

Saliency-driven Variational Retargeting for Historical Maps

Filippo Bergamasco^{1,2}[0000-0001-6668-1556], Arianna Traviglia¹[0000-0002-4508-1540], and Andrea Torsello¹[0000-0001-9189-4924]

¹ DAIS - Università Ca'Foscari, Via Torino 155, 30172 Venice, Italy

² Istituto Italiano di Tecnologia (IIT), Center for Cultural Heritage Technology (CCHT), Via della Libertà 12, 30175 Venice, Italy

{[filippo.bergamasco](mailto:filippo.bergamasco@unive.de), [traviglia](mailto:traviglia@unive.de), [andrea.torsello](mailto:andrea.torsello@unive.de)}@unive.de

Abstract. We study the problem of georeferencing artistic historical maps. Since they were primarily conceived as work of art more than an accurate cartographic tool, the common warping approaches implemented in Geographic Application Systems (GIS) usually lead to an overly-stretched image in which the actual pictorial content (like written text, compass roses, buildings, etc.) is un-naturally deformed. On the other hand, domain transformation of images driven by the perceived salient visual content is a well-known topic known as “image retargeting” which has been mostly limited to a change of scale of the image (ie. changing the width and height) rather than a more general control-points based warping.

In this work we propose a variational image retargeting approach in which the local transformations are estimated to accommodate a set of control points instead of image boundaries. The direction and severity of warping is modulated by a novel tensor-based saliency formulation considering both the visual content and the shape of the underlying features to transform. The optimization includes a flow projection step based on the isotonic regression to avoid singularities and flip overs of the resulting distortion map.

Keywords: Image retargeting; Warping; Historical maps

1 Introduction

Content-aware image warping and retargeting has drawn significant attention so that many studies have been proposed in recent years.

Broadly speaking, the main target application in existing literature has almost exclusively been the change of scale or aspect ratio of images, producing “retargeted” outputs that kept in a tight rectangular image the unmodified perceptually salient features of the source image, while eliminating or distorting the less salient portions. Common to these applications is a constrained rectangular boundary resulting in induced deformations that are mostly axis-aligned. Emblematic of this class of approaches is seam carving [3], which allows to shrink



Fig. 1. Example of an historical map georeferencing process. Left: original map (1682, Sebastiano Alberti, ASV). Right: georeferenced map with respect to 15 manually defined control points overlapped to a satellite view of the area (Google Earth basemap). Map image courtesy of State Archive of Venice.

or expand images by removing or adding whole columns or rows of pixel found through a minimum saliency path from the top to bottom or left to right boundaries of the image.

Existing image retargeting approaches can be classified into two categories: discrete and continuous [10]. Discrete approaches alter the image size by eliminating pixels through cropping or seam carving. Recently, Rubinstein et al. [9] presented a multi-operator algorithm that combines cropping, linear scaling, and seam carving while Pritch et al. [8] remove repeated patterns in homogeneous regions.

Conversely, continuous approaches optimize mapping or warping using smoothness and salient-region preserving constraints to retain perceptual content. Wolf et al. [12] retargeted an image by merging less important pixels to reduce distortion. In this way, the distortion is propagated only along the resizing direction. Wang et al. [11] warp local regions to match optimal scaling factors, thus distributing the distortion in all directions. However, large salient objects may undergo inconsistent deformation throughout their extent. To ease this problem, Zhang et al. [13] and Guo et al. [6] force highly salient objects to undergo similarity transformations when resizing images, resulting in good preservation of local shape. However, inconsistent deformations can occur along elongated structures. Lin et al. [7] propose an approach to preserve both the visually salient objects and structure lines by constraining patches of high saliency to undergo similarity transformations. It is noteworthy to say that all the approaches are designed to give their best on photographic content rather than human-made drawings

In this paper we address a different application for Content-aware image warping: the georeferencing of historical maps (Fig.1). Historical maps provide vital evidence to scholars in disciplines as diverse as history, archaeology, and en-

environmental sciences, which routinely use them to identify and map information on past landscapes that is no more available on the ground and to compare their content with current geomorphological and spatial data within GIS platforms, where historical maps can be georeferenced. A number of tools available in these platforms enable warping images to fit control points to known coordinates; generally, however, these are limited to global context-independent transformations such as polynomial or spline based image warping. These global transformations provide sufficiently accurate results with historical maps produced starting from beginning/middle of 18th century, which were originated using more or less accurate topographic criteria.

Maps created in previous centuries, instead, were generally conceived more like a pictorial depiction of the landscape rather than a rigorous topographical representation of it. In this case, global context-independent transformations tend to destroy the information-rich pictorial content, failing to warp the features over the actual landscape due to the limited degrees of freedom of the transformations available.

A retargeting process enabling the georeferencing of historical maps will differ from traditional retargeting approaches in several key elements. First, the transformation is driven by control points resulting in different boundary conditions: Dirichlet for the control points and Neumann for the actual boundary. Second, salient visual content must be preserved after the transformation. Third, the map has an orientation, so the warp should be smooth and non-decreasing in both horizontal and vertical directions. Fourth, each iconographic element should maintain the original shape and orientation. (eg. the compass rose present in many historical maps should preserve its original direction and be minimally distorted). Finally, maps are rich in linear elements, which can be stretched along their direction, but not orthogonally. Hence, context-based constraint should be directional.

In accordance to these requirements, our approach provides three novel contributions. First, we defined a continuous image retargeting approach based on a set of control points. Second, we consider a tensor-based characterization of the saliency map to favor or penalize strong warping across salient edges. In other words, we generalize the saliency map from a scalar to a tensor field to control both the amount of stretching and its directionality among the image domain. Finally, we avoid the possible fold overs by embedding isotonic regression during the variational map optimization. This enables extreme stretches on possibly small parts of the images, avoiding folds and guaranteeing orientation and monotonicity of the map.

2 Problem formulation

We suppose to have a *source* image I_s defined over a spatial scale-independent domain $\Omega_s \subset \mathbb{R}^2$. Our goal is to transform the image on a new scale-independent domain Ω_t (which can be affinely mapped to world coordinates) according to some point-to-point correspondences $C_p = \{(p_1, q_1), (p_2, q_2), \dots, (p_n \in \Omega_s, q_n \in \Omega_t)\}$

that are manually defined with respect of some visual features that a human operator has identified in I_s . In Fig.1 (Left) we have a typical example of an historical map held at the State Archive of Venice and drawn by geographer Sebastiano Alberti in 1682. To georeference the picture, an algorithm should ideally map each pixel of I_s to a new world coordinate system so that the geographical features (like the rivers, islands, city locations etc.) are mapped to their corresponding places. Since the historical map was not intended for cartographic purposes, local scale of the pictorial content is subject to a certain degree of freedom introduced by the artist. In other words, we do not expect to have a single explicit function to map the whole image into the world reference frame. In this specific example, distances between the cities and the actual shape of the shoreline is not consistent across the image.

Our objective is to find a function $f : \Omega_t \rightarrow \mathbb{R}^2$ defining the displacement of each point from target to source domain. This allows the creation of a *target* image I_t by interpolating all the points of the source image with the mapping function $M(q) = q + f(q)$ such that

$$I_t(q) = I_s(M(q)) \quad (1)$$

and $M(q_i) = p_i \quad \forall (p_i, q_i) \in Cp$ (i.e. all the control points are mapped exactly between source and destination). In computer-vision terms, f is the optical-flow function densely mapping pixels from I_t to I_s .

2.1 Saliency tensor map

The fact that the important pictorial information should be preserved when transforming the source image implies a concept of *perceptually salient content* that must be identified a-priori. In the field of image retargeting, the problem is usually addressed by defining a saliency map $S(p) : \Omega_s \rightarrow [0, 1]$ to quantify the importance of each point of the source image. Such saliency map can be computed automatically from the source image [2] or obtained by directly drawing on it to customize the importance of each pixel [3]. In both the cases, the saliency map is a scalar field over the image domain used to weight the amount of stretching allowed in an area.

Since we require directional context-based constraints, we explicitly model a possible non isotropic response of the saliency map. In other words, the amount of stretching is made not only spatially but also directionally dependent. We start by considering a scalar saliency field S provided by the user to create a rank-2 tensor field T on Ω_t . Such tensor assign to each point $q \in \Omega_t$ a 2×2 matrix defined as:

$$T(q) = k_1 S(M(q)) \mathbf{I} + k_2 H(q), \quad H = W \star (\nabla I_t \nabla I_t^T) \quad (2)$$

where k_1 and k_2 are two parameters to weight the scalar and directional contribution of the saliency, \mathbf{I} is the identity matrix and H is a structure tensor computed over a window W .

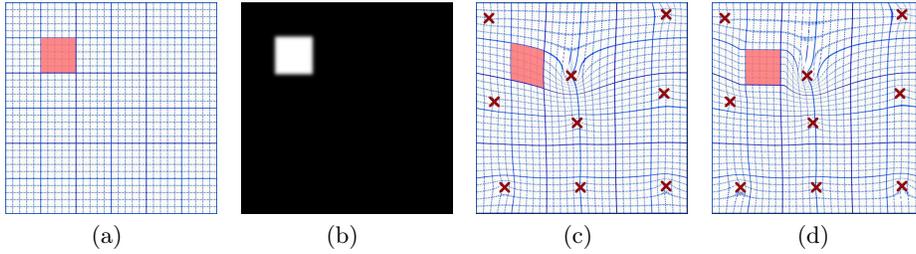


Fig. 2. Effect of the saliency map in our approach. (a) Synthetically generated source image showing a regular grid. (b) A toy-sample saliency map with high values for the pixels corresponding with the coloured square on the source image. (c) Retargeted image with constant saliency (ie. all the pixels have the same visual importance). Retargeted image using the saliency map in (b)

3 Variational functional

We pose our retargeting problem as the minimization of the functional:

$$E(f) = \int_{\Omega_t} \text{trace} \left(\mathbf{J}_f^T T(q) \mathbf{J}_f \right) dq \quad (3)$$

where $\mathbf{J}_f = \begin{pmatrix} \partial f_x / \partial x & \partial f_x / \partial y \\ \partial f_y / \partial x & \partial f_y / \partial y \end{pmatrix}$ is the Jacobian of f . Equation (3) leads to a variational optimization problem in which we set Neumann boundary conditions for the image borders and Dirichlet boundary conditions $f(q_i) = p_i - q_i$ on the control points. The idea is to find a displacement function f between source and target domains with pre-defined values on the control points and for which the first-order partial derivatives (ie. the local stretching caused by the mapping) is small in a quadratic sense. Additionally, the amount of stretching allowed in each direction is modulated by the saliency tensor field computed as in (2). According to the calculus of variations, $E(f)$ has a stationary value if the corresponding Euler-Lagrange equations:

$$\begin{aligned} \frac{\partial E}{\partial f_x} - \frac{\partial}{\partial x} \frac{\partial E}{\partial f_x / \partial x} - \frac{\partial}{\partial y} \frac{\partial E}{\partial f_x / \partial y} &= 0 \\ \frac{\partial E}{\partial f_y} - \frac{\partial}{\partial x} \frac{\partial E}{\partial f_y / \partial x} - \frac{\partial}{\partial y} \frac{\partial E}{\partial f_y / \partial y} &= 0 \end{aligned} \quad (4)$$

are satisfied. This leads to a set of PDEs which are solved numerically as described in the following sections.

3.1 Special cases

In its formulation, our method generalizes other variational warp-based image retargeting techniques [5, 10] but with two main differences. First, the Dirichlet

boundary conditions are set on the control points and not the image boundaries which are free to overshoot or undershoot the source image borders. Second, the saliency map is directional depending on the linear features of the source image itself.

If we set $k_2 = 0$, we eliminate the latter so that the functional (3) can be rewritten as:

$$\int k_1 S(M(q)) \left[\left(\frac{\partial f_x}{\partial x} \right)^2 + \left(\frac{\partial f_x}{\partial y} \right)^2 + \left(\frac{\partial f_y}{\partial x} \right)^2 + \left(\frac{\partial f_y}{\partial y} \right)^2 \right] dq.$$

To simplify, we assume to already have a saliency map $\hat{S}(q)$, in target space, giving a reasonable approximation of $S(M(q))$. According to the Euler-Lagrange equations, the resulting system of PDEs is:

$$\begin{aligned} -2k_1 \hat{S} \Delta f_x &= 0 \\ -2k_1 \hat{S} \Delta f_y &= 0 \end{aligned} \tag{5}$$

where Δ is the Laplace operator. Since the off-diagonal elements of $T(q)$ are equal to 0, each equation depends only on the horizontal and vertical component of f respectively and thus can be computed in parallel as set of elliptic PDEs. The only case in which $\hat{S}(q) = S(M(q))$ is when the saliency is constant everywhere. In this case, the deformation is not saliency aware but only depends on the local deformation of the control points (Fig.2(c)). For all the other cases, we start from an initial approximation of f and we alternate the computation of f and $\hat{S}(q)$ until convergence (See section 4 for details).

4 Numerical solution

To solve the general problem numerically, we consider a discrete approximation of Ω_s and Ω_t into a regular grid with arbitrary resolution, depending on the required level of detail. Consequently, we transform f into a $M \times N \times 2$ tensor and transform the functional (3) to:

$$\begin{aligned} E(f) = \sum_{u=0}^{N-1} \sum_{v=0}^{M-1} \text{trace} \left(\mathbf{J}_f^T \hat{T}(u, v) \mathbf{J}_f \right) + \\ + \sum_{(p,q) \in C_p} \alpha (p - M(q))^2. \end{aligned} \tag{6}$$

To let the computation more stable, we substituted the Dirichlet boundary conditions on the control points with a soft $L2$ constraint on the flow values on the control points. Note that, while with this formulation the exact mapping of control points is not guaranteed, for high values of α the effect is negligible. The corresponding Euler-Lagrange equations are thus:

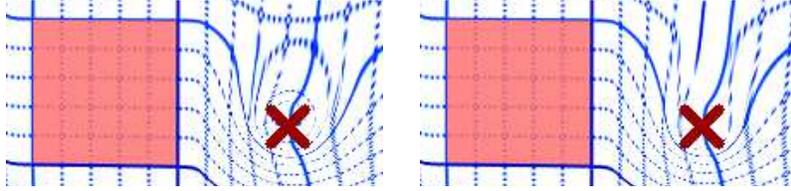


Fig. 3. Effect of the isotonic regression on the resulting distortion map. Left: with no monotonicity constraint singular solutions may cause bubbles around a control point. Right: The isotonic regression step force the resulting mapping to a non-decreasing solution thus removing such artifacts.

$$\begin{aligned} L_x &= 2\delta\alpha(f_x - \hat{f}_x)^2 - 2\hat{T}_{11}\Delta f_x - 2\hat{T}_{12}\Delta f_y = 0 \\ L_y &= 2\delta\alpha(f_y - \hat{f}_y)^2 - 2\hat{T}_{21}\Delta f_x - 2\hat{T}_{22}\Delta f_y = 0 \end{aligned} \quad (7)$$

where δ is a mixture of Kronecker Deltas centered at $q_1 \dots q_n$ and \hat{f} is the known value of f at the control points. We consider an initial approximation of f obtained by interpolating the control points via multiquadric *radial basis functions* and iteratively modify f over time so that $\partial f_x / \partial t = -L_x$ and $\partial f_y / \partial t = -L_y$. We use a discrete approximation of the Laplace operator computed by convolution and forward finite differences for the time differentiation to solve the PDEs via forward Euler scheme. At each step t , we consider the function computed at previous step $f^{(t-1)}$ to obtain \hat{T} according to equation (2).

4.1 Avoiding the fold overs

Variational functional of Eq. (6) pose only constraints on the smoothness of the obtained warp but not on its monotonicity. Before starting the optimization, we compute an initial global affine transformation of the control points so that global rotations or flips of the source domain with respect to target domain are eliminated. After that, we expect that $M(u, v) \leq M(u + 1, v)$ and $M(u, v) \leq M(u, v + 1)$ for each point of the domain. If such constraint is violated, the resulting topology of the retargeted I_t will be no longer consistent with the original image lattice. In particular, we obtain flip overs on the warped image that manifest themselves as bubbles around the control points (See Fig. 3).

To avoid this, during the optimization we alternate the computation of f with a projection to a new space to obtain the closest approximation f^i of f (in $L2$ terms) but with the additional constraint that each row of f_y^i and each column of f_x^i are non-decreasing. This is performed by using the isotonic regression technique [4] applied in parallel to the rows and columns of f . Overall, the numerical optimization is performed by alternating the following steps:

1. Let f^0 be an initial interpolation of the control points using RBF
2. Compute \hat{T} as in eq. (2) using $f^{(t-1)}$ to approximate M and ∇I_t

3. Use forward Euler schema to compute $f^{(t)}$
4. Substitute $f^{(t)}$ with its non-decreasing approximation f^i obtained via isotonic regression
5. Return to step 2 until convergence.

5 Experimental evaluation

We qualitatively tested our approach on three historical maps held at the State Archive of Venice (Italy). The first, (Fig.5, left), drawn by Giovanni Battista Faventini in 1734, represents the costal line of the Marano and Grado’s lagoon (NE Italy) and the mainland comprised between Tagliamento and Isonzo rivers³; the second (Fig.5, right), drawn by unknown artist in 1624 circa, characterizes the town of Grado and surrounding territories up to Terzo and Fiumicello villages⁴. The third map, (Fig. 7, top-left) also sketched by unknown artist in the 17th century, depicts the lower part of the Friuli region up to the marshes surrounding Grado. These maps were chosen as good example of map content rich of iconographic elements and loosely accurate with respect to the real cartographic features.

In all the cases, the saliency map was generated in a semi-supervised way. Specifically, we first computed the Difference of Gaussian of the input images and then re-normalized the result in range $[0 \dots 1]$. Then, we manually painted on top of the saliency map the areas in which interesting iconographic features (like the compass rose or the place names) were present. For reference, saliency maps for the images used in our experiments are shown in Fig.4 and Fig. 7 (top-right) respectively. Note that the intended usage of the saliency is not to define areas of the historical map in which the cartographic information is more accurate but simply to mark the elements (mostly iconographic) for which the local deformation should be minimum. We ran our optimization using $k_1 = 1.0$,

³ Gio. Battista Faventini in 1734, ASV, Senato Rettori, F. 232, d. 1

⁴ Unknown,1624, Dispacci Rettori di Palma, F.21, d. 1 + part

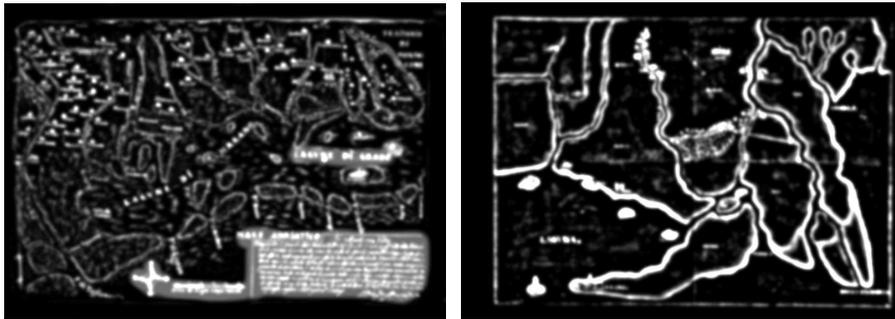


Fig. 4. The saliency map of Faventini 1734 (Left) and Unknown 1624 (Right) used for our experiments

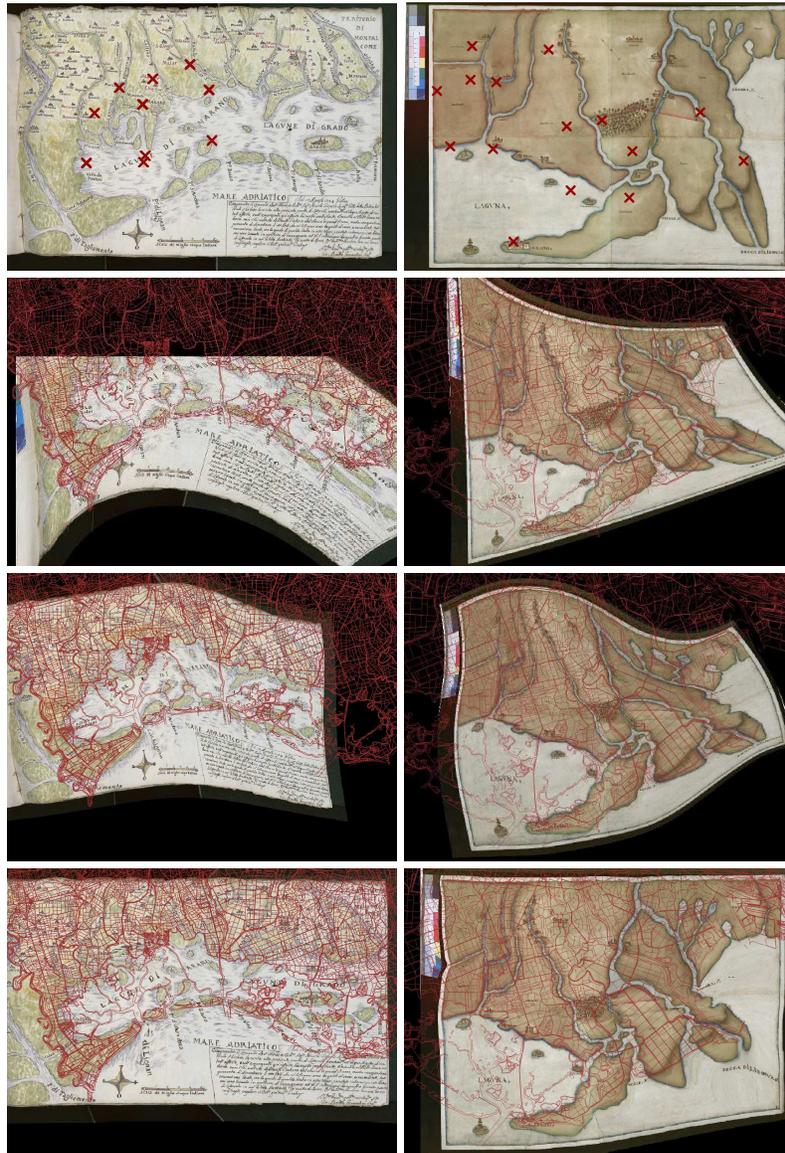


Fig. 5. Comparison between our algorithm and point-based warping methods for two different historical maps (left column: Faventini, 1734; right column: Unknown, 1624). First row: original map with the control points and saliency map superimposed. Second row: Second degree polynomial warping. Third row: Thin-plate spline warping. Fourth row: our proposed method. Note how our retargeting approach enables a better georeferentiation of the input picture while minimizing local distortions of the symbolism, geomorphology and the place names. Images courtesy of State Archive of Venice.

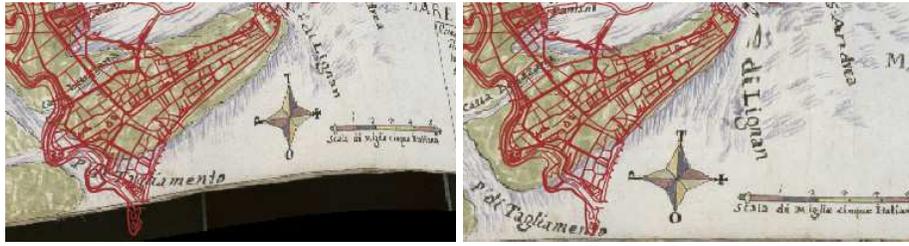


Fig. 6. Detail of the local stretching affecting the compass rose on the bottom of the Faventini 1734 map shown in Fig.5 after processing with TPS (Left) and our method (Right).

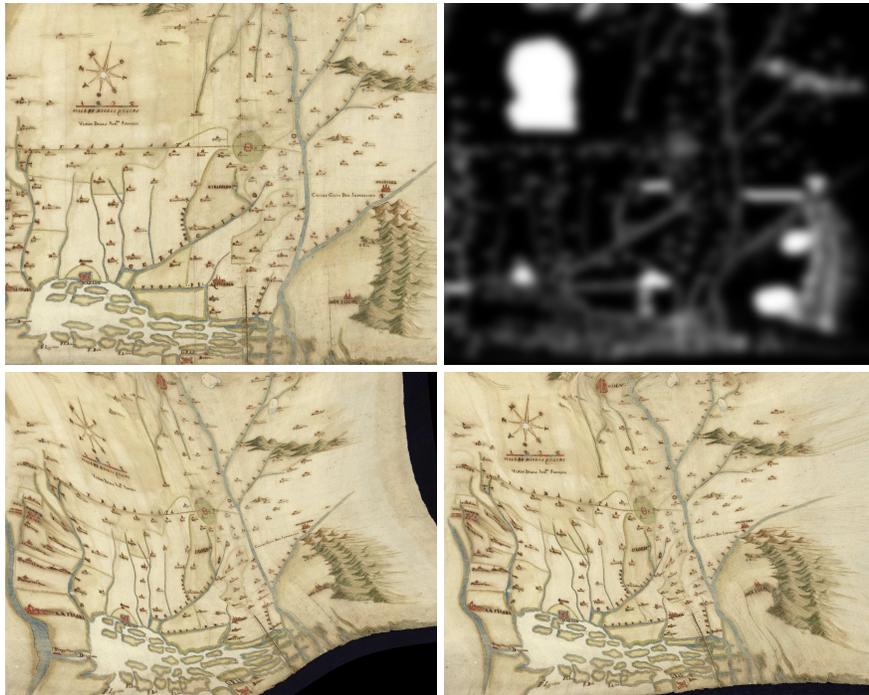


Fig. 7. Retargeting result on a complex pictorial map for which the control point placement causes a highly non-uniform distortion. Top-left: Original image. Top-right: Saliency map in which we manually gave high importance to the pictorial content of the map (see the white blob corresponding to the compass rose). Bottom-left: warping via cubic thin-plate spline. Bottom-right: retargeting computed with our proposed method. Note how our result is less "extreme" while still preserving the correct positioning of the elements. Moreover, the purely pictorial content expressed by the saliency map is retargeted with almost no distortion. Images courtesy of State Archive of Venice (Raccolta Terkuz, n. 34)



Fig. 8. Detail of the local stretching of the upper part of the 1624 (Unknown) map shown in Fig.5 (right column) throughout the optimization. From left to right: retargeted image after 0, 5000, 20000 and 40000 iterations. The saliency tensor allows the rivers to warp more freely along their extent instead of between the banks. This way, the thickness of the linear features is preserved. Note also how the appearance of the two cities evolve toward a low distortion configuration without changing their relative position in the georeferenced image.

$k_2 = 10$, $\alpha = 100$ and a Gaussian 5×5 window W for the computation of structure tensor in all the cases.

In Fig. 5 we compare the result obtained with our technique (bottom row) against two common image warping techniques implemented in QGIS, namely: second degree polynomial and thin plate splines (TPS). We superimposed the actual national base map of the area (red lines) to highlight the overall map fitting accuracy to the geographical features. Since both polynomial and TPS are designed to work with cartographic maps and satellite images, they give the same importance to each pixel of the map. This result in often too exaggerated distortions in the attempt to satisfy each control point constraint. In both the cases, our approach gives a more natural result minimizing the local stretch for areas both far from the control points or subject to low visual saliency.

To better highlight this behavior, in Fig. 6 we show a magnification of the bottom area of the first map comprising both the shoreline and purely iconographic elements (ie. the compass rose and the scale bar). While the shoreline correctly overlaps the actual cartography in both the cases, the shape and orientation of the two iconographic elements are better preserved with our method. In particular, the direction of the compass rose is correctly aligned with the cardinal directions and the scale bar is maintained horizontal and free to strong deformations. This is a peculiar feature of our method, that can exploit a saliency tensor field to selectively limit the distortion on specific parts of the map. In fact, we observe an actual re-positioning of all the high-saliency elements of the map more than a typical warping we obtain with interpolation approaches commonly implemented in GIS. This is further shown in Fig. 7 where we compare our method against cubic thin-plate splines on what we consider the most challenging historical map in our set.

To visualize the evolution of the flow during the optimization, in Fig. 8 we show a magnification of the top area of the map shown in Fig. 5 (Left) at different iterations of our algorithm. The process starts with an RBF interpolation (left-most image) of the displacement computed for each control point. Since the interpolation is not aware of the saliency map, we can observe a strong deforma-

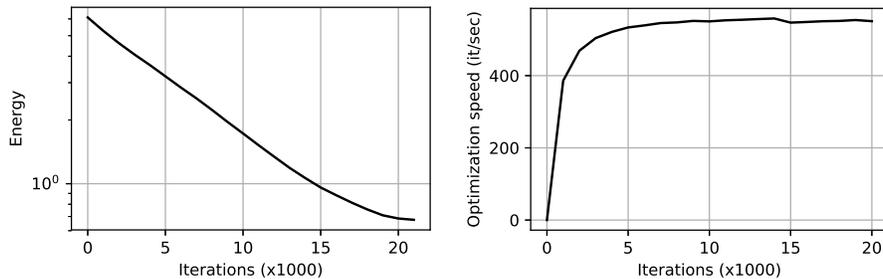


Fig. 9. Left: Evolution of the energy functional (6) throughout the iterations. Right: Optimization speed in iterations per second. Both plots refer to the map (Raccolta Terkuz, n. 34) presented in Fig. 7

tion of the two stylized cities and their related captions. During the iterations, the stretching affecting the two cities is smoothly reduced in accordance to the high value on the saliency map. The effect is particularly visible on the most southern city (Grado), far from any control point and hence implicitly allowed to move. On the other hand, the control point just on top on the north-west city cause a strong deformation of the land eastbound it since no salient content is present. Other interesting observations can be made by analyzing the shape of the rivers among the cities. Consistent with placement of the control points, the intersection of the two main rivers moves easterly during the optimization. However, thanks to the tensor form of the saliency component, the algorithm smoothly modifies the relative angle of the two rivers (and thus their length) without changing their thickness. This is notably evident on the river oriented north-south for which the relative movement is orthogonal to its linear extent. This behavior is remarkably important as it contributes to preserve the overall visual appearance of the original map.

5.1 Implementation details and running times

We designed our method so that it can be implemented in terms of convolutions and simple arithmetical operations among tensors of rank 2 and 3. In practice, we observed that is computationally more efficient to use a very simple variational solver (Forward-Euler) with a small time-step instead of a sophisticated integration scheme that do not scale well when retargeting high resolution maps.

We implemented our method using the Python version of the popular TensorFlow framework [1]. We performed our experiments on a consumer desktop PC, with an Intel i7 CPU and a Nvidia GeForce 1060 GPU. In all the cases, flow function was scaled during the optimization to a rank-3 tensor with shape $(256, 256, 2)$ and the time delta Δ_t for the Forward-Euler scheme was set to $5E-6$. In Fig. 9 we plotted the value of the energy functional (6) during the optimization (left plot) and the number of iterations per second (right plot) while optimizing the map presented in Fig. 7. With the parameters used for our tests,

the whole optimization took 25000 to converge to a minimum at a speed of 520 iterations per second on average. This implies a total time ranging between 40 seconds to a minute to perform a full map retargeting.

6 Conclusion

We proposed a novel technique to perform dense continuous image retargeting driven by a set of control points. It was designed at georeferencing historical maps for which the local distortion of the pictorial content is dynamically adjusted according to a saliency tensor depending both on a saliency map defined by the user and the structure tensor computed on the image itself. This allows part of the map to be warped smoothly to accommodate the control points while preserving the local stretching of purely pictorial content (like the common compass roses) that are seamlessly re-positioned in the resulting image.

Our algorithm involves an optimization of an energy functional based on the calculus of variation that can be efficiently implemented in GPU by means of convolutions and tensor operations, commonly implemented in many frameworks designed for Convolutional Neural Networks and Deep Learning. During the optimization, we included an isotonic regression step to avoid visual artifact in the resulting displacement map. The isotonic regression guarantees that no flip-overs will be present in the retargeted output, even in presence of extreme local stretching that may happen when warping maps not created for cartographic purposes. Qualitative tests performed on original historical maps demonstrate the effectiveness of our approach especially when handling multiple pictorial elements. At the present state, our method still requires human intervention in the definition of the control points and the saliency map. For the near future, we plan to automate the saliency map creation by exploiting some of the existing approaches in the literature and implement our method as a QGIS plugin.

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), <https://www.tensorflow.org/>, software available from tensorflow.org
2. Achanta, R., Süsstrunk, S.: Saliency detection for content-aware image resizing. In: Proceedings of the 16th IEEE International Conference on Image Processing. pp. 1001–1004. ICIP’09, IEEE Press, Piscataway, NJ, USA (2009)
3. Avidan, S., Shamir, A.: Seam carving for content-aware image resizing. *ACM Trans. Graph.* **26**(3) (Jul 2007). <https://doi.org/10.1145/1276377.1276390>
4. Chakravarti, N.: Isotonic median regression: A linear programming approach. *Math. Oper. Res.* **14**(2), 303–308 (May 1989). <https://doi.org/10.1287/moor.14.2.303>, <http://dx.doi.org/10.1287/moor.14.2.303>

5. Gal, R., Sorkine, O., Cohen-Or, D.: Feature-aware texturing. In: Proceedings of the 17th Eurographics Conference on Rendering Techniques. pp. 297–303. EGSR '06, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2006). <https://doi.org/10.2312/EGWR/EGSR06/297-303>, <http://dx.doi.org/10.2312/EGWR/EGSR06/297-303>
6. Guo, Y., Liu, F., Shi, J., Zhou, Z.H., Gleicher, M.: Image retargeting using mesh parametrization. *IEEE Trans. Multimedia* **11**(5), 856–867 (2009), <http://dblp.uni-trier.de/db/journals/tmm/tmm11.html#GuoLSZG09>
7. Lin, S.S., Yeh, I.C., Lin, C.H., Lee, T.Y.: Patch-based image warping for content-aware retargeting. *IEEE Trans. Multimedia* **15**(2), 359–368 (2013), <http://dblp.uni-trier.de/db/journals/tmm/tmm15.html#LinYLL13>
8. Pritch, Y., Kav-Venaki, E., Peleg, S.: Shift-map image editing. In: ICCV. pp. 151–158. IEEE Computer Society (2009), <http://dblp.uni-trier.de/db/conf/iccv/iccv2009.html#PritchKP09>
9. Rubinstein, M., Shamir, A., Avidan, S.: Multi-operator media retargeting. *ACM Trans. Graph.* **28**(3) (2009), <http://dblp.uni-trier.de/db/journals/tog/tog28.html#RubinsteinSA09>
10. Shamir, A., Sorkine, O.: Visual media retargeting. In: ACM SIGGRAPH ASIA 2009 Courses. pp. 11:1–11:13. SIGGRAPH ASIA '09, ACM, New York, NY, USA (2009). <https://doi.org/10.1145/1665817.1665828>
11. Wang, Y.S., Tai, C.L., Sorkine, O., Lee, T.Y.: Optimized scale-and-stretch for image resizing. *ACM Trans. Graph.* **27**(5), 118 (2008), <http://dblp.uni-trier.de/db/journals/tog/tog27.html#WangTSL08>
12. Wolf, L., Guttman, M., Cohen-Or, D.: Non-homogeneous content-driven video-retargeting. In: ICCV. pp. 1–6. IEEE Computer Society (2007), <http://dblp.uni-trier.de/db/conf/iccv/iccv2007.html#WolfGC07>
13. Zhang, G.X., Cheng, M.M., Hu, S.M., Martin, R.R.: A shape-preserving approach to image resizing. *Comput. Graph. Forum* **28**(7), 1897–1906 (2009)