Lagrangian Neural Fields for Infinite-Dimensional Dynamics Modeling

Riccardo Morandi and Noémie Jaquier

***Abstract*— Modeling continuum systems is challenging due to their infinite-dimensional nature, leading most approaches to rely on discretizations that introduce bias and reduce generalization. This paper introduces Lagrangian neural fields, a neural architecture that learns finite-dimensional Lagrangian surrogate models of infinite-dimensional continuum dynamics. By taking a geometric perspective on continuum mechanics and building on recent neural implicit representations, Lagrangian neural fields learns physically-consistent and discretization-agnostic reduced-order models. Our experiment showcase their potential in enabling reliable long-term predictions of infinite-dimensional dynamical systems.**

I. Introduction

Modeling the dynamics of high-dimensional physical systems remains a key challenge in robotics, especially for deformable objects and soft robots [1]. Such systems evolve on highly non-linear, infinite-dimensional configuration spaces that are notoriously hard to model. The existing literature coalesces around the few following key paradigms. **Mesh-based** approaches represent states via grids, particles [2], or graphs [3]–[5], employing convolutional [6] and graph neural networks. While capable of handling irregular domains, these approaches typically disregard fundamental physics laws and suffer from mesh-dependent biases, exhibiting limited generalization across resolutions and geometries.

Physics-informed neural networks incorporate physical priors using Lagrangian [7]–[9] or Hamiltonian formulations [10], enforcing conservation laws and improving data efficiency. However, they are restricted to finite-dimensional systems, often low-dimensional in practice, or rely on explicit spatial discretization for high-dimensional extensions [11], [12]. As for mesh-based approaches, this constrains the learned dynamics to a particular resolution, thereby hindering both transferability and applicability to infinite-dimensional systems. **Operator learning** addresses these limitations by learning mappings between function spaces [13]. The two dominant paradigms in this area, i.e., DeepONets [14] and Fourier neural operators (FNOs) [15], are constrained to specific discretization structures for evaluation and mostly rely on implicit linear dimensionality reduction with limited expressivity [16]. More recent approaches leverage continuous representations from transformers [17], [18] and implicit neural representation (INR) [19]–[21] to alleviate the dependency on predefined discretization. However, classical operator learning often disregards physical laws, relying instead on generic operator approximation principles and universal functional approximation theorems. While this ensures high flexibility, it leads to low data efficiency compared to physics-informed approaches [22].

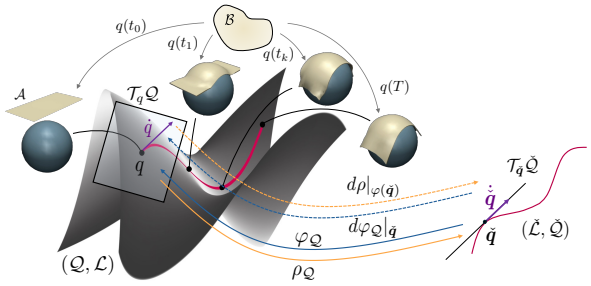

Fig. 1. Geometry of continuum mechanics. Lagrangian neural fields learn a physically-consistent, finite-dimensional surrogate model of the infinite-dimensional continuum dynamics via an embedding φ and a reduction ρ.

Model order reduction (MOR) constructs low-dimensional surrogates, also known as reduced-order models (ROMs), of high-dimensional physical systems [23]. Linear projections-based ROMs [24], [25] are increasingly giving way to non-linear models [26], [27], predominantly using autoencoders [28]–[30] or INRs [20], [31]. Structure-preserving MOR [24] retains the Lagrangian or Hamiltonian structure of the full-order model (FOM) in the ROM, thereby ensuring physical consistency. While most structure-preserving MOR methods are intrusive as they assume a fully-known FOM [25], [28], [30], [32], recent extensions consider unknown FOM dynamic parameters [11], [12], [33]. MOR relies on explicit dimensionality reduction in state space, while operator learning implicitly learns reduced function representations through network architecture.

Taking inspiration from INR-based operator learning [21] and structure-preserving Lagrangian MOR [11], [34], we propose *Lagrangian neural fields*, a neural architecture that learns physically-consistent ROMs of infinite-dimensional dynamical systems. Taking a geometric perspective on continuum mechanics, we represent the dynamics of deformable objects on the infinite-dimensional manifold of embeddings. **Our contributions** are twofold: *(1)* We derive a structure-preserving and discretization-independent ROM for continuum mechanics yielding low-dimensional Lagrangian dynamics; and *(2)* We design an intrusive INR-based ROM that ensures both physical consistency and discretization independence. We validate the proposed Lagrangian neural fields on a toy infinite-dimensional system, showcasing its ability to accurately capture structure-preserving reduced dynamics.

II. Geometry of Continuum Mechanics

Geometric mechanics [35], [36] represents the configuration space as a smooth manifold $\mathcal{Q}$ [37], [38]. The system

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program funded by the Knut and Alice Wallenberg Foundation.

The authors are with the Department of Robotics, Perception and Learning at KTH Royal Institute of Technology, Stockholm, Sweden.

states evolve on the tangent bundle $(q, \dot{q}) \in \mathcal{TQ}$, with $q \in \mathcal{Q}$ and $\dot{q} \in \mathcal{T}_q\mathcal{Q}$ denoting generalized positions and velocities, respectively. Generalized forces $\tau \in \mathcal{T}_q^*\mathcal{Q}$ belong to the cotangent bundle $\mathcal{T}^*\mathcal{Q}$. The system dynamics correspond to vector fields on $\mathcal{TQ}$, whose integral curves are system trajectories. We follow the formulation for infinite-dimensional dynamics, which encompasses the finite-dimensional case [35].

A Lagrangian system is a tuple $(\mathcal{Q}, \mathcal{L})$ of a configuration manifold $\mathcal{Q}$ and a smooth functional $\mathcal{L} : \mathcal{TQ} \to \mathbb{R}$. Its dynamics follow from the Euler-Lagrange equations

$$\frac{d}{dt}\frac{\delta \mathcal{L}}{\delta \dot{q}} - \frac{\delta \mathcal{L}}{\delta q} = 0. \tag{1}$$

A quadratic kinetic energy in $\dot{q}$ induces a Riemannian metric on $\mathcal{Q}$, effectively geometricizing the system's inertia [35].

Continuum mechanics describes deformable systems, such as soft robots and elastic objects, which both translate and deform under applied forces. These systems can be formulated geometrically as Lagrangian systems [39]–[41]. Formally, we represent the reference configuration of an elastic body as an assembly of particles constituting a compact, orientable smooth manifold $\mathcal{B}$ with boundary $\partial\mathcal{B}$. The body deforms in an ambient space, i.e., the real-world, $\mathcal{A} = \mathbb{R}^3$. A configuration of the elastic body corresponds to a map $q : \mathcal{B} \to \mathcal{A}$, assigning a spatial position $x \in \mathcal{A}$ to each material point $X \in \mathcal{B}$, see Fig. 1. Restricting q to the Fréchet manifold of smooth embeddings $\mathcal{Q} = \mathrm{Emb}^\infty(\mathcal{B}, \mathcal{A})$ [42] ensures injectivity, enforcing non-interpenetration of matter [39]. In this setting, tangent and cotangent vectors are velocity fields and force densities, respectively. Continuum dynamics are governed by Lagrangian mechanics [36] and characterized by the Lagrangian functional defined as

$$\mathcal{L}(q, \dot{q}) = T(\dot{q}) - V(q, \dot{q}), \quad \text{with} \tag{2}$$

$$T(q, \dot{q}) = \frac{1}{2}\int_{\mathcal{B}} \omega_0(X)\|\dot{q}(X)\|^2 \, dX, \tag{3}$$

$$V(q) = \int_{\mathcal{B}} W(\nabla_X q(X)) \, dX + \int_{\mathcal{B}} U(q(X)) \, dX, \tag{4}$$

where $\omega_0 : \mathcal{B} \to \mathbb{R}^+$ is the mass density and $W, U : \mathcal{B} \to \mathbb{R}$ are the internal and external potential energy densities. The kinetic energy (3) induces the Riemannian metric

$$m_q(u, v) = \int_{\mathcal{B}} \omega_0(X) \, u(X) \cdot v(X) \, dX, \tag{5}$$

with $\cdot$ denoting the inner product in $\mathcal{A}$, corresponding to the ω_0-weighted L^2 inner product. The Euler-Lagrange equations (1), in strong form, lead to the equations of motion for continuum mechanics [43]. Note that, in the finite-dimensional case, (2)-(5) recover the well-known Lagrangian $\mathcal{L}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = \frac{1}{2}\dot{\boldsymbol{q}}^\mathsf{T}\boldsymbol{M}(\boldsymbol{q})\dot{\boldsymbol{q}} + V(\boldsymbol{q})$, with $q, \dot{q}$ and m_q corresponding to generalized coordinates $\boldsymbol{q}$, velocities $\dot{\boldsymbol{q}}$, and mass-inertia matrix $\boldsymbol{M}(\boldsymbol{q})$, respectively.

III. Lagrangian Neural Fields

In this section, we present Lagrangian neural fields, a neural architecture that learns structure-preserving ROMs of continuum systems by combining geometric MOR with INRs. The overall pipeline is illustrated in Fig. 2.

A. Structure-preserving MOR for Continuum Mechanics

First, we extend structure-preserving Lagrangian MOR to continuum mechanics. Formally, we define a structure-preserving ROM of an infinite-dimensional Lagrangian systems $(\mathcal{Q}, \mathcal{L})$ with $\mathcal{Q} = \mathrm{Emb}^\infty(\mathcal{B}, \mathcal{A})$ as a reduced, finite-dimensional Lagrangian system $(\check{\mathcal{Q}}, \check{\mathcal{L}})$. Following [34], we characterize the ROM by a smooth embedding $\varphi_{\mathcal{Q}} : \check{\mathcal{Q}} \to \mathcal{Q}$ learned such that the image of its lifted map $\varphi : \mathcal{T}\check{\mathcal{Q}} \to \mathcal{TQ}$ approximates the trajectories of the FOM, and an associated reduction map $\rho_{\check{\mathcal{Q}}} : \mathcal{Q} \to \check{\mathcal{Q}}$ (see Fig. 1). As in the finite-dimensional case [24], [34], $\check{\mathcal{L}}$ is constructed as the pullback $\check{\mathcal{L}} = \mathcal{L} \circ \varphi$. Moreover, $\check{\mathcal{Q}}$ and $\varphi_{\mathcal{Q}}$ are obtained such that $\varphi \circ \rho : \mathcal{TQ} \to \varphi(\mathcal{T}\check{\mathcal{Q}})$ defines a projection, i.e.,

$$\rho_{\check{\mathcal{Q}}} \circ \varphi_{\mathcal{Q}} = \mathrm{id}_{\check{\mathcal{Q}}}. \tag{6}$$

Ensuring invariance of $\varphi_{\mathcal{Q}}(\check{\mathcal{Q}})$ under the lift–project operator and consistency of the induced dynamics across projections.

Next, we derive the equations of motions of the reduced system $(\check{\mathcal{Q}}, \check{\mathcal{L}})$ as a function of the infinite-dimensional dynamics of $(\mathcal{Q}, \mathcal{L})$. The pullback $\check{\mathcal{L}}$ implies that the reduced trajectories satisfy the principle of least action $\delta \int \check{\mathcal{L}}(\check{\boldsymbol{q}}, \dot{\check{\boldsymbol{q}}}) dt = 0$, with variations $\delta q = d\varphi_{\mathcal{Q}}|_{\check{\boldsymbol{q}}} \, \delta\check{\boldsymbol{q}}$ constrained to $\varphi(\mathcal{T}\check{\mathcal{Q}})$. Taking variations and integrating in time yields

$$(d\varphi_{\mathcal{Q}}|_{\check{\boldsymbol{q}}})^* \mathcal{F}(\varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})) = \boldsymbol{0}, \tag{7}$$

where $\mathcal{F}(q) \in T_q^*\mathcal{Q}$ is the residual of (1), and $(\cdot)^*$ denotes the adjoint [44]. Expanding (7) yields a reduced Lagrangian composed of the reduced potential $\check{V}(\check{\boldsymbol{q}}) = V \circ \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})$ and the reduced quadratic kinetic energy inducing the pullback metric $\check{m}_{\check{\boldsymbol{q}}}(\check{\boldsymbol{u}}, \check{\boldsymbol{v}}) = m_{\varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}\big(d\varphi_{\mathcal{Q}}|_{\check{\boldsymbol{q}}}\check{\boldsymbol{u}}, d\varphi_{\mathcal{Q}}|_{\check{\boldsymbol{q}}}\check{\boldsymbol{v}}\big)$. In coordinates, the reduced equations of motion take the form

$$\check{\boldsymbol{M}}(\check{\boldsymbol{q}})\ddot{\check{\boldsymbol{q}}} + \check{\boldsymbol{c}}(\check{\boldsymbol{q}}, \dot{\check{\boldsymbol{q}}}) + \check{\boldsymbol{g}}(\check{\boldsymbol{q}}) = \boldsymbol{0}, \quad \text{with} \tag{8}$$

$$\check{\boldsymbol{M}}(\check{\boldsymbol{q}})_{ij} = \int_{\mathcal{B}} \omega_0 \, \frac{\partial \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}{\partial \check{q}_i} \cdot \frac{\partial \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}{\partial \check{q}_j} \, dX, \tag{9}$$

$$\check{\boldsymbol{c}}(\check{\boldsymbol{q}}, \dot{\check{\boldsymbol{q}}})_i = \left[\int_{\mathcal{B}} \omega_0 \, \frac{\partial \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}{\partial \check{q}_i} \cdot \frac{\partial^2 \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}{\partial \check{q}_j \partial \check{q}_k} \, dX\right] \dot{\check{q}}_j \dot{\check{q}}_k, \tag{10}$$

$$\begin{aligned}\check{\boldsymbol{g}}(\check{\boldsymbol{q}})_i = \int_{\mathcal{B}} \Bigg[&\frac{\partial W}{\partial F}(F(\check{\boldsymbol{q}})) : \nabla_X \left(\frac{\partial \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}{\partial \check{q}_i}\right) \\ &+ \nabla U(\varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})) \cdot \frac{\partial \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})}{\partial \check{q}_i}\Bigg] dX.\end{aligned} \tag{11}$$

We emphasize that (8) retains exactly the form of the equations of motion of finite-dimensional mechanical systems.

B. Model Architecture and Training

Next, we learn the mappings $\varphi_{\mathcal{Q}}$ and $\rho_{\check{\mathcal{Q}}}$ from discrete trajectory samples with no assumption on the spatial observation grid $\mathcal{X} = \{X_j\}_{j=1}^{|\mathcal{X}|}$. We consider the intrusive MOR case and assume known continuum dynamics parameters.

We parameterize the embedding $\varphi_{\mathcal{Q}}$ via a modulated SIREN [21], [45], [46], a discretization-independent INR that has shown promising results in both MOR [31] and operator learning [21]. Here, we train a modulated SIREN

$$\varphi_{\mathcal{Q}}(\check{\boldsymbol{q}}, X) = f_{\boldsymbol{\theta}, \boldsymbol{\phi}(\check{\boldsymbol{q}})}(X), \tag{12}$$

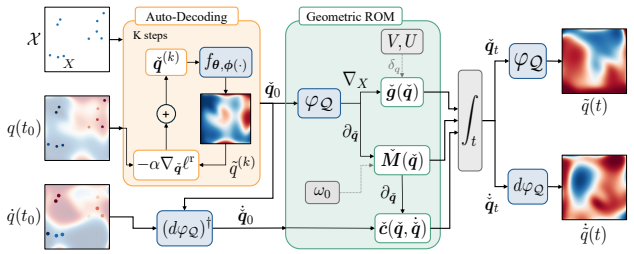

Fig. 2. Flowchart of the forward dynamics of Lagrangian neural fields.

to map material $X \in \mathcal{B}$ and reduced coordinates $\check{\boldsymbol{q}} \in \check{\mathcal{Q}}$ to positions $\boldsymbol{x} \in \mathcal{A}$. Importantly $\varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})$ is analytically differentiable in both $\check{\boldsymbol{q}}$ and X, enabling closed-form expressions of the derivative terms in (9)-(11).

In order to fulfill the projection property (6), we assume the lifted reduction map $\rho = (\rho_{\check{\mathcal{Q}}}, d\rho_{\check{\mathcal{Q}}})$ to be the ω_0-orthogonal projection on $\varphi(\mathcal{T}\check{\mathcal{Q}})$, implicitly defined by

$$\rho_{\check{\mathcal{Q}}}(q) = \operatorname{argmin}_{\check{\boldsymbol{q}} \in \check{\mathcal{Q}}} \left\| q - \varphi_{\mathcal{Q}}(\check{\boldsymbol{q}}) \right\|_{\omega_0}^2, \tag{13}$$

$$d\rho_{\check{\mathcal{Q}}}|_{\varphi_{\mathcal{Q}}(\check{\boldsymbol{q}})} = (d\varphi_{\mathcal{Q}}|_{\check{\boldsymbol{q}}})^{\dagger,\omega_0}, \tag{14}$$

where $(\cdot)^{\dagger,\omega_0}$ denotes the ω_0-weighted generalized inverse. We obtain an approximation $\tilde{\rho}(q, \varphi_{\mathcal{Q}})$ of the reduction $\rho_{\check{\mathcal{Q}}}$ (13) via K auto-decoding steps [47]

$$\check{\boldsymbol{q}}^{(k)} = \check{\boldsymbol{q}}^{(k-1)} - \alpha \nabla_{\check{\boldsymbol{q}}} \ell^{\mathrm{r}}\Big(q, \varphi_{\mathcal{Q}}\big(\check{\boldsymbol{q}}^{(k-1)}\big)\Big), \tag{15}$$

with $\check{\boldsymbol{q}}^{(0)} = \mathbf{0}$, learning rate $\alpha \in \mathbb{R}^+$, and reconstruction loss

$$\ell^{\mathrm{r}}(q_1, q_2) = \frac{1}{|\mathcal{X}|} \sum_{j=1}^{|\mathcal{X}|} \omega_0(X_j) \|q_1(X_j) - (X_j)\|^2. \tag{16}$$

We train the network (12) using the composite loss

$$\ell(q, \varphi_{\mathcal{Q}}) = \ell^{\mathrm{r}}\big(q, \tilde{\rho}(q, \varphi_{\mathcal{Q}})\big) + \ell^{\mathrm{p}}(q, \varphi_{\mathcal{Q}}), \tag{17}$$

where ℓ^{r} promotes reconstruction accuracy and

$$\ell^{\mathrm{p}}(q, \varphi_{\mathcal{Q}}) = \Big\| \tilde{\rho}\big(q, \varphi_{\mathcal{Q}}\big) - \tilde{\rho}\Big(q, \varphi_{\mathcal{Q}}\big(\tilde{\rho}(q, \varphi_{\mathcal{Q}})\big)\Big) \Big\|_{\check{m}}^2 \tag{18}$$

imposes the projection property (6) as a soft constraint.

Training is performed by minimizing (17) with respect to the SIREN parameters, where $\tilde{\rho}(q, \varphi_{\mathcal{Q}})$ is treated as a differentiable function by unrolling and applying backpropagation through the K gradient descent steps (15). During training, the loss is computed on each snapshot independently. At inference, the initial condition is projected onto $\mathcal{T}\check{\mathcal{Q}}$ as $(\check{\boldsymbol{q}}_0, \dot{\check{\boldsymbol{q}}}_0) = (\tilde{\rho}(q(t_0), \varphi_{\mathcal{Q}}), (d\varphi_{\mathcal{Q}}|_{\check{\boldsymbol{q}}_0})^{\dagger,\omega_0} \dot{q}(t_0))$. The reduced dynamics (8) are then integrated forward in time, and the configuration is reconstructed via $\varphi_{\mathcal{Q}}$, as shown in Fig. 2.

IV. Experiments

We evaluate Lagrangian neural fields on a synthetic example. We consider an infinite-dimensional system constructed by lifting the dynamics of a 1-DoF pendulum with angle θ into $\mathcal{Q} = L^2(\Omega)$, $\Omega = [-1, 1]^2$. We construct the field observations as $q(X, t) = u(X)\,\theta(t)$ via a normalized tensor product of Legendre polynomials u on Ω, so that the dynamics effectively lie on a 1-dimensional embedded manifold $\check{\mathcal{Q}} \subset \mathcal{Q}$. We equip this Hilbert manifold with the Lagrangian

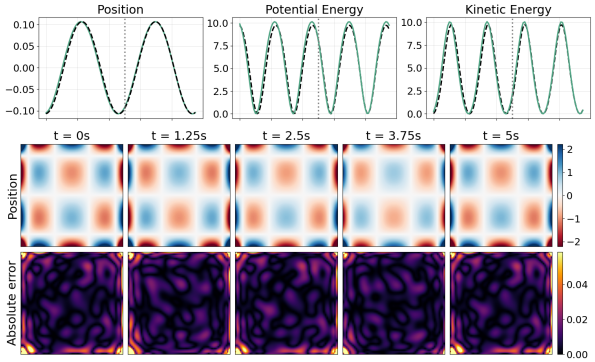

Fig. 3. Experimental results on a synthetic example. *Top*: Comparison between predicted (—) and projected (- -) position, and energies. (:) indicates the training horizon. *Bottom*: predicted fields on Ω and associated absolute errors, both before and after the training horizon of $2.5s$.

functional $\mathcal{L}(q, \dot{q}) = \frac{1}{2} m\ell^2 \|\dot{q}\|_{L^2}^2 + mg\ell \left(1 - \cos\big(\langle q, u\rangle_{L^2}\big)\right)$, where $\langle f, u\rangle_{L^2} = \int_\Omega f(x)u(x)\,dx$. Under the identification $\theta = \langle f, u\rangle_{L^2}$, this functional is exactly equivalent to the 1-DoF pendulum Lagrangian. The ambient space $L^2(\Omega)$ is therefore interpreted as a Hilbert manifold equipped with the weighted L^2 metric induced by the kinetic energy $m_f(u, v) = m\ell^2 \langle u, v\rangle_{L^2}$. As such, this construction yields an analytically-tractable and well-reducible benchmark.

We simulate a pendulum with parameters $m = \ell = 1$ and $g = 9.81$ over $t \in [0, 5]$. We generate 100 timesteps per trajectory, each obtained from independent uniformly-sampled initial conditions. Training data are obtained by lifting the pendulum trajectories using tensorized Legendre polynomials of degree $(3, 4)$ on a 64×64 grid. During training, observations are restricted to 25% of a 32×32 grid, which is fixed for each trajectory but varies across samples. We use the first 50 timesteps for training and the remaining 50 for testing. At test time, predictions are evaluated over all 100 timesteps on the full 64×64 grid, assessing both temporal and spatial extrapolation. The embedding is modeled by a modulated SIREN with latent dimension $d = 1$, depth 2, width 256, and $\nu_0 = 30$, with shift modulation via a 1-layer hypernetwork of width 128. The reduction is performed via 3 auto-decoding steps with learned step size α, as in [21].

Fig. 3 shows representative results. The top panel shows the that latent reduced dynamics evolved in $\check{\mathcal{Q}}$ closely match the projected FOM dynamics, thereby validating the reduced model. Moreover, the accurate recovery of kinetic and potential energies validates its structure-preserving properties. The bottom panel demonstrates high-quality reconstruction on a finer grid than trained on, validating the discretization agnostic character of the learned representation. Finally, evaluations beyond the training time horizon showcases the temporal extrapolation capabilities of the model.

V. Conclusions and Future Work

We presented Lagrangian neural fields, a structure-preserving, discretization-agnostic MOR framework for modeling the infinite-dimensional dynamics of continuum mechanical systems. Future work will focus on extensions to non-conservative forces and to the non-intrusive regime, as well as on validations on deformable objects and soft robots.

REFERENCES

[1] B. Ai, S. Tian, H. Shi, Y. Wang, T. Pfaff, C. Tan, H. I. Christensen, H. Su, J. Wu, and Y. Li, "A review of learning-based dynamics models for robotic manipulation," *Science Robotics*, vol. 10, no. 106, pp. 14–97, 2025.

[2] K. Zhang, B. Li, K. Hauser, and Y. Li, "Particle-grid neural dynamics for learning deformable object models from RGB-D videos," in *Robotics: Science and Systems (R:SS)*, 2025.

[3] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba, "Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids," in *Intl. Conf. on Learning Representations (ICLR)*, 2019.

[4] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *Intl. Conf. on Machine Learning (ICML)*, 2020.

[5] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, "Learning mesh-based simulation with graph networks," in *Intl. Conf. on Learning Representations (ICLR)*, 2020.

[6] S. Chae, J. Shin, S. Kwon, S. Lee, S. Kang, and D. Lee, "PM10 and PM2.5 real-time prediction models using an interpolated convolutional neural network," *Scientific Reports*, vol. 11, no. 1, p. 11952, 2021.

[7] M. Lutter, C. Ritter, and J. Peters, "Deep Lagrangian networks: Using physics as model prior for deep learning," in *Intl. Conf. on Learning Representations (ICLR)*, 2019.

[8] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, "Lagrangian neural networks," in *ICLR Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.

[9] M. Lutter and J. Peters, "Combining physics and deep learning to learn continuous-time dynamics models," *Intl. Journal of Robotics Research*, vol. 42, no. 3, pp. 83–107, 2023.

[10] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," in *Neural Information Processing Systems (NeurIPS)*, 2019.

[11] K. Friedl, N. Jaquier, J. Lundell, T. Asfour, and D. Kragic, "A Riemannian framework for Learning Reduced-order Lagrangian Dynamics," in *Intl. Conf. on Learning Representations (ICLR)*, 2025.

[12] K. Friedl, N. Jaquier, M. Liao, and D. Kragic, "Learning Hamiltonian dynamics at scale: A differential-geometric approach," *arXiv preprint arXiv:2509.24627*, 2025.

[13] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, and A. Anandkumar, "Neural operator: Learning maps between function spaces," *Journal of Machine Learning Research*, vol. 24, no. 89, pp. 1–97, 2023.

[14] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis, "Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators," *Nature Machine Intelligence*, vol. 3, no. 3, pp. 218–229, 2021.

[15] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, "Fourier neural operator for parametric partial differential equations," in *Intl. Conf. on Learning Representations (ICLR)*, 2021.

[16] J. H. Seidman, G. Kissas, P. Perdikaris, and G. J. Pappas, "NOMAD: Nonlinear manifold decoders for operator learning," in *Neural Information Processing Systems (NeurIPS)*, 2022.

[17] Z. Li, K. Meidani, and A. B. Farimani, "Transformer for partial differential equations' operator learning," *Transactions on Machine Learning Research*, 2022.

[18] Z. Hao, Z. Wang, H. Su, C. Ying, Y. Dong, S. Liu, Z. Cheng, J. Song, and J. Zhu, "Gnot: a general neural operator transformer for operator learning," in *Intl. Conf. on Machine Learning (ICML)*, 2023.

[19] Y. Yin, M. Kirchmeyer, J.-Y. Franceschi, A. Rakotomamonjy, and P. Gallinari, "Continuous PDE dynamics forecasting with implicit neural representations," in *Intl. Conf. on Learning Representations (ICLR)*, 2023.

[20] S. Pan, S. L. Brunton, and J. N. Kutz, "Neural implicit flow: a mesh-agnostic dimensionality reduction paradigm of spatio-temporal data," *Journal of Machine Learning Research*, vol. 24, no. 41, pp. 1–60, 2023.

[21] L. Serrano, L. L. Boudec, A. K. Koupaï, T. X. Wang, Y. Yin, J.-N. Vittaut, and P. Gallinari, "Operator learning with neural fields: Tackling PDEs on general geometries," in *Neural Information Processing Systems (NeurIPS)*, 2023.

[22] Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Azizzadenesheli, and A. Anandkumar, "Physics-informed neural operator for learning partial differential equations," *ACM / IMS Journal of Data Science*, vol. 1, no. 3, pp. 1–27, 2024.

[23] W. H. A. Schilders, H. A. Vorst, and J. Rommes, *Model Order Reduction: Theory, Research Aspects and Applications*. Springer, 2008.

[24] S. Lall, P. Krysl, and J. E. Marsden, "Structure-preserving model reduction for mechanical systems," *Physica D: Nonlinear Phenomena*, vol. 184, no. 1-4, pp. 304–318, 2003.

[25] K. Carlberg, R. Tuminaro, and P. Boggs, "Preserving Lagrangian structure in nonlinear model reduction with application to structural dynamics," *SIAM Journal on Scientific Computing*, vol. 37, no. 2, pp. B153–B184, 2015.

[26] J. Barnett and C. Farhat, "Quadratic approximation manifold for mitigating the kolmogorov barrier in nonlinear projection-based model order reduction," *Journal of Computational Physics*, vol. 464, 2022.

[27] H. Sharma, H. Mu, P. Buchfink, R. Geelen, S. Glas, and B. Kramer, "Symplectic model reduction of Hamiltonian systems using data-driven quadratic manifolds," *Computer Methods in Applied Mechanics and Engineering*, vol. 417, p. 116402, 2023.

[28] K. Lee and K. T. Carlberg, "Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders," *Journal of Computational Physics*, vol. 404, p. 108973, 2020.

[29] S. Fresca and A. Manzoni, "POD-DL-ROM: Enhancing deep learning-based reduced order models for nonlinear parametrized PDEs by proper orthogonal decomposition," *Computer Methods in Applied Mechanics and Engineering*, vol. 388, p. 114181, 2022.

[30] S. E. Otto, G. R. Macchio, and C. W. Rowley, "Learning nonlinear projections for reduced-order modeling of dynamical systems using constrained autoencoders," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 33, no. 11, p. 113130, 2023.

[31] P. Y. Chen, J. Xiang, D. H. Cho, Y. Chang, G. A. Pershing, H. T. Maia, M. M. Chiaramonte, K. Carlberg, and E. Grinspun, "CROM: Continuous reduced-order modeling of PDEs using implicit neural representations," in *Intl. Conf. on Learning Representations (ICLR)*, 2023.

[32] J. S. Hesthaven, C. Pagliantini, and N. Ripamonti, "Structure-preserving model order reduction of Hamiltonian systems," *International congress of mathematicians*, vol. 7, pp. 5072–5097, 2022.

[33] H. Sharma and B. Kramer, "Preserving Lagrangian structure in data-driven reduced-order modeling of large-scale dynamical systems," *Physica D: Nonlinear Phenomena*, vol. 462, p. 134128, 2024.

[34] P. Buchfink, S. Glas, B. Haasdonk, and B. Unger, "Model reduction on manifolds: A differential geometric framework," *Physica D: Nonlinear Phenomena*, vol. 468, p. 134299, 2024.

[35] J. E. Marsden and T. S. Ratiu, *Introduction to Mechanics and Symmetry, A Basic Exposition of Classical Mechanical Systems*. Springer, 1999.

[36] J. Marsden and T. J. R. Hughes, *Mathematical foundations of elasticity*. Dover Publications, Inc., 1983.

[37] J. M. Lee, *Introduction to Smooth Manifolds*. Springer, 2012.

[38] J. Lee, *Introduction to Riemannian Manifolds*, ser. Graduate Texts in Mathematics. Springer, 2018.

[39] W. Noll, "A new mathematical theory of simple materials," *Archive for Rational Mechanics and Analysis*, vol. 48, no. 1, pp. 1–50, 1972.

[40] Z. Fiala, "Geometrical setting of solid mechanics," *Annals of Physics*, vol. 326, no. 8, pp. 1983–1997, 2011.

[41] B. Kolev and R. Desmorat, "An intrinsic geometric formulation of hyper-elasticity, pressure potential and non-holonomic constraints," *Journal of Elasticity*, vol. 146, no. 1, pp. 29–63, 2021.

[42] R. S. Hamilton, "The inverse function theorem of nash and moser," *Bulletin of the American Mathematical Society*, vol. 7, no. 1, pp. 65–222, 1982.

[43] F. Irgens, *Continuum Mechanics*. Springer, 2008.

[44] H. Brezis, *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, 2011.

[45] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. Courville, "Film: visual reasoning with a general conditioning layer," in *AAAI Conf. on Artificial Intelligence*, 2018.

[46] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," in *Neural Information Processing Systems (NeurIPS)*, 2020.

[47] S. Tan and M. L. Mayrovouniotis, "Reducing data dimensionality through optimizing neural network inputs," *AIChE Journal*, vol. 41, no. 6, pp. 1471–1480, 1995.