

Active Animations of Reduced Deformable Models with Environment Interactions

ZHERONG PAN, University of North Carolina at Chapel Hill
DINESH MANOCHA, University of North Carolina at Chapel Hill

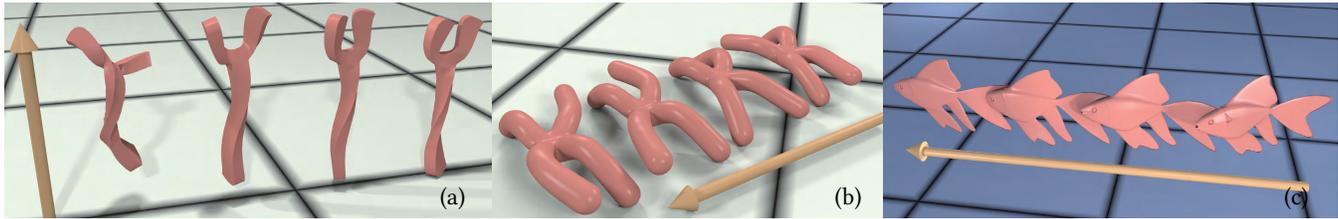


Fig. 1. Active deformable animations automatically generated by our approach: A letter T jumping (a), a spider walking (b), and a fish swimming (c). The reduced configuration spaces of these deformable bodies have 5 – 15 DOFs. We present an efficient spacetime optimization formulation that takes into account physics constraints and environmental interactions.

We present an efficient spacetime optimization method to automatically generate animations for a general volumetric, elastically deformable body. Our approach can model the interactions between the body and the environment and automatically generate active animations. We model the frictional contact forces using contact invariant optimization and the fluid drag forces using a simplified model. To handle complex objects, we use a reduced deformable model and present a novel hybrid optimizer to search for the local minima efficiently. This allows us to use long-horizon motion planning to automatically generate animations such as walking, jumping, swimming, and rolling. We evaluate the approach on different shapes and animations, including deformable body navigation and combining with an open-loop controller for realtime forward simulation.

CCS Concepts: • **Computing methodologies** → **Physical simulation**;

Additional Key Words and Phrases: deformable body, optimal control, locomotion

ACM Reference format:

Zherong Pan and Dinesh Manocha. 2017. Active Animations of Reduced Deformable Models with Environment Interactions. *ACM Trans. Graph.* 1, 1, Article 1 (January 2017), 17 pages.
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Physically-based deformable animation is a well-studied problem in computer graphics and related areas. Early methods such as [Müller and Gross 2004; Terzopoulos et al. 1987] focus on passive animations using numerical simulations. These techniques are widely used to generate plausible simulations of clothes [Bridson et al. 2002], plants [Barbič and Zhao 2011], human tissues [Chentanez et al.

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.

© 2017 Association for Computing Machinery.
0730-0301/2017/1-ART1 \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

2009], etc. Such passive animations are frequently used in movies and games to increase the realism. On the other hand, generating controlled or active deformable body animations [Coros et al. 2012; Kim and Pollard 2011; Tan et al. 2012] is considered more challenging, especially when a deformable body’s movements are governed by physics-based constraints. In such cases, additional control inputs, such as keyframes or rest shapes, need to be determined based on a deformable body’s interactions with the environment in order to generate the animation. This can be computationally challenging for deformable bodies with a high number of degrees of freedom (DOFs). To simplify the problem, previous methods [Hahn et al. 2012; Harmon and Zorin 2013; Kim and James 2011; Liu et al. 2013; Xu and Barbič 2016] partition the deformable body’s DOFs into controlled DOFs and uncontrolled DOFs. In practice, prior techniques specify the trajectories of controlled DOFs using manual keyframes and use physics-based simulation algorithms to generate movements corresponding to uncontrolled DOFs, i.e., the secondary dynamics. Such techniques are widely used for physical character rigging. In general, it is hard to generate controlled deformable body animations without user intervention or specifications. The animators not only need to manually partition the DOFs into controlled DOFs and the uncontrolled DOFs, but they also need to specify the movements of the controlled DOFs.

Main Results: We present a new method for active deformable body animations. The input to our method is a volumetric mesh representation of the body, a specification of the environment, and a high-level objective function that is used to govern the object’s movement. Our algorithm can automatically compute active animations of the deformable body and can generate motions corresponding to walking, jumping, swimming, or rolling, as shown in Figure 1. We compute the animations using a novel spacetime optimization algorithm and formulate the objective function taking into account dynamics constraints as well as various interactions with the environment. These include collisions, frictional contact forces and fluid drag forces. Compared with keyframe-based methods, we use objective functions to control the animation. In practice,

these objective functions are more general and easier for the user to specify. For example, to generate walking animation, user can just specify the target walking speed, instead of manually specifying the walking poses corresponding to different timesteps. Furthermore, our approach can be easily combined with partial keyframe data to provide more user control.

Some of the novel components of our work include:

- A spacetime optimization formulation based on reduced deformable models (Section 4), that takes into account environment interactions.
- A hybrid spacetime optimization algorithm (Section 5), which is more than an order of magnitude faster than previous spacetime optimization methods.
- We combine our spacetime optimization algorithm with dynamic movement primitives (DMP), which have been used in robotics [Schaal 2006]. DMP improves the performance of our algorithm in terms of avoiding suboptimal solutions. Furthermore, we present a two-stage animation framework. During the first stage, we compute the animation trajectories using DMP as a prior. These animations are then tracked and composed together at realtime using DMP as a controller (Section 6).

We demonstrate the benefits of our method by evaluating its performance on different complex deformable bodies with thousands of vertices and 5 – 15 DOFs in different environments (Section 6). For underwater swimming, we use DMP as an open-loop controller to generate realtime swimming animations (Figure 12). For contact-rich locomotion, the optimized animations are tracked at realtime using a feedback controller (Figure 19). Finally, we formulate keyframe-based control as a special case of our method and show animations controlled using partial keyframes and high-level control objectives (Figure 18).

2 RELATED WORK

Our work is inspired by prior work on passive/active deformable body animations and model reduction. In this section, we give a brief overview of related work.

Passive Deformable Body Animation has been an active area of research for more than three decades. The most popular deformable model, especially for deformable bodies without skeletons, is the finite element method (FEM) [Irving et al. 2006; Terzopoulos et al. 1987]. These methods of deformable body modeling have computational complexity that is superlinear in the number of discrete elements, and they therefore are not suitable for interactive applications. Many deformable bodies such as human bodies and animals have embedded rigid skeletons. Robust methods such as [Capell et al. 2002; Hahn et al. 2012; Kim and Pollard 2011] have been proposed to model skeletons’ interactions with soft tissues. In this paper, we use FEM to model a deformable body.

Active Deformable Body Animation is used by animators or artists to direct the animation while satisfying the physics-based constraints. Early works in this area [Barbič et al. 2009, 2012; Bergou et al. 2007; Schulz et al. 2014] try to make the deformable body follow a user-provided animation by applying external forces. However, deformable bodies in real life, such as worms, snakes, and fishes,

can only move themselves by generating internal forces. To respect this property, [Coros et al. 2012; Kim and Pollard 2011; Tan et al. 2012] control virtual deformable bodies to follow a given animation by applying internal forces only. Our work can be considered as complimentary to these methods. We generate animations that can be used as input to these previous methods, with a focus on reduced deformable models. Deformable body control methods can also be categorized based on the underlying user interfaces: [Barbič et al. 2009, 2012; Bergou et al. 2007; Kim and Pollard 2011; Schulz et al. 2014] require the user to specify a set of spacetime keyframes, while [Coros et al. 2012; Tan et al. 2012] and our approach specify the goals for controlling the animation using objective functions. We further show that conventional keyframe-based user-interface can be incorporated into our system using additional objective functions.

Spacetime Optimization Many techniques for deformable body animations are based on spacetime optimization [Witkin and Kass 1988]. Solving these optimization problems can be challenging due to the high-dimensional search space. Some algorithms parameterize the search space using low-dimensional representations such as splines [Hildebrandt et al. 2012] and functional spaces [Mukadam et al. 2016]. Another challenging issue is handling of non-smoothness constraints in the optimization formulation, due to environment interactions corresponding to collisions and contacts. Previous work either use sampling-based methods [Xu and Barbič 2016], complementarity constrained optimizations [Peng et al. 2017], or a smooth variant of the contact model [Mordatch et al. 2012, 2013]. In our work, we handle the high-dimensionality using reduced deformable models. We optimize the environmental forces by using or developing smooth variants of contact and fluid drag force models. Finally, we solve the optimization problem using a hybrid method.

Model Reduction is a practical method for fast deformable body simulations. It is based on the observation that only visually salient deformations need to be modeled. The earliest reduced model is based on Linear Modal Analysis (LMA) [Hauser et al. 2003; Pentland and Williams 1989], which is only accurate for infinitesimal deformations. Methods for non-linear and large deformations have been proposed in [An et al. 2008; Barbič and James 2005; Choi and Ko 2005]. In this paper, we use the rotation-strain space dynamic model [Pan et al. 2015] because it can preserve the key characteristics of deformable bodies with a lower-dimensional configuration space representation. However, our method can also be used with other reduced dynamic models.

3 PROBLEM FORMULATION

In this section, we formulate the problem of generating deformable body animation as a spacetime optimization. Our method searches in the space of deformable body animations with a fixed number of K timesteps. We denote an animation trajectory as: $\mathcal{Q}_K = (q_1 \cdots q_K)$, where each state vector q_i uniquely determines the position of a deformable body at time instance $i\Delta t$, where Δt is a fixed timestep size.

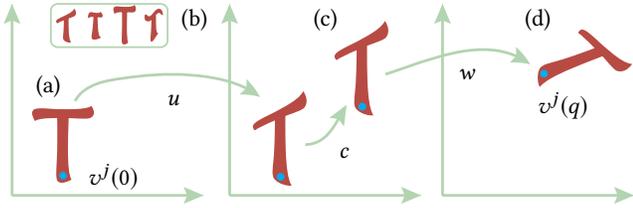


Fig. 2. The deformable body is represented as a triangle mesh (a). With a deformation bases set (b), its position is parameterized by local deformation u , a global rigid translation c (c), and rotation w (d). The Euclidean coordinates of j th vertex v^j (blue dot) can be recovered by transformation function $v^j(q)$.

3.1 Input & Output

In Section 5, we present an efficient optimizer to robustly search for the animation trajectory Q_K by solving the following problem:

$$\operatorname{argmin}_{Q_K, \mathcal{F}, \mathcal{E}} E(Q_K, \mathcal{F}, \mathcal{E}),$$

where, besides the animation trajectory Q_K , we also search for the internal forces (\mathcal{F}) and external forces (\mathcal{E}) on the deformable body at every time instance. Our overall algorithm takes the following components as an input:

- A volumetric mesh representation of the deformable body with V vertices $(v^1 \dots v^V)$.
- A specification of the environment, including type of the environment (in water, or on the ground) and parameters of the environment (e.g., drag force coefficient in water, or contact friction coefficient on the ground).
- The form of high level objective E_{obj} and its parameters. For example, in order for a deformable body to walk to a target position, E_{obj} will penalize the distance between current center of mass c and the target position, and its parameters correspond to the target position's coordinates.

3.2 Objective Function for Spacetime Optimization

By spacetime optimization, we assume that the desired deformable body animation corresponds to the local minima of an objective function. As a result, this objective function must encode all the requirements for a physically correct and plausible deformable body animation. We model these requirements by taking four different energy terms into account in $E(Q_K)$:

$$E(Q_K) = E_{\text{phys}} + E_{\text{obj}} + E_{\text{env}} + E_{\text{hint}}. \quad (1)$$

As outlined in Figure 3, the first term E_{phys} models all the shape changes that occur within the deformable body, i.e., the dynamics that result from the internal forces. It penalizes any violation in the deformable body's equation of motion (Section 4.1.1) and any collisions between different parts of the deformable body (Section 4.1.2). The second term E_{obj} is a task-dependent objective function specified by user (see Section 3.1). The environmental force term E_{env} (Section 4.2) models all the dynamic interactions between the deformable object and the environment, i.e. due to the external forces. It also penalizes any violation in the constraints that the environmental forces, such as frictional contact forces, must satisfy. Finally,

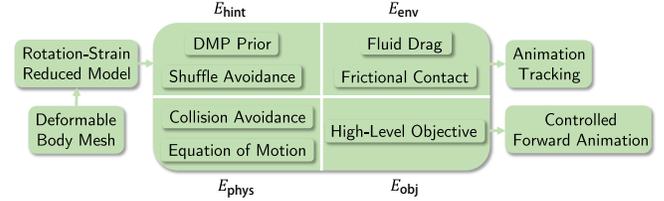


Fig. 3. We highlight various components used in our spacetime optimization algorithm that can also be used for animation tracking and controlled forward animation.

the last term E_{hint} (Section 4.3) guides the optimizer to avoid local minima that may result in less plausible animations.

3.3 Configuration Space Parametrization

Although our method can work with any parametrization of the deformable body's configuration q , different parametrizations result in drastically different computational cost. A straightforward method is to use volumetric meshes with V vertices $(v^1 \dots v^V)$ and define q as all the vertices' Euclidean coordinates. In this case, the dimension of the configuration space, $|q|$, scales linearly with the number of vertices and may be several thousands for moderately complex deformable models. Using optimization algorithms in such a large search space is only practical for very short animations. Indeed, [Bergou et al. 2007; Tan et al. 2012] used this vertex-based parametrization for tracking deformable body animation in a frame-by-frame manner, i.e., $K = 1$.

Instead, we represent the configuration of a deformable body using a rigid-coupled reduced model defined as:

$$q = (u \ c \ w)^T, \quad (2)$$

where u parametrizes the deformable body's non-rigid deformations in its local frame of reference. This is complemented with a rigid transformation in the world coordinates parametrized using a global translation c and rotation w , as illustrated in Figure 2. By using a precomputed dataset of deformation bases, the dimension of local deformation $|u|$ in Equation 2 is usually no more than 20. Moreover, methods such as cubature approximation [An et al. 2008] and fast sandwich transform (FST) [Kim and James 2011] can be used to efficiently recover a vertex v^j 's Euclidean coordinates using the transformation function $v^j(q)$. This transformation function can take a different form depending on the underlying reduced dynamic models. We refer readers to Section 6.2 for more analysis in terms of combining our method with different reduced dynamic models. A widely-known model is the reduced StVK [Barbič and James 2005]. Instead, we use the recently proposed rotation-strain (RS) space dynamic model [Pan et al. 2015] because it achieves comparable results with a lower-dimensional configuration space, i.e., a smaller $|u|$. We provide the details about the computation of $v(q)$ in Appendix C. We denote the reconstructed Euclidean coordinates representation as:

$$\bar{q}(q) = (v^1(q) \ v^2(q) \ \dots \ v^V(q))^T.$$

These formulations make it computationally tractable to numerically optimize a complex nonlinear function $E(Q_K)$. Moreover, Equation 2

is very convenient in terms of formulating our objective functions E_{obj} . For example, we could use a function in c to direct a deformable body to walk to a specific position, or a function in w to specify that a deformable body should stay balanced.

4 OBJECTIVE TERMS

In this section, we present details of the objective function used for spacetime optimization.

4.1 Physics-Based Constraints

The first term E_{phys} penalizes any violation of the equations of motion (EOM), and our formulation is similar to prior work [Barbič et al. 2009, 2012]. In addition, we also penalize any self-penetrations or collisions with static obstacles. Altogether, E_{phys} is represented as:

$$E_{\text{phys}}(\mathbf{Q}_K) = \sum_{i=2}^{K-1} E_{\text{EOM}}(q_{i-1}, q_i, q_{i+1}) + E_{\text{coll}}(q_i) + E_{\text{self}}(q_i).$$

4.1.1 Equations of Motion. Since q_i only represents a deformable body's position, E_{EOM} models the dynamic behavior using 3 consecutive frames. One advantage of this formulation is that we can use a position-based large timestep integrator [Hahn et al. 2012] to formulate our EOM. An implicit-Euler scheme determines q_{i+1} from q_{i-1} , q_i using the following optimization formulation:

$$q_{i+1} = \underset{q}{\operatorname{argmin}} \left[\frac{A(q)^T MA(q)}{2\Delta t^2} + P(q) - \mathcal{F}_i^T u - \mathcal{E}_i^T \tilde{q}(q) \right]$$

$$A(q) \triangleq \tilde{q}(q) - 2\tilde{q}(q_i) + \tilde{q}(q_{i-1}),$$

where M is the mass matrix constructed from the volumetric mesh of the deformable body using FEM, \mathcal{F} represents the internal control forces and \mathcal{E} corresponds to the environmental forces such as gravitational forces, fluid drag forces, and frictional contact forces. P is the elastic potential energy, and we model this energy term using the rotation-strain space linear elastic energy $P(q) = u^T \mathcal{K} u / 2$, where \mathcal{K} is the isotropic stiffness matrix. Even with an arbitrarily large Δt , the above time integrator is always stable. The term E_{EOM} is simply defined as the norm of gradient:

$$E_{\text{EOM}}(q_{i-1}, q_i, q_{i+1}) = \frac{1}{2} \left\| \frac{\partial \tilde{q}(q_{i+1})^T}{\partial q} MA(q_{i+1}) + \frac{\partial [P(q_{i+1}) - \mathcal{F}_i^T u_{i+1} - \mathcal{E}_i^T \tilde{q}(q_{i+1})]}{\partial q} \right\|^2, \quad (3)$$

where we exclude internal forces \mathcal{F}_i and external forces \mathcal{E}_i from our argument list to save space. These forces can be functions of q_i , eliminated analytically, or optimized as separate subproblems.

4.1.2 Collision Avoidance. Collision handling is regarded as a challenging problem in terms of deformable body simulation. In our method, we use two terms to approximately avoid collisions. For collisions with static obstacles, we formulate an energy term as:

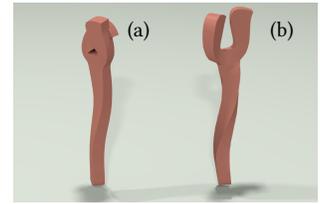
$$E_{\text{coll}}(q) = \frac{C_{\text{coll}}}{2} \sum_{j=1}^V \max(\text{dist}(v^j(q)), 0)^2, \quad (4)$$

where $\text{dist}(v)$ is the signed distance from a vertex's position to static obstacles, which is positive when there is penetration and negative

otherwise. To evaluate this function for vertices, we precompute a signed distance field for the static obstacles.

Handling self-collisions is even more challenging. In order to generate animations such as walking and jumping, many deformable bodies have thin structures that function as legs. Successful handling of self-collisions between such thin structures usually requires continuous collision detection (CCD), as illustrated in Figure 4. We make use of our reduced representation and use an approximate CCD scheme. Given a configuration $q = (u \ c \ w)$ that has self-penetrations, we first search for colliding pairs of vertices by reconstructing $\tilde{q}(q_i)$ from q_i and run a conventional discrete collision detection.

As shown in [Barbič and James 2010], considering only vertex-vertex collisions is enough for plausible handling of self-penetrations in reduced-model deformable body animations. Moreover, we observe that self-penetrations are invariant to the global rigid transformation



$(c \ w)$, so we only look at the local deformation component u of q . Since we already know that the undeformed configuration, i.e., $u = 0$, has no self-collisions, we can use a line-search algorithm in u to find the largest $\eta \in (0, 1]$ such that ηu has no self-collisions. We can then approximate the direction of separation as:

$$d^{ab} \triangleq \frac{v^a(\eta u) - v^b(\eta u)}{\|v^a(\eta u) - v^b(\eta u)\|},$$

this direction is treated as a fixed vector in each iteration of optimization, i.e. its derivatives are ignored. Finally, for each pair of vertices $v^a(u)$ and $v^b(u)$ in collision, we add an energy term:

$$E_{\text{self}}^{ab}(u) = \max((v^b(u) - v^a(u))^T d^{ab}, 0)^2,$$

and E_{self} is then defined as:

$$E_{\text{self}}(u) = \frac{C_{\text{self}}}{2} \sum_{a=1}^V \sum_{b=1}^V E_{\text{self}}^{ab}(u) I^{ab}, \quad (5)$$

where the last I^{ab} is an indicator of whether $v^a(q)$ and $v^b(q)$ are in collision.

4.2 Environmental Force Model

Since we allow only internal forces \mathcal{F} as the control input, a deformable body must make use of external environmental forces \mathcal{E} to move around. We consider two kinds of environmental forces: frictional contact forces and fluid drag forces. The frictional contact forces are used for generating contact-rich animations such as walking, balancing, rolling or jumping. The fluid drag forces are used for underwater swimming.

4.2.1 Frictional Contact Force Model. To model the frictional contact forces, we use contact invariant optimization (CIO) [Mordatch et al. 2012, 2013] and leave external forces \mathcal{E} as an additional optimizable variable. However, \mathcal{E} must satisfy two additional constraints. First, the contact force on vertex v^j , \mathcal{E}^j should lie inside

the frictional cone, we have:

$$\|\mathcal{E}_i^j\| \leq \mu \mathcal{E}_i^j \perp, \quad (6)$$

where \parallel and \perp are the tangent and normal component of the contact force, respectively, and μ is the frictional coefficient. The big advantage of CIO is that it allows the optimizer to jointly search for both contact forces and contact points by introducing the so-called contact-integrity term defined as:

$$E_{\text{env}}(q_i, q_{i-1}, \mathcal{E}_i^j) = C_{\text{env}} \sum_{i=2}^K \sum_{j=1}^V \quad (7)$$

$$(\|\text{dist}(v^j(q_i))\|^2 + \|(v^j(q_i) - v^j(q_{i-1}))\|^2) \|\mathcal{E}_i^j\|^2.$$

This term essentially encourages every external force \mathcal{E} to have maximal velocity dissipation and every contact point to stay on the contact manifold. In this work, we use a slightly different formulation from [Mordatch et al. 2013] and use a quadratic penalty for \mathcal{E} . In this way, the objective function $E(Q_K)$ becomes a quadratic function when we are optimizing only with respect to \mathcal{E} . Together with Equation 6, we can find the optimal \mathcal{E} , given Q_K , by solving a quadratic constrained QP (QCQP) problem. In Equation 7, the function $\text{dist}(\bullet)$ returns the closest distance to the environmental obstacles. We compute this efficiently by precomputing a distance field for the environment and we then use a smoothing algorithm [Calakli and Taubin 2011] so that $\text{dist}(\bullet)$ is C^1 -continuous.

4.2.2 Fluid Drag Force Model. The fluid drag forces, \mathcal{E}_i , are not free variables but functions of q_i, q_{i-1} . [Yuksel et al. 2007] used a quadratic drag force model, which is defined as a summation of forces on each triangular surface patch (v^a, v^b, v^c). Similar to the case with elastic energy, this quadratic drag force model has a corresponding potential energy defined follows:

$$\mathcal{E}_i^T \bar{q} = \mathcal{E}^T(q_{i+1}, q_i) \bar{q}(q_{i+1}) = C_{\text{drag}} \sum_{v^a, b, c} P_{abc}^{\text{drag}} \quad (8)$$

$$P_{abc}^{\text{drag}}(q_i, q_{i-1}) \triangleq \max(N_{abc}^T U_{abc}, 0) U_{abc}^T B_{abc}(q_{i+1}),$$

where we have also approximated the surface patch force as a point force on the barycenter B_{abc} . Here N_{abc} is the area-weighted normal and U_{abc} is the barycenter's relative velocity against fluid, as illustrated in Figure 5. Note that Equation 8 only takes effect when a surface patch is moving towards the fluid body. However, Equation 8 cannot be used by a gradient-based numerical optimizer because the gradient is discontinuous. We propose a continuous model by a slight modification:

$$P_{abc}^{\text{drag}}(q_i, q_{i-1}) \triangleq \max(N_{abc}^T U_{abc}, 0)^2 \frac{N_{abc}^T B_{abc}(q_{i+1})}{\|N_{abc}\|^2 + \epsilon},$$

which is C^1 -continuous, and we set $\epsilon = 10^{-6}$ to avoid degeneracy. This new model only relates drag forces with the normal component of the relative velocity. Since no other constraints or conditions are imposed on \mathcal{E} , we define $E_{\text{env}} = 0$ for the fluid drag model.

4.3 Controller Parametrization and Shuffle Avoidance

The two terms, $E_{\text{phys}}, E_{\text{env}}$, cannot uniquely determine an animation. Therefore, we add two terms that model the prior knowledge in plausible character animations: controller parametrization and shuffle avoidance.

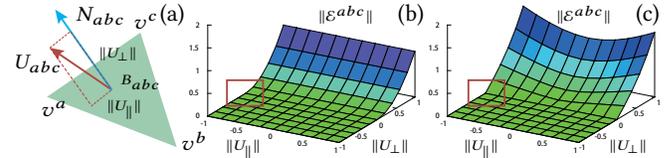


Fig. 5. Fluid drag force is applied on each surface patch (v^a, v^b, v^c). The force strength depends on the surface normal N and relative velocity U (a). We also plot the force strength with respect to tangential relative velocity U_{\parallel} and normal relative velocity U_{\perp} . Our new formulation is C^1 -continuous (b), while the original formulation has a discontinuous gradient (c), especially when the relative velocity is almost tangential (shown with a red rectangle).

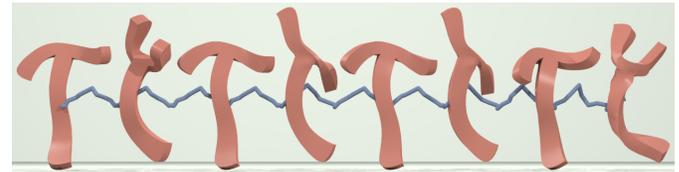


Fig. 6. A letter T jumping forward. With DMP regularization term E_{DMP} , its center of mass (blue) traces out a periodic trajectory.

4.3.1 Periodic and Temporal Smoothness. First, we notice that for several kinds of animations, including walking, swimming, and rolling, the deformable body should move in a periodic manner. Moreover, the desired animation is temporally smooth.

Algorithm 1 Algorithm to update C_{DMP} .

- 1: Evaluate $A \leftarrow \sum_{i=2}^{K-1} E_{\text{EOM}}(q_{i-1}, q_i, q_{i+1})$
- 2: Evaluate $B \leftarrow \sum_{i=2}^{K-1} \|\mathcal{F}_i^j - \text{DMP}(i\Delta t, \mathcal{W}_j)\|^2$
- 3: **if** $B > 0.1A$ **then** ▷ Control input does not match DMP
- 4: $C_{\text{DMP}} \leftarrow 2.1C_{\text{DMP}}$ ▷ Enforce better match
- 5: **end if**
- 6: **if** $B < 0.01A$ **then** ▷ Control input matches DMP
- 7: $C_{\text{DMP}} \leftarrow 0.5C_{\text{DMP}}$ ▷ Allow more animation explorations
- 8: **end if**

To respect this property, we use a general representation: Dynamic Movement Primitives (DMP) [Schaal 2006] to parameterize the control inputs. DMP is a special open-loop controller parametrization that can represent many complex robotic tasks such as tennis playing and walking. DMP is capable of representing both periodic and non-periodic tasks. The latter is useful, e.g., for jumping animations. A periodic DMP controller is defined as:

$$\text{DMP}_p(t, \mathcal{W}) = \sum_{n=1}^N \alpha_n \exp(\beta_n^2 \cos(\tau t - \mu_n)), \quad (9)$$

and a non-periodic DMP controller is defined as:

$$\text{DMP}_{np}(t, \mathcal{W}) = \sum_{n=1}^N \alpha_n \exp(-(\beta_n t - \mu_n)^2) t. \quad (10)$$

Note that DMP can be considered a special kind of one-input-one-output neural network using $\exp()$ and $\cos()$ as the activation function, where \mathcal{N} is the number of neurons in each layer and the

neural-net weights are $\mathcal{W} \triangleq (\alpha_n, \beta_n, \mu_n, \tau)$. In practice, we need one DMP function for each component of \mathcal{F}_i so that the total number of additional variables to be determined is $|\mathcal{W}| \times |\mathcal{F}_i| = |\mathcal{W}| \times |u|$. We denote the DMP for the j th component of \mathcal{F}_i using superscript j . In order to guide the optimizer to look for control inputs that can be represented using DMP, we introduce an additional energy term:

$$E_{\text{DMP}}(\mathcal{F}_i) = \frac{C_{\text{DMP}}}{2} \sum_{j=1}^{|\mathcal{F}_i|} \|\mathcal{F}_i^j - \text{DMP}_{p/np}(i\Delta t, \mathcal{W}_j)\|^2. \quad (11)$$

We simultaneously optimize Q_K and \mathcal{W}_j . We also adaptively adjust the weighting of this term so that E_{DMP} is almost zero after the iterative algorithm converges. As a result, the output of DMP function $\text{DMP}_{p/np}$ matches the required internal control forces \mathcal{F}_i^j exactly and $\text{DMP}_{p/np}$ can be used as an open-loop controller after spacetime optimization. To achieve such exact match between \mathcal{F}_i^j and $\text{DMP}_{p/np}$, we use a simple adaptive penalty method [Boyd et al. 2011]. Specifically, we use Algorithm 1 to adjust C_{DMP} after every iteration of optimization. Our scheme allows the optimizer to quickly explore the space of new animations, while keeping $\|\mathcal{F}_i^j - \text{DMP}(i\Delta t, \mathcal{W}_j)\|^2$ small. Figure 6 illustrates the effect of this heuristic term.

4.3.2 Shuffle Avoidance. As observed in [Mordatch et al. 2013], another artifact due to the lack of internal actuation structure is the shuffling movement across the contact manifold. This means that the contact points are always in close proximity to the solid boundary. To mitigate this artifact, we introduce an additional hint term E_{shuffle} defined as:

$$E_{\text{shuffle}}(q_i, q_{i-1}) = C_{\text{shuffle}} \sum_{j=1}^V \|(v^j(q_i) - v^j(q_{i-1}))\|^2 \exp(-\gamma \text{dist}(v^j(q_i))), \quad (12)$$

where γ is the distance attenuation coefficient. For each vertex v^j , we penalize its tangential velocity attenuated by its distance from static obstacles. In this way, the shuffling artifact is removed by asking a walker to lift its legs to move forward. The effect of this hint term is illustrated in Figure 7. We combine the above two hints and Tikhonov regularization, giving:

$$E_{\text{hint}}(Q_K) = \sum_{i=2}^{K-1} \frac{C_{\text{reg}}}{2} \|\mathcal{F}_i\|^2 + E_{\text{shuffle}}(q_i, q_{i-1}) + E_{\text{DMP}}(\mathcal{F}_i). \quad (13)$$

5 SPACETIME OPTIMIZATION

In this section, we present our efficient, hybrid optimizer to minimize the objective function:

$$\underset{Q_K, \mathcal{F}, \mathcal{E}, \mathcal{W}}{\text{argmin}} E(Q_K, \mathcal{F}, \mathcal{E}, \mathcal{W}), \quad (14)$$

where different subproblem solvers are used for minimizing with respect to each of the 4 free variables: $Q_K, \mathcal{F}, \mathcal{E}, \mathcal{W}$. As a special case, \mathcal{E} is not a free variable for swimming animations using our fluid drag model. Without loss of generality, we consider Equation 14 for presentation. Since the objective function is C^1 -continuous, our first attempt was to use an off-the-shelf implementation of the LBFGS algorithm [Liu and Nocedal 1989]. However, we found out that even for small problems, having very small $|u|$ and K , it takes a large number of iterations to converge. Instead, we present a

Algorithm 2 The hybrid optimization algorithm.

```

1: ▶ Setup multiple initial guesses for DMP periods
2:  $2\pi/\tau \leftarrow 0.2, 0.4, \dots, 5(s)$ 
3: for iteration=0,  $\dots$  do
4:   ▶ Evaluate function values and gradients
5:   Evaluate  $E(Q_K), \frac{\partial E(Q_K)}{\partial Q_K}$ 
6:   Evaluate an approximation to  $\frac{\partial^2 E(Q_K)}{\partial Q_K^2}$  denoted as  $H$ 
7:   ▶ Update trajectory, see [Lourakis 2005] for more details
8:   do
9:      $Q_K^{new} \leftarrow Q_K - (H + dI)^{-1} \frac{\partial E(Q_K)}{\partial Q_K}$ 
10:    if  $E(Q_K^{new}) > E(Q_K)$  then
11:      increase  $d$ 
12:    else
13:      decrease  $d$ 
14:    end if
15:    while  $E(Q_K^{new}) > E(Q_K)$ 
16:       $Q_K \leftarrow Q_K^{new}$ 
17:      ▶ Update contact forces (Equation 3,6,7)
18:      if Using frictional contact model then
19:        for  $i = 2, \dots, K-1$  in parallel do
20:          Update  $\mathcal{E}_i$  using [Todorov 2011]
21:        end for
22:      end if
23:      ▶ Update DMP weights for all initial guess (Equation 11)
24:      if  $\mathcal{W}$  is not updated in the last 10 iterations then
25:        for Each initial guess do
26:          Update DMP weights  $\mathcal{W}$  using 1000
27:          iterations of LBFGS [Liu and Nocedal 1989]
28:        end for
29:        Choose  $\mathcal{W}$  leading to smallest  $E_{\text{DMP}}$ 
30:      end if
31:      if Using frictional contact model then
32:        ▶ Adaptive contact handling (Equation 7, 12)
33:        for all  $1 \leq i \leq K$  and  $1 \leq j \leq V$  do
34:          if  $\text{dist}(v^j(q_i)) > \epsilon_1$  and  $\|\mathcal{E}_i^j\| < \epsilon_2 \max_{i,j} \|\mathcal{E}_i^j\|$  then
35:            Exclude  $v^j(q_i)$  from  $E_{\text{env,shuffle}}$ 
36:          else
37:            Include  $v^j(q_i)$  in  $E_{\text{env,shuffle}}$ 
38:          end if
39:        end for
40:      end if
41: end for

```

novel hybrid optimization algorithm that converges in much fewer iterations.

5.1 Hybrid Optimizer

To accelerate the rate of convergence, we first notice that \mathcal{F} appears only in E_{EOM} , and E_{DMP} as a quadratic function. Therefore, we can solve for \mathcal{F} analytically and eliminate it. We further observe that the other three sets of variables ($Q_K, \mathcal{E}, \mathcal{W}$) appear in the objective function with special structures. The DMP weight vector \mathcal{W} appears only in E_{DMP} and optimizing \mathcal{W} amounts to a small neural-network training problem for which LBFGS is proven to be very effective.

Example	LBFGS(s)	Hybrid(s)
2D Crawling	1534	18
2D Rolling	823	12

Table 1. Performance of LBFGS and our hybrid solver on two examples: 2D worm crawling and 2D ball rolling. Our approach is significantly faster.

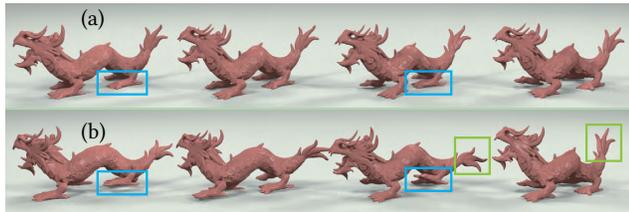


Fig. 7. Frames highlighting the dragon walking trajectory using our approach. In the result without E_{shuffle} (a), the dragon's foot is always in close proximity to the floor plane (blue). The artifact is mitigated using E_{shuffle} (b). Therefore, we observe more secondary dynamics in (b), e.g., around the tail (green).

The external force \mathcal{E} is a quadratic function in both E_{EOM} and E_{env} , and \mathcal{E}_i for each timestep i is separable and can be solved in parallel. Together with constraint Equation 6, finding the optimal \mathcal{E} amounts to solving a QCQP problem, for which special solvers are known. For example, we use a primal interior-point method [Todorov 2011]. We found that solving QCQP is faster than solving QP with a linearized frictional model because it requires fewer constraints and makes use of coherence in the intermediate solutions between the consecutive iterations by allowing warm-starting. We can update these variables \mathcal{W} , \mathcal{E} , and Q_K in an alternate manner. Finally, for trajectory Q_K itself, LBFGS can still be used, but we found that LBFGS does not use gradient information effectively. A large number of gradient evaluations are performed inside the line-search scheme and LBFGS usually chooses a conservative step size. Therefore, we choose the Levenberg-Marquardt (LM) method for updating Q_K . The outline of our method is given in Algorithm 2, and we provide the low-level details in Appendix B. Table 1 shows a comparison between our solver and LBFGS on two small 2D problems.

5.2 Efficient Function Evaluation

The costliest step in our algorithm is the evaluation of the function values, including the gradient and approximate hessian. We combine several techniques to accelerate these evaluations. First, we notice that $v^j(q_i)$, the recovered vertex j 's Euclidean coordinates from reduced representation at timestep i , appears in almost every objective term. Moreover, these values are independent of each other. Therefore, we can compute and store $v^j(q_i)$, $\frac{\partial v^j}{\partial q_i}$, $\frac{\partial^2 v^j}{\partial q_i^2}$ for all $1 \leq i \leq K$ and $1 \leq j \leq V$ in parallel, before each evaluation. We provide some hints for computing the second derivatives $\frac{\partial^2 v^j}{\partial q_i^2}$ in Appendix C. The overhead of these computations is independent of the number of vertices V . We also utilize this information to assemble the hessian. This assembly step can be a computational

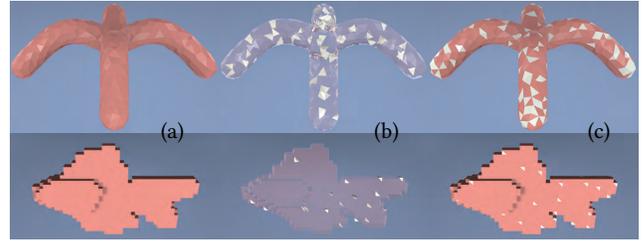


Fig. 8. For the spider (top) and fish (bottom) models (a), we visualize the kinetic cubatures (b) and surface patch cubatures (c). In both cases, only a small fraction of elements need to be considered for the summation. This fraction is 12% for the spider and 0.7% for the fish model.

bottleneck because we have to evaluate the summations over all the vertices that appear in the physics violation term E_{EOM} , in the collision avoidance term $E_{\text{coll,self}}$, in the environmental force term E_{env} , and finally in the shuffle avoidance term E_{shuffle} .

5.2.1 Accelerating the Assembly of $E_{\text{coll,self}}$. For collision avoidance terms, only very few vertices will contribute non-zero values to the objective function. Therefore, we use a bounding volume hierarchy [James and Pai 2004] to update the non-zero terms. This data-structure can be updated solely using reduced representation $|q|$, and the update for different timesteps can be performed in parallel.

5.2.2 Accelerating the Assembly of $E_{\text{env,shuffle}}$. In the previous section, we used cubature approximation to accelerate the fluid drag forces. For frictional contact forces, however, all the vertices in close proximity to the static obstacles will contribute non-zero values to E_{env} and E_{shuffle} . Since these vertices cannot be determined during the precomputation stage, we dynamically update them. Specifically, we remove vertex $v^j(q_i)$ from E_{env} if $\text{dist}(v^j(q_i)) > \epsilon_1$ and $\|\mathcal{E}_i^j\| < \epsilon_2 \max_{i,j} \|\mathcal{E}_i^j\|$. After E_{env} is updated, we update E_{shuffle} accordingly, since E_{shuffle} is also very small for vertices that are far from the static obstacles. These updates can be accelerated using a bounding volume hierarchy.

5.2.3 Accelerating the Assembly of E_{EOM} . [An et al. 2008; Barbič and James 2005] have addressed the problem of accelerating the assembly of E_{EOM} . Specifically, summation over all vertices appears in two places of E_{EOM} highlighted below:

$$\frac{\partial \bar{q}(q_{i+1})^T}{\partial q} MA(q_{i+1}) + \frac{\partial \left[P - \mathcal{F}_i^T u - \mathcal{E}_i^T \bar{q} \right]}{\partial q} (q_{i+1}).$$

We use the cubature approximation [An et al. 2008] to accelerate these two terms. The blue part above corresponds to the kinetic cubature used in [Pan et al. 2015]; see Appendix C.2 for more details. The red part above corresponds to the fluid drag force, which is a summation over all the surface patches. Cubature approximation assumes that:

$$\mathcal{E}^T(q_{i+1}, q_i) \bar{q}(q_{i+1}) = C_{\text{drag}} \sum_{v^{a,b,c}} P_{abc}^{\text{drag}} \approx \sum_{v^{a,b,c} \in \mathcal{T}} C_{abc} P_{abc}^{\text{drag}}, \quad (15)$$

i.e., the sum over all surface patches can be approximated using the

weighted sum of a selected set of surface patches \mathcal{T} . As illustrated in Figure 8, this greatly reduces the computational overhead.

The set \mathcal{T} and weights C_{abc} are computed via dictionary learning. Specifically, we first construct a dataset of D exemplary deformations $\{q_1, q_2, \dots, q_D\}$ according to Gaussian distribution as described as in [An et al. 2008]. Then for each pair of q_i, q_{i+1} , we compute the groundtruth fluid drag force. The problem of finding \mathcal{T} and C_{abc} is formulated as the following sparse coding problem:

$$\begin{aligned} \underset{C_{abc}}{\operatorname{argmin}} \quad & \|C_{abc}\|_0 + \sum_{i=1}^{D-1} \left\| \sum_{\nu=a,b,c} \frac{\partial [C_{\nu abc} P_{abc}^{\text{drag}} - C_{abc} P_{abc}^{\text{drag}}]}{\partial q_{i+1}} \right\|^2 \\ \text{s.t.} \quad & C_{abc} \geq 0, \end{aligned} \quad (16)$$

which is solved using iteratively reweighted Lasso optimization as in [Pan et al. 2015]. We take partial derivative in Equation 16 so that we are measuring error in drag force $\frac{\partial P_{abc}^{\text{drag}}}{\partial q}$, instead of the corresponding potential energy P_{abc}^{drag} . Finally, \mathcal{T} consists of all surface patches with non-zero C_{abc} .

5.3 Robustness to Suboptimal Solutions

It is well-known that spacetime optimization is prone to bad local minima leading to suboptimal solutions, except for simple cases [Barbič et al. 2012]. In our algorithm, there are two energy terms that can result in the computation of bad local minima. One is the contact integrity term, E_{env} , which models the non-smoothness of frictional contacts. The other one is E_{DMP} , which models the trajectory smoothness and periodic movements.

In terms of E_{DMP} , previous methods [RÅijckert and D'Avella 2013; Schaal 2006] use sampling-based methods to search for the global optimum. Since we only use a gradient-based local optimizer, E_{DMP} could result in the computation of a bad local minima. Indeed, we found that our optimizer can have difficulty in terms of finding good DMP parameters \mathcal{W} . At a local minima, several DMP neurons usually have same values of $(\alpha_n, \beta_n, \mu_n)$, values in Equation 9 or Equation 10, meaning that we are wasting parameters. In addition, we found that the period parameter τ can get stuck in a local minima very close to our initial guess. In this section, we introduce some simple modifications to overcome these problems.

We first initialize the phase shift uniformly in the phase space, i.e., $\mu_n = 2\pi n/N$, and we initialize α_n, β_n to very small random values. To avoid the period parameter τ falling into a bad local minima, we use multiple initial guess for τ and run an LBFGS optimization from each initial guess in Line 25 of Algorithm 2. In our experiments, we set $2\pi/\tau = 0.2, 0.4, \dots, 5(\text{s})$ and run LBFGS 25 times very 10 iterations to avoid bad local minima.

After that, we get 25 candidate DMP parameters, \mathcal{W} , and we choose

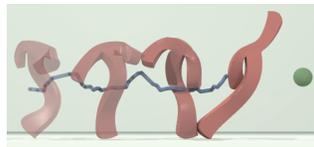


Fig. 9. When user sets a target point (green) too far away (5 meters to the right) and uses very few timesteps (20 in this case), the Letter T chooses to lean itself too much to recover from falling down.

the candidate leading to the smallest E_{DMP} . Such multiple LBFGS optimizations will result in additional computational overhead during the first few iterations of optimization. As the optimizer gets closer to a local minima, τ will converge to a same local minima for several candidates of DMP parameters, and we can merge these candidates into one. In addition, if a certain candidate is never chosen as the best during the last 100 iterations, we remove this candidate from further consideration. In practice, we have only 2 – 3 remaining candidates after 500 iterations.

The approach highlighted above greatly increases the chances that our optimization algorithm computes a good local minima without significant computational overhead. This is because periodic DMP formulation (Equation 9) is guiding the whole trajectory to follow a same gait. When our optimization algorithm finds a useful gait, this information is quickly encoded into the DMP controller and reused to compute the entire trajectory using the E_{DMP} formulation. In order to highlight this feature, we show two swimming trajectories computed using our optimization algorithm. In order to compute the trajectory shown in Figure 10 (a), we initialize the spider pose to $u = c = w = 0$ at all timesteps. While to generate Figure 10 (b), we initialize the spider to a different random pose at every timestep. Moreover, the convergence history of these two optimization schemes are plotted in Figure 11. Our optimizer converges to two different but almost equally effective swimming gaits with very small objective function values. This means that although there are numerous local minima, most of them leads to plausible animations. However, bad local minima can still happen especially in contact-rich animations and we illustrate one such failure case in Figure 9.

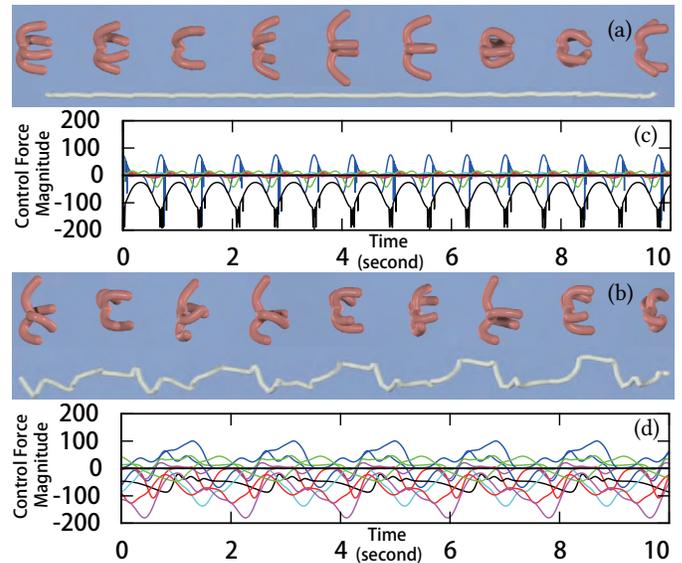


Fig. 10. We show two swimming trajectories optimized using static initialization (a) and random initialization (b). For both trajectories, we plot the locus of the deformable body's center of mass (white curve) and the magnitude of control forces in (c,d). The goal is to move 5 meters to the left after 10 seconds. Our optimizer finds two different but almost equally effective gaits.

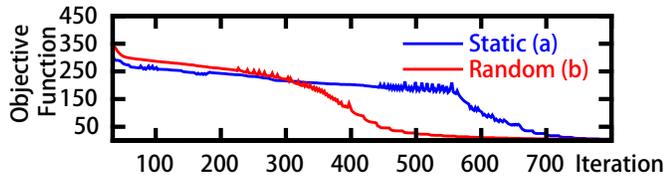


Fig. 11. We plot the convergence history using static initialization (a) and random initialization (b). Both optimizations reduce the objective function to less than 1% of the original value. This plot shows that many local minima of our objective function leads to plausible animations. There are some jittering during the optimization. This is because the adaptive penalty method (Algorithm 1) is adjusting C_{DMP} .

Name	Value	Name	Value
C_{coll} (Equation 4)	10^2	C_{reg} (Equation 13)	10^{-3}
C_{self} (Equation 5)	10^2	ϵ_1 (Section 5.2.2)	$3l$
μ (Equation 6)	0.7	ϵ_2 (Section 5.2.2)	0.01
C_{env} (Equation 7)	$10^1/l^2$	D (Equation 16)	1000
C_{drag} (Equation 8)	10^3	Δt	0.05s
N (Equation 9, 10)	5	Young's modulus	$10^5 Pa$
C_{DMP} (Equation 11)	dynamic	Poisson's ratio	0.48
C_{shfle} (Equation 12)	10^{-1}	Mass density ρ	$1kg/m^3$
γ (Equation 12)	$\log(10)/l$	Gravity	$9.81m/s^2$

Table 2. Parameters.

6 RESULTS

In this section, we highlight the results on complex benchmarks.

Parameter Choices: We use an identical set of parameters listed in Table 2 for all the benchmarks. The coefficient of the physics violation term is 1. Some parameters are related to l , which is the average element size. If a deformable body has volume vol and is discretized using P FEM elements, then $l = (vol/P)^{1/3}$. An exception is the coefficient for E_{DMP} , which is adaptively adjusted within the optimization algorithm.

Benchmarks: We implemented our method in C++ and tested it on many benchmarks using a desktop machine with dual E5-2670 12-core CPU 2.1GHz and 12GB of memory. Given only a volumetric mesh and a definition of the environment, we first precompute the reduced dynamic model using [Pan et al. 2015]. We also precompute the surface cubatures to approximate the fluid drag forces. We use OpenMP to parallelize the function and gradient evaluations and run at most 10000 iterations of optimizations or stop early, if the relative error of $\|\frac{\partial E(Q_K)}{\partial Q_K}\|$ is smaller than 10^{-3} . The setup and computational cost in each benchmark is summarized in Table 3 and analyzed below.

Fish Swimming: Fishes have the simplest deformable bodies and can be used for testing the performance of our method. As illustrated in Figure 1, a fish swims by simply swinging its body, so we use a reduced configuration space of small DOFs: $|u| = 5$, i.e., $|q| = 11$. Under this setting, we command the fish to swim straight forward in a gravityless environment using the following objective:

Example	V/P	$ u $	Pre./PreSF.(min)	$K/\#Q_K$	Opt.(hr)	App.
Fish Swimming (Fig.12a)	2118/7812	5	0.8/0.1	200/3	1.7	DMP
Spider Swimming (Fig.12b)	1054/4033	10	1.2/0.3	200/3	2.5	DMP
Spider Swimming	1054/4033	65	4.6/0.3	200/1	3.9	None
Reduced StVK (Fig.21a)	1054/4033	10	1.2/	200/4	5.2	FB
Dragon Walking (Fig.7)	929/1854	10	1.3/	200/1	2.2	None
Letter T Walking (Fig.19a)	1523/3042	15	1.1/	200/4	4.5	FB
Letter T Walking	1523/3042	65	4.2/	200/1	3.1	None
Reduced StVK (Fig.21b)	1024/640	10	1.1/	100/1	1.1	None
Beam Jumping (Fig.14)	623/1499	10	1.3/	200/1	2.1	None
Dinosaur Walking (Fig.18)	1493/5249	15	0.9/	200/1	1.9	None

Table 3. Benchmark setup and computational overhead. From left to right, number of vertices V /number of FEM elements P , DOFs of local deformation $|u|$, precomputation time for building reduced dynamic model/computing surface patch cubatures, number of frames/number of trajectories, time spent on optimization, and the supported application: DMP means we use DMP as open-loop controller to drive forward simulation, FB means that we use feedback controller to track the animation (both of these are realtime).

$$E_{obj}^{move}(Q_K) = C_{obj}^{move} \sum_{k=2}^K \|c_{k+1} - c_k - \exp(w_{k+1})v_c \Delta t\|^2/2, \quad (17)$$

where $\exp(w_{k+1})$ transforms the velocity to a global frame of reference and v_c is the target swimming speed in a local frame of reference. In addition, we add a balance energy to encourage fixed orientation:

$$E_{obj}^{bal}(Q_K, d) = C_{obj}^{bal} \sum_{k=2}^K \|\exp(w_k)d - d\|^2/2, \quad (18)$$

where d is the balance direction. Here we use $d = g$, the unit gravitational direction. We can even navigate the fish to an arbitrary 3D point by optimizing 3 trajectories: swimming forward, swimming left, and swimming right. For swimming left and right, we add the following objective functions in addition to Equation 17:

$$E_{obj}^{turn}(Q_K, d) = C_{obj}^{turn} \sum_{k=2}^K \|\exp(w_{k+1}) - \exp(\theta d \Delta t) \exp(w_k)\|^2/2, \quad (19)$$

where θd is the target rotating speed, we use $d = g$ again. After the optimization, the DMP function can be used as an open-loop controller to generate controlled forward simulations at real-time framerate. In Figure 12 (a), we wrap our forward simulator into a sampling-based motion planner, RRT* [Karaman and Frazzoli 2011], to navigate the fish to look for food plants.

Spider Swimming: We also evaluated our approach on a more complex model: a four-legged spider. More degrees of freedom are used to allow each leg to move independently, so we use more DOFs: $|u| = 10$, i.e., $|q| = 16$. Again, we optimize to generate 3 trajectories with the same τ and μ_n for all trajectories. However, for the first trajectory we set $v_c = 1, \theta = 0$, and for the other two we set $v_c = 0, \theta = \pm 1$ so that the spider cannot turn itself around while swimming forward. This gives very different gaits for turning and swimming forward. We again use DMP to drive the realtime forward simulator and wrap it into a motion planner, as illustrated in Figure 12 (b).



Fig. 12. We use RRT* to navigate physical swimming characters, the fish (a) and the spider (b), to look for food plants (green). The white line is the locus of the deformable body's center of mass computed using RRT*.

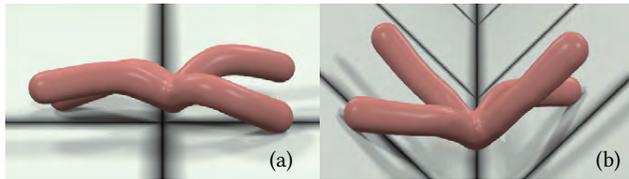


Fig. 13. The spider walking on planar ground (a) and V-shaped ground (b).

Spider Walking: To analyze the walking animation, we use the same spider model and objective Equation 17 but replace the fluid drag force model with the frictional contact force model. However, we observe that this optimization takes approximately twice as many iterations to converge due to the contact-integrity term E_{env} and the shuffle avoidance term $E_{shuffle}$. In Figure 13, we illustrate the walking gaits for two kinds of environments.

Similar to swimming, we allow a user to navigate the spider on the ground by optimizing 4 trajectories: walking left, right, backward, and forward, where the objective function is Equation 19 with corresponding v_c . We then use a feedback controller similar to [Tan et al. 2012] to drive forward simulator. Specifically, we optimize $E(Q_K)$ over one timestep ($K = 3$) with the objective function:

$$E_{obj}^{track}(Q_K) = C_{obj}^{track} (\bar{q}(q_2) - \bar{q}(q^*))^T M (\bar{q}(q_2) - \bar{q}(q^*)) / 2, \quad (20)$$

where q_k^* is the configuration of the tracked trajectory. Due to the efficiency of reduced representation, such short-horizon optimization can be solved at realtime framerates.

Letter T Walking: A more challenging example is Letter T walking, as illustrated in Figure 6. This model has no static stability, so it must keep jumping to move around. Again, we first optimize 4 trajectories and then track these trajectories at realtime to navigate the character.

Beam Jumping: Jumping is an essential component in many animations. To generate these animations, we use the following objective function:

$$E_{obj}^{jump}(Q_K) = C_{obj}^{jump} \|g^T c_{K/2} - h\|^2 / 2 + \quad (21)$$

$$C_{obj}^{jump} \|(\mathbf{I} - gg^T)(c_{K/2} - c_{K/2-1} - \exp(w_{K/2})v_c \Delta t)\|^2 / 2,$$

where the first term specifies the target altitude and the second term specifies the target horizontal velocity v_c so that the character can jump forward. Using different h and v_c , we generate a series of results in Figure 14 and Figure 15 for a small beam, where the beam exhibits huge and varied deformations.

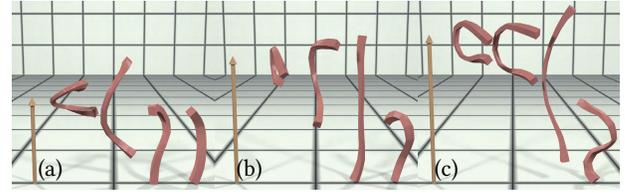


Fig. 14. Different frames during a beam jumping with different target altitudes (yellow arrow); (a): $h = 2$, (b): $h = 3$, and (c): $h = 4$.

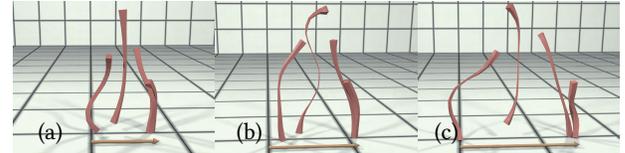


Fig. 15. Different frames during a beam jumping forward with different target distance (yellow arrow); (a): $\|v_c\| = 2$, (b): $\|v_c\| = 2.5$, and (c): $\|v_c\| = 3$ (the target altitude $h = 3$).



Fig. 16. (a): An X-shaped deformable body walking by rolling. (b): The rest shape.

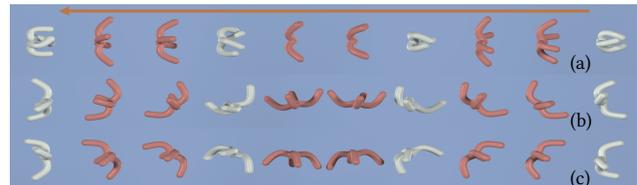


Fig. 17. To compute the navigation path for the spider, we optimize 3 trajectories: swimming forward (a), turning left (b), and turning right (c) (the timestep index increases along the arrow, and the white bodies mark the most deformed configurations). These differences in the gaits can be represented by different DMP parameters α_n and β_n only.

Cross Rolling: As illustrated in Figure 16 (b), we generate a rolling animation for a cross-shaped deformable body by combining the objective functions given in Equation 18 and Equation 19.

6.1 Combining our Algorithm with Partial Keyframe Data

Although our main contribution is a control framework that does not require keyframes, we can easily take keyframes into consideration to provide more flexibility to a user. These keyframes can either be specified fully or partially. A full keyframe specifies a target position for each of the V vertices, while a partial keyframe

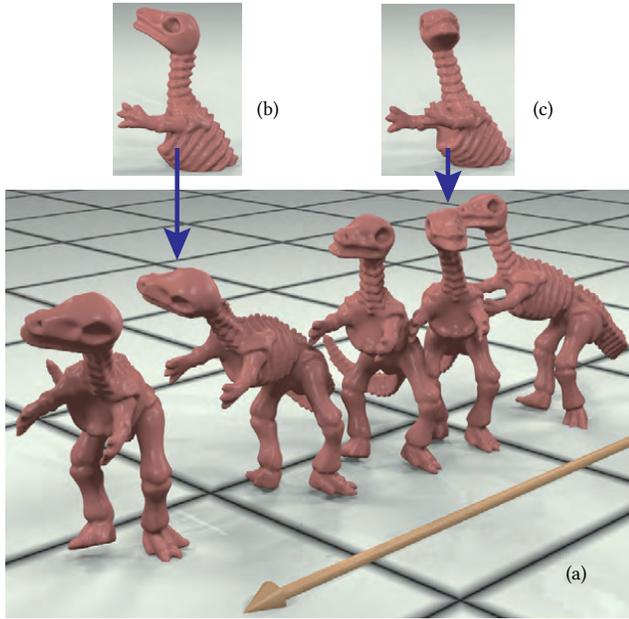


Fig. 18. We show a walking dinosaur guided by both our high-level objectives and user-specified keyframes in (a) so that its head is looking around. And the upper-body partial keyframes are illustrated in (b,c). Each of these keyframes contains 788 of the 1493 vertices.

only specifies a target position for a subset of vertices on the deformable body. For example, in Figure 18 (a), we show a dinosaur walking on the ground with its head swinging periodically to the left and right. The dinosaur's head is guided by a set of M partial keyframes illustrated in Figure 18 (b). The keyframes only specify the head and torso poses and we leave the leg poses to be determined by other objective function terms. We denote these keyframes as $u_1^{\text{key}}, \dots, u_M^{\text{key}}$ specified at timesteps t_1, \dots, t_M . Note that these keyframes only specify the dinosaur's deformable poses u and do not affect the global transformation (c, w). The keyframe guiding is achieved using an additional objective function:

$$E_{\text{obj}}^{\text{key}}(Q_K) = C_{\text{obj}}^{\text{key}} \sum_{i=1}^M \|I(q(u_i^{\text{key}}) - q(u_{t_i}))\|^2 / 2, \quad (22)$$

where I is an importance-weighting matrix allowing the users to specify partial keyframes. In our example, I is a diagonal matrix with diagonal value 1 around the head and torso (788 vertices) and 0 elsewhere.

6.2 Analysis

We summarize objective functions used in all benchmarks in Table 4 and analyze several aspects of our method.

Two-Stage Algorithm: A drawback of our method is that the optimization formulation takes a more complex form, and the resulting optimization algorithm takes longer time than [Barbič et al. 2009]. Fortunately, the DMP function returned by the optimizer can be used as a swimming controller to generate more swimming animation at realtime, as illustrated in Figure 12. This makes our method much more useful than a simple keyframe interpolation.

Example	E_{obj}
Fish Swimming	$E_{\text{obj}} = E_{\text{obj}}^{\text{move}} + E_{\text{obj}}^{\text{turn}} + E_{\text{obj}}^{\text{bal}}, v_c = (1\ 0\ 0)^T, \theta = 0, \pm 1, d = g$
Spider Swimming	$E_{\text{obj}} = E_{\text{obj}}^{\text{move}} + E_{\text{obj}}^{\text{turn}} + E_{\text{obj}}^{\text{bal}}, v_c = (1\ 0\ 0)^T, \theta = 0$ or $v_c = (0\ 0\ 0)^T, \theta = \pm 1, d = g$
Walking	$E_{\text{obj}} = E_{\text{obj}}^{\text{move}} + E_{\text{obj}}^{\text{bal}}, v_c = (\pm 1\ 0\ \pm 1)^T, d = g$
Dinosaur Walking	$E_{\text{obj}} = E_{\text{obj}}^{\text{move}} + E_{\text{obj}}^{\text{bal}} + E_{\text{obj}}^{\text{key}}, v_c = (\pm 1\ 0\ \pm 1)^T, d = g$
Jumping	$E_{\text{obj}} = E_{\text{obj}}^{\text{jump}} + E_{\text{obj}}^{\text{bal}}, h = 2/3/4, v_c = (2/2.5/3\ 0\ 0)^T, d = g$
Rolling	$E_{\text{obj}} = E_{\text{obj}}^{\text{turn}} + E_{\text{obj}}^{\text{bal}}, v_c = (1\ 0\ 0)^T, \theta = 1, d = v_c \times g / \ v_c \times g\ $
	$C_{\text{obj}}^{\text{move}} = C_{\text{obj}}^{\text{turn}} = C_{\text{obj}}^{\text{jump}} = C_{\text{obj}}^{\text{key}} = 10^{-1}$ and $C_{\text{obj}}^{\text{bal}} = 10^{-2}$

Table 4. Objective function used in each benchmark.

However, to generate realtime contact-rich animations, such as walking and jumping, we have to use a feedback controller instead of DMP controller. This is because the contact forces are very sensitive to the discrepancy between forward simulation model and the physics model used in spacetime optimization (model discrepancy).

Multi-Tasking: In order to make the realtime animations directable, we need to simultaneously optimize multiple animation trajectories to allow a motion planner to pick trajectory online. However, if we sequentially run separate optimizations, the generated gaits can be quite different, e.g., the fish might swing its body with different frequencies to swim in different directions. This artifact can be mitigated if we use the same DMP parameters τ and μ_n for all the trajectories to ensure the same period of movement and phase shift, i.e., DMPs differ only in α_n and β_n for different tasks. This idea has been previously used for DMP-based reinforcement learning [Rājickert and D'Avella 2013]. As illustrated in Figure 17, DMP can represent large gait differences using different α_n and β_n only, while the rhythms of the movements are synchronized. We use this strategy in all the navigation examples.

Quality Measure: For jumping animation, we do not require any manual bases design such as basis expansion [Tan et al. 2012]. The reason is that we formulated physics constraints as soft constraints and physics constraints are violated for small tracking errors. To measure the violation to EOM at each frame (q_{i-1}, q_i, q_{i+1}) , we first solve Equation 3 using q_{i-1}, q_i to find a physically correct q_{i+1}^* by calling a conventional deformable body simulator. Next, we recover the vertex positions from q_{i+1}, q_{i+1}^* to compute $\bar{q}(q_{i+1}), \bar{q}(q_{i+1}^*)$. We measure the EOM violation by computing the average distance between $\bar{q}(q_{i+1}), \bar{q}(q_{i+1}^*)$ of V vertices, and comparing it with the average FEM element size l :

$$\sqrt{\|\bar{q}(q_{i+1}) - \bar{q}(q_{i+1}^*)\|^2 / V / l}.$$

According to the plot in Figure 19, the physics violation over the whole trajectory is always less than half of average element size and is neglectable. The physics violation data for other examples can be found in Appendix A. However, manual bases design can sometimes be needed. For example, very different rolling gaits are generated in Figure 16, by restricting the bases to the 2D plane.

Effect of Different Parameters: Instead of using keyframes, the result of our algorithm depends on two sets of parameters. A first set of parameters are listed in Table 2. These parameters are considered internal and not exposed to users. The second set of parameters, which are listed in Table 4, are exposed to users.

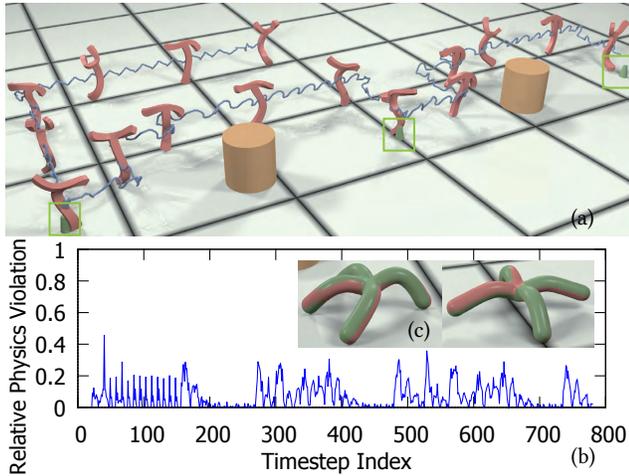


Fig. 19. (a): Navigating the letter T across obstacles (yellow) to reach goal positions (green). (b): A plot of the relative physical violation in (a) with respect to average FEM element size. Throughout the trajectory, the violations to EOM are very small. (c): Two most violated frames in the spider walking trajectory (q_{i+1}^* drawn in green); the differences are indistinguishable.

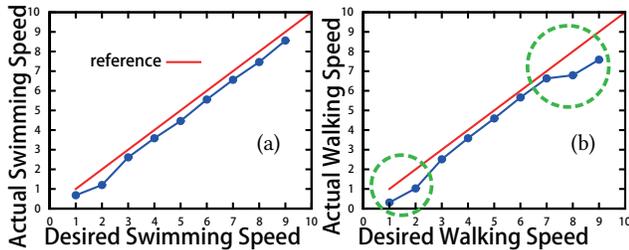


Fig. 20. We generate 9 spider walking trajectories (a) and 9 spider swimming trajectories (b), using different target moving speed v_c . We plot the actual walking/swimming speed compared with the desired speed. These actual values are very close to desired values. However, there are still cases where the discrepancy between actual and desired values are large (green circles).

These parameters have clear meanings such as walking, swimming, or rolling speed. In Figure 20, we highlight the effectiveness of performing animation control. We generated 9 walking/swimming trajectories using different target moving speed v_c . Since we model E_{obj} as a soft penalty, the desired speed cannot be achieved exactly. However, according to Figure 20, the discrepancy between actual and desired moving speeds are very small. Therefore, we expose more parameters to the users compared with keyframe-based methods [Barbič et al. 2009; Schulz et al. 2014], these parameters have intuitive meanings and are helpful for animation control.

We also noticed two cases from Figure 20 (green circles) where the discrepancy between desired and actual moving speed are relatively large. If the desired speed is too small, then our optimizer considers E_{obj} as unimportant and it is given lower importance in order to reduce the residue in other objective terms. If the desired speed is too large, it can result in self-collisions or the optimizer falls into a bad local minima, as shown in Figure 9.

External vs. Internal Control Forces: Theoretically, our algorithm only uses internal forces. To understand this, note that internal forces are forces whose values and torques sum up to zero. In other words, internal forces are forces that cannot change c , w of q . Since a conservative force equals the negative partial derivatives of the corresponding potential energy, any potential energy with zero derivative against c , w corresponds to the internal forces. This is the case with our potential energy $\mathcal{F}^T u$ in Equation 3.

However, since we formulate the physical correctness as a penalty term, E_{EOM} , rather than a hard constraint, there is some residual $E_{EOM} \neq 0$ at the local minima. This residual can be interpreted as a violation of the physical correctness, or as a ghost external force.

If we write $E_{EOM}(q_{i-1}, q_i, q_{i+1}) = \|\mathcal{E}_i^{ghost T} \bar{q}\|$, we are actually controlling the deformable body using both internal force \mathcal{F}_i and an additional ghost external force \mathcal{E}_i^{ghost} , but our objective function is designed to guide the optimizer to search for a solution with minimal ghost external force magnitude.

Using soft penalty instead of hard constraints also allows us to generate realtime deformable body animations by tracking an optimized animation. For example, having the letter-T balanced on a single contact point in Figure 19 is very challenging, which usually requires control over long horizons. However, with soft penalty, we can track the animation by control over only one timestep using Equation 20 as the objective function.

Other Reduced Models: Although we choose [Pan et al. 2015], our method can also work with other reduced models. This can be done by modifying the transformation function $\bar{q}(q)$ and the kinetic energy P in Equation 3. In Appendix D, we analyze the case with two kinds of different but widely used reduced models: LMA [Pentland and Williams 1989] and reduced StVK [Barbič and James 2005]. And two examples are illustrated in Figure 21 (a,b).

A drawback of these alternative models is that they require a higher-dimensional configuration space to achieve similar results as [Pan et al. 2015]. In our experiments, we use $|u| = 65$ and each optimization becomes 3-5 times slower according to Table 1. However, from the plots of DMP control force magnitudes, Figure 21 (c,d), we notice that the optimal \mathcal{F} is actually very sparse. In other words, much computations are wasted on looking for small, unimportant control forces. Such analysis suggests that [Pan et al. 2015] is a better choice.

Although our formulation can also work with fullspace deformable models by replacing \bar{q} with identity function, this approach can be computationally very expensive. As reported in [Pan et al. 2015], using a reduced model accelerates the evaluation of $\bar{q}(q)$ by two orders of magnitude. In our experiments, cubature accelerates the evaluation of fluid drag forces by at least an order of magnitude using Equation 15. Since function evaluation is the major bottleneck of spacetime optimization, we expect it will take weeks or even months to finish an optimization using fullspace models. As a result, it is important to use reduced deformable models for efficiency reasons.

Other (Partial) Keyframe-based Methods: Our method allows a user to specify partial keyframes to control parts of a deformable body that does not involve environmental interactions. A different approach is proposed in [Schulz et al. 2014], which uses partial

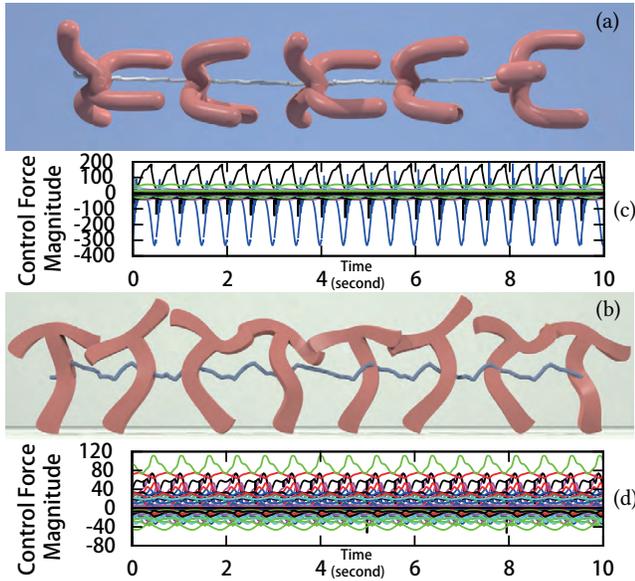
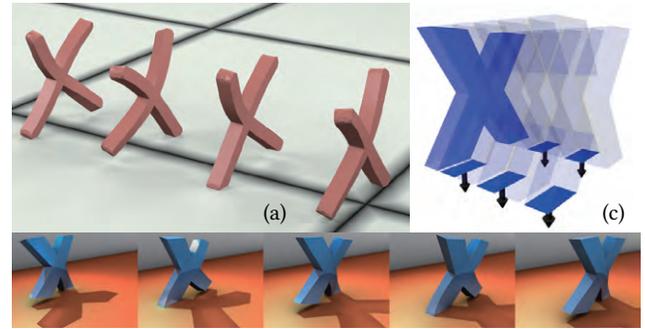


Fig. 21. We computed the animations corresponding to spider swimming (a) and letter T walking (b) using reduced StVK [Barbič and James 2005] as the underlying deformable model. In addition, we plot the control force magnitude for the spider in (c) and for the letter T in (d). We use different curves to plot each \mathcal{F}^j . Since $|u| = 65$, we have 65 curves in each plot. However, most of these curves are centered around zero axis. This means that the magnitude of control forces for most j are very small and \mathcal{F} is quite sparse.

keyframes to manually specify parts that involve environmental interactions. Overall, our approach is complementary and generates different deformations or animations. The two benchmarks correspond to cross-shape walking (Figure 22) and cactus jumping (Figure 23) and we have compared the animations in the video. In both these examples, we encode the objective into the functions shown in Equation 17 and Equation 21. In practice, our approach requires less user input. We also observed that our deformation patterns in jumping animation are quite different from [Schulz et al. 2014].

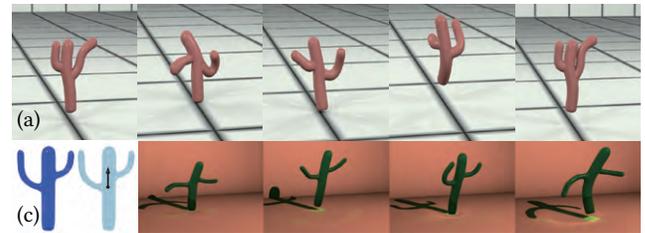
7 LIMITATIONS AND FUTURE WORK

We present a method to automatically generate active animations of reduced deformable bodies, where the user provides a high-level objective and the animation is generated automatically using spacetime optimization. We take into account physics constraints, environmental forces in terms of CIO and fluid drag models, and DMP-based controller parametrization, so that the local minima of our objective function corresponds to a plausible animation. By evaluating objective functions and function gradients in a subspace, the optimization can be accomplished within several hours on a single desktop machine. Although optimization is offline, the results can be used to generate animations at realtime rates. For swimming animations, the optimized DMPs can be used as a controller for forward simulation. Unfortunately, DMP cannot be used as controllers for contact-rich animations. Since DMP is not a feedback controller, model discrepancy can quickly accumulate, leading to



(b): Image courtesy of [Schulz et al. 2014]

Fig. 22. We compare our method with partial keyframe-based method. Our method can generate a cross-shape walking motion (a) similar to the results (b) in [Schulz et al. 2014] obtained using keyframe set (c).



(b): Image courtesy of [Schulz et al. 2014]

Fig. 23. Our method generating a cactus-jumping motion (a) similar to the results (b) in [Schulz et al. 2014] obtained using keyframe set (c). However, our deformation patterns are quite different compared with those in [Schulz et al. 2014].

failures such as falling. In these cases, DMP is just used as a periodic and smoothness prior.

Our approach has some limitations. First, our method inherits all the limitations of the underlying reduced model. For example, current reduced model [Pan et al. 2015] cannot work with user specified skeletons. Working with skeletons is a desirable feature in terms of modelling some animal-like deformable bodies, such as the fish, where deformable tissues are covering skeletal bones. In addition, although our method requires no keyframes or user designs, we still ask the users to choose the form of E_{obj} and their parameters in Section 6. Moreover, without the keyframes, our animations may not exhibit the same level of naturalness as some prior keyframe-based methods [Barbič et al. 2009]. For example, the dinosaur's walking gait is not symmetric in Figure 18 due to the asymmetric rest pose. And it can be difficult to generate the handstand example of [Schulz et al. 2014] in Figure 23. Moreover, our optimizer may get stuck in a bad local minima due to insufficient DOFs of the reduced configuration space, a sub-optimal bases set, or an inappropriate settings of the weights. Furthermore, the inherent limitations of CIO term [Mordatch et al. 2013] for contact modeling and the fluid drag model can also affect our results. For example, we cannot have a deformable body bouncing off the ground since the CIO term only models inelastic contacts. CIO also allows inexact contacts to occur anywhere in the air, not only on

the ground. Finally, like all the optimization-based motion planners, the performance of our method is still governed by a large set of parameters. Some parameters, such as the number of DMPs (N), are determined empirically. We have not evaluated the sensitivity of our method with respect to these parameters.

There are avenues for future work. First, incorporating some body-specific priors can be helpful in several ways. For example, for many muscle-driven deformable bodies, the user might want to parameterize the controller using muscle-tendon units [Wang et al. 2012] to generate more life-like animations. Another part that may benefit from user interactions is the identification of deformation bases in Figure 2. Currently, we identify these components using standard techniques [Pan et al. 2015] that are designed for visual simulation. However, it is not known if a base set for plausible visual simulation is suitable for character locomotion. It is also attractive to consider the optimization method as a general feedback controller, instead of an open-loop controller, for reduced deformable models using reinforcement learning [Peng et al. 2017]. Finally, developing control methods for two-way couple deformable body and articulated body will provide more flexibility to users. A starting point can be [Xu and Barbić 2016].

ACKNOWLEDGEMENT

This research is supported in part by Army Research Office and National Science Foundation.

REFERENCES

- Steven S. An, Theodore Kim, and Doug L. James. 2008. Optimizing Cubature for Efficient Integration of Subspace Deformations. In *ACM SIGGRAPH Asia 2008 Papers (SIGGRAPH Asia '08)*. ACM, New York, NY, USA, Article 165, 10 pages. <https://doi.org/10.1145/1457515.1409118>
- Jernej Barbić, Marco da Silva, and Jovan Popović. 2009. Deformable Object Animation Using Reduced Optimal Control. In *ACM SIGGRAPH 2009 Papers (SIGGRAPH '09)*. ACM, New York, NY, USA, Article 53, 9 pages. <https://doi.org/10.1145/1576246.1531359>
- Jernej Barbić and Doug L. James. 2005. Real-Time Subspace Integration for St. Venant-Kirchhoff Deformable Models. In *ACM SIGGRAPH 2005 Papers (SIGGRAPH '05)*. ACM, New York, NY, USA, 982–990. <https://doi.org/10.1145/1186822.1073300>
- Jernej Barbić and Doug L. James. 2010. Subspace Self-collision Culling. *ACM Trans. Graph.* 29, 4, Article 81 (July 2010), 9 pages. <https://doi.org/10.1145/1778765.1778818>
- Jernej Barbić, Funshing Sin, and Eitan Grinspun. 2012. Interactive Editing of Deformable Simulations. *ACM Trans. Graph.* 31, 4, Article 70 (July 2012), 8 pages. <https://doi.org/10.1145/2185520.2185566>
- Jernej Barbić and Yili Zhao. 2011. Real-time Large-deformation Substructuring. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 91, 8 pages. <https://doi.org/10.1145/1964921.1964986>
- Miklós Bergou, Saurabh Mathur, Max Wardetzky, and Eitan Grinspun. 2007. TRACKS: Toward Directable Thin Shells. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07)*. ACM, New York, NY, USA, Article 50. <https://doi.org/10.1145/1275808.1276439>
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.* 3, 1 (Jan. 2011), 1–122. <https://doi.org/10.1561/22000000016>
- Robert Bridson, Ronald Fedkiw, and John Anderson. 2002. Robust Treatment of Collisions, Contact and Friction for Cloth Animation. *ACM Trans. Graph.* 21, 3 (July 2002), 594–603. <https://doi.org/10.1145/566654.566623>
- F. Calakli and G. Taubin. 2011. SSD: Smooth Signed Distance Surface Reconstruction. *Computer Graphics Forum* 30, 7 (2011), 1993–2002. <https://doi.org/10.1111/j.1467-8659.2011.02058.x>
- Steve Capell, Seth Green, Brian Curless, Tom Duchamp, and Zoran Popović. 2002. Interactive Skeleton-driven Dynamic Deformations. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '02)*. ACM, New York, NY, USA, 586–593. <https://doi.org/10.1145/566570.566622>
- Nuttapong Chentanez, Ron Alterovitz, Daniel Ritchie, Lita Cho, Kris K. Hauser, Ken Goldberg, Jonathan R. Shewchuk, and James F. O'Brien. 2009. Interactive Simulation of Surgical Needle Insertion and Steering. In *ACM SIGGRAPH 2009 Papers (SIGGRAPH '09)*. ACM, New York, NY, USA, Article 88, 10 pages. <https://doi.org/10.1145/1576246.1531394>
- 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 (Jan 2005), 91–101. <https://doi.org/10.1109/TVCG.2005.13>
- Stelian Coros, Sebastian Martin, Bernhard Thomaszewski, Christian Schumacher, Robert Sumner, and Markus Gross. 2012. Deformable Objects Alive! *ACM Trans. Graph.* 31, 4, Article 69 (July 2012), 9 pages. <https://doi.org/10.1145/2185520.2185565>
- Guillermo Gallego and Anthony Yezzi. 2015. A Compact Formula for the Derivative of a 3-D Rotation in Exponential Coordinates. *Journal of Mathematical Imaging and Vision* 51, 3 (March 2015), 378–384. <https://doi.org/10.1007/s10851-014-0528-x>
- Fabian Hahn, Sebastian Martin, Bernhard Thomaszewski, Robert Sumner, Stelian Coros, and Markus Gross. 2012. Rig-space Physics. *ACM Trans. Graph.* 31, 4, Article 72 (July 2012), 8 pages. <https://doi.org/10.1145/2185520.2185568>
- David Harmon and Denis Zorin. 2013. Subspace Integration with Local Deformations. *ACM Trans. Graph.* 32, 4, Article 107 (July 2013), 10 pages. <https://doi.org/10.1145/2461912.2461922>
- Kris K. Hauser, Chen Shen, and James F. O'Brien. 2003. Interactive Deformation Using Modal Analysis with Constraints. In *Graphics Interface*. CIPS, Canadian Human-Computer Communication Society, 247–256. <http://graphics.cs.berkeley.edu/papers/Hauser-IDU-2003-06/>
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2012. Interactive Spacetime Control of Deformable Objects. *ACM Trans. Graph.* 31, 4, Article 71 (July 2012), 8 pages. <https://doi.org/10.1145/2185520.2185567>
- G. Irving, J. Teran, and R. Fedkiw. 2006. Tetrahedral and Hexahedral Invertible Finite Elements. *Graph. Models* 68, 2 (March 2006), 66–89. <https://doi.org/10.1016/j.gmod.2005.03.007>
- Doug L. James and Dinesh K. Pai. 2004. BD-tree: Output-sensitive Collision Detection for Reduced Deformable Models. In *ACM SIGGRAPH 2004 Papers (SIGGRAPH '04)*. ACM, New York, NY, USA, 393–398. <https://doi.org/10.1145/1186562.1015735>
- Sertac Karaman and Emilio Frazzoli. 2011. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research* 30, 7 (2011), 846–894. <https://doi.org/10.1177/0278364911406761>
- Jungwon Kim and Nancy S. Pollard. 2011. Fast Simulation of Skeleton-driven Deformable Body Characters. *ACM Trans. Graph.* 30, 5, Article 121 (Oct. 2011), 19 pages. <https://doi.org/10.1145/2019627.2019640>
- Theodore Kim and Doug L. James. 2011. Physics-based Character Skinning Using Multi-domain Subspace Deformations. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '11)*. ACM, New York, NY, USA, 63–72. <https://doi.org/10.1145/2019406.2019415>
- Dong C. Liu and Jorge Nocedal. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical Programming* 45, 1 (1989), 503–528. <https://doi.org/10.1007/BF01589116>
- Libin Liu, KangKang Yin, Bin Wang, and Baining Guo. 2013. Simulation and Control of Skeleton-driven Soft Body Characters. *ACM Trans. Graph.* 32, 6, Article 215 (Nov. 2013), 8 pages. <https://doi.org/10.1145/2508363.2508427>
- Manolis IA Lourakis. 2005. A brief description of the Levenberg-Marquardt algorithm implemented by levmar. (2005).
- Igor Mordatch, Emanuel Todorov, and Zoran Popović. 2012. Discovery of Complex Behaviors Through Contact-invariant Optimization. *ACM Trans. Graph.* 31, 4, Article 43 (July 2012), 8 pages. <https://doi.org/10.1145/2185520.2185539>
- Igor Mordatch, Jack M. Wang, Emanuel Todorov, and Vladlen Koltun. 2013. Animating Human Lower Limbs Using Contact-invariant Optimization. *ACM Trans. Graph.* 32, 6, Article 203 (Nov. 2013), 8 pages. <https://doi.org/10.1145/2508363.2508365>
- M. Mukadam, X. Yan, and B. Boots. 2016. Gaussian Process Motion planning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. 9–15. <https://doi.org/10.1109/ICRA.2016.7487091>
- Matthias Müller and Markus Gross. 2004. Interactive Virtual Materials. In *Proceedings of Graphics Interface 2004 (GI '04)*. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 239–246. <http://dl.acm.org/citation.cfm?id=1006058.1006087>
- Zherong Pan, Hujun Bao, and Jin Huang. 2015. Subspace Dynamic Simulation Using Rotation-strain Coordinates. *ACM Trans. Graph.* 34, 6, Article 242 (Oct. 2015), 12 pages. <https://doi.org/10.1145/2816795.2818090>
- Xue Bin Peng, Glen Berseth, Kangkang Yin, and Michiel Van De Panne. 2017. DeepLoco: Dynamic Locomotion Skills Using Hierarchical Deep Reinforcement Learning. *ACM Trans. Graph.* 36, 4, Article 41 (July 2017), 13 pages. <https://doi.org/10.1145/3072959.3073602>
- A. Pentland and J. Williams. 1989. Good Vibrations: Modal Dynamics for Graphics and Animation. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 207–214. <https://doi.org/10.1145/74334.74355>
- Elmar RÄijckert and Andrea D'Avella. 2013. Learned parametrized dynamic movement primitives with shared synergies for controlling robotic and musculoskeletal systems. *Frontiers in Computational Neuroscience* 7 (2013), 138. <https://doi.org/10.3389/fncom.2013.00138>

Model	$(X \times Y \times Z)$	l	Model	$(X \times Y \times Z)$	l
Fish	$2 \times 1.2 \times 0.2$	0.071	Dragon	$1.52 \times 0.71 \times 1.03$	0.075
Spider	$1.65 \times 0.51 \times 1.65$	0.079	Beam	$0.2 \times 2 \times 0.2$	0.14
Letter T	$1.00 \times 1.24 \times 0.1$	0.037	Cross	$1 \times 1 \times 0.1$	0.042

Table 5. Model parameters used in our experiment. From left to right, model name, bounding box size $(X \times Y \times Z)$, and average FEM element size l .

- Stefan Schaal. 2006. Dynamic movement primitives—a framework for motor control in humans and humanoid robotics. In *Adaptive motion of animals and machines*. Springer, 261–280.
- Christian Schulz, Christoph von Tycowicz, Hans-Peter Seidel, and Klaus Hildebrandt. 2014. Animating Deformable Objects Using Sparse Spacetime Constraints. *ACM Trans. Graph.* 33, 4, Article 109 (July 2014), 10 pages. <https://doi.org/10.1145/2601097.2601156>
- Jie Tan, Greg Turk, and C. Karen Liu. 2012. Soft Body Locomotion. *ACM Trans. Graph.* 31, 4, Article 26 (July 2012), 11 pages. <https://doi.org/10.1145/2185520.2185522>
- Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. 1987. Elastically Deformable Models. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '87)*. ACM, New York, NY, USA, 205–214. <https://doi.org/10.1145/37401.37427>
- E. Todorov. 2011. A convex, smooth and invertible contact model for trajectory optimization. In *2011 IEEE International Conference on Robotics and Automation*. 1071–1076. <https://doi.org/10.1109/ICRA.2011.5979814>
- Jack M. Wang, Samuel R. Hammer, Scott L. Delp, and Vladlen Koltun. 2012. Optimizing Locomotion Controllers Using Biologically-based Actuators and Objectives. *ACM Trans. Graph.* 31, 4, Article 25 (July 2012), 11 pages. <https://doi.org/10.1145/2185520.2185521>
- Andrew Witkin and Michael Kass. 1988. Spacetime Constraints. In *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '88)*. ACM, New York, NY, USA, 159–168. <https://doi.org/10.1145/54852.378507>
- Hongyi Xu and Jernej Barbic. 2016. Pose-space Subspace Dynamics. *ACM Trans. Graph.* 35, 4, Article 35 (July 2016), 14 pages. <https://doi.org/10.1145/2897824.2925916>
- Cem Yuksel, Donald H. House, and John Keyser. 2007. Wave Particles. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07)*. ACM, New York, NY, USA, Article 99. <https://doi.org/10.1145/1275808.1276501>

A PHYSICS VIOLATION

We provide the physics violation data for all the examples in Table 6. We use the same criterion:

$$\sqrt{\|\bar{q}(q_{i+1}) - \bar{q}(q_{i+1}^*)\|^2 / V / l}.$$

Compared with the average FEM element size l , the error due to physics violation is very small. The parameters for all the volumetric deformable models are provided in Table 5.

B OPTIMIZATION ALGORITHM

In this section we summarize our main Algorithm 2. The skeleton of this algorithm is an LM optimizer. We refer readers to [Lourakis 2005] for a brief introduction of this simple method. In our experiments, LM performs much better than LBFGS for two reasons. First, LM uses $J^T J$ approximation of the hessian which usually leads to a better step size estimation. Second, although LM needs to search for the so-called damping coefficient as in the process of the line-search scheme, it requires only function value evaluations which are much faster than gradient evaluations used by the line-search scheme of LBFGS.

However, a minor problem with using an LM algorithm is that it assumes the objective function is a sum of squares, which is violated by E_{shuffle} . In Appendix B.1, we show that a marginal modification to [Lourakis 2005] will allow us to handle E_{shuffle} .

After every successful LM iteration, we update \mathcal{E} ; we update \mathcal{W} every 10 iterations. Since each update leads to an energy value decrease, the algorithm is guaranteed to converge eventually. Finally, we also update the active vertices contributing to E_{env} and E_{shuffle}

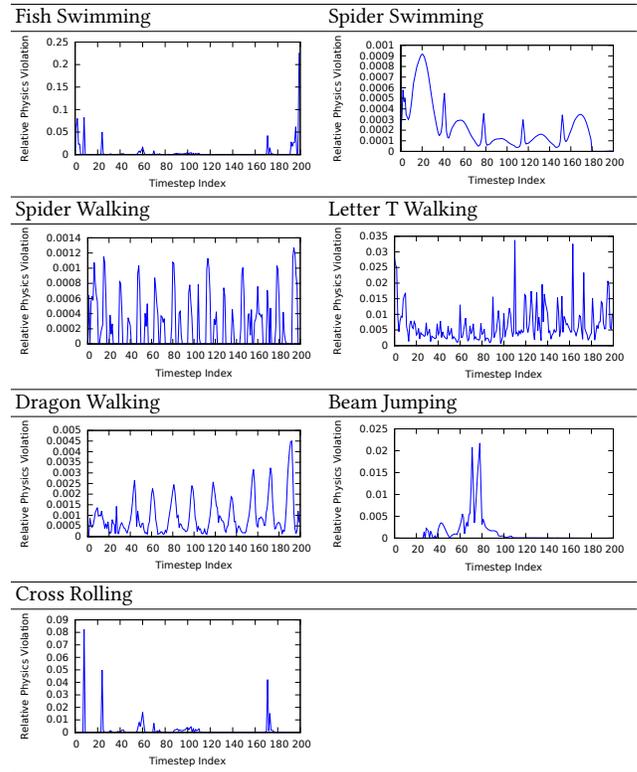


Table 6. The physics violation of all the 7 benchmarks.

in Line 33. A minor bottleneck in this algorithm is the sparse linear system in Line 8, but we notice that the hessian of $E(Q_K)$ is a block-tridiagonal matrix that can be inverted with linear time complexity.

B.1 LM Modification for E_{shuffle}

Conventional LM assumes that the objective function takes the form of $E(Q_K) = f(Q_K)^T f(Q_K) / 2$ where f is a vector of nonlinear terms. This is not the case with E_{shuffle} . However, among the many implementations of LM, the one documented in [Lourakis 2005] is not limited to the above form. If we have a certain approximation of the hessian of $E(Q_K)$ denoted as H (not necessarily in the form of $J^T J$) and use damping coefficient d , then the update to Q_K in our main algorithm is computed as:

$$\Delta Q_K = -(H + dI)^{-1} \frac{\partial E(Q_K)}{\partial Q_K},$$

and the decrease in $E(Q_K)$ after applying the update is:

$$\Delta Q_K^T \frac{\partial E(Q_K)}{\partial Q_K} + \frac{1}{2} \Delta Q_K^T H \Delta Q_K = \frac{1}{2} \Delta Q_K^T \left(\frac{\partial E(Q_K)}{\partial Q_K} - d \Delta Q_K \right),$$

which is exactly the expected function value decrease estimation used in [Lourakis 2005]. This allows us to use a non- $J^T J$ form of approximate hessian for E_{shuffle} . In our case, we approximate the term E_{shuffle} by first order expansion for both $v^j(q_i)$ and $v^j(q_{i-1})$, giving:

$$\begin{aligned}
& E_{\text{shuffle}}(q_i, q_{i-1}) \\
&= \sum_{j=1}^V \left\| (v^j(q_i + \Delta q_i) - v^j(q_{i-1} + \Delta q_{i-1})) \right\|^2 \\
& \exp(-\gamma \text{dist}(v^j(q_i + \Delta q_i))) \\
&\approx \sum_{j=1}^V \left\| (v^j(q_i) + \frac{\partial v^j}{\partial q_i} \Delta q_i - v^j(q_{i-1}) - \frac{\partial v^j}{\partial q_{i-1}} \Delta q_{i-1}) \right\|^2 \\
& \exp(-\gamma \text{dist}(v^j(q_i) + \frac{\partial v^j}{\partial q_i} \Delta q_i)),
\end{aligned}$$

and use the hessian of the last equation above. Conventional $J^T J$ approximation of the hessian is used for all other terms.

B.2 Eliminating Internal Force Terms

In our objective function, \mathcal{F}_i is a quadratic function in E_{EOM} , in the Tikhonov regularization $C_{\text{reg}} \|\mathcal{F}_i\|^2/2$, and in the DMP regularization E_{DMP} . We analytically solve and eliminate \mathcal{F}_i from the optimization as follows:

$$\begin{aligned}
& \min_{\mathcal{F}_i} E_{\text{EOM}}(q_{i-1}, q_i, q_{i+1}) + C_{\text{reg}} \|\mathcal{F}_i\|^2/2 + C_{\text{DMP}} E_{\text{DMP}}(\mathcal{F}_i) \\
&= \min_{\mathcal{F}_i} \frac{1}{2} \left[\|EOM_i - \mathcal{F}_i\|^2 + C_{\text{reg}} \|\mathcal{F}_i\|^2 + C_{\text{DMP}} \|\mathcal{F}_i - DMP_i\|^2 \right] \\
&= \frac{1}{2} \begin{pmatrix} EOM_i^T & DMP_i^T \end{pmatrix} M_{\text{reg}} \begin{pmatrix} EOM_i \\ DMP_i \end{pmatrix} \\
& EOM_i \triangleq (\mathbf{I} \ 0 \ 0) \left[\frac{\partial \bar{q}^T}{\partial q} (q_{i+1}) MA(q_{i+1}) + \frac{\partial [P - \mathcal{E}_i^T \bar{q}]}{\partial q} (q_{i+1}) \right] \\
& DMP_i \triangleq \begin{pmatrix} DMP_{p/np}(i\Delta t, \mathcal{W}_1) \\ DMP_{p/np}(i\Delta t, \mathcal{W}_2) \\ \vdots \\ DMP_{p/np}(i\Delta t, \mathcal{W}_{|\mathcal{F}_i|}) \end{pmatrix} \\
& M_{\text{reg}} \triangleq \frac{1}{1 + C_{\text{reg}} + C_{\text{DMP}}} \begin{pmatrix} C_{\text{reg}} + C_{\text{DMP}} & -C_{\text{DMP}} \\ -C_{\text{DMP}} & C_{\text{reg}} C_{\text{DMP}} + C_{\text{DMP}} \end{pmatrix} \otimes \mathbf{I}.
\end{aligned}$$

C V(Q) AND ITS DERIVATIVES

In this section, we briefly summarize the transformation function from the reduced representation q to a vertex v 's Euclidean coordinates as defined in [Pan et al. 2015]. We then provide some guidance on the computation of $\frac{\partial v(q)}{\partial q}$ and $\frac{\partial^2 v(q)}{\partial q^2}$. Since we use the LM algorithm for space-time optimization, most objective terms only require a first order derivative $\frac{\partial v(q)}{\partial q}$. An exception is the physics violation term E_{EOM} , which requires second order derivatives.

C.1 Transformation Function

We assume that a deformable body is discretized using V vertices and P elements. For each element j , its deformation gradient is denoted as F^j . Vertices' Euclidean coordinates \bar{q} can be reconstructed from F^j through Poisson reconstruction. We abbreviate this linear operator as:

$$\bar{q} = \Delta^{-1} \begin{pmatrix} F^1 T & \dots & F^P T \end{pmatrix}^T.$$

It is well-known that F^j above has a polar decomposition $F^j = R^j S^j$ where R^j is a rotation and S^j is an anisotropic scaling. We can then use the Rodrigues formula on the rotation part to get $R^j = \exp(\tau^j)$, where τ^j is the rotation vector of element j . The rotation-strain (RS) space is defined by the space spanned by all possible τ^j and $S^j - \mathbf{I}$:

$$\text{span}(\text{RS}) = \{(\tau^1, \dots, \tau^P, S^1 - \mathbf{I}, \dots, S^P - \mathbf{I}) | S^j \text{ is SPD}\}.$$

RS representation and Euclidean coordinates are equivalent. However, using RS representation is advantageous because the most visually salient deformations lie in a low-dimensional linear subspace of $\text{span}(\text{RS})$. Therefore, we can use conventional linear dimensionality reduction techniques such as linear modal analysis [Pentland and Williams 1989] in RS space to arrive at the following low-rank approximation:

$$(\tau^1, \dots, \tau^P, S^1 - \mathbf{I}, \dots, S^P - \mathbf{I}) \approx Bu,$$

where B is a set of bases in the RS subspace. The transformation function from u to \bar{q} can then be defined by combining the above equations:

$$\bar{q}(q) = \Delta^{-1} [\exp](Bu)[S](Bu)$$

$$[\exp](Bu) \triangleq \begin{pmatrix} \exp(\tau^1) & & \\ & \ddots & \\ & & \exp(\tau^P) \end{pmatrix} \quad [S](Bu) \triangleq \begin{pmatrix} S^1 \\ \vdots \\ S^P \end{pmatrix}.$$

However, this only encodes deformations in the local frame of reference. To allow arbitrary movement in the global frame of reference, we can superimpose a global translation c and rotation w , giving:

$$\bar{q}(q) = \begin{pmatrix} \exp(w) & & \\ & \ddots & \\ & & \exp(w) \end{pmatrix} \Delta^{-1} [\exp](Bu)[S](Bu) + \begin{pmatrix} c \\ \vdots \\ c \end{pmatrix}.$$

C.2 Kinetic Cubature Acceleration

Merely having the low-rank approximation does not accelerate computation. The evaluation of \bar{q} is still computationally costly because it requires a summation over all the P elements. This procedure can be accelerated using cubature approximation by assuming the following approximation:

$$\Delta^{-1} [\exp](Bu)[S](Bu) = \sum_j \Delta_j^{-1} \exp(w^j) S^j \approx \sum_{j \in \mathcal{T}} C_{\text{RS}}^j \Delta_j^{-1} \exp(w^j) S^j,$$

where Δ_j^{-1} is the block of Δ^{-1} corresponding to the j th element. The weighting C_{RS}^j and the set of cubature elements \mathcal{T} are precomputed using L0-optimization.

C.3 Derivatives

To derive $\frac{\partial v(q)}{\partial q}$ and $\frac{\partial^2 v(q)}{\partial q^2}$, we notice that they can always be written as a long chain of matrix productions, where each matrix is either a constant, a linear function of q , or a linear function of $\exp(w)$, $\frac{\partial \exp(w)}{\partial w}$, and $\frac{\partial^2 \exp(w)}{\partial w^2}$. Equations of this form can be accelerated using fast sandwich transform (FST) [Kim and James 2011], i.e., by precomputing high-order tensors and contracting them with q , $\exp(w)$, $\frac{\partial \exp(w)}{\partial w}$, and $\frac{\partial^2 \exp(w)}{\partial w^2}$ at runtime. The remaining problem is to derive $\frac{\partial \exp(w)}{\partial w}$ and $\frac{\partial^2 \exp(w)}{\partial w^2}$. A closed-form of $\frac{\partial \exp(w)}{\partial w}$

can be found in [Gallego and Yezzi 2015]. We now derive $\frac{\partial^2 \exp(\mathbf{w})}{\partial \mathbf{w}^2}$ below using their notations:

$$\begin{aligned} \mathbf{R} &\triangleq \exp(\mathbf{w}) \quad \frac{\partial \mathbf{R}}{\partial w_i} = [\mathbf{v}_i] \mathbf{R} \\ \frac{\partial^2 \mathbf{R}}{\partial w_i \partial w_j} &= \left[\frac{\partial \mathbf{v}_i}{\partial w_j} \right] \mathbf{R} + [\mathbf{v}_i] \frac{\partial \mathbf{R}}{\partial w_j} \\ \mathbf{v}_i &\triangleq \frac{w_i \mathbf{w} + [\mathbf{w}](\mathbf{I} - \mathbf{R}) \mathbf{e}_i}{\|\mathbf{w}\|^2} \\ &= \bar{w}_i \bar{\mathbf{w}} - \frac{\sin(\theta) [\bar{\mathbf{w}}]^2 + (1 - \cos(\theta)) [\bar{\mathbf{w}}]^3}{\theta} \mathbf{e}_i \\ &= \frac{\theta - \sin(\theta)}{\theta} \bar{w}_i \bar{\mathbf{w}} + \frac{1 - \cos(\theta)}{\theta} [\bar{\mathbf{w}}] \mathbf{e}_i + \frac{\sin(\theta)}{\theta} \mathbf{e}_i, \end{aligned}$$

where we used the identity $\bar{\mathbf{w}} = \mathbf{w} / \|\mathbf{w}\|$, $\theta = \|\mathbf{w}\|$, $[\bar{\mathbf{w}}]^2 = (\bar{\mathbf{w}} \bar{\mathbf{w}}^T - \mathbf{I})$, and $\bar{\mathbf{w}}^T [\bar{\mathbf{w}}] = \mathbf{0}$. Finally, the $\frac{\partial \mathbf{v}_i}{\partial w_j}$ above has the following form:

$$\begin{aligned} \frac{\partial \mathbf{v}_i}{\partial w_j} &= \frac{\theta - \sin(\theta)}{\theta^2} (e_j \bar{w}_i + \bar{w} \delta_{ij} - 2 \bar{w} \bar{w}_i \bar{w}_j) + \\ &\frac{\sin(\theta) - \theta \cos(\theta)}{\theta^2} \bar{w} \bar{w}_i \bar{w}_j + \frac{1 - \cos(\theta)}{\theta^2} [e_j - \bar{w}_j \bar{\mathbf{w}}] \mathbf{e}_i + \\ &\frac{\theta \sin(\theta) + \cos(\theta) - 1}{\theta^2} [\bar{\mathbf{w}}] \mathbf{e}_i \bar{w}_j - \frac{\sin(\theta) - \theta \cos(\theta)}{\theta^2} \mathbf{e}_i \bar{w}_j. \end{aligned}$$

D OTHER REDUCED MODELS

In this section, we analyze the cases where rotation-strain coordinates [Pan et al. 2015] is replaced with either LMA [Pentland and Williams 1989] or reduced StVK [Barbič and James 2005]. As mentioned in Section 6.2, we need to modify both $\bar{q}(q)$ and P in Equation 3. In both cases, $\bar{q}(q)$ takes the following simple form:

$$\bar{q}(q) = \begin{pmatrix} \exp(\mathbf{w}) & & \\ & \ddots & \\ & & \exp(\mathbf{w}) \end{pmatrix} (U\mathbf{u}) + \begin{pmatrix} c \\ \vdots \\ c \end{pmatrix},$$

where U is a set of linear deformation bases. Due to the lack of rotation-strain transformation, we have to introduce more columns to U than B in order to represent nonlinear deformations. If LMA is used, P takes the same quadratic form as that for rotation-strain coordinates, $P(q) = \mathbf{u}^T \mathcal{K} \mathbf{u} / 2$. If reduced StVK is used, $P(u)$ is a quartic function in u whose polynomial coefficients can be precomputed.