Springer, 2015.

617

618 Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter
 619 Battaglia. Learning to simulate complex physics with graph networks. In *International conference
 620 on machine learning*, pp. 8459–8468. PMLR, 2020.

621

622 Aleksei Sholokhov, Yuying Liu, Hassan Mansour, and Saleh Nabi. Physics-informed neural ode
 623 (pinode): embedding physics into models using collocation points. *Scientific Reports*, 13(1):
 10166, 2023.

624

625 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 626 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint
 627 arXiv:2011.13456*, 2020.

628

629 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In Andreas
 630 Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scar-
 631 lett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202
 of *Proceedings of Machine Learning Research*, pp. 32211–32252. PMLR, 23–29 Jul 2023.

632

633 John C Strikwerda. *Finite difference schemes and partial differential equations*. SIAM, 2004.

634

635 Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani,
 636 Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine
 learning. *Advances in Neural Information Processing Systems*, 35:1596–1611, 2022.

637

638 Tapas Tripura and Souvik Chakraborty. Wavelet neural operator: a neural operator for parametric
 639 partial differential equations. *arXiv preprint arXiv:2205.02191*, 2022.

640

641 Ch Tsitouras. Runge–kutta pairs of order 5 (4) satisfying only the first column simplifying assump-
 642 tion. *Computers & mathematics with applications*, 62(2):770–775, 2011.

643

644 Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric par-
 645 tial differential equations with physics-informed deeponets. *Science advances*, 7(40):eabi8605,
 2021.

646

647 Sifan Wang, Jacob H Seidman, Shyam Sankaran, Hanwen Wang, George J Pappas, and Paris
 Perdikaris. Cvit: Continuous vision transformer for operator learning. *arXiv preprint
 648 arXiv:2405.13998*, 2024.

648 Tianshu Wen, Kookjin Lee, and Youngsoo Choi. Reduced-order modeling for parameterized pdes
649 via implicit neural representations. *arXiv preprint arXiv:2311.16410*, 2023.
650

651 Xi Nicole Zhang, Yuan Pu, Yuki Kawamura, Andrew Loza, Yoshua Bengio, Dennis Shung, and
652 Alexander Tong. Trajectory flow matching with applications to clinical time series modelling.
653 *Advances in Neural Information Processing Systems*, 37:107198–107224, 2024.

654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 **A APPENDIX**
703704 **A.1 SPLINE CONSTRUCTION**
705706 This section details the two steps of our spline construction: (i) estimating derivatives at the knots
707 and (ii) piecewise interpolation between knots. We provide a general finite-difference framework for
708 derivative estimation on nonuniform grids, and we give the closed-form quintic Hermite interpolant
709 on each interval.710 The results follow standard numerical references; see, e.g., (Strikwerda, 2004; Hildebrand, 1987).
711712 **A.1.1 FINITE DIFFERENCE STENCIL ON A NONUNIFORM GRID**
713714 Let $f \in C^{m+1}(I)$ on an interval $I \subset \mathbb{R}$ and let $x_0, x_1, \dots, x_m \in I$ be pairwise distinct with point
715 of interest x_0 . Set $\Delta_j := x_j - x_0$ and $h := \max_{0 \leq j \leq m} |\Delta_j|$. Fix $k \in \{0, 1, \dots, m\}$.

716 Define the Vandermonde matrix

717
$$P := [P_{rj}]_{r,j=0}^m, \quad P_{rj} = (x_j - x_0)^r = \Delta_j^r.$$

718

719 Let $e_k \in \mathbb{R}^{m+1}$ denote the k -th standard basis vector. We choose weights $w = (w_0, \dots, w_m)^\top$ by
720 solving

721
$$P w = k! e_k,$$

722 Then the $(m+1)$ -point finite-difference approximation of $f^{(k)}(x_0)$ is
723

724
$$D^{(k)} f(x_0) := \sum_{j=0}^m w_j f(x_j),$$

725
726

727 and the error is controlled by

728
$$f^{(k)}(x_0) - D^{(k)} f(x_0) = \mathcal{O}(h^{m+1-k}).$$

729

730 **A.1.2 QUINTIC HERMITE INTERPOLATION**
731732 Suppose $f \in C^6([t, t'])$ the value, first-, and second-derivatives are s, d, a and s', d', a' at endpoints
733 respectively. Set $h := t' - t$ and $\tau := \frac{x-t}{h} \in [0, 1]$. The unique quintic Hermite interpolant $H(x)$
734 that matches values, first-, and second-derivatives at both endpoints can be written analytically on
735 $[t, t']$ as
736

737
$$H(x) = s H_{00}(\tau) + (d h) H_{10}(\tau) + (a h^2) H_{20}(\tau)$$

738
$$+ s' H_{01}(\tau) + (d' h) H_{11}(\tau) + (a' h^2) H_{21}(\tau),$$

739 where the quintic Hermite basis polynomials are

740
$$H_{00}(\tau) = 1 - 10\tau^3 + 15\tau^4 - 6\tau^5, \quad H_{10}(\tau) = \tau - 6\tau^3 + 8\tau^4 - 3\tau^5,$$

741
742
$$H_{20}(\tau) = \frac{1}{2}(\tau^2 - 3\tau^3 + 3\tau^4 - \tau^5), \quad H_{01}(\tau) = 10\tau^3 - 15\tau^4 + 6\tau^5,$$

743
744
$$H_{11}(\tau) = -4\tau^3 + 7\tau^4 - 3\tau^5, \quad H_{21}(\tau) = \frac{1}{2}(\tau^3 - 2\tau^4 + \tau^5).$$

745 These satisfy the endpoint conditions

746
$$H_{00}(0) = 1, H'_{10}(0) = 1, H''_{20}(0) = 1, \quad H_{01}(1) = 1, H'_{11}(1) = 1, H''_{21}(1) = 1,$$

747 and all other values/derivatives up to order 2 vanish at the opposite endpoint.
748750 **A.1.3 SPLINE SELECTION**
751752 **Spline fitting cost.** The cost of spline fitting is negligible. The finite-difference approximation
753 and spline construction involve only simple algebraic operations (no iterative solvers). They are
754 performed once per trajectory as a preprocessing step and reused throughout training. The computa-
755 tional complexity is $\mathcal{O}(\text{spline order} \times \# \text{time segments} \times \# \text{spatial grid points})$ per trajectory, which
is negligible compared to neural operator backpropagation.

756 **Practical guidelines for choosing splines.**
757

758 1. Given the low cost of spline fitting, we use quintic rather than linear splines as the default,
759 since quintic splines capture acceleration (second derivatives), which is central to many
760 physical laws (e.g., Newton’s second law, wave equations). At the same time, we do not go
761 beyond quintic because higher orders bring diminishing returns in accuracy.

762 2. Linear splines are adequate when (i) time sampling is very dense (e.g., 100% keep rate), or
763 (ii) the underlying dynamics are known to be non-smooth (e.g., shocks or discontinuities),
764 where enforcing higher-order global smoothness may be undesirable.

765 3. Adaptive spline selection is an exciting direction for future work. A dynamic approach
766 could involve learning the optimal spline parameters jointly with the operator, potentially
767 incorporating physics-informed priors to adapt the smoothness.

768
769 **A.2 PROBABILITY-FLOW TRANSPORT**
770

771 This section states the assumption and provides the proof of the probability-flow transport of CFO.
772 It is a direct extension of (Albergo et al., 2023).

773 **Assumption A.1.** *Let $\mathbf{u} = (u(t_0), \dots, u(t_N)) \sim q_{\mathcal{T}}$ on $(\mathbb{R}^d)^{N+1}$. Denote the marginal at t_k by q_k (a density on \mathbb{R}^d). Let $t_0 = 0, t_N = 1$. Assume the following:*

774 **1. (Data marginals)** *For each $k = 0, \dots, N$, $q_k \in C^2(\mathbb{R}^d)$ is strictly positive and has finite Fisher
775 information:*

$$776 \int_{\mathbb{R}^d} \|\nabla \log q_k(x)\|^2 q_k(x) dx < \infty,$$

777 **2. (Spline in time)** *For each fixed \mathbf{u} , $s(\cdot; \mathbf{u})$ is C^2 on $[0, 1]$; it satisfies $s(t_k; \mathbf{u}) = u(t_k)$. Moreover,
778 there exist finite constants M_1, M_2 such that for all $t \in [0, 1]$,*

$$779 \mathbb{E}[\|\partial_t s(t; \mathbf{u})\|^4] \leq M_1, \quad \mathbb{E}[\|\partial_{tt} s(t; \mathbf{u})\|^2] \leq M_2, \quad (10)$$

780 **3. (Noise schedule)** $\gamma \in C^2([0, 1])$ with $\gamma(t_k) = 0$ for all k , $\gamma(t) > 0$ for $t \notin \mathcal{T}$.

781 **Proposition A.1.** *(Extension of Theorem 2.6 in Albergo et al. (2023))*

782 *Under the definitions and Assumption A.1, the random variable $I(t; \mathbf{u}) = s(t; \mathbf{u}) + \gamma(t)Z$ admits a
783 strictly positive (Lebesgue) density $\rho(t, \cdot)$ for every $t \in [0, 1]$. Moreover:*

784 1. $\rho(t_k, x) = q_k(x)$ for $k = 0, \dots, N$.

785 2. *The continuity equation*

$$786 \partial_t \rho(t, x) + \nabla \cdot (\rho(t, x) v^*(t, x)) = 0 \quad \text{for } t \in \bigcup_{k=0}^{N-1} (t_k, t_{k+1}),$$

787 where the velocity field is

$$788 v^*(t, x) := \mathbb{E}[\partial_t I(t; \mathbf{u}) | I(t; \mathbf{u}) = x].$$

789 3. *Consequently, the flow induced by the ODE $\dot{u}(t) = v^*(t, u(t))$ transports the initial law q_0
790 to the knot marginals q_1, \dots, q_N .*

791 *Proof.* Let $g(t, k)$ be the characteristic function of $\rho(t)$. From the independence, we have

$$792 g(t, k) := \mathbb{E}\left[e^{ik \cdot I(t; \mathbf{u})}\right] = \mathbb{E}\left[e^{ik \cdot s(t; \mathbf{u})}\right] \cdot \mathbb{E}\left[e^{ik \cdot \gamma(t)Z}\right] = \mathbb{E}\left[e^{ik \cdot s(t; \mathbf{u})}\right] e^{-\frac{1}{2}\gamma^2(t)|k|^2}.$$

793 From the independence of $s(t; \mathbf{u})$ and Z , the law of $I(t; \mathbf{u})$ is the convolution

$$794 \rho(t, \cdot) = \varphi_{\gamma(t)} * \mu_t,$$

810 where $\mu_t := \text{Law}(s(t; \mathbf{u}))$ is the law of the spline at time t , and $\varphi_\gamma = \mathcal{N}(0, \gamma^2 I_d)$ density.
 811

812 For $t \notin \mathcal{T}$ we have $\gamma(t) > 0$, so $\rho(t, \cdot)$ is smooth and strictly positive by Gaussian convolution.
 813

814 At a knot t_k , since $\gamma(t_k) = 0$ and $s(t_k; \mathbf{u}) = u(t_k)$, we have $\rho(t_k) = \mu_{t_k} = \text{Law}(u(t_k)) = q_k$.
 815

816 For any $t \notin \mathcal{T}$, Under the boundness of spline derivatives (10), we have
 817

$$\begin{aligned} |\partial_t g(t, k)|^2 &= \left| \mathbb{E}[(ik \cdot \partial_t I(t, \mathbf{u}) - \gamma'(t)\gamma(t)|k|^2)e^{ik \cdot I(t; \mathbf{u})}] \right|^2 e^{-\gamma^2(t)|k|^2} \\ &\leq 2 \left(|k|^2 \mathbb{E}[|\partial_t I(t; \mathbf{u})|^2] + |\gamma(t)\gamma'(t)|^2 |k|^4 \right) e^{-\gamma^2(t)|k|^2} \\ &\leq 2 \left(|k|^2 M_1 + 4|\gamma(t)\gamma'(t)|^2 |k|^4 \right) e^{-\gamma^2(t)|k|^2}, \end{aligned}$$

818 we exchange ∂_t and \mathbb{E} by dominated convergence theorem. Hence, for each $p \in \mathbb{N}$,
 819

$$\int_{\mathbb{R}^d} \|k\|^p (|g(t, k)| + |\partial_t g(t, k)|) dk < \infty.$$

820 Denote
 821

$$m(t, k) := \mathbb{E} \left[(\partial_t s(t; \mathbf{u}) + \gamma'(t)z) e^{ik \cdot I(t; \mathbf{u})}, \right] \quad t \notin \mathcal{T},$$

822 thus, $\partial_t g(t, k) = ik \cdot m(t, k)$.
 823

$$\begin{aligned} m(t, k) &= \int_{\mathbb{R}^d} \mathbb{E} \left((\partial_t s(t; \mathbf{u}) + \gamma'(t)z) e^{ik \cdot I(t; \mathbf{u})} | I(t; \mathbf{u}) = x \right) \rho(t, x) dx \\ &= \int_{\mathbb{R}^d} e^{ik \cdot x} \mathbb{E} \left((\partial_t s(t; \mathbf{u}) + \gamma'(t)z) | I(t; \mathbf{u}) = x \right) \rho(t, x) dx \\ &= \int_{\mathbb{R}^d} e^{ik \cdot x} v^*(t, x) \rho(t, x) dx. \end{aligned}$$

830 Thus,
 831

$$\partial_t g(t, k) = ik \cdot \int e^{ik \cdot x} v^*(t, x) \rho(t, x) dx.$$

832 Taking the inverse Fourier transform yields the continuity equation
 833

$$\partial_t \rho(t, x) + \nabla \cdot (v^*(t, x) \rho(t, x)) = 0,$$

834 which holds on each open sub-interval (t_k, t_{k+1}) . As t approaches a knot t_k , we have $\gamma(t) \rightarrow 0$, and
 835 the Gaussian kernel in the convolution deforms into a Dirac delta. This ensures that the density $\rho(t)$
 836 is continuous in time across the knots, with $\lim_{t \rightarrow t_k} \rho(t) = q_k$. Consequently, the flow induced by
 837 the ODE $\dot{u}(t) = v^*(t, u(t))$ transports the initial law q_0 to the knot marginals q_1, \dots, q_N . \square
 838

839 **Proposition A.2.** *The conditional mean velocity $v^*(t, x) := \mathbb{E}[\partial_t I(t; \mathbf{u}) | I(t; \mathbf{u}) = x]$ is the unique
 840 minimizer of*

$$\mathcal{L}(v) = \int_0^T \mathbb{E} \left[\|v(t, x) - \partial_t I(t; \mathbf{u})\|^2 \right] dt, \quad (11)$$

841 *Proof.*
 842

$$\begin{aligned} \mathbb{E} \left[\|v(t, x) - \partial_t I(t; \mathbf{u})\|^2 \right] &= \mathbb{E} \left[\|v(t, x) - \mathbb{E}[\partial_t I(t; \mathbf{u}) | I(t; \mathbf{u}) = x]\|^2 \right] \\ &\quad + \mathbb{E} \left[\|\mathbb{E}[\partial_t I(t; \mathbf{u}) | I(t; \mathbf{u}) = x] - \partial_t I(t; \mathbf{u})\|^2 \right] \\ &\geq \mathbb{E} \left[\|\mathbb{E}[\partial_t I(t; \mathbf{u}) | I(t; \mathbf{u}) = x] - \partial_t I(t; \mathbf{u})\|^2 \right]. \end{aligned}$$

843 The equality holds if and only if $v(t, x) = \mathbb{E}[\partial_t I(t; \mathbf{u}) | I(t; \mathbf{u}) = x]$. \square
 844

864 A.3 EXPERIMENTAL DETAILS
865866 A.3.1 METRICS
867868 We evaluate model accuracy using the relative L^2 norm between the predicted u_{pred} and ground-truth
869 u_{true} trajectories, averaged over all test trajectories:
870

871
$$\text{Relative } L^2 = \frac{1}{N_{\text{test}}} \sum_{k=1}^{N_{\text{test}}} \frac{\left(\sum_{j=1}^{N_t} \sum_{i=1}^{N_x} \left\| u_{\text{pred}}^{(k)}(t_j, x_i) - u_{\text{true}}^{(k)}(t_j, x_i) \right\|^2 \right)^{1/2}}{\left(\sum_{j=1}^{N_t} \sum_{i=1}^{N_x} \left\| u_{\text{true}}^{(k)}(t_j, x_i) \right\|^2 \right)^{1/2}}.$$

872
873
874
875

876 A.3.2 NEURAL NETWORK ARCHITECTURES
877878 We evaluate Multi-Layer Perceptrons (MLPs), U-Nets (Ronneberger et al., 2015), Diffusion Trans-
879 formers (DiTs) (Peebles & Xie, 2023), and Fourier Neural Operators (FNOs) (Li et al., 2023). CFO
880 models receive sinusoidal time embeddings; autoregressive baselines do not.
881882 **MLP** A 5-layer MLP with hidden dims [128, 256, 256, 256, 128] and ReLU activations; time-
883 embedding dimension 16 ($\approx 0.202\text{M}$ parameters).
884885 **FNO** A 1D Fourier Neural Operator with width 128, 5 spectral blocks, and $m = 6$ retained Fourier
886 modes. Each block applies an rFFT, learned complex-mode multiplication with a 1×1 residual conv,
887 and GELU. Inputs are augmented with a normalized grid and lifted to width 128; a two-layer head
888 maps to one output channel ($\approx 0.591\text{M}$ parameters).
889890 **1D U-Net** A minimal 1D U-Net with two scales (`channel_mult = [64, 128]`) using Conv–
891 GroupNorm–Swish blocks, 4×1 stride-2 downsampling, transposed-convolution upsampling, and a
892 1×1 head; time-embedding dimension 64 ($\approx 0.602\text{M}$ parameters).
893894 **2D U-Net** A 2D U-Net with three scales (`channel_mult = [64, 128, 256]`), Conv–GroupNorm–
895 Swish blocks, 4×4 stride-2 downsampling, transposed-convolution upsampling, and a 1×1 head;
896 time-embedding dimension 256 ($\approx 8.1\text{M}$ parameters).
897898 **DiT** A DiT with 8×8 patch embedding, hidden size 384, depth 4, 8 attention heads, and MLP ratio
899 6. Fixed 2D sinusoidal positional embeddings are added to patch tokens; a linear head reconstructs
900 one output channel; time-embedding dimension 256 ($\approx 13.7\text{M}$ parameters).
901
902903 A.3.3 TEACHER-FORCING TRAINING FOR BASELINE MODELS
904905 We summarize the teacher-forcing schemes used to train the autoregressive and neural ODE base-
906 lines.
907908 **Autoregressive baselines.** Let $\{u(t_i)\}_{i=0}^N$ denote a trajectory sampled on a uniform time grid. The
909 AR model F_ϕ is trained as a one-step predictor $\hat{u}(t_{i+1}) = F_\phi(u(t_i))$. At each step the input is the
910 ground-truth state $u(t_i)$ rather than the model prediction (teacher forcing). The training loss is the
911 average one-step error
912

913
$$\mathcal{L}_{\text{AR}}(\phi) = \frac{1}{N} \sum_{i=0}^{N-1} \|F_\phi(u(t_i)) - u(t_{i+1})\|_2^2.$$

914

915 At inference time, we roll out autoregressively by feeding the model’s own prediction back as input
916 from ground-truth initial condition $\hat{u}(t_0) = u(t_0)$:
917

918
$$\hat{u}(t_{i+1}) = F_\phi(\hat{u}(t_i)), \quad i = 0, \dots, N-1.$$

918 **Neural ODE baseline.** For the continuous-time baseline, we parameterize the right-hand side as
 919 a neural network $f_\psi(t, u)$ and define the ODE
 920

$$921 \quad \frac{d}{dt} \tilde{u}(t) = f_\psi(t, \tilde{u}(t)).$$

923 We train this model with a teacher-forcing strategy on short time intervals. For each pair (t_i, t_{i+1})
 924 in a trajectory, we initialize the ODE solver at the ground-truth state $\tilde{u}(t_i) = u(t_i)$ and numerically
 925 integrate to obtain $\tilde{u}(t_{i+1}) = u(t_i) + \int_{t_i}^{t_{i+1}} f_\psi(\tau, \tilde{u}(\tau)) d\tau$. The loss is the average mismatch between
 926 the integrated solution and the next ground-truth snapshot:
 927

$$928 \quad \mathcal{L}_{\text{NODE}}(\psi) = \frac{1}{N} \sum_{i=0}^{N-1} \|\tilde{u}(t_{i+1}) - u(t_{i+1})\|_2^2.$$

931 This setup uses ground-truth states as starting points for each integration interval.
 932

933 A.3.4 LORENZ SYSTEM

934 We consider the chaotic ODE system with spatial-derivative-free dynamics:
 935

$$936 \quad \frac{dx}{dt} = \sigma(y - x), \quad \frac{dy}{dt} = x(\rho - z) - y, \quad \frac{dz}{dt} = xy - \beta z, \quad t \in [0, 5],$$

938 with $(\sigma, \rho, \beta) = (10, 28, 8/3)$.
 939

940 **Dataset.** The initial condition (x_0, y_0, z_0) is sampled uniformly from the box $[-5, 5]^3$. The dataset
 941 comprises 9000 training, 500 validation, and 500 testing trajectories. Each trajectory is sampled at
 942 1001 uniformly spaced time points over a horizon of 5.0s.
 943

944 For the irregular sampling experiments, we create subsampled datasets by randomly keeping 50%
 945 or 25% of the time points for each trajectory, with the selection performed independently for each
 946 trajectory. This procedure is consistent across all equations.
 947

948 **Neural Solver Choice.** Both the CFO and the autoregressive baseline use the same MLP architec-
 949 ture. For CFO models, normalized time $t \in [0, 1]$ is encoded using a 16-band sinusoidal embedding.
 950

951 **Training of CFO and AR.** The model is trained for 200,000 steps using the Adam optimizer with
 952 a batch size of 2048, a learning rate of 10^{-4} , and betas $(\beta_1, \beta_2) = (0.9, 0.99)$, and evaluated every
 953 5000 steps. The noise schedule parameter is set to $\gamma_0 = 10^{-5}$.
 954

955 **Training of Neural ODE.** We use the Tsit5 adaptive Runge–Kutta solver (Tsitouras, 2011) with
 956 absolute tolerance 10^{-6} and relative tolerance 10^{-4} . Backpropagation through the ODE solu-
 957 tion is performed using DiffraX’s `RecursiveCheckpointAdjoint` method (Kidger, 2021) for
 958 memory-efficient gradients. The other training setting matches CFO and AR.
 959

960 Table 4: Endpoint error and empirical order p for the spline-implied velocity on the Lorenz system.
 961 Errors are average relative L^2 at trajectory segment endpoints ($\tau \in \{0, 1\}$) over 50000 segments.
 962 $\Delta t = 0.005$ s.
 963

Spline type	Step size	Error ($\tau=0$)	p ($\tau=0$)	Error ($\tau=1$)	p ($\tau=1$)
Quintic	Δt	5.33×10^{-4}	-	5.32×10^{-4}	-
	$2\Delta t$	2.16×10^{-3}	2.02	2.16×10^{-3}	2.02
	$4\Delta t$	8.63×10^{-3}	2.00	8.63×10^{-3}	2.00
Linear	Δt	2.49×10^{-2}	-	2.48×10^{-2}	-
	$2\Delta t$	4.97×10^{-2}	1.00	4.96×10^{-2}	1.00
	$4\Delta t$	9.92×10^{-2}	1.00	9.88×10^{-2}	1.00

972 A.3.5 BURGERS' EQUATION
973974 Consider the 1D Burgers' equation benchmark investigated in (Wang et al., 2021):
975

976
$$\frac{du}{dt} = \nu \frac{d^2u}{dx^2} - u \frac{du}{dx}, \quad (x, t) \in (0, 1) \times (0, 1],$$

977
$$u(x, 0) = u_0(x), \quad x \in (0, 1),$$

978

979 with the viscosity $\nu = 0.01$.
980981 **Dataset.** The initial condition $u_0(x)$ is generated from a Gaussian Random Field (GRF) \sim
982 $\mathcal{N}(0, 25^2(-\Delta + 5^2 I)^{-4})$ with periodic boundary conditions. The dataset consists of 1000 training,
983 100 validation, and 100 testing trajectories. Each trajectory is sampled at 101 uniformly spaced time
984 points over $[0, 1s]$. The spatial domain $[0, 1]$ is discretized into 100 points.
985986 **Neural Solver Choice.** For CFO, we use the 1D U-Net; results appear in Table 1. To demonstrate
987 flexibility, our ablation study part 5.2 include a FNO variant. The autoregressive baseline mirrors
988 these architectures (MLP, U-Net, FNO) with time embeddings disabled; we report only the best-
989 performing variant—MLP—in Table 1.
990991 Table 5: Relative L_2 Error of autoregressive baselines with different architectures on the 1D Burgers'
992 equation. The MLP architecture, as the best-performing variant, is used for comparison in Table 1.
993

Architecture	MLP	U-Net	FNO
Relative L_2 Error	3.34×10^{-2}	2.56	1.16

998 **Training of CFO and AR.** The model is trained for 60,000 steps using the Adam optimizer with a
999 batch size of 256, a learning rate of 10^{-4} , and betas $(\beta_1, \beta_2) = (0.9, 0.99)$. The model is evaluated
1000 every 5000 steps. The noise schedule parameter is set to $\gamma_0 = 10^{-5}$.
10011003 **Training of Neural ODE.** We use the Tsit5 adaptive Runge–Kutta solver (Tsitouras, 2011) with
1004 absolute tolerance 10^{-6} and relative tolerance 10^{-4} . Backpropagation through the ODE solution
1005 is performed using Diffraix's RecursiveCheckpointAdjoint method (Kidger, 2021) for
1006 memory-efficient gradients. The other training setting matches CFO and AR.
10071008 A.3.6 DIFFUSION-REACTION EQUATION
10091010 We consider the 2D diffusion-reaction equation from PDEBench (Takamoto et al., 2022):
1011

1012
$$\begin{aligned} \partial_t u &= D_u(\partial_{xx}u + \partial_{yy}u) + R_u(u, v), \\ \partial_t v &= D_v(\partial_{xx}v + \partial_{yy}v) + R_v(u, v), \end{aligned}$$

1013

1014 where $u = u(t, x, y)$ is the activator and $v = v(t, x, y)$ denotes the inhibitor on the domain
1015 $(t, x, y) \in (0, 5] \times (-1, 1)^2$. The reaction terms are given by
1016

1017
$$\begin{aligned} R_u(u, v) &= u - u^3 - k - v, \\ R_v(u, v) &= u - v, \end{aligned}$$

1018

1019 with $k = 5 \times 10^{-3}$. The diffusion coefficients are $D_u = 1 \times 10^{-3}$ and $D_v = 5 \times 10^{-3}$.
10201022 **Dataset.** We directly use the dataset provided in PDEBench (Takamoto et al., 2022). The initial
1023 condition is generated as standard normal random noise $u(0, x, y) \sim \mathcal{N}(0, 1.0)$. The dataset con-
1024 sists of 900 training, 50 validation, and 50 testing trajectories of activator and inhibitor (2 channels).
1025 Each trajectory is sampled at 101 uniformly spaced time points over $[0, 5s]$, but the first time step is
truncated to ensure stability. The spatial domain is discretized into a 128×128 grid.
1026

1026 Table 6: Relative L_2 Error of autoregressive baselines with different architectures on the 2D
 1027 Diffusion-Reaction Equation. The U-Net architecture, as the best-performing variant, is used for
 1028 comparison in Table 1.

Architecture	U-Net	DiT
Relative L_2 Error	$(4.23 \pm 0.59) \times 10^{-1}$	$(9.21 \pm 0.10) \times 10^{-1}$

1034 **Neural Solver Choice.** For CFO, we use the 2D U-Net described above; results shown in Table 1.
 1035 To demonstrate flexibility, our ablation study part 5.2 include a DiT variant. The autoregressive
 1036 baseline mirrors these architectures (U-Net, DiT) with time embeddings disabled; we report only
 1037 the best-performing variant—U-Net—in Table 1.

1039 **Training of CFO and AR.** The model is trained for 5000 steps using the Adam optimizer with a
 1040 batch size of 256, a learning rate of 10^{-4} , and betas $(\beta_1, \beta_2) = (0.9, 0.99)$. The model is evaluated
 1041 every 500 steps. The noise schedule parameter is set to $\gamma_0 = 10^{-5}$.

1044 **Training of PDE-Refiner.** For the PDE-Refiner baseline, we follow the training setup from the
 1045 original paper (Lippe et al., 2023) for 2D equations and use the 2D U-Net architecture with a pa-
 1046 rameter count and number of training steps matched to CFO for a fair comparison.

1048 A.3.7 SHALLOW WATER EQUATION

1049 We consider the 2D shallow-water equations from PDEBench (Takamoto et al., 2022):

$$\begin{aligned} \partial_t h + \partial_x(hu) + \partial_y(hv) &= 0, \\ \partial_t(hu) + \partial_x\left(u^2 h + \frac{1}{2} g_r h^2\right) + \partial_y(uvh) &= -g_r h \partial_x b, \\ \partial_t(hv) + \partial_y\left(v^2 h + \frac{1}{2} g_r h^2\right) + \partial_x(uvh) &= -g_r h \partial_y b, \end{aligned}$$

1057 where h is the water depth, (u, v) are the velocities, b is the spatially varying bathymetry, and g_r is
 1058 the gravitational acceleration.

1059 **Dataset.** We use the dataset from PDEBench (Takamoto et al., 2022). The initial condition for
 1060 water depth is generated from a GRF, with initial velocities set to zero. The dataset consists of 900
 1061 training, 50 validation, and 50 testing trajectories of water depth (one channel). Each trajectory is
 1062 sampled at 101 uniformly spaced time points on $[0, 1]$, with the first time step truncated. The spatial
 1063 domain is a 128×128 grid.

1065 **Neural Solver Choice.** For CFO, we use the 2D U-Net; results are shown in Table 1. To demon-
 1066 strate flexibility, our ablation study in Section 5.2 includes a DiT variant. The autoregressive base-
 1067 line mirrors these architectures (U-Net, DiT) with time embeddings disabled; we report only the
 1068 best-performing variant (U-Net) in Table 1 for comparison.

1070 Table 7: Relative L_2 Error of autoregressive baselines with different architectures on the 2D Shallow
 1071 Water Equation. The DiT architecture, as the best-performing variant, is used for comparison in
 1072 Table 1.

Architecture	U-Net	DiT
Relative L_2 Error	3.16×10^{-1}	9.04×10^{-2}

1078 **Training of CFO and AR.** The model is trained for 5000 steps using the Adam optimizer with a
 1079 batch size of 256, a learning rate of 10^{-5} , and betas $(\beta_1, \beta_2) = (0.9, 0.99)$. The model is evaluated
 every 500 steps. The noise schedule parameter is set to $\gamma_0 = 10^{-6}$.

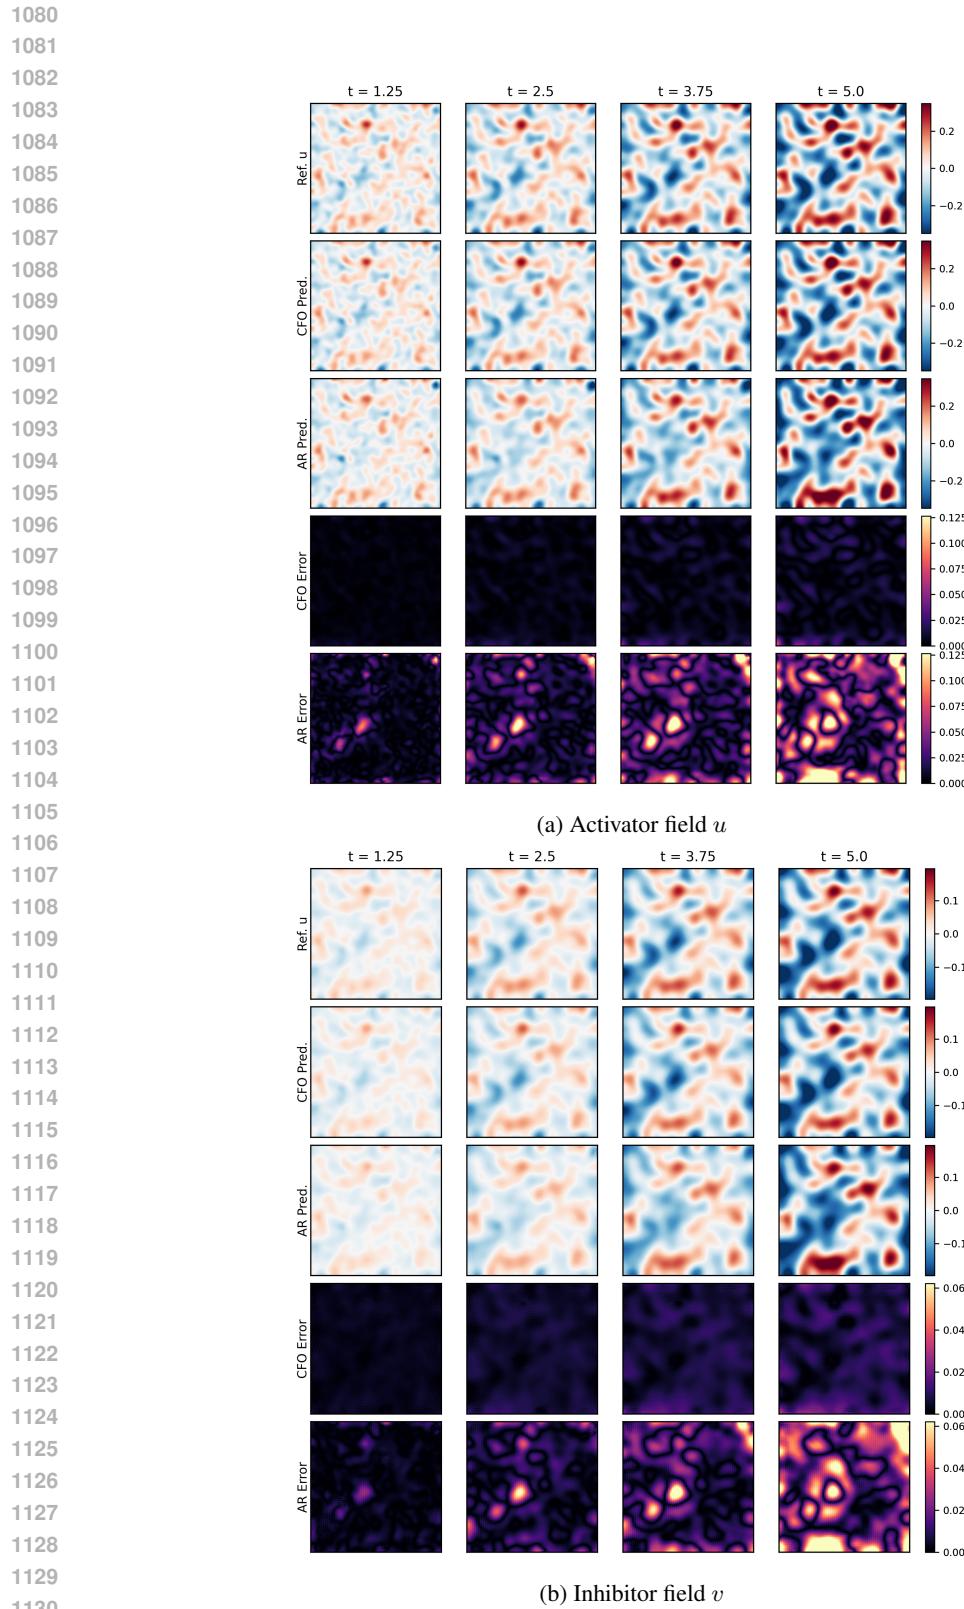


Figure 5: DR trajectory visualization for activator (top) and inhibitor (bottom).

1134
1135
1136
1137
1138
1139

Training of PDE-Refiner. For the PDE-Refiner baseline, we follow the training setup from the original paper (Lippe et al., 2023) for 2D equations and use the 2D U-Net architecture with a parameter count and number of training steps matched to CFO for a fair comparison.

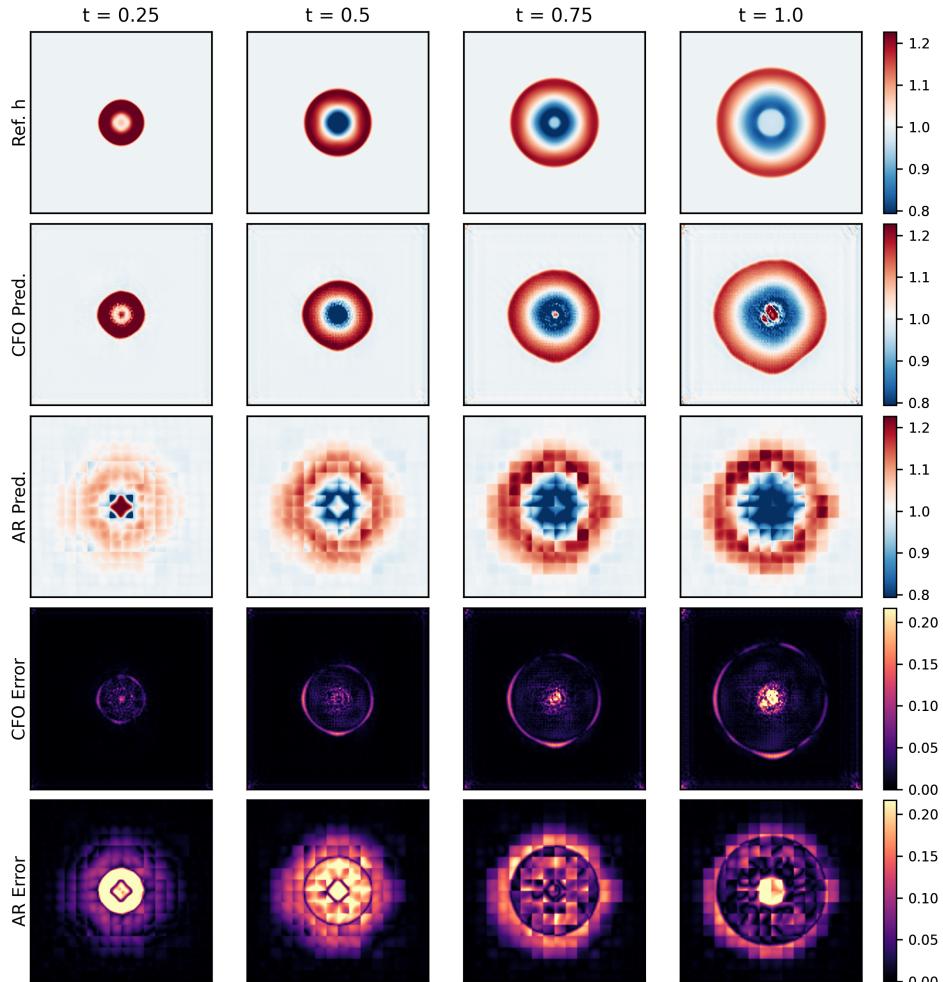
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
11711172
1173
1174
1175

Figure 6: Shallow Water Equation trajectory visualization. Comparison of a long-horizon rollout for the test sample with the highest error. Quintic CFO, trained on 25% irregularly sampled data, captures the dynamics, while the autoregressive baseline (trained on full-resolution data) accumulates significant error.

1176
1177
1178

A.3.8 ABLATION EXPERIMENTS

1179
1180
1181

To complement the ablation study in the main text, we provide additional results and figures in this subsection.

1182
1183
1184
1185
1186
1187

Impact of noise schedule. We sweep the noise magnitude γ_0 . See Figure 7 (SWE on the left; Diffusion–Reaction on the right).

Impact of ODE solvers and function evaluations. We study the tradeoff between accuracy and computation by varying the solver and the number of function evaluations (NFE). Errors generally decrease with NFE, and higher-order solvers achieve lower error at a fixed budget. See the Euler and Heun sweeps in Tables 8 and 9 for final-time relative L_2 errors across tasks. Table 10 compares

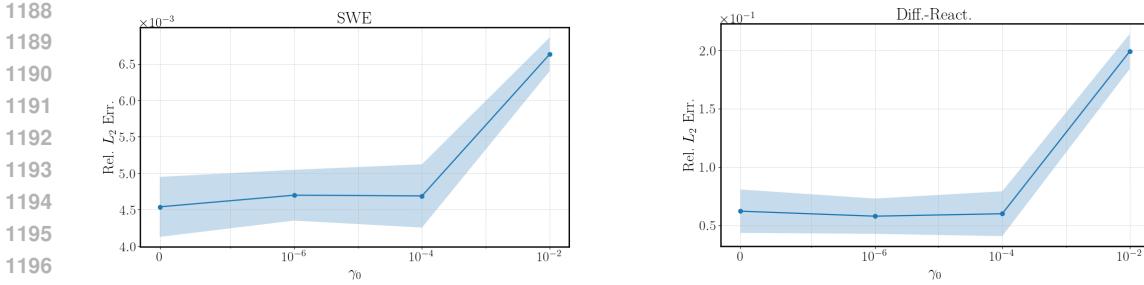


Figure 7: Impact of noise magnitude γ_0 on model performance for SWE (left) and DR (right) benchmarks.

wall-clock inference time and accuracy between the AR baseline and CFO with RK4 solver using 50% of the AR NFE. Figure 8 visualizes the error–runtime across different NFE settings.

Table 8: Influence of number of function evaluations (NFE) on quinticCFO (Euler solver) accuracy (Relative L_2 Error) at the final time step, compared to the autoregressive (AR) baseline.

NFE (% AR)	Lorenz	Burgers'	DR	SWE
AR (100%)	0.1481	0.0647	0.3850	0.1048
50%	0.3032	0.0142	0.0611	0.1345
100%	0.2821	0.0111	0.0603	0.0978
200%	0.2337	0.0099	0.0588	0.0651
400%	0.2076	0.0094	0.0603	0.0389

Table 9: Influence of number of function evaluations (NFE) on quinticCFO (Heun solver) accuracy (Relative L_2 Error) at the final time step, compared to the autoregressive (AR) baseline.

NFE (% AR)	Lorenz	Burgers'	DR	SWE
AR (100%)	0.1481	0.0647	0.3850	0.1048
50%	0.1878	0.0090	0.0694	0.0874
100%	0.1052	0.0090	0.0601	0.0249
200%	0.0705	0.0090	0.0600	0.0076
400%	0.0651	0.0090	0.0594	0.0061

Comparison with other baselines. Figure 9 shows the rescaled training loss for CFO and neural ODE for Lorenz system: CFO converges faster and more stably, especially in the early stage of training.

We also compare to a Conditional flow matching model (with classifier-free guidance) that treats time as a spatial coordinate, without leveraging the ODE structure. Figure 10 visualizes a representative prediction.

Reverse-time inference. We evaluate this the reverse-time inference capabilities of our framework on the dissipative Burgers' equation. Figure 11 visualizes the results.

Temporal extrapolation beyond the training horizon. We evaluate long-horizon temporal extrapolation by training on the first half of each trajectory, $t \in [0, T/2]$, and rolling out over the full horizon, $t \in [0, T]$. Table 11 reports the resulting errors, and Figure 12 shows error accumulation on the four benchmarks, compared to a baseline trained on the full time horizon.

1242	1243	Task	Backbone	Time (ms/sample)		Rel. L^2 Error	
				AR	CFO (50% NFE)	AR	CFO (50% NFE)
1245	Lorenz	MLP	0.16	0.12		1.48×10^{-1}	9.19×10^{-2}
1246	Burgers	1D U-Net	0.68	0.31		1.03×10^1	9.00×10^{-3}
1247	DR	2D U-Net	117.18	49.36		3.74×10^{-1}	8.76×10^{-2}
1248	SWE	2D U-Net	111.83	49.63		3.84×10^{-1}	6.98×10^{-2}

Table 10: Inference runtime (ms/sample) and relative L^2 error at the final time step for the autoregressive baseline and CFO using 50% of the AR NFE, evaluated on a single NVIDIA A6000 GPU (JAX, fp32) with the same spatial backbones.

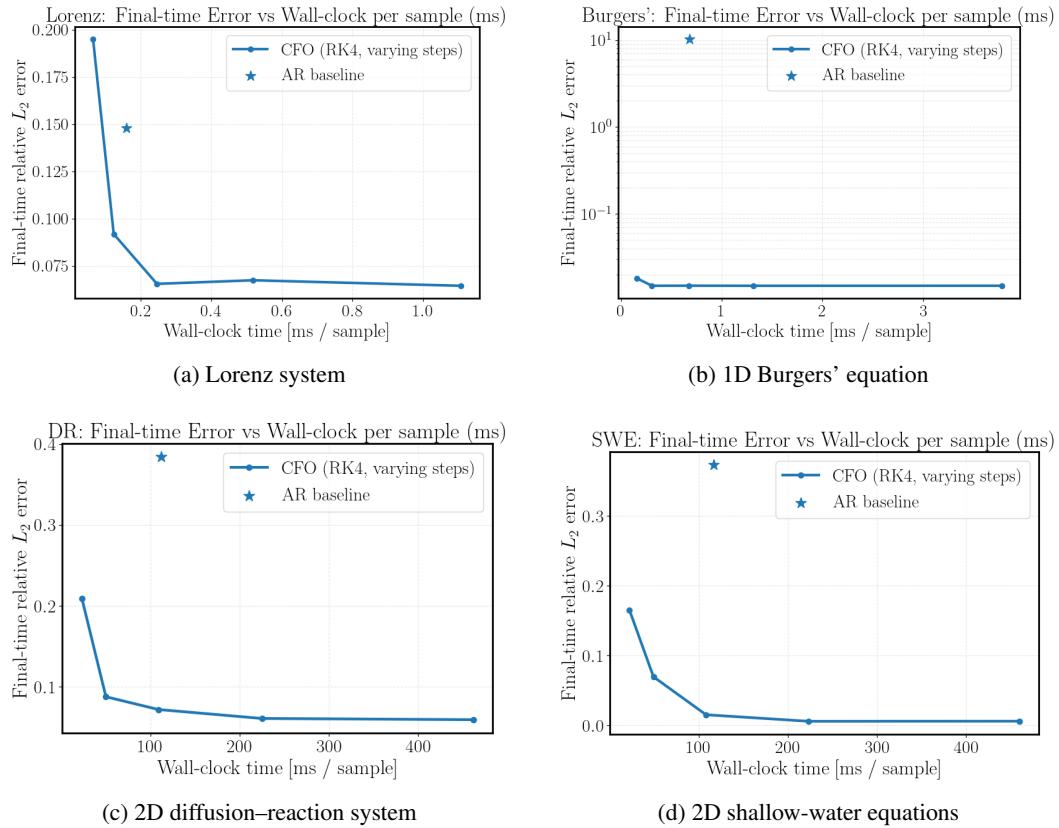


Figure 8: Error versus wall-clock time for CFO with different numbers of function evaluations (NFEs). The star marker denotes the AR baseline.

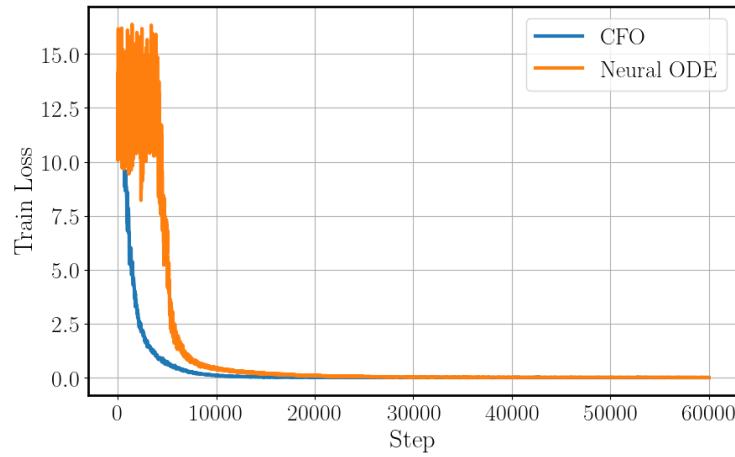
Neural Network Architectures. CFO is largely architecture-agnostic: swapping the spatial backbone (FNO for 1D Burgers; DiT for 2D Diffusion–Reaction and SWE) yields comparable accuracy. See Table 12

More Challenging Examples. We further evaluate CFO on the one-dimensional Kuramoto–Sivashinsky (KS) equation, a fourth-order chaotic nonlinear PDE,

$$u_t + uu_x + u_{xx} + \nu u_{xxxx} = 0.$$

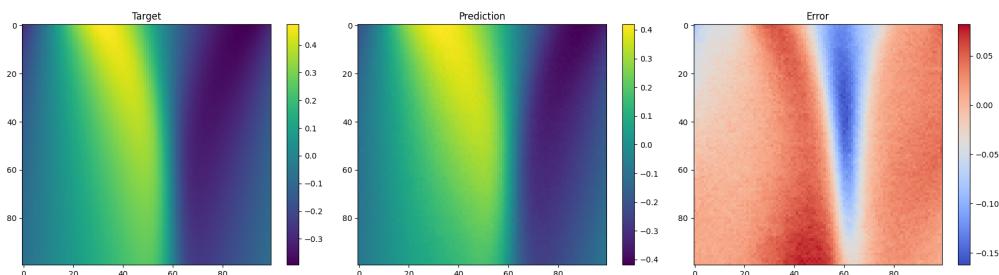
We follow the setting and dataset of (Majid, 2024) (140-step rollout) and use the same neural network architecture (1D U-Net) for all methods. As shown in Table 13, CFO achieves roughly $2\times$ lower error than the neural ODE baseline and $3\times$ lower error than AR, while being vastly more computationally efficient than the neural ODE (comparable training time to AR).

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316



1317
 1318 Figure 9: Training loss versus optimization steps for CFO and the neural ODE baseline on the Lorenz
 1319 system. CFO converges faster and with substantially reduced oscillations, whereas the neural ODE
 1320 exhibits large fluctuations in the early stages before eventually reaching a comparable loss level.
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332

1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341



1342
 1343 Figure 10: Vanilla conditional flow matching model (with classifier-free guidance) on Burgers'
 1344 equation.
 1345
 1346
 1347
 1348
 1349

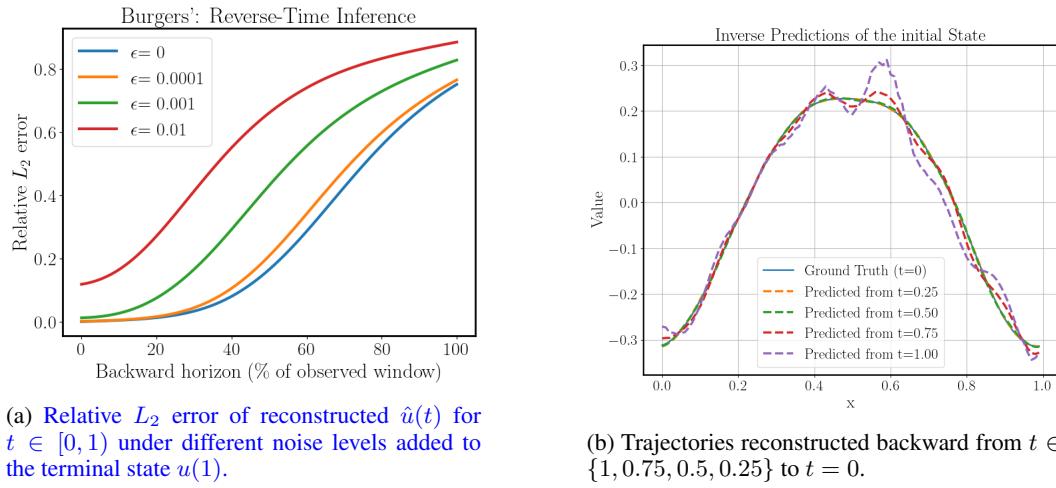


Figure 11: Reverse-time inference for dissipative Burgers' equation.

Equation	Train $[0, T]$	Train $[0, T/2]$, Test $[0, T]$
Lorenz	4.53×10^{-2}	3.67×10^{-2}
Burgers	5.89×10^{-3}	7.94×10^{-3}
DR	4.37×10^{-2}	4.78×10^{-2}
SWE	4.56×10^{-3}	2.28×10^{-2}

Table 11: Relative L_2 error of temporal extrapolation beyond the training horizon.Table 12: CFO performance across (Relative L_2 Error) neural backbones—FNO (1D Burgers) and DiT (2D Diffusion–Reaction, 2D SWE).

Equation	Burgers'	SWE	DR
Rel. L_2 Err.	0.0068	0.0083	0.0821

Latent Space Simulation. In addition to direct application, CFO can be extended to operate in latent spaces. For the PDE in (2), consider an invertible operator \mathcal{E} that maps the function u to a latent representation $\mathcal{E}(u)$.

Define the conjugation of \mathcal{N} as $\tilde{\mathcal{N}} = \mathcal{E}\mathcal{N}\mathcal{E}^{-1}$, then the dynamics in the latent space can be expressed as:

$$\mathcal{E}(u_t) = \tilde{\mathcal{N}}(\mathcal{E}(u)).$$

Specifically, if \mathcal{E} is a time-independent linear operator (e.g., Fourier or wavelet transform), then commuting with the time derivative yields:

$$[\mathcal{E}(u)]_t = \tilde{\mathcal{N}}(\mathcal{E}(u)).$$

In this case, CFO learns the spatial operator $\tilde{\mathcal{N}}$ directly in the latent space. Experiments on the 2D diffusion-reaction and 2D shallow water equations (see Table 14) show that CFO effectively models latent dynamics, achieving performance comparable to direct CFO in the original space.

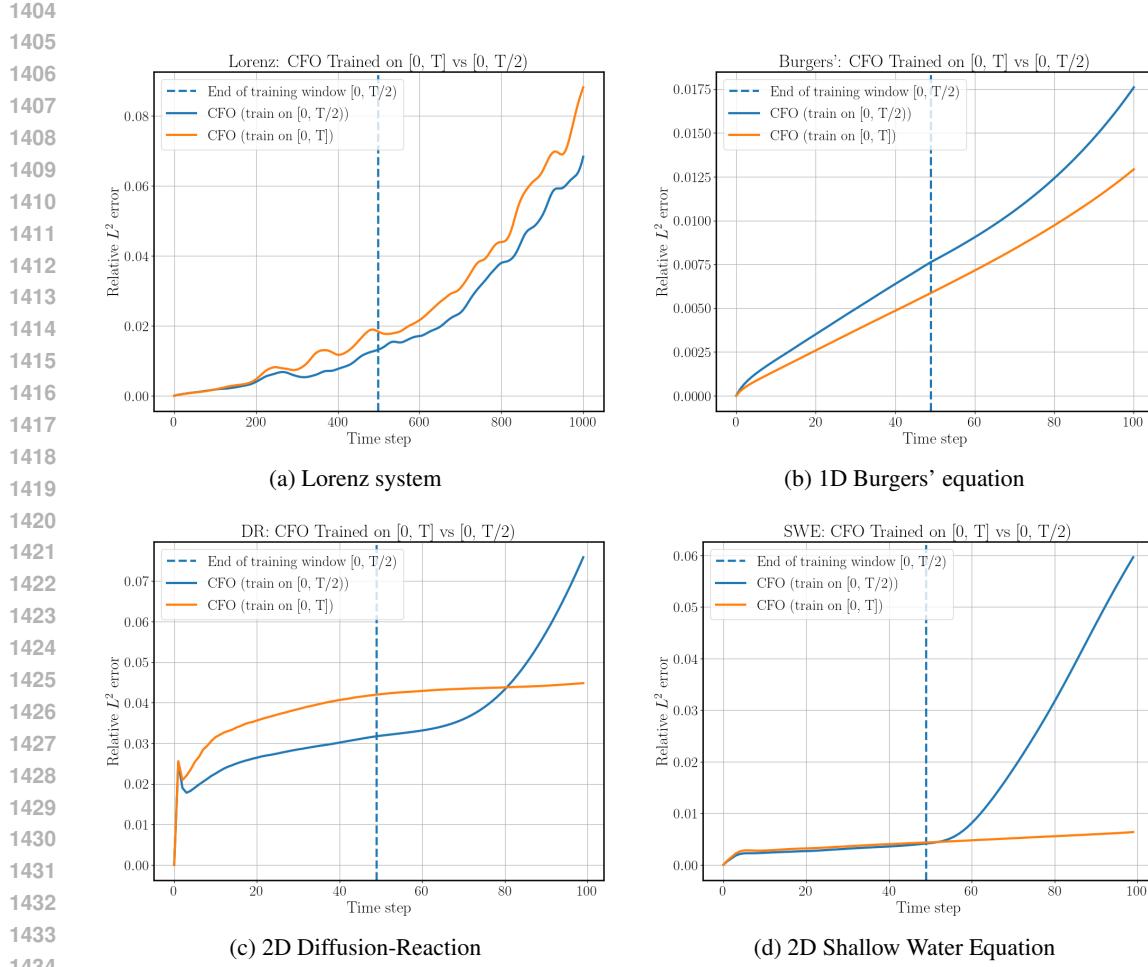


Figure 12: Long-horizon temporal extrapolation: error accumulation on Lorenz, Burgers', DR, and SWE. Models are trained on $t \in [0, T/2]$ and evaluated on rollouts over $t \in [0, T]$.

Method	Rel. L_2 Err.	Training Time
AR	0.317	1 h
Neural ODE	0.197	40 h 53 min
CFO	0.105	1 h 5 min

Table 13: Performance on the 1D Kuramoto–Sivashinsky (KS) equation. CFO attains substantially lower error than AR and Neural ODE while remaining far more computationally efficient than the Neural ODE baseline.

Table 14: CFO in wavelet latent space: Relative L_2 Error for 2D Shallow Water and 2D Diffusion-Reaction equations.

Equation	SWE	DR
Rel. L_2 Err.	0.0053	0.0218