SANM: A Symbolic Asymptotic Numerical Solver with Applications in Mesh Deformation

KAI JIA, MIT CSAIL, USA

(a) Inverse gravity equilibrium with incompressible neo-Hookean material. This figure shows a rest shape that will deform to the original Armadillo under gravity.

(b) Gravity equilibrium with compressible neo-Hookean material. This figure shows the final shape to which an Armadillo will deform under gravity.

(c) Controlled deformation of the Bob model with incompressible neo-Hookean material by fixing the head and moving the tail. Top: rest shape. Bottom: deformed shape. Model created by Keenan Crane.

(d) Controlled deformation via twisting and bending a horizontal bar with the As-Rigid-As-Possible (ARAP) energy.

Fig. 1. We present SANM, an open-source framework that automates and generalizes the Asymptotic Numerical Methods (ANM) to solve symbolically represented nonlinear systems via numerical continuation and higher-order approximations. We apply SANM to static equilibrium problems via continuation on gravity (1a and 1b) and controlled mesh deformation problems with an implicit homotopy formulation via continuation on control handles (1c and 1d). SANM takes a symbolic representation of the nonlinear system from the user (Listing 1) and automatically handles all the complexity of applying ANM.

Solving nonlinear systems is an important problem. Numerical continuation methods efficiently solve certain nonlinear systems. The Asymptotic Numerical Method (ANM) is a powerful continuation method that usually converges faster than Newtonian methods. ANM explores the landscape of the function by following a parameterized solution curve approximated with a high-order power series. Although ANM has successfully solved a few graphics and engineering problems, prior to our work, applying ANM to new problems required significant effort because the standard ANM assumes quadratic functions, while manually deriving the power series expansion for nonquadratic systems is a tedious and challenging task.

This paper presents a novel solver, SANM, that applies ANM to solve symbolically represented nonlinear systems. SANM solves such systems in a fully automated manner. SANM also extends ANM to support many nonquadratic operators, including intricate ones such as singular value decomposition. Furthermore, SANM generalizes ANM to support the implicit homotopy form. Moreover, SANM achieves high computing performance via optimized system design and implementation.

Author's address: Kai Jia, MIT CSAIL, 32 Vassar St, Cambridge, MA, 02139, USA, jiakai@mit.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

0730-0301/2021/8-ART79

<https://doi.org/10.1145/3450626.3459755>

We deploy SANM to solve forward and inverse elastic force equilibrium problems and controlled mesh deformation problems with a few constitutive models. Our results show that SANM converges faster than Newtonian solvers, requires little programming effort for new problems, and delivers comparable or better performance than a hand-coded, specialized ANM solver. While we demonstrate on mesh deformation problems, SANM is generic and potentially applicable to many tasks.

CCS Concepts: • **Computing methodologies** → **Optimization algorithms**; **Mesh geometry models**.

Additional Key Words and Phrases: asymptotic numerical method, nonlinear solving, finite element method, geometry processing

ACM Reference Format:

Kai Jia. 2021. SANM: A Symbolic Asymptotic Numerical Solver with Applications in Mesh Deformation. *ACM Trans. Graph.* 40, 4, Article 79 (August 2021), 16 pages. <https://doi.org/10.1145/3450626.3459755>

1 INTRODUCTION

Solving nonlinear analytic systems (systems that can be locally described by a convergent power series) is at the core of many graphics and engineering applications. Such systems are traditionally solved with Newtonian methods that essentially use a first or second order local approximation. Newtonian methods converge quadratically fast when such an approximation is accurate enough, and the initial

Fig. 2. An example of solving gravity equilibrium with numerical continuation. Starting from an initial shape $\mathbf{x}(0)$, we solve $\mathbf{x}(\lambda)$ for λ ranging from 0 to 1 subject to $\mathbf{f}(\mathbf{x}(\lambda)) + \lambda \mathbf{g} = \mathbf{0}$ where $\mathbf{f}(\cdot)$ computes node elastic forces given node coordinates, and $\mathbf{g}$ is the per-node gravity vector. ANM parameterizes both $\mathbf{x}$ and λ with a , and approximates $\mathbf{x}(a)$ and $\lambda(a)$ with power series. ANM allows easily computing the intermediate equilibrium states almost for free.

guess is sufficiently close. However, these assumptions are often violated in practice, and the convergence is thus much slower [Bonnans et al. 2006].

This paper considers solving the system $\mathbf{f}(\mathbf{x}) + \mathbf{v} = \mathbf{0}$ under a numerical continuation framework [Allgower and Georg 2003]. Here $\mathbf{x} \in \mathbb{R}^n$ is an unknown vector, $\mathbf{f} : \mathbb{R}^n \mapsto \mathbb{R}^n$ is an analytic function, and $\mathbf{v}$ is a constant. Given an initial solution $\mathbf{x}_0$ such that $\mathbf{f}(\mathbf{x}_0) = \mathbf{0}$, numerical continuation methods trace the final solution via solving $\mathbf{x}(\lambda)$ with λ ranging from 0 to 1 subject to $\mathbf{f}(\mathbf{x}(\lambda)) + \lambda \mathbf{v} = \mathbf{0}$. For example, in the static elasticity equilibrium problem, we encode the unknown node coordinates in $\mathbf{x}$, the mapping from node coordinates to node forces in $\mathbf{f}(\cdot)$, and the static external force in $\mathbf{v}$. By setting $\mathbf{x}_0$ as the rest shape, numerical continuation corresponds to gradually increasing the external force while simulating the deformation simultaneously. Figure 2 presents an example.

The Asymptotic Numerical Method (ANM) [Damil and Potier-Ferry 1990] is a numerical continuation method that differs fundamentally from Newtonian approaches by exploring the landscape of the nonlinear system via higher-order approximations. The numerical continuation formulation $\mathbf{f}(\mathbf{x}(\lambda)) + \lambda \mathbf{v} = \mathbf{0}$ defines a solution curve $(\mathbf{x}(\lambda), \lambda)$. However, parameterization of the curve using λ can result in ill-conditioned behavior of $\mathbf{x}(\lambda)$. Instead, ANM parameterizes both $\mathbf{x}$ and λ with a such that $\mathbf{f}(\mathbf{x}(a)) + \lambda(a)\mathbf{v} = \mathbf{0}$. ANM approximates the solution curve with a power series expansion at truncation order N : $\mathbf{x}(a) = \sum_{i=0}^N \mathbf{x}_i a^i$ and $\lambda(a) = \sum_{i=0}^N \lambda_i a^i$. Cochelein [1994] proposes a continuation technique to compute the final solution in a stepwise manner. Specifically, they estimate the valid range of the current approximation as a_r and compute the final solution by iteratively recomputing the approximation at $\mathbf{x}_0 = \mathbf{x}(a_r)$ and $\lambda_0 = \lambda(a_r)$ until $\lambda(a_r)$ reaches 1.

Compared to the widely used Newtonian methods, ANM is able to determine a local representation of the solution curve with a larger range of validity in similar computing time [Cochelein et al. 1994b], and therefore solves the nonlinear system in fewer iterations and often less running time.

The core challenge in applying ANM is to solve the expansion coefficients $\{\mathbf{x}_i\}$ and $\{\lambda_i\}$. The original ANM framework [Damil and Potier-Ferry 1990] analytically solves the coefficients for quadratic functions $\mathbf{f}(\mathbf{x}) = \mathbf{x}_0 + \mathbf{A}\mathbf{x} + \mathbf{x}^\top \mathbf{Q}\mathbf{x}$. Chen et al. [2014] deploys ANM on the inverse deformation problem for 3D fabrication by manually deriving the coefficient solution for the incompressible neo-Hookean elasticity [Bonet and Wood 2008]. They claim their

Taylor coefficient derivation as a major contribution, which is a difficult and laborious task. They have also shown that ANM converges up to orders of magnitudes faster than Newtonian solvers.

To date, however, there is no scalable tool that automates the computation of Taylor coefficients in the general case. The lack of such tools severely limits the application of ANM to new problems. In this work, we show how to solve the coefficients $\{\mathbf{x}_i\}$ and $\{\lambda_i\}$ automatically and efficiently for a symbolically defined function $\mathbf{f}(\cdot)$. More specifically, we devise techniques to establish the connection of Taylor expansion coefficients between $\mathbf{x}(a)$ and $(\mathbf{f} \circ \mathbf{x})(a)$, which in fact conforms to an affine relationship for the highest-order term. We analyze a few operators important for graphics applications, including elementwise analytical functions (such as power and logarithm), matrix inverse, matrix determinant, and singular value decomposition. We speed up the system with batch computing that fits naturally into Finite Element Method (FEM) due to the same computation form shared by all quadrature points. We present a system, called SANM, as an implementation of our techniques.

We deploy SANM to solve the forward and inverse static force equilibrium problems similar to Chen et al. [2014]. In contrast to their manual derivation that only works with incompressible neo-Hookean materials, our system allows easily solving more constitutive models by changing a few lines of code, including the compressible neo-Hookean model that has a logarithm term and the As-Rigid-As-Possible energy that involves a polar decomposition. Our experimental results show that SANM achieves comparable or better performance as the hand-coded, specialized ANM solver of Chen et al. [2014]. We are unaware of efficient alternative methods for the inverse problem other than ANM. For the forward problem, an alternative is to minimize the total potential energy, and we show that SANM exhibits better performance than Newtonian energy minimizers.

We further extend the ANM framework to incorporate implicit homotopy $\mathbf{H}(\mathbf{x}, \lambda) = \mathbf{0}$ where $\mathbf{H} : \mathbb{R}^{n+1} \mapsto \mathbb{R}^n$ admits a one-dimensional solution curve. For instance, we formulate the controlled mesh deformation as an implicit homotopy problem, defined as $\mathbf{H}(\mathbf{x}, \lambda) = \mathbf{f}([\mathbf{x}; \mathbf{x}_c + \lambda \Delta_{\mathbf{x}}])$, where $\mathbf{x}_c$ corresponds to the initial location of control handles, $\Delta_{\mathbf{x}}$ describes their user-specified movement, and $\mathbf{f}(\cdot)$ computes the internal elastic force. The coordinates of unconstrained nodes in the deformed equilibrium state, denoted by $\mathbf{x}^*$, are then governed by $\mathbf{H}(\mathbf{x}^*, 1) = \mathbf{0}$ and can be solved by continuation on λ . In our experiments, SANM runs 1.41 times faster by geometric mean than Newtonian energy minimization methods.

We also demonstrate the robustness and versatility of SANM by twisting and bending a bar to extreme poses, as shown in Figure 1d. Note that ANM and Newtonian minimization methods target different problems and can not replace each other. Section 7.3 further discusses their differences.

To summarize, this paper makes the following contributions:

- (1) We devise analytical solutions for Taylor coefficient propagation through a few nonlinear operators on which ANM has not been applied, including singular value decomposition as a challenging case (Section 5).
- (2) We present a system, SANM, that automatically solves the Taylor expansion coefficients for symbolically defined functions (Section 4). SANM greatly reduces programming effort for adopting ANM-based methods (Section 6.1). SANM adopts generic and FEM-specific optimizations to improve solving efficiency further.
- (3) We present a novel continuation algorithm to reduce accumulated numerical error and approximation error when solving the equational form $f(\mathbf{x}) + \mathbf{v} = \mathbf{0}$ (Section 4.3).
- (4) We extend ANM to handle implicit homotopy and apply it to controlled mesh deformation problems (Section 7.2). Our experiments show that SANM often converges faster than a state-of-the-art Newtonian energy minimizer. Moreover, the numerical continuation framework of SANM directly handles constitutive models that do not support inverted tetrahedrons, which would be challenging for energy minimization methods due to undefined elastic energy at the initial guess.

SANM is available at <https://github.com/jia-kai/SANM>.

2 RELATED WORK

Numerical Optimization: Numerical optimization has been extensively studied, and it is closely related to solving nonlinear systems. For example, we can recast solving $f(\mathbf{x}) = \mathbf{0}$ as minimizing $g(\mathbf{x}) = f(\mathbf{x})^\top f(\mathbf{x})$ and apply generic minimization methods such as the Levenberg–Marquardt algorithm. On the other hand, minimizing $f(\mathbf{x})$ can often be approached via solving $\nabla f(\mathbf{x}) = \mathbf{0}$. For controlled mesh deformation problems, the internal elastic force corresponds to the gradient of the elastic potential energy with respect to node locations. Therefore, one can either directly solve a force equilibrium under Dirichlet boundary conditions (as done by SANM) or minimize the total potential energy to obtain the deformed state. We review the development of As-Rigid-As-Possible (ARAP) energy minimization as an example of improvements on numerical optimizers. Sorkine and Alexa [2007] devises a surface modeling technique by minimizing the ARAP energy via alternating between fitting the rotations and optimizing the locations. Chao et al. [2010] employs a Newton trust region solver to minimize the ARAP energy. Shtengel et al. [2017] accelerates the convergence by computing a positive semidefinite Hessian via constructing a convex majorizer for a specific class of convex-concave decomposable objectives, including the 2D ARAP energy. Smith et al. [2019] presents analytical solutions for the eigensystems of isotropic distortion energies to enable easily projecting the Hessians of 2D and 3D ARAP energies to be positive semidefinite to speed up the convergence. Most optimization methods inherently build on the classic idea of using first or second

order approximations and exploit problem-specific optimization opportunities. This work targets generic nonlinear solving with numerical continuation and uses higher-order approximation.

Mesh Deformation: Mesh deformation control is an important and widely studied problem in graphics. For animation production that only requires plausible but not physically accurate results, the simulation performance can be improved by a variety of approaches such as model analysis [Choi and Ko 2005; Kim and James 2009], skinning [Gilles et al. 2011], and constraint projection [Bender et al. 2014; Bouaziz et al. 2014]. For physically predictive simulations, we need to stick to the formulation derived from continuum mechanics strictly, and the solver convergence rate is often improved by Hessian modification [Kim and Eberle 2020; Shtengel et al. 2017]. In this paper, we choose physically accurate elastic deformation as our target application. We approach the problem by solving a nonlinear system that encodes force equilibrium constraints.

Numerical Continuation Methods: Classic numerical continuation methods include the predictor-corrector method and the piecewise-linear method. Allgower and Georg [2003] provides an introduction to this topic. The basic idea, which is to follow a solution trajectory by taking small steps, has become popular in many applications such as motion planning [Duenser et al. 2020; Yin et al. 2008], MRI reconstruction [Trzasko and Manduca 2008], and drawing assistance [Limpaecher et al. 2013]. These works typically choose a fixed step size or adopt a problem-specific step size schedule in the predictor and use classic first or second order solvers as the corrector. By contrast, asymptotic numerical methods use a higher-order approximation as the predictor without needing a corrector and adaptively choose the step size according to how well the predictor approximates the system.

Asymptotic Numerical Methods: ANM has been applied to solve engineering problems in different domains, including buckling analysis [Azrar et al. 1993; Bouty et al. 2004], vibration analysis [Azrar et al. 2002; Daya and Potier-Ferry 2001], shell and rod simulation [Lazarus et al. 2013; Zahrouni et al. 1999], and inverse deformation problems [Chen et al. 2014]. ANM assumes a quadratic system. An improvement over the standard ANM framework is to increase the range of validity of the approximation via imposing heuristics on the function behavior, such as replacing the power series with a Padé representation [Cochelin et al. 1994a; Elhage-Hussein et al. 2000; Najah et al. 1998]. When adapting ANM to new problems, one typically needs to recast their specific problems into quadratic forms by introducing auxiliary variables and deriving the expansions manually [Guillot et al. 2019]. Abichou et al. [2002] presents a review on adapting ANM for a few nonlinear functions.

Few attempts have been made to automate ANM to handle general nonlinearities. Notably, Charpentier et al. [2008] proposes an automatic differentiation framework, called Diamant, that computes the expansion coefficients by Taylor coefficient propagation via computing higher-order derivatives of the operators. Unfortunately, the Diamant approach is not readily applicable to mesh deformation problems due to the difficulty in computing higher-order derivatives of certain matrix functions such as matrix inverse or determinant used in the constitutive models. Moreover, Diamant is not designed

with high-performance computing in mind. It only works with scalar variables, does not take advantage of the structural sparsity in FEM problems, and is only evaluated on small-sized problems. Lejeune et al. [2012] incorporates the Diamant approach into an object-oriented solver to automate ANM. By contrast, SANM natively works with multidimensional variables and is accelerated with batch computing for large-scale FEM problems. SANM also implements a generic framework for computing the expansion coefficients, which is not limited to the higher-order derivative approach of Diamant and is capable of handling challenging matrix functions.

3 ANM BACKGROUND

This section introduces the asymptotic numerical method. We begin with a toy example of a geometry problem and then formally describe ANM. We first define the notations used in this paper in Table 1.

Table 1. Notation definition

f, x	Scalars or scalar-valued functions
$\mathbf{f}, \mathbf{x}$	Vectors or vector-valued functions
f_i, x_i	A scalar in the vector at given index. For a function $f(\cdot)$, we also use f_i to represent its i^{th} Taylor coefficient, and similarly for $f(\cdot)$ vs f_i and $F(\cdot)$ vs F_i .
$\mathbf{f}_i, \mathbf{x}_i$	A vector in an array of vectors
$F, \mathbf{X}$	Matrices or matrix-valued functions
$F_i, \mathbf{X}_i$	A matrix in an array of matrices
X_{ij}	A coefficient in the matrix at given row and column
$\mathbf{X}_i, \mathbf{X}_j$	The vectors corresponding to the i^{th} row or the j^{th} column in matrix $\mathbf{X}$
$\text{RMS}(\mathbf{x})$	Root-mean-square of $\mathbf{x} \in \mathbb{R}^n$: $\text{RMS}(\mathbf{x}) = \sqrt{\mathbf{x}^\top \mathbf{x} / n}$
$\text{vec}(\mathbf{X})$	Flatten a matrix $\mathbf{X}$ into a column vector by concatenating the columns in $\mathbf{X}$
$\ \mathbf{X}\ $	Frobenius norm of the matrix $\mathbf{X}$, defined as $\sqrt{\text{vec}(\mathbf{X})^\top \text{vec}(\mathbf{X})}$
$\text{diag}(\mathbf{X})$	A vector containing the diagonal coefficients of $\mathbf{X}$
$o(x)$	The little-o notation: $y = o(x)$ if $y/x \rightarrow 0$.

3.1 A Circle-ellipse Intersection Problem

We illustrate ANM with a toy problem that asks for the intersection B of an ellipse e and a circle c as shown in Figure 3. The ellipse intersects the y -axis at $B_0 = (0, -1)$. The circle is centered at $A = (-1, 0)$ and its radius is $|AB| = \sqrt{8}$. To solve the problem with ANM, we first choose a continuation scheme. We start with a smaller concentric circle c_0 with radius $|AB_0| = \sqrt{2}$ and continuously trace its intersection with e while increasing its radius from $\sqrt{2}$ to $\sqrt{8}$. Note that the continuation scheme is problem-specific. While there can be many choices, practical problems typically admit a “natural” choice, such as the external force in static equilibrium problems.

Formally, our goal is to solve (x, y) such that $f_e(x, y) = f_c(x, y) = 0$, where $f_e(\cdot)$ and $f_c(\cdot)$ describe the ellipse e and the circle c respectively:

$$\begin{aligned} f_e(x, y) &= 2x^2 - 5x + y^2 - 4y - 2xy - 5 \\ f_c(x, y) &= (x + 1)^2 + y^2 - 8 \end{aligned} \quad (1)$$

Fig. 3. Our circle-ellipse intersection example problem. ANM uses polynomials to approximate the trace of intersection points between the ellipse and concentric circles of radius from $\sqrt{2}$ to $\sqrt{8}$, which is the arc B_0B .

ANM introduces a variable $\lambda \in [0, 1]$ to represent the continuation. ANM traces the solution curve starting at (x_0, y_0) via varying λ from 0 to 1 while keeping the following equations satisfied:

$$\begin{aligned} f_e(x, y) &= 0\lambda \\ f_c(x, y) + 6 &= 6\lambda \\ (x_0, y_0) &= B_0 = (0, -1) \end{aligned} \quad (2)$$

Geometrically, ANM continuously solves the intersection between the ellipse and a concentric circle with radius $\sqrt{2 + 6\lambda}$. ANM parameterizes the solution curve by a variable a and approximates $x(a)$, $y(a)$, and $\lambda(a)$ with polynomial expansions at truncation order N , with coefficients $\{x_k\}$, $\{y_k\}$, and $\{\lambda_k\}$ to be solved:

$$\begin{aligned} x(a) &= 0 + \sum_{k=1}^N x_k a^k \\ y(a) &= -1 + \sum_{k=1}^N y_k a^k \\ \lambda(a) &= 0 + \sum_{k=1}^N \lambda_k a^k \end{aligned} \quad (3)$$

We iteratively solve the coefficients by introducing the lower-order terms in (3) into (2). We start with (x_1, y_1, λ_1) and introduce $x = x_1 a$, $y = -1 + y_1 a$, and $\lambda = \lambda_1 a$ into (2):

$$\begin{aligned} f_e(x_1 a, y_1 a - 1) &= -(3x_1 + 6y_1)a + o(a) = 0 \\ f_c(x_1 a, y_1 a - 1) + 6 &= (2x_1 - 2y_1)a + o(a) = 6\lambda_1 a \end{aligned} \quad (4)$$

We obtain two linear constraints on the three unknowns by equating the coefficient of a in (4). Let $\mathbf{u}(a) = [x(a); y(a); \lambda(a)]$ denote the solution curve. ANM further identifies the path parameter a as the pseudo-arclength that is the projection of the path along its tangent direction as shown in Figure 4, specifically $a = (\mathbf{u}(a) - \mathbf{u}(0))^\top \mathbf{u}'(0)$, which provides the third constraint for a full-rank system:

$$\begin{cases} -3x_1 - 6y_1 = 0 \\ 2x_1 - 2y_1 - 6\lambda_1 = 0 \\ x_1^2 + y_1^2 + \lambda_1^2 = 1 \end{cases} \quad (5)$$

We require λ_1 to be positive so that $\lambda(a)$ is a locally increasing function at 0, and the solution of (5) is $x_1 = 2/\sqrt{6}$, $y_1 = -1/\sqrt{6}$, and

Fig. 4. ANM solution curve for the circle-ellipse intersection problem. ANM solves the problem in two iterations. The path parameter a is identified by the projection of a point $(x(a), y(a), \lambda(a))$ onto the tangent direction. The length of the projection is $a|\nabla|$ where $\nabla = (x'(0), y'(0), \lambda'(0))$ is the gradient at the beginning of the current iteration.

$\lambda_1 = 1/\sqrt{6}$. We then solve (x_2, y_2, λ_2) by equating the coefficients of a^2 in $f_e(x_2a^2 + 2/\sqrt{6}a, y_2a^2 - 1/\sqrt{6}a - 1) = 0$ and $f_c(x_2a^2 + 2/\sqrt{6}a, y_2a^2 - 1/\sqrt{6}a - 1) = 6(\lambda_2a^2 + 1/\sqrt{6}a)$. Expanding the equations results in two linear constraints for the three unknowns. The third pseudo-arclength constraint is $x_1x_k + y_1y_k + \lambda_1\lambda_k = 0$ for $k \geq 2$. Repeating this step, we can solve the coefficients $\{x_k\}$, $\{y_k\}$, and $\{\lambda_k\}$ that define the polynomials $x(a)$, $y(a)$, and $\lambda(a)$. As will be shown in Section 3.3, the equations for all (x_k, y_k, λ_k) are linear, and the coefficients in these equations are the gradients of $f_e(\cdot)$ and $f_c(\cdot)$ evaluated at (x_0, y_0) .

We then estimate a_r , the range of validity of the polynomial approximations $x(a)$, $y(a)$, and $\lambda(a)$. We iteratively compute a new approximation at $(x(a_r), y(a_r))$ to extend the solution curve until $\lambda(a_r) \geq 1$, and we compute the final solution $(x(a^*), y(a^*))$ with $a^* = \lambda^{-1}(1)$. ANM with truncation order $N = 20$ is able to find the circle-ellipse intersection with a residual (defined as $\sqrt{(f_e^2(x) + f_c^2(x))/2}$) of 2×10^{-6} in two iterations. Figure 4 visualizes the parametric solution curve. The continuation formulation presented in Section 4.3 further reduces the residual to 7×10^{-9} .

An alternative is the Newton-Raphson method, which iteratively computes $\mathbf{x}_{n+1} = \mathbf{x}_n - \mathbf{J}(\mathbf{x}_n)^{-1} \mathbf{f}(\mathbf{x}_n)$, where $\mathbf{f}(\mathbf{x}) = [f_e(\mathbf{x}); f_c(\mathbf{x})]$ and $\mathbf{J}(\mathbf{x})$ is the Jacobian of $\mathbf{f}$ evaluated at $\mathbf{x}$. Starting at $\mathbf{x}_0 = (0, -1)$, the Newton-Raphson method needs four evaluations of the Jacobian to converge to a solution with a residual of 8×10^{-5} .

3.2 ANM Overview

ANM aims to solve the nonlinear system $\mathbf{f}(\mathbf{x}) + \mathbf{v} = \mathbf{0}$ with numerical continuation, where $\mathbf{f} : \mathbb{R}^n \mapsto \mathbb{R}^n$ is an analytic function, and $\mathbf{v}$ is a constant. Starting from an initial solution $(\mathbf{x}_0, \lambda_0)$ such that $\mathbf{f}(\mathbf{x}_0) + \lambda_0\mathbf{v} = \mathbf{0}$, continuation methods compute an approximation to trace the nearby solution curve of $\mathbf{f}(\mathbf{x}) + \lambda\mathbf{v} = \mathbf{0}$. In the general

case, the curve may not be well-conditioned under the parameterization with respect to λ , and it is preferable to consider an arclength parameterization $\mathbf{x}(a)$ and $\lambda(a)$ where a measures the arclength or pseudo-arclength [Allgower and Georg 2003].

ANM approximates $\mathbf{x}(a)$ and $\lambda(a)$ by Taylor expansion at truncation order N such that $\mathbf{f}(\mathbf{x}(a)) + \lambda(a)\mathbf{v}$ should be sufficiently close to zero for small values of a :

$$\begin{aligned} \mathbf{x}(a) &= \sum_{i=0}^N \mathbf{x}_i a^i \\ \lambda(a) &= \sum_{i=0}^N \lambda_i a^i \\ \text{s.t. } \|\mathbf{f}(\mathbf{x}(a)) + \lambda(a)\mathbf{v}\| &= o(a^N) \end{aligned} \quad (6)$$

We require $\lambda_1 > 0$ so that $\lambda(a)$ is locally increasing, and the algorithm makes progress. We then estimate the range of validity a_r such that $\mathbf{x}(a)$ and $\lambda(a)$ are good approximations of the solution when $|a| < a_r$. If $\lambda(a_r) \geq 1$, we can solve a^* such that $\lambda(a^*) = 1$ and compute the final solution $\mathbf{x}^* = \mathbf{x}(a^*)$. Otherwise, when $\lambda(a_r) < 1$, we recompute the power series approximation starting at $\mathbf{x}_0 = \mathbf{x}(a_r)$ and $\lambda_0 = \lambda(a_r)$ and repeat the above steps.

A simple method to estimate a_r , as suggested by Cochelin [1994], builds on the idea that within the range of validity, different orders of approximation should behave similarly:

$$\frac{\|\mathbf{x}(a)_{\text{order } N} - \mathbf{x}(a)_{\text{order } N-1}\|}{\|\mathbf{x}(a)_{\text{order } N} - \mathbf{x}_0\|} < \epsilon \quad (7)$$

which leads to an approximation

$$a_r \approx \left(\epsilon \frac{\|\mathbf{x}_1\|}{\|\mathbf{x}_N\|} \right)^{\frac{1}{N-1}} \quad (8)$$

The equation $\lambda(a^*) = 1$ can be solved by a univariate polynomial root finding algorithm such as Brent's method [Brent 2013].

The remaining part of completing the ANM algorithm is to solve the coefficients $\{\mathbf{x}_i\}$ and $\{\lambda_i\}$ efficiently, which is a core contribution of this paper.

We solve the coefficients $\{\mathbf{x}_i\}$ and $\{\lambda_i\}$ iteratively. Assume we are at the k^{th} iteration, where $\{\mathbf{x}_0, \dots, \mathbf{x}_{k-1}\}$ and $\{\lambda_0, \dots, \lambda_{k-1}\}$ have been solved. With $\mathbf{x}_k \in \mathbb{R}^n$ and $\lambda_k \in \mathbb{R}$ currently unknown, we have the equation:

$$\mathbf{f}(\mathbf{x}(a)) + \lambda(a)\mathbf{v} \approx \mathbf{f}\left(\sum_{i=0}^k \mathbf{x}_i a^i\right) + \left(\sum_{i=0}^k \lambda_i a^i\right)\mathbf{v} \approx \mathbf{0} \quad (9)$$

Assume $\{f_i\}$ are the Taylor coefficients of $\mathbf{f}(\mathbf{x}(a))$:

$$\mathbf{f}(\mathbf{x}(a)) \approx \mathbf{f}\left(\sum_{i=0}^k \mathbf{x}_i a^i\right) = \sum_{i=0}^k f_i a^i + o(a^k) \quad (10)$$

As will be shown in Proposition 1, for an analytic function $\mathbf{f}(\cdot)$, there is an affine relationship between f_k and $\mathbf{x}_k$, specifically $f_k = P(\mathbf{x}_0)\mathbf{x}_k + \mathbf{q}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1})$, where $P(\mathbf{x}_0) \in \mathbb{R}^{n \times n}$ is the slope matrix and $\mathbf{q}(\cdot) \in \mathbb{R}^n$ is the bias vector. By introducing this relationship into the original equation and requiring the coefficient of a^k to be zero, we obtain a linear system that restricts $\mathbf{x}_k$ and λ_k :

$$P(\mathbf{x}_0)\mathbf{x}_k + \lambda_k\mathbf{v} = -\mathbf{q}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}) \quad (11)$$

However, the system has rank n but there are $n+1$ unknowns because the curve behavior with respect to its parameter a is not fully constrained. To obtain a full rank system, Cochelin [1994]

proposes to identify the path parameter a as the pseudo-arclength, similar to other numerical continuation methods [Allgower and Georg 2003]. Pseudo-arclength approximates the arclength of a curve by projecting it onto the tangent space, which constitutes the following constraint:

$$a = (\mathbf{x}(a) - \mathbf{x}_0)^\top \mathbf{x}'(0) + (\lambda(a) - \lambda_0)\lambda'(0) \quad (12)$$

Introduce (6) into (12):

$$\mathbf{x}_i^\top \mathbf{x}_1 + \lambda_i \lambda_1 = \mathbb{1}_{i=1} \quad (13)$$

The unknowns $\mathbf{x}_k$ and λ_k can be solved by combining (11) and (13). The solution is unique if we further require $\lambda'(0) = \lambda_1 > 0$.

3.3 Linearity Between Taylor Coefficients

Before proving the linearity between f_k and $\mathbf{x}_k$ in (10), we introduce an auxiliary definition:

DEFINITION 1. Define $\mathfrak{C}_a^{(k)}[f(a)]$ to be the coefficient of a^k in the Taylor expansion of an analytic function $f : \mathbb{R} \mapsto \mathbb{R}$ such that:

$$f(a) = \sum_{i \geq 0} \mathfrak{C}_a^{(i)}[f(a)] a^i$$

PROPOSITION 1. Let $f : \mathbb{R}^n \mapsto \mathbb{R}$ be an analytic function, and $\mathbf{x}_0, \dots, \mathbf{x}_{k-1}$ known real-valued n -dimensional vectors. Assume coefficients $f_0, \dots, f_k$ satisfy:

$$f\left(\sum_{i=0}^k \mathbf{x}_i a^i\right) = \sum_{i=0}^k f_i a^i + o(a^k)$$

Then we have $f_k = \mathbf{p}(\mathbf{x}_0)^\top \mathbf{x}_k + \mathbf{q}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1})$. Specifically, $\mathbf{p}(\mathbf{x}_0) = \nabla f(\mathbf{x}_0)$ and $\mathbf{q}(\mathbf{x}_0, \dots, \mathbf{x}_{k-1}) = \mathfrak{C}_a^{(k)}\left[f\left(\sum_{i=0}^{k-1} \mathbf{x}_i a^i\right)\right]$.

PROOF. We prove the univariate case for the simplicity of the notations. Our argument also applies to multivariate functions by using the multivariate Taylor theorem.

Let $g(t) = f(x_0 + t) = \sum_{i \geq 0} g_i t^i$ where $g_i = \frac{1}{i!} \frac{d^i g}{dt^i} \Big|_{t=0}$ is the Taylor expansion coefficient and $g_0 = f_0 = f(x_0)$. With $t = \sum_{i=1}^k \mathbf{x}_i a^i$, we rewrite

$$f\left(\sum_{i=0}^k \mathbf{x}_i a^i\right) = g\left(\sum_{i=1}^k \mathbf{x}_i a^i\right) = g_0 + \sum_{i \geq 1} g_i \left(\sum_{j=1}^k \mathbf{x}_j a^j\right)^i$$

To compute f_k , we consider the terms that contribute to a^k in the expansion of the right side hand:

- (1) For terms that contain $\mathbf{x}_k$ and contribute to f_k , it must contain $\mathbf{x}_k a^k$ and no other $\mathbf{x}_i a^i$ for which $i > 0$. There is only one such term, which is $g_1 \mathbf{x}_k a^k$. The slope p for which $f_k = p \mathbf{x}_k + q$ is thus $p = g_1 = f'(x_0)$.
- (2) The bias q consists of terms that do not contain $\mathbf{x}_k$, which can be computed by treating $\mathbf{x}_k$ as zero, or equivalently removing $\mathbf{x}_k a^k$ from t :

$$q = \mathfrak{C}_a^{(k)}\left[f\left(\sum_{i=0}^{k-1} \mathbf{x}_i a^i\right)\right]$$

□

Remarks: We have discussed the case of scalar functions. For a vector function $\mathbf{f} : \mathbb{R}^n \mapsto \mathbb{R}^m$, the slope matrix $\mathbf{P}$ is its Jacobian. Note that $\mathbf{p}(\mathbf{x}_0)$ only depends on the initial point $\mathbf{x}_0$ and remains constant for all orders. This allows us to factorize the coefficient matrix only once to solve all the terms. The result of Proposition 1 is not new. For example, it is a direct consequence of Faà di Bruno's formula [Roman 1980]. Here we have presented a simple proof based on elementary calculus.

3.4 Continuation with Padé Approximation

The original ANM approximates $\mathbf{x}(a)$ and $\lambda(a)$ with Taylor expansions. It has been shown that replacing the Taylor expansions with Padé approximations results in a larger range of validity and thus fewer iterations.

For a scalar function $f : \mathbb{R} \mapsto \mathbb{R}$, its Padé approximation of order $M + N$ approximates the function with a ratio of two polynomials, $P_N(x)$ and $Q_M(x)$ of degrees N and M , respectively, such that $f(x) = \frac{P_N(x)}{Q_M(x)} + o(x^{N+M})$. The polynomials $P_N(x)$ and $Q_M(x)$ are determined by the first $N + M + 1$ Taylor coefficients of f via a set of linear constraints up to a common scaling factor. Although the Padé approximation is constructed from the information contained in the Taylor expansion, for many functions, it has a larger range of validity than the Taylor series, and it sometimes gives meaningful results even when the radius of convergence of the Taylor series is strictly zero [Basdevant 1972].

Cochelin et al. [1994a] proposes to construct a Padé approximation for the vector function $\mathbf{x}(a) = \sum_{i=0}^N \mathbf{x}_i a^i$ in the form:

$$\begin{aligned} P_N(a) &= \mathbf{x}_0 + \sum_{i=1}^{N-1} \frac{D_{N-1-i}(a)}{D_{N-1}(a)} \mathbf{x}_i a^i \\ D_k(a) &= \sum_{i=0}^k d_i a^i \end{aligned} \quad (14)$$

Najah et al. [1998] presents the process of computing the coefficients $\{d_i\}$ from $\{\mathbf{x}_i\}$, which first orthonormalizes $\{\mathbf{x}_i\}$ and then solves $\{d_i\}$ based on the principle that the Padé approximation should share the same lower-order Taylor coefficients. We omit the details here.

We follow the techniques presented in Elhage-Hussein et al. [2000] to determine the range of validity a_p for this Padé approximation:

$$\frac{\|P_N(a_p) - P_{N-1}(a_p)\|}{\|P_N(a_p) - P_N(0)\|} < \epsilon \quad (15)$$

The number a_p is sought via bisection in the range (a_r, r) , where a_r is the range of validity of the Taylor series determined by (8), and r is the smallest positive real root of $D_{N-1}(a)$ that can be found by numerical methods such as Bairstow algorithm [Golub and Robertson 1967]. Note that P_{N-1} is not the first $N - 1$ terms in P_N but the Padé approximation computed from $\{\mathbf{x}_0, \dots, \mathbf{x}_{N-1}\}$. For all the 36 test cases presented in Section 7, Padé approximation uses 1.14 ± 0.57 fewer iterations than the original ANM formulation on average.

Algorithm 1 Taylor coefficient solver in SANM

Input: Computing graph $G = (V_o \cup V_v, E)$
that represents $H : \mathbb{R}^{n+1} \mapsto \mathbb{R}^n$
Input: Initial value $\mathbf{x}_0 \in \mathbb{R}^n$ and $\lambda_0 \in \mathbb{R}$ such that $H(\mathbf{x}_0, \lambda_0) = \mathbf{0}$
Input: Truncation order N
Output: Taylor coefficients $\{\mathbf{x}_i\}$ and $\{\lambda_i\}$ such that
 $\|H(\sum_{i=0}^N \mathbf{x}_i a^i, \sum_{i=0}^N \lambda_i a^i)\| = o(a^N)$

Compute $P \leftarrow \frac{\partial H}{\partial \mathbf{x}}(\mathbf{x}_0, \lambda_0)$ by reverse mode AD on G
Meanwhile, record the Jacobians $\mathbf{p}^u$ for each operator $u \in V_o$
Meanwhile, record $\mathbf{v} \leftarrow \frac{\partial H}{\partial \lambda}(\mathbf{x}_0, \lambda_0)$
Factorize P
for $k \leftarrow 1$ to N **do**
 Compute the biases $\mathbf{q}_k^u$ for $u \in V_o$ in topological order,
 given $\mathbf{x}_0, \lambda_0, \dots, \mathbf{x}_{k-1}, \lambda_{k-1}$ and the rules in Section 5
 Compute $\mathbf{q}_k$ by combining affine transformations $(\mathbf{p}^u, \mathbf{q}_k^u)$
 for $u \in V_o$ in topological order
 Solve $\mathbf{x}_k$ and λ_k according to (13) and (17):

$$\begin{cases} P\mathbf{x}_k + \lambda_k \mathbf{v} = -\mathbf{q}_k \\ \mathbf{x}_k^\top \mathbf{x}_1 + \lambda_k \lambda_1 = \mathbb{1}_{k=1} \end{cases}$$

end for

4 THE SANM FRAMEWORK

This section presents the overall design of SANM and its two novel extensions over ANM: the handling of implicit homotopy $H(\mathbf{x}, \lambda) = \mathbf{0}$ and a formulation to reduce accumulated error when solving the equational form $\mathbf{f}(\mathbf{x}) + \mathbf{v} = \mathbf{0}$. Note that this paper deals with both computing graphs and mesh networks. We use *vertex* to refer to a vertex in a graph and *node* to refer to a node in a polygon mesh.

4.1 Coefficient Propagation on the Computing Graph

SANM represents a nonlinear function $\mathbf{f}(\cdot)$ as a directed acyclic bipartite computing graph composed of predefined operators: $G = (V_o \cup V_v, E)$. The user builds the graph to specify the nonlinear function in their problem symbolically. The vertex set V_o corresponds to a subset of predefined operators offered by SANM: for $u \in V_o$, there is a function f_u that defines the corresponding computing. Each operator takes a fixed number of inputs, where each input can be a higher-order tensor. For example, matrix inverse takes one matrix input, and vector addition takes two vector inputs. The vertex set V_v represents the variables. For a variable $v \in V_v$ and an operator $u \in V_o$, an edge $(v, u) \in E$ if the user passes v as an input of the function f_u . An edge $(u, v) \in E$ if v is an output of the function f_u . Many successful symbolic-numerical systems, such as Theano [Theano Development Team 2016] and Tensorflow [Abadi et al. 2016], have adopted computing graphs to represent user-defined functions.

Besides working with the original ANM formulation $\mathbf{f}(\mathbf{x}) + \lambda \mathbf{v} = \mathbf{0}$, SANM also supports solving the more general implicit homotopy: $H(\mathbf{x}, \lambda) = \mathbf{0}$. Implicit homotopy defines a curve $\mathbf{x}(\lambda)$. Our goal is to solve $\mathbf{x}(1)$ given an initial value $\mathbf{x}(0)$.

Algorithm 2 Continuation framework in SANM

Input: An analytic function $H : \mathbb{R}^{n+1} \mapsto \mathbb{R}^n$
Input: Initial value $\mathbf{x}_0 \in \mathbb{R}^n$ and $\lambda_0 \in \mathbb{R}$ such that $H(\mathbf{x}_0, \lambda_0) = \mathbf{0}$
Input: Target location of the homotopy $\lambda_t > \lambda_0$
Output: The solution $\mathbf{x}^* \in \mathbb{R}^n$ such that $H(\mathbf{x}^*, \lambda_t) = \mathbf{0}$

$k \leftarrow 0$
while $\lambda_k < \lambda_t$ **do**
 Compute Taylor coefficients $\{\mathbf{x}_{k,i}\}$ and $\{\lambda_{k,i}\}$
 using Algorithm 1
 Compute the corresponding Padé approximation $P_N(a)$
 Compute a_r and a_p , the range of validity
 for these two approximations
 $a_m \leftarrow \max(a_r, a_p)$
 Use Taylor or Padé to approximate $\mathbf{f}_\mathbf{x}(a)$ and $f_\lambda(a)$,
 depending on which one achieves the range a_m
 Solve a' such that $f_\lambda(a') = \min(f_\lambda(a_m), \lambda_t)$
 $(\mathbf{x}_{k+1}, \lambda_{k+1}) \leftarrow (\mathbf{f}_\mathbf{x}(a'), f_\lambda(a'))$
 $k \leftarrow k + 1$
end while
 $\mathbf{x}^* \leftarrow \mathbf{x}_k$

Proposition 1 implies that the k^{th} order expansion of the implicit homotopy can be written as:

$$H\left(\sum_{i=0}^k \mathbf{x}_i a^i, \sum_{i=0}^k \lambda_i a^i\right) = \sum_{i=0}^{k-1} H_i a^i + \left(\frac{\partial H}{\partial \mathbf{x}}(\mathbf{x}_0, \lambda_0) \cdot \mathbf{x}_k + \frac{\partial H}{\partial \lambda}(\mathbf{x}_0, \lambda_0) \cdot \lambda_k + \mathbf{q}_k\right) a^k + \dots \quad (16)$$

By requiring the coefficient of a^k to be zero, we obtain the equation

$$\frac{\partial H}{\partial \mathbf{x}}(\mathbf{x}_0, \lambda_0) \cdot \mathbf{x}_k + \frac{\partial H}{\partial \lambda}(\mathbf{x}_0, \lambda_0) \cdot \lambda_k + \mathbf{q}_k = \mathbf{0} \quad (17)$$

Given a user-defined function $H(\cdot)$ and initial values $H(\mathbf{x}_0, \lambda_0) = \mathbf{0}$, SANM computes the slope matrix $P = \frac{\partial H}{\partial \mathbf{x}}(\mathbf{x}_0, \lambda_0)$ via reverse mode Automatic Differentiation (AD) [Baydin et al. 2018]. It then iteratively computes the biases for each order and solves the Taylor coefficients according to (13) and (17). The bias $\mathbf{q}_k$ is computed by merging the affine transformations of individual operators in G . Algorithm 1 summarizes this process.

4.2 Continuation in SANM

SANM aims to solve $\mathbf{x}^*$ such that $H(\mathbf{x}^*, \lambda_t) = \mathbf{0}$ given initial values $H(\mathbf{x}_0, \lambda_0) = \mathbf{0}$. We typically have $\lambda_0 = 0$ and $\lambda_t = 1$. Note that we have changed the meaning of $\mathbf{x}_k$ and λ_k to indicate the initial values at the k^{th} iteration rather than the k^{th} Taylor coefficient as in Section 4.1.

Similar to the standard ANM procedure, at the k^{th} iteration, SANM uses Algorithm 1 to compute the local Taylor expansion coefficients of $\mathbf{x}(a)$ and $\lambda(a)$ near $(\mathbf{x}_k, \lambda_k)$ so that $H(\mathbf{x}(a), \lambda(a)) \approx \mathbf{0}$. SANM then estimates the range of validity a_r using the standard ANM formulation in (8). SANM also computes a Padé approximation from the Taylor coefficients and estimates its range of validity

a_p using techniques outlined in Section 3.4. SANM chooses the approximation with a larger range of validity and computes the next approximation to extend the solution path until $\lambda(a)$ reaches λ_t . Algorithm 2 summarizes this process.

4.3 Reducing Error When Solving The Equational Form

When solving a nonlinear system in the original ANM formulation $f(\mathbf{x}) + \mathbf{v} = \mathbf{0}$, SANM adopts a novel continuation method to reduce accumulated numerical and approximation error. The principle is to modify $\mathbf{v}$ at each iteration to incorporate the residual $f(\mathbf{x}_k) + \lambda_k \mathbf{v} - \mathbf{0}$. Specifically, instead of using a fixed definition $\mathbf{H}(\mathbf{x}, \lambda) = f(\mathbf{x}) + \lambda \mathbf{v}$, we change the definition of $\mathbf{H}$ in Algorithm 2 at the beginning of each iteration:

$$\mathbf{H}_k(\mathbf{x}, \lambda) = f(\mathbf{x}) + (\lambda - \lambda_k)\mathbf{v} - (1 - (\lambda - \lambda_k))f(\mathbf{x}_k) \quad (18)$$

Note that the loop invariant $\mathbf{H}_k(\mathbf{x}_k, \lambda_k) = \mathbf{0}$ is still satisfied. We use $\lambda - \lambda_k$ to represent the progress made in the current iteration. The final solution is found if $\lambda - \lambda_k = 1$. Let $\mathbf{r}_k := f(\mathbf{x}_k) + \mathbf{v}$ denote the residual. In the continuation we solve $\mathbf{H}_k(\mathbf{x}_{k+1}, \lambda_{k+1}) \approx \mathbf{0}$, which implies

$$\mathbf{r}_{k+1} \approx (1 - (\lambda_{k+1} - \lambda_k))\mathbf{r}_k \quad (19)$$

Therefore, the continuation decreases the residual. Updating $\mathbf{H}_k$ at each iteration automatically accounts for numerical error and approximation error. We also change the loop condition to $\|\mathbf{r}_k\| \geq \epsilon$ to achieve the desired solution residual ϵ . This formulation allows us to obtain very accurate solutions such as the low residual RMS shown in Table 4 and Table 5.

5 COMPUTING THE AFFINE TRANSFORMATIONS OF TAYLOR COEFFICIENTS

We present algorithms to determine the affine transformations of the highest-order Taylor coefficient for a few nonlinear operators that are commonly used in graphics applications. As we have shown in Section 4.1, these operators can be combined into a computing graph to define the nonlinearity in the target application. Although Proposition 1 ensures the existence of the linearity being sought in the general case, for certain operators, we can compute the affine transformations directly without resorting to computing higher-order derivatives. In this section, we use f to represent the nonlinear operator under investigation and also the value of this operator with x assumed to be the independent variable. For any variable x , we use $\{x_i\}$ to denote its Taylor coefficients with respect to the path parameter a .

Our goal is to derive an affine relationship between f_k and x_k , assuming that $x_0, \dots, x_{k-1}$ and $f_0, \dots, f_{k-1}$ are known constants.

5.1 Basic Arithmetic Operations

The affine transformations for the four basic arithmetic operations are derived by equating the coefficient of a^k in the expansion of the equation:

- $f = x + y$: By introducing $f = \sum_{k=0}^N f_k a^k$, $x = \sum_{k=0}^N x_k a^k$, and $y = \sum_{k=0}^N y_k a^k$ and equating the coefficient of a^k , we have $f_k = x_k + y_k$.
- $f = x - y$: Similarly, we have $f_k = x_k - y_k$.

- $f = xy$: Similarly, we have $f_k = y_0 x_k + x_0 y_k + \sum_{i=1}^{k-1} x_i y_{k-i}$.
- $f = x/y$: We have $x = fy$, which implies $x_k = f_0 y_k + y_0 f_k + \sum_{i=1}^{k-1} f_i y_{k-i}$, and therefore $f_k = (1/y_0)x_k - (f_0/y_0)y_k - 1/y_0 \sum_{i=1}^{k-1} f_i y_{k-i}$.

5.2 Elementwise Analytic Functions

A function $\mathbf{y} = f(\mathbf{x})$ is said to be elementwise if $\forall i : y_i = f(x_i)$. For such functions, we only need to derive the affine transformation for the univariate case.

The problem of computing the Taylor coefficients of $f(\sum_{i=0}^k x_i a^i)$ given the coefficients $\{x_i\}$ and the Taylor expansion of $f(x_0 + a)$ is known in the literature as the composition problem, for which there exists a fast $O((k \log k)^{3/2})$ algorithm [Brent and Kung 1978]. However, it is not necessarily fast when k is small, and the implementation is complicated.

As shown in Griewank and Walther [2008], for most functions of practical interest, we can find auxiliary functions $a(x)$, $b(x)$, and $c(x)$ whose Taylor coefficients $\{a_i\}$, $\{b_i\}$, and $\{c_i\}$ can be easily computed given $\{x_i\}$, such that

$$b(x)f'(x) - a(x)f(x) = c(x) \quad (20)$$

The coefficients $\{f_i\}$ are then computable in $O(k^2)$ time via a formula involving $\{a_i\}$, $\{b_i\}$, $\{c_i\}$, and $\{x_i\}$. Griewank and Walther [2008] provide a thorough treatment on this subject. We list in Table 2 the result formulas to propagate Taylor coefficients through the elementwise analytic functions currently used in SANM for mesh deformation applications:

Table 2. Expansion coefficient propagation for univariate functions [Griewank and Walther 2008]

$f(x)$	Recurrence for f_k
$\ln(x)$	$\frac{1}{x_0} \left(x_k - \sum_{i=1}^{k-1} \frac{i}{k} x_{k-i} f_i \right)$
x^r	$\frac{1}{x_0} \left(r f_0 x_k + \sum_{i=1}^{k-1} \left(\frac{i}{k} (r+1) - 1 \right) f_{k-i} x_i \right)$

Note that when $f(x) = x^r$ and r is an integer, the recurrence is numerically unstable when $|x|$ is small. In this case, we compute the Taylor coefficients via exponentiation by squaring for polynomials in $O(k^2 \log r)$ time.

5.3 Matrix Inverse

Let $F = X^{-1}$ where $X \in \mathbb{R}^{m \times m}$ is a square matrix. Introduce the power series definition and rearrange the terms:

$$\left(\sum_{i=0}^k F_i a^i \right) \left(\sum_{i=0}^k X_i a^i \right) = I \quad (21)$$

The coefficient of a^k on the left hand side is $\sum_{i=0}^k F_i X_{k-i}$, which must be zero because the right hand side is a constant:

$$\begin{aligned} \sum_{i=0}^k F_i X_{k-i} &= 0 \\ \implies F_k &= -X_0^{-1} X_k X_0^{-1} - \left(\sum_{i=1}^{k-1} F_i X_{k-i} \right) X_0^{-1} \end{aligned} \quad (22)$$

Equation (22) explicitly defines an affine relationship between F_k and X_k .

5.4 Matrix Determinant

Let $g(a) = f(X(a)) = \det\left(\sum_{i=0}^k X_i a^i\right)$ where $X \in \mathbb{R}^{m \times m}$. A straightforward method to compute g_k is to expand the determinant according to the Leibniz formula and compute polynomial products, which incurs exponential complexity $O(m!k^2)$ in terms of the matrix size. Although this suffices for FEM applications in 2D or 3D (with $m = 2$ or $m = 3$), we also present a method with polynomial complexity that is better suited for larger matrices.

The terms containing $X_{kij}a^k$ that contribute to g_k can only be multiplied with $X_{0i'j'}$ where $i' \neq i$ and $j' \neq j$. The multiplier of $X_{kij}a^k$ is in fact C_{ij} , where C is the cofactor of X_0 with C_{ij} defined as the determinant of the remaining matrix by removing the i^{th} row and j^{th} column of X_0 . Therefore:

$$g_k = \text{vec}(C)^\top \text{vec}(X_k) + q_k \quad (23)$$

Computing the slope: To efficiently compute the cofactor matrix C , we use the identity of Cramer's rule:

$$X_0^{-1} = \frac{1}{\det(X_0)} C^\top \quad (24)$$

Computing the matrix inverse incurs numerical stability issues. Instead, we first compute the SVD decomposition $X_0 = U\Sigma V^\top$ and then compute C as:

$$\begin{aligned} C &= \det(X_0) X_0^{-\top} \\ &= \det(U) \cdot \det(\Sigma) \cdot \det(V) \cdot U \Sigma^{-1} V^\top \\ &= \det(U) \cdot \det(V) \cdot U D V^\top \end{aligned} \quad (25)$$

Here $D = \det(\Sigma)\Sigma^{-1}$ is a diagonal matrix and $D_{ii} = \prod_{j \neq i} \Sigma_{jj}$. Such a formulation avoids division of singular values and is stable even for ill-conditioned matrices.

Computing the bias: Similar to the argument in Proposition 1, we have $q_k = \mathfrak{C}_a^{(k)} \left[\det\left(\sum_{i=0}^{k-1} X_i a^i\right) \right]$. This is known as the polynomial matrix determinant problem. We propose an efficient solution using discrete Fourier transform, which has also been discovered by Hromčík and Šebekt [1999]:

- (1) Compute $Y_i = \sum_{l=0}^{k-1} X_l \omega_K^{il}$ for $0 \leq i < K$ with Fast Fourier Transform (FFT), where K is the next power of two after k , and $\omega_K = e^{-j \frac{1}{2\pi K}}$ is a K^{th} root of unity. This step costs $O(k \log km^2)$.
- (2) Compute the determinants $d_i = \det(Y_i)$ for $0 \leq i < K$ in $O(km^3)$ time.
- (3) Use the inverse discrete Fourier transform to compute $q_k = \frac{1}{K} \sum_{i=0}^{K-1} d_i \omega_K^{-ik}$ in $O(k)$ time.

The above method computes the bias q_k in $O((\log k + m)km^2)$ time.

5.5 Singular Value Decomposition

Singular Value Decomposition (SVD) generates three matrices from a single matrix input: $U\Sigma V^\top = X$ where U and V are orthonormal matrices and Σ is a diagonal matrix containing the singular values in

decreasing order. Here we only consider the square case $X \in \mathbb{R}^{m \times m}$. We use $\sigma_i = \Sigma_{ii}$ to represent the singular values.

Although the constitutive models considered in this paper do not directly use SVD, the ARAP energy, defined as $\Psi_{\text{ARAP}}(F) := \frac{\mu}{2} \|F - R\|$, involves a Polar Decomposition (PD) $F = PR$. PD can be computed from SVD via $P = U\Sigma U^\top$ and $R = UV^\top$. SVD is also potentially useful for other applications. Therefore, we first present how to compute SVD in the SANM framework and then discuss extra modifications to compute PD more stably. The ARAP energy actually needs a *rotation-variant* PD to prevent reflections in R by requiring that $\det(R) = 1$. We also discuss how to compute such rotation-variant SVD in SANM.

An obstacle in numerical differentiation of SVD is that when there are identical singular values $\sigma_i = \sigma_j$, the corresponding singular vectors $U_i, U_j, V_i,$ and V_j are not uniquely determined. The Jacobians $\frac{\partial U}{\partial X}$ and $\frac{\partial V}{\partial X}$ in this case are thus undefined since different perturbations on X induce noncontinuous changes in U and V . This case corresponds to a division by zero in the Jacobian computation, which is often circumvented by various numerical tricks [Liao et al. 2019; Papadopoulos and Lourakis 2000; Seeger et al. 2017].

In graphics applications, however, identical singular values occur frequently. For example, the singular values of the deformation gradient matrix in isotropic stretching are all identical. As a remedy, we propose to use an alternative form of SVD that includes $UV^\top$ directly, which we denote by SVD- W :

$$\begin{aligned} U\Sigma U^\top W &= X \\ W &= UV^\top \end{aligned} \quad (26)$$

Note that W is also the rotation matrix in the polar decomposition of X , which is unique when X is invertible and the Jacobian $\frac{\partial W}{\partial X}$ is thus well-defined.

Now we present the derivation of affine transformations from X_k to U_k, Σ_k , and W_k . We are not going to give a final equation because it will be too complex and repeat most of the derivation. Instead, we focus on explaining the overall procedure for deriving these affine transformations.

We start by expanding the product $U\Sigma U^\top W$ and extracting the coefficient of a^k , which should be equal to X_k :

$$\begin{aligned} X_k &= \sum_{\substack{a+b+c+d=k \\ \min(a,b,c,d) \geq 1}} U_a \Sigma_b U_c^\top W_d + E \\ E &= U_k \Sigma_0 U_0^\top W_0 + U_0 \Sigma_k U_0^\top W_0 \\ &\quad + U_0 \Sigma_0 U_k^\top W_0 + U_0 \Sigma_0 U_0^\top W_k \end{aligned} \quad (27)$$

We define $F = U_0^\top E W_0^\top U_0$. Introduce (28) to the right hand side:

$$F = U_0^\top U_k \Sigma_0 + \Sigma_k + \Sigma_0 U_k^\top U_0 + \Sigma_0 U_0^\top W_k W_0^\top U_0 \quad (29)$$

Note that (27) establishes an affine relationship between X_k and E , and $F = U_0^\top E W_0^\top U_0$ is also linear. We now only need to seek the affine transformations from F to U_k, Σ_k and W_k .

Expand the orthogonality constraints $U^T U = I$ and $W^T W = I$:

$$U_0^T U_k + U_k^T U_0 + B_u = 0 \quad (30)$$

$$W_0^T W_k + W_k^T W_0 + B_w = 0 \quad (31)$$

$$\text{where } B_u = \sum_{i=1}^{k-1} U_i^T U_{k-i} \text{ and } B_w = \sum_{i=1}^{k-1} W_i^T W_{k-i}$$

Solving Σ_k : From the constraints (30) and (31), we have

$$\text{diag}(U_0^T U_k) = -\frac{1}{2} \text{diag}(B_u) \quad (32)$$

$$\text{diag}(U_0^T W_k W_0^T U_0) = -\frac{1}{2} \text{diag}(U_0^T B_w U_0) \quad (33)$$

Introducing (32) and (33) into (29) allows us to solve $\text{diag}(\Sigma_k)$ from $\text{diag}(F)$.

Solving W_k : Because $U_0^T U_k \Sigma_0 + \Sigma_k + \Sigma_0 U_k^T U_0$ is symmetric, we can cancel this term in (29) by subtracting F^T from F :

$$F - F^T = \Sigma_0 U_0^T W_k W_0^T U_0 - U_0^T W_0 W_k^T U_0 \Sigma_0 \quad (34)$$

From (31) we have $W_k^T = -B_w W_0^T - W_0^T W_k W_0^T$. Introduce it to (34):

$$F - F^T - U_0^T W_0 B_w W_0^T U_0 \Sigma_0 = \Sigma_0 U_0^T W_k W_0^T U_0 + U_0^T W_k W_0^T U_0 \Sigma_0 \quad (35)$$

Note that (35) is a Sylvester equation in the form $\Sigma_0 M + M \Sigma_0 = A$ with $M = U_0^T W_k W_0^T U_0$. The solution is $M_{ij} = \frac{A_{ij}}{\sigma_i + \sigma_j}$ and $W_k = U_0 M U_0^T W_0$.

Solving U_k : Introducing $U_k^T = -B_u U_0^T - U_0^T U_k U_0^T$ (derived from (30)) and the solutions of W_k and Σ_k into (29) results in another Sylvester equation: $\Sigma_0 U_0^T U_k - U_0^T U_k \Sigma_0 = B$, with the solution $(U_0^T U_k)_{ij} = \frac{B_{ij}}{\sigma_i - \sigma_j}$.

5.5.1 Polar Decomposition Case. If other operators in the computing graph only need the W output matrix of the SVD-W operator, this computation can be further simplified by considering the polar decomposition: $X = PW$.

Note that the symmetry of P implies the symmetry of P_i . Expand the identity $P^T P = XX^T$ and extract the coefficient of a^k :

$$P_k P_0 + P_0 P_k + B_p = X_0 X_k^T + X_k X_0^T \quad (36)$$

$$B_p = \sum_{i=1}^{k-1} (P_i P_{k-i} - X_i X_{k-i}^T) \quad (37)$$

Solving P_k : Introduce $P_0 = U_0 \Sigma_0 U_0^T$ to (36) and multiply both sides with U_0^T and U_0 :

$$U_0^T P_k U_0 \Sigma_0 + \Sigma_0 U_0^T P_k U_0 = U_0^T (X_0 X_k^T + X_k X_0^T - B_p) U_0 \quad (38)$$

We have obtained another Sylvester equation and P_k can be solved similarly to solving W_k from (35).

Solving W_k : We can derive W_k by directly expanding $PW = X$:

$$W_k = U_0 \Sigma_0^{-1} U_0^T \left(X_k - \sum_{i=1}^k P_i W_{k-i} \right) \quad (39)$$

5.5.2 Implementation Notes. Although the SVD-W formulation provides a numerically stable expression to compute the Jacobian $\frac{\partial W}{\partial X}$, the biases W_k still depend on the numerically unstable U_i via (27). The polar decomposition formulation does not suffer from this problem because $P = U \Sigma U^T$ is unique in the presence of equal singular values as long as X is invertible. In the implementation, we transparently switch to the polar decomposition formulation to compute W_k in the SVD-W operator when the outputs U and Σ are not needed by other operators in the computing graph. We adopt the Lorentzian broadening [Liao et al. 2019] $x/y \rightarrow xy/(y^2 + \epsilon)$ with $\epsilon = 10^{-12}$ when computing the divisions in solving the Sylvester equations.

The rotation-variant SVD requires $\det(W) = 1$ so that W is a proper rotation matrix. It is traditionally obtained by negating the last singular value and the corresponding left-singular or right-singular vector if $\det(W) = -1$ [Kim and Eberle 2020]. However, when the last singular value is identical to another singular value, the Jacobian $\frac{\partial W}{\partial X}$ becomes undefined because there are multiple singular vectors for this singular value, and it is arbitrary to negate one of them. From a numerical perspective, we need to compute $1/(\sigma_i + \sigma_j)$ in the Jacobian with some $\sigma_i + \sigma_j = 0$. To improve numerical stability, we modify the rotation-variant SVD computation by grouping identical singular values together and preferring to negate all singular values and vectors in a group of an odd size. In the 3D case, there is an odd number of singular values, and therefore an odd-sized group must exist.

6 THE SANM SYSTEM

We use C++ to implement SANM. This section discusses a few design choices that support efficient mesh deformation applications in SANM.

6.1 API Design

SANM adopts a define-and-run paradigm. The user describes a nonlinear system symbolically and provides initial values and input/output transformations. SANM automatically solves the system using the extended ANM framework described in previous sections.

One of the most important public APIs of SANM is for building the symbolic computing graph to represent the nonlinear system of interest. We adopt an object-oriented design to enable intuitive and efficient computing graph building. We use objects in the program to represent variable vertices in the computing graph. We also provide overloading for common arithmetic operators. Listing 1 shows a SANM code excerpt of expressing the first Piola-Kirchhoff tensors for a few constitutive models. SANM only requires the user to provide the C++ object that represents the whole nonlinear system, the sparse affine transformations on the inputs and outputs (see Section 6.2), and the initial values. The user does not need to modify the solver to work on different tasks.

SANM significantly reduces programming effort for applying ANM. We roughly measure programming effort by the number of lines of C++ code. With SANM, the whole FEM solver for all elastic deformation applications in this paper, including auxiliary functionalities such as mesh input/output and tetrahedron processing, needs only 1.5K lines of code without using external geometry manipulation libraries. The SANM library itself has about 7.5K lines of

Listing 1 A code excerpt for building the first Piola–Kirchhoff stress tensors with three constitutive models using the SANM framework. The symbolic nature of SANM allows easy substitution of different constitutive models. This function is essentially a literal translation of the formulas (43)–(45).

```
SymbolVar pk1(EnergyModel energy_model, const MaterialProperty& material, SymbolVar F) { switch (energy_model) {
  case EnergyModel::NEOHOOKEAN_I: {
    fp_t k = material.bulk_modulus(), mu = material.shear_modulus();
    SymbolVar FTinv = batched_mat_inv_mul(F, {}, true).batched_transpose(), J = F.batched_det(), Ic = F.pow(2).reduce_sum(-1),
    J23 = J.pow(-2. / 3.), t2 = linear_combine({{mu / (-3._fp), J23 * Ic}, {k, J * J}}, {-k, J}), 0) * FTinv,
    P = linear_combine({{mu, J23 * F}, {1._fp, t2}});
    return P;
  }
  case EnergyModel::NEOHOOKEAN_C: {
    fp_t mu = material.shear_modulus(), lambda = material.lame_first();
    SymbolVar FTinv = batched_mat_inv_mul(F, {}, true).batched_transpose(), J = F.batched_det(),
    P = linear_combine({{mu, F}, {-mu, FTinv}, {lambda, J.log() * FTinv}});
    return P;
  }
  case EnergyModel::ARAP: {
    fp_t mu = material.shear_modulus();
    return (F - F.batched_svd_w(true)[2]) * mu;
  }
}
```

code. By contrast, the official ANM implementation of Chen et al. [2014] has 11,499 lines of code for only the ANM numerical solving part (neoHookeanANM.cpp and neoHookeanANMForward.cpp). Their code only supports the incompressible neo-Hookean model, while SANM allows working with multiple constitutive models by changing a few lines of code.

Internally, SANM provides a mechanism to register new operators. Each operator only needs to implement a few functions, such as gradient computing and Taylor coefficient propagation. The SANM framework manipulates the computing graph and orchestrates individual operator functionalities to implement ANM solving. SANM currently has operators to support the constitutive models used in this paper, including arithmetic operators (addition, subtraction, multiplication, division, logarithm, and exponentiation), tensor operators (slicing and concatenation), and batched linear algebra operators (matrix multiplication, inverse, determinant, transpose, and SVD). It is easy to extend SANM to support more operators via the operator registration mechanism.

6.2 Batch Computing

In finite element analysis, we typically carry out computation of some identical form for a set of elements. In the context of 3D mesh deformation, we compute the stress tensors of each tetrahedron from its deformation gradient. This form of computation allows us to exploit modern computer architectures better via batch processing. Specifically, we pack the matrices of all elements into a large tensor on which the operator in the computing graph executes. For example, the deformation gradient of the i^{th} element is computed as $F_i = \text{matmul}(D_{si}, D_{mi}^{-1})$. If we execute the coefficient propagation algorithm discussed in Section 4.1 for each tetrahedral element independently, we will need to invoke the `matmul` operators n times, while each invocation only works on 3×3 small matrices. Instead, we pack all the deformation gradients into a third-order $n \times 3 \times 3$ tensor $\mathbb{F} = [F_1; \dots; F_n]$, and pack $\mathbb{D}_m$ and $\mathbb{D}_s$ similarly. The computing graph then contains a single `batched_matmul` operator that computes matrix multiplication for all elements together:

$\mathbb{F} = \text{batched_matmul}(\mathbb{D}_s, \mathbb{D}_m^{-1})$. In our mesh deformation applications, the unknown vector $\mathbf{x}$ represents node coordinates. We use a sparse affine transformation $\mathbf{A}$ to map the coordinates to shape matrices and another transformation $\mathbf{B}$ to map from stress tensors to nodal force: $f(\mathbf{x}) = \mathbf{B}(\mathbf{P}(\mathbf{F}(\mathbf{A}(\mathbf{x}))))$ where $\mathbf{P}$ represents the Piola–Kirchhoff stress tensors, $\mathbf{F}$ represents the deformation gradients, and $\mathbf{P}(\mathbf{F}(\cdot))$ is computed in a batched manner. The sparse affine transformations $\mathbf{A}$ and $\mathbf{B}$ are provided as input/output transformations to the SANM solver.

Batch computing improves performance by better utilizing the hardware, although it does not reduce computational complexity. Most modern CPUs support Single Instruction Multiple Data (SIMD) parallelism, and GPUs are designed to process large amounts of data in parallel. Without batch computing, such hardware capability can hardly be utilized by the small matrices occurring in finite element analysis. Batch computing also amortizes the overhead of computing graph traversing in Algorithm 1.

SANM supports parallel computing by splitting data on the batch dimension, which is managed by the framework and is oblivious to individual operator implementations. Figure 5 compares solving times achieved with different numbers of threads, which shows that SANM exhibits modest scalability on practical workloads.

6.3 Performance Optimizations

We design SANM with high-performance computing in mind. Here we discuss other optimizations besides batch computing.

Efficient numerical computing primitives: SANM provides an abstraction of numerical computing primitives. The implementations of computing graph operators invoke these primitives instead of directly working on numerical data. This paradigm allows us to separate numerical algorithm description from performance engineering. Currently, we use Eigen [Guennebaud et al. 2010] and Intel Math Kernel Library to implement the computing primitives on CPU with Single Instruction Multiple Data (SIMD) optimizations. We can easily extend SANM to support GPU by adding another GPU backend for the primitives without modifying implementations of

Fig. 5. Comparing the end-to-end solving times with different numbers of threads on the static forward gravity equilibrium problem of the Armadillo model with 221,414 nodes and 696,975 tetrahedrons. SANM exhibits modest scalability using up to 32 threads. The ideal parallelism is the single-thread solving time divided by the number of threads.

operators or the ANM solver. This abstraction also allows SANM to benefit from other research on optimizing tensor computing performance, such as recent related research in deep learning [Chen et al. 2018; Jouppi et al. 2017].

Automatic memory management with copy-on-write: SANM automatically manages tensor memory by reference counting, with eager memory sharing and copy-on-write to simplify programming without sacrificing performance. When a tensor object is copied, only a new reference is stored in the destination. When a tensor object is modified, SANM makes a private copy before modification if the reference count is greater than one.

Sparse affine transformations: We exploit the structural sparsity when computing the Jacobians and the affine transformations in Algorithm 1. For a batch-packed tensor with dimensions $n \times m \times m$, we use an $n \times m^2 \times m^2$ tensor to represent its Jacobian instead of using a full $nm^2 \times nm^2$ matrix because each matrix in the batch is independent of each other. Our sparse Jacobian representation significantly reduces memory usage when n is large. It is also friendly to batch computing. Furthermore, we use a compact $n \times m^2$ matrix to represent Jacobians for elementwise operators.

Special handling of zero tensors: Zero tensors frequently occur, such as being used as the initial accumulation value. SANM retains a special buffer shared by all zero-initialized tensors (also with reference counting and copy-on-write). Thus checking whether a tensor is all zero can be easily done by comparing the buffer address. This design allows implementing a fast path for handling zero inputs in elementary arithmetic operators, such as $x + 0 = x$ and $x \cdot 0 = 0$. This optimization leads to a 7.91% speedup in our experiment.

6.4 Future Performance Optimizations

We discuss other optimizations that are not yet implemented but likely to be helpful. Thanks to the define-and-run paradigm, SANM users can benefit from future optimizations by simply updating their SANM library without modifying their application code.

Computing graph optimization: Since the user symbolically defines the computing graph, SANM can optimize the graph before starting numerical computation. For example, we can simplify arithmetic expressions. We can also fuse arithmetic operators with just-in-time compilation. There is a large body of research on traditional compiler optimization [Lattner and Adve 2004] and recent tensor compiler optimization in deep learning [Lattner et al. 2020] that may benefit future SANM optimizations.

Reducing memory usage: Currently, we store all the intermediate Taylor coefficients in memory, which incurs some memory overhead. A possible improvement is setting up checkpoints on the computing graph and recomputing the Taylor coefficients between checkpoints each time. A good choice of checkpoints might induce little computational cost ($O(1)$ times the original cost) while saving lots of memory ($O(1/\sqrt{N})$ relative memory usage) for a computing graph with a chain of length N [Chen et al. 2016].

7 MESH DEFORMATION APPLICATIONS

We evaluate SANM on a few volumetric mesh deformation problems. We first briefly review the basics of elastic deformation analysis. Readers may refer to Bonet and Wood [2008]; Kim and Eberle [2020]; Sifakis and Barbic [2012] for a more thorough treatment on this subject.

We consider 3D deformation of hyperelastic materials, for which the work done by the stresses during a deformation process only depends on the initial and final state. A constitutive model relates the elastic potential energy density Ψ to the deformation gradient F . Under a piece-wise linear tetrahedral discretization, the deformation gradient is constant within a tetrahedron: $F = D_s D_m^{-1} \in \mathbb{R}^{3 \times 3}$ where D_s is the *deformed shape matrix* (computed from the deformed tetrahedron) and D_m is the *reference shape matrix* (computed from the rest tetrahedron). A shape matrix of a tetrahedron packs the three column vectors from one vertex to the other three. The elastic force $f_{ij} \in \mathbb{R}^3$ exerted by a single tetrahedron i to its j^{th} node is then derived by taking the gradient of the potential energy with respect to the node coordinates: $f_{ij} = -\frac{\partial \Psi(F_i)}{\partial x_{ij}}$. We can obtain the formulation for the internal elastic force at a node i by combining forces exerted by neighboring tetrahedrons, which equals the following:

$$f_i = \sum_{t \in N_i} P(F_t) \bar{n}_{t,i} \quad (40)$$

where N_i is the set of adjacent tetrahedrons containing node x_i , $P(F) := \frac{\partial \Psi(F)}{\partial F}$ is the first Piola–Kirchhoff stress tensor of the constitutive model, and $\bar{n}_{t,i}$ is the outward area-weighted normal vector at node x_i of the tetrahedron t in the undeformed state.

To solve the inverse deformation problem that seeks a rest shape which deforms to a given shape, we introduce the Cauchy stress tensor σ that linearly relates the elastic force to the deformed state:

$$f_i = \sum_{t \in N_i} \sigma(F_t) n_{t,i} \quad (41)$$

$$\sigma(F) = \frac{1}{\det(F)} P(F) F^T \quad (42)$$

where $n_{t,i}$ is the outward area-weighted normal vector at node x_i of the tetrahedron t in the given deformed state.

Table 3. Comparing with the hand-coded, specialized ANM solver of Chen et al. [2014] on gravity equilibrium problems. In this table, inv. and fwd. mean inverse and forward problems respectively, and mt4 or mt6 indicate using 4 or 6 threads for parallel computing. Since the code of Chen et al. [2014] no longer compiles on modern systems, we directly use the data reported in their paper. We use the models included in their open-source release and exclude models with additional external force because their force description file is in a custom binary format. We ran SANM on a server with two Intel Xeon Platinum 8269CY CPUs (2.5GHz - 3.8GHz), while Chen et al. [2014] used a desktop PC with an Intel i7-3770K CPU (3.5GHz - 3.9GHz). We use the same truncation order $N = 20$ as Chen et al. [2014]. The data show that SANM delivers comparable, if not better, performance as a hand-coded and manually optimized ANM solver.

Model	Ours: SANM			Chen et al. [2014]	
	#Iter.	Time	Time (mt4)	#Iter.	Time (mt6)
inv. bar	2	2.95	1.03	2	2.38
inv. plant	2	7.36	2.70	3	7.07
fwd. bar	2	2.90	0.98	3	3.25
fwd. plant	2	7.15	2.69	4	9.27

This paper considers three constitutive models: the compressible neo-Hookean energy (denoted by NC), the incompressible neo-Hookean energy (denoted by NI), and the As-Rigid-As-Possible energy (denoted by ARAP). We list their first Piola–Kirchhoff stress tensors:

$$P_{NC}(F) = \mu(F - F^{-\top}) + \lambda \log(J)F^{-\top} \quad (43)$$

$$P_{NI}(F) = \mu J^{-\frac{2}{3}} \left(F - \frac{1}{3} \|F\| F^{-\top} \right) + \kappa J(J - 1)F^{-\top} \quad (44)$$

$$P_{ARAP}(F) = \mu(F - R) \quad (45)$$

$$J := \det(F)$$

Note that ARAP (45) needs a rotation-variant polar decomposition $F = RS$ such that R is a proper rotation matrix with $\det(R) = 1$. The parameters, λ (Lamé's first parameter), μ (Lamé's second parameter), and κ (the bulk modulus) are all determined by the physical properties of the material.

7.1 Forward and Inverse Static Equilibrium Problems

Given a static external force, we consider the problems of solving the rest shape given the deformed shape (the inverse problem) and solving the deformed shape given the rest shape (the forward problem) similar to Chen et al. [2014]. Formally, let $\bar{x}$ denote the rest shape, x the deformed shape, $f(\bar{x}, x)$ the internal elastic force, and f_{ext} the external force. We solve $f([\bar{x}; x_b], [x; x_b]) + f_{ext} = 0$ given either $\bar{x}$ or x , where x_b contains fixed boundary nodes. This problem naturally fits into the numerical continuation framework by replacing f_{ext} with λf_{ext} . We consider static equilibrium under gravity and set f_{ext} as the per-node gravity.

Chen et al. [2014] has shown that ANM is tens to thousands of times faster than the Levenberg-Marquardt algorithm on the inverse problem, and it is also roughly ten times faster than an implicit Newmark simulation with kinetic damping [Umetani et al. 2011] on the forward problem. Table 3 compares SANM with the hand-coded

and manually optimized ANM solver of Chen et al. [2014], which shows that SANM achieves comparable or better performance while automatically solves the problem. We also evaluate SANM on a large Armadillo model with 221,414 nodes and 696,975 tetrahedrons, which is nearly ten times larger than the models used in Chen et al. [2014]. Figure 5 presents the solving time of the forward problem using different numbers of threads. Figure 2 shows the intermediate states for a few values of λ in the continuation.

7.1.1 Comparison with Newton's methods. We compare with more methods on the forward problem. The minimum total potential energy principle dictates that the equilibrium state is the minimizer of the total potential energy, including the elastic potential energy and the gravitational potential energy. Therefore, an alternative method for the forward problem is to solve $\arg \min_x (\Psi(x) - g^\top x)$ where g is the per-node gravity. We implement Newton's method with backtracking line search to solve the minimization. We also evaluate positive-semidefinite Hessian projection with a state-of-the-art derivation of per-tetrahedron analytic eigensystems for the elastic energy functions [Smith et al. 2019]. Note that the energy minimization method does not apply to the inverse problem due to the lack of corresponding global energy. We also compare with directly minimizing $\|f(x) + g\|$ by the Levenberg-Marquardt algorithm.

We run the experiments on a desktop PC with an AMD Ryzen Threadripper 2970WX CPU. All the implementations are compiled with the same compiler, use the same linear algebra libraries, and use a single thread. Newton's method uses an LU solver, and the projective Newton uses a faster LLT solver due to the guaranteed positive definiteness of projected Hessians. We use techniques described in Section 4.3 to reduce the error of SANM. We set the truncation order $N = 20$. We set the convergence threshold to be $\epsilon = 10^{-10}$ for the RMS of the force residual. We find that energy minimization with Newtonian methods often fails to converge to such a small force residual due to vanishing step sizes near the optimum, and therefore we stop them if either the RMS of the force residual or the change of x in one iteration drops below 10^{-6} . We then use additional Gauss-Newton iterations to fine-tune the solution. Table 4 presents the comparisons, which shows that SANM converges faster than the considered alternative methods in most cases and achieves low residual without additional refinement.

7.2 Controlled Mesh Deformation

We demonstrate implicit homotopy solving on controlled mesh deformation problems. The problem asks for the equilibrium state when specific nodes are moved to given locations. The constrained nodes, called the control handles, are usually specified by a user so that the elastic body can be deformed into the desired pose. This problem is typically solved with an energy minimization framework that minimizes the total elastic potential energy.

We propose an alternative approach with implicit homotopy. Given the rest shape of the body as x_0 , the initial position of control handles as x_c , and the target position of control handles as x_t , we

Table 4. Performance comparison on forward gravity equilibrium problems. We limit the Levenberg-Marquardt algorithm to use no more than 1000 iterations. The SANM speedup is computed by comparing with the fastest correct (i.e., producing no inverted tetrahedrons) alternative method for each problem. The NC and NI materials refer to compressible and incompressible neo-Hookean materials respectively. $\text{RMS}(f)$ refers to the root-mean-square value of force residuals on unconstrained nodes. The bracketed numbers indicate the Gauss-Newton refinement iterations (limited to be 20). Proj. Newton uses per-tetrahedron Hessian projection derived from Smith et al. [2019]. SANM achieves the low residual without extra refinement thanks to the techniques presented in Section 4.3. Bold times mark the fastest methods and italic times mark the second fastest methods. SANM achieves an average speedup of 2.21 by geometric mean.

Mesh	Material	Ours: SANM			Newton			Proj. Newton			Levenberg-Marquardt			SANM Speedup
		#Iter	Time	$\text{RMS}(f)$	#Iter	Time	$\text{RMS}(f)$	#Iter	Time	$\text{RMS}(f)$	#Iter	Time	$\text{RMS}(f)$	
armadillo-s <small>V=13665 F=42288</small>	ARAP	4	10.07	2.7e-13	27(1)	24.92	1.5e-14	73(2)	52.35	1.7e-14	1000	3785.59	4.9e-6	2.47
	NC	2	6.55	5.0e-12	21(1)	17.60	4.3e-14	36(2)	25.94	5.4e-14	1000	3706.22	3.2e-6	2.69
	NI	2	7.16	6.0e-12	21(1)	17.64	4.4e-14	36(2)	26.63	6.8e-14	1000	3415.21	3.2e-6	2.46
bar <small>V=4552 F=19552</small>	ARAP	3	3.74	1.4e-15	24(1)	9.78	1.4e-15	28(2)	9.43	1.6e-15	1000	1439.64	2.2e-6	2.52
	NC	2	2.91	6.8e-15	32(1)	11.41	4.7e-13	44(2)	12.05	3.0e-12	1000	1361.32	3.5e-6	3.92
	NI	2	2.72	6.7e-15	33(1)	11.39	2.6e-14	45(2)	12.11	3.9e-12	1000	1335.08	3.5e-6	4.20
bifur3 <small>V=5093 F=24478</small>	ARAP	2	3.28	1.4e-12	6(1)	3.33	1.4e-12	12(1)	5.12	2.0e-11	182	304.45	1.4e-12	1.02
	NC	1	1.89	1.0e-11	5(1)	2.82	1.0e-11	7(1)	3.14	3.4e-11	336	563.69	1.1e-11	1.49
	NI	1	2.11	9.9e-12	5(1)	2.85	1.0e-11	7(1)	3.10	3.3e-11	336	571.27	1.1e-11	1.35
bob <small>V=7699 F=27577</small>	ARAP	2	4.10	6.8e-15	10(1)	7.35	6.9e-15	17(1)	9.22	5.0e-12	207	563.91	3.4e-11	1.79
	NC	1	2.20	1.0e-11	11(1)	7.22	3.8e-14	14(1)	7.05	1.8e-11	275	672.32	1.1e-12	3.21
	NI	1	2.53	1.1e-13	11(1)	6.92	3.9e-14	14(1)	7.04	2.0e-11	277	705.81	9.9e-13	2.74
human <small>V=25659 F=78067</small>	ARAP	8	41.47	4.1e-12	9(6)	17.55	7.7e-13	91(2)	129.33	1.3e-14	1000	6695.33	7.0e-6*	0.42
	NC	2	13.73	1.1e-13	11(1)	18.10	1.0e-13	14(2)	20.25	6.8e-12	1000	6448.59	9.2e-6	1.32
	NI	2	14.61	1.1e-13	11(1)	19.10	1.0e-13	14(2)	20.14	6.8e-12	1000	6284.69	8.8e-6	1.31
plant <small>V=14842 F=47077</small>	ARAP	6	15.25	4.8e-15	60(2)	57.14	2.0e-12	88(3)	62.59	1.0e-10	1000	2846.89	3.5e-6	3.75
	NC	2	6.66	9.8e-14	49(1)	44.26	3.3e-14	53(3)	39.23	3.4e-11	1000	2586.42	4.3e-6	5.89
	NI	2	7.91	8.5e-14	49(1)	41.99	4.3e-14	52(3)	34.81	7.5e-11	1000	2853.29	4.3e-6	4.40

* The solution contains inverted tetrahedrons.

Table 5. Performance comparison on controlled mesh deformation problems. This table uses similar notations as Table 4. Note that Newton's minimization methods do not work with neo-Hookean energies because the initial guess contains inverted tetrahedrons. SANM uses equation solving presented in Section 4.3 to refine the solution, while Newton's methods use Gauss-Newton iterations for refinement. SANM achieves an average speedup of 1.41 by geometric mean.

Mesh	Ours: SANM (ARAP)			Newton (ARAP)			Proj. Newton (ARAP)			SANM Speedup	SANM (NC)			SANM (NI)		
	#Iter	Time	$\text{RMS}(f)$	#Iter	Time	$\text{RMS}(f)$	#Iter	Time	$\text{RMS}(f)$		#Iter	Time	$\text{RMS}(f)$	#Iter	Time	$\text{RMS}(f)$
armadillo-s	3(1)	6.18	7.0e-15	13(1)	12.12	1.3e-12	10(1)	7.73	2.8e-13	1.25	3(1)	6.98	1.5e-13	4(1)	10.72	6.6e-14
bar2	8(1)	5.89	4.5e-12	23(20)	7.02	1.3e+06*	57(2)	12.13	4.5e-12	2.06	7(1)	5.34	3.6e-11	7(1)	6.24	3.4e-11
bifur3	6(1)	8.56	1.3e-12	43(20)	21.76	5.3e+03*	27(2)	10.39	1.3e-12	1.21	6(1)	9.15	1.1e-11	6(1)	10.45	1.1e-11
bob	3(1)	4.71	1.4e-11	17(11)	10.16	5.6e-15	13(2)	6.97	5.4e-15	1.48	4(1)	6.78	1.1e-12	5(1)	9.54	1.4e-12
human	3(1)	11.42	1.9e-15	9(1)	14.45	3.2e-12	8(1)	10.87	6.1e-11	0.95	4(1)	22.17	1.3e-14	4(1)	22.79	1.4e-14
plant	6(2)	13.59	5.7e-15	28(1)	24.39	9.1e-13	83(3)	60.19	6.8e-13	1.79	5(1)	15.04	1.0e-12	7(2)	21.66	3.7e-14

* The solution contains inverted tetrahedrons.

define an implicit homotopy for the unconstrained nodes $\mathbf{x}$:

$$\mathbf{H}(\mathbf{x}, \lambda) = \mathbf{f}\left(\begin{bmatrix} \mathbf{x}_0 \\ \mathbf{x}_c \end{bmatrix}, \begin{bmatrix} \mathbf{x} \\ \mathbf{x}_c + \lambda(\mathbf{x}_t - \mathbf{x}_c) \end{bmatrix}\right) = \mathbf{0} \quad (46)$$

with $\mathbf{H}(\mathbf{x}_0, 0) = \mathbf{0}$

where $\mathbf{f}(\tilde{\mathbf{x}}, \mathbf{x})$ computes the internal elastic force given the rest shape $\tilde{\mathbf{x}}$ and the deformed shape $\mathbf{x}$. Numerical continuation of λ from 0 to 1 solves the coordinates of unconstrained nodes, and each intermediate configuration is a valid equilibrium state with the control handles on a linear path from $\mathbf{x}_c$ to $\mathbf{x}_t$. Computing an intermediate state at λ_i only requires solving an equation $\lambda(a) = \lambda_i$

and then evaluate $\mathbf{x}(a)$, which incurs negligible extra cost. After solving the homotopy, we refine the solution by solving a static equilibrium problem with zero external force using the improved equation solving presented in Section 4.3 (with truncation order $N = 6$) to reduce the force residual.

We compare SANM with Newton's methods on a few 3D models with manually specified target positions of control handles, such as the deformed Bob shown in Figure 1c. Experimental settings are similar to those in Section 7.1.1. Newton's methods directly minimize $\Psi([\mathbf{x}; \mathbf{x}_t])$ starting at $\Psi([\mathbf{x}_0; \mathbf{x}_t])$. We do not evaluate the Levenberg-Marquardt algorithm since it is too inefficient compared

to others. Note that the initial guess $\mathbf{x} = \mathbf{x}_0$ for Newton's methods induces inverted tetrahedrons. Unfortunately, the neo-Hookean energies do not handle this case (corresponding to $J < 0$ in (43) and (44)), and therefore we do not evaluate Newton's methods on them. Furthermore, a straightforward Newton's method without Hessian projection sometimes fails on the ARAP energy because the inverted tetrahedrons cause indefinite Hessian with more negative eigenvalues that obstruct the optimization progress. By contrast, numerical continuation in SANM ensures a smooth process where no tetrahedron gets inverted, and therefore it also works for neo-Hookean energies. Table 5 presents the results, which shows that SANM is more robust and more efficient in most cases.

We demonstrate the robustness of SANM on a problem with larger deformation. We twist a horizontal bar by 360 degrees and then bend it as shown in Figure 6. Note that the target boundary configuration is indistinguishable from another one which has only the bending but no twisting. SANM naturally handles this case by using a piecewise-linear description of the movement path. However, to apply energy minimization methods, a similar but arguably less principled continuation scheme (such as minimizing the energy in multiple stages) or more complicated initialization strategies are needed to resolve the rotation ambiguity.

Fig. 6. Twisting a horizontal bar by 360 degrees and then bending it, with three constitutive models from left to right: compressible neo-Hookean, incompressible neo-Hookean, and ARAP.

7.3 Discussion

Our experiments show that SANM delivers shorter solving times than Newton's methods. On the 24 comparison experiments (including both forward equilibrium problems and controlled deformation problems), SANM achieves an average speedup of 1.97 by geometric mean compared to the fastest alternative method for each case. On average, SANM spends $23.43\% \pm 1.56\%$ of its running time in the sparse linear solver, while most of the other time is used by Taylor coefficient computation that can be further improved.

SANM and Newton's methods target different problems, and they can not completely replace each other. SANM solves nonlinear systems via numerical continuation, while Newton's methods typically solve minimization problems. Moreover, numerical continuation allows easily computing intermediate equilibrium states almost for free. By comparison, the intermediate states of a Newtonian solver are less interpretable, but such solvers might reduce the energy in early iterations and thus quickly produce visually plausible results.

Energy minimization does not apply to all of our experiment problems. We are unaware of any global energy suitable for the inverse static equilibrium problem. For controlled deformation problems,

finding a proper initial guess for energy minimization becomes non-trivial for certain energies that can not handle inverted elements. We also need to take special care when applying energy minimization to target configurations that involve ambiguity, such as rotations. By contrast, SANM directly and efficiently handles these cases with numerical continuation.

8 CONCLUSION

The asymptotic numerical method is a powerful numerical continuation method for solving nonlinear systems. Prior to our work, a major obstacle of applying ANM was the difficulty in deriving the Taylor coefficients. We have shown that this process can be fully automated and generalized to handle a large family of nonlinearities. We also extend the ANM formulation to handle implicit homotopy. Moreover, we implement an efficient and automatic ANM solver, SANM, that delivers comparable or better performance than a hand-coded, manually optimized, and specialized ANM solver. Although energy minimization targets different problems from SANM in general, we also compare SANM with energy minimization via Newton's methods on a few problems and show that SANM performs favorably.

With our tool, one can explore ANM on many applications in various fields with little effort. SANM can contribute to improvements in numerical solving in many systems. It may also inspire deeper theoretical understanding and further improvement of ANM.

ACKNOWLEDGMENTS

We thank all the anonymous reviewers for providing the detailed review feedback, which has greatly improved the quality of this work. We would also like to thank Changxi Zheng, Wojciech Matusik, Liang Shi, and Martin Rinard for the helpful technical discussions and Lingxiao Li for proofreading. This work was funded by the project "Automatically Learning the Behavior of Computational Agents" (MIT CO 6940111, sponsored by Boeing with sponsor ID #Z0918-5060).

REFERENCES

- Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*. 265–283.
- H Abichou, H Zahrouni, and M Potier-Ferry. 2002. Asymptotic numerical method for problems coupling several nonlinearities. *Computer Methods in Applied Mechanics and Engineering* 191, 51–52 (2002), 5795–5810.
- Eugene L. Allgower and Kurt Georg. 2003. *Introduction to numerical continuation methods*. SIAM.
- L. Azrar, EH Boutyour, and M Potier-Ferry. 2002. Non-linear forced vibrations of plates by an asymptotic-numerical method. *Journal of Sound and Vibration* 252, 4 (2002), 657–674.
- L. Azrar, B Cochelin, N Damil, and M Potier-Ferry. 1993. An asymptotic-numerical method to compute the postbuckling behaviour of elastic plates and shells. *International journal for numerical methods in engineering* 36, 8 (1993), 1251–1277.
- JL Basdevant. 1972. The Padé approximation and its physical applications. *Fortschritte der Physik* 20, 5 (1972), 283–331.
- Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. 2018. Automatic Differentiation in Machine Learning: a Survey. *Journal of Machine Learning Research* 18, 153 (2018), 1–43. <http://jmlr.org/papers/v18/17-468.html>
- Jan Bender, Matthias Müller, Miguel A Otaduy, Matthias Teschner, and Miles Macklin. 2014. A survey on position-based simulation methods in computer graphics. In *Computer graphics forum*, Vol. 33. Wiley Online Library, 228–251.

- Javier Bonet and Richard D. Wood. 2008. *Nonlinear Continuum Mechanics for Finite Element Analysis* (2 ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511755446>
- Joseph-Frédéric Bonnans, Jean Charles Gilbert, Claude Lemaréchal, and Claudia A Sagastizábal. 2006. *Nonlinear optimization: theoretical and practical aspects*. Springer Science & Business Media.
- 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 (2014), 1–11.
- EH Boutyou, H Zahrouni, M Potier-Ferry, and M Boudi. 2004. Bifurcation points and bifurcated branches by an asymptotic numerical method and Padé approximants. *Internat. J. Numer. Methods Engrg.* 60, 12 (2004), 1987–2012.
- Richard P Brent. 2013. *Algorithms for minimization without derivatives*. Courier Corporation.
- Richard P Brent and Hsiang T Kung. 1978. Fast algorithms for manipulating formal power series. *Journal of the ACM (JACM)* 25, 4 (1978), 581–595.
- Isaac Chao, Ulrich Pinkall, Patrick Sanan, and Peter Schröder. 2010. A simple geometric model for elastic deformations. *ACM transactions on graphics (TOG)* 29, 4 (2010), 1–6.
- Isabelle Charpentier, Arnaud Lejeune, and Michel Potier-Ferry. 2008. The diamant approach for an efficient automatic differentiation of the asymptotic numerical method. In *Advances in automatic differentiation*. Springer, 139–149.
- Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. 2018. TVM: An automated end-to-end optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation OSDI 18*, 578–594.
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174* (2016).
- Xiang Chen, Changxi Zheng, Weiwei Xu, and Kun Zhou. 2014. An asymptotic numerical method for inverse elastic shape design. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Min Gyu Choi and Hyeon-Seok Ko. 2005. Modal warping: Real-time simulation of large rotational deformation and manipulation. *IEEE Transactions on Visualization and Computer Graphics* 11, 1 (2005), 91–101.
- Bruno Cochelin. 1994. A path-following technique via an asymptotic-numerical method. *Computers & structures* 53, 5 (1994), 1181–1192.
- Bruno Cochelin, Noureddine Damil, and Michel Potier-Ferry. 1994a. Asymptotic-numerical methods and Padé approximants for non-linear elastic structures. *International journal for numerical methods in engineering* 37, 7 (1994), 1187–1213.
- Bruno Cochelin, Noureddine Damil, and Michel Potier-Ferry. 1994b. The asymptotic-numerical method: an efficient perturbation technique for nonlinear structural mechanics. *Revue européenne des éléments finis* 3, 2 (1994), 281–297.
- Noureddine Damil and Michel Potier-Ferry. 1990. A new method to compute perturbed bifurcations: application to the buckling of imperfect elastic structures. *International Journal of Engineering Science* 28, 9 (1990), 943–957.
- EM Daya and M Potier-Ferry. 2001. A numerical method for nonlinear eigenvalue problems application to vibrations of viscoelastic structures. *Computers & Structures* 79, 5 (2001), 533–541.
- Simon Duenster, Roi Poranne, Bernhard Thomaszewski, and Stelian Coros. 2020. RoboCut: hot-wire cutting with robot-controlled flexible rods. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 98–1.
- Ahmad Elhage-Hussein, Michel Potier-Ferry, and Noureddine Damil. 2000. A numerical continuation method based on Padé approximants. *International Journal of Solids and Structures* 37, 46–47 (2000), 6981–7001.
- Benjamin Gilles, Guillaume Bousquet, Francois Faure, and Dinesh K Pai. 2011. Frame-based elastic models. *ACM transactions on graphics (TOG)* 30, 2 (2011), 1–12.
- Gene H Golub and TN Robertson. 1967. A generalized Bairstow algorithm. *Commun. ACM* 10, 6 (1967), 371–373.
- Andreas Griewank and Andrea Walther. 2008. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.
- Louis Guillot, Bruno Cochelin, and Christophe Vergez. 2019. A generic and efficient Taylor series-based continuation method using a quadratic recast of smooth nonlinear systems. *International Journal for numerical methods in Engineering* 119, 4 (2019), 261–280.
- M Hromčík and M Šebekt. 1999. New algorithm for polynomial matrix determinant based on FFT. In *1999 European Control Conference (ECC)*. IEEE, 4173–4177.
- Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. 2017. In-Datcenter Performance Analysis of a Tensor Processing Unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture* (Toronto, ON, Canada) (ISCA '17). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3079856.3080246>
- Theodore Kim and David Eberle. 2020. Dynamic deformables: implementation and production practicalities. In *ACM SIGGRAPH 2020 Courses*. 1–182.
- Theodore Kim and Doug L James. 2009. Skipping steps in deformable simulation with online model reduction. In *ACM SIGGRAPH Asia 2009 papers*. 1–9.
- Chris Lattner and Vikram Adve. 2004. LLVM: A compilation framework for lifelong program analysis & transformation. In *International Symposium on Code Generation and Optimization*, 2004. CGO 2004. IEEE, 75–86.
- Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. 2020. MLIR: A compiler infrastructure for the end of Moore's law. *arXiv preprint arXiv:2002.11054* (2020).
- Arnaud Lazarus, JT Miller, and Pedro M Reis. 2013. Continuation of equilibria and stability of slender elastic rods using an asymptotic numerical method. *Journal of the Mechanics and Physics of Solids* 61, 8 (2013), 1712–1736.
- Arnaud Lejeune, Fabien Béchet, Hakim Boudaoud, Norman Mathieu, and Michel Potier-Ferry. 2012. Object-oriented design to automate a high order non-linear solver based on asymptotic numerical method. *Advances in Engineering Software* 48 (2012), 70–88.
- Hai-Jun Liao, Jin-Guo Liu, Lei Wang, and Tao Xiang. 2019. Differentiable programming tensor networks. *Physical Review X* 9, 3 (2019), 031041.
- Alex Limpaecher, Nicolas Feltman, Adrien Treuille, and Michael Cohen. 2013. Real-time drawing assistance through crowdsourcing. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–8.
- A Najah, B Cochelin, N Damil, and M Potier-Ferry. 1998. A critical review of asymptotic numerical methods. *Archives of Computational Methods in Engineering* 5, 1 (1998), 31–50.
- Théodore Papadopoulos and Manolis IA Lourakis. 2000. Estimating the jacobian of the singular value decomposition: Theory and applications. In *European Conference on Computer Vision*. Springer, 554–570.
- Steven Roman. 1980. The formula of Fàa di Bruno. *The American Mathematical Monthly* 87, 10 (1980), 805–809.
- Matthias Seeger, Asmus Hetzel, Zhenwen Dai, Eric Meissner, and Neil D Lawrence. 2017. Auto-differentiating linear algebra. *arXiv preprint arXiv:1710.08717* (2017).
- Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahar Z Kovalsky, and Yaron Lipman. 2017. Geometric optimization via composite majorization. *ACM Trans. Graph.* 36, 4 (2017), 38–1.
- Eftychios Sfakias 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*. 1–50.
- Breannan Smith, Fernando De Goes, and Theodore Kim. 2019. Analytic Eigensystems for Isotropic Distortion Energies. *ACM Transactions on Graphics (TOG)* 38, 1 (2019), 1–15.
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, Vol. 4. 109–116.
- Theano Development Team. 2016. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints* abs/1605.02688 (May 2016). <http://arxiv.org/abs/1605.02688>
- Joshua Trzasko and Armando Manduca. 2008. Highly Undersampled Magnetic Resonance Image Reconstruction via Homotopic ℓ_0 -Minimization. *IEEE Transactions on Medical imaging* 28, 1 (2008), 106–121.
- Nobuyuki Umetani, Danny M Kaufman, Takeo Igarashi, and Eitan Grinspun. 2011. Sensitive couture for interactive garment modeling and editing. *ACM Trans. Graph.* 30, 4 (2011), 90.
- KangKang Yin, Stelian Coros, Philippe Beaudoin, and Michiel van de Panne. 2008. Continuation methods for adapting simulated skills. In *ACM SIGGRAPH 2008 papers*. 1–7.
- H Zahrouni, B Cochelin, and M Potier-Ferry. 1999. Computing finite rotations of shells by an asymptotic-numerical method. *Computer methods in applied mechanics and engineering* 175, 1-2 (1999), 71–85.