

Polygonal finite element-based content-aware image warping

Juan Cao, Xiaoyi Zhang, Jiannan Huang
School of Mathematical Sciences
Xiamen University
Juancao@xmu.edu.cn

Yongjie Jessica Zhang
Department of Mechanical Engineering
Carnegie Mellon University
Jessicaz@andrew.cmu.edu

Abstract

Mesh-based image warping techniques typically represent image deformation using linear functions on triangular meshes or bilinear functions on rectangular meshes. This enables simple and efficient implementation, but in turn, restricts the representation capability of the deformation, often leading to unsatisfactory warping results. We present a novel, flexible polygonal finite element (poly-FEM) method for content-aware image warping. Image deformation is represented by high-order poly-FEMs on a content-aware polygonal mesh with a cell distribution adapted to saliency information in the source image. This allows highly adaptive meshes and smoother warping with fewer degrees of freedom, thus significantly extending the flexibility and capability of the warping representation. Benefiting from the continuous formulation of image deformation, our poly-FEM warping method is able to compute the optimal image deformation by minimizing existing or even newly designed warping energies consisting of penalty terms for specific transformations. We demonstrate the versatility of the proposed poly-FEM warping method in representing different deformations and its superiority by comparing it to other existing state-of-the-art methods.

Keywords: Image warping; finite element method; polygonal element; mesh generation

1. Introduction

Due to advances in imaging technology, the acquisition and display of digital images are almost universal. Various display devices are used to view images, such as phones, tablets, monitors, and televisions. Images frequently change size, and should fill the whole screen to achieve an optimal display; screens vary in size. Images also need to be resized in other applications. For example, document display and printing require resizing embedded images to comply with a specified layout. Research into image resizing, also known as image retargeting, has drawn much attention in recent years, and several techniques have been proposed.

Image scaling is the most straightforward method to achieve the image resizing goal. However, scaling often does not produce satisfactory results, as it is oblivious to image content. Another simple method for image retargeting is cropping. Cropping inevitably causes information loss and leads to unpleasant results. To preserve relevant information, especially visually important structures and objects, a more sophisticated class of techniques attempts to resize images in a content-aware fashion. Existing content-aware image retargeting methods can be classified into two general categories: cropping methods and warping methods. In content-aware cropping methods, pixels or regions in an image are removed according to pre-specified criteria. They achieve results with better visual quality than naive cropping. However, important objects may be broken, and artifacts may be introduced as the pixel removal operation highly depends on object detection results, which are often inaccurate.

Warping methods, also referred to as continuous methods, are another popular type of image retargeting technique. Unlike cropping methods, which may discard important contents, warping methods retain important and unimportant contents. To obtain a resized image with important objects preserved, warping methods perform a non-linear deformation that minimizes the distortion of important regions while allowing large distortions in unimportant regions. Mesh-based warping methods construct a mesh on the image domain and obtain the resized image by deforming the mesh. In most previous methods, the warping meshes used for driving the deformation are strictly triangular or quadrilateral. Non-uniform deformation of the image, which is supposed to be a continuous function, is typically represented by piecewise linear functions on warping meshes. However, only allowing a single cell shape in the warping mesh with a linear approximation of the associated functions is too restrictive. This limits the non-uniform image deformation to a relatively small function space, leading to unsatisfactory results.

This paper introduces a novel continuous warping representation and proposes a fully automatic algorithm for content-aware image retargeting. The warp mapping is rep-

resented as a smooth function by high-order generalized barycentric coordinates defined over polygonal meshes. Our representation possesses superior properties, such as supporting highly adaptive meshes, high-order basis functions, and achieving continuity without enforcing additional constraints. Image warping is driven by the deformation of these polygonal meshes determined by a specific energy function. Experimental results show that our algorithm for content-aware image warping achieves a better trade-off between warping quality and mesh size, and greater robustness, than other existing state-of-the-art methods. In summary, our main contributions are:

- (1) A novel poly-FEM-based warping representation for content-aware image retargeting. Image warping is represented by continuously stitched functions with higher-order approximation defined over polygonal meshes on the image domain. This representation includes the piecewise linear representation as a special case, and so can achieve more satisfactory results.
- (2) An efficient and fully automatic framework to warp images. Polygonal meshes for driving the deformation are generated with local density and shape adaptive to feature information. Different warping energy functions can be incorporated and tested easily and consistently to achieve various deformation results.

The remainder of this paper is organized as follows: Section 2 reviews related work. We propose a poly-FEM-based warping representation in Section 3. Section 4 presents the algorithm and implementation for our poly-FEM-based image warping method. Results and comparisons are presented in Section 5. Conclusions, limitations, and suggestions for future work are given in Section 6.

2. Related work

Image retargeting has been extensively studied in computer graphics. In this paper, we focus on content-aware retargeting techniques. From the vast body of literature in the field, we only review references closely related to this paper, and refer the interested reader to [40, 19] for more comprehensive surveys.

2.1. Content-aware cropping

Cropping-based methods discard pixels in unimportant regions and scale or shift the remaining pixels to resize the image. One class of content-aware cropping techniques searches for a cropping rectangle inside which the aggregated importance is maximized [36, 26]. They perform well if the input image contains only one central important object. To deal with images with two or more scattered regions of interest, Setlur *et al.* [33] proposed a method that first removes the regions of interest and inpaints the holes

to generate a background, and then places the cropped objects back. It heavily relies on accurate segmentation of the source image.

Seam carving-based methods form another class of cropping methods, which decrease the image width or height one pixel at a time by removing a seam with least importance and shifting the remaining pixels to compensate for the removed seam [1]. The original seam carving method introduces visible artifacts if the input image contains straight lines or geometric structures. This method is thus enhanced by including line detection to better preserve straight lines [30]. The multi operator (Multiop) method uses several operators for resizing media, including cropping, seam carving, scaling and warping [31]. Patch-based methods achieve retargeting by manipulating patches [3, 27]. As they remove regions or strips, patch-based methods can be considered to generalize cropping-based methods. Generally speaking, cropping-based methods remove pixels from source images, which causes loss of information; hence artifacts can sometimes be observed in the results.

2.2. Content-aware warping

Warping methods scale the source image non-uniformly to preserve important regions. In general, regions with high importance are constrained to distort as little as possible, while unimportant regions are allowed to have relatively large deformation. The image is subdivided into a mesh, whose deformation drives the deformation of the source image. Typically, a triangular mesh [12, 22] or a quadrilateral mesh [44, 43, 17, 20, 34, 5] is used, and image deformation is naturally represented as piecewise linear or bilinear functions on mesh faces, respectively. For example, a piecewise bilinear warping may be computed by iteratively computing optimal local scaling factors for each cell of a quadrilateral mesh according to a significance map [43]. A piecewise linear warping on a triangular mesh can be constructed from an approximation to a prescribed Beltrami representation (BR) [22]. Instead of limiting ourselves to piecewise linear functions on triangular or quadrilateral meshes, we propose a more general representation of the continuous warping that supports high-order continuity and adaptive meshes.

2.3. Deep learning-based methods

Recently, attempts have been made to solve the image retargeting problem using deep learning techniques [6, 38, 48, 7], which are extensions of the methods mentioned above. For example, a weakly- and self-supervised deep convolutional neural network (WSSDCNN) has been proposed for predicting attentive shift maps in [6]. Scaling on grid cells is used to represent image distortion in the deep cyclic image retargeting approach (Cycle-IR) [38]. The multi-operator retargeting is formulated as a Markov decision-making pro-

cess and optimized by reinforcement learning in the semantics and aesthetics aware multi-operator image retargeting (SAMIR) framework [48]. The deep network resizing (DNR) method applies resizing operators, including seam carving and grid-warping, in feature space instead of pixel space [7].

2.4. FEM-based warping

In the computer graphics community, FEM has been applied to applications such as 2D/3D morphing [15, 45] and geometric modeling [2, 16, 47]. Traditional FEM has also been applied to image warping, relying on strictly triangular or rectangular meshes. For example, Gee *et al.* [11] used simple linear elements in medical image warping for registration. Later, a discontinuous Galerkin FEM (DG FEM) with triangular or rectangular elements using power polynomials was applied to the content-aware image warping task [18]. Requirements on element types, to simplify or accelerate the involved computation, restrict the approximation capability of FEM.

2.5. High-order poly-FEM

Poly-FEMs offer several advantages over traditional finite elements in practical applications. Generalized barycentric coordinates (GBCs), such as Wachspress coordinates [41, 25] and mean value coordinates (MVCs) [14], provide suitable bases for linear finite elements on general polygons as generalizations of linear barycentric FEM shape functions. Recently, extensions to higher-order approximations on polygonal elements have also been studied [28, 37, 10]. Higher-order poly-FEMs share attractive properties with piecewise linear poly-FEMs, such as partition of unity, nodal data interpolation, and smoothness. In addition, they provide higher-order reproduction properties.

High-order poly-FEM has been successfully applied to solving partial differential equations [10, 28, 37] and function approximation [4]. We are motivated by these successes to use the high-order poly-FEM method in the image warping problem. Our high-order polygonal element-based method has important advantages over prior mesh-based methods, such as allowing for highly adaptive meshes, smoother and more flexible representations of image deformation using many fewer degrees of freedom (DOFs), and ability to incorporate general deformation energy functions into the framework.

3. Poly-FEM warping representation

To apply high-order poly-FEM to image warping, we first discretize the domain so that the continuous warping map is subdivided into smaller and simpler polygons. Then, poly-FEMs are used to approximate the unknown function over the discretized domain. In this section, we

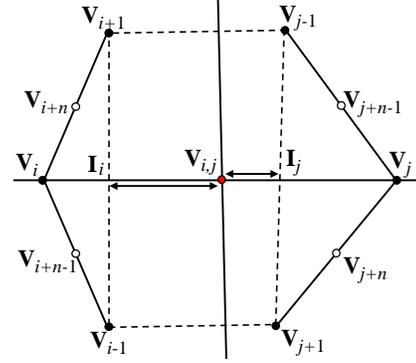


Figure 1. Coefficients.

briefly discuss the poly-FEMs, which we propose as a basis for the discretized image warping representation. We defer discussing the polygonal mesh generation until Section 4.3.2. Note that the discretization operation using high-order elements is independent of the choice of shape functions for the aforementioned high-order GBCs. For simplicity, we describe the image warping representation using the quadratic serendipity elements (QSEs) proposed in [28] as an example.

3.1. Quadratic poly-FEMs

The QSEs in [28] are developed using GBCs, for instance, MVCs. We first introduce some notation; our notation differs slightly from that of [28], in an inessential way. Let Ω be a convex polygon in the plane with n vertices ordered counter-clockwise ($\mathbf{v}_1, \dots, \mathbf{v}_n$), with no more than three consecutive vertices collinear. Each vertex \mathbf{v}_i is associated with a GBC, denoted $\lambda_i(u, v)$. For simplicity, hereinafter we omit the variables in each function, e.g. $\lambda_i(u, v)$ is simplified to λ_i . Let \mathbf{v}_{i+n} denote the midpoint of $\mathbf{v}_i\mathbf{v}_{i+1}$. Let $\mathbf{v}_{i,j}$ denote the midpoint of $\mathbf{v}_i\mathbf{v}_j$ ($i \neq j$). Let \mathbf{I}_i denote the intersection of segments $\mathbf{v}_i\mathbf{v}_j$ and $\mathbf{v}_{i-1}\mathbf{v}_{i+1}$, and \mathbf{I}_j denote the intersection of segments $\mathbf{v}_i\mathbf{v}_j$ and $\mathbf{v}_{j-1}\mathbf{v}_{j+1}$: see Figure 1. Let

$$s = \frac{\|\mathbf{v}_i\mathbf{v}_j\|}{\|\mathbf{v}_i\mathbf{v}_j\| - \|\mathbf{I}_i\mathbf{I}_j\|},$$

and

$$\begin{aligned} c_{i,j}^{i,i} &= -s \left(1 + 2 \frac{\|\mathbf{I}_i\mathbf{v}_{i,j}\|}{\|\mathbf{v}_i\mathbf{v}_j\|} \right), \\ c_{i,j}^{j,j} &= -s \left(1 + 2 \frac{\|\mathbf{I}_j\mathbf{v}_{i,j}\|}{\|\mathbf{v}_i\mathbf{v}_j\|} \right), \\ c_{i,j}^{i,i-1} &= s \frac{\|\mathbf{I}_i\mathbf{v}_{i+1}\|}{\|\mathbf{v}_{i-1}\mathbf{v}_{i+1}\|}, & c_{i,j}^{i,i+1} &= s \frac{\|\mathbf{I}_i\mathbf{v}_{i-1}\|}{\|\mathbf{v}_{i-1}\mathbf{v}_{i+1}\|}, \\ c_{i,j}^{j,j-1} &= s \frac{\|\mathbf{I}_j\mathbf{v}_{j+1}\|}{\|\mathbf{v}_{j-1}\mathbf{v}_{j+1}\|}, & c_{i,j}^{j,j+1} &= s \frac{\|\mathbf{I}_j\mathbf{v}_{j-1}\|}{\|\mathbf{v}_{j-1}\mathbf{v}_{j+1}\|}, \end{aligned} \quad (1)$$

for $i, j = 1, \dots, n$; subscripts are to be taken modulo n .

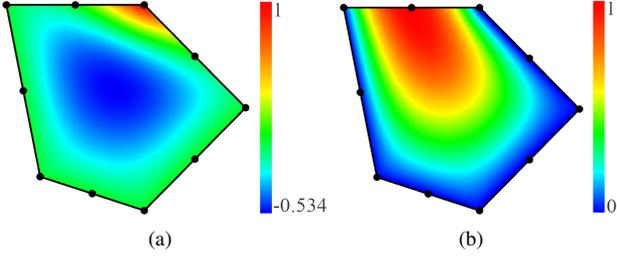


Figure 2. Quadratic serendipity element basis functions. Values of basis functions are color-coded as per the color bars. (a) Quadratic basis associated with a polygon vertex. (b) Quadratic basis associated with the midpoint of a polygon edge.

The quadratic serendipity element basis functions ψ_k associated with vertices or midpoints \mathbf{v}_i , $i = 1, \dots, 2n$, are defined as linear combinations of $\mu_{ij} = \lambda_i \lambda_j$ as follows:

$$\begin{cases} \psi_i = \xi_{i,i} - \xi_{i,i+1} - \xi_{i-1,i}, & i = 1, \dots, n, \\ \psi_{n+i} = 4\xi_{i,i+1}, \end{cases}$$

where

$$\xi_{i,i} = \mu_{ii} + \sum_{j, |i-j|>1} c_{i,j}^{i,i} \mu_{ij},$$

and

$$\xi_{i,i+1} = 2 \left(\mu_{ii+1} + \sum_{j, |i-j|>1} (c_{i,j}^{i,i+1} \mu_{ij} + c_{i+1,j}^{i+1,i} \mu_{i+1,j}) \right).$$

We refer the reader to [28] for further details of the computation of coefficients.

In particular, the QSE basis functions possess all the properties needed for admissible quadratic FEM basis functions:

- Partition of unity: $\sum_{k=1}^{2n} \psi_k = 1$.
- Smoothness: ψ_k is smooth within the domain Ω and is discontinuous across the element boundary.
- Quadratic precision: ψ_k for $k = 1, \dots, 2n$ can reproduce polynomials of up to degree two.
- Nodal interpolation:

$$\psi_k(\mathbf{v}_i) = \begin{cases} 1 & i = k \\ 0 & i \neq k \end{cases}, \quad \forall i, k = 1, \dots, 2n. \quad (2)$$

The n -sided polygonal domain with the associated basis function set $\{\psi_k\}_{k=1}^{2n}$ forms a construction of QSE on Ω using GBCs. Figure 2 shows examples of MVC based QSE basis functions from [28].

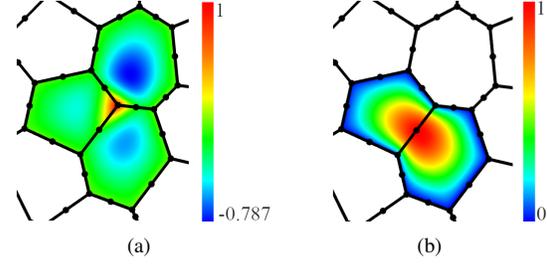


Figure 3. Merged quadratic serendipity element basis functions. Values of basis functions are color-coded as per the color bars. (a) Merged quadratic basis associated with a vertex. (b) Merged quadratic basis associated with an edge midpoint.

3.2. Image warping representation

Consider the rectangular domain $I = [0, m] \times [0, n]$ of a digital image with resolution $m \times n$. The warping map is a continuous vector function $\mathbf{f} = (f_1, f_2)$ that maps a point $(u, v) \in I$ to a warped point $(f_1(u, v), f_2(u, v))$. Now we discretize this continuous model using QSEs in preparation for the optimal warp computation. To the best of our knowledge, this is the first time QSEs have been applied to an image warping representation.

Assume that the domain I has been appropriately discretized into a polygonal mesh M with cells Ω_k for $k = 1, \dots, N$. We construct the QSE basis functions on each cell. Consider two adjacent cells Ω_{k_1} and Ω_{k_2} sharing a common edge $\mathbf{v}_1 \mathbf{v}_2$. The two basis functions $\psi_{k_1}^1$ and $\psi_{k_2}^1$ associated with the same vertex \mathbf{v}_1 (or the same midpoint \mathbf{v}_{n+1}) on Ω_{k_1} and Ω_{k_2} respectively are discontinuous across the edge $\mathbf{v}_1 \mathbf{v}_2$. However, this pair of bases coincides with each other on $\mathbf{v}_1 \mathbf{v}_2$, due to their nodal interpolation property and quadratic precision. To ensure automatic continuity across element boundaries, we collect the basis functions associated with a vertex (or an edge midpoint). Their sum is set as a basis in the final warping representation. Specifically, for two basis functions $\psi_{k_1}^1$ and $\psi_{k_2}^1$ associated with midpoint \mathbf{v}_i on the polygonal mesh, we replace $\psi_{k_1}^1$ and $\psi_{k_2}^1$ with a new basis function B_i which is the sum of these two discontinuous basis functions: $B_i = \psi_{k_1}^1 + \psi_{k_2}^1$. Basis functions associated with mesh vertices are treated in the same fashion. Figure 3 shows examples of the merged bases associated with a vertex and an edge midpoint, respectively.

Consider a polygonal mesh with M vertices and edge midpoints, $\{\mathbf{v}_i\}_{i=1}^M$. We denote the merged basis function associated with a vertex or an edge midpoint by $B_i(u, v)$, $i = 1, \dots, M$. Then the warping map is represented as a linear combination of these basis functions, giving:

$$\mathbf{f}(u, v) = \sum_{i=1}^M B_i(u, v) \mathbf{c}_i, \quad (3)$$

where $\mathbf{c}_i = (c_{i1}, c_{i2})$ are the position vectors of vertices

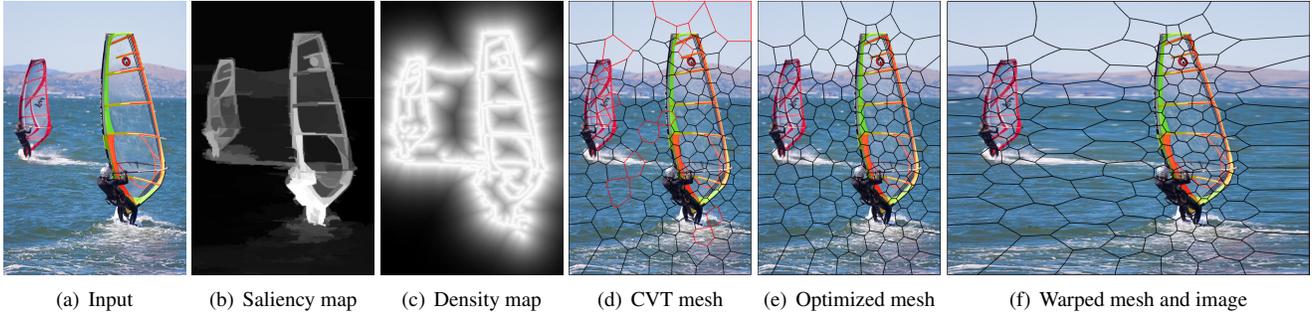


Figure 4. Pipeline of the proposed algorithm. Given an input image (a), a saliency map is generated from object level semantic information (b). A density function is computed (c), and used to generate a content-aware polygonal mesh (d); cells incident to an edge shorter than 5% of the mean edge length are shown in red. This is optimized to remove short edges (e) and finally the warped mesh is computed by optimizing a distortion energy to give the warped image (f).

or edge midpoints \mathbf{v}_i after deformation. The next section describes an image warping framework tailored to compute the coefficients based on our new warping representation. Note that the warping map defined in Equation (3) is naturally continuous along cell edges, benefiting from the basis consolidation mentioned above, whereas for DG FEM bases, additional constraints are needed to enforce continuity of the individually defined quadratic power polynomial functions along the cell edges, to achieve a continuous map [18].

4. Algorithm

Content-aware image warping aims to preserve visually important image regions as much as possible while allowing the visually unimportant image regions to have relatively large distortions. In this section, we present our algorithm for content-aware image warping based on our warping representation.

4.1. Overview

Given an input image I , we aim to find a warping map as represented in Equation (3) that preserves regions with high importance. Most mesh-based image warping techniques boil down to an optimization framework to obtain the warping mesh, differing by particular objective functions. They consist of two steps in general: importance map generation and image retargeting. Here we follow the same workflow. First, we generate a saliency map according to object-level semantic information. Next, a non-uniform polygonal mesh is constructed on the input image based on the importance map. Finally, we compute a warping map defined on the polygonal mesh by optimizing a distortion energy function to achieve image retargeting. The pipeline of our algorithm is shown in Figure 4. Details of each step are given below.

4.2. Saliency map generation

The estimation of visual attention (saliency) has been a fundamental problem in neuroscience, psychology, and computer vision for a long time. Various applications for saliency estimation include object detection and recognition, photo collage, and image compression. Saliency estimation is also the first step of content-aware image retargeting.

In the context of image retargeting, visual attention is the main driving factor determining what information in the image is perceived to be the most important to preserve. Visual attention representations include pixel-level features, such as contours, textural contrast or color features, and higher-level features, such as faces, people, and objects. The saliency map is estimated based on the identification and analysis of several different visual attention factors. Engelke *et al.* investigated the impact of different visual attention representations on content-aware image retargeting [9]. They suggested using object-level regions of interest (ROI) for image retargeting, which gave superior performance in their practical experiments. Accordingly, our saliency map generation is based on object-level semantic information. Here, we adopt the discriminative regional feature integration method [42], which provides a basic fit for our application. This method introduces a regional object-sensitive descriptor and an image-specific backgroundness descriptor to estimate saliency. The final saliency map is obtained by fusing saliency maps computed from the multi-level image segmentation to remedy possible inaccuracies due to unreliable segmentation: see Figure 4(b) for an example saliency map.

It should be pointed out that all content-aware warping methods rely on saliency estimation results. More advanced saliency detection methods could be adopted to achieve better warping results without affecting our overall image warping pipeline; we leave this as a topic for future investigation.

4.3. Polygonal mesh generation

To define the warping map, we need to construct an appropriate polygonal mesh on the input image. Most existing methods for polygonal mesh generation either directly rely on Voronoi diagrams or indirectly exploit the duality of Voronoi diagrams and Delaunay tessellations and their properties. In this section, we adopt a direct approach to generating suitable polygonal mesh. We first study the conditions needed to prevent fold-overs in the warped image. Then we apply the centroidal Voronoi tessellation (CVT) method according to image saliency. The mesh is further optimized to better suit our image warping purpose .

4.3.1 Conditions for foldover-free warping

In image warping, fold-overs should be avoided in the output images as they usually introduce undesirable artifacts. The problem of fold-overs is an essential manifestation of the lack of bijectivity of the warping map defined by both the input and warped polygonal meshes. In this section, we first analyze the conditions to guarantee bijectivity of warping in a single cell, following [32]. Then, we propose criteria for input polygonal mesh generation based on these conditions. We defer the computation of the warping map to Section 4.4.

For simplicity, we now consider the mapping $\mathbf{f} = (f_1, f_2)$ in a single cell Ω_{i^*} . Let

$$J(\mathbf{f}) = \begin{bmatrix} \partial_u f_1 & \partial_v f_1 \\ \partial_u f_2 & \partial_v f_2 \end{bmatrix} \quad (4)$$

be the Jacobian of mapping \mathbf{f} . Then, \mathbf{f} is injective if its Jacobian determinant $\det(J(\mathbf{f}))$ is strictly positive in domain Ω_{i^*} . The linear precision property of quadratic basis functions implies that the image warping map restricted to each cell Ω_{i^*} can be represented as

$$\mathbf{f}(u, v)|_{\Omega_i} = \mathbf{I}(u, v) + \sum_{i=1}^M \psi_i(u, v) \mathbf{d}_i, \quad (5)$$

where $\mathbf{I}(u, v)$ is the identity map and $\mathbf{d}_i = \mathbf{c}_i - \mathbf{v}_i$ is the displacement of vertices or edge midpoints upon warping. Simple algebra reduces the Jacobian of \mathbf{f} to

$$J(\mathbf{f}) = 1 + \sum_i \nabla \psi_i \cdot \mathbf{d}_i + \sum_i \sum_j \partial_u \psi_i \partial_v \psi_j (\mathbf{d}_i \times \mathbf{d}_j). \quad (6)$$

To simplify the notation for later discussion, we rewrite the basis functions using a little algebra as

$$\psi_i = \sum_j A_{i,j} \mu_{i,j}, \quad (7)$$

where coefficients have the form $A_{i,j} = r_{i,j} c_{i,j}^{a,b}$, $r_{i,j}$ are constants independent of the geometry of polygons and $c_{i,j}^{a,b}$

is defined in Equation (1). Let $A = \max\{|A_{i,j}|\}$, $C = \max\{c_{i,j}^{a,b}\}$, $D = \max\{\|\mathbf{d}_i\|\}$ and $M = M_1 + \dots + M_n$ with

$$M_j = \sup_{v \in \Omega_j} \|\nabla \psi_j(v)\|.$$

Then, we have $J(\mathbf{f}) \geq 1 - AMD - A^2 M^2 D^2$, which implies that \mathbf{f} is injective if

$$D < \frac{\sqrt{5} - 1}{2AM}. \quad (8)$$

Note that as the factor s in the coefficients $c_{i,j}^{a,b}$ approaches ∞ , so do the coefficients $A_{i,j}$, as the angles $\mathbf{v}_{i-1} \mathbf{v}_{i+1}$ approach π . Short edges of polygonal cells may also cause a blowup in the coefficients $A_{i,j}$ used to construct ψ_i and the gradient $\nabla \mu_{i,j}$ over the short edges. The problem of extremely large gradients over edges is independent of the generalized barycentric coordinates employed since all quadratic bases are identical on the edges of polygonal meshes. Hence, the mapping on cells with large angles or short edges could fail to meet the sufficient conditions for bijectivity even if the displacements of vertices are very small. In other words, polygonal meshes free of short edges and large angles may allow larger displacements of vertices in defining a bijective warping map. Therefore, large angles and short edges should be avoided in the polygonal meshes. We here use the CVT method to generate large-angle-free polygonal meshes, then use a mesh improvement scheme in [35] to further remove the short edges.

4.3.2 Adaptive mesh generation

The CVT is a special Voronoi tessellation whose seed points coincide with the centroids of the corresponding Voronoi cells. In particular, assume we have a Voronoi diagram with n seed points x_i , for $i = 1, \dots, n$, in the image domain. Then the CVT satisfies:

$$x_i = \frac{\int_{\Omega_i} x \rho(x) dx}{\int_{\Omega_i} \rho(x) dx}, \quad i = 1, \dots, n, \quad (9)$$

where Ω_i is the Voronoi cell of x_i and $\rho(x)$ is a specified density function. A CVT can be computed by Lloyd's relaxation method, which iteratively moves each seed point x_i to the centroid of the corresponding cell Ω_i [8]. When $\rho(x) = 1$, we obtain a polygonal mesh with uniformly distributed polygonal cells. In the following, we obtain content-aware polygonal meshes by carefully choosing the density function.

Note that regions with high saliency in an image undergo small deformation, while regions with low saliency undergo large deformation. We need more polygonal elements in areas with high saliency variance to produce smooth transitions between different deformations. Each vertex or edge

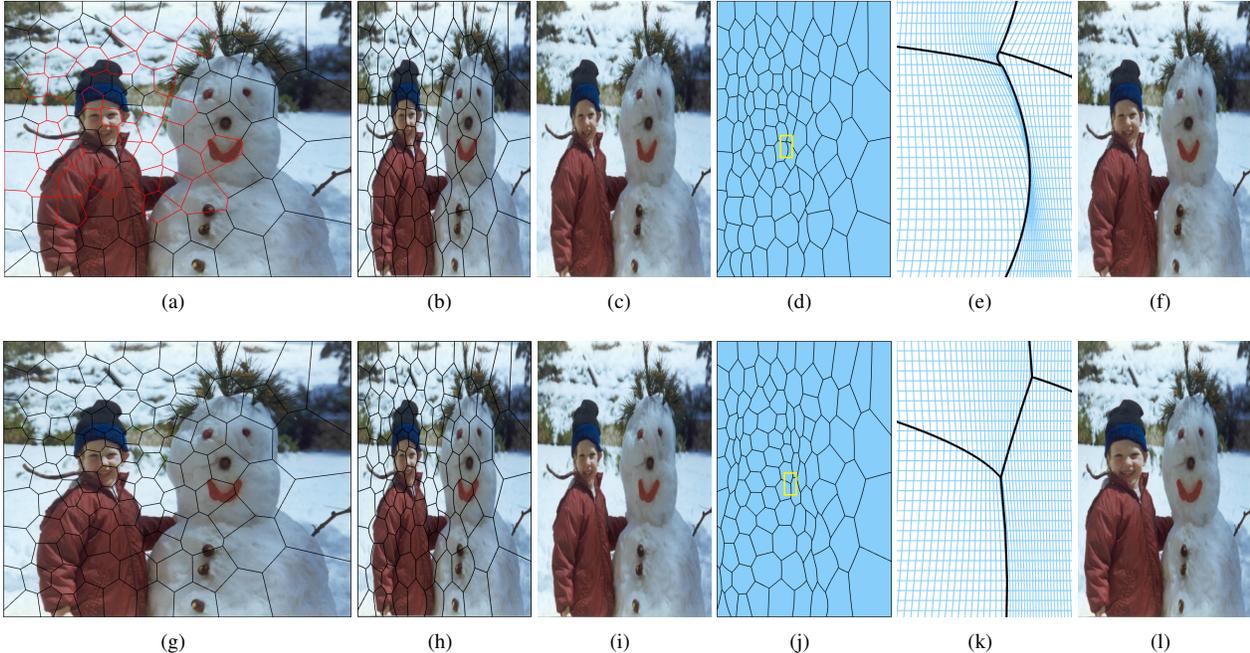


Figure 5. Optimized warping meshes and results for different input polygonal meshes and with different constraints. Above: results for input polygons with short edges. Below: Results for input mesh with short edges removed. (a) Polygonal mesh containing short edges; cells incident to an edge shorter than 5% of the mean edge length are marked in red. (b,c) Deformed mesh and result under sufficient constraints. (d) Deformed polygonal mesh under relaxed constraints. (e) Close up view of a non-bijective region in (d). (f) Warping result using warped polygonal mesh in (d). (g) Polygonal mesh with short edges removed. (h,i) Deformed mesh and results under sufficient constraints. (j) Deformed polygonal mesh under relaxed constraints. (k) Close up view of a region in (j). (l) warping result based on the polygonal mesh in (e).

midpoint of the warped mesh corresponds to a DOF in the warping representation. We can locally increase the resolution of meshes in the corresponding area to introduce more DOFs into the warping representation. We design the density function for the mesh cell distribution as follows:

- (1) We first detect sharp edges in the saliency map using the Sobel edge detector [13]. The density at points on detected edges is set to 1;
- (2) For any other point a distance d away from the detected edges, the density is set to $1/d^3$.

Figure 4(c) visualizes this density function, which leads to the content-aware polygonal mesh in Figure 4(d). Intuitively, more cells are placed in the transitions between the visually important and unimportant areas.

4.3.3 Mesh optimization

Note that a Voronoi vertex is the circumcenter of a Delaunay triangle. To prevent circumcenters of Delaunay triangles from coming close to each other, we relocate the vertices of Delaunay triangles such that circumcenters of triangles are as interior as possible. In particular, we minimize the

squared distances d_t from the circumcenters to incenters of triangles t of the Delaunay triangulation T as follows [35]:

$$E_{se}(\mathbf{v}_1, \dots, \mathbf{v}_m) = \sum_{t \in T} d_t^2 = \sum_{t \in T} R_t(R_t - 2r_t), \quad (10)$$

where R_t and r_t are the circumradius and inradius of triangle t . For a triangle t with edge lengths a, b, c and area A , R_t and r_t are given by $R_t = abc/(4A)$ and $r_t = 2A/(a + b + c)$. Here we use the gradient descent method to minimize the energy in Equation (10).

A polygonal mesh obtained by the CVT method may contain arbitrarily short edges: see Figure 5(a). Our experiments indicate that such short edges may lead to a non-bijective warping map: the resulting deformed pixel grid may locally fold over in regions around short edges: see Figure 5(d,e) for the deformed mesh and a close up view of the warping map, where only the resulting deformed pixel grids are shown to better visualize the map deformation. Ill-shaped polygonal meshes also lead to undesired warping results: see Figure 5(d) for an example. Our experiments indicate that warping maps based on optimized meshes, where all edges shorter than 5% of the mean edge length are removed, achieve bijectivity up to pixel accuracy. Figure 5(g) shows the optimized mesh, corresponding to a

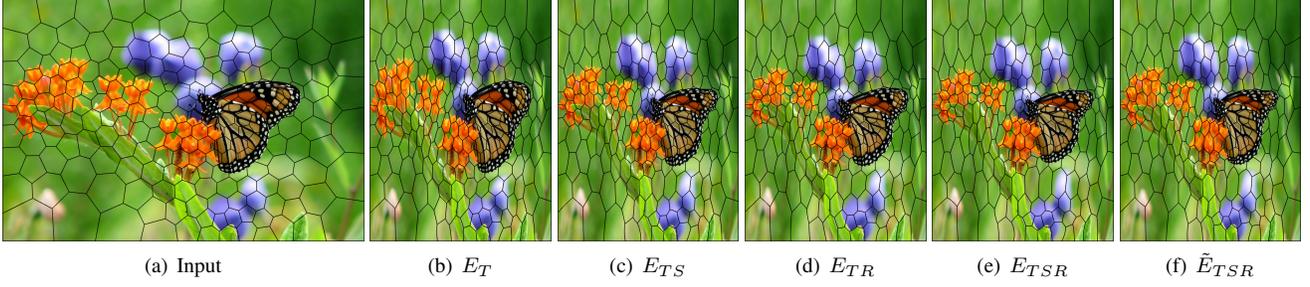


Figure 6. Results using different deformation energy functions; there are 200 elements and 1900 DOFs in our warping representation. (a) Input. (b–f) Warping results by optimizing the energy functions in Eqs. (11, 14–17), respectively.

bijjective warping map, with the deformed pixel grids in Figure 5(j,k) and the superior warping result in Figure 5(l).

It should be pointed out that other methods, such as edge collapse, can also remove short edges. However, this might no longer maintain the Voronoi properties and introduce non-convex elements. The advantage of using a polygonal mesh is that it allows the mesh resolution to be highly adaptive to the content of the source image. We will show later that we can get similar results to previous FEM-based methods using many fewer DOFs, owing to the flexibility of polygonal meshes.

4.4. Deformation energy optimization

With the polygonal mesh in hand, we can now construct the polygonal elements on each cell. Once the basis functions on each cell are defined, the warping representation in Equation (3) is completely determined by the coefficients \mathbf{c}_i . We propose a deformation energy function $E(\mathbf{f})$ to quantify the performance of a specific warping map $\mathbf{f}(u, v)$. The optimal warping map is then determined by minimizing the energy function while considering boundary constraints and other additional constraints, e.g. to preserve lines. For simplicity, we define the deformation energy of a warping map $\mathbf{f}(u, v)$ as

$$E_T(\mathbf{f}) = E(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n) = \int_{\Omega} s(u, v) \|J(\mathbf{f}) - \mathbf{I}\|_F^2 ds, \quad (11)$$

where $s(u, v)$ is the saliency value at point (u, v) given in Section 4.2, $J(\mathbf{f})$ is the Jacobian of the warping \mathbf{f} defined in Equation (4), \mathbf{I} is the 2×2 identity matrix, and $\|\cdot\|_F$ denotes the Frobenius matrix norm. Note that $J(\mathbf{f})$ locally equals \mathbf{I} if $\mathbf{f}(u, v)$ does not distort the image at all. Intuitively, the energy function in Equation (11) allows translations and penalizes all other transformations. When computing the optimal warping, we enforce boundary conditions to ensure that each boundary vertex of the polygonal mesh remains on the boundary after warping. In particular, assume that we resize the image from $m \times n$ to $m' \times n'$, the boundary conditions

are

$$\begin{cases} \mathbf{c}_{i1} = 0 & \text{if } \mathbf{v}_{iu} = 0 \\ \mathbf{c}_{i1} = m' & \text{if } \mathbf{v}_{iu} = m \\ \mathbf{c}_{i2} = 0 & \text{if } \mathbf{v}_{iv} = 0 \\ \mathbf{c}_{i2} = n' & \text{if } \mathbf{v}_{iv} = n \end{cases}. \quad (12)$$

The optimal warp $\mathbf{f}^*(u, v)$ is the minimizer of the energy function in Equation (11):

$$\mathbf{f}^*(u, v) = \operatorname{argmin} E(\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n), \quad (13)$$

under the conditions (12).

Due to our continuous representation of the warping map, we can integrate other particular terms into our objective function to adjust the penalty for different transformations, such as rotation, uniform scaling, and inversion. For example, the deformation energy defined in [21] is written as

$$E_{TS}(\mathbf{f}) = \int_{\Omega} s(u, v) \|J(\mathbf{f}) - \sqrt{\det(J(\mathbf{f}))} \mathbf{I}\|_F^2 ds, \quad (14)$$

which penalizes all transformations other than translation and uniform scaling. In [21] and [46], a rotation-invariant energy, and a distortion energy permitting translation, similarity transform, and rotation are defined as

$$E_{TR}(\mathbf{f}) = \int_{\Omega} s(u, v) (C - I) ds, \quad (15)$$

and

$$E_{TSR}(\mathbf{f}) = \int_{\Omega} s(u, v) (\operatorname{tr}(C) - 2 \det(J(\mathbf{f}))) ds, \quad (16)$$

respectively, where $C = J^T J$. An improvement to the energy in Equation (16) is also designed in [18], which increases the penalty to infinity as the horizontal scaling factor goes towards zero:

$$\tilde{E}_{TSR}(\mathbf{f}) = \int_{\Omega} s(u, v) J^{-2} \|C\|^2 ds. \quad (17)$$



Figure 7. Comparison of linear and quadratic poly-FEM warping. We resize the image to 150% of its original height. Left to right: adaptive polygonal meshes, deformed polygonal meshes, and final warped images. (a–c) Linear poly-FEM warping with 100 polygonal elements and 462 DOFs. (d–f) Quadratic poly-FEM warping with 50 polygonal elements and 460 DOFs. (g–i) Quadratic poly-FEM warping with 30 polygonal elements and 280 DOFs.

4.5. Constrained Energy optimization

The energy functions in Equations (11), (15) and (16) allow for efficient minimization in a single Newton step, while others results in more complex optimization problems. On the other hand, the minimizer of the energies mentioned above may also introduce inversions, creating visible artifacts in the warping results. To optimize all the energies consistently, we adopt the iterative L-BFGS method, taking

into account inversion prevention at each iteration.

Imposition of the sufficient condition (8) on the deformed mesh during energy optimization ensures a bijective map result. However, it may also reduce the flexibility for optimization and lead to unsatisfactory warping results: see Figures 5(b,c) and 5(h,i). Thus, we relax the condition imposed in the optimization. Note also that if the warped polygonal mesh has fold-overs, then the mapping

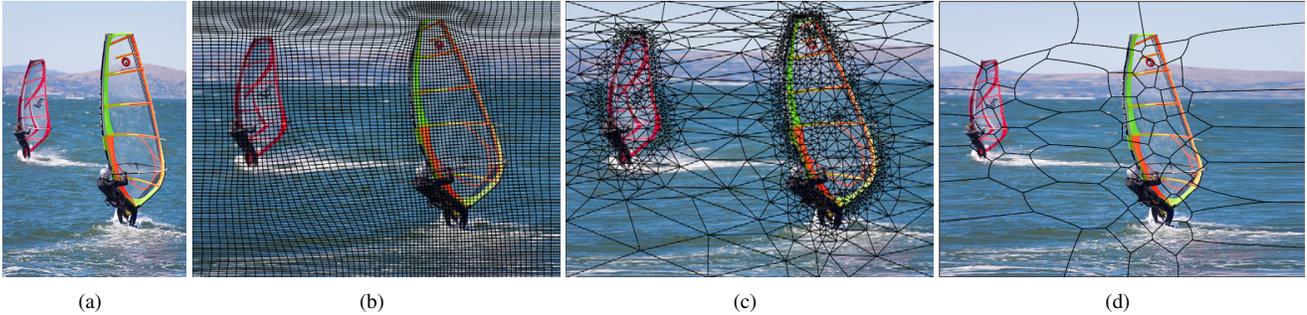


Figure 8. Comparison between different FEM-based warping methods, resizing an image to 200% of the original width. (a) Input. (b,c) Results from the DG FEM-based warping method using quadratic power polynomials on a quadrilateral mesh (6767 DOFs), and an adaptive triangular mesh (1325 DOFs neglecting the number of constraints for gluing elements with 2612 elements) from [18]. (d) Result using the proposed poly-FEM warping method (460 DOFs, 50 elements).



Figure 9. Results of the proposed warping method and manual interactive methods when resizing images to 75% of their original width. (a) Input. (b) Linear scaling result. (c) Result using DNR. (d) Result using BR [22]. (e) Our result using 50 elements (approx. 460 DOFs).

is not-bijective, which suggests an overlapping-free warped polygonal mesh to guarantee a bijective mapping necessarily. Our algorithm, therefore, attempts to satisfy the necessary conditions while maintaining bijectivity. In particular, we check whether the polygonal mesh overlaps after each L-BFGS iteration. If a vertex or midpoint locally introduces an overlap, its position is rolled back to the previous iteration. In detail, we maintain a list of overlapped cells after each L-BFGS iteration. Then we repeatedly remove a cell from the list, and roll back the positions of all the vertices and midpoints associated with the cell to the previous iteration. This rolling back may cause adjacent cells to overlap, and if so they are added to the list. When the list is empty, we stop.

The L-BFGS method involves integrating poly-FEM basis function derivatives. We evaluate these integrals by triangulating each polygonal cell and applying the quadrature rule to the resulting triangles. Based on short-edge free polygonal meshes, this iterative updating and rolling back

optimization strategy for energy minimization never introduced inversion in all examples considered in this paper. Figures 4(e,f) show an example of the deformed polygonal mesh and the final warped image. Figure 6 also shows deformations resulting from optimizing different energy functions using the proposed warping representation.

5. Experimental results

In this section, we demonstrate the versatility of our poly-FEM warping framework for content-aware image warping, providing statistics and comparisons. We compare our method to seventeen image retargeting methods, including five state-of-art methods (WSSCDNN [6], BR [22], Cycle-IR [38]), SAMIR [48] and DNR [7], and twelve classical methods (cropping (CR), scaling (SCL), seam carving (SC) [30], shift-maps (SM) [27], inhomogeneous warping (WARP) [44], scale-and-stretch optimized resizing (SNS) [43], energy-based deformation (LG) [17], Multiop [31], retargeting of streaming video (SV) [20], resizing

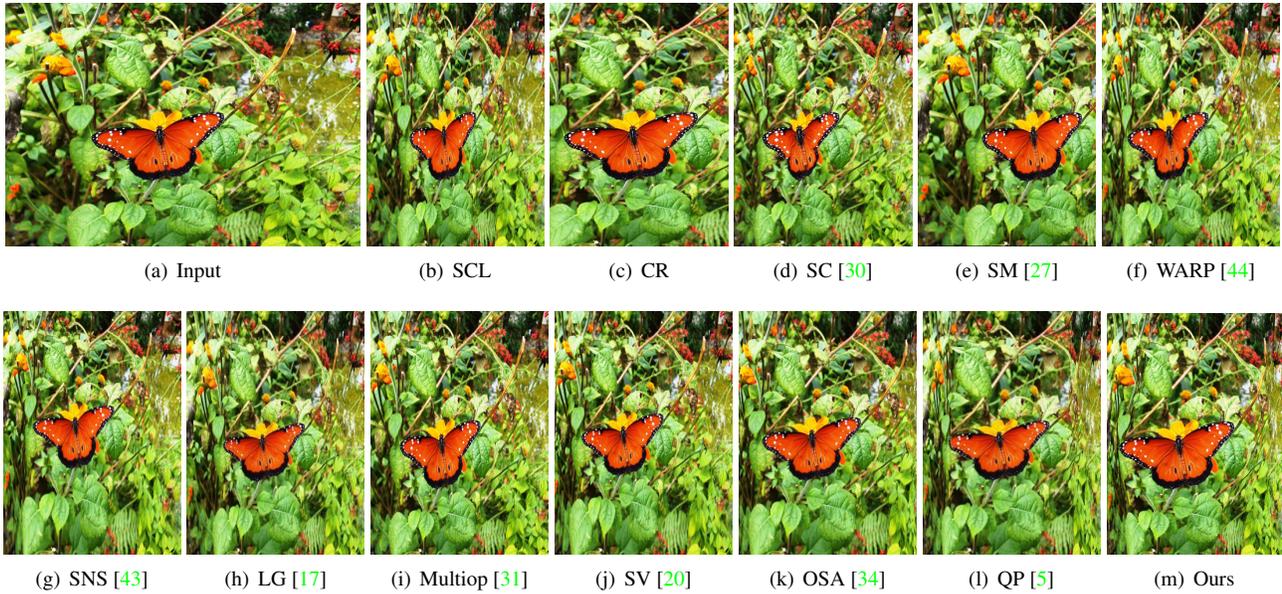


Figure 10. Results from the proposed warping method and 11 classic retargeting methods, resizing an image to 50% of its original width.

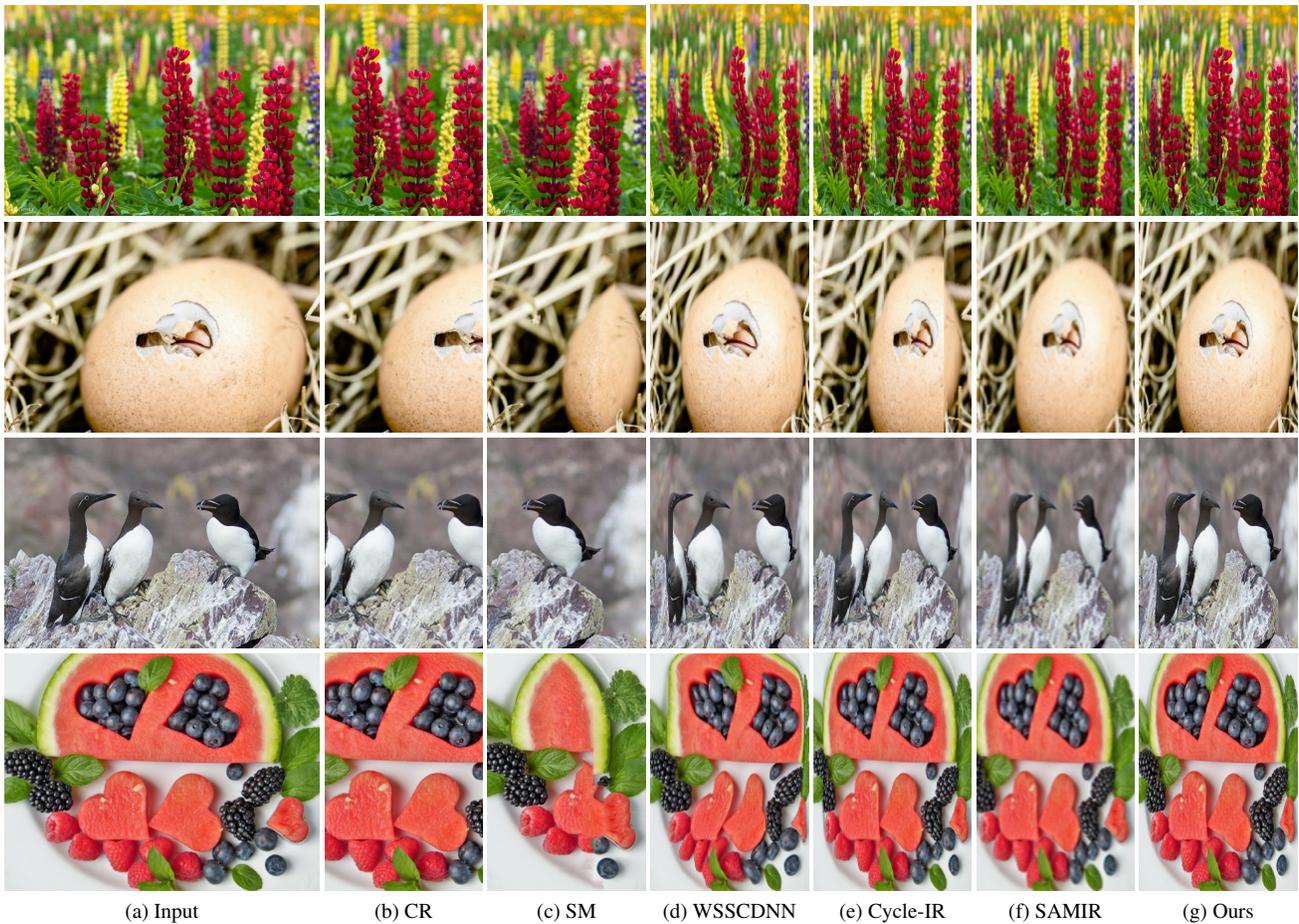


Figure 11. Results of the proposed method and deep learning-based image resizing to 50% of the original width. The retargetability scores of input images from top to bottom are 0.68, 0.56, 0.54 and 0.23, respectively [39].

Table 1. GRNN-OQA scores. A.S. = average score. H.R.. = number of highest-scoring results over all methods. L.R. = number of lowest-scoring results over all methods. N.E. = number of results generated by the corresponding method. Best results are shown in bold.

Method	CR	SCL	SC [30]	SM [27]	WARP [44]	SNS [43]	LG [17]	Multiop [31]
A.S.	0.61002	0.61906	0.61912	0.61804	0.61912	0.61989	0.61962	0.62043
H.R./N.E.	15/69	0/69	1/68	4/66	3/69	2/62	3/66	8/68
L.R./N.E.	1/69	0/69	1/68	1/66	0/69	0/62	0/66	0/68
SV [20]	OSA [34]	QP [5]	WSSCDNN [6]	BR [22]	Cycle-IR [38]	SAMIR [48]	DNR [7]	Ours
0.62036	0.62063	0.61872	0.60524	0.60831	0.61848	0.61935	0.5987	0.62102
2/69	3/68	3/66	0/69	0/5	2/69	4/68	0/19	19/69
0/69	0/68	1/66	46/69	0/5	2/69	0/68	17/19	0/69

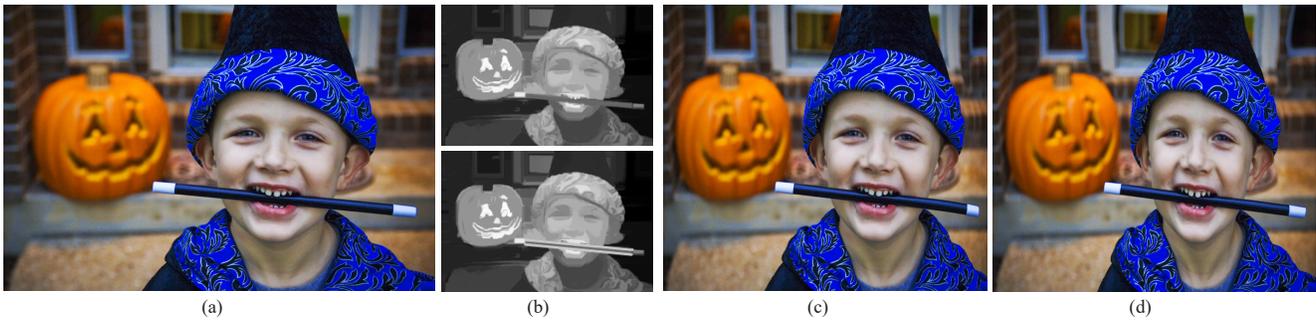


Figure 12. Resizing results to 75% of the original width while preserving straight line structures. (a) Input image. (b) Original saliency map (top) and saliency map with line structures highlighted (bottom). (c) Warping results based on original saliency map. (d) Warping results based on saliency map with line structures.

with object size adjusted (OSA) [34], resizing by quadratic programming (QP) [5] and DG FEM [18]). We use images from the elaborate RetargetMe benchmark [29]. The experimental environment used an Intel Core i5-8300 CPU with 16GB RAM.

5.1. Comparison to the FEM-based method

High-order polygonal elements allow smoother warping while using fewer DOFs. Figures 7(a–c) show examples of warping results using linear polygonal elements on a mesh with 100 elements and 462 DOFs. Figures 7(d–f) show results using quadratic elements with almost the same number of DOFs as results in Figure 7(c). We can observe that the quadratic poly-FEM achieves a much smoother deformation. Unnatural distortion in the result of the linear poly-FEM is highlighted in a red rectangle in Figure 7(c). Figures 7(g–i) show that we can get similar results with even fewer DOFs or cells.

Amongst other warping methods, the most closely related one to ours is the work of [18], where the warping map is defined using a DG FEM using power polynomials on a triangular or quadrilateral mesh. As shown in Figure 8(b), traditional FEM-based warping methods using regular quadrilateral meshes require more DOFs, and a very dense mesh, to obtain satisfactory warping results. On the other hand, warping based on triangular or polygo-

nal meshes achieves a similar result to the structured mesh-based method with far fewer degrees of freedom, as unstructured meshes can provide local refinement and improve the utilization of degrees of freedom: see Figures 8(c,d).

The warping map of the DG-FEM-based method, with individually defined power polynomials on each cell, is discontinuous between elements. Additional DOFs are necessary in the warping function optimization to restore the coupling between elements. If the discontinuities between elements are too large, function values on edges cannot be safely averaged, leading to artifacts. To achieve a visually pleasing result, one must resolve the warping optimization problem on denser meshes. We can observe that the DG FEM-based method has to use a large number of cells to introduce enough DOFs to achieve visually satisfying results: see Figure 8(c). Instead, our method generates naturally smooth warps between elements without specifying any additional constraints. Hence, we can achieve smooth warps with any number of elements: see Figure 8(d).

5.2. Comparison to interactive methods

Automated saliency map generation techniques cannot yet reach the quality of manual methods. To achieve more satisfactory results or avoid failures of automatic algorithms, many works allow user intervention to specify saliency maps. For example, the BR method requires

manual specification of important regions [22]. We compare our proposed poly-FEM warping method to the BR method in Figure 9. Our automatic method generates comparable results to those of the BR method. The DNR method [7] requires manual choice of an appropriate threshold to switch from seam-carving to warping in their multi-operator scheme to achieve good results. Seam removal may introduce artifacts in the final output, especially in less important regions: see Figure 9(c). By contrast, our method achieves visually more satisfactory results.

5.3. Comparison to other methods

We first compare the results of our Poly-FEM method to classic and state-of-the-art retargeting methods on all examples from the most popular benchmark, RetargetMe (with 80 images in total). Due to lack of space, we only show one set of results from all the classic methods in Figure 10 for subjective judgement. We can observe that our approach better avoids unexpected information loss and preserves salient content in this example. To avoid time-consuming and laborious personal quality assessment for the retargeting results of the remaining test images in RetargetMe, we evaluated retargeting quality using the General Regression Neural Network-based Objective Quality Assessment (GRNN-OQA) in [24].

Values from GRNN-OQA are in $[0, 1]$, where a higher value indicates a better quality of the retargeted image. The GRNN-OQA method computes scores which aim to preserve the ranking of subjective scores for same-source results and to provide a reference to compare different-source results. Since each existing retargeting method exhibits its own advantages and limitations, no single method works better than other methods for all the test images. For each retargeting method, we compute the average GRNN-OQA score over RetargetMe. Note that the GRNN-OQA method provides a trained model for ranking resized images' with scaling factors of 0.75 and 0.5. Hence, we only report the average scores when resizing images in this way. The number of times each method ranked highest or lowest in the ranking of retargeting results of the same source image are also reported: see Table 1. Overall, our method is more robust and better than other retargeting methods, in the senses that our method gets the highest average scores, the most highest-scoring results, and the fewest lowest-scoring results.

In [39], the retargetability score in the range $[0, 1]$ is computed to measure the retargeting difficulty of an input image, where a low value indicates a high level of difficulty. Images with retargetability scores between $(0.0; 0.75]$ are suggested for retargeting method assessment. To further evaluate the effectiveness of our proposed method, we conducted experiments on four images with low or moderate retargetability scores: see Figure 11(a). We illustrate

results for a scaling factor of 0.5 using two selected classic methods and the three deep learning-based methods, for subjective judgement. We observe that, due to rich content or geometric structures, the retargeting results present severe artifacts for CR, SM, WSSCDNN, and cycle-IR methods: see Figures 11(b–e). Note that the SAMIR method generates an optimal multi-operator sequence, which needs to scale the image to half the original size as preprocessing. This simple re-scaling of the image leads to blurred results: see Figure 11(f). Compared to these methods, our method better preserves salient content and introduces fewer unnatural artifacts in images with rich content: see Figure 11(g).

5.4. Speed

All our results in this paper used a small number of elements (50–200) and DoFs (450–2000), and the whole computation can be done within 0.5 s. For example, the time to warp an image of size 640×480 to 50% of its original width using 200 elements is 0.417 s; the most time-consuming step is saliency map generation, taking more than 50% of the total time. Our method is slightly more time-consuming than Cycle-IR (0.203 s) and faster than WSSCDNN (0.982 s), SAMIR (25.23 s CPU time), and DNR (60–100 s [7]) when resizing an image of this size.

6. Discussion and conclusion

We have introduced a poly-FEM-based image warping representation and provided a framework for content-aware image warping. The proposed poly-FEM warping method considerably improves the representation power of the warping map, so achieves more satisfactory results than other existing methods while using a small number of DOFs.

Despite the generally promising results shown in the paper, the proposed poly-FEM warping method suffers from two major limitations. First, it relies on saliency detection results, like other content-aware image retargeting techniques. Hence, one of our future works is to investigate more sophisticated and efficient saliency detection methods suitable for content-aware image warping. Second, the proposed method may fail to preserve geometric structures, such as lines, especially when line structures occur in low-saliency regions and the distortion distribution is strongly uneven: see the pen in Figure 12(c). Note that the line structure in low-saliency regions can be well preserved, e.g. the road in Figures 7(f,i) and Figure 9(e), perhaps because the optimal quadratic poly-FEM-based deformation functions are close to linear in the corresponding regions. One possible solution to preserve such line structures is to mark them as salient, resulting in reduced shape deformation in the warping. As shown at the bottom of Figure 12(b), we manually assign higher salient values to line structures by painting lines with more bright colors. These line struc-

tures are well preserved in the final result: see Figure 12(d). We would like to improve the proposed poly-FEM warping method to include automatic line detection and line structure preservation into our framework.

In this paper, optimal warping is computed on a polygonal mesh with a pre-specified number of cells. We may decrease the number of cells without obviously sacrificing deformation quality: see Figures 4(e), 8(d) and 7(d,g) for examples. Hence, we would like to design an automatic method for determining an appropriate number of cells to better balance the mesh size (and so running time) and the warp quality. Benefiting from the Lagrange interpolation property and linear precision property of the high-order basis functions used, the warping map is C^∞ smooth in the interiors of cells and is C^0 stitched along cell boundaries. It would also be possible to include barycentric coordinates with the Hermite interpolation property into our framework, e.g. cubic MVC [23], to get a warping map with higher-order continuity along cell boundaries. In addition, our current framework focuses on image warping. We plan to extend the proposed method to retargeting video by including motion features into the importance map and ensuring consistency of the warping grid.

Acknowledgement

We would like to thank the reviewers for their valuable comments. The research of Juan Cao was supported by the NSFC (Nos. 61872308, 61972327), the NSF of Fujian Province of China (No. 2019J01026), the Fundamental Research Funds for the Central Universities (No. 20720190011), and the Xiamen Youth Innovation Funds (No. 3502Z20206029). Yongjie Jessica Zhang was supported in part by NSF CMMI-1953323 and a Honda grant.

References

- [1] S. Avidan and A. Shamir. Seam carving for content-aware image resizing. *ACM Transactions on Graphics*, 26(3):10–es, July 2007. [2](#)
- [2] C. Bajaj, Y. Zhang, and G. Xu. Physically-based surface texture synthesis using a coupled finite element system. In *International Conference on Geometric Modeling and Processing*, pages 344–357. Springer, 2008. [3](#)
- [3] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3), July 2009. [2](#)
- [4] J. Cao, Y. Xiao, Z. Chen, W. Wang, and C. Bajaj. Functional data approximation on bounded domains using polygonal finite elements. *Computer aided geometric design*, 63:149–163, 2018. [3](#)
- [5] R. Chen, D. Freedman, Z. Karni, C. Gotsman, and L. Liu. Content-aware image resizing by quadratic programming. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, pages 1–8, 2010. [2](#), [11](#), [12](#)
- [6] D. Cho, J. Park, T.-H. Oh, Y.-W. Tai, and I. So Kweon. Weakly-and self-supervised learning for content-aware deep image retargeting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4558–4567, 2017. [2](#), [10](#), [12](#)
- [7] D. Dano, M. Arar, D. Cohen-Or, and A. Shamir. Image resizing by reconstruction from deep features. *Computational Visual Media*, 7:453–466, 2021. [2](#), [3](#), [10](#), [12](#), [13](#)
- [8] Q. Du, V. Faber, and M. Gunzburger. Centroidal Voronoi tessellations: applications and algorithms. *SIAM Review*, 41(4):637–676, 1999. [6](#)
- [9] U. Engelke, J. Wang, and P. Marendy. Perceptual relevance based image retargeting. *IEEE Signal Processing Letters*, 22(6):705–708, June 2015. [5](#)
- [10] M. S. Floater and M.-J. Lai. Polygonal spline spaces and the numerical solution of the Poisson equation. *SIAM Journal on Numerical Analysis*, 54(2):797–824, 2016. [3](#)
- [11] J. C. Gee, D. R. Haynor, M. Reivich, and R. Bajcsy. Finite element approach to warping of brain images. In *SPIE Proceedings*, volume 2167 of *Medical Imaging: Image Processing*, 1994. [3](#)
- [12] Y. Guo, F. Liu, J. Shi, Z.-H. Zhou, and M. Gleicher. Image retargeting using mesh parametrization. *IEEE Transactions on Multimedia*, 11(5):856–867, Aug. 2009. [2](#)
- [13] S. Gupta and S. G. Mazumdar. Sobel edge detection algorithm. *International Journal of Computer Science and Management Research*, 2(2):1578–1583, 2013. [7](#)
- [14] K. Hormann and M. S. Floater. Mean value coordinates for arbitrary planar polygons. *Acm Transactions on Graphics*, 25(4):1424–1441, Oct. 2006. [3](#)
- [15] S. M. Hu, C. F. Li, and H. Zhang. Actual Morphing: A Physical-Based Approach for Blending Two 2D/3D Shapes. In G. Elber, N. Patrikalakis, and P. Brunet, editors, *Solid Modeling*. The Eurographics Association, 2004. [3](#)
- [16] A. Jacobson, E. Tosun, O. Sorkine, and D. Zorin. Mixed finite elements for variational surface modeling. *Computer Graphics Forum*, 29(5):1565–1574, 2010. [3](#)
- [17] Z. Karni, D. Freedman, and C. Gotsman. Energy-based image deformation. In *Proceedings of the Symposium on Geometry Processing*, SGP '09, page 1257–1268, Goslar, DEU, 2009. Eurographics Association. [2](#), [10](#), [11](#), [12](#)
- [18] P. Kaufmann, O. Wang, A. Sorkine-Hornung, O. Sorkine-Hornung, A. Smolic, and M. Gross. Finite element image warping. *Computer Graphics Forum*, 32:31–39, 2013. [3](#), [5](#), [8](#), [10](#), [12](#)
- [19] J. Kiess, S. Kopf, B. Guthier, and W. Effelsberg. A survey on content-aware image and video retargeting. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 14(3):Article 76, 2018. [2](#)
- [20] P. Krähenbühl, M. Lang, A. Hornung, and M. Gross. A system for retargeting of streaming video. *ACM Trans. Graph.*, 28(5):1–10, Dec. 2009. [2](#), [10](#), [11](#), [12](#)
- [21] P.-Y. Laffont, J. Y. Jun, C. Wolf, Y.-W. Tai, K. Idrissi, G. Drettakis, and S.-e. Yoon. Interactive content-aware zooming. In *Proceedings of Graphics Interface 2010*, GI '10,

- page 79–87, CAN, 2010. Canadian Information Processing Society. **8**
- [22] C. P. Lau, C. P. Yung, and L. M. Lui. Image retargeting via Beltrami representation. *IEEE Transactions on Image Processing*, 27(12):5787–5801, Dec. 2018. **2, 10, 12, 13**
- [23] X.-Y. Li, T. Ju, and S.-M. Hu. Cubic mean value coordinates. *ACM Transactions on Graphics*, 32(4):126:1–10, 2013. **14**
- [24] Y. Liu, Y. Han, Z. Ye, and Y. Lai. Ranking-preserving cross-source learning for image retargeting quality assessment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019. **13**
- [25] M. Meyer, A. Barr, H. Lee, and M. Desbrun. Generalized barycentric coordinates on irregular polygons. *J. Graph. Tools*, 7(1):13–22, Nov. 2002. **3**
- [26] M. Nishiyama, T. Okabe, Y. Sato, and I. Sato. Sensation-based photo cropping. In *Proceedings of the 17th ACM International Conference on Multimedia*, page 669–672, New York, NY, USA, 2009. **2**
- [27] Y. Pritch, E. Kav-Venaki, and S. Peleg. Shift-map image editing. In *2009 IEEE 12th International Conference on Computer Vision*, pages 151–158, 2009. **2, 10, 11, 12**
- [28] A. Rand, A. Gillette, and C. Bajaj. Quadratic serendipity finite elements on polygons using generalized barycentric coordinates. *Mathematics of Computation*, 83(290):2691, 2011. **3, 4**
- [29] M. Rubinstein, D. Gutierrez, O. Sorkine, and A. Shamir. A comparative study of image retargeting. *ACM Transactions on Graphics*, 29(6):160:1–160:10, 2010. **12**
- [30] M. Rubinstein, A. Shamir, and S. Avidan. Improved seam carving for video retargeting. *ACM Transactions on Graphics*, 27(3):1–9, 2008. **2, 10, 11, 12**
- [31] M. Rubinstein, A. Shamir, and S. Avidan. Multi-operator media retargeting. *ACM Trans. Graph.*, 28(3), July 2009. **2, 10, 11, 12**
- [32] T. Schneider. *Theory and applications of bijective barycentric mappings*. PhD thesis, Università della Svizzera italiana, 2017. **6**
- [33] V. Setlur, S. Takagi, R. Raskar, M. Gleicher, and B. Gooch. Automatic image retargeting. In *Proceedings of the 4th International Conference on Mobile and Ubiquitous Multimedia*, page 59–68. Association for Computing Machinery, 2005. **2**
- [34] M. Shi, L. Yang, G. Peng, and D. Xu. A content-aware image resizing method with prominent object size adjusted. In *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology, VRST '10*, page 175–176, New York, NY, USA, 2010. Association for Computing Machinery. **2, 11, 12**
- [35] D. Sieger, P. Alliez, and M. Botsch. Optimizing voronoi diagrams for polygonal finite element computations. In S. Shontz, editor, *Proceedings of the 19th International Meshing Roundtable*, pages 335–350, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. **6, 7**
- [36] B. Suh, H. Ling, B. B. Bederson, and D. W. Jacobs. Automatic thumbnail cropping and its effectiveness. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology*, page 95–104, New York, NY, USA, 2003. Association for Computing Machinery. **2**
- [37] N. Sukumar. Quadratic maximum-entropy serendipity shape functions for arbitrary planar polygons. *Computer Methods in Applied Mechanics and Engineering*, 263(263):27–41, 2013. **3**
- [38] W. Tan, B. Yan, C. Lin, and X. Niu. Cycle-IR: deep cyclic image retargeting. *IEEE Transactions on Multimedia*, pages 1–1, 2019. **2, 10, 12**
- [39] F. Tang, W. Dong, Y. Meng, C. Ma, F. Wu, X. Li, and T.-Y. Lee. Image retargetability. *IEEE Transactions on Multimedia*, 22(3):641–654, 2020. **11, 13**
- [40] D. Vaquero, M. Turk, K. Pulli, M. Tico, and N. Gelfand. A survey of image retargeting techniques. In A. G. Tescher, editor, *Applications of Digital Image Processing XXXIII*, volume 7798, pages 328 – 342. International Society for Optics and Photonics, SPIE, 2010. **2**
- [41] E. L. Wachspress. *A rational finite element basis*. Elsevier, 1975. **3**
- [42] J. Wang, H. Jiang, Z. Yuan, M.-M. Cheng, X. Hu, and N. Zheng. Salient object detection: A discriminative regional feature integration approach. *International Journal of Computer Vision*, 123(2):251–268, 2017. **5**
- [43] Y.-S. Wang, C.-L. Tai, O. Sorkine, and T.-Y. Lee. Optimized scale-and-stretch for image resizing. *ACM Trans. Graph.*, 27(5), Dec. 2008. **2, 10, 11, 12**
- [44] L. Wolf, M. Guttman, and D. Cohen-Or. Non-homogeneous content-driven video-retargeting. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–6, 2007. **2, 10, 11, 12**
- [45] H.-B. Yan, S.-M. Hu, and R. Martin. Morphing based on strain field interpolation. *Computer Animation and Virtual Worlds*, 15(3-4):443–452, 2004. **3**
- [46] G.-X. Zhang, M.-M. Cheng, S.-M. Hu, and R. R. Martin. A shape-preserving approach to image resizing. *Computer Graphics Forum*, 28(7):1897–1906, 2009. **8**
- [47] Y. J. Zhang. *Geometric modeling and mesh generation from scanned images*. Chapman and Hall/CRC, 2018. **3**
- [48] Y. Zhou, Z. Chen, and W. Li. Weakly supervised reinforced multi-operator image retargeting. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(1):126–139, 2021. **2, 3, 10, 12**