

Label Space: A Multi-object Shape Representation

James Malcolm¹, Yogesh Rathi², and Allen Tannenbaum¹

¹ Georgia Institute of Technology, Atlanta, GA
{malcolm, tannenba}@ece.gatech.edu

² Brigham and Women's Hospital, Boston, MA
yogesh@bwh.harvard.edu

Abstract. Two key aspects of coupled multi-object shape analysis are the choice of representation and subsequent registration to align the sample set. Current techniques for such analysis tend to trade off performance between the two tasks, performing well for one task but developing problems when used for the other.

This article proposes \mathcal{L}^n *label space*, a representation that is both flexible and well suited for both tasks. We propose to map object labels to vertices of a regular simplex, *e.g.* the unit interval for two labels, a triangle for three labels, a tetrahedron for four labels, etc. This forms a linear space with the property that all labels are equally separated.

On examination, this representation has several desirable properties: algebraic operations may be done directly, label uncertainty is expressed as a weighted mixture of labels, interpolation is unbiased toward any label or the background, and registration may be performed directly.

To demonstrate these properties, we describe variational registration directly in this space. Many registration methods fix one of the maps and align the rest of the set to this fixed map. To remove the bias induced by arbitrary selection of the fixed map, we align a set of label maps to their intrinsic mean map.

1 Introduction

Multi-object shape analysis is an important task in the medical imaging community. When studying the neuroanatomy of patients, clinical researchers often develop statistical models of important structures which are then useful for population studies or as segmentation priors [7,9,10,11,12]. The first step for this problem consists in choosing an appropriate shape descriptor capable of representing its statistical variability.

A common starting point for shape representation is a simple scalar label map, each pixel indicating the object present at that pixel, *e.g.* a one indicating object #1, a two indicating object #2, etc. Many techniques go on to map this entire volume to another space, the value of each pixel contributing to describe the shape. In this new space, arbitrary topologies may be represented, correspondences are naturally formed between pixels, and there are no control points to distribute.

The simplest implicit representation is a binary map where each pixel indicates the presence or absence of the object. Signed distance maps (SDM's) are another example of an implicit representation, each pixel having the distance to the nearest object boundary, a negative distance for points inside the object [7,12].




Fig. 1. Tsai et al. [11] proposed mapping each pixel from object label to a point in a space shaped as a non-regular simplex, each vertex corresponding to an object label. Visualized here for the case of two objects and background, the bottom left background (0,0) is a distance of 1 from both labels top (0,1) and right (1,0), while labels are separated from each other by a distance of $\sqrt{2}$.




Fig. 2. Example configurations for the S^1 hypersphere representation of [2]: three, six, and seven labels (left to right) with background at the center

For the multi-object setting, binary maps may be extended to scalar label maps, each pixel holding a scalar value corresponding to the presence of a particular object; however, this representation is not well suited for algebraic manipulation. For example, if labels are left as scalar values, the arithmetic average of labels with values #1 and #3 would incorrectly indicate the label of value #2, not a mixture of labels #1 and #3.

To address this, mappings of object labels to linear vector spaces were proposed, an approach to which our method is most closely related. The work of Tsai et al. [11] introduced two such representations, each for a particular task. For registration, the authors proposed mapping scalar labels to binary vectors with entries corresponding to labels; a one in an entry indicates the presence of the corresponding label at that pixel location. As an example for the case of two labels and background, Figure 1 visualizes the spatial configuration each pixel is mapped onto. Here the background is at the bottom left origin (0,0) with one label at (1,0) and the other at (0,1). It is also important to note that he goes on to perform registration considering each entry of these vectors separately. For shape analysis, Tsai et al. [11] proposed mapping scalar labels to layered SDM's, in this case each layer giving the signed distance to the corresponding object's interface.

Note that in both vector valued representations described in Tsai et al. [11], each label lies on its own axis and so the dimension of the representation grows linearly with the number of labels, *e.g.* two objects require two dimensions, three objects require three dimensions. To address this spatial complexity, Babalola and Cootes [2,3] propose a lower dimension approximation to replace the binary vectors in registration. By mapping labels to the unit hypersphere S^n , they demonstrate that even configurations involving dozens of labels can be efficiently represented with label locations distributed uniformly on a hypersphere. Figure 2 gives examples for S^1 .

Finally, Pohl et al. [10] indirectly embeds label maps in the logarithm-of-odds space using as intermediate mappings either the binary or SDM representations of [11].




Fig. 3. The first three \mathcal{L}^n label space configurations: a unit interval \mathcal{L}^2 in \mathbb{R} for two labels, a triangle \mathcal{L}^3 in \mathbb{R}^2 for three labels, and a tetrahedron \mathcal{L}^4 in \mathbb{R}^3 for four labels (*left to right*)

Particularly well suited for probabilistic computations, the logarithm-of-odds space is also a field providing closed operations for addition and scalar multiplication. As with the representations of Tsai et al. [11], the dimensionality of the logarithm-of-odds space increases with each additional object. We should also note that the work of [10] did not address registration, but instead assumed an already registered atlas via [8].

Once the representation is settled upon, registration must be performed to eliminate variation due to differences in pose. A common approach is to register the set to a reference image; however, this then introduces a bias to the shape of the chosen reference. Joshi et al. [6] propose unbiased registration with respect the mean sample as a template reference. Assuming a general metric space of transformations, they describe registering a sample set with respect to its intrinsic mean and use the L_2 distance for demonstration. A similar approach uses the minimum description length to measure distance from the intrinsic mean [13]. Instead of registering to a mean template, an alternative approach is to minimize per-pixel entropy. Using binary maps Miller et al. [8] demonstrate that this has a similar tendency toward the mean sample. This approach has also been demonstrated on intensity images [14,15]. Among these energy-based registration techniques, iterative solutions include those that are variational [11,6] and those that use sampling techniques [15].

1.1 Our Contributions

This paper proposes a multi-object implicit representation that maps object labels to the vertices of a regular simplex, going from a scalar label value to a vertex coordinate position in a high dimensional space which we term *label space* and denote by \mathcal{L}^n for n labels. Visualized in Figure 3, this regular simplex is a hyper-dimensional analogue of an equilateral triangle, n vertices capable of being represented in $n - 1$ dimensions ($\mathcal{L}^n \subset \mathbb{R}^{n-1}$). Lying in a linear vector space, this space has several desirable properties: all labels are equally separated in space, addition and scalar multiplication are natural, label uncertainty is expressed as a weighted combination of label vertices, and interpolation is unbiased toward any label including the background.

The proposed method addresses several problems with current implicit mappings. For example, while the binary vector representation of Tsai et al. [11] was proposed for registration, we will demonstrate that it induces a bias sometimes leading to misalignment, and since our \mathcal{L}^n label space representation equally spaces labels, there is no such bias. Additionally, compared to the SDM representation, the proposed method introduces no inherent per-pixel variation across equally labeled regions making it more robust for statistical analysis. Hence, the proposed method better encapsulates the functionality of both representations. Further, the registration energy of Tsai et al. [11] is designed to consider each label independent of the others. In contrast, \mathcal{L}^n label space




Fig. 4. For the S^1 hypersphere configurations of [2], cases such as these yield erroneous results during interpolation. Judged by nearest neighbor, interpolating between two labels resolves to background, ambiguously either background or another label, and finally another label (*left to right*).

jointly considers all labels. We will also demonstrate that, while lowering the spatial demands of the mapping, the hypersphere representation of Babalola and Cootes [2] biases interpolation and can easily lead to erroneous results. The arrangement of our proposed label space incurs no such bias allowing linear combinations of arbitrary labels.

The rest of this paper is organized as follows. Section 2 explores several problems that can develop with the implicit representations described above [2,10,11]. Section 3 then describes the proposed \mathcal{L}^n label space representation documenting several of its properties. Section 4 demonstrates variational registration directly within this representation, and finally in Section 5 we summarize our work.

2 Related Work

2.1 Shape Representation

The signed distance map (SDM) has been used as a representation in several studies [1,7,10,11,12]; however, it may produce artifacts during statistical analysis [4]. For example, small deviations at the interface cause large variations in the surface far away, thus it inherently contains significant per-pixel variation. Additionally, ambiguities arise when using layered signed distance function to represent multiple objects: what happens if more than one of the distance functions indicates the presence of an object? Such ambiguities and distortions stem from the fact SDM's lie in a manifold where these linear operations introduce artifacts [4,5].

Label maps have inherently little per-pixel variation, pixels far from the interface having the same label as those just off the interface. For statistical analysis in the case of one object, Dambreville et al. [4] demonstrated that binary label maps have higher fidelity compared to SDM's. However, for the multi-object setting, the question then becomes one of how to represent multiple shapes using binary maps? What is needed is a richer feature space suitable for a uniform pair-wise separation of labels.

An example of such a richer feature space is that of Babalola and Cootes [2] where labels are mapped to points on the surface of a unit hypersphere S^n placing the background at the center. This is similar to the binary vector representation described by Tsai et al. [11] to spread labels out; however, Babalola and Cootes [2] argue that lower dimensional approximations can be made. They demonstrate that configurations involving dozens of labels can be efficiently represented by distributing label locations

uniformly on the unit hypersphere using as few as three dimensions. Since any label may neighbor the background, the background must be placed at the hypersphere center, equally spaced from all other labels. The fundamental assumption is that pixels only vary between labels that are located near to each other on the hypersphere, so the placement of labels is crucial to avoid erroneous label mixtures. For example, Figure 4 demonstrates that if two labels far from each other are mixed, the result may be attributed erroneously to other labels. Notice in particular that the central placement of the background gets in the way when interpolating across the sphere. Smoothing in Figure 7 also demonstrates these inherent effects of the lower dimensional approximation, effects that cannot be avoided unless the dimension approaches label cardinality.

The logarithm-of-odds representation of Pohl et al. [10] provides the third and final shape representation we compare against. Aside from the normalization requirement for closed algebraic manipulation, the main concern when using this representation is the choice of intermediate mapping, a choice that directly impacts the resulting probabilities. The authors explore the use of both representations from [11]; however, both choices have inherent drawbacks.

For the layered SDM intermediate mapping, Pohl et al. [10] notes that SDM's are a subspace of the logarithm-of-odds space. This means that, while the layered SDM's are exactly the logarithm-of-odds representation, results after algebraic manipulation in the logarithm-of-odds space often yield invalid SDM's (but still valid logarithm-of-odds representations). Using such results, computing probabilities as described in [10] may yield erroneous likelihoods. Notice also, that the generalized logistic function is used to compute probabilities. This introduces additional problems as the use of the exponential ensures that these probabilities will always have substantial nonzero character across the entire domain, even in areas never indicated by the sample set.

Using smoothed binary maps as intermediates also leads to problems. To begin, using binary maps directly would mean probabilities of either zero or one, which in the log domain produce singularities. Smoothing lessens such effects yet results in a loss of fine detail along the interface. Also, Pohl et al. [10] shows examples where after normalization the logarithm-of-odds representation develops artifacts at the interface between objects, an effect which is magnified in the logarithm domain.

2.2 Registration

Tsai et al. [11] propose a binary vector representation specifically for registration. As Figure 1 shows, this representation places labels at the corners of a right-triangular simplex; however, unlike this present work, it is not a regular simplex but has a bias with respect to the background. The background, located at the origin, is a unit distance from any other label, while any two labels, located along a positive axis, are separated by a distance of $\sqrt{2}$. The effect may be seen in registration where there is a bias to misalign labels over the background (penalty 1) rather than over other labels (penalty $\sqrt{2}$).

To demonstrate the effect of this induced bias, consider the example in Figure 5 with black background and two rectangles of label #1, one with strip of label #2 along its top. Using the representation and registration energy of Tsai et al. [11], there are two global minima: the image overlapping and the image shifted up. In the first case, label

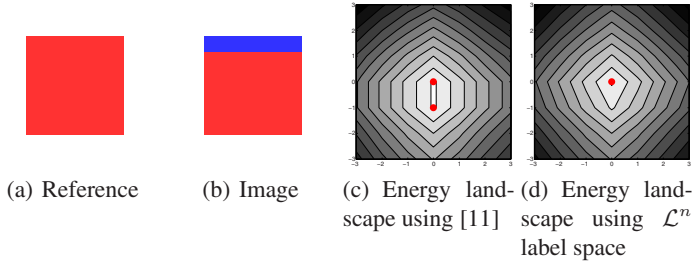


Fig. 5. Alignment of an image with a reference template using the representation of Tsai et al. [11] results in two possible alignments, the shifted one misaligning along both the top and bottom with respect to the reference (*red dots indicate minima*). For just x- and y-translation, isocontours of the energy landscape show the non-unique energy minima in (c).

#1 is misaligned over label #2, while in the second case that a strip of pixels at both the top and bottom are misaligned over the background; that is, because of this bias, there can be twice as many pixels misaligned in the shifted case than in the unshifted. These global minima (indicated by red dots in the energy landscapes) are shown only for translation; considering additional pose parameters further increases the number of local minima in the energy landscape representing misalignments. Also, this is not inherent in the energy, as the same phenomena is observed using the energy in (1). Since all labels are equidistant in the proposed representation, there are fewer minima and hence less chance of misalignment.

3 Label Space

Our goal is to create a robust representation where algebraic operations are natural, label uncertainty is captured, and interpolation is unbiased toward any label. To this end we propose mapping each label to a vertex of a regular simplex; given n labels, including the background, we use a regular simplex which lies in $n - 1$ dimensions and denote this by \mathcal{L}^n (see Figure 3). A regular simplex is an n -dimensional analogue of an equilateral triangle.

In this space, algebraic operations are as natural as vector addition, scalar multiplication, inner products, and norms; hence, there is no need for normalization as in [10]. Label uncertainty is realized as the weighted mixture of vertices. For example, a pixel representing labels #1, #2, and #3 with equal characteristic would simply be the point $p = \frac{1}{3}v_1 + \frac{1}{3}v_2 + \frac{1}{3}v_3$, a point equidistant from those three vertices (see Figure 6). Also, we have that such algebraic operations are unbiased toward any label since all labels



Fig. 6. Proposed \mathcal{L}^3 label space for the case of three labels: a point indicating the equal presence of all three labels (*left*), and a point indicating the unequal mixed presence of just the left and top labels (*right*)

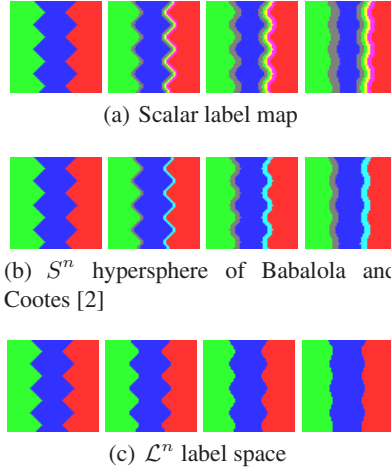


Fig. 7. Progressive smoothing directly on scalar label maps, the hypersphere representation of Babalola and Cootes [2], and \mathcal{L}^n label space. Both the scalar label maps and hypersphere representations develop intervening strips of erroneous labels. Only label space is able to correctly capture the label mixtures during smoothing. The rightmost hypersphere in Figure 4 depicts the S^1 configuration used here in (b).




Fig. 8. Progressive smoothing directly on binary vector representation of Tsai et al. [11] and \mathcal{L}^n label space. Smoothing among several labels in the binary vector representation yields points closer to background (*black*) than any of the original labels. Label space is able to correctly begin to smooth out the sharp corners of the bottom two regions without erroneous introduction of the black background label.

are equally spaced; hence, there is no bias with respect to the background as is found in both [2,11]. Label space is robust to statistical analysis much like binary label maps, a specific case of label space. Additionally, problems encountered in the intermediate representations of [10] are avoided. Specifically, smoothing is unnecessary and so fine detail is retained, and interfaces are correctly maintained.

To demonstrate some of these properties, we performed progressive smoothing using the various representations described: scalar label values, the binary vector representation of Tsai et al. [11], the S^n representation of Babalola et al. [2], and \mathcal{L}^n label space.

In Figure 7, the first experiment has each example beginning on the left with the jagged stripes of labels #5, #7, and #3, respectively. Scalar label values show the appearance of intervening labels #4, #5, and #6 as the original labels blend, and the hypersphere representation shows the appearance of labels #2, #6, and #4 as interpolation is performed across the hypersphere (the hypersphere configuration used here is the rightmost depicted in Figure 4). In Figure 8, the second experiment shows that the smoothing among multiple labels using binary vectors produces points closest to the background (black). In both experiments, only label space correctly preserves the interfaces.

4 Registering to the Mean Map

We demonstrate here the variational registration of a set of maps to their intrinsic mean map, thereby respecting the first order statistics of the sample set. The proposed representation has the advantage of supporting registration directly on the representation. By directly we mean that differentiable vector norms may be used to compare labels.

In this section, we begin with a review of reference-based approaches for rigid registration borrowing the notation of [11]. After demonstrating how a bias can be induced by the choice of reference template, we demonstrate unbiased registration using the mean map as the reference template in the manner of [6]. We conclude with experiments on synthetic maps, the 2D slices from [11] with three labels, and 2D slices with eight labels.

Common approaches to registration begin by fixing one of the maps as a reference and registering the remaining maps to this fixed map. This is done in both [2,11]; however, as Joshi et al. [6] describes, this initial choice biases the spatial statistics of the aligned maps. In Figure 9 we see this effect: as the choice of fixed map is varied, the resulting atlas varies in translation, scale, rotation, and skew (registration was performed as in [11]). To avoid this bias, Joshi et al. [6] describe registration with respect to a reference that best represents the sample set. In addition to avoiding bias, the resulting gradient descent involves far less computation than that proposed in [11] where each map is compared against each other map. Also, since the reference image is a convex combination of the set, there is no fear of the set \tilde{M} shrinking to minimize the energy.

Before presenting the energy used, we first describe the problem borrowing notation from [11]. For the set of label maps $M = \{m_i\}_{i=1}^N$, our goal is to estimate the set of corresponding pose parameters $P = \{\mathbf{p}_i\}_{i=1}^N$ for optimal alignment. We denote as \tilde{m} the label map m transformed by its pose parameters. An advantage of implicit representations over explicit ones is that, once the label maps have undergone




Fig. 9. Label maps from patient MRI data after registration where a different label map has been fixed in each run. The choice of which map to fix can subtly distort measurements and hence the statistical model constructed from the registered set.

this transformation, we can assume direct per-pixel correspondence between maps and use a vector norm to perform comparison. We model pose using an affine model, and so for 2D, the pose parameter is the vector $\mathbf{p} = [x \ y \ s_x \ s_y \ \theta \ k]^T$ corresponding to x-,y- translation, x-,y-scale, in-plane rotation, and shear. Note that this is a fully affine model as compared to the rigid transformation model used in [11]. The transformed map is defined as $\tilde{m}(\tilde{x}, \tilde{y}) = m(x, y)$ where coordinates are mapped according to $[\tilde{x} \ \tilde{y} \ 1]^T = T(\mathbf{p}) [x \ y \ 1]^T$, where $T(\mathbf{p})$ is the decomposable transformation matrix

$$T(\mathbf{p}) = \underbrace{\begin{bmatrix} 1 & 0 & x \\ 0 & 1 & y \\ 0 & 0 & 1 \end{bmatrix}}_{M(x,y)} \underbrace{\begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{R(\theta)} \underbrace{\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{H(s_x, s_y)} \underbrace{\begin{bmatrix} 1 & k & 0 \\ k & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{K(k)}$$

for a translation matrix $M(x, y)$, rotation matrix $R(\theta)$, anisotropic scale matrix $H(s_x, s_y)$, and shear matrix $K(k)$, all for the parameters taken from \mathbf{p} .

As in [6,15], we assume the intrinsic mean map $\tilde{\mu}$ of the sample set to best represent the population. We then attempt to minimize the energy defined as the squared distance between each transformed label map \tilde{m} and this mean map $\tilde{\mu}$ of the set \tilde{M} as it converges:

$$d^2 = \sum_{i=1}^N \|\tilde{m}_i - \tilde{\mu}\|^2, \quad (1)$$

where $\tilde{\mu} = \frac{1}{N} \sum_{i=1}^N \tilde{m}_i$, and while $\|\cdot\|$ may be any differentiable norm, we take it to be the elemental L_2 inner product $\|x\| = \langle x, x \rangle^{1/2} = \int x^2 dx$. Notice how using a vector norm here jointly considers all labels in contrast to the energy proposed by Tsai et al. [11]. Further, since the reference map $\tilde{\mu}$ is intrinsic, there is no concern of the set \tilde{M} shrinking to minimize (1). Hence, there is no need for the normalizing term introduced in [11] which allows for a reduced complexity energy here.

This work uses a variational approach to registration. Specifically we perform gradient descent to solve for the pose parameters minimizing this distance. We find the gradient of this distance, taken with respect to the pose \mathbf{p}_j , to be:

$$\nabla_{\mathbf{p}_j} d^2 = 2 \langle \nabla_{\mathbf{p}_j} \tilde{m}_j, \tilde{m}_j - \tilde{\mu} \rangle. \quad (2)$$

Notice that terms involving other label maps (\tilde{m}_i for $i \neq j$) fall out and that the gradient of the mean contributes nothing. It remains to define $\nabla_{\mathbf{p}_j} \tilde{m}_j$. For the k^{th} element of the pose parameter vector \mathbf{p}_j , using the chain rule produces $\nabla_{\mathbf{p}_j^k} \tilde{m}_j =$

$\left[\frac{\partial \tilde{m}_j}{\partial \tilde{x}} \ \frac{\partial \tilde{m}_j}{\partial \tilde{y}} \ 0 \right] \frac{\partial T(\mathbf{p}_j)}{\partial \mathbf{p}_j^k} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$, where $\frac{\partial T(\mathbf{p}_j)}{\partial \mathbf{p}_j^k}$ is computed for each pose parameter where matrix derivatives are taken componentwise. Finally, gradient descent proceeds by repeated calculation of $\nabla_{\mathbf{p}_j} d^2$ and adjustment of \mathbf{p}_j for each map in the set until convergence.

To illustrate this technique, we first performed alignment of a synthetic 2D set. The unaligned set consists of 15 maps of three labels and background. Figure 10 shows




(a) Example maps from unaligned set



(b) Original

(c) Aligned

Fig. 10. Alignment of a set of 15 synthetic maps with three labels and background. The original and aligned sets are superimposed for visualization.



(a) Perturbed original

(b) Recovered

Fig. 11. From the dataset used by Tsai et al. [11], one map is chosen and perturbed under several transformations, yet registration is able to recover the pose parameters to bring the perturbed versions back to the original chosen map. The perturbations ranged up to translations of 5% of the image, rotational differences of 20° , and scale changes $\pm 5\%$ of the image. The original and aligned sets are superimposed for visualization.

examples from this set as well as the original and aligned sets. For visualization, we created a superimposed map for both the original unaligned set and the aligned set by summing the scalar label values pixelwise and dividing by the number of maps, hence this is the mean scalar map.

We then turned to verifying our method using the 2D data from the study by Tsai et al. [11]. Taking one map from this set, we formed a new set by transforming this map arbitrarily. Restricting ourselves to the rigid rotation pose model used in that study, we formed transformations involving translations of 5% of the image size, rotational differences of 20° , and scale changes of $\pm 5\%$ of the image. Figure 11 shows that the technique successfully recovered the initial map. Figure 12 shows alignment on the entire data set.

Lastly, we performed registration using 2D maps obtained from expert manual segmentation of 33 patient MRI scans involving eight labels and background. Figure 13 shows examples from the original unaligned set as well as the superimposed maps after alignment.




Fig. 12. Alignment of a set of 30 maps used in the study by Tsai et al. [11]. The original and aligned sets are superimposed for visualization.

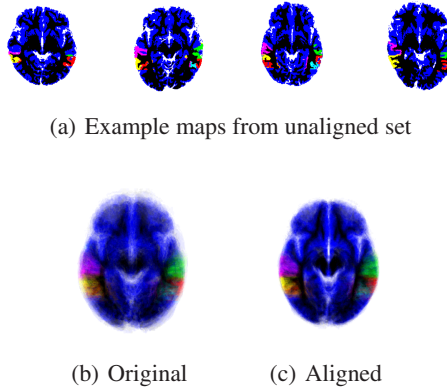


Fig. 13. Alignment of a set of 33 maps with eight labels and background obtained from manual MRI segmentations. The original and aligned sets are superimposed for visualization.

5 Conclusion

This paper describes a new implicit multi-object shape representation. After detailing several drawbacks to current representations, we demonstrated several of its properties. In particular, we demonstrated that algebraic operations may be done directly, label uncertainty is expressed naturally as a mixture of labels, interpolation is unbiased toward any label or the background, and registration may be performed directly.

Modeling shapes in label space does have its limitations. One key drawback to label space is the spatial demand. To address this we are examining lower dimensional approximations much like Babalola and Cootes [2]. Some interpolation issues such as those noted in Figure 4 might be avoided by taking into consideration the empirical presence of neighbor pairings when determining label distribution.

Acknowledgements

This work was supported in part by grants from NSF, AFOSR, ARO, MURI, as well as by a grant from NIH (NAC P41 RR-13218) through Brigham and Women's Hospital. This work is part of the National Alliance for Medical Image Computing

(NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>

References

1. Abd, H., Farag, A.: Shape representation and registration using vector distance functions. In: Computer Vision and Pattern Recognition (2007)
2. Babalola, K., Cootes, T.: Groupwise registration of richly labelled images. In: Medical Image Analysis and Understanding (2006)
3. Babalola, K., Cootes, T.: Registering richly labelled 3d images. In: Proc. of the Int. Symp. on Biomedical Images (2006)
4. Dambreville, S., Rath, Y., Tannenbaum, A.: A shape-based approach to robust image segmentation. In: Int. Conf. on Image Analysis and Recognition (2006)
5. Golland, P., Grimson, W., Shenton, M., Kikinis, R.: Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis* 9, 69–86 (2005)
6. Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23, 150–161 (2004)
7. Leventon, M., Grimson, E., Faugeras, O.: Statistical shape influence in geodesic active contours. In: Computer Vision and Pattern Recognition, pp. 1316–1324 (2000)
8. Miller, E., Matsakis, N., Viola, P.: Learning from one example through shared densities on transforms. In: Computer Vision and Pattern Recognition, pp. 464–471 (2000)
9. Nain, D., Haker, S., Bobick, A., Tannenbaum, A.: Multiscale 3-d shape representation and segmentation using spherical wavelets. *Trans. on Medical Imaging* 26(4), 598–618 (2007)
10. Pohl, K., Fisher, J., Bouix, S., Shenton, M., McCarley, R., Grimson, W., Kikinis, R., Wells, W.: Using the logarithm of odds to define a vector space on probabilistic atlases. *Medical Image Analysis* (to appear, 2007)
11. Tsai, A., Wells, W., Tempany, C., Grimson, E., Willsky, A.: Mutual information in coupled multi-shape model for medical image segmentation. *Medical Image Analysis* 8(4), 429–445 (2003)
12. Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. *Trans. on Medical Imaging* 22(2), 137–154 (2003)
13. Twining, C., Marsland, C., Taylor, S.: Groupwise non-rigid registration: The minimum description length approach. In: British Machine Vision Conf. (2004)
14. Warfield, S., Rexillius, J., Huppi, R., Inder, T., Miller, E., Wells, W., Zientara, G., Jolesz, F., Kikinis, R.: A binary entropy measure to assess nonrigid registration algorithms. In: Niessen, W.J., Viergever, M.A. (eds.) MICCAI 2001. LNCS, vol. 2208, pp. 266–274. Springer, Heidelberg (2001)
15. Zöllei, L., Learned-Miller, E., Grimson, E., Wells, W.: Efficient population registration of 3d data. In: Workshop on Comp. Vision for Biomedical Image Applications (ICCV) (2005)