

LEARNING HEAT-BASED EQUATIONS IN SELF-SIMILAR VARIABLES

Shihao Wang

Department of Mathematics
University of Wisconsin–Madison
Madison, WI 53706, USA
swang2232@wisc.edu

Qipeng Qian

Program of Applied Mathematics
Department of Mathematics
University of Arizona
Tucson, AZ 85721, USA
qqian@arizona.edu

Jingquan Wang

Department of Mechanical Engineering
University of Wisconsin-Madison
Madison, WI 53706, USA
jwang2373@wisc.edu

ABSTRACT

We study solution learning for heat-based equations in self-similar variables (SSV). We develop an SSV training framework compatible with standard neural-operator training. We instantiate this framework on the two-dimensional incompressible Navier-Stokes equations and the one-dimensional viscous Burgers equation, and perform controlled comparisons between models trained in physical coordinates and in the corresponding self-similar coordinates using two simple fully connected architectures (standard multilayer perceptrons and a factorized fully connected network). Across both systems and both architectures, SSV-trained networks consistently deliver substantially more accurate and stable extrapolation beyond the training window and better capture qualitative long-time trends. These results suggest that self-similar coordinates provide a mathematically motivated inductive bias for learning the long-time dynamics of heat-based equations.

1 INTRODUCTION

Neural-network-based methods for solving PDEs have advanced rapidly in recent years. Physics-informed neural networks (PINNs) Raissi et al. (2019) enforce governing equations through the loss, while operator-learning architectures such as DeepONets and Fourier neural operators learn nonlinear maps between function spaces and perform well on parametric benchmarks Li et al. (2021); Lu et al. (2021); Takamoto et al. (2024); Li et al. (2023). However, for time-dependent problems, long-time extrapolation is often fragile: models trained on a finite horizon can accumulate error and suffer from distribution shift under rollout, leading to instability beyond the training window.

Existing approaches for improving long-horizon rollout stability include (i) observation-informed extrapolation that injects PDE structure or sparse measurements into the prediction pipeline Zhu et al. (2023); (ii) refinement/correction mechanisms that iteratively denoise, correct, or project predictions back to physically consistent states Lippe et al. (2023); Huang & Perdikaris (2025); Wei et al. (2025); (iii) learnable time-integration and continuous-time operator formulations that couple operator learning with numerical time stepping or temporally consistent surrogates Nayak & Goswami (2025); Abueidda et al. (2026); Diab & Al Kobaisi (2025); and (iv) rollout-aware training protocols (e.g., recurrent/self-rollout or curriculum schemes) that mitigate exposure bias by aligning training with inference-time dynamics Ye et al. (2025); Ghule et al. (2025). In this paper, however, we take a representation-level approach guided by physical intuition: we exploit the intrinsic scaling structure of the governing dynamics and work in variables that better reflect its long-time behavior. This viewpoint reshapes the learning problem into a form aligned with the system’s underlying physical principles, and is complementary to the above techniques.

In particular, we focus on “heat-based” evolution equations, i.e., systems whose linear part is the heat equation. Prototypical examples include the two-dimensional incompressible Navier-Stokes equations and the one-dimensional viscous Burgers equation. For such systems, long-horizon prediction is especially challenging because the characteristic length scale grows while the amplitude decays, causing the effective scales of the solution to change substantially over time. Since most neural-operator methods are formulated in physical coordinates (x, t) , training on a fixed spatial box and a finite time horizon can capture only a diminishing fraction of the relevant dynamics as t increases.

Self-similar variables (SSV) provide a natural remedy to this mismatch as classical theories Giga et al. (1988); Giga & Kambe (1988); Kato (1994); Gallay & Wayne (2002); Kim & Tzavaras (2001) show that the solutions admit a self-similar long-time behavior. Under the change of variables $\xi = x/\sqrt{t+1}$, $\tau = \log(t+1)$, the heat kernel becomes a stationary profile, and solutions that converge to a self-similar profile become localized with rapidly decaying tails in ξ . This suggests that learning in self-similar variables offers a clear spatial advantage: a fixed sampling window can capture the main part of the solution at large times.

The goal of this paper is to develop a precise self-similar-variable (SSV) formulation for heat-type equations that is compatible with neural-operator training, and to quantify the practical gains of learning in self-similar coordinates rather than in physical coordinates. To demonstrate that the approach applies across different heat-based PDEs, we specialize the SSV formulation to the two-dimensional incompressible Navier-Stokes equations and the one-dimensional viscous Burgers equation. To further show that the effect of the coordinate choice is architecture-independent, we present results using two simple fully connected models: a standard coordinate multi-layer perceptron (MLP) and a factorized fully connected network (FCN).

Our comparisons across architectures show a consistent picture: within the SSV framework, both the plain MLP and the factorized FCN produce significantly more accurate and stable extrapolation than their counterparts trained in physical coordinates. Moreover, we observe that SSV-trained models more reliably reproduce qualitative features that are intrinsic to the underlying dynamics—such as long-time physical patterns and, when present, coarsening/merging behaviors—whereas these effects are often degraded or missed by models trained directly in physical coordinates. This indicates that the main gain comes not from a sophisticated network design but from working in a coordinate system that is intrinsically adapted to the parabolic scaling and self-similar dynamics of the underlying equations. In this sense, our results provide evidence for the potential advantage of SSV-based training and point to a physically motivated strategy for learning solutions of other systems with self-similar asymptotics.

Code to reproduce the experiments is available at <https://github.com/Kagami-319/SSV-Learning>.

2 NETWORKS

In this section, we describe the two network architectures used throughout the study. As a baseline we use MLP, which takes the concatenated coordinate $z \in \mathbb{R}^m$ (either (x, t) or (ξ, τ)) and applies a single fully connected network $f_\theta : \mathbb{R}^m \rightarrow \mathbb{R}$,

$$\mathcal{M}_\theta(z) = f_\theta(z).$$

Our second model FCN is inspired by branch-trunk constructions. Let

$$B_\theta : \mathbb{R} \rightarrow \mathbb{R}^K, \quad T_\theta : \mathbb{R}^m \rightarrow \mathbb{R}^K$$

denote two multilayer perceptrons (the branch and trunk subnetworks). For a time-like coordinate $s \in \mathbb{R}$ (either t or τ) and a space-time coordinate $z \in \mathbb{R}^m$ (either (x, t) or (ξ, τ)), the FCN output is defined by

$$\mathcal{N}_\theta(s, z) = b_\theta + \sum_{k=1}^K B_{\theta,k}(s) T_{\theta,k}(z),$$

where $b_\theta \in \mathbb{R}$ is a bias and $B_{\theta,k}, T_{\theta,k}$ denote the k -th components of B_θ and T_θ .

3 LEARNING THE VORTICITY OF THE 2D NAVIER-STOKES EQUATION

3.1 SYSTEM AND SELF-SIMILAR VARIABLES

In this section, we consider the 2D incompressible Navier-Stokes equations in vorticity form on \mathbb{R}^2 with viscosity normalized to $\nu = 1$:

$$\partial_t \omega + u \cdot \nabla \omega = \Delta \omega, \quad u = K * \omega, \quad K(x) = \frac{1}{2\pi} \frac{x^\perp}{|x|^2}, \quad x^\perp = (-x_2, x_1). \quad (1)$$

The circulation $\Gamma := \int_{\mathbb{R}^2} \omega(x, t) dx$ is conserved. Noting that the scaling $\omega_\lambda(x, t) = \lambda^2 \omega(\lambda x, \lambda^2 t)$, $u_\lambda(x, t) = \lambda u(\lambda x, \lambda^2 t)$ leaves equation 1 invariant, we are inspired to introduce self-similar variables

$$\xi = \frac{x}{\sqrt{t+1}}, \quad \tau = \log(t+1), \quad \omega(x, t) = (t+1)^{-1} \Omega(\xi, \tau), \quad u(x, t) = (t+1)^{-1/2} U(\xi, \tau).$$

A direct change of variable yields

$$\boxed{\partial_\tau \Omega + U \cdot \nabla_\xi \Omega = \mathcal{L} \Omega, \quad U = K * \Omega,} \quad \mathcal{L} := \Delta_\xi + \frac{1}{2} \xi \cdot \nabla_\xi + \text{Id}. \quad (2)$$

The Oseen vortex $G(\xi) = \frac{1}{4\pi} e^{-|\xi|^2/4}$ satisfies $\mathcal{L}G = 0$; moreover $U^G \cdot \nabla G \equiv 0$ by orthogonality of the tangential U^G and radial ∇G , hence $\Omega \equiv \alpha G$ is a steady solution of equation 2 for any circulation α .

3.2 TRAINING IN THE SELF-SIMILAR FRAMEWORK

It was shown in Gallay & Wayne (2002) that, starting from a L^1 initial data, the vorticity will converge to the Gaussian at large time:

Proposition 1 (Long-time Behavior). *Fix $\mu \in (0, \frac{1}{2})$. There exist positive constants r_2 and C such that, for any initial data Ω_0 with $\|\Omega_0\|_2 \leq r_2$, the solution $\Omega(\cdot, \tau)$ of equation 2 satisfies:*

$$\left\| \Omega(\cdot, \tau) - \left(\int_{\mathbb{R}^2} \Omega(\xi) d\xi \right) G(\cdot) \right\|_2 \leq C e^{-\mu\tau}, \quad \tau \geq 0. \quad (3)$$

Thus, by Proposition 1, we shall see

$$|\Omega(\xi, \tau)| \lesssim e^{-c|\xi|^2} \quad \text{for some } c > 0,$$

uniformly for large τ . Therefore the self-similar solution is localized in ξ , and the error outside a moderate self-similar radius decays exponentially. This motivates measuring the approximation error on a fixed self-similar window and ignoring the exponentially small tails. Fix $C > 1$ and denote the 2D window

$$D_C := \{\xi \in \mathbb{R}^2 : |\xi| \leq C\}.$$

For the error field $\mathcal{E}_{\text{ssv}}(\xi, \tau) = \Omega(\xi, \tau) - \widehat{\Omega}(\xi, \tau)$ we define the windowed L^2 error as

$$E_{\text{ssv}}(C; \tau) := \int_{D_C} |\mathcal{E}_{\text{ssv}}(\xi, \tau)|^2 d\xi. \quad (4)$$

3.3 PHYSICAL-SPACE WINDOW (EXPANDING DISK)

In physical coordinates the vorticity field $\omega(x, t)$ spreads diffusively: for the linear heat equation the fundamental solution is the 2D Gaussian

$$G_t(x) = \frac{1}{4\pi t} \exp\left(-\frac{|x|^2}{4t}\right),$$

whose effective support has radius of order \sqrt{t} . For general 2D incompressible Navier-Stokes vorticity, the mild formulation

$$\omega(t) = e^{t\Delta} \omega_0 - \int_0^t e^{(t-s)\Delta} (u \cdot \nabla \omega)(s) ds$$

shows that each Duhamel term is obtained by convolving data with the same heat kernel. Thus, it is natural to regard the core of $\omega(\cdot, t)$ as living on a disk of radius comparable to $\sqrt{t+1}$, and the contribution of the far field as exponentially small in $|x|^2/(t+1)$. This picture is perfectly consistent with the self-similar formulation where we restrict the sampling area inside a fixed disk.

To compare self-similar and physical training in a fair way, we choose the physical window so that it is exactly the image of D_C under the self-similar change of variables. For any $t \geq 0$ and $C > 0$ we define

$$D_{t,C} := \{x \in \mathbb{R}^2 : |x| \leq C\sqrt{t+1}\},$$

which is the expanding disk obtained from D_C by the map

$$x = \sqrt{t+1} \xi. \quad (5)$$

If ξ is sampled uniformly from D_C , then x is sampled uniformly from $D_{t,C}$; the Jacobian of equation 5 is a constant factor $(t+1)$ and does not depend on the direction of ξ . Consequently, the same batch of self-similar samples $\{\xi_j\}_{j=1}^M$ used to train the SSV model induces, via equation 5, a batch of physical points $\{x_j\}_{j=1}^M$ that are uniformly distributed on $D_{t,C}$.

For the physical error field $\mathcal{E}_{\text{phy}}(x, t) := \omega(x, t) - \widehat{\omega}(x, t)$, we therefore measure the continuous windowed L^2 error by

$$E_{\text{phy}}(C; t) := \int_{D_{t,C}} |\mathcal{E}_{\text{phy}}(x, t)|^2 dx. \quad (6)$$

3.4 FAIR COMPARISON PROTOCOL

To compare SSV and physical training in a coordinate-agnostic way we adopt the following protocol.

- 1. Architecture parity.** The SSV and physical models share the same network architecture. For each architecture, we train two independent instances—one in self-similar coordinates and one in physical coordinates—while keeping all hyperparameters identical, including depth, width, optimizer, learning-rate schedule, training time window, batch size, and the total number of training steps.
- 2. Sampling (one-to-one between coordinates).** We draw i.i.d paired samples $\{(\xi_i, \tau_i)\}_{i=1}^M$ by Monte Carlo sampling on the training window

$$\tau_i \sim \text{Unif}[\log(1 + t_{\min}), \log(1 + t_{\max})], \quad \xi_i \sim \text{Unif}(D_C),$$

and then set the corresponding physical variables by $t_i = e^{\tau_i} - 1$ and $x_i = \sqrt{t_i+1} \xi_i$. The linear scaling $\xi \mapsto x$ sends the uniform distribution on D_C to the uniform distribution on $D_{t,C}$. Hence, for each sampled pair (ξ_i, τ_i) used by the SSV model, there is a one-to-one matched physical sample (x_i, t_i) seen by the physical model.

- 3. Loss functions.** Our target is to minimize the *time-averaged windowed* mean-squared errors

$$L_{\text{ssv}}(C) := \mathbb{E}_{\tau} [E_{\text{ssv}}(C; \tau)], \quad L_{\text{phy}}(C) := \mathbb{E}_{\tau} [E_{\text{phy}}(C; t)].$$

In practice, we approximate the above expectations by

$$\widehat{L}_{\text{ssv}} := \frac{1}{M} \sum_{i=1}^M |\widehat{\Omega}(\xi_i, \tau_i) - \Omega(\xi_i, \tau_i)|^2, \quad \widehat{L}_{\text{phy}} := \frac{1}{M} \sum_{i=1}^M |\widehat{\omega}(x_i, t_i) - \omega(x_i, t_i)|^2.$$

- 4. Extrapolation metric (relative MSE).** During extrapolation, we quantify predictive accuracy using the relative mean squared error (MSE) evaluated in the physical domain $D_{t,C}$. Concretely, for a predicted field $\widehat{\omega}$ and the reference solution ω at time t , we report

$$\text{RelMSE}(t) := \frac{\mathbb{E}[(\widehat{\omega}(t, x) - \omega(t, x))^2]}{\mathbb{E}[\omega(t, x)^2]}, \quad (7)$$

where the expectation is taken with respect to the uniform distribution on $D_{t,C}$ (approximated by averaging over the evaluation grid points).

3.5 EXPERIMENTS

Our experiments are designed to evaluate the practical advantage of training in SSV for extrapolation. We first compare SSV-based and physical-coordinate training under a fixed protocol, and then replicate the same study with two different network architectures to show that any observed gain is model-independent.

We use the same initial condition throughout, given by a superposition of two Gaussians. More precisely, we set

$$\omega_0(x, y) = A_1 \exp\left(-\frac{(x - x_1)^2 + (y - y_1)^2}{\sigma_1^2}\right) + A_2 \exp\left(-\frac{(x - x_2)^2 + (y - y_2)^2}{\sigma_2^2}\right),$$

with $A_1 = 1$, $A_2 = 0.6$, $(x_1, y_1) = (-1.5, 0.5)$, $(x_2, y_2) = (1.0, -0.8)$, $\sigma_1 = 1$, and $\sigma_2 = 1.3$, which yields an essentially non-radial initial vorticity field.

The long-time behavior for such initial data is clear. By Proposition 1, as t increases, the two Gaussian components move toward each other and eventually “merge” into a single-Gaussian profile in the long run. After the solution enters this merged regime, learning in SSV becomes comparatively easy, since the target dynamics has largely stabilized in SSV and is close to a time-independent profile. For this reason, we choose the training window $t \in [0, 0.3]$, terminating before the “merging” completes. In this way, extrapolation genuinely tests whether the model has learned the structural feature of the solution, and therefore produces physically faithful long-time predictions, rather than simply fitting the near-steady merged regime.

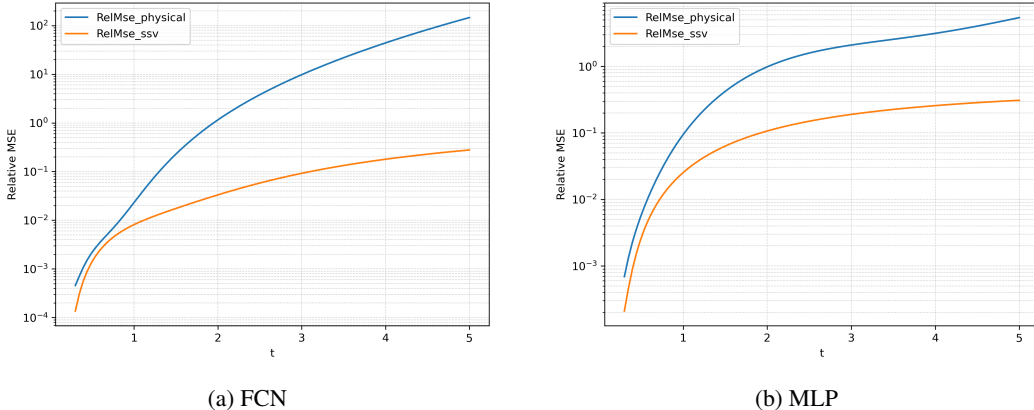


Figure 1: Relative MSE over $t \in [0.3, 5]$ under two architectures.

Figure 1 shows the relative MSE for $t \in [0.3, 5]$ using FCN and MLP respectively; Figure 2 & 3 represent some chosen examples during the extrapolation ($t = 0.5, 1, 1.5$) for the FCN and MLP respectively, where each row corresponds to a fixed time out of the training interval, and the columns display the reference vorticity $\omega(t, x, y)$, the prediction of the physical head, and the prediction of the self-similar head (which has been mapped back to physical variables for visual comparison).

We shall see that, for both architectures and across both short- and long-horizon extrapolation, the self-similar head consistently delivers markedly more accurate predictions than the physical head. More importantly, the SSV-trained models clearly capture the physically relevant “merging” trend. In contrast, the physical head fails to reproduce this qualitative transition, instead developing pronounced artifacts whose discrepancies grow more evident as t increases. Together, these observations support our expectation that learning in the self-similar framework better captures the structure of the solution and reflects the long-time dynamics of the underlying system.

4 VISCOUS BURGERS

In this section, we check the advantages of the SSV framework for the one-dimensional viscous Burgers equation which is as well heat-based, but has a much simpler nonlinearity.

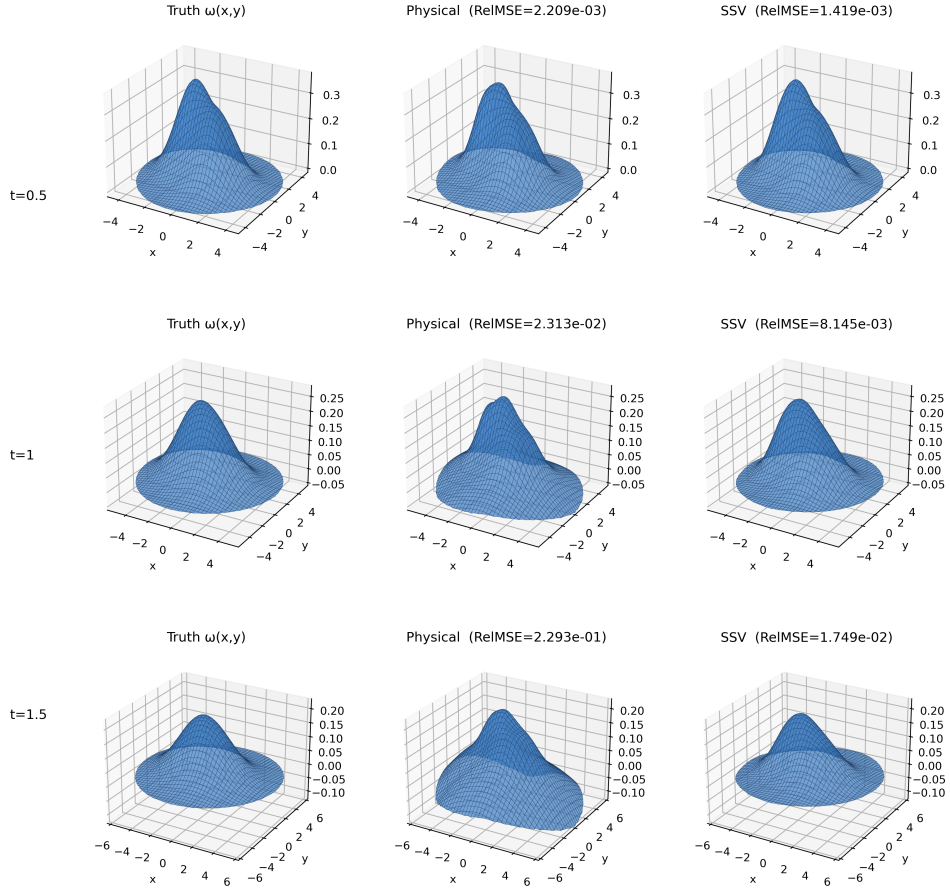


Figure 2: Example of extrapolation to $t = 0.5, 1, 1.5$ under the FCN: truth (left), physical head (middle), and self-similar head mapped to physical space (right).

4.1 EQUATION SETUP

We work on the whole line $x \in \mathbb{R}$ with viscosity normalized to 1 and consider the standard viscous Burgers equation

$$\partial_t u + \frac{1}{2} \partial_x (u^2) = \partial_{xx} u, \quad x \in \mathbb{R}, t > 0. \quad (8)$$

To align this example with the SSV framework of Section 3, we use the same time change

$$\xi = \frac{x}{\sqrt{t+1}}, \quad \tau = \log(t+1), \quad (9)$$

and introduce the self-similar unknown ω by

$$u(x, t) = (t+1)^{-1/2} \omega(\xi, \tau).$$

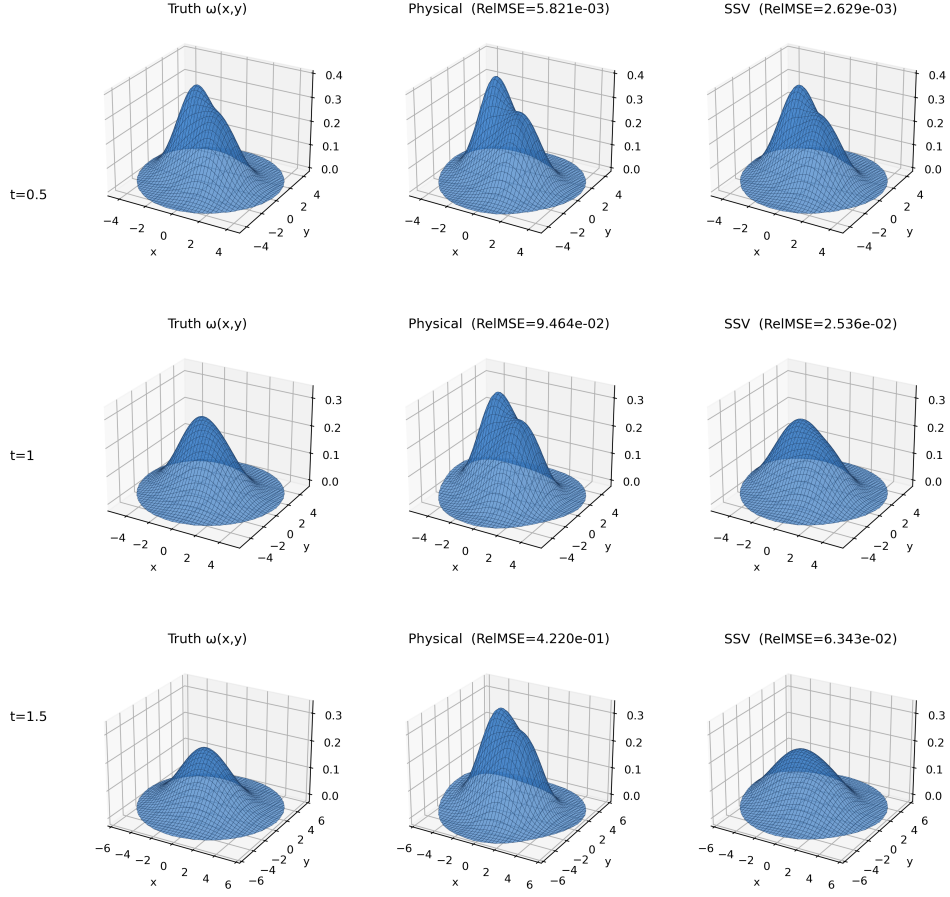
A direct computation shows that ω solves the SSV Burgers equation

$$\partial_\tau \omega + \omega \partial_\xi \omega = \partial_{\xi\xi} \omega + \frac{1}{2} \xi \partial_\xi \omega + \frac{1}{2} \omega. \quad (10)$$

It is well known that, for integrable initial data with nonzero mass, the long-time asymptotics of viscous Burgers are described by the diffusion wave.

Proposition 2 (Long-time Behavior; Kim & Tzavaras (2001)). *Let ω be the solution to the Cauchy problem equation 10 with initial data $\omega_0 \in L^1 \cap L^\infty$, $\partial_\xi \omega_0 \in L^2$, and total mass $\int_{\mathbb{R}} \omega_0(\xi) d\xi = M_0$. Then*

$$\omega(\xi, \tau) \rightarrow \mathcal{G}_{M_0}(\xi) \quad \text{as } \tau \rightarrow \infty, \quad (11)$$

Figure 3: Example of extrapolation to $t = 0.5, 1, 1.5$ under the MLP.

a.e. and in $L^1(\mathbb{R})$, where \mathcal{G}_{M_0} is the diffusion wave given by

$$\mathcal{G}_M(\xi) = \frac{(1 - e^{-M/2}) e^{-\xi^2/4}}{1 - (1 - e^{-M/2}) \frac{1}{\sqrt{\pi}} \int_{-\infty}^{\xi/2} e^{-\zeta^2} d\zeta}, \quad M \in \mathbb{R}. \quad (12)$$

Our experiments therefore probe how well the networks capture this non-Gaussian self-similar profile.

4.2 SAMPLING WINDOWS AND TRAINING LOSS

The SSV and physical sampling windows are one-dimensional analogues of the disks used in Section 3. For a fixed $C > 0$ we set

$$I_C := \{\xi \in \mathbb{R} : |\xi| \leq C\}, \quad I_{t,C} := \{x \in \mathbb{R} : |x| \leq C\sqrt{t+1}\}.$$

Let $u^*(x, t)$ denote the reference solution obtained from a numerical solver, and let $u(x, t)$ and $\omega(\xi, \tau)$ be the predictions of the physical and SSV models, respectively. For each fixed time t we measure the windowed L^2 errors

$$E_{\text{phy}}(C; t) := \int_{I_{t,C}} |u^*(x, t) - u(x, t)|^2 dx,$$

$$E_{\text{ssv}}(C; \tau) := \int_{I_C} |e^{\tau/2} u^*(e^{\tau/2} \xi, e^\tau - 1) - \omega(\xi, \tau)|^2 d\xi.$$

In practice, we form a Monte Carlo training set by drawing i.i.d. paired samples $\{(t, x)\}_{j=1}^N$. Specifically, we sample physical time t_j uniformly on $[t_{\min}, t_{\max}]$ and sample spatial locations x_j by drawing indices uniformly from the reference spatial grid. For the SSV model, each sampled pair (x_j, t_j) is mapped to the corresponding self-similar coordinates. Similar to the Navier-Stokes case, our target is to minimize the time-averaged windowed mean-squared errors

$$L_{\text{SSV}}(C) := \mathbb{E}_t[E_{\text{SSV}}(C; \tau)], \quad L_{\text{phy}}(C) := \mathbb{E}_t[E_{\text{phy}}(C; t)],$$

which are approximated by

$$\widehat{L}_{\text{SSV}} := \frac{1}{N} \sum_{j=1}^N |\widehat{\omega}(\xi_j, \tau_j) - \omega(\xi_j, \tau_j)|^2, \quad \widehat{L}_{\text{phy}} := \frac{1}{N} \sum_{j=1}^N |\widehat{u}(x_j, t_j) - u(x_j, t_j)|^2.$$

4.3 EXPERIMENTS

The initial data is chosen to be a compactly supported bipolar box profile,

$$u(x, 0) = u_0(x) = \begin{cases} 1, & -1 < x < 0, \\ -1, & 0 < x < 1, \\ 0, & |x| \geq 1. \end{cases}$$

Both models are trained on the time window $t \in [0, 0.5]$. We deliberately keep this window short as we did in Section 3, so that in SSV the target remains genuinely time-dependent. Figure 1 reports the relative MSE over $t \in [0.5, 5]$. Figures 5 and 6 further visualize the corresponding extrapolation behavior. In Figures 5 and 6, each panel corresponds to an evaluation time $t \in \{1, 1.5, 2, 2.5\}$ beyond the training horizon. For each time, we plot the reference solution (truth) together with the predictions from the physical and SSV heads, using the same spatial interval for all curves.




Figure 4: Relative MSE over $t \in [0.5, 5]$ under two architectures.

We observe a consistent qualitative pattern, fully analogous to the two-dimensional Navier-Stokes experiments: The SSV framework preserves the shape of the viscous shock profile well beyond the training horizon while the physical framework shows increasingly pronounced distortions as t grows. In other words, when we fix the architecture and training budget, the SSV representation again yields significantly more reliable long-time extrapolation. This one-dimensional example confirms that the advantage of learning in self-similar variables is not an artifact of the Navier-Stokes system, but rather a robust feature of heat-based evolution equations.

5 CONCLUSION

The results in Sections 3 and 4 consistently indicate a robust advantage of SSV training. Models trained in self-similar coordinates not only extrapolate more accurately and stably, but also more

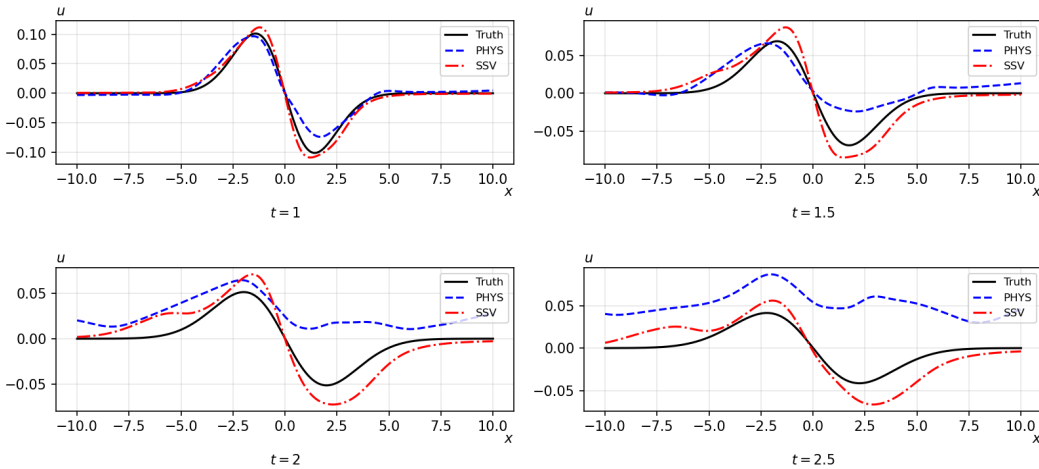


Figure 5: Comparison between the physical FCN and the SSV FCN at $t=1, 1.5, 2, 2.5$.

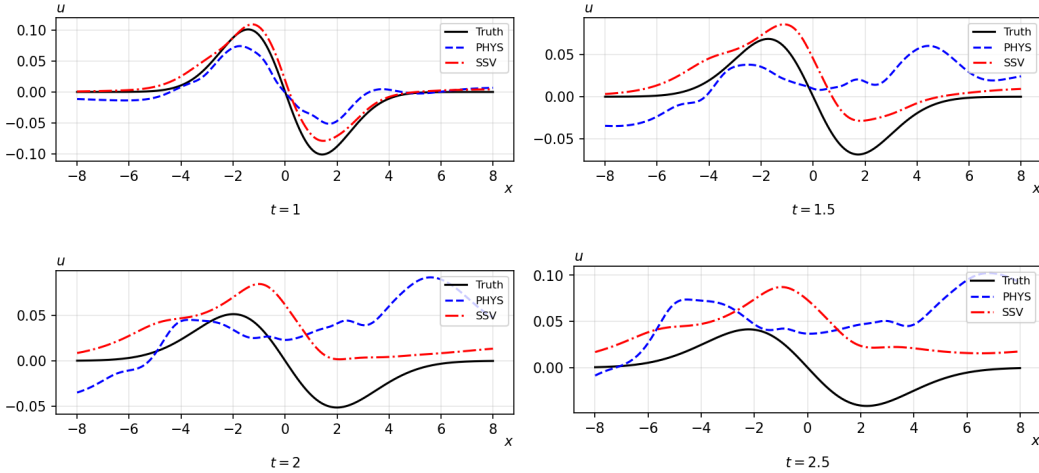


Figure 6: Same experiment as in Figure 5 for MLP.

reliably reproduce qualitative long-time features that are intrinsic to the underlying dynamics. Importantly, this improvement does not primarily come from architectural sophistication; rather, it is driven by an appropriate representation that aligns with the self-similar structure of heat-based evolution. In this sense, self-similar coordinates act as a mathematically motivated inductive bias.

More broadly, our findings suggest a general guideline for PDE operator learning: taking advantage of the physical intuition of the system can be more effective than further architectural engineering. From this perspective, SSV-based training provides a practical route to improved rollout stability and long-time prediction for heat-based systems, and the same strategy extends to other equations with self-similar asymptotics such as the porous medium equation $\partial_t u = \Delta u^m, m \in \mathbb{N}$, which admits self-similar spreading solutions (see Barenblatt (1996)).

Looking ahead, several directions appear natural: extending the framework to broader classes of systems with known scaling laws; developing adaptive windowing strategies for settings where asymptotic localization is weaker; combining SSV-based training with physics-based constraints to further improve robustness; and establishing theory that connects SSV-based sampling and coordinate choices to generalization and rollout stability in operator learning.

REFERENCES

- Diab W. Abueidda, Mbebo Nonna, Panos Pantidis, and Mostafa E. Mobasher. Time resolution independent operator learning. *Computer Methods in Applied Mechanics and Engineering*, 450: 118586, 2026. doi: 10.1016/j.cma.2025.118586.
- G. I. Barenblatt. *Scaling, Self-Similarity, and Intermediate Asymptotics*. Cambridge University Press, Cambridge, 1996.
- Waleed Diab and Mohammed Al Kobaisi. Temporal neural operator for modeling time-dependent physical phenomena. *Scientific Reports*, 15:32791, 2025. doi: 10.1038/s41598-025-16922-5.
- Thierry Gallay and C. Eugene Wayne. Invariant manifolds and the long-time asymptotics of the Navier–Stokes and vorticity equations on \mathbb{R}^2 . *Archive for Rational Mechanics and Analysis*, 163(3):209–258, 2002.
- Lalit Ghule, Akanksh Shetty, and Morgane Bourgeois. A discrete–continuous curriculum learning (dctl) framework for stable long-horizon pde surrogates. In *Machine Learning and the Physical Sciences Workshop, NeurIPS*, 2025.
- Yoshikazu Giga and Tsutomu Kambe. Large time behavior of the vorticity of two-dimensional viscous flow and its application to vortex formation. *Communications in mathematical physics*, 117(4):549–568, 1988.
- Yoshikazu Giga, Tetsuro Miyakawa, and Hirofumi Osada. Two-dimensional Navier-Stokes flow with measures as initial vorticity. *Arch. Rational Mech. Anal.*, 104(3):223–250, 1988. ISSN 0003-9527. doi: 10.1007/BF00281355. URL <https://doi.org/10.1007/BF00281355>.
- Xinquan Huang and Paris Perdikaris. Physicscorrect: A training-free approach for stable neural pde simulations. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2507.02227.
- Tosio Kato. The Navier-Stokes equation for an incompressible fluid in \mathbb{R}^2 with a measure as the initial vorticity. *Differential Integral Equations*, 7(3-4):949–966, 1994. ISSN 0893-4983.
- Yong Jung Kim and Athanasios E. Tzavaras. Diffusive n-waves and metastability in the burgers equation. *SIAM Journal on Mathematical Analysis*, 33(3):607–633, 2001.
- Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations, 2023. URL <https://arxiv.org/abs/2111.03794>.
- Phillip Lippe, Bastiaan S. Veeling, Paris Perdikaris, Richard E. Turner, and Johannes Brandstetter. Pde-refiner: Achieving accurate long rollouts with neural pde solvers, 2023. URL <https://arxiv.org/abs/2308.05732>.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- Dibyajyoti Nayak and Somdatta Goswami. Ti-deeponet: Learnable time integration for stable long-term extrapolation, 2025. URL <https://arxiv.org/abs/2505.17341>.
- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Dan MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning, 2024. URL <https://arxiv.org/abs/2210.07182>.

Hao Wei, Aleksandra Franz, Bjoern List, and Nils Thuerey. Inc: An indirect neural corrector for auto-regressive hybrid pde solvers. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2511.12764.

Zaijun Ye, Chen-Song Zhang, and Wansheng Wang. Recurrent neural operators: Stable long-term pde prediction. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2505.20721.

Min Zhu, Handi Zhang, Anran Jiao, George Em Karniadakis, and Lu Lu. Reliable extrapolation of deep neural operators informed by physics or sparse observations. *Computer Methods in Applied Mechanics and Engineering*, 412:116064, 2023.