

Neural Modes: Self-supervised Learning of Nonlinear Modal Subspaces

Jiahong Wang, Yinwei Du, Stelian Coros, Bernhard Thomaszewski
 ETH Zürich

{wangjiah, duyin, scoros, bthomasz}@ethz.ch



Figure 1. Visualization of our nonlinear subspace on a giraffe model discretized with tetrahedron finite elements. Overlays correspond to different frames from a selected neural mode.

Abstract

We propose a self-supervised approach for learning physics-based subspaces for real-time simulation. Existing learning-based methods construct subspaces by approximating pre-defined simulation data in a purely geometric way. However, this approach tends to produce high-energy configurations, leads to entangled latent space dimensions, and generalizes poorly beyond the training set. To overcome these limitations, we propose a self-supervised approach that directly minimizes the system’s mechanical energy during training. We show that our method leads to learned subspaces that reflect physical equilibrium constraints, resolve overfitting issues of previous methods, and offer interpretable latent space parameters.

1. Introduction

Physics-based simulation is an indispensable tool in science and engineering. Computer graphics has since long relied on simulation to create stunning visual effects and

animations. However, while our ability to create ever more realistic simulations is steadily increasing, physics-based simulations are notoriously time-consuming, and translating techniques to the real-time regime remains a major challenge. One particular strategy to reduce computational burden without sacrificing visual quality is model reduction [13, 20, 35]. These techniques construct low-dimensional subspaces that aim to capture the variability and expressiveness of full-space simulations. However, designing subspaces that strike a balance between visual fidelity and computational complexity is no easy task. While linear vibration modes are ideal for small deformations, they fail to capture motions involving large deformations and rotations.

Instead of approximating nonlinear simulations with a linear subspace, recent work has started to explore machine learning techniques to build nonlinear subspaces from simulation data [8]. While this learning-based approach to subspace simulation has shown its promise, existing methods rely on supervised training, which leads to a number of challenges and limitations. First, supervised learning requires significant amounts of pre-defined and curated training data. Second, this data is typically approximated in a purely geometric way, e.g., by minimizing the reconstruction loss of a deep autoencoder (AE) [18]. This geometric approach, however, ignores the fact that simulation data live on a nonlinear manifold defined through mechanical equilibrium conditions. Consequently, configurations from geometrically-learned subspaces generally exhibit much higher internal energy than the input data, which manifests as stress concentrations and spurious forces that affect the quality and performance of simulations. Another problem is the empirical observation that supervised learning often leads to overfitting and generalizes poorly to configurations outside the training set. We confirm this observation for the case of nonlinear subspace simulation through a series of experiments. Finally, the learned representation of the subspace typically lacks structure and is often strewn with singularities and discontinuities. In particular, latent space dimensions are rarely meaningful (neither from energetic, visual, nor semantic points of view), and smooth trajectories in latent space typically lead to incoherent ge-

ometry.

To address these challenges, we propose a novel self-supervised approach to learning nonlinear subspaces. Rather than building geometry-based subspaces from simulation data, our method learns physics-based subspaces by directly minimizing the energy of the underlying mechanical system. We draw inspiration from Nonlinear Compliant Modes, an extension of linear modes to the large deformation regime proposed by Dünser et al. [7]. Each nonlinear mode is parameterized by a modal coordinate defined as the projection of a given configuration onto the corresponding linear mode. For any modal coordinate value, the corresponding nonlinear mode shape is obtained by minimizing the elastic energy in the space orthogonal to the linear mode. To construct our neural subspaces, we extend Nonlinear Compliant Modes to nonlinear modal spaces that capture the behavior along and across a given set of modes. Crucially, since the map between modal coordinates and nonlinear equilibrium shapes is given by a constrained minimization problem, we can train our neural network to learn the entire solution manifold by minimizing the corresponding energy across parameter space. Once trained, our method allows for real-time exploration and gradient-based navigation of nonlinear modal subspaces. We demonstrate the potential of our approach on a range of examples spanning real-time dynamics simulation, physics-based nonlinear deformations, and keyframe animation. We furthermore analyze the performance of our method through ablation studies and comparisons to baselines, revealing that our self-supervised approach outperforms existing methods based on supervised learning. In summary, our work makes the following contributions:

- We present the first self-supervised approach for learning nonlinear modal subspaces for simulation. Our approach eliminates the need for creating curated data collections, enables end-to-end physics-based training, and has quantified physical accuracy.
- We extend nonlinear compliant modes from single modes to fully coupled modal spaces for exploration.
- We analyze the quality of our learned subspaces and existing alternatives with respect to accuracy, smoothness, and interpretability. While existing data-free methods [32] can lead to mode collapse, we show that our formulation avoids this problem. We further compare the performance of our solution to baselines on a set of subspace simulation examples.

2. Related work

Subspace Simulation. Model reduction or subspace simulation is an established technique with broad applications in many scientific disciplines. Arguably, the most widely used approach is linear modal analysis [31], initially designed for vibration problems. The basic idea of this approach is to de-

fine a low-dimensional linear basis using linear modes, i.e., eigenvectors of the energy Hessian corresponding to low-frequency deformations. Despite its efficiency for small deformations, linear eigenmodes are known to suffer from severe distortion artifacts for large deformations [14].

Various methods aimed to reduce artefacts from linear eigenmodes by augmenting the basis with carefully selected vectors. Some notable examples include modal derivatives [2], higher-order Krylov-type modes [38], higher-order descent directions [14], modal warping that corrects for rotations [5], and rotation strain extrapolation [16]. Nonlinear compliant modes [7] are arguably the approach most closely related to our work. These modes are found by constraining the projection of the system’s configuration onto the linear eigenmodes while minimizing its internal energy in the orthogonal subspace. The method is fully physics-based and has been demonstrated for various analysis tasks on real-world 3D printed examples. However, in their original formulation, nonlinear compliant modes are inherently one-dimensional and do not readily extend to multi-dimensional subspace. We extend the concept of nonlinear compliant modes to complete modal subspaces, enabling exploration and simulation.

Nonlinear Subspaces. While the majority of works focus on modal analysis, alternative strategies for constructing nonlinear subspace have been explored as well. Rig-space approaches solve forward simulation and inverse kinematics tasks in the nonlinear subspaces induced by character rigs [11, 12, 19, 26]. However, these methods are tailored to articulated objects and rely on artist-made rigs.

Example-based approaches [3, 23, 30, 37] build subspace based on a few user-selected example states. Increasing the number of examples to full simulation sequences gives rise to data-driven methods. A widely used approach in this context is to construct linear subspaces from simulation data using principal component analysis (PCA) [24]. Due to the limited expressiveness of linear subspaces, however, recent methods have resorted to neural networks that can directly capture the nonlinearity of the underlying physics.

Neural Physics. Inspired by physics-informed neural networks (PINN) [27], current state-of-the-art methods [6, 15, 34] leverage neural networks for subspace construction and simulation. These methods collect a large number of deformed states and train, e.g., deep autoencoders to obtain a latent manifold. A notable example is latent-space dynamics [8], which achieves physical awareness by collecting training data from a physical simulator. Shen et al. [33] augment latent-space dynamics with higher-order derivatives to accelerate simulation. Romero et al. [28] propose to learn nonlinear neural corrections for linear handle-based subspaces. While these methods have shown the promise of learning-based subspaces, they rely on large amounts of cu-

rated simulation data for training. This data collection is time-consuming, and learned subspaces are biased towards the training data. Additionally, existing methods train networks in a purely geometric way, leading to subspaces that are riddled with high-energy configurations. By contrast, our method learns nonlinear subspaces in a physically principled and fully self-supervised way, eliminating the need for curated simulation data. Our learned subspaces are interpretable and exhibit superior regularity compared to previous methods. Sharp et al. [32] have likewise explored the potential of unsupervised learning for subspace generation. Their method introduces a regularization term to encourage mode orthogonality. However, as we show in our analysis (Sec. 4.4), this regularizer cannot always prevent mode collapse, i.e., a decrease in effective dimensionality of the learned subspaces.

Neural simulation has attracted special attention for cloth simulation [6, 15, 21, 25, 36]. Bertiche et al. [4], Santesteban et al. [29], and Grigorev et al. [9] have investigated self-supervised learning techniques. Whereas these techniques are specialized for garments, our approach applies to a wider range of mechanical systems and materials.

3. Theory

We start by introducing our novel self-supervised learning approach for nonlinear physics-based subspaces in Section 3.1. Since our focus is on learning nonlinear modal subspaces, we briefly summarize linear modes and nonlinear extensions in Section 3.2. In Section 3.3, we present Neural Modes, i.e., learned nonlinear subspaces that are trained in a fully self-supervised way.

3.1. Self-Supervised Subspace Learning

We consider discrete elastic surfaces and solids represented as triangle meshes and tetrahedron meshes, respectively. Let $\mathbf{x} \in \mathcal{R}^{3n}$ and $\mathbf{X} \in \mathcal{R}^{3n}$ denote vectors of deformed and undeformed positions for the n vertices, respectively. Equilibrium states of these discrete systems can be characterized by constrained energy minimization principles in the form of

$$\begin{aligned} \mathbf{x}^*(\phi, \psi) = \arg \min_{\mathbf{x}} E_{\phi}(\mathbf{x}) \\ \text{s.t. } C_{\psi}(\mathbf{x}) = \mathbf{0} \end{aligned} \quad (1)$$

where $E_{\phi}(\mathbf{x})$ is the mechanical energy defined through intrinsic parameters ϕ including material properties and input shape. Furthermore, $C_{\psi}(\mathbf{x})$ is a vector-valued constraint function modeling the boundary conditions, which are defined through a set of extrinsic parameters ψ . The solution set $\mathbf{x}^*(\phi, \psi)$ constitutes a nonlinear subspace of \mathcal{R}^{3n} corresponding to physically-principled equilibrium configurations. Using conventional simulation, sampling from this manifold amounts to solving a constrained minimization

problem for each sample (ϕ_i, ψ_i) , which is computationally expensive. Our approach to this problem is to find a neural network $\mathbf{x}[\theta^*](\phi, \psi)$ that best emulates the equilibrium subspace,

$$\begin{aligned} \mathbf{x}[\theta^*](\phi, \psi) \simeq \arg \min_{\theta} E_{\phi}(\mathbf{x}[\theta](\phi, \psi)) \\ \text{s.t. } C_{\psi}(\mathbf{x}[\theta](\phi, \psi)) = \mathbf{0}. \end{aligned} \quad (2)$$

Instead of optimizing for the unknown equilibrium configuration, a single neural network forward pass instantly provides the solution \mathbf{x}^* . Unlike supervised approaches that aim to learn solution spaces based on simulation data, we find the network weights θ by directly minimizing the mechanical energy across parameter space. To this end, we define a physics-based loss function,

$$L(\theta) = \mathbb{E}_{\phi, \psi} [E_{\phi}(\mathbf{x}[\theta](\phi, \psi)) + \lambda \|C_{\psi}(\mathbf{x}[\theta](\phi, \psi))\|^2], \quad (3)$$

where we used a penalty function with stiffness coefficient λ to enforce constraints. Once trained, the optimal θ give rise to a nonlinear subspace that, instead of approximating simulation data in a geometric sense, reflects the physical principles of the underlying mechanical system. Moreover, this learning process is entirely self-supervised and therefore easy to automate. We focus on thin shells and volumetric solids as discrete mechanical systems and build our subspaces on nonlinear compliant modes [7]. However, our formulation generalizes to a broad range of mechanical models, materials, and subspace parameterizations. The only requirement is that the energy function and constraints are differentiable and weakly convex. In the following, we apply our formulation to nonlinear modal subspaces.

3.2. Modal Subspaces

Linear Modes. Equilibrium configurations of mechanical systems are minimizers of $E(\mathbf{x})$. To construct subspaces that contain physically plausible configurations, we seek directions that minimize the energy for given displacement magnitudes z , i.e.,

$$\min_{\mathbf{u}} E(\mathbf{X} + \mathbf{u}) \quad \text{s.t.} \quad \frac{1}{2}(\mathbf{u}^T \mathbf{u} - z) = 0. \quad (4)$$

For small displacements, we can replace $E(\mathbf{X} + \mathbf{u})$ with a quadratic approximation,

$$\min_{\mathbf{u}} \frac{1}{2} \mathbf{u}^T \mathbf{H} \mathbf{u} \quad \text{s.t.} \quad \frac{1}{2}(\mathbf{u}^T \mathbf{u} - z) = 0, \quad (5)$$

where \mathbf{H} denotes the Hessian of the elastic energy and we have used the fact that both the energy and its gradient vanish at the origin. Considering the Lagrangian of this problem,

$$\mathcal{L} = \frac{1}{2} \mathbf{u}^T \mathbf{H} \mathbf{u} - \frac{\mu}{2} (\mathbf{u}^T \mathbf{u} - z), \quad (6)$$

where μ is a Lagrange multiplier, it becomes evident from the first-order optimality conditions,

$$\begin{aligned} \mathbf{H}\mathbf{u} - \mu\mathbf{u} &= \mathbf{0}, \\ \mathbf{u}^T\mathbf{u} - z &= 0, \end{aligned}$$

that solutions to the constrained optimization problem (5) are eigenvectors \mathbf{e}_i of the Hessian, i.e., $\mathbf{H}\mathbf{e}_i = \mu_i\mathbf{e}_i$. In particular, each eigenvector \mathbf{e}_{i+1} minimizes the system's energy in the $3n - i$ -dimensional subspace orthogonal to all previous eigenvectors \mathbf{e}_i . This observation motivates the choice of constructing linear subspaces from the m eigenvectors—the linear modes—corresponding to the smallest eigenvalues. For small deformations, where the quadratic energy approximation is accurate, linear modes form an energetically optimal subspace in the sense of the mechanical work required to achieve a displacement of given magnitude. For larger deformations, however, linear modes suffer from distortion artifacts that manifest as a steep increase in energy and physically-implausible configurations; see, e.g., Figure 5.

Nonlinear Compliant Modes. To overcome the limitations of linear modes, Dünser et al. [7] define nonlinear compliant modes as

$$\mathbf{n}_i(z) = \arg \min_{\mathbf{u}} E(\mathbf{X} + \mathbf{u}) \quad \text{s.t.} \quad \mathbf{e}_i^T \mathbf{u} = z. \quad (7)$$

Letting $\mathbf{l}_i = z\mathbf{e}_i$ denote the displacement along the linear mode, we can alternatively rewrite Equation (7) as

$$\mathbf{n}_i(z) = \mathbf{l}_i + \arg \min_{\mathbf{y}} E(\mathbf{X} + \mathbf{l}_i + \mathbf{y}) \quad \text{s.t.} \quad \mathbf{l}_i^T \mathbf{y} = 0 \quad (8)$$

where the correction \mathbf{y} is obtained by minimizing the system's energy in the space orthogonal to the linear mode. As demonstrated by Dünser et al., nonlinear compliant modes can achieve large and energetically-optimal deformations in the sense of (7). Nevertheless, the original formulation is restricted to individual modes and does not directly extend to multi-dimensional modal spaces. Furthermore, nonlinear compliant modes are very expensive to compute as evaluating the trajectory at a given parameter location requires the solution of a nonlinear constraint optimization problem. While these limitations seem to rule out nonlinear compliant modes for interactive subspace simulation, we show in the following that our self-supervised learning approach proposed in Section 3.1 can solve both problems in an elegant and efficient way.

3.3. Neural Modes

Our goal is to build learned subspaces on the basis of nonlinear compliant modes. To this end, we first extend the concept by Dünser et al. to m -dimensional modal spaces,

$$\mathbf{n}(\mathbf{z}) = \mathbf{l} + \arg \min_{\mathbf{y}} E(\mathbf{X} + \mathbf{l} + \mathbf{y}) \quad \text{s.t.} \quad \mathbf{l}^T \mathbf{y} = 0, \quad (9)$$

where the linear displacement $\mathbf{l} = \sum_i^m z_i \mathbf{e}_i$ is now a combination of linear modes weighted by the modal coordinates $\mathbf{z} = (z_1, \dots, z_m)$. We then train a neural network to learn the solution to the above optimization problem for any choice of modal coordinates, i.e.,

$$\mathbf{n}[\theta^*](\mathbf{z}) = \mathbf{l} + \mathbf{y}[\theta^*](\mathbf{z}) \quad (10)$$

$$\theta^* = \arg \min_{\theta} E(\mathbf{X} + \mathbf{l} + \mathbf{y}[\theta](\mathbf{z})) \quad (11)$$

$$\text{s.t.} \quad \mathbf{l}^T \mathbf{y}[\theta](\mathbf{z}) = 0.$$

We refer to the resulting nonlinear modal subspace $\mathbf{n}[\theta](\mathbf{z})$ as *Neural Modes*. To train Neural Modes, we use the above expression in (3) and obtain the loss function for self-supervised learning as

$$L(\theta) = \mathbb{E}_{\mathbf{z}} [E(\mathbf{X} + \mathbf{l} + \mathbf{y}[\theta](\mathbf{z})) + \lambda(\mathbf{l}^T \mathbf{y}[\theta](\mathbf{z}))^2] \quad (12)$$

During training, we draw the modal coordinates \mathbf{z} from a uniform distribution with predefined domain. We input \mathbf{z} to the network, and retrieve vertex positions $\mathbf{x} = \mathbf{X} + \mathbf{l} + \mathbf{y}$ from the network output \mathbf{y} . While experimenting with this loss function, we observed that the trained network produces nonzero corrections at the origin, i.e., $\mathbf{y}[\theta](\mathbf{0}) \neq \mathbf{0}$. Since Neural Modes should coincide with linear modes for small displacements, we add a term penalizing corrections at the origin, leading to the modified loss function

$$L_{\text{final}}(\theta) = L(\theta) + \eta \|\mathbf{y}[\theta](\mathbf{0})\|^2. \quad (13)$$

Hyperparameters. The formulation described so far computes Neural Modes for given material and shape parameters. Nevertheless, these parameters significantly affect the resulting nonlinear modes. Fortunately, we can directly learn the effect of these changes by simply adding corresponding parameters to modal coordinates and sampling during training.

3.4. Subspace Dynamics

The goal of subspace simulation is to create compelling motion with physically plausible deformations at a fraction of the cost associated with full-space simulation. Starting from a variational re-formulation of the implicit Euler update equations in full space [23], we have

$$\mathbf{u}_{n+1} = \arg \min_{\mathbf{u}} \frac{1}{2h^2} \|\mathbf{u} - 2\mathbf{u}_n + \mathbf{u}_{n-1}\|_{\mathbf{M}}^2 + E(\mathbf{u}), \quad (14)$$

where h is the step size and $\|\cdot\|_{\mathbf{M}}$ is the metric induced by the mass matrix \mathbf{M} . Following the formulation of Fulton et al. [8], we replace the full-space displacement \mathbf{u} with our neural subspace $\mathbf{n}(\mathbf{z})$ as

$$\mathbf{z}_{n+1} = \arg \min_{\mathbf{z}} \frac{1}{2h^2} \|\mathbf{n}(\mathbf{z}) - 2\mathbf{u}_n + \mathbf{u}_{n-1}\|_{\mathbf{M}}^2 + E(\mathbf{n}(\mathbf{z})), \quad (15)$$

We solve the minimization problem for end-of-step modal coordinates using L-BFGS [22].

4. Results

We implemented all our experiments and network training in PyTorch, taking advantage of automatic differentiation. We use conventional multi-layer perceptrons (MLP) with Gaussian Error Linear Units (GELU) to implement our Neural Modes. We also experimented with Residual Networks (ResNet) but found no improvements in practice. We train our networks on a single RTX 3090 GPU using L-BFGS with line search as the optimizer. We use discrete shells [10] for deformable surfaces and linear tetrahedron elements with a Saint Venant-Kirchhoff material for solids [17]. We summarize our experiment setup in Table 1.

4.1. Performance Comparisons

We quantitatively evaluate our approach on a square shell with a flat rest state. We train a 5-layer MLP to learn a 3-dimensional modal subspace using modal coordinates \mathbf{z} from the domain $[-0.625, 0.625]^3$. We compare two variants of training for our MLP using stochastic and regular (grid-based) sampling, respectively. For stochastic sampling, we choose different modal coordinates \mathbf{z} in each iteration by uniformly sampling the domain $[-0.625, 0.625]^3$. For regular sampling, we sample \mathbf{z} from a 9^3 uniform grid.

We compare our method with Latent-space Dynamics [8], which combines a deep autoencoder (AE) with principal component analysis (PCA). We abbreviate their method as PCA+AE. Our MLP essentially corresponds to the decoder part of the autoencoder used in PCA+AE. For a fair comparison, we use the same MLP as PCA+AE’s decoder and make a symmetric copy for the encoder.

To train PCA+AE, we collect a dataset by regularly sampling \mathbf{z} and solving (9) for each sample. The training, validation, and test sets have resolutions of 9^3 , 7^3 , and 11^3 samples, respectively. We also collect a second training set of the same size by randomly sampling modal coordinates. We train separate instances of PCA+AE on these two training sets (grid, random), and train our MLP using both stochastic and regular sampling. We evaluate and compare all methods on the same test and validation sets.

Elastic Energy. The evaluation results are shown in Table 2. Since nonlinear modes are constrained minimizers of (9), we use average elastic energy as a metric for measuring the accuracy of a subspace. It can be seen from Table 2 that our method has lower average energy than PCA+AE and an order of magnitude smaller errors compared to the ground truth solution. We hypothesize that the lower accuracy of PCA+AE is due to its geometry-based loss function and the associated likelihood of overfitting to the training set. We test this hypothesis in two ways.

First, we train our network again using the same architecture but use the loss function from PCA+AE, which measures the geometric distance to the training data in the L2 sense. We then train on the same regular sampling set, i.e., using the same data as before. As can be seen from Table 2, the resulting network (*L2 supervised*) exhibits significantly higher energy than our Neural Modes. Second, we plot learning curves for all networks in Figure 2. We can see that, while PCA+AE and L2 Supervised variants reduce training L2 loss steadily, their test L2 loss, energy, and stress all increase during training. In contrast, our self-supervised Neural Modes do not suffer from overfitting, steadily reducing the energy and stress of the test set. Finally, we experimented with early stopping, which is a common strategy to mitigate overfitting. We use both L2 loss and elastic energy as stopping criteria and observed somewhat better results for the PCA+AE variants when using energy (see Table 3). Nevertheless, the energy for PCA+AE is still significantly higher than for Neural Modes.

Stress & Nodal Force. We further evaluate the learned subspace by analyzing internal stresses and nodal forces. We compute averaged stress weighted over element areas and maximum element stress in Table 4. We achieve smaller stress and significantly lower errors compared with supervised methods. In addition to stress, we also observe one order of magnitude smaller errors in nodal forces (see Table A1). These results support our hypothesis that the geometric L2 loss is unable to capture the physical principles that govern the manifold of equilibrium configurations. It is furthermore prone to overfitting and can be expected to generalize poorly in practice. By directly minimizing energy during training, our self-supervised method shows much better performance while avoiding overfitting problems.

Deformable Solids. We additionally evaluate the quality of Neural Modes for deformable solids, using the example shown in Figure 5. We stochastically train an MLP to learn a 5-dimensional subspace. Similar to the previous example, we collect 1300 equilibrium states by uniformly sampling modal space and compute ground truth solutions for each sample. The first 900 samples are used for training, the subsequent 100 samples are used for validation, and the last 300 samples are reserved for the test set. We stop training after 2050 epochs as the L2 loss evaluated on the validation set continually increases for the following 500 epochs. The statistics are shown in Table 5. To ensure our evaluation result is not sampled by chance, we break down the test set into subintervals and compare Neural Modes and PCA+AE on each subinterval in addition to the complete set. It can be seen that Neural Modes outperform PCA+AE both for every interval and for the complete test set. This result confirms our observation for deformable surfaces, indicating that Neural Modes consistently outperform PCA+AE by significant margins.

Table 1. Overview of our example setups.

Model	Energy	# of cells	Full dim	Reduced dim	Neural dim	Auxiliary dim	MLP	λ	η	Timestep
Sheet	Discrete Shell	100	363	3	3	0	5x64	1e8	1e7	40 ms
Sheet (aspect ratio)	Discrete Shell	100	363	4	4	1	8x64	1e8	1e7	–
Giraffe	Tetrahedron+Penalty	5424	5094	3	3	0	10x64	1e8	1e7	36 ms
Armadillo	Tetrahedron	26211	21537	11	5	0	7x256	1e8	1e7	35 ms
Bunny	Tetrahedron+Penalty	7585	6360	8	8	0	8x64	1e6	1e3	28 ms

Table 2. Performance comparison on the square sheet example. We compute the elastic energy E and its error ΔE with respect to ground truth on both validation and test sets. The subscript denotes the reduction operator for the batch dimension.

Method	Validation				Test			
	ΔE_{avg}	ΔE_{max}	ΔE_{std}	E_{avg}	ΔE_{avg}	ΔE_{max}	ΔE_{std}	E_{avg}
PCA+AE	1814	22225	2005	18144	2259	27991	2646	24007
PCA+AE (grid)	2188	36168	3626	18088	2649	72834	4664	23217
L2 Supervised	1630	10231	1148	18210	2074	13750	1697	24325
Neural Modes (stoch.)	376	1197	170	17032	386	1154	183	22891
Neural Modes (grid)	199	754	118	16831	245	3062	215	22730

Table 3. We compare early stopping criteria using elastic energy and geometry-based L2 loss as metrics.

Method	Validation		Test	
	L2	Energy	L2	Energy
PCA+AE	18144	17762	24007	23484
PCA+AE (grid)	18088	18042	23217	23051
L2 Supervised	18210	17690	24325	23682

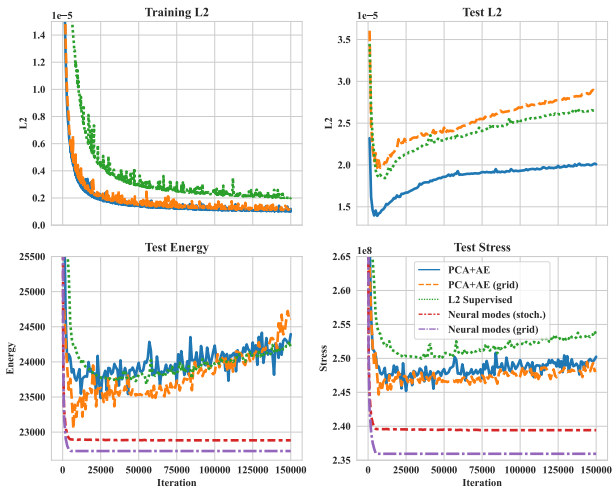


Figure 2. Supervised learning leads to overfitting for all loss functions, i.e., geometry-based L2 distance, average energy E_{avg} , and average stress $|S|_{avg}$. In contrast, our Neural Modes do not suffer from overfitting.

4.2. Latent Space Structure

The previous experiments indicate that our self-supervised learning approach yields subspaces with lower mechanical energy, stress, and nodal force than PCA+AE. Consequently, Neural Modes exhibit substantially lower errors with respect to ground truth values obtained by solving the

underlying constrained minimization problem. We further analyze the structure of the latent spaces with a focus on smoothness, interpretability, and regularity. Specifically, we investigate whether the map from latent space to full space is (1) *smooth* in the sense that smooth motion in latent space leads to smooth interpolated trajectories in geometry space, (2) *regular* in the sense that interpolated geometry is semantically bounded by terminal states, and (3) *interpretable* in the sense that individual latent space dimensions have a clear and distinguishable effect on the full space geometry.

To analyze interpretability, we consider the same example as before and explore the behavior of Neural Modes and PCA+AE when walking along individual latent space coordinate directions. As can be seen from Fig. 3, Neural Modes faithfully reproduce the individual nonlinear modes and preserve symmetry. We also observe physically-plausible deformations when combining multiple latent space dimensions (see Figure A1). The picture looks quite different for PCA+AE. Although each latent space dimension features meaningful shapes, they are interspersed with distorted or crumpled shapes that are physically implausible. We further note that Neural Modes preserve both the origin (i.e., $\mathbf{n}(0) = \mathbf{X}$) and symmetry of the ground truth modes, whereas PCA+AE does neither.

A basic test for regularity is interpolation quality [1]. The median of two latent vectors is expected to yield a semantically meaningful median result and no abrupt changes from terminal states. We visualize interpolation quality in Figures 3 and A1 using evenly spaced sample points. It can be seen that PCA+AE exhibits poor regularity. By contrast, Neural Modes generate semantically meaningful intermediate states, defining a smooth trajectory in geometry space. Unlike PCA+AE, this semantic interpolation property makes Neural Modes attractive for applications such as keyframe animations (see Figure 6).

Table 4. Comparison of internal stress on the square sheet example. We compare average per-element stresses weighted by areas and maximum per-element stress. We compute the Frobenius norm of stress $|S|$ and its error $\Delta|S|$ with respect to ground truth values for each method. The subscript denotes the reduction operator for the batch dimension.

Method ($\times 10^6$)	Weighted by elements			Max over elements				
	$\Delta S _{avg}$	$\Delta S _{max}$	$\Delta S _{std}$	$ S _{avg}$	$\Delta S _{avg}$	$\Delta S _{max}$	$\Delta S _{std}$	$ S _{avg}$
PCA+AE	25	138	17	251	67	431	49	520
PCA+AE (grid)	25	305	29	245	64	784	68	505
L2 Supervised	25	100	14	255	87	350	57	550
Neural Modes (stoch.)	9	36	7	240	22	155	21	486
Neural Modes (grid)	6	24	5	236	16	124	16	479




Figure 3. Visualization of latent space dimensions for Neural Modes and PCA+AE.

Table 5. We evaluate the performance of our method using the model’s average elastic energy evaluated on the test set.

Method ($\times 10^4$)	0~20%	20~40%	40~60%	60~80%	80~100%	all
Ground truth	12.3	9.7	12.5	10.9	10.8	11.3
PCA+AE	106.7	4329.9	216.5	179.5	104.8	987.5
Neural Modes	57.1	58.0	58.9	52.3	51.9	55.8

From these experiments, we conclude that, compared to supervised learning from simulation data (PCA+AE), self-supervised learning of Neural Modes yields well-behaved subspaces with much better smoothness, interpretability, regularity, and structure preservation.

4.3. Material and Shape Parameters

Neural modes are naturally parameterized by modal coordinates, but they generalize to other parameters such as shape and material properties. Figure A2 showcases a simple example where we add aspect ratio as an auxiliary parameter for a rectangular sheet. Starting from a sheet with edge length ratio 2:1, the folding direction smoothly changes from the sheet’s short axis to its diagonal as we decrease the ratio towards 1:1.

4.4. Subspace Efficiency

We additionally compare subspace efficiency with Sharp et al. [32]. Like our method, Sharp et al. learn nonlinear subspaces in an unsupervised way using the mechanical energy $E(\mathbf{x})$ as loss function. Using only energy, however, the optimal network would map all latent variables to the undeformed state $\mathbf{x} = \mathbf{X}$ since $E(\mathbf{X}) = 0$. To




Figure 4. Correlation matrices of latent directions for Armadillo.

avoid this *subspace collapse*, Sharp et al. introduce a regularization term that encourages the network to produce an isometric map from latent space to full space. Although this regularizer avoids subspace collapse, it cannot prevent *mode collapse*, i.e., a loss of effective dimensionality. To show this, we use the method by Sharp et al. to learn a 5D-subspace for the Armadillo model. We compute latent directions $\mathbf{E} = [\mathbf{e}_1 | \dots | \mathbf{e}_5]$ by evaluating the network’s gradient at the origin and compute the correlation matrix $\mathbf{E}^T \mathbf{E}$ (Fig. 4). The eigenvalues of the correlation matrix are (2.3, 1.6, 1.1, 0.0, 0.0), revealing that the subspace is effectively 3-dimensional. Applying the same test to Neural Modes yields first-order orthogonal latent directions and a quasi-diagonal correlation matrix with all eigenvalues close to 1.0 (see Figure 4).

Mode collapse limits the efficiency of subspaces. As shown in Figure A3, while our Neural Modes capture both arm and leg motions, the method by Sharp et al. only learns leg motions with the aforementioned redundancy.

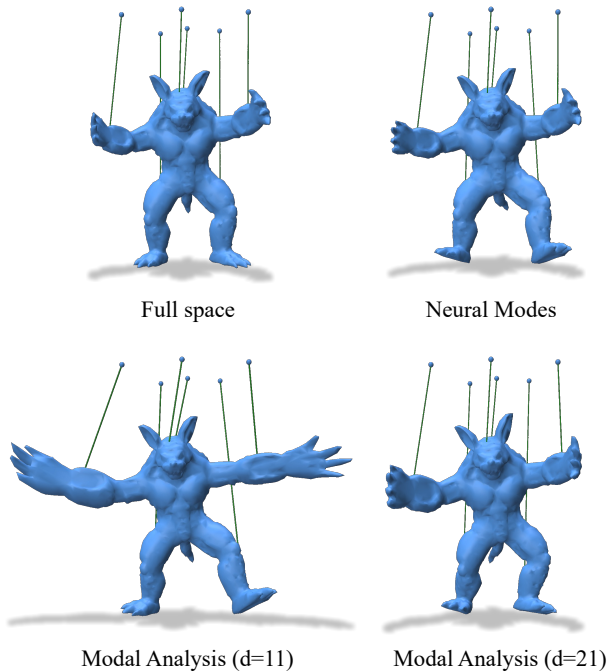


Figure 5. Comparison of 5d Neural Modes, modal analysis, and full-space simulation on an Armadillo puppet.

4.5. Applications

Complementing our quantitative analysis, we additionally evaluate our method on a set of application examples that illustrate potential use cases in modeling and animation.

Subspace Dynamics

As our first example, we consider a subspace simulation problem using the deformable Armadillo puppet shown in Figure 5. The model is driven by six rods whose rest lengths we animate using sinusoidal motion. We train an 11-dimensional subspace, with 5 DoFs corresponding to Neural Modes, and 6 DoFs for rigid transformations. It is evident from Figure 5—but best observed in the accompanying video—that linear modes lead to large distortion artifacts in the arms and legs. This issue is only partly mitigated when adding ten extra basis vectors. Neural Modes, in contrast, produce nonlinear rotations as expected from the ground truth simulation. Moreover, our method achieves real-time performance in ~ 35 ms per time step, whereas the full simulation is two orders of magnitude slower.

We consider two additional examples to illustrate that our method generalizes to arbitrary shapes: a giraffe and a bunny model simulated using Neural Modes subspaces with 3 and 8 dimensions, respectively. Both models run at real-time rates and produce smooth, visually pleasing motion. Please see the accompanying video for animations.

Keyframing. Keyframe animation is a basic tool for generating motion from a sparse set of input poses. Given a set of keyframe poses through modal coordinates, we can




Figure 6. Keyframing of nonlinear sitting-down and standing-up motions using Neural Modes. Asterisks denote keyframes.

generate in-between trajectories with physically plausible deformations by means of interpolation. Figure 6 illustrates this application on the giraffe model. While linear subspaces typically require nonlinear interpolation schemes to achieve acceptable results, our nonlinear subspaces produce physically plausible deformations even for simple linear interpolation of modal coordinates.

5. Conclusion

We presented a novel self-supervised method for learning nonlinear physics-based subspaces for real-time simulation. We demonstrated the potential of our formulation by learning Neural Modes, i.e., physics-based subspaces built on a nonlinear extension of linear eigenmodes. We quantitatively evaluated our method on a set of examples that include thin shells and volumetric solids. Furthermore, we performed an extensive ablation study, revealing that directly minimizing elastic energy leads to significantly better physical accuracy than using simpler geometry-based loss functions. We also showed that self-supervised learning resolves overfitting problems that supervised approaches often struggle with. By visualizing latent spaces, we demonstrated that Neural Modes are interpretable, free from singularities, and well-behaved for arbitrary modal combinations. We further showed that, unlike existing alternatives, our formulation avoids the problem of mode collapse. Finally, we showcase applications of our approach to subspace simulation and keyframing.

Limitations & Future Work.

Our learning method is restricted to differentiable energy functions and constraints. Additionally, we did not investigate simulation acceleration algorithms and relied on L-BFGS for optimization. Future explorations could compute the exact Hessian in the reduced parameter space, which would open the door to more efficient minimization algorithms such as the Newton-Raphson method. While we have focused on subspaces induced by nonlinear modes, our approach generalizes to many other parameterizations. An exciting avenue for future work in this context would be to extend our method to inverse design problems, in, e.g., architecture or engineering design.

References

- [1] Sharjeel Ali and Oliver van Kaick. Evaluation of latent space learning with procedurally-generated datasets of shapes. *Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [2] Jernej Barbič and Doug L. James. Real-time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph.*, 24(3):982–990, 2005. 2
- [3] Jernej Barbič, Marco da Silva, and Jovan Popović. Deformable object animation using reduced optimal control. *ACM Trans. Graph.*, 28(3), 2009. 2
- [4] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Neural cloth simulation. *ACM Trans. Graph.*, 41(6), 2022. 3
- [5] Min Gyu Choi and Hyeong-Seok Ko. Modal warping: Real-time simulation of large rotational deformation and manipulation. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):91–101, 2005. 2
- [6] Y D. Li, Min Tang, Yun Yang, Zi Huang, R F. Tong, Shuang Cai Yang, Yao Li, and Dinesh Manocha. N-cloth: Predicting 3d cloth deformation with mesh-based networks. In *Computer Graphics Forum*, pages 547–558. Wiley Online Library, 2022. 2, 3
- [7] Simon Duenser, Bernhard Thomaszewski, Roi Poranne, and Stelian Coros. Nonlinear compliant modes for large-deformation analysis of flexible structures. *ACM Trans. Graph.*, 42(2), 2022. 2, 3, 4
- [8] Lawson Fulton, Vismay Modi, David Duvinaud, David I. W. Levin, and Alec Jacobson. Latent-space dynamics for reduced deformable simulation. *Computer Graphics Forum*, 2019. 1, 2, 4, 5
- [9] Artur Grigorev, Bernhard Thomaszewski, Michael J Black, and Otmar Hilliges. Hood: Hierarchical graphs for generalized modelling of clothing dynamics. *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [10] Eitan Grinspun, Anil N Hirani, Mathieu Desbrun, and Peter Schröder. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 62–67. Citeseer, 2003. 5
- [11] Fabian Hahn, Sebastian Martin, Bernhard Thomaszewski, Robert Sumner, Stelian Coros, and Markus Gross. Rig-space physics. *ACM Trans. Graph.*, 31(4), 2012. 2
- [12] Fabian Hahn, Bernhard Thomaszewski, Stelian Coros, Robert W. Sumner, and Markus Gross. Efficient simulation of secondary motion in rig-space. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 165–171, New York, NY, USA, 2013. Association for Computing Machinery. 2
- [13] Fabian Hahn, Bernhard Thomaszewski, Stelian Coros, Robert W. Sumner, Forrester Cole, Mark Meyer, Tony DeRose, and Markus Gross. Subspace clothing simulation using adaptive bases. *ACM Trans. Graph.*, 33(4), 2014. 1
- [14] Klaus Hildebrandt, Christian Schulz, Christoph Von Tycowicz, and Konrad Polthier. Interactive surface modeling using modal analysis. *ACM Trans. Graph.*, 30(5), 2011. 2
- [15] Daniel Holden, Bang Chi Duong, Sayantan Datta, and Derek Nowrouzezahrai. Subspace neural physics: Fast data-driven interactive simulation. In *Proceedings of the 18th Annual ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, New York, NY, USA, 2019. Association for Computing Machinery. 2, 3
- [16] Jin Huang, Yiyang Tong, Kun Zhou, Hujun Bao, and Mathieu Desbrun. Interactive shape interpolation through controllable dynamic deformation. *IEEE Transactions on Visualization and Computer Graphics*, 17(7):983–992, 2011. 2
- [17] Theodore Kim and David Eberle. Dynamic deformables: Implementation and production practicalities (now with code!). In *ACM SIGGRAPH 2022 Courses*, New York, NY, USA, 2022. Association for Computing Machinery. 5
- [18] M.A. Kramer. Autoassociative neural networks. *Computers Chemical Engineering*, 16(4):313–328, 1992. Neural network applications in chemical engineering. 1
- [19] Guiqing Li, Yaobin Ouyang, Guodong Wei, Zhibang Zhang, and Aihua Mao. Enhanced rig-space simulation. *Computer Animation and Virtual Worlds*, 27(3-4):250–261, 2016. 2
- [20] Xuan Li, Yu Fang, Lei Lan, Huamin Wang, Yin Yang, Minchen Li, and Chenfanfu Jiang. Subspace-preconditioned gpu projective dynamics with contact for cloth simulation. In *SIGGRAPH Asia 2023 Conference Papers*, New York, NY, USA, 2023. Association for Computing Machinery. 1
- [21] Emmanuel Ian Libao, Myeongjin Lee, Sumin Kim, and Sung-Hee Lee. Meshgraphnetrp: Improving generalization of gnn-based cloth simulation. In *Proceedings of the 16th ACM SIGGRAPH Conference on Motion, Interaction and Games*, New York, NY, USA, 2023. Association for Computing Machinery. 3
- [22] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989. 5
- [23] Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. Example-based elastic materials. In *ACM SIGGRAPH 2011 Papers*, New York, NY, USA, 2011. Association for Computing Machinery. 2, 4
- [24] Ahmed K Noor and Jeanne M Peters. Reduced basis technique for nonlinear analysis of structures. *Aiaa journal*, 18(4):455–462, 1980. 2
- [25] Young Jin Oh, Tae Min Lee, and In-Kwon Lee. Hierarchical cloth simulation using deep neural networks. In *Proceedings of Computer Graphics International 2018*, page 139–146, New York, NY, USA, 2018. Association for Computing Machinery. 3
- [26] Amir Rabbani and Paul Kry. Physik: Physically plausible and intuitive keyframing. In *Proceedings of the 42nd Graphics Interface Conference*, page 153–161, Waterloo, CAN, 2016. Canadian Human-Computer Communications Society. 2
- [27] Maziar Raissi, Paris Perdikaris, and George E 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. 2
- [28] Cristian Romero, Dan Casas, Jesús Pérez, and Miguel Otaduy. Learning contact corrections for handle-based subspace dynamics. *ACM Trans. Graph.*, 40(4), 2021. 2
- [29] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Snug: Self-supervised neural dynamic garments. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8140–8150, 2022. [3](#)
- [30] Christian Schumacher, Bernhard Thomaszewski, Stelian Coros, Sebastian Martin, Robert Sumner, and Markus Gross. Efficient simulation of example-based materials. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 1–8, Goslar, DEU, 2012. Eurographics Association. [2](#)
- [31] Ahmed A Shabana. *Theory of vibration*. Springer, 1991. [2](#)
- [32] Nicholas Sharp, Cristian Romero, Alec Jacobson, Etienne Vouga, Paul G. Kry, David I. W. Levin, and Justin Solomon. Data-free learning of reduced-order kinematics, 2023. [2](#), [3](#), [7](#)
- [33] Siyuan Shen, Yin Yang, Tianjia Shao, He Wang, Chenfanfu Jiang, Lei Lan, and Kun Zhou. High-order differentiable autoencoder for nonlinear model reduction. *ACM Trans. Graph.*, 40(4), 2021. [2](#)
- [34] Edgar Tretschk, Ayush Tewari, Michael Zollhöfer, Vladislav Golyanik, and Christian Theobalt. Demea: Deep mesh autoencoders for non-rigidly deforming objects. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 601–617. Springer, 2020. [2](#)
- [35] Soumianarayanan Vijayaraghavan, Ling Wu, Ludovic Noels, SPA Bordas, Sundar Natarajan, and Lars AA Beex. A data-driven reduced-order surrogate model for entire elastoplastic simulations applied to representative volume elements. *Scientific Reports*, 13(1):12781, 2023. [1](#)
- [36] Tuanfeng Y. Wang, Tianjia Shao, Kai Fu, and Niloy J. Mitra. Learning an intrinsic garment space for interactive authoring of garment animation. *ACM Trans. Graph.*, 38(6), 2019. [3](#)
- [37] Hongyi Xu and Jernej Barbič. Pose-space subspace dynamics. *ACM Trans. Graph.*, 35(4), 2016. [2](#)
- [38] Yin Yang, Dingzeyu Li, Weiwei Xu, Yuan Tian, and Changxi Zheng. Expediting precomputation for reduced deformable simulation. *ACM Trans. Graph.*, 34(6), 2015. [2](#)