Multi-Atlas Image Soft Segmentation via Computation of the Expected Label Value

Iman Aganj and Bruce Fischl

Abstract—The use of multiple atlases is common in medical image segmentation. This typically requires deformable registration of the atlases (or the average atlas) to the new image, which is computationally expensive and susceptible to entrapment in local optima. We propose to instead consider the probability of all possible atlas-to-image transformations and compute the *expected label value (ELV)*, thereby not relying merely on the transformation deemed “optimal” by the registration method. Moreover, we do so without actually performing deformable registration, thus avoiding the associated computational costs. We evaluate our ELV computation approach by applying it to brain, liver, and pancreas segmentation on datasets of magnetic resonance and computed tomography images.

Index Terms—Expected label value (ELV), supervised image segmentation, soft segmentation, atlas, MRI, CT.

I. INTRODUCTION

AUTOMATIC image segmentation is often a central step in medical imaging studies, enabling the analysis of specific regions of interest (ROIs). In supervised segmentation, an algorithm segments a new image using the information derived from a training dataset of images that are accompanied with gold-standard (e.g. manually delineated) ROI labels. Two popular approaches to supervised image segmentation use multiple atlases [1–4] and deep neural networks [5, 6]. In multi-atlas-based segmentation of a new image, atlas images are (or a mean template image is) deformably registered to the new (to-be-segmented) image. The manual labels are then propagated into the new image space using the computed transformations, and fused to form the new labels.

Deformable registration of the atlas images to the new image is computationally very demanding (except for recent deep-learning based approaches [7–10]) and is the bottleneck of atlas-based segmentation. To improve computational efficiency, it has been proposed to use only a subset of atlases [11], albeit at

the price of discarding a portion of the available training data.

The transformation resulting from registration guides label propagation from the atlas to the new image. Being an iterative non-convex optimization, image registration is prone to becoming trapped in local optima, potentially leading to inaccurate propagation of the labels. Moreover, different but equally reasonable transformations may produce similar values for the registration objective function (within its margin of error). Thus, even if the global optimum is found, choosing it as the single correct transformation would disregard valuable information provided by other potentially valid transformations. Such a globally optimal solution is also rarely robust, as it is sensitive to disturbances of or changes to input images, or variations in acquisition parameters. To alleviate this issue, uncertainty in registration has been incorporated into Bayesian segmentation by approximating the marginalization over registration parameters via Markov Chain Monte Carlo techniques [12], which, even though efficiently implemented, would further increase the computational costs. Local measures of uncertainty in deformable registration have also been used to improve the sensitivity of the label propagation in atlas-based segmentation [13, 14].

In this work, we present a new atlas-based method for soft (i.e., fuzzy or probabilistic) segmentation of images, which – instead of attempting to determine a single correct label – produces the expected value of the label at each voxel of the new image, while considering the probability of possible atlas-to-image transformations. This is accomplished without either explicitly sampling from the transformation distribution (which would be intractable) or running the costly deformable registration in training or testing stages. We create a single image from the training data, which we call the *key*. Then, for a new image (after affine alignment, if necessary), we compute the *expected label value (ELV)* map simply via a convolution with the key, which is efficiently performed using the fast Fourier transform (FFT). Our fuzzy ELV map is therefore a

Support for this research was provided by the National Institutes of Health (NIH), specifically the National Institute of Diabetes and Digestive and Kidney Diseases (K01DK101631), the National Institute on Aging (R56AG068261, R01AG022381, R01AG016495), the National Institute for Biomedical Imaging and Bioengineering (P41EB015896, R01EB019956, R01EB023281), and the National Institute for Neurological Disorders and Stroke (R01NS105820, R01NS083534, U01NS086625). Additional support was provided by the BrightFocus Foundation (A2016172S).

Computational resources were provided by NIH Shared Instrumentation Grants (S10RR023401, S10RR019307, S10RR023043, S10RR028832), the O2 High Performance Compute Cluster at Harvard Medical School, the Enterprise

Research Infrastructure & Services at Mass General Brigham (MGB), and the AWS Cloud Credits for Research program.

I. Aganj (iaganj@mgh.harvard.edu, +1-617-724-5652) and B. Fischl (bfischl@mgh.harvard.edu) are with the Athinoula A. Martinos Center for Biomedical Imaging, Radiology Department, Massachusetts General Hospital (MGH), Harvard Medical School, Charlestown, MA 02129, USA, and also with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

B. Fischl has a financial interest in CorticoMetrics, a company whose medical pursuits focus on brain imaging and measurement technologies. B. Fischl’s interests were reviewed and are managed by MGH and MGB according to their conflict of interest policies.

robust combination of labels suggested by atlas-to-image transformations, weighted by a measure of the transformation validity. This soft segmentation can be further used to initiate a subsequent hard-segmentation procedure. We validate our approach through brain, liver, and pancreas segmentation experiments on magnetic resonance (MR) and computed tomography (CT) images.

This article extends our preliminary conference version [15]. In particular, we have improved the method as well as expanded our empirical evaluation by including several new datasets. Moreover, our Matlab toolbox is now publicly available (<https://www.nitrc.org/projects/elv>). In the following, we describe the proposed method in detail (Section II and the appendices) and present experimental results (Section III) along with some concluding remarks (Section IV).

II. METHODS

A. Segmentation from a Single Atlas

Let $I: \mathbb{R}^d \rightarrow \mathbb{R}$ be the d -dimensional image to be segmented, and $J: \mathbb{R}^d \rightarrow \mathbb{R}$ an atlas image with the same contrast as I , for which the manual label of a specific ROI has been provided as $L: \mathbb{R}^d \rightarrow \{0,1\}$.¹ For the new image I , we wish to compute the expected value of the ROI label, $E: \mathbb{R}^d \rightarrow [0,1]$, which is a measure of likelihood of each voxel belonging to the ROI.

In traditional atlas-based image segmentation, the label L is propagated into the space of I as $L \circ \mathbf{T}^*(I, J)$, where the transformation $\mathbf{T}^*(I, J)$ is computed via registration as $\mathbf{T}: \mathbb{R}^d \rightarrow \mathbb{R}^d$ that maximizes the similarity between I and $J \circ \mathbf{T}$.² Here, instead, we propose to compute the *expected value* of the propagated L , while considering a probability for each possible transformation in $\mathbb{T} := \{\mathbf{T}: \mathbb{R}^d \rightarrow \mathbb{R}^d\}$, as follows:

$$\begin{aligned} E &:= \mathbf{E}[L \circ \mathbf{T} | I, J] \\ &= \int_{\mathbb{T}} \Pr(\mathbf{T} | I, J) (L \circ \mathbf{T}) d\mathbf{T}. \end{aligned} \quad (1)$$

Equation (1) computes the ELV as an integral over the space of all transformations, which could be regarded as multiple (theoretically an infinite number of nested) integrals over the space of parameters representing $\mathbf{T}$. For free-form deformation, as considered here, Eq. (1) in fact includes a d -dimensional integral – with respect to the value of $\mathbf{T}(\mathbf{x})$ – for each $\mathbf{x} \in \mathbb{R}^d$. In standard atlas-based segmentation, $\Pr(\mathbf{T} | I, J)$ is considered a Dirac delta, $\delta(\mathbf{T} - \mathbf{T}^*(I, J))$, whereas here we will consider a full probability distribution for it.

Using Bayes' theorem, we can write the probability of the transformation given both the new and atlas images as:

$$\Pr(\mathbf{T} | I, J) \propto \Pr(I, J | \mathbf{T}) \Pr(\mathbf{T}), \quad (2)$$

where the two right-hand-side factors correspond to the image similarity and the transformation regularity, respectively. For

the former, we opt to use the inner product of the image and the transformed atlas, since it is expected to be higher when the two images are well aligned:

$$\Pr(I, J | \mathbf{T}) \propto \int_{\mathbb{R}^d} I(\mathbf{x}) (J \circ \mathbf{T})(\mathbf{x}) d\mathbf{x}. \quad (3)$$

It is, however, well established that the inner product reflects the degree of alignment more effectively when only the *phase* information of the image is included [16, 17], which is how in practice we will proceed, as described in Section II.C. A discussion on our choice of the inner product of phase images as image similarity is provided in Appendix A. In the following, we first consider the case where $\mathbf{T}$ is only a translation.

1) Translation

For a translation, $\mathbf{T}(\mathbf{x}) = \mathbf{x} - \Delta$, the inner product in Eq. (3) becomes the cross-correlation of the image and the atlas, which is commonly used for image alignment [16, 17]:

$$\begin{aligned} \Pr(I, J | \Delta) &\propto \int_{\mathbb{R}^d} I(\mathbf{x}) J(\mathbf{x} - \Delta) d\mathbf{x} \\ &= (I * \bar{J})(\Delta), \end{aligned} \quad (4)$$

where $*$ denotes the convolution operator, and $\bar{J}(\mathbf{x}) := J(-\mathbf{x})$. By assuming a flat prior for the shift (i.e., a constant $\Pr(\Delta)$) and combining Eqs. (1), (2), and (4), the ELV at voxel $\mathbf{y}$ will be:

$$E(\mathbf{y}) \propto \int_{\mathbb{R}^d} (I * \bar{J})(\Delta) L(\mathbf{y} - \Delta) d\Delta, \quad (5)$$

or,

$$\begin{aligned} E &\propto (I * \bar{J}) * L \\ &= I * (\bar{J} * L). \end{aligned} \quad (6)$$

In the second line, we exploited the associativity property of convolution, which leads to the following expression for the ELV:

$$E \propto I * A, \quad (7)$$

where we define and pre-compute the *key*, A , from the atlas, as:

$$A := \bar{J} * L. \quad (8)$$

This convolution operation blurs the flipped atlas image and shifts it in the opposite direction of the ROI, thereby resulting in a key image, A , with the ROI roughly at its center.

Next, we will incorporate deformations in our transformation model.

2) Deformation

To generalize the transformation $\mathbf{T}$ to include deformations,

¹ The gold-standard segmentation may also be a soft label, $L: \mathbb{R}^d \rightarrow [0,1]$.

² We denote vector-valued variables in bold.

we will use the common Tikhonov prior on the regularity of the deformation field as the probability of the transformation:

$$\Pr(\mathbf{T}) \propto R(\mathbf{T}) := e^{-\frac{1}{2\sigma^2} \int_{\mathbb{R}^d} \|\partial \mathbf{T}(\mathbf{z}) - \mathbb{I}\|_F^2 d\mathbf{z}}, \quad (9)$$

where $\partial \mathbf{T}$ is the Jacobian matrix of $\mathbf{T}$, $\mathbb{I}$ is the $d \times d$ identity matrix, and the constant parameter σ represents a prior on the magnitude of the deformations. In Appendix B, we show that the ELV is still computed following Eq. (7), where the key, A , is initially computed as in Eq. (8), but then updated to incorporate the deformation. We show that we can approximate this update by an inhomogeneous blurring of the key, as:

$$A(\mathbf{x}) \leftarrow [A(\mathbf{z}) * G(\mathbf{z}|\mathbf{0}, \sigma^2 \|\mathbf{x}\|_2 \mathbb{I})]_{\mathbf{z}=\mathbf{x}}, \quad (10)$$

where $G(\cdot|\boldsymbol{\mu}, \Sigma)$ represents the Gaussian function with the mean $\boldsymbol{\mu}$ and the co-variance matrix Σ . One can see that the size of the blurring kernel increases with the square root of the Euclidean distance from the center of A – i.e., the region corresponding to the label ROI (see Section II.A.1). Blurring a region in A decreases its contribution to soft segmentation by removing its high-frequency components prior to the convolution in Eq. (7). This means that Eq. (10) takes local deformations into account by giving a smaller weighting to regions in the atlas image that are farther from the ROI, making the information in such far areas less important.

Note that, even without the inhomogeneous blurring (i.e., $\sigma = 0$), the proposed method still encompasses – not just a single optimal, but – all possible translations, each partially contributing to the ELV with a weighting proportional to its degree of aligning. Given that a deformation field itself can be construed as many local translations, each of which would partially contribute to the ELV, one can see that, in some way, deformations are inherently accounted for.

The proposed model accounts for large translations, as well as local deformations, even though we do not run any deformable registration. As for rotation and global scaling, accounting for local deformations covers a small amount of them, and to allow for large amounts, we can initially affinely align the image and the atlas.

B. Multiple Atlases

In case N atlases (affinely aligned in the same space) with manual labels are available, we will write Eq. (1) in the same fashion, as:

$$E := \frac{1}{N} \sum_{i=1}^N \mathbf{E}[L_i \circ \mathbf{T}|I, J_i], \quad (11)$$

where J_i and L_i are the i^{th} pair of atlas and manual-label images, respectively. This will yield similar results as in Eqs. (7) and (10), with the only difference being Eq. (8), now generalized as:

$$A := \frac{1}{N} \sum_{i=1}^N \bar{J}_i * L_i. \quad (12)$$

Note that even in the case of multiple atlases, the key, A , is a *single* image that is pre-computed from the training data.

Given that the to-be-segmented ROI is roughly at the center of the key, the regions whose relative positions with respect to the ROI are consistent across atlases are amplified in the key through constructive addition in Eq. (12), whereas those with relative positions varying from atlas to atlas due to inter-atlas deformations are washed out. Therefore, the relative weighting of the regions that are most informative for segmenting the desired ROI is elevated in the key thanks to aggregation of information from multiple atlases.

C. Implementation

1) Computation in the Fourier Domain

To create the key, A , we first ensure that the N training images are represented roughly in the same space; and if not, we affinely align them. By applying the convolution theorem to Eq. (12), we will then use FFT to initialize A :

$$A = \mathcal{F}^{-1} \left\{ \frac{1}{N} \sum_{i=1}^N \frac{\hat{J}_i^*}{|\hat{J}_i|} \hat{L}_i \right\}, \quad (13)$$

where the *hat* ($\hat{\cdot}$) sign and $\mathcal{F}^{-1}$ represent the Fourier and inverse Fourier transforms, respectively, and $\hat{J}^*$ is the complex conjugate of $\hat{J}$. By only keeping the phase information of the image (i.e., normalizing $\hat{J}_i$ by its magnitude), we create a sharper probability distribution for the aligning transformation in Eq. (3) [16, 17] (see Appendix A). In addition, this has an intensity normalization effect, preventing A from giving a different weighting to an atlas image due to its global intensity scaling. Next, for inhomogeneous blurring of the key (i.e., if $\sigma > 0$), we update A voxel-wise following Eq. (10) by multiplying and summing it with a varying discretized Gaussian kernel.

To segment a new image, I , we first make sure that it is correctly represented in the atlas space (otherwise, affinely align it to the mean atlas image using any existing affine registration tool), and then compute the ELV map from Eq. (7) as follows:

$$E \propto \mathcal{F}^{-1} \left\{ \frac{\hat{I}}{|\hat{I}|} \hat{A} \right\}. \quad (14)$$

Note that $\hat{A}$ is pre-computed from the atlases and kept offline. Different keys can be independently created for different organs for the purpose of multi-organ segmentation, in which case label overlap can be avoided using the soft-max operator [6] on the normalized probability maps of different organs, and via other label-fusion strategies [3, 4] such as STAPLE [18].

For hard segmentation of the organ (or structure) from the map, we threshold the map to keep a voxel subset with the

volume 14% larger than that of an average organ (estimated from the atlases); see Appendix C for the rationale behind this choice.³ We then refine the mask by keeping the largest connected component (CC), as well as the CCs with at least half the volume of the largest CC, and then filling the holes.⁴

2) Second Pass

Once the initial ELV map is obtained, it can be refined by recalculating Eq. (14) while this time prioritizing the initial soft-segmented area. In our experiments, for instance, we used weighted versions of $A(\mathbf{x})$ and $I(\mathbf{x})$, as $A(\mathbf{x})G(\mathbf{x}|\mathbf{0}, s^2\mathbb{I})$ and $I(\mathbf{x})[E(\mathbf{y}) * G(\mathbf{y}|\mathbf{0}, s^2\mathbb{I})]_{\mathbf{y}=\mathbf{x}}$, respectively, where the size of the Gaussian window ($2s$) was chosen to be roughly that of an average organ.

3) Intensity Prior

Given that using the phase image discards some image intensity information, one can further augment the computed ELV volume with image intensities. At a given voxel, the Bayes formula implies:

$$\Pr(L|I, E) = \frac{\Pr(I|L, E) \Pr(L|E)}{\Pr(I|L, E) \Pr(L|E) + \Pr(I|\neg L, E) \Pr(\neg L|E)} \quad (15)$$

where L indicates that the given voxel belongs to the label, with $\neg$ the negation operator, and I and E are the values of the image intensity and the computed ELV at the voxel, respectively. Were it known whether the voxel is included in the label or not, the image intensity would be conditionally independent of the ELV; i.e. $\Pr(I|L, E) = \Pr(I|L)$ and $\Pr(I|\neg L, E) = \Pr(I|\neg L)$. Using the ELV for $\Pr(L|E)$ then leads to:

$$\Pr(L|I, E) = \frac{\Pr(I|L) E}{\Pr(I|L) E + \Pr(I|\neg L) (1 - E)} \quad (16)$$

where $\Pr(I|L)$ and $\Pr(I|\neg L)$ can be approximated by Gaussian functions of the intensity values of I , with their parameters estimated from the atlases (or the image itself using the initial ELV map). We use the acronym ELV+IP to refer to this ELV map that has been modulated by the intensity prior (IP). For E to exhibit the properties of a probability, the ELV map needs to be normalized by its maximum, with any negative values projected to zero. Note that E was created using the *phase* images that had lost some image intensity information, meaning that the full image (I) in $\Pr(L|I, E)$ brings about new information not already included in E .

$\Pr(L|I, E)$ can even substitute for I itself in the computation of the ELV, as – depending on the image contrast – it may better highlight the organ of interest, which is the most informative part of the image for segmentation. In that case, since the ELV map has not yet been computed, we use a constant E in Eq. (16)

equal to the label-to-image volume ratio estimated from the atlases.

Several other post-processing steps are possible after this soft segmentation [1]. If binary segmentation is desired, the ELV map can then be thresholded (see Appendix C and Section II.C.1) or used as a seed region to subsequently initialize an unsupervised hard segmentation algorithm [15, 20].

III. RESULTS AND DISCUSSION

We evaluated our ELV computation method on several medical image databases via leave-one-out cross validation. We inspected the images in each database to make sure that they were represented with the correct orientation in the same space, i.e. they were not misaligned due to dimension permutation or mirroring. Since we did not subsequently see any remaining global scaling and rotation that exceeded the inherent inter-subject variability, we did not affinely register the images.⁵ For each test image in a database, we created the key from the remainder of the images (i.e., labeled atlases) in the database following Eq. (12), computed the label for the test image, and report the Dice overlap coefficient between the computed label and the known label. Since optimizing for σ in Eq. (10) improved the Dice scores only negligibly ($< 1\%$) in our initial benchmarking, we report our results in this section for the simple case with $\sigma = 0$.

As described in Section II.C, we computed the ELV map in two passes, modulated the ELV with the intensity prior in Eq. (15), where we used the atlases to estimate the means (except for the liver; see Section III.B) and standard deviations, and then hard-thresholded the probability maps to create masks. The algorithm chose a single CC from the mask in all experiments, except for 3% and 0.4% of the pancreas and hippocampus segmentation cases, respectively, where two CCs were chosen.

For data homogeneity in the abdominal segmentation experiments, we devised some inclusion criteria for each dataset, which were roughly chosen to minimize the number of excluded subjects (see below). Additional steps to preprocess the abdominal CT images included: smoothing the borders of each image, automatically removing the patient table (via thresholding the image and removing the lower-most one or two CCs), and using the intensity-prior image (Section II.C.3) instead of the abdominal image itself for ELV computation (thereby highlighting the organ amongst all other parts of the image).

A. Brain

We first assessed the ability of our method to imitate FreeSurfer [21] in segmenting brain subcortical structures. We used T₁-weighted MR images of 1224 subjects from the third release in the *Open Access Series of Imaging Studies (OASIS-3)*

³ For data that is not too noisy, the organ size can also be estimated as the inflection point of the curve obtained by sorting the ELV map in descending order. Alternatively, the ELV map can be thresholded with a value optimized from the training data.

⁴ A Markov random field prior on the voxel labels could also be used to encourage spatial regularity [3, 19].

⁵ In our preliminary experiments, initial affine registration did not improve the results.

[22], normalized to the size $256 \times 256 \times 256$ with 1 mm^3 isotropic resolution. We considered the FreeSurfer-generated labels for 12 subcortical structures (left and right thalamus, caudate, putamen, pallidum, hippocampus, and amygdala) as “silver” standard and tried to reproduce the segmentation for each image via the proposed ELV approach. The median, mean, and standard error of the mean (SEM) of the resulting cross-validation Dice scores between the labels generated by ELV+IP and FreeSurfer are shown in Table I. Overall, the Dice score had a median of 0.782 and a mean of $0.766 \pm (\text{SEM}) 0.001$ across subjects and structures. (For the raw ELV, the median was 0.740 and the mean was 0.721 ± 0.001 .) Since no manually delineated labels were used as the gold standard in this experiment, the results merely reveal how faithful the proposed approach is in reproducing FreeSurfer labels, rather than comparing the two methods.

We then externally evaluated our approach by applying the key generated from the OASIS-3 images and FreeSurfer labels to segment two previously-unseen datasets of T1-weighted brain MR images with manual labels. The first dataset, the Internet Brain Segmentation Repository (IBSR) [23], included images from 18 subjects with the slice thickness of 1.5 mm and in-plane pixel sizes ranging from 0.84 mm to 1 mm, which we resampled to 1 mm^3 isotropic voxel size. The second dataset [24] included 1 mm^3 isotropic-voxel images from 39 subjects (28 healthy and 11 with dementia). We ran our ELV segmentation pipeline on each dataset once (without any further finetuning), resulting in median Dice scores of 0.69 (mean: 0.68 ± 0.01) and 0.677 (mean: 0.668 ± 0.004) for the two datasets, respectively. Note that the segmentation error in these results is the aggregate error from both the silver-standard labels used in training and the subsequent ELV segmentation. Further modulation with the intensity prior here deteriorated the results (median ELV+IP Dice: 0.56 and 0.624, respectively); this was not surprising, as the parameters of the image intensity prior had been estimated from another dataset (OASIS-3) with different image acquisition protocols, and the new images were not intensity-harmonized. The ELV itself, however, uses the phase images and is thus not so sensitive to inter-dataset variation in image intensity.

For comparison, in a similar experiment [25], a U-Net type convolutional neural network (CNN) was trained on 581 FreeSurfer-segmented T1-weighted brain images. The authors’ trained model produced mean Dice scores of 0.74 and 0.71 on two manually labeled test datasets. (The authors, however, did not compare the labels that they computed with FreeSurfer-generated labels.)

B. Liver

Next, we used the training dataset of the public *Liver Tumor Segmentation (LiTS) Challenge* [26], which includes abdominal CT images with manually delineated labels for the normal tissue and lesions in the liver, provided by various clinical sites.

⁶ Resampling to a different resolution is only needed if the dimensions of the imaged object (in pixels) and their relative ratios substantially vary across subjects due to varying voxel size.

TABLE I
DICE COEFFICIENTS BETWEEN ELV+IP AND FREESURFER LABELS IN BRAIN

Structure	Hemisphere	Median	Mean	SEM
Thalamus	Left	0.751	0.744	0.001
	Right	0.771	0.764	0.001
Caudate	Left	0.804	0.793	0.002
	Right	0.807	0.796	0.002
Putamen	Left	0.808	0.797	0.002
	Right	0.816	0.802	0.002
Pallidum	Left	0.758	0.685	0.005
	Right	0.769	0.743	0.003
Hippocampus	Left	0.793	0.784	0.001
	Right	0.806	0.799	0.001
Amygdala	Left	0.754	0.742	0.002
	Right	0.758	0.748	0.002
All	Both	0.782	0.766	0.001

Although the subjects were scanned with contrast enhancement for liver lesions, we considered the entire (healthy and lesion) organ label in our experiments. 85 subjects passed our inclusion criteria, mainly the slice thickness being included in the header and no larger than 2 mm. The images were resampled in the space of the first image (where the Dice scores are reported) to $(1.6\text{mm})^3$ isotropic resolution, so they were all of the size $248 \times 248 \times 323$.⁶

We computed the ELV map for each subject, an example of which is illustrated in Fig. 1 (left) for the representative subject (corresponding to the median final Dice score; see below). To create the ELV+IP map, we estimated the standard deviation of the intensities of the liver and the background from the 84 atlas subjects, using the manual labels and their dilated versions (by a sphere of radius 50), respectively. For stability, we estimated the *mean* intensity using the initial ELV mask of the test subject, given that lesion size and intensity varied from subject to subject. Next, we modulated the ELV with the intensity prior (Fig. 1, middle), created a new mask, and further refined it with an updated intensity mean estimated from this mask. For mask preparation, we also performed morphological *opening* with a spherical structuring element with the radius of 2 voxels while keeping the largest CC (i.e., eroding + keeping + dilating), which removed unwanted smaller structures attached to this relatively large ROI.

The cross-validation Dice coefficients between the masks computed from ELV+IP and manual labels (Fig. 1, right) had a median of 0.92 across subjects (mean: 0.91 ± 0.01). Dice scores, false positive rates (FPR), and false negative rates (FNR) are shown in Table II for the ELV and ELV+IP approaches, and a volumetric video of Fig. 1, as well as a video of ELV slices for all subjects are available in the supplementary materials. Among those results with lower (entire-organ) Dice scores, lesion regions were frequently the culprit, as the intensity prior, although generally improving the segmentation, partially excluded some of those regions. The relative portion of the lesion in the liver was significantly negatively correlated with the Dice score generated by the ELV+IP ($r = -0.48$, $p < 10^{-5}$), but not the ELV ($r = -0.01$, $p = 0.93$). Repeating this

Fig. 1. CT image (blue) of the representative subject (i.e., with median segmentation Dice score) in the LiTS dataset. The slice with the largest cross section with the manual label is shown. *Left:* The ELV map of the liver (red; occasional negative values in green). *Middle:* The ELV+IP map (red). *Right:* The resulting binary segmentation (red), the manual label (green), and their overlap (yellow). Intensities have been scaled for better visualization. A 3D video of this figure is available in the supplementary materials.

TABLE II
OVERLAP OF THE AUTOMATIC AND MANUAL LIVER LABELS

Method	Dice Score		FPR	FNR	Runtime (min): Mean $\pm$ SEM	
	Median	Mean $\pm$ SEM	Mean $\pm$ SEM		Training	Testing
ELV	0.87	0.85 $\pm$ 0.01	0.46% $\pm$ 0.03%	8% $\pm$ 1%	31.7 $\pm$ 0.2	0.90 $\pm$ 0.01
ELV+IP	0.92	0.91 $\pm$ 0.01	0.19% $\pm$ 0.02%	9% $\pm$ 1%		1.08 $\pm$ 0.01
SyN deformable: Majority voting	0.80	0.72 $\pm$ 0.02	0.13% $\pm$ 0.01%	37% $\pm$ 2%	300 $\pm$ 4	
SyN deformable: Joint label fusion	0.88	0.82 $\pm$ 0.02	0.07% $\pm$ 0.01%	26% $\pm$ 2%	375 $\pm$ 4	
U-Net	0.90	0.90 $\pm$ 0.01	0.26% $\pm$ 0.02%	8% $\pm$ 1%	(GPU) 131.0 $\pm$ 0.2	(GPU) 12.05 $\pm$ 0.01
LiTS leaderboard	-	0.84 ~ 0.97			-	

experiment with $\sigma = 0.1$ resulted in a negligible change (< 0.0001) in the mean Dice score, which was not statistically significant (paired t -test $p = 0.16$).

To compare our method with state-of-the-art multi-atlas based image segmentation, we then used the `antsJointLabelFusion` function of the Advanced Normalization Tools (ANTs) [27] to segment the liver using both majority voting and joint label fusion [4]. For each test subject, the remaining 84 atlas images were registered to the test image, into whose space labels were then propagated and fused. We chose the *rigid + affine + deformable SyN* as the transformation type. As shown in Table II, joint label fusion (the best of the two) achieved a median Dice score of 0.88 (mean: 0.82 ± 0.02).

Next, we trained a 3D CNN with the U-Net architecture [6] for liver segmentation, which had 2 downsampling layers and 16 initial filters (at the first convolutional layer), and used the Dice coefficient as the objective function. We trained the network for 5 epochs using 3D sample patches of size $132 \times 132 \times 132$ with a mini-batch size of 8. Through 10-fold cross-validation, each subject's label was predicted by averaging the label scores of overlapping patches (10 voxels apart), binarizing the label, and keeping the largest CCs as we did for the ELV (which improved the results). As shown in Table II, the median Dice coefficient was 0.90 (mean: 0.90 ± 0.01). Increasing the number of initial filters to 32 at the

price of a smaller patch size ($100 \times 100 \times 100$) lowered the Dice values.

For comparison, at the time of the submission of this article, the LiTS challenge website [26] reported mean Dice values for the liver on their test data ranging from 0.84 to 0.97 (disregarding the outlier results with mean Dice ≤ 0.35), with many of the methods applying deep learning.

To evaluate the runtimes, we ran the atlas-based methods on a compute cluster, allocating 4 CPU cores, 48GB of memory, and no graphics processing unit (GPU) to each computational job. The ELV method exploited the available 4 cores through Matlab's internal parallelization of FFT, and ANTs did so via parallelization of atlas-to-image registrations (due to our use of the `-c 2 -j 4` options). Each U-Net experiment was run on an Nvidia RTX8000 GPU with 48GB of memory. The runtimes of all methods are shown in Table II. For the ELV+IP, the mean entire-experiment runtime for a test subject was 32.8 ± 0.2 min, of which 31.7 ± 0.2 min were spent in the training step (key creation). Soft segmentation from the key and subsequent hard segmentation, on the other hand, took only 0.86 ± 0.01 min and 0.229 ± 0.001 min, respectively. In contrast, the majority voting and joint label fusion methods of ANTs took 300 ± 4 min and 375 ± 4 min per test subject, respectively. The U-Net took 131.0 ± 0.2 min for training and 12.05 ± 0.01 min to segment a test subject. Given that the U-Net experiment was accelerated via a GPU, its runtime is not meant to be directly compared to those

of the atlas-based methods that did not use GPUs.

C. Pancreas

Lastly, we took a similar approach as in the previous subsection to segment the pancreas in two experiments, using two CT databases from *The Cancer Imaging Archive (TCIA)* acquired at the National Institutes of Health (NIH) Clinical Center [28, 29] (82 subjects) and from the *Memorial Sloan Kettering Cancer Center* [30] (225 subjects; those with slice thickness of 2–3 mm). The labels created from the ELV+IP map in cross-validation had a median Dice score of 0.59 (mean: 0.56 ± 0.02) for the former database and a median Dice score of 0.50 (mean: 0.48 ± 0.01) for the latter database. (For the raw ELV, the median was 0.57 and 0.49, respectively.) Note that the pancreas in the second dataset included lesions (although no significant correlation between Dice and relative lesion load was observed).

The pancreas’ anatomical flexibility and variability in shape, size, and location make it a more challenging organ for segmentation than the liver and the brain subcortical structures, which could explain the lower accuracy of the results by our atlas-based method for this organ. For comparison, recent works on pancreas segmentation applying CNNs to the first (TCIA) dataset report Dice scores as high as 0.83 [29, 31, 32].

Note that, in contrast to mainstream supervised segmentation methods that employ deformable registration or sophisticated trained neural networks, we compute the ELV map via a simple linear convolution operation on the (phase) image. Combining the signal-processing (ELV) and deep-learning (U-Net) approaches into a single neural network that receives both the image and the ELV map may improve the segmentation beyond each of the two approaches, which we propose as part of the future work.

IV. CONCLUSIONS

We have introduced a new approach to supervised soft segmentation, which computes the expected label value (ELV) of a region of interest from an image using a training dataset of annotated atlases. The proposed method does not perform costly deformable registration, thereby also avoiding entrapment in local optima. We have evaluated the performance of our ELV computation technique in segmentation of the brain, the liver, and the pancreas. Future work consists of: using the ELV map to augment the input to a convolutional neural network beyond the image itself, expecting to increase the segmentation accuracy of the better-informed model; inherent inclusion of large rotation and scaling in ELV; and thorough evaluation with metrics other than Dice, such as region volume and surface distances.

V. APPENDICES

A. Inner Product as the Image Similarity Metric

The inner product of the new image I and the transformed atlas image $J \circ T$, which we have proposed as the image

similarity metric in Eq. (3), is closely related to the sum-of-squared-differences (SSD) cost function that is commonly used in image registration:

$$\begin{aligned} SSD &:= \int_{\mathbb{R}^d} (I(\mathbf{x}) - (J \circ T)(\mathbf{x}))^2 d\mathbf{x} \\ &= \int_{\mathbb{R}^d} (I^2(\mathbf{x}) + (J \circ T)^2(\mathbf{x})) d\mathbf{x} \\ &\quad - 2 \int_{\mathbb{R}^d} I(\mathbf{x})(J \circ T)(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (17)$$

In order to establish an equivalence between maximizing our inner-product similarity function and minimizing SSD, it would seem necessary to include in Eq. (3) the term $-\frac{1}{2} \int_{\mathbb{R}^d} (I^2(\mathbf{x}) + (J \circ T)^2(\mathbf{x})) d\mathbf{x}$, which is not necessarily constant with respect to T due to local volume changes in the transformation. The extra terms that such an addition would introduce in $\Pr(T|I, J)$ of Eq. (2), however, can be seen to be independent of the global-translation component of T . Then, since an integral with respect to T can be taken separately with respect to a global translation value and translation-free displacement fields, as in Eqs. (23)–(25), the extra terms in $E(\mathbf{y})$ of Eq. (1) (resulting from the new translation-independent terms in $\Pr(T|I, J)$) would be constant (independent of $\mathbf{y}$), and therefore unnecessary in the computation of the ELV. Consequently, quantifying the similarity between two images as their inner product, as adopted here, corresponds to the common use of the SSD cost function in deformable image registration.

As mentioned in Section II.A, we use only the phase information of the images in Eq. (3), and measure the image similarity with the following inner product:

$$\Pr(I, J|T) = \int_{\mathbb{R}^d} \frac{\hat{I}(\omega) \widehat{J \circ T}(\omega)^*}{|\hat{I}(\omega) \widehat{J \circ T}(\omega)^*|} d\omega. \quad (18)$$

Using only the phase of the images, as in Eq. (18), is more suitable for the estimation of $\Pr(I, J|T)$, as it produces sharper probability distributions [16, 17]. To demonstrate this via an example, let us model the transformation as a simple translation, $T(\mathbf{x}) = \mathbf{x} + \Delta$. The inner product therefore becomes the cross-correlation of the phase images, similar to Eq. (4), with Eq. (18) exhibiting the anticipated normality property, $\int_{\mathbb{R}^d} \Pr(I, J|\Delta) d\Delta = 1$ (although $\Pr(I, J|\Delta)$ can occasionally become negative). Subsequently, in the simplistic case where J is a shifted version of I , i.e. $J(\mathbf{x}) = I(\mathbf{x} - \Delta_0)$, Eq. (18) will lead to $\Pr(I, J|\Delta) = \delta(\Delta - \Delta_0)$, which is the exact desired distribution here.

Lastly, the inner product is zero for non-overlapping I and $J \circ T$, which is a crucial property for the image similarity metric in ELV computation.

B. Incorporation of Deformation

In this appendix, we derive the ELV while accounting for deformations in the transformation (Section II.A.2). By combining Eqs. (1), (2), (3), and (9), the ELV at voxel $\mathbf{y}$ will

be:

$$E(\mathbf{y}) \propto \int_{\mathbb{R}^d} I(\mathbf{x}) d\mathbf{x} \int_{\mathbb{T}} (J \circ \mathbf{T})(\mathbf{x})(L \circ \mathbf{T})(\mathbf{y}) R(\mathbf{T}) d\mathbf{T}. \quad (19)$$

Since $\mathbf{x}$ and $\mathbf{y}$ are fixed in the inner integral, we make the change of variables $\mathbf{T}(\mathbf{z}) = \mathbf{S}(\mathbf{z} - \mathbf{x})$. Note that such a global shift will not change either the regularization, i.e. $R(\mathbf{T}) = R(\mathbf{S})$, or the domain of the inner integral, $\mathbb{T}$. Consequently:

$$\begin{aligned} E(\mathbf{y}) &\propto \int_{\mathbb{R}^d} I(\mathbf{x}) d\mathbf{x} \int_{\mathbb{T}} (J \circ \mathbf{S})(\mathbf{0})(L \circ \mathbf{S})(\mathbf{y} - \mathbf{x}) R(\mathbf{S}) d\mathbf{S} \\ &= \int_{\mathbb{R}^d} I(\mathbf{x}) A(\mathbf{y} - \mathbf{x}) d\mathbf{x}, \end{aligned} \quad (20)$$

or:

$$E \propto I * A, \quad (21)$$

where we define the *key*, A , as:

$$A(\mathbf{x}) := \int_{\mathbb{T}} (J \circ \mathbf{S})(\mathbf{0})(L \circ \mathbf{S})(\mathbf{x}) R(\mathbf{S}) d\mathbf{S}. \quad (22)$$

Next, we write the transformation $\mathbf{S}$ as the sum of a global translation $\mathbf{\Delta} \in \mathbb{R}^d$ and a deformation (displacement) field $\mathbf{u} \in U$:

$$\mathbf{S}(\mathbf{x}) = \mathbf{x} + \mathbf{\Delta} + \mathbf{u}(\mathbf{x}), \quad (23)$$

where $U := \{\mathbf{u}: \mathbb{R}^d \rightarrow \mathbb{R}^d \mid \int_{\mathbb{R}^d} \mathbf{u}(\mathbf{x}) d\mathbf{x} = \mathbf{0}\}$ is the set of translation-free displacement fields. The regularity prior is now:

$$R(\mathbf{S}) = \tilde{R}(\mathbf{u}) := e^{-\frac{1}{2\sigma^2} \int_{\mathbb{R}^d} \|\partial \mathbf{u}(\mathbf{z})\|_F^2 d\mathbf{z}}. \quad (24)$$

We combine the above three equations, and separate the integral over the space of all transformations into an integral over possible translation-free deformations and an integral over possible translations:

$$A(\mathbf{x}) \propto \int_U \tilde{R}(\mathbf{u}) d\mathbf{u} \int_{\mathbb{R}^d} J(\mathbf{\Delta} + \mathbf{u}(\mathbf{0})) L(\mathbf{x} + \mathbf{\Delta} + \mathbf{u}(\mathbf{x})) d\mathbf{\Delta}. \quad (25)$$

Note that this is a linear and invertible change of coordinates, hence $d\mathbf{S} \propto d\mathbf{u} d\mathbf{\Delta}$ (with the ratio independent of $\mathbf{S}$). With $\mathbf{u}$ and $\mathbf{x}$ being constant in the inner integral, we make the change of variables $\mathbf{\Delta} = \tilde{\mathbf{\Delta}} - \mathbf{u}(\mathbf{x}) - \mathbf{x}$, leading to:

$$\begin{aligned} A(\mathbf{x}) &\propto \int_U \tilde{R}(\mathbf{u}) d\mathbf{u} \int_{\mathbb{R}^d} J(\tilde{\mathbf{\Delta}} - \mathbf{u}(\mathbf{x}) + \mathbf{u}(\mathbf{0}) - \mathbf{x}) L(\tilde{\mathbf{\Delta}}) d\tilde{\mathbf{\Delta}} \\ &= \int_U \tilde{R}(\mathbf{u}) A_0(\mathbf{u}(\mathbf{x}) - \mathbf{u}(\mathbf{0}) + \mathbf{x}) d\mathbf{u}, \end{aligned} \quad (26)$$

where A_0 is the key for the translation-only case, introduced in Eq. (8):

$$A_0 := \bar{J} * L. \quad (27)$$

It can be verified that:

$$\lim_{\sigma \rightarrow 0} A \propto A_0. \quad (28)$$

We now analytically estimate the key, A , as a function of A_0 for $\sigma > 0$. Combining Eqs. (24) and (26) leads to:

$$A(\mathbf{x}) \propto \int_U A_0(\mathbf{u}(\mathbf{x}) - \mathbf{u}(\mathbf{0}) + \mathbf{x}) e^{-\frac{1}{2\sigma^2} \int_{\mathbb{R}^d} \|\partial \mathbf{u}(\mathbf{z})\|_F^2 d\mathbf{z}} d\mathbf{u}. \quad (29)$$

For simplicity, let us for now assume that $\mathbf{x}$ lies on the positive half of the first Cartesian coordinate axis, i.e., $\mathbf{x} = x\mathbf{v}_1$, where $\mathbf{v}_1$ is the unit vector in the direction of the first axis, and $x \geq 0$. We also define the line segment $Q_x := \{t\mathbf{v}_1 \mid 0 \leq t \leq x\}$. Accordingly:

$$\mathbf{u}(\mathbf{x}) - \mathbf{u}(\mathbf{0}) = \int_{Q_x} \partial \mathbf{u}(\mathbf{x}') d\mathbf{x}' = \int_0^x \partial_1 \mathbf{u}(t\mathbf{v}_1) dt, \quad (30)$$

where $\partial_1 \mathbf{u}$ is the partial derivative of $\mathbf{u}$ in the direction of $\mathbf{v}_1$. Therefore:

$$\begin{aligned} A(x\mathbf{v}_1) &\propto \int_{\partial U} A_0 \left(x\mathbf{v}_1 + \int_0^x \partial_1 \mathbf{u}(t\mathbf{v}_1) dt \right) \\ &\quad \times e^{-\frac{1}{2\sigma^2} \int_{\mathbb{R}^d} \|\partial \mathbf{u}(\mathbf{z})\|_F^2 d\mathbf{z}} d(\partial \mathbf{u}). \end{aligned} \quad (31)$$

Note that we made further simplifying approximation by integrating over the space of the Jacobian of the deformation, ∂U , instead of the space of the deformation, U , itself.⁷

In Eq. (31), the only values of $\partial \mathbf{u}$ on which A_0 depends are $\partial_1 \mathbf{u}(\mathbf{z})$ for $\mathbf{z} \in Q_x$. Thus, we separate the integral into the product of three integrals, the first one being:

$$\begin{aligned} A(x\mathbf{v}_1) &\propto \int_{\partial_1\{\mathbf{u}: Q_x \rightarrow \mathbb{R}^d\}} A_0 \left(x\mathbf{v}_1 + \int_0^x \partial_1 \mathbf{u}(t\mathbf{v}_1) dt \right) \\ &\quad \times e^{-\frac{1}{2\sigma^2} \int_0^x \|\partial_1 \mathbf{u}(t\mathbf{v}_1)\|_2^2 dt} d(\partial_1 \mathbf{u}), \end{aligned} \quad (32)$$

and the second and third integrals are:

elements that are interdependent given the linear relationship $\nabla \times \partial \mathbf{u} = \mathbf{0}$. As a result, for an exact solution, the integral must be taken with respect to an independent *subset* of the elements of $\partial \mathbf{u}$ that includes the (independent) set $\partial_1 \mathbf{u}(Q_x)$.

⁷ This change of variables (integrating with respect to $\partial \mathbf{u}$ instead of $\mathbf{u}$) is linear due to the linearity of the differential operator ∂ , as well as invertible due to the translation-free constraint on $\mathbf{u}$. We continue with the relaxing assumption that $\partial \mathbf{u}$ has independent elements. Nevertheless, for $d \geq 2$, the variable set $\partial \mathbf{u}$ is redundant and has a larger dimension than $\mathbf{u}$ does, with

$$\begin{aligned} & \int_{\partial_1\{\mathbf{u}:\mathbb{R}^d\setminus Q_x\rightarrow\mathbb{R}^d\}} e^{-\frac{1}{2\sigma^2}\int_{\mathbb{R}^d\setminus Q_x}\|\partial_1\mathbf{u}(\mathbf{z})\|_2^2d\mathbf{z}}d(\partial_1\mathbf{u}) \\ & \times \int_{\partial_{2,\dots,d}U} e^{-\frac{1}{2\sigma^2}\int_{\mathbb{R}^d}\|\partial_{2,\dots,d}\mathbf{u}(\mathbf{z})\|_F^2d\mathbf{z}}d(\partial_{2,\dots,d}\mathbf{u}), \end{aligned} \quad (33)$$

which are integrals of normal distributions and therefore constant, hence not included in the expression for $A(x\mathbf{v}_1)$ in Eq. (32).

Calculation of $A(x\mathbf{v}_1)$ can be made notationally easier by approximating the inner integrals in Eq. (32) as Riemann sums. We divide $[0, x]$ into n equal intervals ($n \rightarrow \infty$), with $dt \approx x/n$, and define:

$$\mathbf{q}_k := \frac{x}{n} \partial_1 \mathbf{u} \left(\frac{k}{n} x \mathbf{v}_1 \right). \quad (34)$$

The integral is now approximated as:

$$\begin{aligned} A(x\mathbf{v}_1) & \propto \int_{\mathbb{R}^{nd}} A_0 \left(x\mathbf{v}_1 + \sum_{k=1}^n \mathbf{q}_k \right) \\ & \times e^{-\frac{1}{2\sigma^2 x/n} \sum_{k=1}^n \|\mathbf{q}_k\|_2^2} d\mathbf{q}_1 \dots d\mathbf{q}_n. \end{aligned} \quad (35)$$

This is, in fact, n consecutive convolutions of A_0 with a d -dimensional Gaussian,

$$A(x\mathbf{v}_1) \propto \left[A_0(\mathbf{z}) * G\left(\mathbf{z} \middle| \mathbf{0}, \sigma^2 \frac{x}{n} \mathbb{I}\right) * \dots * G\left(\mathbf{z} \middle| \mathbf{0}, \sigma^2 \frac{x}{n} \mathbb{I}\right) \right]_{\mathbf{z}=x\mathbf{v}_1}. \quad (36)$$

Given that convolution of n identical Gaussians results in a Gaussian with n times the variance, we have:

$$A(x\mathbf{v}_1) \propto [A_0(\mathbf{z}) * G(\mathbf{z} \middle| \mathbf{0}, \sigma^2 x \mathbb{I})]_{\mathbf{z}=x\mathbf{v}_1}. \quad (37)$$

We now exploit the rotational invariance of the Gaussian in Eq. (37) and that of the Frobenius norm of the Jacobian in Eq. (29), to generalize Eq. (37) for any $\mathbf{x} \in \mathbb{R}^d$:

$$A(\mathbf{x}) \propto [A_0(\mathbf{z}) * G(\mathbf{z} \middle| \mathbf{0}, \sigma^2 \|\mathbf{x}\|_2 \mathbb{I})]_{\mathbf{z}=\mathbf{x}}. \quad (38)$$

Equation (38) is indeed the update presented in Eq. (10). Despite our use of the convolution notation in Eq. (38), A is not computed via an actual convolution, because the co-variance matrix of the Gaussian kernel varies depending on $\mathbf{x}$, where the result of the convolution is evaluated.

C. Volume Threshold

To threshold the computed probability map of the organ, such as the ELV, we sort the values of the map and keep the top v^* voxels, where the optimal v^* needs to be determined. Assuming that the ground-truth label has v_g voxels, we define $l(\eta)$ as the value of the ground-truth label at the top v^{th} voxel, where $v := \eta v_g$. An ideal probability map, whose top v_g voxels are the ground-truth label, is expected to produce the following boxcar function:

$$l_0(\eta) := \begin{cases} 1 & , \quad \eta \in [0 \quad 1] \\ 0 & , \quad \text{o. w.} \end{cases}. \quad (39)$$

In practice, however, the transition to zero at $\eta = 1$ is less sharp due to inaccurately classified voxels, which we approximate with the following inverted sigmoid function:

$$l_\gamma(\eta) := \frac{1 + \gamma}{1 + \gamma \left(1 + \frac{1}{\gamma}\right)^{(1+\gamma)\eta}}, \quad (40)$$

where γ is a nonnegative constant. Note that $\lim_{\gamma \rightarrow 0} l_\gamma$ is the l_0 defined in Eq. (39) for the ideal probability map, and that $\lim_{\gamma \rightarrow \infty} l_\gamma(\eta) = e^{-\eta}$. Furthermore, the normality of l_γ , i.e. $\int_0^\infty l_\gamma(\eta) d\eta = 1$, guarantees the expected property of $\int_0^\infty l_\gamma(v/v_g) dv = v_g$.

Keeping the top v voxels results in a mask that overlaps with the ground-truth label with the following Dice similarity coefficient:

$$\begin{aligned} DSC &= \frac{2 \int_0^v l_\gamma(v'/v_g) dv'}{v + v_g} = \frac{2 \int_0^\eta l_\gamma(\eta') d\eta'}{1 + \eta} \\ &= \frac{2 \log \left(\frac{\gamma + 1}{\gamma + \left(\frac{\gamma}{\gamma + 1}\right)^{(1+\gamma)\eta}} \right)}{\log \left(1 + \frac{1}{\gamma} \right) (1 + \eta)}. \end{aligned} \quad (41)$$

One can verify that, depending on the value of γ , the η_γ^* that maximizes the above Dice score ranges from $\eta_0^* = 1$ to $\eta_\infty^* = -W_{-1}(-e^{-2}) - 2 = 1.146$, where W_k is the branch k of the Lambert W function. Therefore, according to this model, the optimal number of top voxels of the probability map to keep (to maximize Dice) is $v^* = \eta_\gamma^* v_g$. Choosing the nominal value of $\gamma = 1$ results in $\eta_1^* = 1.141$, which led us to keep the top subset of voxels with a volume 14% larger than that of an average organ (Section II.C.1).⁸ Note that subsequent keeping of only the largest CCs in the resulting mask reduces the number of false-positive voxels, which, in our experiments, outweighed the newly introduced false-negative voxels, further increasing the Dice score.

⁸ In our experiments, choosing $\gamma = 1$ consistently improved the results over $\gamma = 0$. The optimal value of γ , however, may vary depending on several factors,

including the segmented organ, image modality and contrast, signal-to-noise ratio, etc.

REFERENCES

- [1] J. E. Iglesias, and M. R. Sabuncu, "Multi-atlas segmentation of biomedical images: A survey," *Medical Image Analysis*, vol. 24, no. 1, pp. 205-219, 2015.
- [2] M. Cabezas, A. Oliver, X. Lladó, J. Freixenet, and M. Bach Cuadra, "A review of atlas-based segmentation for magnetic resonance brain images," *Computer Methods and Programs in Biomedicine*, vol. 104, no. 3, pp. e158-e177, 2011.
- [3] J. E. Iglesias, M. R. Sabuncu, I. Aganj, P. Bhatt, C. Casillas, D. Salat, A. Boxer, B. Fischl, and K. Van Leemput, "An algorithm for optimal fusion of atlases with different labeling protocols," *NeuroImage*, vol. 106, pp. 451-463, 2015.
- [4] H. Wang, J. W. Suh, S. R. Das, J. B. Pluta, C. Craige, and P. A. Yushkevich, "Multi-Atlas Segmentation with Joint Label Fusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 3, pp. 611-623, 2013.
- [5] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. W. M. van der Laak, B. van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60-88, 2017.
- [6] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 234-241.
- [7] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "VoxelMorph: A Learning Framework for Deformable Medical Image Registration," *IEEE Transactions on Medical Imaging*, vol. 38, no. 8, pp. 1788-1800, 2019.
- [8] J. Krebs, H. Delingette, B. Mailhé, N. Ayache, and T. Mansi, "Learning a Probabilistic Model for Diffeomorphic Registration," *IEEE Transactions on Medical Imaging*, vol. 38, no. 9, pp. 2165-2176, 2019.
- [9] B. D. de Vos, F. F. Berendsen, M. A. Viergever, H. Sokootti, M. Staring, and I. Išgum, "A deep learning framework for unsupervised affine and deformable image registration," *Medical Image Analysis*, vol. 52, pp. 128-143, 2019.
- [10] M. A. Morales, D. Izquierdo-Garcia, I. Aganj, J. Kalpathy-Cramer, B. R. Rosen, and C. Catana, "Implementation and Validation of a Three-dimensional Cardiac Motion Estimation Network," *Radiology: Artificial Intelligence*, vol. 1, no. 4, pp. e180080, 2019.
- [11] P. Aljabar, R. A. Heckemann, A. Hammers, J. V. Hajnal, and D. Rueckert, "Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy," *NeuroImage*, vol. 46, no. 3, pp. 726-738, 2009.
- [12] J. E. Iglesias, M. R. Sabuncu, and K. Van Leemput, "Improved inference in Bayesian segmentation using Monte Carlo sampling: Application to hippocampal subfield volumetry," *Medical Image Analysis*, vol. 17, no. 7, pp. 766-778, 2013.
- [13] I. J. Simpson, M. W. Woolrich, and J. A. Schnabel, "Probabilistic segmentation propagation from uncertainty in registration," in *Proceedings of Medical Image Understanding and Analysis*, 2011.
- [14] M. P. Heinrich, I. J. A. Simpson, B. W. Papież, S. M. Brady, and J. A. Schnabel, "Deformable image registration by combining uncertainty estimates from supervoxel belief propagation," *Medical Image Analysis*, vol. 27, pp. 57-71, 2016.
- [15] I. Aganj, and B. Fischl, "Expected label value computation for atlas-based image segmentation," in *Proc. IEEE International Symposium on Biomedical Imaging*, Venice, Italy, 2019, pp. 334-338.
- [16] C. Kuglin, and D. Hines, "The phase correlation image alignment method," in *Proc. Int. Conference Cybernetics Society*, 1975, pp. 163-165.
- [17] J. J. Pearson, D. C. Hines, S. Golosman, and C. D. Kuglin, "Video-Rate Image Correlation Processor," in *21st Annual Technical Symposium*, 1977, pp. 9.
- [18] S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation," *IEEE Transactions on Medical Imaging*, vol. 23, no. 7, pp. 903-921, 2004.
- [19] M. R. Sabuncu, B. T. T. Yeo, K. V. Leemput, B. Fischl, and P. Golland, "A Generative Model for Image Segmentation Based on Label Fusion," *IEEE Transactions on Medical Imaging*, vol. 29, no. 10, pp. 1714-1729, 2010.
- [20] I. Aganj, M. G. Harisinghani, R. Weissleder, and B. Fischl, "Unsupervised medical image segmentation based on the local center of mass," *Scientific Reports*, vol. 8, no. 1, pp. 13012, 2018.
- [21] B. Fischl, "FreeSurfer," *NeuroImage*, vol. 62, no. 2, pp. 774-781, 2012.
- [22] A. F. Fotenos, A. Snyder, L. Gitton, J. Morris, and R. Buckner, "Normative estimates of cross-sectional and longitudinal brain volume decline in aging and AD," *Neurology*, vol. 64, no. 6, pp. 1032-1039, 2005.
- [23] T. Rohlfing, "Image Similarity and Tissue Overlaps as Surrogates for Image Registration Accuracy: Widely Used but Unreliable," *IEEE Transactions on Medical Imaging*, vol. 31, no. 2, pp. 153-163, 2012.
- [24] B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, K. A. van der, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A. M. Dale, "Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain," *Neuron*, vol. 33, pp. 341-355, 2002.
- [25] A. G. Roy, S. Conjeti, D. Sheet, A. Katouzian, N. Navab, and C. Wachinger, "Error Corrective Boosting for Learning Fully Convolutional Networks with Limited Data," *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, pp. 231-239.
- [26] "LiTS - Liver Tumor Segmentation Challenge; <https://competitions.codalab.org/competitions/17094>," 2017.
- [27] B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," *Medical image analysis*, vol. 12, no. 1, pp. 26-41, 2008.
- [28] H. R. Roth, A. Farag, E. B. Turkbey, L. Lu, J. Liu, and R. M. Summers, "Data From Pancreas-CT," The Cancer Imaging Archive, 2016.
- [29] H. R. Roth, L. Lu, N. Lay, A. P. Harrison, A. Farag, A. Sohn, and R. M. Summers, "Spatial aggregation of holistically-nested convolutional neural networks for automated pancreas localization and segmentation," *Medical Image Analysis*, vol. 45, pp. 94-107, 2018.
- [30] "Pancreas Tumor Database of Memorial Sloan Kettering Cancer Center - Medical Segmentation Decathlon; <http://medicaldecathlon.com>," 2018.
- [31] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, "Attention U-Net: Learning where to look for the pancreas," in *Medical Imaging with Deep Learning (MIDL)*, 2018.
- [32] Y. Zhou, L. Xie, W. Shen, Y. Wang, E. K. Fishman, and A. L. Yuille, "A Fixed-Point Model for Pancreas Segmentation in Abdominal CT Scans," *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, pp. 693-701.