

Force-Dual Modes: Subspace Design from Stochastic Forces

OTMAN BENCHEKROUN, University of Toronto, Canada

EITAN GRINSPUN, University of Toronto, Canada

MAURIZIO CHIARAMONTE, Meta Reality Labs, USA

PHILIP ALLEN ETTER, Meta Reality Labs, USA




Fig. 1. We construct simulation subspaces by propagating user-defined or system-implied force distributions through the physical model, yielding Force-Dual modes (FDM), a basis tailored to anticipated actuation patterns. Compared to Linear Modal Analysis (LMA), which assumes spatially uncorrelated forces (*middle*) and produces locking artifacts, our force-aware construction produces simulations that more accurately reflect localized, force-aware and heterogeneous deformation. Despite the dragon having a large mesh of 25,613 vertices and 103,307 tetrahedra, an 11-degree of freedom subspace simulation achieves real-time performance at 448 FPS, while still accurately capturing the dragon's actuated deformation. (*right*).

Designing subspaces for Reduced Order Modeling (ROM) is crucial for accelerating finite element simulations in graphics and engineering. Unfortunately, it's not always clear which subspace is optimal for arbitrary dynamic simulation. We propose to construct simulation subspaces from force distributions, allowing us to tailor such subspaces to common scene interactions involving constraint penalties, handles-based control, contact and musculoskeletal actuation. To achieve this we adopt a statistical perspective on Reduced Order Modelling, which allows us to push such user-designed force distributions through a linearized simulation to obtain a dual distribution on displacements. To construct our subspace, we then fit a low-rank Gaussian model to this displacement distribution, which we show generalizes Linear Modal Analysis subspaces for uncorrelated unit variance force distributions, as well as Green's Function subspaces for low rank force distributions. We show our framework allows for the construction of subspaces that are optimal both with respect to physical material properties, as well as arbitrary force distributions as observed in handle-based, contact, and musculoskeletal scene interactions.

CCS Concepts: • **Computing methodologies** → **Physical simulation**; **Motion processing**.

Additional Key Words and Phrases: Subspace Simulation, Reduced Order Modeling

Authors' Contact Information: Otman Benckekroun University of Toronto, Canada, otman.benckekroun@mail.utoronto.ca; Eitan Grinspun University of Toronto, Canada, eitan@cs.toronto.edu; Maurizio Chiaramonte Meta Reality Labs, USA, mchiaram@meta.com; Philip Allen Etter Meta Reality Labs, USA, paetter@proton.com.

Please use nonacm option or ACM Engage class to enable CC licenses
This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2025 Copyright held by the owner/author(s).
ACM 1557-7368/2025/12-ART
<https://doi.org/10.1145/3763310>

ACM Reference Format:

Otman Benckekroun, Eitan Grinspun, Maurizio Chiaramonte, and Philip Allen Etter . 2025. Force-Dual Modes: Subspace Design from Stochastic Forces. *ACM Trans. Graph.* 44, 6 (December 2025), 14 pages. <https://doi.org/10.1145/3763310>

1 Introduction

Whether it's stretching before your early morning workout routine or tucking yourself into bed for a good night's sleep, humans and animals alike are experts at interacting with the complex, deformable physical world. While we've become extremely good at mimicking this rich physical detail in offline simulations for visual effects, our ability to interact with physically-based virtual environments in real time still remains primitive. This shortcoming stems from the fact that detailed soft body simulation entails solving an extremely high-dimensional optimization problem—where the solution, i.e., the displacement at each point in space, must be computed at every simulation timestep.

A popular approach to make this computation tractable for real-time interaction is Reduced Order Modeling (ROM), where a low-dimensional subspace for the displacements is precomputed, reducing the size of the optimization problem at runtime. However, it is notoriously difficult to construct subspaces that are well-suited to the broad and unpredictable range of interactions encountered during simulation. Each subspace typically captures only a narrow class of motions, and when a simulation strays outside this range, the resulting behavior can be physically implausible. One common artifact is *spooky action at a distance*, where unrelated parts of the




Fig. 2. Interacting with this 8-dimensional Linear Modal Analysis subspace simulation of an elephant by moving its left arm upwards creates *spooky action at a distance*, inducing deformation of the far away head and trunk. By revealing how to provide a better statistical prior to the subspace construction, we can avoid this artifact.

mesh appear to move in a correlated manner due to shared basis vectors (See Fig. 2).

Part of the difficulty is that it is challenging to even characterize what kinds of motions a given subspace will be good at reproducing. Existing approaches often rely on manually selected or heuristically designed training motions, but there is no clear framework for systematically guiding subspace construction to match anticipated interactions. We propose to address this by explicitly modeling the uncertainty inherent in interactive simulation through a statistical perspective, where subspace construction becomes equivalent to low-rank statistical model fitting of physical response data.

But how do we define the prior distribution from which this model fitting is performed? This is a central, unsolved question. While displacements are the primary quantities of interest in simulation, directly specifying distributions over them is unintuitive, as it requires prior knowledge of both material behavior and simulation dynamics. Instead, we advocate for modeling distributions on the forces, which are typically easier to reason about, and then using a linearized dual relationship between forces and displacements to implicitly define a distribution over the latter. This shift enables a principled and flexible method for generating simulation subspaces.

We provide a mathematical framework for treating force distributions as first-class entities in subspace construction. This framework is general and expressive, encompassing prior approaches such as Linear Modal Analysis and Green’s functions, which can now be reinterpreted as arising from implicit statistical assumptions about the underlying force distributions, as shown in Fig. 3. These methods didn’t lack a statistical prior—they simply assumed one without making it explicit, often leading to brittle behavior when the actual forces diverge from those assumptions.

Our framework is the first to support the construction of subspaces that are not only tailored to the physical properties of the simulated material, but also customizable to general force distributions. These distributions can be induced from various sources, such as example trajectories, actuation models, or interactive constraints, allowing users to easily design simulation behaviors that match their specific application needs. As a result, we can robustly simulate complex heterogeneous systems involving diverse force interactions, including contact, spring actuation, moving handle constraints and musculoskeletal actuation.




Fig. 3. Traditional Linear Modal Analysis assumes the bridge to be subject to completely random uncorrelated Gaussian forces. We generalize Linear Modal Analysis by considering arbitrary force distributions, enabling us to construct subspaces better suited to a large variety of force interactions.

2 Related Work

The introduction of Linear Modal Analysis (LMA) to the computer graphics community by Pentland and Williams [1989] brought with it a promise of fully interactive physics simulation through the use of linear subspaces. Unfortunately, this promise did not materialize because the resulting simulation’s quality was noticed to degrade significantly as it underwent large deformation, specifically rotational motion.

2.1 Subspaces for Large Deformation

Much work has been done on attempting to augment Linear Modal Analysis subspaces in order to allow them to account for large rotational deformation.

Modal derivatives [Barbič and James 2005] view the linearity of the subspace itself as the issue, and augment the linear subspace with *derivative modes*, effectively making the subspace quadratic. Alternatively, one can decompose the domain into various substructures, explicitly tracking the rotational motion for each component while also allowing every cluster to have its own separate linear subspace [Barbič and James 2005; Kim and James 2011]. These unfortunately suffer from visual artifacts at substructure interfaces, as shown in Fig. 4. Choi and Ko [2005] instead explicitly account for rotational motion by rotating individual per-vertex components of the subspace itself, effectively doing domain decomposition at the finest per-vertex level. Rotation strain coordinates [Pan et al. 2015; ?] decouple the rotation from the strain at a per-tet level, use a proper exponential integrator to advance the rotational component.

von Tycowicz et al. [2013] derive a column space expansion of a standard LMA subspace that significantly improves the amount of rotational motion it can capture. Skinning eigenmodes [Benckekroun et al. 2023; Trusty et al. 2023] derive a linear subspace based off of linear blend skinning, enabling the capture of large rotational motion. These have since been extended with neural fields to be used for arbitrary geometric discretizations [Chang et al. 2023; Modi et al. 2024].

All these methods only account for better accommodation of large rotational motion, however these methods do not account for the plethora of possible interactions the simulation may encounter at run-time.

2.2 Subspaces for Interaction

Aside from arbitrary large rotational deformation, a simulated soft body may undergo a large variety of other deformations at run-time. What makes these other deformations particularly difficult to account for upon subspace construction is that we cannot a priori

be certain what the exact external forces the body will be exposed to throughout the interactive simulation. Nevertheless, different subspaces exist that make different assumptions on singular specific modes of expected user interaction.

One common approach is to construct subspaces for simulation based off of geometric posing and modeling techniques. Joshi et al. [2007] create a subspace for posing their character from a user-designed coarse cage. Instead of cages, one can also pose their character with skeleton rigs and design subspaces well tailored for smooth rig-based interaction [Jacobson et al. 2011]. Coupling these rig-posing spaces to a physics simulation is also possible [Hahn et al. 2012; Jacobson et al. 2012; Wang et al. 2015].

While these rig-space physics spaces excel at capturing bulk skeletal motion, they frequently lack fine scaled dynamic effects. To address this James and Pai [2002] augment the controlled skeleton rig with modes obtained from LMA, and account for momentum transfer between the two. Because the LMA subspace falls out of date with large deformations, Xu and Barbič [2016] extend this and specifically build multiple LMA subspaces for various poses, and interpolate between them. Instead of standard characters driven by an underlying articulated skeleton, another line of work adds fine-scale physical detail to arbitrary rig animations, by performing simulation in a rig-complementary subspace [Benckekroun et al. 2023; Wu and Umetani 2023; Zhang et al. 2020].

Aside from character posing, another mode of interaction that has been explored for subspace simulation is contact interaction. A characteristic feature of contact simulation is its local response, which is the bane of any LMA-type subspace. To capture the local behavior, Brandt et al. [2018]; Faure et al. [2011]; Lan et al. [2021] create a linear subspace with many sparsely sampled handles [Jacobson et al. 2011], allowing them to resolve the local response characteristic of a contact simulation. Another approach to encourage locality in modal subspaces is to add an L_1 regularization term to the eigenvalue problem objective. Neumann et al. [2014] and Brandt and Hildebrandt [2017] show how to compute a modal analysis-like localized basis by adding such an L_1 regularization term to the eigenvalue problem objective.

Accounting for uncertainty in the contact loads to which a fabricated object may be subjected, Langlois et al. [2016] collect data from a rigid body simulation to construct a force distribution that is then mapped to a distribution on internal stresses. They use this stress field to predict probabilities of a fracture event occurring across the entire object.

To represent arbitrary contacts that are encountered at run-time, many works turn to adaptive schemes that automatically introduce new locally-defined basis functions upon the detection of a new contact event. Harmon and Zorin [2013] augment the subspace with *Boussinesq* functions about the contact point, while others bring back full order degrees of freedom at each vertex when enough local strain is detected as in Mercier-Aubin et al. [2022]; Teng et al. [2015]; Trusty et al. [2024].

Unfortunately, each of these subspaces is tailored to a specific mode of interaction comprising only a small range of the infinite possible interactions during an interactive simulation. For example, muscle actuation is a crucial mode of interaction with a physical environment, with many applications in character animation and




Fig. 4. Localizing deformations via domain decomposition easily leads to ripping or locking artifacts based on the user’s choice of domain coupling stiffness k . Instead, we interpret the same input subdomains as independent regions upon which Gaussian-distributed forces are applied. This correctly allows for deformation to leak between subdomains, leading to a simulation free from these artifacts.

biomechanics [Lee et al. 2018], however to the best of our knowledge, there are no subspace models that are specifically designed for muscle actuation. This leaves practitioners of these applications to make their own simplifications, bounding them to extremely coarse geometry in order to make their simulation tractable [Min et al. 2019].

We propose the first method for constructing simulation subspaces that are aware of arbitrary Gaussian force distributions, including those that arise from contact, handle interaction, as well as muscle, spring and pneumatic actuation.

3 Method

We present Force-Dual modes (FDM), a principled method for constructing simulation subspaces that are tailored to both the physical properties of the material as well as the anticipated force distributions that will actuate the system. To summarize, we propose to build a subspace matrix $\mathbf{B} \in \mathbb{R}^{3n \times m}$ through an eigendecomposition of the matrix:

$$\left(\mathbf{H}^{-1} \Sigma_F^T \Sigma_F \mathbf{H}^{-1} \right) \mathbf{B} = \mathbf{M} \mathbf{B} \mathbf{\Lambda} \quad (1)$$

where $\mathbf{H} \in \mathbb{R}^{3n \times 3n}$ is the physical energy Hessian, $\Sigma_F \in \mathbb{R}^{3n \times 3n}$ is the covariance matrix of the force distribution, $\mathbf{M} \in \mathbb{R}^{3n \times 3n}$ is the mass matrix, and $\mathbf{\Lambda} \in \mathbb{R}^{m \times m}$ is the diagonal eigenvalue matrix.

3.1 Statistical Formalism for Model Reduction

We will now show how we arrive at our final expression Eq. (1) by casting the model reduction problem into a probabilistic framework. While this probabilistic perspective on model reduction has deep roots in engineering—most notably through the Kosambi-Karhunen–Loève (KKL) expansion [Ghanem and Spanos 2003]—it is seldom explored in computer graphics, despite its potential to offer powerful insights for tasks we routinely face, such as evaluating the quality of reduced subspaces, quantifying approximation error under varying conditions, and reasoning about generalization to new inputs or unseen forces.

At a high level, we may frame model reduction as approximating the outcome of a complex system governed by a computationally expensive procedure,

$$U(\omega) = \mathcal{L}(F(\omega)), \quad (2)$$

where $U(\omega) \in \mathbb{R}^n$ denotes the quantity of interest (e.g., displacements), and $F(\omega) \in \mathbb{R}^k$ represents the system's input (e.g., boundary conditions, external forces, prior states). We treat both U and F as random variables dependent on a sample outcome $\omega \in \Omega$, capturing the natural variability and uncertainty present in future simulation configurations. We will suppress ω to simplify notation.

The assumption underlying model reduction is that the output U has a certain low-dimensional structure. For example, in linear model reduction, we approximately express the random structure of U as a small sum of linear components with amplitudes $Z \in \mathbb{R}^m$ and an optional mean term $\mu_U \in \mathbb{R}^n$,

$$U \approx \mathbf{B}Z + \mu_U, \quad (3)$$

where $\mathbf{B} \in \mathbb{R}^{n \times m}$ is the matrix whose m -th column corresponds to factor Z_m , and denotes the mean-zero error in the approximation. Given the basis \mathbf{B} , the optimal reconstruction of the factors Z in the \mathbf{M} -norm is given by the projection

$$Z = \mathbf{B}^\dagger (U - \mu_U), \quad (4)$$

where \mathbf{B}^\dagger denotes the pseudo-inverse,

$$\mathbf{B}^\dagger = (\mathbf{B}^T \mathbf{M} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{M}^T. \quad (5)$$

The objective of linear model reduction is to appropriately choose \mathbf{B} such that the approximation error is minimized in an appropriate metric \mathbf{M} ,

$$\min_{\mathbf{B}} \mathbb{E} \|U - \mathbf{B} \mathbf{B}^\dagger U\|_{\mathbf{M}}^2. \quad (6)$$

This problem, and related formulations, are common across many disciplines in engineering, statistics, and the sciences. In statistics, this framework is known as Principal Component Analysis (PCA) [Hotelling 1933]; in engineering, when reconstructing simulation snapshots, it is referred to as the Proper Orthogonal Decomposition (POD) [Sirovich 1987]. Its continuous analog appears in stochastic mechanics as the KKL expansion [Ghanem and Spanos 2003].

A unique solution to the problem does not exist unless one also applies an additional constraint. The complete optimization problem is therefore

$$\min_{\mathbf{B}} \mathbb{E} \|U - \mathbf{B} \mathbf{B}^\dagger U\|_{\mathbf{M}}^2 \quad \text{subject to} \quad \mathbf{B}^T \mathbf{M} \mathbf{B} = \mathbf{I}, \quad (7)$$

where the \mathbf{M} -orthogonality constraint ensures that the basis vectors are orthonormal with respect to the mass matrix. It can be shown (see App. A) that the optimal choice of \mathbf{B} are the *top* m eigenvectors of $\Sigma_U \mathbf{M}$, i.e.,

$$\Sigma_U \mathbf{M} \mathbf{B} = \mathbf{B} \Lambda, \quad (8)$$

where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_m)$ are the eigenvalues. And if Σ_U is invertible, then this is equivalent to the *bottom* m generalized eigenvectors of equation

$$\Sigma_U^{-1} \mathbf{B} = \mathbf{M} \mathbf{B} \Lambda^{-1}. \quad (9)$$

To summarize, once we know a covariance Σ_U for the quantity of interest, we can generate an error-minimizing subspace for simulation by solving the generalized eigenvalue problem Eq. (9). Different approaches—whether they acknowledge this explicitly or not—make different assumptions about this distribution.

Data-driven methods, such as POD, estimate Σ_U from high-resolution simulation snapshots. While accurate, this process is often expensive, tedious, and problematic in interactive contexts. For example, it is paradoxical to predict the distribution of *interactive* forces that

will occur during interactive use when only relying on *offline* simulations.

Other approaches circumvent this by having users build smooth, localized weights [Jacobson et al. 2011; Wang et al. 2015]. While this allows users to encode some intuitive features, such as the area of influence of interaction, it fails to capture critical aspects that are not easily conceptualized—such as boundary conditions, material heterogeneity, or anticipated contact regions (See Fig. 5).

A third common choice is to derive Σ_U from the inverse Hessian of an elastic energy function, which corresponds to Linear Modal Analysis (LMA) [Pentland and Williams 1989]. This approach is appealing due to its physical grounding, but it lacks any mechanism for incorporating prior knowledge about expected forces. For instance, the elastic energy Hessian does not encode information about where contact might occur, rendering it blind to one of the most important sources of localized deformation.

In this work, we propose a different approach: we construct Σ_U not directly, but as the result of propagating uncertainty in the input forces F through the system operator \mathcal{L} . In other words, we first model the statistics of F —which are often more intuitive and accessible to the user—and then transform those statistics through the system, imbuing them with non-intuitive physical information, to produce the resulting Σ_U .

As we will show, many interactions in a scene naturally induce plausible statistical models for F . These in turn give rise to a physically meaningful and computationally effective choice of covariance for the solution space. While users may struggle to guess the appropriate statistics of Σ_U directly, they often have a much better intuition for the kinds of forces or interactions their system will encounter.

3.2 Subspaces from Forces Distributions

We will now derive our main contribution, a derivation of the covariance matrix Σ_U from a distribution of forces. We formally define the process \mathcal{L} for our application of elastodynamic simulation, and start at the standard variational perspective that casts simulation as the solution to a minimization problem,

$$U^* = \underset{U}{\operatorname{argmin}} \mathcal{V}(U) + C(U), \quad (10)$$

where \mathcal{V} represents a physical energy (e.g., elastic energy, kinetic energy, constraints) and C denotes an external influence on the system (e.g. from contact, actuation, or other user interaction). We make a first order approximation of our external potential to expose the external forces $F = -\nabla C(U)$ acting on our system,

$$U^* = \underset{U}{\operatorname{argmin}} \mathcal{V}(U) - U^T F. \quad (11)$$

Our goal is to relate the randomness in the distribution of F , to the distribution on U^* .

Assuming U is close to the rest state of the system at which \mathcal{V} is minimized ($U = 0$), one may perform a second order Taylor expansion of \mathcal{V} to arrive at the quadratic approximation, $\mathcal{V}(U) \approx \frac{1}{2} U^T \mathbf{H} U$, simplifying our variational problem further,

$$U^* = \underset{U}{\operatorname{argmin}} \frac{1}{2} U^T \mathbf{H} U - U^T F, \quad (12)$$




Fig. 5. Interpreting user input as a force prior leads to a subspace that is physically aware, allowing us to fully bend the chicken wing backwards, without bending the bone. Instead, attempting to construct a subspace by directly prescribing a distribution on displacements leads to subspaces that are not material aware, and lead to locking during simulation, and a disrespect of the physical material properties.

where \mathbf{H} is the physical energy Hessian at $U = 0$. The solution to the above quadratic energy is then given by,

$$U^* = \mathbf{H}^{-1}F. \quad (13)$$

Assuming F is a multivariate Gaussian random variable $F \sim \mathcal{N}(\mu_F, \Sigma_F)$, the affine transformation of such a random variable is well known, and induces a Gaussian distribution on $U \sim \mathcal{N}(\mu_U, \Sigma_U)$, in terms of the parameters of the distribution of F ,

$$\mu_U = \mathbf{H}^{-1}\mu_F, \quad (14)$$

$$\Sigma_U = \mathbf{H}^{-1}\Sigma_F\mathbf{H}^{-1}, \quad (15)$$

where we've arrived at a distribution on displacements in terms of a distribution on forces.

3.2.1 Linear Modal Analysis as a Special Case LMA is a common technique used in structural analysis and model reduction that assumes that the most likely directions in which one will experience deformations is those with the minimum contribution to elastic energy, and therefore uses the lowest eigenvectors of the matrix \mathbf{H} as a reduced basis,

$$\mathbf{H}\mathbf{B}_{\text{modal}} = \mathbf{M}\mathbf{B}_{\text{modal}}\mathbf{\Gamma}, \quad (16)$$

where \mathbf{H} , in this case is the elastic energy Hessian (also known as the stiffness matrix), $\mathbf{\Gamma}$ is a diagonal matrix of modal frequencies and \mathbf{M} is the mass matrix. Within our Force-Dual framework, we show in App. A that by choosing our force covariance $\Sigma_F = \mathbf{M}$ to be the mass matrix, which arises from continuous white noise applied over our domain, we recover the exact same subspace as the one prescribed by LMA. To the best of our knowledge, we are the first to formally characterize the exact force distribution assumed by LMA.

This analysis offers an alternative perspective on LMA that makes it clear why the resulting modes are almost always globally supported — the forces F are assumed to be uncorrelated and almost always nonzero everywhere. This global force prior assumption introduces artifacts that have plagued model reduction techniques for decades, such as *spooky action at a distance*.

3.2.2 Green's Functions as a Special Case On the other extreme, instead of assuming random uncorrelated forces everywhere on the mesh, Green's functions form a subspace spanned by a small set of force-responses, a technique which has been used in structural analysis for thin plates and beams [Kim et al. 2013]. While this perspective is useful when the *exact* force impulses are known, the resulting subspace degrades in quality significantly when observed forces lie outside of its assumption, as shown in Fig. 6, which uses a Green's function subspace constructed by sampling localized normal forces at random points on the left half of the heart. We show in App. C that our Force-Dual modes framework also recovers Green's functions subspaces as a special case where the force distributions come from a set of randomly actuated forces, of smaller rank than the subspace dimension.

We view both LMA and Green's function spaces as opposite extremes of force prior assumptions. At one end, LMA corresponds to maximal uncertainty: spatially uncorrelated forces of equal variance applied everywhere. At the other end, Green's function spaces correspond to minimal uncertainty: the set of possible force directions is known and fixed.

Our framework allows us to find a smooth middle ground between these two extremes, as shown in 6, where our Force-Dual modes method constructs a prior of random uncorrelated normally directed forces on the left half of the heart. This lets us capture the local detail with more fidelity than LMA, without suffering from the sharp jarring deformations exhibited by the Green's function subspace, with the same number of degrees of freedom as both.

3.3 Constructing Generalized Force Distributions

Having established that our method corresponds to LMA in the case where $F \sim \mathcal{N}(\mathbf{0}, \mathbf{M})$, we now take our framework beyond this narrow assumption and consider a more general distribution $F \sim \mathcal{N}(\mu_F, \Sigma_F)$. As we will see, many modes of interaction are much better described by a different prior than the one assumed of LMA. Properly making use of such priors is key in designing subspaces well suited to arbitrary forms of user interaction.




Fig. 6. Green's functions describe responses from specific sampled sharp impulse forces, introducing artifacts when the observed forces are different from the samples. LMA constructs subspaces assuming spatially uncorrelated forces everywhere on the shape, struggling to accomodate specific local deformation. Our Force-Dual framework generalizes both of these approaches, and allows us to construct a subspace assuming a more realistic distribution of normal contact forces defined on the surface, resulting in a smooth localized response at run-time.

Let us delve into two overarching strategies for constructing generalized force distributions. We briefly begin with specifying nonuniform variances, before moving on to explore the construction of low-rank force distributions in greater detail.

3.3.1 Variance Specification The simplest generalization for user interaction involves modifying only the diagonal components of Σ_F with additional structure. For example, by making the force distribution spatially localized, we can trivially generate subspaces that are localized as well. As an example, Fig. 12 shows that by localizing the force variance about the left wing of the bat, the resulting first three skinning eigenmodes reflect the locality of the deformation.

On top of just locality, however, our subspaces are also aware of properties of the PDE operator such as smoothness, material heterogeneity and imposed boundary conditions. This is because our formulation makes use of the energy Hessian \mathbf{H} in defining our final distribution,

If the domain is already split up into regions in which a user will interact, as is done in substructuring [Kim and James 2011], it's very easy to define subspaces based off of these regions by localizing a force distribution within that region as shown Fig. 4. Our subspace allows deformation to smoothly leak between adjacent substructures without the need to tune auxiliary parameters such as interface stiffness.

Furthermore we find it is relatively easy to craft and *paint* non-uniform force distributions that *just work*, as we enable the user to focus on selecting the components of the shape they plan on handling independently, without having to worry about complex simulation details. Fig. 5 shows a user crafting subspaces with various rig-weight crafting algorithms such as Bounded Biharmonic Weights [Jacobson et al. 2011], Biharmonic Coordinates [Wang et al. 2015] and manual weight painting (guided by diffusion) with common geometry processing operations such as smoothing and diffusion. All of these methods allow the bones to bend, failing to account for their prescribed near-rigid material stiffness. By taking the exact same manually painted weights, and instead interpreting them as load variances, we can use our Force-Dual mode framework to obtain a Force-Dual subspace well suited to the heterogeneous flapping motion of the chicken wing.



Fig. 7. The locality of the input force distribution is reflected in the resulting force-dual skinning eigenmodes.



Fig. 8. A contact patch being used to specify a distribution on displacements results in an extremely localized deformation, deforming only the mattress as the ogre sits. Interpreting the contact patch instead as specifying a force distribution allows the resulting simulation to also capture global components of the deformation, such as the bending of the stiff frame Eq. (21). In this example, contact simulation is performed using a standard quadratic spring penalty with a single mode, and the ogre is approximated with a spherical collider.

3.3.2 Linear Low-Rank Force Distributions Some simulation interactions imply their own specific force distribution, and capturing this force interaction is crucial for the quality of the resulting simulated interaction. We assume these force interactions are linearly parameterized by a set of random actuation coefficients $A \in \mathbb{R}^r$,

$$F = \mathbf{D}A, \quad (17)$$

where the columns of $\mathbf{D} \in \mathbb{R}^{3n \times r}$ span the space of possible forces to which a simulation may be subjected. It is then easier to reason about a Gaussian distribution on the actuation coefficients $A \sim \mathcal{N}(\mu_A, \Sigma_A)$, which induces a force distribution of

$$F \sim \mathcal{N}(\mathbf{D}\mu_A, \mathbf{D}\Sigma_A\mathbf{D}^T). \quad (18)$$

In the following, we walk through several strategies for constructing linear low-rank force distributions.

Point Handle Interaction A common interaction mode is to pin vertices and encourage them towards a target position via a quadratic spring-like objective.

$$C(U) = \alpha \|\mathbf{S}^T U - A\|_N^2 \quad (19)$$

with α being the strength of the objective, $\mathbf{S} \in \mathbb{R}^{3n \times r}$ is the selection matrix selecting vertex displacements from U , $\mathbf{N} = \mathbf{S}^T \mathbf{M} \mathbf{S}$ is the diagonal mass-matrix containing the mass of the selected vertices, and $A \in \mathbb{R}^r$ are the target displacements our selection should take. We again treat A as normally distributed random variables.

The above objective leads to the forces

$$F = -\frac{\partial C}{\partial U} = \alpha \mathbf{S} \mathbf{N} A = \mathbf{D}A \quad (20)$$

where $\mathbf{D} \in \mathbb{R}^{3n \times r}$ is a mass-weighting selection matrix selecting vertex forces from F . The form of the equation above is such that we can use Eq. (18) to define the final force distribution. Note that if $A \sim \mathcal{N}(0, \mathbf{I})$, the resulting Σ_F results in a diagonal distribution, where the only non-zero entries are those of the handle vertices,




Fig. 9. By leveraging our pneumatic actuation force prior, we correctly recover the expected behavior of the first mode, inflating the air pockets and bending the soft actuated gripper. In contrast, the first mode using the LMA force prior results in a completely uncharacteristic motion of the gripper.

weighed by their mass and the penalty strength α . Figs. 2, 5, 14, depict simulations involving point handle interaction with an elephant, chicken wing, and pegasus, respectively.

Contact A contact event j applies a contact force at surface vertex i as modelled by

$$F_i = \sum_j |C| v_i w_{ij} [\mathbf{n}_j \ \mathbf{t}_j \ \mathbf{r}_j] \begin{bmatrix} A_{n_j} \\ A_{t_j} \\ A_{r_j} \end{bmatrix}, \quad (21)$$

where v_i is the mass of vertex i , w_{ij} is a user crafted contact patch weight, defining the region on the mesh over which this force is applied. The unit vectors \mathbf{n}_j , \mathbf{t}_j and \mathbf{r}_j are the contact normals, and two tangent vectors defining the frame of contact j . The values A_{n_j} , A_{t_j} , A_{r_j} define the strength of each component of the contact force, which we treat as our normally distributed random actuations. With this simple framework, we are able to model localized contact response with physically aware-properties as shown in Fig. 8, where a user has crafted a contact patch for a single contact frame w_j . Note that by using the Force-Dual modes framework, we get the expected buckling of the bed under the ogre's weight, whereas making use of just the user-crafted contact patch results in the overly localized mattress deforming, with the rest of the bed remaining immobile.

The above equation can be simplified into our canonical linear form, $F = DA$, with $D \in \mathbb{R}^{3n \times 3|C|}$ where,

$$D_{3i:3i+3, 3j:3j+3} = w_{ij} [\mathbf{n}_j \ \mathbf{t}_j \ \mathbf{r}_j] \text{ and} \\ A_{3j:3j+3} = [A_{n_j} \ A_{t_j} \ A_{r_j}]^T.$$

From here, we can obtain the final distribution on forces from Eq. (18).

Although truly robust contact with large deformations usually requires an extremely non-linear incremental potential, making use of a linearized contact force approximation such as our own within the construction of our subspace greatly improves the run-time simulation quality, especially when compared to LMA and sparsely sampled Green's functions, as shown in Fig. 10 and Fig. 6.

Pneumatic Actuation Pneumatic Elastic Actuation is an extremely popular actuation strategy in soft robotics [Xavier et al. 2022], where




Fig. 10. The ball collides against the teddy bear in a variety of different configurations. Our method allows us to construct load distributions for each contact configuration, and detect which of the distributions is active, allowing us to better capture the localized deformation induced by the contact at run-time. Using the standard force prior assumed with Skinning Eigenmodes leads to spooky action at a distance.

air is pumped into an air pocket such that, upon inflation, causes an elastic force response corresponding to a desired bending motion. Unfortunately, because the interaction of the air with the elastic material is crucial, modeling this interaction requires detailed expensive elastic simulation that cannot be carried out in real-time. We show that by leveraging the force prior admitted by pneumatic actuation, we can build a subspace in which we can simulate such soft pneumatic actuation scenarios as shown in Fig. 9.

Specifically, assuming $|P|$ air pockets, an air pocket $j \in \{0, 1, \dots, |P|-1\}$, applies a force to every surface vertex i contouring that air pocket

$$F_i = \begin{cases} v_i \mathbf{n}_i A_j & i \in |V_j| \\ \mathbf{0} & \text{otherwise,} \end{cases} \quad (22)$$

where \mathbf{n}_i is the area averaged normal at vertex i , and v_i is the Voronoi surface area of vertex i , and $|V_j|$ are all the surface vertices incident on pocket j . It is then easy to recover our canonical low-rank force subspace, $F = DA$, with $D \in \mathbb{R}^{3n \times |P|}$ where

$$D_{3i:3i+3, j} = \begin{cases} v_i \mathbf{n}_i & i \in |V_j| \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (23)$$

Muscle Actuation The capacity to simulate active musculoskeletons, such as those of animals and humans, remains the unfortunate computational bottleneck for many character control tasks with soft agents, as arise in soft robotics, medicinal geometry, or character animation. This limits such applications to extremely coarse geometry [Lee et al. 2018; Min et al. 2019]. We show in Fig. 11 that once again, by leveraging a simple force prior admitted from such actuated systems, we can enable interactive simulation of musculoskeletons that previously would have exhibited jarring visual physical artifacts. We adopt a simple and popular quadratic muscle actuation model [Modi et al. 2021; Teran et al. 2003] that defines an active musculoskeletal energy. Assuming we are given a set of $|T|$ active elements for our musculoskeletal system, we have

$$C(U) = \frac{1}{2} \sum_j |T| v_j \mathbf{d}_j^T \mathbf{F}_j^T \mathbf{F}_j \mathbf{d}_j A_j, \quad (24)$$

where A_j is the per-element actuation signal, v_j is the per-element volume, $\mathbf{d}_j \in \mathbb{R}^3$ is the per-element muscle fiber direction and \mathbf{F}_j is



Fig. 11. We can generate subspaces well suited for heterogeneous material biomechanical muscle actuation. Standard skinning eigenmodes are not aware of the intended user interaction, and cannot accommodate the simulated flexing of the bicep without visual artifacts. Both simulations use four skinning eigenmodes.

the element deformation gradient. The resulting force vector that arises, to first order, is given by

$$F = -\frac{\partial C}{\partial U} = -\sum_t \frac{\partial C}{\partial F_j} : \frac{\partial F_j}{\partial U}, \quad (25)$$

$$= -\sum_j v_j F_j d_j d_j^T : \frac{\partial F_j}{\partial U} A_j = DA, \quad (26)$$

where $D \in \mathbb{R}^{3|V| \times |T|}$ is a force actuation matrix mapping from active element actuation $A \in \mathbb{R}^{|T|}$ to vertex forces and $\frac{\partial F_j}{\partial U} \in \mathbb{R}^{3 \times 3 \times 3|V|}$ is the standard per-element Jacobian operator [Sifakis and Barbič 2012].

Spring Actuation A similar type of actuation strategy popular in character animation and robotics are spring-based actuators [El-Atab et al. 2020; Geijtenbeek et al. 2013; Tan et al. 2012]. These springs change their rest length during simulation in order to help the character or robot achieve a desired motion.

Starting from the external potential for a set of r mass springs,

$$C(U) = ||l(U) - l_0||^2, \quad (27)$$

where $l \in \mathbb{R}^r$ and $l_0 \in \mathbb{R}^r$ are both vectors of per-spring edge lengths. Taking the negative gradient with respect to U reveals the induced force every timestep,

$$F = \frac{\partial l}{\partial U}^T (l_0 - l(U)), \quad (28)$$

where we've revealed $(l_0 - l(U))$, the changing quantity that actuates our system. We choose this quantity to be modeled as a random set of actuation parameters, retrieving the canonical form required of Eq. (18),

$$F = DA, \quad (29)$$

where $D = \frac{\partial l}{\partial U}^T \in \mathbb{R}^{3n \times r}$ is the edge-length Jacobian matrix mapping from spring actuation parameters $A \in \mathbb{R}^r$ to vertex forces. Fig. 1 and Fig. 18 show the real-time inverse kinematics control of a subspace simulation of a dragon. Not leveraging this spring-actuator force prior leads to stiff and erroneously global motion that should be unachievable with the given spring actuator arrangement.

3.4 Gaussian Mixtures for Multi-Modal Interaction

Until now, we have assumed that a given force distribution prior is applicable throughout a simulation. However, simulations often




Fig. 12. Our force-dual skinning eigenmodes reflect both the heterogeneous material properties of the material as well as the actuation distribution flexing the muscle fibers of the bicep.

have different regimes of actuation. For example, an elastic system may be subject to a large localized contact force at one point in time, but later on may be dominated by more global inertial forces. We can model such disparate modes of actuation by choosing a mixture of Gaussians for our force distribution,

$$F = \sum_{k=1}^K \chi_k F_k, \quad (30)$$

where F_k is the k -th Gaussian in our mixture model and the $\chi_k \in \{0, 1\}$ are mask variables such that the k -th Gaussian is only active when $\chi_k = 1$. The most common way to define a mixture is to let χ_k be the indicator for an independent categorical random variable $C \in \{1, \dots, K\}$,

$$\chi_k = \begin{cases} 1 & C = k \\ 0 & \text{otherwise.} \end{cases} \quad (31)$$

It is common to denote the probability distribution of C as a vector π such that $\pi_k = \mathbb{P}(C = k)$. The statistics of the mixture F now read

$$\begin{aligned} \mu_F &= \sum_{k=1}^K \pi_k \mu_k, \\ \Sigma_F &= \sum_{k=1}^K \pi_k (\Sigma_k + (\mu_k - \mu_F)(\mu_k - \mu_F)^T), \end{aligned} \quad (32)$$

where μ_k and Σ_k are the means and covariances of the components F_k .

On top of allowing us to design subspaces for multiple modes of actuation, a mixture formulation allows us to have a probabilistically motivated way of adaptively selecting a subspace. In particular, if we knew which component F_k was active during the simulation, we could select a basis of modes tailored for that subspace and solve for the amplitudes of those modes on-the-fly.

Although a direct observation of the force sample $F = \hat{F}$ does not uniquely identify which mixture component generated it, Bayes' Rule provides a principled way to infer the posterior probability of a component's responsibility for the observation during runtime,

$$\mathbb{P}(C = k | F = \hat{F}) = \frac{\mathbb{P}(F \in dF(\hat{F}) | C = k) \mathbb{P}(C = k)}{\mathbb{P}(F \in dF(\hat{F}))}.^1 \quad (33)$$

¹Here, the notation $F \in dF(\hat{F})$ is necessary because technically the probability $F = \hat{F}$ is zero, so the above expression must be taken in the limit of F being in an infinitesimal interval of size dF around \hat{F} .

Taking logarithms and expanding,

$$\begin{aligned} \log \mathbb{P}(C = k \mid F = \hat{F}) \\ = -\frac{1}{2}(\hat{F} - \mu_k)^T \Sigma_k^{-1}(\hat{F} - \mu_k) - \frac{1}{2} \log \det(\Sigma_k) - \log c, \end{aligned} \quad (34)$$

where c is a constant. Therefore, from an observation \hat{F} of forces acting on the object during runtime, we can adaptively select the most likely component that corresponds to this observation by finding the k that maximizes Eq. (34). This allows us to choose our subspaces adaptively during the simulation by using a basis computed from the covariance matrix of the k -th component. While truly evaluating Eq. (34) is expensive as it requires inverting the $3n \times 3n$ covariance matrix of each k -th component of the force distributions, we often don't require computing the full inverse. Because many of the observed interaction forces are usually only applied over a small sub-region of the domain, it's possible to marginalize the distributions on F to focus on these sub-regions. The resulting mean and covariance matrices of each Gaussian component can then be restricted to that given sub-region and quickly queried for real-time evaluation of Eq. (34).

Fig. 13 shows how given a selection of three load distributions defined on the seal, as well as an observed force, we can quantify how likely it was that the observed force was sampled from each of the load distributions, as marginalized over the subdomain on which the force is applied. We make use of this ability to efficiently choose which subspace to use for simulation at run-time, for a scene with multimodal contact interaction Fig. 10, and a scene with multimodal click interaction Fig. 14. This allows us to capture, using only a maximum of 4 Skinning Eigenmodes at any point in time, a much more realistic and interactive simulation than using a fixed subspace of the same size.

3.5 Implementation

The Generalized Eigenvalue Problem described by Eq. (8), when solved naively, quickly becomes intractably expensive due to the dense computation required of \mathbf{H}^{-1} . To address this challenge, we adopt two different strategies for constructing our subspace depending on the structure of Σ_F .

The simplest case lies when Σ_F is of full rank and either diagonal or block diagonal (with blocks small enough to invert quickly). In this scenario, we instead solve the equivalent problem of Eq. (9), which finds the smallest m -eigenvectors of Σ_U^{-1} . Recalling our expression for the covariance on displacements $\Sigma_U^{-1} = \mathbf{H}\Sigma_F^{-1}\mathbf{H}$, the GEVP becomes

$$\mathbf{H}\Sigma_F^{-1}\mathbf{H} = \mathbf{M}\mathbf{B}\mathbf{\Lambda}^{-1}. \quad (35)$$




Fig. 13. By leveraging Bayes' rule Eq. (34), we identify which force distribution was most likely to generate an observed sample force.




Fig. 14. A user clicks on the pegasus to assign a control handle, and controls the pegasus by moving the control handle with a spring-like energy. By adaptively picking the subspace based off the observed interaction force, our method allows us to efficiently and accurately capture the bend of the wing and the bow of the head, without the scaling or stretching artifacts that would arise from using a static Skinning Eigenmode subspace.

Because the inverse Σ_F^{-1} remains diagonally sparse and is fast to compute ($O(n)$) and \mathbf{H} and \mathbf{M} are both sparse as well, we can solve this problem with an out-of-the-box sparse GEVP solver, namely eigs from scipy [Virtanen et al. 2020].

Another common scenario arises when Σ_F is not diagonal, and may contain frequent off-diagonal terms. In this scenario, one saving grace occurs if Σ_F is low-rank. In this event, solving for the subspace is still feasible if we have a Cholesky-like decomposition on Σ_F and \mathbf{M} of the form

$$\Sigma_F = \mathbf{L}\mathbf{L}^T \quad \text{and} \quad \mathbf{M} = \mathbf{N}\mathbf{N}^T,$$

where $\mathbf{L} \in \mathbb{R}^{3n \times r}$ and $\mathbf{N} \in \mathbb{R}^{3n \times 3n}$ is sparse. We return to the original form of Eq. (8) and introduce a change of variables on $\tilde{\mathbf{B}} = \mathbf{N}^T \mathbf{B}$.

With this in hand, we compute our subspace by performing a singular value decomposition on the matrix

$$\mathbf{N}\mathbf{H}^{-1}\mathbf{L} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \quad (36)$$

where we retain the top m left-singular vectors and m singular values in order to form our subspace matrix and eigenvalues

$$\mathbf{B} = \mathbf{N}^{-1}\mathbf{U} \quad \text{and} \quad \mathbf{\Lambda} = \mathbf{S}^2. \quad (37)$$

This strategy is particularly advantageous when \mathbf{L} is thin ($r < 3n$), as it only requires the solution to r systems of equations sharing the same system matrix $\mathbf{H}^{-1}\mathbf{L}$, followed by a second set of linear solves $\mathbf{B} = \mathbf{N}^{-1}\mathbf{U}$. Fortunately in our case, we always assume a lumped diagonal mass matrix, which makes both identifying \mathbf{N} and inverting it trivial. Additionally, in all our examples with specialized force distributions, we obtain the Cholesky-like factor for free, as it is given by the structure of the canonical force distribution itself and we can safely set $\mathbf{L} = \mathbf{D}$.

4 Experiments and Discussion

4.1 Subspace Simulation Solvers

Our method is extremely general, and can be made to accelerate many standard solvers popular in computer graphics.

By default, all examples shown in this paper employ a standard Newton solver with backtracking line search [Li et al. 2019], as this can be used to minimize any smooth energies with a first and second derivative, such as those that arise in contact [Li et al. 2020] or muscle actuation [Modi et al. 2021]. To accelerate the evaluation of the reduced-space Hessian and gradient, we make use of a standard cubature scheme integration scheme [An et al. 2008], where cubature points and weights are selected according to the method of Trusty et al. [2023].

For certain convex energies, it's possible to use a Projective Dynamics (PD) [Bouaziz et al. 2014; Brandt et al. 2018] solver to solve the optimization problem. These solvers avoid recomputing the Hessian and drastically accelerate simulation, as shown in Table. 1. For examples involving actuated spring systems, we use this solver with an ARAP [Kim and Eberle 2022] elastic energy for the volumetric deformation and the fast mass spring energy [Liu et al. 2013] for the spring actuator deformation. To accelerate the evaluation of the local step for the ARAP term, we make use of the clustering scheme of Jacobson et al. [2012].

4.2 Run-Time Analysis

Table. 1 shows a table with run-time analysis of some of the scenes presented in this paper, using both the Newton solver, or the PD solver described above, showing in brackets on the last column the speedup factor we achieve when compared to the full space. Our subspace simulations are consistently orders of magnitude faster than their full space counterparts, allowing for hard-real time interaction.

4.3 Ablation on Spatial Decay of Load Variance

The more representative a force distribution to the true distribution of the observed forces, the more accurate the resulting simulation. Fig. 16 explores the convergence of a subspace simulation to full space solution as the assumed force prior becomes more representative of the actual imposed load.

The ground truth solution is defined as the response to a localized load applied on the left arm of the bear. We construct 4 different Force-Dual LMA subspaces, all assuming force priors that vary in their spatial decay and analyze how well they reconstruct the ground truth response as we vary the number of modes. Specifically, the load variance of each component of the force distribution at a point \mathbf{x} is given by $\sigma(\mathbf{x}) = 1 - (\frac{1}{\exp(-\alpha(\|\mathbf{x}-\mathbf{y}\|-r)^2)})$, where we set $\alpha = 10$ for a sharply decaying variance beyond radius r , which in turn controls the spatial locality of the variance.

For each radius value c , we compute the reconstruction error,

$$E_c = \|\mathbf{u}^* - \mathbf{u}_c\|_{\mathbf{M}}^2, \quad (38)$$

between the ground truth solution $\mathbf{u}^* = \mathbf{H}^{-1}\mathbf{f} \in \mathbb{R}^{3n}$ and the reduced-space solution at that radius value $\mathbf{u}_c = \mathbf{B}_c(\mathbf{B}_c^T\mathbf{H}\mathbf{B}_c)^{-1}\mathbf{B}_c^T\mathbf{f} \in \mathbb{R}^{3n}$.




Fig. 16. We show the reconstruction error $E_c = \|\mathbf{u}^* - \mathbf{u}_c\|_{\mathbf{M}}^2$ between the ground truth solution and reduced-space solutions as a function of the number of modes. Our Force-Dual method with appropriately localized force priors achieves orders of magnitude lower reconstruction error compared to LMA, which assumes spatially uncorrelated forces. The error decreases as the assumed force distribution more closely matches the actual applied load.

We show that even with a modestly localized force prior, we observe errors an order of magnitude smaller than the naive LMA baseline. Achieving lower error at low mode counts is especially important because the solution to a reduced space linear system evolves cubically with the size of the subspace, and simply adding more modes is rarely a sustainable strategy. The inset shows the solution time in seconds for a single reduced linear solve, of which frequently many need to be performed every non-linear simulation iteration.




Fig. 15. Computational cost scaling with number of modes.

every non-linear simulation iteration.

4.4 Comparison to Data-Driven Techniques

Another option of estimating distributions on displacements is through data driven techniques such as the Proper Orthogonal Decomposition [Sirovich 1987], where offline simulations are recorded and the subspace is then constructed from performing Principal Component Analysis (PCA) on the interaction responses.

Fig. 17 shows that, in the limit of exhaustive samples, the two approaches agree on the same subspace at the same subspace, but ours arrives at it directly whereas a data-driven technique needs a lengthy training process to arrive at it. Unfortunately, these data-driven techniques are ultimately dependent on the data given to them, and collecting these samples can be an expensive, tedious process.

Our Force-Dual subspace construction framework shows that in the case where we can describe the force distributions with a Gaussian, we don't need to sample this process whatsoever, and can arrive directly at the closed form expression for the subspace.

Table 1. Run-time analysis for some large scenes presented in this paper. Here, m denotes the number of subspace modes and k denotes the number of force distribution components.

Mesh	#Vertices	# Tets	m	k	Solver	Max iters	Full Sim. (sec/step)	FDM Sim. (sec/step)	FDM Recon. (sec/step)	FDM Total (sec/step)
Dragon (Fig. 1)	25,613	103,307	11	30	PD	10	7.6×10^0	2.1×10^{-3}	1.3×10^{-4}	$2.4 \times 10^{-3} (\times 3166)$
Teddy (Fig. 10)	21,209	89,942	24	50	Newton	3	6.7×10^0	2.3×10^{-2}	2.6×10^{-3}	$2.7 \times 10^{-2} (\times 248)$
Pegasus (Fig. 14)	15,281	58,401	24	200	Newton	1	2.8×10^0	3.2×10^{-2}	1.9×10^{-3}	$3.4 \times 10^{-2} (\times 82.4)$




Fig. 17. We compare our Force-Dual method to a data-driven PCA subspace. The left plot shows the mean alignment error between the first ten modes ($|\sum_{i=1}^{10} |\mathbf{B}_{FDM_i}^T \mathbf{B}_{PCA_i}|$) and the reconstruction error between the data-driven PCA modes and our Force-Dual Modes converge as more training samples are collected. The right plot shows the cost of data-driven PCA (dense SVD) scales with the number of samples, while our method requires no sampling. Our method immediately recovers the true principal components of the desired distribution without the need for expensive data collection.

4.5 Non-linear Motion

Our method, like Linear Modal Analysis, is developed on the linearized version of the variational form of the physical system. Despite this, our method can leverage decades of work in engineering and animation that enhances linear model reduction to better accommodate large non-linear deformation.

Fig. 9 shows we can easily apply the Rotation Strain Coordinates [Huang et al. 2011] interpolation framework to properly reconstruct the large deformation associated with the soft pneumatic gripper. Our framework requires no additional changes to be made to the Rotation Strain Coordinates algorithm, as they simply take as input a linear actuation and properly integrate the rotational component the actuation attempts to describe.

Furthermore, for many experiments involving large deformation, such as Fig. 4, Fig. 5, Fig. 14 and Fig. 16, we extend our framework and develop a Force-Dual version of Skinning Eigenmodes [Benckroun et al. 2023].

To achieve this, we collapse the vector distributions on $U \in \mathbb{R}^{3n}$ into a scalarized distribution on $\phi \in \mathbb{R}^n$ by simply summing the individual components on U

$$\phi = U_x + U_y + U_z. \quad (39)$$

The statistics on ϕ are now described by

$$\begin{aligned} \Sigma_\phi &= \Sigma_{U_{xx}} + \Sigma_{U_{yy}} + \Sigma_{U_{zz}}, \\ \mu_\phi &= \mu_{U_x} + \mu_{U_y} + \mu_{U_z}. \end{aligned}$$




Fig. 18. A single geometry may have arbitrary force distributions acting upon it. Ignoring this valuable prior and assuming the generic LMA force distribution leads to an extremely stiff simulation, where the dragon is unable to properly move its front leg (top), or head (bottom), without erroneously moving its entire body. Properly taking these distributions into account is crucial for realistic control tasks.

We can then identify the principal components of this distribution, and use those as skinning weights to construct the same linear blend skinning subspaces as defined in [Benckroun et al. 2023]. For low-rank force distributions, we still proceed as before in section Sec. 3.5, assembling the factor $\mathbf{K} = \mathbf{M}\mathbf{H}^{-1}\mathbf{L}$. However we instead perform an SVD on its scalarized \mathbf{K}_ϕ ,

$$\mathbf{K}_\phi = \mathbf{K}_x + \mathbf{K}_y + \mathbf{K}_z, \quad (40)$$

where $\mathbf{K}_d \in \mathbb{R}^{n \times r}$ is the component of \mathbf{K} corresponding to dimension d only. This extension recovers the same GEVP as developed by Benckroun et al. [2023] when the assumed forces are Gaussian and uncorrelated $\Sigma_F = \mathbf{M}$.

5 Inverse Kinematics on Spring Actuators

We are the first to show a system capable of real-time inverse kinematics on soft deformable characters of arbitrary mesh resolution with spring-like actuators, as shown in Fig. 1 and Fig. 18. To achieve this we use a Levenberg-Marquardt solver [Nocedal and Wright 1999] to find the spring actuations, taking one step of this solver every simulation timestep.

We find that being explicit about properties of your load distribution upon subspace construction is *critical* in physical control problems. Because naive subspaces like LMA couple motion of far away vertices together, using them to solve control tasks allow them to *cheat*.

For example, in Fig. 18, the inverse kinematics attempts to find the spring actuations of the head, or the front leg, that move a control point to a target location. Using an LMA subspace lets the dragon attempt to satisfy this constraint by unnaturally moving its back legs and body forward, two locations where it does not have spring actuations and therefore should have no control over. Our Force-Dual modes perspective emphasizes the importance of explicitly stating the underlying force distribution assumed of the subspace, resulting in a controlled simulation that naturally actuate the parts of the mesh the musculature can control.

6 Limitations & Conclusion

We are extremely excited about many future directions towards which this research might lead. In particular, the theory we have developed assumes linearized dynamics, making it easy to transfer between force distributions and displacement distributions. For simulations that exhibit extreme nonlinearity, as is very common in highly articulated systems undergoing contact [Sharp et al. 2023], we hope to develop a nonlinear version of this same theory, which would allow us to craft nonlinear subspaces sensitive to nonlinear actuation distributions.

Although we have primarily applied our framework to elastodynamics, the variational theory underpinning it is very general, and could be used to accurately accelerate many other simulations subjected to well known force distributions, such as fluid control [Chen et al. 2024; Cui et al. 2018].

Finally, the biggest bottleneck of soft-body character control remains the computational cost of the simulation. We think our Force-Dual subspace construction framework will be essential for making these previously impossible control problems feasible, and we look forward to attempting to use it on more complex control tasks such as swimming [Min et al. 2019] or walking [Tan et al. 2012].

In conclusion, we have shown that by leveraging a priori the expected distribution of forces a simulation will undergo, we can greatly improve the physical fidelity of interactive subspace simulations. To achieve this, we have presented a derivation of the subspace construction problem from a statistical perspective, allowing us to reason about fitting subspaces to specific force priors.

Our work fits within a broader trend in computer graphics toward embracing stochastic and statistical interpretations of traditional algorithms, as exemplified by recent work in stochastic geometry processing [?]. Our statistical perspective on subspace construction extends this trend to physics simulation, where we explicitly model the uncertainty inherent in interactive simulation through force distributions.

Not only does this derivation generalize prior subspace construction techniques, it allows us to design novel subspaces well suited for various common scene interactions previously impossible, such as handle-based control, contact, musculoskeletal actuation and pneumatic actuation and spring actuation.

Acknowledgments

We would like to thank John Hancock and Xuan Dam for administrative and systems support, David I.W. Levin and Silvia Sellán, Sarah Fernandez and Ty Trusty for illuminating discussions and Gabrielle Browne for proofreading.

References

- Steven S. An, Theodore Kim, and Doug L. James. 2008. Optimizing cubature for efficient integration of subspace deformations. *ACM Trans. Graph.* 27, 5, Article 165 (Dec. 2008), 10 pages. <https://doi.org/10.1145/1409060.1409118>
- Jernej Barbič and Doug L. James. 2005. Real-Time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph.* 24, 3 (July 2005), 982–990. <https://doi.org/10.1145/1073204.1073300>
- Otman Benchekroun, Jiayi Eris Zhang, Siddhartha Chaudhuri, Eitan Grinspun, Yi Zhou, and Alec Jacobson. 2023. Fast Complementary Dynamics via Skinning Eigenmodes. *ACM Trans. Graph.* 42, 4, Article 106 (July 2023), 21 pages. <https://doi.org/10.1145/3592404>
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: fusing constraint projections for fast simulation. *ACM Trans. Graph.* 33, 4, Article 154 (July 2014), 11 pages. <https://doi.org/10.1145/2601097.2601116>
- Christopher Brandt, Elmar Eisemann, and Klaus Hildebrandt. 2018. Hyper-reduced projective dynamics. *ACM Trans. Graph.* 37, 4, Article 80 (July 2018), 13 pages. <https://doi.org/10.1145/3197517.3201387>
- Christopher Brandt and Klaus Hildebrandt. 2017. Compressed vibration modes of elastic bodies. *Computer Aided Geometric Design* 52–53 (2017), 297–312. <https://doi.org/10.1016/j.cagd.2017.03.004>
- Yue Chang, Peter Yichen Chen, Zhecheng Wang, Maurizio M. Chiamonte, Kevin Carlberg, and Eitan Grinspun. 2023. LiCROM: Linear-Subspace Continuous Reduced Order Modeling with Neural Fields. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 111, 12 pages. <https://doi.org/10.1145/3610548.3618158>
- Yixin Chen, David Levin, and Timothy Langlois. 2024. Fluid Control with Laplacian Eigenfunctions. In *ACM SIGGRAPH 2024 Conference Papers*. 1–11.
- Min Gyu Choi and Hyeon-Seok Ko. 2005. Modal Warping: Real-Time Simulation of Large Rotational Deformation and Manipulation. *IEEE Transactions on Visualization and Computer Graphics* 11, 1 (Jan. 2005), 91–101. <https://doi.org/10.1109/TVCG.2005.13>
- Qiaodong Cui, Pradeep Sen, and Theodore Kim. 2018. Scalable laplacian eigenfluids. *ACM Trans. Graph.* 37, 4, Article 87 (July 2018), 12 pages. <https://doi.org/10.1145/3197517.3201352>
- Nazek El-Atab, Rishabh B Mishra, Fhad Al-Modaf, Lana Joharji, Aljohara A Alsharif, Haneen Alamoudi, Marlon Diaz, Nadeem Qaiser, and Muhammad Mustafa Hussain. 2020. Soft actuators for soft robotic applications: A review. *Advanced Intelligent Systems* 2, 10 (2020), 2000128.
- François Faure, Benjamin Gilles, Guillaume Bousquet, and Dinesh K. Pai. 2011. Sparse meshless models of complex deformable solids. *ACM Trans. Graph.* 30, 4, Article 73 (July 2011), 10 pages. <https://doi.org/10.1145/2010324.1964968>
- Thomas Geijtenbeek, Michiel van de Panne, and A. Frank van der Stappen. 2013. Flexible Muscle-Based Locomotion for Bipedal Creatures. *ACM Transactions on Graphics* 32, 6 (2013).
- Roger G Ghanem and Pol D Spanos. 2003. *Stochastic finite elements: a spectral approach*. Courier Corporation.
- Fabian Hahn, Sebastian Martin, Bernhard Thomaszewski, Robert Sumner, Stelian Coros, and Markus Gross. 2012. Rig-space physics. *ACM Trans. Graph.* 31, 4, Article 72 (July 2012), 8 pages. <https://doi.org/10.1145/2185520.2185568>
- David Harmon and Denis Zorin. 2013. Subspace integration with local deformations. *ACM Trans. Graph.* 32, 4, Article 107 (July 2013), 10 pages. <https://doi.org/10.1145/2461912.2461922>
- Harold Hotelling. 1933. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology* 24, 6 (1933), 417.
- Jin Huang, Yiyang Tong, Kun Zhou, Hujun Bao, and Mathieu Desbrun. 2011. Interactive Shape Interpolation through Controllable Dynamic Deformation. *IEEE Transactions on Visualization and Computer Graphics* 17, 7 (July 2011), 983–992. <https://doi.org/10.1109/TVCG.2010.109>
- Alec Jacobson, Ilya Baran, Ladislav Kavan, Jovan Popović, and Olga Sorkine. 2012. Fast automatic skinning transformations. *ACM Trans. Graph.* 31, 4, Article 77 (July 2012), 10 pages. <https://doi.org/10.1145/2185520.2185573>
- Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. 2011. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph.* 30, 4, Article 78 (July 2011), 8 pages. <https://doi.org/10.1145/2010324.1964973>
- Doug L. James and Dinesh K. Pai. 2002. DyRT: dynamic response textures for real time deformation simulation with graphics hardware. *ACM Trans. Graph.* 21, 3 (July 2002), 582–585. <https://doi.org/10.1145/566654.566621>
- Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. 2007. Harmonic coordinates for character articulation. *ACM Trans. Graph.* 26, 3 (July 2007), 71–es. <https://doi.org/10.1145/1276377.1276466>
- Kwangkeun Kim, Adrian G Radu, XQ Wang, and Marc P Mignolet. 2013. Nonlinear reduced order modeling of isotropic and functionally graded plates. *International Journal of Non-Linear Mechanics* 49 (2013), 100–110.

- Theodore Kim and David Eberle. 2022. Dynamic deformables: implementation and production practicalities (now with code!). In *ACM SIGGRAPH 2022 Courses* (Vancouver, British Columbia, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 7, 259 pages. <https://doi.org/10.1145/3532720.3535628>
- Theodore Kim and Doug L. James. 2011. Physics-based character skinning using multi-domain subspace deformations. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Vancouver, British Columbia, Canada) (SCA '11). Association for Computing Machinery, New York, NY, USA, 63–72. <https://doi.org/10.1145/2019406.2019415>
- Lei Lan, Yin Yang, Danny Kaufman, Junfeng Yao, Minchen Li, and Chenfanfu Jiang. 2021. Medial IPC: accelerated incremental potential contact with medial elastics. *ACM Trans. Graph.* 40, 4, Article 158 (July 2021), 16 pages. <https://doi.org/10.1145/3450626.3459753>
- Timothy Langlois, Ariel Shamir, Daniel Dror, Wojciech Matusik, and David I. W. Levin. 2016. Stochastic structural analysis for context-aware design and fabrication. *ACM Trans. Graph.* 35, 6, Article 226 (Dec. 2016), 13 pages. <https://doi.org/10.1145/2980179.2982436>
- Seunghwan Lee, Ri Yu, Jungnam Park, Mridul Aanjaneya, Eftychios Sifakis, and Jehee Lee. 2018. Dexterous manipulation and control with volumetric muscles. *ACM Trans. Graph.* 37, 4, Article 57 (July 2018), 13 pages. <https://doi.org/10.1145/3197517.3201330>
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Trans. Graph.* (SIGGRAPH) 39, 4, Article 49 (2020).
- Minchen Li, Ming Gao, Timothy Langlois, Chenfanfu Jiang, and Danny M. Kaufman. 2019. Decomposed optimization time integrator for large-step elastodynamics. *ACM Trans. Graph.* 38, 4, Article 70 (July 2019), 10 pages. <https://doi.org/10.1145/3306346.3322951>
- Tiantian Liu, Adam W. Bargteil, James F. O'Brien, and Ladislav Kavan. 2013. Fast simulation of mass-spring systems. *ACM Trans. Graph.* 32, 6, Article 214 (Nov. 2013), 7 pages. <https://doi.org/10.1145/2508363.2508406>
- Alexandre Mercier-Aubin, Paul G. Kry, Alexandre Winter, and David I. W. Levin. 2022. Adaptive Rigidification of Elastic Solids. *ACM Trans. Graph.* 41, 4, Article 71, 11 pages. <https://doi.org/10.1145/3528223.3530124>
- Sehee Min, Jungdam Won, Seunghwan Lee, Jungnam Park, and Jehee Lee. 2019. SoftCon: simulation and control of soft-bodied animals with biomimetic actuators. *ACM Trans. Graph.* 38, 6, Article 208 (Nov. 2019), 12 pages. <https://doi.org/10.1145/3355089.3356497>
- V. Modi, L. Fulton, A. Jacobson, S. Sueda, and D.I.W. Levin. 2021. EMU: Efficient Muscle Simulation in Deformation Space. *Computer Graphics Forum* 40, 1 (2021), 234–248. <https://doi.org/10.1111/cgf.14185> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14185
- Vismay Modi, Nicholas Sharp, Or Perel, Shinjiro Sueda, and David I. W. Levin. 2024. Simplicitis: Mesh-Free, Geometry-Agnostic Elastic Simulation. *ACM Trans. Graph.* 43, 4, Article 117 (July 2024), 11 pages. <https://doi.org/10.1145/3658184>
- T. Neumann, K. Varanasi, C. Theobalt, M. Magnor, and M. Wacker. 2014. Compressed manifold modes for mesh processing. In *Proceedings of the Symposium on Geometry Processing* (Cardiff, United Kingdom) (SGP '14). Eurographics Association, Goslar, DEU, 35–44. <https://doi.org/10.1111/cgf.12429>
- Jorge Nocedal and Stephen J Wright. 1999. *Numerical optimization*. Springer.
- Zherong Pan, Hujun Bao, and Jin Huang. 2015. Subspace dynamic simulation using rotation-strain coordinates. *ACM Trans. Graph.* 34, 6, Article 242 (Nov. 2015), 12 pages. <https://doi.org/10.1145/2816795.2818090>
- A. Pentland and J. Williams. 1989. Good vibrations: modal dynamics for graphics and animation. In *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques* (SIGGRAPH '89). Association for Computing Machinery, New York, NY, USA, 215–222. <https://doi.org/10.1145/74333.74355>
- Nicholas Sharp, Cristian Romero, Alec Jacobson, Etienne Vouga, Paul Kry, David I.W. Levin, and Justin Solomon. 2023. Data-Free Learning of Reduced-Order Kinematics. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 40, 9 pages. <https://doi.org/10.1145/3588432.3591521>
- Eftychios Sifakis and Jernej Barbič. 2012. FEM Simulation of 3D Deformable Solids: A Practitioner's Guide to Theory, Discretization and Model Reduction. In *ACM SIGGRAPH 2012 Courses* (Los Angeles, California, USA). Association for Computing Machinery, New York, NY, USA, 20:1–20:50. <https://doi.org/10.1145/2343483.2343501>
- Lawrence Sirovich. 1987. Turbulence and the dynamics of coherent structures. I. Coherent structures. *Quarterly of applied mathematics* 45, 3 (1987), 561–571.
- Jie Tan, Greg Turk, and C. Karen Liu. 2012. Soft body locomotion. *ACM Trans. Graph.* 31, 4, Article 26 (July 2012), 11 pages. <https://doi.org/10.1145/2185520.2185522>
- Yun Teng, Mark Meyer, Tony DeRose, and Theodore Kim. 2015. Subspace condensation: full space adaptivity for subspace deformations. *ACM Trans. Graph.* 34, 4, Article 76 (July 2015), 9 pages. <https://doi.org/10.1145/2766904>
- J. Teran, S. Blemker, V. Ng Thow Hing, and R. Fedkiw. 2003. Finite volume methods for the simulation of skeletal muscle. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (San Diego, California) (SCA '03). Eurographics Association, Goslar, DEU, 68–74.
- Ty Trusty, Otman Benckekroun, Eitan Grinspun, Danny M. Kaufman, and David I.W. Levin. 2023. Subspace Mixed Finite Elements for Real-Time Heterogeneous Elastodynamics. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 112, 10 pages. <https://doi.org/10.1145/3610548.3618220>
- Ty Trusty, Yun (Raymond) Fei, David Levin, and Danny Kaufman. 2024. Trading Spaces: Adaptive Subspace Time Integration for Contacting Elastodynamics. *ACM Trans. Graph.* 43, 6, Article 227 (Nov. 2024), 16 pages. <https://doi.org/10.1145/3687946>
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, António H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Christoph von Tycowicz, Christian Schulz, Hans-Peter Seidel, and Klaus Hildebrandt. 2013. An efficient construction of reduced deformable objects. *ACM Trans. Graph.* 32, 6, Article 213 (Nov. 2013), 10 pages. <https://doi.org/10.1145/2508363.2508392>
- Yu Wang, Alec Jacobson, Jernej Barbič, and Ladislav Kavan. 2015. Linear subspace design for real-time shape deformation. *ACM Trans. Graph.* 34, 4, Article 57 (July 2015), 11 pages. <https://doi.org/10.1145/2766952>
- Yuhan Wu and Nobuyuki Umetani. 2023. Two-Way Coupling of Skinning Transformations and Position Based Dynamics. *Proc. ACM Comput. Graph. Interact. Tech.* 6, 3, Article 47 (Aug. 2023), 18 pages. <https://doi.org/10.1145/3606930>
- Matheus S Xavier, Charbel D Tawk, Ali Zolfagharian, Joshua Pinski, David Howard, Taylor Young, Jiewen Lai, Simon M Harrison, Yuen K Yong, Mahdi Bodaghi, et al. 2022. Soft pneumatic actuators: A review of design, fabrication, modeling, sensing, control and applications. *IEEE Access* 10 (2022), 59442–59485.
- Hongyi Xu and Jernej Barbič. 2016. Pose-space subspace dynamics. *ACM Trans. Graph.* 35, 4, Article 35 (July 2016), 14 pages. <https://doi.org/10.1145/2897824.2925916>
- Jiayi Eris Zhang, Seungbae Bang, David I. W. Levin, and Alec Jacobson. 2020. Complementary dynamics. *ACM Trans. Graph.* 39, 6, Article 179 (Nov. 2020), 11 pages. <https://doi.org/10.1145/3414685.3417819>

A Deriving Optimal Subspace reconstruction

We can derive a mass orthonormal PCA/KKL as the subspace matrix $\mathbf{B} \in \mathbb{R}^{n \times m}$ that minimizes the reconstruction error for the stochastic process \mathbf{U} .

$$\argmin_{\mathbf{B}} \mathbb{E} [\|\mathbf{U} - \hat{\mathbf{U}}\|_{\mathbf{M}}^2] \quad s.t. \quad \mathbf{B}^T \mathbf{M} \mathbf{B} = \mathbf{I} \quad (41)$$

Where $\hat{\mathbf{U}} = \mathbf{B}^+ \mathbf{U} = \mathbf{B} \mathbf{B}^T \mathbf{M} \mathbf{U}$ is the projected approximation of the stochastic process. Note, we assume without loss of generality that \mathbf{U} is mean-zero. First, consider the term inside the expectation,

$$\|\mathbf{U} - \hat{\mathbf{U}}\|_{\mathbf{M}}^2 = \text{tr} \left(\mathbf{U}^T \mathbf{M} \mathbf{U} + \hat{\mathbf{U}}^T \mathbf{M} \hat{\mathbf{U}} - 2 \hat{\mathbf{U}}^T \mathbf{M} \mathbf{U} \right). \quad (42)$$

We drop the constant term $\mathbf{U}^T \mathbf{M} \mathbf{U}$ and continue,

$$\begin{aligned} \|\mathbf{U} - \hat{\mathbf{U}}\|_{\mathbf{M}}^2 &\cong \text{tr} \left(\mathbf{U}^T \mathbf{M} \mathbf{B} \underbrace{\mathbf{B}^T \mathbf{M} \mathbf{B}}_{\mathbf{I}} \mathbf{B}^T \mathbf{M} \mathbf{U} - 2 \mathbf{U}^T \mathbf{M} \mathbf{B} \mathbf{B}^T \mathbf{M} \mathbf{U} \right) \\ &= \text{tr} \left(\mathbf{U}^T \mathbf{M} \mathbf{B} \mathbf{B}^T \mathbf{M} \mathbf{U} - 2 \mathbf{U}^T \mathbf{M} \mathbf{B} \mathbf{B}^T \mathbf{M} \mathbf{U} \right) \\ &= \text{tr} \left(\mathbf{B}^T \mathbf{M} \mathbf{U} \mathbf{U}^T \mathbf{M} \mathbf{B} \right) - 2 \text{tr} \left(\mathbf{B}^T \mathbf{M} \mathbf{U} \mathbf{U}^T \mathbf{M} \mathbf{B} \right) \\ &= - \text{tr} \left(\mathbf{B}^T \mathbf{M} \mathbf{U} \mathbf{U}^T \mathbf{M} \mathbf{B} \right). \end{aligned} \quad (43)$$

Taking expectations and interchanging with trace gives

$$\mathbb{E} \|\mathbf{U} - \hat{\mathbf{U}}\|_{\mathbf{M}}^2 \cong - \text{tr} \left(\mathbf{B}^T \mathbf{M} \mathbb{E} [\mathbf{U} \mathbf{U}^T] \mathbf{M} \mathbf{B} \right) = - \text{tr} \left(\mathbf{B}^T \mathbf{M} \Sigma_{\mathbf{U}} \mathbf{M} \mathbf{B} \right). \quad (44)$$

Introducing Lagrange multipliers Λ for the constraint, the Lagrangian of eq. (41) is written

$$\mathcal{L} = -\text{tr}(\mathbf{B}^T \mathbf{M} \Sigma_U \mathbf{M} \mathbf{B}) + \text{tr}(\Lambda(\mathbf{B}^T \mathbf{M} \mathbf{B} - \mathbf{I})). \quad (45)$$

The variation of the Lagrangian is given by

$$\begin{aligned} \delta \mathcal{L} &= -2 \text{tr}((\delta \mathbf{B})^T \mathbf{M} \Sigma_U \mathbf{M} \mathbf{B}) + 2 \text{tr}(\Lambda((\delta \mathbf{B})^T \mathbf{M} \mathbf{B})) \\ &= -2 \text{tr}((\delta \mathbf{B})^T (\mathbf{M} \Sigma_U \mathbf{M} \mathbf{B} - \mathbf{M} \mathbf{B} \Lambda)), \end{aligned} \quad (46)$$

where we have used the symmetry of trace while taking variations. Now, the only way $\delta \mathcal{L}$ can be zero for all variations $\delta \mathbf{B}$ is for $\mathbf{M} \Sigma_U \mathbf{M} \mathbf{B} - \mathbf{M} \mathbf{B} \Lambda$ to identically be zero. Thus, the critical points of the optimization satisfy the truncated generalized eigenvalue problem

$$\mathbf{M} \Sigma_U \mathbf{M} \mathbf{B} = \mathbf{M} \mathbf{B} \Lambda, \quad (47)$$

with Λ the largest eigenvalues of $\mathbf{M} \Sigma_U \mathbf{M}$, and \mathbf{B} their corresponding eigenvectors. This can be further simplified by multiplying both sides by \mathbf{M}^{-1} ,

$$\Sigma_U \mathbf{M} \mathbf{B} = \mathbf{B} \Lambda. \quad (48)$$

While this tells us that \mathbf{B} should satisfy the above eigenvalue equation, it doesn't tell us *which* eigenvectors we should take. In order to answer that question, we plug eq. (48) back into the objective to see that

$$\mathbb{E} \|\mathbf{U} - \hat{\mathbf{U}}\|_{\mathbf{M}}^2 \cong -\text{tr}(\mathbf{B}^T \mathbf{M} \mathbf{B} \Lambda) = -\text{tr} \Lambda. \quad (49)$$

Hence, in order to minimize the objective, we want to choose the top m eigenvectors of the generalized eigenvalue problem eq. (48).

B Deriving Modal Analysis from Force-Dual Modes

The force distribution has $\Sigma_F = \mathbf{M}$. Hence, we arrive at $\Sigma_U = \mathbf{H}^{-1} \mathbf{M} \mathbf{H}^{-1}$ from eq. (15). The corresponding eigenvalue equation for the optimal basis \mathbf{B} is the lowest eigenvectors of

$$\mathbf{H} \mathbf{M}^{-1} \mathbf{H} \mathbf{B} = \mathbf{M} \mathbf{B} \Lambda^{-1}. \quad (50)$$

Observe the modes $\mathbf{B}_{\text{modal}}$ satisfy this equation, since,

$$\begin{aligned} \mathbf{H} \mathbf{M}^{-1} (\mathbf{H} \mathbf{B}_{\text{modal}}) &= \mathbf{H} \mathbf{M}^{-1} \mathbf{M} \mathbf{B}_{\text{modal}} \mathbf{\Gamma} \\ &= \mathbf{H} \mathbf{B}_{\text{modal}} \mathbf{\Gamma} = \mathbf{M} \mathbf{B}_{\text{modal}} \mathbf{\Gamma}^2, \end{aligned} \quad (51)$$

with $\mathbf{\Gamma}^2 = \Lambda^{-1}$. Hence, modal analysis arises naturally as the optimal linear basis for this force distribution.

$$\alpha \mathbf{H} \mathbf{B} = \mathbf{M} \mathbf{B} \Lambda^{-1}. \quad (52)$$

which we observe is the same as the modal analysis equations eq. (16) but with $\mathbf{\Gamma} = \alpha^{-1} \Lambda^{-1}$.

C Deriving Green's Functions from Force-Dual Modes

We consider the special case where the force prior is low-rank Gaussian, $\mathbf{F} \sim \mathcal{N}(0, \mathbf{F} \mathbf{F}^T)$, where $\mathbf{F} \in \mathbb{R}^{3n \times r}$ is a tall skinny matrix of r different applied forces. This force distribution induces a displacement distribution with covariance $\Sigma_U = \mathbf{G} \mathbf{G}^T$, where $\mathbf{G} = \mathbf{H}^{-1} \mathbf{F}$. The Force-Dual framework constructs a subspace by solving the generalized eigenvalue problem $\Sigma_U \mathbf{M} \mathbf{B} = \mathbf{B} \Lambda$, which becomes

$$\mathbf{G} \mathbf{G}^T \mathbf{M} \mathbf{B} = \mathbf{B} \Lambda.$$

Examining the left hand side: Since for any matrix \mathbf{A} , $\text{Col}(\mathbf{G} \mathbf{A}) \subset \text{Col}(\mathbf{G})$, therefore $\text{Col}(\text{LHS}) \subset \text{Col}(\mathbf{G})$. Turning to the RHS, since Λ is invertible, $\text{Col}(\mathbf{B}) = \text{Col}(\mathbf{B} \Lambda) = \text{Col}(\text{RHS})$. Finally, since $\text{Col}(\text{LHS}) = \text{Col}(\text{RHS})$, therefore $\text{Col}(\mathbf{B}) = \text{Col}(\text{LHS}) \subset \text{Col}(\mathbf{G})$.