

The Generalized Log-Ratio Transformation: Learning Shape and Adjacency Priors for Simultaneous Thigh Muscle Segmentation

Shawn Andrews and Ghassan Hamarneh

Abstract—We present a novel probabilistic shape representation that implicitly includes prior anatomical volume and adjacency information, termed the generalized log-ratio (GLR) representation. We demonstrate the usefulness of this representation in the task of thigh muscle segmentation. Analysis of the shapes and sizes of thigh muscles can lead to a better understanding of the effects of chronic obstructive pulmonary disease (COPD), which often results in skeletal muscle weakness in lower limbs. However, segmenting these muscles from one another is difficult due to a lack of distinctive features and inter-muscular boundaries that are difficult to detect. We overcome these difficulties by building a shape model in the space of GLR representations. We remove pose variability from the model by employing a presegmentation-based alignment scheme. We also design a rotationally invariant random forest boundary detector that learns common appearances of the interface between muscles from training data. We combine the shape model and the boundary detector into a fully automatic globally optimal segmentation technique. Our segmentation technique produces a probabilistic segmentation that can be used to generate uncertainty information, which can be used to aid subsequent analysis. Our experiments on challenging 3D magnetic resonance imaging data sets show that the use of the GLR representation improves the segmentation accuracy, and yields an average Dice similarity coefficient of 0.808 ± 0.074 , comparable to other state-of-the-art thigh segmentation techniques.

Index Terms—Statistical Shape Analysis, Muscle Segmentation, Probabilistic Segmentation, Edge Detection, Uncertainty, COPD

I. INTRODUCTION

An important precursor to any anatomical volume measurement or shape analysis from a medical image is segmentation. Manual segmentation in 3D medical images is extremely time consuming, tedious, and suffers from inter- and intra-operator variability. Highly automated segmentation is important for studies involving a large cohort of subjects to reduce manual labour and variability, and to improve the efficiency in analyzing large groups of data. When manually segmenting an image, a clinician often must rely on their prior knowledge of anatomy in order to distinguish different structures; automated segmentation methods must somehow encode similar anatomical information to achieve adequate accuracy. For example, Yazdanpanah et al. and Garvin et al. [1], [2], encoded spatial relationships between retinal layers

via level sets with concentric circle shape priors. Nosrati et al. and Delong and Boykov [3], [4] encoded containment and exclusion constraints in level sets and graph cut segmentation frameworks, respectively, and applied their methods to cardiac, bone, microscopy, and other segmentation tasks.

In this paper, we propose training statistical shape models using a novel segmentation representation that implicitly includes prior anatomical volume and adjacency information, termed the *generalized log-ratio* (GLR) representation. The GLR representation is designed to incorporate statistically meaningful localized uncertainty information through the use of probabilistic labels [5]. Encoding uncertainty in automated segmentations aids in subsequent user analysis and identification of segmentation errors [6], [7]. We demonstrate the efficacy of the GLR representation by applying it to task of knee extensor and flexor (thigh) muscle segmentation on a data set including patients with chronic obstructive pulmonary disease (COPD).

In patients with COPD, skeletal muscle weakness is common [8]–[12]. Thigh muscles are often the most affected due to lack of use [11]. Reduced muscle mass is among several factors that lead to reduced force production, also including changes in the muscle contractile apparatus and neuromuscular activation [8]. The relative effects of these factors are an ongoing area of study [12]. Thus, size and shape measurements are required to study the contribution of muscle mass reduction to force loss. Recently [13], a non-uniform distribution of atrophy and size changes was found across knee extensors and flexors in patients with COPD, which may be reflective of localized factors such as denervation, limited recruitment, or atrophy of specific muscle fibers, rather than systemic factors contributing to muscle atrophy. Further, classifiers trained on 3D shape descriptors of thigh muscles have been able to detect COPD [14]. Information regarding the specific muscles or muscle regions that are most atrophied together with functional assessment will enable therapeutic interventions to be targeted to the affected regions [13]–[15] rather than the prescription of a generalized approach that may prove ineffectual.

Magnetic resonance imaging (MRI) can be used to generate volumetric images, from which muscle can be distinguished from the surrounding regions and volume and shape properties of individual muscles can be estimated [16] (Fig. 1). However, there are two issues that complicate the task of thigh muscle segmentation from MRI. The first issue is that each of the thigh muscles have similar intensity and texture, so they cannot be differentiated using local image features. The second issue is

S. Andrews and G. Hamarneh are with the Medical Image Analysis Lab, School of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby, B.C., Canada, V5A 1S6. (email: {sda56, hamarneh}@sfu.ca)

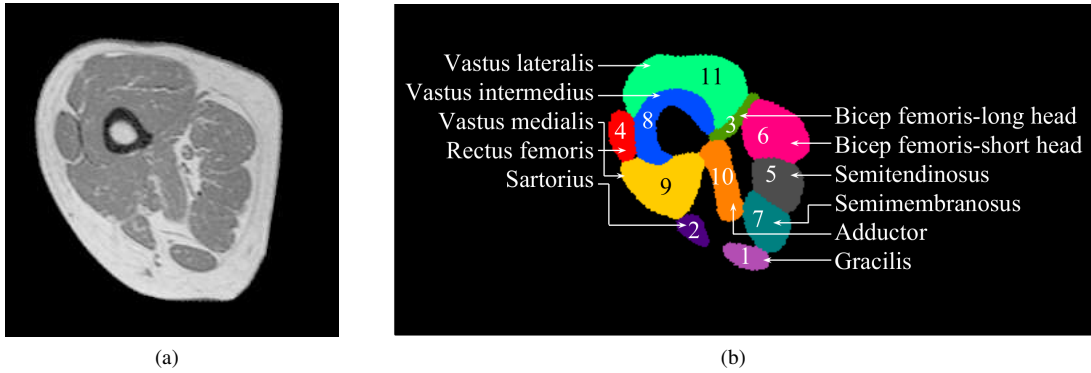


Fig. 1. (Color Figure) An example slice of a thigh MRI and the corresponding labeled ground truth.

that the boundaries of intermuscular fat between neighboring muscles often has thickness on the order of the imaging resolution, and tends to be obscured in many areas by imaging noise and partial volume effects. Further, the texture and *intramuscular* fat in the muscles can create intensity gradients greater than nearby intermuscular boundaries (Fig. 1).

A diverse set of techniques for overcoming these issues have been proposed in the muscle segmentation literature. Gilles et al. deformably register template meshes of muscles and other thigh anatomy to a novel image, matching model edges to boundaries in the images and employing regularization terms designed to model real anatomy [17], [18]. Other works have been successful in employing atlas-based segmentation to muscle segmentation (including certain thigh muscles), registering novel MRI data to an already segmented atlas [19]–[21]. Essafi et al. propose a wavelet-based encoding for calf muscles represented using landmarks, which provides a hierarchical encoding of shape variability [22]. Wang et al. capture muscle shape variability using the relative positions of small cliques of nearby landmark points, and segment novel images by training random forest classifiers for the landmark locations based on image features [23]. The above methods rely on matching explicit surface models to muscle boundaries and can suffer from difficult to detect or false boundaries that often occur in thigh MRI data. Further, none of these methods guarantee a globally optimal solution, and thus can be sensitive to improper initialization.

Baudin et al. employ several techniques for implicit segmentation of thigh muscles that achieve excellent results. First, they present a technique to automatically detect voxels inside the various muscles utilizing a boundary map calculated from local intensity variance, and use these voxels as “seeds” to guide a subsequent segmentation [24]. However, the detection of enough boundaries to ensure seeds lie in the interior of muscle regions may be difficult. Baudin et al. also present a technique for encoding shape models learnt from training segmentations into a random walker (RW) segmentation [25] framework, which benefits from RW’s robustness to missing edges [26], [27]. These works exhibit the greatest similarity to our proposed approach, as the segmentation step is convex and it provides a probabilistic solution. However, the MRI acquisition scheme employed in these works is designed for

intermuscular boundary detection [28], whereas our data exhibits many false edges from intramuscular fat (Fig. 4, Fig. 14) which can cause errors in RW based approaches due to their dependence on local gradient information for regularization. Also, they employ a deformable registration step prior to the construction of their shape models. Deformable registration is difficult to perform optimally, and may remove some of the shape variability that we wish to capture in our PCA shape model in order to be able to use it to differentiate COPD and healthy muscle structure.

In this work, we build a fully automatic thigh segmentation technique that is robust to low quality MRI data. To ensure consistent and predictable automatic results, we wish to ensure our algorithm converges to a globally optimal solution. This is in contrast to many deformable model and atlas-based techniques [18], [20], which tend to use non-convex energy functions and are prone to getting stuck in local minima (but often can be more efficient to locally optimize). Given that a fully automated segmentation technique should be designed to facilitate subsequent manual analysis, our segmentation technique is designed to encode statistically meaningful uncertainty information that can be used to help identify segmentation errors or image pathologies.

Our segmentation technique consists of four main steps:

- 1) We define the GLR representation in the context of thigh muscle segmentation to encode muscle size and adjacency information and train a statistical shape model over the space of GLR representations, capturing the relative thigh muscle shapes and locations.
- 2) We introduce a technique for *presegmenting* images into bone, fat, and muscle classes, with all knee extensor and flexor muscles combined under one label. We use this presegmentation to perform an anatomical alignment of the images, removing pose and size variance so only the variance in higher order shape moments remains.
- 3) We train a random forest classifier to robustly detect all intermuscular boundary locations in the images while discarding false boundaries from intramuscular fat. We devise features oriented towards the nearest boundary when training, adding rotational invariance to the classifier.
- 4) Information from the shape model, the presegmentation,

and the boundary likelihoods are incorporated into a strictly convex energy functional that is globally minimized using a primal-dual technique to produce a probabilistic segmentation.

The remainder of the paper is laid out as follows. In Sec. II, we describe our generalized log-ratio segmentation representation and discuss its applicability to thigh muscle segmentation. In Sec. III, we describe our random forest muscle boundary detection scheme and discuss its benefits. In Sec. IV, we provide our presegmentation and anatomical alignment steps, and develop a strictly convex energy functional including the learnt shape model and muscle boundary locations. In Sec. V, we examine the accuracy, robustness, and usefulness of our segmentation technique.

II. GENERALIZED LOG-RATIO SEGMENTATION REPRESENTATIONS

A. Log-Ratio Transformations

While our goal in this paper is fully automated segmentation, verification of results by a human expert is often necessary for clinical accuracy. Further, the difficulty in detecting intermuscular boundaries in the images means that the location for these boundaries may have to be chosen somewhat arbitrarily in some locations. To address these issues, we design our algorithm to incorporate statistically meaningful, spatially localized uncertainty information into the segmentation.

Multi-label uncertainty information can be encoded using a *probabilistic labeling* of the image. For an image with K labels, a probabilistic labeling is represented by a function $\mathbf{q} = [q_1, \dots, q_k]^T : \Omega \rightarrow \mathbb{P}_K$, where Ω is the set of n voxels and \mathbb{P}_K is the unit simplex:

$$\mathbb{P}_K = \left\{ \mathbf{q} \in \mathbb{R}^K \mid q_i > 0, \sum_{i=1}^K q_i = 1 \right\}. \quad (1)$$

For $x \in \Omega$, $q_\ell(x)$ represents the probability that x is assigned label ℓ . A probabilistic labeling allows, among other things, the ambiguity of intermuscular boundaries to be quantified.

To guide our segmentation, we build a statistical shape model by applying principal component analysis (PCA) to a set of M K -label manually generated training segmentations $\{\mathbf{q}^1, \dots, \mathbf{q}^M\}$, with $\mathbf{q}^i : \Omega \rightarrow \mathbb{P}_K$, extracting the mean shape and modes of greatest variation. In previous works, PCA was applied directly to training segmentations by treating probabilistic labelings as vectors in $\mathbb{R}^{n \cdot K}$ [26], [29]. However, PCA assumes that data lies in an unconstrained vector space. Thus, these shape models include invalid probabilities that require projection back onto the simplex. This projection will discard uncertainty information, since many projections may be onto the vertices of the simplex, with probability 1 for one of the labels and 0 for the others.

Fortunately, techniques have been established that avoid the above issue. The general idea is to map the probabilistic labels at each voxel to an unconstrained vector space, perform statistical analysis, and then map the results back to the simplex. To describe how best to do this, we first introduce some basic concepts in simplicial geometry.

The simplex \mathbb{P}_K is a $(K - 1)$ dimensional vector space, with vector addition and scalar multiplication given by the perturbation and power transformation operations, respectively [30], [31]. For $\mathbf{p}, \mathbf{q} \in \mathbb{P}_K$, perturbation is defined as

$$\mathbf{p} \oplus \mathbf{q} = \mathcal{C}([p_1 q_1, \dots, p_K q_K]), \quad (2)$$

and for a scalar α , power transformation is defined as

$$\alpha \odot \mathbf{p} = \mathcal{C}([p_1^\alpha, \dots, p_K^\alpha]). \quad (3)$$

Here $\mathcal{C}(\cdot)$ is the *normalization* operator, ensuring a vector sums to 1. The perturbation operation can be interpreted as an application of Bayes rule to discrete distributions (i.e. probabilistic labelings), combining a prior distribution (\mathbf{q}) and a likelihood function (\mathbf{p}) into a posterior distribution ($\mathbf{q} \oplus \mathbf{p}$). The power transformation can be used to control the certainty of a probabilistic labeling [32].

Several techniques exist that bijectively map \mathbb{P}_K to \mathbb{R}^{K-1} while mapping (2) and (3) to Euclidean vector addition and scalar multiplication, respectively. The most prominent of these transformations is the LogOdds transformation, introduced by Pohl et al. [32]:

$$\text{LogOdds}(\mathbf{q}) = \left[\log\left(\frac{q_1}{q_K}\right), \dots, \log\left(\frac{q_{k-1}}{q_K}\right) \right]^T, \quad (4)$$

where $\mathbf{q} \in \mathbb{P}_K$ and $\log(\cdot)$ is the natural logarithm.

While the LogOdds transformation is symmetric between the first $k-1$ labels, it is asymmetric in the last label, so changes in the probability of this label have a greater influence on the result of the transformation than the other labels. A symmetrized analogue to the LogOdds transformation is the ILR transformation, used previously in medical image analysis applications [31]. However, a more general asymmetric transformation would allow us control how much influence different labels have on the transformation.

In this paper, we introduce a family of transformations that generalize the LogOdds function, known as *log-ratio* (LR) transformations. An LR transformation of the unit simplex is a function $\text{LR} : \mathbb{P}_K \rightarrow \mathbb{R}^{K-1}$ given by

$$\text{LR}(\mathbf{q}) = A \log(\mathbf{q}), \quad A \mathbf{1} = \mathbf{0}, \quad (5)$$

where $\mathbf{q} \in \mathbb{P}_K$. $A = [\mathbf{a}^1, \dots, \mathbf{a}^K]$ is a $(K - 1) \times K$ matrix, $\mathbf{1}$ and $\mathbf{0}$ are vectors of all 1's and 0's, respectively, and $\log(\cdot)$ is understood to be applied component-wise to \mathbf{q} . The LogOdds transformation in (4) is equivalent to (5) with $A = [I, -\mathbf{1}]$, where I is the identity matrix. The condition on A ensures that vector addition and scalar multiplication correspond to perturbation and power transformation, respectively (see Appendix A). The condition also results in a simple form for the inverse transformation, given by (see Appendix A)

$$\text{LR}^{-1}(\phi) = \mathcal{C}(\exp(A^+ \phi)), \quad (6)$$

where $\phi \in \mathbb{R}^{K-1}$ is an LR-vector, $\exp(\cdot)$ is the natural exponential function, (understood to be applied component-wise), and A^+ is the pseudoinverse of A (i.e. $AA^+A = A$). We now demonstrate how the matrix A can be chosen to encode prior anatomical volume and adjacency information directly into a segmentation representation.




Fig. 2. An illustration of how the GLR transformation can be used to ensure small or thin structures are properly segmented. An image region may be segmented correctly into 3 regions (a) or incorrectly into 2 regions (b). This example assumes a regularization term penalizing the total change in the segmentation. Using no prior adjacency information such as the ILR transformation may result in the middle region being skipped, as the segmentation incurs too large of a regularization penalty (c). If it is known from the training data that regions 1 and 3 do not share boundaries with each other, but rather with the thin region 2, the GLR transformation can be constructed so that transitioning from region 1 to region 3 through region 2 does not incur any additional cost compared to skipping region 2 (d). Such situations occur in thigh muscle segmentation (e).




Fig. 3. (Color Figure) An example of the columns of A for the LogOdds, ILR, and GLR transformations for $K = 3$ labels. For the GLR transformation, the columns are chosen to satisfy a predefined pairwise distance function D .

An LR transformation is a weighted sum of the columns of A , with weights given by the log of the probabilities, so there is a correspondence between the columns of A and the probabilistic labels. To make this correspondence clearer, let $\mathbf{r}^i = \mathcal{C}(\mathbf{e}^i + \epsilon \cdot \mathbf{1}) \in \mathbb{P}_K$ for $i \in \{1, \dots, K\}$, where \mathbf{e}^i is the i^{th} standard basis vector and ϵ is a small scalar, so \mathbf{r}^i represents a probabilistic labeling where label i is very likely. Taking advantage of the fact that $A \cdot \mathbf{1} = \mathbf{0}$ (see Appendix A),

$$\text{LR}(\mathbf{r}^i) = A \log(\mathcal{C}(\mathbf{e}^i + \epsilon \cdot \mathbf{1})) \quad (7)$$

$$= A \log\left(\frac{1}{\epsilon} \cdot \mathbf{e}^i + \mathbf{1}\right) \quad (8)$$

$$\approx \log\left(\frac{1}{\epsilon}\right) \mathbf{a}^i. \quad (9)$$

Thus, the magnitude of $\text{LR}(\mathbf{r}^i)$ is proportional to $\|\mathbf{a}^i\|$, the Euclidean norm of \mathbf{a}_i , and the distance between $\text{LR}(\mathbf{r}^i)$ and $\text{LR}(\mathbf{r}^j)$ is proportional to $\|\mathbf{a}^i - \mathbf{a}^j\|$. By adjusting the value of $\|\mathbf{a}^i - \mathbf{a}^j\|$, we can encode pairwise label distances into an LR transformations that determine how much the transformation is affected by changes from “likely label i ” to “likely label j ”. There are two ways that this pairwise distance will affect our segmentation algorithm:

- To regularize a segmentation we penalize large gradient values in the segmentation’s components (see Sec. IV-D for more details), and the gradient values between an image region that is “likely label i ” and an image region that is “likely label j ” will be proportional to $\|\mathbf{a}^i - \mathbf{a}^j\|$. Increasing (decreasing) $\|\mathbf{a}^i - \mathbf{a}^j\|$ will deter (encourage) boundaries between labels i and j in the segmentation. This can ensure, for example, that a thin label is not

neglected by discouraging its neighboring labels from sharing a boundary.

- Increasing $\|\mathbf{a}^i - \mathbf{a}^j\|$ will result in greater variance in LR space near boundaries between labels i and j , and thus these boundaries will have a greater influence in PCA.

In LogOdds, the first $(K - 1)$ columns of A are symmetric (i.e. equidistant from each other), but the last column is not, with $\|\mathbf{a}^K - \mathbf{1}\| = \sqrt{K - 1}$ (Fig. 3a). Since the last label is usually chosen to represent the “background” label, foreground-background boundaries have a greater effect on the PCA modes calculated in LogOdds space than foreground-foreground boundaries [31]. Since accurately modeling intermuscular boundaries is key to accuracy in thigh muscle segmentation, the LogOdds transformation is not appropriate for this task.

The isometric log-ratio (ILR) transformation is another LR transformation that has been applied to medical image segmentation, and uses a matrix A with all K columns symmetric (i.e. equidistant) (Fig. 3b) [30], [31], [33], [34]. This results in all boundaries being treated equally, so the ILR transformation is an appropriate choice when lacking prior knowledge about the segmentation labels.

In thigh muscle segmentation, however, we want our shape model to focus on capturing intermuscular boundaries correctly, as the muscle-background boundaries are relatively easy to detect from image gradients. Implicit PCA shape models often capture the variability in larger structures more accurately than smaller structures, simply because larger structures contain more voxels, which increases their overall influence on the PCA modes. We can offset this bias towards larger structures by constructing our LR transformation so that voxels on the boundaries of smaller structures have a greater influence on the PCA modes. In the following section, we present a technique for choosing the matrix A to encode predefined pairwise label distances, resulting in a *generalized log-ratio* (GLR) transformation (Fig. 3c).

B. Setting Label Transition Penalties

In order to better capture the shapes of smaller structures, we want to choose a distance function between labels that encodes the relative volumes and adjacencies of the structures. At a high level, the idea is to construct a distance function between labels that varies inversely with the surface area

of the shared boundaries of the structures corresponding to those labels. Smaller structures will have smaller boundaries, and thus greater distances and increased influence on a PCA shape model. From the training segmentations, we calculate the average number of voxels in the boundary between labels i and j (approximately the average surface area), $\sigma_{i,j}$ (with $\sigma_{i,i} = 0$). We define σ as the matrix of all surface areas and $\bar{\sigma}$ as the mean across all pairs of labels:

$$\bar{\sigma} = \frac{2}{K \cdot (K-1)} \sum_{i=1}^K \sum_{j=i+1}^K \sigma_{i,j}. \quad (10)$$

We want to define a squared-distance matrix D with components varying inversely to the corresponding components of σ ; for example we use

$$D_{i,j} = \frac{\bar{\sigma}}{\sigma_{i,j} + \bar{\sigma}} \text{ if } i \neq j, \quad D_{i,i} = 0. \quad (11)$$

The smaller the boundary shared by a pair of labels, the greater the distance between them will be under the GLR transformation, discouraging transitions between them in the segmentation (Fig. 2).

Given such a $K \times K$ matrix D of pairwise squared distances between labels, we would like to choose A such that $\|\mathbf{a}^i - \mathbf{a}^j\|^2 = D_{i,j}$. It is known that D encodes a Euclidean distance if and only if $T = -\frac{1}{2}CDC$ is positive semidefinite, where $C = I - \frac{1}{K}\mathbf{1}\mathbf{1}^\top$ is a *centering* matrix [35]. If T is positive semidefinite, then the desired A is given by $T = A^\top A$. The distance encoded in D from (11) may not be Euclidean, in which case T will have negative eigenvalues. Only the relative values of the distances in (11) are meaningful (i.e. which labels are closer and which are farther), so we want to alter these distances to make them Euclidean while maintaining this relative ordering. Let ψ be the smallest eigenvalue of T . If $\psi < 0$, we add (-2ψ) to the off diagonal components of D , which will ensure a Euclidean distance (see Appendix B) and maintain the relative distance ordering.

For the remainder of this paper, the GLR transformation uses the distance function specified in (11). We will use the GLR transformation in our segmentation framework, detailed in Sec. IV, but first we will introduce another component of our segmentation framework: a robust statistically-based technique for identifying intermuscular boundaries.

III. INTERMUSCULAR BOUNDARY DETECTION

Accurate detection of intermuscular boundaries in thigh MRI data is confounded by several factors. First, the layer of intermuscular fat dividing the muscle tissue can often be of a sub-voxel thickness, resulting in partial volume effects obscuring the boundary. Second, *intramuscular* fat and other textures inside muscle regions can create many false edges (Fig. 4). Third, the intermuscular boundaries vary greatly in both thickness and appearance, even within a single slice of an MRI (Fig. 4d).

Wang et al. account for the difficulty in locating muscle boundaries by adopting a random forest (RF) classifier [23], [36], trained from existing segmentations, to attempt to locate

landmark points. We expand on this technique and train an RF classifier to identify all intermuscular boundaries and distinguish them from false boundaries. RFs are able to efficiently model complex unknown distributions, including distributions with many modes, and thus may be able to account for the large variety in intermuscular boundary appearances.

Our RF boundary classifier builds upon our previous technique for intermuscular boundary detection [37], where we calculated the direction and magnitude of maximum curvature of the intensity values at each voxel. However, the variety of boundary appearances (Fig. 4d) limits the ability of the curvature alone to detect them all, and intramuscular fat results in false boundaries. We increase the “recall”, detecting a larger variety of boundary appearances, by using a larger set of filters beyond just the curvature. We increase the “precision”, detecting less false boundaries, by training a RF classifier to distinguish intermuscular fat from intramuscular fat.

Instead of only looking at the curvature of the image, we use a set of 20 3D directional Gabor filters, designed to capture a range of boundary widths and appearances, oriented in 18 different directions, for a total of $20 \times 18 = 360$ filters. A Gabor filter in the direction of the unit vector \mathbf{v}' is given by

$$g(\mathbf{v}) = \exp\left(-\frac{\|\mathbf{v}\|^2}{2\sigma^2}\right) \cos\left(2\pi \frac{\mathbf{v}'^\top \mathbf{v}}{\gamma}\right) \quad (12)$$

$$\text{or } g(\mathbf{v}) = \exp\left(-\frac{\|\mathbf{v}\|^2}{2\sigma^2}\right) \sin\left(2\pi \frac{\mathbf{v}'^\top \mathbf{v}}{\gamma}\right). \quad (13)$$

We use $\sigma^2 \in \{0.5, 1, 2, 4, 8\}$ and $\gamma \in \{0.25, 1\}$ for both \sin and \cos to get our 20 filters for a given direction. Since the out-of-plane resolution is much greater than the in-plane resolution (see Sec. V-A) and since the thigh muscle boundaries are usually oriented perpendicular to the image plane, the 18 directions are all chosen in-plane, aligned with the angles $\{10^\circ, 20^\circ, \dots, 180^\circ\}$. Further, we found that increasing the number of features did not have a significant influence on accuracy, likely due to the limited amount of training data.

We use the 360 filter responses to construct a feature vector for each voxel (with each set of 20 filters for each direction grouped together), and then use these features to train a RF classifier from the training data. We use 100 trees in the RF, each with a depth of 7 and 20 features randomly chosen per tree and each node splitting on a single feature. The trade-off in these parameters is computation time versus accuracy; we found increasing these parameters did not increase accuracy.

If we append the filter responses in a fixed order to create a length 360 feature vector, boundaries that are similar in appearance but oriented differently will have completely different feature vectors, complicating the training process. Instead, to construct the feature vectors for training, we shift the filter responses so the first set of 20 filters are oriented together towards the nearest boundary (calculated using the training segmentations), with the remaining filters in clockwise order, as illustrated in Fig. 5. For voxels on a boundary, this boundary alignment results in the first 20 feature components corresponding to responses for filters oriented across the boundary, thus adding rotational invariance to the classifier and effectively increasing the amount of training data




Fig. 4. An example showing the false boundaries created by intramuscular fat (b), with segmentation (i.e. true boundaries) shown in red for comparison (c) and the zoomed versions of some true boundaries (green boxes) and false boundaries (red) (d), demonstrating the difficulty in distinguishing them.

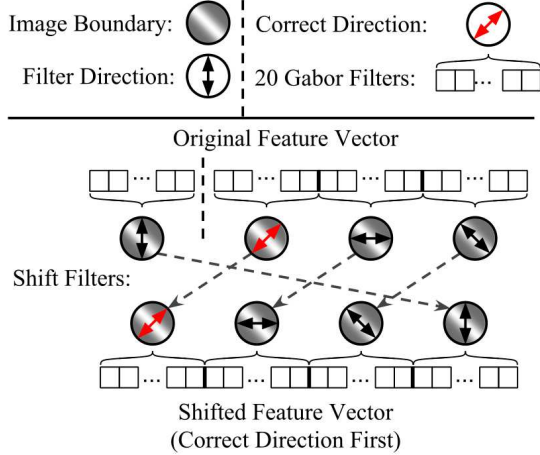


Fig. 5. An illustration of how the filters are shifted at each voxel to make the feature vector used when training the RF. The filters corresponding to the red direction should give the largest response. The feature vectors at each location are constructed by ordering the filter responses so that the direction with the largest expected response comes first. When applying the RF, each of the 18 possible orderings of the filters are tested.

available. Incorporating all of the filters into the RF is useful for detecting false boundaries, which may have large filter responses in multiple different directions.

Since we are focused on detecting intermuscular boundaries (as mentioned, other boundaries are easy to detect), we only train on voxels that are labeled as muscle in the training segmentations. Intermuscular boundary voxels are assigned target value 1, voxels adjacent to these are also assigned target value 1 (since a near miss will still provide decent segmentation accuracy and this increases the amount of positive training data), and other muscle voxels are assigned target value 0.

Given a novel image, we again extract the 360 filter responses at each voxel, but we are not able to shift these filters to align with the nearest boundaries, since these boundaries are not known. Instead, we apply the RF classifier to 18 different feature vectors, generated from all 18 possible shifts of the filter responses, which gives us a boundary likelihood for each voxel/direction pair. This results in a boundary likelihood function $h : \Omega \rightarrow [0, 1]$ and corresponding boundary normal vector field $\mathbf{d}^1 : \Omega \rightarrow \mathbb{S}^2$ (where \mathbb{S}^2 is the unit sphere). An example of the boundary likelihood function h is shown in Fig. 6. We describe how we incorporate h and \mathbf{d}^1 into our segmentation algorithm in Sec. IV-D.




Fig. 6. An example of the boundary likelihoods found using our RF classifier. In (b), lighter shading corresponds to higher likelihoods. (c) shows the true boundaries overlaid on the slice.

IV. CONVEX SEGMENTATION FRAMEWORK

We formulate our segmentation algorithm as an energy minimization problem. We represent segmentations as functions $\phi : \Omega \rightarrow \mathbb{R}^{K-1}$, where $\phi(x)$ is a GLR vector as defined in Sec. II-B. Since Ω is a discrete set of n voxels, segmentations can be interpreted as vectors in $\mathbb{R}^{n \cdot (K-1)}$, useful when performing PCA. We define our energy functional as a weighted sum of terms:

$$E(\phi) = E_{\text{rgn}}(\phi) + w_1 E_{\text{shp}}(\phi) + w_2 E_{\text{bnd}}(\phi), \quad (14)$$

where w_1 and w_2 are scalar weights. The first term incorporates local regional image information into the segmentation by performing a *presegmentation*, as described in Sec. IV-A. The second term, incorporating the information from a PCA shape model, is described in Sec. IV-C. The third term forces the segmentation boundaries to conform to the intermuscular boundaries found in Sec. III, and is described in Sec. IV-D.

A. Presegmentation

While it is difficult to distinguish thigh muscles from each other, it is much easier to distinguish muscle tissue from other anatomical structures, even by intensity alone, as seen in Fig. 7a. We take advantage of this to perform a *presegmentation* and identify which voxels are muscle (regardless of which particular muscle) to guide the segmentation. While we are only eventually interested in segmenting the muscles, it is useful to segment the fat tissue and the femur bone for the purpose of anatomical alignment (detailed in Sec. IV-B). Since the *presegmentation* does not require a shape model, it can be used to align the training images prior to shape model construction.

To perform the *presegmentation*, we note that the thigh images exhibit three main intensity classes (Fig. 7a), “dark”

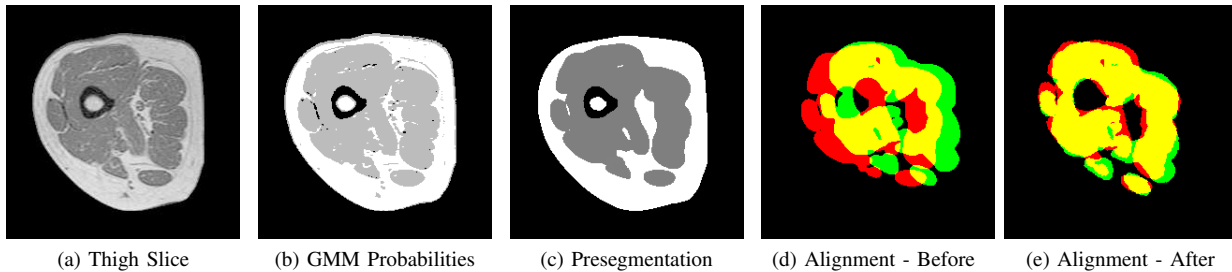


Fig. 7. An example illustrating the GMM probabilities and presegmentation. In (b), we see the thresholded probabilities for GMM on “light”, “gray”, and “dark” labels. In (c), we see the results of the thresholded presegmentation of the image. In (d) we see the overlay of the presegmented muscle labels from two images (red and green, overlap in yellow), and in (e) we see the same muscle labels overlaid after being aligned.

(background, cortical bone), “gray” (muscle), and “light” (fat, bone marrow). We assign a probability for each of these classes to each voxel by fitting a Gaussian mixture model (GMM) to the voxel intensities. Since we know roughly the means and variances of the intensities of these three classes from training data, we can provide a good initialization to the GMM, and because the three classes are so distinct, we have found the GMM to converge to the same distribution when the initialization is perturbed. The result is a length 3 probability vector, mapped to \mathbb{R}^2 using the ILR transformation¹, assigned to each voxel, $\pi_{\text{GMM}} : \Omega \rightarrow \mathbb{R}^2$.

To account for artifacts in the presegmentation due to noise and partial volume effects (Fig. 7b), we incorporate a regularization step:

$$\pi_{\text{opt}} = \arg \min_{\pi} E_{\text{preseg}}(\pi) \quad (15)$$

$$E_{\text{preseg}}(\pi) = \sum_{x \in \Omega} \|\pi(x) - \pi_{\text{GMM}}(x)\|^2 + w_3 \sum_{x \in \Omega} \sum_{i=1}^2 \|\nabla \pi_i(x)\|, \quad (16)$$

where w_3 is a scalar weight and ∇ is the spatial gradient operator. The second term in (16) is a standard total variation (TV) term on the components of ϕ , which can be globally optimized using primal-dual techniques [38] and creates a smooth presegmentation (Fig. 7c).

Identifying the femur bone is useful both for the alignment of images described in Sec. IV-B and to provide anatomical context to the shape model described in Sec. IV-C. To extract the bone from the presegmentation, we use the following postprocessing technique, which we have empirically found to give consistently accurate results:

- 1) Convert π_{opt} to \mathbb{P}_3 and threshold to create a crisp (non-probabilistic) segmentation;
- 2) Find the convex hull of the “gray” (muscle) label;
- 3) Label the largest connected component of the “dark” region inside this convex hull as cortical bone;
- 4) Find the “light” region voxels inside the cortical bone region and label it as bone marrow.

Following this procedure, we have an estimate of the location of muscle tissue (“gray” regions), bone, and fat (the remaining “light” regions, see Fig. 7c). Incorporating bone and fat into

the segmentation allows us to encode which muscles border bone and which border fat in the GLR transformation.

Using the probabilities from π_{opt} and the bone and fat locations found using the above procedure, we construct a 4 label probabilistic segmentation (muscle, bone, fat, and background). We convert this to a 14 label segmentation by “spreading” the muscle probability at each voxel evenly among all 11 muscles labels, and apply the GLR transformation to the resulting probabilities, denoting the result π_{rgn} . For the remainder of the paper, we fix the number of labels as $K = 14$.

We define the regional energy term as:

$$E_{\text{rgn}}(\phi) = - \sum_{x \in \Omega} \langle \phi(x), \pi_{\text{rgn}}(x) \rangle, \quad (17)$$

where $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product. Here, we use an inner product instead of a distance term (as in (16)) since we do not want $\phi(x)$ to be close to $\pi_{\text{rgn}}(x)$ (which has a maximum muscle probability of $1/11$) but to only follow its guidelines as to what voxels are and are not muscle.

B. Anatomical Alignment

Besides acting as a regional prior for the segmentation, we use the presegmentation π_{opt} to guide the alignment of images for the purpose of building an anatomical shape prior. Since our shape model is poor at capturing variability in lower order shape moments such as pose and size, but is strong at capturing variability in higher shape moments, we preprocess the images by applying a linear transform to each slice to remove pose and size variability. In Sec. V-A, the images are cropped from just above the knee to just below the hip using a supervised method, so the slices in each image already correspond. Further, the out-of-plane resolution is much greater than the in-plane resolution, so out-of-plane rotations result in significant image blurring. For these reasons, we focus on in-plane alignment. This leaves 5 degrees of freedom to eliminate for each slice: 2 for translation, 1 for rotation, and 2 for scaling.

The translation is determined by aligning the centroid of the bone to a common location. The rotation and scaling are determined by fitting an ellipse to the detected muscle region, rotating the image around the bone so the major axes of the ellipses align, and scaling the slices anisotropically along the ellipse axes so the muscles are roughly the same size (Fig. 7e). This procedure may result in adjacent image slices being

¹We do not have strong prior knowledge to incorporate into GLR.

transformed differently and no longer being aligned, so we take each of the 5 transformation parameters for each slice and fit to them a linear function of the slice index, so the transformations vary linearly across slices.

C. Building Muscle Shape Models

Once the images have been aligned, we capture the variability in the higher order shape moments by constructing a PCA shape model using training segmentations.

The training segmentations are crisp (non-probabilistic), so must be converted to probabilistic segmentations in order to be mapped to GLR space. We use the technique introduced by Pohl et al., calculating the signed distance maps (SDMs) of the foreground labels and converting them to probabilistic segmentations using the LogOdds transformation [32]. We then map each probabilistic segmentation to GLR-space and perform PCA to find the $m < M$ modes of greatest variance, with m chosen to capture 90% of the variability in the M training segmentations. This gives us the mean shape μ , the $(n \cdot (K - 1)) \times m$ matrix of shape modes Γ , and the $m \times m$ diagonal matrix of variances $\Lambda = \text{diag}([\lambda_1, \dots, \lambda_m]^\top)$, where $\lambda_1 \geq \dots \geq \lambda_m$.

We decompose $\phi = \phi_{\text{PC}} + \phi_{\text{orth}} + \mu$, where ϕ_{PC} lies in the span of the shape modes Γ and ϕ_{orth} is orthogonal to each of the shape modes ($\Gamma^\top \phi_{\text{orth}} = \mathbf{0}$). We divide our shape energy term into two parts to separately penalize the two components of ϕ . The first part is given by a squared Mahalanobis distance:

$$E_{\text{PC}}(\phi) = (\phi_{\text{PC}})^\top \Gamma \Lambda^{-1} \Gamma^\top (\phi_{\text{PC}}) \quad (18)$$

$$= (\phi - \mu)^\top \Gamma \Lambda^{-1} \Gamma^\top (\phi - \mu). \quad (19)$$

We note (19) does not penalize ϕ_{orth} . A common practice is to constrain $\|\phi_{\text{orth}}\| = 0$ [29], [33]. In this paper, to allow for more flexible segmentations, we introduce a second shape energy term penalizing ϕ_{orth} :

$$E_{\text{orth}}(\phi) = \|\phi_{\text{orth}}\|^2 = \|(\phi - \mu) - \Gamma \Gamma^\top (\phi - \mu)\|^2. \quad (20)$$

We combine these to get

$$E_{\text{shp}}(\phi) = E_{\text{PC}}(\phi) + \lambda_m^{-1} E_{\text{orth}}(\phi), \quad (21)$$

where E_{orth} is multiplied by λ_m^{-1} to ensure that variations out of the span of the shape modes are penalized at least as much as variations within the shape modes.

D. Boundary-Based Term

In Sec. III, we described how to train and apply a rotationally invariant RF classifier that, when applied to a novel image, provides a boundary likelihood function $h : \Omega \rightarrow [0, 1]$ and a vector field $\mathbf{d}^1 : \Omega \rightarrow \mathbb{S}^2$ normal to the boundary tangent plane at each voxel. We define $\mathbf{d}^2, \mathbf{d}^3 : \Omega \rightarrow \mathbb{S}^2$ as two vector fields that, along with \mathbf{d}^1 , form orthonormal bases at each voxel (thus $\mathbf{d}^2(x)$ and $\mathbf{d}^3(x)$ span the plane tangent to the detected boundary at $x \in \Omega$). To incorporate h and $D(x) = [\mathbf{d}^1(x), \mathbf{d}^2(x), \mathbf{d}^3(x)]$ into an energy term, we penalize gradients in the segmentation's components in directions other than orthogonal to the boundary, and penalize changes orthogonal to the boundary according to the likelihood h . Specifically,




Fig. 8. A graphical illustration of the dual formulation of E_{bnd} in (23). The example is shown in 2D instead of 3D for clarity. (a) shows the p corresponding to a $\nabla \phi^i(x)$ parallel to $\mathbf{d}^1(x)$. (b) shows the p corresponding to a $\nabla \phi^i(x)$ perpendicular to $\mathbf{d}^1(x)$. (c) shows the p for an arbitrary $\nabla \phi^i(x)$, with the strength of the penalty being proportional to the length of the projection of p onto $\nabla \phi^i(x)$.

defining the diagonal matrix $H(x) = \text{diag}([(1 - h(x)), 1, 1])$, we define our energy term as

$$E_{\text{bnd}}(\phi) = \sum_{x \in \Omega} \sum_{i=1}^{K-1} \|H(x)^{-1} \cdot D(x)^\top \cdot \nabla \phi^i(x)\|. \quad (22)$$

We ensure $(1 - h(x)) > \epsilon > 0$ for numerical stability. By projecting $\nabla \phi^i(x)$ onto D , we penalize the component of $\nabla \phi^i(x)$ parallel to $\mathbf{d}^1(x)$ by a factor of only $(1 - h(x)) < 1$ as much as the component of $\nabla \phi^i(x)$ in the span of $\{\mathbf{d}^2(x), \mathbf{d}^3(x)\}$, encouraging the segmentation to change labels across the detected muscle boundaries.

This modified total variation term can be efficiently globally minimized using established primal dual techniques [38], [39]:

$$E_{\text{bnd}}(\phi) = \sum_{x \in \Omega} \sum_{i=1}^{K-1} \max_{\mathbf{p}^i(x) \in C(x)} \mathbf{p}^i(x)^\top \nabla \phi^i(x) \quad (23)$$

$$C(x) = \left\{ \mathbf{p} \in \mathbb{R}^3 \left| \left(\frac{\mathbf{p}^\top \mathbf{d}^1(x)}{1 - h(x)} \right)^2 + \left(\mathbf{p}^\top \mathbf{d}^2(x) \right)^2 + \left(\mathbf{p}^\top \mathbf{d}^3(x) \right)^2 \leq 1 \right. \right\}. \quad (24)$$

The dual variables $P = [\mathbf{p}^1, \dots, \mathbf{p}^{K-1}]$ are constrained to the ellipsoids $C(x)$ at each voxel, which is flattened along the direction $\mathbf{d}^1(x)$ by a factor of $1/(1 - h(x))$ (Fig. 8). We employ a primal-dual proximal point method, which can guarantee convergence given small enough step size [40], [41] (see Appendix C).

V. RESULTS

A. Data

A sample of patients was obtained with moderate to severe COPD on the basis of the Global Initiative for Chronic Obstructive Lung Disease guidelines [42] and of greater than 50 years of age. The healthy control group were age greater than 50 years, free of lung disease and non-smokers.

A 1.5T MRI scanner (1.5T Horizon Echospeed Scanner; GE Healthcare, Milwaukee, WI) was used to acquire 5 mm-thick, contiguous, axial slices from the anterior superior iliac spine to the tibial plateau while the subjects lower extremities were strapped to a foam block to minimize movement. Images were T1-weighted magnetic resonance (echo time, 8 ms; repetition time, 650 ms) with a 40 cm² field of view and

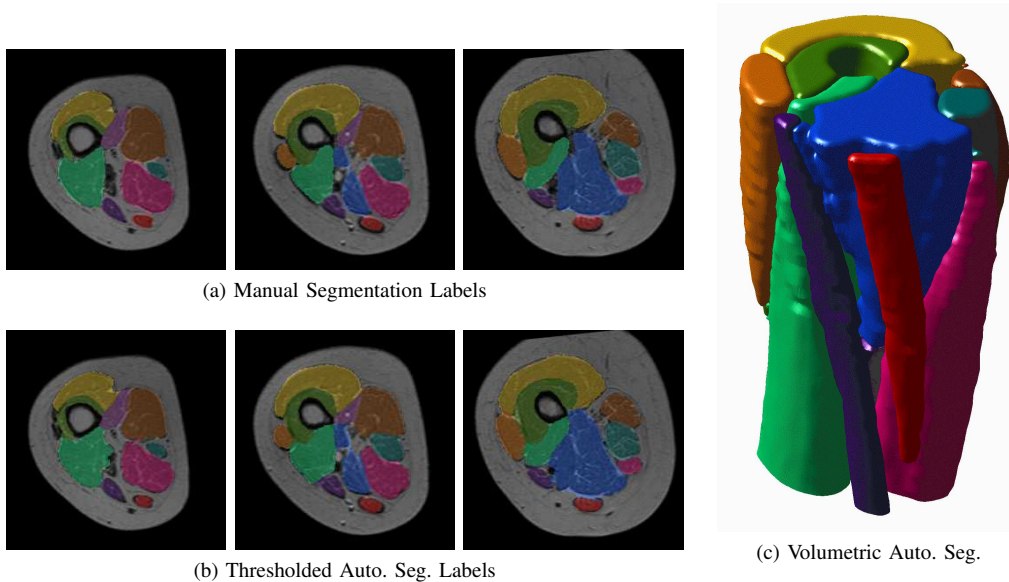


Fig. 9. A comparison of manual and automated segmentation results.

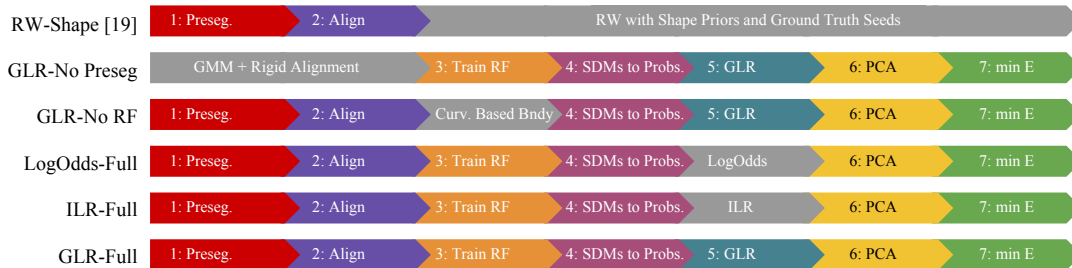


Fig. 10. A summary of our method and the various techniques we compare against. Colored blocks correspond to steps from Algorithm 1, and gray blocks correspond to one of our contributions being replaced by another technique.

Algorithm 1 An Overview of the Segmentation Framework

- 1: Construct a presegmentation of all the images (training and testing) (see Sec. IV-A, Fig. 7).
- 2: Use the presegmentation to align all the images (see Sec. IV-B, Fig. 7e).
- 3: Train a RF boundary classifier (see Sec. III, Fig. 6).
- 4: Convert the training segmentations to probabilistic segmentations by applying LogOdds to the SDMs of the labels (see Sec. IV-C).
- 5: Use boundary overlaps of the labels in the training segmentations to choose an appropriate GLR transformation, and apply it to the probabilistic training segmentations (see Sec. II-B).
- 6: Construct a PCA shape model (see Sec. IV-C).
- 7: For each testing image, construct and minimize the appropriate energy functional (14) to find the optimal probabilistic segmentation.

a 512×384 pixel matrix (in-plane resolution, 0.78×0.78 mm). For each subject, two sets of images, one for the upper and another for the lower thigh regions, were collected in immediate succession without a change in the subjects

position. A landmark (vitamin E capsule) at the midthigh (half the distance between the anterior inferior iliac spine and the superior margin of patella) was identified on the MRI scans to register the two images into a single image of the entire thigh. The vitamin E landmark facilitated the identification of overlap between the two sets of images. The MRI scan yielded a total of approximately 100 slices for each participant, which were merged into a single 3D image using the Merge module in the Amira 3.1 software package (Mercury Computer Systems, Inc, Chelmsford, MA), with the coincidence of the landmark verified after merging. ITK-SNAP 1.6.0.1 [43] was used for the manual slice-by-slice segmentation of individual muscles of the left thigh from the merged axial MRIs. Segmentation was performed by a physical therapist with expert anatomic knowledge. The images were semi-automatically cropped down to $250 \times 250 \times 40$, so that only the left thigh was contained, and the slices started just above the knee and extended 5×40 mm to the upper thigh, so that the muscles of interest were all almost entirely contained. Differences in thigh length are left to be accounted for by the shape model.

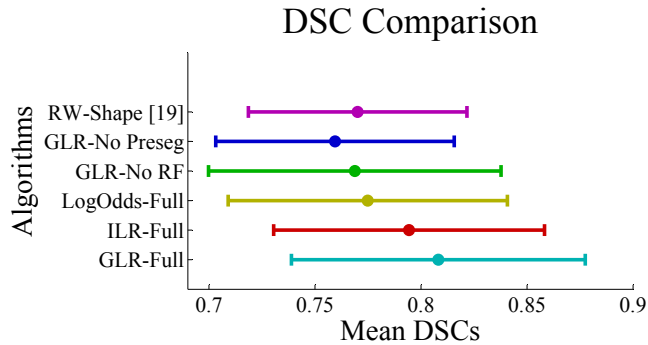


Fig. 11. A comparison of GLR-Full to each of the other 5 methods in Fig. 10, summarized for each label and each image. We note that the GLR-Full results have a higher mean DSC and also lower variance than the other methods, indicating more consistent results across all muscles. The RW based algorithm has a higher DSC than some of the other techniques, but a larger variability, because of instances where false boundaries near the seeds skewing the results.

Muscle	DSC
1	0.72 \pm 0.24
2	0.71 \pm 0.27
3	0.70 \pm 0.16
4	0.75 \pm 0.20
5	0.80 \pm 0.16
6	0.89 \pm 0.08
7	0.85 \pm 0.11
8	0.79 \pm 0.06
9	0.93 \pm 0.06
10	0.81 \pm 0.11
11	0.86 \pm 0.03

TABLE I
THE MEAN DSC FOR EACH MUSCLE IN GLR-FULL.

B. Muscle Segmentation

We test our automated segmentation technique by randomly taking 20 images for training and use the remaining 20 for testing. We choose 10 of the training images from the COPD patients and 10 from the healthy patients, ensuring our trained models incorporate both healthy and pathological examples.

Our parameters w_1 , w_2 , and w_3 were set in the training stage using only the training data. For w_1 and w_2 , we found the best combination of weights, in terms of average DSC, out of $20 \times 20 = 400$ possibilities when running our segmentation algorithm on each of the training images in a leave-one-out approach, giving $w_1 = 5$ and $w_2 = 50$. w_3 was set by finding the value that provided the best match between the presegmented muscles and the manually segmented muscles in the training data, giving $w_3 = 10$. By setting the parameters using only the training data, we avoid biasing the weights.

We perform the 7 steps in Algorithm 1. Our main contributions correspond to steps 1 and 2 (presegmentation and alignment), step 3 (RF boundary classifier), and step 5 (the GLR segmentation representation). We will evaluate the importance of each of these contributions by removing them one at a time from the segmentation algorithm and comparing the change in segmentation accuracy, quantified using the Dice similarity coefficient (DSC) with the manual segmentation.

We remove steps 1 and 2 by using the GMM probabilities instead of the presegmentation to construct E_{reg} in (17) and

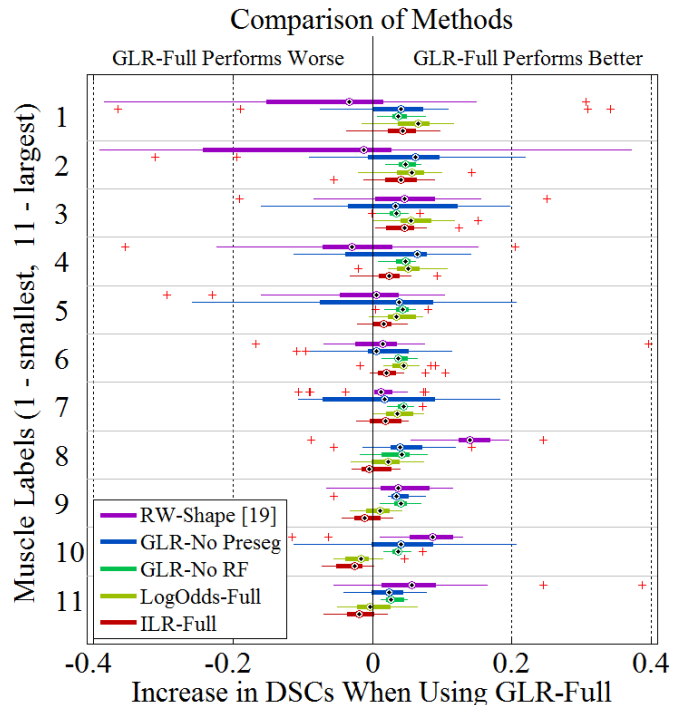


Fig. 12. A comparison of the various algorithms, where positive values indicate GLR-Full performs better. We note that the GLR-Full algorithm typically provides higher DSC values compared to the other algorithms that do not include all of our contributions, with the exceptions being RW-Shape Prior and ILR-Full. RW-Shape performs better on muscles with strong intermuscular boundaries and minimal intramuscular fat (e.g. label 1), but performs significantly worse on muscles with weak intermuscular boundaries (e.g. label 8). ILR-Full performs better on some of the larger muscles, as is expected, because the GLR transformation was designed to ensure the smaller muscles were captured better by the segmentation. The results are summarized in Fig. 11.

by performing a rigid alignment of the foreground voxels instead of the alignment described in Sec. IV-B, denoting the resulting algorithm “GLR-No Preseg”. We remove step 3 by using the existing curvature-based boundary function described at the beginning of Sec. III [37], denoting the resulting algorithm “GLR-No RF”. We remove the use of the GLR transformation in step 5 by instead using the LogOdds and ILR transformations, denoting the resulting algorithms “LogOdds-Full” and “ILR-Full” respectively. We denote our full algorithm “GLR-Full”. Fig. 10 provides a summary.

We evaluate our method against the recent thigh segmentation techniques of Baudin et al. [27]. The technique of Baudin et al. achieves an excellent mean DSC of 0.86 ± 0.07 , though their data differs significantly from ours. Their images exhibit stronger intermuscular boundaries and less false boundaries from intramuscular fat (Fig. 14). These complications are discussed in more detail in Sec. VI; in this section we compare to their technique by running an implementation of it on our data. Their technique incorporates a PCA shape model into a random walker segmentation formulation (RW), so we use the same training and testing sets to build their shape model. We set the free parameters in their algorithm similarly to how we set our w_1 , w_2 , and w_3 , using only the training data. Baudin et al. also introduce a technique for automatically

finding seed voxels in the centers of muscles, based on detecting image gradients [24]. We found the false boundaries from intramuscular fat in our images made placing seeds correctly difficult, but to ensure a fair comparison, we use the manual segmentations of the testing images to provide the RW algorithm with 5 to 15 seeds for each muscle (depending on the muscle volume) spread along their 3D medial axes (note these seeds are not used when running our method).

The comparative segmentation results of the methods from Fig. 10 are summarized in Fig. 11, where we see that GLR-Full in general outperforms the other methods, both in terms of mean DSC and variance in DSC, due to segmenting all muscles well. The DSC values for GLR-Full grouped by muscle are seen in Table I. More detailed results for individual muscles are seen in Fig. 12. The muscle labels are ordered by mean volume across the images, with label 1 corresponding to the smallest muscle (Fig. 1).

RW provides good results in the presence of missing edges, which was a motivation for its use in thigh segmentation, where intermuscular boundaries can be difficult to detect. However, RW regularizes the segmentation using local image gradients, so in the presence of numerous false edges, such as those resulting from intramuscular fat, RW can give erroneous segmentation results. We see this reflected in Fig. 11, as RW-Shape Prior performs well on muscles with strong intermuscular boundaries and minimal intramuscular fat (e.g. label 1), but performs poorly on muscles with weak intermuscular boundaries and intramuscular fat (e.g. label 8).

From Fig. 11, we also see that ILR-Full outperforms GLR-Full on some of the larger muscles, which is to be expected, since there must be a trade-off as to which muscles are better modeled by the LR transformation. However, GLR-Full exhibits less variance between the DSC scores by focusing on the more difficult smaller muscles, an important property because the shapes of smaller muscles are also important for COPD treatment planning [13].

Another the recent thigh segmentation technique is that of Gilles and Magnenat-Thalmann [18]. Gilles and Magnenat-Thalmann also focus on modeling other anatomy, such as ligaments and cartilages. They report a mean distance between manual and automatic segmentations (without user input) of 1.58 ± 1.92 mm, whereas the mean distance found from our automatic segmentation after thresholding is slightly less, at 1.54 ± 0.67 mm.

VI. DISCUSSION

1) *Segmentation Uncertainty*: Our probabilistic segmentation technique actively attempts to maintain meaningful probabilities at each voxel. As such, we should be able to extract useful segmentation uncertainty information from these probabilities. We follow Saad et al. [7] and calculate uncertainty as 1 minus the difference between the highest and second highest label probabilities, a measure of how close the segmentation was to assigning a different label to a voxel. These uncertainty values should correlate well with segmentation error; more specifically, the set of voxels with the highest uncertainty should have both high “error precision”

(i.e. most of them should be erroneous) and high “error recall” (i.e. most of the error should be contained in this set). We evaluate these assertions in Fig. 13.

This evaluation of uncertainty could have many uses. It would be useful if our automatically generated segmentations were used as prior terms in an interactive segmentation algorithm such as RW. Uncertainty could be used to direct the user towards areas of high uncertainty to provide input [44]. These areas of high uncertainty may indicate a pathology inconsistent with the shape model or errors made in the automated segmentation.

2) *COPD Classification*: A unique component of our data is the inclusion of COPD patients. While inclusion of COPD data complicates the segmentation process, through increased intramuscular fat and shape variability, it also provides an opportunity to evaluate our shape model by attempting to distinguish COPD and healthy images. We build two new shape models, one from the 10 COPD training images and one from the 10 healthy training images, and then construct the energy term E_{shp} from (21) for both. We evaluate these two shape terms at each of the automatically generated segmentations from the testing set, the intuition being that the we should be able to determine if a patient has COPD symptoms based on which shape energy gives a lower value (indicating a better shape match). We found that, of the 10 COPD and 10 healthy images from our testing set, 6 of the COPD images and 9 of the healthy images were properly classified using this technique. This is not a statistically significant result (due to the small sample size), but indicates that potential differences in higher order shape moments are being captured by these shape models.

3) *Comparison to Other Works*: Due to the lack of publicly available thigh muscle data and differences in imaging modalities between related publications, comparison to competing methods proves difficult. The works of Baudin et al. [24], [26], [27] are perhaps closest to this work algorithmically, but the data they use exhibits significantly different properties (Fig. 14).

They use a 3T scanner and a 3-point Dixon sequence that combines several acquisitions in order to improve the strength of intermuscular boundaries [28]. Unfortunately, MRI methods designed for muscle detection are not always available, as is the case for our COPD data.

Due to the stronger intermuscular boundaries and less false edges (Fig. 14), their RW-based algorithm is well-suited to the task (and performs better than our method would on similar data). However, as demonstrated in Sec. V, our method performs better on our data. We note that, due to the intensity inhomogeneities in the image from Fig. 14b, our presegmentation method would have to be updated, e.g. to correct for these inhomogeneities ahead of time [45]. Fortunately, the boundaries between muscle, fat, and bone are still clear, so developing an appropriate presegmentation technique should not be a large obstacle. Future work toward establishing standardized testing data for this challenging problem could ameliorate these difficulties and speed research progress.

4) *Limitations*: Our segmentation framework has several limitations that result from trade-offs made to keep our algo-

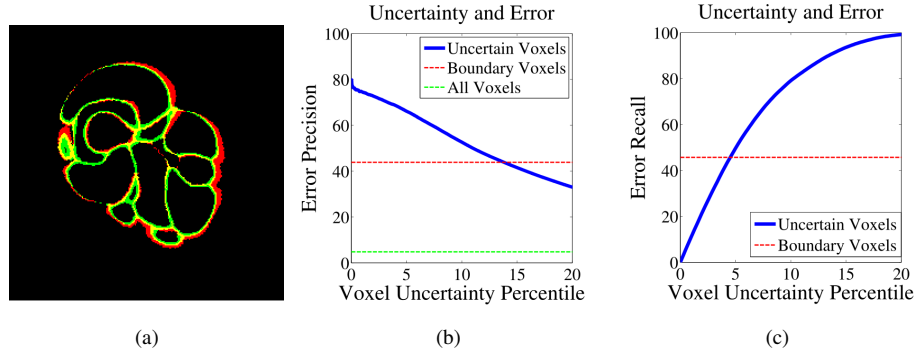


Fig. 13. (Color Figure) An analysis of uncertainty information and error. (a) shows the voxels of top 5% most uncertainty (green), the erroneous voxels (red), and their intersection (yellow). (b) and (c) show the “error precision” (yellow/(yellow + green)) and “error recall” (yellow/(yellow + red)) for percentiles of uncertainty across all images. A general strategy for identifying errors in a crisp segmentation might be to simply take boundary voxels as being likely erroneous, so we present this as a baseline in (b) and (c). We see that the top 5% of voxels have about 20% higher precision but the same recall as the boundary voxels, indicating uncertainty is useful for error detection.




Fig. 14. Our data (a) compared to data used by Baudin et al. (b).

rithm fully automatic and to accurately model thigh muscle variability.

The GLR transformation only imposes soft adjacency constraints on the labels, rather than hard constraints that would be imposed by deformable models [17], [18] or atlas-based segmentation [20]. Such soft constraints are necessary for probabilistic segmentations, and allow pathologies to be more easily segmented.

Because our shape model is poor at capturing pose variability, it would be difficult to apply our technique to data sets involving limb articulation without advanced pre-processing of the images to remove articulation variability.

Our segmentation technique is designed to run automatically without any human interaction, and thus accuracy rather than speed was our primary goal. Using unoptimized MATLAB code run on a machine with 2 Quad Core Intel Xeon 2.33 GHz CPUs, our segmentation method took 50 ± 4.3 minutes per image to run, compared to the RW with shape model algorithm which took 13 ± 1.2 minutes per image. Gilles et al. report segmentation run times in the 2 to 8 minute range for their different algorithms [17], [18]. Interactive algorithms that do not use shape models but instead rely on user input run even faster, for example standard RW with user input seeds requires only 2.5 ± 0.2 minutes to run. Requiring that a user interact with the algorithm means that time spent computing a

segmentation is the valuable time of a human expert, whereas an automatic algorithm run before any user input is given is only spending relatively cheap computation time.

Our shape model is also fairly memory intensive, but it uses memory comparable to other techniques that model shape variability, such as RW-Shape. If memory is an issue, it can be mitigated by using fewer PCA modes and possibly downsampling the images when constructing the shape model.

VII. CONCLUSION

In this paper, we introduce the GLR transformation for representing probabilistic segmentations. Our GLR transformation implicitly encodes label transition penalties, similar to penalties encoded explicitly by other techniques [35]. Once the GLR transformation has been applied, the penalties are automatically enforced by standard gradient and distance terms. The GLR transformation can be designed to ensure statistical shape models capture variability in smaller structures accurately. In the future, we will explore other techniques for incorporating prior information into the GLR transformation.

We used our GLR representation as a basis for a novel fully automatic thigh muscle segmentation scheme. The GLR representation and corresponding statistical shape model improve the segmentation accuracy of smaller muscles that are usually more difficult to detect. In the context of thigh muscle segmentation, we provide several other contributions. Our anatomically-based alignment scheme leverages a relatively easy to solve subproblem (presegmentation) to perform an automatic rigid alignment. Non-rigid alignments often introduce user-dependence or free parameters into the algorithm, and it is not clear what shape variability will remain to be captured by PCA shape models. Our RF boundary detection technique learns the appearance of both true and false boundaries.

In the future, we will extend the RF to learn boundaries between specific pairs of labels, and include this information in our energy term. In order to do this, we will explore including more complex filters into the training process. Another direction of future research will be to combine RW like energy terms that perform well for muscles with strong intermuscular

boundaries with our energy terms, which focus on muscles with weak boundaries, and switch between these types of terms in a spatially adaptive manner [46], [47]. We will also explore extending other popular segmentation frameworks to this challenging problem.

Our preliminary results indicate that statistical shape models can capture variability in the higher order shape moments between healthy and COPD thigh muscles. For future work, we will investigate how these shape models are useful for subsequent analysis of the COPD effects on these muscles, leading to better diagnosis and treatment planning.

REFERENCES

- [1] A. Yazdanpanah, G. Hamarneh, B. Smith, and M. Sarunic, "Intra-retinal layer segmentation in optical coherence tomography using an active contour approach," in *MICCAI*, 2009, pp. 649–656.
- [2] M. K. Garvin, M. D. Abramoff, R. Kardon, S. R. Russell, X. Wu, and M. Sonka, "Intraretinal layer segmentation of macular optical coherence tomography images using optimal 3-D graph search," *IEEE TMI*, vol. 27, no. 10, pp. 1495–1505, 2008.
- [3] M. Nosrati, S. Andrews, and G. Hamarneh, "Bounded labeling function for global segmentation of multi-part objects with geometric constraints," in *ICCV*, 2013, pp. 2032–2039.
- [4] A. Delong and Y. Boykov, "Globally optimal segmentation of multi-region objects," in *ICCV*, 2009, pp. 285–292.
- [5] I. Bloch, "Fuzzy spatial relationships for image processing and interpretation: a review," *Image and Vision Computing*, vol. 23, no. 2, pp. 89–110, 2005.
- [6] A. Saad, G. Hamarneh, and T. Moller, "Exploration and visualization of segmentation uncertainty using shape and appearance prior information," *IEEE Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1366–1375, 2010.
- [7] A. Saad, T. Möller, and G. Hamarneh, "ProbExplorer: Uncertainty-guided exploration and editing of probabilistic medical image segmentation," in *Computer Graphics Forum*, vol. 29, no. 3. Wiley Online Library, 2010, pp. 1113–1122.
- [8] D. Saey, R. Debigaré, P. LeBlanc, M. J. Mador, C. H. Côté, J. Jobin, and F. Maltais, "Contractile leg fatigue after cycle exercise: a factor limiting exercise in patients with chronic obstructive pulmonary disease," *American journal of respiratory and critical care medicine*, vol. 168, no. 4, pp. 425–430, 2003.
- [9] A. M. Schols, R. Broekhuizen, C. A. Weling-Scheepers, and E. F. Wouters, "Body composition and mortality in chronic obstructive pulmonary disease," *The American Journal of Clinical Nutrition*, vol. 82, no. 1, pp. 53–59, 2005.
- [10] J. Vestbo, E. Prescott, T. Almdal, M. Dahl, B. G. Nordestgaard, T. Andersen, T. I. Sørensen, and P. Lange, "Body mass, fat-free body mass, and prognosis in patients with chronic obstructive pulmonary disease from a random population sample: findings from the Copenhagen City Heart Study," *American journal of respiratory and critical care medicine*, vol. 173, no. 1, pp. 79–83, 2006.
- [11] J. Seymour, M. Spruit, N. Hopkinson, S. Natanek, W.-C. Man, A. Jackson, H. Gosker, A. Schols, J. Moxham, M. Polkey *et al.*, "The prevalence of quadriceps weakness in COPD and the relationship with disease severity," *European Respiratory Journal*, vol. 36, no. 1, pp. 81–88, 2010.
- [12] F. Maltais, M. Decramer, R. Casaburi, E. Barreiro, Y. Burelle, R. Debigaré, P. R. Dekhuijzen, F. Franssen, G. Gayan-Ramirez, J. Gea *et al.*, "An official American thoracic society/European respiratory society statement: Update on limb muscle dysfunction in chronic obstructive pulmonary disease," *American Journal of Respiratory and Critical Care Medicine*, vol. 189, no. 9, pp. e15–e62, 2014.
- [13] B. HajGhanbari, G. Hamarneh, N. Changizi, A. Ward, and W. Reid, "MRI-based 3D shape analysis of thigh muscles: Patients with chronic obstructive pulmonary disease versus healthy adults," *Academic Radiology*, vol. 18, no. 2, pp. 155–166, 2011.
- [14] H. Mirzaalian, G. Hamarneh, B. HajGhanbari, and W. D. Reid, "3D shape analysis of the knee extensor and flexor muscles in patients with COPD using mesh projection-based features," in *MICCAI MeshMed*, 2011, pp. 1–8.
- [15] B. Tracy, F. Ivey, E. J. Metter, J. Fleg, E. Siegel, and B. Hurley, "A more efficient magnetic resonance imaging-based strategy for measuring quadriceps muscle volume," *Medicine & Science in Sports & Exercise*, vol. 35, no. 3, p. 425, 2003.
- [16] C. Engstrom, G. Loeb, J. Reid, W. Forrest, and L. Avruch, "Morphometry of the human thigh muscles. a comparison between anatomical sections and computer tomographic and magnetic resonance images," *Journal of anatomy*, vol. 176, p. 139, 1991.
- [17] B. Gilles and D. K. Pai, "Fast musculoskeletal registration based on shape matching," in *MICCAI*, 2008, pp. 822–829.
- [18] B. Gilles and N. Magnenat-Thalmann, "Musculoskeletal MRI segmentation using multi-resolution simplex meshes with medial representations," *Medical Image Analysis*, vol. 14, no. 3, pp. 291–302, 2010.
- [19] A. Gubern-Mérida, M. Kallenberg, R. Martí, and N. Karssemeijer, "Segmentation of the pectoral muscle in breast MRI using atlas-based approaches," in *MICCAI*, 2012, pp. 371–378.
- [20] E. Ahmad, M. H. Yap, H. Degens, and J. S. McPhee, "Atlas-registration based image segmentation of MRI human thigh muscles in 3D space," in *SPIE Medical Imaging*. International Society for Optics and Photonics, 2014, pp. 90371L–90371L.
- [21] B. Ibragimov, J. L. Prince, E. Z. Murano, J. Woo, M. Stone, B. Likar, F. Pernuš, and T. Vrtovec, "Segmentation of tongue muscles from super-resolution magnetic resonance images," *Medical Image Analysis*, 2014.
- [22] S. Essafi, G. Lings, J. Deux, A. Rahmouni, G. Bassez, and N. Paragios, "Wavelet-driven knowledge-based MRI calf muscle segmentation," in *ISBI*, 2009, pp. 225–228.
- [23] C. Wang, O. Teboul, F. Michel, S. Essafi, and N. Paragios, "3D knowledge-based segmentation using pose-invariant higher-order graphs," *MICCAI*, pp. 189–196, 2010.
- [24] P.-Y. Baudin, N. Azzabou, P. G. Carlier, and N. Paragios, "Automatic skeletal muscle segmentation through random walks and graph-based seed placement," in *IEEE ISBI*, 2012, pp. 1036–1039.
- [25] L. Grady, "Random walks for image segmentation," *IEEE TPAMI*, vol. 28, no. 11, pp. 1768–1783, 2006.
- [26] P.-Y. Baudin, N. Azzabou, P. G. Carlier, N. Paragios *et al.*, "Manifold-enhanced segmentation through random walks on linear subspace priors," in *British Machine Vision Conference*, 2012.
- [27] P.-Y. Baudin, N. Azzabou, P. G. Carlier, and N. Paragios, "Prior knowledge, random walks and human skeletal muscle segmentation," in *MICCAI*. Springer, 2012, pp. 569–576.
- [28] P.-Y. Baudin, "Graph-based segmentation of skeletal striated muscles in NMR images," Ph.D. dissertation, Center for Computer Vision, Ecole Centrale Paris, 2013.
- [29] D. Cremers, F. Schmidt, and F. Barthel, "Shape priors in variational image segmentation: Convexity, Lipschitz continuity and globally optimal solutions," in *CVPR*, 2008, pp. 1–6.
- [30] J. Egozcue, V. Pawlowsky-Glahn, G. Mateu-Figueras, and C. Barceló-Vidal, "Isometric logratio transformations for compositional data analysis," *Mathematical Geology*, vol. 35, no. 3, pp. 279–300, 2003.
- [31] S. Andrews, N. Changizi, and G. Hamarneh, "The isometric log-ratio transform for probabilistic multi-label anatomical shape representation," *IEEE TMI*, Accepted, 2014.
- [32] K. Pohl, J. Fisher, S. Bouix, M. Shenton, R. McCarley, W. Grimson, R. Kikinis, and W. Wells, "Using the logarithm of odds to define a vector space on probabilistic atlases," *Medical Image Analysis*, vol. 11, no. 5, pp. 465–477, 2007.
- [33] S. Andrews, C. McIntosh, and G. Hamarneh, "Convex multi-region probabilistic segmentation with shape prior in the isometric log-ratio transformation space," in *ICCV*, 2011, pp. 2096–2103.
- [34] N. Changizi and G. Hamarneh, "Probabilistic multi-shape representation using an isometric log-ratio mapping," in *MICCAI*, 2010, pp. 563–570.
- [35] J. Lellmann, F. Becker, and C. Schnorr, "Convex optimization for multi-class image labeling with a novel family of total variation based regularizers," in *ICCV*, 2009, pp. 646–653.
- [36] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [37] S. Andrews, G. Hamarneh, A. Yazdanpanah, B. HajGhanbari, and W. Reid, "Probabilistic multi-shape segmentation of knee extensor and flexor muscles," in *MICCAI*, 2011, pp. 651–658.
- [38] A. Chambolle, V. Caselles, M. Novaga, D. Cremers, T. Pock *et al.*, "An introduction to total variation for image analysis," *Lecture Notes for the summer school on sparsity in Linz, Austria, August*, 2009.
- [39] E. Strelakovsky and D. Cremers, "Generalized ordering constraints for multilabel optimization," in *ICCV*, 2011, pp. 2619–2626.
- [40] R. T. Rockafellar, "Augmented Lagrangians and applications of the proximal point algorithm in convex programming," *Mathematics of operations research*, vol. 1, no. 2, pp. 97–116, 1976.
- [41] T. Pock, T. Schoenemann, G. Graber, H. Bischof, and D. Cremers, "A convex formulation of continuous multi-label problems," in *ECCV*, 2008, pp. 792–805.

- [42] L. National Heart and B. Institute, "Global initiative for obstructive lung disease. Global strategy for the diagnosis, management and prevention of chronic obstructive pulmonary disease." pp. 6–8, 2004.
- [43] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig, "User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability," *Neuroimage*, vol. 31, no. 3, pp. 1116–1128, 2006.
- [44] A. Top, G. Hamarneh, and R. Abugharbieh, "Active learning for interactive 3D image segmentation," in *MICCAI*, 2011, pp. 603–610.
- [45] J. G. Sled, A. P. Zijdenbos, and A. C. Evans, "A nonparametric method for automatic correction of intensity nonuniformity in MRI data," *IEEE TMI*, vol. 17, no. 1, pp. 87–97, 1998.
- [46] J. Rao, R. Abugharbieh, and G. Hamarneh, "Adaptive regularization for image segmentation using local image curvature cues," in *ECCV 2010*. Springer, 2010, pp. 651–665.
- [47] S. R. Thiruvankadam, T. F. Chan, and B.-W. Hong, "Segmentation under occlusions using selective shape prior," in *Scale Space and Variational Methods in Computer Vision*. Springer, 2007, pp. 191–202.

APPENDIX A LOG-RATIO CALCULATIONS

An LR transformation for $\mathbf{q} \in \mathbb{P}_K$ is given by

$$\text{LR}(\mathbf{q}) = A \log(\mathbf{q}) \quad (\text{A.1})$$

$$A\mathbf{1} = \mathbf{0}, \quad (\text{A.2})$$

where A is a $(K-1) \times K$ matrix. Recall that $\log(\cdot)$ of a vector denotes component-wise application. LR transformations of this form map perturbation and power transformation to vector addition and scalar multiplication. That is, for $\mathbf{p}, \mathbf{q} \in \mathbb{P}_K$ and $\alpha \in \mathbb{R}$,

$$\text{LR}(\mathbf{p} \oplus \mathbf{q}) \quad (\text{A.3})$$

$$= \text{LR}(\mathcal{C}([p_1 q_1, \dots, p_K q_K]^\top)) \quad (\text{A.4})$$

$$= A \log \left(\frac{1}{Z} [p_1 q_1, \dots, p_K q_K]^\top \right) \quad (\text{A.5})$$

$$= A([\log(p_1) + \log(q_1), \dots, \log(p_K) + \log(q_K)]^\top) - \log(Z)A\mathbf{1} \quad (\text{A.6})$$

$$= A \log(\mathbf{p}) + A \log(\mathbf{q}) - \mathbf{0}, \quad \text{from (A.2)} \quad (\text{A.7})$$

$$= \text{LR}(\mathbf{p}) + \text{LR}(\mathbf{q}) \quad (\text{A.8})$$

$$\text{LR}(\alpha \odot \mathbf{q}) \quad (\text{A.9})$$

$$= \text{LR} \left(\mathcal{C}([q_1^\alpha, \dots, q_K^\alpha]^\top) \right) \quad (\text{A.10})$$

$$= A \log \left(\frac{1}{Z} [q_1^\alpha, \dots, q_K^\alpha]^\top \right) \quad (\text{A.11})$$

$$= A(\alpha[\log(q_1), \dots, \log(q_K)]^\top) - \log(Z)A\mathbf{1} \quad (\text{A.12})$$

$$= \alpha A \log(\mathbf{q}) - \mathbf{0}, \quad \text{from (A.2)} \quad (\text{A.13})$$

$$= \alpha \text{LR}(\mathbf{q}). \quad (\text{A.14})$$

Z is the normalizing scalar associated with $\mathcal{C}(\cdot)$ that ensures a vector sums to 1.

From (A.2), A has 0 as an eigenvalue with corresponding right eigenvector $\mathbf{1}$. We assume A has no other zero eigenvalues, i.e. its columns are all linearly independent. If A^+ is the pseudo-inverse of A , then

$$A^+ A = I - \frac{\mathbf{1}\mathbf{1}^\top}{K}, \quad (\text{A.15})$$

where I is the identity matrix. This allows us to define the inverse of a LR transformation:

$$\text{LR}^{-1}(\text{LR}(\mathbf{q})) \quad (\text{A.16})$$

$$= \mathcal{C}(\exp(A^+ \text{LR}(\mathbf{q}))) \quad (\text{A.17})$$

$$= \mathcal{C}(\exp(A^+ A \log(\mathbf{q}))) \quad (\text{A.18})$$

$$= \mathcal{C} \left(\exp \left(\left(I - \frac{\mathbf{1}\mathbf{1}^\top}{K} \right) \log(\mathbf{q}) \right) \right) \quad (\text{A.19})$$

$$= \mathcal{C} \left(\begin{bmatrix} \exp(\log(q_1)) / \exp(\mathbf{1}^\top \log(\mathbf{q})/K) \\ \vdots \\ \exp(\log(q_K)) / \exp(\mathbf{1}^\top \log(\mathbf{q})/K) \end{bmatrix} \right) \quad (\text{A.20})$$

$$= \mathcal{C} \left(\begin{bmatrix} \frac{q_1}{\beta} \\ \vdots \\ \frac{q_K}{\beta} \end{bmatrix} \right), \quad \text{where } \beta = \exp \left(\frac{\mathbf{1}^\top \log(\mathbf{q})}{K} \right) \quad (\text{A.21})$$

$$= \mathbf{q}, \quad \text{since } \mathcal{C}(\cdot) \text{ removes scaling factors.} \quad (\text{A.22})$$

Expanding the derivation of (9) from Sec. II-A, we let $\mathbf{r}^i = \mathcal{C}(\mathbf{e}^i + \epsilon \cdot \mathbf{1}) \in \mathbb{P}_K$ for $i \in \{1, \dots, K\}$, where \mathbf{e}^i is the i^{th} standard basis vector and ϵ is a small scalar, and we let $\delta_{i,j}$ be the Kronecker delta. Then

$$\text{LR}(\mathbf{r}^i) = A \log(\mathcal{C}(\mathbf{e}^i + \epsilon \cdot \mathbf{1})) \quad (\text{A.23})$$

$$= A \log \left(\frac{1}{Z} \left(\frac{1}{\epsilon} \cdot \mathbf{e}^i + \mathbf{1} \right) \right) \quad (\text{A.24})$$

$$= A \log \left(\frac{1}{\epsilon} \cdot \mathbf{e}^i + \mathbf{1} \right) - A \cdot (\log(Z) \cdot \mathbf{1}) \quad (\text{A.25})$$

$$= A \log \left(\frac{1}{\epsilon} \cdot \mathbf{e}^i + \mathbf{1} \right) \quad (\text{A.26})$$

$$= [\mathbf{a}^1, \dots, \mathbf{a}^K] \begin{bmatrix} \log(\delta_{1,i}/\epsilon + 1) \\ \vdots \\ \log(\delta_{K,i}/\epsilon + 1) \end{bmatrix} \quad (\text{A.27})$$

$$= \sum_{k=1}^K \log(\delta_{k,i}/\epsilon + 1) \mathbf{a}^k \quad (\text{A.28})$$

$$= \log \left(\frac{1 + \epsilon}{\epsilon} \right) \mathbf{a}^i \quad (\text{A.29})$$

$$\approx \log \left(\frac{1}{\epsilon} \right) \mathbf{a}^i. \quad (\text{A.30})$$

APPENDIX B EUCLIDEAN LABEL TRANSITION PENALTIES

Let D be a symmetric $K \times K$ matrix with 0's on the diagonal, calculated from label transition penalties using (11). As described in Sec. II-B, if $T = -\frac{1}{2}CDC$ is positive semidefinite, where $C = I - \frac{1}{K}\mathbf{1}\mathbf{1}^\top$, D encodes a Euclidean distance metric. Further, if $T = A^\top A$ and $A = [\mathbf{a}^1, \dots, \mathbf{a}^K]$, then $\|\mathbf{a}^i - \mathbf{a}^j\|^2 = D_{i,j}$.

If T is not positive semi-definite, we advocate adding the positive constant (-2ψ) to the off-diagonal components of D , where $\psi < 0$ is the smallest eigenvalue of T . This technique ensures a Euclidean distance matrix and preserves the relative ordering of the label transition penalties. We denote the matrix $\hat{D} = D - 2\psi(\mathbf{1}\mathbf{1}^\top - I)$ and $\hat{T} = -\frac{1}{2}C\hat{D}C$ and show that \hat{D}

corresponds to a Euclidean distance by showing \hat{T} is positive definite. First, we establish several properties of C :

$$C\mathbf{1} = \mathbf{1} - \frac{1}{K}\mathbf{1}\mathbf{1}^\top \mathbf{1} \quad (\text{B.31})$$

$$= \mathbf{1} - \frac{K}{K}\mathbf{1} \quad (\text{B.32})$$

$$= \mathbf{0} \quad (\text{B.33})$$

$$CC = I - \frac{2}{K}\mathbf{1}\mathbf{1}^\top + \frac{1}{K^2}\mathbf{1}\mathbf{1}^\top \mathbf{1}\mathbf{1}^\top \quad (\text{B.34})$$

$$= I - \frac{2}{K}\mathbf{1}\mathbf{1}^\top + \frac{K}{K^2}\mathbf{1}\mathbf{1}^\top \quad (\text{B.35})$$

$$= C. \quad (\text{B.36})$$

Now, since $C\mathbf{1} = \mathbf{0}$, $\mathbf{1}$ is a zero eigenvalue of T corresponding to the eigenvector $\mathbf{1}$. Let $T = Q\Psi Q^\top$ be the eigenvector decomposition of T , with Q a matrix of orthonormal eigenvectors and Ψ a diagonal matrix of eigenvalues. Then

$$\hat{T} = -\frac{1}{2}C(\hat{D})C \quad (\text{B.37})$$

$$= -\frac{1}{2}(CDC - 2\psi C\mathbf{1}\mathbf{1}^\top C - 2\psi CC) \quad (\text{B.38})$$

$$= -\frac{1}{2}(CDC + 2\psi C) \quad (\text{B.39})$$

$$= Q\Psi Q^\top - \psi I + \frac{\psi}{K}\mathbf{1}\mathbf{1}^\top \quad (\text{B.40})$$

$$= Q(\Psi - \psi I)Q^\top. \quad (\text{B.41})$$

$Q(\Psi - \psi I)Q^\top$ is an eigenvector decomposition of \hat{T} , and since all of the diagonal components of $(\Psi - \psi I)$ are greater than or equal to zero (the smallest component of Ψ is ψ), \hat{T} is positive semi-definite.

APPENDIX C PRIMAL DUAL PROXIMAL POINT

In Sec. IV, we solve an optimization problem of the form

$$\min_{\phi} \max_{P \in C_P} E_{\text{cnv}}(\phi) + E_{\text{bnd-tot}}(\phi, P) \quad (\text{C.42})$$

$$C_P = \left\{ P = [\mathbf{p}^1, \dots, \mathbf{p}^{K-1}] \mid \mathbf{p}^i(x) \in C(x), \quad (\text{C.43}) \right.$$

$$\left. \forall i \in \{1, \dots, K-1\}, x \in \Omega \right\}, \quad (\text{C.44})$$

where $C(\cdot)$ are the convex sets defined in (24). E_{cnv} encompasses the convex regional and shape energy terms, and can be expressed as a quadratic:

$$E_{\text{cnv}}(\phi) = \frac{1}{2}\phi^\top M\phi + \mathbf{b}^\top \phi, \quad (\text{C.45})$$

where M is a positive semi-definite matrix and $\mathbf{b} \in \mathbb{R}^{n \cdot (K-1)}$. $E_{\text{bnd-tot}}$ equals E_{bnd} with the implicit maximization P removed:

$$E_{\text{bnd-tot}}(\phi, P) = \sum_{x \in \Omega} \sum_{i=1}^{K-1} \mathbf{p}^i(x)^\top \nabla \phi^i(x) \quad (\text{C.46})$$

$$= \sum_{x \in \Omega} \sum_{i=1}^{K-1} (\text{div } \mathbf{p}^i(x)) \phi^i(x), \quad (\text{C.47})$$

where div is the divergence operator [38].

We solve (C.42) using a proximal point method [40], [41]. The idea is to iteratively optimize for the primal variables ϕ and the dual variables P , adding a soft constraint to prevent the solutions from changing too much in a given step. Given solutions $\phi_j = [\phi_j^1, \dots, \phi_j^{K-1}]$ and $P_j = [\mathbf{p}_j^1, \dots, \mathbf{p}_j^{K-1}]$ at iteration j , the primal step is

$$\phi_{j+1} = \arg \min_{\phi} E_{\text{cnv}}(\phi) + E_{\text{bnd-tot}}(\phi, P_j) + \frac{\|\phi - \phi_j\|^2}{2\tau_1}. \quad (\text{C.48})$$

With a fixed P_j , $E_{\text{bnd-tot}}(\phi, P_j)$ is linear in ϕ (from (C.47)), so (C.48) is a convex quadratic and can be globally optimized.

The dual step is:

$$P_{j+1} = \arg \max_{P \in C_P} E_{\text{bnd-tot}}(\phi_{j+1}, P) - \frac{\|P - P_j\|^2}{2\tau_2} \quad (\text{C.49})$$

$$\|P - P_j\|^2 = \sum_{i=1}^{K-1} \|\mathbf{p}^i - \mathbf{p}_j^i\|^2. \quad (\text{C.50})$$

With a fixed ϕ_{j+1} , $E_{\text{bnd-tot}}(\phi_{j+1}, P)$ is linear in P (from (C.46)), so (C.49) is a concave quadratic. C_P is a hyper-ellipse, and thus a convex set that can be projected onto. Thus the dual step can also and can be globally optimized.

This algorithm converges for small enough τ_1 and τ_2 [41]; empirically we found $\tau_p = \tau_q = 0.1$ provided fast convergence.