The Shape Collapse Problem in Image Registration

Oguz C. Durumeric[1], Ipek Oguz[2], and Gary E. Christensen[2,3]

Departments of [1]Mathematics, [2] Electrical and Computer Engineering and [3] Radiation Oncology, The University of Iowa

Abstract. This paper investigates the shape collapse problem in non-rigid image registration. The shape collapse problem is the situation when an appendage of a deforming object does not overlap with the target shape and collapses to a set of zero measure during the registration process. The dual problem occurs when a new appendage grows out of the object to match the target shape. In both cases, the estimated correspondence between the source and target objects is often undesirable. The shape collapse problem is caused by deforming the moving image in the gradient direction of the similarity cost and affects both small and large deformation registration algorithms. Minimizing a registration cost function by following the similarity-cost gradient drives the registration to a local energy minimum and does not permit an increase in energy to ultimately reach a lower energy state. Furthermore, once an object collapses locally, it has zero measure under the similarity cost in this region and is permanently stuck in a local minimum. This paper presents a criterion for detecting image regions that will collapse if the similarity cost gradient direction is followed during optimization. This criterion is based on the skeletal points of the moving image in the symmetric difference of the original two binary images. Experimental results are presented that demonstrate that the shape collapse problem can be detected before registration.

Keywords: image registration, skeleton, shape collapse, collapse points, predicting regions of collapse

1 Introduction

A common problem that affects both small and large-deformation image registration algorithms is when a part of an object collapses during the registration process. The problem can be illustrated by registering images of hands in which four of five fingers overlap (see Fig. 1). The shape collapse occurs during registration when the index finger in the source image collapses into a thin set. At the same time, a new finger grows out of the palm of the hand to match the finger in the target hand. The biologically relevant transformation in this case would be instead to shift the finger to match its pose in the target image. This solution is difficult to reach with greedy optimization algorithms often used in image registration, because collapsing the index finger in the source image to a

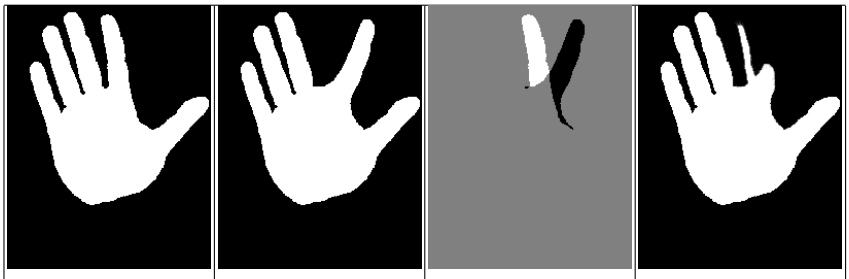

(a) Moving image | (b) Target image | (c) Difference before registration | (d) Deformed moving image

Fig. 1. Example of an undesirable shape collapse during image registration.

thin set reduces the registration similarity cost immediately, as opposed to rotating the finger which may increase the cost initially. Note that Fig. 1-d shows the final result for a small deformation image registration and only a partial deformation for a large deformation registration. Additionally, once a region of source image collapses during the registration process, it remains collapsed, i.e. a collapsed region in the source image corresponds to a local minimum of the registration cost function.

We are not aware of others that have studied the shape collapse problem in image registration. The distance function from smooth submanifolds and its cut loci have been thoroughly studied by many in differential geometry. See [2] for basic introduction and a review. The notion of skeleton generalizes the notion of cut locus when the boundary is not differentiable. Skeletonization and medial representations have been thoroughly studied; a good overview can be found in [3].

In this paper we present a novel method for a priori detection of object points where shape collapse is likely to occur and predicting the collapse loci. We provide several examples where our algorithm successfully predicts collapse loci. Detecting collapse loci is important for the alignment of sulci in brain image registration, which is similar to the alignment of fingers in the hand example discussed above. In this case, once we detect collapsing sulci, we can devise solutions to this problem by, for instance, enforcing the sulcus to retract into or sliding along the cortex rather than collapsing. This may also be used as a potential method to detect different folding patterns between the brains. The approach to handle collapsing is different based on the application domain, and as such, is beyond the scope of this paper.

2 Mathematical Formulation

The ideas presented in this paper are independent of space dimension. For example, the area-decreasing discussion below becomes volume-decreasing in 3-D. Likewise, the definition of skeleton below allows for curves and sheets in 3-D.

Let Ω be a subset of $\mathbb{R}^n$, typically with $n = 2$ or 3. Let $I_i : \Omega \to \mathbb{R}$, $i = 1, 2$ be two binary images to be registered, where I_1 denotes the moving image and I_2 denotes the target image. We denote the foreground of I_i by $V_i \subset \Omega$ for $i = 1, 2$. We assume that the V_i is compact. For the sake of simplicity, we will assume that the boundary of V_i is piecewise C^1 closed curve when $n = 2$. In the $n = 3$ case, the boundary of V_i is assumed to be a piecewise-smooth surface that is a union of finitely many surfaces, curves and points (e.g., the boundary of a cube).

The collapse problem occurs when V_1 has an appendage which is not included in V_2, and the growth problem occurs when V_2 has an appendage which is not included in V_1. The collapse and growth behaviors look different if they are observed only from the point of view of the foreground. However, the growth of the foreground can be studied as a certain type of collapsing of the background.

Definition 1. *Let V be a subset of a metric space (Ω, d), and $B_r(p)$ denotes the open metric balls $\{x \in \Omega : d(x, p) < r\}$. A closed ball $\overline{B} \subset V$ is called a maximal ball of V, if for every closed ball $\overline{B}'$, $\overline{B} \subseteq \overline{B}' \subseteq V$, one has $\overline{B} = \overline{B}'$. The set $\left\{q \in V : \exists r > 0, \overline{B}_r(q) \text{ is maximal ball of } V\right\}$ is defined to be the skeleton $S(V)$ of V by maximal balls [4].*

Suppose Ω is a metric space. Let p be a point of Ω, and let V be any subset of Ω. The interior of V, denoted $\mathrm{Int}(V)$, is the union of all open subsets of Ω contained in V. The exterior of V, denoted $\mathrm{Ext}(V)$, is the union of all open subsets of Ω contained in $\Omega - V$. The boundary of V, denoted by ∂V, is the set of all points of Ω that are in neither $\mathrm{Int}(V)$ nor $\mathrm{Ext}(V)$.

We first discuss the collapsing behavior in the simple case of $V_2 = \varnothing$ and no regularization. Since we have a binary image, a greedy algorithm for the similarity-cost gradient follows the direction that decreases the area in the fastest way. We study this area-decreasing flow on the interior of V_1, and the distance function $f : V_1 \to [0, \infty)$ to the boundary ∂V_1, given by $f(x) = dist(x, \partial V_1) = \inf\{|x - y| : y \in \partial V_1\}$. For a piecewise C^1 boundary, f is C^1 except on a set (containing the skeleton and the boundary) of measure zero, and its gradient ∇f (when it exists) is perpendicular to the level sets of f.

The fastest area decreasing occurs by deforming the interior of V_1 along ∇f until the skeleton is reached at that direction, since following the gradient one goes into the interior deeper in the fastest way. The superlevel sets $sl_r = \{x \in V_1 : f(x) \geq r\}$ are the stages of this deformation in the continuous category, as justified below.

For every $q \in S(V_1)$, there is a unique maximal ball $\overline{B}_r(q)$ of V_1 centered at q, and the set of points along the boundary ∂V_1 associated to q is $A_\partial(q) := \overline{B}_r(q) \cap \partial V_1$. The associated set $A_\partial(q)$ is nonempty. It usually has two or more points, but it can be one point, such as the focal point of a boundary curve (or a surface for $n = 3$) at a strict local maximum of curvature (or principal curvature for $n = 3$). For every $p \in \partial V_1$, we can also define the associated set $A_S(p) = \{q \in S(V_1) : p \in A_\partial(q)\}$, which is the set of points along the skeleton associated with p. If ∂V_1 is C^2 about p, then $A_S(p)$ contains one point. Let $L(p, q)$ denote the line segment with end points p and q, and $L^\circ(p, q) = L(p, q) - \{p, q\}$. For

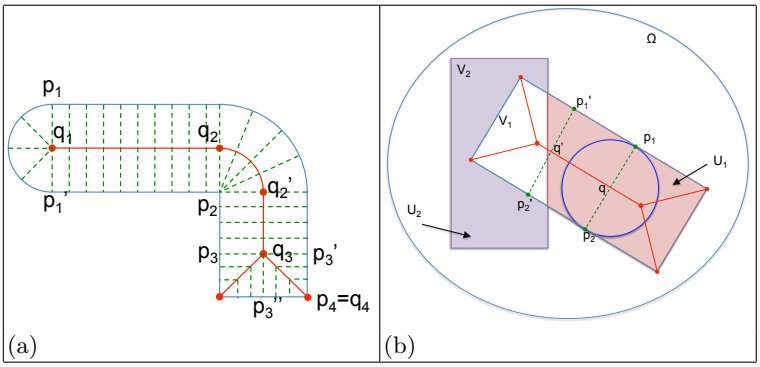

Fig. 2. *(a)* Possible configurations of the skeletal and boundary points. A skeletal point can correspond to one (q_4), two (q_2), three (q_3) or more (q_1) boundary points. Similarly, a boundary point can have one (p_4) or many (p_2) associated skeletal points. *(b)* Collapsing and non-collapsing points. p_1 is a collapsing point because it belongs to the same maximal ball $B_r(q)$ as p_2, and both p_1 and p_2 are in the non-overlap region U. In contrast, p_1' is not a collapsing point because it only shares a maximal ball ($B_r(q')$) with p_2', but p_2' is not in U.

simplicity, we will use 2-D examples to illustrate these concepts throughout this paper, even though the extension to n-D using line segments is straightforward. Figure 2 illustrates the possible configurations of boundary and skeletal points and their associations, as well as the notation used.

Proposition 1. *The set of line segments* $\{L^\circ(p,q) : q \in S(V_1)$ *and* $p \in A_\partial(q)\}$ *is a partition of* $V_1 - (\partial V_1 \cup S(V_1))$.

Proof. We want to show that every point $a \in V_1$ belongs to a $L(p,q)$ where $p \in \partial V_1$, $q \in S(V_1)$ and they are associated with each other. The point a has a closest point $p \in \partial V_1$ at a distance $f(a) = c$ and hence $\overline{B}_c(a) \subseteq V_1$. Then $\overline{B}_c(a) \subseteq \overline{B}_r(q) \subseteq V_1$ for a maximal ball $\overline{B}_r(q)$. If $p \in \overline{B}_c(a) \subseteq \overline{B}_r(q)$ is an interior point of $\overline{B}_r(q)$ then it cannot be on ∂V_1. Hence $\partial\overline{B}_r(q)$ and $\partial\overline{B}_c(a)$ are tangential at p, and thus $a \in L(p,q)$. Therefore V_1 is the union of all $L(p,q)$ with $q \in S(V_1)$ and $p \in A_\partial(q)$.

For $p_i \in \partial V_1$ and $q_i \in S(V_1)$ with $p_i \in A_\partial(q_i)$, for $i = 1, 2$, if $L(p_1,q_1) \cap L(p_2,q_2) \neq \varnothing$ then either $p_1 = p_2$ or $q_1 = q_2$, i.e. these segments can only intersect along the boundary or the skeleton unless they are identical. Because, an intersection at the interior of the line segments would contradict the triangle inequality and p_i being the closest point of ∂V_i to q_i. □

By Proposition 1 and its notation, if $a \in L^\circ(p,q)$, then p is the unique closest point of ∂V_1 to a. This justifies that $f^{-1}(r)$ is the set of all points a_{pq} on $L(p,q)$ with $\|a_{pq} - p\| = r$ as long as $\|q - p\| \geq r$, as q varies through all of $S(V_1)$ and p through all of $A_\partial(q)$. The fastest area decreasing flow on the interior of V_1 is along the line segments $L^\circ(p,q)$ with unit speed until the skeleton is reached

along each segment. The flow is not definable along the skeleton. However, away from the skeleton, sl_r are the stages of this deformation, and $L(p,q)$ is parallel to $\nabla f(a_{pq})$ if it exists.

We remark that this flow needs to be considered on the interior of V_1, since it is possible that $A_S(p)$ may contain more than one point if ∂V_1 is not differentiable at $p \in \partial V_1$ (e.g. p_2 in Fig. 2-a). In that case, the flow is not defined at p, but it is well defined on the union of $\{L^\circ(p,q) : \forall q \in A_S(p)\}$. Rather than letting ∂V_1 follow the flow, we use the boundary of sl_r, namely $f^{-1}(r)$. Using sl_r has the additional advantage of describing the area reduction process accurately around the skeleton. This is because $S(sl_r) = S(V_1) \cap sl_r$ and the flow along the $L(p,q)$ reaches different parts of the skeleton at different times and stops. Therefore it is possible that sl_r may break into components. In this case, the flow is not continuous at the skeleton; however, the area-reduction is continuous.

In the continuous area-reduction case and with $V_2 = \varnothing$, the above process (without regularization) will reduce V_1 to $\varnothing$ since it will remove the skeleton along the way. By using small (but not zero) regularization and discrete image registration, we see that the skeleton actually remains. The deforming forces towards a point q on the skeleton are from different directions along $L(p,q)$ for several $p \in A_\partial(q)$. The regularization will reduce the effect of these forces by averaging and slow down the area reduction. Since the skeleton has measure zero and the deformations we use for the image registration are discrete, removing the skeleton has no or little gain in the cost function, and therefore a skeleton remains. Since the deformations are not going exactly along sl_r, the skeleton reached through image registration cannot be expected to be same as $S(V_1)$, but it is a very good approximation since the above mentioned factors are not in effect until the flow comes near $S(V_1)$.

The next step is to consider the case in which the target image contains a foreground object, i.e., $V_2 \neq \varnothing$. We begin with some definitions.

Definition 2. *Let V_1 and V_2 be subsets of Ω representing the foreground objects of images I_1 and I_2, respectively. The overlap of the foreground objects is denoted as $W = V_1 \cap V_2$ and the symmetric difference region (non-overlap) $U = U_1 \cup U_2$ where $U_1 = V_1 - V_2 = \{x \in V_1 : x \notin V_2\}$ and $U_2 = V_2 - V_1$.*

Definition 3. *A point $p_1 \in \partial V_1$ is a collapsing point for $q \in S(V_1)$, if $\exists p_2 \in \partial V_1$ such that $p_1 \neq p_2$, $\{p_1, p_2\} \subset A_\partial(q)$, and $L(p_1,q) \cup L(p_2,q) \subset U_1$. The notation* CLPS($q$) *denotes all collapsing points (if any) along ∂V_1 associated with $q \in S(V_1)$ and* CLPS(V_1) *denotes the union of* CLPS(q) *for all $q \in S(V_1)$.*

Figure 2-b illustrates these definitions. In this example, p_1 is a collapsing point because there exists a p_2 satisfies all three conditions: $p_1 \neq p_2$, $\{p_1, p_2\} \subset A_\partial(q)$, and $L(p_1,q) \cup L(p_2,q) \subset U_1$. In contrast, p_1' is not a collapsing point because the only point in ∂S_1 that shares a maximal ball with it is p_2', and $p_2' \notin U_1$. During the registration of the image of S_1 into the image of S_2, p_1 p_2 will move towards q on the skeleton of S_1 whereas p_1' will move towards S_2 and p_2' will stay in S_2.

Regions of foreground reduction of the moving image occur on U_1 and regions of foreground growth occurs on U_2. We can study U_1 and U_2 separately, since the deformation algorithms are local. We start by discussing the area-decreasing flow away from skeleton $S(V_1)$ in the continuous case without regularization again. If $p_1 \in \partial V_1 \subset V_1$ but $p_1 \notin U_1$, that is $p_1 \in V_2$, then it is already in the target set, there will be no area-decreasing flow near p_1. If $p_1 \in U_1 \cap \partial V_1$, then the area-decreasing flow on the union of $\{L^\circ(p,q) : \forall q \in A_S(p)\}$ starts as described before for p near p_1, and continues along $L^\circ(p,q)$ until it comes close to either (i) a $q \in A_S(p)$ when $L^\circ(p,q)$ is away from the overlap W, or (ii) the flow comes close to a point of ∂V_2 provided that $L^\circ(p,q)$ enters V_2 at most once.

For the cases when $p_1 \in U_1 \cap \partial V_1$, $q \in A_S(p_1)$ and $L(p_1,q) \subset U_1$, but p_1 is not a collapsing point due to the rest of $A_\partial(q)$ being in W, there will be one-sided flows towards q at the start. However, since these one-sided flows will alter the skeleton of the remaining set, the flow needs to be readjusted according to the new skeleton, a case that will not happen when $V_2 = \varnothing$. All flows near W will have this adjustment.

At the collapsing points (nonempty) $\mathrm{CLPS}(q) \subset A_\partial(q)$, we expect that the area decreasing flow to reduce the area tending towards q from the several directions $\{L(p,q) : \forall p \in \mathrm{CLPS}(q)\}$, and some type of collapse occur at q.

If there exists an open set Z bounded away from W such that for all $q \in Z \cap S(V_1)$ one has $\mathrm{CLPS}(q) = A_\partial(q)$ and for all $p \in \mathrm{CLPS}(q)$, $L(p,q) \subset Z$ then the area decreasing procedure in Z will behave the same as $V_2 = \varnothing$. Understanding how this works allows us to predict the outcome with a small regularization factor. A part of a skeleton will be reached through image registration for the collapsing of U_1 within Z, it cannot be expected to be same as a part of $S(V_1)$, but it is a good approximation since the above mentioned factors when $V_2 = \varnothing$ are not in effect until the flow comes near this particular portion of $S(V_1)$.

The dual problem of growth in the foreground of the moving image is a reduction in the background. Hence, we can predict the behavior of the growth as well by using the same proceedure. Revisiting Fig. 1, we can now compare the collapse and growth of the index finger. The collapse of the index finger U_1 is mostly sideways flowing towards the vertical part of the skeleton $S(V_1)$ through U_1, by becoming thinner and thinner. But the growth of the finger in U_2 is a reduction of a part of the complement of V_1 (not V_2!). It has to follow the level sets of the distance function to ∂V_1 on the complement of V_1. These levels sets restricted to U_2 are parallel to the index finger in the moving image and therefore U_2 grows in the perpendicular direction to these level sets. The growth in U_2 keeps the maximal width since there is no skeleton of the background of the moving image about that area.

3 Discussion

Figure 3 presents a simple 2D registration example that illustrates the shape collapse problem for parts of the foreground and background of the moving image. The moving image (Fig. 3-a) is a 128x128 pixel, binary image that has

value 0 for the background and 1 for the foreground. The target image (Fig. 3-b) differs from the moving image by a translation of the lower rectangle appendage. The difference image (Fig. 3-c) demonstrates the amount of overlap between the appendage of the moving and target image.

Figure 3 shows the skeletal points for the foreground (d) and background (e) of the moving image. The skeletal points of the moving shape are used to predict the location where a potential collapse of the moving image will occur. Note that the designation of foreground and background objects are for convenience of presentation and could have just been referred to as the white and black object of the moving image. Regions of shape collapse only occur in the non-overlap region between the moving and target images. Fig. 3-g and 3-h shows the skeletal points masked by the region of overlap since these are the only skeletal points that we are concerned with.

As discussed in Section 2, we will detect the object boundary points where the region of collapse may happen when the images are registered. The set of boundary points where the shape will collapse can only occur in the region of non-overlap between the moving and target image. The foreground and background boundary points that occur in the symmetric difference region correspond to a super set of the boundary collapse points and are shown in Fig.3-f and 3-i. The foreground and background boundary collapse points are shown in Fig. 3-j and 3-k. The collapse points are defined as the boundary points in the symmetric difference region that correspond to the skeletal points in the symmetric difference region. The predicted foreground and background collapse points are shown superimposed on the moving image in Fig. 3-l. The skeletal points associated with boundary collapse points are a rough estimate of where the moving image will collapse when registered to the target image, as shown in Figure 4.

Fig. 4 shows that the foreground and the background of the moving image collapsed following the predictions of the boundary collapse points. The foreground collapses to a thin strip of white foreground (square callout) and the background collapses to a thin strip of black background (circle callout). Note that the regions of collapse shown in Fig. 4 are small and hard to see. It is because regions of collapse are hard to see that they are often ignored.

Figure 5 shows the result of registering the moving image to the target image for the example in Fig. 3. The image registration algorithm[1] minimized the objective function $C(u)$ which consists of the weighted sum of a similarity cost (sum of squared intensity differences (SSD)) and a regularization term (membrane model), as described by the following equations:

$$C(u) = \frac{\alpha}{2}\int_{\Omega}\left(I_1(x+u(x)) - I_2(x)\right)^2 dx + \frac{\beta}{2}\int_{\Omega}||Lu||^2 dx \quad \text{Registration Cost}$$

$$\delta C(u(x)) = \alpha\left(I_1(x+u(x)) - I_2(x)\right)\cdot \nabla I_1|_{x+u(x)} + \beta L^2 u \quad \text{Gradient of Cost}$$

$$u^{(i+1)}(x) = u^{(i)}(x) - \Delta\delta C(u(x)) \quad \text{Gradient Descent}$$

The gradient descent used a step size of $\Delta = 0.95$, a weight of $\alpha = 1$ for the SSD term and a weight of $\beta = 0.1$ for the membrane regularization term.

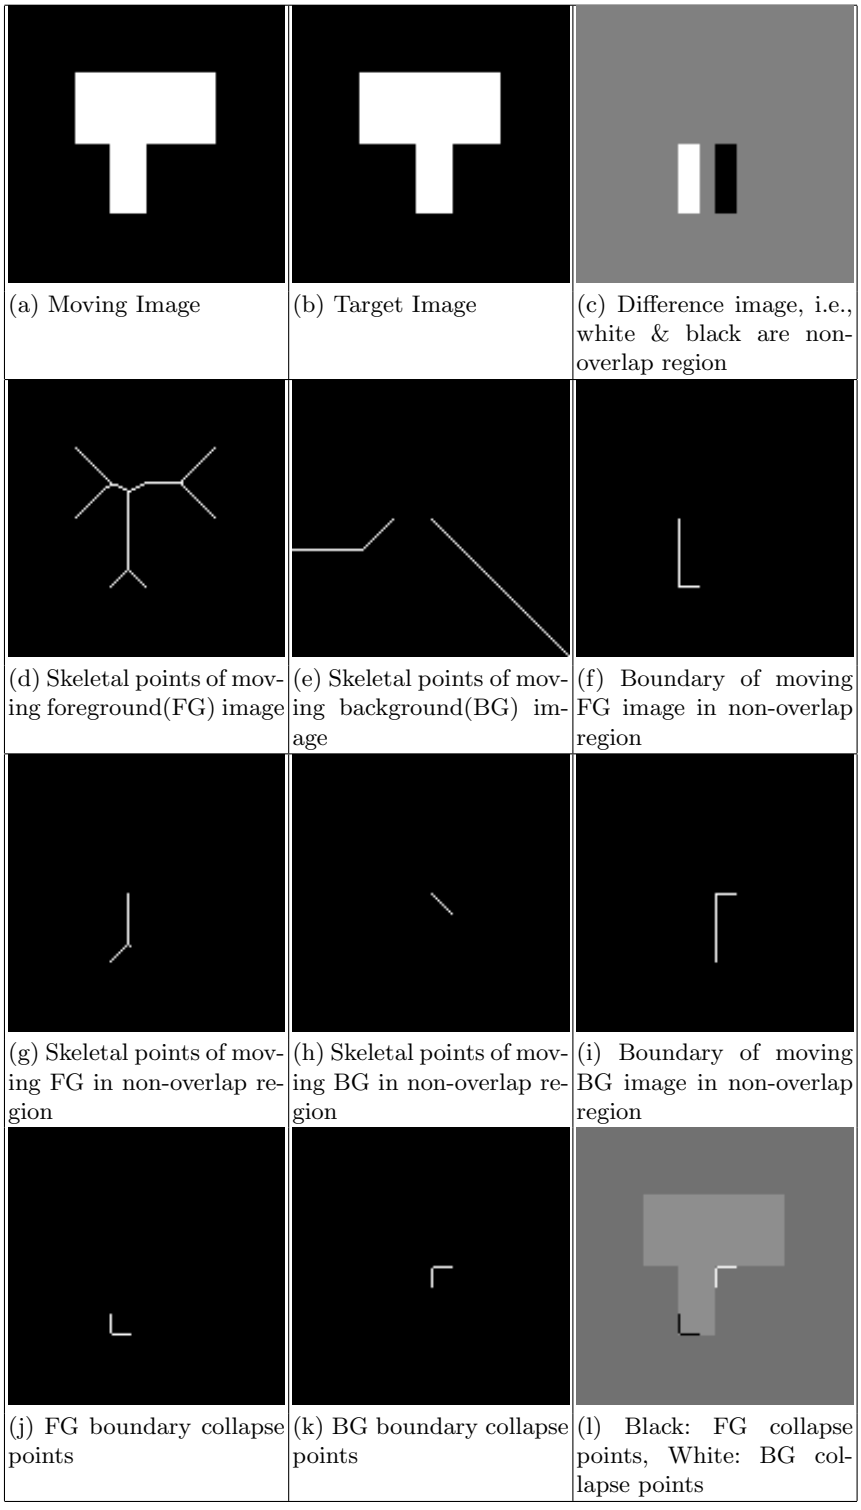

(a) Moving Image | (b) Target Image | (c) Difference image, i.e., white & black are non-overlap region

(d) Skeletal points of moving foreground(FG) image | (e) Skeletal points of moving background(BG) image | (f) Boundary of moving FG image in non-overlap region

(g) Skeletal points of moving FG in non-overlap region | (h) Skeletal points of moving BG in non-overlap region | (i) Boundary of moving BG image in non-overlap region

(j) FG boundary collapse points | (k) BG boundary collapse points | (l) Black: FG collapse points, White: BG collapse points

Fig. 3. Detection of boundary collapse points for overlapping appendage experiment. Both the foreground and background collapses in this example.

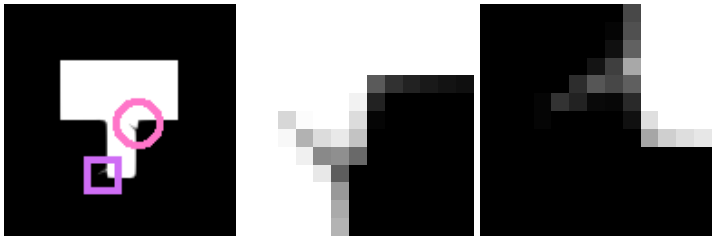

Fig. 4. The result of transforming the moving image into the shape of the target image shown in the top row of Figure 3. The middle panel magnifies the collapse region of the background (circle inset). The right panel magnifies the collapse region of foreground (square inset).

The operator L in the regularization term is a linear differential, self-adjoint operator and corresponds to $L^2u = \nabla^2 u$ for the membrane regularization model and ∇^2 is the Laplacian. For this experiment, the displacement field $u(x)$ was represented as a 128x128 vector field. Bilinear interpolation was used to compute the deformation of the moving image. 1000 iterations of gradient descent were used to minimize the registration cost.

Figure 5 illustrates the time progression of how using a greedy optimization method causes the regions of the moving image to collapse. The arrows in Figure 5-a show the direction of the force generated by the gradient of the SSD cost function. The white arrows show that there is a force pushing right and a force pushing up. The combination of these two forces causes the collapse of the foreground object at the lower left corner of the object. The black arrows show a force pushing right and another pushing down. The result of these two forces causes the background of moving image to collapse at the inside corner of the object. Note that once a region of the image collapses, there is no way for a greedy optimization method to "un-collapse" the region. The reason for this is that the collapsed region has zero measure in the SSD cost function and therefore has little to no cost. Thus, there is no incentive (i.e., no gradient force) for the greedy optimization method to retract or otherwise fix the collapsed region.

Figure 6 shows the x- and y-displacement fields associated with this registration experiment. Notice that the x-displacement field is not symmetric in the vertical direction, i.e., the left black region appears to be shifted up compared to the black region on the right. The vertical asymmetry of the x-displacement field results from the shape collapse of the moving image. The asymmetry of the x-displacement field corresponds to the nonzero regions of the y-displacement field in the regions of image collapse. The nonzero regions of the y-displacement field indicate a problem with the estimated transformation for this experiment. It is reasonable to assume that the vertical appendage of the moving image only needs to be shifted to the right to match the corresponding structure in the target image. Thus, any nonzero displacement in the y-direction indicates a counter-intuitive registration result.

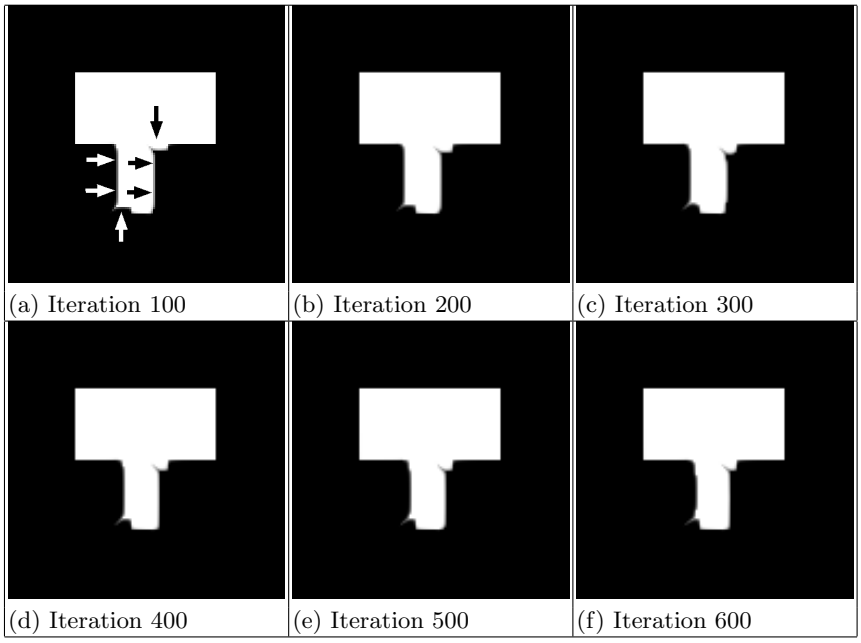

Fig. 5. Progression of deformation from moving to target image. The arrows show in panel (a) the direction and location of the deformation. The white arrows show where the foreground of the moving image collapses and the black arrows show where the background of the moving image collapses.

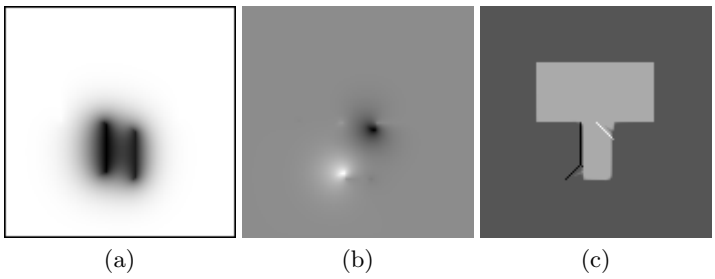

Fig. 6. Displacement fields for the image registration of overlapping appendage experiment. (a) X-displacement, (b) Y-displacement, (c) foreground and background skeletal points overlaid on deformed image predict regions of collapse

The displacement fields shown in Figure 6 demonstrate that the regions of collapse should not be ignored because they lead to poor correspondence in the regions of collapse. Ignoring the poor correspondence in collapsed regions may have a major impact on the conclusions drawn in the applications using such transformations. These include mapping brain function, computing average

shapes, computing shape statistics, computing mechanical properties of the lung, adaptive radiotherapy planning, and computing cumulative dose in radiation therapy to mention a few application areas.

Figure 6-c shows the skeletal points of the foreground and background objects from the original moving image within the symmetric difference region, superimposed on the registration result. Notice that the skeletal points do a good job of predicting the locations of the moving image collapse. The skeletal points do not predict the actual locations of the moving image collapse due to the shape change of the moving image during registration and the effect of regularization.

The minimum value of the Jacobian for the estimated transformation for this experiment was 0.0853. The fact that the minimum value of the Jacobian was positive indicates that the estimated transformation remained a diffeomorphism throughout the optimization procedure. This fact illustrates that the collapse problem occurs even when the transformation is a diffeomorphism.

Figure 7 shows three additional examples of predicting collapse points. The first row of this figure illustrates collapse points occurring when a subregion of an object retracts. The second row illustrates collapse points occurring when a subregion of an object expands. The final row shows the cortical ribbon extracted from a clinical MR image with $1 \times 1\text{mm}^2$ resolution. The cortical segmentation was adapted to limit the example to illustrate four retraction and two expansion regions around sulci. This example is important to show the scale of the expected effect size of the collapsing problem for neuroimaging applications.

4 Conclusions

This paper presented the shape collapse problem for image registration which is a problem for both small and large deformation image registration algorithms. A mathematical justification was presented for why the collapsing problem occurs for binary images that each contain a single corresponding object. Collapsing boundary points were defined mathematically and were shown to predict before the registration, which boundary points in the moving image will collapse during an iterative image registration procedure that minimize the registration cost at each iteration. Furthermore, the foreground and background skeletons of the moving image in the symmetric difference region were shown to closely predict the collapse loci in the moving image.

References

1. Christensen, G.E.: Deformable Shape Models for Anatomy. Ph.D. thesis, Washington University, St. Louis, MO. (Aug 1994)
2. Do Carmo, M.: Riemannian Geometry. Birkhauser, Boston (1992)
3. Siddiqi, K., Pizer, S.: Medial Representations: Mathematics, Algorithms and Applications. Springer Publishers (2008)
4. Sonka, M., Hlavac, V., Boyle, R.: Image Processing, Analysis, and Machine Vision. Thomson Engineering, Toronto, Canada, 3rd edn. (2008)

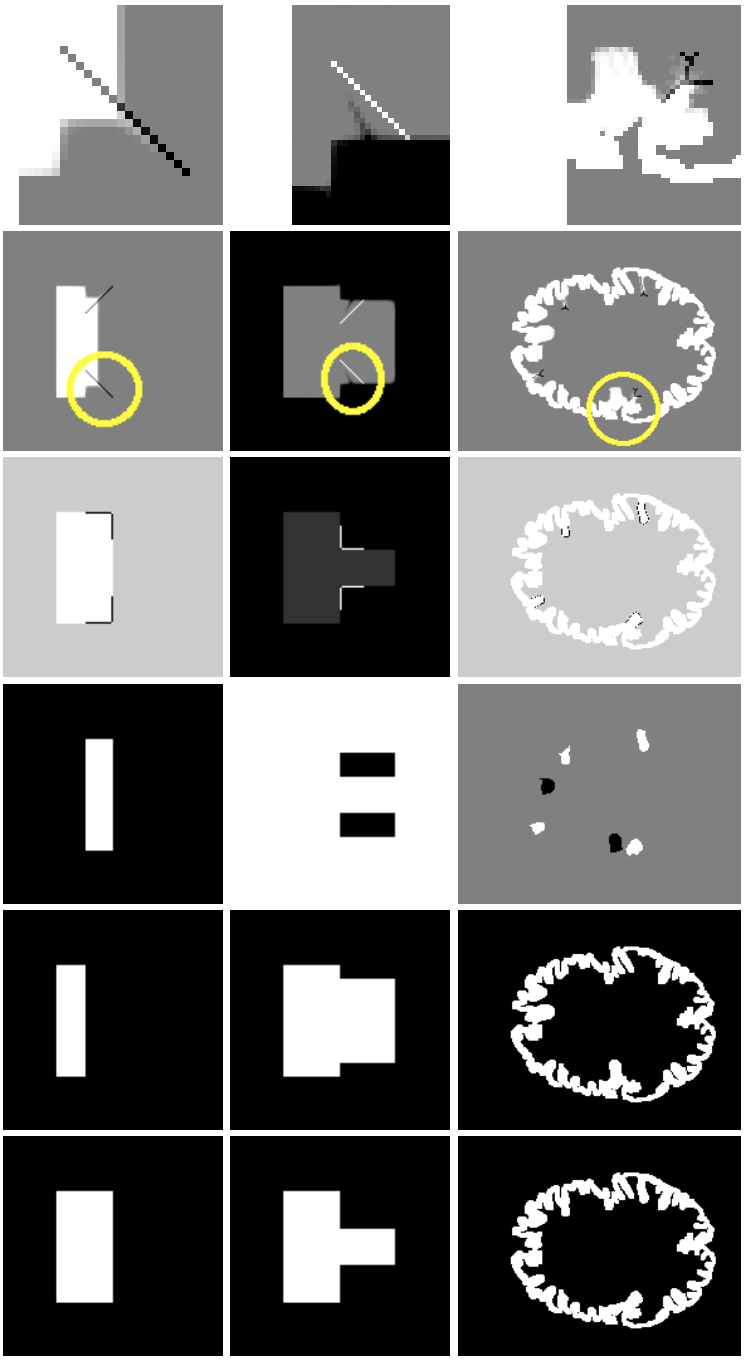

Fig. 7. Examples. Each row of this figure shows an additional example of collapsing point detection. The columns from left to right show the moving image, target image, difference image before registration, deformed collapse points on the moving image, and the foreground and background skeletal points overlaid on the deformed image.