Published in final edited form as:

Med Image Anal. 2011 December ; 15(6): 877–884. doi:10.1016/j.media.2011.05.001.

An atlas-navigated optimal medial axis and deformable model algorithm (NOMAD) for the segmentation of the optic nerves and chiasm in MR and CT images

Jack H. Noble and **Benoit M. Dawant**

Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235, USA

Abstract

In recent years, radiation therapy has become the preferred treatment for many types of head and neck tumors. To plan the procedure, vital structures, including the optic nerves and chiasm, must be identified using CT/MR imagery. In this work we present a novel method for automatically localizing the optic nerves and chiasm using a tubular structure localization algorithm in which a statistical model and image registration are used to incorporate a priori local intensity and shape information. The method results in mean Dice coefficients of 0.8 when compared to manual segmentations over ten test cases, which suggests that our method is more accurate than existing techniques developed for the segmentation of these structures.

Keywords

optic nerves; chiasm; optimal paths; radiation therapy; IMRT; model-based segmentation

1 INTRODUCTION

Segmentation is an important step in the radiation therapy process. Modern IMRT (Inverse Modulated Radiation Therapy) treatment techniques require both the delineation of structures to be irradiated and of structures to be spared, so that an optimal radiation delivery plan can be computed unique to each patient. This segmentation process is time consuming and error-prone. The optic nerves and optic chiasm are sensitive structures in the head that are especially difficult to segment. Figure 1 shows these structures in MR and CT images. The arrows in these images show regions where contrast is lacking in either the CT images (orange, dashed lines), the MR images (yellow, thin arrowhead), or both (blue, solid line). A number of methods have been proposed in recent years to segment these structures automatically. Bekes et al., (2008) have proposed a geometric model-based method for semi-automatic segmentation of the eye balls, lenses, optic nerves and optic chiasm in CT images. They report quantitative sensitivity and specificity results from STAPLE (Warfield

© 2011 Elsevier B.V. All rights reserved.

Corresponding Author Info: Jack Noble (jack.h.noble@vanderbilt.edu) Vanderbilt University Dept. of Electrical Engineering and Comp. Sci. VU Station B, #1662, 2301 Vanderbilt Pl., 37235 Ph: 615-322-8478 Fax: 615-343-5459.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

An earlier version of this paper was printed in The Proceedings of the SPIE Conf. on Medical Imaging, 2009, Vol 7259, pp. 725916-1 —725916-10.

et al., 2004) of approximately 77% and 95% for the optic nerves and approximately 65% and 94% for the chiasm. Qualitatively, they report a general lack of consistency with the results they obtain for the nerves and chiasm. Atlas-based methods are commonly used for image segmentation (see for instance Fischl et al., 2002; and Dawant et al., 1999). In this type of approach, a transformation is computed between an atlas image, where the structures of interest have been manually delineated, and a target image using non-rigid registration techniques (see for instance Rueckert et al., 1998; Rhode et al., 2003; Rohlfing and Maurer, 2003; Dinggang and Davatzikos, 2002). The manual delineations in the atlas can then be projected through the registration transformation onto the target volume, thus automatically identifying those structures. Several atlas-based approaches have been proposed to segment the optic nerves and chiasm for radiation therapy (D'haese et al., 2003; Gensheimer et al., 2007; Isambert et al., 2008). Although atlas-based approaches have been reported to produce good results for most head structures, these approaches have had only limited success for segmentation of the optic nerves and chiasm with Dice similarity coefficients (DSC) ranging from 0.39 to 0.78 for the nerves and 0.41 to 0.74 for the chiasm.

In this work, we present a new and different method to attack the problem. First, because it is more and more common in clinical practice to acquire both MR and CT volumes for treating patients with brain tumors, we use both modalities. Second, as opposed to previous methods, which localize the optic nerves and chiasm independently, we extend the optic nerves past the chiasm and we consider the optic nerve and the contra-lateral optic tract as a single structure. Third, we do not explicitly segment the optic chiasm. Instead, the chiasm is found as the intersection of the two optic nerve/optic tract structures, which is its true anatomic definition. This is illustrated in Figure 2. The first structure of interest (SOI) is made of the left optic nerve and right optic tract (left SOI), the second structure is made of the right optic nerve and left optic tract (right SOI). The intersection between these two structures is the optic chiasm. This approach has two main advantages. First, we do not need to find the optic chiasm directly, which is the most difficult structure to segment. Second, we can view the segmentation of the optic nerves and tracts as the segmentation of two independent tubular structures.

There is a large body of literature that has been published on the segmentation of tubular structures in medical images. For instance, Feng et al., (2004) propose a tubular structure deformable model. The model attempts to iteratively converge to the correct solution by minimizing an energy functional. For the energy function to be effective, the structure of interest must be enhanced, which requires a strong contrast between the structure and the background. Yim et al., (2001) propose another deformable model approach, but it relies on a manual initialization of the axis of the structure. Manniesing et al., (2007) propose a level-set based surface evolution algorithm. The authors construct a speed function which allows faster propagation along the axis of the structure. The evolving front can then be iteratively skeletonized and propagated. The speed function that the authors propose relies on contrast between the structure of interest and the background. Wesarg and Firle (2004) use an iterative border detection scheme to approximate the surface of the structure and estimate the centerline. It is another example of approaches that assume the structure has been enhanced in the image. Santamaria-Pang et al., (2007) propose a probabilistic wave propagation approach to localize the structure's centerline, but assume that the structure of interest is enhanced in the image. Olabarriaga et al., (2003) propose a minimal cost path approach for segmenting the coronary artery. Their approach relies on an artery enhancement preprocessing procedure that assumes the artery is brighter than the background. Hanssen et al., (2004) propose a semi-automatic approach for extracting nerve centerlines using fast marching methods. They also assume that the intensity value of the nerve contrasts with the background in the region of the nerve. Li and Yezzi (2007) propose a 4-D minimal cost path approach to vessel segmentation. The method assumes consistent

intensity characteristics along the length of the structure. These approaches are examples of the state of the art of general approaches for tubular structure segmentations (Lesage et al., (2009) present a more comprehensive review of tubular structure segmentation methods). These types of approaches are unlikely to work for the optic nerves and chiasm due to variation in intensity along the structure's length and the lack of contrast with the background. In this article, we use a method similar to the one we have proposed recently to segment the facial nerve and chorda tympani in CT images (Noble et al., 2008). These are other examples of tubular structures that are difficult to segment because of their size (diameter on the order of 2-3 voxels for the facial nerve and smaller for the chorda) and lack of contrast with surrounding structures. This approach, which we have called NOMAD, for atlas-navigated optimal medial axis and deformable model, combines an optimal path algorithm with a-priori information provided by the registration of a model to the image to be segmented. The novelty of the approach is that it is able to provide the optimal path algorithm with spatially-varying, local, atlas-based intensity and shape priors. The work presented in this paper shows not only that the methods we propose produce very good results for the problem at hand but also that the method we propose is generally applicable to tubular structures for which a reasonable model can be created.

The rest of the paper is organized as follows. Section 2 first presents the general approach and the specific model we use for the application discussed herein. Section 3 presents our results. These results and suggestions for future work are discussed in Section 4.

2 MATERIALS AND METHODS

In this section we first describe the various data sets used. The construction of structure models and how they are used to segment images are discussed in Sections 2.2. Finally, our methods for evaluation, selection of various parameters, and implementation details are discussed in Sections 2.3-5.

2.1 Data

A model training set of four, a parameter training set of ten, and a testing set of ten CT/MR pairs were acquired with IRB approval and used in this study. The CTs were acquired on four different scanners at 120-130 kVp and exposure time of 350-1750 mAs with voxel sizes ranging from 0.5 mm³ isotropic to 1.0 × 1.0 × 3.0 mm³. Six contrast enhanced T₁ weighted MRs were acquired on 1.5 T scanners with voxel size of 1.0 × 1.0 × 1.2 mm³. The remaining eighteen MRs, ten of which compose the testing set, were non-contrast T₁ weighted acquired on 1.5-3 T scanners with voxel size 1.0 mm³ isotropic. Each CT/MR pair was rigidly registered using standard techniques.

2.2 General approach

The method we have developed requires a model of the structures to be segmented. Because we are interested in tubular structures, these models typically include two components. The first captures information related to the properties of the structure's centerline such as the intensity of a voxel located at a specific location along the centerline, or the orientation of the centerline along its length. The second captures information related to the structure's diameter. Once a model is built as discussed in the following sub-section, it is used to segment the structure in other image volumes. This is done as follows. First the image volume in which the model has been defined is registered to a new volume. Using the registration transformation, the model's centerline is projected onto the volume. This is illustrated in Figure 3 for a very simple example: the segmentation of a dark structure on a bright background. The model centerline is shown in (a). In (b), the image to be segmented is shown. The dotted curve represents a structure in the image that is close in shape and

intensity to the model, but it is affected by noise. The straight line that closes the dotted curve represents a structure in the image that we do not wish to segment, but since it is of approximately the correct intensity and connects the starting and ending points, it would represent an issue to most tubular structure segmentation algorithms. The result of the registration is displayed in (c). Note that the registration is inaccurate. Registration error is modeled here because, in general, the registration is imperfect. If the registration was perfect, it would permit the segmentation of the structure using an atlas-based method alone. Once the model is projected onto the image to be segmented, a cost matrix is created by computing, at each voxel, the difference between the voxel's characteristics, e.g., intensity, and the characteristics of the closest model voxel. Finding correspondence between the image voxels and the model is shown in (d). For our simple problem, the resulting intensity based cost image is shown in (e). We also use the model to provide a priori information about the orientation of the centerline as a type of shape cost. This information is useful for branching structures or structures that lack contrast with surrounding structures. This curve orientation information is represented in vector form across the image as shown in (f). Once the costs are defined, starting and ending points are chosen. This is also done via the registration transformation, i.e., a set of starting and ending points are defined in the atlas, and projected onto the other images. Next, a minimum cost path algorithm finds the optimal path between starting and ending points. For this simple example, the path finding algorithm would give us a structure centerline result as shown in (g). Finally, once the centerline is localized, the complete structure is segmented using a level set approach. The speed function used in this step also makes use of *a priori* information provided by the model.

2.2.1 Model Creation—Because the left and right structures we segment are mirror images of each other, we need only to build a model for one of them. We create a model for the left SOI, and a right SOI can be added to the model by reflecting the image across the mid-sagittal plane. Our model consists of three quantities for each point along the centerline: (1) the orientation of the centerline curve E_O , (2) the width of the structure E_W and (3) an intensity vector $\vec{E}_I$ that captures the intensity of voxels surrounding each centerline point. Intensity models are built for the MR and CT images separately but the MR model is used on some sections of the structure while the CT model is used on the others as discussed later.

The models are built as follows.

1. All corresponding MR and CT image volumes are registered to each other using a standard mutual-information rigid-body registration method (Maes et al., 1997; Wells et al., 1996).
2. The structures of interest are segmented manually in each of these MR/CT pairs and the centerline of these structures is extracted using a thinning method. Here we have used the method presented by Bouix et al., 2005.
3. One MR/CT set is selected as the reference set, which will be referred to as the atlas; the other volumes will be referred to as training volumes.
4. A correspondence between points along the centerlines is established. This is done in several steps
 - a. All the training volumes are registered to the atlas using an affine transformation to correct for differences in scale and orientation. Registration was performed by optimizing twelve parameters (translation, rotation, scaling, and skew) using Brent's line search and Powell's direction set algorithm (Press et al., 1992) to maximize the mutual information (Maes et al., 1997; Wells et al., 1996) between the MR

volumes. Parameters are first optimized to convergence by computing the mutual information over the full volumes at a sampling factor of 3 in each dimension, then further optimized by computing the mutual information over the region of interest (ROI) containing the optic nerves and chiasm, which has been manually identified in the atlas. In our experience, this approach is faster and more accurate than a full-resolution registration of the whole head. Note that the MR, rather than the CT, was used for all inter-patient registrations, because we found that transformations computed based on MR data were more accurate in the region of interest.

- b. Registration is refined in the ROI using the non-rigid intensity-based adaptive bases algorithm we have proposed in the past (Rhode et al., 2003), which uses normalized mutual information as a similarity measure (Studholme et al., 1999).
- c. Using the transformations computed in steps (a) and (b), the points on the atlas centerline are projected onto each of the training volumes.
- d. For each model point, the corresponding point is chosen to be the point closest to this projected point on the training volume's centerline.

This procedure results in a set of points $\{\vec{q}_i\}_{i=1}^N$ corresponding to each $\vec{q}_0$ point on the model centerline, where N is the number of training volumes. Once correspondence is established, the three quantities we used in this application are computed.

E_O and E_W are computed as the average curve orientation and average structure width over the corresponding centerline points, where curve orientation is estimated using central differences of the curve points. The intensity vector is computed with the method illustrated in Figure 4 (this figure shows a 2D case but the procedure can readily be expanded to 3D). Panel (a) in the figure shows the pixel of interest, marked with a blue dot, and the surrounding patch of intensities. In panel (b), a red circle is drawn around the point of interest, defining a neighborhood. Here we use a circle of radius R_O (in 3D we define a sphere). Then, every point within that neighborhood is visited following a pre-determined sequence (as shown in (b)), and an intensity vector that contains the pixel (voxel) intensities is created. This process is applied to the atlas centerline point as well as its set of corresponding training centerline points, resulting in N intensity vectors for each atlas centerline point, which are averaged. One common problem with MR images is the variability in intensity characteristics between volumes. To address this issue we use the following procedure. The values of the average intensity vector are first remapped into sequential integer values with the lowest value remapped to zero, the second to 1, etc., to produce the vector $\vec{w}$. An example of this intensity remapping is shown in panel (c). This representation retains the information about the relative intensities of the voxels surrounding the point of interest while discarding exact intensity information. It will thus not be impacted by intensity differences between volumes as long as these intensities can be related by a monotonic function. Finally, the feature vector $\vec{w}$ is obtained by normalizing $\vec{w}$ to length 1. The similarity between two points can be compared simply by computing the Euclidean distance between their respective feature vectors, where a smaller distance indicates greater similarity. The parameter R_O controls the locality of the measurement.

2.2.2 Structure segmentation—Once the models are built, new images can be segmented. In Figure 5, we present a flow chart describing our segmentation approach. In this figure, circles indicate an image (roman letters) or a transformation (greek letters). Rectangles represent operations. T is the target image set, i.e., the CT/MR pair that needs to be segmented, and A is the atlas image set. Using the methods described in Section 2.2.1, the

target image T is first affinely registered to the atlas volume. This results in image T^* and an affine transformation τ^a . A non-rigid registration is then computed between T^* and A over the ROI to produce transformation τ^n . Note that segmentation of the structures will be performed on the image T^* . This rescales and reorients the target image set to match the atlas and permits using the *a priori* shape and intensity information in the model. The segmentation method then proceeds as follows: The transformation τ^n is used to associate each voxel in T^* with expected feature values. Using these expected feature values, a cost function, and a set of curve endpoints, the SOI centerlines are computed using a minimum cost path algorithm. Once the centerline is found, the complete structure is extracted using the level set approach. When the structures have been localized in T^* , they are projected back to T , completing the process. The rest of this section details this approach.

Minimum cost path algorithms require two pieces of information (1) starting and endings points and (2) a cost associated with transitions from each point to its neighbors. To compute costs based on *a priori* information, we use the following approach. First, the atlas centerline points are projected from the atlas to image T^* using τ^n . Next, for each voxel in T^* , the closest projected point is found. The expected feature values for each voxel in T^* are then set to be the feature values of the closest projected point. The cost function, which depends on these expected features, can then be computed anywhere in the image. This approach compensates for small registration errors. Assume, for instance, that one centerline atlas point is projected near, but not on, the structure centerline in T^* at point n . Suppose also that the closest point on the true structure centerline in T^* is c . Using our scheme, the intensity-based cost associated with point n will be higher than the cost associated with point c , assuming that the structure characteristics in T^* are similar to those in the model, which is the fundamental assumption of our approach. A minimum cost path algorithm will thus tend to pass through point c rather than point n , which would be on the path obtained if only a non-rigid registration algorithm was used.

The terms we have included in our cost function are based on the expected intensity feature vectors and the expected curve orientation. For a given location, the feature vector is either based on MR or CT data. Mimicking a typical physician, the optic nerve is localized using CT data while the chiasm is localized using MR data. This approach is advantageous due to the structure contrast observed in the two modalities for these structures (see Figure 1). To automatically decide which image should be used, a point at the optic nerve/chiasm junction along the atlas centerline was chosen so that all image voxels corresponding to atlas points towards the optic nerve from this point use the CT features while voxels that lie towards the chiasm use MR features. The cost for feature vectors is defined by

$$\beta \left(\frac{\|\vec{E}_1 - \widehat{w}\|}{\sqrt{2}} \right)^\gamma, \quad (1)$$

where $\widehat{w}$ is the feature vector of the point of interest and β and γ are weighting and sensitivity parameters. The maximum distance between any two normalized ranking vectors composed of all positive values can be shown to be $\sqrt{2}$. Thus, the constant factor of $\sqrt{2}$ in Eq. (1) is used to scale the range of the function to the interval [0,1] prior to scaling by β and exponentiation by γ . The curve orientation cost is computed by

$$\left(\left(1 - \frac{\vec{t} \cdot \vec{E}_o}{\|\vec{t}\| \|\vec{E}_o\|} \right) / 2 \right)^\delta, \quad (2)$$

where $\vec{t}$ is the direction of the voxel transition (i.e., up, left, right, etc.). This term favors transitions in the direction predicted by the model. The inner fraction in Eq. (2) corresponds to the cosine of the angle between $\vec{t}$ and $\vec{E}_0$, which ranges from -1 to 1. Thus, the constant factor of 2 in Eq. (2) scales the function into the interval [0,1] prior to exponentiation by δ . The total cost for transitioning to any particular voxel in the image is the sum of the intensity and orientation terms in Eqs. (1) and (2).

Finally, endpoints are needed. Instead of using a single starting and ending point in our optimization, we include many and simply use the starting point and ending point that result in the optimal path. Specifically, we include all voxels on a disc centered at the endpoint given by non-rigid registration, perpendicular to the curve, and with the radius equal to the maximum distance from the endpoint to its corresponding points in the training images. This is the radius necessary to ensure that registration error in the endpoint regions can be handled by the segmentation algorithm. These starting and ending points and the cost function are thus the inputs necessary to apply a minimal cost path algorithm.

A minimal cost path algorithm (Dijkstra, 1959) can be used to efficiently find a curve s that minimizes the energy equation

$$\int_0^L C(s(l)) dl, \quad (3)$$

where C is the energy function, by treating it as a problem of finding a minimal cost path between two nodes, or endpoints, in a graph $G=\{V,E\}$. V is the set of nodes, and E is the set of edges connecting neighbor nodes, each of which is associated with a cost. The algorithm is used to find the set of edges connecting two endpoints that are associated with the minimum total cost. This set of edges represents the curve s that minimizes Eqn. (3). In this application, V is composed of the set of voxels, E connects neighboring voxels in a 26-connected neighborhood, and the cost of each edge is the sum of Eqns. (1) and (2). The minimal cost path algorithm is then used to find the curve that minimizes Eqn. (3), which should correspond to the structure centerline.

The segmented structure centerline is used to initialize a standard level set-based geometric deformable model (Sethian, 1999). In this approach, the surface is the zero-level set of an embedding function, ϕ , which obeys the evolution equation

$$\phi_t + F|\nabla\phi| = 0, \quad (4)$$

where F is the speed function, i.e., the function that specifies the speed at which the surface evolves along its normal direction. The speed function we have designed is computed as

$$E_w \left(1 - \frac{\|\vec{E}_r - \vec{w}\|}{\sqrt{2}} \right), \quad (5)$$

where E_w , $\vec{E}_r$, and $\vec{w}$ are the expected width, expected feature vector, and actual feature vector at the point of interest. The expected width term ensures this function is larger where the structure is expected to be wider. The intensity feature vector term increases function values where the intensity features are more similar to the expected centerline intensity features. Thus, function values tend to be relatively large at the structure centerline and slowly decrease inside the structure away from the centerline. The intensity feature terms

ensure that function values drop more drastically at the edge of the structure. To initialize the level set algorithm, an initial level set distance map is computed, treating the voxels corresponding to the extracted structure centerline as interior points and all other voxels as exterior points. Traditionally with level set approaches, Eq. (4) is solved iteratively until the level set front stabilizes and stops. This is not possible here because structure contrast is very weak. To address this issue, the speed function we have designed, Eq. (5), is a strictly expanding function, and we solve Eq. (4) only once over time Δt , where Δt is used as a parameter to control the total expansion of the zero-level set.

2.3 Segmentation evaluation

Gold standards for the training and testing sets were obtained from manual delineations of the structures of interest. For the training sets, the contours were first drawn by a student rater then corrected by two experienced radiation oncologists concurrently. Two independent sets of delineations were obtained for the ten MR/CT pairs in the testing set. These delineations were created first by two separate student raters. Then, one set was corrected by an experienced radiation oncologist, and the other was corrected by an experienced radiologist. To evaluate accuracy, the automatic segmentations are compared to these manual delineations using the DSC (Dice similarity coefficient (Dice, 1945). Mean and maximum surface errors are also reported. The point of transition of the optic nerves into the chiasm is somewhat subjective. In fact, the optic nerves and chiasm are simply particular regions of the single optic pathway structure, and they are treated with equal priority in radiotherapy planning. Because of this, and because there are no visual cues of this transition between nerves and chiasm in MR/CT, the manual delineations and the automatic segmentations represent the optic nerves and chiasm as the single optic pathway structure. To evaluate accuracy separately for the optic nerves and chiasm, we use an approach similar to Gensheimer et al., 2007. After segmentation, the coronal plane that appropriately separates the optic nerves and chiasm was identified manually for each volume tested. Then, for each volume, error was measured for the optic nerves by comparing the portion of the automatic and manual segmentations that lies anterior to this plane. Similarly for the chiasm, error is measured by comparing the portion of the segmentations posterior to this plane.

2.4 Parameter Selection

For each parameter, an acceptable value was first chosen by visually observing the behavior of the algorithm as a function of parameter values. Once an acceptable value was found, it was modified in the direction that increased the mean DSC over the training volumes until a value was found at which the error clearly increased. The final value was chosen away from this point in the generally flat error region preceding it. Once all final values were chosen, the DSC for each volume was recorded. Next, the sensitivity of the algorithm to parameter values was analyzed. This was done by sequentially modifying parameter values in 10% increments around their selected value until the DSC dropped below 0.5 on any training volume.

2.5 Implementation

The presented methods were implemented in C++ and tested on an Intel Core 2 Duo 2.4 GHz, Windows XP-based computer. To compute the minimal cost path, we implemented Dijkstra's algorithm. Pseudocode of our implementation is shown in Algorithm 1. First, the set of candidate seed nodes are initialized and inserted into the priority queue. m is arbitrarily assigned to one of these seed voxels. The search is conducted inside a while-loop. At each iteration, the neighbors of m that have not been visited are found. For each free neighbor, m is recorded as its parent node, and its cost is computed as the sum of the cost associated with m and the cost of the arc between itself and m . The neighbor is then inserted

into the queue with priority equal to its cost. Next, m is assigned to be the node popped from the head of the queue, which is the node associated with the minimal cost. m is marked as having been visited, and if m is not an end node, the next iteration begins. This breadth-first searching scheme ensures that once an end node is popped from the head of the queue, the path of minimal cost, P , can be recorded by tracing the marked parent node pointers from the end node back to the seed associated with that path. Note that all of the candidate seed points are included in the optimization, and the final seed point chosen by the algorithm is the one associated with the path of minimal cost.

3 RESULTS

The final values of the various parameters used and the results of sensitivity analysis are listed in Table 1. As can be seen from the table, the most sensitive parameter R_O can be changed by 40% before causing DSC in any volume to drop below 0.5 in our testing set.

Segmentation of the optic nerves in a test volume requires approximately 20 minutes. DSC between the automatic segmentations and the manually delineated training and testing sets are shown in the left graph of Figure 6. This figure also shows the inter-rater results and the mean results reported by other articles. The distribution range, median, mean, and one standard deviation from the mean are indicated by the black I-bar, green line, red line, and blue box, respectively. The purple circles, red x's, and blue squares indicate approximate results reported by Gensheimer et al., 2007, D'haese et al., 2003, and Isambert et al., 2008, respectively. Bekes et al., 2008 do not report DSC. As can be seen in the figure, mean DSC achieved for the testing set are just below 0.8 for both the optic nerves and at 0.8 for the chiasm. The distributions of the testing set results are similar to those of the training set. All experiments result in DSC that are larger than mean results reported by other methods, with the exception of the optic nerve results achieved by Gensheimer et al., 2007, where comparable results are achieved. Shown in the middle and right graphs of the figure are plots of the mean and maximum surface errors in mm. Error distributions are indicated similarly to the DSC plot. As can be seen from the figures, the automatic segmentations result in mean errors of approximately 0.5 mm (about half a voxel), maximum overall errors of approximately 4.0 mm (about 4 voxels), and typical maximum errors of less than 1.0 mm (less than a voxel). As can also be seen in the plots, the agreement between the gold standard segmentations, which were manually delineated carefully in a laboratory setting, is slightly better than the agreement between the automatic and gold standard segmentations.

Renderings of several cases are shown in Figure 7. From top to bottom, the figure contains MR data (non-contrast datasets on left, contrast enhanced datasets on right), CT data, and a 3D rendering of the structures of interest. In the 2D panes, the automatic segmentation is shown as the green solid contour and manual segmentation is shown with the purple dotted contour. To generate the 2D views, a thin-plate spline transformation was computed that warps the medial axes of the structures to lie in a plane, and then the image data and contours were passed through this transformation. This was done so that a cross section of the entire structure could be viewed in one 2D plane. In the 3D views, the automatically segmented optic nerves (red) and chiasm (blue) are shown. All results appear to be qualitatively accurate and reasonable.

4 DISCUSSION

The method we propose requires selecting a total of 5 parameters. Although presently this has required an initial manual adjustment, our study shows that our method is not extremely sensitive to the selection of these parameters. Any of these parameters can be changed by more than 40% before the DSC with respect to manual segmentations drops below 0.5. The

model has been constructed using 4 CT/MR pairs of training volumes with MRs that were acquired with a single protocol (non-contrast 1.5 T), parameters trained using 10 other image sets with MRs acquired on various scanners according to varying protocols (four non-contrast 1.5 T and six contrast-enhanced 1.5 T), and tested on 10 other image sets with MRs acquired on various scanners according to varying protocols (contrast-enhanced 1.5-3 T). This shows that our method is robust to changes in intensity characteristics, including variations not seen in the training set. Our results also show that our method is accurate. Typically, a DSC of 0.8 is considered good (Zijdenbos et al., 1994), although the DSC is typically more unforgiving on very thin structures such as the optic nerves and chiasm. With our approach, we achieve mean DSC of approximately 0.8 with an overall minimum DSC of 0.74.

To the best of our knowledge, these results suggest that the method presented in this article is the most accurate method to date to accomplish automatic segmentation of the optic nerves and chiasm. The segmentation error we have obtained is comparable to the inter-rater difference observed when contours are delineated without time constraint in a laboratory setting. Error in manual delineation obtained under these conditions can be considered a lower bound. The DSC measure is sensitive to factors such as image resolution, and the only truly fair approach for comparing the performance of two algorithms is to test them on the same dataset. Although it was not possible in this study to use the same datasets as those used in previous studies, the consistently higher DSC we achieve as indicated by the distributions in Figure 6 suggest that the methods presented in this paper outperform previously presented approaches for localizing these structures. This claim is supported by the fact that Bekes et al., 2008, D'haese et al., 2003, and Isambert et al., 2008 conclude that their automatic segmentations of the optic nerves and chiasm are not sufficiently accurate for radiotherapy planning. Similarly, Gensheimer et al., 2007 report dissatisfaction with their chiasm segmentation results. Their qualitative assessments support our quantitative comparison, which does suggest that the method we propose is more accurate.

Also, to the best of our knowledge, this article presents the first attempt to segment the optic nerves and chiasm as a union of two tubular structures. In our experience, and based on the results of other research (Bekes et al., 2008; D'haese et al., 2003; Gensheimer et al., 2007; Isambert et al., 2008), purely atlas-based methods have proven to be ineffectual for this problem. Treating the optic nerves and chiasm as two tubular structures allows the use of many general tubular structure localization algorithms, such as those based on optimal path finding approaches. However, due to lack of contrast and changing characteristics along the structures length, typical optimal path-based approaches would likely be ineffectual. The solution lies in providing the optimal path finding algorithm with *a priori* local intensity and shape feature information, which is the approach used by the NOMAD algorithm.

Finally, the analytical form of the cost function we have designed for this application is quite general in applicability. The approach we have described can be used to build feature models and segment tubular structures in other applications by simply providing a training set and doing some parameter tuning. We are currently working to apply this idea to a variety of nerves and vessels in the head and neck.

Acknowledgments

This research has been supported by grants R01EB006193 from the National Institute of Biomedical Imaging and Bioengineering and F31DC009791 from the National Institute of Deafness and Other Communication Disorders. The content is solely the responsibility of the authors and does not necessarily represent the official views of these institutes. The authors would like to thank Dr. Kenneth Nierman and Dr. Edwin Donnelly for reviewing and editing manual segmentations.

REFERENCES

- Bekes G, Máté E, Nyúl L, Kuba A, Fidrich M. Geometrical model-based segmentation of the organs of sight on CT images. *Med. Phys.* 2008; 35:735. [PubMed: 18383695]
- Bouix S, Siddiqi K, Tannenbaum A. Flux driven automatic centerline extraction. *Medical Image Analysis.* 2005; 9:209–221. [PubMed: 15854842]
- D'Haese P, Duay V, Li R, du Bois d'Aische A, Merchant T, Cmelak A, Donnelly E, Niermann K, Macq B, Dawant B. Automatic segmentation of brain structures for radiation therapy planning. *Proceedings of SPIE: Medical Imaging 2003: Image Processing.* 2003; 5032:517–526.
- Dawant B, Hartmann SL, Thirion JP, Maes F, Vandermeulen D, Demaerel P. Automatic 3D Segmentation of Internal Structures on the Head in MR Images Using a Combination of Similarity and Free Form Transformations: Part I, Methodology and Validation on Normal Subjects. *IEEE Trans. Med. Imaging.* 1999; 18:909–916. [PubMed: 10628950]
- Dice LR. Measures of the amount of ecologic association between species. *Ecology.* 1945; 26:297–302.
- Dijkstra EW. A note on two problems in connexion with graphs. *Numerische Mathematik.* 1959; 1:269–271.
- Dinggang S, Davatzikos C. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imag.* 2002; 21:1421–1439.
- Feng J, Ip HHS, Cheng SH. A 3D Geometric Deformable Model for Tubular Structure Segmentation. *MMM.* 2004; 10:174–180.
- Fischl B, Salat D, Busa E, Albert M, Dieterich M, Haselgrove C, van der Kouwe A, Killiany R, Kennedy D, Klaveness S, Montillo A, Makris N, Rosen B, Dale A. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron.* 2002; 33:341–355. [PubMed: 11832223]
- Gensheimer M, Cmelak A, Niermann K, Dawant B. Automatic delineation of the optic nerves and chiasm on CT images. *Proceedings of SPIE: Medical Imaging 2007: Image Processing.* 2007; 6512:41.
- Hanssen N, Burgielski Z, Jansen T, Lievin M, Ritter L, von Rymon-Lipinski B, Keeve E. Nerves - level sets for interactive 3D segmentation of nerve channels. *IEEE International Symposium on Biomedical Imaging: Nano to Macro.* 2004; 1:201–204.
- Isambert A, Dhermain F, Bidault F, Commowick O, Bondiau P, Malandain G, Lefkopoulos D. Evaluation of an atlas-based automatic segmentation software for the delineation of brain organs at risk in a radiation therapy clinical context. *Radiother. Oncol.* 2008; 87(1):93–9. [PubMed: 18155791]
- Lesage D, Angelini E, Bloch I, Funka-Lea G. A review of 3D vessel lumen segmentation techniques: Models, features, and extraction schemes. *Medical Image Analysis.* 2009; 13:819–845. [PubMed: 19818675]
- Li H, Yezzi A. Vessels as 4D Curves: Global Minimal 4D Paths to Extract 3D Tubular Surfaces and Centerlines. *IEEE Trans. Med. Imag.* 2007; 26(9):1213–1223.
- Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P. Multimodality image registration by maximization of mutual information. *IEEE Trans Med Imag.* 1997; 16:187–198.
- Manniesing R, Viergever MA, Niessen WJ. Vessel Axis Tracking Using Topology Constrained Surface Evolution. *IEEE Trans. Med. Imag.* 2007; 26:309–316.
- Noble J, Warren FM, Labadie RF, Dawant BM. Automatic segmentation of the facial nerve and chorda tympani in CT images using spatially dependent feature values. *Medical Physics.* 2008; 35(12):5375–5384. [PubMed: 19175097]
- Olabarriaga SD, Breeuwer M, Niessen WJ. Minimum cost path algorithm for coronary artery central axis tracking in CT data. *Lecture Notes in Comp. Sci.-MICCAI.* 2003; 2879:687–694.
- Press, WH.; Flannery, BP.; Teukolsky, SA.; Vetterling, WT. *Numerical Recipes in C.* 2nd ed.. Cambridge University Press; Cambridge, U. K.: 1992. p. 412-419.
- Rohde GK, Aldroubi A, Dawant BM. The adaptive bases algorithm for intensity-based nonrigid image registration. *IEEE Trans. Med. Imag.* 2003; 22:1470–1479.

- Rohlfing T, Maurer CR Jr. Nonrigid image registration in shared-memory multiprocessor environments with application to brains, breasts, and bees. *IEEE Trans. on Info. Tech. in Biomed.* 2003; 7:16–25.
- Rueckert D, Sonoda L, Hayes C, Hill DLG, Leach MO, Hawkes DJ. Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images. *IEEE Trans. Med. Imag.* 1998; 18:712–721.
- Santamaría-Pang A, Colbert CM, Saggau P, Kakadiaris IA. Automatic centerline extraction of irregular tubular structures using probability volumes from multiphoton imaging. *Lecture Notes in Comp. Sci.-MICCAI.* 2007; 10:486–494.
- Sethian, J. *Level Set Methods and Fast Marching Methods.* 2nd ed.. Cambridge University Press; Cambridge, U. K.:
- Studholme C, Hill DLG, Hawkes DJ. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition.* 1999; 32(1):71–86.
- Warfield S, Zou K, Wells W. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans Med Imaging.* 2004; 23(7):903–921. [PubMed: 15250643]
- Wells WM III, Viola P, Atsumi H, Nakajima S, Kikinis R. Multi-modal Volume Registration by Maximization of Mutual Information. *Medical Image Analysis.* 1996; 1:35–51. [PubMed: 9873920]
- Wesarg S, Firlle EA. Segmentation of vessels: The corkscrew algorithm. *Progress in Biomedical Optics and Imaging - Proceedings of SPIE.* 2004; 5370:1609–1620.
- Yim PJ, Cebal JJ, Mullick R, Marcos HB, Choyke PL. Vessel surface reconstruction with a tubular deformable model. *IEEE Trans. Med. Imag.* 2001; 20:1411–1421.
- Zijdenbos AP, Dawant BM, Margolin R. Morphometric Analysis of White Matter Lesions in MR Images: Method and Validation. *IEEE Transactions on Medical Imaging.* 1994; 13(4):716–724. [PubMed: 18218550]

Research Highlights

> We test a tubular structure localization algorithm. > A statistical model incorporates a priori local intensity and shape information. > Segmentation of the optic nerves and chiasm results in dice coefficients of 0.8. > Results suggest that our approach is more accurate than existing techniques.

Figure 1.

Transverse slices containing the optic nerves (anterior contours) and chiasm (posterior contour) in MR (top row) and CT (bottom row). The arrows indicate areas where the structures are not well contrasted in MR (yellow, thin arrowhead), CT (orange, dashed line), or both (blue, solid line).

Figure 2.
The optic nerves, chiasm, and tracts represented as two intersecting tubular structures.

Color figure for web version

Grayscale figure for printed version

Figure 3.

2D example of the segmentation approach. (a) The model centerline. (b) A target image to be segmented. (c) The atlas centerline (red) registered to the target image. (d) Blue lines show correspondences found from points in the target image to the atlas centerline. (e) Intensity based cost image with blue=low and red=high cost. (f) Illustration of the vectors that drive the shape cost. (g) Centerline segmentation based on the intensity and shape cost terms.

Color figure for web version

Grayscale figure for printed version

Figure 4.

Construction of an intensity feature vector. (a) shows the pixel of interest surrounding patch of intensities. The sequence of neighboring pixels is shown in (b). In (c), a vector of the remapped intensities of the neighborhood is shown.

Figure 5.
Flow chart of the segmentation process

Color figure for web version

Grayscale figure for printed version

Figure 6.
Dice coefficients and mean/max surface errors of the automatic segmentation results

Color figure for web version

Grayscale figure for printed version

Figure 7.
Renderings of the automatic segmentation results

Table 1

Parameter values and sensitivities

Parameter	R_0 (mm)	β	γ	δ	Δt
Value	3	2	0.05	4	0.35
Sensitivity (%)	40	90	100	60	50

Algorithm 1

Search algorithm for path of minimal cost

```

 $O := \{\text{seed nodes}\}; \text{cost}(\{\text{seed nodes}\}) = 0$ 
priority_queue.insert( $\{\text{seed nodes}\}, 0$ )
 $m := \text{priority\_queue.pop\_head\_node}()$ 
while  $m \neq \text{end node}$ 
     $\{n\}_i = \text{neighbor\_nodes}(m)$ 
    for  $\forall i \text{ s.t. } n_i \notin O$ 
        prev_node( $n_i$ ) =  $m$ ;  $\text{cost}(n_i) = \text{arc\_cost}(m, n_i) + \text{cost}(m)$ 
        priority_queue.insert( $n_i, \text{cost}(n_i)$ )
    endfor
     $m := \text{priority\_queue.pop\_head\_node}()$ 
     $O := O \cup m$ 
endwhile
 $P := \{m\}$ 
while  $\exists \text{ prev\_node}(m)$ 
     $m := \text{prev\_node}(m)$ 
     $P := P \cup m$ 
endwhile

```