Locking-Proof Tetrahedra

N.B. When citing this work, cite the original published paper.

(article starts on next page)

Locking-Proof Tetrahedra

MIHAI FRÂNCU, ARNI ASGEIRSSON, and KENNY ERLEBEN, University of Copenhagen
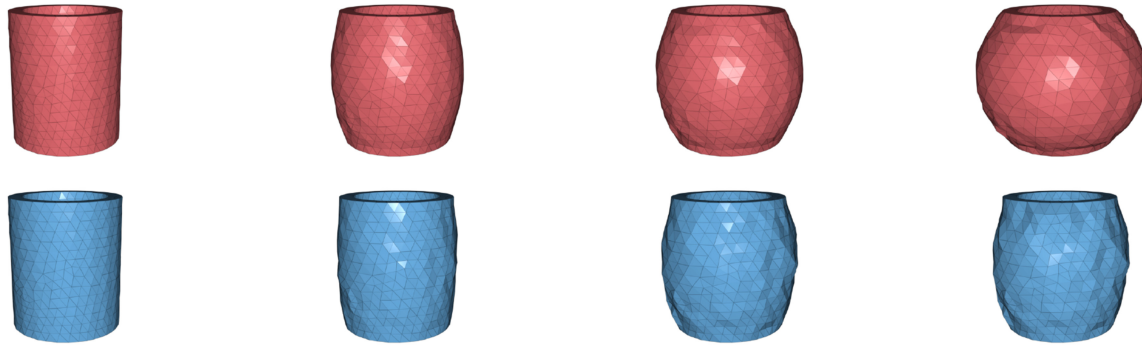MADS J. L. RØNNOW, Chalmers University of Technology

Fig. 1. A simple scenario for studying volumetric locking: an incompressible rubber tube made out of silicone ($E = 90$ KPa, $v = 0.499$, $\rho = 1070$ kg/m^3) is inflated by internal pressure (10 KPa). This is the kind of air chamber found in most pneumatic soft robots. We simulate the inflation using a linear tetrahedral mesh of 3,628 elements, a quasi-static approach (10 steps), and a Newton solver (5 iterations). Snapshots are taken at corresponding times. Our mixed FEM method (red) using a Neo-Hookean distortional material shows significantly more deformation than standard Neo-Hookean FEM (blue), which locks.

The simulation of incompressible materials suffers from locking when using the standard finite element method (FEM) and coarse linear tetrahedral meshes. Locking increases as the Poisson ratio v gets close to 0.5 and often lower Poisson ratio values are used to reduce locking, affecting volume preservation. We propose a novel mixed FEM approach to simulating incompressible solids that alleviates the locking problem for tetrahedra. Our method uses linear shape functions for both displacements and pressure, and adds one scalar per node. It can accommodate nonlinear isotropic materials described by a Young's modulus and any Poisson ratio value by enforcing a volumetric constitutive law. The most realistic such material is Neo-Hookean, and we focus on adapting it to our method. For $v = 0.5$, we can obtain full volume preservation up to any desired numerical accuracy. We show that standard Neo-Hookean simulations using tetrahedra are often locking, which, in turn, affects accuracy. We show that our method gives better results and that our Newton solver is more robust. As an alternative, we propose a dual ascent solver that is simple and has a good convergence rate. We validate these results using numerical experiments and quantitative analysis.

CCS Concepts: • **Computing methodologies → Physical simulation**; • **Computer systems organization** → *Robotics*;

Additional Key Words and Phrases: Finite element method (FEM), mixed FEM, incompressible, locking, nonlinear materials, constrained dynamics

1 INTRODUCTION

Ever since the introduction of soft body simulation to computer graphics by Terzopoulos et al. [1987], the physical fidelity has increased proportionally to demand and hardware improvements. This is even more the case nowadays when we are gradually moving these simulation techniques to areas like fabrication, soft robotics, and bio-medical applications.

Our motivation stems mainly from the design and digital prototyping of soft robots. In such scenarios, it is very important to use nonlinear incompressible materials that allow for large deformations without losing volume. Many authors use standard finite element method (FEM) methods for simulating such materials [Bern et al. 2019; Coevoet et al. 2017]. But early on, Roth et al. [1998] stressed the need for true and accurate volumetric physical models.

The simulation of highly deforming soft bodies using FEM can pose a lot of challenges. In graphics, Stomakhin et al. [2012] note that "the simulation of such large deformation problems with a Lagrangian mesh is notoriously unstable and error-prone." This is usually due to numerical ill-conditioning and difficulties of solving the nonlinear system. This is especially true for nearly

incompressible materials, which are "notoriously difficult to simulate robustly and accurately" [Smith et al. 2018].

An often overlooked problem is the phenomenon of *locking*. Locking manifests in much smaller displacements than expected when using Poisson ratios ν very close to 0.5. Tan et al. [2012] report using Poisson ratios of no more than 0.45 in order to circumvent locking artifacts at the price of losing physical accuracy. Only a handful of authors have addressed locking in graphics [Irving et al. 2007; Kaufmann et al. 2009; Patterson et al. 2012].

The problem of first order FEM is that it always exhibits locking. This is unavoidable as it is inherent in the mathematical model. Many authors analyzed this rigorously in engineering and proposed various fixes: e.g., reduced integration and mixed formulations [Zienkiewicz et al. 2005], assumed strain or B-bar method [Hughes 2012], discontinuous Galerkin [Wihler 2006], or higher order interpolation [Heisserer et al. 2008].

We build on top of this existing work in graphics and engineering, and offer a relatively cheap solution for removing locking from linear tetrahedral meshes. We show that there are important cases where locking cannot be ignored any longer as it really affects the accuracy of the results. A good example is the use of the Neo-Hookean material for simulating soft robots or biological tissue. In such cases, locking is mostly glossed over in graphics and fabrication, despite the focus on accurate volume preservation [Smith et al. 2018]. We show, in this article, that this can have a dramatic impact on the accuracy of digital prototyping.

We propose a novel version of the *mixed* finite element method that removes locking without resorting to higher order elements. The mixed attribute denotes the fact that we are solving for both the displacement and the pressure field as unknowns. We set mixed FEM as a standard for high quality when simulating incompressible materials and in general for any Poisson ratio value. We show that our formulation is equivalent to that of constrained dynamics, e.g., articulated bodies, frictional contact [Bender et al. 2014], incompressible fluids [Bridson 2015], cloth [Bender et al. 2017], strain-limiting [Wang et al. 2010], or constraint-based FEM [Servin et al. 2006].

Our most significant contributions can be summarized as:

— Handling of locking for linear tetrahedral meshes;
— A general procedure for turning nonlinear volumetric constitutive laws into compliant constraints;
— A thorough analysis on choosing the distortional elastic energy focused on nonlinear Neo-Hookean materials;
— A Newton solver that can support $\nu = 0.5$ and handles volumetric stiffness better than its standard FEM counterpart;
— A dual ascent solver that is simple, converges well, and can reuse existing standard FEM solvers.

2 VOLUMETRIC LOCKING

There are multiple types of locking: volumetric, membrane, or shear locking. In this work, we only deal with *volumetric locking*. All locking types share the feature that they are parameter dependent, and locking manifests whenever that parameter goes to infinity. In essence, all locking types are due to the geometric discretization. In this respect, they can be associated with the accuracy error and numerical stiffness artifacts for coarse meshes. Usually, these

associated effects go away with the increase of mesh resolution. On the other hand, locking is a more particular phenomenon as it increases its effect with stiffness. In other words, a convergence study on the mesh resolution may show an increase in accuracy and decrease in numerical stiffness, but locking may still occur [Babuška and Suri 1992]—see Figure 2 for an illustration. In simple terms, this is basically due to having too few degrees of freedom (DOFs) compared to the number of constraints we are imposing.

A clear case of volumetric locking is illustrated in 2D in the inset to the right and occurs when we are imposing incompressibility constraints to the elements together with boundary conditions. The black nodes are fixed, and only the white node can move. In order to preserve the area of each triangle, the node can only move vertically for the blue triangle and only horizontally for the red triangle, thus resulting in deadlock. Note that these constraints do not have to be explicit and can be modulated by a stiffness factor. A classic example is the use of a volume penalty term in the elastic energy with the stiffness factor being the *bulk modulus K*.

We briefly state the necessary condition for preventing volumetric locking [Zienkiewicz et al. 2005]:

$$n_p \leq n_x, \tag{1}$$

where n_p is the number of pressure (or volumetric) DOFs and n_x is the number of position DOFs. Hence, the number of pressure DOFs should always be less or equal to the one for freely deforming nodal positions. The inset above shows that even for the equality case (i.e., $n_p = n_x = 2$) locking can still occur and is highly dependent on boundary conditions.

The condition in Equation (1) only makes sense in the context of mixed FEM, as it makes explicit the discretization of both the displacement and pressure fields. But the reasoning can be transferred over to standard FEM by realizing that n_p is in fact the number of volume constraints or the number of points where we enforce the constitutive law. For linear tetrahedral elements, strain is constant per element and, thus, the stress and the pressure too. This means that n_p is equal to the number of elements that can go as high as $4n$ or more, where n is the number of nodes, and the system becomes over-constrained, i.e., violating Equation (1). The same happens for explicit mixed formulations that enforce incompressibility per element (see inset above for an example). More mathematical details are given in Section 8.1 and Appendix A. The fact that mixed FEM allows us to choose different shape functions for the displacement and pressure fields is key to preventing locking. Roughly speaking, the extra DOFs $n_p - n_x$ allow the body to deform. And by carefully choosing the shape functions, we can control whether locking happens or not, which is not possible for standard FEM.

Similarly to Irving et al. [2007], we choose to store pressure DOFs per node instead in order to prevent locking, i.e., $n_p = n < 3n = n_x$. This amounts to adding more DOFs to the discretized system than in standard linear FEM. This still does not completely guarantee accurate solutions, as such elements require additional stabilization in the pressure term [Al Akhrass et al. 2012; Malkus and Olsen 1984; Sani et al. 1981]. In order to satisfy the tight theoretical Ladyzhenskaya–Babuška–Brezzi (LBB) condition, one needs to use at least quadratic displacements
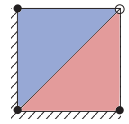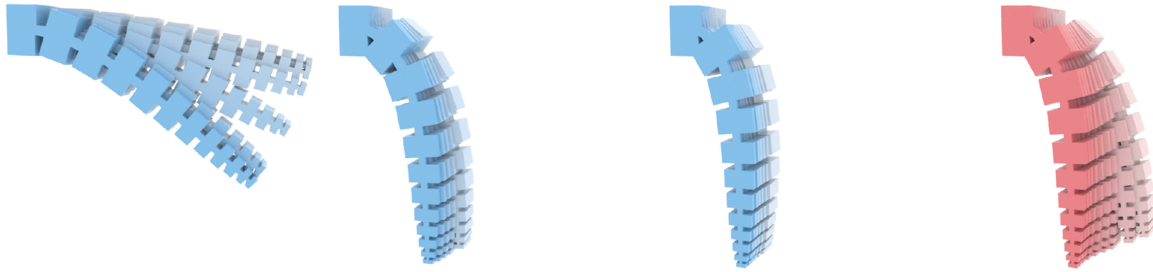
Fig. 2. Convergence test for standard FEM (blue) and our method (red). From left to right: first, second, and third order for standard FEM. The meshes were refined by reducing the edge length. In the case of linear standard FEM, the deformation appears to keep increasing unbounded. Our method converges and gives results in accord with cubic, which we regard as the ground truth.

and linear pressures [Boffi et al. 2013]. But this comes at a higher cost than using standard quadratic tetrahedra, which also alleviate locking. Still, the LBB condition is a sufficient condition, and it is not required in this case to ensure non-locking, albeit desirable. We show in this article that the linear-linear (in displacement and pressure, respectively) approach for locking offers very good results in practice and out-performs standard FEM without resorting to higher-order interpolation.

3 RELATED WORK

3.1 Finite Element Method

There is a massive amount of computer graphics work involving the finite element method [Manteaux et al. 2017; Nealen et al. 2006]. Many classic text books introduce all the important concepts of FEM in the context of linear elasticity [Bathe 2006; Hughes 2012; Zienkiewicz et al. 2005]. The Newton method is the work horse in mechanical engineering for solving nonlinear FEM [Bonet and Wood 1997; Wriggers 2008]. Sifakis and Barbic [2012] summarize the theory for nonlinear FEM and present reduced-order models in the context of computer graphics.

It is worth mentioning the linear corotational method [Müller and Gross 2004], which was very popular and still is for simulating large realistic deformation. Parker and O'Brien [2009] adapted it for real-time interactive simulation. Further improvements addressed inversion handling [Irving et al. 2004; Schmedding and Teschner 2008; Stomakhin et al. 2012], polar decomposition [Civit-Flores and Susín 2014; Kugelstadt et al. 2018], matrix factorization [Hecht et al. 2012] or nonlinearity [Chao et al. 2010; McAdams et al. 2011].

For simulating general nonlinear materials, a lot of effort has been put into the eigen-analysis of the energy Hessian in order to fix indefiniteness and use the conjugate gradient method [Kim et al. 2019; Smith et al. 2018, 2019; Teran et al. 2005]. Advanced techniques like higher-order elements [Bargteil and Cohen 2014; Weber et al. 2011] or multi-resolution solvers [Weber et al. 2015; Zhu et al. 2010] were also employed in computer graphics. To address the issues with volume mesh generation, some authors embedded the original surface mesh into a lattice of hexahedra with special quadrature rules [McAdams et al. 2011; Nesme et al. 2009; Patterson et al. 2012].

Projective dynamics introduced a new perspective on elasticity based on constraint projection and a faster solver [Bouaziz et al. 2014; Dinev et al. 2018; Liu et al. 2017; Narain et al. 2016].

3.2 Constrained Dynamics

Constrained dynamics has proven to be a powerful approach for rigid bodies with frictional contact [Bender et al. 2014]. Very fast and robust solvers were developed in this field and are currently being applied to real-time interactive simulation, robotics, and engineering (e.g., Bullet, Havok, PhysX, MuJoCo). Regularization and soft compliant constraints have paved the way to robust FEM simulation [Servin et al. 2006; Todorov 2014]. *Position based dynamics* (PBD) is another popular method for simulating soft bodies using constraints [Bender et al. 2017; Müller et al. 2007]. More recently, constraint-based FEM has been extended to PBD [Frâncu and Moldoveanu 2017; Macklin et al. 2016]. Constrained dynamics for elastic bodies now covers full implicit time integration, nonlinear materials, and complementarity based contact [Macklin et al. 2019]. Volumetric *strain limiting* techniques [Perez et al. 2013; Wang et al. 2010] used constraints to enhance stiff elastic simulations. Tournier et al. [2015] emphasized the important role of the geometric stiffness matrix in stabilizing constraints, including for elastic simulations.

3.3 Volume Preservation and Locking

A lot of interest has been given in graphics to volume conservation for soft body simulation. We distinguish two approaches: volume constraints using PBD [Bender et al. 2017] and the penalty energy term [Kikuuwe et al. 2009; Martin et al. 2011; Teschner et al. 2004; Wang et al. 2019]. In fluid dynamics, incompressibility is enforced using a pressure projection ensuring a divergence free velocity field [Bridson 2015]. The same Eulerian technique can be applied to solids [Brandt et al. 2019; Zhu et al. 2010]. Irving et al. [2007] devised a similar linearized pressure projection in the Lagrangian setting and exposed the perils of locking. In the context of the *material point method* (MPM), Stomakhin et al. [2014] used a fluid-like approach close to ours by splitting the corotational energy into a deviatoric part and a pressure part, the latter enforcing a stiff equation of state. Patterson et al. [2012] proposed both a penalty and a nonlinear mixed FEM approach, both offering remedies for locking. Kaufmann et al. [2009] used a discontinuous Galerkin method that is also locking free.

But these articles stop short of telling the full story of locking, which we have started in Section 2. Just to give an idea, having constant or averaged pressure hexahedral elements as in Patterson et al. [2012] may prevent locking, but this is no longer the case for

tetrahedral elements, as pointed out by Irving et al. [2007]. Even using linear pressure elements for tetrahedra is not the best option, but going higher order is often not an affordable solution.

Locking is much better known in the engineering community. Zienkiewicz et al. [2005] provide detailed reasons as to why locking afflicts the standard *irreducible* form of FEM. There are many approaches to removing locking as already mentioned. Some of them can be applied directly to the irreducible form without many changes. Increasing the order of the polynomials helps to avoid the problem. Babuška and Suri [1992] prove this for linear elasticity and that increasing the resolution of the mesh does not alleviate the problem. Heisserer et al. [2008] extend the result to nonlinear Neo-Hookean materials. This works mainly due to the increase of deformation DOFs per element. The downside is that one is required to go as high as order four to prevent locking. Reduced integration uses counter-intuitively less quadrature points, thereby reducing the number of times the underlying constitutive law is enforced. Discontinuous Galerkin (or non-conforming) elements give up on the continuity requirements, resulting in sparser matrices and do not lock [Wihler 2006]. The assumed strain [Krysl and Zhu 2008] or the B-bar method [Hughes 2012] is another relatively simple method to remove locking from standard FEM.

3.4 Mixed Formulation

Reduced integration and assumed strain are shown to be equivalent to the *mixed formulation* by Zienkiewicz et al. [2005] and others. The theory behind mixed FEM is quite elaborate and can explain why all standard FEM suffers from locking [Boffi et al. 2013]. This is why mixed formulations are recommended in practice for nonlinear incompressible materials [Bonet and Wood 1997; Maas et al. 2012; Wriggers 2008]. Solving the mixed formulation always results in a *saddle point problem* [Benzi et al. 2005]. After close inspection, this boils down to solving a discrete constrained dynamics problem. And this is one of our main reasons for choosing it in favor of the other solutions.

Roth et al. [1998] were among the first to use the mixed formulation of linear elasticity for facial animation with "accurate volumetric physics." The method of Irving et al. [2007] can be shown to be a particular case of mixed FEM using linear shape functions for both displacements and pressure: their div and grad operators correspond to our constraint Jacobian matrix and its transpose. They apply an operator *splitting* approach like in fluid dynamics and use arbitrary internal elastic energies. For the former, Batty and Bridson [2008] have shown that splitting can introduce unwanted errors. And pertaining to the latter, we show in this article that the distortional energy needs in fact to be carefully chosen. Patterson et al. [2012] use the mixed formulation for nonlinear hexahedral elements and briefly sketch a recipe for choosing distortional energies without focusing much on locking aspects or the equivalence to constrained dynamics.

Although the linear-linear approach is not considered a stable solution for locking in the engineering community, considerable effort has been put into improving it, given its low cost: the mini-element [Hughes 2012], the bubble function [Zienkiewicz et al. 2005], orthogonal sub-grid scales [Cervera et al. 2003], average nodal pressure [Bonet and Burton 1998], or the 10-node compos-
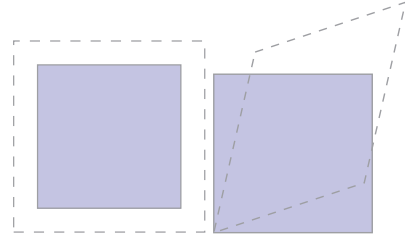

Fig. 3. Volumetric (left) and distortional (right) deformations of a square represented as dashed lines.

ite element [Ostien et al. 2016]. We improve on previous work and show that, in practice, the linear-linear approach for nonlinear materials and tetrahedra produces satisfactory results in terms of locking artifacts.

4 INCOMPRESSIBLE MIXED FEM

In this section, we lay down the theoretical ground work for the mixed formulation and present our contributions regarding locking. We invite the reader to consult [Bonet and Wood 1997] for a proper theoretical introduction on nonlinear continuum mechanics if needed. Alternatively, the course by Sifakis and Barbic [2012] introduces most of the same concepts to a computer graphics audience.

Here, we will limit ourselves to clarifying some of the notation. We denote by $\mathbf{x}$ spatial (or deformed) coordinates and by $\mathbf{X}$ material (or undeformed) coordinates. The *deformation* vector field is given by the mapping $\mathbf{x} \equiv \boldsymbol{\varphi}(\mathbf{X})$ and its Jacobian $\mathbf{F} \equiv \nabla \boldsymbol{\varphi} = \partial \mathbf{x} / \partial \mathbf{X}$ is called the *deformation gradient*. See the 2D illustration on the right. The *displacement* vector field is $\mathbf{u} \equiv \mathbf{x} - \mathbf{X}$ and the displacement gradient is $\nabla \mathbf{u} = \mathbf{F} - \mathbf{I}$, where $\mathbf{I}$ is the identity matrix.

The right Cauchy-Green tensor is defined as $\mathbf{C} = \mathbf{F}^{\mathsf{T}}\mathbf{F}$ and the Green strain tensor as $\mathbf{E} \equiv (\mathbf{C} - \mathbf{I})/2$. The linearization of $\mathbf{E}$ is the small Cauchy strain $\boldsymbol{\varepsilon} \equiv (\nabla \mathbf{u} + \nabla \mathbf{u}^T)/2$. These are used to define *isotropic hyperelastic* materials for which the energy density Ψ depends only on the deformation gradient $\mathbf{F}$.

The Cauchy stress $\boldsymbol{\sigma}$ is a spatial tensor corresponding to the physical forces acting in the system. The first Piola-Kirchoff (PK1) stress tensor $\mathbf{P} = J\boldsymbol{\sigma}\mathbf{F}^{-T}$ is work conjugate to $\mathbf{F}$, i.e., $\mathbf{P} \equiv \partial \Psi / \partial \mathbf{F}$.

4.1 Incompressibility

Another known fact from continuum mechanics is the formula relating the spatial differential volume dv to the material one dV,

$$dv = J\,dV, \tag{2}$$

where $J = \det(\mathbf{F})$. This means that for incompressible materials, we need to have $J = 1$.

In general, we can consider a split between the *volumetric* and the *distortional* deformation. A distortional or isochoric deformation is one that conserves volume, as depicted in Figure 3. The deformation gradient in this case should be $\hat{\mathbf{F}} = J^{-1/3}\mathbf{F}$ such that $\hat{J} = \det\hat{\mathbf{F}} = 1$. At the stress level we know that the trace of the
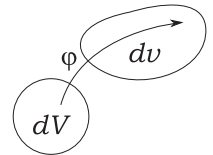
distortional Cauchy stress should be zero so that the internal forces acting on an infinitesimal element have the form $\nabla \cdot \sigma_d + \nabla \bar{p}$, where $\bar{p} = \text{tr}(\sigma)/3$ is the *pressure* and $\sigma = \sigma_d + \bar{p}\mathbf{I}$. Note that $\text{tr}(\sigma_d) = 0$ does not imply that the trace of the distortional Green strain or PK1 tensor should be zero [Bonet and Wood 1997]. The condition $\text{tr}(\varepsilon) = 0$ is only an approximation in the linear regime.

The split at the energy level can be expressed as

$$\Psi = \Psi_d + \Psi_v, \tag{3}$$

where the function $\Psi_d(\mathbf{C})$ must be homogeneous of order 0, i.e., depend only on $\hat{\mathbf{F}}$ [Bonet and Wood 1997]. This split results in two separate constitutive laws for the two deformation modes. In the case of full incompressibility, Ψ_v can be interpreted as an indicator function enforcing the hard incompressibility constraint $J = 1$. Materials that adhere to the strict split in Equation (3) are called *uncoupled materials* [Maas et al. 2012; Simo and Taylor 1991].

4.2 Penalty Formulation

In practice, a common solution for enforcing incompressibility is by defining Ψ_v as a term penalizing volume change. Usually, this is done using a quadratic term

$$\Psi_v = \frac{\kappa}{2}\Phi(\mathbf{F})^2, \tag{4}$$

where κ is a volumetric stiffness factor and the $\Phi(\mathbf{F}) = 0$ level set marks the volume preserving configurations. We call Φ the *volumetric strain* function.

Assumption 1. *One obvious choice, which we will prefer, is*

$$\Phi(J) \equiv J - 1 = 0. \tag{5}$$

Another example is $\Phi(J) \equiv \log J = 0$ —see Neo-Hookean below and Section 6. The logarithm term enforces an infinite barrier at $J = 0$ such that the body never inverts or degenerates to a point.

As Smith et al. [2018] point out, there are two problems associated with uncoupled materials. The first one is that using the negative powers or the logarithm of J can cause numerical issues for degenerate and inverted elements. The second one is that such materials are not *rest stable*, i.e., do not have zero energy and gradient for $J = 1$.

Take, for example, the standard Neo-Hookean material energy

$$\Psi_{NH} = \frac{\mu}{2}(I_1 - 3) - \mu \log J + \frac{\lambda}{2}(\log J)^2, \tag{6}$$

where $I_1 = \text{tr}(\mathbf{C})$ is the first invariant of $\mathbf{C}$, and μ and λ are the Lamé parameters. Note that λ becomes infinity when $\nu = 0.5$. This is an example of a *coupled material* as the μ terms depend on invariants of $\mathbf{F}$ rather than $\hat{\mathbf{F}}$ and are clearly not purely distortional. On the other hand, it is rest stable. In addition, the logarithm squared can be replaced by $(J - 1)^2$ [Ogden 1997] so that it fits under Assumption 1. This is a clear case showing that the split in Equation (3) does not always hold for practical nonlinear materials, especially when allowing for compression. Actually, there is some leeway in choosing the split—we touch upon it here but defer further discussion to Section 5.

Patterson et al. [2012] identify Ψ_v as being the λ term in the linear elasticity, corotational, Saint Vennant-Kirchoff (StVK), and Neo-Hookean energies, resulting in coupled materials. This makes sense as this term expresses volume preservation in different ways

for each material. For linear elasticity this is $\Phi(\mathbf{F}) = \text{tr}(\mathbf{F} - \mathbf{I}) = \text{div } \mathbf{u} = 0$, similar to a divergence free velocity field in fluids. And similar approximations are given for corotational and StVK. Instead, we prefer to use volumetric energies that depend solely on J, which is the true nonlinear measure of volumetric deformation. The dependence solely on J also permits easy extraction of the constraint function for volumetric energies of the form $\Psi_v = \kappa\Phi(J)^2/2$. For instance, the energy $\Psi_v = \kappa(J - 1)/2$ [Wang and Yang 2016] will have the constraint function $\Phi(J) = \sqrt{J - 1} = 0$.

Assumption 2. *Thus, we will assume, in this article, that volumetric energies always depend only on J and are quadratic in $\Phi(J)$.*

4.3 Pressure-displacement Mixed Formulation

As already hinted under Equation (3), another approach is to enforce hard constraints. The most common way is to use Lagrange multipliers, which are equivalent to the pressure field $\bar{p}$ under the assumptions of Section 4.1 [Bonet and Wood 1997]. This fact is well known in the simulation of incompressible fluids [Bridson 2015].

Again, in practice, we might want to allow for some compression as we did in the previous section. The mixed formulation allows us to take any volumetric penalty energy and convert it into a soft constraint.

Starting from the definition of the Cauchy stress $\sigma_v = \bar{p}\mathbf{I}$ for pure dilation, we get that $\bar{p}$ and J are work conjugate variables, i.e., their generalized dot product gives us the work of the pressure forces. From the assumption that the volumetric energy is of the form $\Psi_v(J)$, we get the constitutive law

$$\bar{p}(J) \equiv \frac{d\Psi_v}{dJ} = \Psi_v'(J). \tag{7}$$

See Appendix B and Stomakhin et al. [2014] for more details. Similarly, for the distortional part, we use the conjugate pair $\mathbf{P}$ and $\mathbf{F}$. Thus, $\bar{p}$ and $\mathbf{P}$ play the role of generalized forces, and J and $\mathbf{F}$ are deformation measures. Since we are not using $\hat{\mathbf{F}}$, $\mathbf{P}$ may not be purely distortional and the pressure $\bar{p}$ no longer corresponds to the total pressure field, i.e., $\text{tr}(\sigma_d) \neq 0$.

We generalize these known facts to any volumetric strain function $\Phi(J)$. Applying the chain rule to Equation (7), we get $\bar{p} = \Psi_v'(\Phi)\Phi'(J)$ and introduce the notation $p \equiv \Psi_v'(\Phi) = \bar{p}/\Phi'(J)$. In other words, we have replaced the conjugate pair $\bar{p}$ and J by p and Φ. The continuous soft constraint is obtained by inverting the new $p = p(\Phi)$ constitutive law:

$$\Theta(J, p) = \Phi(J) - \Gamma(p) = 0, \tag{8}$$

where Γ is the inverse of the Ψ_v' function. In the particular case of Assumption 2, we have $p = \kappa\Phi$ and $\Gamma(p) = p/\kappa$, which hold for any constraint function $\Phi(J)$.

Under these assumptions, we can derive the *variational* or *weak form* of the differential algebraic equations of motion using our adaptation of the nonlinear mixed formulation [Wriggers 2008]

$$\int_\Omega \rho\delta\mathbf{x} \cdot \ddot{\mathbf{x}}\, dV + \int_\Omega \mathbf{P} : \delta\mathbf{F}\, dV + \int_\Omega p\,\delta\Phi\, dV$$
$$- \int_\Omega \delta\mathbf{x} \cdot \mathbf{b}\, dV - \int_{\partial\Omega} \delta\mathbf{x} \cdot \mathbf{t}\, dA = 0, \tag{9}$$
$$\int_\Omega \delta p\,[\Phi - \Gamma(p)]dV = 0. \tag{10}$$

(a) $\nu = 0.49$ (b) $\nu = 0.499$ (c) $\nu = 0.4999$ (d) $\nu = 0.49999$
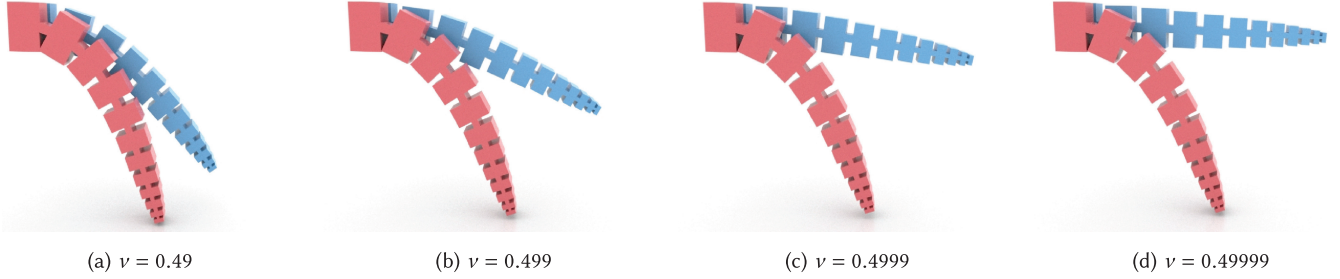
Fig. 4. Spine robot arm bending under gravity with the left end fixed. Comparing our method (red) to the standard Neo-Hookean method (blue) with $E = 266$ KPa, $\rho = 1070$ kg/m^3, 1,778 tetrahedra, $g = -9.8$ m/s^2, and increasing Poisson ratios ν. For both, we used 10 Newton iterations and 10 quasi-static steps. It is clear that the standard method locks more and more for higher ν, whereas our method does not.

where ρ is the density, **b** and **t** are external loads, and Ω and $\partial\Omega$ represent the material body domain and its boundary. The operator δ denotes taking the *variation* [Sifakis and Barbic 2012], and $\delta\mathbf{x}$ and δp are *test functions*. The operator ":" denotes the *double contraction* between two second-order tensors [Bonet and Wood 1997]. Note that, in general, the pressures $\bar{p}$ are only identical to the Lagrange multipliers p when $\Phi'(J) = 1$, the case of Assumption 1. We will denote p as pressures further on, but they are better interpreted as Lagrange multipliers enforcing Equation (8), especially for coupled materials.

The standard finite element method as well as the mixed form are procedures for spatial discretization of such weak form equations. For this, we use finite elements Ω_e and sample both the deformation and the pressure fields using shape functions: $\mathbf{x} \approx \sum_i N_\mathbf{x}^i(X)\mathbf{x}^i = N_\mathbf{x}\mathbf{x}$ and $p \approx \sum_j N_p^j(X)p^j = N_p p$, respectively, where the super-scripts denote the discrete interpolation points of each element. As it is customary [Erleben et al. 2005; Zienkiewicz et al. 2005], we have aggregated the shape functions into shape function matrices in order to condense further notation. Keep in mind that all derived quantities like $\mathbf{F}$, J, or derivatives of Ψ_υ and Φ can now be interpolated from $\mathbf{x}$. From here on, we will use sans-serif Latin and small bold Greek letters for discrete matrix and vector quantities (as opposed to serif Latin and capital Greek letters for the continuum).

One of our contributions is to identify, in the integral terms of Equation (10), the per-element *discrete volumetric strain function*:

$$\boldsymbol{\phi}_e(\mathbf{x}) = \int_{\Omega_e} N_p^\mathsf{T}\Phi(\mathbf{x})\, dV, \tag{11}$$

and the *discrete compliance function* (or constitutive law):

$$\boldsymbol{\gamma}_e(p) = \int_{\Omega_e} N_p^\mathsf{T}\Gamma(p)\, dV = \int_{\Omega_e} N_p^\mathsf{T}(\Psi_\upsilon')^{-1}\, dV. \tag{12}$$

This constraint-based interpretation is novel to our knowledge and applies, in general, for any volumetric energy, element type, and interpolation order.

By expanding the $\delta\Phi$ term in Equation (9), we identify the per-element Jacobian matrix:

$$G_e(\mathbf{x}) = \frac{\partial\boldsymbol{\phi}_e}{\partial\mathbf{x}} = \int_{\Omega_e} N_p^\mathsf{T}\frac{\partial\Phi}{\partial\mathbf{x}}\, dV. \tag{13}$$

By converting all tensor quantities to matrices and vectors, and by assembling them for all elements, we obtain the spatially dis-

cretized equations of motion:

$$M\ddot{\mathbf{x}} + G(\mathbf{x})^\mathsf{T}p - f_d(\mathbf{x}) - f = 0, \tag{14}$$

$$\boldsymbol{\theta}(\mathbf{x}, p) \equiv \boldsymbol{\phi}(\mathbf{x}) - \boldsymbol{\gamma}(\mathbf{x}, p) = 0, \tag{15}$$

where M is the mass matrix (lumped or not), f is the external force, and f_d is the distortional elastic force. Our contribution is the realization that Equations (14)–(15) are, in fact, equivalent to a constrained dynamics system with constraint function $\boldsymbol{\phi}$, Jacobian matrix G, Lagrange multipliers p, and regularization term $\boldsymbol{\gamma}$. It is now time to recall that given a proper discretization of the displacement and pressure fields that satisfy Equation (1), we can achieve locking-free simulations by solving the above differential and algebraic equations.

4.4 Linearization

Sooner or later, we will have to linearize the constraint Equation (15) in order to solve the nonlinear discrete system it is part of. This linearization has the form

$$G(\mathbf{x})\Delta\mathbf{x} - C\Delta p + g = 0, \tag{16}$$

where g is the constant part of the Taylor expansion of $\boldsymbol{\theta}$ and the global *compliance matrix* C is assembled from local blocks

$$C_e = \frac{\partial\boldsymbol{\gamma}_e}{\partial p} = \int_{\Omega_e} N_p^\mathsf{T}\Gamma'(p)N_p\, dV = \int_{\Omega_e} N_p^\mathsf{T}(\Psi_\upsilon'')^{-1} N_p\, dV. \tag{17}$$

Using Assumption 2, we always have a linear relation $\Gamma = p/\kappa$ and, therefore, the compliance matrix is constant, i.e.,

$$C_e = \frac{1}{\kappa}\int_{\Omega_e} N_p^\mathsf{T}N_p\, dV. \tag{18}$$

This means that Equation (15) becomes

$$\boldsymbol{\theta}(\mathbf{x}, p) = \boldsymbol{\phi}(\mathbf{x}) - Cp, \tag{19}$$

which we will use throughout the rest of the article.

This linearization result can be shown to be equivalent to the one in Patterson et al. [2012] using a hexahedral element with a single averaged pressure node. A similar result can be found in Roth et al. [1998] and Roth [2002] for the case of linear elasticity and high-order tetrahedral elements. Our results in Equations (17) and (18) are an extension of the previous work as they apply to nonlinear incompressibility and arbitrary shape functions.

4.5 Time Integration

We choose to use the implicit Euler integration because of its simplicity and stability, which make it very popular [Baraff and Witkin 1998]. One could use Newmark or BDF-2 instead as they dissipate less energy or any other preferred implicit or explicit integrator. The implicit Euler time discretization of Equation (14) is

$$\sigma(x, p) = \frac{1}{h^2} M(x - x^n - hv^n) + G(x)^\mathsf{T} p - f - f_d(x) = 0, \quad (20)$$

where the super-script n denotes the current timestep of length h. Note the identification $x \equiv x^{n+1}$ and $p \equiv p^{n+1}$ for the unknown configuration at the end of the timestep. As the constraint gradient $G(x)$ is evaluated at the end of the timestep, just as the pressures p, we conclude that the constraint forces are integrated in a fully implicit manner; this, in turn, ensures the stability in time of the system [Goldenthal et al. 2007; Tournier et al. 2015].

We employ the Newton method for nonlinear system solving and consider the constant compliance matrix assumption in Equation (18) for most practical cases. The resulting discretized and linearized equations are

$$\begin{bmatrix} \frac{1}{h^2} M + K_d + K_g & G^\mathsf{T} \\ G & -C \end{bmatrix} \begin{pmatrix} \Delta x \\ \Delta p \end{pmatrix} = \begin{pmatrix} -\sigma(x_k, p_k) \\ -\theta(x_k, p_k) \end{pmatrix}, \quad (21)$$

where k denotes the current Newton iteration, G is evaluated at x_k, and

$$K_d = -\frac{\partial f_d(x_k)}{\partial x} \text{—the distortional stiffness matrix,} \quad (22)$$

$$K_g = \frac{\partial^2 \phi(x_k)}{\partial x^2} p_k \text{—the geometric stiffness matrix.} \quad (23)$$

Note that the geometric stiffness matrix K_g is the by-product of the fact that we are dealing with nonlinear constraints, and the Jacobian G is varying with x.

By omitting the acceleration term and time integration completely, we can obtain the static case. This corresponds to Equations (20) and (21) without the terms multiplied by $1/h^2$ [Sifakis and Barbic 2012]:

$$\begin{bmatrix} K & G^\mathsf{T} \\ G & -C \end{bmatrix} \begin{pmatrix} \Delta x \\ \Delta p \end{pmatrix} = \begin{pmatrix} -G^\mathsf{T} p_k + f + f_d(x_k) \\ -\theta(x_k, p_k) \end{pmatrix}, \quad (24)$$

where $K = K_d + K_g$. As you can see the static formulation is the same as in Patterson et al. [2012], but the dynamics case in Equation (21) is different, as we identify the pressures p as Lagrange multipliers, as is the case for all constrained dynamics applications [Bender et al. 2014, 2017; Macklin et al. 2019].

4.6 Linear Tetrahedral Elements

In this article, we choose to use only linear tetrahedral meshes as they are widely used in computer graphics. Extensions to higher order can be made using numerical quadrature or analytically using Bernstein-Bézier polynomials as in Roth et al. [1998], Roth [2002], and Frâncu et al. [2019], and similar to standard FEM [Bargteil and Cohen 2014; Weber et al. 2011, 2015].

If we assume the stress is constant over an element (as in standard FEM), then pressure is constant, too, over the element, i.e., $N_p \equiv 1$. We call these order zero or piece-wise constant pressure shape functions. This gives us very simple formulas for $J_e = v_e(x)/V_e$, $\phi_e(x) = v_e(x) - V_e$, and $C_e = V_e/\kappa \, I$, which are constant
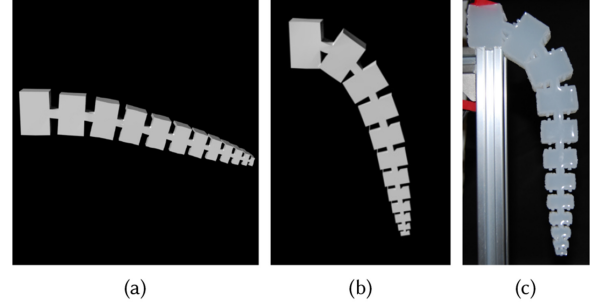

(a)   (b)   (c)

Fig. 5. Spine robot arm hanging under gravity with one side fixed; simulated using our method with values $E = 150$ KPa, $\nu = 0.4999$, $\rho = 1{,}070$ kg/m^3, 1,778 tetrahedra, and $g = -9.8$ m/s^2: (a) zero-order pressure, (b) first-order pressure, (c) real life photograph. For both simulations, we used 10 Newton iterations and 10 quasi-static steps. The first-order pressure simulation took about 38 s, while for constant pressure, it was less—in this case, 30 s. The former locks much less, even for such a coarse mesh.

over the element e. Using the triple product formula for the volume of a tetrahedron

$$v_e(x) \equiv \frac{1}{6}(x_1 - x_0) \cdot [(x_2 - x_0) \times (x_3 - x_0)], \quad (25)$$

we can evaluate the four Jacobian components:

$$G_1^\mathsf{T} = \frac{1}{6}(x_2 - x_0) \times (x_3 - x_0), \quad (26a)$$

$$G_2^\mathsf{T} = \frac{1}{6}(x_3 - x_0) \times (x_1 - x_0), \quad (26b)$$

$$G_3^\mathsf{T} = \frac{1}{6}(x_1 - x_0) \times (x_2 - x_0), \quad (26c)$$

and $G_4 = -G_1 - G_3 - G_3$. This makes sense, as the pressure maintaining the volume of the element will act with forces normal to the opposite faces and proportional to their area [Parker and O'Brien 2009; Teran et al. 2003].

But constant pressure elements are not the best choice for locking, as visible in Figure 5. As described in Section 2, we choose to use linear shape functions for both deformation and pressure, i.e., $N_p = N_x$ are the four barycentric coordinates inside a tetrahedron. Looking at Equation (11), this means we now have four volumetric constraints per tetrahedron located at the nodes (co-located with the pressures and displacements). These four constraint components are equal, and we denote them by $\phi_e^i = (v_e(x) - V_e)/4$, where $i = 1..4$ is a local node index. Remember that we are solving for global constraint functions located at the nodes. These are obtained by assembling together all the per-element constraint functions. In the end, for each global node, we get a constraint function of the form

$$\phi_j = \sum_{e \in \mathcal{I}(j)} \phi_e^{\mathrm{map}(e,j)}, \quad (27)$$

where $\mathcal{I}(j)$ is the set of incident elements to a global node j and map$(e, .)$ is a mapping from global indices to local indices in a given element e. A similar assembly process can be applied for the global Jacobian matrix G and the compliance matrix C. On the right, we have illustrated the process and notation in 2D.

This is essentially the one-ring approach of Irving et al. [2007], where their discrete divergence and gradient operators correspond to the Jacobian matrix G and its respective transpose. However,
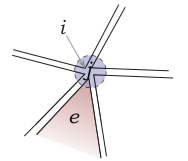
their volume constraint function is linearized and the discrete operators are constant. Their solver is much akin to a divergence free fluid solver corresponding to linear elasticity. They employ an equation splitting approach that yields an approximation of the system matrix in the pressure projection step, as similarly done by Perez et al. [2013]. The true matrix, as we will see later, is the one in Equation (38). Finally, their method is built solely on linear-linear mixed tetrahedral elements and do not offer a recipe for extension for other interpolation schemes, as we do in Section 4.3.

To sum up, we have introduced new linear tetrahedral elements that are more resilient to locking. We call them *locking proof tetrahedra*. The reason behind this claim is that the necessary condition in Equation (1) holds very well in practice—see the results in Section 9.

5 DISTORTIONAL ENERGY CHOICE

The choice of the distortional energy type must be consistent with the type of material being simulated; but as we have seen for both uncoupled and coupled materials, the choice of the split may be arbitrary at times and up to the choice of the user. Irving et al. [2007] suggest that their method can be used with arbitrary nonlinear material models. The danger is that stiff volumetric components in the model can introduce back locking. Patterson et al. [2012] simply choose to use the μ term in the energy or make $\lambda = 0$, as this is the term that contains the problematic $1 - 2\nu$ denominator. Bonet and Wood [1997] present a rigorous way of separating the distortional energy.

Most materials already have a volume preservation term or approximation thereof, as hinted in Section 4.2, e.g., corotational is simply a linearization of Neo-Hookean. Therefore, we are basically enforcing incompressibility twice if using such standard materials in the mixed formulation: once through the constraints and once through the penalty term.

Due to all these reasons, we came up with a few guiding principles when choosing a distortional energy: it should not contain a stiff volumetric term; it should handle singularities and inversion well and, ideally, be rest stable. Of course, some of these could be relaxed depending on the context, especially in the $\nu \approx 0.5$ regime. Still, we found out that in order to support any value of ν, then the distortional energy together with the soft constraints should form a *coupled compressible* material.

The first energy we tried was simply the deviatoric component of linear elasticity, which can be found in many references [Roth et al. 1998; Zienkiewicz et al. 2005]

$$\Psi_d^L = \frac{\mu}{2}\boldsymbol{\varepsilon} : \boldsymbol{\varepsilon} - \frac{2\mu}{3}\text{tr}(\boldsymbol{\varepsilon})^2. \tag{28}$$

This is an an uncoupled material that contrasts the shear only energy $\mu(\boldsymbol{\varepsilon} : \boldsymbol{\varepsilon})/2$ used by Patterson et al. [2012].

For the nonlinear case, we tried to obtain a distortional material that is still based on the classical Young's modulus E and Poisson ratio ν (or the Lamé parameters λ and μ). We settled on the popular Neo-Hookean material and tried out several options.

One of them is the energy due to Mooney [Smith et al. 2018]
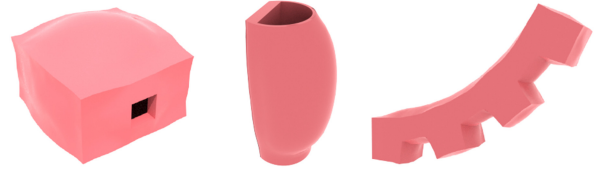
$$\Psi_d^M = \frac{\mu}{2}(I_1 - 3). \tag{29}$$

Fig. 6. Pneumatic soft robots inflated by internal pressure. From left to right: "cushion," "bubble," "finger."

This works well as an uncoupled material, especially in the incompressible regime, being a particular case of the Mooney-Rivlin material [Maas et al. 2012].

The exact uncoupled nonlinear distortional energy is calculated in Bonet and Wood [1997]

$$\Psi_d^{BW} = \frac{\mu}{2}(J^{-2/3}I_1 - 3). \tag{30}$$

Despite its apparent simplicity, the PK1 tensor has a complicated formula and even more so for its differential. By analogy with the Mooney material and assuming J is constant, we simplified it to

$$\mathbf{P}_d^{BW} \approx \mu J^{-2/3}\mathbf{F}. \tag{31}$$

The Ψ_d^{BW} distortional energy is obtained based on the assumption $J \approx 1$. We would like to work with any type of material in our formulation, so this assumption does not hold unless $\nu \approx 0.5$. In addition, Smith et al. [2018] present a nice discussion of why the term $J^{-2/3}$ is not so great to have, despite its theoretical appeal. Another important note they make is about the *rest stability* condition for the total energy, i.e., the rest point corresponds to $J = 1$ [Ogden 1997]. Most coupled Neo-Hookean formulations satisfy this condition. Due to all these reasons, we choose our Neo-Hookean distortional material as follows

$$\Psi_d^{NH} = \frac{\mu}{2}(I_1 - 3) - \mu \log J. \tag{32}$$

This distortional energy appears in most Neo-Hookean variants and is the same as in Patterson et al. [2012]. For the volumetric part, they use $\kappa = \lambda$. When the Poisson ratio is zero, such as for cork, we have $\kappa = 0$ or the absence of volumetric constraints, which reflects the true behavior of the material. But most nonlinear elasticity literature uses the bulk modulus instead, as it measures the resistance to compression and correctly aligns with the linearized elasticity in Equation (28)—see Maas et al. [2012], Section 6.1.6 in Ogden [1997], or Section 6.5.3 in Bonet and Wood [1997]. Therefore, we chose to use $\kappa = K$, even if, for $\nu = 0$, we get $\kappa = E/3$.

The nice thing about this last distortional material is that it can be used in standard FEM without a volume term and does not deform in the perpendicular directions to the compression/stretching. It does this by coupling the volumetric deformation to the distortional one through the shear modulus μ. Given this and the fact that its first and second-order derivatives are easy to compute, we chose it as our preferred distortional material when simulating Neo-Hookean materials. For some practical uses of our method using this material, see Figure 6.

For any general nonlinear material, one can follow the recipe for the incompressible regime from Bonet and Wood [1997]. For verification, one can linearize the distortional energy and make sure

that it matches the form in Equation (28). Or one can use the distortional split in Patterson et al. [2012], e.g., for corotational or StVK. We could apply the linearization strategy to StVK and obtain a formula similar to the linear distortional energy, i.e., with the small Cauchy strain ε replaced by the Green strain $\mathbf{E}$. On the other hand, the full StVK material seems to benefit from the mixed formulation, as it notoriously misbehaves under compression [Sifakis and Barbic 2012]. This amounts to an extension of StVK that was done previously using a penalty potential [Kikuuwe et al. 2009]. Given they do not fall under Assumption 2, we have not done many experiments with these other materials besides Neo-Hookean.

6 LOGARITHMIC CONSTRAINTS

Although we focused on the Neo-Hookean material and the soft $J = 1$ constraint, our methods do not lack in generality. As hinted before, most nonlinear materials contain powers of $J - 1$ or $\log J$ in their volumetric energy expression. For the distortional part, one can choose any material of choice according to the guidelines presented in the previous section.

One obvious extension of our method is the use of the $\Phi(J) \equiv \log J$ constraint. Therefore, we will re-derive our mixed formulation using the new constraint and see what equations change. The continuous soft constraint is $\Theta(J, p) = \log J - p/\kappa$. As you can see, the definition of the Lagrange multipliers p ensures that we always have the same form of Equation (8) under Assumption 2. This means that we can use most of the remaining formulas in the article without modification. The only caveat now is that the computed Lagrange multipliers have an absorbed J factor, and we need to divide by it in order to obtain the real pressures $\bar{p} = p/J$. The discrete constraint function is now $\phi_e^i(x) = V_e(\log v_e(x) - \log V_e)/4$ per local node. This means we need to apply the chain rule for computing the constraint gradient, which, in the end, amounts to scaling the relations in Equation (26) by $1/J$, and similarly for the second derivative in the geometric stiffness matrix.

7 CONNECTION TO OPTIMIZATION THEORY

The saddle point problem in Equations (9)–(10) can be written as a variational minimization problem $\delta\Pi(\mathbf{x}, p) = 0$, where Π is the total energy functional of the distortional, constraint, and external forces. In finite dimensions, after discretization, we can define the energy

$$E(\mathbf{x}) = \frac{1}{h^2}\|\mathbf{x} - \mathbf{x}^n - h\mathbf{v}^n\|_M^2 + U_d(\mathbf{x}) - \mathbf{x}^T\mathbf{f}, \qquad (33)$$

where U_d is the discrete distortional elastic energy, and the hard constraints $\phi(\mathbf{x})$ together with the Lagrangian $\mathcal{L}(\mathbf{x}, \mathbf{p}) = E(\mathbf{x}) + \mathbf{p}^T\phi(\mathbf{x})$. We introduce the *perturbed Lagrangian* $\mathcal{L}_p(\mathbf{x}, \mathbf{p}) = \mathcal{L}(\mathbf{x}, \mathbf{p}) - \frac{1}{2}\mathbf{p}^T\mathbf{C}\mathbf{p}$ [Zienkiewicz et al. 2005], which amounts to regularizing the constraints. In this light, the variational problem amounts in the discrete case to finding a saddle point

$$(\mathbf{x}, \mathbf{p})^{n+1} = \arg\min_{\mathbf{x}}\max_{\mathbf{p}} \mathcal{L}_p(\mathbf{x}, \mathbf{p}). \qquad (34)$$

If compliance is zero, then this is simply a minimization of the total energy with hard incompressibility constraints

$$\mathbf{x}^{n+1} = \arg\min_{\phi(\mathbf{x})=0} E(\mathbf{x}). \qquad (35)$$

It can be shown that the saddle point of the perturbed Lagrangian is equivalent to the unconstrained minimization of the total energy using the penalty formulation in Equation (4)

$$\mathbf{x}^{n+1} = \arg\min_{\mathbf{x}} E(\mathbf{x}) + \frac{\kappa}{2}\|\phi(\mathbf{x})\|^2. \qquad (36)$$

But discretizing directly the strain field and minimizing only with respect to $\mathbf{x}$ leads to standard FEM and locking; so we need to preserve the saddle point structure and the Lagrange multipliers $\mathbf{p}$. Thus, the equivalence above between the mixed and standard formulations is only theoretical, while, numerically, the latter is ill-conditioned and prone to locking. Therefore, we cannot use an off-the-shelf optimization solver and resort to the Newton method instead to solve the discrete nonlinear saddle point in Equations (14)–(15). These equations are also known as the Karush-Kuhn-Tucker (KKT) optimality conditions associated with $\mathcal{L}_p$. In the end, we get the series of sparse indefinite linear systems in Equation (21).

8 SOLVING THE PROBLEM

The saddle point Newton solver was implemented in C++ using the Pardiso LDLT decomposition from the Eigen library [Guennebaud et al. 2010] with Math Kernel Library (MKL) integration [Wang et al. 2014]. We implemented another version using the *conjugate residuals* (CR) algorithm, which is an iterative solver well suited for indefinite systems [Saad 2003]. We also had success with the *conjugate gradient* (CG) method, but both of them required diagonal pre-conditioning. This is because the constraint equation needs to be scaled several orders of magnitude to match the force balance equation. Both CR and CG are matrix-free solvers, which enabled us to use the force differentials from Sifakis and Barbic [2012]. For the direct solver, we assembled the stiffness matrix by computing the force differentials on basis vectors.

For standard FEM simulations with nonlinear materials, we used mostly our own implementation and the FEBio software suite [Maas et al. 2012] in order to compare results. We also used PolyFEM [Schneider et al. 2019a] to cross-validate some of the results. For mesh generation, we used TetGen [Si 2015], NetGen [Schöberl 2009], Wildmeshing [Hu et al. 2018], and other sources. For vizualization and comparisons, we used the VTK file format and ParaView [Ahrens et al. 2005]. Our implementation is based in large part on Sifakis and Barbic [2012]. For solving the linear system in the Newton iterations we mostly opted for a direct solver using the Pardiso LU decomposition. Often times, especially for stiff systems, the Newton solver can fail if the initial guess is too far from the solution. Line search strategies can help but may find an inverted solution or fail altogether. The quasi-static approach where the loads and boundary conditions are applied slowly is a good remedy as we are approaching the true solution with better and better guesses. However, for both the quasi-static and the dynamic case, one has to tweak the timestep in order to achieve stable simulations.

We found that our Newton solver behaves very well for very high Poisson ratios (e.g., $\nu = 0.499$) for which the standard FEM with Neo-Hookean materials struggles to converge. For many of the experiments presented in this article, we had to reduce the timestep at least by half when doing standard Neo-Hookean
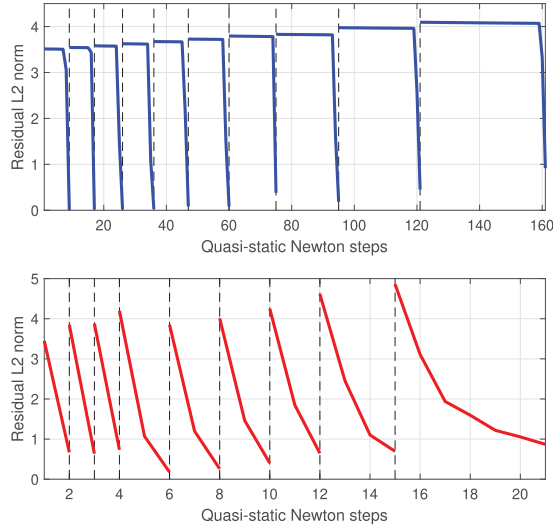
Fig. 7. Convergence plot of the Newton solver with a residual threshold of 1 over a series of 10 quasi-static steps. The accumulated number of iterations for each step is shown on the horizontal axis. Simulation scenario: "bubble" soft robot inflated by $p = 4$ KPa ($\nu = 0.499$, $E = 100$ KPa, $\rho = 1{,}070$ kg/m^3, 5,193 tetrahedra). Our method (red) requires less iterations than standard Neo-Hookean (blue), which has a poorer convergence behavior.

incompressible simulations. For our method, we often found that back-tracking in the Newton step is seldom needed and the convergence rate is optimal for large timesteps, whereas the opposite happens for standard FEM. In Figure 7, we give an example of such a case where the standard FEM Newton solver takes very small steps at the beginning before finally converging. Although it usually helps, lowering the timestep value in this case had the inverse effect of increasing the number of iterations even higher. As you can see from the same figure, our method can use less Newton iterations and less function and Jacobian evaluations to achieve the same deformation, making it potentially faster despite the larger size of the problem. For this particular scenario, a Newton step costs around 0.36 s for standard FEM and 0.65 s for our method. However, the last quasi-static step was almost 15 s for standard Neo-Hookean and under 4 s for our method. Keep in mind that our simulators were not optimized, and the real ratios between timings may vary. Simulations were run on a desktop PC with a Intel i7-3770 CPU with four physical cores. The computation of the stresses and stiffness matrices were distributed among cores using OpenMP.

For more details about performance, see Table 1. As expected, our method takes roughly double the time of standard FEM when the Newton solver is running in similar conditions. The exception comes when the latter cannot handle incompressibility well and requires more Newton iterations or more quasi-static steps. We hope that through some of the approaches described below, we will be able to reduce this performance gap in the future.

8.1 Static Condensation

As an alternative to solving the saddle point problem, one can use the *Schur complement* to reduce the size of the problem in

Equation (21). The drawback is that we lose sparsity and need to back-solve for the rest of the unknowns. There are actually two Schur complements we could take. The first one we investigate uses the inverse of the compliance matrix

$$A_1 = \frac{1}{h^2}M + K + G^T C^{-1} G. \tag{37}$$

This is called *static condensation* in the finite element static analysis literature [Bathe 2006; Roth 2002]. It, in fact, very much resembles standard FEM, in the sense that we are dealing with only one equation ($A_1 u = f$) with a single n_x by n_x stiffness matrix. The difference is that the volumetric component $G^T C^{-1} G$ is constructed in such a way that guarantees non-locking—see Appendix A. In fact, the solver can now be seen as a penalty method corresponding to the formulation in Equation (36), only that now it is locking free due to the special mixed discretization of the stiffness matrix. The downside is that it does not allow for $\nu = 0.5$ and it grows unbounded for $\nu \approx 0.5$. Still, when using a constant compliance matrix like we do, a major optimization is to pre-compute its inverse.

8.2 Boundary Conditions and Contacts

We have used the static condensation approach for Dirichlet boundary conditions (BCs) and frictionless contacts. Usually, when marking some of the nodes as prescribed boundary conditions, they can be removed from the numerical system. Another way is to consider them as constraints and enforce them through Lagrange multipliers or a penalty factor (i.e., hard or soft constraints). This permitted us to dynamically add or remove fixed nodes and constrain them in certain directions only, e.g., the contact normals. We implemented these constraints by extending Equation (21) with the extra compliant constraints and a tweakable stiffness parameter. We then chose to apply static condensation to these terms in order to hide the corresponding Lagrange multipliers. This permitted us to apply the same technique to standard FEM. For a full listing of the simulation loop pseudo-code, see Algorithm 1. Note that the mass matrix M and compliance matrix C are computed at the beginning of the simulation.

A more correct way of enforcing contacts is by using the nonpenetration complementarity conditions. The KKT form can be extended to accommodate these conditions, which stem from inequality constraints. Coulomb friction can be added, too, resulting in a *nonlinear complementarity problem*. However, in this article, we chose to consider the contacts as bilateral (or equality) soft constraints for the duration of a frame, and then use the static condensation approach as it can also be applied to standard FEM.

8.3 Dual Solver

The other Schur complement is the one used more often in constrained dynamics where the upper left block is inverted instead. The advantage is that for fully incompressible materials, the compliance becomes zero and the Schur complement does not grow to infinity, but remains constant [Servin et al. 2006], i.e.,

$$A_2 = G \left(\frac{1}{h^2}M + K \right)^{-1} G^T + C. \tag{38}$$

We will focus on the incompressible case in this section, although compliance can be easily added back to A_2 and to the constraints.

Table 1. Time Measurements in Seconds for the Various Quasi-Static Experiments Presented in
This Article for Both the Standard Method and Ours

Experiment	Figure	E (KPa)	v	p (KPa)	Steps	Tolerance	Standard (s)	Mixed (s)
Inflated tube	1	90	0.499	10	10	0.01	9	19
Hanging spine 1	4	266	0.499	-	10	0.01	3	4
Hanging spine 2	5	150	0.4999	-	20/10	0.01	2	2.5
Cushion	6	90	0.499	10	10	0.01	2	3
Bubble	6	100	0.499	4	10	0.2	12	27
Finger	6	100	0.499	10	10	0.1	28	62
Compressed box	12(b)	66	0.4999	15	10	0.1	102	12
Stretched cylinder	13	66	0.49/0.4999	-	80/20	0.01	106	78
Twisted box	14	66	0.45	-	20/10	0.01	68	60

The figures are rounded and taken from the best run of the Newton solver with line search. Young's moduli (in KPa) and Poisson ratios, as well as the tolerance for the residual are given for reference. For the last two experiments we had to use more quasi-static steps in order to make standard FEM work, while for the stretched cylinder we also had to lower the Poisson ratio to 0.49. Except from a few cases, our method takes more time as we are solving a larger system.

ALGORITHM 1: Simulation loop

Require: current positions x, undeformed positions X, matrices M, C
Ensure: new positions x, pressures p
 1: Compute external forces f
 2: Compute Jacobian matrices of Dirichlet BCs/contacts
 3: **for** iter = 1:maxNewtonIter **do**
 4:     Assemble distortional stiffness matrix K_d
 5:     Assemble geometric stiffness matrix K_g
 6:     Assemble volumetric Jacobian matrix G
 7:     Add Neuman BCs contributions to f
 8:     Compute right hand side (RHS) from σ and ϕ
 9:     Add contributions due to contact to RHS
 10:     Form the KKT system matrix
 11:     Add the contact/Dirichlet BCs condensed stiffness matrices
 12:     Solve the linear system
 13:     Update positions x and pressures p
 14: **end for**
 15: Update velocities (for time integration)

As you can see, the Newton step becomes quite costly, as matrix inversion is an expensive operation. Instead, we have developed a dual algorithm that does not take Newton steps and replaces matrix inversion with one single linear system solve. The algorithm is a form of the *dual ascent* method [Boyd et al. 2011], also known as *nonlinear Uzawa* [Aujol 2009]

$$x_{k+1} = \arg\min_x \mathcal{L}(x, p_k), \tag{39}$$

$$p_{k+1} = p_k + \omega\phi(x_{k+1}). \tag{40}$$

The first step in Equation (39) is a primal solver, which amounts to solving the standard FEM problem with a distortional material and the current constraint forces. The second step in Equation (40) is a simple gradient ascent step of length ω on the dual variables. In this context, the constraint function ϕ plays the role of the gradient of the dual objective function. Choosing a value can be hard, in general; if it is too small, a convergence can be very slow. We chose to use a steepest descent step [Shewchuk 1994] by exploiting, at every iteration, the information in the dual Hessian A_2

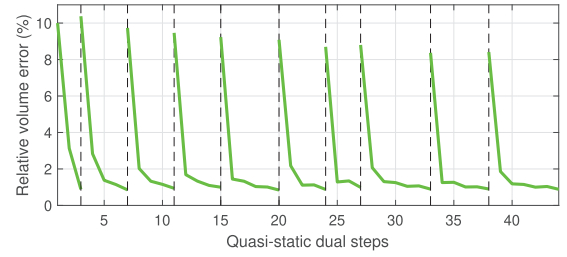$$\omega = \frac{\phi^T\phi}{\phi^T A_2 \phi}. \tag{41}$$

Fig. 8. Convergence plot of the dual solver over a series of 10 quasi-static steps. The residual was computed as the L1 norm of the volume error divided by the total initial volume. The accumulated number of iterations for each step is shown on the horizontal axis. Simulation scenario: "bubble" soft robot inflated by $p = 4$ KPa ($E = 100$ KPa, $\rho = 1,070$ kg/m^3, 5,193 tetrahedra). We used a Neo-Hookean material with $v = 0.4$ in an augmented Lagrangian approach to incompressibility. The saddle point Newton solver (not shown here) achieves an error threshold of 1% in one iteration, while our dual method (green) does it in at most six iterations.

This gives us a decent convergence for this class of solvers compared to Newton, as can be seen from Figure 8. In order to compute the matrix-vector product, we exploited the structure of A_2: we first solved for z in $(M/h^2 + K)z = G^T\phi$ and then computed $A_2\phi = Gz$. This avoids the matrix inversion and allows us to reuse an existing matrix factorization from the primal step if available.

Still, this method in its current form struggled to converge for some difficult scenarios like a stretching or twisting cantilever. The fix comes from using a standard material with a moderate Poisson ratio value (around 0.4). This corresponds to using an *augmented Lagrangian* in Equation (39) and the theoretical solution is still the same one corresponding to the hard incompressibility constraints assumed above [Boyd et al. 2011]. In effect, this makes the iterations less aggressive and the solver converge smoothly to a volume preserving solution. Note that this is not exactly the augmented Lagrangian method [Nocedal and Wright 2006], as the volumetric stiffness and the step ω are not the same.

We found this method to be working very well in practice, and it can also be used for enforcing contacts and boundary conditions. Its main advantage is that one can reuse an existing standard FEM solver by simply adding the constraint forces G^Tp as

external forces; and with a very simple loop, one can turn any material into an incompressible one without locking. Of course, our method calls the primal solver several times and is therefore more expensive. But if one already has a GPU primal solver [Wang and Yang 2016], the whole dual solver can be easily ported to GPU by adding an external loop containing the steps in Equations (39), (40), and (41).

Another advantage of the dual method is that we can now use projection on the dual variables for contact, i.e. impose that the contact Lagrange multipliers are in $\mathbb{R}_+$. This corresponds to the aforementioned complementarity formulation of contact and paves way for adding Coulomb friction. We have tested this projected gradient ascent approach and found it to be working well. But as already stated, we chose to mainly use the condensation approach instead, and leave complementarity-based frictional contact for future work.

The matrix inversion problem has been faced by many other authors who found various ways of dealing with it: Tournier et al. [2015] solve only one linear saddle point problem; Andrews et al. [2017] and Macklin et al. [2019] approximate the geometric stiffness matrix; and Duriez [2013] computes the exact matrix inverse asynchronously. Many choose to invert only the mass matrix [Perez et al. 2013], which is very easy to do as it is usually diagonal or block-diagonal. This can be seen as a form of splitting the differential equations as in Irving et al. [2007].

9 RESULTS

To investigate locking, we let a cantilever of size $0.1 \times 0.1 \times 1$ meters bend under gravity with one end fixed. For a very coarse mesh, we can observe that our method bend about 1 cm more under the same load than standard Neo-Hookean. Locking becomes more evident if we instead let our "spine" soft robot behave as a hanging cantilever, as seen in Figure 5. We increased the Poisson ratio while comparing our method with standard Neo-Hookean. By inspecting Figure 4 and Table 2, we can observe that noticeable locking occurs for the standard Neo-Hookean material already at $\nu = 0.49$. Increasing the Poisson ratio further greatly increases the amount of visible locking while our method bends without locking. As expected, the zero-order pressure elements in Figure 5 behave similarly to the standard Neo-Hookean ones. A plot of the potential energy evolution in time in Figure 9 shows that our method reaches its minimum for a larger deflection and has a wider range of movement. The data was gathered using dynamic simulation with an implicit time integrator.

In order to prove that our mixed formulation does indeed reach the correct deformation, we performed a convergence test and compared to higher-order simulations. We performed the quadratic and cubic simulations in FEBio. As you can see in Figure 2, the second-order standard FEM simulation is still locking, while cubic is deforming more. Unfortunately, we could not use higher orders in FEBio, but given the agreement between cubic and our method, we infer that they are both converging toward the ground truth. Keep in mind that there is still no full guarantee that cubic is not locking, but it seems this is not happening for the Poisson ratio of $\nu = 0.4999$. In Figure 10, we drew a convergence plot using the vertical deflection of the spine model. One can see that both

Table 2. Deflection in Centimeters of the Endpoint of the Spine Robot Arm Bending Under Gravity

ν	Our method (cm)	Neo-Hookean (cm)
0.49	14.6	11.2
0.499	14	6.3
0.4999	13.8	2.4
0.49999	13.8	0.8

Comparison of our method to standard Neo-Hookean with $E = 266$ KPa, $\rho = 1,070$ kg/m^3, 1,778 tetrahedra and increasing Poisson ratios. Neo-Hookean goes quickly to zero while our method stays fairly constant.
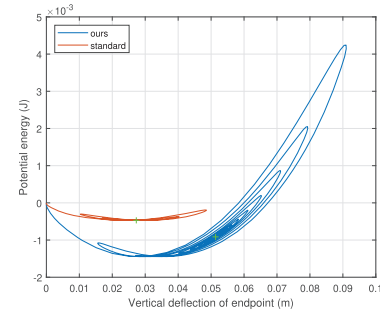

Fig. 9. Plot of the total potential energy of the spine robot arm as a function of the vertical deflection ($E = 266$ KPa, $\nu = 0.49$, $\rho = 1,070$ kg/m^3, $g = -1.5$ m/s^2). Samples are gathered using an implicit dynamics simulation with a timestep of 16 ms and 3 Newton iterations. The green crosshairs mark the minimum (equilibrium) points. Our method achieves larger deflections and oscillates more.

our method and standard cubic FEM are converging toward the same value, while quadratic tetrahedra and hexahedra (hex20) are lagging behind. Moreover, our method obtains very large deformations for coarse meshes, especially compared to standard linear FEM (Figure 2), even if it has a lower convergence rate. In general, we are able to achieve similar results to standard cubic FEM for roughly the same number of DOFs, which dictate the size of the system to solve. Of course, the corresponding cubic meshes have much fewer elements with a larger diameter, but our method requires no numerical quadrature; matrices are sparser and linear FEM methods are generally simpler to implement. To give a concrete example, using our method on a mesh with maximum edge length of $l = 0.02$, we can achieve comparable results to cubic running on $l = 0.04$. The former has 2,914 nodes and 11,618 tetrahedra; and the latter has 8,730 nodes and 1,382 elements. Despite the higher number of elements that can affect computation times, we still have less than half the number of DOFs to solve for.

The inflating tube is a popular thought experiment in mechanical engineering [Herrmann 1965; Simo and Taylor 1991], and it has applications in pneumatic actuation of soft robots [Ilievski et al. 2011]. We fixed the tube by its upper and lower margins and then applied an internal pressure while considering the external one to be zero. Our method showed a larger radius of deformation compared to standard FEM as observed in Figure 1. Note that the bumps on the surface are due to the coarse tetrahedralization and the very high Poisson ratio, and were also observed in FEBio. In Figure 11, we plotted the average radius of the tube as a function of ν. Our
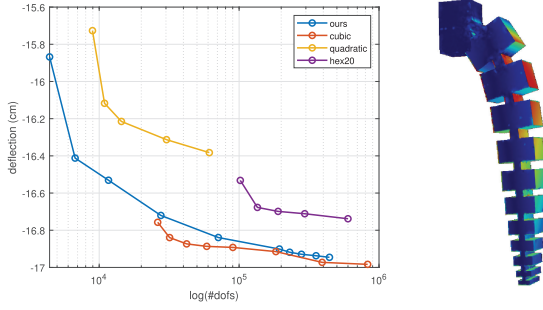
Fig. 10. Convergence plot for the spine model ($E = 266$ KPa, $\nu = 0.4999$, $\rho = 1{,}070$ kg/m^3, $g = -9.8$ m/s^2). On the horizontal axis, we measure the number of DOFs on a logarithmic scale. This number is proportional to the number of nodes in the mesh by a factor of three for standard FEM and four for our method. On the vertical axis, we plot the lowest y coordinate of the spine cantilever. Our method achieves convergence for a value close to the cubic one. On the right, we show the result of our method for the finest mesh ($l = 0.0065$) and the Hausdorff distance to the corresponding finest mesh for the cubic ($l = 0.015$). The two meshes are very close to each other.
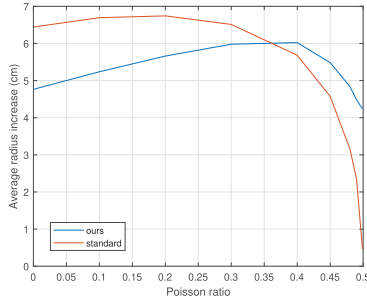

Fig. 11. Plot of the increase in the average radius of the inflated tube in Figure 1 as a function of ν. Orange: standard Neo-Hookean simulation; blue: our method. Standard FEM locks close to $\nu = 0.5$, i.e., it has a zero radius increase.

method maintains a high value of the radius even for large Poisson ratios, demonstrating our simulation is suffering less from locking.

For more incompressibility analysis, we devised an experiment where we compress an elastic cube. As can be seen in Figure 12, no noticeable difference is measured between our method and the standard Neo-Hookean material simulated in FEBio. A noticeable difference does, however, occur when increasing the pressure. When $\nu \geq 0.4999$, we are unable to achieve stable simulation with standard Neo-Hookean for pressure values higher than 15 KPa and need to use an uncoupled Mooney-Rivlin material. Therefore, our method allows for stable simulations using high pressure values for incompressible materials.

Volume preservation only happens when $\nu = 0.5$. This usually is impossible to set for standard simulators, and values around 0.49 are used instead. Smith et al. [2018] report a 4.3% volume loss for a cylinder made of over 300 thousand elements with $\nu = 0.49$. Irving et al. [2007] reported volume errors of under 1% and even 0.1% using free falling bodies with contact and moderate stretching. Our method supports $\nu = 0.5$ as the compliance matrix then becomes zero. We may still use a value of 0.4999 or similar for regularization. We used the stretched cylinder test case and

(a) $p = 15$ KPa (b) $p = 15$ KPa
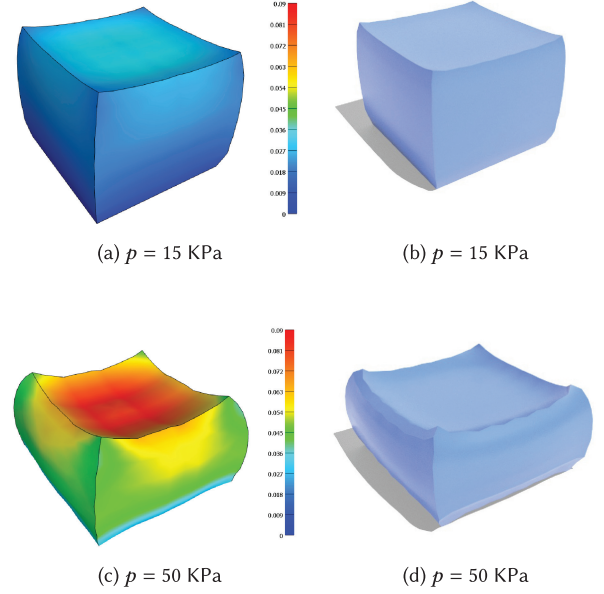

(c) $p = 50$ KPa (d) $p = 50$ KPa

Fig. 12. A rubber cube of size $0.2 \times 0.2 \times 0.2$ m compressed vertically by a traction force ($E = 66$ KPa, $\nu = 0.4999$, 10,368 tetrahedra): (a) standard Neo-Hookean in FEBio with $p = 15$ KPa, (b) our method with $p = 15$ KPa, (c) uncoupled Mooney-Rivlin material in FEBio with $p = 50$ KPa, and (d) our method with $p = 50$ KPa.

obtained a volume change of 0.015%, notably better than 1.8% for standard Neo-Hookean with $\nu = 0.49$ as shown in Figure 13. Recall that simulating standard Neo-Hookean FEM with higher Poisson ratios is very demanding on the numerical solver. We also consider volume preservation of the rubber block and compress it from above by applying a fixed amount of displacement. The standard Neo-Hookean method loses 0.14% of its volume while our method virtually preserves the volume.

Compared to hard constraint approaches (e.g., strain limiting), our method has physical meaning for the whole range of Poisson ratios through the use of compliant constraints. We exemplify using a cube that is twisted. We twist the cube first by 90 degrees and then by 180 using a Poisson ratio of 0.45 as seen in Figure 14. In this case, we cannot expect volume conservation, but the constitutive law should hold. We check this by computing the expected volume based on the current numerical estimates of the pressure and comparing the expected volume to the current volume. This is actually the residual of the constraint equation in the Newton solver. We are able to bring this measure very close to zero (on the order of 10^{-8}). For $\nu = 0.45$, this means 1.5% relative to the initial volume. Thus, our method is able to reproduce complex deformations without any visible artifacts, like the popular twisted cube, and give expected results even for compressible materials. We used our standard FEM simulator to check that the deformation looks correct.

In terms of using the logarithmic volume constraint, we have not seen major differences to the normal linear constraint in Equation (5). This applies to the numerical methods, too, as we have not seen any improvements in convergence rate due to the nonlinear constraint function. But indeed, for high compression, the
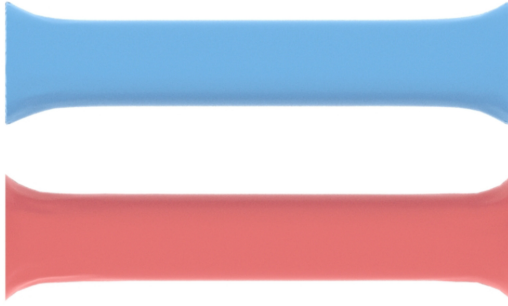
Fig. 13. Cylinder stretched twice its initial length ($E = 66$ KPa, $\rho = 1,070$ kg/m^3, and 7,594 tetrahedra). For our method (in red, $\nu = 0.4999$, 20 quasi-static steps), the volume change is only 0.015%, whereas for standard Neo-Hookean (in blue, $\nu = 0.49$, 100 quasi-static steps), the volume change is 1.8%. We used five Newton iterations for both. Even if the deformation looks almost the same, our method loses much less volume.
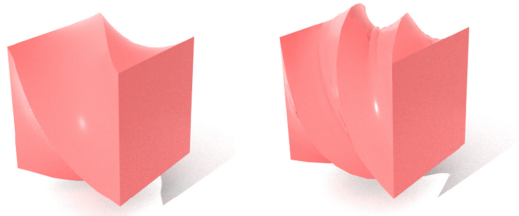

Fig. 14. Rubber cube twisted by 90 degrees (left) and 180 degress (right) using our method with: $E = 66$ KPa, $\nu = 0.45$, $\rho = 1,070$ kg/m^3, and 20,250 tetrahedra (15 divisions per side). Our method is able to reproduce such complex deformation while obeying the volumetric constitutive law. Simulated using 10 quasi-static steps and 5 Newton iterations.

logarithmic goes much faster toward infinity, and we have seen more volume preservation when using it. For example, for the scenario in Figure 12(d), we obtained a 13% volume loss using the logarithmic constraint function for $\nu = 0.4$ compared to 14.5% using the linear constraint function. But this is expected, as the Neo-Hookean material behaves similarly when switching between the logarithmic and the linear function in the quadratic penalty term. In the end, any constraint function that gives zero when $J = 1$ should work just the same for $\nu \approx 0.5$.

10 LIMITATIONS

The linear-linear mixed formulation (also known as P1-P1) always has a non-locking solution even if it does not satisfy the LBB condition—see, for example, Section 8.3.1 of Boffi et al. [2013]. This result follows from G having a non-trivial null space—see Appendix A. It does suffer, though, from inaccuracy and instability in the pressure solution. This is due to spurious modes of the pressure, i.e., oscillatory solutions, which manifest as a checkerboard pattern in the pressure map. This is a known problem of mixed interpolations using the same order for both fields and unstable mixed elements in general [Sani et al. 1981]. In Figure 15, we show that our method does suffer from the analog of checkerboarding for tetrahedra, but we are willing to ignore it as we are mainly interested in accurate displacement solutions, and the spurious pressure modes do not affect the displacements. Notice in the figure that standard
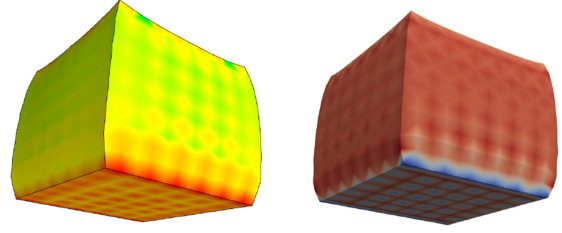

Fig. 15. The scenarios from Figure 12(a) and (b) both manifest checkerboarding in the pressure map. The map on the left was obtained in FEBio, while on the right, we used Paraview to vizualize the pressure field computed using our mixed formulation.

FEM suffers from the same problem, but it can also lock. On the other hand, the checkerboard patterns do not always appear and are heavily dependent on the mesh and boundary conditions, as described by the engineering literature [Bathe 2001].

In the context of linear elasticity, Malkus and Olsen [1984] and others proved under certain assumptions that mixed schemes that do not satisfy the LBB condition but satisfy the constraint count condition in Equation (1) are still convergent in the displacement field, even if the pressure field is wrong. This is because the former condition is sufficient, whereas the latter is necessary. Still, some authors report that the problem can become ill-posed in the case of full incompressibility or non-zero prescribed displacements [Chapelle and Bathe 1993; Silvester and Thatcher 1986].

In practice, we did notice some hard cases for the solver when using a Poisson ratio of exactly 0.5 or when inflating an air chamber at high pressure. Such examples include the cushion and the bubble robots from Figure 6 for which the Newton solver had trouble converging. On the other hand, some other incompressible cases worked just fine, like the stretched cylinder in Figure 13. Such edge cases may require more advanced solvers like continuation methods for Newton and element inversion handling.

If accurate pressures are needed, one can use filtering to smooth out the solution field [Malkus and Olsen 1984; Sani et al. 1981; Simo and Taylor 1991] or stabilized elements like the MINI element [Zienkiewicz et al. 2005]. Of course, one can always increase the order of the displacements to P2-P1. But this would be more expensive than standard FEM with second-order tetrahedra (P2), which is already performing well in terms of locking [Schneider et al. 2019b].

In terms of generality, it may seem that Assumption 2 limits it considerably. However, we did not encounter practical energies in other forms, even if assuming Mooney-Rivlin or Ogden materials or the more general Valanis-Landel hypothesis [Xu et al. 2015]. Thus, a triaxial energy function always falls under the conditions of Assumption 2. We can even replace J in Equation (8) with a different measure of volumetric change, as pointed out by Patterson et al. [2012], without having to change the ensuing method. However, this makes little sense outside the scope of linear and corotational elasticity.

If one wants to use elaborate volumetric energy functions in their original form without having to invert the constitutive law, then we consider the three-field mixed formulation (also known as *mean dilatation procedure*) as a better alternative [Bonet and Wood 1997; Simo and Taylor 1991]. This formulation includes the

volumetric constitutive law $p = \Psi'_v(\Phi)$ as an explicit constraint and circumvents the need for Assumption 2 altogether.

11 CONCLUSION

In this article, we have devised a new way to alleviate locking in linear tetrahedral elements using the mixed formulation of nonlinear FEM. We are using linear elements for both displacements and pressure and thus need to store an extra scalar value for the pressure per node. Although we employ only a necessary condition for locking-free behavior, we have shown through experiments that our method handles locking much better than standard Neo-Hookean simulations. Hence, we call our simulation meshes as made of *locking-proof tetrahedra*.

Despite making an objective of not changing the setting of linear tetrahedra much in this article, it would be interesting to further investigate stabilized linear pressure elements or higher order mixed and standard elements, including hexahedra. Another desideratum is to make the solver safer and more robust toward element inversion. Other future work will include handling frictional contact using complementarity. Collision detection should also be extended from analytic cases, i.e., primitives to triangle meshes or signed distance fields.

It is also possible to use static condensation to hide the pressure values and not store them. This would also reduce the solving time, and we expect a similar reduction from the Schur complement or dual ascent solvers. Overall, our method has all the elements of a general multibody simulator, and we would like to pursue this avenue further and focus on efficiently solving the nonlinear saddle point problem.

APPENDICES

A LOCKING PROOF FOR STANDARD FEM

Standard (or irreducible) finite element analysis usually results in linearized equations of the form

$$K_{\text{irr}} u = f. \tag{42}$$

In the linearized case. the discrete volumetric constraint is

$$\phi(u) \approx Gu = 0, \tag{43}$$

where G is the constraint Jacobian introduced in Section 4.3. As discussed in Section 4.2, we can use a quadratic penalty energy to enforce the constraint. Here, we apply the constraint regularization directly to Equation (43) and obtain a static analysis problem similar to Equation (24) from Section 4.5

$$\begin{bmatrix} K & G^T \\ G & -\frac{1}{\kappa}\mathbf{I} \end{bmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \tag{44}$$

After applying static condensation, i.e., the Schur complement from Section 8.1, we have

$$(K + \kappa G^T G)u = f. \tag{45}$$

If we consider the infinitely stiff case, i.e., $\kappa \to \infty$ (which corresponds to $v = 0.5$), we get

$$(G^T G)u = 0. \tag{46}$$

As Zienkiewicz et al. [2005] make the point, we do not want the solution of this equation to be only $\mathbf{u} = 0$ (i.e., the matrix must have a

non-trivial null space). In other words, even if we have a very large κ and a non-zero load f, we do not want to obtain a solution very close to zero, i.e., $u \approx 0$, as this would amount to locking. Therefore, we need the matrix $G^T G$ to be *singular*! Or, in other words, G has to be full rank for solvability, but still have a non-trivial null space rich enough to represent a non-locking solution [Boffi et al. 2013].

This condition is satisfied in the mixed formulation, as we are building the rectangular matrix G explicitly (in contrast to standard FEM) and we can make it as "thin" as we want. That means that G is a $n_p \times n_u$ matrix, and we impose the condition $n_p \leq n_u$, which is exactly the non-locking necesssary condition in Equation (1).

For linear elasticity, it turns out that the smallest eigenvalue of $G^T G$ is given by the ratio of shear to bulk modulus:

$$\frac{G}{K} = \frac{3}{2}\frac{1 - 2v}{1 + v}. \tag{47}$$

One can see that for $v = 0.5$, this ratio becomes 0; therefore, the matrix is singular. But as we can never really input $v = 0.5$ into K_{irr}, we will never actually get 0, i.e., $G/K \approx 0$. Hence, the matrix $G^T G$ will always be non-singular in the irreducible form, and the simulation will always lock (there is no in-between).

B VOLUMETRIC WORK

In order to prove that $\bar{p}$ and J are work conjugate, we start from the principle of virtual power applied to pure dilations, i.e., $\sigma = \bar{p}\mathbf{I}$. The variation of volumetric work per unit volume and unit time is

$$\delta w = \int_{\Omega'} -(\nabla \cdot \sigma) \cdot \delta\mathbf{v}\,dv = \int_{\Omega'} -\nabla\bar{p} \cdot \delta\mathbf{v}\,dv = \int_{\Omega'} \bar{p}(\nabla \cdot \delta\mathbf{v})\,dv, \tag{48}$$

where $\mathbf{v} = \dot{\mathbf{x}}$ is the velocity vector, $\delta\mathbf{v}$ its variation, and Ω' is the spatial integration domain. Note that we have omitted boundary terms.

The variation of the volume change J is $\delta J(\mathbf{u}) = J(\nabla \cdot \mathbf{u})$ [Bonet and Wood 1997]. If we switch over to the material integration domain using Equation (2), we get

$$\delta w = \int_{\Omega} J\bar{p}(\nabla \cdot \delta\mathbf{v})\,dV = \int_{\Omega} \bar{p}\delta J\,dV, \tag{49}$$

which proves that p and J are work conjugate variables.

Taking the argument further, we can show that p and Φ are also work conjugate. That is, given $\delta\Phi = \Phi'(J)\delta J$ and the definition of p, we can write

$$\delta w = \int_{\Omega} \frac{\bar{p}}{\Phi'}\Phi'\delta J\,dV = \int_{\Omega} p\delta\Phi\,dV, \tag{50}$$

as seen in Equation (9).

ACKNOWLEDGMENTS

REFERENCES

James Ahrens, Berk Geveci, and Charles Law. 2005. Paraview: An end-user tool for large data visualization. *The Visualization Handbook* 717, 8 (2005), 717–731.

Dina Al Akhrass, Sylvain Drapier, Julien Bruchon, and Sébastien Fayolle. 2012. Stabilized finite element methods to deal with incompressibility in solid mechanics in finite strains. In *European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS 2012)*.

Sheldon Andrews, Marek Teichmann, and Paul G. Kry. 2017. Geometric stiffness for real-time constrained multibody dynamics. *Computer Graphics Forum* 36, 2 (2017), 235–246.

Jean-François Aujol. 2009. Some first-order algorithms for total variation based image restoration. *Journal of Mathematical Imaging and Vision* 34, 3 (2009).

Ivo Babuška and Manil Suri. 1992. Locking effects in the finite element approximation of elasticity problems. *Numer. Math.* 62, 1 (1992), 439–463.

David Baraff and Andrew Witkin. 1998. Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'98)*. 43–54.

Adam W. Bargteil and Elaine Cohen. 2014. Animation of deformable bodies with quadratic Bézier finite elements. *ACM Transactions on Graphics (TOG)* 33, 3, Article 27 (2014), 10 pages.

Klaus-Jürgen Bathe. 2001. The inf–sup condition and its evaluation for mixed finite element methods. *Computers & Structures* 79, 2 (2001), 243–252.

Klaus-Jürgen Bathe. 2006. *Finite Element Procedures*. Prentice Hall.

Christopher Batty and Robert Bridson. 2008. Accurate viscous free surfaces for buckling, coiling, and rotating liquids. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 219–228.

Jan Bender, Kenny Erleben, and Jeff Trinkle. 2014. Interactive simulation of rigid body dynamics in computer graphics. *Computer Graphics Forum* 33, 1 (2014), 246–270.

Jan Bender, Matthias Müller, and Miles Macklin. 2017. A survey on position based dynamics. In *EUROGRAPHICS 2017 Tutorials*.

Michele Benzi, Gene H. Golub, and Jörg Liesen. 2005. Numerical solution of saddle point problems. *Acta Numerica* 14 (2005), 1–137.

James Bern, Pol Banzet, Roi Poranne, and Stelian Coros. 2019. Trajectory optimization for cable-driven soft robot locomotion. In *Robotics: Science and Systems (RSS 2019)*.

Daniele Boffi, Franco Brezzi, and Michel Fortin. 2013. *Mixed Finite Element Methods and Applications*. Vol. 44. Springer.

Javier Bonet and A. J. Burton. 1998. A simple average nodal pressure tetrahedral element for incompressible and nearly incompressible dynamic explicit applications. *Communications in Numerical Methods in Engineering* 14, 5 (1998), 437–449.

Javier Bonet and Richard D. Wood. 1997. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge University Press.

Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: Fusing constraint projections for fast simulation. *ACM Transactions on Graphics (TOG)* 33, 4, Article 154 (2014), 11 pages.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.

Christopher Brandt, Leonardo Scandolo, Elmar Eisemann, and Klaus Hildebrandt. 2019. The reduced immersed method for real-time fluid-elastic solid interaction and contact simulation. *ACM Transactions on Graphics (TOG)* 38, 6, Article 191 (2019), 16 pages.

Robert Bridson. 2015. *Fluid Simulation for Computer Graphics*. AK Peters/CRC Press.

Miguel Cervera, Michele Chiumenti, Quino Valverde, and Carlos Agelet de Saracibar. 2003. Mixed linear/linear simplicial elements for incompressible elasticity and plasticity. *Computer Methods in Applied Mechanics and Engineering* 192, 49 (2003), 5249–5263.

Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. 2010. A simple geometric model for elastic deformations. *ACM Transactions on Graphics (TOG)* 29, 4, Article 38 (2010), 6 pages.

Dominique Chapelle and Klaus-Jürgen Bathe. 1993. The inf-sup test. *Computers & Structures* 47, 4–5 (1993), 537–545.

Oscar Civit-Flores and Antonio Susín. 2014. Robust treatment of degenerate elements in interactive corotational FEM simulations. In *Computer Graphics Forum*, Vol. 33. 298–309.

Eulalie Coevoet, Adrien Escande, and Christian Duriez. 2017. Optimization-based inverse model of soft robots with contact handling. *IEEE Robotics and Automation Letters* 2, 3 (2017), 1413–1419.

Dimitar Dinev, Tiantian Liu, Jing Li, Bernhard Thomaszewski, and Ladislav Kavan. 2018. FEPR: Fast energy projection for real-time simulation of deformable objects. *ACM Transactions on Graphics (TOG)* 37, 4, Article 79 (2018), 12 pages.

Christian Duriez. 2013. *Real-time Haptic Simulation of Medical Procedures Involving Deformations and Device-tissue Interactions*. Habilitation à diriger des recherches. Université des Sciences et Technologie de Lille - Lille I.

Kenny Erleben, Jon Sporring, Knud Henriksen, and Henrik Dohlmann. 2005. *Physics-based Animation*. Charles River Media.

Mihai Frâncu, Arni Asgeirsson, and Kenny Erleben. 2019. High fidelity simulation of corotational linear FEM for incompressible materials. In *Motion, Interaction and Games (MIG 2019)*.

Mihai Frâncu and Florica Moldoveanu. 2017. Position based simulation of solids with accurate contact handling. *Computers & Graphics* 69 (2017), 12–23.

Rony Goldenthal, David Harmon, Raanan Fattal, Michel Bercovier, and Eitan Grinspun. 2007. Efficient simulation of inextensible cloth. *ACM Transactions on Graphics (TOG)* 26, 3, Article 49 (2007), 8 pages.

Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. Retrieved from http://eigen.tuxfamily.org.

Florian Hecht, Yeon Jin Lee, Jonathan R. Shewchuk, and James F. O'Brien. 2012. Updated sparse Cholesky factors for corotational elastodynamics. *ACM Transactions on Graphics (TOG)* 31, 5, Article 123 (2012), 13 pages.

Ulrich Heisserer, Stefan Hartmann, Alexander Düster, and Zohar Yosibash. 2008. On volumetric locking-free behaviour of p-version finite elements under finite deformations. *Communications in Numerical Methods in Engineering* 24, 11 (2008), 1019–1032.

Leonard R. Herrmann. 1965. Elasticity equations for incompressible and nearly incompressible materials by a variational theorem. *AIAA Journal* 3, 10 (1965), 1896–1900.

Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral meshing in the wild. *ACM Transactions on Graphics* 37, 4, Article 60 (2018), 14 pages.

Thomas J. R. Hughes. 2012. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Courier Corporation.

Filip Ilievski, Aaron D. Mazzeo, Robert F. Shepherd, Xin Chen, and George M. Whitesides. 2011. Soft robotics for chemists. *Angew. Chem International Edition* 50, 8 (2011), 1890–1895.

Geoffrey Irving, Craig Schroeder, and Ronald Fedkiw. 2007. Volume conserving finite element simulations of deformable models. *ACM Transactions on Graphics (TOG)* 26, 3, Article 13 (2007), 6 pages.

Geoffrey Irving, Joseph Teran, and Ronald Fedkiw. 2004. Invertible finite elements for robust simulation of large deformation. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 131–140.

Peter Kaufmann, Sebastian Martin, Mario Botsch, and Markus Gross. 2009. Flexible simulation of deformable models using discontinuous Galerkin FEM. *Graphical Models* 71, 4 (2009), 153–167.

Ryo Kikuuwe, Hiroaki Tabuchi, and Motoji Yamamoto. 2009. An edge-based computationally efficient formulation of Saint Venant-Kirchhoff tetrahedral finite elements. *ACM Transactions on Graphics (TOG)* 28, 1, Article 8 (2009), 13 pages.

Theodore Kim, Fernando De Goes, and Hayley Iben. 2019. Anisotropic elasticity for inversion-safety and element rehabilitation. *ACM Transactions on Graphics (TOG)* 38, 4, Article 69 (2019), 15 pages.

Petr Krysl and B. Zhu. 2008. Locking-free continuum displacement finite elements with nodal integration. *International Journal for Numerical Methods in Engineering* 76, 7 (2008), 1020–1043.

Tassilo Kugelstadt, Dan Koschier, and Jan Bender. 2018. Fast corotated FEM using operator splitting. *Computer Graphics Forum* 37, 8 (2018), 149–160.

Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-Newton methods for real-time simulation of hyperelastic materials. *ACM Transactions on Graphics (TOG)* 36, 3, Article 116a (2017), 16 pages.

Steve A. Maas, Benjamin J. Ellis, Gerard A. Ateshian, and Jeffrey A. Weiss. 2012. FEBio: Finite elements for biomechanics. *Journal of Biomechanical Engineering* 134, 1 (2012), 10.

Miles Macklin, Kenny Erleben, Matthias Müller, Nuttapong Chentanez, Stefan Jeschke, and Viktor Makoviychuk. 2019. Non-smooth Newton methods for deformable multi-body dynamics. *ACM Transactions on Graphics (TOG)* 38, 5, Article 140 (2019), 20 pages.

Miles Macklin, Matthias Müller, and Nuttapong Chentanez. 2016. XPBD: Position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games (MIG 2016)*. 49–54.

David S. Malkus and Elwood T. Olsen. 1984. Obtaining error estimates for optimally constrained incompressible finite elements. *Computer Methods in Applied Mechanics and Engineering* 45, 1–3 (1984), 331–353.

Pierre-Luc Manteaux, Christopher Wojtan, Rahul Narain, Stéphane Redon, François Faure, and Marie-Paule Cani. 2017. Adaptive physically based models in computer graphics. *Computer Graphics Forum* 36, 6 (2017), 312–337.

Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. 2011. Example-based elastic materials. *ACM Transactions on Graphics (TOG)* 30, 4, Article 72 (2011), 8 pages.

Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. 2011. Efficient elasticity for character skinning with contact and collisions. *ACM Transactions on Graphics (TOG)* 30, 4, Article 37 (2011), 12 pages.

Matthias Müller and Markus Gross. 2004. Interactive virtual materials. In *Proceedings of Graphics Interface 2004*. 239–246.

Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position based dynamics. *Journal of Visual Communication and Image Representation* 18, 2 (2007), 109–118.

Rahul Narain, Matthew Overby, and George E. Brown. 2016. ADMM ⊇ projective dynamics: Fast simulation of general constitutive models. In *Symposium on Computer Animation (SCA 2016)*. 21–28.

Andrew Nealen, Matthias Müller, Richard Keiser, Eddy Boxerman, and Mark Carlson. 2006. Physically based deformable models in computer graphics. *Computer Graphics Forum* 25, 4 (2006), 809–836.

Matthieu Nesme, Paul G. Kry, Lenka Jeřábková, and François Faure. 2009. Preserving topology and elasticity for embedded deformable models. *ACM Transactions on Graphics (TOG)* 28, 3, Article 52 (2009), 9 pages.

Jorge Nocedal and Stephen Wright. 2006. *Numerical Optimization*. Springer.

Raymond W. Ogden. 1997. *Non-linear Elastic Deformations*. Courier Corporation.

Jakob T. Ostien, James W. Foulk, Alejandro Mota, and Michael Veilleux. 2016. A 10-node composite tetrahedral finite element for solid mechanics. *International Journal for Numerical Methods in Engineering* 107, 13 (2016), 1145–1170.

Eric G. Parker and James F. O'Brien. 2009. Real-time deformation and fracture in a game environment. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2009)*. 165–175.

Taylor Patterson, Nathan Mitchell, and Eftychios Sifakis. 2012. Simulation of complex nonlinear elastic bodies using lattice deformers. *ACM Transactions on Graphics (TOG)* 31, 6, Article 197 (2012), 10 pages.

Alvaro G. Perez, Gabriel Cirio, Fernando Hernandez, Carlos Garre, and Miguel A. Otaduy. 2013. Strain limiting for soft finger contact simulation. In *Proceedings of the 2013 World Haptics Conference (WHC)*. 79–84.

Samuel H. M. Roth. 2002. *Bernstein-Bézier Representations for Facial Surgery Simulation*. Ph.D. Dissertation.

Samuel H. M. Roth, Markus H. Gross, Silvio Turello, and Friedrich R. Carls. 1998. A Bernstein-Bézier based approach to soft tissue simulation. *Computer Graphics Forum* 17, 3 (1998), 285–294.

Yousef Saad. 2003. *Iterative Methods for Sparse Linear Systems*. Vol. 82. SIAM.

Robert L. Sani, Philip M. Gresho, Robert L. Lee, and David F. Griffiths. 1981. The cause and cure (?) of the spurious pressures generated by certain FEM solutions of the incompressible Navier-Stokes equations: Part 1. *International Journal for Numerical Methods in Fluids* 1, 1 (1981), 17–43.

Ruediger Schmedding and Matthias Teschner. 2008. Inversion handling for stable deformable modeling. *The Visual Computer* 24, 7 (2008), 625–633.

Teseo Schneider, Jérémie Dumas, Xifeng Gao, Denis Zorin, and Daniele Panozzo. 2019a. Polyfem. Retrieved from https://polyfem.github.io/.

Teseo Schneider, Yixin Hu, Xifeng Gao, Jeremie Dumas, Denis Zorin, and Daniele Panozzo. 2019b. A large scale comparison of tetrahedral and hexahedral elements for finite element analysis. *arXiv preprint* (2019).

Joachim Schöberl. 2009. NETGEN–4. X. *RWTH Aachen University, Germany* (2009).

Martin Servin, Claude Lacoursiere, and Niklas Melin. 2006. Interactive simulation of elastic deformable materials. In *Proceedings of the 2006 Annual SIGRAD Conference; Special Theme: Computer Games*.

Jonathan R. Shewchuk. 1994. *An Introduction to the Conjugate Gradient Method without the Agonizing Pain*. Technical Report.

Hang Si. 2015. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software* 41, 2, Article 11 (2015), 36 pages.

Eftychios Sifakis and Jernej Barbic. 2012. FEM simulation of 3D deformable solids: A practitioner's guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses*. Article 20, 50 pages.

David J. Silvester and Ronald W. Thatcher. 1986. The effect of the stability of mixed finite element approximations on the accuracy and rate of convergence of solution when solving incompressible flow problems. *International Journal for Numerical Methods in Fluids* 6, 11 (1986), 841–853.

Juan C. Simo and Robert L. Taylor. 1991. Quasi-incompressible finite elasticity in principal stretches. Continuum basis and numerical algorithms. *Computer Methods in Applied Mechanics and Engineering* 85, 3 (1991), 273–310.

Breannan Smith, Fernando De Goes, and Theodore Kim. 2018. Stable neo-hookean flesh simulation. *ACM Transactions on Graphics (TOG)* 37, 2, Article 12 (2018), 15 pages.

Breannan Smith, Fernando De Goes, and Theodore Kim. 2019. Analytic eigensystems for isotropic distortion energies. *ACM Transactions on Graphics (TOG)* 38, 1, Article 3 (2019), 15 pages.

Alexey Stomakhin, Russell Howes, Craig Schroeder, and Joseph M. Teran. 2012. Energetically consistent invertible elasticity. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2012)*. 25–32.

Alexey Stomakhin, Craig Schroeder, Chenfanfu Jiang, Lawrence Chai, Joseph Teran, and Andrew Selle. 2014. Augmented MPM for phase-change and varied materials. *ACM Transactions on Graphics (TOG)* 33, 4, Article 138 (2014), 11 pages.

Jie Tan, Greg Turk, and C. Karen Liu. 2012. Soft body locomotion. *ACM Transactions on Graphics (TOG)* 31, 4, Article 26 (2012), 11 pages.

Joseph Teran, Sylvia Blemker, Victor Ng Thow Hing, and Ronald Fedkiw. 2003. Finite volume methods for the simulation of skeletal muscle. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2003)*. 68–74.

Joseph Teran, Eftychios Sifakis, Geoffrey Irving, and Ronald Fedkiw. 2005. Robust quasistatic finite elements and flesh simulation. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2005)*.

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*, Vol. 21. 205–214.

Matthias Teschner, Bruno Heidelberger, Matthias Muller, and Markus Gross. 2004. A versatile and robust model for geometrically complex deformable solids. In *Proceedings of the Computer Graphics International, 2004*. 312–319.

Emanuel Todorov. 2014. Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in MuJoCo. In *Proceedings of the 2014 IEEE International Conference on Robotics and Automation (ICRA)*. 6054–6061.

Maxime Tournier, Matthieu Nesme, Benjamin Gilles, and François Faure. 2015. Stable constrained dynamics. *ACM Transactions on Graphics (TOG)* 34, 4, Article 132 (2015), 10 pages.

Bohan Wang, George Matcuk, and Jernej Barbič. 2019. Hand modeling and simulation using stabilized magnetic resonance imaging. *ACM Transactions on Graphics (TOG)* 38, 4, Article 115 (2019), 14 pages.

Endong Wang, Qing Zhang, Bo Shen, Guangyong Zhang, Xiaowei Lu, Qing Wu, and Yajuan Wang. 2014. Intel math kernel library. In *High-Performance Computing on the Intel® Xeon Phi*. Springer, 167–188.

Huamin Wang, James O'Brien, and Ravi Ramamoorthi. 2010. Multi-resolution isotropic strain limiting. *ACM Transactions on Graphics (TOG)* 29, 6, Article 156 (2010), 10 pages.

Huamin Wang and Yin Yang. 2016. Descent methods for elastic body simulation on the GPU. *ACM Transactions on Graphics (TOG)* 35, 6, Article 212 (2016), 10 pages.

Daniel Weber, Thomas Kalbe, André Stork, Dieter Fellner, and Michael Goesele. 2011. Interactive deformable models with quadratic bases in Bernstein–Bézier-form. *The Visual Computer* 27, 6 (2011), 473–483.

Daniel Weber, Johannes Mueller-Roemer, Christian Altenhofen, André Stork, and Dieter Fellner. 2015. Deformation simulation using cubic finite elements and efficient p-multigrid methods. *Computers & Graphics* 53 (2015), 185–195.

Thomas Wihler. 2006. Locking-free adaptive discontinuous Galerkin FEM for linear elasticity problems. *Mathematics of Computation* 75, 255 (2006), 1087–1102.

Peter Wriggers. 2008. *Nonlinear Finite Element Methods*. Springer.

Hongyi Xu, Funshing Sin, Yufeng Zhu, and Jernej Barbič. 2015. Nonlinear material design using principal stretches. *ACM Transactions on Graphics (TOG)* 34, 4 (2015).

Yongning Zhu, Eftychios Sifakis, Joseph Teran, and Achi Brandt. 2010. An efficient multigrid method for the simulation of high-resolution elastic solids. *ACM Transactions on Graphics (TOG)* 29, 2, Article 16 (2010), 18 pages.

Olek C. Zienkiewicz, Robert L. Taylor, and Jian Z. Zhu. 2005. *The Finite Element Method: Its Basis and Fundamentals*. Elsevier.