

DA Wand: Distortion-Aware Selection using Neural Mesh Parameterization

Richard Liu
University of Chicago

Noam Aigerman
Adobe Research

Vladimir G. Kim
Adobe Research

Rana Hanocka
University of Chicago



Figure 1. DA Wand enables interactive decaling through a conditional selection of local *distortion-aware* patches. Our method identifies large patches within both developable and high-curvature regions which induce a low distortion parameterization.

Abstract

We present a neural technique for learning to select a local sub-region around a point which can be used for mesh parameterization. The motivation for our framework is driven by interactive workflows used for decaling, texturing, or painting on surfaces. Our key idea is to incorporate segmentation probabilities as weights of a classical parameterization method, implemented as a novel differentiable parameterization layer within a neural network framework. We train a segmentation network to select 3D regions that are parameterized into 2D and penalized by the resulting distortion, giving rise to segmentations which are distortion-aware. Following training, a user can use our system to interactively select a point on the mesh and obtain a large, meaningful region around the selection which induces a low-distortion parameterization. Our code¹ and project² are publicly available.

1. Introduction

Many interactive workflows for decaling, texturing, or painting on a 3D mesh require extracting a large surface patch around a point that can be mapped to the 2D plane with low distortion. Unlike *global* parameterization approaches that map the entire mesh to 2D while introducing as few cuts as possible [1, 7–9, 17, 35, 38, 41, 43, 44, 53, 54, 56]), this work focuses on segmenting a *local sub-region* around a point of interest on a mesh for parameterization [25, 36, 45, 46, 60, 69, 70, 73]. Local parameterizations are advantageous in certain modeling settings because they are in-

herently user-interactive, can achieve lower distortion than their global counterparts, and are computationally more efficient. To date, however, techniques for extracting a surface patch that is amenable to local parameterization have largely relied on heuristics balancing various compactness, patch size, and developability priors [18].

This work instead takes a data-driven approach to learn *distortion-aware* local segmentations that are optimal for local parameterization. Our proposed framework uses a novel *differentiable parameterization layer* to predict a patch around a point and its corresponding UV map. This enables *self-supervised* training, in which our network is encouraged to predict area-maximizing and distortion-minimizing patches through a series of carefully constructed priors, allowing us to sidestep the scarcity of parameterization-labeled datasets.

We name our system the **Distortion-Aware Wand (DA Wand)**, which given an input mesh and initial triangle selection, outputs soft segmentation probabilities. We incorporate these probabilities into our parameterization layer by devising a weighted version of a classical parameterization method - *LSCM* [35] - which we call *wLSCM*. This adaptation gives rise to a *probability-guided parameterization*, over which a distortion energy can be computed to enable *self-supervised* training. We prove a theorem stating that the *wLSCM* UV map converges to the *LSCM* UV map as the soft probabilities converge to a binary segmentation mask, which establishes the direct relation between probabilities and binary segmentation in the parameterization context.

Reducing the distortion of the UV map and maximizing the segmentation area are competing objectives, as UV distortion scales monotonically with patch size. Naively summing these objectives leads to poor optimization with

¹<https://github.com/threedle/DA-Wand>

²<https://threedle.github.io/DA-Wand/>

undesirable local minima. We harmonize these objectives by devising a novel *thresholded-distortion loss*, which penalizes triangles with distortion above some user-prescribed threshold. We additionally encourage compactness through a *smoothness loss* inspired by the graphcuts algorithm [26].

We design a novel near-developable segmentation dataset to initialize the weights of our segmentation network, with an automatic generation algorithm which can be run out of the box. We then train this network end-to-end on a dataset of unlabelled natural shapes using our parameterization layer with distortion and compactness priors.

We leverage a MeshCNN [15] backbone to learn directly on the input triangulation which enables sensitivity to sharp features and a large receptive field which enables patch growth. Moreover, by utilizing *intrinsic* mesh features as input, our system remains invariant to rigid-transformations.

DA Wand allows a user to interactively select a triangle on the mesh and obtain a large, meaningful region around the selection which can be UV parameterized with low distortion. We show that the neural network is able to leverage global information to extend the segmentation with minimal distortion gain, in contrast to existing heuristic methods which stop at the boundaries of high curvature regions. Our method can produce user-conditioned segmentations at interactive rates, beating out alternative techniques. We demonstrate a compelling interactive application of DA Wand in Fig. 1, in which different regions on the sorting hat mesh are iteratively selected and decaled. We show additional example textures in the supplemental.

2. Related Work

2.1. Global Cutting & Parameterization

Surface parameterization has been a long-standing geometry processing problem. An extensive line of work has dealt with *global* surface parameterization, in which a 3-dimensional surface, discretized as a triangle mesh with fixed connectivity, is mapped to the plane [1, 7–9, 17, 35, 38, 41, 43, 44, 53, 54, 56]). Most of these methods assume a disk topology surface, and some generate seams on a heuristic basis to ensure this. A few seamless approaches exist that instead introduce point singularities. [11, 22, 30, 39, 58].

Though most parameterization works assume a disk topology input, the target domain for many texturing and fabrication applications is often a closed mesh without boundaries, and segmenting a mesh into topological disks in an optimal way is a non-trivial problem. Thus, many works have focused on the sub-problem of cutting a closed surface into disks with desirable properties for the target application – typically developability and minimal-length seams [13, 18, 34, 35, 37, 48, 50, 52, 63, 71]. Some works jointly optimize for the parameterization and cutting objectives [31, 40], taking advantage of the close relationship be-

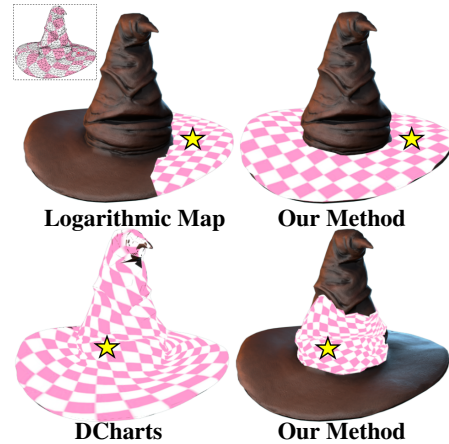


Figure 2. Segmentations from our baseline results conditioned on the same selections (stars) as our method in Fig. 1, with resulting UV textures. To produce the Logarithmic Map segmentation, we select a low-distortion patch from the full map (dashed inset).

tween seam placement and parameterization distortion.

2.2. Local Cutting & Parameterization

A few past works have examined the *local* problem of segmenting and parameterizing a surface patch which encompasses a point of interest. Exponential Maps (more accurately referred by its inverse *logarithmic map*, which is how we will refer to this method moving forward) is a cornerstone work in this domain [46], and subsequent local parameterization works have primarily extended the insights of the original method [36, 45, 60, 61, 69, 70, 73]. All of the local parameterization works to date either do not rely on a local segmentation of the domain, grow out a fixed-radius geodesic circle, or apply some form of local shape analysis to iteratively floodfill the selection. This last class of chart growing techniques based on surface analysis is also performed by some global mesh decomposition methods, such as DCharts [18], which makes these methods amenable for adaptation to the local segmentation setting.

Our work is primarily interested in this *local segmentation for parameterization* problem. To our knowledge, however, our work is the first to (1) make explicit the objective of a *minimal-distortion, large* local segmentation and (2) apply data-driven techniques to solve this problem. Our approach overcomes the limitations of the aforementioned heuristic methods, which preclude their use in certain texturing applications, shown in Fig. 2.

2.3. Mesh Segmentation

Our work also has a foot in the space of general mesh segmentation techniques, in which a mesh is partitioned into disjoint parts according to some set of criteria. One prominent class of segmentation techniques involves *semantic segmentation*, where a mesh is decomposed into meaning-

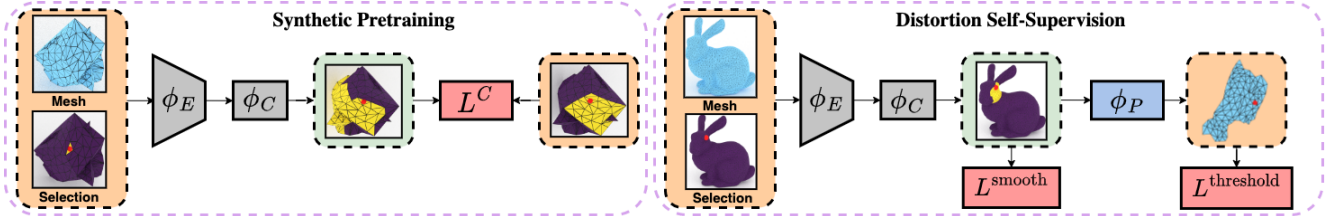


Figure 3. **Overview.** The encoder ϕ_E maps a **mesh** and a **selection triangle** to edge features which are passed into ϕ_C to output a local **segmentation**. The training is split into two phases. We have a **pretrain phase** where we jointly train our encoder and classifier network on a synthetic dataset of shapes with ground truth segmentations and randomly sampled initial selections. We supervise with cross entropy loss L^C . In our main **distortion self-supervision phase**, we finetune ϕ_E and ϕ_C on a dataset of natural shapes, where *no ground truth segmentations exist*. We introduce a differentiable parameterization layer ϕ_P which computes a **weighted UV embedding** for the network-predicted segmentation on the fly during training. The UV map is sent into a distortion function of choice D , which is used in the thresholded-distortion loss $L^{\text{threshold}}$. The smoothness loss L^{smooth} enforces a compactness prior on the network prediction, which helps curtail non-contiguous segmentations.

ful parts [3, 5, 12, 14, 21, 27–29, 55, 62, 65, 66, 68, 72, 74, 75]. Another prominent class tackles the disk topology objective discussed in the previous section [13, 18, 35, 52], [34, 37, 48, 50, 63, 71]. A closely related set of methods perform sharp feature identification, where the surface geometry is analyzed to identify boundaries of minimal principal curvature, and the mesh is segmented along these boundaries [2, 4, 10, 19, 20, 23, 24, 32, 64, 72, 77]. Recently, learning-based methods have been introduced for each of these objectives [14–16, 21, 27, 47, 66, 67], but to our knowledge there does not exist a learning-based work that targets our objective of local distortion-aware segmentation.

3. Method

We assume as input a triangle mesh M with faces F and vertices V , and a selected triangle t^* . Our goal is to define a binary segmentation of the mesh faces \bar{F} which contains t^* . Additionally, \bar{F} needs to satisfy three competing objectives: i) enclose a region that can be UV-mapped with low distortion; ii) cover an as-large-as-possible mesh area; iii) have a compact boundary, which we define as a boundary enclosing a single connected region on the mesh surface. We denote by $U \in \mathbb{R}^{|V| \times 2}$ the computed UV mapping assigning a 2D coordinate to each vertex.

Our framework predicts a *soft* segmentation, in which each triangle is assigned a probability designating whether it belongs to the segmented region. The key contribution that drives our framework is a *differentiable parameterization layer* (Sec 3.1), which computes a UV parameterization conditioned on the segmentation probabilities. To train this architecture we introduce novel distortion and compactness objectives favoring large contiguous segments with low-distortion UV parameterizations (Sec 3.2).

Training is conducted in two stages: i) in pre-training, we leverage a novel synthetic dataset (Sec 3.3), with a ground truth decomposition into near-developable regions, and train our segmentation network with strong supervi-

sion. ii) we then train the entire system end-to-end with distortion self-supervision (enabled by the parameterization layer) on a dataset of unlabelled natural shapes (Sec 3.4). An overview of the system is shown in Fig. 3.

3.1. Probability-Guided Parameterization

Our goal with the parameterization layer is to devise a fast, differentiable mapping from the segmentation W to a low-distortion UV map. To predict the segmentation, we leverage MeshCNN, a robust mesh segmentation architecture [15]. The network predicts a soft segmentation W which assigns a weight w_t to each triangle t . This soft segmentation can be clamped to a binary segmentation (e.g. through rounding) to define a subregion \bar{F} of the mesh. The challenge is then, how to devise a parameterization that accounts for the soft segmentation and smoothly approximates the parameterization of the subregion \bar{F} ?

Our main observation is that we can consider a classic parameterization technique - Least Square Conformal Maps (LSCM) [35] - and define a *weighted* version of it, dubbed $wLSCM$, such that as the weights approach binary values $\{0, 1\}$ which define a hard segmentation \bar{F} , $wLSCM$ converges to the (non-weighted) LSCM parameterization of \bar{F} . We achieve this such that $wLSCM$ can still be implemented as a straightforward differentiable layer.

The original LSCM method considers the vertices of each triangle (v_1, v_2, v_3) , and their image under the (unknown) UV map, (u_1, u_2, u_3) , and considers the affine transformation representing that map, which satisfies

$$A_t v_i + \delta_t = u_i, \quad i \in \{1, 2, 3\}, \quad (1)$$

where $A_t \in \mathbb{R}^{2 \times 2}$, $\delta_t \in \mathbb{R}^2$ are the linear and translational components of the map, respectively (see [35]).

LSCM then aims to minimize *conformal error*, which measures the change of the triangle angles resulting from the map. It considers the closest similarity matrix $S(A_t)$

(a matrix with no conformal error), and measures the least-squares error between A_t and the similarity matrix:

$$E_{\text{LSCM}} = \sum_{t \in \mathbf{F}} \|A_t - S(A_t)\|^2. \quad (2)$$

LSCM computes the UV parameterization \mathbf{U} which minimizes this energy, which reduces to solving a sparse linear system. We propose *relaxing* the segmentation problem by adding a weighting term using the weights W predicted by the segmentation network:

$$E_{\text{wLSCM}} = \sum_{t \in \mathbf{F}} w_t \|A_t - S(A_t)\|^2. \quad (3)$$

The minimizer of this energy can still be obtained via a linear solve. Hence, our differentiable parameterization layer consists of plugging the predicted soft segmentation W into the wLSCM energy (Eq. (3)), and solving the linear system to get the UV parameterization \mathbf{U} . Fig. 4 illustrates how weights affect the parameterization. Importantly, as the weights converge to a binary segmentation, the resulting wLSCM UV parameterization converges to the LSCM parameterization of the segmented region:

Theorem 1 *Let $\bar{\mathbf{F}} \subset \mathbf{F}$ be a subset of triangles of the mesh, which comprises one connected component. Let W be non-negative weights assigned to the triangles s.t. the weights of $\bar{\mathbf{F}}$ are non-zero. Let \mathbf{U}_W be the minimizer of Eq. (3) w.r.t W . Then \mathbf{U}_W , restricted to $\bar{\mathbf{F}}$, is a well-defined, continuous function of W . Furthermore, if the non-zero weights of W are all equal to 1, then \mathbf{U}_W restricted to $\bar{\mathbf{F}}$ is exactly equal to the (non-weighted) LSCM parameterization of $\bar{\mathbf{F}}$.*

The theorem has a straightforward proof via continuity of polynomial roots. See the appendix for the full proof.

Segmentation network architecture. For the soft segmentation prediction, we adapt a 12-layer (6 down, 6 up) MeshCNN [15] to define an edge encoder, ϕ_E . Each layer has 3 residual convolution blocks with hidden dimension 16. This fully-convolutional backbone enables a large receptive field with relatively few parameters, while the graph-based convolutions allow for the fine-grained edge understanding necessary to segment along noisy geometric features. We do not incorporate pooling in the lower layers, as we find it oversmooths the learned features. ϕ_E takes as input batched pairs of meshes and point selections (M, x) , and computes dihedral angles, the heat kernel signature (HKS) [59], and a one-hot vector of selection triangle incidence as edge features.

A second per-triangle classifier ϕ_C predicts the per-triangle weights W . It is implemented as a 3-layer multi-layer perceptron (MLP) with dropout and a final sigmoid activation layer to map scalars to soft probabilities.

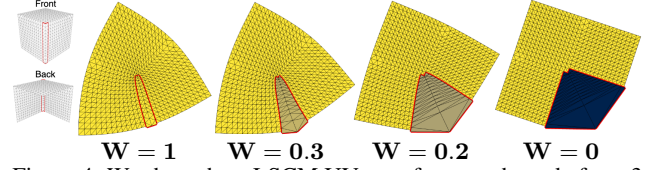


Figure 4. We show the wLSCM UV map for a mesh made from 3 sides of a cube, with decreasing weights W for the region outlined in red. The wLSCM UV map of the non-outlined region when $W=0$ is exactly the UV map of unweighted LSCM over the same region, as proved in Theorem 1.

3.2. Distortion and Compactness Priors

To train the network, we carefully design losses which balance across our three competing patch objectives of size, distortion, and compactness.

Thresholded-distortion loss. We aim to produce a UV parameterization with low isometric distortion, defined via the As-Rigid-As-Possible (ARAP) energy (D_{arap}) [33].

$$D_{\text{arap}}(x_t) = \sum_{i=0}^2 \cot(\theta_t^i) \|(x_t^i - x_t^{i+1}) - L_t(v_t^i - v_t^{i+1})\|^2 \quad (4)$$

where v_t^i is the i -th vertex of triangle t in local coordinates, θ_t^i is the angle opposite the edge (x_t^i, x_t^{i+1}) , and L_t is the “best-fit” rotation matrix as described in [33].

Typically the distortion of a UV parameterization grows monotonically with the size of the selection area. This puts in direct contrast our objectives of large patch area and low-distortion parameterization. We resolve this by formulating an objective that aims to maximize the number of triangles with distortion under a *threshold*. We write this objective in terms of a novel *thresholded-distortion loss*:

$$L^{\text{threshold}}(D) = \sum_t \frac{A_t}{\sum_t A_t} \left(1 - e^{-(D_t/\gamma)^\alpha}\right), \quad (5)$$

where A_t is the area of triangle t , γ acts as a soft threshold on the per-triangle distortion D_t , and α is a hyperparameter which controls the sensitivity to the threshold. This loss acts as a weighted count of triangles with distortion below γ . In all our experiments, we set $\gamma = 0.01$ and $\alpha = 5$.

Compactness priors. In order to achieve our last objective of a compact segmentation boundary, we introduce two priors into network training. First, we introduce a *smoothness prior*, in the form of a *smoothness loss* derived from the well-known graphcuts algorithm [26].

$$L^{\text{smooth}} = \frac{1}{|\mathbf{E}|} \sum_{t_1, t_2 \in \mathbf{E}} -\omega * \log\left(\frac{\theta_{t_1, t_2}}{\pi}\right) |w_{t_1} - w_{t_2}| \quad (6)$$

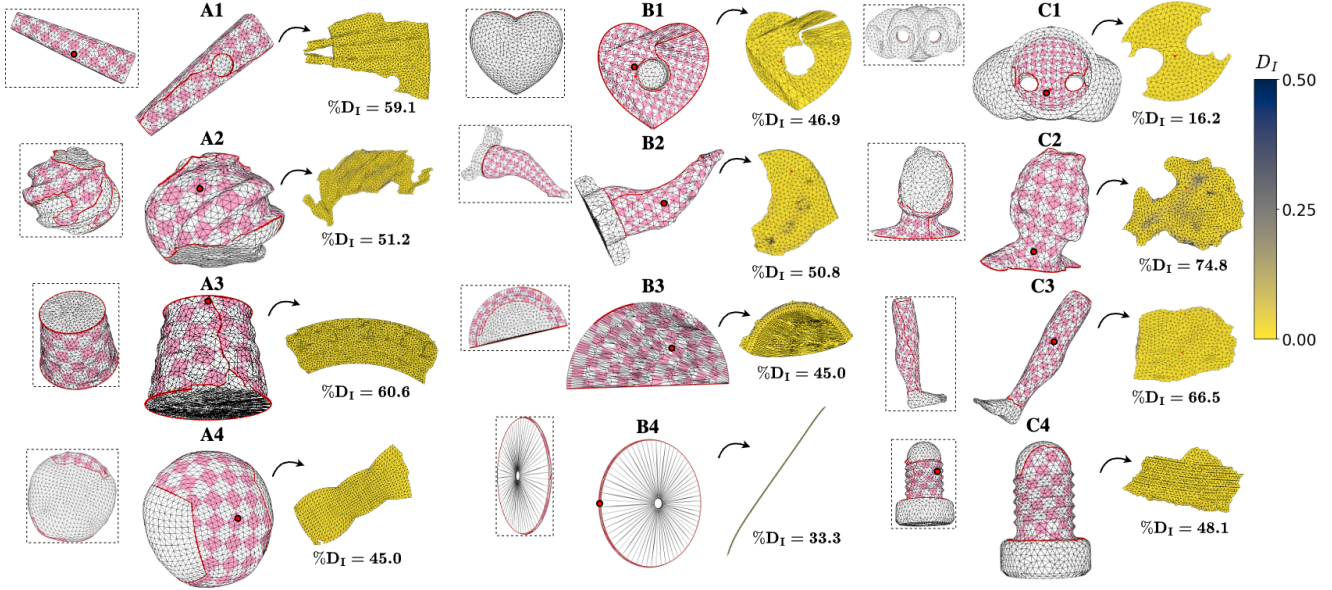


Figure 5. Conditional segmentations predicted by our method. We show the corresponding UVs computed with SLIM, colored by the isometric distortion metric I . Selection points are marked in red, and $\%D_I$ is reported beneath each UV.

where \mathbf{E} is the set of mesh edges $\mathbf{E} \subset \mathbf{F} \times \mathbf{F}$, θ_{t_1, t_2} is the dihedral angle between triangles t_1, t_2 which share an edge, and ω is an energy scaling term. In all our experiments, we set $\omega = 0.1$.

Lastly, LSCM does not perform as well with disconnected domains. To guarantee the segmentation boundary encloses a single connected region, we add a *floodfill prior* during training, in which we identify all triangles with probabilities ≥ 0.5 , and take the largest contiguous patch of these triangles which contains the source triangle t^* . We mask out the weights of the other triangles, generating the floodfill weights w^* . We pass the floodfill weights into the parameterization layer and pass the resulting UV into our threshold loss. In order to discourage spurious weights outside of the compact segmentation boundary, we penalize the original probabilities by passing them into our smoothness loss $L^{\text{smooth}}(w)$. Incorporating all of these terms, our final loss function during end-to-end training is

$$L = L^{\text{threshold}}(w^*) + L^{\text{smooth}}(w) \quad (7)$$

3.3. Near-Developable Segmentation Pretraining

In practice we need to initialize our system to produce reasonable segmentations which can be further improved through our end-to-end training. To this end, we pretrain our network to segment developable (flattenable with no distortion) patches with boundaries at sharp creases.

As there is no existing dataset of meshes with labeled distortion-aware decompositions, we build a synthetic one. We consider a collection of geometry primitives – cones, cubes, cylinders, spheres, and tetrahedrons, with known de-

velopable decompositions (excluding the sphere). To generate training shapes we sample a subset of primitives, apply additional non-rigid deformations (twist, bend, taper, stretch), and compute a CSG union on these perturbed primitives to obtain a watertight mesh. For each mesh, we generate an additional augmented mesh by sampling 10% of the vertices and jittering them along their normals. We apply Laplacian smoothing to the resulting mesh. See the supplemental for more details and pseudocode.

For each of our geometry primitives, we take their natural developable decomposition as the starting ground truth (i.e. cubes are segmented into planes, cylinders into the body and two caps, etc.). To create ground truth segmentation labels we retain correspondences to the original primitives, and use them to segment the perturbed mesh. Note, however, that our augmentations do not preserve developability, though developable segmentation is typically too strong of a constraint for practical applications outside of fabrication. Instead we ensure *near-developability* of target segmentations by parameterizing each augmented segmentation with SLIM [42], and discarding all patches with isometric or conformal distortion above a threshold of 0.05. Our measures of isometric and conformal distortion are $D_I = \max(\sigma_1, 1/\sigma_2)$ and $D_C = (\sigma_1 - \sigma_2)^2$, respectively, where σ_1 and σ_2 are the singular values of the Jacobian of the UV mapping in decreasing order.

To generate training samples we simulate user selection by randomly sampling 5% of the triangles (up to a maximum of 20) on valid near-developable patches (excluding regions corresponding to spheres), to use as initial selec-

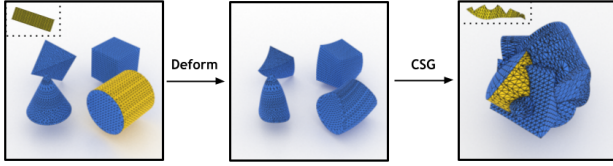


Figure 6. Generation of our “near-developable” synthetic dataset: we draw random geometric primitives, deform them, and apply a CSG union to obtain our final synthetic shape, while transferring the correspondence to the original segmentation (yellow). UVs of the corresponding segmentations are shown in the insets.

tions t^* . The dataset construction procedure is illustrated in Fig. 6. Our training dataset consist of 87 meshes with 7,912 initial selections, and our test set consists of 17 meshes with 1,532 initial selections. We pretrain our edge encoder network ϕ_E and classifier network ϕ_C with L2 supervision using Adam with learning rate 10^{-3} for 150 epochs.

3.4. Distortion Self-Supervised Training

In the second phase of training, we train the system end-to-end with the differentiable parameterization layer over a dataset of natural shapes taken from the Thingi10k dataset [76]. We download a subset of meshes from Thingi10k, filtering for meshes which are between $3K$ and $20K$ faces, are manifold, and have one connected component. We remesh the shapes using the isotropic explicit remeshing tool from Meshlab [6] so that the meshes have between $5K$ and $12K$ edges. For each remeshed object, we sample 1% of the surface triangles to use as training anchor samples, up to a maximum of 50 faces. Our resulting training set consists of 315 meshes with 6,300 initial selections, and a validation set containing 86 meshes with 1,720 initial selections.

Similar to the pretraining phase, we batch input pairs of meshes and initial triangle selections. To avoid distribution shift, we additionally draw training samples from the synthetic dataset 50% of the time and apply strong supervision to those predictions. We perform end-to-end training using Adam with learning rate 10^{-3} for 100 epochs.

3.5. Inference and Patch Postprocessing

During inference, we compute UV maps with SLIM [42] instead of LSCM, as the former directly minimizes isometric distortion. Thus the isometric distortion of the LSCM UV map is an *upper bound* of the isometric distortion achieved by the SLIM UV map. We avoid using SLIM in the parameterization layer during training as it is an iterative and heavy optimization technique. However, we consider this to be interesting follow-up work. As SLIM assumes a disk-topology patch, we run floodfill as before (Sec. 3.2) to get a single connected component, and apply a seam-cutting algorithm by iteratively choosing the longest boundary loop

and its nearest boundary loop, and cutting between them along the shortest path until the patch is disk topology.

Lastly, in order to straighten the segmentation’s boundary and reduce jagged features, we run the graphcuts algorithm [26] over our model predictions, where the graph’s node potentials are set to W and the edge weights are set to $-\log(\frac{\theta}{\pi})$ where θ is the dihedral angle between two faces.

4. Experiments

In this section we quantitatively and qualitatively evaluate the distortion-aware selections produced by our method and a few baselines. Since some methods only focus on segmentation heuristics, we post-process all segmentations using SLIM [42] to get the resulting UV parameterizations.

We describe our evaluation metrics and datasets (Sec. 4.1), discuss segmentations produced with our method (Sec. 4.1), compare to baseline alternatives (Sec. 4.3), and run ablations (Sec. 4.4).

4.1. Experimental Setup

Datasets. We evaluate our method and baselines on three datasets. First, we use our Near-Developable Segmentation Dataset, as it is the only dataset that has ground truth parameterization-aware segmentations. We evaluate on a held-out test set composed of 47 meshes and 1,532 (face, segmentation) samples. Second, we create a dataset of natural shapes from Thingi10k [76] models, filtered to have medium face count, manifoldness and a single connected component, pre-processed as described in Sec. 3.4, leading to 110 test shapes with 2,125 samples (in addition to the 401 shapes used during training). These models are mostly originally designed for 3D printing, and cover a wide range of geometric and topological features associated with real objects. Third we use the Mesh Parameterization Benchmark dataset [51] composed of 337 3D models originally designed for digital art. We filter for meshes in the same way as with Thingi10k, which leaves us with 32 models and 269 sampled points. While this data is annotated with ground truth artist UV segmentations, global parameterization is a different objective from distortion-aware local segmentation (as discussed in Sec. 2), so we reserve the comparison to the supplemental.

Evaluation Metrics. Our work is motivated by applications that require maximizing patch area, while keeping distortion under some user-prescribed bound. Our Near-Developable Segmentation Dataset is the only dataset with relevant labels, for which we report standard metrics – accuracy, mean average precision and F1 score.

When ground truth segmentations are not available, we propose a novel metric $\%D_I^\lambda$, which we define as the percentage of the faces parameterized with isometric distortion below threshold λ : $D_I^\lambda = |\{t : D_I(t) < \lambda\}|/|\mathbf{F}|$, where the

numerator enumerates faces t with distortion below λ , normalized by the total faces in \mathbf{F} . We use the following energy for our isometry measure $D_I = (\max(\sigma_1, \frac{1}{\sigma_2}) - 1)^2$, where (σ_1, σ_2) are the singular values of the Jacobian of the UV parameterization [57]. We set $\lambda = 0.05$, as we find that this distortion level leads to no significant visual distortion. A higher percentage $\%D_I^\lambda$ indicates a larger, more distortion-aware segmentation. We report $\%D_I^\lambda$ for varying λ in the supplemental to show that our conclusions are robust.

4.2. DA Wand Results

We run our method on all three datasets and show examples in Fig. 5. Our method handles meshes with complex geometric features and topological variations, consistently obtaining a large patch which can be parameterized with little distortion. We handle near-developable surfaces of mechanical parts by splitting them into correct primitives such as cylinders, planes, and annuli, even in presence of high frequency noise (A3, B1, B3, C1, C4). Global analysis enabled by MeshCNN allows our method to extend segments far beyond the selection (A1, B2, B4, C3). Our method also discovers large low-distortion patches on organic shapes with no apparent decomposition into primitives (A2, A4, C2), thanks to distortion-guided training.

4.3. Baselines

	Acc \uparrow	mAP \uparrow	F1 \uparrow
LogMap	0.54	0.19	0.24
DCharts	0.85	0.37	0.47
DA Wand (Ours)	0.91	0.8	0.82

Table 1. Accuracy (Acc), Mean average precision (mAP), and F1 score (F1) reported for each respective method on the ground truth segmentations in our synthetic test dataset.

Though none of the prior methods discussed in Sec. 2 directly address the distortion-aware selection problem, we identify two methods that can be adapted to our problem.

First is the quasi-developable segmentation technique, DCharts, originally devised for global parameterization-aware segmentation [18]. The original method uses Lloyd iterations to optimize over multiple charts, which we adapt to single-chart segmentation conditioned on an initial triangle. Starting with the selected triangle, we greedily add triangles with the smallest DCharts energy, defined at each triangle with respect to our patch as the product of the fitting (developability), compactness, and straight boundary energies. As in the original work, we use the same weights ($\alpha = 1.0, \beta = 0.7, \gamma = 0.5$), and in every iteration grow out the patch by adding adjacent triangles with fitting error under the threshold $F_{max} = 0.2$. We maintain a priority queue of triangles adjacent to the chart boundary based on the DCharts energy, and recompute the fitting energy af-

	Synthetic		Thing10k		Mesh Param.		Time (s)
	$\%D_I \uparrow$	N	$\%D_I \uparrow$	N	$\%D_I \uparrow$	N	$t_{seg} \downarrow$
LogMap	6.5	236	7.5	266	5.0	179	0.15
DCharts	14.3	676	15.1	770	1.6	2754	32.92
DA Wand (Ours)	19.9	951	19.9	1031	14.3	1281	0.11

Table 2. Median $\%D_I$ and N reported for each method across the three benchmark datasets. We also report the median time to segmentation (t_{seg}) in seconds across all our test samples.

ter popping a triangle from the queue to ensure it meets the threshold for the current patch. Following the original work, after each iteration (when the queue becomes empty) we re-estimate the developability proxy vector N_C and angle θ , which are used for the fitting energy, and re-initialize the queue. We terminate once there are no more valid triangles even after re-estimating the developability proxy.

Second, we adapt a local parameterization technique to guide segmentation. Specifically, we run a version of the logarithmic map implemented with the Vector Heat Method [49] which computes a local parameterization of the whole mesh centered at the selected triangle. We mark all triangles that are parameterized with isometric distortion D_I below 0.05 (our evaluation threshold) as valid, and use the largest contiguous patch of valid triangles containing the selection as the segmentation.

For consistency, we postprocess all baseline results with the graphcuts [26] and floodfill algorithms as in Sec. 3.5. The resulting segmentation is guaranteed to have a disk topology, which we parameterize using SLIM [42].

While D-Charts leverages similar developability and compactness priors to our work, these priors are not context-aware, as they only incorporate local patch and triangle information, and the relative importance of these priors have to be manually tuned. We illustrate this in Fig. 7, where the parameters need to be significantly altered for DCharts to produce the expected segmentation on a simple cylinder body. However, these same parameter settings swings the bias in the other direction, so that it segments only the curved region in the ashtray mesh from Thing10k (second row) while ignoring the large planar patch right next to it. By virtue of its learned inductive bias over many types of surfaces, our method can identify both segmentations without any additional parameter tuning or training.

Our other baseline derives segmentations with the logarithmic map, which makes them trivially distortion-aware. A key limitation, however, is that the logarithmic map parameterizes the entire surface, so cannot ignore regions that introduce heavy distortion (e.g., high-curvature geometry near the selection). We illustrate this in Fig. 8, where on a simple cylinder, the Exponential Maps baseline segmentation creeps into the top cap, which introduces enough distortion to prevent the segmentation from reaching the far end of the cylinder body. On the other hand, our method iden-

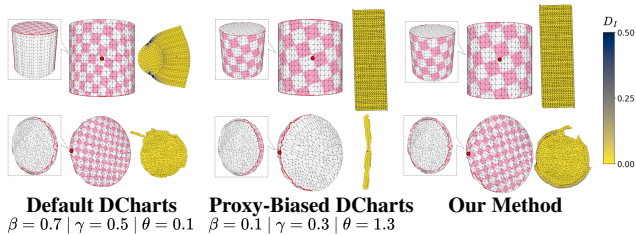


Figure 7. Segmentations predicted by DCharts with default parameters (**Default DCharts**), DCharts biased towards the developability proxy energy (**Proxy-Biased DCharts**), and DA Wand (**Our Method**). The parameter values are given for each DCharts configuration. We show results on a cylinder mesh (top row), and an ashtray mesh from our Thingi10k dataset (bottom row). UV embeddings on the right are colored by isometric distortion D_I .

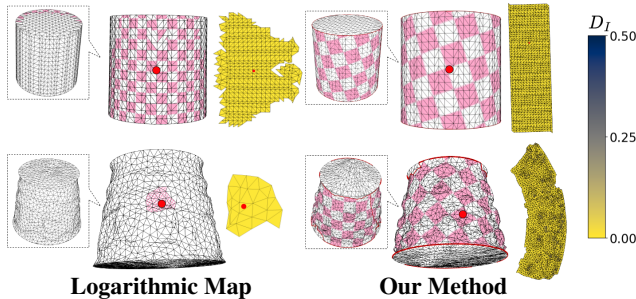


Figure 8. Segmentations predicted on a cylinder mesh (top row) and a noisy cylinder from our Thingi10k dataset (bottom row), for the Logarithmic Map baseline and our method. UV embeddings on the right are colored by isometric distortion D_I .

ties the ideal cylindrical segmentation on both the simple cylinder and the noisy sample from the Thingi10k dataset.

Tab. 1 reports the classification performance of each method with respect to the ground truth segmentations for our synthetic test set. Unsurprisingly, our method outperforms the baselines on the segmentation metrics. We report our thresholded-distortion metric $\%D_I^{0.05}$ on all three datasets in Tab. 2. We also report the number of segmented triangles (N) and time to segmentation (t_{seg}) in seconds. Our method outperforms the baselines by a wide margin on all three datasets. Our segmentation time is 25% faster than the logarithmic maps baseline (LogMap) and several orders of magnitude faster than DCharts. Please see the supplemental for additional qualitative comparisons and time analysis.

4.4. Ablations

We justify our design choices through a series of ablations (Tab. 3). Training our method without the compactness priors (b, c) leads to larger but higher distortion patches. Note that the pre-training segmentations (d) are nearly half the size of our final segmentations. On the other hand, removing the pre-training step (f) or not incorporating synthetic dataset samples during distortion self-supervision training (e) significantly degrades the results, which illus-

	Synthetic		Thingi10k		Mesh Param.	
	$\%D_I \uparrow$	N	$\%D_I \uparrow$	N	$\%D_I \uparrow$	N
Our Method (a)	19.9	951	19.9	1031	14.3	1281
No floodfill (b)	15.4	1056	16.5	1113	12.3	997
No smoothness loss (c)	10.1	1405	11.6	1500	6.8	1220
No distortion training (d)	10.1	478.5	10.0	647	3.5	647
No mixed training (e)	4.5	2125	10.2	1729.2	2.3	886
No pre-training (f)	4.3	2561	7.0	2537	1.2	1512

Table 3. Median $\%D_I$ and N for our method and different ablation settings across our three benchmark datasets.

trates the importance of this initialization step.

5. Discussion and Future Work

We presented a framework which enables training a neural network to learn a data-driven prior for distortion-aware mesh segmentation. Key to our method is a *probability-guided* mesh parameterization layer, which is differentiable, fast, and theoretically grounded. Our wLSCM formulation can be viewed as a way to *relax* the problem of segmenting a local domain which induces a low distortion parameterization. Our experiments demonstrate that our method is an effective technique for producing distortion-aware local segmentations, conditioned on a user-selected triangle.

Our method holds a few limitations. First, the use of MeshCNN [15] leads to dependence on the triangulation, in the form of a limited receptive field. Second, the output of our method is not guaranteed to be disk topology (necessary to produce a valid UV map), and is post-processed using a simple cutting algorithm to achieve this property.

We see many future directions for our work. Instead of using distortion as guidance, other differentiable quantities, such as curvature or visual losses, can be used to guide our segmentation through similar custom differentiable layers. Additionally, we wish to extend our work to predict a segmentation through edges instead of faces, which would enable seam generation within the segmented region.

6. Acknowledgements

This work was supported in part by gifts from Adobe Research. We are grateful for the AI cluster resources, services, and staff expertise at the University of Chicago. We thank Nicholas Sharp for his assistance with the Polyscope demo and Logarithmic Map code, Georgia Shay for her assistance with the parameterization benchmark, and Vincent Lagrassa for his research assistance. We thank Haochen Wang and the members of 3DL for their helpful comments, suggestions, and insightful discussions.

References

- [1] Noam Aigerman and Yaron Lipman. Orbifold Tutte embeddings. *ACM Transactions on Graphics*, 34(6):190:1–190:12, Oct. 2015. 1, 2

- [2] Shmuel Asafi, Avi Goren, and Daniel Cohen-Or. Weak convex decomposition by lines-of-sight. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing*, SGP '13, pages 23–31, Goslar, DEU, July 2013. Eurographics Association. 3
- [3] Marco Attene, Bianca Falcidieno, and Michela Spagnuolo. Hierarchical mesh segmentation based on fitting primitives. *The Visual Computer*, 22(3):181–193, Mar. 2006. 3
- [4] Chanderjit L. Bajaj and Tamal K. Dey. Convex Decomposition of Polyhedra and Robustness. *SIAM Journal on Computing*, 21(2):339–364, Apr. 1992. Publisher: Society for Industrial and Applied Mathematics. 3
- [5] Filippo Bergamasco, Andrea Albarelli, and Andrea Torsello. A graph-based technique for semi-supervised segmentation of 3D surfaces. *Pattern Recognition Letters*, 33(15):2057–2064, Nov. 2012. 3
- [6] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. *Eurographics Italian Chapter Conference*, page 8 pages, 2008. Artwork Size: 8 pages ISBN: 9783905673685 Publisher: The Eurographics Association. 6
- [7] S. Claici, M. Bessmeltsev, S. Schaefer, and J. Solomon. Isometry-Aware Preconditioning for Mesh Parameterization. *Computer Graphics Forum*, 36(5):37–47, Aug. 2017. 1, 2
- [8] Mathieu Desbrun, Mark Meyer, and Pierre Alliez. Intrinsic Parameterizations of Surface Meshes. *Computer Graphics Forum*, 21(3):209–218, Sept. 2002. 1, 2
- [9] Michael S. Floater. Mean value coordinates, 2003. 1, 2
- [10] Mukulika Ghosh, Nancy M. Amato, Yanyan Lu, and Jyh-Ming Lien. Fast approximate convex decomposition using relative concavity. *Computer-Aided Design*, 45(2):494–504, Feb. 2013. 3
- [11] Mark Gillespie, Boris Springborn, and Keenan Crane. Discrete conformal equivalence of polyhedral surfaces. *ACM Transactions on Graphics*, 40(4):103:1–103:20, July 2021. 2
- [12] Aleksey Golovinskiy and Thomas Funkhouser. Consistent segmentation of 3D models. *Computers & Graphics*, 33(3):262–269, June 2009. 3
- [13] Xianfeng Gu, Steven J. Gortler, and Hugues Hoppe. Geometry images. *ACM Transactions on Graphics*, 21(3):355–361, July 2002. 2, 3
- [14] Kan Guo, Dongqing Zou, and Xiaowu Chen. 3D Mesh Labeling via Deep Convolutional Neural Networks. *ACM Transactions on Graphics*, 35(1):1–12, Dec. 2015. 3
- [15] Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. MeshCNN: A Network with an Edge. *ACM Transactions on Graphics*, 38(4):1–12, July 2019. arXiv: 1809.05910. 2, 3, 4, 8
- [16] Yong He, Hongshan Yu, Xiaoyan Liu, Zhengeng Yang, Wei Sun, Yaonan Wang, Qiang Fu, Yanmei Zou, and Ajmal Mian. Deep Learning based 3D Segmentation: A Survey. *arXiv:2103.05423 [cs]*, Mar. 2021. arXiv: 2103.05423. 3
- [17] Kai Hormann, Bruno Levy, and Alla Sheffer. Mesh Parameterization: Theory and Practice. *ACM SIGGRAPH 2007 Papers - International Conference on Computer Graphics and Interactive Techniques*, 2, Jan. 2008. 1, 2
- [18] Dan Julius, Vladislav Kraevoy, and Alla Sheffer. D-charts: Quasi-developable mesh segmentation. In *In Computer Graphics Forum, Proc. of Eurographics*, pages 581–590, 2005. 1, 2, 3, 7
- [19] Bert Jüttler, Mario Kapl, Dang-Manh Nguyen, Qing Pan, and Michael Pauley. Isogeometric segmentation: The case of contractible solids without non-convex edges. *Computer-Aided Design*, 57:74–90, Dec. 2014. 3
- [20] Oliver Van Kaick, Noa Fish, Yanir Kleiman, Shmuel Asafi, and Daniel Cohen-OR. Shape Segmentation by Approximate Convexity Analysis. *ACM Transactions on Graphics*, 34(1):4:1–4:11, Dec. 2015. 3
- [21] Evangelos Kalogerakis, Aaron Hertzmann, and Karan Singh. Learning 3D mesh segmentation and labeling. *ACM Transactions on Graphics*, 29(4):102:1–102:12, July 2010. 3
- [22] Liliya Kharevych, Boris Springborn, and Peter Schröder. Discrete conformal mappings via circle patterns. *ACM Transactions on Graphics*, 25(2):412–438, Apr. 2006. 2
- [23] Duck Hoon Kim, Il Dong Yun, and Sang Uk Lee. A new shape decomposition scheme for graph-based representation. *Pattern Recognition*, 38(5):673–689, May 2005. 3
- [24] Oscar Kin-Chung Au, Youyi Zheng, Menglin Chen, Pengfei Xu, and Chiew-Lan Tai. Mesh Segmentation with Concavity-Aware Fields. *IEEE Transactions on Visualization and Computer Graphics*, 18(7):1125–1134, July 2012. Conference Name: IEEE Transactions on Visualization and Computer Graphics. 3
- [25] Iasonas Kokkinos, Michael M. Bronstein, Roe Litman, and Alex M. Bronstein. Intrinsic shape context descriptors for deformable shapes. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 159–166, June 2012. ISSN: 1063-6919. 1
- [26] Vladimir Kolmogorov and Ramin Zabih. What Energy Functions Can Be Minimized via Graph Cuts? *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 26(2):13, 2004. 2, 4, 6, 7
- [27] Alon Lahav and Ayellet Tal. MeshWalker: deep mesh understanding by random walks. *ACM Transactions on Graphics*, 39(6):263:1–263:13, Nov. 2020. 3
- [28] Yu-Kun Lai, Shi-Min Hu, Ralph R. Martin, and Paul L. Rosin. Rapid and effective segmentation of 3D models using random walks. *Computer Aided Geometric Design*, 26(6):665–679, Aug. 2009. 3
- [29] Yu-Kun Lai, Qian-Yi Zhou, Shi-Min Hu, and Ralph R. Martin. Feature sensitive mesh segmentation. In *Proceedings of the 2006 ACM symposium on Solid and physical modeling*, SPM '06, pages 17–25, New York, NY, USA, June 2006. Association for Computing Machinery. 3
- [30] Mo Li, Qing Fang, Wenqing Ouyang, Ligang Liu, and Xiaoming Fu. Computing sparse integer-constrained cones for conformal parameterizations. *ACM Transactions on Graphics*, 41(4):1–13, July 2022. 2
- [31] Minchen Li, Danny M. Kaufman, Vladimir G. Kim, Justin Solomon, and Alla Sheffer. OptCuts: joint optimization of surface cuts and parameterization. *ACM Transactions on Graphics*, 37(6):1–13, Jan. 2019. 2
- [32] Jyh-Ming Lien and Nancy M. Amato. Approximate convex decomposition of polygons. In *Proceedings of the twentieth*

- annual symposium on Computational geometry, SCG '04, pages 17–26, New York, NY, USA, June 2004. Association for Computing Machinery. 3
- [33] Ligang Liu, Lei Zhang, Yin Xu, Craig Gotsman, and Steven J. Gortler. A Local/Global Approach to Mesh Parameterization. *Computer Graphics Forum*, 27(5):1495–1504, July 2008. 4
- [34] Victor Lucquin, Sébastien Deguy, and Tamy Boubekeur. SeamCut: interactive mesh segmentation for parameterization. In *SIGGRAPH Asia 2017 Technical Briefs*, pages 1–4, Bangkok Thailand, Nov. 2017. ACM. 2, 3
- [35] Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. Least squares conformal maps for automatic texture atlas generation. *ACM Transactions on Graphics*, 21(3):362–371, July 2002. 1, 2, 3
- [36] Eivind Lyche Melvæ and Martin Reimers. Geodesic Polar Coordinates on Polygonal Meshes. *Computer Graphics Forum*, 31(8):2423–2435, 2012. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2012.03187.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2012.03187.x). 1, 2
- [37] Jun Mitani and Hiromasa Suzuki. Making papercraft toys from meshes using strip-based approximate unfolding. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, pages 259–263, New York, NY, USA, Aug. 2004. Association for Computing Machinery. 2, 3
- [38] Patrick Mullen, Yiyang Tong, Pierre Alliez, and Mathieu Desbrun. Spectral Conformal Parameterization. *Computer Graphics Forum*, 27(5):1487–1494, July 2008. 1, 2
- [39] Ashish Myles and Denis Zorin. Global parametrization by incremental flattening. *ACM Transactions on Graphics*, 31(4):109:1–109:11, July 2012. 2
- [40] Roi Poranne, Marco Tarini, Sandro Huber, Daniele Panozzo, and Olga Sorkine-Hornung. Autocuts: simultaneous distortion and cut optimization for UV mapping. *ACM Transactions on Graphics*, 36(6):1–11, Nov. 2017. 2
- [41] Michael Rabinovich, Tim Hoffmann, and Olga Sorkine-Hornung. Discrete Geodesic Nets for Modeling Developable Surfaces. Technical Report arXiv:1707.08360, arXiv, July 2017. arXiv:1707.08360 [cs] type: article. 1, 2
- [42] Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. Scalable Locally Injective Mappings. *ACM Transactions on Graphics*, 36(2):16:1–16:16, Apr. 2017. 5, 6, 7
- [43] Pedro V. Sander, John Snyder, Steven J. Gortler, and Hugues Hoppe. Texture mapping progressive meshes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 409–416, New York, NY, USA, Aug. 2001. Association for Computing Machinery. 1, 2
- [44] Rohan Sawhney and Keenan Crane. Boundary First Flattening. *ACM Transactions on Graphics*, 37(1):1–14, Jan. 2018. arXiv: 1704.06873. 1, 2
- [45] R. Schmidt. Stroke Parameterization. *Computer Graphics Forum*, 32(2pt2):255–263, 2013. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12045](https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12045). 1, 2
- [46] Ryan Schmidt, Cindy Grimm, and Brian Wyvill. Interactive decal compositing with discrete exponential maps. *ACM Transactions on Graphics*, 25(3):605–613, July 2006. 1, 2
- [47] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. DiffusionNet: Discretization Agnostic Learning on Surfaces. *ACM Transactions on Graphics*, 41(3):1–16, June 2022. 3
- [48] Nicholas Sharp and Keenan Crane. Variational surface cutting. *ACM Transactions on Graphics*, 37(4):1–13, Aug. 2018. 2, 3
- [49] Nicholas Sharp, Yousuf Soliman, and Keenan Crane. The Vector Heat Method. *ACM Transactions on Graphics*, 38(3):24:1–24:19, June 2019. 7
- [50] Idan Shatz, Ayellet Tal, and George Leifman. Paper craft models from meshes. *The Visual Computer*, 22(9):825–834, Sept. 2006. 2, 3
- [51] Georgia Shay, Justin Solomon, and Oded Stein. A Dataset and Benchmark for Mesh Parameterization, Aug. 2022. 6
- [52] A. Sheffer and J.C. Hart. Seamster: inconspicuous low-distortion texture seam layout. In *IEEE Visualization, 2002. VIS 2002.*, pages 291–298, Oct. 2002. 2, 3
- [53] Alla Sheffer, Bruno Lévy, Maxim Mogilnitsky, and Alexander Bogomyakov. ABF++: fast and robust angle based flattening. *ACM Transactions on Graphics*, 24(2):311–330, Apr. 2005. 1, 2
- [54] Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahr Z. Kovalsky, and Yaron Lipman. Geometric optimization via composite majorization. *ACM Transactions on Graphics*, 36(4):38:1–38:11, July 2017. 1, 2
- [55] Oana Sidi, Oliver van Kaick, Yanir Kleiman, Hao Zhang, and Daniel Cohen-Or. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. *ACM Transactions on Graphics*, 30(6):1–10, Dec. 2011. 3
- [56] Jason Smith and Scott Schaefer. Bijective parameterization with free boundaries. *ACM Transactions on Graphics*, 34(4):1–9, July 2015. 1, 2
- [57] O. Sorkine, D. Cohen-Or, R. Goldenthal, and D. Lischinski. Bounded-distortion piecewise mesh parameterization. In *IEEE Visualization, 2002. VIS 2002.*, pages 355–362, Boston, MA, USA, 2002. IEEE. 7
- [58] Boris Springborn, Peter Schröder, and Ulrich Pinkall. Conformal equivalence of triangle meshes. *ACM Transactions on Graphics*, 27(3):1–11, Aug. 2008. 2
- [59] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Computer Graphics Forum*, 28(5):1383–1392, July 2009. 4
- [60] Qian Sun, Long Zhang, Minqi Zhang, Xiang Ying, Shi-Qing Xin, Jiazhi Xia, and Ying He. Texture brush: an interactive surface texturing interface. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '13*, pages 153–160, New York, NY, USA, Mar. 2013. Association for Computing Machinery. 1, 2
- [61] Ryo Suzuki, Tom Yeh, Koji Yatani, and Mark D. Gross. Autocomplete Textures for 3D Printing, Mar. 2017. arXiv:1703.05700 [cs] version: 1. 2

- [62] Oliver van Kaick, Andrea Tagliasacchi, Oana Sidi, Hao Zhang, Daniel Cohen-Or, Lior Wolf, and Ghassan Hamarneh. Prior Knowledge for Part Correspondence. *Computer Graphics Forum*, 30(2):553–562, 2011. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2011.01893.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2011.01893.x). 3
- [63] Charlie Wang. Computing Length-Preserved Free Boundary for Quasi-Developable Mesh Segmentation. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):25–36, Jan. 2008. Conference Name: IEEE Transactions on Visualization and Computer Graphics. 2, 3
- [64] Hao Wang, Tong Lu, Oscar Kin-Chung Au, and Chiew-Lan Tai. Spectral 3D mesh segmentation with a novel single segmentation field. *Graphical Models*, 76(5):440–456, Sept. 2014. 3
- [65] Yunhai Wang, Minglun Gong, Tianhua Wang, Daniel Cohen-Or, Hao Zhang, and Baoquan Chen. Projective analysis for 3D shape segmentation. *ACM Transactions on Graphics*, 32(6):192:1–192:12, Nov. 2013. 3
- [66] Zhige Xie, Kai Xu, Ligang Liu, and Yueshan Xiong. 3D Shape Segmentation and Labeling via Extreme Learning Machine. *Computer Graphics Forum*, 33(5):85–95, 2014. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12434](https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12434). 3
- [67] Haotian Xu, Ming Dong, and Zichun Zhong. Directionally Convolutional Networks for 3D Shape Segmentation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2717–2726, Oct. 2017. ISSN: 2380-7504. 3
- [68] Kai Xu, Honghua Li, Hao Zhang, Daniel Cohen-Or, Yueshan Xiong, and Zhi-Quan Cheng. Style-content separation by anisotropic part scales. In *ACM SIGGRAPH Asia 2010 papers, SIGGRAPH ASIA '10*, pages 1–10, New York, NY, USA, Dec. 2010. Association for Computing Machinery. 3
- [69] Songgang Xu and John Keyser. Texture mapping for 3D painting using geodesic distance. In *Proceedings of the 18th meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '14*, page 163, New York, NY, USA, Mar. 2014. Association for Computing Machinery. 1, 2
- [70] Songgang Xu, Hang Li, and John Keyser. Field-Aware Parameterization for 3D Painting. In Marina Gavrilova, Jian Chang, Nadia Magnenat Thalmann, Eckhard Hitzer, and Hiroshi Ishikawa, editors, *Advances in Computer Graphics*, Lecture Notes in Computer Science, pages 131–142, Cham, 2019. Springer International Publishing. 1, 2
- [71] Hitoshi Yamauchi, Stefan Gumhold, Rhaleb Zayer, and Hans-Peter Seidel. Mesh segmentation driven by Gaussian curvature. *The Visual Computer*, 21(8):659–668, Sept. 2005. 2, 3
- [72] Dong-Ming Yan, Wenping Wang, Yang Liu, and Zhouwang Yang. Variational mesh segmentation via quadric surface fitting. *Computer-Aided Design*, 44(11):1072–1082, Nov. 2012. 3
- [73] Dongbo Zhang, Zheng Fang, Xuequan Lu, Hong Qin, Antonio Robles-Kelly, Chao Zhang, and Ying He. Deep Patch-Based Human Segmentation. In Haiqin Yang, Kitsuchart Papsupa, Andrew Chi-Sing Leung, James T. Kwok, Jonathan H. Chan, and Irwin King, editors, *Neural Information Processing*, Lecture Notes in Computer Science, pages 229–240, Cham, 2020. Springer International Publishing. 1, 2
- [74] Huijuan Zhang, Chong Li, Leilei Gao, Sheng Li, and Guoping Wang. Shape segmentation by hierarchical splat clustering. *Computers & Graphics*, 51:136–145, Oct. 2015. 3
- [75] Juyong Zhang, Jianmin Zheng, and Jianfei Cai. Interactive Mesh Cutting Using Constrained Random Walks. *IEEE Transactions on Visualization and Computer Graphics*, 17(3):357–367, Mar. 2011. Conference Name: IEEE Transactions on Visualization and Computer Graphics. 3
- [76] Qingnan Zhou and Alec Jacobson. Thingi10K: A Dataset of 10,000 3D-Printing Models. *arXiv preprint arXiv:1605.04797*, 2016. 6
- [77] Yixin Zhuang, Ming Zou, Nathan Carr, and Tao Ju. Anisotropic geodesics for live-wire mesh segmentation. *Computer Graphics Forum*, 33(7):111–120, Oct. 2014. 3