

# Image warp preserving content intensity

Enrico Segre\*

**Abstract.** An accurate method for warping images is presented. Differently from most commonly used techniques, this method guarantees the conservation of the intensity of the transformed image, evaluated as the sum of its pixel values over the whole image or over corresponding transformed subregions of it. Such property is mandatory for quantitative analysis, as, for instance, when deformed images are used to assess radiances, to measure optical fluxes from light sources, or to characterize material optical densities. The proposed method enforces area resampling by decomposing each rectangular pixel in two triangles, and projecting the pixel intensity onto half pixels of the transformed image, with weights proportional to the area of overlap of the triangular half-pixels. The result is quantitatively exact, as long as the original pixel value is assumed to represent a constant image density within the pixel area, and as long as the coordinate transformation is diffeomorphic. Implementation details and possible variations of the method are discussed.

**Key words.** Warping, Area resampling, Image distortion, Photometry

**AMS subject classifications.** 68U10, 65D18, 54H30

**1. Introduction.** Many scientific procedures which make quantitative use of the image content, in fields which range from microscopy to astronomy, require images to be transformed and remapped onto deformed coordinates system. Typical applications include the correction of geometrical aberrations produced by imaging systems, the mutual registration of scenes recorded with different optical systems or from different points of view [44, 23]; stitching together different images with partial overlap among themselves [13], also referred to as multi-frame joint image registration; the fusion and proper coadding of different images of the same source fields [42]. In medical imaging and computational anatomy in particular, to name another application, cross image registration is always required to properly compare features of compliant soft tissue. Furthermore, dynamic mapping of image sequences over deforming templates was used with expressive intent in yesteryears, in a procedure called “morphing” [38, 27] where the appearance of one object was transformed smoothly into that of another (e.g. a human figure into an animal) by means of gradual deformation and blending.

For our purposes, we consider two dimensional images, generically represented as two dimensional arrays of values of the intensity over Cartesian grids. The methods for determining the appropriate geometrical transformation between the source and the target coordinates are varied and sophisticated [14, 24], depend on the task, and are not themselves of concern of this paper. Such methods may make use of functional relations between the coordinate systems known a priori, or may rely on the identification of common landmark features appearing in the images [3, 20], either known from supervised annotation, from model fit or from trained deep learning (e.g. [41, 43]). In computational anatomy, for instance, diffeomorphic flow is assumed between source and target images, and LDDMM [2] in a number of variants is very popular. A large body of literature exists on these methods, which do not need to be reviewed here.

---

\*Physics Core Facilities, Weizmann Institute of Science, Rehovot, Israel, ([enrico.segre@weizmann.ac.il](mailto:enrico.segre@weizmann.ac.il))

Once the functional transformation which maps the two systems of coordinates is established, a “best” way of transforming also the image values is sought. The acceptance of “best” is sometimes subjective and in many cases depends on the application: it may refer to a cosmetically pleasing result, to an optimal way of representing and preserving sharp level transitions in the destination image, or to the suppression of moiré or aliasing artifacts. Optimal ways of prefiltering and resampling of transformed images by means of interpolation are well described in literature [26, 1, 12] and implemented in widely adopted software libraries (e.g. OpenCV [4], scikit [35], ImageMagick [18]) as well as in graphic applications. In other cases, like in superresolution imaging and reconstruction [22], the recovery of realistic, underresolved image details is achieved relying on a priori subscale models, or optimal use of information resulting from multiple low resolution images belonging to a sequence. In this paper, in contrast, we describe a procedure which is purely intended to preserve the brightness of the image content across the transformation, even when the images are not Nyquist sampled, without invoking any help from the image data itself or from a priori knowledge of structure lost by the process of image formation. In simple terms, we exactly redistribute the whole intensity content of the source pixels over the target raster. The procedure is linear, and amounts to the determination of a reweighting matrix which projects the pixel values from the source to the destination image, and most importantly, depends only on the coordinate transform and not on the image data itself. As such, some variants of the procedure can be devised from the basic scheme, including one which provides a stable alternative to image interpolation without ad hoc filtering.

Our procedure implements a rigorous area resampling. The concept is known even from earlier literature, but does not seem to have received adequate attention, probably because of its higher computational cost which hinders its applications, and does not seem to have been pursued in the general case. Early attempts include that of [9], which proposes a fast implementation, based on a scanline decomposition. Scanline approaches treat the deformation of a raster image by carrying on some of the intensity content from one pixel to its adjacent in scan order, and are not proven to be exact for transformations beyond simple shears. A simpler version of Fant’s algorithm, applied only to raster resizing, goes under the name of pixel mixing [31] and was probably implemented in the open first by the *pamscale* function of the netpbm package [28]. The thesis [5] generalizes the problem, introducing the term “imaging-consistent integrating resampler”, taking into account also the point spread function of the imager, and blurring in image formation. The algorithm proposed there, though, still falls within the category of separable, scanline approaches, with a single accumulation register providing intensity remainders carried over from one pixel to the next. Another double pass, scanline algorithm is that of [16]. The seminal thesis [17] discusses the problem, and gives a partial solution in terms of adaptive local deformation of circular neighborhoods. This method is also implemented in the popular software package ImageMagick [33]. Another cognate approach proposed, and employed specifically for oversampling stacks of dithered astronomical images while preserving photometry, is “Drizzle” [11], but it relies on empirical factors, and treats both source and destination pixels as squares. “Drizzling” estimates pixel area overlaps using a sort of a Montecarlo approach, where the randomness is provided by inter-image pixel shifts. As a procedure, it some offers other advantages like the possibility of assigning individual quality weights to each contributing pixel; still it is not general for arbitrary deformations.

In summary, the existing literature on area resampling concentrated on the search for “efficient” variants of the method, or which seem to lack generality or exactness, when specifically looking at the preservation of the photometric intensity. It is our intention to discuss here a rigorous procedure, and its implementation.

The paper is organized as follows: section 2 describes the geometrical principle of pixel remapping, section 3 outlines the algorithm used, section 4 demonstrates it, section 5 discusses some variations, section 6 shows the advantage of area resampling in photometric measurements, and section 7 concludes and outlines future perspectives. Computational details are included in the appendices: the convention adopted for barycentric coordinates is given in appendix A; appendix B discusses the problem of intersecting triangles, and code performance is reported in appendix C.

**2. Exact area resampling by pixel triangulation.** We start from an intensity image, given as a set of  $N_1 \times M_1$  pixel values  $I_1(i, j)$ , representing the cumulative value of some quantity (for instance, the number of photons impinging the area of an individual photosensitive element), integrated over the rectangular pixel  $p_{ij}$ , defined as the rectangle  $x_i \leq x < x_{i+1}$ ,  $y_j \leq y < y_{j+1}$ , for  $1 \leq i \leq N_1$  and  $1 \leq j \leq M_1$ . For simplicity we will treat here an equispaced coordinate grid,  $x_i = x_1 + (i - 1) \cdot \Delta x$  and  $y_j = y_1 + (j - 1) \cdot \Delta y$ , though the procedure can be easily generalized to non-equispaced plaid grids. We consider an a priori given bijective and differentiable coordinate transformation  $(X, Y) = f(x, y)$ . We assume that the underlying intensity density  $i_1(x, y)$  inside the pixel  $p_{ij}$  is uniform, and that  $\int_{p_{ij}} i_1(x, y) dx dy = I_1(i, j)$ . Therefore,  $i_1(x, y) = I_1(i, j) / \mathcal{A}[p_{ij}]$ , where  $\mathcal{A}[p_{ij}] = \Delta x \Delta y$  is the area of the pixel.

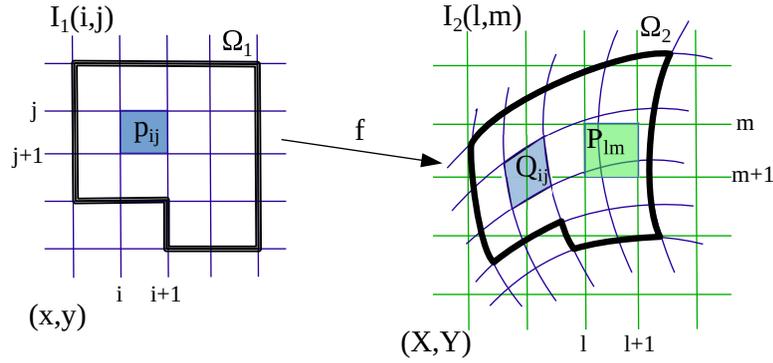
Our goal is to produce a new image of  $N_2 \times M_2$  pixels, transforming the set of values  $I_1$  into a new set  $I_2(l, m)$  on a new equispaced grid  $\{(X_l, Y_m)\}$ , with  $1 \leq l \leq N_2$ ,  $1 \leq m \leq M_2$  and spacing  $\Delta X$ ,  $\Delta Y$ , in such a way that the cumulative intensity within any closed contour is preserved by the transformation:

$$(2.1) \quad \int_{\Omega_1} i_1(x, y) dx dy = \int_{\Omega_2} i_2(X, Y) dX dY$$

for any region  $\Omega_2 = f(\Omega_1)$ , and assuming an underlying transformed intensity density  $i_2$  in the destination image. A natural way of achieving this property is to consider the quadrilateral  $Q_{ij} = \{f(x_i, y_j), f(x_{i+1}, y_j), f(x_{i+1}, y_{j+1}), f(x_i, y_{j+1})\}$ , which approximates (to second order in  $\Delta x$ ,  $\Delta y$ ) the transform of the rectangular pixel  $p_{ij}$ , identified by the set of its four vertices  $\{(x_i, y_j), (x_{i+1}, y_j), (x_{i+1}, y_{j+1}), (x_i, y_{j+1})\}$  (see Fig. 2.1). Save for singular or extreme deformations and coarse griddings which are of little practical interest, we can tacitly assume that  $Q_{ij}$  remains a convex quadrilateral (concavity would imply a change of sign of the Jacobian of the transformation, violating the assumption of diffeomorphic transformation). For shorthand, we write  $Q_{ij} \simeq f(p_{ij})$ , applying  $f()$  to polygons and contours as well as to individual points. Locally, this scalar density would be transformed as

$$(2.2) \quad i_2(X, Y) = J_f i_1(x, y) = \left| \begin{array}{cc} \frac{\partial X}{\partial x} & \frac{\partial X}{\partial y} \\ \frac{\partial Y}{\partial x} & \frac{\partial Y}{\partial y} \end{array} \right| i_1(x, y),$$

so that, to second order, (2.1) is satisfied for  $Q_{ij}$  and hence for any region composed of sets of pixels of image 1. Within the same approximation, we assume that the Jacobian  $J_f$  is



**Figure 2.1.** Pixel transform  $p_{ij} \rightarrow Q_{ij}$  from the source to the destination image space, and mapping of closed contours  $\Omega_1 \rightarrow \Omega_2$  including specific pixel groups.

constant within  $p_{ij}$  and the density  $i_2$  constant within  $Q_{ij}$ . This position allows us to reduce the change of integration variable in Eq. (2.1) into a problem of decomposition of pixel areas: each fraction of  $Q_{ij}$  projected on the destination grid, will contribute to the target intensity proportionally to its fractional area only.

The idea of considering the shape change of the pixel in the transformation is not new (see for example of the procedure described in §15.5 of [36]); however, our development is different in that we do not invoke arbitrary interpolations for the reconstruction of the destination image. The intensity of the rectangular pixel

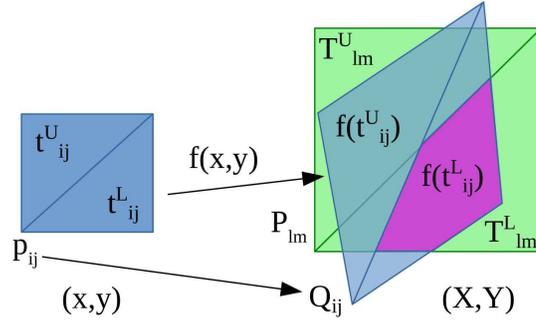
$$P_{lm} = \{(X_l, Y_m), (X_{l+1}, Y_m), (X_{l+1}, Y_{m+1}), (X_l, Y_{m+1})\}$$

on the target image is expressed as a sum of contributions

$$(2.3) \quad I_2(l, m) = \sum_{\text{overlaps}} I_2^{ij}(l, m) = \sum_{\text{overlaps}} \frac{\mathcal{A}[P_{lm} \cap Q_{ij}]}{\mathcal{A}[Q_{ij}]} I_1(i, j),$$

from each of the transformed pixels  $Q_{ij}$  of image 1 partially overlapping with  $P_{lm}$  in image 2. The subset of indices  $i, j$  to be taken into account is indicated here generically as “overlaps”; a criterion for selecting them will be formulated in the following. In (2.3),  $\mathcal{A}[\ ]$  indicates the area of the resulting polygon. The procedure involves therefore two steps: 1) for any given destination pixel  $P_{lm}$  identify the set of original pixels  $p_{ij}$  whose transform  $Q_{ij}$  overlaps with it, and 2) determine the polygonal shape of each intersection and compute its area.

The intersection of two convex quadrangles can be, in general, a polygon with anything between three and eight sides. Algorithms for the intersection of generic polygons exist in reputable computer geometry packages (e.g. in CGAL [10]), but their generality comes as a hindrance for our specialized case, requiring peculiar organized data structures, and is not necessarily optimal for a fast calculation. We prefer to simplify the task one step further. We divide both the origin pixel  $p_{ij}$  and the destination pixel  $P_{lm}$  in two triangles, splitting the quadrangles arbitrarily along one of their two diagonals, for instance  $t_{ij}^U = \{(x_i, y_j), (x_{i+1}, y_{j+1}), (x_i, y_{j+1})\}$  and  $t_{ij}^L = \{(x_i, y_j), (x_{i+1}, y_j), (x_{i+1}, y_{j+1})\}$ , and analogously  $T_{lm}^U$  and  $T_{lm}^L$  (Fig. 2.2).



**Figure 2.2.** Decomposition of a source pixel  $p_{ij} = t_{ij}^U \cup t_{ij}^L$  in two triangles in the origin and in the destination space,  $Q_{ij} = f(p_{ij}) = f(t_{ij}^U) \cup f(t_{ij}^L)$ , and its overlap with the destination pixel  $P_{lm} = T_{lm}^U \cup T_{lm}^L$ . The intersection  $T_{lm}^L \cap f(t_{ij}^L)$ , in this case a pentagon, is highlighted in purple for illustration.

The task of identifying intersections between  $N_1 \times M_1$  origin and  $N_2 \times M_2$  destination pixels, required by Eq. (2.3), is split in that of finding the intersections between four times as many triangular half pixels. This is still non trivial, as there are no less than 17 possible, topologically different ways of intersecting two triangles (Figure B.1), excluding degenerate cases, as further discussed in appendix B, but is definitely a simpler task than for quadrangles.

The same partial intensity  $I_1(i, j)/2$  is assigned to each of the two original triangles. Eq. (2.3) therefore expands into

$$(2.4) \quad I_2(l, m) = \sum_{\text{overlaps}} \left[ \frac{\mathcal{A}[T_{lm}^U \cap f(t_{ij}^U)] + \mathcal{A}[T_{lm}^L \cap f(t_{ij}^U)]}{2\mathcal{A}[f(t_{ij}^U)]} + \frac{\mathcal{A}[T_{lm}^U \cap f(t_{ij}^L)] + \mathcal{A}[T_{lm}^L \cap f(t_{ij}^L)]}{2\mathcal{A}[f(t_{ij}^L)]} \right] I_1(i, j)$$

Formally, the transformation between pixel intensities from the one to the other image can be written as

$$(2.5) \quad I_2(l, m) = \sum_{i,j=1}^{N_1, M_1} B_{lm,ij} I_1(i, j),$$

where  $B_{lm,ij}$  is the incidence matrix detailing which fraction of  $Q_{ij}$  intersects with  $P_{lm}$ . In typical cases, in which pixels of the original and of the destination rasters are of comparable sizes, this matrix is very sparse.

For transformations in which a source pixel  $p_{ij}$  is completely mapped on pixels on the the destination raster, the property  $\sum_{lm} B_{lm,ij} = 1$  holds. Conversely,  $\sum_{ij} B_{lm,ij}$  gives a discretized representation of  $J_f^{-1}$  on the destination raster.

As an aside, once the transformation from  $I_1$  to  $I_2$  has been computed according to Eq. (2.5), its inverse can be obtained directly inverting the sparse matrix  $B$ , for which numerical techniques are well studied. This may be more advantageous than using the inverse

coordinate mapping, if using an algorithm like the one described in the next section, which exploits the fact that the source image grid is cartesian.

**3. Algorithm layout.** As resulting from (2.5), the intensity transformation between the two rasters reduces to a simple matrix multiplication, once the matrix elements  $B_{lm,ij}$  are computed. To this extent, the steps sketched in inset 3.1 are required.

---

**Algorithm 3.1** Intensity transformation between  $I_1$  and  $I_2$

---

For each pixel of the source image, i.e. iterating on  $i$  and  $j$ :

1. the coordinates of the vertices of each original hemipixel  $t_{i,j}^{L,U}$ , are transformed with  $f$
2. the set of hemipixels  $\{T_{lm}^{L,U}\}$  which have a non empty overlap with  $f(t_{i,j}^{L,U})$  is determined
3. the area of the intersections between each of the triangles of this set and each  $f(t_{i,j}^{L,U})$  in turn, is found.
4. the areas  $\mathcal{A}\left[f\left(t_{i,j}^U\right)\right]$  and  $\mathcal{A}\left[f\left(t_{i,j}^L\right)\right]$  are computed.
5. the relevant contributions are summed to construct the matrix  $B_{lm,ij}$ ,

Finally, the image  $I_2$  is obtained by (2.5).

---

Step 1 is the simple evaluation of a given function  $f$  of the coordinates, and does not need to be described here. In the terminology of image processing, we are using naturally here a *forward* mapping between source and destination image.

For step 2, all triangles  $T$  which have at least one vertex within the bounding box  $[\min_X(Q_{ij}), \max_X(Q_{ij})] \otimes [\min_Y(Q_{ij}), \max_Y(Q_{ij})]$  are selected ( $\otimes$  denoting the Cartesian product of the two intervals). Since the triangular half pixels  $T$  are defined on a structured grid, they can be indexed in such a way that the criterion is translated to a simple choice of indices, involving integer arithmetics. It is algorithmically simpler to use this simplified criterion, which may sometimes include additional disjoint triangles, than to refine the search to the subset of triangles which have an actual intersection. The condition for a positive overlap is not as simple as for instance the requirement that vertices of  $T_{lm}$  fall internally to  $f(t_{ij})$  or viceversa (figure B.1 provides many counter examples).

The computation of overlap areas is more involved, and performed at step 3. For that, we make due use of barycentric coordinates [6, 8] to reference the position of a point within a given triangle  $T$ . In barycentric coordinates, the position of any point  $B$  in the plane is determined by a triple of real numbers  $(b_1, b_2, b_3)$ . This system has several properties that come to advantage for topological tests. With proper normalization,  $B$  can be said to be internal to  $T$  if all the three numbers  $b$  are positive;  $B$  falls on a side of  $T$  if one of the three  $b$  is null, and coincides with a vertex of  $T$  if two  $b$  are simultaneously null. Intersection points between two segments (in our case, sides of  $T$  and of  $f(t)$ ) are easily computed from the barycentric coordinates of their extremes (equation A.4). Since a segment and the side of a triangle intersect only if the relevant barycentric coordinate of the extremes have opposite signs, inspection of the signs can also be used as a flag to avoid unnecessary computation of non existing crossings. Details are in Appendix A.

Two possible algorithms for computing the areas of the triangle intersections needed for step 3 are described in detail in Appendix B. In our approach we make use of the one described

in B.4, which, albeit possibly slightly less efficient, is of much simpler implementation.

For step 4 the area  $A$  is elementary obtained from the vertex coordinates, computing the outer product of two side vectors, whereas step 5 is mechanic.

**4. Warping examples and evaluation.** We provide an example of the area resampling method using an 8 bit monochrome,  $512 \times 512$  pixels test image (`boat.512`, from [37]). The image coordinates are defined so that  $x_1 = y_1 = 0$  and  $x_{512} = y_{512} = 1$  ( $y$  increasing downwards). For the sake of illustration we take, as warping transformation,

$$(4.1) \quad \begin{pmatrix} X \\ Y \end{pmatrix} = f(x, y) = \begin{pmatrix} x + \frac{3 \sin(2\pi y)}{20} \\ y - \frac{3 \sin(\pi x)}{20} \end{pmatrix},$$

which induces no deformation on the sides of the unit square, and has Jacobian comprised between  $0.45 < J_f < 1.65$ . To quantify the numerical error in the preservation of intensity of the warped image, we compute the total intensity discrepancy

$$(4.2) \quad \delta = \frac{\sum_{i,j=1}^{N_1, M_1} I_1(i, j) - \sum_{l,m=1}^{N_2, M_2} I_2(l, m)}{\sum_{i,j=1}^{N_1, M_1} I_1(i, j)}.$$

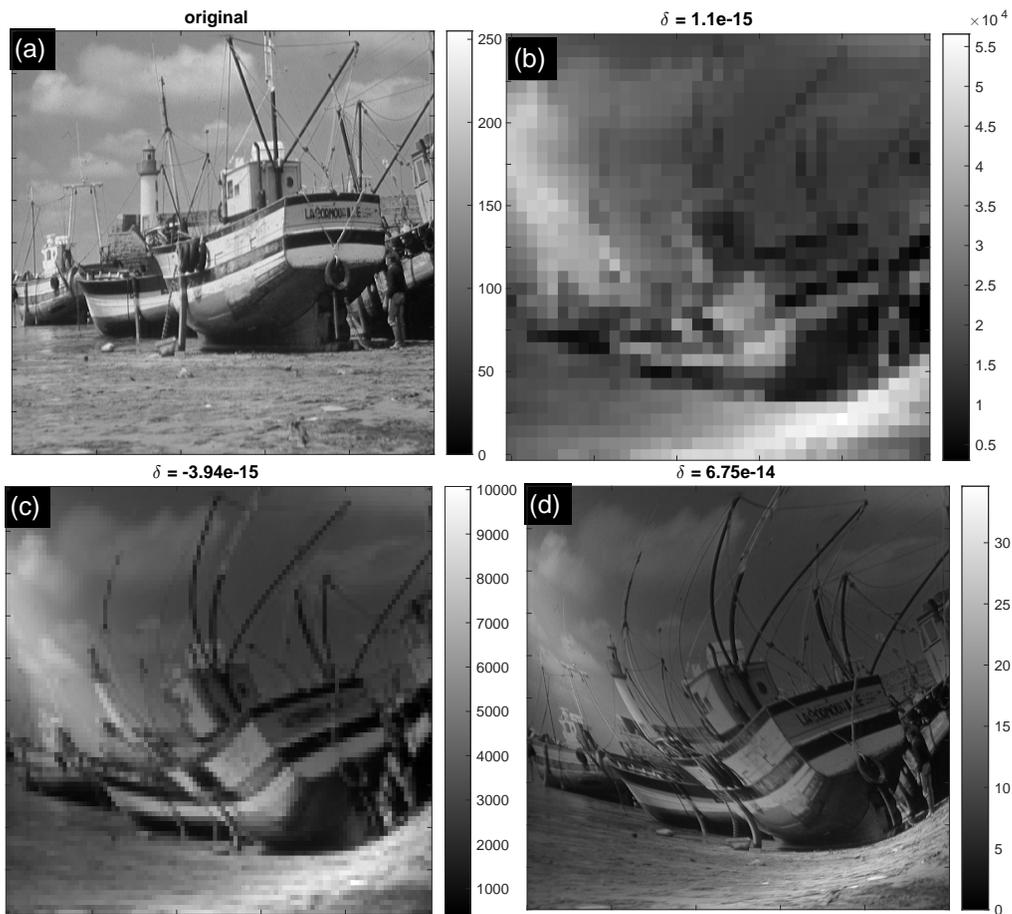
The result of warping is shown in Figure 4.1. The warp has been computed at various completely arbitrary resolutions, under and oversampling the image, to show the generality of the procedure. The calculation is performed in double precision floating point. The resulting  $\delta$ , reported over each warped image, are barely over numerical precision.

**5. Extensions.** Alternative forms of the matrix element  $B_{lm,ij}$  can be devised, giving different weights to the deformed pixel overlaps. The form of (2.4) distributes the available intensity separately on each destination hemipixel. Its effect can be appreciated in figure (5.1) for high downsampling ratio and non-affine pixel deformations, for which  $\mathcal{A} \left[ f \left( t_{ij}^L \right) \right]$  is significantly different from  $\mathcal{A} \left[ f \left( t_{ij}^U \right) \right]$ . Two other choices are presented in the following.

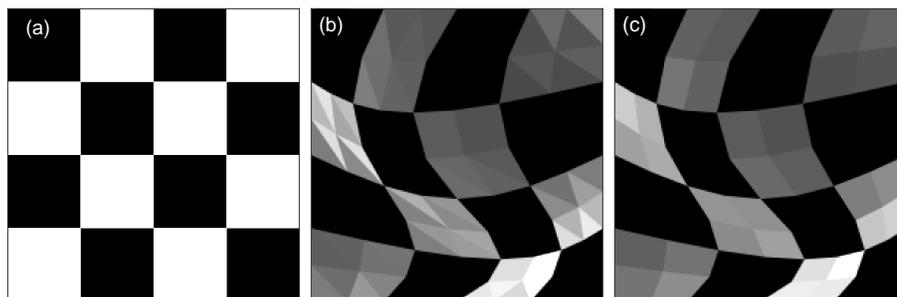
**5.1. Pixel uniformity.** Grouping together the two hemipixels, we can recast the total intensity of the original rectangular pixel onto the destination pixel, as actually prescribed by Eq. (2.3):

$$(5.1) \quad B_{lm,ij}^q = \frac{\mathcal{A} \left[ T_{lm}^U \cap f \left( t_{ij}^U \right) \right] + \mathcal{A} \left[ T_{lm}^L \cap f \left( t_{ij}^U \right) \right] + \mathcal{A} \left[ T_{lm}^U \cap f \left( t_{ij}^L \right) \right] + \mathcal{A} \left[ T_{lm}^L \cap f \left( t_{ij}^L \right) \right]}{\mathcal{A} \left[ f \left( t_{ij}^U \right) \right] + \mathcal{A} \left[ f \left( t_{ij}^L \right) \right]}.$$

In other words, the contributions of the transformed half pixels  $f \left( t_{ij}^U \right)$  and  $f \left( t_{ij}^L \right)$  are weighted with a cumulative factor, which is the average of the two denominators in (2.4). The effect of this choice is illustrated in Fig. 5.1. It may be argued that this weighting is more natural, as it preserves the original quadrangular pixel identity, rather than splitting it arbitrarily along one of its two diagonals.



**Figure 4.1.** Warped boat at different resolutions: a) original image,  $512 \times 512$ ; b)  $41 \times 36$ ; c)  $105 \times 87$ , d)  $1757 \times 1876$ . The values of  $\delta$  are reported over each warped image. Color bars at the side of each panel show how the intensity range of the image is inversely proportional to its resolution, so that the sum of the pixel values remains constant.



**Figure 5.1.** Effect of hemipixel vs. full pixel weighting on a highly oversampled transformation. a) Original  $8 \times 8$  pixels image, in which each square of the checker pattern occupies  $2 \times 2$  pixels. b) Warp to  $200 \times 200$  pixels using Eq. (2.4). The different brightness of the halves of highly stretched pixels is apparent. c) Warp to  $200 \times 200$  pixels using Eq. (5.1), which averages the contributions of the two original halves. The grayscale map of a) is different than that of b) and c) in order to stretch the visual contrast.

**5.2. Weighted area interpolation.** In the intensity-preserving resampling illustrated before, we assign a contribution of pixel  $I_1(i, j)$  to the destination pixel  $P_{lm}$  which is proportional to the inverse area of  $Q_{ij}$ , i.e. we take into account the local stretch or contraction caused by the warp. If instead we weight the contribution according to the destination area covered, normalizing over the area of the destination pixel, as in

$$(5.2) \quad B_{lm,ij}^a = \frac{\mathcal{A} [T_{lm}^U \cap f(t_{ij}^U)] + \mathcal{A} [T_{lm}^L \cap f(t_{ij}^U)] + \mathcal{A} [T_{lm}^U \cap f(t_{ij}^L)] + \mathcal{A} [T_{lm}^L \cap f(t_{ij}^L)]}{\mathcal{A} [T_{lm}^L] + \mathcal{A} [T_{lm}^U]},$$

we achieve a form of area averaging and resampling. For a Cartesian grid, obviously  $\mathcal{A} [T_{lm}^L] + \mathcal{A} [T_{lm}^U] = \Delta X \Delta Y$ . When undersampling, i.e. when several transformed pixels fall into a single destination pixel, their intensity values are averaged with a weight proportional to the area which they occupy on the destination; when oversampling, i.e. as a single deformed pixel covers more than a destination pixel, that destination pixel is assigned the same intensity of the source. This form can be thus seen as a value preserving warping, rather than an intensity preserving warping, and can be compared to other interpolation techniques in use in image processing.

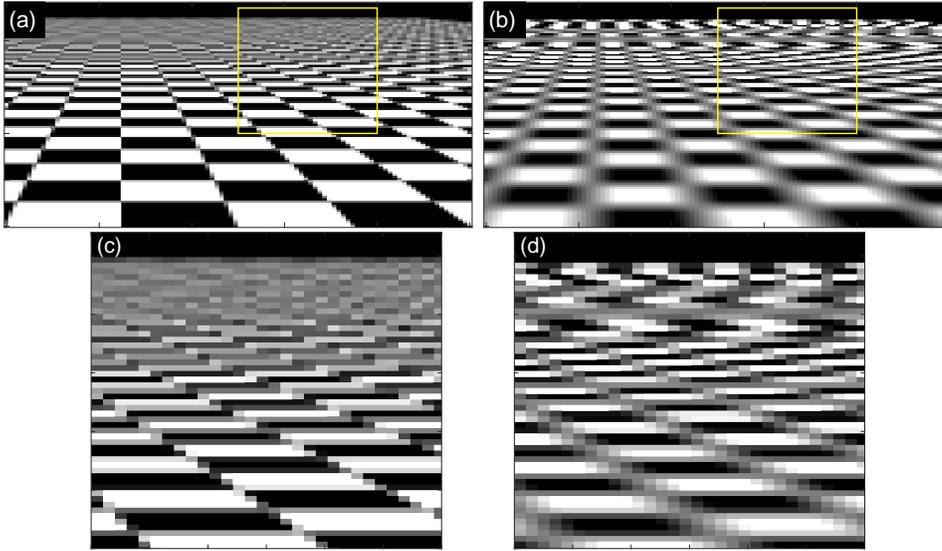
**5.2.1. Comparison with resampling interpolation.** Figure 5.2 provides a visual comparison of the merits of the weighted area interpolation based on Eq. (5.2) versus the commonly used bilinear interpolation, as a reference. Other more sophisticated interpolators, like higher order polynomial (e.g. bicubic, spline), Lanczos, or edge preserving (Akima), etc. could be compared as well, without affecting the main result. While many more interpolation methods are known in literature, the comparison with the simplest baseline algorithm is justified by the fact that the present area resampling recipe is only dependent on the geometry of the coordinate transform, not on the image data itself, nor on any assumed or a priori knowledge about the structure of the image. Other data-dependent interpolators (e.g. Takeda’s kernel regression [32], not to mention even more elaborate techniques based on deep learning) may produce more “realistic” results on the perceptual point of view, or even behave well as image denoisers (which implies a discrimination between an underlying image model and the superimposed corrupting noise). Here we merely report about the own merits of the weighted area resampler, without claiming that it is outperforming other image reconstruction techniques.

As an example of warping, we use the perspective transformation

$$(5.3) \quad \begin{pmatrix} X \\ Y \end{pmatrix} = g(x, y) = \begin{pmatrix} a + (x - a) \frac{d-b}{y-b} \\ c \left(1 + \frac{d-b}{y-b}\right) \end{pmatrix},$$

which describes the projection of an image on the  $xy$  plane on the vertical plane  $x = X$ ,  $y = d$ ,  $Y = z$ , from the viewpoint  $x = a$ ,  $y = b$ ,  $z = c$ .

To compare the two, we apply (2.5) and (5.2) using areas of triangles transformed from the  $(x, y)$  to the  $(X, Y)$  space, i.e. using a *direct* coordinate transform. For the interpolation instead, we exploit a more efficient, customary implementation which evaluates the image values on the regular  $(X, Y)$  destination grid by looking up and interpolating values on  $g^{-1}(X, Y)$ .



**Figure 5.2.** Comparison between area resampling (left) and unfiltered bilinear interpolation resampling (right). The original image is a checker pattern defined on  $128 \times 64$  pixels, which is projected with (5.3-5.4) onto a raster of  $100 \times 100$  rectangular pixels. The lower panels (c) and (d) show a zoom-in of the area enclosed within the yellow boxes. Differences are appreciated as absence of smoothing at the checker boundaries on the “near” side of the checker, and as reduction of aliasing artifacts at the “far” end.

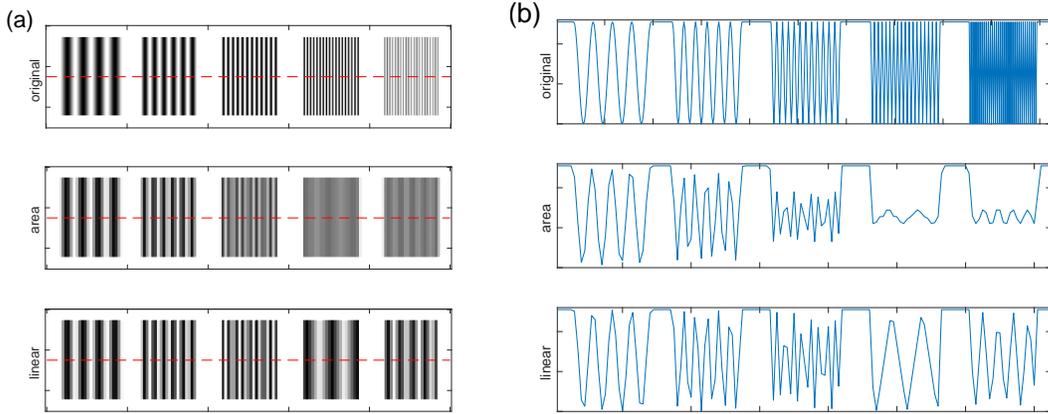
In other words, we perform an *inverse* pixel lookup. The transformation (5.3) has analytical inverse

$$(5.4) \quad \begin{pmatrix} x \\ y \end{pmatrix} = g^{-1}(X, Y) = \begin{pmatrix} \frac{2ac - cX - aY}{c - Y} \\ \frac{bY + ce - 2bc}{c - Y} \end{pmatrix},$$

with  $a = \frac{1}{4}$ ,  $b = -\frac{1}{10}$ ,  $c = \frac{1}{2}$ ,  $d = 0$ , and  $-5 < J_g < -\frac{5}{1331}$  for  $0 < y < 1$ . To stress the essential differences between the two methods, no dealiasing filter prior to interpolation is applied.

Area resampling somehow averages many original image values falling on the destination pixels, whereas the traditional resampling only picks up one value or averages few neighbors, sampling those which fall close to an interpolation point. Area resampling provides thus a smoother result than interpolation when undersampling images, and can be less prone to aliasing, since in that case it automatically behaves as an adaptive box filter. Interpolation is affected by aliasing, which is usually cured by low pass prefiltering; however, for a general warp transformation this filtering has to be local [17], complicating matters. An example of aliasing reduction is shown in Figure 5.3, where a global rescaling ratio and a simple periodic pattern evidence the different amplitude of the aliased spatial component. The amount of alias suppression, though, is entirely dependent on the particular local downsampling ratio, and on the original image content.

On the other hand, when oversampling, the weighted area transform produces images which are sharper and more faithful to pixel edges, since the destination value picked in that case represents well the original pixel value, rather than being an interpolation between nearest neighbors.



**Figure 5.3.** (a) A  $512 \times 100$ , 8 bit image containing sinusoidal test bars, with periods of 20, 12, 6, 4 and 2 pixels, (top) downsampled to  $145 \times 80$  using area resampling (center) and linear interpolation (the pixel aspect ratio is varied so that all images have the same coordinate span). (b) intensity profiles of the three images along the center line, red dashed lines in (a). For these bar spatial frequencies and resampling ratio, the amplitude of the aliased spatial frequency of the rightmost patterns is reduced using area downsampling.

**6. Example: source photometry.** Preservation of the intensity is an essential property when photometry is performed on deformable images. We show in figure 6.1 an illustrative example with synthetic data. To make our point, we consider the warp of a high resolution image to a lower resolution, and different ways to estimate, in the transformed image, the original intensity of each source. An original  $400 \times 400$  pixels image is created, simulating well separated, randomly placed sources with a gaussian peak profile with  $\sigma = 0.01L$ , where  $L = 1$  is the size of the square image. Each source  $k$ , centered at  $(x_k^s, y_k^s) \in ]0, 1[ \otimes ]0, 1[$ , contributes to the pixel  $I_1(i, j)$  with intensity

$$(6.1) \quad I_1^k(i, j) = \frac{1}{4} \left[ \operatorname{erf} \left( \frac{x_i - x_k^s + \Delta x}{\sigma} \right) - \operatorname{erf} \left( \frac{x_i - x_k^s}{\sigma} \right) \right] \left[ \operatorname{erf} \left( \frac{y_j - y_k^s + \Delta y}{\sigma} \right) - \operatorname{erf} \left( \frac{y_j - y_k^s}{\sigma} \right) \right].$$

With this integral definition, the total contribution of each source is normalized to the value  $s_k = 1$ . The image is then warped and downsampled to  $50 \times 50$  using the transformation

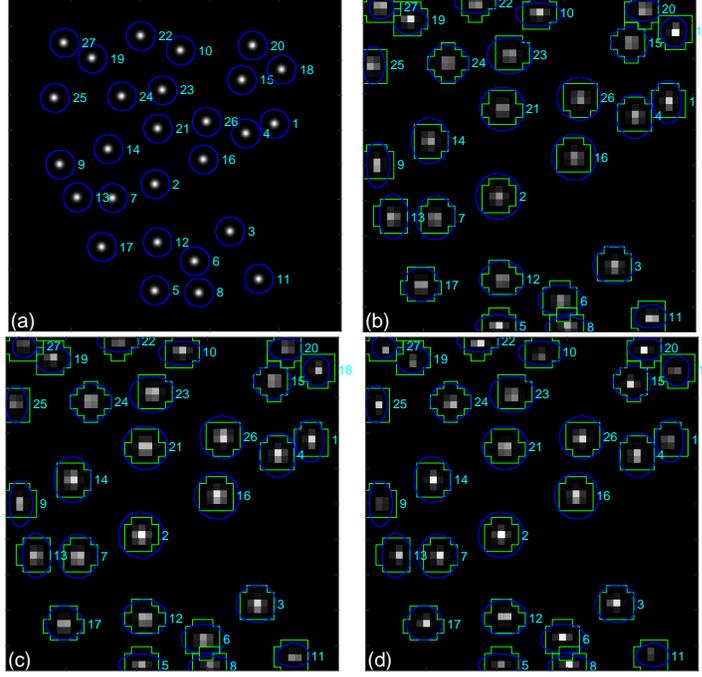
$$(6.2) \quad \begin{pmatrix} X \\ Y \end{pmatrix} = f(x, y) = \begin{pmatrix} \frac{1 - \cos(\pi x)}{2} \\ \frac{1 - \cos(\pi y)}{2} \end{pmatrix}$$

which expands the original image away from its center, compressing it at the edges. Differently than (4.1), this mapping has a closed inverse form

$$(6.3) \quad \begin{pmatrix} x \\ y \end{pmatrix} = f^{-1}(X, Y) = \begin{pmatrix} \frac{1}{2} - \frac{\sin^{-1}(1-2X)}{\pi} \\ \frac{1}{2} - \frac{\sin^{-1}(1-2Y)}{\pi} \end{pmatrix}.$$

The Jacobian of the direct transformation is  $J_f = \frac{1}{4}\pi^2 \sin(\pi x) \sin(\pi y)$ , whereas that of its inverse is

$$J_f^{-1} = \frac{1}{\pi^2 \sqrt{X(X-1)Y(Y-1)}}.$$



**Figure 6.1.** Photometry of synthetic sources in a warped image: a) original image  $I_1$ ,  $400 \times 400$ . The intensity  $\tilde{s}_k^{orig}$  of each dot is evaluated summing up all pixel values in neighborhoods of radius  $4\sigma$ , whose contour is plotted in blue. b) Image  $I_2$ , warped and downsampled to  $50 \times 50$  using Eq. (5.1). New pixel neighborhoods of radius  $4\sigma$  around the displaced source centers are plotted in green, along with the deformed contours of the original neighborhoods, in blue. c) Image  $I_2^v$ , warped to  $50 \times 50$  using the weight of Eq. (5.2). d)  $I_2^{interp}$ , warped resampling to  $50 \times 50$  of  $I_1$ , using bilinear interpolation.

While this mapping is not exceedingly representative of the transformations used in practice to correct imaging defects (which are often modeled by polynomial functions), the existence of an analytical inversion formula instead of an approximation to it, allows a fair comparison with traditional resampling image interpolations, which are easily performed using the inverse map.

To estimate a posteriori the contribution of the sources, all intensity values in a pixel neighborhood of radius  $4\sigma$  are summed, and the result is compared to the nominal unit intensity of the peak. To simulate possible pitfalls of the process, when applied to real images, we add some real world methodological errors. In figure 6.1a, we compute the pixel sums even when some of the neighborhoods receive overlapped contributions from more than a source, or sources fall near the margins of the image and contribute incompletely to the total. In figure 6.1b, we estimate the intensity of each source by summing pixel values within a circular pixel neighborhood centered on the transformed source position, rather than transforming the shape of the initial neighborhood. In figure 6.1c and d, we interpolate the image, compute the neighborhood sums, compensating for the area changes by either multiplying them by  $J_f^{-1}(x_k^s, y_k^s)$  computed merely at the source center, or by multiplying the local intensity by the local value of the Jacobian. In total we compare six different estimators of the intensity of the source, assuming  $(x_k^s, y_k^s)$  known a priori:

1. On the original image, we compute  $\tilde{s}_k^{\text{orig}} = \sum_{4\sigma} I_1(i, j)$  over the pixels within a distance  $4\sigma$  from the center  $(x_k^s, y_k^s)$  (within the blue contours in figure 6.1a)
2. On the area warped image, we compute  $\tilde{s}_k^{\text{area warp}} = \sum_{4\sigma} I_2(l, m)$ , summing this time the intensities of the destination pixels which fall within  $4\sigma$  from the *transformed* center  $(X_k^s, Y_k^s)$  (within the green contours in figure 6.1b). These summation neighborhoods may differ from the transformed original ones (compare green and blue lines); by using this evaluation, we want to assess the error involved, which is presumed small given the rapid decay and good separation of the peaks.
3. Using instead the area resampled image  $I_2^w$  of figure 6.1c, obtained using the weighted equalization of (5.2), we compute

$$\tilde{s}_k^{\text{area resampled}} = \frac{\Delta X \Delta Y}{\Delta y \Delta y} \sum_{4\sigma} J_f^{-1}(l, m) \cdot I_2^w(l, m)$$

Here  $J_f(l, m)$  is the value of the Jacobian evaluated at the center of each pixel of the destination image.

4. Using the area resampled image of figure 6.1c, we compute

$$\tilde{s}_k^{\text{area resampled/center}} = J_f^{-1}(X_k^s, Y_k^s) \frac{\Delta X \Delta Y}{\Delta y \Delta y} \sum_{4\sigma} I_2^w(l, m)$$

using the value of the Jacobian evaluated at the transformed position of the center of the source alone.

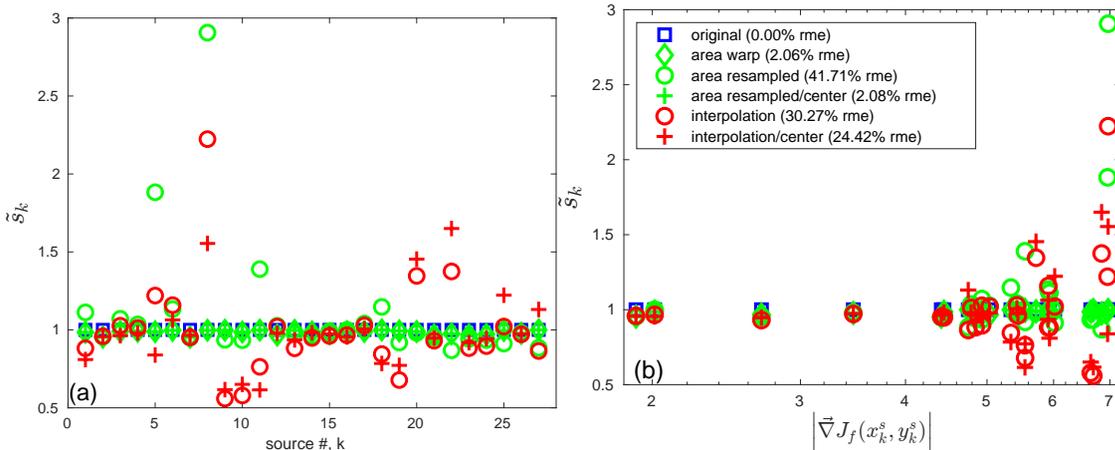
5. Using instead the warped and interpolated image  $I_2^{\text{interp}}(l, m)$  of figure 6.1d, we compute

$$\tilde{s}_k^{\text{interpolation}} = \frac{\Delta X \Delta Y}{\Delta y \Delta y} \sum_{4\sigma} J_f^{-1}(l, m) \cdot I_2^{\text{interp}}(l, m)$$

6. Using the interpolated image  $I_2^{\text{interp}}(l, m)$ , we compute

$$\tilde{s}_k^{\text{interpolation/center}} = J_f^{-1}(X_k^s, Y_k^s) \frac{\Delta X \Delta Y}{\Delta y \Delta y} \sum_{4\sigma} I_2^{\text{interp}}(l, m)$$

Figure 6.2 summarizes the results of the various estimators. Discrepancies of the different  $\tilde{s}_k$  with respect to the nominal value can be ascribed to sums over circular neighborhoods which deviate from the actual warped profile of the peak, as well as to overlapping peak tails (which are minimal in our example), but are notably due to the use of Jacobian factors evaluated pointwise. Figure 6.2b shows in particular how errors significantly increase for sources at the periphery of the domain, where  $J$  varies more rapidly, as quantified by  $|\vec{\nabla} J_f|$ , which is easily computed analitically. We have chosen here deliberately an extreme case, in which  $J$  is not constant across the domain, and a downsampling of a factor 8, in order to exacerbate the loss of information due to point-based resampling. It is clear that naive estimators based on the interpolated image can lead to misestimation, with root mean square errors  $\varepsilon = \sqrt{\langle (\tilde{s}_k - s_k)^2 \rangle}$  of the order of a quarter of the nominal peak intensity itself. In contrast, even a naive estimation, using mere undeformed circular neighborhoods, on the area warped image ( $\tilde{s}_k^{\text{area warp}}$  of point 2) produces much more accurate results.



**Figure 6.2.** Source intensities  $\tilde{s}_k$  estimated according to different methods. The root mean square error of each method is reported in the caption. In (a), the datapoints are sorted according to the source number  $k$ , for identification on Figure 6.1; in (b) the inferred source strengths are plotted versus the value of the gradient of the Jacobian at the center of the source. The plot shows how, while area warp provides a reasonable photometric estimation everywhere in the warped domain, methods 5 and 6 succeed only where  $J_f$  varies slowly. The relative better performance of method 6 over 5 is probably due to the symmetry of the original source, and the better representativeness of the central value of  $J_f$ .

**7. Conclusions and future outlook.** We have described a rigorous method for warping images, which by construction preserves the cumulative brightness of their features, and thus is suitable for photometric measurements on the transformed image. In doing so we dwelt on the computational geometry problem of finding the area of the intersection of two triangles. Despite its geometrical simplicity, we were not aware of a viable and robust algorithm for it available openly, and we provide one. We showed that a slight variation of the procedure can instead preserve the local values of intensity, and thus be directly compared with traditional implementations of warping, based on resampling the deformed image at gridpoints. Our method remains an area resampling method also in this application, and thus has implicit different filtering properties, not requiring for instance a preliminary antialiasing filter, and preserving sharper edges in case of severe oversampling.

The method has been showcased on monochrome images, but its extension to multichannel (e.g. color) images would be trivial, and in its simplest conception would amount to the computation of equation (2.5) independently for each channel.

The algorithms proposed are computationally more demanding than plain resampling ones, and in this work we have not pursued their highest possible efficiency. Future work could concentrate on developing faster implementations of them. Being prone to parallelization (see Appendix C), a GPU implementation of the algorithm can be envisioned. Once proved viable, the implementation of the present method in different programming languages and its inclusion into popular software packages like those mentioned in the introduction can be advocated for.

As a further development, the procedure could be adapted to non-rectangular source pixels, which can be in any case be decomposed into constituent triangles. The core of the method would remain the same, the only differences would be in devising an indexing for the trian-

gulation of the shaped pixels. Two use cases come in mind: for one, real physical imagers, notably CMOS sensor chips, have by architectural necessity photosensitive areas which cover partially the rectangular pixel cell [39]. Our procedure would provide an area-consistent way of resampling their measurements on differently gridded or deformed coordinates. As for another application, our method could be used where by design the image pixels are not arranged over a Cartesian grid at all, like for instance in hexagonal image processing [21].

**Appendix A. Barycentric coordinates.** Given the triangle  $T = \{\vec{u}, \vec{v}, \vec{w}\}$ , defined by the plane coordinates of its three vertices, and a point  $\vec{x}$ , we define:

$$(A.1) \quad \begin{aligned} 2A &= 2[(v_y - w_y)(u_x - w_x) - (v_x - w_x)(u_y - w_y)] \\ s &= (\vec{x} - \vec{w}) \cdot (v_y - w_y, v_x - w_x) \\ t &= (\vec{x} - \vec{w}) \cdot (u_y - w_y, u_x - w_x), \end{aligned}$$

where  $A$  is the signed area of the triangle  $T$ , positive or negative depending on the clockwise order of the vertices. Our convention is to define the (unscaled) barycentric coordinates of point  $\vec{x}$  as

$$(A.2) \quad \vec{b} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = \text{sign}(A) \begin{pmatrix} s \\ t \\ 2A - s - t \end{pmatrix}.$$

For any point  $\vec{x}$  internal to  $T$ ,  $0 \leq b_k \leq 2A$ , for any  $k = \{1, 2, 3\}$ .

If one component  $b_k = 0$ , the point  $\vec{x}$  lies on a side of  $T$ . If two components  $b_k$  are simultaneously zero, the point  $\vec{x}$  is simultaneously on two sides of  $T$ , i.e. coincides with a vertex of the reference triangle. All three components of  $\vec{b}$  can be null only for the degenerate case of a triangle with three coinciding vertices.

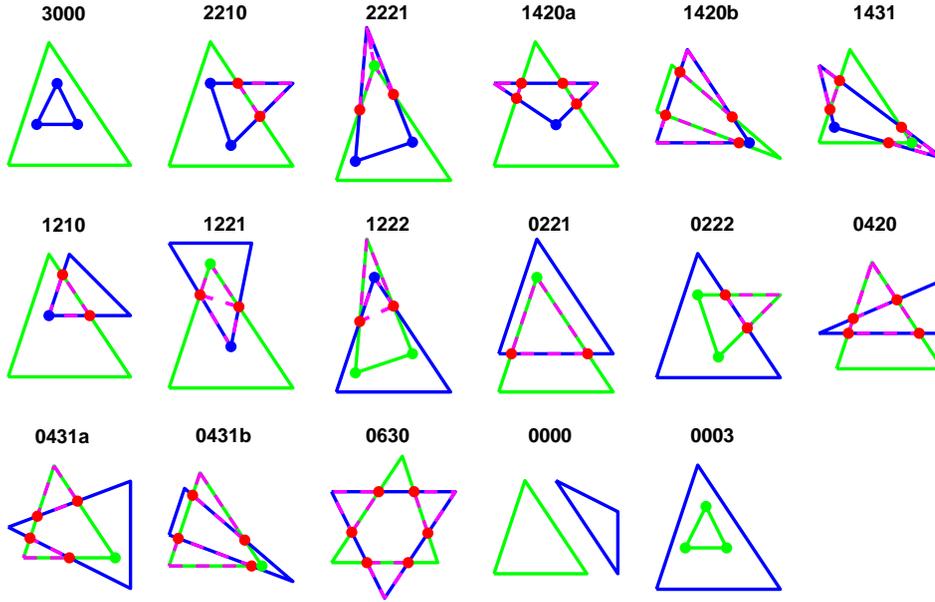
The inversion relation giving  $\vec{x}$  from its barycentric coordinates is

$$(A.3) \quad \vec{x} = \frac{b_1 \vec{u} + b_2 \vec{v} + b_3 \vec{w}}{2|A|}.$$

A segment  $\{\vec{x}_1, \vec{x}_2\}$ , whose extremes have barycentric coordinates  $\vec{b}^1$  and  $\vec{b}^2$  with respect to  $T$ , intersects the  $k$ -th side of the triangle  $T$  if  $b_k^1$  and  $b_k^2$  have opposite signs. The barycentric coordinates  $\vec{b}^c$  of the intersection point are then

$$(A.4) \quad \vec{b}^c = \frac{b_k^2 \vec{b}^1 + b_k^1 \vec{b}^2}{b_k^2 - b_k^1}.$$

**Appendix B. Area of the intersection of two triangles.** The problem of intersecting triangles in two and three dimensions has received due attention in computer graphics, being fundamental in a number of applications which involve triangulation of domains, like collision detection or intersection of triangulated surfaces [7], and indeed literature on it is available (e.g. [40, 15, 25, 34, 29]). However, in these works at most the conditions for the detection of planarity and intersection of two triangles are given, but not an explicit algorithm computing the overlap area of planar triangles, which we need here. Its derivation is discussed in this appendix.



**Figure B.1.** The seventeen topologically different ways for two triangles  $T_1$  (blue) and  $T_2$  (green) to intersect. Blue and green dots indicate respectively vertices of  $T_1$  internal to  $T_2$  and viceversa; red dots intersections between sides. The numeric label above each couple reflects the classification explained in the text. The area of the intersection polygon can always be computed as a sum or a difference of constituent triangetlets (up to four), which are highlighted by pink dashed lines.

**B.1. Enumeration of possible cases.** The possible ways in which two triangles  $T_1$  and  $T_2$  can overlap and intersect can be classified according to topological properties. Being triangles always convex, their intersection is always a convex polygon. A side of  $T_2$  can intersect zero, one or two sides of  $T_1$ . A side of  $T_2$  with one vertex internal and one external to  $T_1$  implies a single intersection with one side of  $T_1$ , while both vertices internal or external could both grant either zero or two intersections. The different overlap cases can be labeled according to: the number  $v_1$  of vertices of  $T_1$  falling inside  $T_2$  (which can be 0, 1, 2, or 3); the number  $i_s$  of intersections between sides of the two triangles (0, 2, 4 or 6); the number  $s_2$  of sides of  $T_2$  intersected, and, to remove ambiguities, the number  $v_2$  of vertices of  $T_2$  falling into  $T_1$ . All the possible cases are depicted in Fig. B.1. The label above each subfigure is derived from these numbers,  $v_1 i_s s_2 v_2$ . A further distinction is necessary for some of the cases with  $i_s = 4$ : two topologically different arrangements are possible with the same classification numbers, hence **1420a** and **1420b**, **0431a** and **0431b**.

As for areas, clearly  $\mathcal{A}[T_1 \cap T_2] = \mathcal{A}[T_2 \cap T_1]$ , and commutativity would reduce the number of topologically different cases to 11 (five of the seventeen cases are topologically invariant for the exchange of the two members, like e.g. **1431**, **1221**; the other twelve have each one their dual, like **2210** and **0222**, counted only once). In our computation, though, the two member triangles have different roles, and we must in principle treat the all cases as distinct. In any case, the area of the intersection polygon can always be computed as a sum or difference of at most four smaller triangetlets, formed either by the vertices or by the intersection points of the member sides.

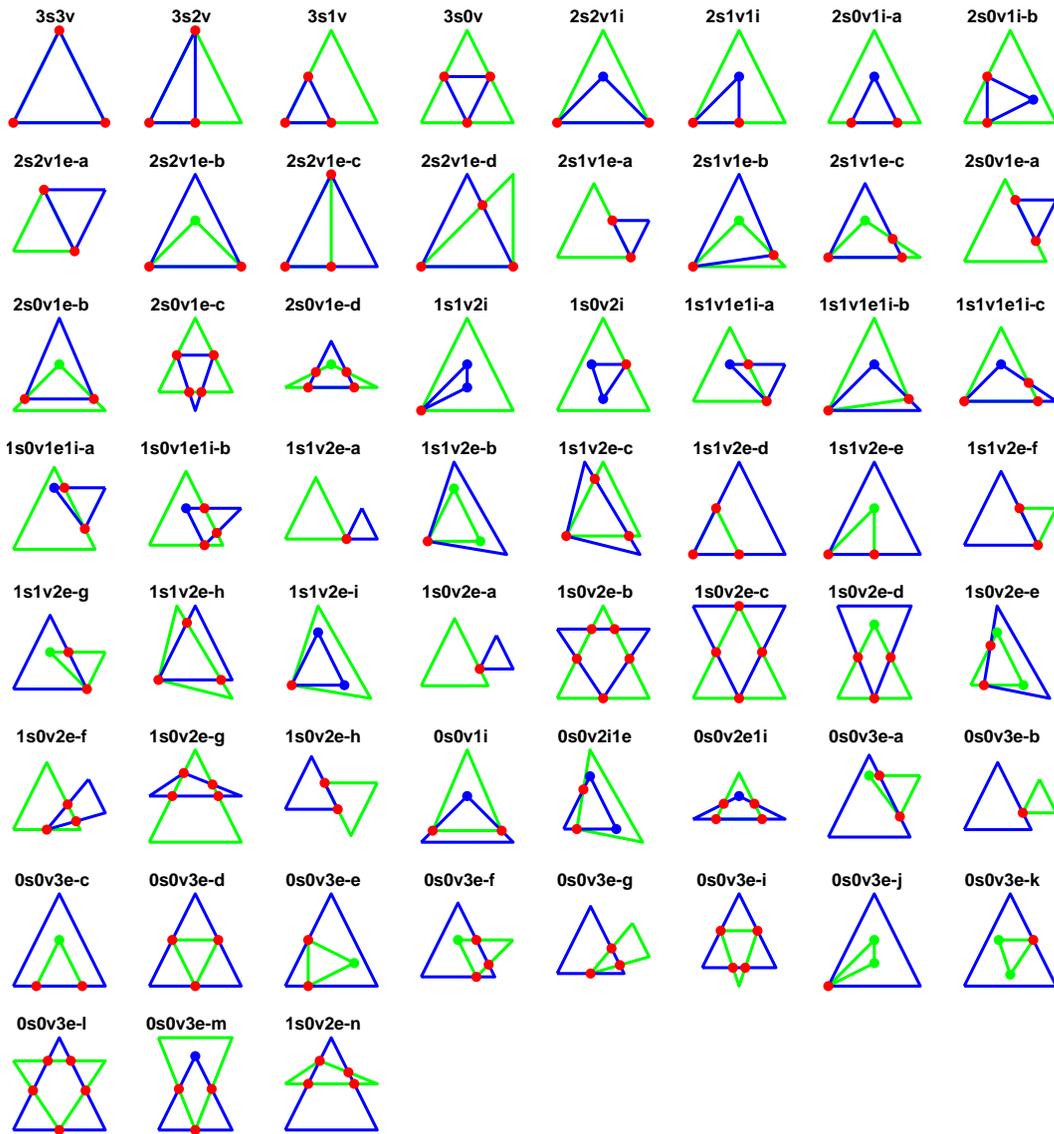
We note that an early analysis of the problem was given in the report [30], though without an explicit computation algorithm and without considering degenerate cases.

**B.2. Degeneracy and numerical precision.** Problems arise for triangles which have some colinear side, or simply vertices of one triangle falling on the sides of the other. Figure B.2 displays 59 topologically different configurations, and is possibly not even exhaustive of them. Such cases have to be treated with care in the numerical computation, because their identification requires the simultaneous satisfaction of more than a single equality condition. For instance, if one vertex of  $T_1$  falls onto a side of  $T_2$ , three conditions are to be true: the vertex must belong to the side of the second triangle, and two sides of the first triangle must intersect that side, both exactly in that point. For our purposes, one of the triangles will be a half pixel in the destination image, and the other a warped half pixel of the original image. For a generic functional mapping between  $(x, y)$  and  $(X, Y)$ , expressed by algebraic or transcendental functions, such cases may be extremely rare; however, colinear points will be very frequent for transformations like grid sub or oversampling by integer factors, or rotation by notable angles, which are indeed among the most typical test cases. A reliable algorithm for warping must be capable to treat them adequately as well.

At numerical precision, due to the propagation of truncation errors, exact equality conditions may often be violated; moreover, they may be violated in a way which is topologically inconsistent. For example, it may result numerically that one vertex of  $T_1$  falls on a side of  $T_2$ , whereas the sides of  $T_1$  originating from that vertex may not appear to intersect  $T_2$ , or, to intersect it in numerically different points than the vertex in question. Additionally, extreme image warpings around singular points of the coordinate transformation can produce pathological triangles with nearly-colinear vertices. The numerical identification of intersections between sides can also produce bogus results for them.

A way of coping with nearly degenerate cases would be to evaluate the algebraic conditions which define whether a point is interior or exterior to a triangle, or whether two segments intersect, within preassigned tolerances, larger than typical truncation errors [19, 8]. The problem in that is that the resolution of the ambiguity usually requires a compatible set of decisions for more than a single test. If for instance one of the barycentric coordinates of a point is found to be nearly zero, implying that the point lies on the side of the test triangle, it is not possible to consider systematically that point as interior or exterior to the triangle, by itself. The decision depends on which intersections among sides should be counted or not, in order to reconduce the limit case to one of the 17 basic ones of figure B.1.

**B.3. Topological approach.** If it was not for the possible degeneration, we could compute the area of  $T_1 \cap T_2$  in each of the 17 cases mentioned in section B.1 identifying in each configuration the composing triangetlets (up to four) which are to be considered. Such triangetlets are highlighted by pink dashed lines in Fig. B.1. In this approach, the case is first identified counting the number of points of  $T_1$  internal to  $T_2$ , and then determining the number and the coordinates of the intersections only among the sides required. In some of the cases, the knowledge of the internal points makes the computation of some intersections among sides unnecessary, saving operations. For instance, in case **2210** it is known a priori that one side of  $T_1$  is all contained in  $T_2$ , and doesn't intersect any of its sides. Then, with conditional code which treats each case differently, the relevant triangetlets are singled out and their areas



**Figure B.2.** 59 topologically different degenerate cases of triangles with one or more vertices of  $T_1$  (blue) falling on sides of  $T_2$  (green), or viceversa. A possible classification, hinted in the labels, may count the number of vertices of  $T_1$  falling on sides of  $T_2$  (suffix  $s$ ), the number of common vertices (suffix  $v$ ), the number of other vertices of  $T_1$  internal or external to  $T_2$  (suffixes  $i$  or  $e$ ); but these indicators alone are not exhaustive (hence the lowercase letters appended to the label).

added. This may be, for some of the cases, computationally more economical than with the approach described in the next section. A sample Matlab implementation of this method is given in the file `areaTriangleIntersection.m` included in the Supplementary Material (directory `Triangles/topologicalIntersection/`). However, the procedure is incomplete if degenerate cases are not treated, identified and cast into one of the 17 basic patterns. While this is possible, it requires detailed case-by-case code, branching through all possibilities (see

the function `intersectTriangles2.m` in the same directory). That would lead to a code complicated and difficult to be maintained, outweighing the minimal performance gain which could result from it.

**B.4. Brute force approach.** To avoid the classification of all potentially degenerate triangle overlap cases, we adopt a simpler approach. We start noting that the vertices of the polygon  $T_1 \cap T_2$  are points which always belong to at least one of three sets: the vertices of  $T_1$  which are internal to  $T_2$ , the vertices of  $T_2$  internal to  $T_1$ , and the intersections between sides of  $T_1$  and  $T_2$ . First, a list of such points (at most  $12 = 3 + 3 + 6$ ) is compiled using barycentric coordinates for the computations. In the degenerate cases, some of these points can appear in more than one of the subsets: for example a point of  $T_1$  on a side of  $T_2$  can appear both as an internal point and as the intersection of two different sides of  $T_1$  with  $T_2$ . The list is therefore pruned, eliminating duplicate points which coincide within a given numerical tolerance. The area of the polygon is then computed from the pruned list of  $N$  vertices  $\{\vec{x}_k\} = \{(x_k, y_k)\}$  according to algorithm B.1, which is robust to nearly coincident or colinear vertices, and economical in terms of operations.

---

**Algorithm B.1** area of the intersection polygon

---

if  $N < 3$ ,  $\mathcal{A}[\{\}] = 0$

if  $N = 3$ ,

$$\mathcal{A}[\{\vec{x}_1, \vec{x}_2, \vec{x}_3\}] = \frac{|(x_3 - x_1)(y_2 - y_1) - (x_2 - x_1)(y_3 - y_1)|}{2}$$

if  $N > 3$ :

1. the center point  $\vec{x}_M = \frac{1}{N} \sum_{k=1}^N \vec{x}_k$  is computed
2. ray angles from the center are computed,  $\varphi_k = \tan^{-1} \frac{y_M - y_k}{x_M - x_k}$
3. the set  $\{\vec{x}_k\}$  is sorted in order of increasing  $\varphi_k$
4.  $\vec{x}_1$  is taken as a pivot, and for  $3 \leq k \leq N$  the  $N - 2$  triangle areas

$$\mathcal{A}[\{\vec{x}_1, \vec{x}_{k-1}, \vec{x}_k\}] = \frac{(x_k - x_1)(y_{k-1} - y_1) - (x_{k-1} - x_1)(y_k - y_1)}{2}$$

are computed and summed.

---

Sorting the vertices in cyclic order is required for  $N > 3$  (steps 1–3 of algorithm B.1), since the pruned list of vertices is not guaranteed to be ordered, by construction. This adds a computational cost of two divisions by  $N$  (step 1),  $N$  evaluations of `atan2()` (step 2) and a sort operation of a list of four to six floating numbers (step 3). Conditional code to list the internal points and the intersections in a proper order, on the other hand, would be more convoluted. As an aside, step 4 involves only  $2(N - 2)$  multiplications, whereas the standard shoelace algorithm for computing the area of the polygon would require  $2N$ . The cyclic order of the vertices, and the convexity of the polygon itself, guarantee that  $(x_k - x_1)(y_{k-1} - y_1) - (x_{k-1} - x_1)(y_k - y_1) > 0$  for all  $k$ . There would be other algorithmic possibilities to achieve cyclic sorting, avoiding the evaluation of `atan2()`; (compare for instance the function `polygonArea2.m` given in Supplementary Material, which uses it, with `polygonArea3.m`, which does not, both in directory `Triangles/`); however, as for Matlab is

concerned, the first option is faster).

The advantage of this procedure is that a single threshold identification criterion is applied to the list of points obtained, without the need of pondering the compatibility between resolutions of internality and crossing in the limit cases. Conversely, the effect of including or excluding some nearly coincident polygon vertices, dependent on the threshold value chosen, amounts only to adding or not some nearly null area contributions.

This algorithm is implemented in the file `areaTriangleIntersection2.m` included in the Supplementary Material (directory `Triangles/`), and has been used to process the images shown in this paper.

**Appendix C. Computational performance.** The algorithms described in this paper have been implemented in Matlab, and the code is provided as Supplementary Material of this paper. Coding has paid some amount of attention to efficiency and good programming practice, but ultimate performance has not been sought for itself.

Some CPU times of the algorithm presented, vs. the much faster resampling interpolation of the inverse map, are presented in Table C.1. Timings were obtained on a 12 core Intel Xeon® W-2135 CPU with 3.70GHz clock, using Matlab 2020a. Tests runs on square images defined on unit square coordinates, similar to that reported in section 5.2.1 were executed at different resolutions, and the average of 10 area warp iterations and 1000 interpolation iterations is recorded. The sin mapping of Eq. (6.2) was chosen, having an analytical inverse and being bijective on the whole unit square.

**Table C.1**

*Timings for area warping vs. resampling.*

$I_1$ resolution	$I_2$ resolution	area warp, ms	$\#B_{lm,ij}^a \neq 0$	bilinear resampling, ms
$64 \times 64$	$100 \times 100$	251.6	26244	0.313
$64 \times 64$	$1000 \times 1000$	4594	1127844	10.75
$512 \times 512$	$100 \times 100$	3540	372100	2.43
$512 \times 512$	$1000 \times 1000$	10748	2280100	13.43

These timings are to be taken only as somewhat indicative of a general trend. In our implementation of the area warping, the core routine for computing  $\mathcal{A}[T_{lm} \cap f(t_{ij})]$  using eq. (A.4) and Algorithm B.1 is compiled into a mex file, for higher efficiency, but all other parts of the code are executed in a loop by the Matlab interpreter. This includes the index referencing to  $t_{ij}$  and  $T_{lm}$ , as well as the computation of  $f(t_{ij})$ . On the other hand, the outer loop on  $i, j$  of algorithm 3.1 described in §3 can be easily parallelized, using Matlab’s `parfor` construct. In contrast, when we look at traditional resampling, we are comparing with a single call of `interp2`, which is certainly well optimized internally, and involves much less operations. It is thus not too constraining, at this stage, to observe that the area warp procedure is some three orders of magnitude slower than the usual technique. Moreover, the computational effort for the area warp is expected to be dependent on the amount of stretching and destination domain coverage caused by the particular mapping  $f$ , which affects the sparsity of the resulting matrix  $B_{lm,ij}^a$ , and even by the number of sides of each individual triangle intersection generated.

**Acknowledgments.** The author wishes to thank Eran Ofek for fruitful discussion on the content of this paper.

## REFERENCES

- [1] A. AMANATIADIS AND I. ANDREADIS, *A survey on evaluation methods for image interpolation*, Measurement Science and Technology, 20 (2009), p. 104015, <https://doi.org/10.1088/0957-0233/20/10/104015>.
- [2] M. F. BEG, M. I. MILLER, A. TROUVÉ, AND L. YOUNES, *Computing large deformation metric mappings via geodesic flows of diffeomorphisms*, International Journal of Computer Vision, 61 (2005), pp. 139–157, <https://doi.org/10.1023/B:VISI.0000043755.93987.aa>.
- [3] T. BEIER AND S. NEELY, *Feature-based image metamorphosis*, SIGGRAPH Comput. Graph., 26 (1992), pp. 35–42, <https://doi.org/10.1145/142920.134003>.
- [4] G. BRADSKI, *The OpenCV Library*, Dr. Dobb's Journal of Software Tools, (2000).
- [5] M.-C. CHIANG, *Imaging-consistent warping and super-resolution*, PhD thesis, Columbia University, 1998, <https://www.proquest.com/docview/304435854>.
- [6] H. S. M. COXETER, *Introduction to Geometry*, Wiley, 2nd ed. ed., 1969.
- [7] A. H. ELSHEIKH AND M. ELSHEIKH, *A reliable triangular mesh intersection algorithm and its application in geological modelling*, Engineering with Computers, 30 (2014), pp. 143–157, <https://doi.org/10.1007/s00366-012-0297-3>.
- [8] C. ERICSON, *Real-Time Collision Detection*, The Morgan Kaufmann series in interactive 3D technology, CRC Press, London, 2005, <https://learning.oreilly.com/library/view/real-time-collision-detection/9781558607323/>.
- [9] K. M. FANT, *A nonaliasing, real-time spatial transform technique*, IEEE Computer Graphics and Applications, 6 (1986), pp. 71–80, <https://doi.org/10.1109/mcg.1986.276613>.
- [10] E. FOGEL, O. SETTER, R. WEIN, G. ZUCKER, B. ZUKERMAN, AND D. HALPERIN, *2D regularized boolean set-operations*, in CGAL User and Reference Manual, CGAL Editorial Board, 5.3 ed., 2021, <https://doc.cgal.org/5.3/Manual/packages.html#PkgBooleanSetOperations2>.
- [11] A. S. FRUCHTER AND R. N. HOOK, *Drizzle: A method for the linear reconstruction of undersampled images*, Publications of the Astronomical Society of the Pacific, 114 (2002), pp. 144–152, <https://doi.org/10.1086/338393>.
- [12] P. GETREUER, *Linear Methods for Image Interpolation*, Image Processing On Line, 1 (2011), pp. 238–259, [https://doi.org/10.5201/ipol.2011.g\\_lmii](https://doi.org/10.5201/ipol.2011.g_lmii).
- [13] D. GHOSH AND N. KAABOUCH, *A survey on image mosaicing techniques*, Journal of Visual Communication and Image Representation, 34 (2016), pp. 1–11, <https://doi.org/10.1016/j.jvcir.2015.10.014>.
- [14] C. A. GLASBEY AND K. V. MARDIA, *A review of image-warping methods*, Journal of Applied Statistics, 25 (1998), pp. 155–171, <https://doi.org/10.1080/02664769823151>.
- [15] P. GUIGUE AND O. DEVILLERS, *Fast and robust triangle-triangle overlap test using orientation predicates*, Journal of Graphics Tools, 8 (2003), pp. 25–32, <https://doi.org/10.1080/10867651.2003.10487580>.
- [16] D. HAN, *Real-time digital image warping for display distortion correction*, in Image Analysis and Recognition, M. Kamel and A. Campilho, eds., Berlin, Heidelberg, 2005, Springer Berlin Heidelberg, pp. 1258–1265, [https://doi.org/10.1007/11559573\\_152](https://doi.org/10.1007/11559573_152).
- [17] P. S. HECKBERT, *Fundamentals of texture mapping and image warping*, master's thesis, Dept. of Electrical Engineering and Computer Science, University of California, Berkeley, 1989, <http://www.cs.cmu.edu/~ph/texfund/texfund.pdf>.
- [18] IMAGEMAGICK DEVELOPMENT TEAM, *ImageMagick*, 2021, <https://imagemagick.org> (accessed 2021-01-04). Version 7.0.10.
- [19] C. JULES, *Accurate point in triangle test*, 2014, <http://totologic.blogspot.com/2014/01/accurate-point-in-triangle-test.html>.
- [20] S. LEE, G. WOLBERG, AND S. Y. SHIN, *Polymorph: morphing among multiple images*, IEEE Computer Graphics and Applications, 18 (1998), pp. 58–71, <https://doi.org/10.1109/38.637304>.
- [21] L. MIDDLETON AND J. SIVASWAMY, *Hexagonal Image Processing*, Springer-Verlag, 2005, <https://doi.org/10.1007/1-84628-203-9>.

- [22] P. MILANFAR, ed., *Super-Resolution Imaging*, CRC Press, first ed., 2011, <https://doi.org/10.1201/9781439819319>.
- [23] J. MODERSITZKI, *Numerical methods for image registration*, Numerical mathematics and scientific computation, Oxford University Press, Oxford, 2004.
- [24] J. MODERSITZKI, *FAIR: Flexible Algorithms for Image Registration*, Society for Industrial and Applied Mathematics, 2009, <https://doi.org/10.1137/1.9780898718843.ch3>.
- [25] T. MÖLLER, *A fast triangle-triangle intersection test*, Journal of Graphics Tools, 2 (2004), pp. 25–30, <https://doi.org/10.1080/10867651.1997.10487472>, <http://web.stanford.edu/class/cs277/resources/papers/Moller1997b.pdf>.
- [26] J. A. PARKER, R. V. KENYON, AND D. E. TROXEL, *Comparison of interpolating methods for image resampling*, IEEE Transactions on Medical Imaging, 2 (1983), pp. 31–39, <https://doi.org/10.1109/TMI.1983.4307610>.
- [27] R. J. RADKE, *Computer Vision for Visual Effects*, Cambridge University Press, 2012, <https://doi.org/10.1017/CBO9781139019682>.
- [28] M. REINELT, *pamscale*, in *Netpbm User Manual*, 2020, <http://netpbm.sourceforge.net/doc/pamscale.html> (accessed 2021-10-03).
- [29] C. L. SABHARWAL AND J. L. LEOPOLD, *A triangle-triangle intersection algorithm*, Computer Science & Information Technology (CS & IT), (2015), <https://doi.org/10.5121/csit.2015.51003>, <https://www.airccj.org/CSCP/vol5/csit54203.pdf>.
- [30] M. L. SAMPOLI, *An automatic procedure to compute efficiently the intersection of two triangles*, Tech. Report 465, Università di Siena, Dipartimento di Scienze Matematiche ed Informatiche, 2004, [https://www.researchgate.net/publication/265809721\\_An\\_Automatic\\_Procedure\\_to\\_Compute\\_Efficiently\\_the\\_Intersection\\_of\\_Two\\_Triangles](https://www.researchgate.net/publication/265809721_An_Automatic_Procedure_to_Compute_Efficiently_the_Intersection_of_Two_Triangles) (accessed 2022-04-27).
- [31] J. SUMMERS, *ImageWorsener - Pixel Mixing*, 2012, <http://entropymine.com/imageworsener/pixelmixing/>.
- [32] H. TAKEDA, S. FARSIU, AND P. MILANFAR, *Kernel regression for image processing and reconstruction*, IEEE Transactions on Image Processing, 16 (2007), pp. 349–366, <https://doi.org/10.1109/TIP.2006.888330>.
- [33] A. THYSSEN, *ImageMagick v6 examples – distorting images*, 2012, [http://www.imagemagick.org/Usage/distorts/#area\\_resample](http://www.imagemagick.org/Usage/distorts/#area_resample) (accessed 2020-9-13).
- [34] O. TROPP, A. TAL, AND I. SHIMSHONI, *A fast triangle to triangle intersection test for collision detection*, Computer Animation and Virtual Worlds, 17 (2006), pp. 527–535, <https://doi.org/10.1002/cav.115>, <https://cs.nyu.edu/exact/pap/mesh/fast-tri-tri-intersect2006.pdf>.
- [35] S. VAN DER WALT, J. L. SCHÖNBERGER, J. NUNEZ-IGLESIAS, F. BOULOGNE, J. D. WARNER, N. YAGER, E. GOUILLART, T. YU, AND THE SCIKIT-IMAGE CONTRIBUTORS, *scikit-image: image processing in Python*, PeerJ, 2 (2014), p. e453, <https://doi.org/10.7717/peerj.453>.
- [36] L. VELHO, A. FRERY, AND J. GOMES, *Warping and Morphing*, Springer London, London, 2009, pp. 387–412, [https://doi.org/10.1007/978-1-84800-193-0\\_15](https://doi.org/10.1007/978-1-84800-193-0_15).
- [37] A. G. WEBER, *The USC-SIPI image database: Version 6*, 2018, [http://sipi.usc.edu/database/SIPI\\_Database.pdf](http://sipi.usc.edu/database/SIPI_Database.pdf) (accessed 2021-10-03).
- [38] G. WOLBERG, *Digital image warping*, IEEE Computer Society Press, Los Alamitos, CA, 1990.
- [39] O. YADID-PECHT AND R. ETIENNE-CUMMINGS, *CMOS Imagers*, 2004, <https://doi.org/10.1007/b117398>.
- [40] F. YAMAGUCHI, *A unified approach to interference problems using a triangle processor*, SIGGRAPH Comput. Graph., 19 (1985), pp. 141–149, <https://doi.org/10.1145/325165.325224>.
- [41] X. YANG, R. KWITT, M. STYNER, AND M. NIETHAMMER, *Quicksilver: Fast predictive image registration — a deep learning approach*, NeuroImage, 158 (2017), pp. 378–396, <https://doi.org/10.1016/j.neuroimage.2017.07.008>.
- [42] B. ZACKAY AND E. O. OFEK, *How to COAAD images. I. optimal source detection and photometry of point sources using ensembles of images*, The Astrophysical Journal, 836 (2017), p. 187, <https://doi.org/10.3847/1538-4357/836/2/187>.
- [43] S. ZHANG, P. X. LIU, M. ZHENG, AND W. SHI, *A diffeomorphic unsupervised method for deformable soft tissue image registration*, 120, p. 103708, <https://doi.org/10.1016/j.compbimed.2020.103708>.
- [44] B. ZITOVÁ AND J. FLUSSER, *Image registration methods: a survey*, Image and Vision Computing, 21 (2003), pp. 977–1000, [https://doi.org/10.1016/S0262-8856\(03\)00137-9](https://doi.org/10.1016/S0262-8856(03)00137-9).