

# Fast Complementary Dynamics via Skinning Eigenmodes

OTMAN BENCHEKROUN, University of Toronto, Canada

JIAYI ERIS ZHANG, Stanford University, U.S.A.

SIDDHARTHA CHAUDHURI, Adobe Research, India

EITAN GRINSPUN, University of Toronto, Canada

YI ZHOU, Adobe Research, U.S.A.

ALEC JACOBSON, University of Toronto and Adobe Research, Canada



▶0m40s

Fig. 1. The simple rig motions of 26 underwater sea creatures are augmented with our *real-time* secondary dynamics. The full scene of 330,563 mesh vertices and 1,293,625 tetrahedra runs at over 60 fps. Throughout our paper, the ▶ indicates a corresponding clip in the supplemental video.

We propose a reduced-space elastodynamic solver that is well suited for augmenting rigged character animations with secondary motion. At the core of our method is a novel deformation subspace based on Linear Blend Skinning

Authors' addresses: Otman Benckroun, otman.benckroun@mail.utoronto.ca, University of Toronto, Canada; Jiayi Eris Zhang, eriszhang@stanford.edu, Stanford University, U.S.A.; Siddhartha Chaudhuri, sidch@adobe.com, Adobe Research, India; Eitan Grinspun, eitan@grinspun.com, University of Toronto, Canada; Yi Zhou, yizho@adobe.com, Adobe Research, U.S.A.; Alec Jacobson, alecjacobson@gmail.com, University of Toronto and Adobe Research, Canada.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2023 Association for Computing Machinery.

0730-0301/2023/7-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

that overcomes many of the shortcomings prior subspace methods face. Our skinning subspace is parameterized entirely by a set of scalar weights, which we can obtain through a small, material-aware and rig-sensitive generalized eigenvalue problem. The resulting subspace can easily capture rotational motion and guarantees that the resulting simulation is rotation equivariant. We further propose a simple local-global solver for linear co-rotational elasticity and propose a clustering method to aggregate per-tetrahedra non-linear energetic quantities. The result is a compact simulation that is fully decoupled from the complexity of the mesh.

Additional Key Words and Phrases: Linear Blend Skinning, Secondary Motion, Complementary Dynamics

## ACM Reference Format:

Otman Benckroun, Jiayi Eris Zhang, Siddhartha Chaudhuri, Eitan Grinspun, Yi Zhou, and Alec Jacobson. 2023. Fast Complementary Dynamics via Skinning Eigenmodes. *ACM Trans. Graph.* 1, 1 (July 2023), 20 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>




Fig. 2. Our reduced complementary dynamics model can reproduce rich visual details at a fraction of the cost of the original method (10,000 vertices, 42,205 tets).

## 1 INTRODUCTION

Virtual reality, video games, and digital art increasingly make use of controllable animated characters. Such characters should provide *interactive* responses to user input in order to communicate the action and emotion of the moment. On the other hand, their motion must be rich with *realistic* visual details, which is what brings these characters to life.

The most widely used construct for authoring animations are deformation *rigs* mapping low-dimensional parameters (e.g., skeleton or cage positions) to static geometric deformations. Managing these rigs can quickly overwhelm an artist. Simple rigs are easy to animate, but their deformations lack detail; complex rigs provide rich, fine-grained detail, but they are daunting to manipulate.

Zhang et al. [2020] offer a way out of this dilemma; they employ a physics simulation whose role is to supplement rig motions with secondary dynamics that are orthogonal — in the algebraic sense — to the motion of the rig. The resulting *complementary dynamics* are exactly those motions not producible by the rig itself.

Unfortunately, complementary dynamics is poorly suited to interactive applications for two reasons. First, enforcing rig complementarity introduces significant computational overhead. Second, the runtime of the method grows with the mesh resolution. Adding secondary motion even on a modestly sized mesh requires computation that lags far from real-time rates. As shown in Fig. 2, a complementary dynamics animation on a modestly sized mesh of 10,000 vertices, 42,000 tetrahedra requires 3 seconds of compute time for each frame. Attempting to accelerate this method by embedding a high resolution display mesh inside of a coarse simulation (see Fig. 3) comes at the cost of visible cage artifacts.

To accelerate any high-dimensional optimization problem, a popular approach is to solve the problem in a low-dimensional, representative *subspace*. The *de-facto* subspace for elastodynamics in graphics is the one spanned by the first few eigenvectors of the elastic energy Hessian [Pentland and Williams 1989]. We call these eigenvectors *displacement modes* because—for elasticity—they correspond to the set of least energy-incurring infinitesimal displacements about the rest state.

Our use case exposes the pitfalls of this classical subspace. First, it is well known that this subspace does not represent rotational deformations, leading to warping or shearing artifacts [Choi and Ko 2005]. Second and (as we will show) *distinctly*, this subspace induces simulations lacking *rotation equivariance*. In a nutshell, a rotation equivariant optimization produces a rotated version of the same minimizer when the problem geometry is rotated. The absence of this

important property leads to frame-dependent artifacts. This is particularly noticeable in our application, where local rotations form a primary degree of freedom of the interactive rig. As the user interacts with the rig, the secondary physics exhibit non-physical dampened motion, and the mesh gets stuck in local minima as shown in Fig. 4 and Fig. 11.

To address these challenges, we propose *skinning eigenmodes* for reduced simulation. Inspired by linear blend skinning, the subspace spanned by our skinning modes yields rotation equivariant secondary elastodynamics. Indeed, we *prove* that it meets the necessary and sufficient conditions for doing so. Our subspace is fully parameterized by a compact set of *skinning weights*, which we derive through a physically motivated, material aware and simple to implement generalized eigenvalue problem.

The formulation of our skinning modes as the solution to an eigenvalue problem has many advantages. A user can easily explore the cost versus richness tradeoff of the resulting dynamics by simply truncating the eigenspace. We also benefit from a large array of work that aims at promoting different qualities from eigen problems, such as enforcing locality and sparsity in our modes [Brandt and Hildebrandt 2017; Houston 2017; Nasikun et al. 2018], or enforcing homogeneous linear equality constraints [Golub 1973]. In particular we benefit from the latter to make our skinning subspace orthogonal to the input rig-space. This effectively makes our modes rig-sensitive, allowing them to more efficiently capture the space of secondary motions available given an input rig.

However, skinning modes alone are insufficient for achieving realistic-looking real-time dynamics. This is because secondary elastodynamics look best with non-linear elastic models. Simulating such materials, even with a deformation subspace, *still* requires the computation of per-tetrahedron quantities [Choi and Ko 2005]. This once again ties our runtime complexity to the resolution of the mesh. To accommodate this, we adopt clustering [Jacobson et al. 2012], which approximates these per-tet quantities to per-cluster ones. We also extend Jacobson et al. [2012]’s local-global solver to approximate a co-rotational elastic potential. The result is an iterative solver that harmonizes well with our subspace and allows rich non-linear secondary motions in a simulation step that is *entirely* decoupled from the mesh resolution. Additionally, by virtue of being based on *linear* blend skinning, projecting from our low dimensional subspace to the full space can be done efficiently in the vertex shader.

Fig. 1 demonstrates our reduced elastodynamics augmenting a large scene with secondary motion in *real-time*. We further demonstrate the success of our method through a variety of comparisons to the original (offline) complementary dynamics and against alternative real-time acceleration methods. We additionally highlight how our skinning modes can be used for a wide range of standard deformation tasks (not just complementary dynamics). Finally, we show successful application of our real-time method to scenarios with rapidly changing rig-input, such as rigid-body enrichment, VR puppetry, and rigged character secondary effects.

## 2 RELATED WORK

There are many choices for accelerating (full-space) elastodynamic simulation, such as Position Based Dynamics [Müller et al. 2007],




Fig. 3. A coarse embedded simulation groups motion across vertices that are close in Euclidean space. Instead, our subspace groups motion across vertices that share elastic energy properties (cage constructed via [Sellán et al. 2021]).

Projective Dynamics [Bouaziz et al. 2014], or Mixed FEM [Trusty et al. 2022]. Elastodynamic simulations usually require the solution to a linear system of equations. A simple method of accelerating them is to just use a faster solver such as Multigrid methods [Liu et al. 2021], Newton and Quasi-Newton methods [Liu et al. 2017] and matrix pre-factorizations such as Cholesky decomposition. Yet another option is to reformulate the constitutive equations to allow for an efficient boundary-only discretization [James and Pai 1999; Sugimoto et al. 2022]. All these methods scale in complexity with the final mesh resolution and scale *at best* linearly with the complexity of the mesh. This puts an unnecessary burden on the user, forcing them to choose between fast elastodynamics with coarse meshes, or slow elastodynamics with fine meshes. To obtain sub-linear rates the simulation needs to be carried out in a low-dimensional representative subspace.

## 2.1 Subspace Simulation

The most popular subspace for most graphics tasks is the one composed of the first few eigenvectors of the energy Hessian [Pentland and Williams 1989], which we call displacement modes.

Unfortunately, displacement modes are not well suited for large displacements, and in particular they struggle to capture rotational motion [Barbič and James 2005; James and Pai 2002]. Much work over the last two decades aims at remedying this short-coming for subspace simulation, a goal we share. On top of this however, we also expose another desired property that displacement modes do not generally satisfy: they lead to simulations that are not *rotation equivariant*.

Barbič and James [2005] introduce modal derivatives, later extended to geometric modelling [Hildebrandt et al. 2011] and structural analysis [Duenser et al. 2022]. This methodology provides extra derivative modes, whose goal is to correct the primary subspace as it falls out of date with large deformations. While these methods provide vast improvements over traditional displacement modes, the added derivative modes do not represent rotational motion, nor do they ensure rotation equivariance in the resulting simulation.

To accommodate rotational motion, modes can be warped with best fit aggregate rotational motion [Choi and Ko 2005] or skinning motion [Kry et al. 2002]. In a similar vein, sub-structuring (also called domain decomposition) [Barbič and Zhao 2011; Kim and James 2011], separates a shape into independent regions, each with their own local linear subspace. The rotational motion is tracked externally and is




Fig. 4. Using a vanilla rig-complementary displacement subspace (red) can lead to kinks and local minima when a user rotates the rig. Our rig-complementary skinning eigenmodes (cyan) are rotation equivariant and are much better suited for accommodating this type of motion.

used to update the subspace for each region. Instead of updating the quality of the linear basis, Rotation Strain coordinates [Huang et al. 2011] attempt to *fix* the rotation-lacking motion of the subspace simulation at the end of each time-step via a non-linear projection step. Unfortunately all these methods require per timestep "fixes", which result in a costly simulation step that limits the richness of the dynamics available for real-time interaction.

The data driven neural subspaces of Zheng et al. [2021] show promise for real-time applications, but suffer from artifacts when applied on meshes they are not trained on. They also provide no guarantee of any energy conserving properties desired of an elasticity subspace.

Von Tycowicz [2013] expand traditional displacement modes with each of the  $d \times d$  entries of a linear map. Their result also leads to a rotation spanning subspace that guarantees simulation rotation equivariance. However, their subspace cannot be trivially made to accommodate homogeneous linear equality constraints, such as the one needed to impose rig-complementarity. As a result, this subspace is inefficient for real-time complementary dynamics.

## 2.2 Subspace Simulation via Skinning Modes

Linear blend skinning is a popular rigging method used to easily pose characters. It has also widely been used as a subspace for deformation, with prior works varying in how they compute their skinning weights. While prior methods acknowledge that skinning subspaces produce high quality rotational deformation, we further motivate our use of this subspace by identifying that it *guarantees* rotation equivariance, a key invariant for elastica that pure displacement subspaces break.

Rohmer et al. [2021] propose a method that prescribes secondary motion directly on the user-provided primary rig. They derive physics-like behavior that aims to mimic common squash and stretch motions. For more physically motivated inverse kinematics, Jacobson et al. [2012] use the user provided skinning rig as a subspace for minimizing an elasto-static energy. Rig-Space Physics [Hahn et al. 2012] aim to add secondary motion in a rig subspace, a similar goal to ours. These works all rely on an artist to specify the rig subspace for secondary motion themselves. Hahn et al. [2013] mitigates this by fitting the rig subspace to a target simulation using a least-squares solve. However, this still requires a previously constructed rig animation as input. In contrast, we propose a method that derives skinning weights entirely from the rest-pose geometry of our shape.




Fig. 5. Our skinning eigenmodes’ closure under rotations ensures that our subspace complementary dynamics simulations are rotation equivariant. A user experiences the same dynamics independent of the rig orientation. The same cannot be said for a complementary subspace built from displacement modes + modal derivatives (see Section 2.3).

One way to achieve such weights is by requiring them to be smooth. To this end, Gilles et al. [2011]; Wang et al. [2015] derive skinning weights as a solution to a Laplace and bi-Laplace equation respectively. Similarly, Lan et al. [2020, 2021] use bounded bi-harmonic weights [Jacobson et al. 2011] to create smooth, non-negative skinning weights. Instead, Brandt et al. [2018, 2019] compute weights with simple truncated radial basis functions to obtain similarly smooth and locally supported bases. While smoothness is an attractive quality for deformation, it is ill suited to represent motions for more complex materials with large heterogeneities. Faure et al. [2011] address this by sampling source points for their weight computation via a *compliance* (inverse stiffness) weighed sampling of their shape, for which they rely on a Voronoi tessellation of their shape. Additionally, just like von Tycowicz et al. [2013], it is unclear how to enforce the complementarity constraint on the construction of these skinning weights. In contrast, our weights are derived from a generalized eigenvalue problem on a weight-space Hessian, which allows them to reflect material properties and accommodate the rig-complementarity constraint without requiring additional discretization.

### 2.3 Simulations Embedded in Rotating Frames

A common solution for fixing some of the global rotational artifacts that occur in free-flying elastodynamic simulation, (as shown in Fig. 14 or Fig. 11) is to embed the elastodynamic simulation in a rotating frame [Terzopoulos and Witkin 1988]. Here, the elastic response is computed in a rest frame, while the rotation is tracked explicitly via a rigid body simulator [Terzopoulos and Witkin 1988]. The two simulations are then coupled with forces that arise out of angular momentum [James and Pai 2002], and the rotation is then used to transform the deformed rest state every time-step. Unfortunately, this approach doesn’t easily generalize for complementary dynamics since the rigid motions are not orthogonal to the rig, which is why the artifacts present in Fig. 5 would remain. For rigs that have a global rotation, one could use that as the rotational degree of freedom [James and Pai 2002], unfortunately most rigs do not directly have such a global rotation.

## 3 BACKGROUND: COMPLEMENTARY DYNAMICS

Complementary dynamics provides a methodology for augmenting rigged animations with the detailed elastodynamics [Zhang et al. 2020]. They split the total displacement field  $\mathbf{u} \in \mathbb{R}^{n(d)}$  for a mesh with  $n$  vertices in  $d$ -dimensional space into an artist prescribed component  $\mathbf{u}^r$ , and a physical component  $\mathbf{u}^c$ :

$$\mathbf{u} = \mathbf{u}^r + \mathbf{u}^c. \quad (1)$$

The artist-prescribed component is obtained from a rig, which takes as input low dimensional rig parameters  $\mathbf{p} \in \mathbb{R}^p$  that are exposed to the user for interactive manipulation, and maps them to high dimensional rig displacement  $\mathbf{u}^r$ :

$$\mathbf{u}^r = f_{\text{rig}}(\mathbf{p}). \quad (2)$$

The physical motion on the other hand is obtained from a physics simulation, which can be formulated as an energy minimization problem:

$$\mathbf{u}^c = \underset{\mathbf{u}^c}{\text{argmin}} E(\mathbf{u}^c + \mathbf{u}^r + \mathbf{x}_0) \quad \text{s.t. } \mathbf{J}^T \mathbf{u}^c = \mathbf{0}, \quad (3)$$

We introduce the rig Jacobian  $\mathbf{J} = \frac{\partial f_{\text{rig}}(\mathbf{p})}{\partial \mathbf{p}} \in \mathbb{R}^{n(d) \times p}$ . Note that complementary dynamics is completely agnostic to the elasto-dynamic energy used  $E(\cdot)$ . Where it differs from a regular elasto-dynamic energy minimization is the specification of the complementarity constraint  $\mathbf{J}^T \mathbf{u}^c = \mathbf{0}$ . This constraint enforces that the physical motion must not be in the space of motions producible by the rig,  $\text{Col}(\mathbf{J})$ . Zhang et al. [2020] enforce the complementarity constraint through Lagrange multipliers and the resulting energy is minimized using Newton’s method, requiring the frequent solve of the following KKT system:

$$\begin{bmatrix} \mathbf{H} & \mathbf{J} \\ \mathbf{J}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} d\mathbf{u}^c \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} -\mathbf{g} \\ \mathbf{0} \end{bmatrix}, \quad (4)$$

where  $\mathbf{H} \in \mathbb{R}^{n(d) \times n(d)}$  and  $\mathbf{g} \in \mathbb{R}^{n(d)}$  are the elasto-dynamic energy Hessian & gradient,  $d\mathbf{u}^c \in \mathbb{R}^{n(d)}$  is the Newton search direction, and  $\boldsymbol{\lambda} \in \mathbb{R}^p$  collects the Lagrange multipliers enforcing the complementarity constraint.

Iteratively solving this system is too expensive in real-time applications for two main reasons:

- (1) the system scales with mesh resolution, and
- (2) the constraints  $\mathbf{J}^T \mathbf{u}^c = \mathbf{0}$  are typically dense, even if  $\mathbf{H}$  is sparse.

## 4 SKINNING EIGENMODES

Our goal is to derive a suitable linear subspace so that full-space complementary displacements  $\mathbf{u}^c \in \mathbb{R}^{n(d)}$  may be approximated with a smaller  $m$ -dimensional linear subspace:

$$\mathbf{u}^c \approx \mathbf{B}\mathbf{z}, \quad (5)$$

where the columns of  $\mathbf{B} \in \mathbb{R}^{n(d) \times m}$  form the subspace basis, and the vector  $\mathbf{z} \in \mathbb{R}^m$  are the reduced degrees of freedom optimized at run-time.

To apply subspace reduction to the complementary dynamics problem (Eq. (3)), we would like our subspace  $\mathbf{B}$  to simultaneously: deal well with large (global and local) rotations, well-approximate the space of low-energy displacements, and accommodate the rig-complementarity constraints. We make use of a linear blend skinning subspace basis for

deformation  $\mathbf{B}_{\text{lbs}}$  [Gilles et al. 2011] and demonstrate in the following sections how we meet these three desirable criteria.

Linear blend skinning represents displacements as a weighted summation of  $m$  affine transformations applied to a shape's rest positions. The  $i$ th vertex on the shape is displaced via

$$\mathbf{u}_i = \sum_{b=1}^m w_{ib} T_b \mathbf{X}_i, \quad (6)$$

where  $w_{ib} \in \mathbb{R}$  is the weight of the  $b$ th transformation at vertex  $i$ ,  $T_b \in \mathbb{R}^{d \times (d+1)}$  is the  $b$ th transformation, and  $\mathbf{X}_i \in \mathbb{R}^{d+1}$  is the  $i$ th vertex's rest position in homogeneous coordinates. This equation is linear in  $T$  and so it follows that it may be rearranged so that the degrees of freedom in  $T$  are collected in a single vector  $\mathbf{z} = \text{vec}(T) \in \mathbb{R}^k$  with  $k = d(d+1)m$  and the weights  $w$  and rest positions  $\mathbf{X}$  form the columns of a matrix  $\mathbf{B}_{\text{lbs}} \in \mathbb{R}^{n(d) \times k}$ :

$$\mathbf{B}_{\text{lbs}} = \mathbf{I}_d \otimes ((\mathbf{1}_m^T \otimes \mathbf{X}) \odot (\mathbf{W} \otimes \mathbf{1}_{d+1}^T)), \quad (7)$$

where  $\mathbf{W} \in \mathbb{R}^{n \times m}$  is a matrix with columns collecting each transformation's weights and  $\mathbf{X} \in \mathbb{R}^{n \times (d+1)}$  collects homogeneous rest positions in rows. Since the rest positions are generally given, the only variables in our subspace design are the weights  $\mathbf{W}$ . We now propose a method for choosing  $\mathbf{W}$  to ensure that weights span low-energy motions *and* satisfy the complementarity constraints by construction. We defer discussion of how our choice of linear blend skinning subspace directly ensures good rotational properties (see Sec. 5).

Our first step follows the process for standard modal subspaces. We approximate our elastodynamic energy with a Taylor expansion about the rest state truncated to second order terms,

$$\begin{aligned} E(\mathbf{u} + \mathbf{x}_0) &= E_0 + \mathbf{u}^T \mathbf{g}_0 + \frac{1}{2} \mathbf{u}^T \mathbf{H}_0 \mathbf{u} + O(\|\mathbf{u}\|^3), \\ E(\mathbf{u} + \mathbf{x}_0) &\approx \frac{1}{2} \mathbf{u}^T \mathbf{H} \mathbf{u}, \end{aligned} \quad (8)$$

where we have dropped the subscript for the elastic energy Hessian in the second line for readability. Without loss of generality, we have assumed zero elastic energy and vanishing elastic forces at rest ( $E_0 = 0$ ,  $\mathbf{g}_0 = \mathbf{0}$ ).

We arrive at a standard modal subspace by adding a (mass-) orthogonality constraint  $\mathbf{B}^T \mathbf{M} \mathbf{B} = \mathbf{I}$ ; substituting the subspace  $\mathbf{u} = \mathbf{B} \mathbf{z}$ ; assuming  $\mathbf{z} \sim \mathcal{D}$  are sampled from an as-of-yet arbitrary distribution, and minimize the expected value of the energy over  $\mathcal{B}$ :

$$\mathbf{B}_{\text{disp}} = \underset{\mathbf{B}^T \mathbf{M} \mathbf{B} = \mathbf{I}}{\text{argmin}} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\mathbf{z}^T \mathbf{B}^T \mathbf{H} \mathbf{B} \mathbf{z}] \quad (9)$$

$$= \underset{\mathbf{B}^T \mathbf{M} \mathbf{B} = \mathbf{I}}{\text{argmin}} \text{tr}(\mathbf{B}^T \mathbf{H} \mathbf{B} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\mathbf{z} \mathbf{z}^T]). \quad (10)$$

We further assume that  $\mathbf{z}$  are independent and identically distributed (i.i.d.) samples of a normal distribution, then  $\mathbb{E}_{\mathbf{z} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)} [\mathbf{z} \mathbf{z}^T] = \mathbf{I}$ , and

$$\mathbf{B}_{\text{disp}} = \underset{\mathbf{B}^T \mathbf{M} \mathbf{B} = \mathbf{I}}{\text{argmin}} \text{tr}(\mathbf{B}^T \mathbf{H} \mathbf{B}). \quad (11)$$

The optimal  $\mathbf{B}_{\text{disp}}$  may be found relatively efficiently with a generalized eigenvalue solver that supports large sparse matrices. The columns of  $\mathbf{B}_{\text{disp}}$  can be directly interpreted as minimal-energy *displacement* eigenmodes.

Our *skinning* eigenmodes follows a similar derivation but we replace the optimization over  $\mathbf{B}$  with an optimization over the weights  $\mathbf{W}$ . While Eq. (7) may appear to define  $\mathbf{B}_{\text{lbs}}$  as a complicated function of  $\mathbf{W}$ , it is *linear* in and separable over the columns of  $\mathbf{W}$ . Thus, we may rewrite it as

$$\mathbf{B}_{\text{lbs}} = [\mathbf{A}_{i,j} \quad \dots \quad \mathbf{A}_{d,(d+1)}] (\mathbf{I}_{d(d+1)} \otimes \mathbf{W}), \quad (12)$$

where we introduce  $\mathbf{A}_{i,j} \in \mathbb{R}^{n(d) \times n}$ , our weight-space skinning Jacobians. These map contributions of each weight for all  $d(d+1)$  affine parameters to the final skinning Jacobian.

We derive  $\mathbf{A}_{i,j}$  for the  $d = 2$  and  $d = 3$  in Appendix B.1, but for clarity show the result for  $d = 3$  here:

$$\begin{aligned} \mathbf{A}_{1,1} &= P_x \bar{X} & \mathbf{A}_{1,2} &= P_x \bar{Y} & \mathbf{A}_{1,3} &= P_x \bar{Z} & \mathbf{A}_{1,4} &= P_x \\ \mathbf{A}_{2,1} &= P_y \bar{X} & \mathbf{A}_{2,2} &= P_y \bar{Y} & \mathbf{A}_{2,3} &= P_y \bar{Z} & \mathbf{A}_{2,4} &= P_y \\ \mathbf{A}_{3,1} &= P_z \bar{X} & \mathbf{A}_{3,2} &= P_z \bar{Y} & \mathbf{A}_{3,3} &= P_z \bar{Z} & \mathbf{A}_{3,4} &= P_z \end{aligned} \quad (13)$$

where the  $\mathbf{P}_* \in \mathbb{R}^{3n \times n}$  selection matrices concatenate to form the identity matrix  $\mathbf{I}_{3n} = [P_x P_y P_z]$  and  $\bar{X}, \bar{Y}, \bar{Z} \in \mathbb{R}^{n \times n}$  are diagonal matrices containing the the  $x, y$  and  $z$  rest position values.

Following a similar procedure as before, we add a weight space orthogonality constraint  $\mathbf{W}^T \mathbf{M}_w \mathbf{W} = \mathbf{I}$  and assume a generic distribution  $\mathcal{D}$  on our sampling of  $\mathbf{z} \sim \mathcal{D}$  to obtain

$$\mathbf{W} = \underset{\mathbf{W}^T \mathbf{M}_w \mathbf{W} = \mathbf{I}}{\text{argmin}} \text{tr}(\mathbf{B}_{\text{lbs}}^T \mathbf{H} \mathbf{B}_{\text{lbs}} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\mathbf{z} \mathbf{z}^T]). \quad (14)$$

We now need to make assumptions on the distribution of  $\mathbf{z}$ , as these now correspond to flattened affine matrix parameters and so have some structure to their distribution. Specifically we assume that parameters belonging to different affine matrices are i.i.d. with respect to each other, but generally allow for intradependence between parameters belonging to the same affine matrix, as measured by the covariance matrix  $\mathbf{C} \in \mathbb{R}^{d(d+1) \times d(d+1)}$ .

$$\mathbf{W} = \underset{\mathbf{W}^T \mathbf{M}_w \mathbf{W} = \mathbf{I}}{\text{argmin}} \text{tr} \left( (\mathbf{I}_{d(d+1)} \otimes \mathbf{W})^T \mathbf{A}^T \mathbf{H} \mathbf{A} (\mathbf{I}_{d(d+1)} \otimes \mathbf{W}) (\mathbf{I}_m \otimes \mathbf{C}) \right).$$

Expanding out all the Kronecker products and leveraging that the trace is just a sum of diagonal entries :

$$\mathbf{W} = \underset{\mathbf{W}^T \mathbf{M}_w \mathbf{W} = \mathbf{I}}{\text{argmin}} \text{tr} \left( \mathbf{W}^T \left( \sum_i^{d(d+1)} \sum_j^{d(d+1)} (\mathbf{A}_i^T \mathbf{H} \mathbf{A}_j) c_{ij} \right) \mathbf{W} \right). \quad (15)$$

Leading to the weight-space optimization:

$$\mathbf{W} = \underset{\mathbf{W}^T \mathbf{M}_w \mathbf{W} = \mathbf{I}}{\text{argmin}} \text{tr}(\mathbf{W}^T \mathbf{H}_w \mathbf{W}) \quad (16)$$

where  $\mathbf{M}_w \in \mathbb{R}^{n \times n}$  is the weight-space mass matrix (we use the diagonal lumped mass matrix) and where we call  $\mathbf{H}_w \in \mathbb{R}^{n \times n}$  the *weight-space* elastic energy Hessian.

It is important to note this is overly determined for  $\mathbf{W}$ ; The same set of weights are used to specify  $d(d+1)$  different types of affine motions: scales, shears and translations. As a result, the set of weights that leads to optimal translations may not be the same set of weights that lead to optimal scales or shears. We can change which of these parameters we prioritize by modifying our affine parameter covariance matrix  $\mathbf{C}$ .

We choose to prioritize translations, neglecting shears and scales entirely, which are poorly suited for deformation subspaces. The logic is that shears and scales are origin-dependent. This leads the



Fig. 6. We generate a linear blend skinning subspace for secondary motion parameterized by a set of skinning weights. Each weight  $i$  shown is independently normalized to lie between  $[-1, 1] \text{abs}(\mathbf{W}_i)$  and centered around 0. (Top) Weights generated by solving the unconstrained generalized eigenvalue problem on a weight-space elasticity Hessian. (Bottom) Secondary skinning weights that satisfy the weight-space complementarity constraint and are orthogonal to our rig space. These are naturally rig-aware, leading to higher frequency motion.

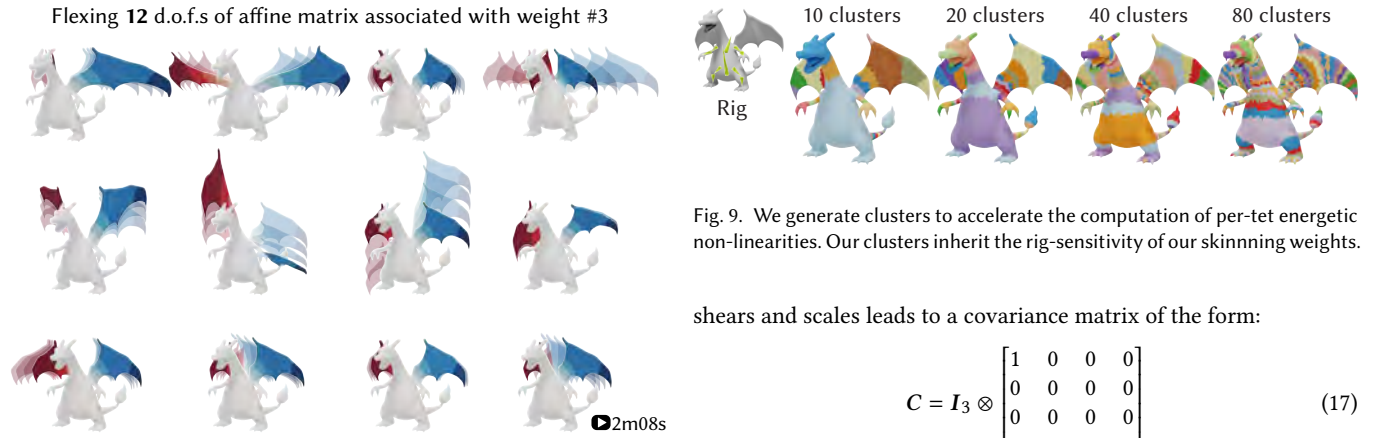


Fig. 7. One secondary linear blend skinning weight could produce 12 different motions, corresponding to 12 d.o.f.s of an affine matrix. We showcase this by flexing those associated with weight #3.




Fig. 8. Prioritizing scaling and shearing (middle left and middle right) provides weights that are unnaturally centered around the origin. For this reason, we prioritize translations (right).

optimization in Eq. (16) to see vertices far from the origin as *stiffer* than vertices that are close to it, resulting in weights that are unnaturally concentrated around the origin, and decay far away from it as shown in Fig. 8.

For  $d = 3$ , taking i.i.d. samples from the standard normal distribution of each of the three translation parameters, while neglecting


shears and scales leads to a covariance matrix of the form:

$$C = I_3 \otimes \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (17)$$

which very conveniently leads to a simplified weight space Hessian:

$$\mathbf{H}_w = \mathbf{P}_x^T \mathbf{H} \mathbf{P}_x + \mathbf{P}_y^T \mathbf{H} \mathbf{P}_y + \mathbf{P}_z^T \mathbf{H} \mathbf{P}_z. \quad (18)$$

The inset, unburdened by notation, more clearly shows the simplicity of deriving this final weight-space Hessian; just take the diagonal blocks for each dimension of the Hessian and sum them up. For co-rotational elasticity with homogeneous materials,  $\mathbf{H}_w$  is proportional to the mesh’s cotangent Laplacian matrix. Whereas heterogeneous materials distributions affect  $\mathbf{H}_w$  non-trivially and thus also our optimal weights  $\mathbf{W}$ .



With these matrices defined, our optimal skinning eigenmodes are solutions to Eq. (16), found efficiently via a generalized eigenvalue solver. Each individual skinning eigenmode — as a linear blend skinning weight — corresponds to  $d(d + 1)$  degrees of freedom and may be used to generate  $d(d + 1)$  different motions, as shown in Fig. 7 for  $d = 3$ .

**4.0.1 Weight Space Complementarity Constraint.** At run-time, our secondary-effect displacements should satisfy  $J^T \mathbf{u}^c = \mathbf{0}$ , where recall  $J \in \mathbb{R}^{3n \times p}$  is the current rig Jacobian. In our subspace, this becomes  $J^T \mathbf{B}_{\text{lbs}} \mathbf{z} = \mathbf{0}$ . Without knowledge of  $J$  a priori, our optimized skinning eigenmodes will, in general, not admit non-trivial solutions. Even if they did, enforcing this constraint at run-time leads to a more difficult constrained optimization problem. Fortunately, our formulation above as a generalized eigenvalue problem allows us to easily add constraints to our skinning weights, thus ensuring that our modes admit non-trivial solutions but also implicitly satisfy the constraint allowing us to remove it entirely at run-time.

Zhang et al. [2020] define  $f_{\text{rig}}(\mathbf{p})$  generically. For real-time applications, we will assume that  $f_{\text{rig}}$  is linear (single affine handle, linear blend skinning, blendshapes, etc.) and thus has a constant rig Jacobian  $J$ . Given  $J$ , the constraint we need to add is

$$J^T \mathbf{B}_{\text{lbs}} = \mathbf{0} \quad (19)$$

To express this in terms of  $\mathbf{W}$ , we can again make use of our weight-space skinning Jacobian matrices from Eq. (13) (not to be confused with  $J$ ) and expand the constraint to act on each weight. This leads to a series of constraints that our weights need to satisfy:

$$J^T \mathbf{A}_{i,j} \mathbf{W} = \mathbf{0} \in \mathbb{R}^{p \times m} \quad \forall i \in \{1, \dots, d\}, j \in \{1, \dots, d+1\} \quad (20)$$

We can stack all our constraint matrices  $J^T \mathbf{A}_{i,j}$ :

$$\begin{bmatrix} J^T \mathbf{A}_{1,1} \\ \vdots \\ J^T \mathbf{A}_{d,d+1} \end{bmatrix} \mathbf{W} = \mathbf{J}_w \mathbf{W} = \mathbf{0} \in \mathbb{R}^{p(d)(d+1) \times m} \quad (21)$$

where we call  $\mathbf{J}_w \in \mathbb{R}^{p(d)(d+1) \times n}$  our weight-space complementarity constraint matrix.

We can incorporate this constraint in a standard generalized eigenvalue problem by solving instead

$$\begin{bmatrix} \mathbf{H}_w & \mathbf{J}_w^T \\ \mathbf{J}_w & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{W} \\ \boldsymbol{\mu} \end{bmatrix} = \begin{bmatrix} \mathbf{M}_w & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{W} \\ \boldsymbol{\mu} \end{bmatrix} \Lambda. \quad (22)$$

Fig. 6 shows how our derived skinning weights change to accommodate the rig-complementarity constraint.

## 5 BUT WHAT ABOUT ROTATIONS?

Interesting elastodynamic effects exhibit rotations: both global, where the entire shape rotates in space, and local, where part or parts of the shape rotate relative to the rest/each other. Rotations are notoriously difficult for previous linear subspaces. For example, it is well known that displacement modes ( $\mathbf{B}_{\text{disp}}$  defined in Eq. (11)) struggle to represent local rotational motion (see Fig. 12 and Fig. 13 (Left)) [Barbič and James 2005; Barbič and Zhao 2011; Choi and Ko 2005; Huang et al. 2011]).

Our use case reveals that issues with rotations go beyond this and can be more insidious. In the following discussion, we assume that the elastic potential  $E$  is rotation invariant. That is,  $E(\mathbf{u} + \mathbf{x}_0) = E((\mathbf{R} \otimes \mathbf{I})(\mathbf{u} + \mathbf{x}_0))$ , where multiplying by  $(\mathbf{R} \otimes \mathbf{I}) \in \mathbb{R}^{n(d) \times n(d)}$  applies the same rotation  $\mathbf{R} \in SO(d)$  to all vertices.

## 5.1 Rotation Spanning vs. Closure Under Rotations

Rotational problems may be categorized into two separate issues.

First, does a given subspace span rotations? By global rotation spanning, we mean there always exists some subspace parameters to reproduce any rotational displacement. If  $\mathbf{x}_0$  are the rest positions then

$$\exists \mathbf{z} \text{ such that } (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}_0 - \mathbf{x}_0 = \mathbf{B} \mathbf{z} \quad \forall \mathbf{R} \in SO(d). \quad (23)$$

For free-flight objects, failing to span global rotations means the subspace will unnaturally deform in an attempt to minimize  $E$  rather than rotate. Unfortunately, this problem also occurs locally, too. For example, if the arms of a character bend in opposite directions, failure to span these local rotations will disturb (by introducing local shears and stretches to attempt to minimize  $E$ ) or prevent (by the minimization of  $E$  detesting such scales and shears) the desired deformation. For example, Barbič and James [2005] emphasize how displacement modes  $\mathbf{B}_{\text{disp}}$  induce scaling and shearing artifacts when approximating rotations and bending deformations.

Second, does a given subspace induce a rotationally equivariant simulation? Treat the simulation as a map from problem specification parameters (e.g., forces, rig displacements, rest positions) to optimal (full-space) displacements. Rotation equivariance means that any rotated version of the problem results in a correspondingly rotated solution:

$$\forall \mathbf{R} \in SO(d), \quad \forall \mathbf{x} \in \mathbb{R}^{n(d)}, \quad \mathbf{B} \underset{z}{\text{argmin}} E(\mathbf{B} \mathbf{z} + (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}) = (\mathbf{R} \otimes \mathbf{I}) \mathbf{B} \underset{z}{\text{argmin}} E(\mathbf{B} \mathbf{z} + \mathbf{x}), \quad (24)$$


where — without loss of generality — we lump problem specification parameters into the vector  $\mathbf{x} \in \mathbb{R}^{n(d)}$ .

A subspace simulation lacking rotation equivariance may experience unpredictably different deformations under rotations. This is especially problematic in a complementary dynamics setting where the entire object or large subpart may rotate due to the user rig. Users will expect analogously rotated secondary effects and be surprised by behavior that depends on global or local rotations coming from the rig.

Fig. 14 shows how this shortcoming expresses itself as overly energetic deformation, while Fig. 11 showcases some of the kinky local minima that can easily arise under simple rotations. The root of this problem is shown didactically in the inset: a single displacement mode describes a completely different type of motion if its underlying shape rotates.

Building on this intuition, we prove (see App. A.1) that a linear subspace simulation is rotation equivariant if and only if the subspace basis is *closed under rotations*:

$$\forall \mathbf{R} \in SO(d) \text{ and } \mathbf{z} \in \mathbb{R}^m \exists \mathbf{w} \in \mathbb{R}^m, \text{ such that } (\mathbf{R} \otimes \mathbf{I}) \mathbf{B} \mathbf{z} = \mathbf{B} \mathbf{w}. \quad (25)$$



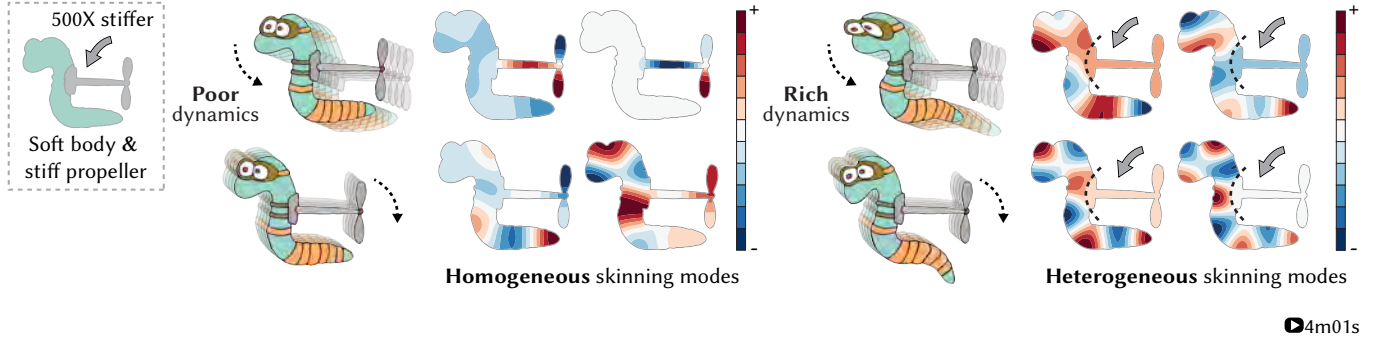


Fig. 10. A material-aware subspace more efficiently captures the space of motions available to our simulation. This directly leads to richer dynamics.




Fig. 11. We compute the subspace at rest (top-left). A user rotates the mesh and perturbs the system with a random initial deformation. Using displacement modes creates jarring local minimum artefacts in a rotated frame. Our skinning modes find the global minimum effortlessly, obtaining the same rest state than if we had embedded the simulation in a rigid frame.

## 5.2 Displacement Modes Simulations Are Fragile Under Rotations

Displacement modes ( $B_{\text{disp}}$  defined in Eq. (11)) [Pentland and Williams 1989] and many of their improvements (e.g., [Barbič and James 2005]) are neither rotation spanning nor closed under rotations. Rotations are a full-spectrum displacement, so any (reasonable) truncated elastic eigenspace will fail to span arbitrary global rotations (see Fig. 13 (Left)). While — as discussed above — global rotation spanning has an easy fix, much effort has been made to improve local rotation spanning such as data-driven PCA bases [Kry et al. 2002], modal derivatives [Barbič and James 2005], sub-structuring [Barbič and Zhao 2011], or splitting the simulation into rigid and deformable components [Terzopoulos and Witkin 1988].

Nevertheless, large local rotations may still be problematic (see Fig. 12). Displacement modes — except if truncated to just null modes or completely non-truncated — are not closed under rotations (see counterexamples in Fig. 13 (Right) and Fig. 11).

## 5.3 Skinning Eigenmodes Are Robust to Rotations

In contrast, skinning eigenmodes are both rotation spanning and closed under rotations (see Fig. 13). When the complementarity constraint is absent, the first skinning eigenmode will be a constant function thus spanning all affine motions including rotations. When used




Fig. 12. No matter how hard a user tries, the eyes of this reduced elastodynamic alien will never bend when using a small displacement mode subspace (middle). Our skinning subspace (right) enables the rotational motion. Both results use 60 degrees of freedom.

for fast complementary dynamics, the rig typically contains global rotations so we explicitly (and purposefully) avoid global rotation spanning. We do still want and indeed observe local rotation spanning (see Fig. 12). We can easily show that the linear blend skinning — and thus also skinning subspaces — are closed under rotations. Given some rotation  $R \in \mathcal{SO}(d)$ , rotating linear blend skinning's output is equivalent to rotating all of the input transformations:

$$R \sum_{b=1}^m w_{ib} T_b X_i = \sum_{b=1}^m w_{ib} R T_b X_i. \quad (26)$$

Any rotation of its output is producible by its input, as required for a rotation equivariant subspace simulation. This fact was similarly utilized in previous works albeit in different settings [Jacobson et al. 2011; Langer and Seidel 2008; Wang et al. 2015]. There are many prior subspaces [Faure et al. 2011; Gilles et al. 2011; von Tycowicz et al. 2013; Wang et al. 2015] that do not explicitly mention rotation equivariance as a criterion for the subspace simulation. In hindsight, leveraging the machinery of Appendix 34, we can see that since these prior subspaces are closed under rotations, those methods also maintain a rotation equivariant simulation.

## 6 CLUSTERING

For many rich rotational dynamic properties, an elastic energy usually requires the use of per-tetrahedron non-linear operations [Kim and Eberle 2020]. Despite our use of a subspace for our displacement degrees of freedom, these energies *still* demand per-tet computation.



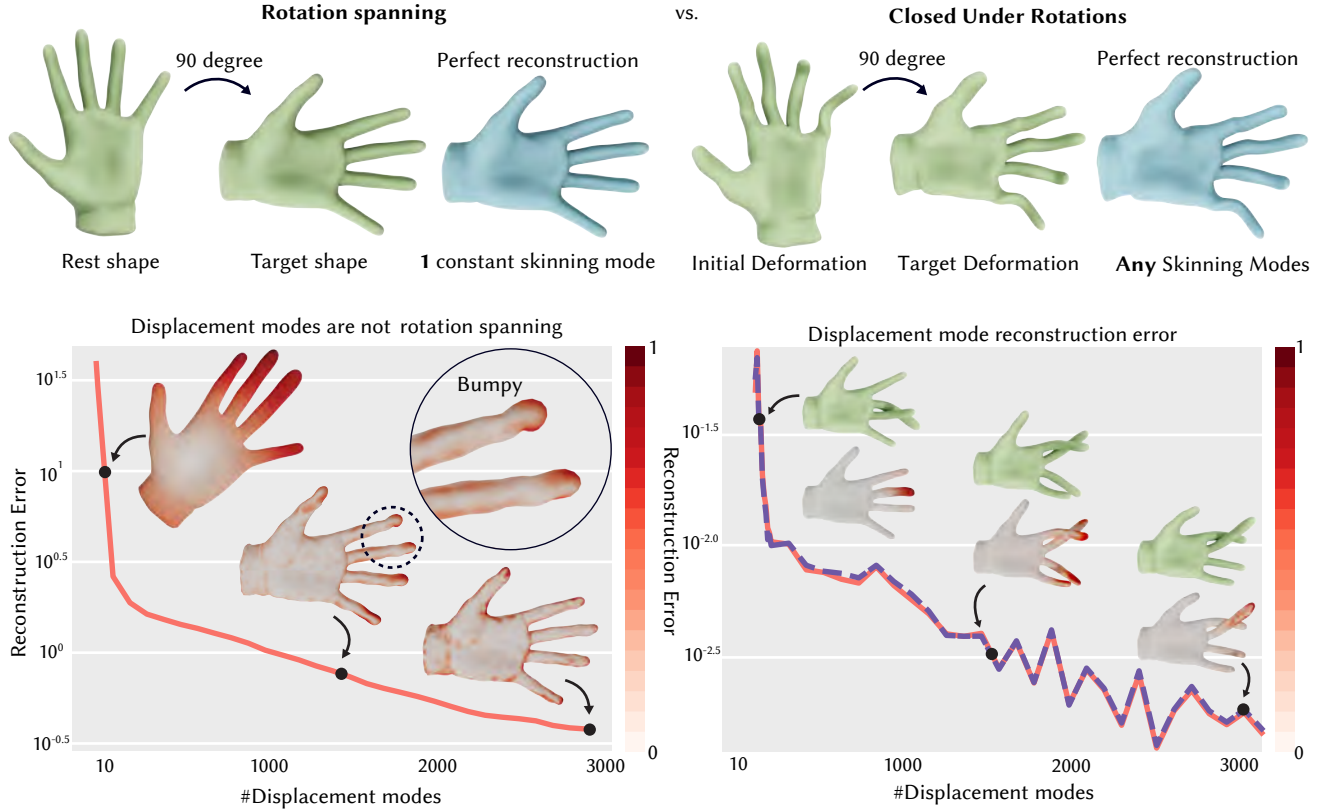


Fig. 13. **Rotation Spanning vs. Closure Under Rotations.** Given some initial shape, we optimize for optimal displacements that minimize the squared distance between each vertex position and its rotated target. (Left) Our skinning modes are **rotation spanning**, as modulated by the skinning weights. With a single constant skinning weight, we can perfectly reconstruct (by least squares projection) any rotation of the rest shape. Displacement modes do not span rotational motion, even with excessively abundant modes. (Right-red) Our skinning modes are **closed under rotations**, so any deformation in the span of those can be reconstructed (by least squares projection) under the same set of modes even if the mesh is arbitrarily rotated by a user. The same cannot be said for even an impractically large number of displacement modes even if augmented with affine degrees of freedom (right-purple).

Our simulations are not yet truly decoupled from the resolution of the mesh.

Many methods in the past sidestep this issue by making use of *curvature* [An et al. 2008; Barbič and James 2005; Teng et al. 2014], where the deformation quantities at each tetrahedra are estimated as a sum of weighed contributions from a sparse set of pre-determined sample tetrahedra. These samples and their interpolation weight are usually obtained through a data-fitting phase. We instead opt for a clustering [Jacobson et al. 2012] scheme, which estimates these non-linearities through  $r$  clusters of tetrahedra. Instead of requiring training data, we compute our clusters via a deformation prior, obtaining them via a  $k$ -means clustering [Arthur and Vassilvitskii 2007] of our subspace secondary skinning weights  $\mathbf{W}$ . For a mesh with  $t$  tetrahedra, we use  $r$  clusters to build the grouping matrix  $\mathbf{G} \in \mathbb{R}^{r \times t}$ :

$$G_{ij} = \begin{cases} \frac{V_j}{\sum_{q \in C_i} V_q} & \text{if } j \in C_i \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

where  $V_j$  is the volume of the  $j$ -th tetrahedron, while  $C_i$  contains the indices of the tetrahedra belonging to cluster  $i$ . This grouping

matrix computes cluster quantities via a mass weighed averaging of its member tetrahedra. We use these grouping matrices whenever we require per-tet evaluations in our solver as outlined in Sec. 7.

We build our clusters via a  $k$ -means++ [Arthur and Vassilvitskii 2007], using features obtained from our secondary weights. Specifically, we average our vertex weights onto each tetrahedra, then we scale each of our weights  $w_i$  by the inverse-squared of its associated eigenvalue  $\lambda_i$ , obtained in Eq. (22). Fig. 9 shows clustering results for a sample mesh with different rigs. Because our skinning subspace is the only feature we use for building our clusters, our clusters inherit many of the properties of the skinning modes, including being material-sensitive, rig-sensitive. This allows our clusters to provide a higher resolution in parts of the mesh more likely to exhibit secondary motion. Because our skinning modes can be global, this results in clusters that could also be global. We detect these in a post processing step and *split* clusters into their individual separate components.




Fig. 14. Simulations making use of displacement modes (left) are not rotation equivariant. Augmenting displacement modes with modal derivatives (middle), helps in representing higher quality deformations, but still create different high energy deformations under rotations. Our skinning modes (right) allow us to capture both local rotational motion and maintain rotation equivariance and lead to the lowest energy even without explicitly tracking a rotating frame [Terzopoulos and Witkin 1988] each timestep (bottom).

## 7 LOCAL-GLOBAL SOLVER

We make use of a simple local-global solver for the minimization of our elastodynamic energy. This solver leverages our reduced space degrees of freedom and our clusters to *never* make use of any full-space operations. This solver allows us to avoid recomputing expensive quadratic energy Hessians each timestep while still accommodating frequently desired elastic non-linearities, such as rotations.

### 7.1 Full-Space Local-Global Solver

We assume energies that can be decomposed as

$$E(\mathbf{u}) = \Psi(\mathbf{u}) + \Phi(\mathbf{u}),$$

where  $\Psi(\mathbf{u})$  is quadratic in  $\mathbf{u}$  and  $\Phi$  is a non-linear in  $\mathbf{u}$  and piecewise constant over tets. We then need to identify the source of the non-linearities in  $\Phi$ ,  $\mathbf{R}$ , as auxiliary degrees of freedom

$$\mathbf{u}, \mathbf{R} = \underset{\mathbf{u}, \mathbf{R}}{\operatorname{argmin}} \Psi(\mathbf{u}) + \Phi(\mathbf{u}, \mathbf{R}).$$

We split elastodynamic optimization into two steps [Sorkine and Alexa 2007]. The local step optimizes for the non-linear sources  $\mathbf{R}$  in our energy: holding the primary degrees of freedom fixed,

$$\mathbf{R}_i = \underset{\mathbf{R}}{\operatorname{argmin}} \Phi(\mathbf{u}_{i-1}, \mathbf{R}). \quad (28)$$

The global step optimizes for the primary degrees of freedom, holding the auxiliary ones fixed:

$$\mathbf{u}_i = \underset{\mathbf{u}}{\operatorname{argmin}} \Psi(\mathbf{u}) + \Phi(\mathbf{u}, \mathbf{R}_i). \quad (29)$$

These two steps are repeated iteratively until convergence..

### 7.2 Hyper-Reduced Local-Global Solver

To ensure our solver never performs any full space operations, we make use of our subspace  $\mathbf{u} \approx \mathbf{Bz}$  and our clusters  $\Phi \approx \tilde{\Phi}$ , where the tilde denotes clustered quantities. While before  $\Phi$  was piecewise

constant over *tets*, it is now piecewise constant over *clusters*, where clustered physical quantities, (like mass, Lamé parameters, deformation gradients, etc.) have been averaged over all the tets in each cluster. Our hyper-reduced energy becomes:

$$E(\mathbf{u}) = \Psi(\mathbf{z}) + \tilde{\Phi}(\mathbf{z}).$$

Exposing the per-cluster source of non-linearities as their own auxiliary degrees of freedom  $\tilde{\mathbf{R}}$ :

$$E(\mathbf{z}) = \Psi(\mathbf{z}) + \tilde{\Phi}(\mathbf{z}, \tilde{\mathbf{R}})$$

Our hyper-reduced local step becomes

$$\tilde{\mathbf{R}}_i = \underset{\tilde{\mathbf{R}}}{\operatorname{argmin}} \tilde{\Phi}(\mathbf{z}_{i-1}, \tilde{\mathbf{R}}), \quad (30)$$

Because the non-linearity is piece-wise constant over clusters, this per-cluster optimization can additionally be solved in parallel over all clusters, providing further acceleration. while the hyper-reduced global step becomes

$$\mathbf{z}_i = \underset{\mathbf{z}}{\operatorname{argmin}} \Psi(\mathbf{z}) + \tilde{\Phi}(\mathbf{z}, \tilde{\mathbf{R}}_i). \quad (31)$$

Depending on the specific energy, different optimization schemes need to be adopted for both the local step and the global step. The only criteria we need to maintain on any of these solvers is that they *never* perform any full-space operations. Appendix Sec. C.1 and Sec. C.2 show how to implement our reduced local-global solver for elastic energies such as ARAP [Chao et al. 2010] and Linear Co-Rotational Elasticity [McAdams et al. 2011].

## 8 IMPLEMENTATION

Most of our implementation is in C++ with some calls to Matlab routines for building the subspace. All timings reported were performed on a Dell XPS i9-12900HK with a 2.90 GHz processor and 64GB of RAM, equipped with an NVIDIA GeForce 3050Ti graphics card.

### 8.1 Subspace Construction

Algorithm 1 provides pseudocode for the construction of our subspace. We require as input the rest state, an elastic energy Hessian, and a user-defined homogeneous linear equality constraint matrix  $\mathcal{J}$ , in our case the complementarity constraint matrix. So far, for conciseness, we have stated that our complementarity constraint matrix was equal to the rig jacobian  $\mathcal{J} = \mathbf{J}$ . However, as described by [Zhang et al. 2020], to make our complementarity constraint sensitive to the mesh resolution, it is also weighed by a mass-matrix. Additionally the constraint matrix is weighed by a momentum-leaking matrix  $\mathbf{D}$ . This diagonal matrix with entries ranging from 0 to 1 specifies where momentum can leak from the rig to the mesh. By default, we compute these diagonal entries from a fast diffusion of the surface to the interior, and renormalize the values to lie between 0 and 1, which easily gives us standard *follow-through* and *anticipation* effects [Zhang et al. 2020]. For additional control, a user can also customize and easily provide their own scalar field designed by a method of their choice. The final complementarity constraint matrix is  $\mathcal{J} = \mathbf{D}\mathbf{M}\mathbf{J}$ .

For the GEVP solver, we use Matlab's eigs. We also remove redundant rows from our weight space complementarity constraint quickly using Matlab's `ref`.

**Algorithm 1:** Subspace Construction : given rest positions  $X$ , tet indeces  $T$ , an elastic energy hessian  $H$  a complementarity constraint matrix  $\mathcal{J}$ , compute a subspace composed of  $m$  skinning weights  $W$  and  $r$  cluster labels  $l$  for our tets.

```

Function computeSubspace( $X, T, H, \mathcal{J}, m, r$ ):
   $H_w \leftarrow$  weightSpaceHessian( $H$ )
   $M_w \leftarrow$  massmatrix( $X, T$ )
   $J_w \leftarrow$  weightSpaceConstraint( $\mathcal{J}, X$ )
   $\mathcal{H} \leftarrow \begin{bmatrix} H_w & J_w^T \\ J_w & 0 \end{bmatrix}$   $\mathcal{M} \leftarrow \begin{bmatrix} M_w & 0 \\ 0 & 0 \end{bmatrix}$ 
   $W, V \leftarrow$  eigs( $\mathcal{H}, \mathcal{M}, m$ )
   $W_t \leftarrow$  averageOntoTets( $WV^{-2}, T$ )
   $l \leftarrow$  kmeans( $W_t, r$ )
   $l \leftarrow$  splitComponents( $T, l$ )
  return  $W, l$ 

Function weightSpaceConstraint( $\mathcal{J}, X$ ):
   $A \leftarrow []$ 
  for  $i \leftarrow 0$  to  $\dim$  do
    for  $j \leftarrow 0$  to  $\dim + 1$  do
       $A_{i,j} \leftarrow$  weightSkinningJacobian( $X, i, j$ ) Eq. (21)
       $A \leftarrow \begin{bmatrix} A \\ \mathcal{J}^T A_{i,j} \end{bmatrix}$ 
   $J_w \leftarrow$  removeRedundantRows( $A$ )
  return  $J_w$ 

```

## 8.2 Simulation Step

Algorithm 2 provides an overview for the simulation step of our proposed method. We color in blue matrix products that can be computed *once* at the start of the simulation, cached and then called at run-time. We color in red vectors that do not change throughout the multiple iterations of a single local-global solve. Products with these red vectors can be computed once at the start of the timestep step and reused throughout the timestep. To solve the linear system for the Quasi-Newton search direction, we make use of a precomputed Cholesky Factorization. In the case of vanishing Poisson ratio, where Co-Rotational Elasticity becomes ARAP, we can safely remove the line search.

The individual constituents of these matrices are defined as follows:

- $B = B_{lbs}$  is the linear blend skinning Jacobian matrix for our secondary weights  $W$ , as computed by Eq. (7).
- $G_9 \in \mathbb{R}^{9r \times 9t}$  is our exploded cluster grouping matrix, computable from our per-cluster labels  $l$  and our rest geometry, that performs a mass-weighted averaging of the 9 deformation gradient quantities for all tetrahedra belonging to a cluster.
- $K \in \mathbb{R}^{9t \times 3n}$  is a standard vector gradient operator, which maps displacements of vertices to deformation gradients on tetrahedra [Kim and Eberle 2020].
- $L = K^T \mathcal{U} V K \in \mathbb{R}^{3n \times 3n}$  is the heterogeneous Laplacian operator. ( $\mathcal{U}$  and  $V$  are  $9t \times 9t$  diagonal matrices containing the first Lamé parameter and volume respectively for each tetrahedron on their diagonal).
- $M \in \mathbb{R}^{3n \times 3n}$  is the vector mass matrix.
- $J \in \mathbb{R}^{3n \times p}$  is the rig Jacobian for our *primary* rig.

**Algorithm 2:** performs one simulation step of our reduced Complementary Dynamics using a linear Co-Rotated elasticity model

```

Function simulationStep( $p, z_{hist}, p_{hist}, f_{ext}$ ):
   $z \leftarrow 0$ 
  while not converged do
     $\frac{\partial \tilde{\phi}}{\partial f} \leftarrow$  localStep( $z, p$ )
     $z_{next} \leftarrow$  globalStep( $z, p, z_{hist}, p_{hist}, f_{ext}, \frac{\partial \tilde{\phi}}{\partial f}$ )
     $z \leftarrow z_{next}$ 
  return  $z_{next}$ 

Function localStep( $z, p$ ):
   $\tilde{f} \leftarrow G_9 K B z + G_9 K J_0 p$ 
   $\tilde{F} \leftarrow$  reshape( $\tilde{f}, [r, \dim, \dim]$ )
   $\frac{\partial \tilde{\phi}}{\partial F} \leftarrow$  zeros( $r, \dim, \dim$ )
  for  $i \leftarrow 0$  to  $r$  do
     $F_i \leftarrow F[i, :, :]$ 
     $R_i \leftarrow$  PolarSVD( $F_i$ )
     $\frac{\partial \tilde{\phi}}{\partial F}[i] \leftarrow -m_i \mu_i R_i + m_i \frac{\lambda_i}{2} R_i tr(R_i^T F_i - I)$ 
   $\frac{\partial \tilde{\phi}}{\partial f} \leftarrow$  vec( $\frac{\partial \tilde{\phi}}{\partial F}$ )
  return  $\frac{\partial \tilde{\phi}}{\partial f}$ 

Function globalStep( $z, z_{hist}, p, p_{hist}, f_{ext}, \frac{\partial \tilde{\phi}}{\partial f}$ ):
   $f_{elastic} \leftarrow B^T L J p + (G_9 K B)^T \frac{\partial \tilde{\phi}}{\partial f}$ 
   $f_{inertia} \leftarrow \frac{1}{h^2} (B^T M J (p - p_{hist}) - B^T M B z_{hist})$ 
   $f \leftarrow f_{elastic} + f_{inertia} + f_{ext}$ 
   $A \leftarrow B^T L B + \frac{1}{h^2} B^T M B$ 
   $dz \leftarrow A dz = -f - A z$ 
   $\alpha \leftarrow$  lineSearch( $dz$ )
   $z_{next} \leftarrow z + \alpha dz$ 
  return  $z_{next}$ 

```

We can leverage the fact that our subspace just performs linear blend skinning to project our low dimensional output  $z_{next}$  to the full space entirely on the GPU. We implement this projection step in the vertex shader, where we load the initial skinning weights as vertex attributes, and then pass the updated  $z_{next}$  as a uniform in each draw call.

## 9 EXPERIMENTS & DISCUSSION

We evaluate the effectiveness of our skinning subspace for deformation by comparing it to a suite of common deformation subspaces. In our qualitative comparisons, we consider equal dimensional subspaces (e.g. 48 displacement modes = 12 degrees of freedom \* 4 skinning eigenmodes). While this puts CPU computation on equal footing, we emphasize that our skinning eigenmodes use significantly less GPU memory in a vertex shader implementation (12×).

### 9.1 Comparison to Modal Derivatives

Modal derivatives [Barbič and James 2005] augment the primary displacement subspace with a second set of modes that aim to correct the primary modes as they fall out of date with large deformations. This




Fig. 15. Our skinning subspace (right) can easily represent twisting motions, something displacement modes struggle with, even with localized rotation fitting.

effectively makes the subspace *quadratic* instead of linear. Unfortunately, modal derivatives are not closed under rotations. We show how this expresses itself in an elastic simulation in Fig. 14. We compare 3 subspace simulations: 60 displacement modes, 11 displacement modes augmented with 49 modal derivatives, and 10 skinning modes (corresponding to 60 total d.o.f.s in 2D). While modal derivatives greatly enrich the space available to the elasto-dynamics, they provide unintuitively different deformations under rotations. By contrast, our subspace allows for the lowest energy at equilibrium, maintains a consistent equilibrium shape (and energy) no matter the orientation.

Fig. 5 shows how this shortcoming manifests itself in our use case of fast complementary dynamics. In this example, a user applies small repeated rotations about the y-axis on a single affine handle controlling the moray eel, mimicking a swimming motion. Rig momentum is allowed to leak from the rig to the eel to induce secondary motion. Then, the user rotated the eel about the z-axis 90 degrees, making the eel swim sideways. As a user performs this rotation they would expect the resulting simulation to be identical up to a rotation as well, a property our 14 (168 d.o.f.s) skinning modes guarantee. Unfortunately the same cannot be said of a simulation using a modal derivative subspace of slightly greater size than ours. In this example, 17 primary modes augmented with a full set of 153 modal derivatives (for a total of 170 d.o.f.s) were used to construct the modal derivative subspace. Like our skinning eigenmodes, was also constrained to satisfy the complementarity constraint by construction.

Both methods benefit from performing a full-space projection step (to update the visualization) on the GPU. As a linear blend skinning subspace, we can perform this step at a much smaller memory cost for the vertex shader. For a full space projection step with  $k = d(d + 1)m$  degrees of freedom, modal derivatives would require  $k\text{-vec}^4$ 's in vertex shader memory, whereas ours would only require  $m$ .

## 9.2 Comparison to Sub-Structuring and Modal Warping

Perhaps the most intuitive way ensuring our subspace is closed under rotations is by warping the subspace manually with a single best fit rotation. This corresponds to a simplified version of both sub-structuring and modal warping [Barbič and Zhao 2011; Choi and Ko 2005; Kim and James 2011]. While this indeed would fix the simulation rotation equivariance problem, it comes at the cost of not being able to represent rotational motion within each sub-structure, as shown didactically in Fig. 15. Increasing the number of sub-structures would mitigate this, but still requires the user to warp the subspace associated with each vertex every timestep, a full space operation. Accelerating this update at each using the Fast Sandwich Transform (FST) proposed by [Kim and James 2011] is certainly possible for terms *linear* in




Fig. 16. Rotation Strain coordinates [Huang et al. 2011] are not well suited for reconstructing shapes undergoing rotational motion, as is commonly imposed by control rigs. Our subspace can fit this complex localized rotational motion with only 10 skinning modes.




Fig. 17. Adopting ICCM [Brandt and Hildebrandt 2017] allows us to reconstruct the highly localized reference deformation (left) using a least squares fit, a task in which traditional global modes struggle. We build our local modes (bottom) with a sparsity regularization parameter  $\eta = 100$  and only fit the first 5 modes for each.

**B.** However, it requires the reassembly of all linear pre-computed matrices involving  $\mathbf{B}$ , summing the contributions of each of the 9 rotation parameters *for each* new sub-structure. Albeit accelerated compared to a naive update, the FST update nevertheless becomes the bottleneck of our method, limiting the richness of the overall dynamics we can get for real-time rates. The FST also incurs a memory cost; in 3D, the aforementioned precomputed matrices must be stored separately 9 times *for each* new rotation cluster. On top of this, the FST would not allow us to efficiently update terms that are non-linear with respect to  $\mathbf{B}$ , such as our system matrix or our Cholesky prefactorization.

By contrast, our skinning subspace can represent twisting motions with a single static subspace computed once in a precomputation phase and *never* updated again.

## 9.3 Rotation Spanning vs. Rotation Equivariance

Fig. 13 highlights the difference between rotation spanning and closure under rotations. The left side shows a least squares fitting task carried out in a subspace, attempting to fit a rotation of the rest geometry. If the subspace does not span rotations, like standard displacement modes, it will not be able to reconstruct the target rotated rest geometry. Adding rotational degrees of freedom [Terzopoulos and Witkin




Fig. 18. Our rig-complementary subspace (right) composed of 14 skinning modes provide richer dynamics than an equal sized unconstrained subspace(left).

1988] by definition fix this problem and provides an equivalent solution to using a single constant skinning mode as our subspace. The right side of this figure does *not* attempt to fit rotations, but attempts to fit a subspace displacement applied on top of a user-rotated frame. Specifically it shows that without closure under rotations, how well you reconstructing this displacement becomes a frame dependant process. Even if we added an extra set of affine degrees of freedom to our subspace (with error shown in purple), these would provide negligible help.

#### 9.4 Comparison To Rotation Strain Coordinates

Rotation strain coordinates [Huang et al. 2011] fix linearized deformation artifacts that arise out of simulations using linear subspaces by projecting it to an energetically favorable pose. They decompose the tetrahedron’s deformation gradient  $F = r + s + I$  into a symmetric shear part  $s$  and antisymmetric linearized rotation part  $r$ . The linearized rotation is then projected to a full rotation matrix via an exponential map  $R = \exp(r)$ , from which the deformation gradient is reconstructed as  $F' = R(s + I)$ . The vertex positions  $x$  are determined by minimizing  $\|Kx - f'\|_M$ , where  $f' = \text{vec}(F')$ . As Fig. 16 shows however, this formulation depends on the input deformation to be composed mainly of shear motions (as a normal displacement subspace would). If a user manipulating a shape through a rig created a rotation on the neck of the giraffe, Rotation Strain coordinates lead to an undesired reconstruction of the shape, rendering them less suitable for reconstruction over rig-like motion. In fact, we can verify that for an input 2D rotational motion about the  $z$ -axis of  $\frac{\pi}{2}$  leads to a target deformation gradient

$$F' = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

which will act to collapse neighboring vertices to lie solely along the  $z$ -axis. Furthermore, the Rotation Strain coordinates position fitting step may undo any constraints enforced in earlier steps, including our complementarity constraint. Identifying how to optimally enforce such constraints in the rotation strain coordinates pipelines remains an interesting avenue for future work.


#### 9.5 Localized Skinning Subspaces

Thanks to our generalized eigenvalue formulation, we can greatly benefit from prior work in localizing and sparsifying PCA/Modal analysis subspaces [Houston 2017; Melzi et al. 2018; Nasikun et al. 2018]. As an example, we can promote sparsity and locality in our modes to by applying the *iterated convexification for compressed modes* (ICCM) [Brandt and Hildebrandt 2017] algorithm. This amounts to simply adding an L1 regularization term to our minimization in Eq. (22). The result is a subspace that can faithfully capture local deformations as shown in Fig. 17. The skinning modes’ closure under rotation, critical for producing rotation equivariant results, naturally inherits the locality presented by the skinning weights: observe that Eq. (26) is a superposition of the closure of every mode  $b$  in isolation:

$$w_{ib} (RT_b) \begin{bmatrix} x_{0i} \\ 1 \end{bmatrix} = R \left( w_{ib} T_b \begin{bmatrix} x_{0i} \\ 1 \end{bmatrix} \right). \quad (32)$$

#### 9.6 Difficulty Capturing High Frequency Motion

Every secondary skinning weight in our subspace is associated with 12 affine degrees of freedom in 3D, as shown in Fig. 7 (and 6 in 2D). While this is the key to many of the rotational qualities of our subspace, it also means that our subspace re-



quires many degrees of freedom to represent certain high frequency motions. The inset shows a floppy bar being pulled at its left extremity, creating an elastic wave that ripples from left to right. A traditional displacement subspace can capture such high frequency motion, whereas our skinning subspace suffers from global artifacts with the same number of degrees of freedom.

#### 9.7 Comparison to Coarsening Meshes

A common approach to accelerating elasto-dynamics for real-time applications is to embed a fine mesh inside a coarse one, solve the elasto-dynamics on the coarse mesh, and then map the final motion back to the fine mesh via the embedding.

This corresponds to a very specific subspace  $B_{embed}$ ; one that is very sparse, highly localized, can represent rotations, and maintains simulation rotation equivariance. Unfortunately, its construction from a coarse mesh embedding groups fine scale features that are close in Euclidean space whose motion should not be correlated. This leads to visible embedding artifacts as shown in Fig. 3. To compare, our subspace couples motion of vertices as measured by the elastic energy.

#### 9.8 Discussion on Material Sensitivity

An important feature that distinguishes our skinning weights from those of Gilles et al. [2011]; Lan et al. [2020]; Wang et al. [2015] is that our skinning weights are *material-aware*. We show that this directly leads to richer dynamics when dealing with subspace simulations of heterogeneous materials, as shown in Fig. 10.

Because our final weight space Hessian in Eq. (18) does not capture any coupling interactions between the dimensions of our simulation





Fig. 19. Our skinning subspace can easily capture volume-preserving effects for non-zero poisson ratio, and maintains such high quality deformations under any rotations of the rest frame.

it is invariant to changes in the Poisson ratio, which appear in the off-diagonal entries of the full space hessian  $H$ . Fortunately, our skinning subspace can form non-uniform scales and shears, which allow it to excellently capture volume preserving effects with very few skinning modes, as shown in Fig. 19.

### 9.9 Importance of Constraining Subspace

One big difference between our skinning subspace and prior methods [Brandt et al. 2018; Faure et al. 2011; von Tycowicz et al. 2013] is that we can impose homogeneous equality constraints on our skinning weights by construction. Without this property, we’d have to impose our constraint at run-time. A user is burdened with having to select a subspace with more degrees of freedom than the constraint set in order to avoid an over-constrained system. To make things worse, the constraints for our primary rig can quickly climb in dimensionality; for a linear blend skinning control rig, every new bone brings with it  $d(d+1)$  new columns in our complementarity constraint matrix  $J$ . Fig. 18 shows a primary control rig composed of 14 bones (168 constraints). In order for a simulation to avoid being over-constrained, the user would then have to prescribe over 14 skinning modes (168) degrees of freedom. In contrast, our subspace frees the user from this dilemma, and allows the user to pick the dimensionality of the subspace without worrying about the well-posedness of the simulation.

### 9.10 Subspace Ablation



The inset shows an ablation study on the number of modes (top) and the number of clusters (bottom). For both ablations, we compare the displacement incurred after a single timestep’s displacement to a reference full space displacement, and compute the L2 error  $\|\mathbf{u} - \mathbf{u}_{ref}\|^2$ . For the ablation on the modes, we fix the number of clusters to 3000. For the ablation on the clusters, we fix the number of skinning modes to 200.




Fig. 20. We observe volume preserving effects made possible by accommodating co-rotational elasticity.

Table 1. Computing 10 of our skinning modes (120 d.o.f.s) is much faster than computing 120 displacement modes (120 d.o.f.s) subspace of an equivalent size.


Mesh	#Vertices	Disp.(s)	Ours(s)
Elephant	7842	0.782	<b>0.135</b>
Bulldog	31368	6.76	<b>0.830</b>
XYZ Dragon	99813	62.3	<b>5.14</b>
King Ghidora	294033	143.4	<b>14.07</b>

### 9.11 Evaluating Co-Rotational Elasticity

We generalize the clustering based local global solver of Jacobson et al. [2012] to work with a more general material model of Linear Co-Rotational Elasticity [McAdams et al. 2011]. We show a simple example in Fig. 19 that shows how our subspace simulation with 50 clusters and 10 skinning modes can incorporate smooth volume preserving effects.

Fig. 20 highlights that these volume preserving effects have a noticeable visual impact on the secondary dynamics. Unfortunately, such non-linear effects aren’t free, a non-zero Poisson’s ratio incurs the cost of tacking on a line search at the end of our solver’s global step [Liu et al. 2017]. Evaluating the energy for the line search requires re-computing per-cluster deformation gradients, a  $O(m)$  operation that corresponds to the slowest part in our simulation pipeline for large modes and large clusters. While this does sometimes hinder the speed of our simulation step, we have found that there are large portions of the mode/cluster parameter space that easily accommodate this computational overhead while still providing rich dynamics.

### 9.12 Mode/Cluster Pareto-Front



We evaluate how our per-timestep cost changes with the size of our subspace skinning modes and the number of clusters used. The inset highlights the 30 FPS real-time pareto front. For each trial, the same timestep was run 100 times until convergence on an ARAP elasticity model.

### 9.13 Computing Our Modes

Because our weight space Hessian is  $n \times n$ , compared to the traditional full space  $n(d) \times n(d)$  Hessian, solving for our subspace is

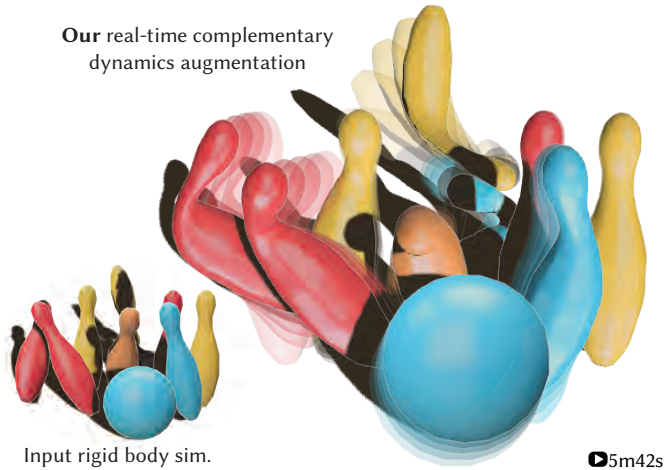



Fig. 21. Augmenting a real-time rigid body simulation with complementary dynamics in real-time.

naturally faster for higher dimensions, as shown in Table 1. Applying the complementarity constraint on our modes means augmenting the weight space Hessian  $H_w$  in our weight computation with our weight-space complementarity constraint  $J_w \in \mathbb{R}^{n \times p(d)(d+1)}$ . This makes mode computation naturally more computationally expensive for more complicated rigs. The inset shows how changing the number of rig parameters in a control rig affects how long it takes to compute 10 skinning modes on the elephant mesh.

## 10 RESULTS

### 10.1 Rigid Body Augmentation

Rigid body simulations are key ingredients in many video games and virtual reality applications. A rigid body can also be interpreted as a very simple specific case of a linear blend skinning, with only a single bone of constant weight equal to 1 across all the vertices. We exploit this analogy to augment real-time rigid-body simulations with secondary effects (Fig. 21 shows screenshots from an interactive bowling game, where the pins are stylized with secondary effects.)



to augment real-time rigid-body simulations with secondary effects (Fig. 21 shows screenshots from an interactive bowling game, where the pins are stylized with secondary effects.)

### 10.2 Augmenting Mixamo Characters with Secondary Motion

Animators often design large suites of rig animations, often having to browse and edit many of these in rapid succession. Adding on to the tediousness of this process, they also have to foresee how these animations will look when augmented with secondary motion. Our method allows a user to browse through any number of rig-animations and immediately observe the secondary secondary dynamics in real-time. We



Fig. 22. Augmenting and browsing various Mixamo character animations in real-time.

demonstrate this for a character downloaded from the Mixamo [2023] website in Fig. 22.

### 10.3 Secondary Dynamics for Digital Avatars

Digital avatars are key to expressing personality and style in virtual environments. We show that our Fast Complementary Dynamics can be plugged into very simple camera based pose trackers made available by mediapipe [Lugaresi et al. 2019] to breath life into these digital avatars. Fig. 23 demonstrates augmenting human faces tracked by mediapipe’s Faces solution with secondary motion. Similarly, Fig. 24 augment skeletons tracked by mediapipe’s Pose solution with real-time secondary motion. In practice the bottleneck of this application became mediapipe’s own Pose solution, requiring us to separate the tracker from the dynamics and place them into two separate threads.



Fig. 23. Augmenting mediapipe’s face tracking with real-time secondary motion on digital avatars.

## 11 LIMITATIONS AND FUTURE WORK

Compared to the original complementary dynamics [Zhang et al. 2020], we assume that the user’s rig is linear. While non-linear rigs are less common in real-time settings, it would still be interesting to accommodate them (e.g., [Kavan et al. 2008], perhaps via low-rank updates to the reduced system matrices in Alg. 2 [Hahn et al. 2012]). Our reduction techniques are in many ways agnostic to the choice of elastic potential. In future work, we would like to explore beyond co-rotational models. Our skinning eigenmodes easily facilitate rig-orthogonality constraints for complementary dynamics; their other



Fig. 24. Augmenting mediapipe's posetracking with real-time secondary motion on digital avatars.

good qualities suggest they could be useful beyond our target application, applied to more general elasticity problems (e.g., structural analysis or inverse design).

We leave comprehensive collision handling as future work. While rigid-body augmentation provides a drop-in heuristic for collision effects for relatively stiff or fast moving objects, it would be interesting to explore more accurate methods, perhaps combining our contributions with cubature techniques [Harmon and Zorin 2013]. For real-time VR avatar, applications we found the bottleneck for quality lies in the tracking and mapping of user movements to primary rig controls. Our secondary effects would immediately inherit any improvements in these active research areas.

We derive skinning eigenmodes by considering distributions of translations. This sidesteps the issue of determining an origin and scale associated with each mode. It would be interesting to treat these as variables (to achieve theoretical optimality over distributions of general affine transformations). This appears non-trivial and irreducible to a generalized eigenvalue problem

We have presented a novel subspace for deformation that is well suited for augmenting real-time rig animations with secondary, complementary motion. Our computation is well-balanced between fast small iterations on the CPU and memory-efficient, standard-pipeline vertex shaders on the GPU. In future work, we are interested in not just adding features to our elasticity effects, but also considering complementary dynamics more broadly into domains such as fluid simulation, electrostatics, and crowds. We hope that our work also serves as general recipe for translating complementary dynamics to real-time scenarios.

## ACKNOWLEDGMENTS

This project is funded in part by NSERC Discovery (RGPIN2017–05235, RGPIN-2021-03733, RGPAS–2017–507938) New Frontiers of Research Fund (NFRFE–201), the Ontario Early Research Award program, the Canada Research Chairs Program, a Sloan Research Fellowship, the DSI Catalyst Grant program and gifts by Adobe Inc. Otman Benchekroun was funded by an NSERC CGS-M scholarship, while Jiayi Eris Zhang is funded by a Stanford Graduate Fellowship. We thank David I.W. Levin, Danny M. Kaufman, Doug L. James, Ty Trusty, Sarah Kushner

and Silvia Sellán for insightful conversations. We thank Alejandra Baptista Aguilar for early discussion and adoption in 3D environments, and for providing some of the models in the paper. We thank Silvia Sellán, Selena Ling, Shukai Chen, Aravind Ramakrishnan and Kinjal Parikh for proofreading. We thank Xuan Dam and John Hancock for technical and administrative assistance throughout the course of this project.

## REFERENCES

- Adobe Systems Inc. 2023. Mixamo. <https://mixamo.com>.
- 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>
- David Arthur and Sergei Vassilvitskii. 2007. k-means++: the advantages of careful seeding. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms* (New Orleans, Louisiana). Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1027–1035.
- David Baraff and Andrew Witkin. 1998. Large Steps in Cloth Simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 43–54. <https://doi.org/10.1145/280814.280821>
- Jernej Barbič and Doug L. James. 2005. Real-Time Subspace Integration for St. Venant-Kirchhoff Deformable Models. *ACM Trans. Graph.* 24, 3 (jul 2005), 982–990. <https://doi.org/10.1145/1073204.1073300>
- Jernej Barbič and Yili Zhao. 2011. Real-time Large-deformation Substructuring. *ACM Trans. on Graphics (SIGGRAPH 2011)* 30, 4 (2011), 91:1–91:7.
- 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 (jul 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 (03 2017). <https://doi.org/10.1016/j.cagd.2017.03.004>
- Christopher Brandt, Leonardo Scandolo, Elmar Eisemann, and Klaus Hildebrandt. 2019. The Reduced Immersed Method for Real-Time Fluid-Elastic Solid Interaction and Contact Simulation. *ACM Trans. Graph.* 38, 6, Article 191 (nov 2019), 16 pages. <https://doi.org/10.1145/3355089.3356496>
- Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. 2010. A Simple Geometric Model for Elastic Deformations. 29, 4, Article 38 (jul 2010), 6 pages. <https://doi.org/10.1145/1778765.1778775>
- Min Gyu Choi and Hyeong-Seok Ko. 2005. Modal warping: real-time simulation of large rotational deformation and manipulation. *IEEE Transactions on Visualization and Computer Graphics* 11, 1 (2005), 91–101. <https://doi.org/10.1109/TVCG.2005.13>
- Simon Duenser, Bernhard Thomaszewski, Roi Poranne, and Stelian Coros. 2022. Nonlinear Compliant Modes for Large-Deformation Analysis of Flexible Structures. *ACM Trans. Graph.* 42, 2, Article 21 (nov 2022), 11 pages. <https://doi.org/10.1145/3568952>
- 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 (jul 2011), 10 pages. <https://doi.org/10.1145/2010324.1964968>
- Benjamin Gilles, Guillaume Bousquet, Francois Faure, and Dinesh K. Pai. 2011. Frame-Based Elastic Models. *ACM Trans. Graph.* 30, 2, Article 15 (apr 2011), 12 pages. <https://doi.org/10.1145/1944846.1944855>
- Gene H. Golub. 1973. Some Modified Matrix Eigenvalue Problems. *SIAM Rev.* 15, 2 (1973), 318–334. <https://doi.org/10.1137/1015032> arXiv:<https://doi.org/10.1137/1015032>
- Fabian Hahn, Sebastian Martin, Bernhard Thomaszewski, Robert Sumner, Stelian Coros, and Markus Gross. 2012. Rig-Space Physics. *ACM Trans. Graph.* 31, 4, Article 72 (jul 2012), 8 pages. <https://doi.org/10.1145/2185520.2185568>
- Fabian Hahn, Bernhard Thomaszewski, Stelian Coros, Robert W Sumner, and Markus Gross. 2013. Efficient simulation of secondary motion in rig-space. In *Proceedings of the 12th ACM SIGGRAPH/eurographics symposium on computer animation*. 165–171.
- David Harmon and Denis Zorin. 2013. Subspace Integration with Local Deformations. *ACM Trans. Graph.* 32, 4, Article 107 (jul 2013), 10 pages. <https://doi.org/10.1145/2461912.2461922>
- Klaus Hildebrandt, Christian Schulz, Christoph Von Tycowicz, and Konrad Polthier. 2011. Interactive Surface Modeling Using Modal Analysis. *ACM Trans. Graph.* 30, 5, Article 119 (oct 2011), 11 pages. <https://doi.org/10.1145/2019627.2019638>
- Kevin Houston. 2017. Sequentially-Defined Compressed Modes via ADM. In *Proceedings of the Symposium on Geometry Processing: Posters* (London, United Kingdom) (SGP '17). Eurographics Association, Goslar, DEU, 1–2. <https://doi.org/10.2312/sgp.20171201>
- 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 (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 (2012), to appear.
- Alec Jacobson, Ilya Baran, Jovan Popovic, and Olga Sorkine. 2011. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph.* 30, 4 (2011), 78.
- Doug James and Dinesh K. Pai. 1999. ARTDEFO: accurate real time deformable objects. *Proc. of ACM SIGGRAPH 99*, 65–72. <https://doi.org/10.1145/311535.311542>
- 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 (jul 2002), 582–585. <https://doi.org/10.1145/566654.566621>
- Ladislav Kavan, Steven Collins, Jiri Zara, and Carol O’Sullivan. 2008. Geometric Skinning with Approximate Dual Quaternion Blending. *ACM Trans. Graph.* 27, 4 (2008), 105.
- Theodore Kim and David Eberle. 2020. Dynamic Deformables: Implementation and Production Practicalities (*SIGGRAPH ’20*). Association for Computing Machinery, New York, NY, USA, Article 23, 182 pages. <https://doi.org/10.1145/3388769.3407490>
- Theodore Kim and Doug L. James. 2011. Physics-Based Character Skinning Using Multi-Domain Subspace Deformations (*SCA ’11*). Association for Computing Machinery, New York, NY, USA, 63–72. <https://doi.org/10.1145/2019406.2019415>
- Paul G. Kry, Doug L. James, and Dinesh K. Pai. 2002. EigenSkin: Real Time Large Deformation Character Skinning in Hardware. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (San Antonio, Texas) (SCA ’02)*. Association for Computing Machinery, New York, NY, USA, 153–159. <https://doi.org/10.1145/545261.545286>
- Lei Lan, Ran Luo, Marco Fratarcangeli, Weiwei Xu, Huamin Wang, Xiaohu Guo, Junfeng Yao, and Yin Yang. 2020. Medial Elastics: Efficient and Collision-Ready Deformation via Medial Axis Transform. *ACM Trans. Graph.* 39, 3, Article 20 (apr 2020), 17 pages. <https://doi.org/10.1145/3384451>
- 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 (jul 2021), 16 pages. <https://doi.org/10.1145/3450626.3459753>
- Torsten Langer and Hans-Peter Seidel. 2008. Higher Order Barycentric Coordinates. *Comput. Graph. Forum* 27, 2 (2008), 459–466.
- Hsueh-Ti Derek Liu, Jiayi Eris Zhang, Mirela Ben-Chen, and Alec Jacobson. 2021. Surface Multigrid via Intrinsic Prolongation. *ACM Trans. Graph.* 40, 4, Article 80 (jul 2021), 13 pages. <https://doi.org/10.1145/3450626.3459768>
- Tiantian Liu, Adam W. Bargteil, James F. O’Brien, and Ladislav Kavan. 2013. Fast Simulation of Mass-Spring Systems. *ACM Transactions on Graphics* 32, 6 (Nov. 2013), 209:1–7. <http://cg.cis.upenn.edu/publications/Liu-FMS> Proceedings of ACM SIGGRAPH Asia 2013, Hong Kong.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-Newton Methods for Real-Time Simulation of Hyperelastic Materials. *ACM Trans. Graph.* 36, 4, Article 116a (jul 2017), 16 pages. <https://doi.org/10.1145/3072959.2990496>
- Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. 2019. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172* (2019).
- Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. 2011. Efficient Elasticity for Character Skinning with Contact and Collisions. In *ACM SIGGRAPH 2011 Papers* (Vancouver, British Columbia, Canada) (*SIGGRAPH ’11*). Association for Computing Machinery, New York, NY, USA, Article 37, 12 pages. <https://doi.org/10.1145/1964921.1964932>
- Simone Melzi, Emanuele Rodolà, Umberto Castellani, and Michael M Bronstein. 2018. Localized manifold harmonics for spectral shape analysis. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 20–34.
- Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position Based Dynamics. *J. Vis. Commun. Image Represent.* 18, 2 (apr 2007), 109–118. <https://doi.org/10.1016/j.jvcir.2007.01.005>
- Ahmad Nasikun, Christopher Brandt, and Klaus Hildebrandt. 2018. Fast Approximation of Laplace-Beltrami Eigenproblems. *Computer Graphics Forum* 37 (08 2018), 121–134. <https://doi.org/10.1111/cgf.13496>
- A. Pentland and J. Williams. 1989. Good Vibrations: Modal Dynamics for Graphics and Animation. *SIGGRAPH Comput. Graph.* 23, 3 (jul 1989), 207–214. <https://doi.org/10.1145/74334.74355>
- Damien Rohmer, Marco Tarini, Niranjan Kalyanasundaram, Faezeh Moshfeghifar, Marie-Paule Cani, and Victor Zordan. 2021. Velocity Skinning for Real-time Stylized Skeletal Animation. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 549–561.
- Silvia Sellán, Jack Luong, Leticia Mattos Da Silva, Aravind Ramakrishnan, Yuchuan Yang, and Alec Jacobson. 2021. Breaking Good: Fracture Modes for Realtime Destruction. *CoRR abs/2111.05249* (2021). <https://arxiv.org/abs/2111.05249>
- Olga Sorkine and Marc Alexa. 2007. As-Rigid-As-Possible Surface Modeling. In *Geometry Processing*, Alexander Belyaev and Michael Garland (Eds.). The Eurographics Association. <https://doi.org/10.2312/SGP/SGP07/109-116>
- Ryusuke Sugimoto, Christopher Batty, and Toshiya Hachisuka. 2022. Surface-Only Dynamic Deformables using a Boundary Element Method. *Computer Graphics Forum* (2022). <https://doi.org/10.1111/cgf.14625>
- Yun Teng, Miguel A. Otaduy, and Theodore Kim. 2014. Simulating Articulated Subspace Self-Contact. *ACM Trans. Graph.* 33, 4, Article 106 (jul 2014), 9 pages. <https://doi.org/10.1145/2601097.2601181>
- Demetri Terzopoulos and Andrew Witkin. 1988. Physically Based Models with Rigid and Deformable Components. *IEEE Comput. Graph. Appl.* 8, 6 (nov 1988), 41–51. <https://doi.org/10.1109/38.20317>
- Ty Trusty, Danny M. Kaufman, and David I. W. Levin. 2022. Mixed Variational Finite Elements for Implicit, General-Purpose Simulation of Deformables. *CoRR abs/2202.00183* (2022). <https://arxiv.org/abs/2202.00183>
- 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 Barbic, and Ladislav Kavan. 2015. Linear Subspace Design for Real-Time Shape Deformation. *ACM Trans. Graph.* 34, 4 (2015).
- Jiayi Eris Zhang, Seungbae Bang, David I.W. Levin, and Alec Jacobson. 2020. Complementary Dynamics. *ACM Transactions on Graphics* (2020).
- Mianlun Zheng, Yi Zhou, Duygu Ceylan, and Jernej Barbic. 2021. A deep emulator for secondary motion of 3d characters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5932–5940.

## A PROOFS

### A.1 A subspace simulation is rotation equivariant if and only if the subspace is closed under rotations

Consider the minimization of a rotation invariant elastodynamic energy  $E(\mathbf{u} + \mathbf{x})$ , where  $\mathbf{u}, \mathbf{x} \in \mathbb{R}^{3n}$ . By rotation equivariance we mean that

$$\forall \mathbf{R}, \forall \mathbf{x}, \underset{\mathbf{u}}{\operatorname{argmin}} E(\mathbf{u} + (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}) = (\mathbf{R} \otimes \mathbf{I}) \underset{\mathbf{v}}{\operatorname{argmin}} E(\mathbf{v} + \mathbf{x}).$$

Here and henceforth the domains for  $\mathbf{R}$  and  $\mathbf{x}$  shall be the rotations  $SO(3)$  and the positions  $\mathbb{R}^{3n}$ , respectively.

In a full space simulation, the above holds. For a subspace simulation where  $\mathbf{u} = \mathbf{B}\mathbf{z}$ , and  $\operatorname{Col}(\mathbf{B}) \subset \mathbb{R}^{3n}$ , the equivalent statement,

$$\forall \mathbf{R}, \forall \mathbf{x}, \underset{\mathbf{z}}{\operatorname{argmin}} E(\mathbf{B}\mathbf{z} + (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}) = (\mathbf{R} \otimes \mathbf{I}) \underset{\mathbf{w}}{\operatorname{argmin}} E(\mathbf{B}\mathbf{w} + \mathbf{x}), \quad (33)$$

no longer holds for all subspaces. We shall prove that the above holds if and only if our subspace is closed under rotations,

$$\forall \mathbf{R} \in SO(3), \forall \mathbf{z} \in \mathbb{R}^m, \exists \mathbf{w} \in \mathbb{R}^m, \mathbf{B}\mathbf{z} = (\mathbf{R} \otimes \mathbf{I}) \mathbf{B}\mathbf{w}. \quad (34)$$

**A.1.1 Forward Proof.** Given that our subspace satisfied closure under rotations (Eq. (34)), we prove that Eq. (33) holds.

**Theorem:** (34)  $\rightarrow$  (33). **Proof:** Consider (without loss of generality) some rotation  $\mathbf{R}$  and position  $\mathbf{x}$ . Let  $\mathbf{z}^*$  be the minimizer of the left hand side of Eq. (33):

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} E(\mathbf{B}\mathbf{z} + (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}) \quad (35)$$

By closure of our subspace under rotations,  $\exists \mathbf{w}^*, \mathbf{B}\mathbf{z}^* = (\mathbf{R} \otimes \mathbf{I}) \mathbf{B}\mathbf{w}^*$ . The minimum energy is thus

$$E^* = E(\mathbf{B}\mathbf{z}^* + (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}) = E((\mathbf{R} \otimes \mathbf{I}) \mathbf{B}\mathbf{w}^* + (\mathbf{R} \otimes \mathbf{I}) \mathbf{x}) = E(\mathbf{B}\mathbf{w}^* + \mathbf{x})$$

where (from left to right) the equalities use that  $\mathbf{z}^*$  is a minimizer,  $\mathbf{B}$  is closed under rotations, and  $E(\cdot)$  is rotation invariant, respectively. Consequently,  $\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} E(\mathbf{B}\mathbf{w} + \mathbf{x})$ , which we substitute along with (35) into (34) to complete the proof.  $\square$

**A.1.2 Backward Proof.** Given our subspace simulation is rotation equivariant (Eq. (33)), we prove that Eq. (34) holds.

*Lemma.*

$$\forall \mathbf{R}^* \in SO(3), \forall \mathbf{z}^* \in \mathbb{R}^m, \quad (36)$$

$$\exists \mathbf{x}^*, \underset{z}{\operatorname{argmin}} E(\mathbf{B}z + (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{x}^*) = \mathbf{z}^*. \quad (37)$$

**Proof:** Let  $\tilde{z}$  be some minimizer of  $E(\mathbf{B}z)$  by construction. To satisfy the Lemma we choose  $\mathbf{x}^* = (\mathbf{R}^* \otimes \mathbf{I})^{-1} \mathbf{B}(\tilde{z} - \mathbf{z}^*)$ , because

$$\begin{aligned} \underset{z}{\operatorname{argmin}} E(\mathbf{B}z + (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{x}^*) &= \\ \underset{z}{\operatorname{argmin}} E(\mathbf{B}z + (\mathbf{R}^* \otimes \mathbf{I}) (\mathbf{R}^* \otimes \mathbf{I})^{-1} \mathbf{B}(\tilde{z} - \mathbf{z}^*)) &= \\ \underset{z}{\operatorname{argmin}} E(\mathbf{B}(z + \tilde{z} - \mathbf{z}^*)) &. \end{aligned} \quad (38)$$

Thus  $w = z + \tilde{z} - z^*$  minimizes  $E(\mathbf{B}w)$ , and so does  $w = \tilde{z}$  (by construction). Since  $w = w$ , then  $z = z^*$  minimizes  $E(\mathbf{B}z + (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{x}^*)$ .  $\square$

*Theorem:* (33)  $\rightarrow$  (34). **Proof:** For any given  $\mathbf{R}^*$ ,  $\mathbf{z}^*$ , we produce  $\mathbf{w}^*$  satisfying

$$\mathbf{B}z^* = (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{B}w^*.$$

Pick a specific  $\mathbf{R}^*$ ,  $\mathbf{z}^*$ . Using the Lemma, we pick  $\mathbf{x}^*$  such that  $\mathbf{z}^*$  minimizes  $E(\mathbf{B}z + (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{x}^*)$ . By the rotation equivariance of our minimization,

$$\begin{aligned} \mathbf{B} \underset{z}{\operatorname{argmin}} E(\mathbf{B}z + (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{x}^*) &= (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{B} \underset{w}{\operatorname{argmin}} E(\mathbf{B}w + \mathbf{x}^*) \\ \mathbf{B}z^* &= (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{B} \underset{w}{\operatorname{argmin}} E(\mathbf{B}w + \mathbf{x}^*) \\ &= (\mathbf{R}^* \otimes \mathbf{I}) \mathbf{B}w^*, \end{aligned}$$

where we have chosen  $w^* = \underset{w}{\operatorname{argmin}} E(\mathbf{B}w + \mathbf{x}^*)$ .  $\square$

## B MATRIX DERIVATIONS

### B.1 Weight Space Skinning Jacobian Products

The product  $\mathbf{J}^T \mathbf{B}_{lbs}$  frequently arises in our derivations. Here  $\mathbf{J} \in \mathbb{R}^{n(d) \times c}$  can be a generic matrix and  $\mathbf{B}_{lbs} \in \mathbb{R}^{n(d) \times d(d+1)(m)}$  is our Linear Blend Skinning Jacobian matrix, which depends linearly on some set of weights  $\mathbf{W} \in \mathbb{R}^{n(d) \times m}$ . We leverage the relationship between  $\mathbf{B}_{lbs}$  and  $\mathbf{W}$  to rewrite the entries of the matrix product  $\mathbf{J}^T \mathbf{B}_{lbs}$  entirely in terms of the weights:

$$\mathbf{B}_{lbs} = (\mathbf{I}_d \otimes ((\mathbf{1}_m^T \otimes [\mathbf{V} \mathbf{1}_n]) \otimes (\mathbf{W} \otimes \mathbf{1}_{d+1}^T))) \underset{\mathcal{B}}{\quad}$$

We assume the 3D case  $d = 3$ , and restate the final result 2D. The full product whose entries we wish to relate linearly to our weights  $\mathbf{W}$  is

$$\mathbf{J}^T (\mathbf{I}_3 \otimes \mathcal{B}) \quad (39)$$

Separating the multiplying matrix into its 3 sets of columns:

$$\mathbf{J}^T = [\mathbf{J}_x^T \mid \mathbf{J}_y^T \mid \mathbf{J}_z^T] \quad (40)$$

Expanding the rows of this product:

$$[\mathbf{J}_x^T \mathcal{B} \mid \mathbf{J}_y^T \mathcal{B} \mid \mathbf{J}_z^T \mathcal{B}] \quad (41)$$

We focus on one set of columns  $i \in \{x, y, z\}$  at a time:

$$\begin{aligned} \mathbf{J}_i^T \mathcal{B} &= \mathbf{J}_i^T ((\mathbf{1}_m^T \otimes [\mathbf{X} \mathbf{1}_n]) \otimes (\mathbf{W} \otimes \mathbf{1}_4^T)) \in \mathbb{R}^{j \times (4)m} \\ &= \mathbf{J}_i^T ((\mathbf{1}_m^T \otimes [\tilde{x} \tilde{y} \tilde{z} \mathbf{1}_n]) \otimes ([\mathbf{w}_1 \mathbf{w}_2 \dots \mathbf{w}_m] \otimes \mathbf{1}_4^T)) \end{aligned}$$

We recall the definition of the Linear Blend Skinning Jacobian where we split

$\mathbf{X} = [\tilde{x} \tilde{y} \tilde{z}]$  into its individual columns. Expanding the two Kronecker products:

$$\mathbf{J}_i^T ([\underbrace{\tilde{x} \tilde{y} \tilde{z} \mathbf{1}_n}_1 \mid \dots \mid \underbrace{\tilde{x} \tilde{y} \tilde{z} \mathbf{1}_n}_b] \otimes [\underbrace{\mathbf{w}_1 \mathbf{w}_1 \mathbf{w}_1 \mathbf{w}_1}_1 \mid \dots \mid \underbrace{\mathbf{w}_m \mathbf{w}_m \mathbf{w}_m \mathbf{w}_m}_m])$$

Distributing the products further:

$$\mathbf{J}_i^T [\underbrace{([\tilde{x} \tilde{y} \tilde{z} \mathbf{1}_n] \otimes [\mathbf{w}_1 \mathbf{w}_1 \mathbf{w}_1 \mathbf{w}_1])}_1 \mid \dots \mid \underbrace{([\tilde{x} \tilde{y} \tilde{z} \mathbf{1}_n] \otimes [\mathbf{w}_m \mathbf{w}_m \mathbf{w}_m \mathbf{w}_m])}_m]$$

Moving the matrix product inside, and looking at one block  $q \in [1, \dots, m]$  at a time:

$$\begin{aligned} \mathbf{J}_i^T ([\tilde{x} \tilde{y} \tilde{z} \mathbf{1}_n] \otimes [\mathbf{w}_q \mathbf{w}_q \mathbf{w}_q \mathbf{w}_q]) \\ = \mathbf{J}_d^T ([\tilde{x} \otimes \mathbf{w}_q \mid \tilde{y} \otimes \mathbf{w}_q \mid \tilde{z} \otimes \mathbf{w}_q \mid \mathbf{1}_n \otimes \mathbf{w}_q]) \end{aligned}$$

Rewriting the individual Hadamard products as diagonal matrix multiplication where we introduce the capitalized overbar notation  $\bar{X}, \bar{Y}, \bar{Z}$  to indicate the diagonal matrices whose diagonal entries are  $\tilde{x}, \tilde{y}, \tilde{z}$  respectively.

$$\begin{aligned} &= \mathbf{J}_d^T ([\bar{X} \mathbf{w}_q \mid \bar{Y} \mathbf{w}_q \mid \bar{Z} \mathbf{w}_q \mid \mathbf{w}_q]) \\ &= [\mathbf{J}_d^T \bar{X} \mathbf{w}_q \mid \mathbf{J}_d^T \bar{Y} \mathbf{w}_q \mid \mathbf{J}_d^T \bar{Z} \mathbf{w}_q \mid \mathbf{J}_d^T \mathbf{w}_q] \end{aligned}$$

Applying this for all weights  $b$ :

$$[\mathbf{J}_i^T \bar{X} \mathbf{W} \mid \mathbf{J}_i^T \bar{Y} \mathbf{W} \mid \mathbf{J}_i^T \bar{Z} \mathbf{W} \mid \mathbf{J}_i^T \mathbf{W}]$$

Finally, applying this to all 3 dimensions of  $i$ :

$$\begin{aligned} \mathbf{J}^T \mathbf{B}_{lbs} &= [\mathbf{J}_x^T \bar{X} \mathbf{W} \mid \mathbf{J}_x^T \bar{Y} \mathbf{W} \mid \mathbf{J}_x^T \bar{Z} \mathbf{W} \mid \mathbf{J}_x^T \mathbf{W} \mid \dots \\ &\quad \mathbf{J}_y^T \bar{X} \mathbf{W} \mid \mathbf{J}_y^T \bar{Y} \mathbf{W} \mid \mathbf{J}_y^T \bar{Z} \mathbf{W} \mid \mathbf{J}_y^T \mathbf{W} \mid \dots \\ &\quad \mathbf{J}_z^T \bar{X} \mathbf{W} \mid \mathbf{J}_z^T \bar{Y} \mathbf{W} \mid \mathbf{J}_z^T \bar{Z} \mathbf{W} \mid \mathbf{J}_z^T \mathbf{W}] \end{aligned} \quad (42)$$

We can derive our weight Skinning Jacobians by rewriting the entries  $\mathbf{B}_{lbs}$  alone in terms of the weights  $\mathbf{W}$  as a special case of the above result. If we choose  $\mathbf{J}$  to be the identity, and slice out the 3 sets of columns belonging to the identity, we obtain our dimensional selection matrices  $\mathbf{J} = \mathbf{I}_{3n} = [\mathbf{P}_x \mathbf{P}_y \mathbf{P}_z]$ . We can then rewrite :

$$\begin{aligned} \mathbf{B}_{lbs} &= [\mathbf{P}_x \bar{X} \mathbf{W} \mid \mathbf{P}_x \bar{Y} \mathbf{W} \mid \mathbf{P}_x \bar{Z} \mathbf{W} \mid \mathbf{P}_x \mathbf{W} \mid \dots \\ &\quad \mathbf{P}_y \bar{X} \mathbf{W} \mid \mathbf{P}_y \bar{Y} \mathbf{W} \mid \mathbf{P}_y \bar{Z} \mathbf{W} \mid \mathbf{P}_y \mathbf{W} \mid \dots \\ &\quad \mathbf{P}_z \bar{X} \mathbf{W} \mid \mathbf{P}_z \bar{Y} \mathbf{W} \mid \mathbf{P}_z \bar{Z} \mathbf{W} \mid \mathbf{P}_z \mathbf{W}] \end{aligned} \quad (43)$$

Where we identify 12 weight-Space Skinning Jacobian matrices:

$$\begin{aligned} \mathbf{A}_{1,1} &= \mathbf{P}_x \bar{X} & \mathbf{A}_{1,2} &= \mathbf{P}_x \bar{Y} & \mathbf{A}_{1,3} &= \mathbf{P}_x \bar{Z} & \mathbf{A}_{1,4} &= \mathbf{P}_x \\ \mathbf{A}_{2,1} &= \mathbf{P}_y \bar{X} & \mathbf{A}_{2,2} &= \mathbf{P}_y \bar{Y} & \mathbf{A}_{2,3} &= \mathbf{P}_y \bar{Z} & \mathbf{A}_{2,4} &= \mathbf{P}_y \\ \mathbf{A}_{3,1} &= \mathbf{P}_z \bar{X} & \mathbf{A}_{3,2} &= \mathbf{P}_z \bar{Y} & \mathbf{A}_{3,3} &= \mathbf{P}_z \bar{Z} & \mathbf{A}_{3,4} &= \mathbf{P}_z \end{aligned} \quad (44)$$

We can derive the 2D case of these matrices by following the equivalent steps as above, but with one less set of columns for  $\mathbf{I}$  and one less column in our rest positions  $\mathbf{X}$ :

$$\begin{aligned} \mathbf{A}_{1,1} &= \mathbf{P}_x \bar{X} & \mathbf{A}_{1,2} &= \mathbf{P}_x \bar{Y} & \mathbf{A}_{1,3} &= \mathbf{P}_x \\ \mathbf{A}_{2,1} &= \mathbf{P}_y \bar{X} & \mathbf{A}_{2,2} &= \mathbf{P}_y \bar{Y} & \mathbf{A}_{2,3} &= \mathbf{P}_y \end{aligned} \quad (45)$$

## C HYPER-REDUCED LOCAL GLOBAL SOLVERS FOR DIFFERENT ELASTIC ENERGIES

### C.1 Hyper-Reduced ARAP

The full space ARAP elastic energy can efficiently be written as:

$$E_{ARAP}(\mathbf{x}) = \sum_t^k m_t \mu_t \|F_t - R_t\|_F^2, \quad (46)$$

where  $F_t$  is our per-tet deformation gradient (a linear function of our degrees of freedom) and  $R_t$  is a best-fit rotation on the deformation gradient, a non-linear function of our degrees of freedom. Expanding the square Frobenius norm, we can identify a quadratic component and non-linear component to our energy:

$$\begin{aligned} E_{ARAP}(\mathbf{x}) &= \underbrace{\sum_t^k m_t \mu_t \text{tr}(F_t^T F_t)}_{\Psi} + \underbrace{\sum_t^k -2m_t \mu_t \text{tr}(F_t^T R_t)}_{\Phi} \\ &= \Psi(\mathbf{u}) + \Phi(\mathbf{u}). \end{aligned}$$

We can rewrite this energy in its hyper-reduced form by recalling  $\mathbf{x} = \mathbf{Bz}$  and approximating  $\Phi \approx \tilde{\Phi}$ . We expand our clustered non-linear component to the elastic energy:

$$\begin{aligned} E_{ARAP}(z) &= \Psi(z) + \underbrace{\sum_c^r -2m_c \mu_c \text{tr}(F_c^T R_c)}_{\tilde{\Phi}} \\ &= \Psi(z) + \tilde{\Phi}(z). \end{aligned} \quad (47)$$

We then expose the source of the non-linearities, the best fit per-cluster rotation matrices, as auxiliary degrees of freedom.

$$E_{ARAP}(z, \tilde{R}) = \Psi(z) + \tilde{\Phi}(z, \tilde{R}). \quad (48)$$

The hyper reduced local step becomes:

$$\begin{aligned} \tilde{R} &= \underset{\tilde{R}}{\text{argmin}} \tilde{\Phi}(z, \tilde{R}) \\ &= \underset{\tilde{R}}{\text{argmin}} \sum_c^k m_c \mu_c \text{tr}(F_c^T R_c) \\ &= \text{PolarSVD}(\tilde{F}) \end{aligned}$$

which is found through polar decomposition of the per-cluster deformation gradient  $\tilde{F}$ , obtainable via a reshape operation of  $\tilde{f} = \mathbf{x}^T \mathbf{K}^T \mathbf{G}$ . The hyper reduced global step minimizes

$$\begin{aligned} z &= \underset{z}{\text{argmin}} \Psi(z) + \tilde{\Phi}(z, \tilde{R}) \\ &= \underset{z}{\text{argmin}} z^T \mathbf{B}^T \mathbf{K}^T \mathbf{U} \mathbf{V} \mathbf{K} \mathbf{B} z + \tilde{\Phi}(z, \tilde{R}) \\ &= \underset{z}{\text{argmin}} z^T \mathbf{B}^T \mathbf{L} \mathbf{B} z + \tilde{\Phi}(z, \tilde{R}), \end{aligned}$$

where we introduce the heterogeneous Laplacian matrix  $\mathbf{L} = \mathbf{K}^T \mathbf{U} \mathbf{V} \mathbf{K}$  for brevity. Here,  $\mathbf{K} \in \mathbb{R}^{9k \times 3n}$  is the vector gradient operator over our mesh, mapping deformed coordinates to per-tet deformation gradients,  $\mathbf{V} \in \mathbb{R}^{9k \times 9k}$  is a diagonal matrix of tetrahedron volumes,  $\mathbf{U} = (\mathbf{I}_9 \otimes \text{diag}(\boldsymbol{\mu})) \in \mathbb{R}^{9k \times 9k}$  is a diagonal containing first lamé parameter  $\mu_t$  for each tet  $t$  along its diagonal.

For ARAP, the global-step optimization is *exactly* quadratic in the degrees of freedom  $z$  (because  $\tilde{\Phi}(z, \tilde{R})$  is linear in  $z$ ), and can directly be found every local-global iteration via a single prefactorizable system solve:

$$\begin{aligned} \mathbf{0} &= 2\mathbf{B}^T \mathbf{L} \mathbf{B} z + \frac{\partial \mathbf{x}}{\partial z} \frac{\partial f}{\partial \mathbf{x}} \frac{\partial \tilde{f}}{\partial f} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}} \\ &= \mathbf{B}^T \mathbf{L} \mathbf{B} z + \frac{1}{2} \mathbf{B}^T \mathbf{K}^T \mathbf{G} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}} \\ \mathbf{B}^T \mathbf{L} \mathbf{B} z &= -\frac{1}{2} \mathbf{B}^T \mathbf{K}^T \mathbf{G} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}}. \end{aligned}$$

Above, we made use of the linear relationships  $f = \mathbf{Kx}$ ,  $\mathbf{x} = \mathbf{Bz}$ , and  $\tilde{f} = \mathbf{G}_9 f$  to derive the chained partial derivatives  $\frac{\partial \mathbf{x}}{\partial z}$ ,  $\frac{\partial f}{\partial \mathbf{x}}$ ,  $\frac{\partial \tilde{f}}{\partial f}$ . The last unknown  $\frac{\partial \tilde{\Phi}}{\partial \tilde{f}}$  changes each local-global iteration and can efficiently be computed during the local step when looping through each cluster:

$$\frac{\partial \tilde{\Phi}}{\partial \tilde{f}} = \text{vec} \left( \begin{bmatrix} \vdots \\ -\frac{\partial}{\partial \tilde{F}_c} m_c \mu_c \text{tr}(\tilde{F}_c^T \tilde{R}_c) \\ \vdots \end{bmatrix} \right) = \text{vec} \left( \begin{bmatrix} \vdots \\ -m_c \mu_c \mathbf{R}_c \\ \vdots \end{bmatrix} \right) \quad \text{vec} \left( \begin{bmatrix} \vdots \\ -m_c \mu_c \mathbf{R}_c \\ \vdots \end{bmatrix} \right) \quad (49)$$

### C.2 Hyper-Reduced Co-Rotational Elasticity

A discrete linear co-rotational elastic energy is defined as

$$E_{CoRot}(\mathbf{u}) = \sum_t^k m_t \mu_t \|F_t - R_t\|_F^2 + m_t \frac{\lambda_t}{2} \text{tr}^2(R_t^T F_t - \mathbf{I})$$

Where  $F_t$  is our per-tet deformation gradient (a linear function of our degrees of freedom) and  $R_t$  is a best-fit rotation on the deformation gradient, a non-linear function of our degrees of freedom. Expanding the square Frobenius norm, we can identify a quadratic component and non-linear component to our energy:

$$\begin{aligned} E_{CoRot}(\mathbf{u}) &= \underbrace{\sum_t^k m_t \mu_t \text{tr}(F_t^T F_t)}_{\Psi} + \underbrace{\sum_t^k -2m_t \mu_t \text{tr}(F_t^T R_t) + m_t \frac{\lambda_t}{2} \text{tr}^2(R_t^T F_t - \mathbf{I})}_{\Phi} \\ &= \Psi(\mathbf{u}) + \Phi(\mathbf{u}). \end{aligned}$$

We can rewrite this energy in its hyper-reduced form by recalling  $\mathbf{x} = \mathbf{Bz}$  and approximating  $\Phi \approx \tilde{\Phi}$ . We expand our clustered non-linear component to the elastic energy:

$$\begin{aligned} E_{CoRot}(z) &= \Psi(z) + \underbrace{\sum_c^C -2m_c \mu_c \text{tr}(F_c^T R_c) + m_c \frac{\lambda_c}{2} \text{tr}^2(R_c^T F_c - \mathbf{I})}_{\tilde{\Phi}} \\ &= \Psi(z) + \tilde{\Phi}(z). \end{aligned}$$

We then expose the source of the non-linearities, the best fit per-cluster rotation matrices, as auxillary degrees of freedom.

$$E_{CoRot}(z, \tilde{\mathbf{R}}) = \Psi(z) + \tilde{\Phi}(z, \tilde{\mathbf{R}}).$$

The hyper reduced local step becomes:

$$\begin{aligned} \tilde{\mathbf{R}} &= \underset{\tilde{\mathbf{R}}}{\operatorname{argmin}} \tilde{\Phi}(z, \tilde{\mathbf{R}}) \\ &= \underset{\tilde{\mathbf{R}}}{\operatorname{argmin}} \sum_c^C m_c \mu_c \operatorname{tr}(F_c^T \mathbf{R}_c) + m_c \frac{\lambda_c}{2} \operatorname{tr}^2(\mathbf{R}_c^T F_c - \mathbf{I}) \\ &= \operatorname{PolarSVD}(\tilde{F}) \end{aligned}$$

which is found through polar decomposition of the per-cluster deformation gradient  $\tilde{F}$ , obtainable via a reshape operation of  $\tilde{f} = \mathbf{x}^T \mathbf{K}^T \mathbf{G}$ .

The hyper reduced global step minimizes:

$$\begin{aligned} z &= \underset{z}{\operatorname{argmin}} \Psi(z) + \tilde{\Phi}(z, \tilde{\mathbf{R}}) \\ &= \underset{z}{\operatorname{argmin}} z^T \mathbf{B}^T \mathbf{K}^T \mathcal{U} \mathbf{V} \mathbf{K} \mathbf{B} z + \tilde{\Phi}(z, \tilde{\mathbf{R}}) \\ &= \underset{z}{\operatorname{argmin}} z^T \mathbf{B}^T \mathbf{L} \mathbf{B} z + \tilde{\Phi}(z, \tilde{\mathbf{R}}) \end{aligned}$$

where we introduce the heterogeneous Laplacian matrix  $\mathbf{L} = \mathbf{K}^T \mathcal{U} \mathbf{V} \mathbf{K}$  for brevity. Here,  $\mathbf{K} \in \mathbb{R}^{9k \times 3n}$  is the vector gradient operator over our mesh, mapping deformed coordinates to per-tet deformation gradients,  $\mathbf{V} \in \mathbb{R}^{9k \times 9k}$  is a diagonal matrix of tetrahedron volumes,  $\mathcal{U} = (\mathbf{I}_9 \otimes \operatorname{diag}(\boldsymbol{\mu})) \in \mathbb{R}^{9k \times 9k}$  is a diagonal containing first lamé parameter  $\mu_t$  for each tet  $t$  along its diagonal.

The above optimization can be performed without recomputing full space second order terms with an iterative Quasi-Newton method, where each timestep we solve for the search direction  $z_{j+1} = z_j + \alpha dz$ :

$$\begin{aligned} \mathbf{0} &= 2\mathbf{B}^T \mathbf{L} \mathbf{B} dz + 2\mathbf{B}^T \mathbf{L} \mathbf{B} z_j + \frac{\partial f}{\partial z} \frac{\partial \tilde{f}}{\partial f} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}} \\ &= \mathbf{B}^T \mathbf{L} \mathbf{B} dz + \mathbf{B}^T \mathbf{L} \mathbf{B} z_j + \frac{1}{2} \mathbf{B}^T \mathbf{K}^T \mathbf{G} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}} \\ \mathbf{B}^T \mathbf{L} \mathbf{B} dz &= -\mathbf{B}^T \mathbf{L} \mathbf{B} z_j - \frac{1}{2} \mathbf{B}^T \mathbf{K}^T \mathbf{G} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}} \end{aligned}$$

where we have ignored any of the second order terms buried  $\frac{\partial \tilde{\Phi}}{\partial \tilde{f}}$ , which would normally incur a full space update in the system matrix every iteration. Because the search direction is no longer *perfect*, we make use of a back-tracking line search to figure out the step size we take along that search direction  $z_{j+1} = z_j + \alpha dz$ .

Above, we made use of the linear relationships  $f = \mathbf{K}x$ ,  $x = \mathbf{B}z$ , and  $\tilde{f} = \mathbf{G}_9 f$  to derive the chained partial derivatives  $\frac{\partial x}{\partial z}$ ,  $\frac{\partial f}{\partial x}$ ,  $\frac{\partial \tilde{f}}{\partial f}$ . The last unknown  $\frac{\partial \tilde{\Phi}}{\partial \tilde{f}}$  changes each local-global iteration and can efficiently be computed during the local step when looping through each cluster:

$$\begin{aligned} \frac{\partial \tilde{\Phi}}{\partial \tilde{f}} &= \operatorname{vec} \left( \begin{bmatrix} \frac{\partial}{\partial \tilde{F}_c} - m_c \mu_c \operatorname{tr}(\tilde{F}_c^T \tilde{\mathbf{R}}_c) + m_c \frac{\lambda_c}{2} \operatorname{tr}^2(\mathbf{R}_c^T F_c - \mathbf{I}) \\ \vdots \\ \vdots \\ \vdots \end{bmatrix} \right) \\ &= \operatorname{vec} \left( \begin{bmatrix} \vdots \\ -m_c \mu_c \mathbf{R}_c + m_c \frac{\lambda_c}{2} \mathbf{R}_c \operatorname{tr}(\mathbf{R}_c^T F_c - \mathbf{I}) \\ \vdots \end{bmatrix} \right) \end{aligned} \quad (50)$$

### C.3 Adding Hyper-Reduced Inertia

To add dynamics to our simulation, we add a Kinetic energy term to the elastic energies mentioned in the previous section. We employ a Backward Euler [Baraff and Witkin 1998; Liu et al. 2013] time integrator:

$$E_{Kinetic} = \frac{1}{2h^2} (\mathbf{x} - \mathbf{y})^T \mathbf{M} (\mathbf{x} - \mathbf{y})$$

Where  $\mathbf{y}$  represents position history  $\mathbf{x}_{hist} = 2\mathbf{x}_{curr} - \mathbf{x}_{prev}$  and  $h$  is the length of the timestep chosen.

We show we can represent this energy fully in our reduced space, without having to go to the full space. First we directly rewrite the final displacement in terms of our reduced complementary displacement, and our rigged position.  $\mathbf{x} = \mathbf{B}z + \mathbf{J}\mathbf{p}$

$$E_{Kinetic}(z, \mathbf{p}) = \frac{1}{2h^2} (\mathbf{B}z + \mathbf{J}\mathbf{p} - \mathbf{x}_{hist})^T \mathbf{M} (\mathbf{B}z + \mathbf{J}\mathbf{p} - \mathbf{x}_{hist})$$

Expanding this out and only keeping the terms that depend on our degrees of freedom  $z$  leads to:

$$= \frac{1}{2h^2} z^T \mathbf{B}^T \mathbf{M} \mathbf{B} z + \frac{1}{h^2} (z^T \mathbf{B}^T \mathbf{M} \mathbf{J} \mathbf{p} - z^T \mathbf{B}^T \mathbf{M} \mathbf{x}_{hist}) \quad (51)$$

Writing  $\mathbf{x}_{hist}$  in terms of our subspace and control rig:

$$\mathbf{y} = \mathbf{B}z_{hist} + \mathbf{J}\mathbf{p}_{hist}$$

Where  $z_{hist} = 2z_{curr} - z_{prev}$  and  $\mathbf{p}_{hist} = 2\mathbf{p}_{curr} - \mathbf{p}_{prev}$  are our reduced space complementary and rig displacement histories respectively.

Plugging the above into our reduced energy so far:

$$= \frac{1}{2h^2} z^T \mathbf{B}^T \mathbf{M} \mathbf{B} z + \frac{1}{h^2} z^T \mathbf{B}^T \mathbf{M} \mathbf{J} (\mathbf{p} - \mathbf{p}_{hist}) - \frac{1}{h^2} z^T \mathbf{B}^T \mathbf{M} \mathbf{B} z_{hist} \quad (52)$$

Adding this kinetic energy to the prior elastic energies, then minimizing would simply mean adding  $\frac{1}{h^2} \mathbf{B}^T \mathbf{M} \mathbf{B}$  to the quadratic terms, and then  $f_{inertia} = \frac{1}{h^2} z^T \mathbf{B}^T \mathbf{M} \mathbf{J} (\mathbf{p} - \mathbf{p}_{hist}) - \frac{1}{h^2} z^T \mathbf{B}^T \mathbf{M} \mathbf{B} z_{hist}$  to the linear terms, as shown in Algorithm 2.