Local Deformation for Interactive Shape Editing

Honglin Chen
Columbia University
New York City, NY, USA
honglin.chen@columbia.edu

Changxi Zheng
Columbia University
New York City, NY, USA
cxz@cs.columbia.edu

Kevin Wampler
Adobe Research
Seattle, WA, USA
kwampler@adobe.com

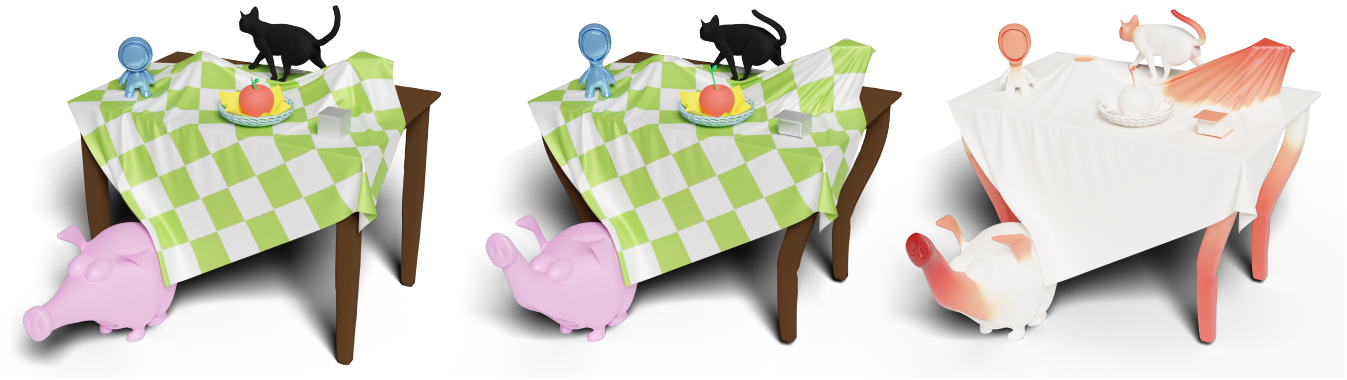

Figure 1: Our method enables the user to edit shapes in an interactive and physically plausible way. The edit is *local*, meaning that the user can focus on one region of the complex scene without worrying about inadvertent changes elsewhere. To visualize the locality, in the rightmost figure we highlight the regions where the vertex displacement is larger than 10^{-3} in red. (Undeformed scene shapes thanks to [Zhang et al. 2022])

ABSTRACT

We introduce a novel regularization for localizing an elastic-energy-driven deformation to only those regions being manipulated by the user. Our local deformation features a natural region of influence, which is automatically adaptive to the geometry of the shape, the size of the deformation and the elastic energy in use. We further propose a three-block ADMM-based optimization to efficiently minimize the energy and achieve interactive frame rates. Our approach avoids the artifacts of other alternative methods, is simple and easy to implement, does not require tedious control primitive setup and generalizes across different dimensions and elastic energies. We demonstrates the effectiveness and efficiency of our localized deformation tool through a variety of local editing scenarios, including 1D, 2D, 3D elasticity and cloth deformation.

CCS CONCEPTS

• **Computing methodologies** → *Mesh geometry models.*

KEYWORDS

Local control, shape deformation, elasticity, sparsity, ADMM.

SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA

ACM ISBN 979-8-4007-0159-7/23/08…$15.00
https://doi.org/10.1145/3588432.3591501

ACM Reference Format:
Honglin Chen, Changxi Zheng, and Kevin Wampler. 2023. Local Deformation for Interactive Shape Editing. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 6–10, 2023, Los Angeles, CA, USA.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3588432.3591501

1 INTRODUCTION

Local deformation is a core component in modeling and animation. In a localized deformation, only the parts of the shape near where the user is currently manipulating move—everything else stays still, ensuring that the user can focus entirely on one region of the shape without worrying about inadvertent changes elsewhere. However, existing localized deformation tools tend to have practical impediments for interactive design: they are either too slow to run, unaware of the geometry, introduce artifacts, or require a careful control point setup.

When global deformation is acceptable, a widely useful approach is to solve for the deformation by minimizing an elastic energy defined over the shape, subject to positional constraints derived from the user's input. This paradigm has many advantages. The deformation accounts for the geometry of the shape, generalizes well to 2D, 3D, and cloth, and the elastic energy can be used to model a wide range of both real-world and stylized materials. Unfortunately elastic energy minimization is by its nature global, and jointly solves for all the degrees of freedom in the shape. This necessitates a rigging step to "pin down" certain aspects of the deformation lest the optimizer move them. This limits the applicability of such methods to situations where a suitable rig is available, or the region of influence for a deformation is known in advance.

We seek to combine the advantages of elastic energy minimization with the locality of sculpting-style tools. In doing so, we also want the locality of a deformation to be *automatic*, *natural* and *efficient*. The locality of an edit should be automatic in that the region of influence (ROI) of the deformation should scale automatically depending on the size of the desired deformation. In addition, although we do not require a rig, in cases where a rig or other constraints have been placed on the shape, the ROI needs to automatically adapt to them. We further want the notion of locality to be *natural* in that it adjusts to both the geometry of the shape and the elastic energy driving its deformation, where changes to either will lead to a fitting change in the ROI. Finally, we require that the method be fast enough to run in real time.

We achieve this with the following contributions:

- We introduce a novel deformation regularizer, called a smoothly clamped ℓ_1 (SC-L1) loss which augments an elastic energy with a notion of locality. SC-L1 regularization is simple to implement, and avoids artifacts of previous methods.
- We enable real-time localized deformation with an ADMM-based optimization algorithm for SC-L1-regularized deformation which is significantly faster than prior work using a group lasso regularizer.

We illustrate the utility of SC-L1 regularization on a wide range of examples, including multiple different elastic energies, 1D curves, 2D and 3D meshes, and cloth. This provides a localized deformation tool which avoids artifacts of other regularizers, is easy to implement, generalizes across different dimensions and material models, and performs fast enough to run in real time.

2 RELATED WORK

Shape deformation algorithms in computer graphics have typically fallen into one of two categories, which we will call the *direct* approach and the *optimization* approach. In the direct approach, the shape's deformation is an explicit function of the user's input, often either by modifying some high-level parameterization or by applying a pre-specified deformation field. In the optimization approach, the shape deformation is an indirect product of the user's input combined with an elastic energy, and the deformation itself is only known after an optimization process has converged to minimize this energy. These two approaches differ in how (and often whether) they enforce the locality of a deformation.

2.1 Locality in High-Level Parameterizations

A time-honored and common approach for localized deformation is the direct manipulation of high-level parameters. From the onset, a notion of locality is baked into these parameters, which can directly encode the shape itself, often using splines [Hoschek et al. 1993] or a "rig" to control a shape's deformation, as with cage-based generalized barycentric coordinates [Joshi et al. 2007; Lipman et al. 2008], linear-blend skinning [Jacobson et al. 2011; Magnenat-Thalmann et al. 1989; Wang et al. 2015], lattice deformers [Coquillart 1990; Sederberg and Parry 1986], wire curves [Singh and Fiume 1998], or learned skinning weights [Genova et al. 2020], to list a few.

Approaches of this nature, although popular in computer graphics, have a few disadvantages. Firstly, since the locality is baked into the parameterization, it cannot easily adapt based on changes in the deformation. Secondly, without an optimization step, even in cases where a capable elastic energy is within reach there is no way to incorporate it into the deformation.

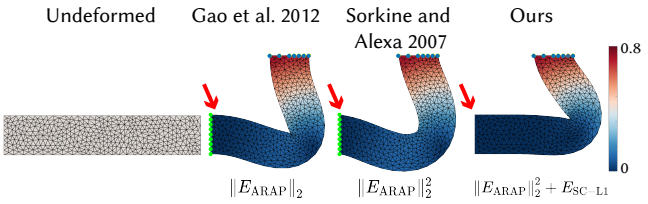

Figure 2: Directly applying different norms to an elastic energy leads to global deformation, and thus the user needs to carefully set up additional fixed handles (green points) to keep the shapes from freely moving around in the space.

2.2 Locality in Deformation Fields

Instead of providing localized deformation via pre-chosen parameters, it is also possible to define a localized deformation field, then apply it to a shape. These methods often follow a "sculpting" metaphor, and include simple move, scale, pinch, and twist edits as well as more sophisticated operations [Cani and Angelidis 2006]. Recently, De Goes and James [2017] introduced *regularized Kelvinlets*, which provides real-time localized volumetric control based on the regularized closed-form solutions of linear elasticity. These closed-form solutions were later extended to handle dynamic secondary motions [De Goes and James 2018], sharp deformation [de Goes and James 2019] and anisotropic elasticity [Chen and Desbrun 2022].

These approaches allow for local deformation with real-time feedback. However, as they are designed for digital sculpting, these methods usually require the user to explicitly pick the falloff of the brushes. Furthermore, these methods are usually based on Euclidean distance, unaware of the shape's geometry. In contrast, our method is shape-aware and enables automatic dynamic region of influence with interactive feedback.

2.3 Localized Optimization via an ROI

The most natural formulation of optimization-based deformation editing is by solving globally for the entire shape's deformation at once [Shtengel et al. 2017; Smith et al. 2019; Zhu et al. 2018]. Nevertheless, there are methods attempting to enforce locality in the shape optimization process. Previous methods of this sort have often computed the ROI of a manipulation as a preprocessing step, then restricting the optimization to only move parts of the shape within this ROI. The ROI can be taken as an input [Alexa 2006] or based on a small amount of user markup [Kho and Garland 2005; Luo et al. 2007; Zimmermann et al. 2007]. Other methods combine deformation energies with handle-based systems, including skeleton rigs [Hahn et al. 2012; Jacobson et al. 2012; Kavan and Sorkine 2012] and cages [Ben-Chen et al. 2009].

Unfortunately, in many contexts, the ROI is hard or impossible to know in advance. This is particularly the case when constraints are involved, or where it is not known in advance if a deformation will be small (best fitting a small ROI) or large (best fitting a large ROI). In addition, the correct ROI may also depend on the elastic energy driving the deformation, thus difficult to account for when the ROI calculation is decoupled in a separate step.

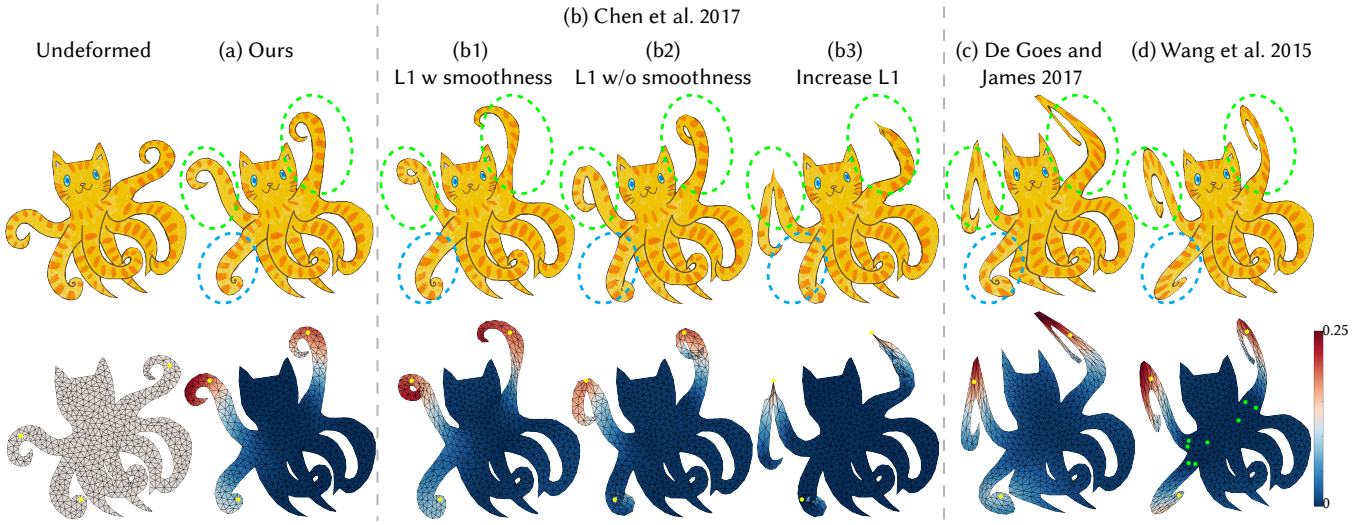

Figure 3: Our method (a) produces *local* deformation which is smooth, shape-aware and has automatically adaptive region of influence without the need of additional handles. Here we highlight all the handles that have been moved in yellow. For the other alternatives, l_1-norm-based method (b) contains artifacts due to the use of l_1 norm (see the green highlighted regions); regularized Kelvinlets technique (c) is based on Euclidean distance and unaware of the geometry (see the blue highlighted regions); and biharmonic coordinates approach (d) requires a careful placement of additional handles (green) to explicitly control the region of influence. Note that the use of l_1-norm (b) leads to either small global motions and minor artifacts (b2, smaller l_1-norm) or no global motions but significant artifacts (b3, larger l_1-norm). Thus [Chen et al. 2017] mitigates it with a smoothing regularizer (b1), but this changes the elastic energy and it no longer deforms like a localized ARAP.

2.4 Localized Optimization via Sparsity Norms

Sparsity-inducing norms, such as the smooth ℓ_0 norm, have been widely applied to many domains and problems including medical image reconstruction [Xiang et al. 2022], sparse component analysis [Mohimani et al. 2007] and UV mapping [Poranne et al. 2017]. To allow an adaptive ROI while preserving the benefits of optimization-based deformation, a few works have adopted a sparsity-inducing norm, typically a $\|x\|_1$ or $\|x\|_2$ norm, in their energy, which is then minimized by Alternating Direction Method of Multipliers (ADMM) [Boyd et al. 2011; Peng et al. 2018; Zhang et al. 2019] or Augmented Lagrangian Method (ALM) [Bertsekas 1996]. We refer the readers to [Xu et al. 2015] for a survey on sparsity in geometry modelling and processing.

Several methods of this variety rely on sparsity-inducing norm formulated as a sum of $\|x\|_2$ norms, referred to as $\ell_{2,1}$ or ℓ_1/ℓ_2 norms, or a *group lasso* penalty. They have been applied in a preprocessing phase to compute sparse deformation modes for interactive local control [Brandt and Hildebrandt 2017; Deng et al. 2013; Neumann et al. 2013]. However, the deformation is limited by the linear deformation modes and thus struggles with large deformation.

Another class of methods adds a sparsity-induced regularization to an elastic energy optimization to achieve local deformation. Gao et al. [2012] applied different ℓ_p sparsity norms to the as-rigid-as-possible energy [Sorkine and Alexa 2007] to create various deformation styles. Recently, Chen et al. [2017] used $\ell_{2,1}$ regularization on vertex positions to locally control the deformation. However, their direct use of the $\ell_{2,1}$ norm will create artifacts when the control point is not on the boundary of the shape (see Fig. 3(b) and Fig. 13(b)). Moreover, their method requires one ADMM solve in each global iteration, which renders the optimization less efficient and slow in runtime. Also, their framework is limited to 2D deformation with ARAP energy only, while our framework generalizes across dimensions and a variety of energy models.

Our algorithm, inspired by sparsity-seeking regularizers such as that used by [Chen et al. 2017; Fan and Li 2001], addresses these shortcomings. In particular, we propose a simple and novel sparsity-inducing norm that eliminates artifacts arising from the $\ell_{2,1}$ norm, and our efficient optimization scheme leads to interactive performance.

3 OVERVIEW

The main idea of our method is to use a novel regularization term to produce local deformation with a dynamic region of influence (ROI). Our method takes a triangle/tetrahedral mesh (or a 1D polyline) and a set of selected vertices as control handles as input. The output of our method is a deformed shape where the deformation is both *local* and *natural* and the ROI is automatically adaptive to the deformation. Here the "locality" implies that a handle only dominates its nearby areas without affecting the regions far away.

3.1 SC-L1 Regularization

We suggest a sparsity-inducing regularization term to produce natural local deformation. This regularizer is applied per-vertex to $\mathbf{V}_i - \widetilde{\mathbf{V}}_i$ to bias each vertex deformed position $\mathbf{V}_i$ to exactly match its initial rest position $\widetilde{\mathbf{V}}_i$ except in isolated regions of the shape. The most obvious choice for this regularization would be to enforce sparsity either with an ℓ_1-norm, or with a *group lasso* / $\ell_{2,1}$ *regularization* defined as $\sum_i \|\mathbf{V}_i\|_2$ as in [Chen et al. 2017].

However, direct use of a $\ell_{2,1}$ regularization term leads to artifacts, due to the fact that the $\ell_{2,1}$-norm competes with the elastic energy by dragging *all* the vertices towards their original positions. This results in undesired distortion near the deformation handles (see Fig. 3(b) and Fig. 13(b)), making the deformed region look unnatural. Previous $\ell_{2,1}$-based methods [Chen et al. 2017] focus specifically on ARAP-like deformation, and attempt to alleviate this artifact by adding a Laplacian smoothness term and a weighting term based on biharmonic distance. Unfortunately, as seen in Fig. 3 at the ends of the octocat's tentacles, artifacts can arise in regions quite close to the deformation handles, and there is not necessarily any setting of these parameters which alleviates these artifacts without oversmoothing the entire deformation.

Inspired by "folded concave" losses in statistical regression [Fan and Li 2001; Zhang 2010] and the use of the $\ell_{2,1}$ loss in deformation [Chen et al. 2017], we propose using a smoothly clipped group ℓ_1-norm as our locality-inducing regularization. We call it a SC-L1 loss for "smoothly clamped ℓ_1 loss" (see the inset).

We adopt a simple implementation for our SC-L1 loss:

$$\|\mathbf{x}\|_{\text{SC-L1}} = \begin{cases} \|\mathbf{x}\|_2 - \frac{1}{2s}\|\mathbf{x}\|_2^2 & \|\mathbf{x}\|_2 < s \\ \frac{1}{2}s & \|\mathbf{x}\|_2 \geq s \end{cases}, \tag{1}$$

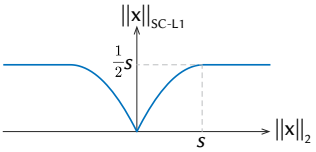

where s is the threshold distance beyond which the regularizer is disabled (this is equivalent to a group-variant of the MCP loss in [Zhang 2010], but we use the term "SC-L1" to emphasize that the min-imax concave property is not critical for localized deformation). This function is continuously differentiable and piecewise smooth, and admits a proximal shrinkage operator free of local minima. Near the origin the SC-L1 loss function acts like the group ℓ_1-norm, which drives $\|\mathbf{x}\|_{\text{SC-L1}}$ towards 0 in a sparse-deformation-seeking manner. When $\|\mathbf{x}\|_2 \geq s$, the SC-L1 loss function value is a constant and has no penalty on $\|\mathbf{x}\|_2$. For a detailed comparison between our SC-L1 loss function and other alternatives, please see Sec.5 of the supplementary material.

3.2 Local Deformation Energy

We denote $\mathbf{V}$ as a $|\mathbf{V}| \times d$ matrix of vertex positions at the deformed state, and $\widetilde{\mathbf{V}}$ as a $|\mathbf{V}| \times d$ matrix containing rest state vertex positions.

The total energy for our local deformation is as follows:

$$\underset{\mathbf{V}}{\text{minimize}} \quad \underbrace{E(\mathbf{V})}_{\text{Elasticity}} + \sum_{i \in V} \underbrace{w a_i \|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}}_{\text{Locality}}, \tag{2a}$$

$$\text{s.t.} \quad \mathbf{V}_s = \mathbf{p}_s, \quad \text{(position constraint)} \tag{2b}$$

$$\mathbf{A}_k \mathbf{V}_t = \mathbf{b}_k, \forall t \in \mathcal{S}_k \quad \text{(affine constraint)} \tag{2c}$$

The first term is an elasticity energy of choice, and can be selected independent of the locality regularization. The second term is the novel "SC-L1 loss" term on the vertex position changes, which measures the locality of the deformation. a_i is the barycentric vertex area of the i-th vertex, which ensures the consistency of the result across different mesh resolutions for the same constant w. To enable more user control, position constraints and optional affine constraints can be added on selected vertices to achieve different deformation effects. s denotes the indices of the vertices with the position constraint, and we call these vertices "handles". $\mathcal{S}_k$ is the k-th set of vertex indices where an affine constraint is added. For simplicity, we omit the position constraints and affine constraints in the discussion below, as they can be easily intergrated to the system by removing the corresponding degrees of freedom and using Lagrange multiplier method (see Eq(1) in [Wang et al. 2015]).

ALGORITHM 1: Three-block ADMM Overview

Input: A triangle or tetrahedral mesh $\widetilde{\mathbf{V}}, \mathbf{T}$
Output: Deformed vertex positions $\mathbf{V}$
$\mathbf{V} \leftarrow \widetilde{\mathbf{V}}$
while *not converged* **do**
  $\mathbf{X}_i \leftarrow \mathit{local\_step\_X}(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 1
  $\mathbf{Z}_i \leftarrow \mathit{local\_step\_Z}(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 2
  $\mathbf{V} \leftarrow \mathit{global\_step}(\mathbf{X}_i, \mathbf{Z}_i, \widetilde{\mathbf{V}})$ ▷ global step
  $\mathbf{U}_i \leftarrow \mathit{dual\_update}(\mathbf{Z}_i, \mathbf{V}_i, \widetilde{\mathbf{V}}_i)$ ▷ dual update 1
  $\mathbf{W}_i \leftarrow \mathit{dual\_update}(\mathbf{X}_i, \mathbf{V}_i, \widetilde{\mathbf{V}}_i)$ ▷ dual update 2
end

Inspired by the local-global strategy in [Brown and Narain 2021], our local deformation energy (Eq. 2) can be rewritten as:

$$\underset{\mathbf{V}, \{\mathbf{X}_j\}}{\text{minimize}} \quad \sum_j \underbrace{E(\mathbf{X}_j)}_{\text{Elasticity}} + \sum_{i \in V} \underbrace{w a_i \|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}}_{\text{Localness}}, \tag{3a}$$

$$\text{s.t.} \quad \mathbf{X}_j = \text{sym}(\mathbf{D}_j \mathbf{V}), \forall j, \tag{3b}$$

where $\mathbf{D}_j$ is the selection matrix for edges of the j-th vertex or element. $\text{sym}(\mathbf{F})$ denotes the symmetric factor $\mathbf{S}$ computed using the polar decomposition $\mathbf{F} = \mathbf{RS}$, where $\mathbf{F}$ is the deformation gradient. Thus $\mathbf{X}_j$ is the symmetric factor of deformation gradient of the j-th vertex or element. (Note that in general $\text{sym}(\mathbf{F}) \neq \frac{1}{2}(\mathbf{F} + \mathbf{F}^\top)$.) The goal of using sym() here is to ensure the local coordinates $\mathbf{X}_j$ are invariant to rotations as well as translations. For details, please see [Brown and Narain 2021].

4 OPTIMIZING WITH ADMM

A natural way to minimize this energy is to use the alternating direction method of multipliers [Boyd et al. 2011] for the sparsity term, and to use a local-global update strategy for the elasticity term. However, previous $\ell_{2,1}$-based methods [Chen et al. 2017] apply these two strategies separately in their local and global steps, resulting in an inefficient optimization scheme. As they only support 2D ARAP energy, we discuss further in Sec. 4.1.

We propose a new way to efficiently minimize such energies in Eq. 3 by combining the sparsity-targeted ADMM with the elasticity-focused local-global strategy. We minimize our energy (Eq. 3) using a three-block alternating direction method of multipliers scheme [Boyd et al. 2011] following the local-global update strategy. Our first local step, finding the optimal symmetric factor $\mathbf{X}_i$ of the deformation gradient, can be formulated as a minimization problem on its singular values. Our second local step, minimizing the SC-L1 loss term for each $\mathbf{Z}_i$, can be solved using a shrinkage step. Our global step, updating vertex positions $\mathbf{V}$, is achieved by solving a linear system. We provide an overview of our three-block ADMM scheme in Alg. 1.

4.1 Example I: Local ARAP Energy

We begin by considering how to minimize an as-rigid-as-possible (ARAP) energy [Sorkine and Alexa 2007] when combined with a SC-L1 loss regularizer. With an ARAP elastic energy, the total energy (Eq. 3) for our local deformation is as follows:

$$\underset{\mathbf{V},\{\mathbf{R}_i\}}{\text{minimize}} \quad \sum_{i\in V} \underbrace{\frac{1}{2}\|\mathbf{R}_i\mathbf{D}_i - \widetilde{\mathbf{D}}_i\|^2_{\mathbf{W}_i}}_{\text{ARAP}} + \underbrace{wa_i\|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}}_{\text{Localness}}, \tag{4}$$

where $\mathbf{R}_i$ is a $d\times d$ rotation matrix, W_i is a $|\mathcal{N}(i)|\times|\mathcal{N}(i)|$ diagonal matrix of cotangent weights, $\widetilde{\mathbf{D}}_i$ and $\mathbf{D}_i$ are $3\times|\mathcal{N}(i)|$ matrices of "spokes and rims" edge vectors of the i-th vertex at the rest and deformed states respectively. $\|\mathbf{X}\|^2_{\mathbf{W}_i}$ denotes $\text{Tr}(\mathbf{X}^\top\mathbf{W}_i\mathbf{X})$. Here we use $\mathbf{R}_i$ to denote $\mathbf{X}_i$, since we drive the deformation gradient towards a rotation matrix in ARAP energy.

Previous method [Chen et al. 2017] optimizes the $\ell_{2,1}$ version of Eq. 4 in a less efficient way. Their local step optimizes over per-vertex rotation $\mathbf{R}_i$ and their global step minimizes over vertex positions $\mathbf{V}$ using a two-block ADMM scheme. This leads to an expensive optimization with a full ADMM optimization in *each* global step, making their method too slow for interactive usage. In contrast, applying our new three-block ADMM scheme to the local ARAP energy results in a much more efficient solver, which is one ADMM optimization itself (see the inset, where the blue regions denote an ADMM optimization). We further show the pseudocode of our three-block ADMM for local ARAP energy in Suppl. Alg.1.

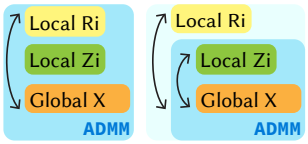

More concretely, by setting $\mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i$, we can further rewrite Eq. 4 as

$$\underset{\mathbf{V},\{\mathbf{R}_i\},\mathbf{Z}}{\text{minimize}} \quad \sum_{i\in V} \frac{1}{2}\|\mathbf{R}_i\mathbf{D}_i - \widetilde{\mathbf{D}}_i\|^2_{\mathbf{W}_i} + wa_i\|\mathbf{Z}_i\|_{\text{SC-L1}}, \tag{5a}$$

$$\text{s.t.} \quad \mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i, \quad \forall i. \tag{5b}$$

The above minimization problem can be solved efficiently using the following ADMM update steps:

$$\mathbf{R}_i^{k+1} \leftarrow \underset{\mathbf{R}_i\in\text{SO}(3)}{\arg\min} \ \frac{1}{2}\|\mathbf{R}_i\mathbf{D}_i - \widetilde{\mathbf{D}}_i\|^2_{\mathbf{W}_i} \tag{6a}$$

$$\mathbf{Z}_i^{k+1} \leftarrow \underset{\mathbf{Z}_i}{\arg\min} \ wa_i\|\mathbf{Z}_i\|_{\text{SC-L1}} + \frac{\rho}{2}\|\mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i + \mathbf{U}_i^k\|_2^2 \tag{6b}$$

$$\mathbf{V}^{k+1} \leftarrow \underset{\mathbf{V}}{\arg\min} \ \mathbf{W}(\mathbf{V}^\top\mathbf{L}\mathbf{V} - \mathbf{B}^\top\mathbf{V}) + \frac{\rho}{2}\|\mathbf{V} - \widetilde{\mathbf{V}} - \mathbf{Z}^k + \mathbf{U}^k\|_2^2 \tag{6c}$$

$$\mathbf{U}_i^{k+1} \leftarrow \mathbf{U}_i^k + \mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i^{k+1} \tag{6d}$$

Here ρ is a fixed penalty parameter. For a detailed derivation of the ADMM update, please see Sec. 2 of the supplementary material.

The various steps in this ADMM-based algorithm are computed as follows:

For **updating** $\mathbf{R}_i$, local step 1 (Eq. 6a) is an instance of the Orthogonal Procrustes problem, which can be solved in the same way as the rotation fitting step in [Sorkine and Alexa 2007]. The optimal $\mathbf{R}_i$ can be computed as $\mathbf{R}_i^{k+1} \leftarrow \mathcal{V}_i\mathcal{U}_i^\top$ from the singular value decomposition of $\mathbf{M}_i = \mathcal{U}_i\Sigma_i\mathcal{V}_i^\top$, where $\mathbf{M}_i = \mathbf{D}_i\widetilde{\mathbf{D}}_i^\top$.

For **updating** $\mathbf{Z}_i$, following the derivation of the proximal operator of SC-L1 loss in Sec. 1 of the supplementary material, our local step 2 (Eq. 6b) is solved using a SC-L1 loss-specific shrinkage step:

$$\mathbf{Z}_i^{k+1} \leftarrow \mathcal{S}^k_{wa_i}\left(\mathbf{V}_i - \widetilde{\mathbf{V}}_i + \mathbf{U}_i\right) \tag{7}$$

$$\mathcal{S}_{wa_i}(\mathbf{x}) = \begin{cases} \left(\frac{\rho s - wa_i s/\|\mathbf{x}\|_2}{\rho s - wa_i}\right)_+ \mathbf{x}, & \text{if } \|\mathbf{x}\|_2 \le s \\ \mathbf{x}, & \text{otherwise} \end{cases} \tag{8}$$

To avoid local minima in the shrinkage step, this assumes ρ is set to satisfy $\rho > \frac{\max(wa_i)}{s}$ (see Sec. 1 of the supplementary material).

For **updating** $\mathbf{V}$, the global step (Eq. 6c) can be achieved by solving a linear system:

$$(\mathbf{L} + \rho\mathbf{I})\mathbf{V} = \mathbf{B} + \rho(\widetilde{\mathbf{V}} + \mathbf{Z}^k - \mathbf{U}^k), \tag{9}$$

where the Laplacian $\mathbf{L}$ and $\mathbf{B}$ are defined in the same way as the global step (Eq. 9) in [Sorkine and Alexa 2007]. For fixed ρ an efficient implementation is obtained by precomputing and storing the Cholesky factorization of $\mathbf{L} + \rho\mathbf{I}$.

4.2 Example II: Local Neo-Hookean Energy

Our local deformation scheme can be further extended to physics-based elasticity energies, e.g., Neo-Hookean energy. Using the Neo-Hookean energy as our elasticity energy and following the framework of [Brown and Narain 2021], the optimization problem in Eq. 3 can be written as follows:

$$\underset{\mathbf{V},\{\mathbf{X}_j\},\{\mathbf{Z}_i\}}{\text{minimize}} \quad \sum_{j\in T} \underbrace{E_{\text{nh}}(\mathbf{X}_j)}_{\text{Neo-Hookean}} + \sum_{i\in V} \underbrace{wa_i\|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}}_{\text{Locality}}, \tag{10a}$$

$$\text{s.t.} \quad \mathbf{X}_j = \text{sym}(\mathbf{D}_j\mathbf{V}), \forall j, \tag{10b}$$

where T denotes all the elements.

Similarly, by introducing $\mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i$, we can minimize our local Neo-Hookean energy using ADMM. For a detailed derivation, please see Sec. 3 of the supplementary material.

The ADMM update (Alg. 1) for the above minimization problem is as follows:

$$\mathbf{X}_j^{k+1} \leftarrow \underset{\mathbf{X}_j}{\arg\min} \ E_{\text{nh}}(\mathbf{X}_j) + \frac{\gamma}{2}\|\text{sym}(\mathbf{D}_j\mathbf{V}) - \mathbf{X}_j + \mathbf{W}_j\|_2^2 \tag{11a}$$

$$\mathbf{Z}_i^{k+1} \leftarrow \underset{\mathbf{Z}_i}{\arg\min} \ wa_i\|\mathbf{Z}_i\|_{\text{SC-L1}} + \frac{\rho}{2}\|\mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i + \mathbf{U}_i^k\|_2^2 \tag{11b}$$

$$\mathbf{V}^{k+1} \leftarrow \underset{\mathbf{V}}{\arg\min} \ \sum_{j\in T}\frac{\gamma}{2}\|\text{sym}(\mathbf{D}_j\mathbf{V}) - \mathbf{X}_j + \mathbf{W}_j\|_2^2 + \sum_{i\in V}\frac{\rho}{2}\|\mathbf{V} - \widetilde{\mathbf{V}} - \mathbf{Z}^k + \mathbf{U}^k\|_2^2 \tag{11c}$$

$$\mathbf{W}_j^{k+1} \leftarrow \mathbf{W}_j^k + \text{sym}(\mathbf{D}_j\mathbf{V}) - \mathbf{X}_j \tag{11d}$$

$$\mathbf{U}_i^{k+1} \leftarrow \mathbf{U}_i^k + \mathbf{V}^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i^{k+1} \tag{11e}$$

Here ρ and γ are fixed penalty parameters.

The local step 2 (updating $\mathbf{Z}_i$) and the global step (updating $\mathbf{V}$) can be solved in the same way as the local ARAP energy (see Sec. 4.1).

For **updating** $\mathbf{X}_i$, local step 1 (Eq. 11a) can be solved by performing the energy minimization on the singular values of $\text{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$. This is a proximal operator of E_{nh} at the rotation-invariant $\text{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$.

ALGORITHM 2: Three-block ADMM for local ARAP Energy

Input: A triangle or tetrahedral mesh $\widetilde{\mathbf{V}}, \mathbf{T}$
Output: Deformed vertex positions $\mathbf{V}$
$\mathbf{V} \leftarrow \widetilde{\mathbf{V}}$
while *not converged* **do**
  $\mathbf{R}_i \leftarrow \mathit{local\_step\_X}(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 1
  $\mathbf{Z}_i \leftarrow \mathit{local\_step\_Z}(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 2
  $\mathbf{V} \leftarrow \mathit{global\_step}(\mathbf{X}_i, \mathbf{Z}_i, \widetilde{\mathbf{V}})$ ▷ global step
  $\mathbf{U}_i \leftarrow \mathit{dual\_update}(\mathbf{Z}_i, \mathbf{V}_i, \widetilde{\mathbf{V}}_i)$ ▷ dual update 1
end

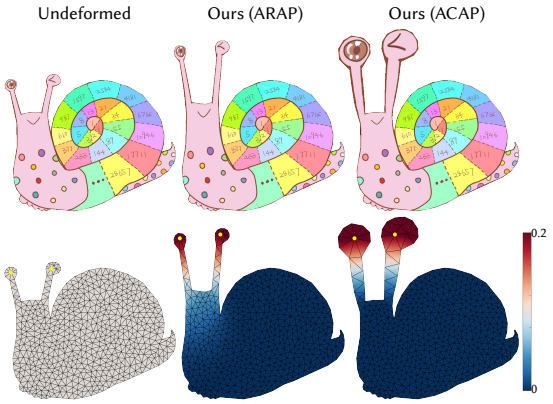

Figure 4: Our method automatically choose a natural ROI based on the elastic energy in use. Here we use the same parameter settings for both the local ARAP and local ACAP energies. With the local ACAP energy, we have a smaller ROI than the case of local ARAP energy as the ACAP energy allows for local scaling.

Let us denote the proximal operator of E_{nh} and the singular value decomposition of $\text{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$ as:

$$\text{prox}_{E_{\text{nh}}}(\mathbf{X}_j) = E_{\text{nh}}(\mathbf{X}_j) + \frac{\gamma}{2}\|\text{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j - \mathbf{X}_j\|_2^2 \tag{12}$$

$$\text{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j = \mathcal{U}_j\Sigma_j\mathcal{V}_j^\top \tag{13}$$

where γ is the augmented Lagrangian parameter for $\mathbf{X}_j$.

As shown by [Brown and Narain 2021], we can compute the optimal $\mathbf{X}_j$ as:

$$\Sigma_j^{k+1} \leftarrow \arg\min_{\Sigma_j} \text{prox}_{E_{\text{nh}}}(\mathcal{U}_j\Sigma_j\mathcal{V}_j^\top) \tag{14}$$

$$\mathbf{X}_j \leftarrow \mathcal{U}_j\Sigma_j^{k+1}\mathcal{V}_j^\top \tag{15}$$

Specifically, one can compute the SVD of $\text{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$ and perform the minimization of $\text{prox}_{E_{\text{nh}}}$ only on its singular values, while keeping singular vectors unchanged. The above optimization of singular values Σ_j^{k+1} can be performed using an L-BFGS solver.

4.3 Extension to Other Elastic Energies

Our algorithm can easily generalize across different dimensions and material models. Switching the material model only requires a change on the minimization problem in the local step 1 $\arg\min_X$, which can be optimized over the singular values of the symmetric factor $\mathbf{X}_i$ of the deformation gradient.

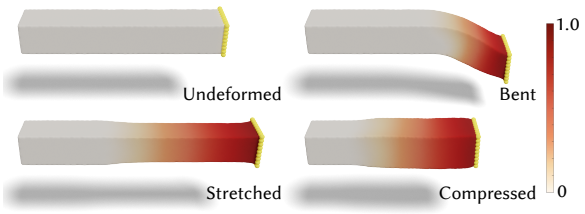

Figure 5: Given the *same* handle offset magnitude, a "natural" ROI size also depends on the way the handles are moved, which is more complex than simply growing the ROI proportional to the handle displacement. Here we moved the handles (yellow) with the *same* offset magnitude 1.0 towards the bottom, left and right respectively, resulting different ROIs for the *same* handle offset magnitude.

4.3.1 As-Conformal-As-Possible Energy. For editing tasks where users intend to locally scale the geometry while preserving the texture, it's desirable to constrain the angle preservation, or conformality (see Fig. 4 and Fig. 10). We can adapt the ARAP energy to as-conformal-as-possible (ACAP) energy [Bouaziz et al. 2012] by allowing local scaling:

$$E_{\text{ACAP}}(\mathbf{V}) = \sum_{k\in T}\sum_{i,j\in\mathcal{N}(k)} \frac{w_{ij}}{2}\|s_k\mathbf{R}_k\widetilde{\mathbf{d}}_{ij} - \mathbf{d}_{ij}\|_2^2 \tag{16}$$

where s_k is a scalar controlling the scaling of the local patch and can be computed analytically (see Sec.4 of the supplementary material).

4.3.2 Cloth. Our method also generalizes to higher co-dimensional settings, such as deformable thin sheets and cloth in $\mathbb{R}^3$. We model the cloth deformation using ARAP elasticity (Eq. 4), hard strain limiting, and quadratic bending resistance [Bergou et al. 2006].

4.3.3 1D Polyline. Our algorithm can be also extended to the local editing of 1D polyline in vector graphics. The deformation of a polyline can be modeled using the ARAP energy (Eq. 4) with uniform weights.

5 RESULTS

We evaluate our method by comparing it against existing local deformation tools and showcasing its extension to various elastic energies. All the colormaps in our figures visualize the vertex displacement with respect to the rest shape. The accompanying video also includes several animation examples generated using our local deformation tool.

We implement a 2D version of our method in MATLAB with gptoolbox [Jacobson et al. 2018a], and a 3D version in C++ with libigl [Jacobson et al. 2018b] based on the WRAPD framework [Brown and Narain 2021]. We also implement the 2D version of our method in C++ for runtime evaluation and comparison. Benchmarks are performed using a MacBook Pro with an Apple M2 processor and 24GB of RAM for 3D and a Windows desktop with an i9-9900K 3.60 GHz CPU for 2D. Table 2 in the supplementary material shows the performance statistics and relevant parameters of all our examples.

Quality. We compare our methods against other local editing tools, including **i)** ℓ_1-based deformation [Chen et al. 2017], **ii)** regularized Kelvinlets [De Goes and James 2017] and **iii)** biharmonic

coordinates [Wang et al. 2015]. Among them, the use of $\ell_{2,1}$-norm regularization [Chen et al. 2017] causes artifacts (see Fig. 3-b and Fig. 13-b). Regularized Kelvinlets technique [De Goes and James 2017] deforms a shape based on Euclidean distances, thus not shape-aware, creating artifacts when two disjoint parts are close in Euclidean space but far away geodesically (see the blue region in Fig. 3-c and the teeth area in Fig. 13-c). Methods based on biharmonic coordinates, such as [Wang et al. 2015], usually require careful placement of additional fixed control points to pre-determine the ROI. The latter two methods do not minimize any elastic energy in the deformation process, and thus their deformations are more susceptible to shape distortion (see Fig. 13). We additionally compare our method against the *sparse* deformation method [Gao et al. 2012], which directly introduces a sparsity-induced norm in ARAP energy. As shown in Fig. 2, the resulting deformation is *sparse* but not *local*, thus requiring the setup of additional fixed constraints.

In contrast, our method produces the deformation which is local, natural, and shape-aware; it automatically adapts the ROI without the need for careful control primitive setup.

Efficacy. We illustrate the ROI adaptation of our method in different situations: It adapts to different energy models—for example, the local ACAP has a smaller ROI than the local ARAP energy as the former allows for local scaling (see Fig. 4). It also adapts to different extents of deformation. As Fig. 8 demonstrates, the ROI gradually increases as the deformation of the bar becomes larger.

One can configure the local deformation style by choosing various elastic energy models. For example, the local Neo-Hookean energy leads to deformation that preserves volume, while the local ARAP energy is volume agnostic (see Fig. 11). The deformation can be further tuned by introducing additional affine constraints—for instance, to enable the character to wave hands (Fig. 12) or the crocodile to open its mouth (Fig. 13)—in a natural way.

Performance. In terms of performance, our solver is able to efficiently minimize the energy at interactive rate, while the method of [Chen et al. 2017] is too slow to run in realtime. Because their method only supports 2D ARAP energy, in Table 1 of the supplementary material, we evaluate the runtime of our method (using both the SC-L1 loss and $\ell_{2,1}$ loss) and [Chen et al. 2017] on a 2D ARAP local energy and across different mesh resolutions and deformations. Measured with the same convergence threshold, our method runs orders of magnitude faster than [Chen et al. 2017], achieving roughly 1000× speedup for small deformation and 100× speedup for large deformation.

Extensibility. Our method can easily generalize to other dimensions and material models, such as the ACAP deformation, cloth deformation, and 1D polyline deformation (see Fig. 10 and Fig. 4). The local ACAP energy enables local scaling and better preserves the texture around the deformed region. In Fig. 7, the user can interactively edit a polyline and naturally recovers its rest shape, which is a desirable feature by the users. In Fig. 6, to deform a cloth in a physically plausible way, the deformation locality is particularly useful, as otherwise a local edit of the cloth may cause a global change leading to unexpected intersections with other objects. Lastly, to demonstrate our method in a more complex scenario, an editing session involving multiple objects and clothes is shown in Fig. 1.

6 CONCLUSION & FUTURE WORK

We describe a regularization based on an "SC-L1 loss" which provides an effective and simple to implement tool for localizing an elastic energy driven deformation to only those regions of a shape being manipulated by a user. The region of influence induced by our method naturally adapts to the geometry of the shape, the size of the deformation, and the elastic energy being used. Furthermore SC-L1 regularization is generic enough to be applied to a wide range of shapes and elastic energies, including 1D, 2D, 3D and cloth finite element deformation, and is fast enough to be used in real-time. Our proposed approach offers several benefits for shape manipulation: It avoids undesired movement in far-off regions of a shape when only one part is being moved by the user, it allows parts of a shape to be deformed with direct manipulation without a pre-rigging step, and avoids the visual artifacts of previous work.

There remain several issues related to localized shape deformation not addressed by our method. Firstly our regularization is applied independently per-vertex, which makes it difficult to apply to splines, NURBS, or even meshes with highly irregular element sizes, which we mark an important direction for future work. In addition, since we use an ADMM method in the optimization, our approach suffers from the common shortcomings of applying ADMM to non-convex energies, including lack of convergence guarantees and slow convergence when high precision is required. Exploration of other optimization algorithms alleviating these issues is another useful future direction. Finally, although it is out of scope for our work here, we note in particular the usefulness of incorporating localized elastic energy deformation into sculpting workflows for artists. This involves a number of facets: Choosing the correct elastic energy to achieve an artistic effect, providing an intuitive UI to adjust the scale of the ROI (for instance by adjusting w and s, see Fig. 9 and the supplementary video), and in ensuring that our tool integrates well with other sculpting tools. This is particularly useful when handling large "freeform" deformations, as the elastic energy will tend to fight against such deformations, making other tools more suitable. One simple idea for this is to simply reset the rest shape after each click-and-drag, since each deformation step is then independent of the others, and one could switch between our method and others at each step. We have found this mode of interaction to be useful even when only using our method, as it leads to a simple sculpting-style interface, and we include some examples in the supplementary video.

ACKNOWLEDGMENTS

This work is funded in part by National Science Foundation (1910839). We especially thank George E. Brown for sharing the WRAPD implementation and the help with setting up experiments. We thank Jiayi Eris Zhang and Danny Kaufman for sharing the undeformed scene geometry; Lillie Kittredge and Huy Ha for proofreading; Rundi Wu for the help with rendering; all the artists for sharing the 2D and 3D models and anonymous reviewers for their helpful comments and suggestions.

REFERENCES

Marc Alexa. 2006. Mesh Editing Based on Discrete Laplace and Poisson Models. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) *(SIGGRAPH '06)*. Association for Computing Machinery, New York, NY, USA, 51–59.

Mirela Ben-Chen, Ofir Weber, and Craig Gotsman. 2009. Variational Harmonic Maps for Space Deformation. *ACM Trans. Graph.* 28, 3, Article 34 (Jul 2009), 11 pages.

Miklos Bergou, Max Wardetzky, David Harmon, Denis Zorin, and Eitan Grinspun. 2006. A Quadratic Bending Model for Inextensible Surfaces. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing* (Cagliari, Sardinia, Italy) *(SGP '06)*. Eurographics Association, Goslar, DEU, 227–230.

D.P. Bertsekas. 1996. *Constrained Optimization and Lagrange Multiplier Methods*. Athena Scientific, Nashua.

Sofien Bouaziz, Mario Deuss, Yuliy Schwartzburg, Thibaut Weise, and Mark Pauly. 2012. Shape-Up: Shaping Discrete Geometry with Projections. *Comput. Graph. Forum* 31, 5 (Aug 2012), 1657–1667.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.* 3, 1 (Jan 2011), 1–122.

Christopher Brandt and Klaus Hildebrandt. 2017. Compressed vibration modes of elastic bodies. *Computer Aided Geometric Design* 52-53 (2017), 297–312. Geometric Modeling and Processing 2017.

George E. Brown and Rahul Narain. 2021. WRAPD: Weighted Rotation-aware ADMM for Parameterization and Deformation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 40, 4 (8 2021).

Marie-Paule Cani and Alexis Angelidis. 2006. Towards Virtual Clay. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) *(SIGGRAPH '06)*. Association for Computing Machinery, New York, NY, USA, 67–83. https://doi.org/10.1145/1185657.1185676

Jiong Chen and Mathieu Desbrun. 2022. Go Green: General Regularized Green's Functions for Elasticity. In *ACM SIGGRAPH 2022 Conference Proceedings* (Vancouver, BC, Canada) *(SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 6, 8 pages.

Jiaxu Chen, Long Zhang, Xiaoxu Li, Bo Zhang, and Zhongfu Ye. 2017. Locally controlled as-rigid-as-possible deformation for 2D characters. *Computer Animation and Virtual Worlds* 28, 6 (2017), e1750. e1750 cav.1750.

Sabine Coquillart. 1990. Extended Free-Form Deformation: A Sculpturing Tool for 3D Geometric Modeling. *SIGGRAPH Comput. Graph.* 24, 4 (Sep 1990), 187–196.

Fernando De Goes and Doug L. James. 2017. Regularized Kelvinlets: Sculpting Brushes Based on Fundamental Solutions of Elasticity. *ACM Trans. Graph.* 36, 4, Article 40 (Jul 2017), 11 pages.

Fernando De Goes and Doug L. James. 2018. Dynamic Kelvinlets: Secondary Motions Based on Fundamental Solutions of Elastodynamics. *ACM Trans. Graph.* 37, 4, Article 81 (Jul 2018), 10 pages.

Fernando de Goes and Doug L. James. 2019. Sharp Kelvinlets: Elastic Deformations with Cusps and Localized Falloffs. In *Proceedings of the 2019 Digital Production Symposium* (Los Angeles, California) *(DigiPro '19)*. Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages.

Bailin Deng, Sofien Bouaziz, Mario Deuss, Juyong Zhang, Yuliy Schwartzburg, and Mark Pauly. 2013. Exploring Local Modifications for Constrained Meshes. *Computer Graphics Forum* 32, 2pt1 (2013), 11–20.

Jianqing Fan and Runze Li. 2001. Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties. *J. Amer. Statist. Assoc.* 96, 456 (2001), 1348–1360.

Lin Gao, Guoxin Zhang, and Yu-Kun Lai. 2012. Lp shape deformation. *Science China Information Sciences* 55 (05 2012).

Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. 2020. Local Deep Implicit Functions for 3D Shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4857–4866.

Fabian Hahn, Sebastian Martin, Bernhard Thomaszewski, Robert Sumner, Stelian Coros, and Markus Gross. 2012. Rig-Space Physics. *ACM Trans. Graph.* 31, 4, Article 72 (Jul 2012), 8 pages.

Josef Hoschek, Dieter Lasser, and Larry L. Schumaker. 1993. *Fundamentals of Computer Aided Geometric Design*. A. K. Peters, Ltd., USA.

Alec Jacobson et al. 2018a. gptoolbox: Geometry Processing Toolbox. http://github.com/alecjacobson/gptoolbox.

Alec Jacobson, Ilya Baran, Ladislav Kavan, Jovan Popović, and Olga Sorkine. 2012. Fast Automatic Skinning Transformations. *ACM Trans. Graph.* 31, 4, Article 77 (Jul 2012), 10 pages.

Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. 2011. Bounded Biharmonic Weights for Real-Time Deformation. *ACM Trans. Graph.* 30, 4, Article 78 (Jul 2011), 8 pages.

Alec Jacobson, Daniele Panozzo, et al. 2018b. libigl: A simple C++ geometry processing library. https://libigl.github.io/.

Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. 2007. Harmonic Coordinates for Character Articulation. *ACM Trans. Graph.* 26, 3 (Jul 2007), 71–es.

Ladislav Kavan and Olga Sorkine. 2012. Elasticity-Inspired Deformers for Character Articulation. *ACM Trans. Graph.* 31, 6, Article 196 (Nov 2012), 8 pages.

Youngihn Kho and Michael Garland. 2005. Sketching Mesh Deformations. *ACM Trans. Graph.* 24, 3 (Jul 2005), 934.

Yaron Lipman, David Levin, and Daniel Cohen-Or. 2008. Green Coordinates. *ACM Trans. Graph.* 27, 3 (Aug 2008), 1–10.

Qiong Luo, Bo Liu, Zhan-Guo Ma, and Hong-Bin Zhang. 2007. Mesh Editing in ROI with Dual Laplacian. In *Computer Graphics, Imaging and Visualisation (CGIV 2007)*. 195–199. https://doi.org/10.1109/CGIV.2007.57

N. Magnenat-Thalmann, R. Laperrière, and D. Thalmann. 1989. Joint-Dependent Local Deformations for Hand Animation and Object Grasping. In *Proceedings on Graphics Interface '88* (Edmonton, Alberta, Canada). Canadian Information Processing Society, CAN, 26–33.

G. Hosein Mohimani, Massoud Babaie-Zadeh, and Christian Jutten. 2007. Fast Sparse Representation Based on Smoothed l0 Norm. In *Independent Component Analysis and Signal Separation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 389–396.

Thomas Neumann, Kiran Varanasi, Stephan Wenger, Markus Wacker, Marcus Magnor, and Christian Theobalt. 2013. Sparse Localized Deformation Components. *ACM Trans. Graph.* 32, 6, Article 179 (Nov 2013), 10 pages.

Yue Peng, Bailin Deng, Juyong Zhang, Fanyu Geng, Wenjie Qin, and Ligang Liu. 2018. Anderson Acceleration for Geometry Optimization and Physics Simulation. *ACM Trans. Graph.* 37, 4, Article 42 (Jul 2018), 14 pages. https://doi.org/10.1145/3197517.3201290

Roi Poranne, Marco Tarini, Sandro Huber, Daniele Panozzo, and Olga Sorkine-Hornung. 2017. Autocuts: Simultaneous Distortion and Cut Optimization for UV Mapping. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH ASIA)* 36, 6 (2017).

Thomas W. Sederberg and Scott R. Parry. 1986. Free-Form Deformation of Solid Geometric Models. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '86)*. Association for Computing Machinery, New York, NY, USA, 151–160.

Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahar Z. Kovalsky, and Yaron Lipman. 2017. Geometric Optimization via Composite Majorization. *ACM Trans. Graph.* 36, 4, Article 38 (jul 2017), 11 pages. https://doi.org/10.1145/3072959.3073618

Karan Singh and Eugene Fiume. 1998. Wires: A Geometric Deformation Technique. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 405–414.

Breannan Smith, Fernando De Goes, and Theodore Kim. 2019. Analytic Eigensystems for Isotropic Distortion Energies. *ACM Trans. Graph.* 38, 1, Article 3 (Feb 2019), 15 pages. https://doi.org/10.1145/3241041

Olga Sorkine and Marc Alexa. 2007. As-Rigid-as-Possible Surface Modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (Barcelona, Spain) *(SGP '07)*. Eurographics Association, Goslar, DEU, 109–116.

Yu Wang, Alec Jacobson, Jernej Barbič, and Ladislav Kavan. 2015. Linear Subspace Design for Real-Time Shape Deformation. *ACM Trans. Graph.* 34, 4, Article 57 (Jul 2015), 11 pages. https://doi.org/10.1145/2766952

Jianhong Xiang, Hao Xiang, Linyu Wang, and Yu Zhong. 2022. Medical Image Reconstruction Method Based on Smooth L0 Norm. https://doi.org/10.21203/rs.3.rs-1535219/v1

Linlin Xu, Ruimin Wang, Juyong Zhang, Zhouwang Yang, Jiansong Deng, Falai Chen, and Ligang Liu. 2015. Survey on sparsity in geometric modeling and processing. *Graphical Models* 82 (2015), 160–180.

Cun-Hui Zhang. 2010. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics* 38, 2 (2010), 894 – 942. https://doi.org/10.1214/09-AOS729

Juyong Zhang, Yue Peng, Wenqing Ouyang, and Bailin Deng. 2019. Accelerating ADMM for Efficient Simulation and Optimization. *ACM Trans. Graph.* 38, 6, Article 163 (nov 2019), 21 pages. https://doi.org/10.1145/3355089.3356491

Jiayi Eris Zhang, Jèrèmie Dumas, Yun (Raymond) Fei, Alec Jacobson, Doug L. James, and Danny M. Kaufman. 2022. Progressive Simulation for Cloth Quasistatics. *ACM Trans. Graph.* 41, 6, Article 218 (2022).

Yufeng Zhu, Robert Bridson, and Danny M. Kaufman. 2018. Blended Cured Quasi-Newton for Distortion Optimization. *ACM Trans. on Graphics* (2018).

Johannes Zimmermann, Andrew Nealen, and Marc Alexa. 2007. SilSketch: Automated Sketch-Based Editing of Surface Meshes. In *EUROGRAPHICS Workshop on Sketch-Based Interfaces and Modeling*, Michiel van de Panne and Eric Saund (Eds.). The Eurographics Association. https://doi.org/10.2312/SBM/SBM07/023-030

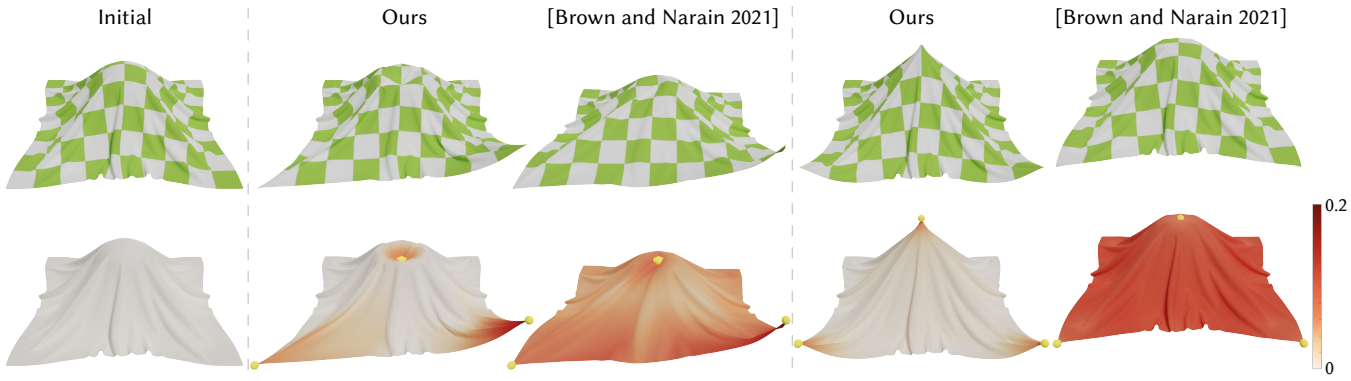

Figure 6: Our method enables the user to locally edit a simulated cloth in an interactive fashion, without the need of rerunning the simulation. In contrast, quasi-static deformations generated by directly moving the control points [Brown and Narain 2021] have global effects, easily deviating the edit from the initial shape in a distinct way. Bottom row: To show the extent to which the handles affect cloth vertex positions, we colormap the vertex displayment from the input cloth. The handles moved by the user are in yellow. (Undeformed geometry thanks to [Zhang et al. 2022])

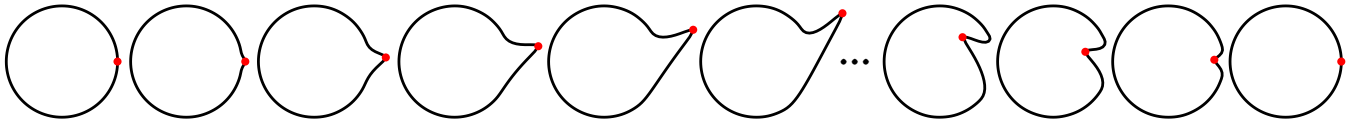

Figure 7: Our method also generalizes to 1D polyline editing. Our method enables the deformed shape to naturally return to the rest shape when they are close enough.

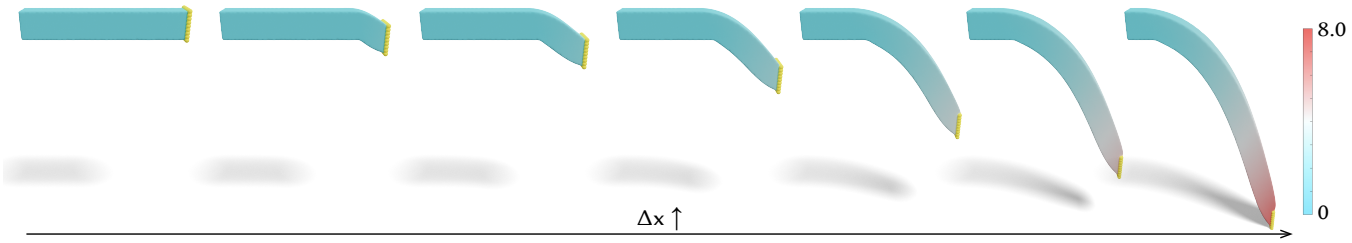

Figure 8: Our algorithm naturally enables an adaptive region of influence under different deformations. Only the vertices (yellow) on the rightmost end are selected as handles. From left to right, the offset Δx of the rightmost end is 0.0, 0.5, 1.0, 2.0, 4.0, 6.0 and 8.0 respectively.

Figure 9: We show how the ROI (yellow) changes when adjusting s or w, while keeping the rest the same. Here the handles are highlighted in red. In general, we use w to control the scale of the ROI, and set s to a small factor of the size of the shape and then leave it alone.

Figure 10: As-conformal-as-possible (ACAP) energy encourages conformality, thus better preserving the color texture by allowing local scaling, while the as-rigid-as-possible (ARAP) energy favors preserving the rigidity. Here the handle (in blue) is placed around the eye and an affine constraint is added to the red region.

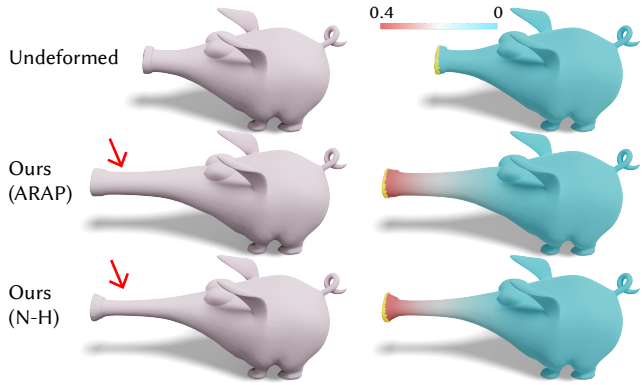

Figure 11: ARAP v.s. Neo-Hookean energy with SC-L1 loss. Only the vertices (yellow) on the nose of the pig are selected as handles. Note that Neo-Hookean version exhibits the volume-preservation property.

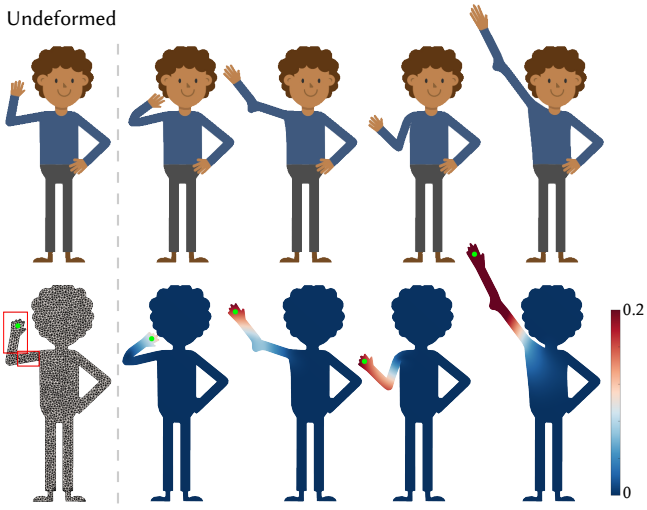

Figure 12: Our method supports adding both positional constraints (green) and affine constraints (red) on vertices. The region of influence in the deformed shape naturally adapts to the resulting deformation.

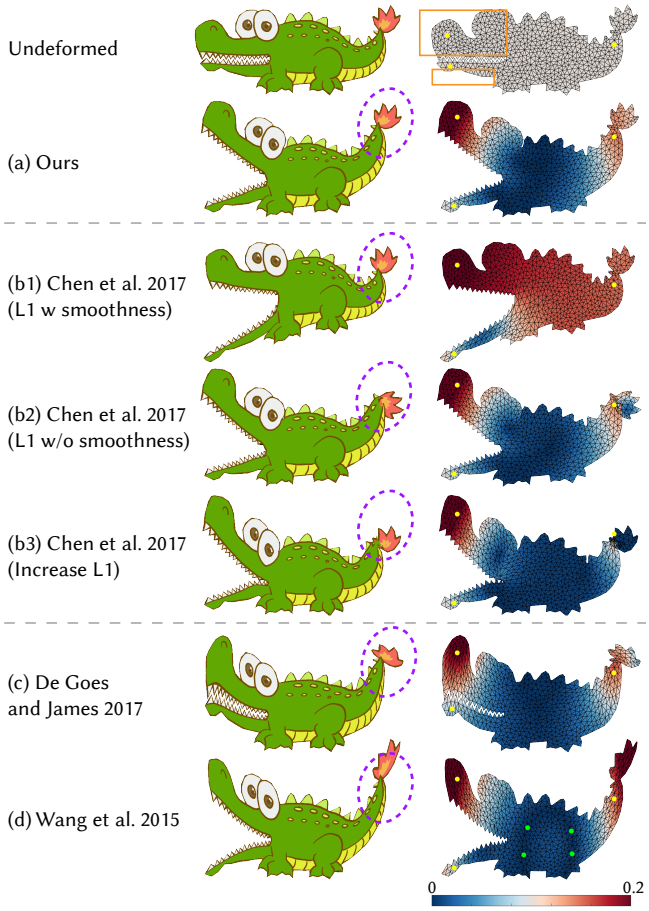

Figure 13: Our local deformation (a) is natural and shape-aware and supports an adaptive ROI without the need of additional handles. Here we highlight all the handles that have been moved in yellow, all the additional fixed handles in green and the affine constraint regions in orange. Other local editing alternatives either introduces artifacts (see the flame on the tail in b and d), is unaware of the geometry (see the teeth region in c) or requires additional handles setup (d). More specifically, [Chen et al. 2017] uses a group lasso penalty, which yields hard to avoid artifacts (b). They therefore mitigate it with a smoothing regularizer, at the expense of fidelity in rotations (b1). Removing the smoothing leads to other artifacts and still has global motion (b2). Further increasing the group lasso penalty removes the global motion but amplifies the artifacts (b3).

Supplementary Material for Local Deformation for Interactive Shape Editing

Honglin Chen
Columbia University
New York City, NY, USA
honglin.chen@columbia.edu

Changxi Zheng
Columbia University
New York City, NY, USA
cxz@cs.columbia.edu

Kevin Wampler
Adobe Research
Seattle, WA, USA
kwampler@adobe.com

ACM Reference Format:
Honglin Chen, Changxi Zheng, and Kevin Wampler. 2023. Supplementary Material for Local Deformation for Interactive Shape Editing. *ACM Trans. Graph.* 1, 1 (May 2023), 5 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 ADDITIONAL RESULTS

We report the runtime of our method and [Chen et al. 2017] on a 2D ARAP local energy and across different mesh resolutions and deformations in Table 1. We also provide the timings and experiment setup of all our 3D examples in Table 2 and the convergence plot in Fig. 1. In addition to the results in Sec. 5, we include an additional example with one or multiple control points in Fig. 2.

Table 1: Timings comparing our ARAP ADMM formulation in Sec.4.1 with standard local-global ARAP and [Chen et al. 2017] for moving the top-right arm of the octocat in Fig.3 of the main text, meshed at three different resolutions. Wall-clock time is measured by averaging over 100 random deformations and uses the (relatively tight) convergence criteria from [Chen et al. 2017]. All timings are generated by single-threaded C++ implementations, written as equivalently as possible between the methods. For an apples-to-apples comparison these timings include the E_{smooth} energy from [Chen et al. 2017] in the SC-L1 loss and ARAP elastic energies.

	ours (w. SC-L1 loss)	ARAP	[Chen et al. 2017]
# triangles	small deformation		
200	0.00087s	0.00015s	0.11s
1000	0.0055s	0.0010s	0.59s
10000	0.058s	0.017s	2.05s
# triangles	large deformation		
200	0.0053s	0.00032s	0.022s
1000	0.0078s	0.0018s	0.17s
10000	0.095s	0.027s	2.03s

Authors' addresses: Honglin Chen, Columbia University, New York City, NY, USA, honglin.chen@columbia.edu; Changxi Zheng, Columbia University, New York City, NY, USA, cxz@cs.columbia.edu; Kevin Wampler, Adobe Research, Seattle, WA, USA, kwampler@adobe.com.

0730-0301/2023/5-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

2 DERIVATION OF PROXIMAL OPERATOR FOR SC-L1 LOSS

For simplicity, we start the derivation with scalar function $f(v) = |v|_{\text{SC-L1}}$. However, the following derivation can be easily extended to vector function $f(\mathbf{x}) = \|\mathbf{x}\|_{\text{SC-L1}}$ by replacing the $|\cdot|$ with $\|\cdot\|$.

For scalar v,

$$f(v) = \begin{cases} |v| - \frac{1}{2s}v^2 & |v| < s \\ \frac{1}{2}s & |v| \geq s \end{cases} \tag{1}$$

The gradient of $f(v)$ is

$$f'(v) = \begin{cases} \operatorname{sgn}(v) - \frac{1}{s}v & |v| < s \\ 0 & |v| \geq s \end{cases} \tag{2}$$

The proximal operator of $f(v)$ can be evaluated as

$$\underset{v}{\operatorname{argmin}} \quad f(v) + \frac{\rho}{2}(v - x)^2 \tag{3}$$

The minimum can be found by setting the derivative to be 0, thus we have

$$0 = f'(v) + \rho(v - x) \tag{4}$$

$$0 = \rho(v - x) + \begin{cases} \operatorname{sgn}(v) - \frac{1}{s}v & |v| < s \\ 0 & |v| \geq s \end{cases} \tag{5}$$

$$0 = \begin{cases} \left(\rho - \frac{1}{s}\right)v + \operatorname{sgn}(v) - \rho x & |v| < s \\ \rho(v - x) & |v| \geq s \end{cases} \tag{6}$$

The location of the minumum v depends on whether or not $\rho > \frac{1}{s}$.

Case 1: $\rho \leq \frac{1}{s}$

When $\rho \leq \frac{1}{s}$, then the proximal operator may have up to two local minima—one at the origin and one in the flat region where $|x| \geq s$. We arbitrarily prefer the minima at the origin if it exists, giving:

$$\mathcal{S}(x) = \begin{cases} 0 & |x| < s \\ x & |x| \geq s \end{cases} \tag{7}$$

Case 2: $\rho \geq \frac{1}{s}$

When $\rho \geq \frac{1}{s}$, then there is only one (global) minimum, so the optimal v is:

$$\mathcal{S}(x) = \begin{cases} 0 & |x| \leq \frac{1}{\rho} \\ \frac{\rho s x - s \cdot \operatorname{sgn}(x)}{\rho s - 1} & \frac{1}{\rho} < |x| < s \\ x & |x| \geq s \end{cases} \tag{8}$$

for a vector x this can be further simplied to be

$$\mathcal{S}(x) = \begin{cases} \left(\frac{\rho s - s/\|x\|_2}{\rho s - 1}\right)_+ x, & \text{if } \|x\|_2 \leq s \\ x, & \text{otherwise} \end{cases} \tag{9}$$

Table 2: Performance of and parameters for our algorithm on all the 3D examples. All wall-clock timings are reported in seconds, physical parameters are reported with appropriate units. Timing experiments are measured with 8-thread parallelization and PARDISO support [Alappat et al. 2020; Bollhöfer et al. 2019, 2020]. |V| is the number of vertices, |T| is the number of elements. ρ is the ADMM penalty parameter and s is the SC-L1 loss threshold. Local 1, Local 2 and Global is the total runtime for the local step 1, the local step 2 and the global step respectively. Total is the total runtime for our ADMM algorithm. Residual is the final residuals of our energy, defined as $|E_{\text{curr}} - E_{\text{prev}}|/(E_{\text{curr}} + 1.0)$, where E is the total energy and adding 1.0 is to avoid dividing by 0. We use $w = 1e^4$ as the weight of the SC-L1 loss for all the 3D examples. The Young's modulus is $1e^7$ Pa and the poisson's ratio is 0.45 for all our Neo-Hookean examples. We use different ρ for ARAP and Neo-Hookean energy to handle the different magnitude of the energies due to the physical parameters in Neo-Hookean energy. In general we set s to be 0.01 times the max extent of the mesh's bounding box, but other values for s also work well. For interactive frame rates, we generate all the results using a fixed 500 ADMM iterations or when its objective residual reaches $1e^{-11}$. As showen by Fig. 1, all of examples converged within 500 iterations. For the bar example in Fig.7 and Fig.13 of the main text, we report its average runtime.

Example	|V|	|T|	Material	ρ	s	Residual	Local 1(s)	Local 2(s)	Global(s)	Total(s)
Cloth 1(Fig. 6)	14,641	28,800	Cloth	$5e^3$	0.4	$3.93e^{-5}$	1.58	0.04	1.76	3.38
Cloth 2(Fig. 6)	14,641	28,800	Cloth	$5e^3$	0.4	$2.09e^{-5}$	1.52	0.03	1.74	3.29
Cloth 1(Suppl Fig.3)	14,641	28,800	Cloth	$3e^6$	0.01	$5.04e^{-5}$	1.95	0.05	1.88	3.89
Cloth 2(Suppl Fig.3)	14,641	28,800	Cloth	$3e^6$	0.01	$3.54e^{-7}$	1.60	0.06	1.71	3.37
Cloth(Fig. 1)	46,977	92,928	Cloth	$3e^6$	0.01	$3.26e^{-4}$	6.06	0.10	5.41	11.57
Pig(Fig. 1)	10,416	43,749	NH	$5e^5$	0.01	$3.58e^{-3}$	5.85	0.03	1.61	7.49
Pig 1(Fig. 11)	10,416	43,749	ARAP	$1e^{-1}$	0.01	$3.84e^{-7}$	2.40	0.03	1.59	4.03
Pig 2(Fig. 11)	10,416	43,749	NH	$5e^5$	0.01	$3.94e^{-5}$	5.56	0.03	1.58	7.18
Cat(Fig. 1)	9,078	36,115	NH	$5e^3$	0.01	$7.81e^{-6}$	4.24	0.02	1.30	5.56
Robot(Fig. 1)	6,601	26,024	NH	$1e^3$	0.04	$6.88e^{-4}$	3.91	0.04	0.86	4.81
Cube(Fig. 1)	4,913	20,480	NH	$1e^5$	1.0	$1.14e^{-5}$	2.77	0.01	2.17	4.95
Apple(Fig. 1)	1,088	4,072	NH	$1e^3$	0.01	$1.69e^{-3}$	0.60	0.01	0.52	1.13
Table(Fig. 1)	6,483	21,864	ARAP	$2e^{-4}$	0.01	$3.41e^{-11}$	1.46	0.04	1.03	2.53
Bar(Fig. 8)	7,011	34,557	ARAP	$1e^{-1}$	0.1	$1.28e^{-4}$	2.89	0.08	2.69	5.65
Bar(Fig. 5)	7,011	34,557	NH	$5e^5$	0.1	$6.66e^{-5}$	5.74	0.03	1.63	7.41

3 DERIVATION OF LOCAL ARAP ENERGY

With as-rigid-as-possible (ARAP) energy [Sorkine and Alexa 2007] as our elastic energy, the total energy for our local deformation is as follows:

$$\underset{\mathbf{V},\{\mathbf{R}_i\}}{\text{minimize}} \quad \sum_{i\in V}\sum_{j\in\mathcal{N}(j)} \underbrace{\frac{w_{ij}}{2}\|\mathbf{R}_i\widetilde{\mathbf{d}}_{ij} - \mathbf{d}_{ij}\|_2^2}_{\text{ARAP}} + \underbrace{wa_i\|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}}_{\text{Localness}}, \tag{10}$$

where $\mathbf{R}_i$ is a $d \times d$ rotation matrix, w_{ij} is the cotangent weight, $\mathbf{d}_{ij} = [\mathbf{V}_j - \mathbf{V}_i]^\top$ and $\widetilde{\mathbf{d}}_{ij} = [\widetilde{\mathbf{V}}_j - \widetilde{\mathbf{V}}_i]^\top$ are the edge vectors between vertices i, j at the deformed and rest states respectively. $\mathcal{N}(j)$ denotes the neighboring edges of the i-th vertex in the style of "spokes-and-rims". Here we use $\mathbf{R}_i$ to denote $\mathbf{X}_i$, since we drive the deformation gradient towards a rotation matrix in ARAP energy.

Our local ARAP deformation energy (Eq. 10) can be rewritten as

$$\underset{\mathbf{V},\{\mathbf{R}_i\}}{\text{minimize}} \quad \sum_{i\in V} \frac{1}{2}\|\mathbf{R}_i\mathbf{D}_i - \widetilde{\mathbf{D}}_i\|_{\mathbf{W}_i}^2 + wa_i\|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}, \tag{11}$$

where W_i is a $|\mathcal{N}(i)| \times |\mathcal{N}(i)|$ diagonal matrix of cotangent weights, $\widetilde{\mathbf{D}}_i$ and $\mathbf{D}_i$ are $3 \times |\mathcal{N}(i)|$ matrices of "spokes and rims" edge vectors of the i-th vertex at the rest and deformed states respectively. $\|\mathbf{X}\|_{\mathbf{W}_i}^2$ denotes $\text{Tr}(\mathbf{X}^\top\mathbf{W}_i\mathbf{X})$.

By setting $\mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i$, we can further rewrite Eq. 11 as

$$\underset{\mathbf{V},\{\mathbf{R}_i\},\mathbf{Z}}{\text{minimize}} \quad \sum_{i\in V} \frac{1}{2}\|\mathbf{R}_i\mathbf{D}_i - \widetilde{\mathbf{D}}_i\|_{\mathbf{W}_i}^2 + wa_i\|\mathbf{Z}_i\|_{\text{SC-L1}}, \tag{12a}$$

$$\text{s.t.} \quad \mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i, \quad \forall i \tag{12b}$$

The ADMM update for the above minimization problem is as follows:

$$\mathbf{R}_i^{k+1} \leftarrow \underset{\mathbf{R}_i\in\text{SO(3)}}{\arg\min} \ \frac{1}{2}\|\mathbf{R}_i\mathbf{D}_i - \widetilde{\mathbf{D}}_i\|_{\mathbf{W}_i}^2 \tag{13a}$$

$$\mathbf{Z}_i^{k+1} \leftarrow \underset{\mathbf{Z}_i}{\arg\min} \ wa_i\|\mathbf{Z}_i\|_{\text{SC-L1}} + \frac{\rho}{2}\|\mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i + \mathbf{U}_i^k\|_2^2 \tag{13b}$$

$$\mathbf{V}^{k+1} \leftarrow \underset{\mathbf{V}}{\arg\min} \ \mathbf{W}(\mathbf{V}^\top\mathbf{L}\mathbf{V} - \mathbf{B}^\top\mathbf{V}) + \frac{\rho}{2}\|\mathbf{V} - \widetilde{\mathbf{V}} - \mathbf{Z}^k + \mathbf{U}^k\|_2^2 \tag{13c}$$

$$\mathbf{U}_i^{k+1} \leftarrow \mathbf{U}_i^k + \mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i^{k+1} \tag{13d}$$

Here ρ is a fixed penalty parameter.

Updating R

Local step 1 (Eq. 13a) is an instance of the Orthogonal Procrustes problem, which can be solved in the same way as the rotation fitting

ALGORITHM 1: Three-block ADMM for ARAP energy

Input: A triangle or tetrahedral mesh $\widetilde{\mathbf{V}}, \mathbf{T}$
Output: Deformed vertex positions $\mathbf{V}$
$\mathbf{V} \leftarrow \widetilde{\mathbf{V}}$
while *not converged* **do**
  $\mathbf{R}_i \leftarrow local\_step\_X(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 1
  $\mathbf{Z}_i \leftarrow local\_step\_Z(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 2
  $\mathbf{V} \leftarrow global\_step(\mathbf{X}_i, \mathbf{Z}_i, \widetilde{\mathbf{V}})$ ▷ global step
  $\mathbf{U}_i \leftarrow dual\_update(\mathbf{Z}_i, \mathbf{V}_i, \widetilde{\mathbf{V}}_i)$ ▷ dual update 1
end

step in [Sorkine and Alexa 2007]:

$$\mathbf{R}_i^{k+1} \leftarrow \underset{\mathbf{R}_i \in SO(3)}{\arg\max}\ \mathrm{Tr}\left(\mathbf{R}_i \mathbf{M}_i\right) \tag{14}$$

$$\mathbf{M}_i = \mathbf{D}_i \widetilde{\mathbf{D}}_i^{\top} \tag{15}$$

One can derive the optimal $\mathbf{R}_i$ from the singular value decomposition of $\mathbf{M}_i = \mathcal{U}_i \Sigma_i \mathcal{V}_i^{\top}$:

$$\mathbf{R}_i^{k+1} \leftarrow \mathcal{V}_i \mathcal{U}_i^{\top} \tag{16}$$

Updating Z

Local step 2 (Eq. 13b) can be solved with a shrinkage step (see derivation in App. 2):

$$\mathbf{Z}_i^{k+1} \leftarrow S_{wa_i}^k \left(\mathbf{V}_i - \widetilde{\mathbf{V}}_i + \mathbf{U}_i\right) \tag{17}$$

$$\mathcal{S}_{wa_i}(\mathbf{x}) = \begin{cases} \left(\frac{\rho s - wa_i s/\|\mathbf{x}\|_2}{\rho s - wa_i}\right)_+ \mathbf{x}, & \text{if } \|\mathbf{x}\|_2 \le s \\ \mathbf{x}, & \text{otherwise} \end{cases} \tag{18}$$

where ρ is set to satisfy $\rho > \frac{\max(wa_i)}{s}$ (see App. 2).

Updating V

The global step (Eq. 13c) can be achieved by solving a linear system:

$$(\mathbf{L} + \rho \mathbf{I})\mathbf{V} = \mathbf{B} + \rho(\widetilde{\mathbf{V}} + \mathbf{Z}^k - \mathbf{U}^k), \tag{19}$$

which can be precomputed for fixed ρ.

4 DERIVATION OF LOCAL NEO-HOOKEAN ENERGY

Our local deformation scheme can be further extended to physics-based elasticity energies, e.g., Neo-Hookean energy. Using the Neo-Hookean energy as our elasticity energy, the optimization problem can be written as follows:

$$\underset{\mathbf{V},\{\mathbf{X}_j\},\{\mathbf{Z}_i\}}{\text{minimize}} \quad \sum_{j\in T} \underbrace{E_{\text{nh}}(\mathbf{X}_j)}_{\text{Neo-Hookean}} + \sum_{i\in V} \underbrace{wa_i \|\mathbf{V}_i - \widetilde{\mathbf{V}}_i\|_{\text{SC-L1}}}_{\text{Localness}}, \tag{20a}$$

$$\text{s.t.} \quad \mathbf{X}_j = \mathrm{sym}(\mathbf{D}_j \mathbf{V}), \forall j, \tag{20b}$$

where T denotes all the elements.

By introducing $\mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i$, Eq. 20 can be further rewritten into:

$$\underset{\mathbf{V},\{\mathbf{X}_j\},\{\mathbf{Z}_i\}}{\text{minimize}} \quad \sum_{j\in T} E_{\text{nh}}(\mathbf{X}_j) + \sum_{i\in V} wa_i \|\mathbf{Z}_i\|_{\text{SC-L1}}, \tag{21a}$$

$$\text{s.t.} \quad \mathbf{X}_j = \mathrm{sym}(\mathbf{D}_j \mathbf{V}), \forall j, \tag{21b}$$

$$\mathbf{Z}_i = \mathbf{V}_i - \widetilde{\mathbf{V}}_i, \quad \forall i \tag{21c}$$

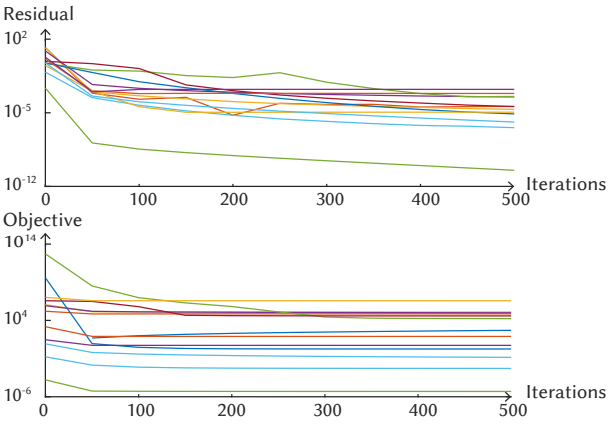

Figure 1: Convergence plot. We demonstrate the convergence of our ADMM by plotting residuals ($|E_{\text{curr}} - E_{\text{prev}}|/(E_{\text{curr}}+1.0)$) and objective values of all our examples in Table 2 in the main text over the ADMM iterations. As can be seen from the plot, all our 3D examples converge within 500 ADMM iterations.

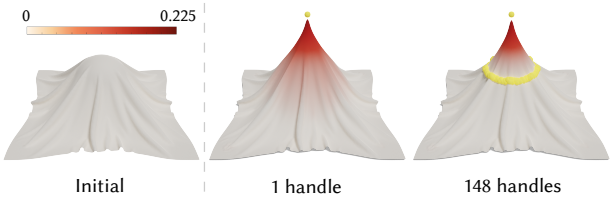

Figure 2: Our local deformer allows the user to use just a single control point, without the need to "pin down" the mesh. It also supports adding more control points to achieve more detailed control.

The augmented Lagrangian for this problem (Eq. 21) is:

$$L(\mathbf{V}; \{\mathbf{X}_i\}, \{\mathbf{W}_i\}; \{\mathbf{Z}_i\}, \{\mathbf{U}_i\}) = \sum_{j\in T} E_{\text{nh}}(\mathbf{X}_j) + \sum_{i\in V} wa_i \|\mathbf{Z}_i\|_{\text{SC-L1}} + \sum_{j\in T} \frac{\gamma}{2} \|\mathrm{sym}(\mathbf{D}_j \mathbf{V}) - \mathbf{X}_j + \mathbf{W}_j\|_2^2 + \sum_{i\in V} \frac{\rho}{2} \|\mathbf{V} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i + \mathbf{U}_i\|_2^2. \tag{22}$$

Therefore, the ADMM update for the above minimization problem is as follows:

$$\mathbf{X}_j^{k+1} \leftarrow \underset{\mathbf{X}_j}{\arg\min}\ E_{\text{nh}}(\mathbf{X}_j) + \frac{\gamma}{2} \|\mathrm{sym}(\mathbf{D}_j \mathbf{V}) - \mathbf{X}_j + \mathbf{W}_j\|_2^2 \tag{23a}$$

$$\mathbf{Z}_i^{k+1} \leftarrow \underset{\mathbf{Z}_i}{\arg\min}\ wa_i \|\mathbf{Z}_i\|_{\text{SC-L1}} + \frac{\rho}{2} \|\mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i + \mathbf{U}_i^k\|_2^2 \tag{23b}$$

$$\mathbf{V}^{k+1} \leftarrow \arg\min_{\mathbf{V}}\ \sum_{j\in T} \frac{\gamma}{2} \|\mathrm{sym}(\mathbf{D}_j \mathbf{V}) - \mathbf{X}_j + \mathbf{W}_j\|_2^2 + \sum_{i\in V} \frac{\rho}{2} \|\mathbf{V} - \widetilde{\mathbf{V}} - \mathbf{Z}^k + \mathbf{U}^k\|_2^2 \tag{23c}$$

$$\mathbf{W}_j^{k+1} \leftarrow \mathbf{W}_j^k + \mathrm{sym}(\mathbf{D}_j \mathbf{V}) - \mathbf{X}_j \tag{23d}$$

$$\mathbf{U}_i^{k+1} \leftarrow \mathbf{U}_i^k + \mathbf{V}_i^{k+1} - \widetilde{\mathbf{V}}_i - \mathbf{Z}_i^{k+1} \tag{23e}$$

Here ρ and γ are fixed penalty parameters.

Compared with the ADMM steps for the local ARAP energy (Eq. 13), the local step 2 ($\arg\min_{\mathbf{Z}_i}$) and the global step ($\arg\min_{\mathbf{V}}$)

ALGORITHM 2: Three-block ADMM for Neo-Hookean energy

Input: A triangle or tetrahedral mesh $\widetilde{\mathbf{V}}, \mathbf{T}$
Output: Deformed vertex positions $\mathbf{V}$
$\mathbf{V} \leftarrow \widetilde{\mathbf{V}}$
while *not converged* **do**
  $\mathbf{X}_i \leftarrow \mathit{local\_step\_X}(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 1
  $\mathbf{Z}_i \leftarrow \mathit{local\_step\_Z}(\mathbf{V}, \widetilde{\mathbf{V}})$ ▷ local step 2
  $\mathbf{V} \leftarrow \mathit{global\_step}(\mathbf{X}_i, \mathbf{Z}_i, \widetilde{\mathbf{V}})$ ▷ global step
  $\mathrm{U}_i \leftarrow \mathit{dual\_update}(\mathrm{Z}_i, \mathrm{V}_i, \widetilde{\mathrm{V}}_i)$ ▷ dual update 1
  $\mathbf{W}_i \leftarrow \mathit{dual\_update}(\mathbf{X}_i, \mathrm{V}_i, \widetilde{\mathrm{V}}_i)$ ▷ dual update 2
end

for the local Neo-Hookean energy can be solved in the same way as the local ARAP energy (see App. 3).

Updating X

Local step 1 (Eq. 23a) can be solved by performing the energy minimization on the singular values of $\mathrm{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$. This is a proximal operator of E_{nh} at rotation-invariant $\mathrm{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$.

Let us denote the proximal operator of E_{nh} and the singular value decomposition of $\mathrm{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$ as:

$$\mathrm{prox}_{E_{\mathrm{nh}}}(\mathbf{X}_j) = E_{\mathrm{nh}}(\mathbf{X}_j) + \frac{\gamma}{2}\|\mathrm{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j - \mathbf{X}_j\|_2^2 \tag{24}$$

$$\mathrm{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j = \mathcal{U}_j\Sigma_j\mathcal{V}_j^\top \tag{25}$$

As shown by [Brown and Narain 2021], we can compute the optimal $\mathbf{X}_j$ as:

$$\Sigma_j^{k+1} \leftarrow \arg\min_{\Sigma_j} \ \mathrm{prox}_{E_{\mathrm{nh}}}(\mathcal{U}_j\Sigma_j\mathcal{V}_j^\top) \tag{26}$$

$$\mathbf{X}_j \leftarrow \mathcal{U}_j\Sigma_j^{k+1}\mathcal{V}_j^\top \tag{27}$$

Specifically, one can compute the SVD of $\mathrm{sym}(\mathbf{D}_j\mathbf{V}) + \mathbf{W}_j$ and perform the minimization of $\mathrm{prox}_{E_{\mathrm{nh}}}$ only on its singular values, while keeping singular vectors unchanged. The above optimization of the singular values Σ_j^{k+1} can be performed using an L-BFGS solver.

5 OPTIMAL SCALING s_k FOR E_{ACAP}

Adopting the as-conformal-as-possible energy E_{ACAP} as our elastic energy requires solving an instance of isotropic orthogonal Procrustes problem in the local step 1 (updating $\mathbf{R}_k$). Following the derivation in [Schönemann and Carroll 1970], one can compute its analytical solution as follows: The optimal rotation $\mathbf{R}_k$ can be computed the same way as the local step 1 (updating $\mathbf{R}_k$) for the ARAP energy, and the optimal scaling s_k can be computed analytically as

$$s_k = \frac{\mathrm{Tr}\left(\mathbf{W}_k\widetilde{\mathbf{D}}_k^\top\mathbf{R}_k\mathbf{D}_k\right)}{\mathrm{Tr}\left(\mathbf{W}_k\mathbf{D}_k^\top\mathbf{D}_k\right)} \tag{28}$$

When assembling the matrices $\mathbf{B}$ for the global step, one needs to replace $\mathbf{R}_k$ with $s_k\mathbf{R}_k$.

6 COMPARISON WITH OTHER SC-L1 LOSS ALTERNATIVES

Our SC-L1 loss is designed with the following desirable properties in mind:

(1) Behaves like a ℓ_1 loss near zero.
(2) Is constant far from zero.
(3) Admits an efficient proximal shrinkage operator free of local minima.

We will consider each of these properties in turn, examining the motivation and comparing our loss to other alternatives.

(1) Behaves like a ℓ_1 loss near zero. This property is what gives our method the ability to localize deformation. The non-differentiable ℓ_1-like "vertex" of the SC-L1 loss at the origin is critical in this, and we refer the reader to [Hastie et al. 2015] Sec. 2.2 for more information as to why. Using a loss which is smooth at the origin will invariably result in small amounts of global motion, with the degree of global motion being roughly related to the magnitude of the function's Hessian at the origin, so the smoother the loss, the more global motion should be expected. Examples of such smooth losses are the smoothed-ℓ_1 or Huber loss [Huber 1964] or smoothed-ℓ_0 losses [Mohimani et al. 2007; Xiang et al. 2022], and we avoid these losses due to their inability to eliminate global deformations.

(2) Constant far from zero. This property has the effect of disabling the locality-inducing force for parts of a deformed shape which have moved beyond a certain radius, and serves to remove artifacts resulting from parts of a deformed shape being pulled toward a rest position which is now far away. Fig.3(b) and Fig.10(b) in the main text illustrate the types of artifacts which can be observed using a loss without this property.

(3) Efficient proximal shrinkage operator free of local minima. It is convenient for the loss' proximal shrinkage operator to be uniquely defined. Recalling that this operator is defined as the minimizer(s) of Eq. 3, the condition that it be uniquely defined is equivalent to requiring that Eq. 3 has a unique minimum. A sufficient condition for this is that $f(v) + \frac{\rho}{2}(v-x)^2$ is strictly convex, or equivalently that $\rho > -\frac{d^2}{dv^2}f(v)$. We thus prefer a loss function which has a bounded negative second derivative, and the lower this bound the lower the value of ρ which can be used to convexify it. Within distance s from the origin (excluding the origin itself), our loss has a constant negative value for $\frac{d^2}{dv^2}f(v)$, making it in some sense the minimal function satisfying the three desired properties. Other losses are of course possible, but would require a higher value of ρ for the same radius of effect.

Alternative losses. A natural question is whether there might be other locality-inducing losses which could be used in place of our SC-L1 loss, perhaps even a generic spline-based loss which could be shaped in detail. Although we have not directly explored this, we expect that there would be many possible options provided they exactly or sufficiently approximately satisfy the three desired properties outlined above. That said, in our experiments we have found that it's normally useful to set s to be relatively small to limit artifacts. In this situation the specific shape of the loss is not particularly important since it only has an impact within a visually small region, and our simple formulation seems to work well.

REFERENCES

Christie Alappat, Achim Basermann, Alan R. Bishop, Holger Fehske, Georg Hager, Olaf Schenk, Jonas Thies, and Gerhard Wellein. 2020. A Recursive Algebraic Coloring Technique for Hardware-Efficient Symmetric Sparse Matrix-Vector Multiplication.

ACM Trans. Parallel Comput. 7, 3, Article 19 (June 2020), 37 pages. https://doi.org/10.1145/3399732

Matthias Bollhöfer, Aryan Eftekhari, Simon Scheidegger, and Olaf Schenk. 2019. Large-scale Sparse Inverse Covariance Matrix Estimation. *SIAM Journal on Scientific Computing* 41, 1 (2019), A380–A401. https://doi.org/10.1137/17M1147615

Matthias Bollhöfer, Olaf Schenk, Radim Janalik, Steve Hamm, and Kiran Gullapalli. 2020. State-of-the-Art Sparse Direct Solvers. (2020), 3–33. https://doi.org/10.1007/978-3-030-43736-7_1

George E. Brown and Rahul Narain. 2021. WRAPD: Weighted Rotation-aware ADMM for Parameterization and Deformation. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 40, 4 (8 2021).

Jiaxu Chen, Long Zhang, Xiaoxu Li, Bo Zhang, and Zhongfu Ye. 2017. Locally controlled as-rigid-as-possible deformation for 2D characters. *Computer Animation and Virtual Worlds* 28, 6 (2017), e1750. e1750 cav.1750.

Trevor Hastie, Robert Tibshirani, and Martin Wainwright. 2015. *Statistical Learning with Sparsity: The Lasso and Generalizations.* Chapman & HallCRC.

Peter J. Huber. 1964. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics* 35, 1 (1964), 73–101. http://www.jstor.org/stable/2238020

G. Hosein Mohimani, Massoud Babaie-Zadeh, and Christian Jutten. 2007. Fast Sparse Representation Based on Smoothed $\ell 0$ Norm. In *Independent Component Analysis and Signal Separation.* Springer Berlin Heidelberg, Berlin, Heidelberg, 389–396.

P. H. Schönemann and Robert M. Carroll. 1970. Fitting one matrix to another under choice of a central dilation and a rigid motion. *Psychometrika* 35 (1970), 245–255.

Olga Sorkine and Marc Alexa. 2007. As-Rigid-as-Possible Surface Modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (Barcelona, Spain) *(SGP '07).* Eurographics Association, Goslar, DEU, 109–116.

Jianhong Xiang, Hao Xiang, Linyu Wang, and Yu Zhong. 2022. Medical Image Reconstruction Method Based on Smooth L0 Norm. https://doi.org/10.21203/rs.3.rs-1535219/v1