Beyond Equality: Inequality-Constrained Flow Matching for Physically Admissible Solutions

Alfred Wei Lun Leong [1] **Yatao Bian** [1] **Bryan Kian Hsiang Low** [1]

Abstract

Generative models for PDE solutions must produce *physically admissible* samples, yet state-of-the-art physics-constrained samplers enforce only equality constraints ($h(u) = 0$) and ignore the inequality constraints ($g(u) \leq 0$) that equally govern admissibility—e.g., entropy conditions and invariant-interval bounds. We introduce Inequality-aware Physics-Constrained Flow Matching (I-PCFM), a framework that *jointly handles equality and inequality constraints* at inference time without retraining. Within I-PCFM we systematically compare three general-purpose projection strategies grounded in classical constrained optimization—(A) slack-variable reformulation, (B) log-barrier augmentation, and (C) active-set projection—together with (D) Frank-Wolfe guidance, a projection-free strategy with closed-form linear-minimization oracles (LMOs) for linear constraints and a Riemannian extension for nonlinear settings. Across four PDE benchmarks, including the 1D Heat and 1D Reaction-Diffusion equations, we establish three key findings: (1) equality-only projection *degrades* inequality satisfaction by 5–20$\times$ relative to unconstrained generation, (2) among general-purpose strategies, active-set projection attains both the highest joint feasibility (97% on Reaction-Diffusion) and the lowest runtime, and (3) Frank-Wolfe guidance matches active-set projection's feasibility while running up to 40$\times$ faster. The code can be found at `https://anonymous.4open.science/r/ipcfm-icml-workshop/`.

[1]Department of Computer Science, National University of Singapore, Singapore. Correspondence to: Alfred Wei Lun Leong <alfred_leong@u.nus.edu>.

Proceedings of the AI4Physics Workshop at the 43rd International Conference on Machine Learning (AI4Physics@ICML 2026), Seoul, South Korea. 2026. Copyright 2026 by the author(s).

1. Introduction

Generative models for partial differential equations (PDEs) have emerged as a powerful tool for scientific simulation, offering efficient sampling, uncertainty quantification, and the ability to model multimodal solution distributions (Kerrigan et al., 2023; Huang et al., 2024). However, one fundamental requirement in such a setting is that the generated solutions must be *physically admissible*. Physically admissible solutions refer to points that obey the physical constraints of a system, and distinguish realistic physical scenarios from purely mathematical solutions that could violate physical laws. Importantly, physical admissibility is not a soft constraint that can be traded off against other objectives, as a solution that is not physically admissible is of no use from a scientific standpoint, regardless of how well it fits the data distribution.

Many physical admissibility conditions arise naturally as *inequality* constraints. In reaction-diffusion systems, species concentrations must remain within an invariant interval, bounded within biologically valid ranges (Murray & Murray, 2003). In thermodynamics, entropy must be non-decreasing across shock discontinuities (the entropy condition) (LeVeque, 1992). In fluid dynamics, density and pressure must remain strictly positive throughout the domain (Zhang & Shu, 2011; LeVeque, 1992). Such inequality constraints are necessary conditions for a solution to correspond to a physically realizable state of the system.

The recent Physics-Constrained Flow Matching (PCFM) framework (Utkarsh et al., 2025) represents the current state of the art for enforcing *hard* constraints in pretrained flow matching models. PCFM operates at inference time, projecting intermediate flow states onto the constraint manifold through a Gauss-Newton solver, and achieves machine-precision satisfaction of equality constraints of the form $h(u) = 0$. It subsumes prior work such as ECI (Cheng et al., 2025) and substantially outperforms soft-penalty approaches like DiffusionPDE (Huang et al., 2024) and D-Flow (Ben-Hamu et al., 2024). However, PCFM is fundamentally restricted to equality constraints, and extending it to accommodate inequality constraints is an important open problem, as acknowledged by the authors themselves.

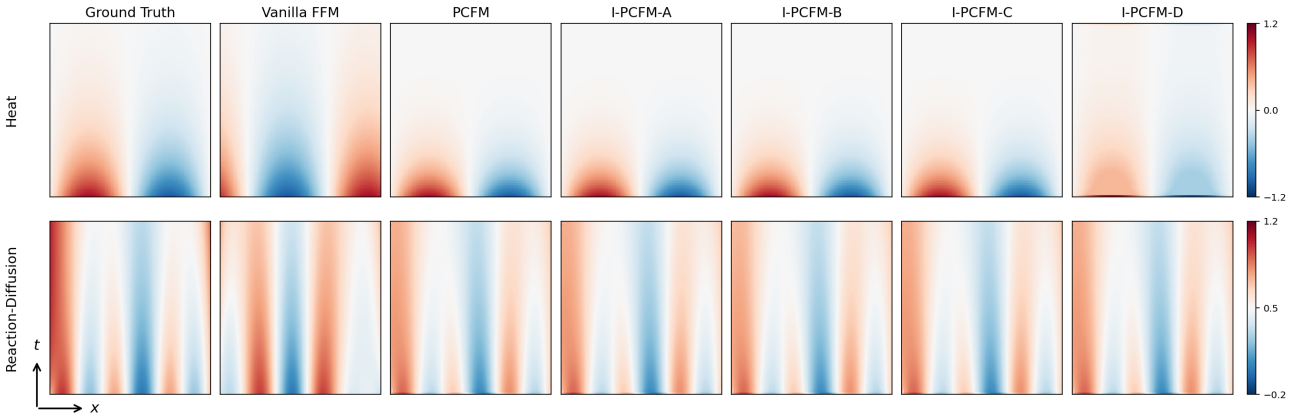

Figure 1. Generated samples for Heat (top) and Reaction-Diffusion (bottom). On Heat, Vanilla-FFM produces reconstructions that deviate noticeably from the ground truth, whereas constrained methods recover it faithfully. The same pattern holds on Reaction-Diffusion.

Table 1. **Constraint satisfaction.** CE-Ineq measures the inequality violation magnitude; Joint-Feas. is the joint-feasibility rate (fraction of samples for which all equality and inequality constraints are satisfied). Best results per benchmark are **bolded**. I-PCFM-C and I-PCFM-D show strong performance on Heat and Reaction-Diffusion, demonstrated by low CE-Ineq and high Joint-Feas. values.

	Heat		**Reaction-Diffusion**	
Method	CE-Ineq ↓	Joint-Feas. ↑	CE-Ineq ↓	Joint-Feas. ↑
Vanilla FFM	7.2×10^{-7}	0.00	0.082	0.00
PCFM	7.1×10^{-4}	0.94	1.628	0.30
I-PCFM-A	1.6×10^{-4}	0.95	0.008	0.40
I-PCFM-B	5.7×10^{-4}	0.94	1.875	0.27
I-PCFM-C	3.6×10^{-5}	**1.00**	0.0005	**0.97**
I-PCFM-D	**0**	**1.00**	$\mathbf{6.2\times10^{-12}}$	**0.97**

Extending PCFM to inequality constraints is non-trivial: equality constraints define a *manifold* $\{u : h(u) = 0\}$ onto which one can project via Gauss-Newton, whereas inequality constraints $g(u) \leq 0$ carve out a *feasible region* that has both an interior and a boundary. A single Newton step that drives all residuals to zero is no longer appropriate; instead, active constraints (at the boundary) and inactive constraints (in the interior) must be distinguished and treated differently.

To address this, we introduce **Inequality-aware Physics-Constrained Flow Matching (I-PCFM)**, a framework that *jointly enforces equality and inequality constraints* at inference time without retraining. We propose and systematically compare three general projection strategies, each grounded in a classical constrained-optimization paradigm: (A) a slack variable reformulation that converts inequalities to equalities; (B) a log-barrier augmentation that penalizes proximity to the constraint boundary; and (C) an active-set projection that identifies and enforces only the binding constraints at each flow step. We further introduce (D) Frank-Wolfe guidance, a projection-free strategy for linear constraint settings that replaces projection with cheap linear minimization oracle (LMO) calls, and handles nonlinear constraints via an interleaved Newton restoration. Across our four PDE benchmarks, our experiments (Table 3) yield three key findings: **(1)** equality-only projection *degrades* inequality satisfaction (5–20× worse than unconstrained generation), **(2)** among the three general-purpose strategies, active-set projection attains both the highest joint feasibility (97% on Reaction-Diffusion) and the lowest runtime, and **(3)** Frank-Wolfe guidance matches active-set projection's feasibility while running up to 40× faster.

Our primary contributions are:

- **Framework.** We propose I-PCFM, a unified inference-time framework that extends PCFM to jointly handle both equality and inequality constraints in pretrained flow matching models for PDE solutions.
- **Systematic comparisons.** We introduce and compare four principled inequality-enforcement strategies—slack-variable reformulation, log-barrier augmentation, active-set projection, and Frank-Wolfe guidance—bridging classical constrained optimization with generative flow matching.
- **Empirical insights.** Across the four PDE benchmarks, we show that equality-only projection actively degrades inequality feasibility, whereas active-set projection emerges as a robust general-purpose strategy that dominates in both feasibility and runtime. Frank-Wolfe guidance further matches active-set feasibility at substantially lower cost.

2. Related Work

Flow Matching. Flow matching (Lipman et al., 2023) is a simulation-free framework for training continuous normalizing flows by regressing vector fields of conditional probabil-

ity paths. Functional Flow Matching (FFM) (Kerrigan et al., 2023) extends flow matching to function spaces, enabling generative modeling of PDE solutions on discretized grids, which is the backbone model underlying both PCFM and our work. Guidance-based approaches (Feng et al., 2025) offer an alternative route for controlled generation by steering the learned velocity field, but do not enforce hard constraints. Our method is agnostic to the specific flow matching variant and complementary to guidance: any pretrained FFM model can be equipped with I-PCFM projection strategies at inference time.

Constraint Enforcement in Generative Models for PDEs. Existing approaches span a spectrum from soft to hard enforcement. *Gradient-based (soft):* DiffusionPDE (Huang et al., 2024) and D-Flow (Ben-Hamu et al., 2024) steer sampling via backpropagation through PDE operators but achieve only approximate satisfaction at high computational cost. *Gradient-free zero-shot:* ECI (Cheng et al., 2025) applies a mixing-based correction that satisfies constraints exactly, yet is restricted to simple linear, non-overlapping constraints. *Hard equality:* PCFM (Utkarsh et al., 2025) generalizes ECI to arbitrary nonlinear equality constraints via Gauss-Newton projection at each flow step, attaining machine-precision satisfaction—but is limited to $h(u) = 0$ and cannot handle $g(u) \leq 0$. *Chance-constrained:* CCFM (Liang et al., 2025) recasts projection as a chance-constrained program and incidentally employs a hinge-residual active set for inequality handling; however, it does not systematically compare alternative inequality strategies. A complementary line of work enforces physics during *training* rather than inference. Physics-Based Flow Matching (PBFM) (Baldan et al., 2026) embeds PDE residuals into the training loss via temporal unrolling, and Physics-Constrained Fine-Tuning (PCFT) (Tauberschmidt et al., 2026) minimizes weak-form PDE residuals in a post-training stage. Both avoid inference-time projection cost but require (re-)training and provide only soft PDE-residual minimization rather than hard constraint satisfaction. Our work fills the gap left by PCFM: we *systematically study* inference-time inequality enforcement in flow matching, comparing four classical paradigms (slack, barrier, active-set, Frank-Wolfe) within a single unified framework.

Constrained Generation in Diffusion and Flow Models. A separate line of work modifies the generative process itself to enforce constraints. Riemannian Score-Based Generative Modelling (Bortoli et al., 2022) defines diffusion on known manifolds, and Riemannian Flow Matching (Chen & Lipman, 2024) extends flow matching to general Riemannian geometries; both require the constraint set to admit a known manifold structure (e.g., spheres, rotation groups). Projected Diffusion Models (Christopher et al., 2024) enforce constraints via per-step projection in image and design domains but do not target PDE-specific physics constraints. CASAL (Blanke et al., 2026) proposes a split augmented Langevin sampler for strictly constrained generation with conservation laws, but requires a primal-dual formulation rather than the lightweight projection framework of PCFM. In contrast, I-PCFM targets arbitrary nonlinear inequality constraints that need not define a smooth manifold, operates entirely at inference time without retraining, and applies post-hoc to any pretrained flow matching model.

Inequality-Constrained Optimization. Our three strategies draw from classical constrained optimization: slack variable reformulations convert inequalities to equalities for reuse of equality solvers (Bertsekas, 1997); log-barrier (interior point) methods add a penalty that diverges at the constraint boundary (Boyd & Vandenberghe, 2004); and active-set methods enforce only the binding constraints as equalities, which is efficient when the active set is sparse (Nocedal & Wright, 2006). While well-studied in optimization, their adaptation to inference-time inequality enforcement in flow-based generative models is, to our knowledge, novel.

Frank-Wolfe Methods. The Frank-Wolfe (FW) algorithm (Frank & Wolfe, 1956) solves constrained convex optimization by replacing projection with a linear minimization oracle (LMO) over the feasible set. More recently, Jaggi (2013) established sublinear convergence rates and highlighted FW's projection-free structure as well-suited to constraint sets with cheap LMOs (polytopes, ℓ_1-balls, trace-norm balls). Subsequent work tightened convergence through away-step and pairwise variants (Lacoste-Julien & Jaggi, 2015) and extended FW to block-coordinate (Lacoste-Julien et al., 2013) and stochastic settings. Of particular relevance to our work, Weber & Sra (2021) introduced Riemannian Frank-Wolfe for optimization over geodesically convex sets on manifolds, combining a tangent-space LMO with a retraction step—the template our Riemannian variant adopts to handle the nonlinear equality manifold of Reaction-Diffusion. To our knowledge, FW has not previously been used as an inference-time feasibility correction in generative modeling; we exploit its projection-free structure and per-iterate feasibility certificate to enforce inequality constraints during constrained flow-matching sampling without the iterative Newton-Schur solvers.

3. I-PCFM: Inequality-aware Physics-Constrained Flow Matching

3.1. Problem Setup

We adopt the same problem setup as Utkarsh et al. (2025). We assume access to a pretrained Functional Flow Matching (FFM) model with learned vector field $v_\theta(u, \tau)$, where $\tau \in [0, 1]$ is the flow time. This model transports samples from

Table 2. Constraint structure across the four PDE benchmarks. ✓ = linear, ✗ = nonlinear. All initial-condition (IC) and inequality constraints are linear; nonlinearity arises only in the dynamics equality for Reaction-Diffusion and Burgers.

Benchmark	Equality constraints	Inequality constraint
Heat	IC ✓, mass conservation ✓	Max-principle box ✓
Navier-Stokes	IC ✓, mass conservation ✓	Data-driven box ✓
Reaction-Diffusion	IC ✓, mass conservation ✗	Invariant interval $[0, 1]$ ✓
Burgers	IC ✓, Godunov-flux unroll ✗	Oleinik condition ✓

a prior $u_0 \sim \pi_0$ to a target distribution of PDE solutions $u_1 \sim \pi_1$. Given a set of k equality constraints $h(u) = 0$ and m inequality constraints $g(u) \leq 0$, our goal is to generate samples that satisfy all constraints exactly at the final state u_1, operating entirely at inference time without retraining.

In PCFM (Utkarsh et al., 2025), equality constraints are addressed through a three-step procedure at each flow substep $\tau \to \tau + \delta\tau$: (i) a forward shoot to the terminal state u_1, (ii) a Gauss-Newton projection onto the linearized equality constraint manifold, (iii) a reverse update via the optimal transport displacement interpolant to propagate the correction back to $\tau + \delta\tau$, and (iv) a relaxed constraint correction to balance satisfying constraint and preserving local flow alignment. A final full projection at $\tau = 1$ ensures machine-precision constraint satisfaction. Our framework, I-PCFM, preserves this overall structure and modifies either the Gauss-Newton projection step or the relaxed constraint correction step to accommodate *inequality constraints*.

3.2. Inequality Projection Framework

At each flow substep, I-PCFM must solve the equality- and inequality-constrained projection:

$$\min_u \|u - u_1\|^2 \quad \text{subject to} \quad h(u) = 0, \quad g(u) \leq 0 \quad (1)$$

where u_1 is the terminal state obtained from the forward shoot. This is a constrained nonlinear least-squares problem, and we introduce three strategies for solving it. All four strategies share the same outer loop as PCFM and the main difference lies in the execution of the inner projection step or the relaxed correction step.

3.3. Strategy A: Slack Variable Reformulation

The slack variable approach converts each inequality constraint $g_i(u) \leq 0$ into an equality constraint by introducing an auxiliary variable $s_i \geq 0$:

$$g_i(u) + s_i = 0, \quad s_i \geq 0, \quad i = 1, \ldots, m. \quad (2)$$

This yields an augmented constraint vector $h^+(u, s) = [h(u);\, g(u) + s] = 0$ over the extended variable (u, s). The projection is then solved by alternating between two steps. First, a Gauss-Newton step on the equality $h^+(u, s) = 0$ is applied, using PCFM's existing Schur complement solver. This step is essentially unchanged from PCFM, and the solver operates on the augmented Jacobian $J^+ = [J_h,\, 0;\, J_g,\, I]$, where I is the identity block corresponding to the slack variables. Second, a non-negativity projection is applied to the slack variables: $s \leftarrow \max(s, 0)$. These two steps are repeated until convergence.

Strategy A has several attractive properties. It requires minimal modification to the PCFM implementation, since the core Schur complement solver is reused without change. The computational cost per Newton step scales as $\mathcal{O}((k + m)^3)$ for the Schur complement solve, the same asymptotic cost as PCFM for equality constraints alone. However, convergence may be slow when constraints are nearly active ($g_i(u) \approx 0$), since small perturbations in u can cause the slack variables to oscillate near zero.

3.4. Strategy B: Log-Barrier Augmentation

Rather than modifying the constraint structure, strategy B augments PCFM's relaxed correction objective with a log-barrier term that penalizes proximity to the constraint boundary:

$$u_{\tau'} = \arg\min_u \|u - \hat{u}_{\tau'}\|^2 + \lambda \|h(u + \gamma v_\theta(u, \tau'))\|^2 - \mu \sum_{i=1}^{m} \log(-g_i(u)), \quad (3)$$

where $\mu > 0$ is the barrier parameter, λ is the relaxed penalty weight, and $\gamma = 1 - \tau'$ is the flow-time-dependent extrapolation coefficient. The log-barrier term diverges as any inequality constraint approaches its boundary, penalizing the generative trajectory for approaching the boundary of the feasible region during sampling. Rather than projecting onto or near the constraint boundary as in strategies A and C, strategy B discourages the trajectory from reaching it in the first place. However, this soft enforcement means that the barrier reduces, but does not eliminate the possibility of constraint violation at intermediate steps.

We use a decreasing schedule for the barrier parameter μ across flow steps, providing only soft enforcement of inequalities during sampling; exact constraint satisfaction relies on a final Gauss-Newton projection at $\tau = 1$ onto the active set $\mathcal{A} = \{i : g_i(u_1) \geq -\varepsilon\}$, treating active constraints as equalities $g_i(u) = 0$ and leaving inactive ones ($g_i(u_1) < -\varepsilon$) untouched since they are already satisfied. Strategy B is fully differentiable through the barrier term and

integrates into PCFM's autodiff-compatible batched solver, but it requires careful tuning of μ and its schedule, as a small barrier term is ineffective in enforcing constraints, while a large barrier term can greatly distort the flow trajectory away from the true solution.

3.5. Strategy C: Active-Set Projection

For the active-set projection approach, we identify the set of *active* inequality constraints ($g_i(u_1) \geq -\varepsilon$ for a small tolerance $\varepsilon > 0$) at each flow step, and treat these as equality constraints, applying the Gauss-Newton projection exactly as it does for equality constraints. Formally, the active set is $\mathcal{A}(u_1) = \{i : g_i(u_1) \geq -\varepsilon\}$, and the projection step is:

$$u_{\text{proj}} = u_1 - J_{\mathcal{A}}^{\top}(J_{\mathcal{A}}J_{\mathcal{A}}^{\top})^{-1}g_{\mathcal{A}}(u_1), \tag{4}$$

where $J_{\mathcal{A}} = \nabla g_{\mathcal{A}}(u_1)$ is the Jacobian restricted to the active constraints and $g_{\mathcal{A}}(u_1)$ is the corresponding residual vector. After the projection, feasibility is verified for all constraints. If any previously inactive constraint has become violated after the projection step, it is added to the active set and the projection is repeated. This active-set expansion procedure continues until a fully feasible point is found.

Strategy C inherits PCFM's exact constraint satisfaction guarantee for the active constraints. When the active set is small, strategy C is the most computationally efficient of the three strategies. However, it is potentially unstable when the active set changes rapidly across flow steps, which can occur near constraint boundaries in highly nonlinear systems.

3.6. Strategy D: Frank-Wolfe Guidance

We explore Frank Wolfe methods (projection-free) to solve constraint projection iteratively. At iteration k, with current iterate $u^{(k)}$ and gradient $\nabla\phi(u^{(k)}) = u^{(k)} - u_1$ of the squared-distance objective, a linear-minimization oracle (LMO) selects a feasible vertex $v^{(k)} \in \mathcal{C}$:

$$\begin{aligned} v^{(k)} &= \arg\min_{v\in\mathcal{C}} \langle \nabla\phi(u^{(k)}),\, v\rangle, \\ u^{(k+1)} &= u^{(k)} + \gamma_k\big(v^{(k)} - u^{(k)}\big), \end{aligned} \tag{5}$$

with step size $\gamma_k \in [0, \gamma_k^{\max}]$ chosen by closed-form line search. Because every iterate is a convex combination of feasible vertices, $u^{(k)} \in \mathcal{C}$ throughout: projection is replaced by a sequence of LMO calls, and each iterate carries a feasibility certificate.

Linear Frank-Wolfe variant. When the constraints are affine in u, as for Heat and Navier-Stokes benchmarks, where IC, mass conservation, and the max-principle box are all linear—$\mathcal{C}$ reduces to a polytope $\{u : Eu = b,\ u_{\text{lo}} \leq u \leq u_{\text{hi}}\}$ and the LMO admits a closed-form solution. We extend the standard box-LMO with the affine equality $Eu = b$, yielding an *augmented LMO* that jointly enforces IC, mass conservation, and the box at every iterate.

Riemannian Frank-Wolfe variant. When an equality constraint is nonlinear, as in the Reaction-Diffusion benchmark, we run *feasible Riemannian* FW on the manifold $\mathcal{M} = \{u : h(u) = 0\}$. The FW direction $d^{(k)} = v^{(k)} - u^{(k)}$ is projected onto the tangent space $T_{u^{(k)}}\mathcal{M} = \ker J_h(u^{(k)})$ via $d_{\tan}^{(k)} = \big(I - J_h^{\top}(J_hJ_h^{\top})^{-1}J_h\big)\, d^{(k)}$, and a Newton retraction $u \leftarrow u - J_h^{\top}(J_hJ_h^{\top})^{-1}h(u)$ restores the iterate to $\mathcal{M}$ after each step; a $\gamma_k^{\max}$ line search against the box constraints maintains inequality feasibility throughout.

Every iterate of I-PCFM-D, which refers to I-PCFM with the FW variant, lies in $\mathcal{C}$, so early termination is safe. In the linear case, the augmented LMO satisfies all equalities and the box *simultaneously* to machine precision in a single sort. The Riemannian variant trades some equality precision satisfaction (via Newton retraction) for guaranteed inequality feasibility at every iterate. However, the principal limitation is LMO tractability: the linear variant requires affine equalities and an axis-aligned box, while the Riemannian variant additionally requires $J_hJ_h^{\top}$ to be well-conditioned. Burgers exposes three failures: (i) Oleinik's inequality is a non-axis-aligned polytope, eliminating the closed-form LMO, (ii) the k-step Godunov-flux unroll makes $J_hJ_h^{\top}$ severely ill-conditioned near shocks, causing Cholesky failures in our experiments, and (iii) most fundamentally, the equality and inequality encode the same underlying entropy condition, so FW alternates between coupled manifolds rather than converging.

4. Experiments

4.1. Experimental Setup

Task and Data. We evaluate I-PCFM on four PDE benchmarks: 1) the **1D Heat equation** $u(x,t)$, $\partial_t u = \nu, \partial_x^2 u$, on the space-time domain $(x,t) \in [0, 2\pi] \times [0,1]$ discretized on a 100×100 grid, 2) the **2D Navier-Stokes equation** in vorticity form $\omega(x,y,t)$, $\partial_t\omega + \mathbf{u}\cdot\nabla\omega = \nu, \Delta\omega + f$ with $\nu = 10^{-3}$, on the space-time domain $(x,y,t) \in [0,2\pi]^2 \times [0,1]$ discretized on a $32 \times 32 \times 25$ grid, 3) the **1D Reaction-Diffusion (Fisher-KPP) equation** $u(x,t)$, $\partial_t u = \nu\,\partial_x^2 u + \rho\, u(1-u)$, on the space-time domain $(x,t) \in [0,1] \times [0,1]$ discretized on a 128×100 grid, and 4) the **1D inviscid Burgers' equation** $u(x,t)$, $\partial_t u + \partial_x\big(\frac{u^2}{2}\big) = 0$, on the space-time domain $(x,t) \in [0,1]^2$ discretized on a 101×101 grid. Further details on the experimental setup can be found in Appendix A.

Constraints. For each PDE we impose two classes of constraints on generated trajectories: equality constraints $h(u) = 0$, which encode the initial condition and conservation laws, and inequality constraints $g(u) \leq 0$, which

Table 3. Comparison on the four PDE benchmarks. Best results per column within each benchmark are **bolded**. Joint-Feas. is reported at threshold 10^{-3}. For Time(s), Vanilla FFM is excluded from the best-time comparison since it performs no constraint enforcement; the best constrained method is bolded instead. I-PCFM-D results on Burgers are omitted (denoted "—") as the method fails to converge. I-PCFM-B and I-PCFM-C on Burgers use best hyperparameters from the sweeps (Tables 5, 6): $\mu_0 = 10^{-4}$ and $\varepsilon = 10^{-2}$ respectively.

Method	CE(IC)↓	CE(CL)↓	CE-Ineq↓	Ineq-Feas.↑	Joint-Feas.↑	MMSE↓	SMSE↓	Time(s)↓
Heat								
Vanilla FFM	6.351	0.0385	7.2×10^{-7}	**1.000**	0.000	$\mathbf{3.6\times10^{-4}}$	$\mathbf{6.5\times10^{-4}}$	0.12
PCFM	8.4×10^{-8}	1.3×10^{-6}	7.1×10^{-4}	0.940	0.940	3.9×10^{-4}	0.0021	0.42
I-PCFM-A (Slack)	1.3×10^{-5}	3.5×10^{-6}	1.6×10^{-4}	0.950	0.950	3.9×10^{-4}	0.0021	7.14
I-PCFM-B (Barrier)	4.0×10^{-5}	1.8×10^{-4}	5.7×10^{-4}	0.940	0.940	3.8×10^{-4}	0.0021	8.13
I-PCFM-C (Active-Set)	1.3×10^{-5}	3.3×10^{-6}	3.6×10^{-5}	**1.000**	**1.000**	3.9×10^{-4}	0.0021	6.82
I-PCFM-D (FW)	$\mathbf{1.2\times10^{-12}}$	$\mathbf{6.8\times10^{-7}}$	**0**	**1.000**	**1.000**	0.0040	0.0056	**0.17**
Navier-Stokes								
Vanilla FFM	18.975	3.472	0.0054	0.990	0.000	**0.1408**	0.0504	0.18
PCFM	6.7×10^{-7}	2.3×10^{-6}	0.0267	0.990	0.990	0.4851	0.1449	1.12
I-PCFM-A (Slack)	6.7×10^{-7}	2.7×10^{-6}	0.0068	0.990	0.990	0.4850	0.1448	9.25
I-PCFM-B (Barrier)	7.4×10^{-4}	0.0048	0.0261	0.990	0.050	0.4850	0.1448	9.82
I-PCFM-C (Active-Set)	6.6×10^{-5}	2.4×10^{-5}	7.2×10^{-7}	**1.000**	**1.000**	0.4850	0.1448	8.28
I-PCFM-D (FW)	$\mathbf{3.0\times10^{-9}}$	$\mathbf{2.2\times10^{-6}}$	**0**	**1.000**	**1.000**	0.4994	**0.0438**	**0.33**
Reaction-Diffusion								
Vanilla FFM	4.886	0.0501	0.0816	0.530	0.000	0.0090	0.0047	0.14
PCFM	2.9×10^{-8}	$\mathbf{5.4\times10^{-7}}$	1.628	0.300	0.300	7.2×10^{-4}	$\mathbf{7.1\times10^{-4}}$	**0.74**
I-PCFM-A (Slack)	6.5×10^{-6}	1.6×10^{-5}	0.0077	0.400	0.400	0.0020	7.3×10^{-4}	8.87
I-PCFM-B (Barrier)	1.2×10^{-5}	4.7×10^{-4}	1.875	0.270	0.270	$\mathbf{4.0\times10^{-4}}$	7.2×10^{-4}	10.85
I-PCFM-C (Active-Set)	6.5×10^{-6}	1.6×10^{-5}	5.0×10^{-4}	0.970	**0.970**	0.0022	7.5×10^{-4}	7.69
I-PCFM-D (FW)	$\mathbf{3.6\times10^{-11}}$	0.0073	$\mathbf{6.2\times10^{-12}}$	**1.000**	**0.970**	0.0022	7.6×10^{-4}	22.59
Burgers								
Vanilla FFM	3.410	0.0578	**0.1489**	0.000	0.000	0.0426	0.0658	0.12
PCFM	$\mathbf{1.1\times10^{-6}}$	$\mathbf{3.9\times10^{-4}}$	0.7496	0.000	0.000	0.0292	**0.0086**	**6.88**
I-PCFM-A (Slack)	1.4×10^{-4}	5.2×10^{-4}	0.2094	0.000	0.000	0.0245	0.0090	34.86
I-PCFM-B (Barrier)	0.0024	0.0018	0.5572	0.140	0.140	0.0244	0.0094	31.03
I-PCFM-C (Active-Set)	2.7×10^{-3}	0.018	0.2958	**0.570**	**0.560**	**0.0209**	0.0141	17.64
I-PCFM-D (FW)	—	—	—	—	—	—	—	—

enforce pointwise admissibility. The equality residual $H(u)$ combines the initial condition (IC) with one or more conservation laws per PDE. For Heat, H_{Heat} adds linear mass conservation (199 constraints per sample). For 2D Navier–Stokes in vorticity form, H_{NS} adds mass conservation (1,048). For Reaction–Diffusion, H_{RD} adds nonlinear mass conservation with a reaction source and boundary-flux correction (227). For Burgers, H_{Burgers} adds a sequence of local Godunov-flux conservation updates and global nonlinear mass conservation (2,181). The inequality constraints encode each PDE's pointwise admissibility conditions: the linear max-principle box $u_0^{\min} \leq u(x,t) \leq u_0^{\max}$ for Heat (19,800 inequalities); a data-driven max-principle box $|\omega(x,y,t)| \leq 1.5 \cdot \max|\omega_{\text{test}}|$ for Navier–Stokes (49,152); the Fisher–KPP invariant interval $0 \leq u(x,t) \leq 1$ for Reaction–Diffusion (25,344); and the Oleinik entropy condition for Burgers (10,000). Full definitions appear in Appendix C.

Baselines and Methods. We compare five methods in our experiments. **Vanilla FFM** generates samples from the pretrained flow matching model with no constraint enforcement, serving as a lower bound on constraint satisfaction. **PCFM** applies the original Physics-Constrained Flow Matching framework (Utkarsh et al., 2025), enforcing only equality constraints via Gauss-Newton projection. Our four proposed methods are **I-PCFM-A** (slack variable reformulation), **I-PCFM-B** (log-barrier augmentation), **I-PCFM-C** (active-set projection), and **I-PCFM-D** (Frank-Wolfe guidance), each enforcing both equality and inequality constraints as described in Section 3. Note that we do not report I-PCFM-D results for Burgers as the method fails to converge (see Section 3.6).

All methods integrate the learned vector field with explicit Euler over $\tau \in [0,1]$ using $n_{\text{steps}} = 100$ ($\Delta\tau = 10^{-2}$). Initial noise u_0 is drawn with a fixed random seed at the start of every run.

4.2. Metrics

We report several metrics to characterize constraint satisfaction, solution quality, and computational cost. **CE(IC)** and **CE(CL)** measure the squared-norm constraint residual for the initial condition and conservation law equality constraints. **CE-Ineq** measures the inequality violation magnitude, averaged over all generated samples. **Ineq-Feasibility**

and **Joint-Feasibility** report the fraction of generated samples that satisfy *all inequality constraints* and *all constraints* (equality and inequality), respectively, within a threshold of 10^{-3}. **MMSE** and **SMSE** measure the pointwise mean squared errors of the mean and standard deviation respectively, between the generated samples and ground-truth trajectories. **Time (s/sample)** reports the wall-clock time per sample on a single NVIDIA L40 GPU, averaged over 100 test samples.

4.3. Main Results

Table 3 reports all metrics for the six methods on four benchmarks (see Appendix B for more information). Comparisons of example trajectories from the different methods can be found in Appendix F.

Equality-only projection harms inequality satisfaction. Despite achieving the best equality satisfaction on both benchmarks (lowest CE(IC) and CE(CL)), PCFM's inequality violation is *worse than Vanilla FFM on all tasks* (e.g., CE-Ineq jumps from 0.082 to 1.628 on Reaction-Diffusion, $20\times$ worse). The gap extends to feasibility: on Reaction-Diffusion, PCFM's inequality feasibility rate drops from 0.530 under Vanilla FFM to 0.300. Projecting samples onto the equality-constraint manifold can push them out of the inequality-feasible region, a failure mode that equality-only projection cannot handle. This directly motivates I-PCFM, which *jointly handles equality and inequality constraints*.

I-PCFM-C dominates among general strategies. Among the three general-purpose strategies in Table 3, the active-set variant consistently achieves the highest joint feasibility, even attaining 0.560 on the challenging Burgers. This dominance comes from how I-PCFM-C allocates its projection budget: at each flow step, it identifies the currently violated constraints and explicitly handles them at the Gauss-Newton projection step. The slack and barrier variants, by contrast, apply their correction uniformly across all inequalities, including the inactive ones, which consume optimization effort without moving the trajectory closer to the feasible region. This selective projection is also reflected in runtime as I-PCFM-C is the fastest variant out of the three across benchmarks. This demonstrates that concentrating the Newton-Schur solve on the active subset both improves feasibility and reduces computational cost. Further details on the limitations of I-PCFM-A and I-PCFM-B are provided in Appendix D. We also provide hyperparameter sensitivity analysis in Appendix E.

Frank-Wolfe guidance dominates where it applies. I-PCFM-D achieves the strongest constraint satisfaction across every PDE where it applies, with effectively exact inequality satisfaction (CE-Ineq ≈ 0) and perfect joint feasibility (1.000) on Heat and Navier-Stokes, and 0.970 joint feasibility on Reaction-Diffusion. On the linear-constraint benchmarks (Heat and Navier-Stokes), it does so at 25–$40\times$ the speed of I-PCFM-C (0.17s vs 6.82s on Heat; 0.33s vs 8.28s on Navier-Stokes), because of the augmented LMO's closed-form solution. On Reaction-Diffusion, where the equality is nonlinear, the Riemannian variant attains the lowest CE-Ineq among all methods (6.2×10^{-12}) and matches I-PCFM-C's joint feasibility, but at $\sim 3\times$ the runtime (22.59s vs 7.69s) due to the per-iterate inversion of $J_h J_h^\top$ for tangent projection and Newton retraction. I-PCFM-D is omitted on Burgers, whose constraint coupling and ill-conditioning at shocks cause both Frank-Wolfe variants to fail to converge. These results suggest I-PCFM-D should be preferred whenever its closed-form LMO is available.

4.4. Active-Set Dynamics of I-PCFM-C

To understand how inequality constraints interact with the flow, we track the active-set size $|\mathcal{A}(\tau)|$ for I-PCFM-C ($\varepsilon = 10^{-3}$) on the Burgers benchmark at every flow step $\tau \in [0, 1]$ across 100 test samples, out of the 10,000 total entropy constraints per sample. The results are shown in Figure 2.

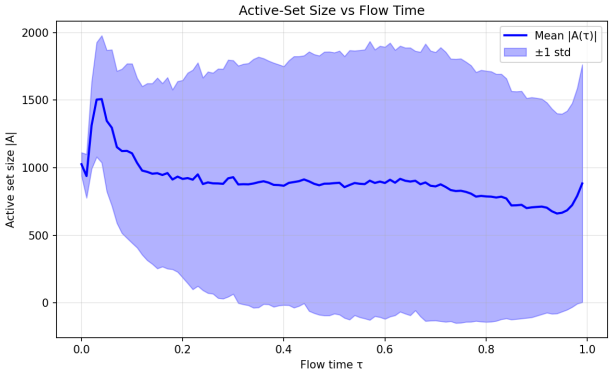

Figure 2. Active-set size $|\mathcal{A}(\tau)|$ over flow time τ for I-PCFM-C ($\varepsilon = 10^{-3}$) across 100 test samples. The shaded region shows ±1 standard deviation.

The active set decreases throughout the flow. The mean active-set size peaks at $|\mathcal{A}| \approx 1500$ at $\tau = 0.04$, which means that only 15% of all entropy constraints are active after flow onset (Figure 2). The mean decreases steadily throughout the flow time. By $\tau = 0.5$ (midpoint of flow), $|\mathcal{A}|$ has dropped to roughly 9% and reached a minimum of 6.6% at $\tau = 0.95$. This steady decrease makes I-PCFM-C's active set restriction effective, as it avoids projection against the full 10,000 inequality constraints at every flow step. As the cost of each Newton-Schur solve is cubic in $|\mathcal{A}|$, the cost savings of I-PCFM-C enables it to run much faster than the slack and barrier variants (Table 3).

Another interesting observation is that the mean increases

sharply after $\tau \approx 0.95$. One possible reason is that the equality projection, which becomes tighter as $\tau \to 1$ could inadvertently introduce new entropy violations near shock locations. This is consistent with the finding from Section 4.3 that equality-only projection harms inequality satisfaction.

Active-set size varies widely across samples. The sample-to-sample spread in $|\mathcal{A}|$ grows sharply after flow onset and exceeds the mean itself by mid-flow (Figure 2, shaded band). Since the Newton-Schur solve scales as $|\mathcal{A}|^3$, this translates to a large per-step compute difference across trajectories. Some shock-rich samples dominate the projection budget but most remain cheap. This variance in active-set size explains the feasibility values in Table 3, as I-PCFM-C can resolve the active set of inequalities on well-behaved trajectories but fails on outliers whose local $|\mathcal{A}|$ is too large for a single-step Newton projection to resolve.

5. Discussion & Future Work

5.1. Limitations

Our work has several limitations. First, the three general-purpose I-PCFM strategies incur substantial computational overhead relative to equality-only PCFM (10–15$\times$ on Reaction-Diffusion, 2.5–5$\times$ on Burgers). Second, strategy B is a weaker method by design, in comparison to the other general-purpose strategies A and C. It handles inequalities in PCFM's relaxed correction step via a log-barrier, rather than in the Gauss-Newton projection itself, which caps the quality of constraint enforcement achievable regardless of μ_0. Finally, strategy D fails when the equality and inequality constraints are not independent. For Burgers, the Godunov-flux equality and Oleinik inequality both encode the same entropy condition, and the Godunov residual is non-smooth precisely at shocks—exactly where the inequality binds—so the Newton retraction onto the equality manifold degenerates and the FW-Newton loop cycles without converging.

5.2. Conclusion & Future Work

We introduced **Inequality-aware Physics-Constrained Flow Matching (I-PCFM)**, a framework that extends Physics-Constrained Flow Matching to jointly handle equality and inequality constraints at inference time without retraining. We proposed and systematically compared three general projection strategies, each grounded in a classical constrained-optimization paradigm—slack reformulation, log-barrier augmentation, and active-set projection. We also introduced Frank-Wolfe guidance, a projection-free strategy with closed-form LMOs for linear constraints and a Riemannian extension for nonlinear settings. Across four PDE benchmarks, we establish three key findings. First, equality-only projection actively harms inequality satisfaction, and inequality constraints must be handled explicitly to achieve physical admissibility. Second, among the three general-purpose strategies, active-set projection dominates on feasibility and computational efficiency. Lastly, Frank-Wolfe guidance matches active-set projection's feasibility while running up to $40\times$ faster.

Several directions for future work follow naturally. First, extending the evaluation to higher-dimensional PDEs and richer inequality structures (e.g., non-axis-aligned polytopes, chance constraints) would test whether the results observed in this work generalize to other settings. Second, extending I-PCFM to settings where constraints must be learned from data—for instance, when the governing physics is partially unknown—would broaden its applicability to inverse problems and data-driven scientific discovery. More broadly, we view I-PCFM as a step toward generative models for the physical sciences in which physical admissibility across both equalities and inequalities is enforced as a hard requirement, and not as a soft objective.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Use of Generative AI

Claude (Anthropic) was used as a coding assistant as part of this work, to help implement and debug the code.

References

Baldan, G., Liu, Q., Guardone, A., and Thuerey, N. Physics vs distributions: Pareto optimal flow matching with physics constraints, 2026. URL https://arxiv.org/abs/2506.08604.

Ben-Hamu, H., Puny, O., Gat, I., Karrer, B., Singer, U., and Lipman, Y. D-flow: Differentiating through flows for controlled generation, 2024. URL https://arxiv.org/abs/2402.14017.

Bertsekas, D. P. *Nonlinear Programming*. Athena Scientific, 1997.

Blanke, M., Qu, Y., Shamekh, S., and Gentine, P. Strictly constrained generative modeling via split augmented langevin sampling, 2026. URL https://arxiv.org/abs/2505.18017.

Bortoli, V. D., Mathieu, E., Hutchinson, M., Thornton, J., Teh, Y. W., and Doucet, A. Riemannian score-based generative modelling, 2022. URL https://arxiv.org/abs/2202.02763.

Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, 2004.

Chen, R. T. Q. and Lipman, Y. Flow matching on general geometries, 2024. URL https://arxiv.org/abs/2302.03660.

Cheng, C., Han, B., Maddix, D. C., Ansari, A. F., Stuart, A., Mahoney, M. W., and Wang, Y. Gradient-free generation for hard-constrained systems, 2025. URL https://arxiv.org/abs/2412.01786.

Christopher, J. K., Baek, S., and Fioretto, F. Constrained synthesis with projected diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Feng, R., Yu, C., Deng, W., Hu, P., and Wu, T. On the guidance of flow matching, 2025. URL https://arxiv.org/abs/2502.02150.

Frank, M. and Wolfe, P. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956. doi: https://doi.org/10.1002/nav.3800030109. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800030109.

Huang, J., Yang, G., Wang, Z., and Park, J. J. Diffusionpde: Generative pde-solving under partial observation, 2024. URL https://arxiv.org/abs/2406.17763.

Jaggi, M. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In Dasgupta, S. and McAllester, D. (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 427–435, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL https://proceedings.mlr.press/v28/jaggi13.html.

Kerrigan, G., Migliorini, G., and Smyth, P. Functional flow matching, 2023. URL https://arxiv.org/abs/2305.17209.

Lacoste-Julien, S. and Jaggi, M. On the global linear convergence of frank-wolfe optimization variants. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, pp. 496–504, Cambridge, MA, USA, 2015. MIT Press.

Lacoste-Julien, S., Jaggi, M., Schmidt, M., and Pletscher, P. Block-coordinate frank-wolfe optimization for structural svms. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML'13, pp. I–53–I–61. JMLR.org, 2013.

LeVeque, R. J. *Numerical methods for conservation laws (2. ed.).* Lectures in mathematics. Birkhäuser, 1992. ISBN 978-3-7643-2723-1.

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

Liang, J., Sun, Y., Samaddar, A., Madireddy, S., and Fioretto, F. Chance-constrained flow matching for high-fidelity constraint-aware generation. *arXiv preprint arXiv:2509.25157*, 2025.

Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling, 2023. URL https://arxiv.org/abs/2210.02747.

Murray, J. D. and Murray, J. D. *Mathematical biology: II: spatial models and biomedical applications*, volume 18. Springer, 2003.

Nocedal, J. and Wright, S. J. *Numerical Optimization*. Springer, 2nd edition, 2006.

Tauberschmidt, J., Fellenz, S., Vollmer, S. J., and Duncan, A. B. Physics-constrained fine-tuning of flow-matching models for generation and inverse problems, 2026. URL https://arxiv.org/abs/2508.09156.

Utkarsh, U., Cai, P., Edelman, A., Gomez-Bombarelli, R., and Rackauckas, C. V. Physics-constrained flow matching: Sampling generative models with hard constraints, 2025. URL https://arxiv.org/abs/2506.04171.

Weber, M. and Sra, S. Riemannian optimization via frank-wolfe methods, 2021. URL https://arxiv.org/abs/1710.10770.

Zhang, X. and Shu, C.-W. Positivity-preserving high order discontinuous Galerkin schemes for compressible Euler equations with source terms. *Journal of Computational Physics*, 230(4):1238–1248, February 2011. doi: 10.1016/j.jcp.2010.10.036.

A. Experimental Setup Details

Dataset for Heat equation. Training data consists of 5,000 trajectories at resolution 100×100 on the periodic space-time domain $(x, t) \in [0, 2\pi) \times [0, 1]$. Because the 1D heat equation $\partial_t u = v\, \partial_x^2 u$ with sinusoidal initial conditions admits a closed-form solution, trajectories are computed analytically as $u(x, t) = \sin(x + \varphi)\, e^{-vt}$ rather than via a numerical solver. The initial condition is $u_0(x) = \sin(x + \varphi)$ with phase $\varphi \sim \mathcal{U}[0, \pi]$; periodic boundary conditions are inherited from the torus domain. The diffusivity is sampled per trajectory as $v \sim \mathcal{U}[1, 5]$. The test set contains 1,000 trajectories generated from the same distribution with held-out random seeds. We evaluate on a fixed subset of $N = 100$ test trajectories (seed 42) for all experiments.

Dataset for Navier–Stokes equation. Training data consists of $100 \times 100 = 10{,}000$ trajectories from 100 initial vorticity fields paired with 100 forcing functions, generated numerically by a pseudo-spectral solver with Crank–Nicolson time stepping on the viscous term and 2/3-rule dealiasing of the nonlinear advection. Native generation resolution is $64 \times 64 \times 50$ on the periodic unit torus $(x, y) \in [0, 1)^2$ at viscosity $\nu = 10^{-3}$ (Reynolds number $\mathrm{Re} = 1{,}000$) over a physical time horizon $T = 49$; we coarsen to $32 \times 32 \times 25$ by every-other subsampling for training and evaluation, with model time normalized to $t \in [0, 1]$. Initial vorticity fields are drawn from a Gaussian random field $\omega_0 \sim \mathrm{GRF}(\alpha = 2.5,\ \tau = 7)$, and the forcing is a deterministic Kolmogorov-like sinusoid $f_k(x, y) = 0.1\sqrt{2}\, \sin\big(2\pi(x + y) + \varphi_k\big)$ with phase $\varphi_k = (\pi/2) \cdot k/99$ for $k = 0, \ldots, 99$. The data-driven max-principle bound is set to $B = 1.5 \cdot \max |\omega_{\text{test}}|$ (≈ 3.18 on our split). The test set contains $10 \times 100 = 1{,}000$ trajectories from the same solver and forcing parametrization; we evaluate on a random subset of $N = 100$ for all experiments.

Dataset for Reaction-Diffusion (Fisher-KPP) equation. Training data consists of $80 \times 80 = 6{,}400$ trajectories generated at resolution 128×100 on $(x, t) \in [0, 1]^2$ via a splitting scheme that integrates the logistic reaction exactly and the diffusion with centred finite differences under Neumann boundary fluxes $g_L \sim \mathcal{U}(0, 0.05)$ and $g_R \sim \mathcal{U}(-0.05, 0)$. Each initial condition u_0 is sampled as a sum of three random sinusoidal modes, optionally masked by a smooth bump, and rescaled so that $u_0(x) \in [0, 1]$. The reaction rate and diffusivity are fixed at $\rho = 0.01$ and $\nu = 0.005$. The test set contains $30 \times 30 = 900$ trajectories generated with a held-out random seed. We use a random subset of $N = 100$ test trajectories in the main experiment.

Dataset for Burgers' equation. Training data consists of $80 \times 80 = 6{,}400$ solution trajectories generated from 80 initial conditions paired with 80 boundary conditions, produced numerically with a Godunov flux scheme at resolution 101×101 on the space-time domain $(x, t) \in [0, 1]^2$. The initial condition is a shifted sigmoid $u_0(x) = 1/(1 + \exp((x - p)/\varepsilon))$ with $\varepsilon = 0.02$ and shift $p \sim \mathcal{U}(0.2, 0.8)$; the left Dirichlet boundary value satisfies $u_{\text{bc}} \sim \mathcal{U}(0, 1)$ and the right boundary is zero-gradient (Neumann). The test set contains $30 \times 30 = 900$ trajectories generated with a held-out random seed. We use a random subset of $N = 100$ test trajectories for all experiments.

Model. For all tasks, we train a Functional Flow Matching (FFM) model (Kerrigan et al., 2023) from scratch for 20,000 iterations (batch size 256; Adam with $\eta = 3 \times 10^{-4}$, $(\beta_1, \beta_2) = (0.9, 0.999)$; plateau scheduler with factor 0.5 and patience 10; gradient-norm clipping at 100). The vector-field network is a Fourier Neural Operator (Li et al., 2020) with $n_{\text{modes}} = (32, 32)$, 4 Fourier layers, 64 hidden channels, 256 projection channels, and 32-dimensional sinusoidal time embeddings ($\sim$9.0 M parameters). The FFM Gaussian-process prior uses kernel length scale 10^{-3} and variance 1.

Sampling. All methods integrate the learned velocity field with a first-order Euler scheme using 100 steps over the flow time $\tau \in [0, 1]$ (distinct from the physical PDE time t). I-PCFM applies Newton projection at each step with a single linearised step and solve tolerance 10^{-4}.

B. Additional Configurations for I-PCFM Methods

All I-PCFM strategies share the following solver settings: one Newton step per flow substep, a Tikhonov regularizer $\epsilon_{\text{solve}} = 10^{-4}$ on the Gram system $(JJ^\top + \epsilon_{\text{solve}} I)$, a trust region capping each correction at $10\times$ the current state norm, and a least-squares fallback whenever the Gram system is numerically singular. Strategy-specific defaults are: I-PCFM-A: slack threshold $= 0.05$; I-PCFM-B: $\mu_0 = 10^{-4}$ for Burgers and $\mu_0 = 10^{-3}$ for Reaction-Diffusion, decay rate $\beta = 3$ (so $\mu(\tau) = \mu_0 \exp(-\beta\tau)$), barrier weight $\lambda = 10^{-2}$, 10 inner gradient steps of size 10^{-2}, $\epsilon_{\text{barrier}} = 10^{-6}$; I-PCFM-C: active-set tolerance $\varepsilon = 10^{-2}$ for Burgers and $\varepsilon = 10^{-3}$ for Reaction-Diffusion.

I-PCFM-D uses two variants depending on the equality structure. The linear variant is used for Heat and Navier-Stokes, where the augmented LMO solves IC, mass conservation, and the box jointly per FW iteration via a closed-form continuous-knapsack sort. The Riemannian variant is used for Reaction-Diffusion, running Riemannian FW on the equality manifold with feasibility-preserving Newton retraction and IC pinning. Both variants use $K_{\text{FW}} = 50$ inner FW iterations per correction, FW duality-gap tolerance 10^{-7}, a joint (equality + inequality) early-stop tolerance of 10^{-4}, $K_{\text{Newton}} = 8$ Newton steps in the retraction loop, and equality-system Tikhonov $\epsilon_{\text{eq}} = \epsilon_{\text{solve}} = 10^{-4}$; the number of outer alternating-projection rounds is $K_{\text{outer}} = 30$ for Heat and Navier-Stokes and $K_{\text{outer}} = 15$ for Reaction-Diffusion.

C. PDE details

C.1. Heat equation

We consider the 1D linear heat equation on the periodic domain $x \in [0, 2\pi)$, $t \in [0, 1]$:

$$\partial_t u = \nu\, \partial_x^2 u, \qquad \nu > 0, \tag{6}$$

with periodic boundary conditions and diffusivity ν sampled per trajectory.

Equality constraints. Both terms in (6) are linear in u, and on the periodic domain the boundary contributions vanish. Integrating over the spatial domain yields global mass conservation:

$$\frac{d}{dt} m(t) = 0, \qquad m(t) := \int_0^{2\pi} u(x,t)\, dx, \tag{7}$$

i.e., $m(t) = m(0)$ for all $t > 0$. Combining with the initial condition gives the equality residual

$$H_{\text{Heat}}(u) = \begin{pmatrix} u(x,0) - u_0^\star(x) \\ m(t) - m(0) \end{pmatrix} = 0, \tag{8}$$

where the first component pins the initial condition and the second enforces global mass conservation. Both components are linear in u. This yields $n_x + (n_t - 1) = 100 + 99 = 199$ equality constraints per sample.

Inequality constraints. The 1D heat equation satisfies the maximum principle: the solution is bounded pointwise by the extrema of the initial data,

$$u_0^{\min} \le u(x,t) \le u_0^{\max}, \qquad u_0^{\min} := \min_x u_0(x), \;\; u_0^{\max} := \max_x u_0(x), \tag{9}$$

for all $x \in [0, 2\pi)$ and $t > 0$. We enforce this pointwise as a two-sided box constraint,

$$g_{i,n}^{\text{upper}}(u) = u[i,n] - u_0^{\max} \le 0, \qquad g_{i,n}^{\text{lower}}(u) = u_0^{\min} - u[i,n] \le 0, \tag{10}$$

for $i = 0, \ldots, n_x - 1$ and $n = 1, \ldots, n_t - 1$, yielding $2\, n_x(n_t - 1) = 19{,}800$ pointwise inequalities per sample. As with Reaction-Diffusion, the $t = 0$ slice is excluded because the IC is already pinned by the first component of $H_{\text{Heat}}(u)$.

C.2. Navier-Stokes equation

We consider the 2D incompressible Navier-Stokes equation in vorticity-streamfunction form on the periodic torus $(x, y) \in [0, 1)^2$, $t \in [0, 1]$:

$$\partial_t \omega + (\mathbf{u} \cdot \nabla)\omega = \nu\, \nabla^2 \omega + f(x,y), \qquad \nu = 10^{-3}, \tag{11}$$

with periodic boundary conditions, divergence-free velocity $\mathbf{u}$ recovered from ω via the streamfunction, and Kolmogorov-like deterministic forcing $f(x,y) = 0.1\sqrt{2}\, \sin\big(2\pi(x+y) + \varphi\big)$.

Equality constraints. Integrating (11) over the periodic torus, the advective term and the viscous Laplacian both vanish (divergence theorem with periodic BCs), and the Kolmogorov forcing has zero mean: $\int_0^1\!\int_0^1 \sin(2\pi(x+y) + \varphi)\, dx\, dy = 0$ for any φ. Total vorticity is therefore conserved:

$$\frac{d}{dt} m(t) = 0, \qquad m(t) := \int_{[0,1)^2} \omega(x,y,t)\, dx\, dy. \tag{12}$$

Combining with the initial vorticity field gives the equality residual

$$H_{\text{NS}}(\omega) = \begin{pmatrix} \omega(x, y, 0) - \omega_0^\star(x, y) \\ m(t) - m(0) \end{pmatrix} = 0, \tag{13}$$

where the first component pins the initial vorticity over the $n_x \times n_y$ spatial grid and the second enforces conservation of total vorticity at every subsequent time step. Both components are linear in ω. This yields $n_x\, n_y + (n_t - 1) = 32 \cdot 32 + 24 = 1{,}048$ equality constraints per sample.

Inequality constraints. The 2D Navier-Stokes equation does not admit a strict parabolic maximum principle, but the vorticity remains bounded over the flow regimes considered. We enforce a data-driven two-sided box on the vorticity magnitude,

$$g^{\text{upper}}_{i,j,n}(\omega) = \omega[i, j, n] - B \le 0, \qquad g^{\text{lower}}_{i,j,n}(\omega) = -\omega[i, j, n] - B \le 0, \tag{14}$$

where $B = c \cdot \max |\omega_{\text{test}}|$ with safety factor $c = 1.5$ and the maximum taken over the $N = 100$ evaluation trajectories ($B \approx 3.18$ on our split). For $i = 0, \ldots, n_x - 1$, $j = 0, \ldots, n_y - 1$, and $n = 1, \ldots, n_t - 1$, this yields $2\, n_x\, n_y\, (n_t - 1) = 2 \cdot 32 \cdot 32 \cdot 24 = 49{,}152$ pointwise inequalities per sample. As before, the $t = 0$ slice is excluded because the IC is already pinned by the first component of $H_{\text{NS}}(\omega)$.

C.3. Reaction-Diffusion equation

We consider the 1D Reaction-Diffusion (Fisher-KPP) equation on $x \in [0, 1]$, $t \in [0, 1]$:

$$\partial_t u = \nu\, \partial_x^2 u + \rho\, u(1 - u), \qquad (\rho, \nu) = (0.01, 0.005), \tag{15}$$

with Neumann boundary fluxes $g_L(t)$ and $g_R(t)$ prescribed at the left and right boundaries.

Equality constraints. Because the reaction term is nonlinear in u, global mass is not automatically conserved. Integrating (15) over the spatial domain and using the Neumann boundary conditions yields

$$\frac{d}{dt} m(t) = \rho \int_0^1 u(1 - u)\, dx + g_L(t) - g_R(t), \qquad m(t) := \int_0^1 u(x, t)\, dx. \tag{16}$$

Integrating (16) in time from 0 to t gives the equality residual

$$H_{\text{RD}}(u) = \begin{pmatrix} u(x, 0) - u_0^\star(x) \\ m(t) - m(0) - \rho \displaystyle\int_0^t \int_0^1 u(1 - u)\, dx\, ds - \int_0^t \big[g_L(s) - g_R(s)\big]\, ds \end{pmatrix} = 0, \tag{17}$$

where the first component pins the initial condition and the second enforces mass balance with the nonlinear reaction source and boundary-flux transport. The boundary fluxes g_L, g_R are recovered from u at each step using a 5-point one-sided finite-difference stencil. This yields $n_x + (n_t - 1) = 128 + 99 = 227$ equality constraints per sample.

Inequality constraints. The Reaction-Diffusion equation preserves the invariant interval $u(x, t) \in [0, 1]$ for all $t > 0$, provided the initial and boundary data lie in $[0, 1]$. We enforce this pointwise as a two-sided box constraint,

$$g^{\text{upper}}_{i,n}(u) = u[i, n] - 1 \le 0, \qquad g^{\text{lower}}_{i,n}(u) = -\, u[i, n] \le 0, \tag{18}$$

for $i = 0, \ldots, n_x - 1$ and $n = 1, \ldots, n_t - 1$, yielding $2\, n_x (n_t - 1) = 25{,}344$ pointwise inequalities per sample. The $t = 0$ slice is excluded because the initial condition is already pinned by the first component of $H_{\text{RD}}(u)$.

C.4. Inviscid Burgers' equation

We consider the 1D inviscid Burgers' equation on $x \in [0, 1]$, $t \in [0, 1]$:

$$\partial_t u + \partial_x \left(\tfrac{u^2}{2} \right) = 0, \tag{19}$$

with convex flux $f(u) = u^2/2$, Dirichlet boundary condition $u(0, t) = u_{\text{bc}}$ prescribed on the left, and zero-gradient (Neumann) condition on the right.

Equality constraints. Burgers' equation is a scalar conservation law: integrating (19) over the spatial domain and applying the boundary conditions gives the mass balance

$$\frac{d}{dt}m(t) = f(u(0,t)) - f(u(1,t)), \qquad m(t) := \int_0^1 u(x,t)\,dx, \tag{20}$$

whose time integral pins the global mass of the solution up to the net boundary flux. Locally, conservation is enforced by the discrete Godunov update: given the state at time n, the update at time $n+1$ is

$$R_{i,n}^{(k)}(u) = u_i^{n+1} - u_i^n + \frac{\Delta t}{\Delta x}\Big[F(u_i^n, u_{i+1}^n) - F(u_{i-1}^n, u_i^n)\Big] = 0, \qquad n = 0, \ldots, k-1, \tag{21}$$

with Godunov numerical flux

$$F(u_L, u_R) = \begin{cases} \min\{\frac{1}{2}u_L^2,\ \frac{1}{2}u_R^2\} & \text{if } u_L \le u_R, \\ \frac{1}{2}u_L^2 & \text{if } u_L > u_R \text{ and } (u_L + u_R)/2 > 0, \\ \frac{1}{2}u_R^2 & \text{if } u_L > u_R \text{ and } (u_L + u_R)/2 \le 0. \end{cases} \tag{22}$$

Combining the initial condition, k unrolled local updates, and the global mass balance gives the equality residual

$$H_{\text{Burgers}}(u) = \begin{pmatrix} u(x,0) - u_0^\star(x) \\ R_{\text{Flux}}^{(k)}(u) \\ m(t) - m(0) + \displaystyle\int_0^t \big[f(u(1,s)) - f(u(0,s))\big]\,ds \end{pmatrix} = 0, \tag{23}$$

where the three components pin the initial condition, enforce the local Godunov dynamics over a short horizon, and enforce global mass balance, respectively. We use $k = 20$ unrolled steps, balancing short-horizon dynamics enforcement against Newton-projection cost. This yields $n_x + k(n_x - 2) + (n_t - 1) = 101 + 20 \cdot 99 + 100 = 2{,}181$ equality constraints per sample.

Inequality constraints. For the inviscid Burgers' equation, equality constraints alone admit infinitely many weak solutions whenever shocks form. The physically correct one—coinciding with the vanishing-viscosity limit $\nu \to 0^+$ of $\partial_t u + \partial_x(u^2/2) = \nu\,\partial_x^2 u$—is selected by the Oleinik entropy condition, which we enforce pointwise:

$$g_{i,n}(u) = u[i{+}1, n] - u[i,n] - \frac{\Delta x}{n\,\Delta t} \le 0, \qquad i = 0, \ldots, n_x - 2, \;\; n = 1, \ldots, n_t - 1, \tag{24}$$

yielding $(n_x - 1)(n_t - 1) = 10{,}000$ pointwise inequalities per sample. The condition is imposed from $n = 1$ onwards since the bound scales as $\Delta x/(n\,\Delta t)$ and is undefined at $t = 0$; at that slice, the solution is already pinned by the IC equality.

C.5. Summary

Table 4 summarizes the scale and type of both constraint classes across the four benchmarks. Equality constraints are linear in u for Heat and Navier-Stokes—the IC pinning and conservation laws are both linear in the state variable—but nonlinear for Reaction-Diffusion (logistic reaction term $\rho\,u(1-u)$) and Burgers (quadratic flux $f(u) = u^2/2$ entering both the Godunov update and the boundary mass balance). All four inequality classes are linear: Heat and Navier-Stokes use a max-principle box; Reaction-Diffusion uses the Fisher-KPP invariant interval $[0, 1]$; and Burgers uses the Oleinik entropy condition.

PDE	Equality type (# eq.)	Inequality type (# ineq.)
Heat	Linear (199)	Max-principle box, linear (19,800)
Navier-Stokes	Linear (1,048)	Data-driven box, linear (49,152)
Reaction-Diffusion	Nonlinear (227)	Invariant interval, linear (25,344)
Burgers	Nonlinear (2,181)	Oleinik entropy, linear (10,000)

Table 4. Summary of constraint structure for each benchmark PDE. Counts are per sample.

D. Discussion on Slack and Barrier methods

From Table 3, neither slack nor barrier strategy matches the active-set variant's feasibility. I-PCFM-A (Slack) reaches 0.400 on Reaction-Diffusion and zero feasibility on Burgers: its slack reformulation adds one variable per inequality ($10{,}000$–$49{,}152$ extra degrees of freedom), diluting the Newton step across all constraints rather than the active ones. I-PCFM-B (Barrier) performs worse, with high CE-Ineq on all benchmarks, and achieving only 0.140 feasibility on Burgers. This is a structural disadvantage as I-PCFM-B runs equality-only Gauss-Newton at the projection step and only handles inequalities at the subsequent relaxed correction step via a log-barrier. I-PCFM-A and I-PCFM-C, by contrast, incorporate inequality handling directly into the Gauss-Newton projection itself, which inherits the projection's hard-satisfaction guarantee at every flow step. Interestingly, both methods perform relatively well on the Heat equation, which could be an easy benchmark for constraint satisfaction.

E. Hyperparameter Sensitivity Analysis

We sweep the barrier coefficient μ_0 for I-PCFM-B and the active-set threshold ε for I-PCFM-C to characterize each method's sensitivity to its respective hyperparameter and identify practical operating ranges. We run these experiments only on the Burgers' equation.

Table 5. I-PCFM-B sensitivity to initial barrier coefficient μ_0 on Burgers.

μ_0	CE(IC) ↓	CE(CL) ↓	CE-Ineq ↓	Ineq-Feas. ↑	Joint-Feas. ↑	MMSE ↓	SMSE ↓	Time (s/sample)
10^{-4}	2.4e-3	**1.8e-3**	**0.557**	**0.140**	**0.140**	**0.0244**	0.0094	31.03
10^{-3}	2.4e-3	**1.8e-3**	0.573	0.040	0.040	0.0247	**0.0093**	30.04
10^{-2}	**1.5e-3**	2.2e-3	1.443	0.000	0.000	0.0321	0.0098	54.19
10^{-1}	5.1e-3	0.091	9.338	0.000	0.000	0.0303	0.0509	92.71

Table 6. I-PCFM-C sensitivity to active-set threshold ε on Burgers.

ε	CE(IC) ↓	CE(CL) ↓	CE-Ineq ↓	Ineq-Feas. ↑	Joint-Feas. ↑	MMSE ↓	SMSE ↓	Time (s/sample)
10^{-4}	1.3e-3	**3.6e-3**	0.408	0.440	0.440	0.0248	**0.0125**	**16.39**
10^{-3}	**9.6e-4**	5.4e-3	0.418	0.530	0.530	0.0246	**0.0125**	16.28
10^{-2}	2.7e-3	0.018	**0.296**	**0.570**	**0.560**	**0.0209**	0.0141	17.64
10^{-1}	0.540	182.1	12.55	0.000	0.000	0.0761	1.055	33.19

I-PCFM-B is fragile to μ_0. Table 5 shows that I-PCFM-B is sensitive to the initial barrier coefficient μ_0. Peak feasibility of 0.140 occurs at $\mu_0 = 10^{-4}$, and increasing it leads to a collapse in feasibility. Larger μ_0 values also make the log-barrier numerically ill-conditioned near the feasible boundary, inflating per-sample runtime by $3\times$, from 31s at $\mu_0 = 10^{-4}$ to 93s at $\mu_0 = 10^{-1}$, with no corresponding improvement in any metric. Even at its best setting, I-PCFM-B reaches only a quarter of the feasibility I-PCFM-C achieves under its own best setting, reinforcing that the barrier's poor performance is structural rather than tuning-dependent.

I-PCFM-C degrades gracefully with ε. I-PCFM-C exhibits a monotonic, well-behaved tradeoff across $\varepsilon \in [10^{-4}, 10^{-2}]$, as shown in Table 6. Feasibility improves smoothly from 0.440 to 0.570 as the threshold loosens, while runtime remains relatively constant. The method degrades catastrophically only at $\varepsilon = 10^{-1}$, where the threshold is so loose that nearly every constraint enters the active set, ill-conditioning the Newton solve: CE(CL) diverges to 182 and feasibility drops to zero. But, in contrast to I-PCFM-B, the active-set variant provides a usable operating range that spans three orders of magnitude of ε, with the best setting ($\varepsilon = 10^{-2}$) winning on CE-Ineq, MMSE, Ineq-Feasibility and Joint-Feasibility.

F. Example Trajectories

We present qualitative trajectory visualizations for three representative test samples for each benchmark, comparing all methods across flow time $\tau \in [0, 1]$ (Figures 3, 4, 5 for Heat equation, Figures 6, 7, 8 for Navier-Stokes equation, Figures 9, 10, 11 for Reaction-Diffusion equation and Figures 12, 13, 14 for Burgers' equation). Each figure shows space-time heatmaps of the solution $u(x, t)$ at selected flow snapshots, with the ground-truth trajectory in the top row for reference. The following rows are in the order of Vanilla FFM, PCFM, I-PCFM-A, I-PCFM-B, I-PCFM-C, and I-PCFM-D (if available). These examples illustrate how the different strategies behave under varying trajectories and conditions.

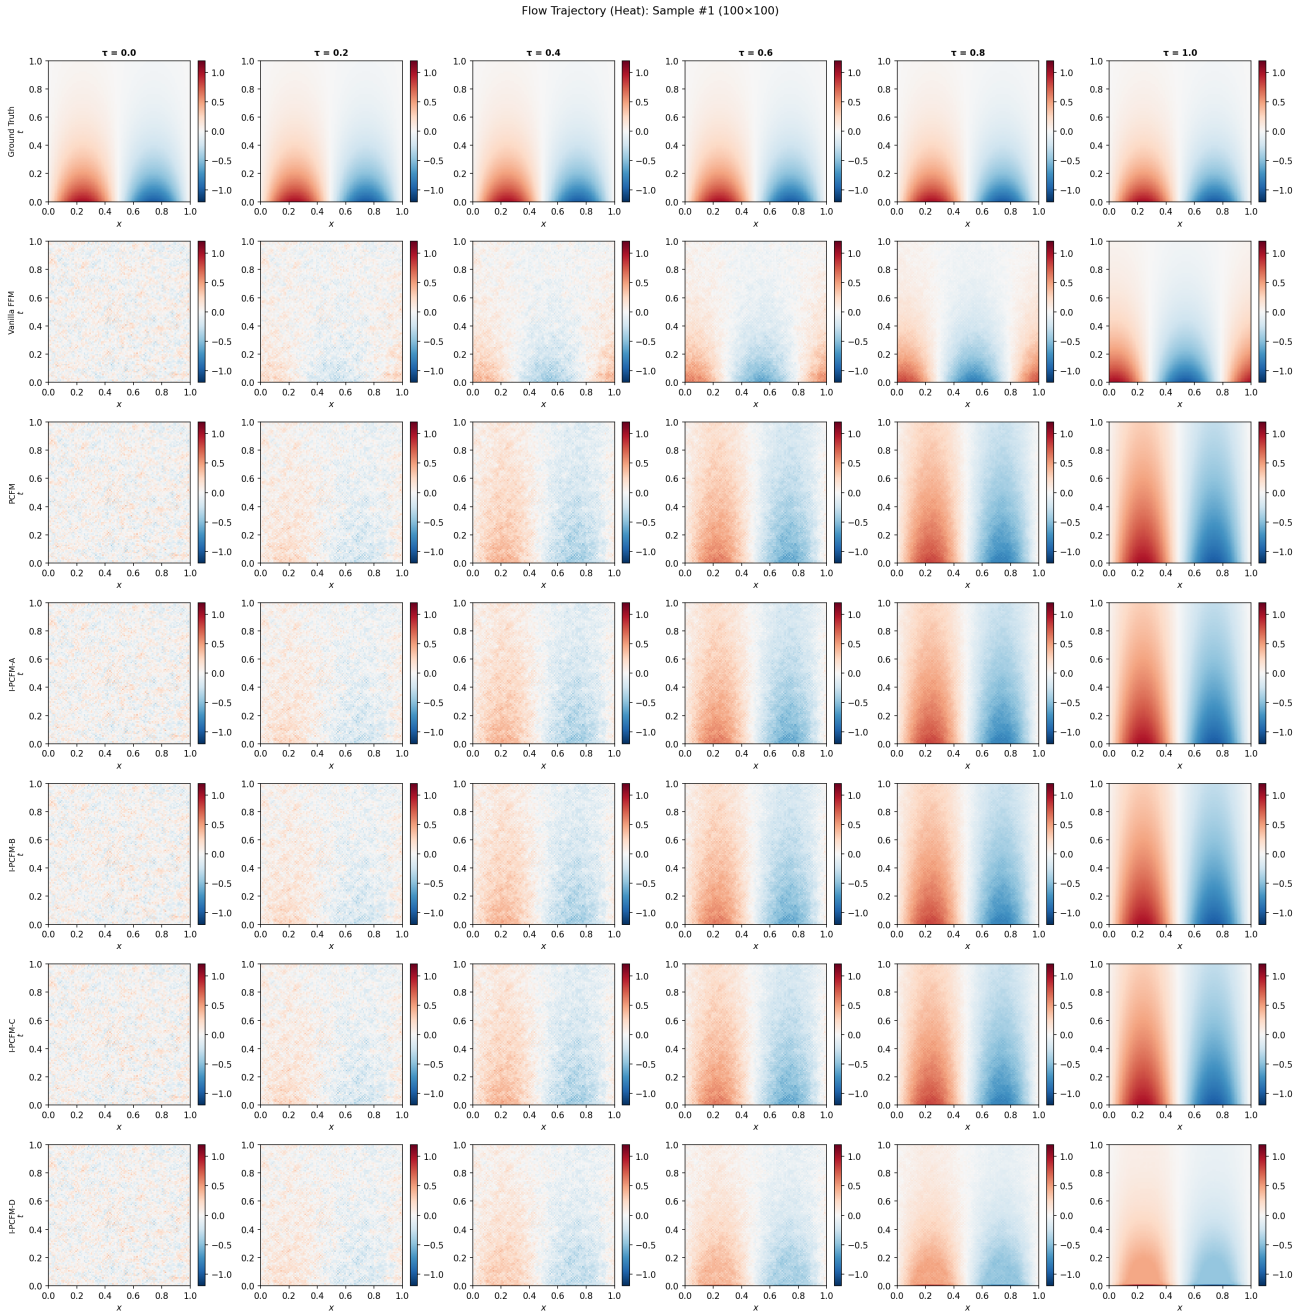

Figure 3. Flow-time evolution of an example Heat trajectory. Constraint-enforcing methods achieve better performance visually, as compared to Vanilla FFM.

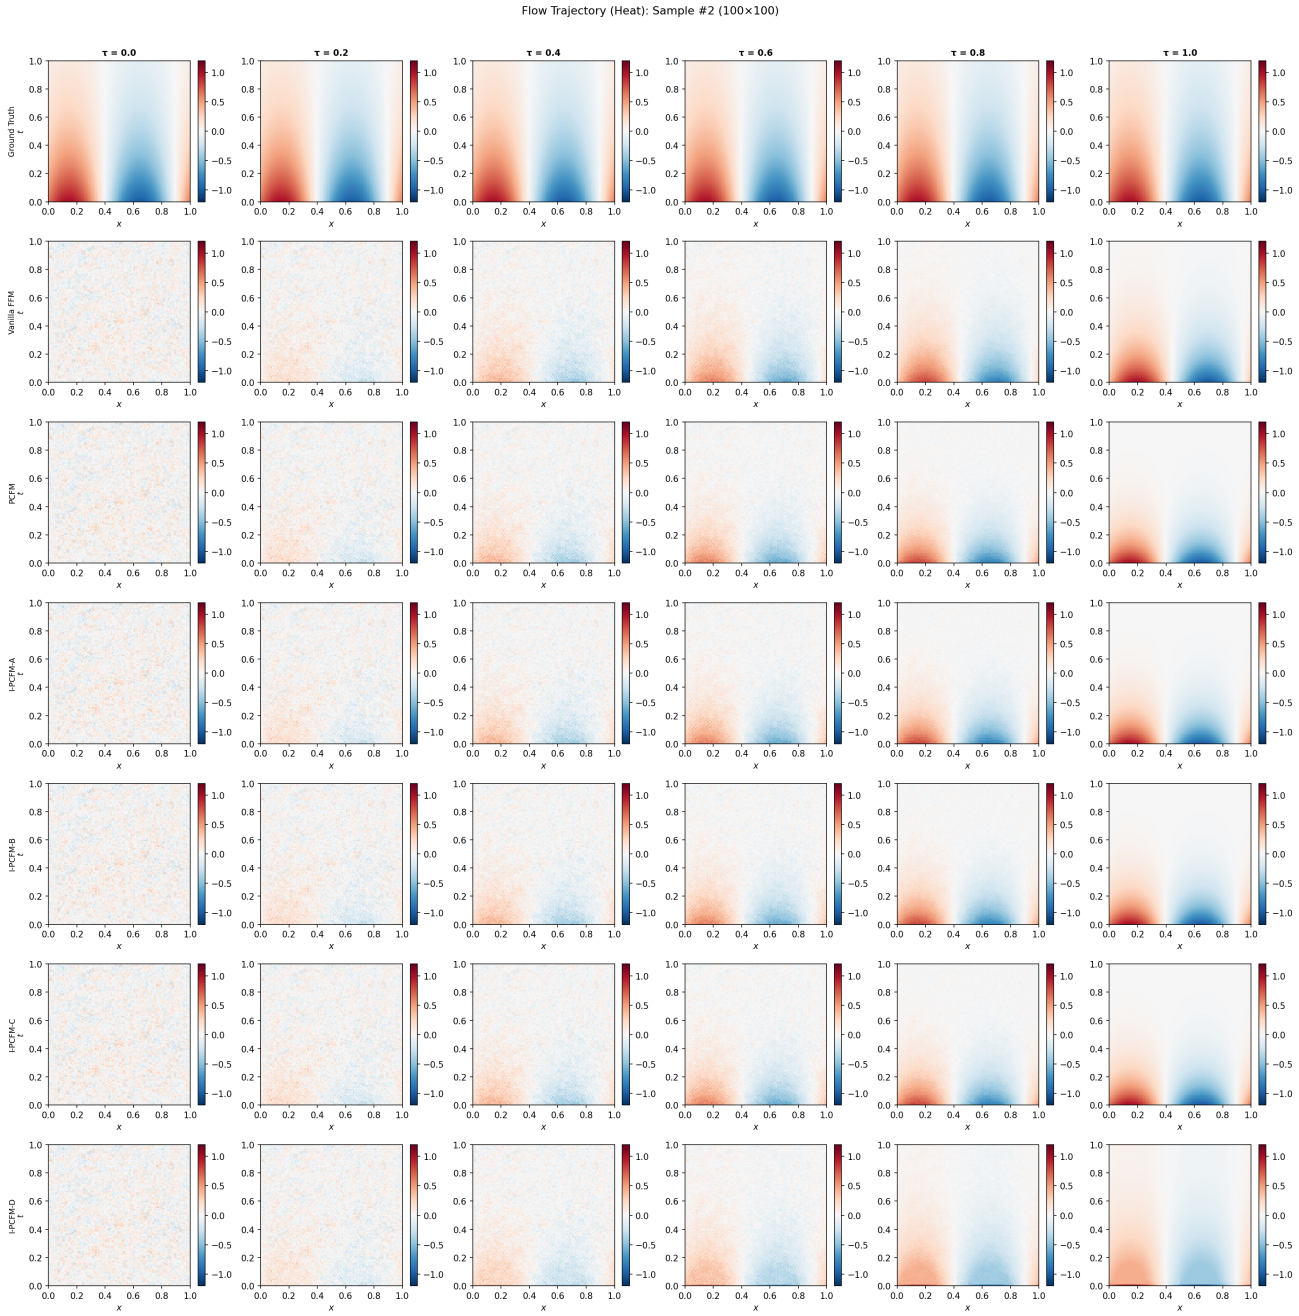

Figure 4. Flow-time evolution of an example Heat trajectory.

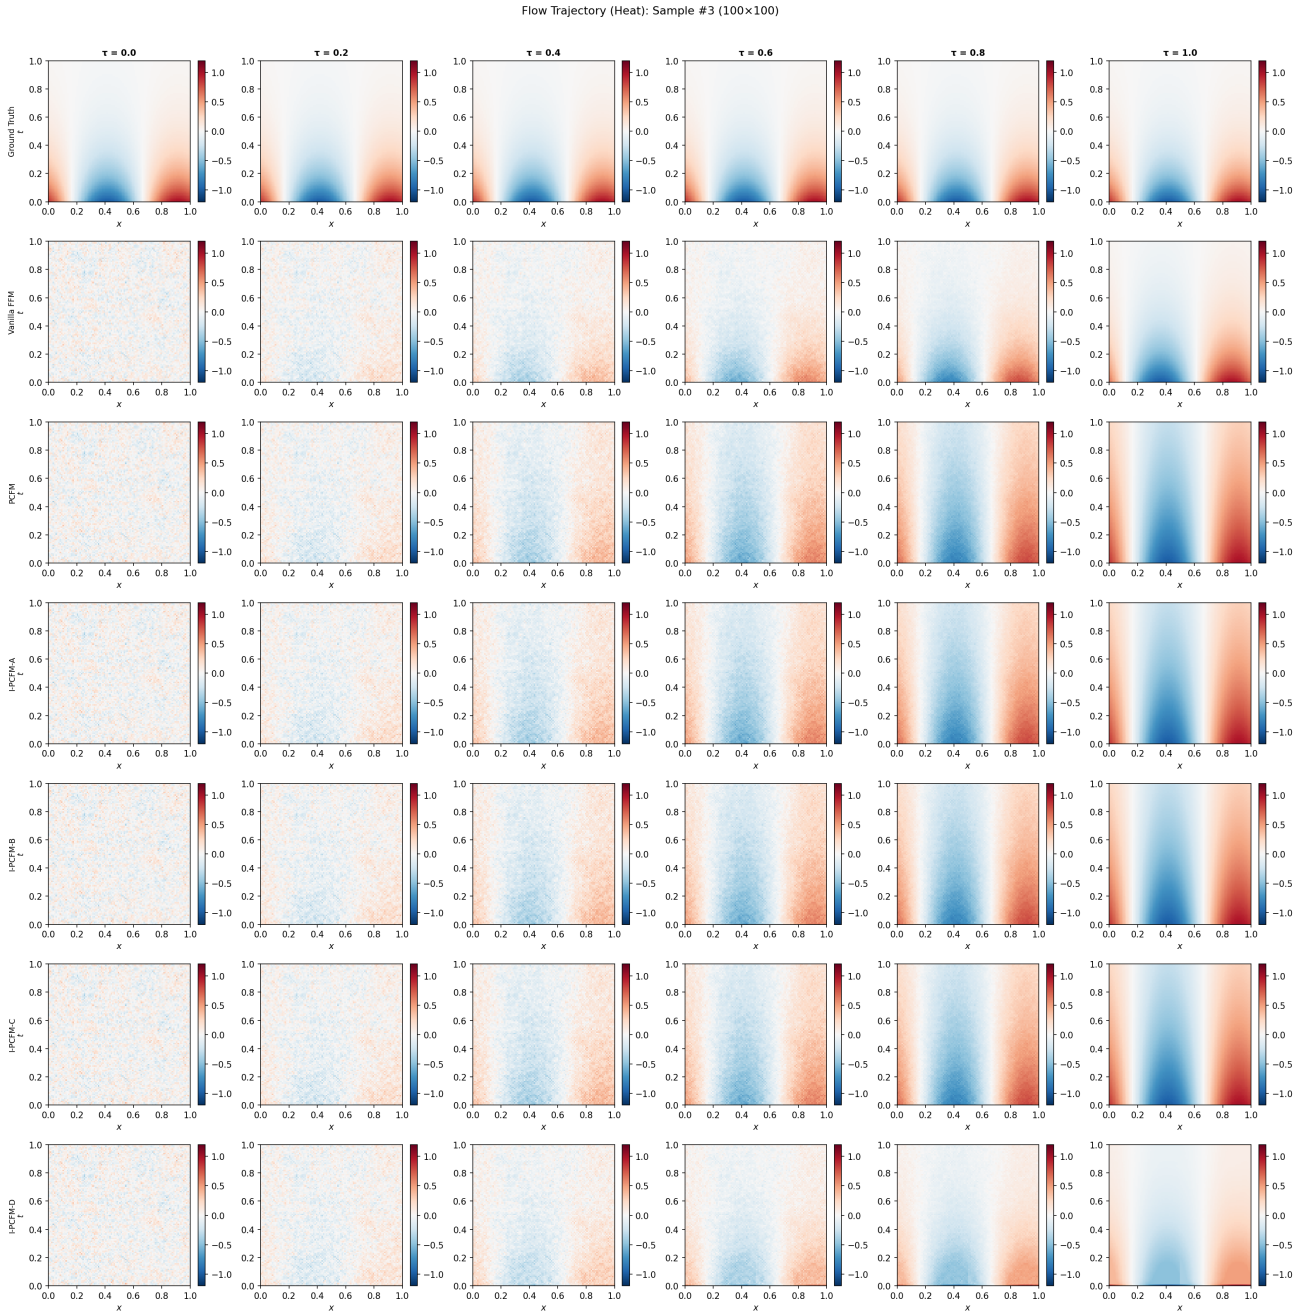

Figure 5. Flow-time evolution of an example Heat trajectory. Constraint-enforcing methods achieve better performance visually, as compared to Vanilla FFM.

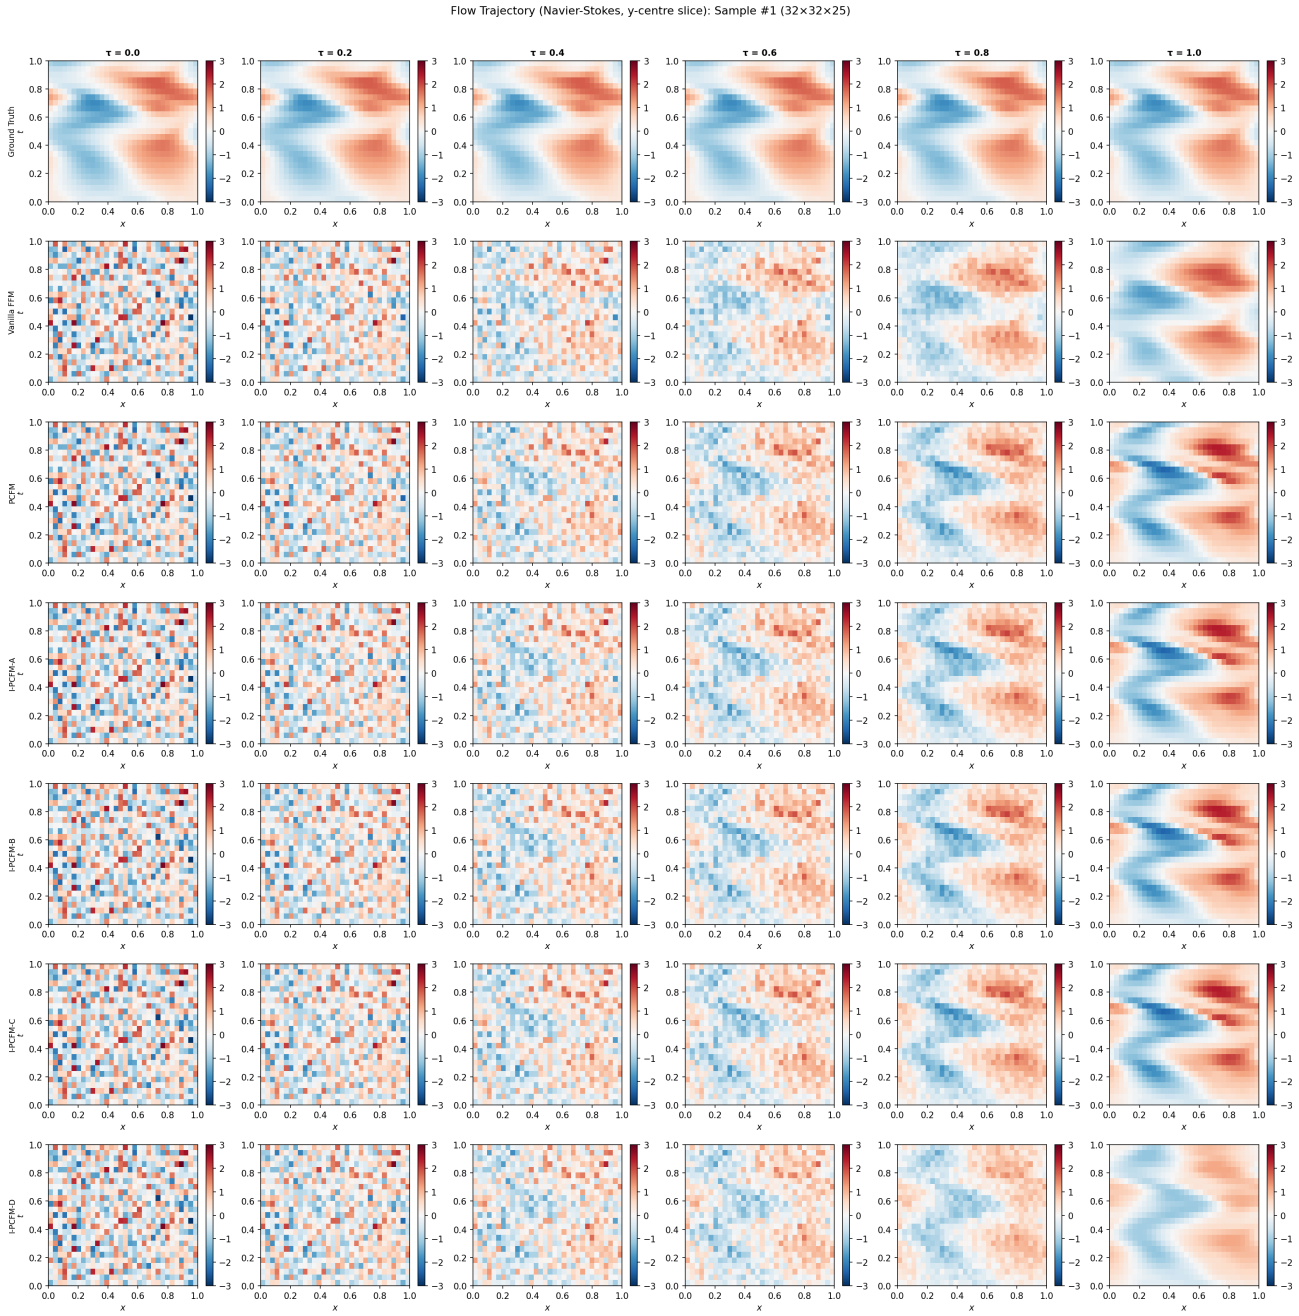

Figure 6. Flow-time evolution of an example Navier-Stokes trajectory.

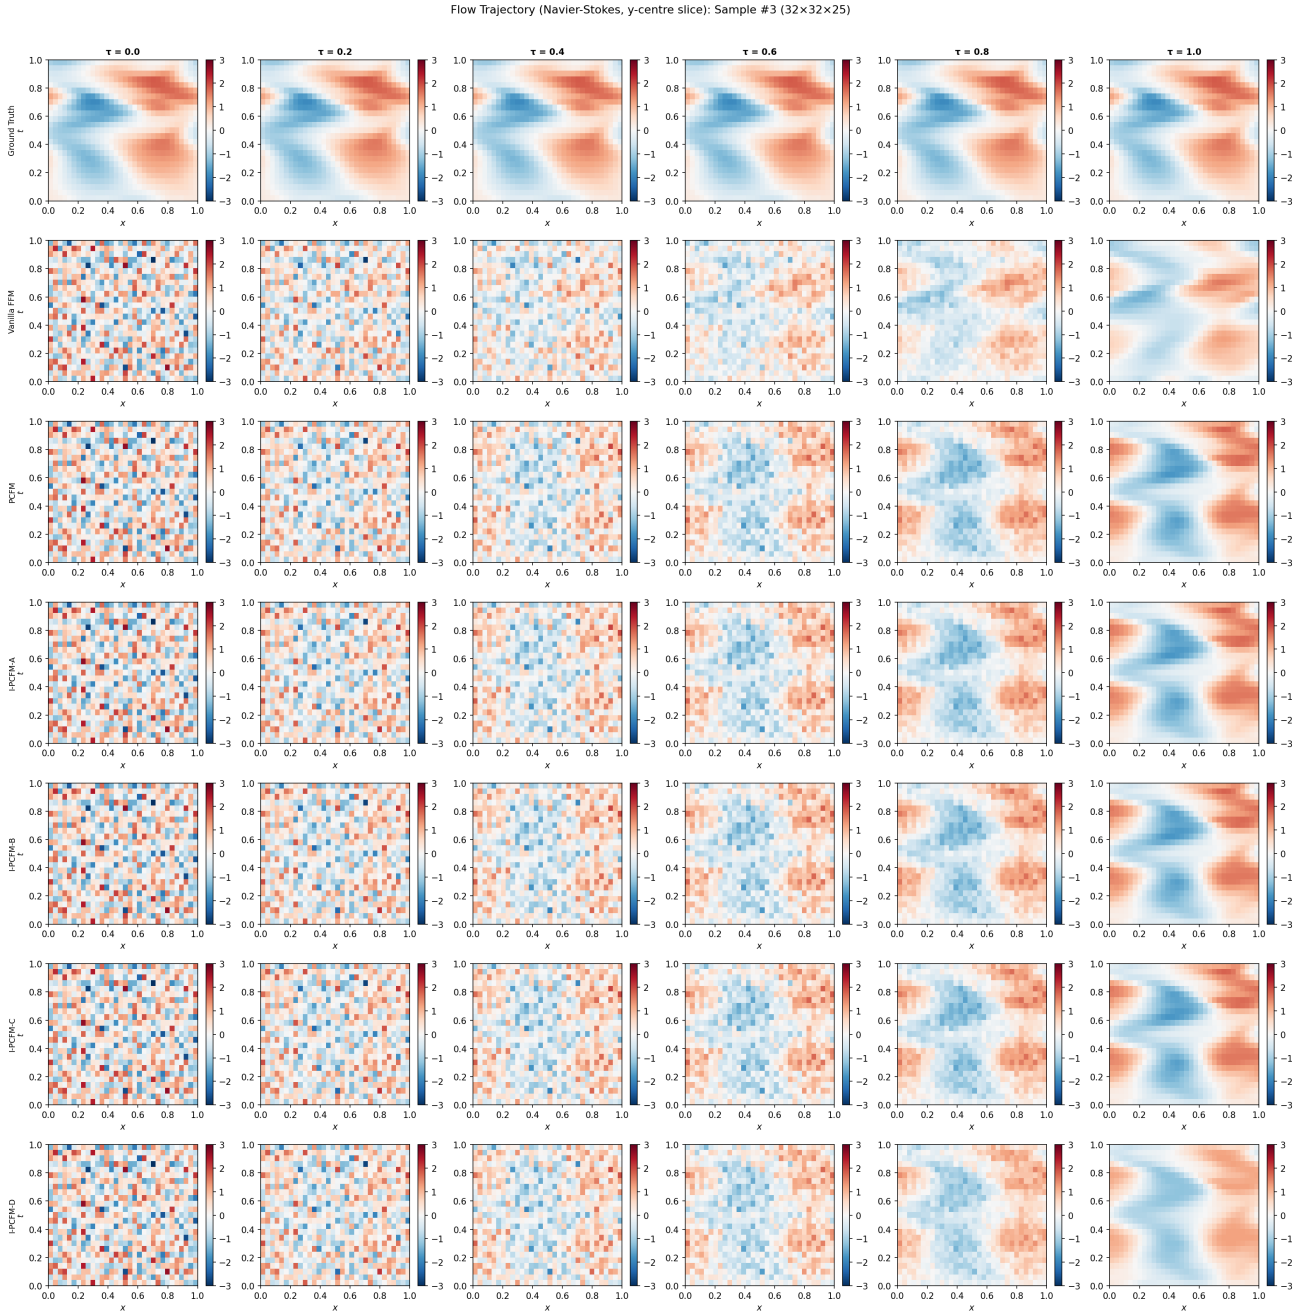

Figure 7. Flow-time evolution of an example Navier-Stokes trajectory. Constraint-enforcing methods achieve better performance visually, as compared to Vanilla FFM.

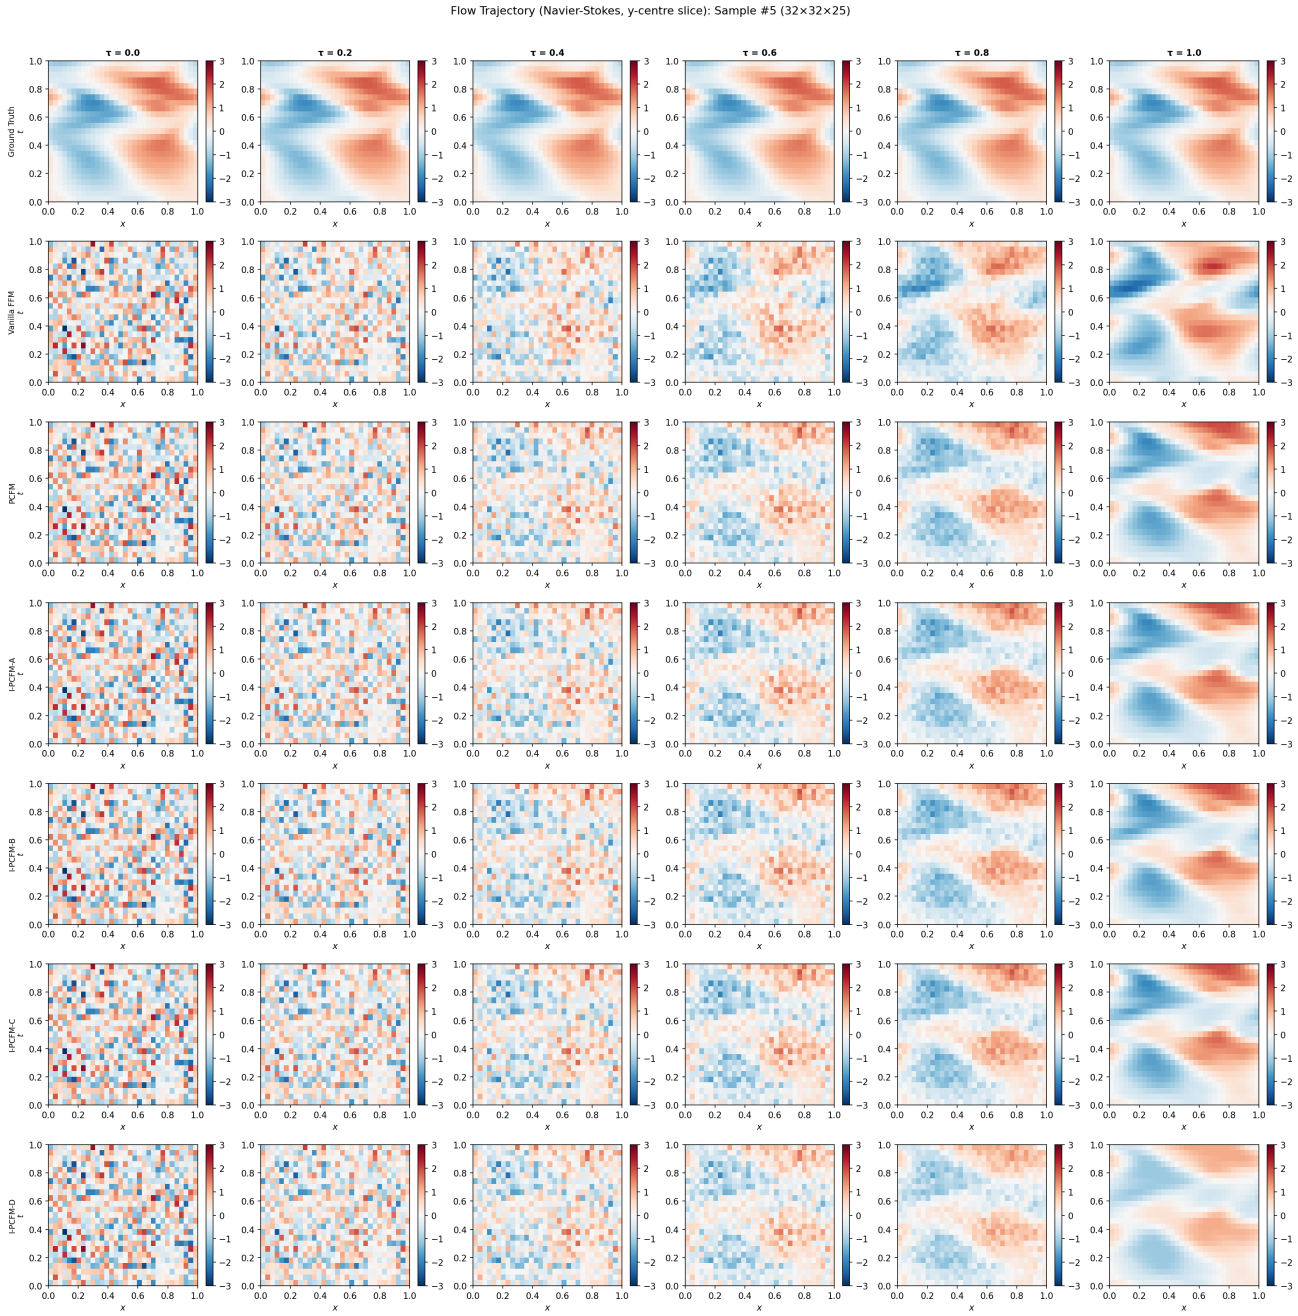

Figure 8. Flow-time evolution of an example Navier-Stokes trajectory.

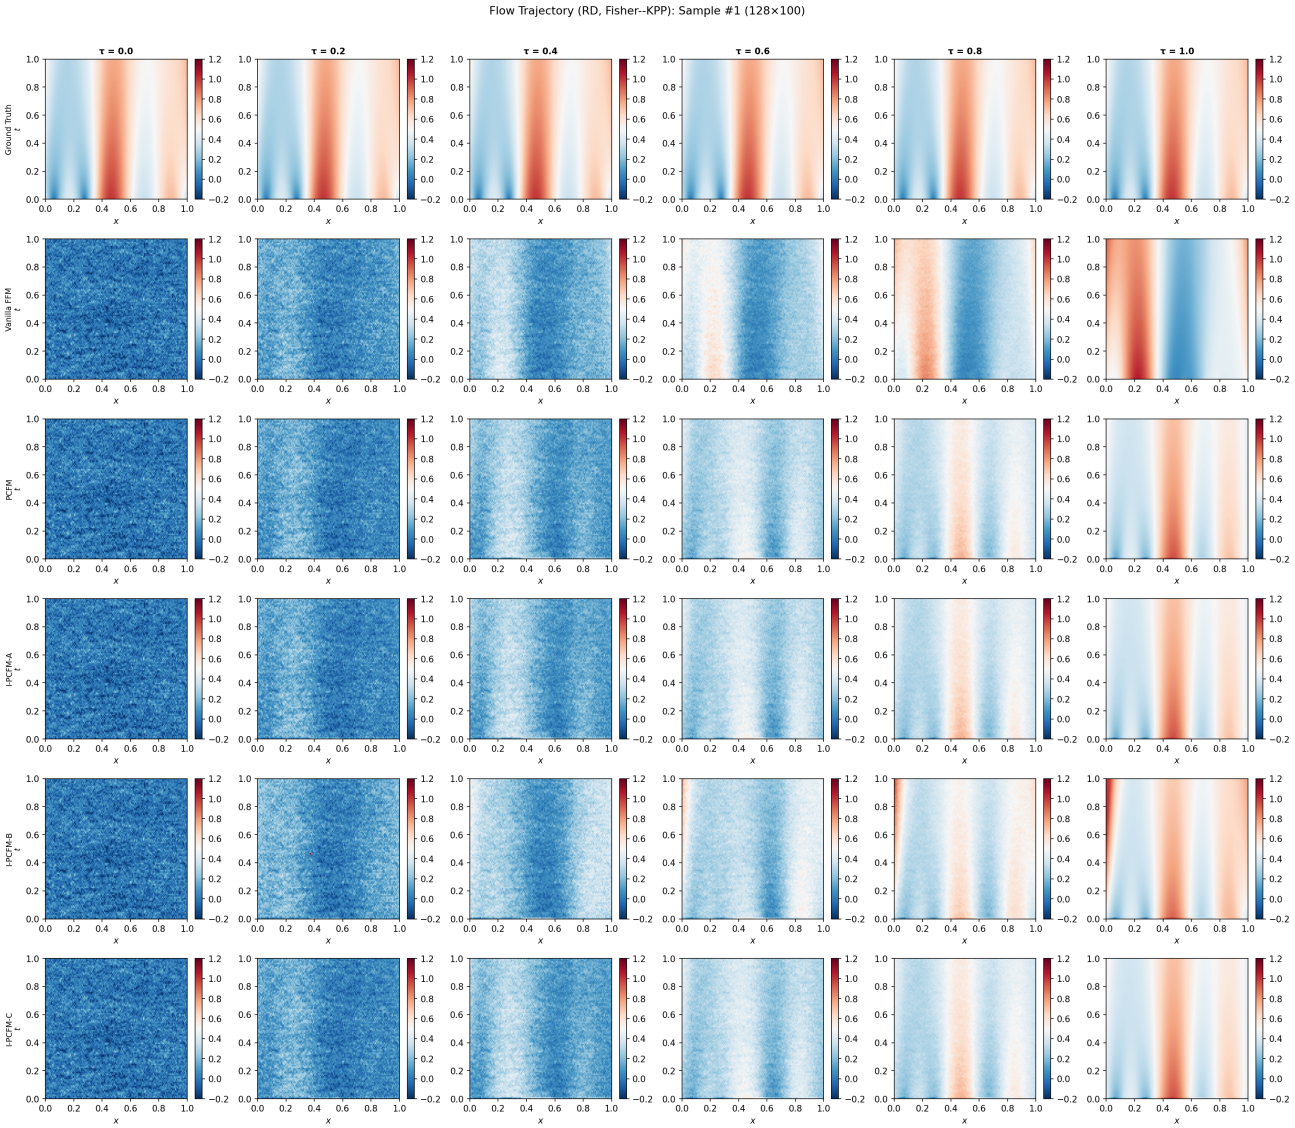

Figure 9. Flow-time evolution of an example Reaction-Diffusion (Fisher-KPP) trajectory. Constraint-enforcing methods achieve better performance visually, as compared to Vanilla FFM.

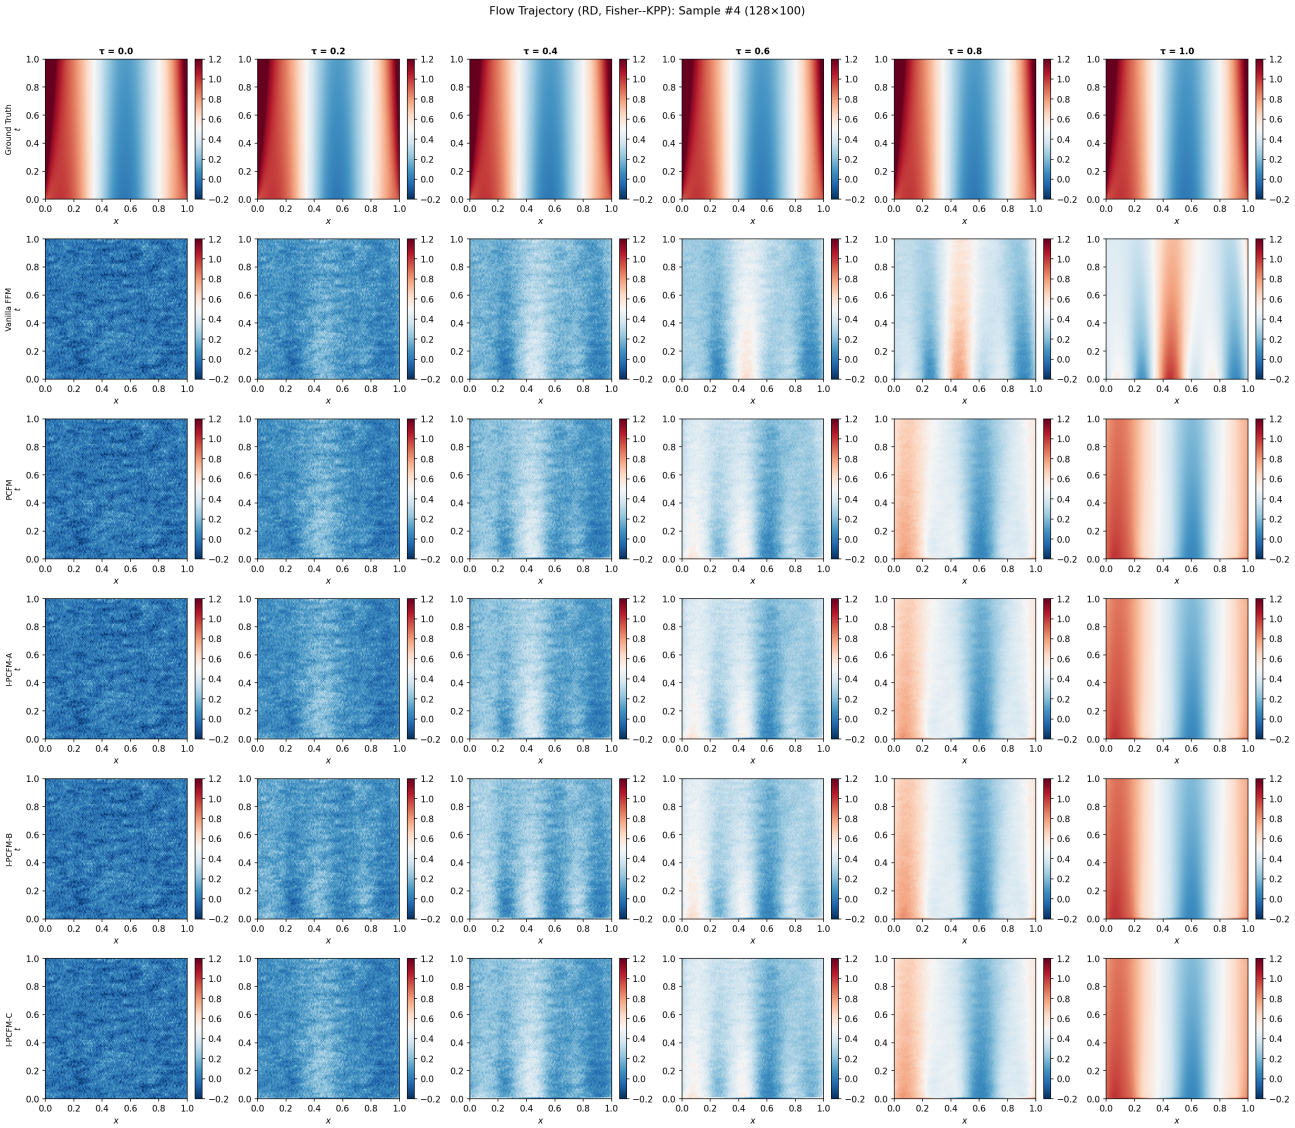

Figure 10. Flow-time evolution of an example Reaction-Diffusion (Fisher-KPP) trajectory. Constraint-enforcing methods achieve better performance visually, as compared to Vanilla FFM.

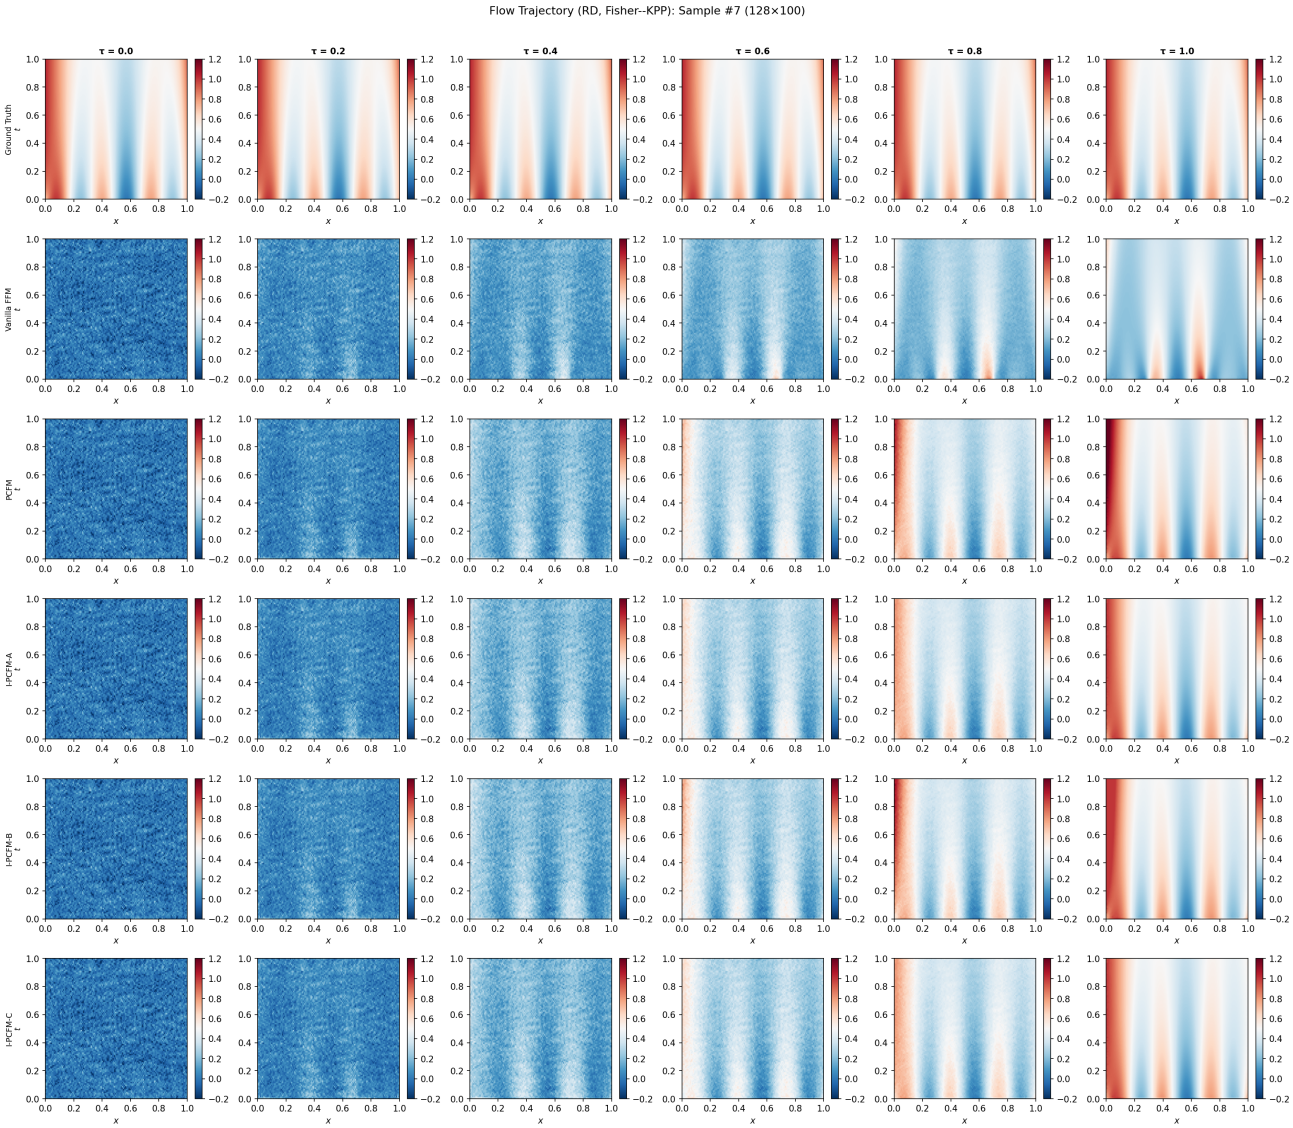

Figure 11. Flow-time evolution of an example Reaction-Diffusion (Fisher-KPP) trajectory. Constraint-enforcing methods achieve better performance visually, as compared to Vanilla FFM.

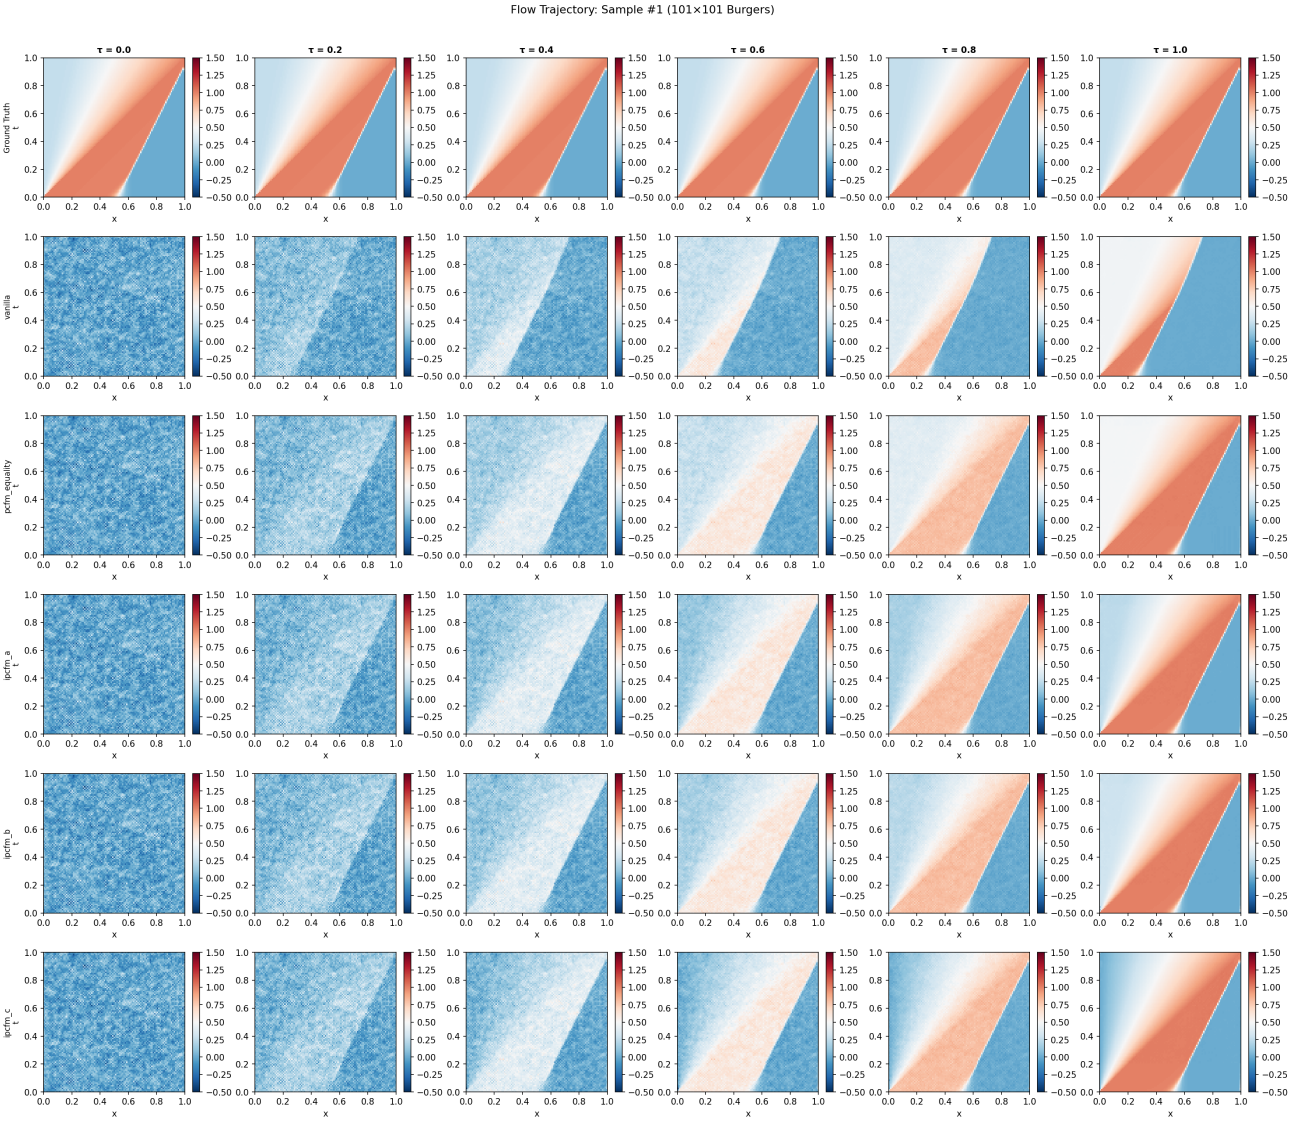

Figure 12. Flow-time evolution of an example Burgers' trajectory. I-PCFM methods improve quality of the generated sample compared to Vanilla FFM and PCFM.

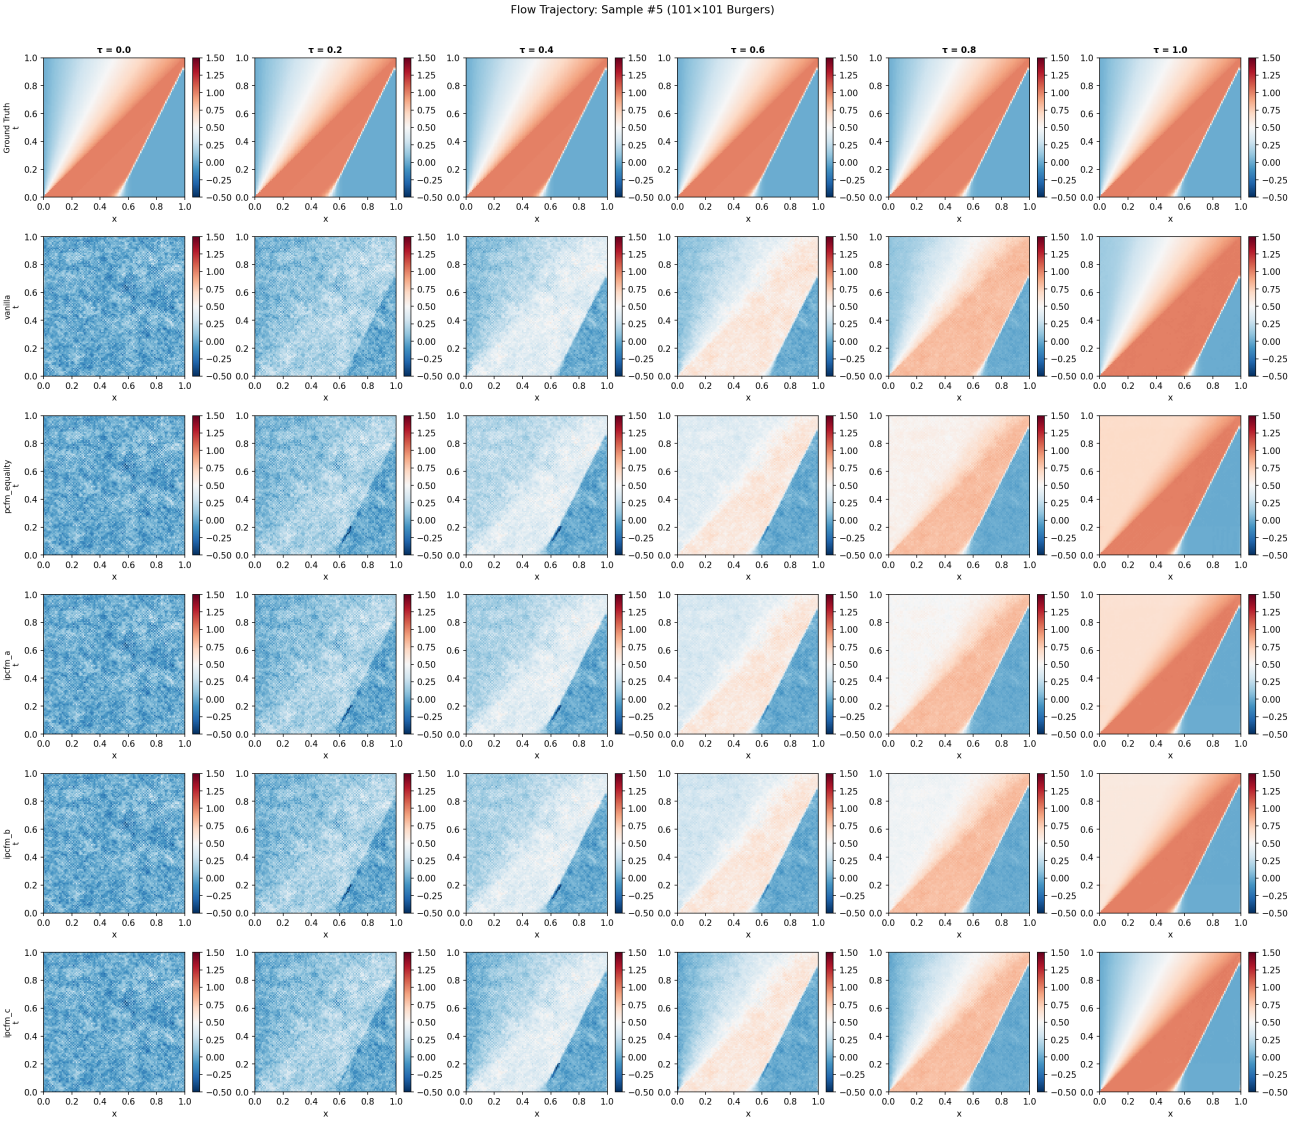

Figure 13. Flow-time evolution of an example Burgers' trajectory. Out of the constraint-enforcing methods, only I-PCFM-C's generated sample recovers the ground-truth faithfully.

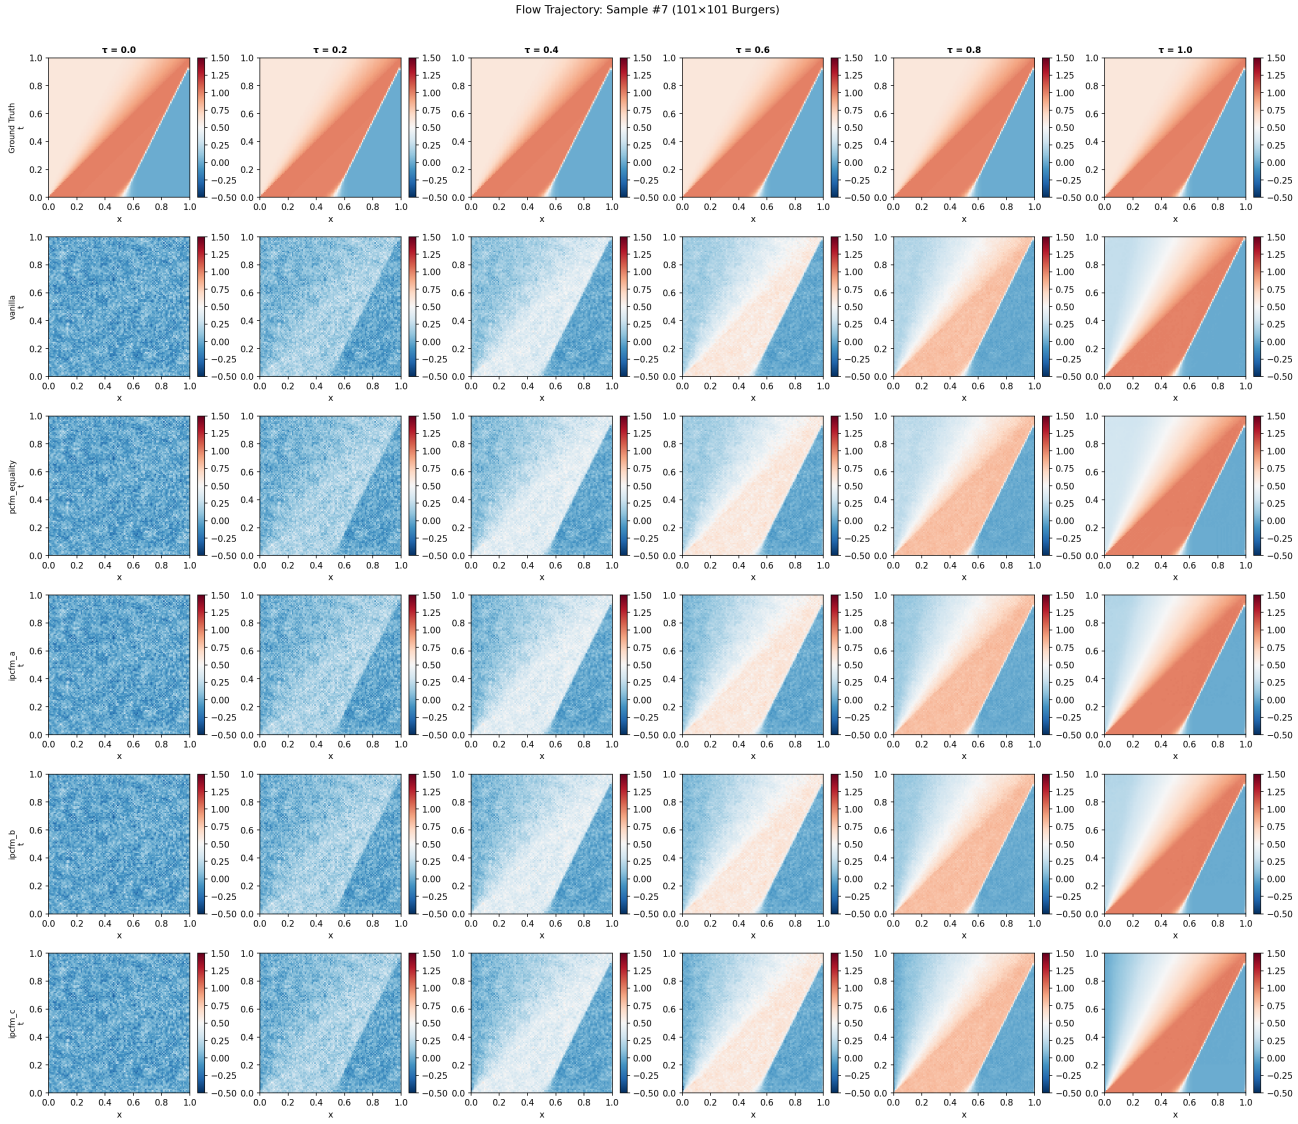

Figure 14. Flow-time evolution of an example Burgers' trajectory. In this example, all methods do not recover the ground-truth visually.