14th International Symposium on Turbulence and Shear Flow Phenomena (TSFP14)
Heidelberg, Germany, July 28–31, 2026

IDENTIFYING EFFICIENT ROUTES TO LAMINARIZATION USING REDUCED ORDER MODELS AND CONSTRAINED OPTIMIZATION

Jake Buzhardt
Department of Chemical & Biological Engineering
University of Wisconsin – Madison
1415 Engineering Dr, Madison, WI, 53706
buzhardt@wisc.edu

Michael D. Graham
Department of Chemical & Biological Engineering
University of Wisconsin – Madison
1415 Engineering Dr, Madison, WI, 53706
mdgraham@wisc.edu

INTRODUCTION

Understanding the transition from laminar flow to turbulence and, conversely, the return from turbulence to laminar flow in wall-bounded shear flows remains one of the central challenges in fluid dynamics. The problem of transition to turbulence has been extensively studied from linear, nonmodal, and fully nonlinear perspectives, including the minimal-seed framework reviewed by Kerswell (2018). Relaminarization has received comparatively less attention, despite its clear relevance for flow control and drag reduction. A dynamical-systems perspective suggests that relaminarization should be governed by the geometry of the basin boundary separating turbulent and laminar behavior, and that trajectories which return most efficiently to laminar flow can reveal the structures organizing that boundary.

In the present work, we focus on the *minimal seed for relaminarization*: the smallest perturbation of a turbulent trajectory that crosses from the turbulent basin of attraction into the laminar basin. This object is the natural complement to the better-known minimal seed for transition. In contrast to the transition problem, however, the relaminarization problem is posed in a chaotic region of state space, making direct trajectory optimization substantially more challenging. To address this issue, we formulate the problem as a constrained nonlinear optimization and use the multi-step penalty approach of Chung & Freund (2022) to stabilize the optimization over chaotic trajectories. In recent work, we introduced the concept of the minimal seed for relaminarization, developed an optimization-based framework for computing it, and applied that framework to a low-dimensional model of transitional shear flow to identify minimal seeds and demonstrate its potential utility for control (Buzhardt & Graham, 2026).

The framework is first demonstrated in the nine-mode Moehlis–Faisst–Eckhardt (MFE) model of sinusoidal shear flow (Moehlis *et al.*, 2004). In this setting, we identify relaminarizing perturbations, show how the optimized trajectories cross the edge of chaos, and relate the resulting behavior to an edge periodic orbit. We then consider the broader goal motivating this work: extending these ideas to much higher-dimensional transitional shear flows. For this purpose, we present reduced-order modeling results for transitional plane Couette flow and plane Poiseuille flow using convolutional autoencoders and latent-space dynamics models. Together, these results illustrate both the constrained-optimization approach for identifying efficient routes to laminarization and the reduced-order modeling tools needed to extend that approach to DNS-informed systems.

The paper therefore has two aims. First, it presents a constrained-optimization framework for identifying efficient routes to laminarization in a reduced model where the full dynamical-systems analysis can be carried out. Second, it assesses the current readiness of learned ROMs of transitional wall-bounded shear flows for extending this program beyond low-dimensional test systems.

PROBLEM FORMULATION AND OPTIMIZATION METHOD

The central question considered here is: *What is the smallest perturbation of a turbulent trajectory that leads to relaminarization?* A schematic of the corresponding geometry in state space is shown in Fig. 1. Let $\mathbf{x}$ denote the state of the system, let $\mathbf{x}_T$ be a reference point on a turbulent attractor, and let d_T and d_L denote distances from the turbulent reference point and laminar state, respectively. For a fixed initial perturbation magnitude d_0, the first step is to solve

$$\begin{aligned} \min_{\mathbf{x}_0} \quad & d_L^2(\tau), \\ \text{s.t.} \quad & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}),\ \mathbf{x}(0) = \mathbf{x}_0,\ t \in (0,\tau), \\ & d_T(0) = d_0, \end{aligned} \tag{1}$$

where τ is a finite optimization horizon. For initial conditions inside the laminar basin of attraction, the optimized final distance $d_L(\tau)$ tends toward zero for sufficiently large τ. For initial conditions outside that basin, $d_L(\tau)$ remains finite. The second step is therefore to sweep over d_0 and identify the smallest perturbation magnitude for which the optimized trajectory relaminarizes.

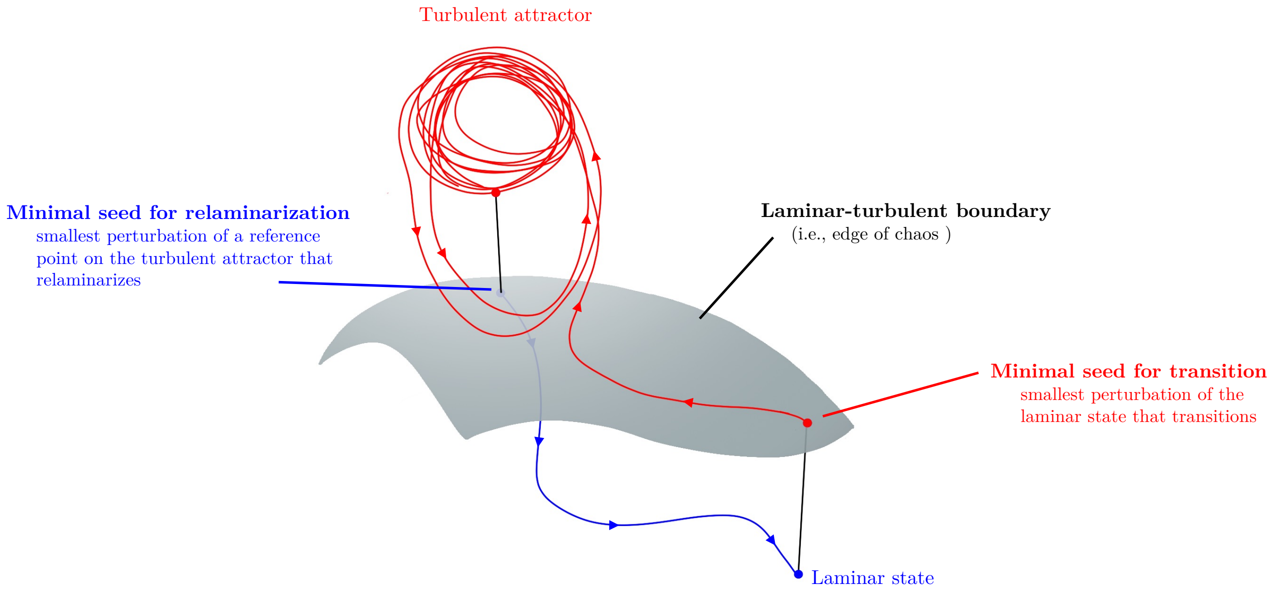

Figure 1. Schematic depiction of the minimal seed for transition, the minimal seed for relaminarization, and the laminar–turbulent basin boundary. This figure can likely be reused directly from the abstract.

The primary difficulty is that this optimization is performed in a chaotic region of state space. Standard adjoint-based optimization becomes ill-conditioned because sensitivities with respect to initial conditions grow exponentially in time. We therefore adopt the multi-step penalty method introduced by Chung & Freund (2022). In this approach, a long trajectory is broken into shorter segments, each of which is optimized more stably; continuity between successive segments is then imposed progressively through a quadratic penalty. This procedure regularizes the sensitivity growth sufficiently to enable gradient-based optimization over chaotic trajectories.

Several practical choices affect the quality of the computed minimal seed. First, the horizon τ must be long enough that laminarizing and non-laminarizing optimized trajectories are clearly distinguishable. Second, the choice of norm used in d_T and d_L determines how "closeness" is measured in state space and should be stated explicitly for each system. Third, once an approximate relaminarizing perturbation is found, local edge tracking can be used to locate nearby invariant solutions and to interpret the basin-boundary geometry associated with the computed seed.

MFE MODEL: MINIMAL SEEDS FOR RELAMINARIZATION

We first apply the framework to the nine-mode MFE model (Moehlis *et al.*, 2004), which provides a convenient testbed for studying transitional dynamics and basin-boundary structure. Figure 2 summarizes the main optimization result. The optimized terminal distance from laminar, $d_L(\tau)$, is shown as a function of the initial perturbation magnitude $d_T(0)$ for several optimization horizons. As discussed previously, a sharp decrease in optimized $d_L(\tau)$ signals that the edge of chaos has been crossed and therefore provides an estimate of the minimal seed for relaminarization. Increasing the optimization horizon makes this transition sharper and improves the separation between relaminarizing and non-relaminarizing trajectories.

Time traces for trajectories on either side of the basin boundary are also shown in Fig. 2. Starting from nearby initial conditions, one trajectory undergoes a direct return toward the laminar state, whereas the other remains transiently turbulent. Edge tracking initialized from this pair indicates convergence toward a periodic orbit on the edge. In the MFE system, this edge orbit appears to organize the separation between relaminarizing and non-relaminarizing behavior: trajectories remain close for an extended period and then depart along different directions associated with the stable and unstable manifolds of the edge state.

A complementary view is given in Fig. 3, where relaminarizing trajectories computed from several turbulent reference points are shown in an input–dissipation projection. These trajectories emphasize that the optimal route to laminarization need not be monotone in any simple measure of turbulence intensity. Instead, some trajectories initially move away from the laminar state or circulate around the turbulent region before ultimately relaminarizing. This counterintuitive behavior highlights why an optimization-based approach is needed: efficient laminarization pathways are not generally obvious from simple energetic intuition.

Taken together, these results show that the constrained-optimization framework can identify efficient relaminarizing perturbations in a reduced chaotic model and can be used to reveal the invariant structures controlling the associated basin-boundary geometry. In this way, the MFE system provides a clear demonstration of the relaminarization methodology and of the dynamical-systems mechanisms that underlie the optimized routes to laminar flow. For full details of the method and extended results on the MFE system, the reader is referred to (Buzhardt & Graham, 2026).

LEARNED REDUCED-ORDER MODELS FOR TRANSITIONAL SHEAR FLOWS

The broader objective of this work is to apply the minimal seed for relaminarization framework in higher-dimensional

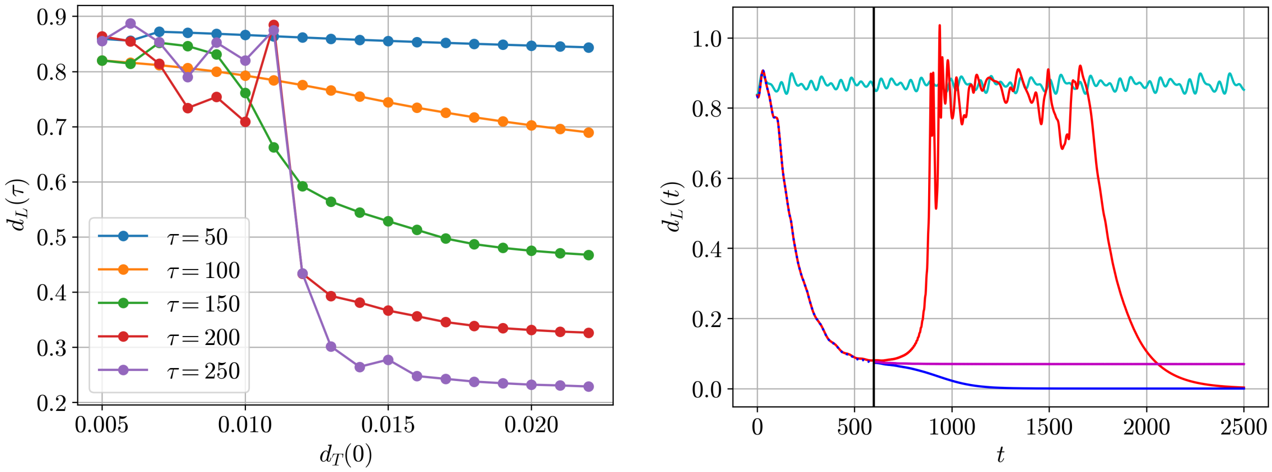

(a) (b)

Figure 2. Minimal seed for relaminarization: (a) Optimized final distance from the laminar state, $d_L(\tau)$ for varying initial perturbation magnitudes, $d_T(0)$ from the reference point on the turbulent attractor. The sharp decrease in $d_L(\tau)$ indicates that the basin boundary has been crossed. (b) Time series of the distance to laminar, d_L for two trajectories on either side of the basin boundary near the minimal seed for relaminarization (laminarizing ——, turbulent ——). Also shown for reference are trajectory on the chaotic attractor (——) and the edge periodic orbit (——), for $t > 600$.

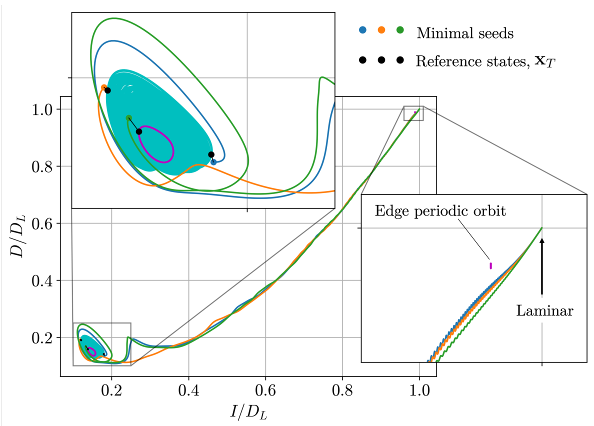

Figure 3. Minimal seeds for relaminarization and laminarizing trajectories from multiple reference points in the turbulent region depicted in the input-dissipation projection of the state-space, normalized by the laminar value, D_L. The minimal seed trajectories vary in the vicinity of the chaotic attractor, but eventually follow a common path to laminar, passing near the periodic orbit on the edge. Also shown are a trajectory on the chaotic attractor (——), a periodic orbit in the turbulent region, and the periodic orbit on the edge (——). Other colored curves show the laminarizing trajectory from each minimal seed. Insets show magnified views of the region around the turbulent attractor and the edge periodic orbit.

systems such as the dynamical systems resulting from direct numerical simulation of fluid flows. As a step toward this goal, we present progress toward the development of data-driven reduced-order models for transitional plane Couette flow and plane Poiseuille flow. These models provide low-dimensional representations of the flow fields and form the basis for extending the relaminarization optimization beyond reduced chaotic models such as MFE.

Data generation

The training data for the learned reduced-order models are obtained from three-dimensional direct numerical simulations performed in Dedalus, an open-source spectral PDE solver implemented efficiently in Python (Burns *et al.*, 2020). In both cases, the flow is represented on a Fourier–Chebyshev–Fourier grid, corresponding to periodic directions in the streamwise and spanwise coordinates and a Chebyshev discretization in the wall-normal coordinate. The reduced-order models are trained on snapshots of the mean-subtracted velocity field sampled on this DNS grid. Subtracting the mean flow focuses the representation on the fluctuating component of the dynamics, rather than allocating model capacity to the dominant stationary mean profile.

For plane Couette flow, the simulations are carried out in a domain of size $L_x = 1.75\pi$, $L_y = 2$, and $L_z = 1.2\pi$, using $48 \times 65 \times 48$ grid points in the streamwise, wall-normal, and spanwise directions, respectively. The Reynolds number is $Re = 400$, based on the wall speed and channel half-height. The dataset consists of 11 trajectories, each evolved for 11000 time units with snapshots saved every 1 time unit. To remove initial transients, the first 1000 time units of each trajectory are discarded, leaving 10000 retained snapshots per trajectory and 110,000 snapshots in total.

For plane Poiseuille flow, the simulations are carried out in a domain of size $L_x = \pi$, $L_y = 2$, and $L_z = \pi/2$, using $48 \times 81 \times 48$ grid points. The Reynolds number is $Re = 1800$, based on the laminar centerline velocity and channel half-height. The simulations are performed in the constant-mass-flux formulation, with a forcing adjusted to maintain the prescribed streamwise flux. The dataset again consists of 11 trajectories evolved for 10500 time units with snapshots saved every 1 time unit. In this case, the first 1000 time units are omitted, leaving 9500 retained snapshots per trajectory and 104,500 snapshots in total. An 80/20 split is then used for training and testing in both cases.

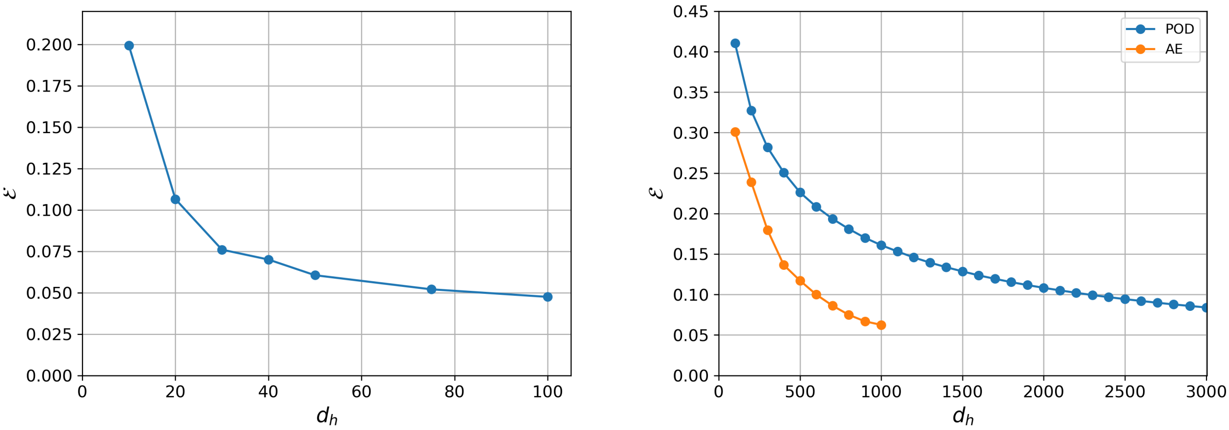

Figure 4. Reconstruction error versus latent dimension for the learned autoencoder representations of (a) plane Couette flow and (b) plane Poiseuille flow. Reconstruction error from a projection onto d_h POD modes is also shown for comparison for the plane Poiseuille flow case.

Autoencoder representations

For both systems, we trained 3D convolutional autoencoders to map the mean-subtracted velocity field to a low-dimensional latent representation and to reconstruct the full state from that representation. The input and output of the network are both full three-component velocity snapshots on the DNS grid, while the bottleneck layer provides a compact coordinate representation of the flow on the underlying attractor. The principal design parameter is the latent dimension, which determines the degree of compression and is varied systematically to assess the dimension required for accurate state representation.

The encoder and decoder are implemented as 3D convolutional neural networks with residual blocks acting on the full velocity field. In the Couette-flow model, the encoder compresses the spatial representation through three convolutional downsampling stages, followed by additional convolutional processing at reduced resolution and a final linear map to the latent vector. In the first downsampling stage, downsampling is applied only in the streamwise and spanwise directions, while the wall-normal resolution is retained; subsequent downsampling stages reduce resolution in all three directions. Residual blocks are included at each stage to enhance expressivity and improve optimization. For the Couette-flow model, the encoder channel dimensions progress from 3 to 32, 64, 128, 32, and finally 8 before flattening and linear projection to the latent vector, and the decoder reverses this progression. The decoder mirrors the encoder structure by first mapping the latent vector back to a low-resolution feature tensor and then reconstructing the full velocity field through a sequence of interpolation and convolutional refinement stages. The Poiseuille-flow model follows the same general design, but with increased representational capacity and a substantially larger latent dimension.

Several aspects of the architecture are chosen to reflect the geometry of the underlying flow. The convolutional layers use periodic padding in the streamwise and spanwise directions, consistent with the periodicity of the computational domain, and zero padding in the wall-normal direction to reflect the presence of physical boundaries. Training is performed by minimizing a relative mean-squared reconstruction error, with a weighting chosen so that it is consistent with the underlying spatial discretization. Optimization is performed using AdamW and a one-cycle learning-rate schedule.

In both plane Couette and plane Poiseuille flow, the reconstruction quality is strong, with relative reconstruction errors of approximately $5-6\%$ on held-out test data. This level of performance compares favorably with many previously reported autoencoder-based ROM results for transitional or weakly turbulent shear flows. Figure 4 shows the reconstruction error as a function of latent dimension for the two systems. The error metric reported here is the ensemble-averaged relative error

$$\mathcal{E} = \left\langle \frac{\|\hat{\mathbf{v}} - \mathbf{v}\|}{\|\mathbf{v}\|} \right\rangle,$$

where $\mathbf{v}$ is the mean-subtracted velocity field, $\hat{\mathbf{v}}$ is its reconstruction by the autoencoder, $\|\cdot\|$ denotes the spatial L_2 norm of the velocity field, evaluated discretely using quadrature weights to account for the nonuniform grid, and $\langle\cdot\rangle$ denotes an ensemble average. For plane Couette flow, the error appears to level off at latent dimension roughly 50, with only small improvements for larger latent dimensions. For plane Poiseuille flow, by contrast, the same error metric appears to level off only around latent dimension roughly 700–800. This large disparity suggests that, in the present representation, plane Poiseuille flow requires a substantially higher-dimensional latent space than plane Couette flow at the chosen parameter values.

Latent-space dynamics models

Having obtained accurate low-dimensional state representations, the next step in the construction of the reduced order model is to model the time evolution on the learned manifold. Following the DManD framework (Linot & Graham, 2023), we model the dynamics of the latent variable from the trained autoencoders using a neural ODE. Due to the large dimension required for the plane Poiseuille flow, we have chosen to focus these efforts on developing dynamics models for the plane Couette flow system.

The latent dynamics model is represented as a neural ODE of the form $\dot{h} = g_\theta(h)$, where $h \in \mathbb{R}^{d_h}$ denotes the latent state. In the present work, the vector field g_θ is parameterized by a fully connected neural network with three hidden

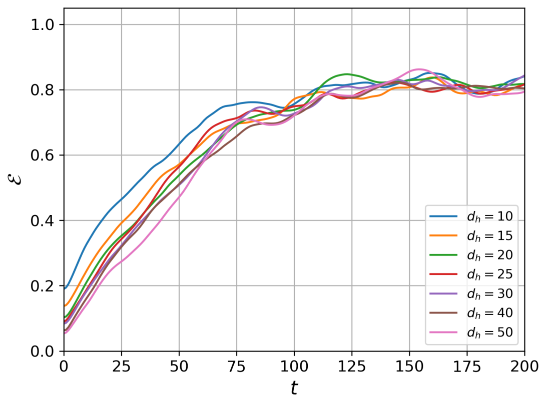

Figure 5. Ensemble averaged relative error in predictions by the DManD model on the plane Couette flow system for varying latent dimension, d_h

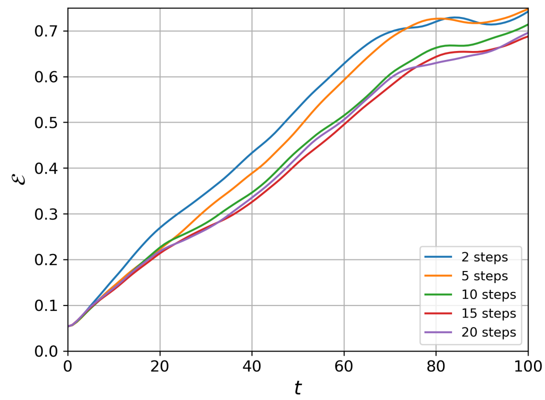

Figure 6. Comparison of prediction horizon used in neural ODE training on the short-time predictive performance of DManD models.

layers of width 1024, GELU activations, and a final linear output layer mapping back to $\mathbb{R}^{d_h}$. Training is performed on short trajectory segments in latent space using a mean-square-error loss between the predicted and true latent trajectories over the training horizon

$$J(\theta_g) = \frac{1}{mK} \sum_{i=1}^{m} \sum_{k=0}^{K-1} \left\| \hat{\mathbf{h}}_{i,k} - \mathbf{h}_{i,k} \right\|_2^2, \tag{2}$$

where m is the batch size, K is the number of time steps in the training prediction horizon, $\mathbf{h}_{i,k}$ is the true latent state at the k-th time step of segment i, and $\hat{\mathbf{h}}_{i,k}$ is the corresponding neural-ODE prediction obtained by integrating

$$\frac{d\mathbf{h}}{dt} = g(\mathbf{h}; \theta_g), \qquad \hat{\mathbf{h}}_{i,0} = \mathbf{h}_{i,0}, \tag{3}$$

forward over the training time interval. For the Couette-flow models shown here, we use a batch size of 40 and the AdamW optimizer with learning rate 10^{-4} and weight decay 10^{-5}. A short linear warmup of the learning rate is followed by cosine-annealing decay over the remainder of training. Model selection is based on the validation loss computed on held-out latent-space trajectory segments.

Figure 5 shows the ensemble averaged relative error, $\mathcal{E}$ over time for the DManD models at varying latent dimension, d_h. For this system, the Lyapunov timescale has previously been approximated to be $\tau_L = 48$ time units (Inubushi *et al.*, 2015; Linot & Graham, 2023), so the prediction horizon shown in Fig. 5 spans approximately $4\tau_L$. These errors are computed by integrating latent trajectories forward in time from a test initial condition, decoding the latent trajectory, and evaluating the error in the prediction on the full velocity field. Here we see that increasing the latent dimension does allow for improved prediction at least up to $d_h = 50$, but we note that the relative improvement with an increase in the latent dimension becomes considerablt smaller beyond $d_h = 20$. Additionally, after the initial error growth, the error for each of the models appears to saturate after roughly $2\tau_L$.

In Fig. 6, we examine the effect of varying the length of prediction horizon, K, used in training on the short-time prediction performance of these models. These results indicate that increasing the training prediction horizon up to 10 steps (of one time unit) improves the short-time predictions of the model, with very little improvement achieved beyond that by increasing the training horizon to 15 or 20 timesteps.

Incorporating known physics

In addition to the purely data-driven DManD models described above, we also consider the physics-informed DManD (PI-DManD) framework recently proposed by Guo and Graham (Guo & Graham, 2025). The starting point is the Manifold Galerkin (MG) construction, in which the known full-order model is projected onto the autoencoder manifold coordinates. If $\mathbf{u}$ denotes the full state, $\mathbf{h} = E(\mathbf{u})$ the latent coordinates, and $D(\mathbf{h})$ the decoder map, then the latent-space vector field induced by the full-order model f_{FOM} is

$$\frac{d\mathbf{h}}{dt} = \mathbf{g}_m(\mathbf{h}), \qquad \mathbf{g}_m(\mathbf{h}) = \frac{\partial E}{\partial \mathbf{u}} f_{\text{FOM}}(\mathbf{u}). \tag{4}$$

Because direct evaluation of $\mathbf{g}_m$ requires repeated decoding and evaluation of the full-order model, it is generally too expensive to serve directly as a practical reduced-order model. Following Guo & Graham (2025), we therefore train a neural network approximation to this projected vector field, denoted MG-N,

$$\frac{d\mathbf{h}}{dt} = \mathbf{g}_{mn}(\mathbf{h}), \tag{5}$$

which provides a physics-informed latent-space model at much lower cost. For a given trained encoder, E, and stored evaluations of $f_{\text{FOM}}(\mathbf{u})$, training $\mathbf{g}_{mn}$ is a straightforward supervised learning problem.

Two hybrid variants are then constructed by combining MG-N with the data-driven DManD approach. In PI-DManD-C, a neural ODE correction term $\mathbf{g}_c$ is added to the MG-N vector field,

$$\frac{d\mathbf{h}}{dt} = \mathbf{g}_{mn}(\mathbf{h}) + \mathbf{g}_c(\mathbf{h}), \tag{6}$$

so that the correction network learns to account for discrepancies arising from autoencoder error, approximation error in

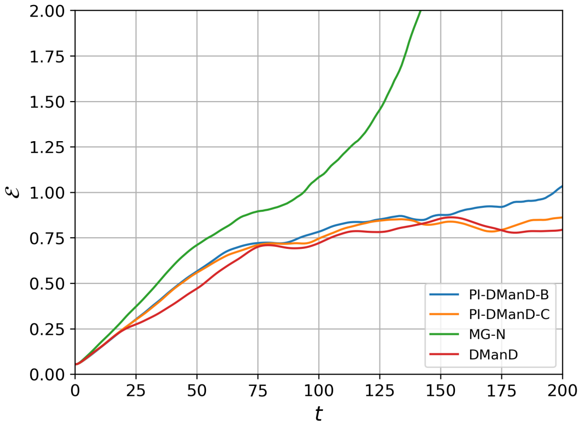

Figure 7. Comparison of predictions between DManD model and physics-informed reduced order models, PI-DManD-B, PI-DManD-C, and MG-N for the the Couette flow dataset with latent dimension $d_h = 50$.

the MG-N fit, and possible error in the full-order model itself. In PI-DManD-B, by contrast, the MG-N model is used as an initialization for training the latent-space neural ODE. In this case, the learned vector field $\mathbf{g}_B$ is initialized with the weights of the MG-N model and then fine-tuned using the latent trajectory data. Thus, in the present implementation, PI-DManD-B may be viewed as a variant in which the projected physics provide an informed starting point for subsequent data-driven training. For full-details of PI-DManD, readers are referred to Guo & Graham (2025).

In the present work, we compare the purely data-driven DManD models against MG-N, PI-DManD-B, and PI-DManD-C for the plane Couette-flow latent dynamics. This comparison is intended to assess whether incorporating the known governing equations through manifold projection improves predictive performance relative to a purely data-driven latent-space model. Fig. 7 shows a comparison of the predictive performance for each of the physics-informed DManD variants with a standard DManD implementation for the Couette-flow dataset with autoencoder latent dimension $d_h = 50$. Among these models, MG-N performs the worst: its error is already larger at short times, and at longer times the predictions become unstable while the errors of the DManD and PI-DManD models saturate. This behavior is not surprising, since MG-N is trained only to match the instantaneous projected vector field, whereas DManD, PI-DManD-B, and PI-DManD-C are trained directly on short latent-space trajectory rollouts. By contrast, PI-DManD-B and PI-DManD-C perform comparably to one another and remain close to the performance of the baseline DManD model, though in the present case both are slightly worse than DManD. One possible explanation is that, for this Couette-flow system, the overall predictive accuracy is already limited to a significant extent by autoencoder reconstruction error. In the work of Guo & Graham (2025), the physics-informed latent-dynamics models were tested in settings where the autoencoder reconstruction errors were substantially smaller than those obtained here, so differences in the latent-dynamics model could be reflected more directly in the rollout accuracy.

DISCUSSION AND CONCLUSIONS

We have presented a constrained-optimization framework for identifying efficient routes to laminarization in chaotic shear-flow dynamics and demonstrated it in the nine-mode MFE model. The optimization reveals the optimal perturbation of a given reference point in the turbulent region of state space, identifies the nearest point on the edge of chaos to that reference point, and supports the interpretation that a periodic edge orbit organizes the nearby laminarizing and non-laminarizing trajectories for the MFE system.

We have also reported progress toward extending this framework to higher-dimensional transitional shear flows through learned reduced-order models. Convolutional autoencoders provide accurate low-dimensional representations for both plane Couette and plane Poiseuille flow, with latent-dimension sweeps indicating markedly different effective complexity in the two systems. For plane Couette flow, latent-space dynamics models yield encouraging predictive performance, while the physics-informed variants considered here have not shown a clear advantage over the baseline DManD model in the present tests. We believe that this is at least partly because any benefit of incorporating the known physics is limited by errors in the learned manifold representation and the associated projection of the dynamics onto that manifold.

Taken together, these results demonstrate the relaminarization methodology in a reduced chaotic model and help clarify the challenges involved in extending it to larger systems. The present results therefore provide a concrete foundation for future efforts to compute minimal seeds for relaminarization in realistic wall-bounded shear flows.

REFERENCES

Burns, Keaton J, Vasil, Geoffrey M, Oishi, Jeffrey S, Lecoanet, Daniel & Brown, Benjamin P 2020 Dedalus: A flexible framework for numerical simulations with spectral methods. *Physical Review Research* **2** (2), 023068.

Buzhardt, Jake & Graham, Michael D 2026 Identifying efficient routes to laminarization: an optimization approach. *Physical Review Fluids* **11** (2), 023902.

Chung, Seung Whan & Freund, Jonathan B. 2022 An optimization method for chaotic turbulent flow. *Journal of Computational Physics* **457**, 111077.

Guo, Alex & Graham, Michael D 2025 Blending data and physics for reduced-order modeling of systems with spatiotemporal chaotic dynamics. *arXiv preprint arXiv:2507.21299* .

Inubushi, Masanobu, Takehiro, Shin-ichi & Yamada, Michio 2015 Regeneration cycle and the covariant lyapunov vectors in a minimal wall turbulence. *Physical review E* **92** (2), 023022.

Kerswell, Rich R 2018 Nonlinear nonmodal stability theory. *Annual Review of Fluid Mechanics* **50**, 319–345.

Linot, Alec J. & Graham, Michael D. 2023 Dynamics of a data-driven low-dimensional model of turbulent minimal Couette flow. *Journal of Fluid Mechanics* **973**, A42.

Moehlis, Jeff, Faisst, Holger & Eckhardt, Bruno 2004 A low-dimensional model for turbulent shear flows. *New Journal of Physics* **6** (1), 56.