A Mori-Zwanzig-Inspired Neural Closure for Domain-Decomposed ROMs of Coupled Physical Systems

Keyu Xu*, Qian Wang*
Corresponding author: qian.wang@csrc.ac.cn
* Beijing Computational Science Research Center, Beijing 100193, China.

Abstract:
Projection-based reduced-order models (ROMs) provide efficient surrogates for repeated high-fidelity simulations, but their accuracy can deteriorate in domain-decomposed reduced-order models (DDROMs) when local reduced coefficients are affected by unresolved modes and interfacial interactions with neighboring subdomains. This work studies a closure-enhanced DDROM (Closure-DDROM) framework for coupled physical systems. For each subdomain, a local proper orthogonal decomposition (POD) basis is constructed, and Galerkin projection is used to obtain the Markovian reduced dynamics. The discrepancy between the projected high-fidelity reduced dynamics and the Markovian local model is then formulated as a finite-memory closure term inspired by the Mori–Zwanzig formalism and learned from reduced-order histories using recurrent neural networks (RNNs). The learned closure is incorporated into the online reduced-order rollout to compensate for the missing reduced dynamics. The framework is assessed using an advection-diffusion-reaction problem and a two-dimensional vorticity-streamfunction flow problem. In both cases, the Closure-DDROM framework improves the accuracy of the baseline local POD-Galerkin DDROM and better preserves the high-fidelity dynamics. These results indicate that residual-based finite-memory closures offer a promising mechanism for improving local POD-Galerkin DDROMs, while further investigation is needed to assess parameter generalization, long-time stability, and interface consistency.

Keywords: Reduced Order Model, Domain Decomposition, Neural Closure, Proper Orthogonal Decomposition.

1. Introduction

High-fidelity simulations of time-dependent partial differential equations (PDEs) remain computationally demanding, especially when repeated simulations are required for design exploration, parametric studies, uncertainty quantification, or optimization. Projection-based reduced-order models (ROMs) reduce this cost by approximating the full-order state in a low-dimensional trial space and evolving a smaller set of generalized coordinates. Proper orthogonal decomposition (POD) is one of the standard data-driven tools for constructing such spaces from high-fidelity snapshot data [1, 2, 3].

For large-scale coupled systems and geometrically complex domains, a single global reduced basis may be inefficient or difficult to interpret. Domain-decomposed reduced-order models (DDROMs) address this issue by constructing local reduced spaces on subdomains or components and coupling the resulting reduced systems through interface conditions or reduced transmission models [4, 5, 6, 7]. Recent work has considered several strategies for coupling local reduced models, including flux-based coupling, composite reduced bases, Schur-complement formulations, and boundary response maps [8, 9, 10].

Despite these developments, local POD-Galerkin ROMs may still exhibit accuracy loss during online prediction. In particular, aggressive modal truncation removes dynamically relevant information, while the coupling of independently reduced subdomains can introduce additional residual effects at interfaces. These unresolved contributions appear in the reduced equations as a discrepancy between the projected full-order dynamics and the Markovian reduced-order dynamics obtained by local Galerkin projection.

For long or nonlinear rollouts, this discrepancy can accumulate and degrade the predicted reduced trajectory [7].

The Mori–Zwanzig (MZ) formalism provides a useful interpretation of such unresolved effects by decomposing the evolution of resolved variables into Markovian, memory, and noise contributions [11, 12, 13]. In reduced-order modeling, this viewpoint motivates closure terms that compensate for the influence of truncated or unresolved modes. Data-driven closure models, including recurrent neural network closures and latent-space closure models, have therefore been investigated as practical approximations of history-dependent reduced dynamics [14, 15, 16, 17].

In this work, we develop a closure-enhanced DDROM (Closure-DDROM) framework based on local POD-Galerkin projection. The baseline model is obtained by constructing local POD bases on each subdomain and applying Galerkin projection to the corresponding local full-order operators. The closure model is trained to approximate the right-hand-side residual discrepancy between the projected full-order dynamics and the unclosed DDROM dynamics. Motivated by the finite-memory interpretation of the MZ formalism, this residual closure is represented by a gated recurrent unit (GRU) network that takes recent reduced-coordinate histories as input [18]. The trained closure is then inserted into the online reduced-order rollout, while the intrusive Galerkin structure of the baseline DDROMs is retained.

The main contributions of this study are threefold. First, we formulate a residual-based neural closure for local POD-Galerkin DDROMs in which the learned term corrects the reduced dynamics rather than replacing the full reduced dynamics. Second, we clarify the offline and online roles of the closure model, including the construction of the residual training target and its use during time integration. Third, we assess the resulting Closure-DDROM framework on two decomposed test problems: an advection-diffusion-reaction equation with Robin-type transmission conditions and a two-dimensional multi-injector flow with strong injector-injector coupling. The numerical results compare the proposed closure-enhanced model with the standard unclosed DDROM baseline.

2. Methodology

2.1 Reduced basis method

We consider a semi-discrete full-order model (FOM) of the form

$$\begin{cases} \dfrac{\mathrm{d}u_h(t)}{\mathrm{d}t} = f(u_h(t);\omega), \\ u_h(0) = u_{h,0}, \end{cases} \tag{1}$$

where $u_h(t) \in \mathbb{R}^N$ is the full-order state vector and ω denotes the physical or geometrical parameters. In projection-based model order reduction, the FOM state is approximated in an affine low-dimensional trial space,

$$u_h(t) \approx \bar{u} + V\alpha(t), \tag{2}$$

where $\bar{u} \in \mathbb{R}^N$ is a reference state, $V \in \mathbb{R}^{N\times r}$ is a reduced basis matrix, and $\alpha(t) \in \mathbb{R}^r$ is the vector of reduced coefficients, with $r \ll N$. The basis V is obtained by POD and is assumed to be orthonormal:

$$V^\top V = I_r. \tag{3}$$

Substituting the reduced approximation into the FOM and applying Galerkin projection gives

$$V^\top V \frac{d\alpha(t)}{dt} = V^\top f(V\alpha(t) + \bar{u};\omega). \tag{4}$$

Since $V^\top V = I_r$, the standard POD-Galerkin reduced-order model is

$$\frac{\mathrm{d}\alpha(t)}{\mathrm{d}t} = V^\top f(V\alpha(t) + \bar{u}; \omega) \tag{5}$$

2.2 Domain-decomposed local reduced-order model

To handle complex geometries and localized phenomena, the global computational domain Ω is partitioned into N_s subdomains, $\Omega = \bigcup_{i=1}^{N_s} \Omega_i$. Instead of extracting a global basis, the high-fidelity snapshot data is decomposed accordingly, and local POD bases V_i are extracted independently for each subdomain. By treating each subdomain as an independent system, the local full-order state $u_{h,i}$ evolves as

$$\frac{\mathrm{d}u_{h,i}}{\mathrm{d}t} = F_i\left(u_{h,i}, \{u_{h,j}\}_{j\in\mathcal{N}_i}; \omega\right), \qquad i = 1, \ldots, N_s, \tag{6}$$

where $\mathcal{N}_i$ is the set of neighboring subdomains. The operator F_i includes both the interior discretized dynamics on Ω_i and the coupling terms induced by the interface conditions.

The subdomain coupling is enforced through transmission conditions on the shared interfaces. Let Γ_{ij} denote the interface between Ω_i and Ω_j, and let k be the Schwarz iteration index. A Robin-type interface condition can be written in the abstract form

$$\mathcal{B}_{ij}\left(u_{h,i}^{(k)}\right) = \mathcal{B}_{ji}\left(u_{h,j}^{(k-1)}\right), \qquad \text{on } \Gamma_{ij}, \tag{7}$$

where $\mathcal{B}_{ij}$ denotes a Robin trace operator combining Dirichlet and Neumann information. Upon convergence of the Schwarz iterations, the neighboring states are incorporated into the local operator F_i. For each subdomain, we construct a local POD basis $V_i \in \mathbb{R}^{N_i\times r_i}$ and approximate

$$u_{h,i}(t) \approx \bar{u}_i + V_i\alpha_i(t), \tag{8}$$

where $\alpha_i(t) \in \mathbb{R}^{r_i}$. Applying Galerkin projection locally gives the baseline domain-decomposed ROM (DDROM),

$$\frac{\mathrm{d}\alpha_i}{\mathrm{d}t} = R_i(\alpha_i, \{\alpha_j\}_{j\in\mathcal{N}_i}; \omega), \tag{9}$$

with

$$R_i\left(\alpha_i, \{\alpha_j\}_{j\in\mathcal{N}_i}; \omega\right) = V_i^\top F_i(\bar{u}_i + V_i\alpha_i, \{\bar{u}_j + V_j\alpha_j\}_{j\in\mathcal{N}_i}; \omega). \tag{10}$$

2.3 Mori-Zwanzig formalism

The Mori-Zwanzig (MZ) formalism provides a rigorous theoretical framework for understanding the effect of unresolved variables on the resolved dynamics. By treating each subdomain as an independent system with exact interface conditions, the closure exclusively accounts for the unresolved scales lost during the local POD truncation.

Using projection operators $\mathcal{P}$ (onto the resolved subspace) and $\mathcal{Q} = \mathcal{I} - \mathcal{P}$ (onto the unresolved subspace), the MZ identity rewrites the exact evolution of a resolved observable (such as the POD coefficients) into three distinct components:

$$\frac{d\alpha(t)}{dt} = \underbrace{e^{t\mathcal{L}}\mathcal{P}\mathcal{L}\alpha(0)}_{\text{Markov term}} + \underbrace{e^{t\mathcal{Q}\mathcal{L}}\mathcal{Q}\mathcal{L}\alpha(0)}_{\text{Noise term}} + \underbrace{\int_0^t e^{(t-s)\mathcal{L}}\mathcal{P}\mathcal{L}e^{s\mathcal{Q}\mathcal{L}}\mathcal{Q}\mathcal{L}\alpha(0)\,ds}_{\text{Memory term}}, \tag{11}$$

where $\mathcal{L}$ is the Liouville operator of the underlying system.

In the limiting case where the full-order trajectory remains on the resolved POD manifold, the MZ Markovian contribution reduces to the standard POD-Galerkin right-hand side. However, the projected FOM dynamics generally contains contributions that are not represented by the local Galerkin dynamics R_i. These contributions may arise from unresolved POD modes, projection residuals, and coupling effects introduced by the independent reduction of neighboring subdomains. The MZ formalism provides a theoretical interpretation of such unresolved effects by decomposing the resolved dynamics into Markovian, noise, and memory terms [11, 12, 13]. In the present work, we use this viewpoint as motivation for a finite-memory residual closure rather than as an exact evaluation of the MZ memory kernel. We augment the baseline DDROM as

$$\frac{\mathrm{d}\alpha_i(t)}{\mathrm{d}t} = R_i(\alpha_i, \{\alpha_j\}_{j\in\mathcal{N}_i}) + \mathcal{M}_i(t;\omega). \tag{12}$$

where $\mathcal{M}_i$ denotes the closure correction for subdomain Ω_i.

2.4 Artificial neural network memory model

Guided by the Mori-Zwanzig formalism, we approximate the time-dependent closure term using a recurrent neural network (RNN) [19]. Specifically, a Gated Recurrent Unit (GRU) is employed to capture finite-memory effects [20, 18], offering an efficient alternative to Long Short-Term Memory (LSTM) [21] networks with fewer trainable parameters.

For a memory length of q, the input at time t_n is the local reduced-coeffients history

$$\alpha_i^{n-q+1:n} = \left(\alpha_i^{n-q+1}, \dots, \alpha_i^n\right). \tag{13}$$

The learned closure is then written as

$$\mathcal{M}_i(t_n;\omega) \approx \mathcal{M}_{i,\theta}\left(\alpha_i^{n-q+1:n};\omega\right), \tag{14}$$

where θ denotes the trainable network parameters.

Fundamentally, the neural network is trained to learn the missing right-hand side dynamics. Thus, the exact training target is formulated as the residual discrepancy between the projection of the full-order dynamics and the unclosed Markovian evaluation,

$$\mathcal{M}_i^*(t_n;\omega) = V_i^\top F_i(u_h;\omega) - R_i(\alpha_i(t_n), \{\alpha_j(t_n)\}_{j\in\mathcal{N}_i};\omega). \tag{15}$$

The internal hidden state of the GRU is updated recursively. For an input vector x_n (drawn from the sequence α_i) at time t_n and the preceding hidden state h_{n-1}, the gating operations are computed as follows:

$$z_n = \sigma(W_z x_n + U_z h_{n-1} + b_z), \tag{16}$$
$$r_n = \sigma(W_r x_n + U_r h_{n-1} + b_r), \tag{17}$$
$$\widetilde{h}_n = \tanh(W_h x_n + U_h(r_n \odot h_{n-1}) + b_h), \tag{18}$$
$$h_n = (1 - z_n) \odot h_{n-1} + z_n \odot \widetilde{h}_n. \tag{19}$$

Here, z_n and r_n are the update and reset gates, respectively, $\widetilde{h}_n$ is the candidate hidden state, $\sigma(\cdot)$ is the sigmoid activation, and $\odot$ denotes element-wise multiplication. The closure output is obtained from the final hidden state through an affine map,

$$\mathcal{M}_{i,\theta} = W_o h_n + b_o. \tag{20}$$

The network parameters are trained by minimizing the mean-squared error between the predicted closure and the residual target over the training snapshots,

$$\mathcal{J}_i(\theta) = \frac{1}{N_{\text{train}}} \sum_{n \in \mathcal{T}_{\text{train}}} \left\| \mathcal{M}_{i,\theta} \left((\alpha_i^{\text{FOM}})^{n-q+1:n}; \omega \right) - \mathcal{M}_i^*(t_n; \omega) \right\|_2^2. \tag{21}$$

In the numerical experiments, the optimization is initialized with Adam [22] for the first 90% of training, performing fast mini-batch exploration, and then refined using full-batch L-BFGS [23] for the final 10% of epochs to drive the gradient residuals to a strict local minimum.

After training, the online closed DDROM is advanced as

$$\frac{d\alpha_i(t_n)}{dt} = R_i \left(\alpha_i^n, \{\alpha_j^n\}_{j \in \mathcal{N}_i}; \omega \right) + \mathcal{M}_{i,\theta} \left(\alpha_i^{n-q+1:n}; \omega \right). \tag{22}$$

3. Numerical results

This section evaluates the proposed Closure-DDROM against the standard local POD-Galerkin DDROM. The comparison focuses on reconstruction accuracy at the final simulation time and on the spatial distribution of the absolute error. For a reconstructed field u_r and the corresponding FOM reference u_h, the final-time error is evaluated using

$$e = u_h(T) - u_r(T), \qquad \varepsilon_{\text{rel}} = \frac{\|e\|_2}{\|u_h(T)\|_2}.$$

The mean absolute error and maximum absolute error are also reported over the spatial degrees of freedom. Unless otherwise stated, both the standard DDROM and the Closure-DDROM use the same local POD dimensions, time integrator, and domain-decomposition interface treatment.

3.1 Advection-diffusion-reaction equation

The first test case is a decomposed linear advection-diffusion-reaction (ADR) problem on the unit square. This example is used to assess whether the learned closure can reduce the final-time error of a local POD-Galerkin DDROM in a setting where the subdomains have different transport behavior. The governing equation on each subdomain $s \in \{A, B\}$ is

$$\frac{\partial u_s}{\partial t} - \mu_s \Delta u_s + \boldsymbol{\beta}_s \cdot \nabla u_s + \sigma_s u_s = 0, \qquad \mathbf{x} \in \Omega_s. \tag{23}$$

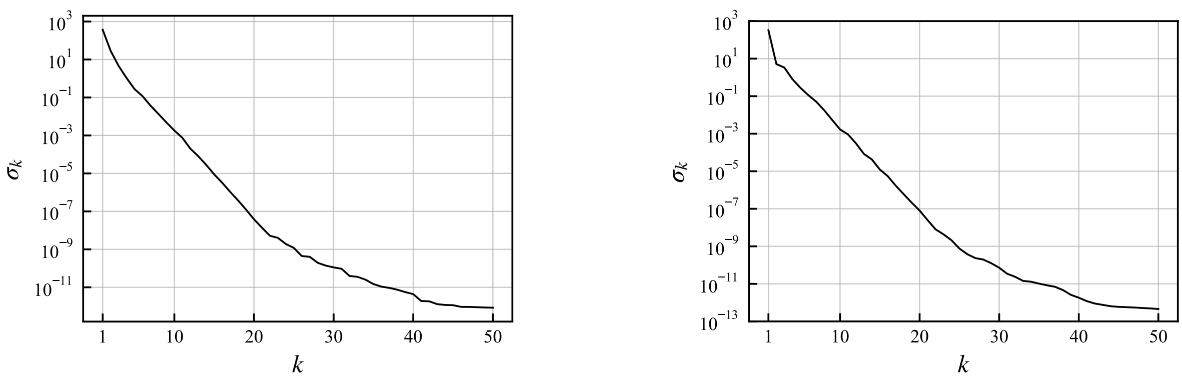

(a) Subdomain A (b) Subdomain B

Figure 1: Decay of POD singular values for the local bases in the ADR test case.

The global domain is partitioned into two subdomains sharing the interface $\Gamma = \{0.6\} \times [0, 1]$. The physical parameters dictate distinct transport behaviors: Subdomain A is advection-dominated ($\mu_A = 1/3$, $\boldsymbol{\beta}_A = (3, 2)^\top$, $\sigma_A = 1$), whereas Subdomain B is purely diffusive ($\mu_B = 1$, $\boldsymbol{\beta}_B = (0, 0)^\top$, $\sigma_B = 0$).

The FOM is discretized on a uniform Cartesian mesh with $h_x = h_y = 0.01$. Time integration is performed using a fully implicit Euler method with $\Delta t = 2.5 \times 10^{-4}$ up to $T = 0.05$. The subdomain coupling is enforced by a Schwarz alternating procedure with Robin-type transmission conditions, using a maximum of 100 iterations per time level and a convergence tolerance of 10^{-8} [24, 25]. The Robin parameter is set to $\lambda = 10$ on both sides of the interface. At cross-points where the interface intersects the global boundary, overlapping nodal contributions are averaged with weight 0.5 during assembly.

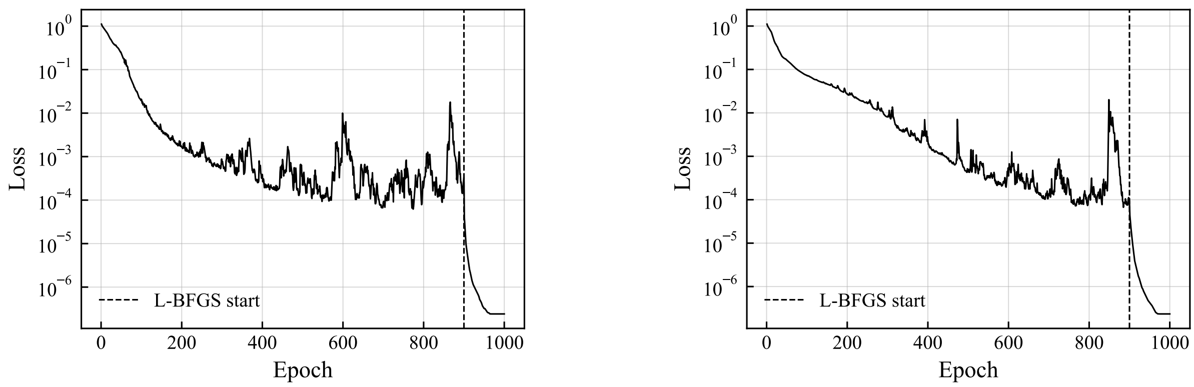

(a) Subdomain A (b) Subdomain B

Figure 2: Training loss histories for the local reduced-order models in the ADR test case.

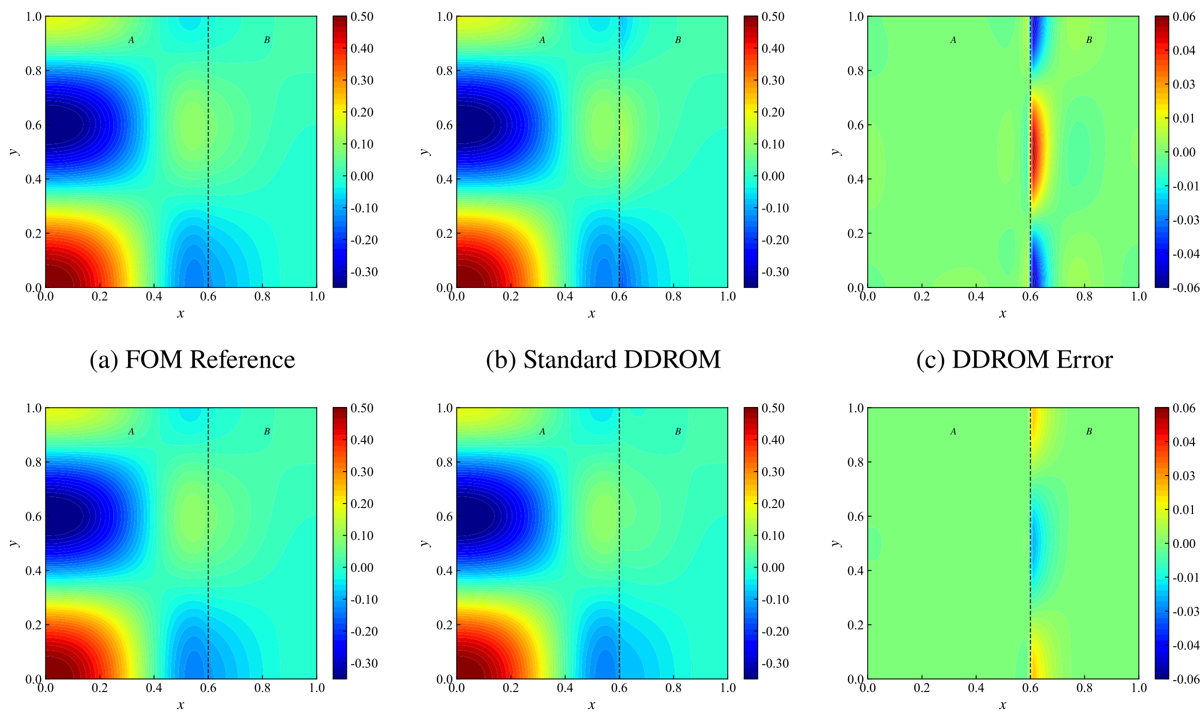

(a) FOM Reference (b) Standard DDROM (c) DDROM Error

(d) FOM Reference (e) Closure-DDROM (f) Closure-DDROM Error

Figure 3: Comparison of reconstructed spatial fields and absolute error distributions for the ADR system at $T = 0.05$.

There are 201 snapshots in total. Figure 1 shows the singular value decay of the local POD snapshot matrices. For the reduced representations, the POD latent dimension is truncated at $r = 5$ for both subdomains. The GRU closure model uses a hidden dimension of 32 and a memory length of $q = 5$. All 201 snapshots are used for training, with a batch size of 64, a single layer, and 1000 epochs. The convergence of the neural network training process is confirmed by the loss histories depicted in Figure 2.

Figure 3 compares the final-time reconstructed fields and absolute error maps. The corresponding quantitative errors are summarized in Table 1. The standard DDROM gives a relative L_2 error of 4.96×10^{-2}, while the Closure-DDROM reduces this value to 1.74×10^{-2}. The maximum absolute error is also reduced from 5.73×10^{-2} to 1.84×10^{-2}. These results indicate that, for this ADR test case, the learned residual closure improves the final-time reconstruction accuracy relative to the unclosed DDROM baseline.

Table 1: Quantitative error metrics for the ADR.

Model	Mean Abs. Error	Max Abs. Error	L_2 Error	Relative L_2
Standard DDROM	2.60×10^{-3}	5.73×10^{-2}	7.11×10^{-1}	4.96×10^{-2}
Closure-DDROM	9.29×10^{-4}	1.84×10^{-2}	2.49×10^{-1}	1.74×10^{-2}

3.2 Coupled Multi-Injector Flows

The second numerical application focuses on the more complex three-injector flow configuration shown in Figure 4a. A plane Poiseuille flow profile with an average velocity of $u_{\text{avg}} = 1.0$ is prescribed at the inlet boundary, defined as $u_{\text{inlet}} = 6\hat{y}(1-\hat{y})$, where $\hat{y}$ is the vertical coordinate normalized by the single injector height h_{injector}. No-slip conditions ($u_{\text{wall}} = 0$) are imposed along all solid boundaries, while a standard zero-gradient outflow condition ($\partial_x \mathbf{u} = 0$) is applied at the outlet. The high-fidelity model employs the vorticity-streamfunction formulation solved by second-order finite differences and explicit Runge-Kutta time integration. Consistent boundary conditions for the vorticity ω and streamfunction ψ are enforced through the physical velocity constraints $\omega = \nabla \times \mathbf{u}$ and $\mathbf{u} = (\partial_y \psi, -\partial_x \psi)^\top$. Additionally, a multigrid solver is utilized for the Poisson equation to accelerate convergence.

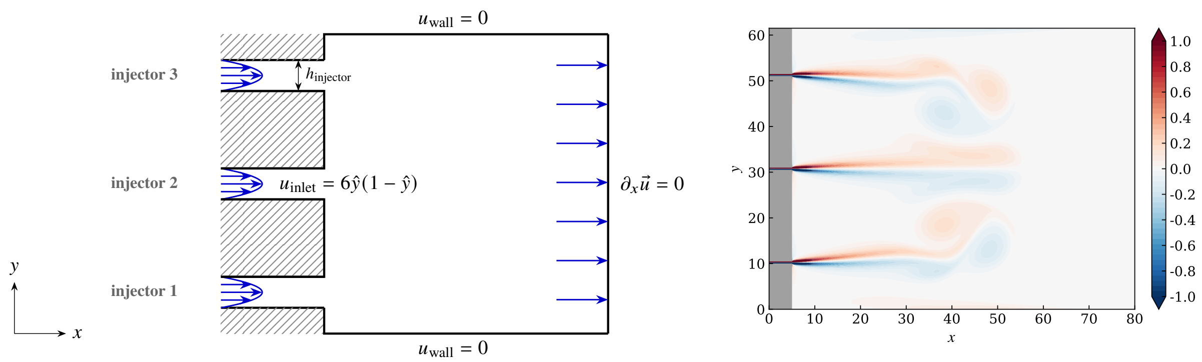

(a) Three-injector flow configuration.

(b) High-fidelity vorticity field at $t = 275.0$ s.

Figure 4

For the reduced-order modeling procedure, POD snapshots are collected from the developed flow regime over the time interval $t \in [275.0, 375.0]$ s. The initial snapshot time, shown in Figure 4b, is selected after the injector-induced vortical structures and their mutual interactions have become sufficiently developed. A total of 401 snapshots were extracted to construct the basis. The reduced-order approximation for vorticity ω and streamfunction ψ is given by:

$$\omega_h(\mathbf{x}, t) \approx \omega_r(\mathbf{x}, t) = \bar{\omega}(\mathbf{x}) + \sum_{i=1}^{r} \alpha_i(t)\phi_i^{\omega}(\mathbf{x}), \tag{24}$$

$$\psi_h(\mathbf{x}, t) \approx \psi_r(\mathbf{x}, t) = \bar{\psi}(\mathbf{x}) + \sum_{i=1}^{r} \alpha_i(t)\phi_i^{\psi}(\mathbf{x}). \tag{25}$$

Streamfunction using the same time-dependent coefficient as vorticity, since the vorticity and streamfunction are related by the Poisson equation. To construct the local reduced-order bases, the POD singular value decay for each of the three independent subdomains is analyzed (see Figure 5).

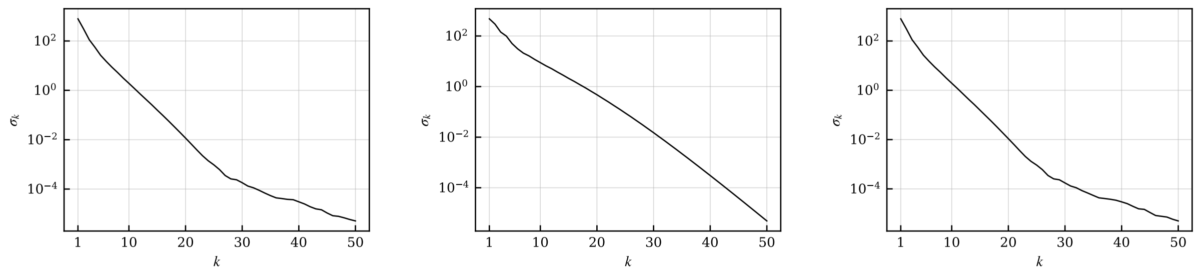

(a) Subdomain A (b) Subdomain B (c) Subdomain C

Figure 5: Decay of POD singular values for the three subdomains in the multi-injector flow.

Based on this decay, the latent dimension is truncated at $r = 5$ across all subdomains. To robustly capture the non-Markovian interfacial dynamics, the GRU closure networks utilize a hidden dimension of 64 and a memory sequence length of $q = 8$. Figure 5 shows the POD singular value decay for the three local subdomains. The models are trained for 1000 epochs, and the training loss histories are shown in Figure 6.

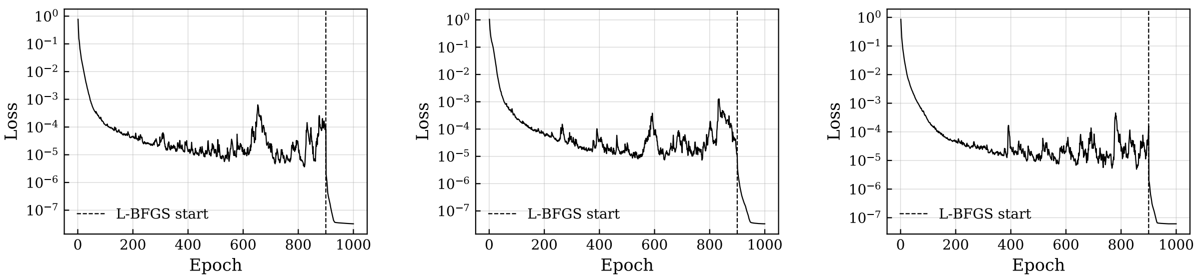

(a) Subdomain A (b) Subdomain B (c) Subdomain C

Figure 6: Training loss histories for the multi-injector sub-models.

The necessity of the proposed closure formulation becomes particularly evident in the multi-injector configuration. The final-time reconstructed fields and absolute error maps in Figure 7 show that the standard DDROM suffers from rapid error accumulation, with unphysical oscillations primarily localized around the vortical structures.

This behavior is further quantified in Table 2. The unclosed DDROM yields a relative L_2 error of 4.24×10^{-1} and a maximum absolute error of 1.40, whereas the Closure-DDROM reduces these errors to 3.16×10^{-2} and 1.22×10^{-1}, respectively. This corresponds to an error reduction of approximately one order of magnitude in both metrics. These results demonstrate that, for the multi-injector case, the learned closure term significantly improves the final-time reconstruction accuracy and preserves the main physical structures of the high-fidelity trajectory. A trajectory-level error analysis will be useful in future

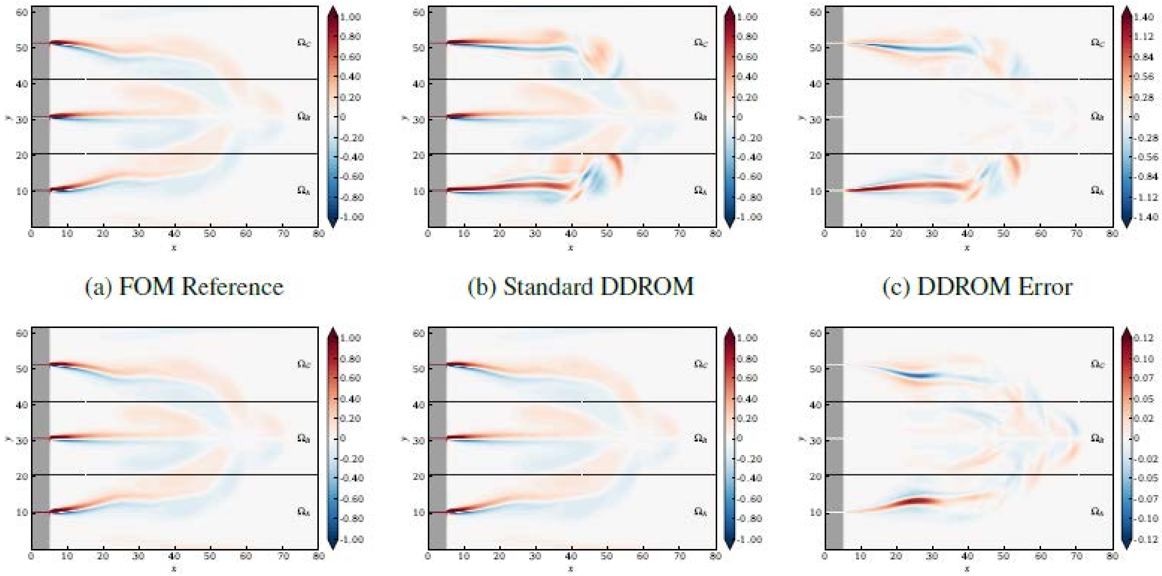

(a) FOM Reference (b) Standard DDROM (c) DDROM Error

(d) FOM Reference (e) Closure-DDROM (f) Closure-DDROM Error

Figure 7: Reconstructed flow fields and absolute error maps for the multi-injector configuration.

work to further assess whether this improvement is maintained throughout the full online rollout.

Table 2: Quantitative error metrics for the multi-injector flow.

Model	Mean Abs. Error	Max Abs. Error	L_2 Error	Relative L_2
Standard DDROM	3.71×10^{-2}	1.40×10^{0}	2.85×10^{2}	4.24×10^{-1}
Closure-DDROM	3.02×10^{-3}	1.22×10^{-1}	2.13×10^{1}	3.16×10^{-2}

4. Conclusion

This work investigated a closure-enhanced domain-decomposed reduced-order modeling framework for local POD-Galerkin ROMs. The baseline DDROM was constructed by applying Galerkin projection independently on local POD spaces, while the closure term was trained to approximate the residual discrepancy between the projected full-order dynamics and the unclosed reduced dynamics. Motivated by the Mori–Zwanzig interpretation of unresolved effects, the closure was represented by a GRU model that uses finite histories of reduced coordinates during online prediction.

The approach was assessed on two decomposed numerical examples. For the advection-diffusion-reaction problem, the Closure-DDROM reduced the final-time relative L_2 error from 4.96×10^{-2} to 1.74×10^{-2} compared with the standard DDROM. For the multi-injector vorticity-streamfunction flow, the relative L_2 error decreased from 4.24×10^{-1} to 3.16×10^{-2}. The corresponding error maps show that the closure correction reduces the spatial error levels in the tested online rollouts.

The predictive capability and algorithmic robustness of the framework were validated through both linear and non-linear transport systems. In the ADR case, the closure-augmented model achieved a substantial enhancement in reconstruction accuracy over the baseline. More importantly, in the highly non-linear multi-injector flow problem, where the standard unclosed DDROM suffered from rapid error accumulation and noticeable perturbations within the vortex structures, the proposed Closure-DDROM successfully suppressed these unphysical oscillations. By stabilizing the interfacial state exchanges, the closed system accurately tracked the high-fidelity trajectory and preserved the essential vortex dynamics.

These results suggest that residual-based neural closures provide a practical way to improve the accuracy of local POD-Galerkin DDROMs without replacing the intrusive reduced operator. At the same time, the present study is limited to the reported test configurations. Future work will focus on the theoretical and numerical characterization of the proposed closure formulation. Key issues include the long-time stability of the closure-enhanced reduced system, parameter generalization to unseen physical or geometrical configurations, and the sensitivity of the method to the memory length, POD dimension, and neural-network architecture. A stricter treatment of interface consistency is also needed to better control inter-subdomain information exchange. Finally, extensions to three-dimensional flow configurations and broader comparisons with alternative closure, stabilization, and interface-coupling strategies will be important steps toward a more general and reliable domain-decomposed reduced-order modeling framework.

Acknowledgements

This work was partially supported by the National Natural Science Foundation of China (Grant 12372284).

References

[1] Lawrence Sirovich. Turbulence and the dynamics of coherent structures. i. coherent structures. *Quarterly of Applied Mathematics*, 45(3):561–571, 1987.

[2] Alfio Quarteroni, Andrea Manzoni, and Federico Negri. *Reduced Basis Methods for Partial Differential Equations: An Introduction*. Springer, Cham, 2016.

[3] Jan S Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.

[4] Joan Baiges, Ramon Codina, and Sergio Idelsohn. A domain decomposition strategy for reduced order models. application to the incompressible navier–stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 267:23–42, 2013.

[5] Sabrina Kelbij Star, Benjamin Sanderse, Giovanni Stabile, Gianluigi Rozza, and Joris Degroote. Reduced order models for the incompressible navier-stokes equations on collocated grids using a 'discretize-then-project'approach. *International Journal for Numerical Methods in Fluids*, 93(8):2694–2722, 2021.

[6] Tommaso Taddei, Xuejun Xu, and Lei Zhang. A non-overlapping optimization-based domain decomposition approach to component-based model reduction of incompressible flows. *Journal of Computational Physics*, 509:113038, 2024.

[7] Shenhui Ruan, Andreas G Class, and Gianluigi Rozza. Introduction to the combination of reduced order models and domain decomposition: State of the art and perspectives. *arXiv preprint arXiv:2601.09623*, 2026.

[8] Cheng Huang, Karthik Duraisamy, and Charles Merkle. Component-based reduced order modeling of large-scale complex systems. *Frontiers in Physics*, 10:900064, 2022.

[9] Amy de Castro, Pavel Bochev, Paul Kuberry, and Irina Tezaur. Explicit synchronous partitioned scheme for coupled reduced order models based on composite reduced bases. *Computer Methods in Applied Mechanics and Engineering*, 417:116398, 2023.

[10] Niccolò Discacciati and Jan S Hesthaven. Model reduction of coupled systems based on non-intrusive approximations of the boundary response maps. *Computer Methods in Applied Mechanics and Engineering*, 420:116770, 2024.

[11] Hazime Mori. Transport, collective motion, and brownian motion. *Progress of Theoretical Physics*, 33(3):423–455, 1965.

[12] Robert Zwanzig. Memory effects in irreversible thermodynamics. *Physical Review*, 124(4):983–992, 1961.

[13] Dror Givon, Raz Kupferman, and Ole H Hald. Existence proof for orthogonal dynamics and the

mori-zwanzig formalism. *Israel Journal of Mathematics*, 145(1):221–241, 2005.

[14] Qian Wang, Nicolò Ripamonti, and Jan S Hesthaven. Recurrent neural network closure of parametric pod-galerkin reduced-order models based on the mori-zwanzig formalism. *Journal of Computational Physics*, 410:109402, 2020.

[15] Romit Maulik, Arvind Mohan, Bethany Lusch, Sandeep Madireddy, Prasanna Balaprakash, and Daniel Livescu. Time-series learning of latent-space dynamics for reduced-order model closure. *Physica D: Nonlinear Phenomena*, 405:132368, 2020.

[16] Abhinav Gupta and Pierre FJ Lermusiaux. Neural closure models for dynamical systems. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 477(2252), 2021.

[17] Yen Ting Lin, Yifeng Tian, Danny Perez, and Daniel Livescu. Regression-based projection for learning mori–zwanzig operators. *SIAM Journal on Applied Dynamical Systems*, 22(4):2890–2926, 2023.

[18] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

[19] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

[20] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734, 2014.

[21] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[23] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.

[24] Barry F. Smith, Petter E. Bjørstad, and William D. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, Cambridge, 1996.

[25] Alfio Quarteroni and Alberto Valli. *Domain decomposition methods for partial differential equations*. Oxford University Press, 1999.