

Adaptive Block Coordinate Descent for Distortion Optimization

Alexander Naitat¹, Yufeng Zhu² and Yehoshua Y. Zeevi¹

¹Viterbi Faculty of Electrical Engineering, Technion - Israel Institute of Technology

²Facebook Reality Labs

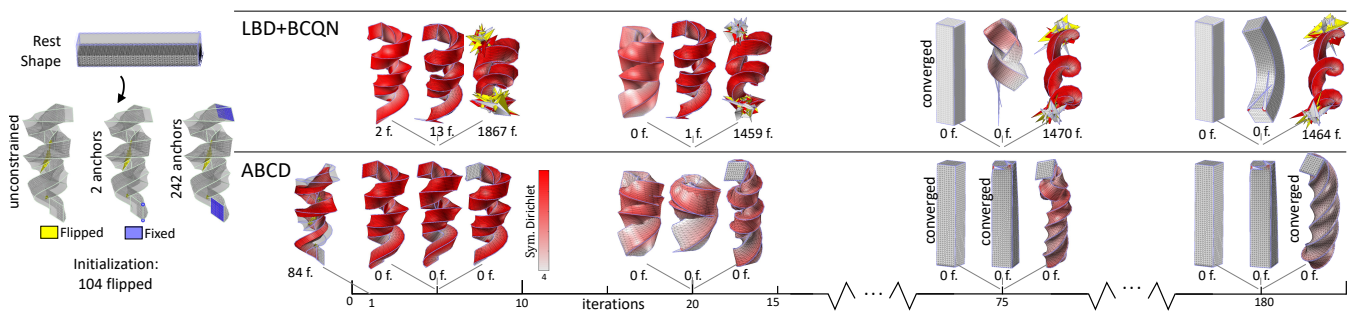


Figure 1: Comparison of the alternating combination of LBD and BCQN (top) with ABCD(BCQN) (bottom) in 3D. The same initialization is used for the three shown trials: unconstrained problem, two fixed anchors, two fixed endpoints (the second problem of Fig. 18). We report the number of inverted elements below each snapshot.

Abstract

We present a new algorithm for optimizing geometric energies and computing positively-oriented simplicial mappings. Our major improvements over the state-of-the-art are: (i) introduction of new energies for repairing inverted and collapsed simplices; (ii) adaptive partitioning of vertices into coordinate blocks with the blended local-global strategy for more efficient optimization; (iii) introduction of the displacement norm for improving convergence criteria and for controlling block partitioning. Together these improvements form the basis for the Adaptive Block Coordinate Descent (ABCD) algorithm aimed at robust geometric optimization. ABCD achieves state-of-the-art results in distortion minimization, even under hard positional constraints and highly distorted invalid initializations that contain thousands of collapsed and inverted elements. Starting with an invalid non-injective initial map, ABCD behaves as a modified block coordinate descent up to the point where the current mapping is cleared of invalid simplices. Then, the algorithm converges rapidly into the chosen iterative solver. Our method is very general, fast-converging and easily parallelizable. We show over a wide range of 2D and 3D problems that our algorithm is more robust than existing techniques for locally injective mapping.

Keywords: geometric optimization, inversion-free mapping, deformation, mesh parametrization, tetrahedral mesh.

1. Introduction

Computing injective mappings with low distortions on triangulated domains is a fundamental problem in computer graphics, geometrical modeling, and physical simulations. This problem often results in non-convex and non-linear optimization of geometric energies defined in a finite element manner on triangular and tetrahedral meshes.

The existing solutions to the above problem typically fall into two major categories: (1) *map fixers*: algorithms focused on the

injectivity of maps; (2) *core-solvers*: iterative descent algorithms focused on minimizing rotation-invariant energies. Map fixers are aimed at repairing simplices that have folded over (*foldovers*) and changed their orientation. The goal of map fixers is to repair all foldovers of a non-injective map and restraining its geometric distortions into a finite range. Core solvers start from a one-to-one initialization and ensure that optimization results remain one-to-one, at least locally. Efficient designs of map fixers and core-solvers have been extensively studied by the computer graphics community. But due to their inherent differences the treatment re-




Figure 2: Free boundary parametrization with a noisy initialization (intentionally distorted Tutte map without flips). The 1-ring GD denotes the gradient descent in blocks of single vertices; each block receives a single GD iteration per cycle depicted in the plot.

mained separate: one would either provide a framework to solve the foldover problem or the energy minimization problem. Nevertheless, both aspects of the problem play crucial role in finding a good solution to a given geometric optimization. In this paper, we propose the Adaptive Block Coordinate Descent (ABCD) algorithm that successfully tackles both problems within the same routine. Our algorithm achieves state-of-the-art results in distortion minimization even with highly distorted invalid initializations that contain thousands of degenerate and inverted elements. Consequently, compared with core-solvers, our algorithm is much more robust. At the same time, compared with recent map fixer methods, our algorithm achieves superior results, making some previously intractable problems applicable. The core contributions of our algorithm are:

1. Introduction of **new distortion measures** designed for repairing inverted and degenerate simplices, as well as **modification of classical energies** to allow an efficient integration of these energies into the proposed framework.
2. An **adaptive partitioning** of vertices into blocks of varying sizes according to the geometric configuration of the problem.
3. An **adaptively blended local-global strategy** providing robustness to highly distorted elements, while allowing employment of state-of-the-art solvers for inducing global deformations.
4. A **cured alternating optimization strategy** designed to best exploit contributions 1-3 and provide enough flexibility for our algorithm to be combined with other popular methods. This stage include number of novel sub-steps, such as **an enhanced line search filtering**.
5. Introduction of **displacement norm** for improving convergence criteria and for controlling the local-global blending.

2. Problem Formulation and Overview

We compute inversion-free simplicial maps that induce low distortions. Let $f: M \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a piecewise affine mapping defined on a mesh M with a m -dimensional simplex set \mathbf{C} and a vertex set \mathbf{V} , and let $\mathbf{x} \in \mathbb{R}^{|\mathbf{V}|^m}$ be the column stack of vertex positions under f . We express a distortion energy of the mapping $f[\mathbf{x}]$ as a weighted

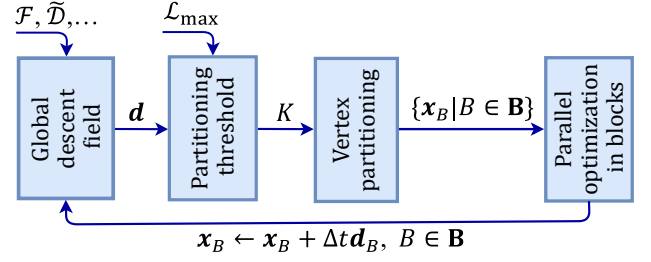


Figure 3: ABCD main stages: 1) Computing descent field of alternately selected distortions $\mathcal{D}_j = \mathcal{F}, \tilde{\mathcal{D}}$, defined by (15) and (10); 2) Estimating partitioning threshold K ; 3) Vertex partitioning into blocks \mathbf{B} by Algorithm 1; 4) Parallel optimization of (2) in coordinate blocks $B \in \mathbf{B}$.

sum over simplices

$$E(f[\mathbf{x}]) = \sum_{c \in \mathbf{C}} w(c) \mathcal{D}(f_c[\mathbf{x}]), \quad (1)$$

where $w(c)$ are simplex weights, usually equal to $\text{volume}(c)$, f_c is f 's component on c and \mathcal{D} is a given distortion measure invariant to rotations and translations of source and target vertices. Denote by df_c the Jacobian of f_c modulo a rigid transformation of c from \mathbb{R}^n to \mathbb{R}^m . Then, we address the following distortion minimization problem for $2 \leq m \leq n \leq 3$:

$$\underset{\mathbf{x}}{\text{argmin}} \quad E(f[\mathbf{x}]), \quad (2)$$

$$\text{such that: } \det(df_c) > 0, \quad \forall c \in \mathbf{C}, \quad (3)$$

$$A\mathbf{x} = \mathbf{z}, \quad (4)$$

where (3) guarantees that f is free of inversions and (4) are the given positional constraints.

Numerous first and second order techniques for solving (2) and its unconstrained version have been developed. A typical geometric solver updates target coordinates \mathbf{x} via

$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + \Delta t \mathbf{d}(\nabla_{\mathbf{x}} E, \nabla_{\mathbf{x}}^2 E), \quad (5)$$

where \mathbf{d} is the descent direction $\langle \mathbf{d}, \nabla_{\mathbf{x}} E \rangle < 0$, expressed as a function of the gradient and the Hessian

$$\nabla_{\mathbf{x}} E = \frac{\partial E(\mathbf{x}(t))}{\partial \mathbf{x}}, \quad \nabla_{\mathbf{x}}^2 E = \frac{\partial^2 E(\mathbf{x}(t))}{\partial \mathbf{x}^2}. \quad (6)$$

In first order methods, $\mathbf{d} = \mathbf{d}(\nabla E)$ and it is computed, in general, by preconditioning the gradient with a sparse *proxy matrix*, i.e., in unconstrained problems \mathbf{d} is a solution of a sparse linear system

$$H\mathbf{d} = -\nabla_{\mathbf{x}} E. \quad (7)$$

If \mathbf{x}^i are vertex coordinates at iteration i , then updating \mathbf{x} via (7) is equivalent to minimizing the following quadratic proxy:

$$E_H(\mathbf{x}) = E(\mathbf{x}^i) + (\mathbf{x} - \mathbf{x}^i)^\top \nabla_{\mathbf{x}} E + (\mathbf{x} - \mathbf{x}^i)^\top H(\mathbf{x} - \mathbf{x}^i) / 2. \quad (8)$$

For example, $H = I$ in the gradient descent (GD) [NSZ18; FLG15]. In Sobolev gradient descent (SGD), and in the related Accelerated Quadratic Proxy (AQP) [KGL16], H is chosen to be the rest mesh Laplacian. The method of Scalable Locally Injective Mappings (SLIM) [RPPSH17] extends the local-global parametrization

[LX*08] to general distortions by introducing a weighted proxy in the global step, which is equivalent to approximating H by reweighted Laplacian. The geometric approach of [CBSS17] introduces the Killing operator of discrete vector fields as an isometry-aware preconditioner (AKVF).

Quasi-Newton methods [ZBK18; LBK17; SS15] optimize a quadratic proxy constructed from gradients of both the current and previous iterations, i.e., $\mathbf{d} = \mathbf{d}(\nabla_{\mathbf{x}}E^i, \nabla_{\mathbf{x}}E^{i-1}, \dots)$. In particular, Zhu et al. [ZBK18] had proposed Blended Cured Quasi-Newton (BCQN) strategy of a gradual blending between AQP [KGL16] and L-BFGS [SS15]. This approach benefits both from the rapid progress of AQP at the first iterations, and having the super-linear convergence of L-BFGS in the vicinity of the optimal point.

Similarly, second order methods are based on “Newton” update step in (5), where \mathbf{d} is the function of both ∇E and $\nabla^2 E$. Since, the problem (2) is highly non-convex, these solvers compute positive semidefinite approximations of $\nabla^2 E$, most commonly, using the Hessian diagonal $H = \text{diag}(\nabla^2 E(\mathbf{x}))$ in Jacobi gradient descent [WY16]; by projecting $m(m+1) \times m(m+1)$ blocks of $\nabla^2 E$ into PSD cone for each simplex in Projected Newton (PN) methods [LKK*18; LBK16; TSIF05]. Specifically, in 2D, positive semidefinite H can be obtained via complex problem formulation [GSC18; CW17], or by using the Composite Majorization (CM) [SPSH*17] technique and the related closed-form expression of 2×2 Jacobian singular values.

Besides the aforementioned core-solvers, there are techniques intended for accelerating existing optimization algorithms. For instance, the update step in AQP is equipped with a Nesterov-like acceleration [Nes83] that approximates $\mathbf{x}(t)$ in (5) by an affine combination of the results from the current and previous iterations, $\mathbf{x}(t) = (1 + \theta)\mathbf{x}^t - \theta\mathbf{x}^{t-1}$. Progressive parametrization (PP) [LYNF18] accelerates core-solvers by decomposing an initial map into intermediate mappings with bounded singular values. Anderson acceleration [PDZ*18] adopts numerical analysis techniques for geometric optimization by treating solvers as fixed-point iterations. The method is intended for alternating local-global algorithms, such as [BDS*12; BML*14], where in the local step vertex coordinates are projected into varying sets of positional constraints.

Despite the abundance of techniques, the vast majority of the existing solvers are iterative algorithms that require an inversion-free initialization of f in (2). Moreover, in most methods $f[\mathbf{x}^i]$ needs to be kept free of foldovers for each iteration $i \geq 0$. In view of the above limitation, different strategies have been proposed to keep f satisfying (3) during the optimization. These strategies include: designing distortions with flip barriers [SS15; FLG15], inversion-aware line search [SS15], the recently proposed barrier-aware line search filtering [ZBK18], and employment of scaffold meshes [LKK*18; JSP17].

For unconstrained 2D distortion optimization, such as parametrization of disc-topology surfaces with free boundaries, the required positively oriented initialization f^0 is readily available via the Tutte’s embedding [Tut63] and its variants [Flo03]. However, computing feasible f^0 may still be challenging in number of geometry processing applications, including 2D and 3D shape deformations with positional constraints, shape matching

and volumetric parametrizations. Moreover, Tutte’s mapping to unconvex regions cannot guarantee (3), while mapping complex shapes to a disc and other convex domains usually yields huge isometric distortions. Thus, producing f^0 free of foldovers often costs many more additional iterations for a typical parametrization algorithm. It therefore seems that mapping to domains that match the structure of a source mesh should be a much better starting point if the algorithm can fix occasional foldovers. Although there is a number of recent studies on embedding triangle mesh onto non-convex planar domains, our algorithm has unique properties: it operates both in 2D and 3D; it employs hard positional constraints, and therefore, unlike Boundary First Flattening (BFF) [SC18], it interpolates the exactly prescribed boundary shape; ABCD updates only vertex coordinates and, in contrast to the recently introduced Progressive Embedding (PE) [SJZP19], our algorithm does not involve mesh re-triangulation.

Among the recent studies that explicitly address unfeasible initializations are: bounding distortion mappings (BD and LBD) [AL13; KABL15] and simplex assembly (SA) [FL16]. BD solves the quadratic problem of constructing an optimal projection of f on the bounded distortion space, whereas LBD enhances this strategy by linearization of the BD space constraints and pre-factorization of the obtained KKT matrix. SA [FL16] optimizes significantly larger problem than (2), since its objective variables are linear and translation components of f_c represented, per simplex, by $n \times n + m$ unknowns instead of commonly employed m coordinates per vertex ($n \geq m$ and, in general, $|\mathbf{C}|$ is much larger than $|\mathbf{V}|$ according to Euler formula).

The method of [KABL14] can start with trivial maps and converge to prescribed positional constraints, but it suffers from slow performance of semidefinite programming (SDP). Weber et al. [WMZ12] introduced the Least Square Beltrami (LSB) energy for computing extremal quasi-conformal mappings. Optimization of LSB can be initialized with foldovers, since the energy is finite on inverted triangles. However, this method is limited to 2D and it does not extend to other popular energies, such as Symmetric Dirichlet energy \mathcal{D}_{iso} [SS15] or AMIPS [FLG15].

The Autocuts algorithm [PTH*17], designed for a user-assisted global parametrization, can handle non-injective initializations because it treats the input mesh as a triangle soup. Nevertheless, the algorithm requires a consistent initial orientation of deformed triangles. Moreover, this method operates on disassembled simplices which, similarly to SA, significantly increases number of objective variables. Although the underlining homotopy approach for multi-objective optimization performs well on middle resolution meshes, Autocuts supports only soft positional constraints in 2D and does not extend to volumetric domains.

We propose a novel algorithm, Adaptive Block Coordinate Descent (ABCD), aimed at robust distortion optimization. By robustness we mean that our algorithm is capable of solving challenging problems that previously available methods struggle with. These problems include cases of highly distorted, non-locally-injective, noisy initializations, and mapping with complex positional constraints. Our algorithm combines the best of core-solvers (CM, PN, BCQN, AKVF, SLIM) and map fixers (BD, LBD, SA) into unified Algorithm 3. We show that, over a wide range of test cases,




Figure 4: A typical ABCD work-flow: initialization with flips (in yellow), fixer, optimizer and final convergence. Dashed lines along descent directions depict length of line search intervals (22).

ABCD achieves superior results than mere cascade or alternating combinations of core-solvers and map fixers. Moreover, since our algorithm is very general, it can be further improved by employing novel acceleration techniques, such as PP [LYNF18].

For standard unconstrained problems our algorithm behaves exactly as an integrated core-solver (BCQN, AKVF, PN, etc.), except for a very short period, typically 2-3 first iterations, for which computational cost is negligible. In the presence of highly distorted invalid initializations, ABCD passes through two phases: the first one of cleaning foldovers, where distortions are optimized in coordinate blocks of varying sizes, and the second stage of optimizing a positively oriented map, during which ABCD typically behaves as a global solver. Our algorithm exhibits such flexibility because, at each cycle, we adaptively modify coordinate blocks and switch between different energies. Unlike the common notion of optimization theory, where “adaptiveness” is referring to a proper selection of block coordinate frames, we consider, instead, an adaptive strategy of the block partitioning and an adaptive blending of local-global solvers.

3. Modifying Distortion Measures

Distortion measures considered in (2) are rotation invariants which, according to [RPPSH17; NSZ18], can be expressed as functions of signed singular values $\sigma_1, \dots, \sigma_m$ of the Jacobian df_c :

$$\mathcal{D}(f_c) = \mathcal{D}(\sigma_1, \dots, \sigma_m). \quad (9)$$

We assume w.l.o.g. that only σ_1 can be negative in the signed SVD of the Jacobian, $df_c = U \text{diag}(\sigma_1, \dots, \sigma_m) V^\top$. Thus, on a non-degenerate simplex c , $\text{sign}(\sigma_1) = \text{sign}(\det(df_c))$.

The most common energies, considered in geometric optimization, are measures of isometric (length-wise) [CPSS10; SS15; FLG15; NSZ18] and conformal (angle-wise) [HG00; LPRM02; DMA02; FLG15] distortions. Usually, due to existence of a *barrier term*, these measures become infinite when $\det(df_c) \leq 0$. While this prohibits appearance of new foldovers, it also prevents optimization from fixing the existing ones. Therefore, we aim to modify existing energies in such a way that both tasks of fixing invalid simplices (inverted or collapsed) and preventing generation of new ones are possible within the same procedure.

Denote by s_c the orientation of simplex c at the iteration $(i-1)$, i.e., $s_c = \text{sign}(\det df_c^{i-1})$; then we modify the original measure \mathcal{D}

as follows:

$$\tilde{\mathcal{D}}(f_c) = \begin{cases} \mathcal{D}(\sigma_1, \dots, \sigma_m) & \sigma_1 > 0, s_c > 0, \\ \infty & \sigma_1 \leq 0, s_c > 0, \\ \min \mathcal{D} & s_c \leq 0, \end{cases} \quad (10)$$

where $\min \mathcal{D}$ is the absolute minimum of distortion \mathcal{D} . Namely, at each iteration we filter out the contribution of simplices that were previously inverted, while on valid simplices $\tilde{\mathcal{D}}$ equals the original measure plus the barrier term in order to prevent new foldovers. This approach enables an efficient alternating optimization of classical rotation-invariant energies and of measures designed for fixing invalid simplices (*map fixer measures*) without badly interfering with each other. So that our algorithm alternates between the two phases: minimization of map fixer measures and optimization of distortions, constructed by (10). We refer to these two phases as the *fixer* and *optimizer* phases, respectively. Since orientation of a zero-volume simplex is not well-defined, we consider inverted simplices (flips) and collapsed simplices as the two separate classes of illegally deformed elements. Next, we describe how map fixer measures are devised to correct each of the two illegal deformations:

Inverted simplices: Assuming the positive orientation of the source elements, a simplex c is inverted under f if

$$\det(df_c) = \prod \sigma_j < 0. \quad (11)$$

Rather than using binary measures, or infinite barrier terms that can easily block optimization progress, we penalize both the number of foldovers and their target volumes (areas). Since foldover volumes are negative and proportional to $\det df_c$, we consider the following *flip penalty* measure:

$$\mathcal{D}_{\text{flip}}(\sigma, \Lambda) = \begin{cases} \Lambda - \prod \sigma_j & \prod \sigma_j \leq 0, \\ 0 & \text{else,} \end{cases} \quad (12)$$

where $\Lambda \geq 0$ is the uniform penalty cost of an inverted simplex.

Collapsed simplices. Similarly to $\mathcal{D}_{\text{flip}}$, the *collapsed simplex penalty* is the cost function of nearly zero-volume target elements

$$\mathcal{D}_{\text{collapse}}(\sigma, \Lambda) = \begin{cases} \Lambda & \exists j : |\sigma_j| < \varepsilon, \\ 0 & \text{else,} \end{cases} \quad (13)$$

where ε is the singularity threshold. Unlike (12), $\mathcal{D}_{\text{collapse}}$ contains no smooth terms, and thus we provide a separate definition of its pseudo gradient by devising how σ should be modified to unfold a degenerate simplex. Specifically, if $\dim(\text{Ker}(df_c)) = k$, then $\sigma_{m-k+1}, \dots, \sigma_m = 0$, and thus $\partial \mathcal{D}_{\text{collapse}} / \partial \sigma_j$ should be non-zero for $m-k+1 \leq j \leq m$. Consequently, we set gradient entries as follows:

$$j = 1, \dots, m : \frac{\partial \mathcal{D}_{\text{collapse}}}{\partial \sigma_j} = \begin{cases} -1 & |\sigma_j| \leq \varepsilon, \\ 0 & \text{else.} \end{cases} \quad (14)$$

For processing all invalid elements with a single measure, we combine (12) and (13) into the *invalid simplex penalty*:

$$\mathcal{F}(\sigma, \Lambda) = \begin{cases} \mathcal{D}_{\text{flip}}(\sigma, \Lambda) & \sigma_1 < 0, \sigma_m \geq \varepsilon, \\ \mathcal{D}_{\text{collapse}}(\sigma, \Lambda) & \text{else,} \end{cases} \quad (15)$$

where we set a higher priority for penalizing flips, since minimizing (12) is a more complex geometric problem than unfolding col-

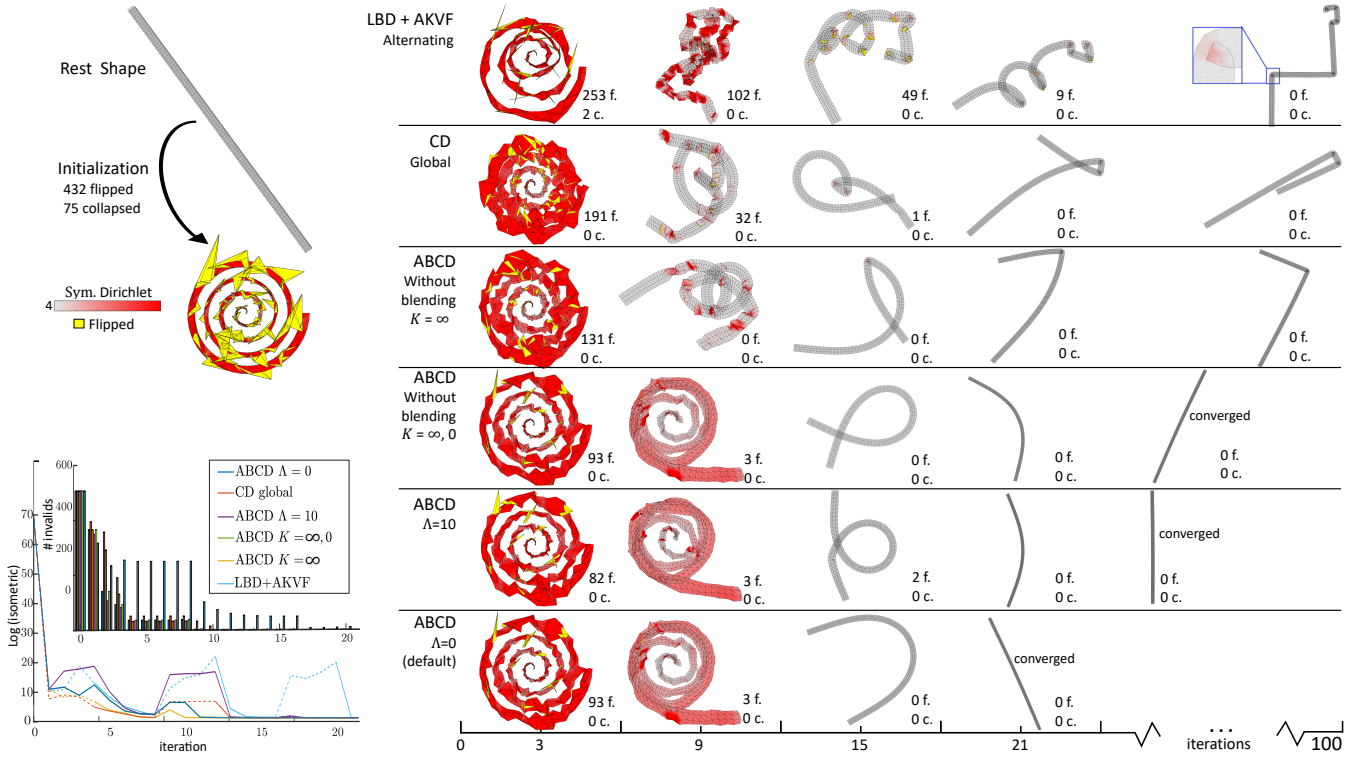


Figure 5: Comparing ABCD(AKVF) (bottom) with other methods and with simplified variants of ABCD(AKVF). **Top to bottom:** alternating optimization of LBD and AKVF; global ABCD version (CD); ABCD without blending that includes constant $K = \infty$ and constantly alternating thresholds $K = \infty, 0$; ABCD with $\mathcal{F}(\sigma, 10)$ and $\mathcal{F}(\sigma, 0)$. Each snapshot is annotated with the number of flipped and collapsed triangles, denoted by ‘f.’ and ‘c.’, respectively.

lapsed elements. Measure (15) is design to behave well in GD. In particular, GD minimization of \mathcal{F} either flips back foldovers, or shrinks them to minimize magnitudes of their volumes (Fig. 4).

Note that despite \mathcal{F} having non-differentiable points and separate definitions of energy values and energy gradients, map fixer measures are fully compatible with the geometric optimization framework. Particularly, Algorithm 3 receives \mathcal{D} and $\nabla \mathcal{D}$ in separate inputs and it can handle non-differentiable measures as long as these two inputs are consistent.

At first glance, it seems reasonable to set $\Lambda \gg 0$ in (15), since we are interested in minimizing the total number of invalid elements, rather than merely decreasing their total volume. However, the smoothness of $E_{\mathcal{F}} = \sum w(c) \mathcal{F}(\sigma, \Lambda)$ decreases in Λ . If Λ is large enough, then $E_{\mathcal{F}}$ is approximately equal to the number of invalid simplices times Λ . This quantity has a combinatorial structure, and thus minimizing $E_{\mathcal{F}}$ with common geometric solvers is not efficient for large Λ . Therefore, on average, the most significant progress is achieved by setting Λ to be a tiny positive number. Our experiments with different values of Λ are depicted in Fig. 5 and Fig. 24 in the supplemental.

By definition, $\nabla \mathcal{F}$ is non-zero only on vertices of invalid elements, often distributed in patterns of “islands” surrounded by

positively oriented simplices (see second row of Fig. 17). Following this observation, we suggest to construct block partitioning in which vertices sharing only inverted simplices and the rest of vertices are put in separate blocks. As discussed in the next section, our adaptive coordinate descent strategy has a number of apparent advantages over purely global or local approaches.

4. Local Versus Global Optimization

While most of optimization techniques are focused on approximating (1) with proxy energies, or on designing gradient preconditioner in (5), we adopt a new viewpoint and address a different aspect of the problem: How to divide objective variables \mathbf{x} into blocks to attain more efficient and robust optimization?

The vast majority of existing geometric solvers are *global methods* that update entire set of target coordinates simultaneously. Among the few purely *local methods* are block coordinate descent (BCD) strategies with constant small blocks [HG00; FLG15; FL16; NSZ18]. Note that our terminology here is different from the common notation of local-global solvers, such as [PDZ*18; RPPSH17; LX*08], where an update (5) is performed in the global stage, while auxiliary properties per simplex are computed in the local one. We

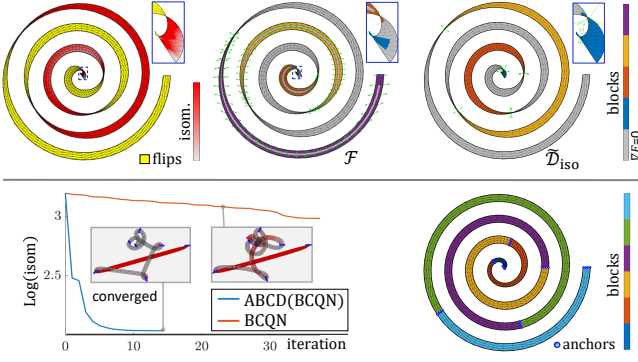


Figure 6: *Top, from left to right:* Unconstrained swirl planar deformation, initialized with foldovers, block partitioning in ABCD(GD) for \mathcal{F} and \mathcal{D}_{iso} energies, respectively. **Bottom:** We compare BCQN and ABCD(BCQN) optimizations (left) of the constrained swirl deformation, initialized without flips (right).

Algorithm 1 Vertex partitioning into blocks

Input:

- Mesh (\mathbf{V}, \mathbf{C}) with source and target coordinates \mathbf{y} and \mathbf{x} .
- Descent field \mathbf{d} and partitioning threshold K .

Output:

 Vertex blocks \mathbf{B} .

- 1: $\mathbf{E} \leftarrow \emptyset$.
 - 2: **for** $v \in \mathbf{V}$ **do**
 - 3: **for** $u \in \text{Neighbours}(v)$ **do**
 - 4: Add $\mathbf{E} \leftarrow \mathbf{E} \cup \{u, v\}$ **if** $\mathcal{L}_{uv}(\mathbf{d}, \mathbf{x}, \mathbf{y}) \leq K$ and $u \sim v$.
 - 5: **end for**
 - 6: **end for**
 - 7: $(\mathbf{V}, \mathbf{E}) \leftarrow$ Graph of nodes \mathbf{V} and edges \mathbf{E} .
 - 8: $\mathbf{B} \leftarrow$ Connected components of (\mathbf{V}, \mathbf{E}) .
-

consider both the locality and the globality with respect to the selection of \mathbf{x} 's coordinate blocks in (5).

To highlight the importance of the raised question, consider a toy model, depicted in Fig. 2, where a simple 1-ring block gradient descent [HG00; NSZ18] easily outperforms much more advanced global solvers. In general, the global strategy works well when descent direction \mathbf{d} is a smooth vector field. However, a noisy initialization produces chaotic gradients. Therefore, in this case, a global descent update of (5) can easily get stuck in the line search stage, since neighboring vertices may be forced to move in completely random directions. In particular, repairing initializations with many flipped and collapsed simplices often introduces highly non-smooth gradients and thus cannot be effectively handled by a global approach alone. Although [ZBK18] proposed some line search enhancements, where components of the global descent field \mathbf{d} are rescaled to increase the line search range, the global update of \mathbf{x} along \mathbf{d} may still be subjected to an adverse coupling of vertices in highly distorted initializations. Moreover, a sufficiently high noise level in vertex position often leads to poor preconditioning due to extremely ill-conditioned approximate Hessian. Hence, this issue

cannot be solved solely by filtering line search directions or using other related approaches.

Another motivation for using an intelligent block partitioning is derived from observing standard optimization problems with Dirichlet positional constraints. This type of constraints often result in a natural block partition, where inner block boundaries pass through fixed vertices (*anchors*) with vanishing descent directions. Similar configurations occur in minimizing map fixer measures, since they have non-zero gradients only on vertices adjacent to invalid elements. Optimizing one block of vertices has no effect on vertices in another block if the boundary between the two blocks is locked (see Fig. 6 bottom). Therefore, in such cases, it is more computationally efficient to employ BCD inherent parallelism for processing different blocks of vertices in parallel.

However, local methods are not effective in dealing with problems where the decrease in a distortion energy can be achieved only by modifying large groups of vertices, at once. These cases often include complex shape deformations, such as the one depicted in Fig. 6, and surface parametrization with standard harmonic initializations. Apparently, as illustrated by Fig. 5, distortion minimization can have mixed cases, where neither purely global, nor local approaches are suitable, and thus some blending of two opposite strategies should be adopted.

5. Adaptive Block Partitioning

In this section, we propose an *adaptive block partitioning* strategy that solves problem (2) separately over blocks of target vertices $\{\mathbf{x}_B | B \in \mathbf{B}^i\}$, where the subscript in \mathbf{x}_B denotes the column stack of coordinates of vertices in B and the superscript in \mathbf{B}^i denotes that the block partitioning is recomputed for each iteration i . Similarly, we denote by \mathbf{d}_v , \mathbf{x}_v and \mathbf{y}_v , the search direction, target and source coordinates of a single vertex $v \in \mathbf{V}$, respectively.

Our partitioning scheme is designed to disconnect *static vertices* $\mathbf{V}_0 = \{v \in \mathbf{V} | \mathbf{d}_v = \mathbf{0}\}$ from the rest of the mesh and to group non-static (*free*) vertices $\mathbf{V} \setminus \mathbf{V}_0$ into edge-connected blocks with a sufficiently smooth descent field. More precisely, assume that $\mathbf{d}(\nabla E, \nabla^2 E)$ is the descent direction and denote by $u \sim v$ that u and v are neighboring vertices such that both are in \mathbf{V}_0 or in $\mathbf{V} \setminus \mathbf{V}_0$. Then, we divide vertices in such a way that each block B is a connected set with respect to edges $\{\{u, v\} | u \sim v\}$ and it satisfies:

$$\forall u, v \in B : \|\mathbf{d}_u - \mathbf{d}_v\|_2 \leq K \|\mathbf{x}_u - \mathbf{x}_v\|_2, \quad (16)$$

where K is the adjustable threshold value that controls maximal deviation in descent direction. Since “geometric information” cannot propagate through static vertices, a block surrounded by edges $\{u, v\}$, $u \sim v$, can be updated independently of the rest of the mesh. Motivated by this observation, we construct a *connectivity graph* (\mathbf{V}, \mathbf{E}) in Algorithm 1 by iterating over vertices $u \sim v$ and measuring the following quantity:

$$\mathcal{L}_{uv} = \|\mathbf{d}_v - \mathbf{d}_u\|_2 \|\mathbf{y}_v - \mathbf{y}_u\|_2 \|\mathbf{x}_v - \mathbf{x}_u\|_2^{-1}. \quad (17)$$

The term $\|\mathbf{y}_v - \mathbf{y}_u\|_2$ is chosen so that it normalizes distances in the target domain. As specified in Algorithm 1, $\{u, v\}$ is an edge of \mathbf{E} if $u \sim v$ and $\mathcal{L}_{uv} \leq K$, and blocks \mathbf{B}^i are connected components of

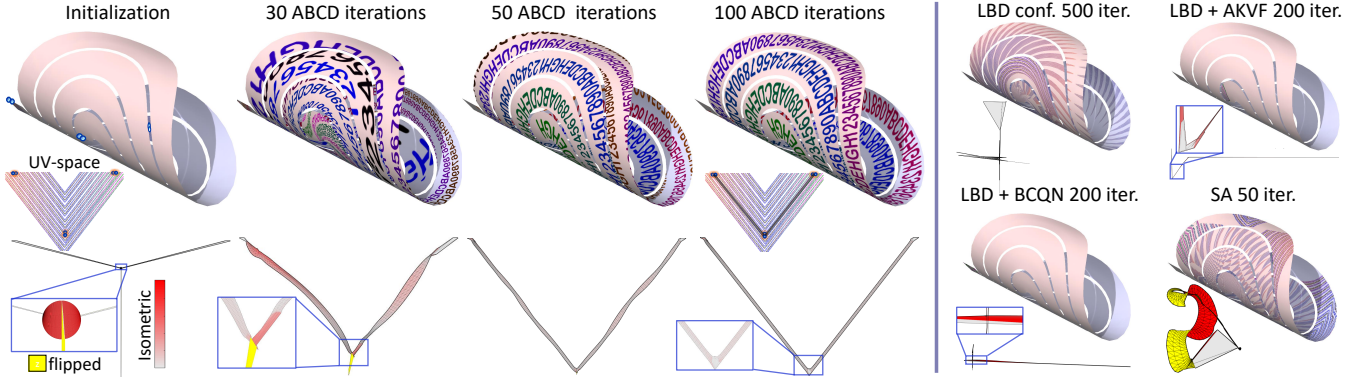


Figure 7: Constrained parametrization of a space fitting curve. The source and target anchor positions are shown in the top-left corner. Free vertices are initialized by Tutte embedding. ABCD(AKVF) results (left) are compared with other methods (right): standalone conformal LBD, alternating combinations of LBD with AKVF and BCQN, and SA. All the methods, except LBD, optimize isometric distortions.

(\mathbf{V}, \mathbf{E}). Static blocks, obtained in Algorithm 1, include anchors and vertices that share only simplices with zero distortion energy.

Our default approach to controlling vertex partitioning in Algorithm 1 is to set $K = \infty$. In this case, blocks \mathbf{B}^i are connected components of the graph $(\mathbf{V}, \{\{u, v\} \mid u \sim v\})$. If B_1, \dots, B_N are free blocks, obtained in this partitioning scheme, then optimizing separately these blocks is more effective than using a global solver because BCD methods update target coordinates in (5) with different step sizes $\Delta t_{B_1}, \dots, \Delta t_{B_N}$. This statement is proved in Appendix A. Note that if there are no inversions, then setting $K = \infty$, in most cases, results in a global optimization of unconstrained vertices. Next, we discuss another strategy in which parameter K is modified to control further partitioning of free vertices into smaller blocks.

5.1. Local-global blending

In continuous settings, condition (16) is equivalent to requiring \mathbf{d} to be a K -Lipschitz (continuous) function of \mathbf{x} in B . The block selection rule (16) is motivated by the well known fact that optimization of K -Lipschitz functions converges fast for small K . However, a fast convergence of distortion energy in small blocks does not guarantee the same for the total energy (1), since it may take time for “information” to propagate from one block to another. Therefore, parameter K should be well tuned to obtain an optimal trade-off between the number of blocks and the smoothness of \mathbf{d} , obtained in each block. A simple approach to control BCD is to alternate between number of constant thresholds parameters. For instance, one can alternate between global optimization and the finest partitioning, achieved by setting $K = 0$. However, this strategy is often wasteful, since, as illustrated by Fig. 5, a significant part of the obtained iterations can be ineffective or even counterproductive. Moreover, as explained in Section 4, using purely global approach often results in slow progress at the beginning of the optimization, due to the presence of highly distorted elements.

We propose to estimate an optimal parameter K for vertex partitioning by means of the gradual blending between two opposite strategies: 1) partitioning into connected components of ‘ \sim ’,

obtained if $K \geq \mathcal{L}_{\max} = \max_{u \sim v} \mathcal{L}_{uv}$; 2) fine partitioning into small blocks, obtained for a low value of K , which we denote by \mathcal{L}_{\min} . We refer to the proposed approach as *local-global blending* (LGB). Starting with $K^1 = \mathcal{L}_{\max}$ and $K^2 = \mathcal{L}_{\min}$ at the first and second iterations, we steadily update partitioning thresholds K^i for next iterations according to the observed progress toward achieving an optimal number of coordinate blocks.

Assume that \mathcal{P}^1 and \mathcal{P}^2 are estimates of the optimization progress achieved at the first two iterations, respectively. Then, $\mathcal{R}^{1/2} = \mathcal{P}^1 / \mathcal{P}^2$ is the rate by which the first iteration outperforms the second one. Our blending procedure is designed to either increase or decrease the initial coefficients K^1, K^2 in such a way that their values are changed in proportion to $\mathcal{R}^{1/2}$. The same process is repeated for each successive pair of parameters (K^{2s-1}, K^{2s}) , $s \geq 2$. That is, we compute K^3 and K^4 according to Algorithm 2 at the third iteration, then we apply the algorithm again at the fifth iteration to compute (K^5, K^6) as a function of $\mathcal{R}^{3/4}, K^3$ and K^4 . This process is repeated until the blending termination criterion is reached (see Section 7.4).

It is difficult to predict how the non-linear optimization (2) evolves by observing energy values or energy derivatives alone. In fact, a slight modification of a single nearly-collapsed simplex may consume more energy than a complex deformation of the entire shape. We therefore, estimate optimization progress \mathcal{P}^i by observing both the decrease in the relative energy

$$\Delta E^i = \frac{E(\mathbf{x}^{i-1}) - E(\mathbf{x}^i)}{E(\mathbf{x}^{i-1})}, \quad (18)$$

and the magnitude of non-rigid motion by which target vertices are moved. We refer to the latter quantity as *displacement norm*, defined via Frobenius norm as follows:

$$\text{Disp}^i = \|\mathbf{x}^i - \mathbf{x}^{i-1} - \text{Proj}_{\text{Ker}(\mathbf{x}^{i-1})}(\mathbf{x}^i - \mathbf{x}^{i-1})\|_{\text{Fro}}, \quad (19)$$

where $\text{Ker}(\mathbf{x}^{i-1})$ denotes the linear space of rigid transformations of the target shape [CBSS17].

Fig. 8 depicts optimization progress and plots partitioning

Algorithm 2 Local-global blending (LGB)**Input:**

- Partitioning thresholds K^1, K^2 of the previous two iterations.
- Global-local performance ratio \mathcal{R} .
- \mathcal{L}_{\max} = maximum of coefficients in (17).

Output: Thresholds K^3, K^4 for the next two iterations.

- 1: **if** $\mathcal{R} \geq 1$ **then**
- 2: $(K^3, K^4) \leftarrow (\mathcal{L}_{\max}, K^1) - \frac{1}{\mathcal{R}}(\mathcal{L}_{\max} - K^1, K^1 - K^2)$.
- 3: **else**
- 4: $(K^3, K^4) \leftarrow (K^2 + \mathcal{R}(K^1 - K^2), \mathcal{R}K^2)$.
- 5: **end if**

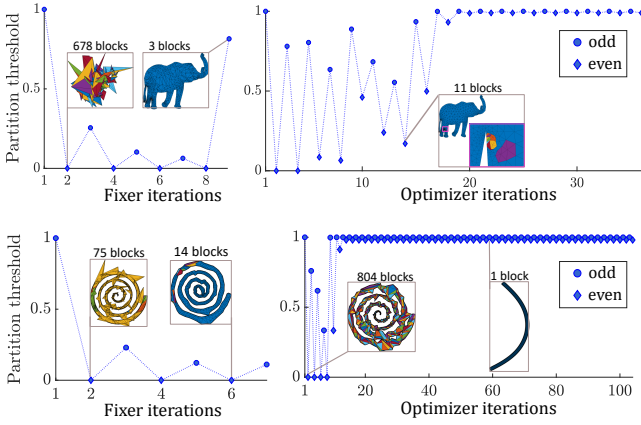


Figure 8: Normalized partitioning thresholds (K / \mathcal{L}_{\max}) computed by Algorithm 2 for distortions \mathcal{F} (left) and $\tilde{\mathcal{D}}_{iso}$ (right). Snapshots are colored according to vertex partitioning.

thresholds, computed by Algorithm 2. Further details of the local-global blending are presented in Appendix B.

6. Cured Alternating Optimization

In this section, we address the problem of optimizing multiple distortion measures. The main challenge in implementing ABCD stems from the fact that the algorithm copes with minimization of two potentially competing measures — distortion measures in a fixer phase (e.g., \mathcal{F}) vs distortion measures in optimization phase (e.g., isometric and conformal distortions) — reducing one might temporarily increase the other measure. Denote by E_1 and E_2 distortion energies of these two measures, obtained according to (1). One has to be cautious about how to combine E_1 and E_2 into a single optimization scheme. Roughly speaking, there are three options to be considered:

- Processing measures in a *cascade*, i.e., optimizing the first energy till convergence before moving to the next one;
- Combining all measures into a single objective. For example, optimizing a sum of energies

$$E_{\text{single}}(\mathbf{x}) = \alpha_1 E_1(\mathbf{x}) + \alpha_2 E_2(\mathbf{x}); \quad (20)$$

- Processing measures in an *alternating* manner. That is, repeatedly cycle through E_1 and E_2 .

The cascade approach is not robust enough because, often, there are inversions that cannot be repaired by fixer alone without deforming valid simplices.

Often, minimizing a single objective (20) is not productive, since vertex positions can be locked easily in the regions of vanishing or inconsistent descent directions. Although alternating optimization of different energies has no guaranties to avoid similar difficulties, we found that the alternating scheme works best. Figs. 7, 17 and 18 show that this option is much more robust than the others. In fact, in all these trials, we have not encountered any instances in which the optimization got stuck. Even for hard constrained problems with randomized initializations, optimization proceeds until convergence with no trouble. We posit that the reason for such a remarkable performance of the algorithm is the combination of distortion enhancements, introduced in Section 3, and our adaptive block partitioning method. In particular, our modified distortion measures significantly reduce any possible instabilities due to competing measures of the different optimization phases. In addition, our partitioning strategy avoids vertices being interlocked by allowing a separate processing of valid and invalid simplices.

It is important to mention that our algorithm is based on *inexact* block coordinate minimization of each individual measure E_j . That is, we optimize each block for a small number of successive iterations and repeat this process before the algorithm proceeds to the next measure in the queue. This is in contrast with the approach in which E_j , in its turn, is minimized until convergence or up to the point where the progress is lost. Advantages of the inexact approach were observed in number of recent studies [TRG16; FLG15]. When optimizer and map fixer distortions are alternatively optimized using the inexact approach, these measures support each other in a straightforward manner.

When the minimization of (20) is getting stuck, some recent papers propose to try to smoothen the objective more and more, until the optimization is able to resume again. This strategy may work well in Autocuts algorithm [PTH*17], since it is flexible enough to split problematic regions into disconnected sets. However, our tests show that the single objective approach is much less effective in seamless parametrization, and in parametrization with given cuts, as well as in more general 2D and 3D applications, where neighboring simplices cannot be disassembled. The uniqueness of ABCD approach is in preserving target mesh topology during the entire optimization. In fact, dividing the vertex set into blocks in our algorithm can be considered as a “weak” separation between simplices, which, unlike the standard simplex disassembling, does not duplicate even a single vertex.

7. ABCD Optimization Procedure

We are finally in position to present the full ABCD implementation, whose pseudo code is given in Algorithm 3. This implementation integrates all the introduced strategies into four major steps, illustrated schematically in Fig. 3. In addition to these stages, our procedure includes a number of sub-steps described below. Some

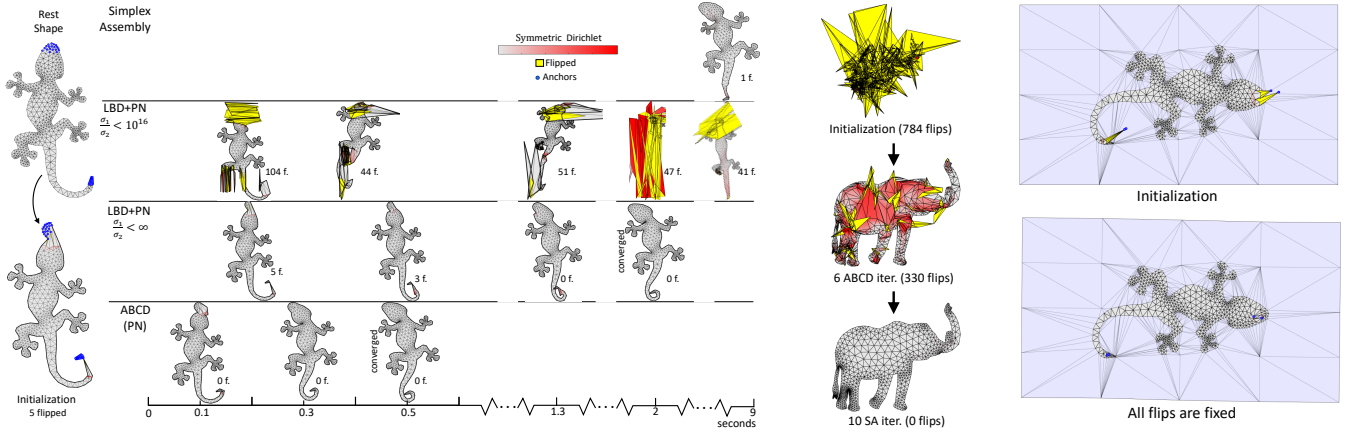


Figure 9: *Left:* Comparing ABCD(PN) with LBD+PN and SA for a planar problem, initialized with flips. We used the following parameters for LBD: finite unbounded condition number $\sigma_1/\sigma_2 < \infty$ and $\sigma_1/\sigma_2 < 10^{16}$. *Middle:* Running SA with first six ABCD iterations to solve the failure case of SA from Fig. 17. *Right:* Using scaffold meshes to induce a globally injective map in ABCD(PN). Showing the two stages: anchors are moved to initialize the problem (top), all inversions are repaired (bottom).

of these sub-steps are unique for our approach, whereas others are well known methods that we have adjusted to our needs.

7.1. Energy sequence specifications

To achieve its goal, the list of energies in ABCD should contain both an invalid simplex penalty (15) and an optimizer distortion \tilde{D} , defined according to (10). In our experiments, we minimize energies \mathcal{F} and \tilde{D} and pick a small number of successive iterations n_j for each measure to exploit best the alternating optimization. We always begin with \mathcal{F} to provide a better starting point for the optimizer.

Not that the exact mix of the energies can be very general and there is also some flexibility in setting optimization parameters. We explore a more general configurations for ABCD in Table 1 and in Fig. 24 in supplemental. According to these results, our algorithm behaves consistently across different settings.

7.2. Core-solver specifications

Each energy E_j , can, in principle, be minimized by a different solver in Algorithm 3. So we can have S_1 and S_2 solvers, where each solver should meet the following two criteria: (i) it optimizes rotation invariant distortions; (ii) it is a line search based, i.e., it modifies vertex positions along the obtained descent direction. Although very general solvers can be used, it is important to keep in mind two things related to their usage:

- Often, at the beginning, iterations include one-ring vertex blocks in which expensive core-solvers lose their efficiency. After experimenting with various scenarios and estimating an average time required for solving (7) in first and second order methods, we suggest GD as a solver of choice for optimizing a single vertex block. Note that, in general, small blocks disappear soon after all elements attain positive orientation.

- We also recommend to use GD for minimizing \mathcal{F} . A simple GD actually works better for this measure than more advanced solvers, both when optimizing over small and large blocks. The reason for this is that \mathcal{F} is specifically designed to fit GD, whereas state-of-the-art solvers, such as CM or SLIM, are designed to minimize isometric and conformal distortions. Moreover, in GD, vertices with vanishing gradients belong to static blocks, since they have a zero descent direction. As a result, ABCD fixer with GD updates only vertices that share invalid elements. Therefore, if only a small fraction of simplices are flipped, then fixer iterations run much faster with GD than with other solvers.

7.3. Enhanced line search filtering

The line search sub-step is aimed at modifying vertex coordinates for the next iteration along the given block descent field, i.e., $\mathbf{x}_B^{i+1} = \mathbf{x}_B^i + \Delta t_B^i \mathbf{d}_B^i$. We use the Armijo backtracking search to estimate a sufficiently good minimizer

$$\Delta t_B^i = \operatorname{argmin}_{t \in [0, T_{\max}]} E(\mathbf{x}_B^i + t \mathbf{d}_B^i). \quad (21)$$

We compute T_{\max} in (21) differently for map fixer measures and for optimizer distortions, defined by (10). For optimizer distortions, we employ the filtering of [SS15] over valid elements

$$T_{\max} = (1 - \delta^+) t_B^+, \quad (22)$$

$$t_B^+ = \min_{c \in \mathbf{C}^+(B)} \{\infty, t > 0 \mid \det(df_c[t]) = 0\}, \quad (23)$$

where $\mathbf{C}^+(B)$ denotes the set of positively oriented simplices that share vertices in B , $df_c[t]$ is the Jacobian of c , induced by mapping vertices in B to $\mathbf{x}_B^i + t \mathbf{d}_B^i$, and $\delta^+ > 0$ is a small number that we choose to avoid numerical errors on nearly-collapsed elements.

Denote by $\mathbf{C}^-(B)$ the set of inverted simplices adjacent to vertices in B . Then, for map fixer distortions, such as \mathcal{F} , the length

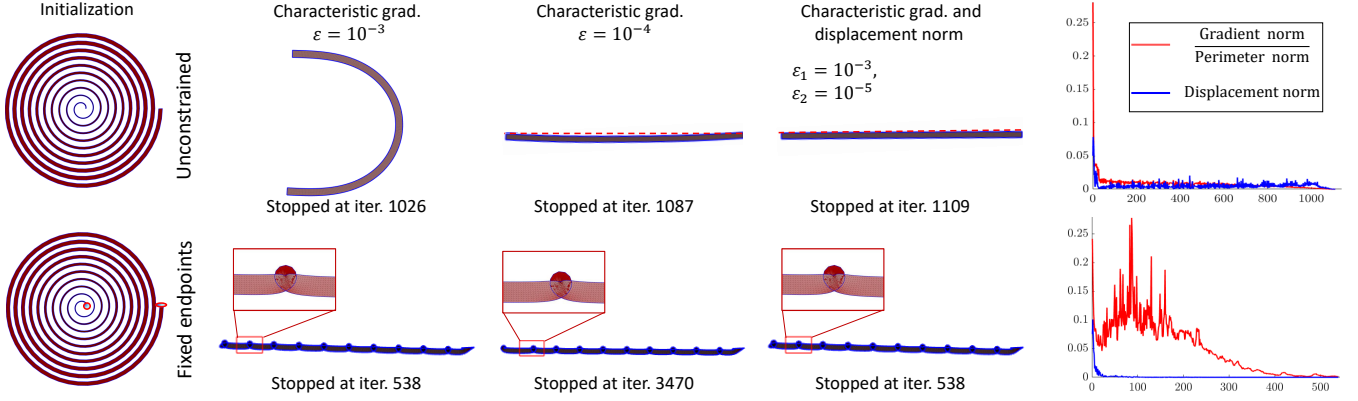


Figure 10: Characteristic gradient termination criteria [ZBK18] is compared against enhanced criteria (24). We test these criteria with BCQN solver and \mathcal{D}_{iso} distortion for unconstrained swirl deformation (top) and for the same problem with fixed endpoints (bottom).

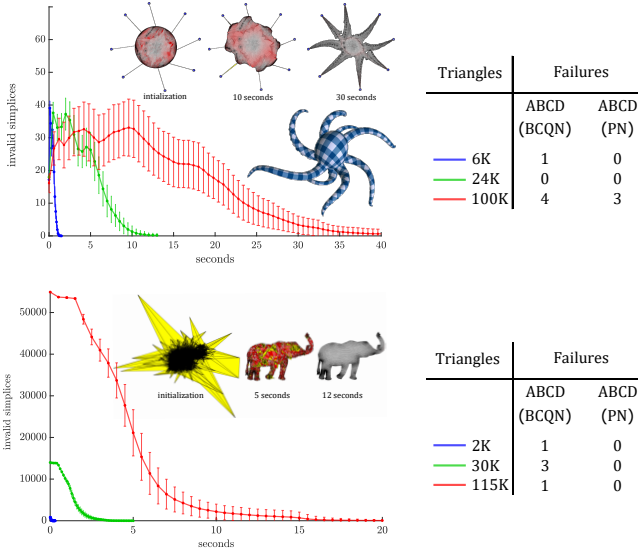


Figure 11: Testing ABCD on randomly generated deformations that include 600 trials of constrained parametrization with randomly placed anchors (top) and 600 trials of planar shape deformation with fixed boundary and randomized initialization of interiors (bottom). We tested 100 problems with Eigen library for three resolution levels: 2K, 30K and 115K triangles for the elephant model; 6K, 24K, 100K for the octopus. **Left:** Average number of flips in ABCD(PN) as a function of the runtime (vertical bars represent the standard deviation). **Right:** the number of failures encountered in a hundred trials of ABCD(BCQN) and ABCD(PN).

of the search interval is set to $T_{\max} = (1 - \delta^-)t_B^-$, where t_B^- is the minimum of the right side of (23), computed over $\mathbf{C}^-(B)$, and δ^- is a negative number. Here we use $\delta^- < 0$ instead of δ^+ because often a minimizer of $E_{\mathcal{F}}(\mathbf{x}_B^l + t\mathbf{d}_B^l)$ lies outside the interval $[0, t_B^-]$. See Fig. 4 for an illustration of the line search step.

7.4. Termination criteria

Unlike global solvers with a single objective, ABCD has several termination criteria — to stop processing of a current block, a current measure and the entire optimization, and the special criteria for terminating or resetting the LGB. A successful operation of ABCD should encounter two phases: the first one of cleaning foldovers and the final one of $\tilde{\mathcal{D}}$ convergence. The first phase is completed when \mathcal{F} vanishes, and further processing is devoted solely to $\tilde{\mathcal{D}}$. Thus, the termination of an individual measure is claimed when its average value reaches the optimum. The final convergence criteria is met by either reaching the maximal iteration number (i_{\max}), or by arriving at sufficiently accurate solution. To recognize such solution, we employ both the characteristic gradient (perimeter) norm [ZBK18] and the average of the proposed displacement norm (19). That is, we stop if the following two criteria are met:

$$\sum_j \|\nabla E_j\| < \varepsilon_1 \sum_j \text{Char}(E_j), \quad \frac{\text{Disp}^i(\mathbf{x})}{\Sigma w(c)} < \varepsilon_2, \quad (24)$$

where $\text{Char}(E_j)$ denotes the characteristic gradient norm of E_j . Although using $\text{Char}(E_i)$ is more robust than a direct gradient tolerance, we found that characteristic norm can still lead to incorrect termination claims if ε_1 is too high, or to redundant iterations if the threshold is too small. As exemplified by Fig. 10, we can solve failure cases of the perimeter norm criterion by adding the displacement norm check, since it provides a simple, yet reliable, way for ensuring the absence of any further progress.

We use (24) for terminating optimization in blocks, too. However, we bound block iterations by a small number b_j , to follow our paradigm of inexact optimization. To minimize redundant operations in non-optimally selected blocks, we start with $b_j = 1$ and increment it at each cycle. Note that the impact of b_j disappears once block partitioning reaches a steady state. We stop updating parameters K by Algorithm 2 after the number of blocks in even iterations becomes equal to that of odd iterations. However, we always track the progress over last iterations and, if there is no decrease in a distortion energy or in the number of invalid elements, then the thresholds K are reset to their initial values.

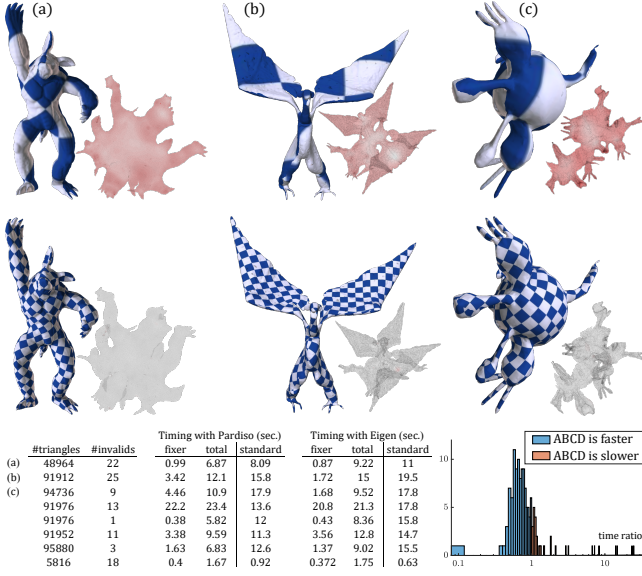


Figure 12: Using conformal flattening as a starting point for isometric parametrization. We depict few samples from our experiment in which BFF conformal maps contain flipped or collapsed triangles. **Top:** BFF maps. **Middle:** Isometric parametrization, obtained by ABCD(PN), initialized with BFF. **Bottom:** The table compares the speed of our method ('total') with the speed of the standard isometric parametrization with PN solver, initialized by Tutte ('standard'). The histogram depicts runtime ratio of our method and the standard one — when the ratio is less than one our method is faster.

7.5. Implementation

We have written the main body of Algorithm 3 in Matlab for attaining a highly customizable interface that supports multiple core-solvers (GD, AKVF, BCQN, PN) and different variants of ABCD. Critical parts of the code were implemented in C++. We reimplemented AKVF and modified existing BCQN code to fully integrate these core-solvers with ABCD. The parallel version of our code supports PN solver and it is based on [CM83] graph coloring algorithm. We have tested Eigen and Pardiso libraries to solve (7) in ABCD(PN). We report runtimes for ABCD(PN) in Figs. 9, 14, 12, 15, 16 and in Figs. 26, 21 and 24 in the supplemental.

8. Results

Since our algorithm unifies map fixers and core-solvers into a single framework, we compare our results with methods from both categories. To make a fair comparison, we examine different core-solvers for ABCD (BCQN, PN, AKVF, GD) and alternating combinations of state-of-the-art map fixers and core-solvers (e.g. LBD+BCQN, LBD+AKVF, LBD+PN, SA). To compare iterations of these methods with ABCD, we consider global solvers as BCD applied on a single block containing the entire vertex set. We rearrange alternating iterations of global solvers into the hierarchical structure of Algorithm 3, and then count overall number of b_j loops (line 9 of Algorithm 3).

Algorithm 3 Adaptive block coordinate descent (ABCD)

Input:

- Source mesh (\mathbf{V}, \mathbf{C}) and vertex initialization \mathbf{x}^0 .
- Distortions \mathcal{D}_1 (fixer) and \mathcal{D}_2 (optimizer).
- Core solvers S_1 and S_2 with n_1 and n_2 successive iterations.

- 1: Initialize block iterations (b_1, b_2) and thresholds (K_1, K_2) .
- 2: **repeat** {starting with $i = 0$ }
- 3: **for** $j = 1, 2$ **do** $\{n_j$ times each $\}$
- 4: $E_j \leftarrow \mathcal{D}_j$ energy of $f[\mathbf{x}^i]$ according to (1).
- 5: **if** $\text{average}(E_j) \leq \min \mathcal{D}_j$ **continue**.
- 6: $\mathbf{d} \leftarrow$ descent direction of E_j attained by S_j .
- 7: $\{\mathbf{x}_B^i\}_{B \in \mathbf{B}} \leftarrow$ BlockPartitioning(K, \mathbf{d}) via Algorithm 1.
- 8: **for** $B \in \mathbf{B}$ in parallel **do**
- 9: $\mathbf{x}_B^{i+1} \leftarrow$ solve (2) in B by S_j with b_j max iterations.
- 10: **end for**
- 11: $i \leftarrow i + 1$.
- 12: Update b_j and K_j . [without blending use $K_j = \infty$]
- 13: **end for**
- 14: **until** (24) is true for each E_j or $i \geq i_{\max}$.

We test meshes with different numbers of elements and summarize our results in Figs. 5, 7, 11, 13, 17 for 2D problems and in Figs. 1, 18 and 16 for volumetric problems. We annotate our results as follows: ABCD(S) denotes that GD and S are the fixer and optimizer solvers, respectively; $S_1 + S_2$ refers to the alternating combination of S_1 and S_2 methods in which iterations are divided between S_1 and S_2 in the same way as in our algorithm. We use ARAP distortion $\bar{\mathcal{D}}_{\text{ARAP}}$ in Fig. 16 and $\bar{\mathcal{D}}_{\text{iso}}$ distortion in other ABCD trials, where $\bar{\mathcal{D}}_{\text{iso}} = \sigma_1^2 + \sigma_1^{-2} + \dots + \sigma_m^2 + \sigma_m^{-2}$ and $\bar{\mathcal{D}}_{\text{ARAP}} = (\sigma_1 - 1)^2 + \dots + (\sigma_m - 1)^2$.

First, we compare LGB strategy with the block partitioning scheme in which the partitioning threshold is constant, $K = \infty$. We conclude that ABCD with LGB is more robust, but it has a higher computation cost at the beginning of the optimization. The impact on the algorithm speed is not significant on meshes with a few dozens of thousands of simplices or less. However, in higher resolution, we recommend either to increase the value of \mathcal{L}_{\min} or to use ABCD without LGB.

Next, we compare ABCD with LBD + AKVF in Fig. 5, to show that our strategies support each other, and, when combined together, attain the most significant progress. These and other related experiments with shape deformations and constrained parametrization indicate that ABCD is more likely to avoid poor local minima, where positively oriented maps are non-locally injective (see Fig. 5 top-right). Although we can not guarantee avoiding cases of poor local minima, our method is very general and it can be further modified for inducing globally injective mapping. This subject is briefly discussed in Appendix C and our experiment with scaffold meshes is depicted in Fig. 9 (right).

Methods from Fig. 5 perform similarly over a wide range of unconstrained and constrained problems in 2D and 3D. Furthermore, as shown in Fig. 11, we have tested ABCD on 600 models with randomly generated initializations and encountered only few failures, while competing methods (LBD, SA) appear to fail constantly

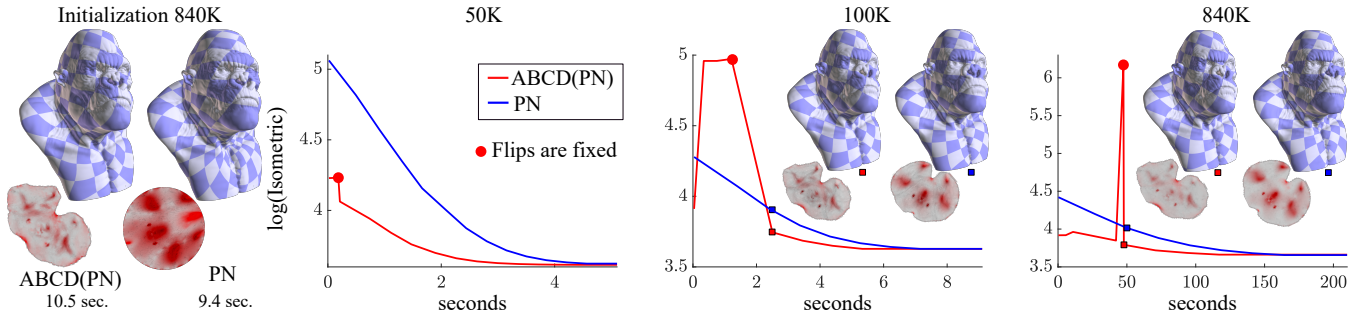


Figure 13: We compare PN parametrization, initialized by Tutte embedding, and ABCD(PN) parametrization, initialized by mapping surfaces into UV-map of a decimated mesh. Specifically, we parametrize a mesh with 2K triangles and use the deformation embedding method [NZ19] for initializing the problem in a higher resolution (see Fig. 25 in supplemental). We run 12 iterations for each mesh resolution and depict \tilde{D}_{iso} as a function of the runtime. In these examples, we used Eigen library for linear algebra.

Problem in 2D	#triangles	#invalid	#vertices / #anchors	timing to fix		iterations to fix		timing total		iterations total	
				(no LGB)	(LGB)	(no LGB)	(LGB)	(no LGB)	(LGB)	(no LGB)	(LGB)
Gecko	1238	5	804 / 45	0.04	0.121	3	7	0.17	0.25	13	22
Elephant	1796	412	1105 / 788	0.203	0.173	11	9	0.32	0.35	18	28
Octopus UV	5986	18	3924 / 8	0.469	0.313	13	9	0.91	0.89	20	27
Grid 100x100	20000	4149	10201 / 0	1.105	1.571	9	7	2.38	3.11	15	20
Octopus UV	23944	18	13833 / 8	2.022	4.757	15	49	2.58	5.82	17	57
Elephant	28736	14038	15193 / 1648	2.178	5.819	12	38	4.33	7.10	21	47
Octopus UV	95776	19	51609 / 8	9.414	13.779	20	45	14.82	17.80	27	55
Gorilla UV	100000	7	50379 / 0	0.463	0.473	3	3	2.89	3.10	6	8
Elephant	114944	54900	59121 / 3296	7.403	7.023	12	14	14.06	13.08	20	28
Gorilla UV	839092	20	420408 / 0	5.736	5.828	3	3	34.85	37.41	7	10

Problem in 3D	#tets	#invalid	#vertices / #anchors	timing to fix		iterations to fix		timing total		iterations total	
				(no LGB)	(LGB)	(no LGB)	(LGB)	(no LGB)	(LGB)	(no LGB)	(LGB)
Twisted bar	12000	6961	2541 / 0	1.65	1.65	8	8	3.18	3.18	14	14
Bar bending into spiral	30000	104	6171 / 0	0.9	0.9	6	6	3.36	3.36	10	10
Bar bending into spiral	30000	104	6171 / 2	0.91	0.91	6	6	4.41	4.41	12	12
Bar bending into spiral	30000	104	6171 / 242	1.03	1.03	7	7	9.52	9.52	21	21
Twisted wrench	50122	120	14798 / 7942	2.16	2.16	9	9	5.52	5.52	13	13
Bended wrench	50122	322	14798 / 4792	18.2	18.2	49	49	23.7	23.7	53	53
Bended wrench (ARAP)	50122	322	14798 / 4792	7.04	7.04	20	20	11.8	11.8	25	25
Armadillo	56917	405	15791 / 662	12.4	12.4	24	24	20.1	20.1	29	29
Armadillo (ARAP)	56917	405	15791 / 662	7.27	7.27	18	18	14.6	14.6	23	23
Dinosaur	58191	141	16115 / 680	22.4	22.4	58	58	28.4	28.4	60	60

Figure 14: Reporting numbers of iterations and seconds until all invalid elements are fixed, and until the final convergence (24) of ABCD(PN) with $\epsilon_1 = 10^{-3}$ and $\epsilon_2 = 10^{-2}$. In 2D, we report both the results of ABCD without LGB and of ABCD with LGB for $\mathcal{L}_{\min} = 0.1 \mathcal{L}_{\max}$. We used Pardiso linear solver to optimize \tilde{D}_{iso} (default) and \tilde{D}_{ARAP} distortions. All experiments were timed on the four-core i7-8565U CPU.

on every trial. These tests show that ABCD is well-scalable and can recover from extremely distorted meshes of a high resolution. We compared ABCD(PN) for shape deformation problems with other related methods. As demonstrated in Figs. 9 (left), 15 and 16, ABCD is both faster and more reliable than other techniques for computing positively-oriented simplicial maps. Also, our experiments with different core-solvers show that ABCD(PN) is more robust than ABCD(S) for any tested first order solver S.

We examine failure cases of LBD and SA fixers in both 2D (Figs. 17, 9) and 3D (Figs. 18, 16), and find that these methods can, in general, handle only maps with relatively small fractions of inverted and collapsed simplices. These conclusions are exemplified for highly distorted initializations of the volumetric bar and planar model in Figs. 18 and 17. Moreover, as illustrated in Fig. 1, adding positional constraints reduces significantly LBD capabilities. Note that our experiments reveal that running LBD and state-of-the-art optimizers in the alternating manner is more effective than using standalone LBD or combining these methods into a single cascade.

We tested our algorithm for computing isometric parametrizations with fixed anchors and initial mappings to non-convex domains. Some of these scenarios are presented in Figs. 7 and 17. To construct the starting point, we, first, compute the standard Tutte embedding, and then move anchors to their prescribed positions.

We experiment with alternative initialization schemes for surface parametrization. In Figs. 13 and 17 (second row), we start

with Tutte mapping onto the non-convex region obtained by parameterizing a decimated mesh. That is, we initialize a problem by mapping a rest surface into the image of its UV-map, computed in a lower resolution (Section 12 in supplemental). As demonstrated by Fig. 13, this initialization scheme can accelerate UV-map computations for meshes of a high resolution.

In Figs. 12 and 26 (supplemental), we compare the BFF conformal initialization scheme with the standard isometric parametrization. Specifically, we ran BFF conformal parametrization on 600 meshed from the dataset of [LYNF18] and collected 190 maps with invalid triangles. These maps were used to initialize ABCD(PN). In this experiment, we have only one mesh where ABCD(PN) failed to repair invalid elements in BFF, while Tutte embedding produced flipped or collapsed triangles in 20% of the trials. We tested the remaining meshes with valid Tutte maps and found that in 84% of the trials our method ran faster than the standard parametrization, initialized by Tutte embedding.

As observed in our experiment, Tutte embedding has certain implementation-related limitations and, for challenging meshes, it can fail even if the target domain is convex. In particular, as demonstrated in [SJZP19], Tutte embedding may produce invalid triangles due to precision loss in the floating point arithmetic. We tested ABCD(PN) and PE on a challenging Tutte failure example, introduced in [SJZP19]. As demonstrated in Fig. 15, ABCD ran faster and converged to a mapping with lower isometric distortion.

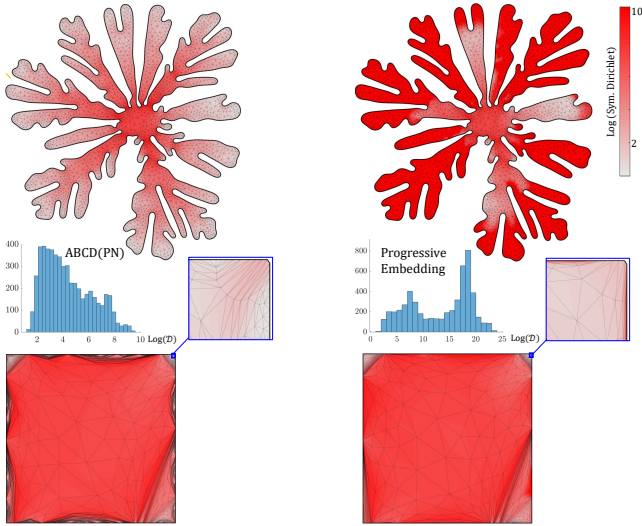


Figure 15: Mapping of the Hele-Shaw polygon [SJZP19] (top) onto the square (bottom) using ABCD(PN) (left) and the PE method (right). The initial Tutte embedding produces 41 flipped and 845 degenerate triangles. It takes 25 seconds to fix all invalid triangles by ABCD(PN), while PE runs 60 seconds. Histograms show final isometric distortions of triangles in a logarithmic scale.

According to our proposition in Appendix A, we presume that even for valid initializations, our vertex partitioning strategy can improve a core-solver performance, provided that positional constraints disconnect the mesh. See an example of such scenario at the bottom of Fig. 6.

To summarize, we have demonstrated over a variety of geometric problems that ABCD exhibits superior performance compared to existing methods in terms of its speed, robustness to inverted and highly distorted simplices, and it can deal with more complex positional constraints. In particular, our algorithm can start with the proposed conformal initialization scheme to speed up isometric parametrization for challenging meshes.

9. Conclusion and Future Work

Until recently, a lack of locally-injective initialization was one of the major concerns for running geometric optimization. Our algorithm resolves this issue for a wide range of scenarios. In this, ABCD broadens the frontiers of tractable geometric processing problems and we believe that it can contribute a lot to modern design tools.

We suggest to further explore conformal flattening and Tutte mapping onto non-convex regions as an alternative starting point for accelerating parametrizations. We believe that these methods can be integrated with the recently-proposed acceleration techniques [LYNF18; PDZ*18] to attain even faster performance. Our algorithm is built upon a number of heuristics that work well empirically. However, we cannot prove that there is a guaranty to repair all invalid elements and reach the optimality. A supplementary

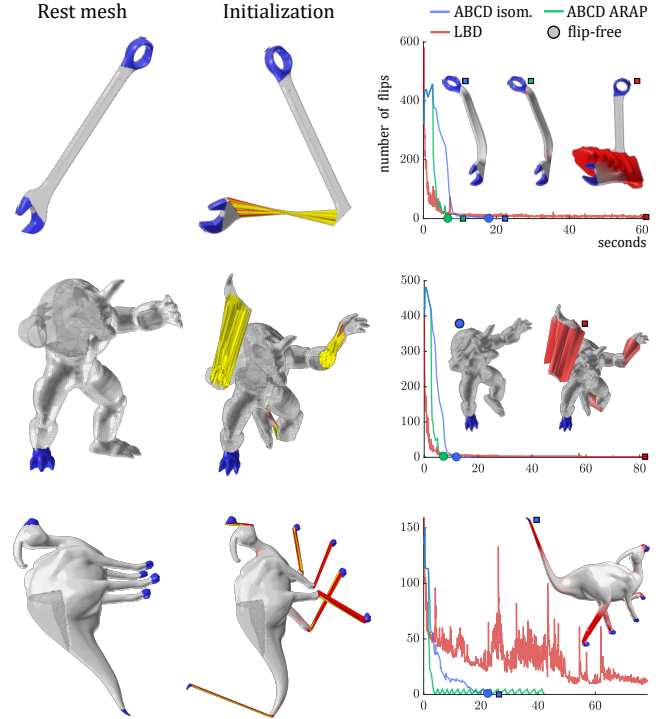


Figure 16: ABCD(PN) with \tilde{D}_{iso} and \tilde{D}_{ARAP} distortions versus LBD on tetrahedral meshes. We bound condition number in LBD by $\sigma_1/\sigma_3 < 10^2$. **Left to right:** Rest mesh, initialization with anchor points, marked in blue, and numbers of flipped tets per second.

mathematical analysis is required to better understand our empirical results and consider conditions for optimality.

Although a naive GD map fixer works well in our experiments, a more sophisticated approach can potentially yield even better results. For example, gradient preconditioning of many first-order solvers provides a better descent direction than a naive GD because preconditioning methods take into account the intrinsic metric of the deformation space. Unfortunately, at the moment it is not clear what is the right preconditioner for map fixer optimization steps. Considering a metric approach oriented at minimizing this class of geometric measures and finding the right preconditioner may be a promising direction for future research.

Acknowledgments. We thank Gregory Naitzat for proof-reading and giving valuable comments, and the anonymous reviewers for their constructive feedback. This research has been supported in part by the Ollendorff Minerva Center.

Appendix A. Proof of BCD Superiority

Proposition. Let S be a global core-solver, initialized by \mathbf{x}^0 . Denote by \mathbf{x} and $\bar{\mathbf{x}}$ results of running a single iterations of S and of ABCD(S) with $K \geq \mathcal{L}_{max}$, respectively. If the exact line search is used in (5) to minimize energy E , then $E(\bar{\mathbf{x}}) \leq E(\mathbf{x})$.

Proof. Suppose w.l.o.g. that there are two free blocks B_1, B_2 in

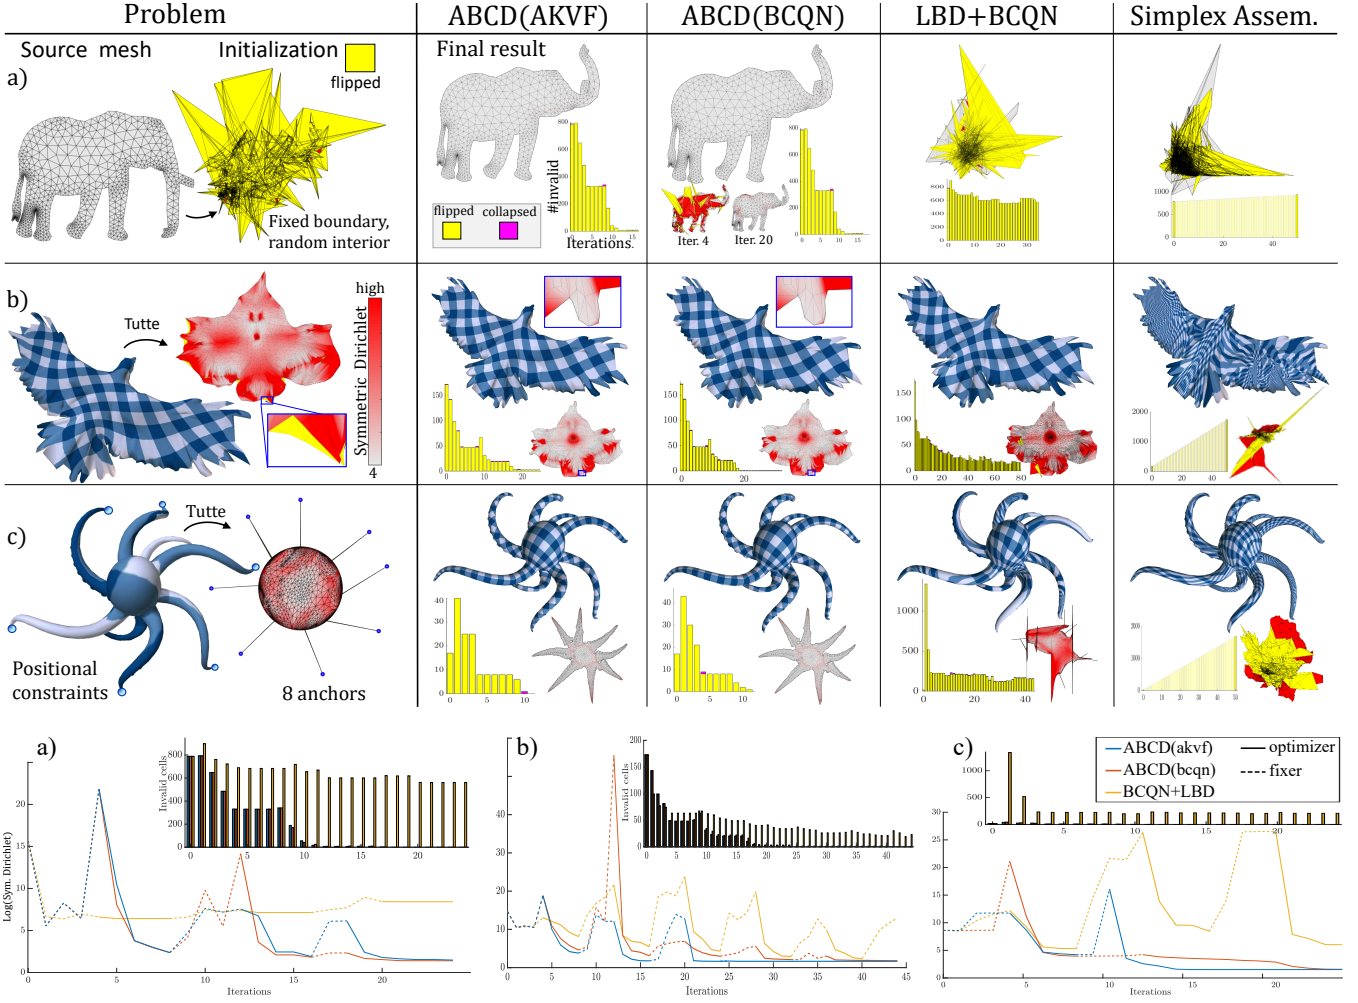


Figure 17: ABCD with first order core-solvers versus other related methods in 2D. Each element in the table shows the final result of the corresponding method and invalid triangle numbers per iteration. We plot Symmetric Dirichlet energy in logarithmic scale for each problem (bottom), where solid curves and dashed lines denote optimizer and fixer iterations, respectively (LBD is considered as a fixer). Note that SA results before the last iteration are our estimates, derived from the final output of SA.

ABCD and denote the remaining static vertices by $B_0 = \mathbf{V} \setminus (B_1 \cup B_2)$. Then, there are no mesh edges $\{u, v\}$ for $u \in B_1$ and $v \in B_2$ because, by our assumption on block partitioning, B_1 is disconnected from B_2 with respect to ‘ \sim ’ relation. As a result, we can represent E as a sum of energies, defined over disjoint simplices:

$$E = E_{B_0} + E_{B_1} + E_{B_2}; \quad E_{B_i} = \sum_{c \in \mathbf{C}(B_i)} w(c) \mathcal{D}(df_c),$$

where $\mathbf{C}(B_i)$, $i = 1, 2$, are simplices that share at least one vertex from B_i and $\mathbf{C}(B_0)$ are the rest of the elements. Denote by \mathbf{d} and $\bar{\mathbf{d}}$ the descent directions attained in S and $\text{ABCD}(S)$, respectively. Define $E[\mathbf{z}] = E(\mathbf{x}^0 + \mathbf{z})$, then the resulting energy in $\text{ABCD}(S)$ is

$$E(\bar{\mathbf{x}}) = \min_{\Delta_1} E_{B_1}[\Delta_1 \bar{\mathbf{d}}_{B_1}] + \min_{\Delta_2} E_{B_2}[\Delta_2 \bar{\mathbf{d}}_{B_2}] + \underbrace{E_{B_0}}_{\text{const.}}, \quad (25)$$

while running a single iteration of S yields

$$E(\mathbf{x}) = \min_{\Delta} (E_{B_1}[\Delta \mathbf{d}_{B_1}] + E_{B_2}[\Delta \mathbf{d}_{B_2}] + E_{B_0}). \quad (26)$$

If $\bar{\mathbf{d}}_{B_i} = \mathbf{d}_{B_i}$ for $i = 1, 2$, then the proposition’s inequality (25) \leq (26) is proven because ‘min’ is sub-additive and E_{B_0} is constant by the definition (only free vertices are updated in S and $\text{ABCD}(S)$).

Note that H in (7) satisfies the following: (i) $H_{uv} \neq 0$ only if u and v share a common simplex c ; (ii) H_{uv} depends only on \mathbf{x}_p , where p are neighbors of c ’s vertices; (iii) $H_{uv} = 0$ for any $u \in B_1$ and $v \in B_2$, since there are no edges between B_1 and B_2 . Properties (i)-(iii) imply that, by reordering vertices, (7) can be rewritten as

$$\begin{pmatrix} H_{B_0} & * & * \\ * & H_{B_1} & 0 \\ * & 0 & H_{B_2} \end{pmatrix} \begin{pmatrix} 0 \\ \mathbf{d}_{B_1} \\ \mathbf{d}_{B_2} \end{pmatrix} = - \begin{pmatrix} \nabla E_{B_0} \\ \nabla E_{B_1} \\ \nabla E_{B_2} \end{pmatrix}, \quad (27)$$

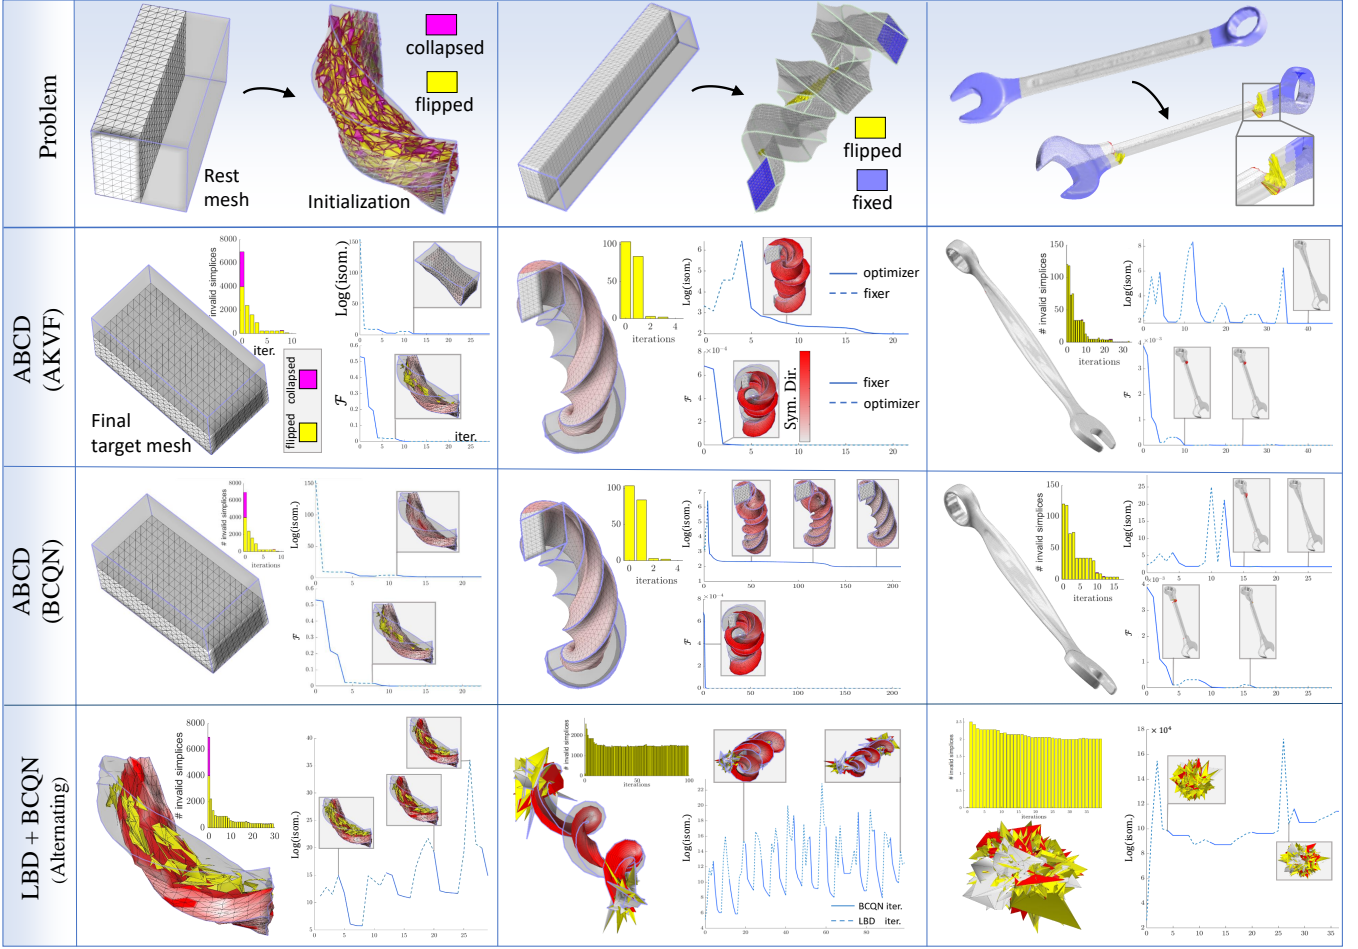


Figure 18: ABCD versus alternating combination of LBD and BCQN on tet meshes. We show number of invalid tets and energy values per iteration. We plot energy E_j in solid lines for iterations in which E_j is minimized and in other iterations E_j is shown in dashed lines.

where each H_{B_i} is a block of H , computed for vertices in B_i . Then, we conclude by (27) that $\tilde{\mathbf{d}}_{B_i} = \mathbf{d}_{B_i}$ for $i = 1, 2$. For N free blocks, B_1, \dots, B_N , we repeat the proof with $B_0 = \mathbf{V} \setminus (B_1 \cup B_2 \dots \cup B_N)$ and show that $\tilde{\mathbf{d}}_{B_i} = \mathbf{d}_{B_i}$ for $i \geq 1$.

Appendix B. Blending Procedure

Denote odd (global) and even (local) iterations by $g(i) = 2i - 1$ and $l(i) = 2i$ for $i \geq 1$, respectively. We estimate the performance ratios of k global-local iterations by

$$\mathcal{R}^i = \frac{1}{\min\{k, i\}(1 + \mu)} \sum_{s=[i-k+1]}^i \frac{\Delta E^{g(s)}}{\Delta E^{l(s)}} + \mu \frac{\text{Disp}^{g(s)}}{\text{Disp}^{l(s)}}, \quad (28)$$

where $\lceil \bullet \rceil = \max\{\bullet, 1\}$, $\mu > 0$ is a constant blending parameter and the denominator value is chosen for obtaining $\mathcal{R}^i = 1$ for equally performed global-local iterations. We compute (28) and adjust partitioning thresholds for the next cycle via

$$(K^{2i+1}, K^{2i+2}) = \text{Blend}(K^{2i-1}, K^{2i}, \mathcal{R}^i, \mathcal{L}_{\max}), \quad (29)$$

where ‘Blend’ is the procedure described in Algorithm 2.

Appendix C. Map Injectivity

Although orientation preserving maps are not always one-to-one, there is the following simple condition for guaranteeing the global bijectivity [AL13]: a positively oriented simplicial map $f : M \rightarrow \Omega$ is a global bijection if f maps bijectively ∂M onto $\partial \Omega$. As a result, ABCD produces a globally injective map upon its successful run if the optimized boundary is non self-intersecting. We have experimented with techniques [JSP17; SS15] to avoid self-intersections between boundary edges. These methods meet the necessary conditions, listed in Section 7.2, and thus can be integrated with ABCD.

In a toy model example, demonstrated in Fig. 9 (right), we attain a globally injective map by constructing a scaffold mesh [JSP17] and by running Algorithm 3 with the line search step modified to prevent boundary self-intersections [NSZ18].

References

- [AL13] AIGERMAN, NOAM and LIPMAN, YARON. “Injective and bounded distortion mappings in 3D”. *ACM Trans. Graph. (TOG)* (2013) 3, 15.
- [BDS*12] BOUAZIZ, SOFIEN, DEUSS, MARIO, SCHWARTZBURG, YULIY, et al. “Shape-Up: Shaping Discrete Geometry with Projections”. *Computer Graphics Forum*. 2012, 1657–1667 3.
- [BML*14] BOUAZIZ, SOFIEN, MARTIN, SEBASTIAN, LIU, TIAN, et al. “Projective dynamics: fusing constraint projections for fast simulation”. *ACM Transactions on Graphics (TOG)* 33.4 (2014), 154 3.
- [CBSS17] CLAICI, S, BESSMELTSEV, M, SCHAEFER, S, and SOLOMON, J. “Isometry-Aware Preconditioning for Mesh Parameterization”. *Computer Graphics Forum*. Vol. 36. Wiley Online Library. 2017, 37–47 3, 7.
- [CM83] COLEMAN, THOMAS F and MORÉ, JORGE J. “Estimation of sparse Jacobian matrices and graph coloring blems”. *SIAM journal on Numerical Analysis* 20.1 (1983), 187–209 11.
- [CPSS10] CHAO, ISAAC, PINKALL, ULRICH, SANAN, PATRICK, and SCHRÖDER, PETER. “A simple geometric model for elastic deformations”. *ACM transactions on graphics (TOG)* 29.4 (2010), 38 4.
- [CW17] CHEN, RENJIE and WEBER, OFIR. “GPU-accelerated locally injective shape deformation”. *ACM Transactions on Graphics (TOG)* 36.6 (2017), 214 3.
- [DMA02] DESBRUN, MATHIEU, MEYER, MARK, and ALLIEZ, PIERRE. “Intrinsic parameterizations of surface meshes”. *Computer Graphics Forum*. Vol. 21. 3. Wiley Online Library. 2002, 209–218 4.
- [FL16] FU, XIAO-MING and LIU, YANG. “Computing inversion-free mappings by simplex assembly”. *ACM Trans. Graph. (TOG)* 35.6 (2016), 216 3, 5.
- [FLG15] FU, XIAO-MING, LIU, YANG, and GUO, BAINING. “Computing locally injective mappings by advanced MIPS”. *ACM Transactions on Graphics (TOG)* 34.4 (2015), 71 2–5, 8.
- [Flo03] FLOATER, MICHAEL. “One-to-one piecewise linear mappings over triangulations”. *Mathematics of Computation* 72.242 (2003) 3.
- [GSC18] GOLLA, BJÖRN, SEIDEL, HANS-PETER, and CHEN, RENJIE. “Piecewise linear mapping optimization based on the complex view”. *Computer Graphics Forum*. Vol. 37. 7. Wiley Online Library. 2018 3.
- [HG00] HORMANN, KAI and GREINER, GÜNTHER. *MIPS: An efficient global parametrization method*. Tech. rep. ERLANGEN-NUERNBERG UNIV (GERMANY) COMPUTER GRAPHICS GROUP, 2000 4–6.
- [JSP17] JIANG, ZHONGSHI, SCHAEFER, SCOTT, and PANOZZO, DANIELE. “Simplicial complex augmentation framework for bijective maps”. *ACM Trans. Graph. (TOG)* 36.6 (2017), 186 3, 15.
- [KABL14] KOVALSKY, SHAHAR Z, AIGERMAN, NOAM, BASRI, RONEN, and LIPMAN, YARON. “Controlling Singular Values with Semidefinite Programming”. *ACM Trans. Graph. (TOG)* 33.4 (2014), 68 3.
- [KABL15] KOVALSKY, SHAHAR Z, AIGERMAN, NOAM, BASRI, RONEN, and LIPMAN, YARON. “Large-scale bounded distortion mappings”. *ACM Trans. Graph* 34.6 (2015), 191 3.
- [KGL16] KOVALSKY, SHAHAR Z, GALUN, MEIRAV, and LIPMAN, YARON. “Accelerated quadratic proxy for geometric optimization”. *ACM Trans. Graph. (TOG)* 35.4 (2016), 134 2, 3.
- [LBK16] LIU, TIAN, BOUAZIZ, SOFIEN, and KAVAN, LADISLAV. “Towards real-time simulation of hyperelastic materials”. *arXiv preprint arXiv:1604.07378* (2016) 3.
- [LBK17] LIU, TIAN, BOUAZIZ, SOFIEN, and KAVAN, LADISLAV. “Quasi-newton methods for real-time simulation of hyperelastic materials”. *ACM Trans. Graph. (TOG)* 36.4 (2017), 116a 3.
- [LKK*18] LI, MINCHEN, KAUFMAN, DANNY M, KIM, VLADIMIR G, et al. “OptCuts: joint optimization of surface cuts and parameterization”. *SIGGRAPH Asia 2018 Technical Papers*. ACM. 2018, 247 3.
- [LPRM02] LÉVY, BRUNO, PETITJEAN, SYLVAIN, RAY, NICOLAS, and MAILLOT, JÉRÔME. “Least squares conformal maps for automatic texture atlas generation”. *Acm transactions on graphics (tog)*. Vol. 21. 3. ACM. 2002, 362–371 4.
- [LX*08] LIU, LIGANG, 0021, LEI ZHANG, XU, YIN, et al. “A Local/Global Approach to Mesh Parameterization”. *Computer Graphics Forum* 27.5 (2008), 1495–1504 3, 5.
- [LYNF18] LIU, LIGANG, YE, CHUNYANG, NI, RUIQI, and FU, XIAO-MING. “Progressive parameterizations”. *ACM Transactions on Graphics (TOG)* 37.4 (2018), 41 3, 4, 12, 13, 22.
- [Nes83] NESTEROV, YURII E. “A method for solving the convex programming problem with convergence rate $O(1/k^2)$ ”. *Dokl. Akad. Nauk SSSR*. Vol. 269. 1983, 543–547 3.
- [NSZ18] NAITSAT, ALEXANDER, SAUCAN, EMIL, and ZEEVI, YEHOShUA Y. “Geometry-based distortion measures for space deformation”. *Graphical Models* 100 (2018), 12–25 2, 4–6, 15.
- [NZ19] NAITSAT, ALEXANDER and ZEEVI, YEHOShUA Y. “Multi-resolution approach to computing locally injective maps on meshes”. *ACM SIGGRAPH 2019 Posters*. ACM. 2019, 87 12, 18.
- [PDZ*18] PENG, YUE, DENG, BAILIN, ZHANG, JUYONG, et al. “Anderson acceleration for geometry optimization and physics simulation”. *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–14 3, 5, 13.
- [PTH*17] PORANNE, ROI, TARINI, MARCO, HUBER, SANDRO, et al. “Autocuts: simultaneous distortion and cut optimization for UV mapping”. *ACM Trans. Graph. (TOG)* 36.6 (2017), 215 3, 8.
- [RPPSH17] RABINOVICH, MICHAEL, PORANNE, ROI, PANOZZO, DANIELE, and SORKINE-HORNUNG, OLGA. “Scalable locally injective mappings”. *ACM Transactions on Graphics (TOG)* (2017) 2, 4, 5.
- [SC18] SAWHNEY, ROHAN and CRANE, KEENAN. “Boundary first flattening”. *ACM Transactions on Graphics (ToG)* 37.1 (2018), 5 3, 17, 19.
- [SJZP19] SHEN, HANXIAO, JIANG, ZHONGSHI, ZORIN, DENIS, and PANOZZO, DANIELE. “Progressive embedding”. *ACM Transactions on Graphics (TOG)* 38.4 (2019), 32 3, 12, 13.
- [SPSH*17] SHTENDEL, ANNA, PORANNE, ROI, SORKINE-HORNUNG, OLGA, et al. “Geometric optimization via composite majorization”. *ACM Trans. Graph* 36.4 (2017), 38 3.
- [SS15] SMITH, JASON and SCHAEFER, SCOTT. “Bijective parameterization with free boundaries”. *ACM Trans. Graph. (TOG)* 34.4 (2015), 70 3, 4, 9, 15.
- [TRG16] TAPPENDEN, RACHAEL, RICHTÁRIK, PETER, and GONDZIO, JACEK. “Inexact coordinate descent: complexity and preconditioning”. *Journal of Optimization Theory and Applications* 170.1 (2016) 8.
- [TSIF05] TERAN, JOSEPH, SIFAKIS, EFTYCHIOS, IRVING, GEOFFREY, and FEDKIW, RONALD. “Robust quasistatic finite elements and flesh simulation”. *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*. ACM. 2005, 181–190 3.
- [Tut63] TUTTE, WILLIAM THOMAS. “How to draw a graph”. *Proceedings of the London Mathematical Society* 3.1 (1963), 743–767 3.
- [WMZ12] WEBER, OFIR, MYLES, ASHISH, and ZORIN, DENIS. “Computing Extremal Quasiconformal Maps”. *Computer Graphics Forum* 31.5 (2012), 1679–1689 3.
- [WY16] WANG, HUAMIN and YANG, YIN. “Descent methods for elastic body simulation on the GPU”. *ACM Trans. Graph. (TOG)* 35.6 (2016), 212 3.
- [ZBK18] ZHU, YUFENG, BRIDSON, ROBERT, and KAUFMAN, DANNY M. “Blended cured quasi-newton for distortion optimization”. *ACM Transactions on Graphics (TOG)* 37.4 (2018), 40 3, 6, 10, 20.

Adaptive Block Coordinate Descent for Distortion Optimization

Alexander Naitzat¹, Yufeng Zhu² and Yehoshua Y. Zeevi¹¹Viterbi Faculty of Electrical Engineering, Technion - Israel Institute of Technology²Facebook Reality Labs

Supplementary Document

In this document we provide some additional results and discuss ABCD limitations and certain technical details related to an optimal usage of our algorithm. We adopt abbreviations and notations introduced in the paper. In particular, \mathcal{D}_j denotes a SVD-based distortion measures; its energy (1) is denoted by E_j and S_j refers to a core-solver chosen to optimize E_j in ABCD.

10. Additional Results

We provide the following additional results:

- Fig. 26 presents results of conformal initialization scheme for all 190 meshes collected from the Progressive Parametrization dataset. These results are summarized in Fig. 12.
- In Fig. 24, we tested how hyper-parameters affect ABCD algorithm. Each trial, tested in Fig. 24, had a different set of hyper-parameters that control the following processes: division between global and block iterations, line search, LGB and evaluation of the measure \mathcal{F} . These parameters are listed in Table 1. Note that fixer and optimizer stages have different hyper-parameters.
- In Fig. 19, we experimented with as killing as possible constrained parametrization. In this figure, we tested how ABCD behaves under positional constraints that enforce negative orientation of target simplices.
- In Fig. 21, we used ABCD(PN) to fix failures in Tutte embedding and parametrize space fitting curves with different initialization schemes. We compared the two initialization methods: Tutte embedding and the BFF mapping [SC18]. According to our results, the proposed BFF initialization scheme outperforms significantly the standard approach with Tutte embedding. Starting with BFF, conformal embedding produces no flips in these experiments, whereas Tutte embedding produces a small number of foldovers due to the presence of extremely distorted triangles. Furthermore, conformal initializations are very close to the final isometric parametrization.
- Fig. 23 compares iterations of ABCD(BCQN) and ABCD(PN) methods that we used in our randomized experiments (Fig. 11).

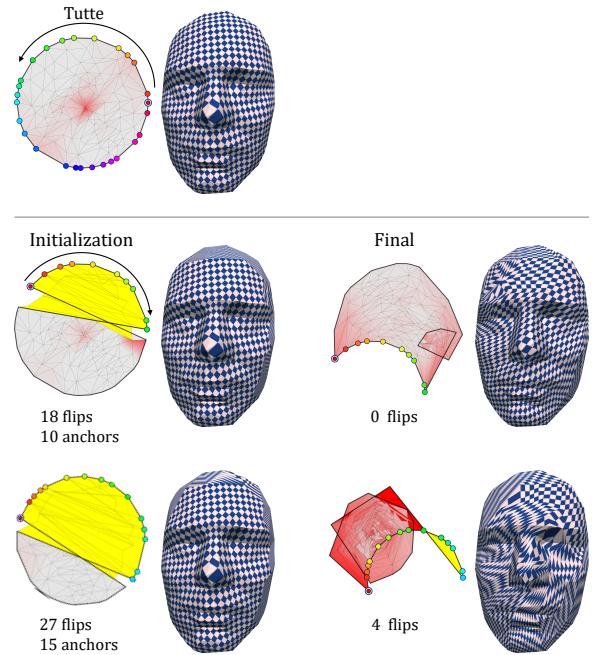


Figure 19: Running a stress test of ABCD(PN) to identify positional constraints inconsistent with a positively oriented UV-map. To find failure cases in the algorithm, we deform Tutte mapping and fix positions of distorted vertices. **Top:** Positively-oriented Tutte mapping into a disc; **Middle:** We flip positions of 10 successive boundary anchors. **Bottom:** We flip positions of 15 successive boundary anchors.

11. Failures and Challenging Cases

Because we use hard positional constraints, our algorithm cannot repair inversions if given positional constraints, $A\mathbf{x} = \mathbf{z}$, contradict the orientation requirement, $\forall c \in \mathbf{C} : \det df_c > 0$.

Obviously, if $A \subset \mathbf{V}$ is a set of anchors, then the restriction of an initial map to A has to be an orientation-preserving map. In Fig. 19, we use this property of initial maps to test how many inconsistently placed anchors can be handled by our algorithm. We start

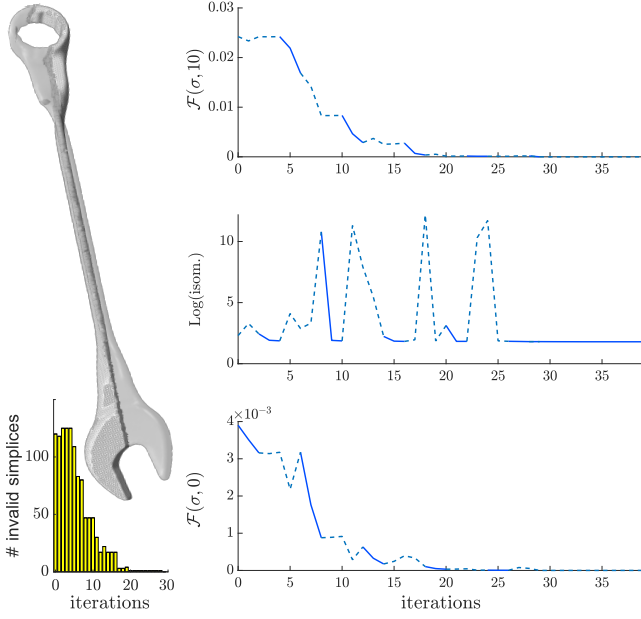


Figure 20: Here we solve the third problem from Fig. 18 by minimizing three distortions: $\mathcal{F}(\sigma, 0)$, \tilde{D}_{iso} and $\mathcal{F}(\sigma, 10)$. These distortions are optimized according to Algorithm 4 with $n_j = 2$ and the results are visualized in the same way as in Fig. 18.

with a positively oriented initialization and add gradually anchors to revert counter clockwise orientation of the boundary. Our algorithm succeeds to repair inversion for almost a half of inconsistently placed boundary vertices (Fig. 19 middle) and it fails when more than half of the boundary is reverted to the clockwise orientation. Note that, in the unconstrained version of the problem, ABCD(PN) repairs all the inversions without any difficulty.

To resolve failures induced by inconsistent anchor positions, one should use soft positional constraints, instead of hard ones. That is, one should replace (4) with a cost function that penalizes anchor deviations from their prescribed positions. However, the subject on soft constraints is beyond the scope of this work because we have restricted objective functions in (2) to be SVD-based energies.

In some challenging cases, ABCD may struggle to repair inversions due to the presence of extreme differences in sizes of target elements. In particular, running BFF on some meshes from Fig. 26 produces conformal mapping with many “poles” — areas where target elements around inversions are shrunk to extremely small areas. It is difficult to set a proper singularity threshold ϵ in (13) that works well on these meshes. Setting ϵ too low may lead to numerical issues in the optimizer stage, setting it too high may false-positively identify some proper triangles as invalid elements. Nevertheless, we succeeded to resolve these scenarios from Fig. 26, but it took significantly more time for our algorithm to converge. Not that most of these meshes cannot be processed by standard methods because, in these cases, Tutte map tends to produce similar poles with inverted triangles. See Section 8 for our discussion on Tutte embedding and precision lost issues.

Algorithm 4 ABCD with LGB for q measures

Input:

- Source mesh (\mathbf{V}, \mathbf{C}) and vertex initialization \mathbf{x}^0 .
- Distortions $\{\mathcal{D}_j\}_{j=1}^q$ and their gradients $\{\nabla \mathcal{D}_j\}_{j=1}^q$.
- Core solvers $\{S_j\}_{j=1}^q$ with iteration numbers $\{n_j\}_{j=1}^q$.

```

1:  $i \leftarrow 0$ .
2: for  $j = 1, \dots, q$  do
3:    $G_j \leftarrow 1, L_j \leftarrow 0, b_j \leftarrow 1, \text{state}(j) \leftarrow \text{'global'}$ .
4: end for

5: repeat
6:   for  $j = 1, \dots, q$  do  $\{n_j$  times each  $\}$ 
7:      $E_j \leftarrow \mathcal{D}_j$  energy of  $f[\mathbf{x}^i]$  according to (1).
8:     if  $\text{average}(E_j) \leq \min \mathcal{D}_j$  continue.
9:     if  $\text{state}(j) = \text{'local'}$  then
10:       $K \leftarrow L_j$ .
11:     else
12:       $K \leftarrow G_j$ .
13:     end if
14:      $\mathbf{d} \leftarrow$  descent direction of  $E_j$  attained by  $S_j$ .
15:      $\{\mathbf{x}_B^i\}_{B \in \mathbf{B}} \leftarrow \text{BlockPartition}(K, \mathbf{d})$  via Algorithm 1.
16:     for  $B \in \mathbf{B}$  in parallel do
17:        $\mathbf{x}_B^{i+1} \leftarrow$  solution of (2) in  $B$  using  $S_j$  with  $b_j$  max iterations.
18:     end for
19:     if  $\text{state}(j) = \text{'local'}$  then
20:        $(G_j, L_j) \leftarrow \text{Blend}(G_j, L_j)$  via Algorithm 2.
21:       Update  $b_j$ .
22:     end if
23:     Switch  $\text{state}(j)$  between ‘local’ and ‘global’.
24:   end for
25:    $i \leftarrow i + 1$ .
26: until (24) is true for each  $E_j$  or  $i = i_{\max}$ .

```

We found that if the parameter Λ in (13) is slightly increased, then it takes less ABCD iterations to process meshes with poles.

12. Initialization with Decimated UV-Map

We provide some explanations on the *embedding deformation* method [NZ19]. We used this method in Fig. 13 to construct a better starting point for surface parametrization.

Assume that Y is a decimated mesh of a mesh X , \mathbf{y}^0 is the Tutte embedding of Y and \mathbf{y} is a parametrization of Y , computed by initializing isometric optimization with \mathbf{y}^0 . The embedding deformation method is a simple approach aimed at mapping X onto the region contained in $\partial \mathbf{y}$ with only a small number of foldovers. Assume that \mathbf{x}^0 is Tutte embedding of X to a disc $(1 - \epsilon)D$. Then, the embedding deformation method includes the following steps:

1. Construct the simplicial mapping $g : \mathbf{y}^0 \mapsto \mathbf{y}$.
2. For each vertex v of X , find simplex $s(v)$ of Y such that \mathbf{x}_v^0 is contained in the initial mapping of $s(v)$, i.e., \mathbf{x}_v^0 is located in the triangle with coordinates $\mathbf{y}_u^0, u \in s(v)$.
3. Map each vertex v of X to $\mathbf{x}_v = g_{s(v)}(\mathbf{x}_v^0)$.

The above steps are illustrated by Fig. 25. Note that we scale slightly \mathbf{x}^0 by a factor of $(1 - \epsilon)$ to avoid encountering numerical issues on the mesh boundary.

13. GD for Map Fixer Measures

To justify our definitions in (14), assume, first, that we have only a single simplex c and observe how updating \mathbf{x} via GD for distortion \mathcal{D} affects df_c . We can represent df_c as a linear function of singular values σ , acting on its right singular vectors:

$$df_c[\sigma] : \mathbf{v}_j \mapsto \sigma_j \mathbf{u}_j, \quad (30)$$

where \mathbf{u}_j and \mathbf{v}_j are columns of U and V matrices that appear in the signed SVD of df_c , $df_c = U \text{diag}(\sigma_1, \dots, \sigma_m) V^\top$. Since GD modifies df_c non-rigidly, its rotation components U and V are unchanged. Thus, the GD update step has a particularly simple form with respect to (30)

$$\sigma(t + \Delta t) = \sigma(t) - \Delta t \nabla_{\sigma} \mathcal{D}. \quad (31)$$

Choose $\mathcal{D} = \mathcal{D}_{\text{collapse}}$, then $-\nabla \mathcal{D}_{\text{collapse}}$ is a decent direction of $\mathcal{D}_{\text{collapse}}$ and $\sigma(t + \epsilon)$ is its nearest minimum obtained in (31). If a number of simplices c_1, \dots, c_N share a common vertex v , then the contribution of each gradient $\nabla \mathcal{D}_{\text{collapsed}}(\sigma(df_c))$ is summed up to form the energy gradient at v . Hence, repairing collapsed simplices with $\mathcal{D}_{\text{collapse}}$ reduces unnecessary deformation of valid elements, since (31) modifies df_c only on degenerate simplices along their collapsed axes. Similarly, only vertices of inverted elements are modified in GD optimization of $\mathcal{D}_{\text{flip}}$.

14. Algorithm Generalization


In general, our algorithm can optimize an arbitrary list of distortion measures. In particular, ABCD can function well with more than two measures, $\{\mathcal{D}_j\}_{j=1}^q$, if these measures do not contradict each other and there is at least one map fixer and at least one optimizer distortion (10). In this case, we can use Algorithm 4 that presents ABCD generalization for q measures. This algorithm uses the notion of the local and global iterations and other related notations, introduced in Appendix B.

Fig. 20 illustrates an example of running Algorithm 4 with three distortion measures.

15. Configurations for Local-Global Blending

If LGB is enabled in Algorithm 3, then we advise to set an even number of successive iterations n_j for each energy E_j to prevent the global-local cycle of one energy from being interrupted by iterations of another one. These interruptions adversely affect the local-global blending because the mutual impact of different measures can be mixed up in (28), leading to incorrect performance estimates.

We also suggest to start with the ‘‘global-local’’ order in LGB. That is, we suggest to set $K^1 = \mathcal{L}_{\text{max}}$ and $K^2 = \mathcal{L}_{\text{min}}$ in the first call to Algorithm 2. This order of partitioning thresholds leads to better results because at the beginning global optimization tends to deform the entire shape of a target mesh, while local optimization are more concentrated on improving shapes of individual simplices.



#triangles	#invalids		initialization time		timing to fix		iterations to fix		timing total		iterations total	
	BFF	Tutte	BFF	Tutte	BFF	Tutte	BFF	Tutte	BFF	Tutte	BFF	Tutte
32228	0	1	0.498	0.094	0	0.926	0	8	0.805	1.915	11	23
51898	0	10	0.749	0.133	0	0.727	0	3	1.473	7.582	11	54
128912	0	32	1.983	0.384	0	25.31	0	41	3.233	35.87	11	84
515648	0	62	7.37	1.594	0	230.4	0	151	13.5	241.8	11	164

Figure 21: Unconstrained parametrization of Hilbert curve of different resolutions. We use ABCD(PN) to fix inversions in Tutte maps and compare parametrization, initialized by Tutte, with the parametrization initialized by BFF [SC18] conformal mapping. We run these tests with Pardiso library on four-core i7 CPU and report runtimes in seconds.

As a result, the progress of BCD step at the first iteration is often lost at the next step if we start with $K^1 = \mathcal{L}_{\text{min}}$, $K^2 = \mathcal{L}_{\text{max}}$ instead.

16. Quasi-Newton Solvers

For an optimal integration with ABCD algorithm, the parameters for a quasi-newton solver need to be saved when it completes its iterations in the current cycle, and reloaded when we resume its iterations in the next cycle. In this way, one can achieve a better continuity in optimization process between different cycles, which experimentally yields faster convergence. For instance, if S_j is the L-BFGS solver, then the values of \mathbf{x}_B^j and ∇E_B^j should be stored over a number of successive iterations. This data should be restored when S_j resumes its operation, since to write a secant equation one would like to consider the current gradient, and gradients from previous run of the preceding optimization cycle of S_j , rather than use gradients from iterations of a different measure.

17. Note on Enhanced Line Search

We added the following simple modification to the Armijo back tracking line search: we store coordinates $\mathbf{x}'_B(t)$ that induced the highest decrease in block energy (21) during the line search. If the Armijo-Goldstein condition is not satisfied within the maximal number of the line search iterations, then block coordinates for the next iteration are set to $\mathbf{x}'_B(t)$.

This enchantment is aimed at improving map fixer performance by making it less sensitive to line search parameters. In particular, in some configurations, map fixer can repair inversions but, at the same time, it may struggle to attain a sufficient decrease in \mathcal{F} to satisfy the Armijo rule. In these scenarios, our method improves the fixer performance.

Note that core-solver S can modify a solution of the descent equation (7) to get more progress in the line search. We consider

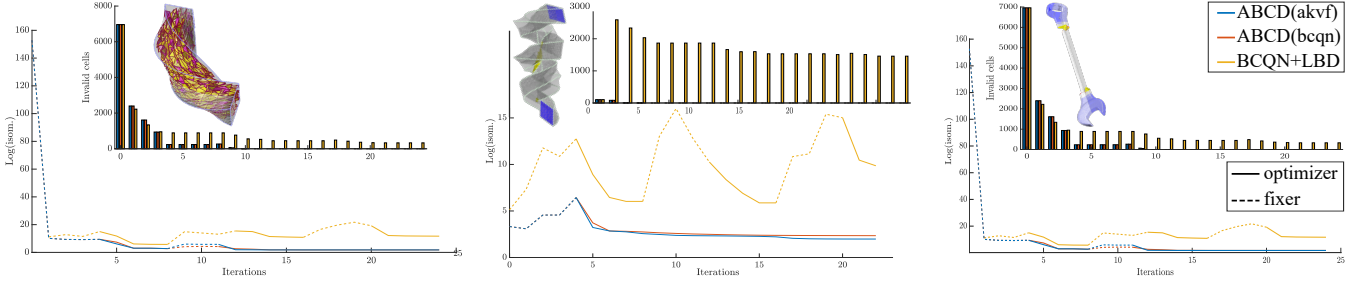


Figure 22: Comparing ABCD(BCQN), ABCD(AKVF) and LBD+BCQN from Fig. 18. Here we plot number of inverted tetrahedrons and isometric energy per iteration in the same way as in Fig. 17.

Table 1: ABCD parameters are specified for each trial, tested in Fig. 24. We use the following notations: b'_j refers to the maximal number of block iterations; ‘Search’ refers to the line search filtering range, described in Section 7.3; $\mathcal{L}'_{\min} = \mathcal{L}_{\min} / \mathcal{L}_{\max}$ and $\mathcal{L}'_{\min} = \dots$ denotes trials in which LGB is disabled.

Trial	Fixer parameters					Optimizer parameters			
	Iterations	Search	LGB	Measure		Iterations	Search	LGB	
	n_j	b'_j	$1 - \delta^-$	\mathcal{L}'_{\min}	Λ	n_j	b'_j	$1 - \delta^+$	\mathcal{L}'_{\min}
1	4	4	1.1	...	0	4	1	0.5	...
2	2	1	1.1	...	0	1	1	0.5	...
3	1	1	1.1	...	0	1	1	0.5	...
4	4	10	1.1	...	0	4	1	0.5	...
5	2	10	1.1	...	0	2	1	0.5	...
6	1	10	1.1	...	0	1	1	0.5	...
7	4	10	1.1	...	0	4	4	0.5	...
8	2	10	1.1	...	0	1	4	0.5	...
9	1	10	1.1	...	0	1	4	0.5	...
10	4	10	1.5	...	0	4	4	0.5	...
11	4	10	2.5	...	0	4	4	0.5	...
12	2	10	10	...	0	2	4	0.5	...
13	4	4	1.1	0.05	0	2	2	0.5	0.05
14	4	4	1.1	0.1	0	2	2	0.5	0.1
15	4	10	1.1	0.2	0	2	2	0.5	0.2
16	4	4	1.1	0.01	0	2	2	0.5	...
17	4	4	1.1	0.1	0	2	2	0.5	...
18	4	4	1.1	0.2	0	2	2	0.5	...
19	4	4	1.1	...	0.001	2	2	0.5	...
20	4	4	1.1	...	0.1	2	2	0.5	...
21	4	4	1.1	...	1	2	2	0.5	...

these modifications as the part of the descent direction computation, employed in ABCD(S) algorithm. For example, the barrier-aware filter on search directions [ZBK18] is the part of BCQN solver, and therefore it is also employed in ABCD(BCQN).

18. Weighted Simplex Displacement Norm

We found that using the coordinates of simplex centers in (19), instead of vertices, yields better scalability, since $\Delta \mathbf{x}^i$ is more likely to accumulate numerical errors. We define a simplex-based displacement norm as follows:

$$\text{Disp}_{\mathbf{C}}^i = \frac{1}{\bar{\epsilon}} \|\text{diag}(\mathbf{w})(\Delta \mathbf{C}^i - \text{Proj}_{\text{Ker}(\mathbf{C}^{i-1})}(\Delta \mathbf{C}^i))\|_{\text{Fro}}, \quad (32)$$

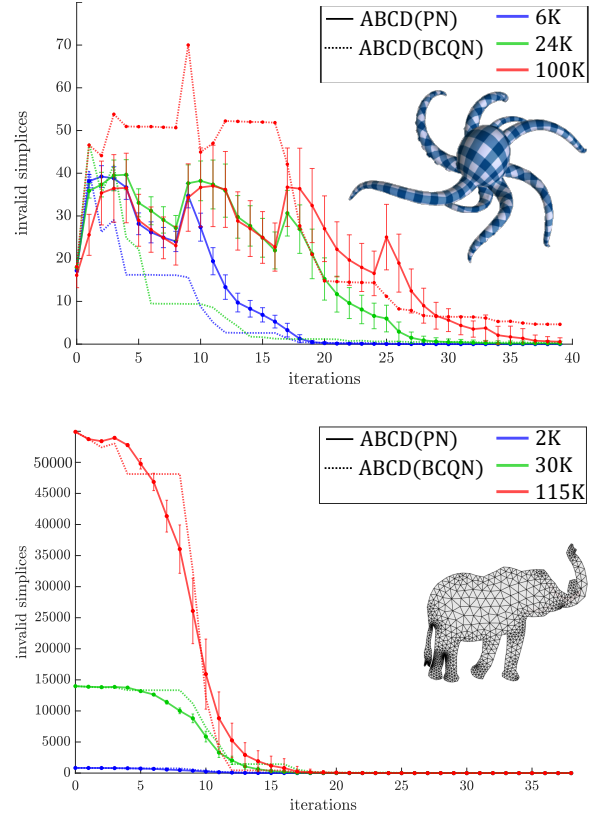


Figure 23: Showing an average number of invalid triangles per iteration, obtained in the experiments from Fig. 11. Vertical bars show the standard deviation.

where \mathbf{w} is the stack of weights $w(c)$ from (1), $\mathbf{C}^i \in \mathbb{R}^{|\mathbf{C}| \times m}$ are simplex target centers in i^{th} iteration, $\Delta \mathbf{C}^i = \mathbf{C}^i - \mathbf{C}^{i-1}$ and $\bar{\epsilon}$ is an approximate average length of edges in the target domain.

