

Using the logarithm of odds to define a vector space on probabilistic atlases

Kilian M. Pohl^{a,b,*}, John Fisher^a, Sylvain Bouix^c, Martha Shenton^c,
Robert W. McCarley^d, W. Eric L. Grimson^a, Ron Kikinis^b, William M. Wells^{a,b}

^a *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*

^b *Surgical Planning Laboratory, Harvard Medical School and Brigham and Women's Hospital, 75 Francis Street, Boston, MA 02115, USA*

^c *Psychiatry Neuroimaging Laboratory, Department of Psychiatry, Brigham and Women's Hospital, Harvard Medical School, 1249 Boylston Street, Boston, MA 02215, USA*

^d *Department of Psychiatry, VA Boston Healthcare System and Harvard Medical School, 1400 VFW Parkway, West Roxbury, MA 02132, USA*

Received 25 May 2006; received in revised form 22 May 2007; accepted 4 June 2007

Available online 22 June 2007

Abstract

The logarithm of the odds ratio (LogOdds) is frequently used in areas such as artificial neural networks, economics, and biology, as an alternative representation of probabilities. Here, we use LogOdds to place probabilistic atlases in a linear vector space. This representation has several useful properties for medical imaging. For example, it not only encodes the shape of multiple anatomical structures but also captures some information concerning uncertainty. We demonstrate that the resulting vector space operations of addition and scalar multiplication have natural probabilistic interpretations.

We discuss several examples for placing label maps into the space of LogOdds. First, we relate signed distance maps, a widely used implicit shape representation, to LogOdds and compare it to an alternative that is based on smoothing by spatial Gaussians. We find that the LogOdds approach better preserves shapes in a complex multiple object setting. In the second example, we capture the uncertainty of boundary locations by mapping multiple label maps of the same object into the LogOdds space. Third, we define a framework for non-convex interpolations among atlases that capture different time points in the aging process of a population.

We evaluate the accuracy of our representation by generating a deformable shape atlas that captures the variations of anatomical shapes across a population. The deformable atlas is the result of a principal component analysis within the LogOdds space. This atlas is integrated into an existing segmentation approach for MR images. We compare the performance of the resulting implementation in segmenting 20 test cases to a similar approach that uses a more standard shape model that is based on signed distance maps. On this data set, the Bayesian classification model with our new representation outperformed the other approaches in segmenting subcortical structures.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Logarithm of odds; Logistic function; Anatomical variations; Shape representation; Distance map; Bayesian modeling; Interpolation; Statistical classification

1. Introduction

Statistical shape representation in medical imaging applications is an important and challenging problem. Many anatomical structures, such as the right superior temporal gyrus shown in Fig. 1, have ambiguous boundaries in MR images as their intensity profiles are nearly indistinguishable from their neighbors. This causes variations among expert

* Corresponding author. Address: Surgical Planning Laboratory, Harvard Medical School and Brigham and Women's Hospital, 75 Francis Street, Boston, MA 02115, USA.

E-mail address: pohl@csail.mit.edu (K.M. Pohl).

URL: www.csail.mit.edu/~pohl (K.M. Pohl).

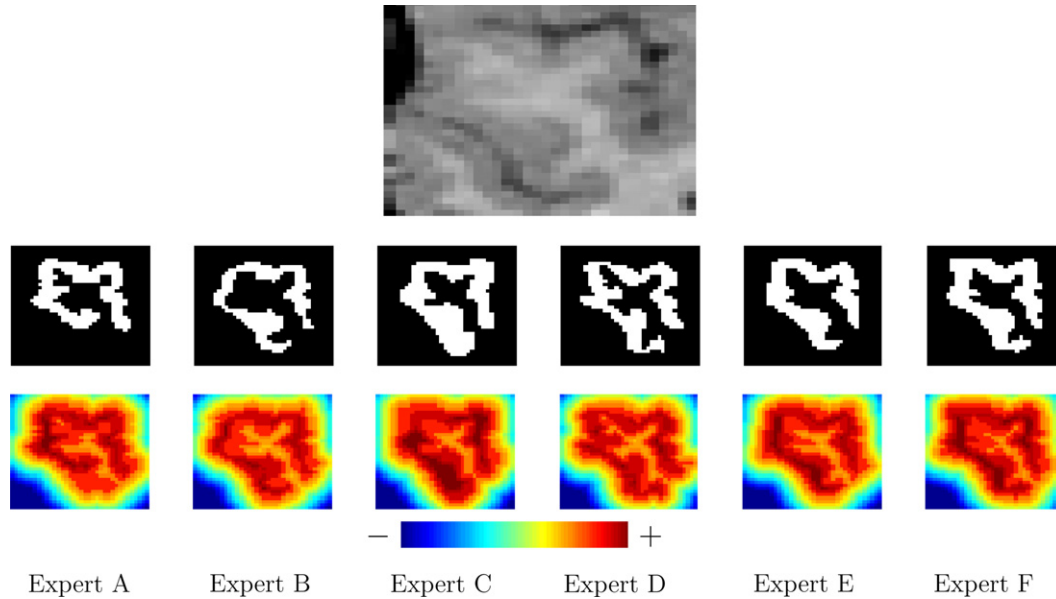


Fig. 1. The MR image on the top shows the right superior temporal gyrus. The corresponding segmentations by six experts (A–F) are shown below. Significant difference between the segmentations are visible. The third row shows the corresponding signed distance maps that can capture the boundary of each segmentation but not the uncertainty about the boundary location across the raters.

segmentations of the same structure (Fig. 1, second row). Capturing the intrinsic properties of the shape in this context is difficult, as flexible and efficient statistical models are needed to represent both the properties of the shape and its variability.

The approaches to shape representation in this context may be broadly categorized as *explicit* or *implicit*. In an explicit framework, the shape is represented by a set of, usually connected, primitives (e.g., points, triangles, medial nodes) that model the object. The model is then augmented with statistical information capturing the variability of shapes within a class or population. This approach has been used in many applications, including point distribution models (Bookstein, 1996), active shape models (Cootes et al., 1994), medial representations (Styner et al., 2004; Pizer et al., 2003), or spherical harmonics (Brechtbühler et al., 1995). The explicit approach directly represents the surface of the shape, however, implementations may require significant application-specific attention, for example to generate suitable meshes for representing an object.

In the implicit category, level set functions are used to model anatomical structures. For example, Signed Distance Maps (SDMs) represent shape by defining the boundary as a zero-level set and the values of the remaining voxels by their shortest (usually Euclidean) distance to the boundary (Leventon, 2000; Tsai et al., 2003; Yang et al., 2004; Pohl et al., 2006; Golland et al., 2005; Kohlberger et al., 2006). The representation is simple to implement, and to some extent, obviates the need for establishing correspondences among objects, though the density of the representation may incur additional computational burden.

A single SDM for an object, however, does not encode variability such as segmentation disagreements among experts. A popular solution is to perform Principal Com-

ponents Analysis (PCA) over a set of example SDMs and then fit a statistical model to the resulting PCA coefficients (Cootes et al., 1998). The resulting model can be used in a variety of tasks including segmentation (Leventon, 2000; Tsai et al., 2003; Yang et al., 2004), or longitudinal studies of shape variation (Kohlberger et al., 2006). The primary advantage of the approach is that it projects high-dimensional SDMs into a lower-dimensional representation that provides efficient statistical modeling and inference. One major drawback of this approach as noted by Golland et al. (2005) is that it is not obvious how to impose a vector space structure on SDMs (i.e. defining vector operations that are closed under the set of SDMs). For example, if we interpret SDMs as vectors of real numbers then the addition of two SDMs generally does not lead to an SDM. This is usually dealt with by projecting samples from the distribution given by the PCA coefficients back onto the manifold of valid SDMs (Golland et al., 2005).

In this paper, we present a new shape representation, called LogOdds, that embeds SDMs in a vector space and relates them to Probabilistic Atlases (PA) that define a probability of a label being present throughout the image domain. For the probability p of a binary variable, the LogOdds (also called logit) is the logarithm of the ratio between the probability p and its complement $1 - p$. It is a well established technique in areas such as neural networks (Minsky and Papert, 1988), economics (McFadden, 1973), and logistic regression (Giudici, 2003).

We relate LogOdds to the certainty of objects' boundaries in images. Like SDMs, LogOdds encode the boundary of the shape via a zero-level set that now represents the set of voxels with the highest uncertainty of being assigned to fore- or background. Unlike SDMs, the rest of the space is defined by the logarithm of the odds of a

structure to be present at that location under the assumption that voxels in an image are independently distributed. This relationship with the odds of the presence of an anatomical label provides a natural way to capture boundary uncertainty. Importantly, the space of LogOdds is closed under addition and scalar multiplication, and as such, it can be used for efficient and straightforward statistical modeling and inference of shape.

We note that a variety of models in medical imaging depend on the logarithm transformations, such as the scalar logarithm on the determinants of deformation tensors (Ashburner and Friston, 2000), and the matrix logarithm on tensors (Arsigny et al., 2007). Our work was inspired by these approaches, although, to our knowledge, it is the first time LogOdds are utilized in the context of shape description.

This article is organized as follows: in Section 2, we provide the mathematical definition and properties of LogOdds, as well as their relationship with PA. In Section 3, we discuss examples for mapping label maps into the LogOdds space. First, we show how a single label map can be transformed into a LogOdds representation. Second, we show how the uncertainty associated with manual tracings can be captured in the LogOdds space. Third, we show how the vector space properties of LogOdds can be used to create a continuous PA of an aging brain. In Section 4, we incorporate our new shape model into an Expectation-Maximization (EM) segmentation algorithm (Wells et al., 1996). The shape model is obtained by performing PCA on LogOdds maps of manually-segmented structures of the brain. At this point we assume that the trainings set consists of aligned segmentations so that PCA captures the variability within the label maps after alignment. Twenty subjects are segmented using our LogOdds shape model, an SDM based shape model and a precomputed PA as suggested by Van Leemput et al. (1999), Pohl et al. (2004). The quality of each segmentation technique is evaluated against manual segmentations by human experts. Overall, the LogOdds shape model helps to achieve higher accuracy than the other two representations.

2. LogOdds and its properties

In this section we review important properties of the (multinomial) LogOdds representation. Medical imaging often makes use of PA of anatomical structures using discrete distributions. Discrete distributions are defined with respect to random variables that take on several discrete values. Unfortunately, these discrete distributions are difficult to combine (e.g. to create statistical models) as the space of discrete distributions is not closed under addition and scalar multiplications. One can, however, establish a one-to-one mapping between these distributions and the space of LogOdds, which is a vector space. One can thus perform standard arithmetic or statistics in this space and map the results back into the space of discrete distribu-

tions. The next few sections explain how this mapping is defined. Moreover, we show how one can define *probabilistic addition and scalar multiplication* operators that induce a useful vector space structure on discrete distributions.

2.1. An introduction to LogOdds

LogOdds are an example of a class of functions that map the space of discrete distributions (Kendall and Buckland, 1976) to Euclidean space. Let \mathbb{P}_M be the open probability simplex (the space of discrete distributions) for M labels:

$$\begin{aligned} \mathbb{P}_M &\triangleq \left\{ p \mid p = (p_1, \dots, p_M) \in (0, 1)^M \wedge \sum_{i=1, \dots, M} p_i = 1 \right\} \\ &= \left\{ p \mid p = \left(p_1, \dots, p_{M-1}, 1 - \sum_{i=1, \dots, M-1} p_i \right) \in (0, 1)^M \right\}. \end{aligned}$$

Note that \mathbb{P}_M is an $M - 1$ dimensional space as the M th entry can be computed from the first $M - 1$ entries. Furthermore, the space is open avoiding the degenerative distributions that are certain about the assignments. For the specific case of $M = 2$, \mathbb{P}_2 is the Bernoulli distribution $\mathbb{P} \triangleq \{ p \mid p \in (0, 1) \}$ (Evans et al., 2000). Many binary classification problems use the Bernoulli distribution where p represents the probability that a voxel belongs to a particular anatomical structure and its complement $\bar{p} = 1 - p$ the probability of the voxel being in the background.

The multinomial LogOdds function $\text{logit}(\cdot) : \mathbb{P}_M \rightarrow \mathbb{R}^{M-1}$ of a discrete distribution $p \in \mathbb{P}_M$ is defined as the logarithm of the ratio between the i -th and last entry of p :

$$[\text{logit}(p)]_i \triangleq \log \left(\frac{p_i}{p_M} \right),$$

with $i \in \{1, \dots, M - 1\}$. For the Bernoulli distribution, this function simplifies to the logarithm of the ratio between the probability p and its complement:

$$\text{logit}(p) \triangleq \log \left(\frac{p}{1-p} \right) = \log p - \log(1-p).$$

The inverse of the log odds function $\text{logit}(\cdot)$ is the generalized logistic function¹

$$[\mathcal{P}(t)]_i \triangleq \begin{cases} \frac{e^i}{Z}, & \text{for } i \in \{1, \dots, M - 1\} \\ \frac{1}{Z}, & \text{if } i = M \end{cases},$$

where $Z \triangleq 1 + \sum_{j=1, \dots, M-1} e^{t_j}$ is the normalization factor.

Let \mathbb{L}_{M-1} be the $M - 1$ dimensional space of LogOdds induced from \mathbb{P}_M :

$$\mathbb{L}_{M-1} \triangleq \{ \text{logit}(p) \mid p \in \mathbb{P}_M \}.$$

We note that \mathbb{L}_{M-1} is equivalent to a $(M - 1)$ D real vector space and is thus a vector space.

¹ Note, that for $M = 2$, $\mathcal{P}(\cdot)$ is also called the Sigmoid function.

2.2. The relationship between LogOdds and PA

The function $\text{logit}(\cdot)$ and its inverse comprise a homeomorphism between \mathbb{P}_M and \mathbb{L}_{M-1} so that we can borrow the vector space structure on \mathbb{L}_{M-1} to induce one on \mathbb{P}_M .

2.2.1. Addition in \mathbb{P}_M

The *probabilistic addition* $p_1 \oplus p_2$ in \mathbb{P}_M is constructed by mapping p_1 and p_2 into the LogOdds space, performing the addition between $\text{logit}(p_1)$ and $\text{logit}(p_2)$, and then mapping the result back into the \mathbb{P}_M via the logistic function. We can show that this operation is equivalent to a normalized multiplication of two discrete probabilities within \mathbb{P}_M :

$$p_1 \oplus p_2 \triangleq \mathcal{P}(\text{logit}(p_1) + \text{logit}(p_2)) = \frac{1}{\sum_{i=1, \dots, M} p_{1i} \cdot p_{2i}} (p_{11} \cdot p_{21}, \dots, p_{1M} \cdot p_{2M}). \tag{1}$$

Note that *probabilistic addition* \oplus is closed in \mathbb{P}_M so that (\mathbb{P}_M, \oplus) forms an Abelian group with the zero element being the uniform distribution $(\frac{1}{M}, \dots, \frac{1}{M})$. The additive inverse of a discrete probability $p \in \mathbb{P}_M$ is its complement \bar{p} , defined as $[\bar{p}]_i \triangleq \frac{1}{\sum_{j=1, \dots, M} \frac{1}{p_{1j}}}$, for all $i \in \{1, \dots, M\}$.

We end this section by discussing the relationship between *probabilistic addition* and Bayes' rule. Bayes' rule may be written as follows:

$$P(A = i|B) = \frac{L(A = i|B)P(A = i)}{\sum_{j=1, \dots, M} L(A = j|B)P(A = j)},$$

where $L(A = i|B) \triangleq P(B|A = i)$ is the *likelihood* of the observations conditioned on the data. When viewed as a function of i , this is probably not a probability, because, for example, it may not integrate to one. We can, if we choose, normalize the likelihood via $Z \triangleq \sum_{j=1, \dots, M} L(A = j|B)$ so that it does integrate to one, and this does not change the resulting posterior probability. In this case we may use the result of Eq. (1) to carry out the arithmetic of Bayes' rule as

$$\begin{aligned} \left[\frac{L(A|B)}{Z} \oplus P(A) \right]_i &= \frac{P(A = i) \cdot \frac{P(B|A=i)}{Z}}{\sum_{k=1, \dots, M} P(A = k) \cdot \frac{P(B|A=k)}{Z}} \\ &= \frac{P(A = i|B)}{\sum_{k=1, \dots, M} P(A = k|B)} = P(A = i|B). \end{aligned}$$

To summarize, we may obtain the LogOdds of the posterior probability of a label, given an image, by adding the LogOdds of a label-wise prior to the LogOdds of the normalized likelihood.

2.2.2. Scalar multiplication in \mathbb{P}_M

To induce a vector space structure on \mathbb{P}_M , we also need to define a scalar multiplication operator. As with the probabilistic addition, the *probabilistic scalar multiplication* $\alpha \otimes p$ in \mathbb{P}_M^n is defined as the logistic function of the product between α and the LogOdds $\text{logit}(p)$:

$$\alpha \otimes p \triangleq \mathcal{P}(\alpha * \text{logit}(p)) = \frac{1}{\sum_{i=1, \dots, M} p_i^\alpha} (p_1^\alpha, \dots, p_M^\alpha).$$

It can be shown that this is equivalent to exponentiating the discrete distribution with α and normalizing it. The technique of exponentiating and normalizing probabilities is frequently used in areas such as Markov Random fields (Besag, 1986). We now have constructed the vector space $(\mathbb{P}_M, \oplus, \otimes)$ with the identity element of the scalar multiplication being 1, the complement of p is defined as $\bar{p} = -1 \otimes p$. By construction, this vector space is equivalent to $(\mathbb{L}_{M-1}, +, *)$.

Fig. 2 shows the effect of *probabilistic scalar multiplication* on a typical PA used in imaging. A PA captures the probability of a label being present at voxels throughout the image domain (assuming voxels are independently distributed). In Fig. 2, the PA $A \in \mathbb{P}^n$ represents a circle with uncertainty associated with its boundary and the actual contour composed of the voxels x with $A_x = 0.5$. When this atlas is multiplied by α , the slope of the PA in the region of the boundary is changed (see Fig. 2 second row). When the absolute value of α is greater than 1, the slope gets steeper, but it gets smoother when $|\alpha| < 1$. The steepness of the slope also corresponds to the certainty within the boundary location as shown by the local Entropy (Shannon, 1948) plots of Fig. 2. Thus, α can be used to control the certainty of a location of a boundary.




Fig. 2. Displaying the impact of $\alpha \in \mathbb{R}$ on the results of the *probabilistic scalar multiplication* \otimes with a PA. The first row shows a 2D PA. The result of the operation with $\alpha = 0.5$ and $\alpha = 2$ are shown in the second row. When α is small the slope of the PA is gentle, indicating higher uncertainty of the boundary location as also shown by the graph of the corresponding entropy in the third row. When α is large the slope steepens and the entropy is characterized by a thinner ridge.

This completes our discussion on how the homeomorphism between LogOdds \mathbb{L}_{M-1}^n and discrete distributions \mathbb{P}_M^n induces a vector space structure on \mathbb{P}_M^n . Probabilistic addition and scalar multiplication operators can be defined and used to perform statistical computations in \mathbb{P}_M^n . We note that any other invertible function from \mathbb{P}_M^n to $\mathbb{R}^{(m-1) \times n}$ could have been used to induce a vector space structure on \mathbb{P}_M^n . However, if we use the *logit*-function, the vector operations have a particularly useful meaning. Probabilistic addition closely relates to Bayes' rule, and the scalar multiplication corresponds to the certainty of boundary locations in a PA.

3. Representing shapes via LogOdds

We now apply the LogOdds technique to represent shapes given: a single label map, a set of label maps of the same structure traced by different experts, and a set of label maps of an aging population. First, we interpret SDMs, a frequently used implicit shape representation in medical imaging, as an element of the LogOdds space. This allows us to transform SDMs to PAs through the generalized logistic function $\mathcal{P}(\cdot)$. We also present an alternative approach consisting of first creating a PA via Gaussian smoothing and then transforming it into a LogOdds through *logit*(\cdot). These representations are quite similar in the case of a simple Bernoulli distribution, but interesting differences can be observed in more complex discrete distributions. Then, we encode a set of expert segmentations of anatomical structures and their uncertainty within LogOdds maps. The final example involves the interpolation of longitudinal data capturing the progression of schizophrenia in eight patients.

In the remainder of this article, we use data provided by the Psychiatry Neuroimaging Laboratory, Brigham and Women's Hospital, Harvard. MR images are acquired with a 1.5-T General Electric scanner (GE Medical Systems, Milwaukee) and a contiguous spoiled gradient-recalled pulse sequence (repetition time = 35 ms, echo time = 5 ms, one repetition, 45° nutation angle, 24-cm field of view, number of excitations = 1, matrix = 256 × 256 [192 phase-encoding steps] × 124). Voxels were 0.9375 × 0.9375 × 1.5 mm. Data are formatted in the coronal plane and analyzed as 124 coronal 1.5-mm-thick slices.

We use the following mathematical conventions throughout this article:

- The 3D image domain is composed of n voxels with index 1 to n . The index of a voxel represents the order in which it appears by stacking the columns of the image domain on top of each other (such as Tsai et al., 2001). A volume X over the image domain is therefore seen as an n -dimensional vector, where X_j denotes the j th vector entry of X .
- $X^{(i)}$ represents a vector that is linked to the specific instance i . For example, if $X \triangleq \{X^{(A)}, \dots, X^{(F)}\}$ is a set of segmentations generated by experts A to F then $X^{(C)}$ corresponds to the segmentation of expert C .

- $\mathbb{B}_M^n \triangleq \{1, \dots, M\}^n$ is the space of label maps with M labels.
- \mathbb{L}_{M-1}^n represents the corresponding LogOdds space that captures the shapes of the label maps. Elements of \mathbb{L}_{M-1}^n are called LogOdds maps.
- \mathbb{P}_M^n represents the space of PAs (Probabilistic Atlases) that define a probability of a label being present throughout the image domain. This space assumes that the voxels of the image domain are independently distributed in \mathbb{P}_M^n .

3.1. Signed distance maps, LogOdds and probabilistic atlases

LogOdds maps in \mathbb{L}_1^n define the boundary of a shape as a zero-level set function. One subset of maps in \mathbb{L}_1^n are SDMs that also conform to the Eikonal equation (Rauch, 1991) with uniform speed. Thus, SDMs can always be interpreted as LogOdds maps, but the reverse is in general not true. The corresponding signed distance map transformation $\mathcal{D}(\cdot) : \mathbb{B}_2^n \rightarrow \mathbb{L}_1^n$ can be seen as a direct mapping between binary maps and the LogOdds space. Transformed to \mathbb{P}_2^n , these maps define probabilistic atlases, where voxels inside the object are represented by probabilities higher than 0.5 for the foreground and voxels outside with probabilities higher than 0.5 for the background. In the case of discrete data, the mapping $\mathcal{D}_{M-1} : \mathbb{B}_M^n \rightarrow \mathbb{L}_{M-1}^n$ is defined by combining a set of SDMs into a vector $\mathcal{D}_{M-1}(B) = (\mathcal{D}(B_1), \dots, \mathcal{D}(B_{M-1}))$, where $B_j \in \mathbb{B}_2^n$ is the binary map corresponding to label j in B . This mapping is illustrated for a few label maps in the second row of Fig. 3. Within the LogOdds space, SDMs not only represent shapes but also define a set of PAs in \mathbb{P}_M^n (see middle row² of Fig. 4), which are generated via the generalized logistic function $\mathcal{P}(\cdot)$. These PAs characterize the inside of an object with higher probabilities than voxels outside the object.

The proper choice of PAs depends on the application. For example, one can generate PAs directly via Gaussian smoothing of each binary map B_j . This results in maps with values between zero and one. We then have to normalize the resulting maps to create a discrete distribution map, called GAUSS, that sums to one at each voxel location. An example of such a PA is shown in the last row of Fig. 4, where we used a Gaussian filter of standard deviation 10. The corresponding LogOdds (last row of Fig. 3) are generated via the logarithm of odds function *logit*(\cdot).

In the case of Bernoulli distributions, these representations are very similar (first two columns of Figs. 3 and 4). However, once we turn to discrete distributions, stronger differences appear. The PAs generated from SDMs show distortions at the interface between the two circles (see second row of Fig. 4 (Light Gray and Dark Gray)). These distortions decrease when the distance between the two objects increases. This suggests that SDMs may not

² For the binary maps, we focused on the PAs of the foreground as the PAs of the background are simply their complement.




Fig. 3. The first row shows two binary maps and a multicategorical label map. The corresponding SDMs are shown in the second row. The contours nicely preserve the original shape. The third row shows the LogOdds map defined by the *logit* function of the Gaussian smoothed binary maps (GAUSS) (see third row of Fig. 4). These maps are very similar to the SDMs for the binary maps. For the label map of the two circles (Light Gray and Dark Gray), however, the corresponding contours of the LogOdds maps are influenced by the neighboring circle.

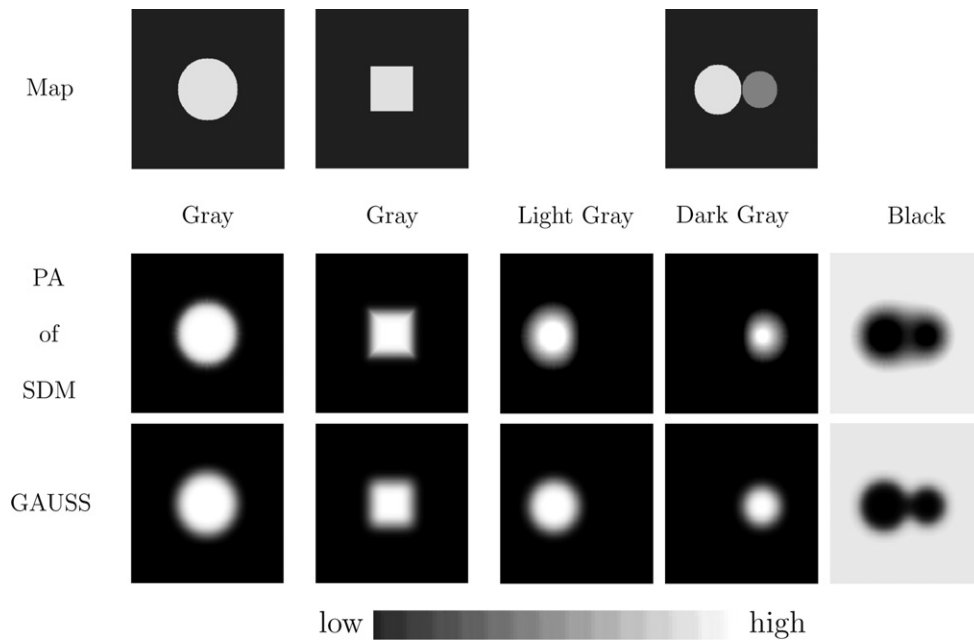


Fig. 4. The first row shows the binary and label maps of Fig. 3. The second and third rows are the PAs generated from the SDMs in Fig. 3 (second row) and the PA defined by Gaussian smoothing of the binary maps (GAUSS). While the contours of GAUSS preserve the original shape, the PAs generated from SDMs do not for the label map of the two circles (Light Gray and Dark Gray). Thus, in this example, SDMs are not well suited for capturing uncertainty about boundary location.

be the best LogOdds representation of discrete data for representing close objects. Of note, if linear operations are performed on the SDMs, the result will be a LogOdds

but likely not an SDM, making its interpretation in terms of shape difficult. For Gaussian PAs on the other hand, the two distributions are not impacted by each other

regardless of distance between the objects or the setting of the Gaussian filter. Thus, Gaussian PAs seem to better capture the shape of the two circles (see third row of Fig. 4 (Light Gray and Dark Gray)) as they were directly designed for the space of discrete probabilities.

The main advantage of the LogOdds framework is that if one wants to perform standard statistical analysis on shape, *any* form of PA can be used, as long as it is made up of valid discrete distributions.

3.2. Defining rater-specific LogOdds maps

For many anatomical structures, manually tracing its boundary from an MR image is a serious challenge, as the MR signal does not provide enough contrast to clearly see the outline of the structure. This causes variations among expert segmentations of the same structure (see Fig. 1). In this section, we show how PAs and thus LogOdds can be constructed to capture this variability.

Our representation makes use of the STAPLE algorithm presented in Warfield et al. (2006). The method takes the set of binary maps traced by experts and turns them into SDMs. It then estimates a *reference SDM* via STAPLE based on the agreement between the experts' SDMs. We compute the performance parameters of each expert \mathcal{E} by calculating the mean, $\mu^{(\mathcal{E})}$, and the variance, $\sigma^{(\mathcal{E})}$, of the voxel-wise difference between the SDM of the expert, $\mathcal{D}^{(\mathcal{E})}$, and the reference SDM, $\mathcal{D}^{(\mathcal{R})}$. The mean indicates an overall over- or under-estimation of the size of the structure by the expert, and the variance captures his ability to trace consistently. Given this model, the probability of a distance value $\mathcal{D}_x^{(\mathcal{E})}$ at voxel x for expert \mathcal{E} is defined by the Gaussian distribution $\mathcal{N}(\mathcal{D}_x^{(\mathcal{E})} - \mathcal{D}_x^{(\mathcal{R})} - \mu^{(\mathcal{E})}; \sigma^{(\mathcal{E})})$.

Assuming that the voxels in the image are independently distributed, the probability of a voxel x being inside the object given $\mathcal{D}^{(\mathcal{E})}$, $\mu^{(\mathcal{E})}$ and $\sigma^{(\mathcal{E})}$ is then defined as:

$$\begin{aligned} P(\mathcal{D}_x^{(\mathcal{R})} \geq 0 | \mathcal{D}_x^{(\mathcal{E})}, \mu^{(\mathcal{E})}, \sigma^{(\mathcal{E})}) \\ &= \int_{y \geq 0} \mathcal{N}(\mathcal{D}_x^{(\mathcal{E})} - y - \mu^{(\mathcal{E})}, \sigma^{(\mathcal{E})}) dy \\ &= \int_{y \leq \mathcal{D}_x^{(\mathcal{E})} - \mu^{(\mathcal{E})}} \mathcal{N}(y - \mu^{(\mathcal{E})}, \sigma^{(\mathcal{E})}) dy = \Phi\left(\frac{\mathcal{D}_x^{(\mathcal{E})} - \mu^{(\mathcal{E})}}{\sigma^{(\mathcal{E})}}\right), \end{aligned}$$

where $\Phi(y) \triangleq \frac{1}{2}[1 + \text{erf}(\frac{y}{\sqrt{2}})]$ is the Gaussian cumulative distribution and $\text{erf}(\cdot)$ the error function.

We can now interpret the map defined by the voxel entries $P(\mathcal{D}_x^{(\mathcal{R})} \geq 0 | \mathcal{D}_x^{(\mathcal{E})})$ as a rater-specific PA embedded in \mathbb{P}_2^m . A natural definition of the mapping function \mathcal{T} of a binary map $B^{(\mathcal{E})}$ drawn by an expert to LogOdds space would therefore be the LogOdds of the conditional probability:

$$\begin{aligned} \mathcal{T}(B^{(\mathcal{E})}, \mu^{(\mathcal{E})}, \sigma^{(\mathcal{E})})_x &\triangleq \text{logit}(P(\mathcal{D}_x^{(\mathcal{R})} \geq 0 | \mathcal{D}_x^{(\mathcal{E})}, \mu^{(\mathcal{E})}, \sigma^{(\mathcal{E})})) \\ &= \log \left[\frac{1 + \text{erf}\left(\frac{\mathcal{D}_x^{(\mathcal{E})} - \mu^{(\mathcal{E})}}{\sqrt{2}\sigma^{(\mathcal{E})}}\right)}{1 - \text{erf}\left(\frac{\mathcal{D}_x^{(\mathcal{E})} - \mu^{(\mathcal{E})}}{\sqrt{2}\sigma^{(\mathcal{E})}}\right)} \right]. \end{aligned} \quad (2)$$

Fig. 5 shows the graph of the Gaussian cumulative functions computed from manual segmentations of the superior temporal gyrus by six experts (see Fig. 1). The perfect segmentation would result in a step function (see light gray curve in Fig. 5). All experts seem to perform similarly, except for expert D who is not only far from the truth but also unreliable, and expert E who shows great consistency and high accuracy. Similar observations can be made by looking at the corresponding LogOdds maps to the right of the graph, as the slope of the map indicates the overall accuracy of the expert. The map of expert E is close to a binary map indicating a high degree of agreement to the reference standard, whereas expert D shows a much softer LogOdds map.

3.3. Defining a time continuous atlas based on a finite number of samples

Neuroscientists often carry out longitudinal studies to better understand the aging process of a population. These studies are frequently defined by a set of subjects that have been scanned at different time points. We now explore the use of the LogOdds function for interpolating longitudinal data between time points.

This example is based on a longitudinal data set consisting of eight schizophrenic patients. Each patient has been scanned three times with an average separation of 14 months between the first and second scan, and an average




Fig. 5. The performance of each of the six expert segmentations is represented by a Gaussian cumulative function (left graph). The ideal distribution is a step function (shown in light gray). Their corresponding log odds map generated using Eq. (2) can be observed on the right. Dark blue and dark red indicates high certainty that the voxel is assigned to the background and foreground respectively. All other colors represent statistical uncertainty about the assignment of the voxel.

23 months separation between the second and third scan. For each time point, we generate a PA by first aligning the 24 MR images towards their central tendency using affine transformations computed from the population registration framework of (Zöllei et al., 2005). Afterward, we segment the gray matter from each aligned MR image using an atlas-based EM segmentation algorithm (Pohl et al., 2004). We then compute the PA of an anatomical structure at a given time point based on the overlap of the eight corresponding gray matter segmentations. The first row of Fig. 7 shows a sample slice of the PA at three time points.

We can interpolate among atlases in either PA or LogOdds space. However, one should not perform linear operations directly on PAs, unless one restricts these operations to convex combinations (CC). As longitudinal data are generally composed of a few time points, only a very limited set of CCs can be applied to this type of data. Moreover, computing CC of PAs in the original space does not preserve the characteristic of certain Gaussian distributions over space as shown in Fig. 6. In this example, the PA of a population is defined by a Gaussian distribution with mean A at time point 0 and B at time point 1. We therefore expect an interpolation between time point 0 and 1 to preserve the “hump” characterizing the distribution (see top graph, right column). This shape disappears at time point 0.5, when computing the CC of the two distributions within the PA space (see middle graph, right column). We can address this issue, however, by mapping the PAs into the LogOdds space and performing the CC there (see bottom graph, right column).

Most importantly, LogOdds give us the ability to use non-convex interpolation techniques. For longitudinal data, this provides us with a much richer class of interpolation functions than are available through CCs. An area of particular interest is the region around the thalamus (Fig. 7a) where an increase in volume can be observed over time in the PA. The second row of Fig. 7 shows a magnified

version of the PA of the thalamus region (a), and the corresponding linear (b) and quadratic spline interpolations (c). In the graphs, the z -axis and the symbolize the probability, the x -axis is the time axis, and the y -axis represents the row of voxels highlighted by the black line in Fig. 7a. Unlike (b), the quadratic interpolation is differentiable over time, which could enable us to extract additional parameters from the data such as the rate of change in the aging process of the population.

This completes the discussion of three examples that show the advantages of this representation over existing technologies. We first related SDMs to LogOdds and compared them to an alternative that was based on smoothing by spatial Gaussians. We then captured the uncertainty of boundaries by combining multiple segmentations of the same image to one LogOdds map. The last example described a framework for increasing the temporal resolution of PAs by interpolating the atlases within the LogOdds space.

4. Including a deformable shape atlas into a Bayesian classifier

In this section, we evaluate the power of LogOdds by studying PAs in the context of segmentation. We first build a statistical shape atlas by performing Principal Component Analysis (PCA) on a set of LogOdds (Section 4.1). We then present three different ways to transform this atlas into a PA to be used in an EM segmentation algorithm (Section 4.2). One approach is to define the PA as the mean shape described by the atlas, another is to interpret it as a level set of the shape defined by the mean and variations of the PCA, the final design is to transform the whole PCA of LogOdds into a discrete distribution through the logistic function. The performance of each model is evaluated by segmenting the caudate nucleus and thalamus in 20 data sets. The accuracy of each automatic segmentation is com-

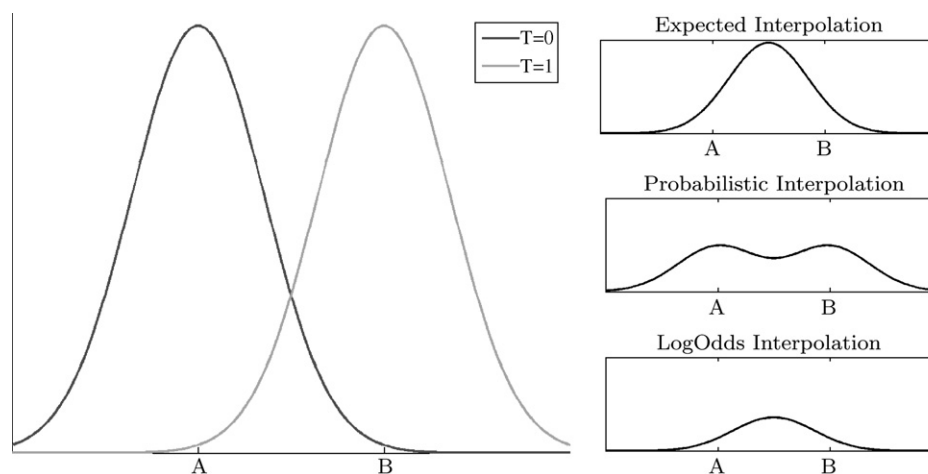


Fig. 6. The graph to the left shows the probabilistic atlas of a population at time point 0 and 1. The atlas is characterized by a Gaussian distribution in space with mean A at time point 0 and B at time point 1. The result of the convex combination of these distributions at time point 0.5 resembles a multimodal distribution in \mathbb{P}_2^n and a normal distribution in \mathbb{L}^n .




Fig. 7. The first row shows a sample slice of an interpolation of a longitudinal schizophrenia study. Each image represents a PA of the gray matter at a specific point in time of the study. low probability of the gray matter. The second row shows the PA of the thalamus (a) with black indicating the voxels that are interpolated over three time point in (b) and (c). Graph (b) was produced by linear interpolation, while the smoother quadratic spline interpolation is shown in (c).

puted by its overlap to experts' segmentations. In this experiment, the logistic function of the PCA representation consistently achieves better results.

4.1. Generating a deformable atlas via PCA on \mathbb{L}_{M-1}^n

We generate the deformable atlas by first turning a set of k manual segmentations $\{B^{(1)}, \dots, B^{(k)}\}$ into SDMs $\{\mathcal{F}^{(1)}, \dots, \mathcal{F}^{(k)}\}$. Note that the segmentations have been aligned to each other through registering the corresponding images via (Zöllei et al., 2005). At this stage, our representation is very similar to that of (Tsai et al., 2003; Leventon, 2000; Yang et al., 2004; Golland et al., 2005). However, we interpret these SDMs as LogOdds maps and can thus embed the representation within a vector space. As such, we can perform PCA on the training set $\{\mathcal{F}^{(1)}, \dots, \mathcal{F}^{(k)}\}$ (Gentle, 1998). The deformable atlas is now defined by the eigenvector or modes of variation matrix U and $\bar{\mathcal{F}} \triangleq (\bar{\mathcal{F}}_1^T \dots \bar{\mathcal{F}}_{M-1}^T)^T$, where $\bar{\mathcal{F}}_a$ is the mean vector of LogOdds of label a .

The deformable atlas $(\bar{\mathcal{F}}, U)$ encodes shapes $\mathcal{F}(\theta) \in \mathbb{L}_{M-1}^n$ within that atlas space by the expansion coefficient θ with $\mathcal{F}(\theta) = \bar{\mathcal{F}} + U \cdot \theta$. We refer to the LogOdds maps of a label a defined by θ as

$$\mathcal{F}(\theta)_a = \bar{\mathcal{F}}_a + U_a \cdot \theta, \quad (3)$$

where matrix U_a is the entry in U corresponding to structure a . We define $U_{a,i} \in \mathbb{L}_1^n$ as the i th eigenvector of structure a .

4.2. A shape-based segmenter for MR images

We now make use of the shape atlas by integrating it into a voxel-based classifier for the segmentation of ana-

tomical structures with weakly visible boundaries in MR images. The method is an extension of a class of voxel-based classification methods (Wells et al., 1996; Van Leemput et al., 1999; Kapur, 1999; Marroquin et al., 2003; Pohl et al., 2004). These classifiers simultaneously estimate the image inhomogeneity β and determine the underlying label map \mathcal{T} based on the observed MR image \mathcal{I} . We extend this model by incorporating the expansion coefficients θ of the PCA model into the estimation process. We provide a full derivation of the new algorithm in Appendix A.

Briefly, the resulting EM-implementation is defined by two steps. The *Expectation-Step (E-Step)* computes the weights for each structure a and voxel x . The weights are defined by the product between the label map probability $P(\mathcal{T}_x(a) = e_a | \theta')$ conditioned on the shape parameter θ' and the intensity probabilities $P(\mathcal{I}_x | \mathcal{T}_x(a) = e_a, \beta'_x)$ conditioned on voxel x being assigned to label a with image inhomogeneity β'_x :

$$\mathcal{W}_x(a) \propto P(\mathcal{I}_x | \mathcal{T}_x = e_a, \beta'_x) \cdot P(\mathcal{T}_x = e_a | \theta'). \quad (4)$$

The *Maximization Step (M-Step)* estimates the inhomogeneities β' and shape θ' based on the weights \mathcal{W}_x . The estimates for the shape parameters are the solution to the following MAP estimation problems:

$$\theta' \leftarrow \arg \max_{\theta} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{T}_x = e_a | \theta) + \log P(\theta). \quad (5)$$

The solution of Eq. (5) depends on the definition of the conditional probability $P(\mathcal{T}_x = e_a | \theta)$ that captures the relationship between the shape parameters θ and the label

map \mathcal{T} . We now present three different interpretations of $P(\mathcal{T}_x = e_a|\theta)$ based on the shape atlas described in Section 4.1.

The first probabilistic model of $P(\mathcal{T}_x = e_a|\theta)$ is based on the assumption that the deformable shape atlas represents a family of level set functions whose zero level sets are the object's boundary. This representation reduces $\mathcal{F}(\theta)_a$ to a binary map via the Heaviside function:

$$\mathcal{H}(v) \triangleq \begin{cases} 1, & \text{if } v \geq 0 \\ 0, & \text{otherwise} \end{cases}$$

A natural definition of the conditional probability is therefore

$$P_{\mathcal{H}}(\mathcal{T}_x = e_a|\theta) \triangleq \begin{cases} \frac{\mathcal{H}(\mathcal{F}(\theta)_{a,x})}{Z}, & \text{for } a \in \{0, \dots, M-1\} \\ \frac{1}{Z}, & a = M \end{cases} \quad (6)$$

with normalization factor $Z = 1 + \sum_{a'=1, \dots, M-1} \mathcal{H}(\mathcal{F}(\theta)_{a',x})$ and M being the background. In general, Eq. (6) defines a distribution over space with a very steep slope along the boundary of the structure. This causes the EM implementation to overemphasize the shape atlas compared to the image data, when computing the weights in the E-Step (see Eq. (4)).

The second probabilistic model interprets the result of the PCA as members of the LogOdds space \mathbb{L}_{M-1}^n . According to Section 2.1, the inverse of the multinomial *logit* of the LogOdds entry $\mathcal{F}(\theta)_{a,x}$ defines the probability that voxel x is assigned to label a so that

$$P_{\mathcal{P}}(\mathcal{T}_x = e_a|\theta) \triangleq \mathcal{P}(\mathcal{F}(\theta))_{a,x} = \begin{cases} \frac{e^{\mathcal{F}(\theta)_{a,x}}}{Z}, & \text{for } a \in \{1, \dots, M-1\} \\ \frac{1}{Z}, & \text{if } i = M \end{cases} \quad (7)$$

We define the normalization $Z \triangleq 1 + \sum_{a'=1, \dots, M-1} e^{\mathcal{F}(\theta)_{a',x}}$. The conditional probability is now characterized with a more gradual slope along the boundary of the structures, allowing for more flexibility in determining the contour of the object.

The third probabilistic model is similar to the previous one, but the shape parameter $\theta = 0$ is fixed. Thus, the PA for this model is defined as the multinomial *logit* of the mean LogOdds function $\bar{\mathcal{F}}$

$$P_N(\mathcal{T}_x = e_a) \triangleq \mathcal{P}(\mathcal{F}(0))_{a,x} = \mathcal{P}(\bar{\mathcal{F}})_{a,x} = \begin{cases} \frac{e^{\bar{\mathcal{F}}_{a,x}}}{Z}, & \text{for } a \in \{1, \dots, M-1\} \\ \frac{1}{Z}, & \text{if } i = M \end{cases} \quad (8)$$

with $Z \triangleq 1 + \sum_{a'=1, \dots, M-1} e^{\bar{\mathcal{F}}_{a',x}}$. This model ignores the modes of variations of the atlas and is thus quite similar to static PAs already proposed by Van Leemput et al. (1999) and Pohl et al. (2006).

We constructed three different EM implementations that only differ in the mapping of the shape atlas to PAs that is captured by the definition of $P(\mathcal{T}_x = e_a|\theta)$. The first and

third mapping (Eqs. (6) + (8)) were influenced by representations of shape variations commonly used in the literature and the second mapping (Eq. (7)) embodied the contribution of this paper. In the next section, we measure the robustness of the three implementations to provide the reader with a meaningful comparison of our LogOdds representation to well-established techniques. We note, however, that many other definitions of $P(\mathcal{T}_x = e_a|\theta)$ are possible.

4.3. The accuracy of Bayesian classifier

We now investigate the impact of the three implementations on the accuracy of the segmentation algorithm. We base the implementation EM- \mathcal{H} on the level set representation captured by Eq. (6), the implementation EM- \mathcal{P} on our new shape representation as defined by Eq. (7), and the implementation EM-N (see also Pohl et al., 2004) on the static PA as described by Eq. (8). All three implementations segment the caudate nucleus and the thalamus in 20 test cases. These two structures are of special interest for evaluation as they are characterized by very blurry boundaries in MR images. They are also characterized by different types of shapes, the thalamus being very round and the caudate more elongated.

We determine the quality of the automatic segmentations by comparing them to manual segmentations using a measure of overlap, the Dice coefficient (Dice, 1945). The score resembles the overlap between manual and automatic segmentation with a higher score given to those automatic segmentations that have greater overlap to the manual ones. The graph in Fig. 8 shows the average Dice measures and standard error for the three implementations. If we interpret the manual segmentations as the gold standard then the average Dice score represents the accuracy of the implementation. The standard error is a measure of reliability with a small error indicating low fluctuation in performance of the approach. The discrepancy in performance is especially striking between EM- \mathcal{H}

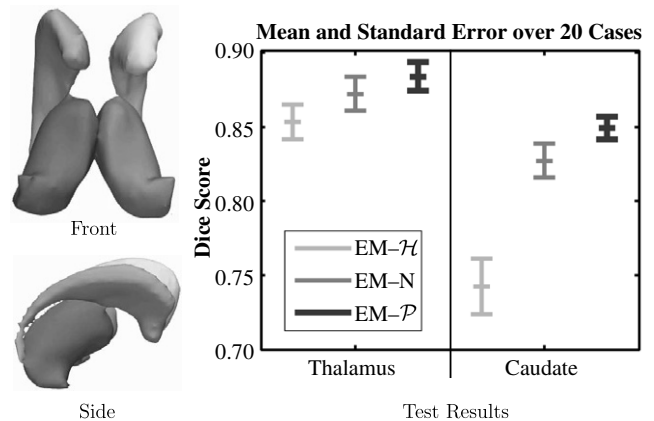


Fig. 8. Different views of a 3D model of the thalamus (dark gray) and the caudate (light gray). The model is based on a segmentation generated by EM- \mathcal{P} . The graph to the right summarizes the results of our experiment. For both structures EM- \mathcal{P} performs much better than EM- \mathcal{H} and EM-N.

and EM- \mathcal{P} . For both structures, EM- \mathcal{P} achieves a significantly higher average score (thalamus: $88.4 \pm 0.8\%$, caudate: $84.9 \pm 0.8\%$; mean DICE score \pm standard error) than EM- \mathcal{H} (thalamus: $85.3 \pm 1.2\%$, caudate: $74.3 \pm 1.6\%$) as the range of scores defined by the means and standard errors of each structure do not overlap between the two implementations. EM- \mathcal{P} also performs much better than EM-N (thalamus: $87.3 \pm 1.2\%$, caudate: $82.7 \pm 1.2\%$) with a significantly higher score for the caudate and better average score and lower standard error for the thalamus. We also note that EM-N achieves a much higher average score than EM- \mathcal{H} .

We think there are several reasons why EM- \mathcal{P} is more accurate than the other algorithms, although their main drawback is that they do not properly capture variations within the population of shapes. EM-N only incorporates a static model of the shape variations, and is thus restricted to impose a shape constraint that only models the mean of the population. For the thalamus, this limitation only slightly reduces the accuracy of the approach (in comparison to EM- \mathcal{P}) because the shape of the structure hardly changes across the healthy population. The opposite, however, is true for the caudate, where the accuracy of the approach is much lower than reported for EM- \mathcal{P} . The caudate's elongated shape wraps around the ventricles, which can greatly differ in size.

The Heaviside function induces a shape model that is too strong compared to the image model, even when modes of variations are taken into account. The intensity information is especially important for segmenting the caudate as the elongated shape is more easily determined by the clearly visible boundary to the neighboring ventricles. Since the approach largely ignores the intensity information, it produces less reliable results than EM- \mathcal{P} . We believe EM- \mathcal{P} to be a more flexible approach as it captures more variability and uncertainty, and thus only imposes the shape model when boundaries are weakly visible in the MR image.

In summary, we presented a statistical framework for the segmentation of anatomical structures in MR images. The framework is guided by the low-dimensional PCA shape model of Section 4.1 as the shape representation is turned into PA. We derived three different implementations that only differed within the probabilistic model. We ran each implementation on 20 test cases segmenting the thalamus and caudate. The segmentation algorithm based on our new representation EM- \mathcal{P} performs much better than EM- \mathcal{H} that is based on a level set representation, or EM-N that uses a “conventional” PA such as described by Van Leemput et al. (1999) and Pohl et al. (2006).

5. Discussion and conclusion

In this article, we use LogOdds maps to induce a vector-space structure on PAs (Probabilistic Atlases). LogOdds not only provide us with a framework to perform probabilistic addition and scalar multiplication but can also be

interpreted as implicit representations of the shape of objects. We also show that Signed Distance Maps (SDMs) can be viewed as a subset of the space of LogOdds maps so that the corresponding PAs can model shape. We demonstrate that LogOdds based on SDMs may not be the best implicit shape representation, especially when dealing with multi-categorical data, and propose an alternative model based on Gaussian smoothing.

We provide example applications in which LogOdds are used as shape representation, to (i) capture uncertainty among expert raters, (ii) build a time continuous atlas, and (iii) incorporate shape priors into a brain segmentation algorithm. The performance of this representation is evaluated by comparing it with other shape-driven segmentation algorithms. The LogOdds model consistently outperforms competing techniques.

This article did not discuss a set of criteria for choosing the optimal conversion of label maps into LogOdds. Based on our experience, these criteria will depend on the application as well as the training data. For example, if the training set consists of segmentations representing small spatial variations, combining the segmentations into one LogOdds map similar to Section 3.2. is probably better than deriving a PCA model, such as was done in Section 4.1. If we had such a set of criteria we could then test the accuracy of the mapping more directly than we proposed in Section 4.3.

We note that modeling shape through LogOdds, although very powerful, has its limitations. Our framework is built with the assumption that voxels in an image are independently distributed. This is likely not the case in highly structured data often observed in medical images. This should be addressed to make the method more powerful. Moreover, like most implicit representations, LogOdds do not explicitly capture variations within shape positions and orientations. Thus, our current model assumes proper alignment of the segmentations in the training data. The interpretation and application of LogOdds capturing spatial variation before alignment is a topic we would like to explore in the future. Explicit representations, such as m-reps or point distribution models, capture this information more naturally and may provide a more intuitive interpretation of shape. However, they may require a relatively high degree of customization to apply the corresponding statistical models to existing applications. The LogOdds, on the other hand, require very little adaptation, as a set of (pre-aligned) segmentations is all that is needed to build a statistical shape model. The ease of use and initialization of our framework makes it a very attractive model for applications in which shape plays an important role.

Acknowledgements

This research was supported by the Department of Veterans Affairs Merit Awards, the Brain Science Foundation, and grants from the US Army (SBIR W81XWH-04-C0031), NIH (R01 MH 40799, K05 MH 70047, R01 MH

50747, NIBIB NAMED U54 EB005149, NCRR mBIRN U24-RR021382, NCRR NAC P41-RR13218, NINDS R01-NS051826, U41-RR019703, NIAAA R01 AA016748-01), and NSF (JHU ERC CISST). We thank Mark Dreesicke from the Psychiatry Neuroimaging Laboratory, Harvard Medical School, for organizing the longitudinal study data and Torsten Rohlfing for his helpful comments.

Appendix A. Deriving a PCA-based EM implementation

It is generally difficult to determine a solution within a model accurately representing the relationship between β , θ , \mathcal{T} , and \mathcal{I} . If, however, the label map \mathcal{T} would be known then the estimation of β and θ would be simplified. This dependency is known in the machine-learning community as an incomplete data problem. A popular algorithm for estimating solutions within incomplete data problems is the expectation-maximization algorithm (EM) (McLachlan and Krishnan, 1997).

In the EM framework, the label map \mathcal{T} defines the unknown data, \mathcal{I} represents the observed data, and the parameter space consists of θ and β . At each iteration, the method improves the estimates (β', θ') of the true solution (β, θ) by solving the following Maximum A Posteriori estimation problem (MAP)

$$(\beta', \theta') \leftarrow \arg \max_{\beta, \theta} E_{\mathcal{T}|\mathcal{I}, \beta', \theta'}(\log P(\beta, \theta, \mathcal{T}|\mathcal{I})). \quad (\text{A.1})$$

The expected value is defined as $E_{A|B}(f(C)) \triangleq \sum_A P(A|B) f(C)$. We use the notation \sum_A as the sum over all possible values of A . We note that the above equation describes an existing class of EM segmentation algorithms (Wells et al., 1996; Van Leemput et al., 1999; Kapur, 1999; Marroquin et al., 2003; Pohl et al., 2004) when leaving out the shape parameter θ .

We further formalize the label map \mathcal{T} in order to continue our discussion of Eq. (A.1). The label map $\mathcal{T} = (\mathcal{T}_1, \dots, \mathcal{T}_M)$ is composed of the indicator random vector $\mathcal{T}_x \in \{e_1, \dots, e_N\}$, where x represents a voxel on the image grid. The vector e_a is zero at every position but a , where its value is one. For example, if $\mathcal{T}_x = e_a$ then voxel x is assigned to the structure a .

It was shown in Pohl et al. (2005, 2006) that Eq. (A.1) simplifies to

$$(\beta', \theta') \leftarrow \arg \max_{\beta, \theta} \sum_x \sum_a E_{\mathcal{T}_x|\mathcal{I}, \beta', \theta'}(\mathcal{T}_x(a)) \cdot [\log P(\mathcal{I}|\mathcal{T}_x = e_a, \beta) + \log P(\theta|\mathcal{T}_x = e_a)] + (\log P(\beta) + \log P(\theta)) \quad (\text{A.2})$$

that is the sum over all structures and voxels of the addition of two terms. The first term is composed of the product of the expected value of the label map, $E_{\mathcal{T}_x|\mathcal{I}, \beta', \theta'}(\mathcal{T}_x(a))$, and the sum between the log likelihood of the inhomogeneity, $\log P(\mathcal{I}|\mathcal{T}_x = e_a, \beta)$, and log probability of the label map conditioned on the shape, $\log P(\theta|\mathcal{T}_x = e_a)$. The second term is the addition of the log prior of the image inhomogeneity, $\log P(\beta)$, and the shape, $\log P(\theta)$.

The EM algorithm solves Eq. (A.2) in two steps. The *Expectation-Step (E-Step)* computes the weights $\mathcal{W}_x(a) \triangleq E_{\mathcal{T}_x|\mathcal{I}, \beta', \theta'}(\mathcal{T}_x(a))$ for each structure a and voxel x . As shown in Pohl et al. (2005), the weights are the product between the label map probability $P(\mathcal{T}_x(a) = e_a|\theta')$ conditioned on the shape parameter θ' and the intensity probabilities $P(\mathcal{I}_x|\mathcal{T}_x(a) = e_a, \beta'_x)$ conditioned on voxel x being assigned to a and the image inhomogeneity β'_x :

$$\mathcal{W}_x(a) \propto P(\mathcal{I}_x|\mathcal{T}_x = e_a, \beta'_x) \cdot P(\mathcal{T}_x(a) = e_a|\theta').$$

The *Maximization Step (M-Step)* estimates the inhomogeneities β' and shape θ' based on the weights \mathcal{W}_x . The estimates are determined as the solution of Eq. (A.2) that defines the following two MAP estimation problems:

$$\beta' \leftarrow \arg \max_{\beta} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{I}|\mathcal{T}_x = e_a, \beta) + \log P(\beta), \quad (\text{A.3})$$

$$\theta' \leftarrow \arg \max_{\theta} \sum_x \sum_a \mathcal{W}_x(a) \cdot \log P(\mathcal{T}_x = e_a|\theta) + \log P(\theta). \quad (\text{A.4})$$

We note that Eq. (A.3) was originally presented by Wells et al. (1996). Since then a variety of models with closed-form solutions have been proposed in the literature (Van Leemput et al., 1999; Marroquin et al., 2003; Ashburner and Friston, 2005). For this implementation, we choose the model by Wells et al. (1996) that defines $P(\mathcal{I}_x|\mathcal{T}_x = e_a, \beta'_x)$ by the Gaussian distribution $\mathcal{N}(\beta'_x + \mu_a, \Upsilon_a)$. (μ_a, Υ_a) captures the mean and variance of the intensity distribution of the structure a (see also Pohl et al., 2006).

Appendix B. Supplementary material

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.media.2007.06.003.

References

- Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2007. Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications* 29, 328–347.
- Ashburner, J., Friston, K., 2000. Voxel-based morphometry – the methods. *NeuroImage* 11, 805–821.
- Ashburner, J., Friston, K., 2005. Unified segmentation. *NeuroImage* 26 (3), 839–851.
- Besag, J., 1986. On the statistical analysis of dirty pictures. *Journal of the Royal Society Series B* 48 (3), 259–302.
- Bookstein, F., 1996. Landmark methods for forms without landmarks: morphometrics of group differences in outline shape. *Medical Image Analysis* 1 (3), 225–243.
- Brechbühler, C., Gerig, G., Kübler, O., 1995. Parametrization of closed surfaces for 3-D shape description. *Computer Vision and Image Understanding* 71, 154–170.
- Cootes, T., Hill, A., Taylor, C., Haslam, J., 1994. The use of active shape models for locating structures in medical imaging. *Imaging and Vision Computing* 12 (6), 335–366.
- Cootes, T., Edwards, G., Taylor, C., 1998. Active appearance model. In: *European Conference on Computer Vision (ECCV)*, vol. 2, pp. 484–498.

- Dice, L., 1945. Measure of the amount of ecological association between species. *Ecology* 26 (3), 297–302.
- Evans, M., Hastings, N., Peacock, B., 2000. *Statistical Distributions*, Third ed. Wiley, Chapter 4: Bernoulli Distribution, pp. 31–33.
- Gentle, J., 1998. *Numerical Linear Algebra for Applications in Statistics*. Springer.
- Giudici, P., 2003. *Applied Data Mining: Statistical Methods for Business and Industry (Statistics in Practice)*. John Wiley.
- Golland, P., Grimson, W., Shenton, M., Kikinis, R., 2005. Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis* 9, 69–86.
- Kapur, T., 1999. Model based three dimensional medical imaging segmentation. Ph.D. thesis, Massachusetts Institute of Technology.
- Kendall, M.G., Buckland, W.R., 1976. *A Dictionary of Statistical Terms*. Longman Group.
- Kohlberger, T., Cremers, D., Rousson, M., Ramaraj, R., Funke-Lea, G., 2006. 4-D shape priors for a level set segmentation of the left myocardium in SPECT sequences. In: *Medical Image Computing and Computer-Assisted Intervention*. Vol. 4190 of Lecture Notes in Computer Science, pp. 92–100.
- Leventon, M.E., 2000. Statistical models in medical image analysis. Ph.D. thesis, Massachusetts Institute of Technology.
- Marroquin, J., Santana, E., Botello, S., 2003. Hidden Markov measure field models for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25, 1380–1387.
- McFadden D., 1973. *Frontiers in Economics. Conditional Logit Analysis of Quantitative Choice Behavior*.
- McLachlan, G.J., Krishnan, T., 1997. *The EM Algorithm and Extensions*. John Wiley.
- Minsky, M.L., Papert, S.A., 1988. *Perceptrons*, Second ed. MIT Press.
- Pizer, S.M., Gerig, G., Joshi, S., Aylward, S.R., 2003. Proceedings of the IEEE: Special Issue on Emerging Medical Imaging Technology, Multiscale medial shape-based analysis of image objects. Vol. 91, pp. 670–679.
- Pohl, K., Bouix, S., Kikinis, R., Grimson, W., 2004. Anatomical guided segmentation with non-stationary tissue class distributions in an expectation-maximization framework. In: *IEEE International Symposium on Biomedical Imaging*, pp. 81–84.
- Pohl, K., Fisher, J., Kikinis, R., Grimson, W., Wells, W., 2005. Shape based segmentation of anatomical structures in magnetic resonance images. In: *IEEE International Conference on Computer Vision*. Vol. 3765 of Lecture Notes in Computer Science. Springer, pp. 489–498.
- Pohl, K.M., Fisher, J., Grimson, W., Kikinis, R., Wells, W., 2006. A Bayesian model for joint segmentation and registration. *NeuroImage* 31 (1), 228–239.
- Rauch, J., 1991. *Partial Differential Equations*. Graduate Texts in Mathematics, Springer.
- Shannon, C., 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27 (379–423), 623–656.
- Styner, M., Lieberman, J.A., Pantazis, D., Gerig, G., 2004. Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis* 8 (3), 197–203.
- Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W., Willsky, A., 2001. Model-based curve evolution technique for image segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. I-463–I-468.
- Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W., Willsky, A., 2003. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging* 22 (2), 137–154.
- Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging* 18 (10), 897–908.
- Warfield, S., Zou, K., Wells, W., 2006. Validation of image segmentation by estimating rater bias and variance. *Medical Image Computing and Computer-Assisted Intervention* 4190, 839–847.
- Wells, W., Grimson, W., Kikinis, R., Jolesz, F., 1996. Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging* 15, 429–442.
- Yang, J., Staib, L.H., Duncan, J.S., 2004. Neighbor-constrained segmentation with level set based 3D deformable models. *IEEE Transactions on Medical Imaging* 23 (8), 940–948.
- Zöllei, L., Learned-Miller, E., Grimson, W., Wells, W., 2005. Efficient population registration of 3D data. In: *IEEE International Conference on Computer Vision*. Vol. 3765 of Lecture Notes in Computer Science, pp. 291–301.