
Automatic Segmentation of Head and Neck CT Images by GPU-Accelerated Multi-atlas Fusion

Xiao Han, Lyndon S. Hibbard, Nicolette O'Connell, and Virgil Willcut

July 23, 2009

CMS Software, Elekta Inc., 13723 Riverport Drive, Maryland Heights, MO 63043, USA

Abstract

Treatment planning for high precision radiotherapy of head and neck (H&N) cancer patients requires accurate delineation of critical structures. Manual contouring is tedious and often suffers from large inter- and intra-rater variability. In this paper, we present a fully automated, atlas-based segmentation method and apply it to tackle the H&N CT image segmentation problem in the MICCAI 2009 3D Segmentation Grand Challenge. The proposed method employs a multiple atlas fusion strategy and a hierarchical atlas registration approach. We also exploit recent advancements in GPU technology to accelerate the deformable atlas registration and to make multi-atlas segmentation computationally feasible in practice. Validation results on the eight clinical datasets distributed by the MICCAI workshop showed that the proposed method gave very accurate segmentation of the mandible and the brainstem, with a volume overlap close to or above 90% for most subjects. These results suggest that our method is clinically applicable, accurate, and may significantly reduce manual labor and improve contouring efficiency.

1 Introduction

Head and neck (H&N) cancer is the sixth most common cancer in the world and accounts for more than 40,000 new cases annually in the United States and about half a million cases worldwide. Due to the complex geometry of the various target volumes and organs at risk (OARs), it is desirable to apply intensity-modulated radiotherapy (IMRT) for their treatment. The successful implementation of IMRT, however, requires accurate delineation of target volumes and OARs on planning computed tomography (CT) images. Manual contouring, although still the common standard for high quality segmentation in clinics, is tedious, time-consuming, and suffers from large intra- and inter-rater variability. Tools for automated segmentation are thus needed.

Fully automated segmentation of CT images has been proven to be a very challenging problem due to image noise and other artifacts, as well as limited image contrast for most soft-tissue structures. Among the few existing methods proposed in the literature for the automated segmentation of H&N CT images, atlas-based techniques have been the most popular choice largely due to their ability to incorporate prior anatomical knowledge about structure shapes and their geometric relationship into the segmentation process [1, 2, 3, 4, 5, 6]. In such a method, a novel subject image is first aligned to an atlas image – an image that has the structures of interest already labeled – through image registration and the atlas labels are then mapped to the subject image using the computed image transformation. This approach has two critical components that determine the quality of the segmentation results, namely the atlas registration method and the atlas

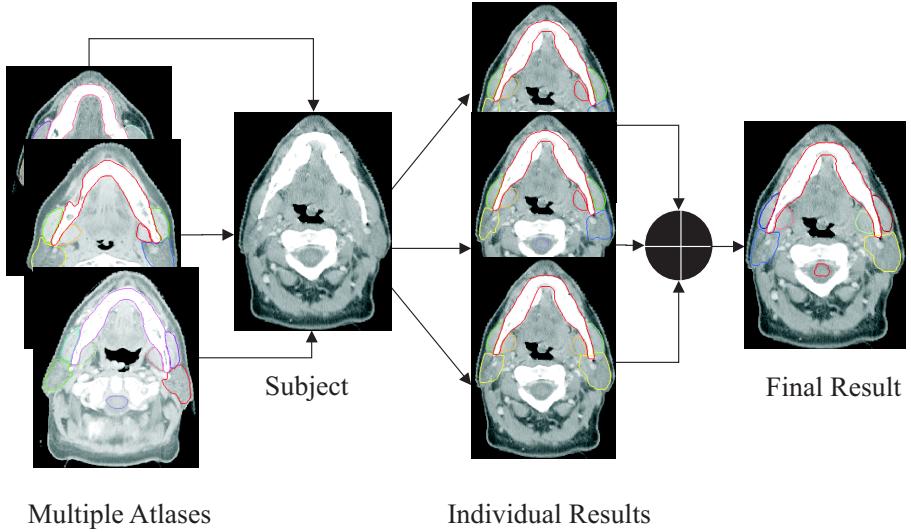


Figure 1: Overall workflow of H&N CT image segmentation using multi-atlas fusion.

selection or construction strategy. Different atlas-based methods most often differ by the registration method they apply and the type of atlases used.

The method we propose in this paper follows a similar framework as we developed earlier [5], which consists of a hierarchical atlas registration method and a multiple atlas fusion strategy. We introduce here a few modifications to the previous method to further improve computation speed and accuracy, and to adapt it to the specific data of this MICCAI workshop. In particular, we have added an extra atlas registration step to account for head-pose differences among different patients. We also proposed using a deformable model method to further improve the segmentation accuracy of the mandible. In addition, we have implemented the dense atlas registration on commodity graphics cards (GPUs) in order to take advantage of the superb computation power of modern GPUs and to make multi-atlas segmentation computationally feasible in practice. The GPU-acceleration can easily offer a speed-up of more than $25\times$, thus allowing segmentation using multiple atlases be finished within the time-frame of a single atlas segmentation on high-speed CPU.

In the following, we first briefly summarize the proposed multi-atlas segmentation framework and then describe our atlas registration method in more details. Finally, we present the validation results based on the test datasets distributed by the MICCAI 2009 Head and Neck Auto-Segmentation Challenge workshop.

2 Method

2.1 Overview

The overall framework of our proposed method is summarized in Fig. 1, where the segmentation of a new subject is computed by applying multiple atlases separately and then combining the individual segmentation results. As demonstrated in our previous work [5], this multi-atlas fusion strategy usually provides much better accuracy than using a single atlas due to the high anatomical variability among different subjects in the head and neck region. Another popular strategy to make use of multiple atlases is to first compile an average atlas from the multiple training data and then apply it as a single atlas [4]. But from our experience and as also observed by others [7], an average atlas usually lacks of detailed anatomical information due to cross-

subject averaging, thus limiting the segmentation accuracy that can be achieved. In this work, we apply the multiple individual atlases independently in the beginning to get multiple segmentations of the same subject, which are then combined in a multi-classifier framework to get a final unique segmentation [8, 7]. Following [7], we use the STAPLE (*Simultaneous Truth And Performance Level Evaluation*) algorithm introduced by Warfield et al. [8] as the segmentation fusion method. The STAPLE method was originally proposed for the estimation of the unknown ground truth from multiple manual segmentations of the same image, but the same principle can be applied to the fusion of multiple auto-segmentation results.

The multi-atlas fusion strategy still requires a reliable atlas registration method in order to get good segmentation result for each individual atlas. In our previous work [5], we have developed a hierarchical atlas registration method for accurate and robust atlas registration. A major consideration in the previous design is to gradually increase the degrees-of-freedom of the image transformation model in order to handle large inter-subject anatomical variations more robustly. We adopt a similar strategy in this work but introduce some modifications. Most importantly, we add an extra step to explicitly correct for head-pose differences between the atlas and the subject after the initial linear registration. We also add a postprocessing step to refine the mapped atlas structures using a deformable surface model method. Due to computation speed and memory considerations, we replace the *local correlation coefficients* (LCC) image similarity metric used in the previous work with a mutual information-based metric. We note that since only two structures (the mandible and the brain-stem) are involved in this study, the *poly-smooth* non-linear registration step in the previous method is no longer necessary and is skipped for all experiments reported in this paper. The final single-atlas registration and structure mapping method consists of four major steps, as described in more details in the following subsections.

2.2 Linear Registration

The first linear registration step aims to correct for global differences in position, orientation, and size between an atlas and a subject image. We use a 9 degrees of freedom linear registration model where the free parameters include three translation, three rotation, and three scaling factors. To find the optimal transformation parameters, we maximize the global mutual information (MI) between the atlas and subject images using a multi-resolution stochastic gradient-descent optimization scheme [9]. This scheme exploits the fact that the linear registration is an over-constrained problem since the number of transformation parameters is far less than the number of voxels in the image domain. Thus, at each iteration, the stochastic scheme uses only a very small subset (100 in our experiments) of random image samples to estimate the image similarity metric and the transformation parameters, which leads to a highly efficient algorithm that usually takes only a few seconds even without GPU-acceleration.

2.3 Head-pose Correction

The previous linear registration step aligns two images (the atlas and the subject) globally. Unfortunately, it cannot account for head-pose differences between different subjects caused by the articulated motion of the head with respect to the body, as illustrated in Fig. 2. The head-pose differences typically lead to large misalignment of structures in the head region (e.g., the brainstem and the skull) even after the initial linear registration and cause difficulty for subsequent non-linear atlas registration. Thus, we propose to address this problem explicitly by designing an algorithm to correct for head-pose differences between the atlas and the subject images.

For typical H&N CT images, the head pose differences appear mainly as a front-back rotation of the head within the mid-sagittal plane. We thus propose to model the head-pose alignment using a two-component

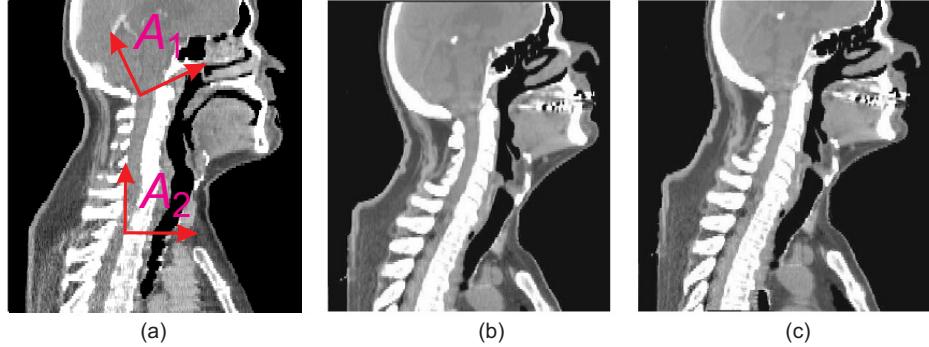


Figure 2: Illustration of head-pose alignment: (a) atlas image; (b) subject image after global linear registration; and (c) subject image after head-pose alignment.

poly-rigid transformation model (cf. [10]) defined within the mid-sagittal plane of the atlas image, as illustrated in Fig. 2(a). In this model, the image transformation T is parameterized by two rigid components:

$$T(\mathbf{x}) = \sum_{i=1}^2 w_i(\mathbf{x}) A_i(\mathbf{x}), \quad (1)$$

where $A_i(\mathbf{x})$ denotes either one of the two rigid transformation components, and w_i is the corresponding weighting function. The value of w_i at each image point \mathbf{x} is chosen to be inversely proportional to the distance of \mathbf{x} to the center or *anchor point* of each transformation component. Following [10], a smooth and invertible approximation of T can be computed using a Log-Euclidean framework.

In this work, the anchor points for the two rigid components are chosen automatically based on the location of the labeled brainstem in the atlas image. In particular, the anchor point for the first component is set at the bottom of the brainstem in the mid-sagittal plane, whereas the anchor point for the second component is chosen by shifting the first anchor point vertically in the mid-sagittal plane towards the bottom of the image (cf. Fig. 2(a)). In all experiments presented later, we empirically set the second anchor point to be 10 cm away from the first one. We note that the mid-sagittal plane location can be easily estimated based on the labeled brainstem.

The main parameters to estimate for the head-pose model are the rotation angle and the translation parameters for the first rigid component. We estimate these parameters based on the matching of image points belonging to boundaries of bones in the head region, in particular the skull and the mandible. Voxels belonging to the skull boundary can be detected using simple intensity thresholding. We note that a complete segmentation of the skull is not necessary at this step. We also assume that the mandible is already labeled in the atlas image. The reason that we only use boundary voxels of bones for matching at this step is because these points have good contrast in CT images and their correspondences are thus easier and more robust to compute.

For each chosen atlas point, its correspondence in the subject image is found using a block-matching procedure, where a block-wise mutual information is used as the local point-wise image similarity measure [1, 5]. We then apply a least-square fit to estimate the rotation angle and the translation parameters. We emphasize again that the computation is performed in 2D within the mid-sagittal plane (the same transformation is then applied to each sagittal plane of the 3D volume). Thus, the computation is fast and usually takes only about a couple seconds.

After parameters of the first rigid component are estimated, parameters of the second component are assigned correspondingly. In particular, the translation parameters are set equal to that of the first component, and the rotation angle is simply set to zero, since the goal of this step is to model the relative rotation of the head with respect to the neck and the body regions. An example result is shown in Fig.2(c). As can be seen, after the head-pose alignment, the subject image resembles more closely to the atlas image, which makes the following non-linear atlas registration a lot easier.

2.4 Dense Mutual-Information Deformable Registration

In the final atlas registration step, we aim to compute a dense deformable registration to align detailed structures of the atlas and the subject images. We adopt the MI as the image similarity measure due to its ability to handle inter-subject intensity and contrast variations. MI is also faster to compute than the LCC metric we previously used in [5]. The MI of two images I and J measures the degrees of dependence between I and J and can be approximated as (cf. [9]):

$$\text{MI}(I, J) = \frac{1}{N} \sum_{\mathbf{x}} \log \left[\frac{p_{I,J}(I(\mathbf{x}), J(T(\mathbf{x})))}{p_I(I(\mathbf{x})) p_J(J(T(\mathbf{x})))} \right], \quad (2)$$

where N denotes the total number of image points within the overlapped region of the two images and T denotes the image transformation model. The functions $p_{I,J}(\cdot, \cdot)$ and $p_I(\cdot)$, $p_J(\cdot)$ denote the joint and marginal image intensity distributions respectively.

We adopt a non-parametric transformation model at this step where the image transformation is modeled directly as a vectorial displacement field \mathbf{U} , such that $T(\mathbf{x}) = \mathbf{x} + \mathbf{U}(\mathbf{x})$ for every image point \mathbf{x} (cf. [12]). Such a dense deformation field model is more flexible than a parameterized model, and can take into account all image information and allow detecting fine anatomical differences.

The solution for the optimal deformation field can be found using a compositive, explicit local search scheme that we recently introduced in [13]. In particular, the deformation field \mathbf{U} is updated iteratively according to the following equation:

$$\mathbf{U}^n = \mathbf{U}^{n-1} \circ (\mathbf{Id} + \mathbf{u}^n) + \mathbf{u}^n, \quad (3)$$

where \mathbf{Id} denotes the identity transformation and “ \circ ” denotes transformation composition. \mathbf{u}^n is a local update field that maximizes the MI between the original atlas image I and the deformed subject image $J^n = J \circ (\mathbf{Id} + \mathbf{U}^n)$ at each iteration, in particular,

$$\begin{aligned} \text{MI}(I, J^n) &\circ (\mathbf{Id} + \mathbf{u}^n) \\ &= \frac{1}{N} \sum_{\mathbf{x}} \log \left[\frac{p_{I,J^n}(I(\mathbf{x}), J^n(\mathbf{x} + \mathbf{u}^n(\mathbf{x})))}{p_I(I(\mathbf{x})) p_{J^n}(J^n(\mathbf{x} + \mathbf{u}^n(\mathbf{x})))} \right] \\ &\triangleq \frac{1}{N} \sum_{\mathbf{x}} \text{SMI}(I(\mathbf{x}), J^n(\mathbf{x} + \mathbf{u}^n(\mathbf{x}))). \end{aligned} \quad (4)$$

In the above equation, $\text{SMI}(\cdot, \cdot)$ denotes a point-wise MI measure, a notation initially introduced in [14].

As proposed in [13], the optimal update $\mathbf{u}^n(\mathbf{x})$ at each atlas image location \mathbf{x} can be found by searching the local neighbors of \mathbf{x} in the deformed subject image J^n and finding the neighbor \mathbf{x}' that maximizes $\text{SMI}(I(\mathbf{x}), J^n(\mathbf{x}' = \mathbf{x} + \mathbf{u}^n(\mathbf{x})))$. This simple explicit neighbor search scheme avoids the difficulty in determining the optimal time step size as needed for a traditional gradient-descent optimization approach. It also eliminates the need to compute derivatives of either the images or the image intensity distributions. A

symmetric version of the local update scheme was designed in [13] that further improves the convergence rate of the registration method.

This dense deformable registration is the most time-consuming part of the whole atlas-based segmentation method. To improve the computation efficiency and make multi-atlas segmentation feasible in practice, we have implemented the above proposed method on NVIDIA GPUs using the NVIDIA CUDA programming model [15]. As we have observed, the GPU-based implementation easily offers a speed-up of more than $25\times$ comparing against a modern CPU [13]. Details of the GPU implementation can be found in [13].

2.5 Refinement of Mandible Segmentation using Deformable Surface Model

Although model-based segmentation is usually considered as an alternative approach with respect to atlas-based methods, combining the strengths of both can lead to much improved accuracy. For example, with proper initialization, a deformable model method can easily produce accurate segmentation for structures with good contrast and well-defined boundaries, whereas a pure atlas-based method may be hindered by large shape differences between the subject image and a given atlas. Of all the structures in the head and neck region, the mandible is a structure well suited for model-based segmentation due to its high contrast against surrounding soft tissues. We thus designed a refinement step that applies a deformable surface model method to further improve the initial mandible segmentation as produced by atlas registration and structure mapping. At this refinement step, the atlas image is ignored and the initial mandible surface is deformed to align better with high image gradient points in the subject image that are likely to be boundary points of the mandible.

Mathematically, we denote the initial mandible boundary surface by a vector-valued parametric function $\mathbf{x}_0(r,s) = [x_0(r,s), y_0(r,s), z_0(r,s)]$, which gives the position of each surface point (indexed by (r,s)) in the subject image's coordinate system. We model the surface deformation using a surface deformation field $\mathbf{d}(r,s)$, such that $\mathbf{x}(r,s) = \mathbf{x}_0(r,s) + \mathbf{d}(r,s)$ gives the final location of each surface point. The surface refinement is then formulated as finding the optimal solution, \mathbf{d}_{opt} , of the following energy minimization problem problem:

$$\mathbf{d}_{\text{opt}} = \underset{\mathbf{d}}{\operatorname{argmin}} E(\mathbf{d}) = \underset{\mathbf{d}}{\operatorname{argmin}} \int \int \exp(-\|\nabla J(\mathbf{x}(r,s))\|^2) dr ds + \int \int \|\nabla \mathbf{d}(r,s)\|^2 dr ds. \quad (5)$$

In the above equation, the first term drives the surface to image locations with high intensity gradients and the second term enforces the regularity or smoothness of the surface displacement field. Regularizing the surface deformation instead of the surface itself allows capturing fine details of the object shape whereas at the same time enforcing the deformed surface to stay close to the initial shape. The solution is computed iteratively with a gradient-descent scheme. In general, only local refinement is necessary and this step usually takes less than ten seconds. Fig. 3 shows one example for the mandible refinement.

3 Results

To evaluate accuracy of the proposed method, we applied it to the 8 test datasets provided by the Princess Margaret Hospital for the MICCAI 2009 Head and Neck Auto-Segmentation Challenge workshop. The datasets contain non-contrast-enhanced CT scans with a typical resolution of $1 \times 1 \times 2 \text{ mm}^3$.

We used the 10 training datasets also distributed by the workshop as the atlases, for which expert manual segmentations of the mandible and the brainstem were made available. As described in Section 2, for each test dataset, 10 individual segmentation results were first computed using each of the 10 atlases, and the

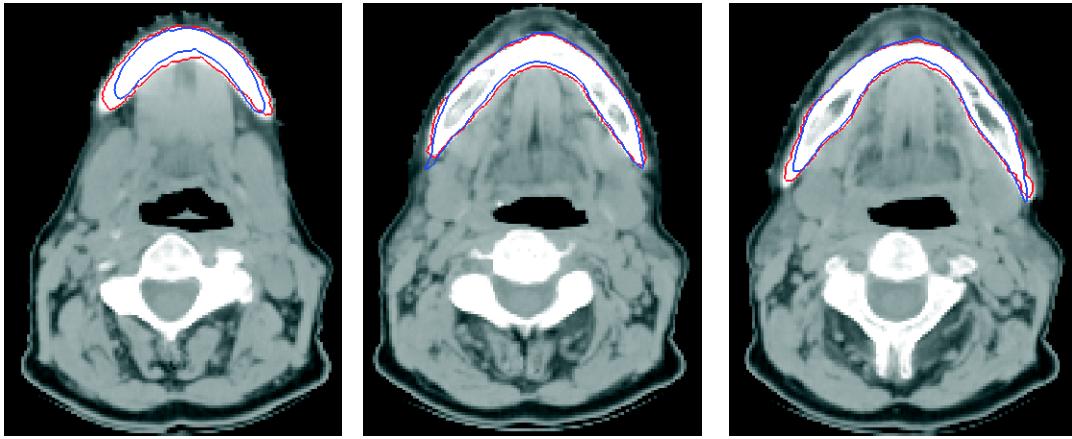


Figure 3: Illustration of deformable model based refinement of mandible segmentation. The images show three axial cross-sections for one subject. The blue curves are the initial results before refinement and the red curves are the final results.

results are then fused using the STAPLE algorithm to get the final auto-segmentation result. The fusion was performed independently for the mandible and the brainstem. The final auto-segmentation results for all test subjects were submitted to the organizers of the workshop and evaluated independently by the organizers against manual expert segmentation.

The experiments were carried out on a HP xw8400 desktop computer equipped with an Intel Xeon Quad-core 2.66 GHz CPU and a NVIDIA GTX 280 graphics card. The computation time was about 1 minute for each dataset if a single atlas is used and the total computation time was about 10 minutes when all 10 atlases were used.

Figs. 4 and 5 illustrate the segmentation results for the mandible and the brainstem respectively for one test dataset. The figures were generated by an independent reviewer (an organizer of the workshop). As we can see from the figures, the auto-segmentation results match the manual segmentation very well for both structures.

For quantitative evaluation, the auto- and manual- segmentations were compared in a slice-by-slice fashion for each dataset and the following quantitative measures were computed for each slice: the symmetric Hausdorff distance (HD) and the Dice similarity coefficient. A volumetric Dice similarity coefficient was also computed for each dataset to assess the overall volume overlap. Details of these evaluation criteria and their computation can be found in [16].

The overall statistics of the quantitative measures are summarized in Tables 1-2 for the mandible results and Tables 3-4 for the brainstem. From these tables, it can be seen that the median slice-wise Hausdorff distance is mostly less than 3 mm for the mandible and less than 3.6 mm for the brainstem. The median slice-wise Dice coefficient is mostly above 0.88 for the mandible and above 0.87 for the brainstem. The total volume overlap for most subjects is close to or above 0.9 for both the mandible and the brainstem. Thus, our method provided very accurate segmentation for both structures.

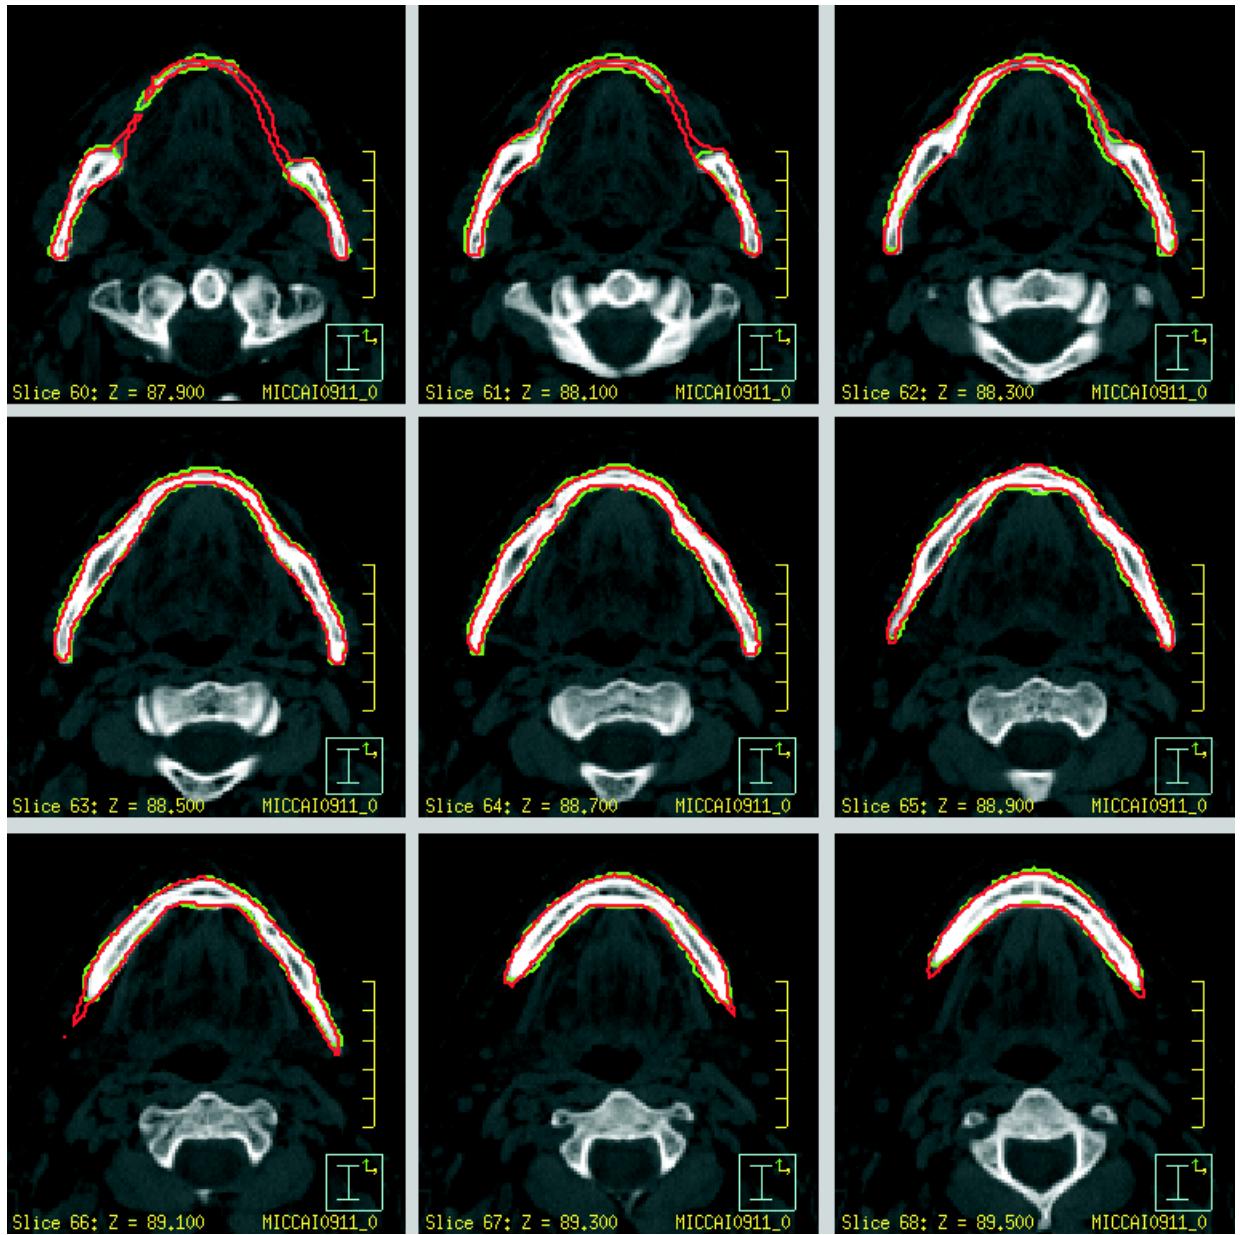


Figure 4: Snapshots of the mandible segmentation result for one test subject. Red curves indicate the auto-segmentation result and green curves are the corresponding manual segmentation.

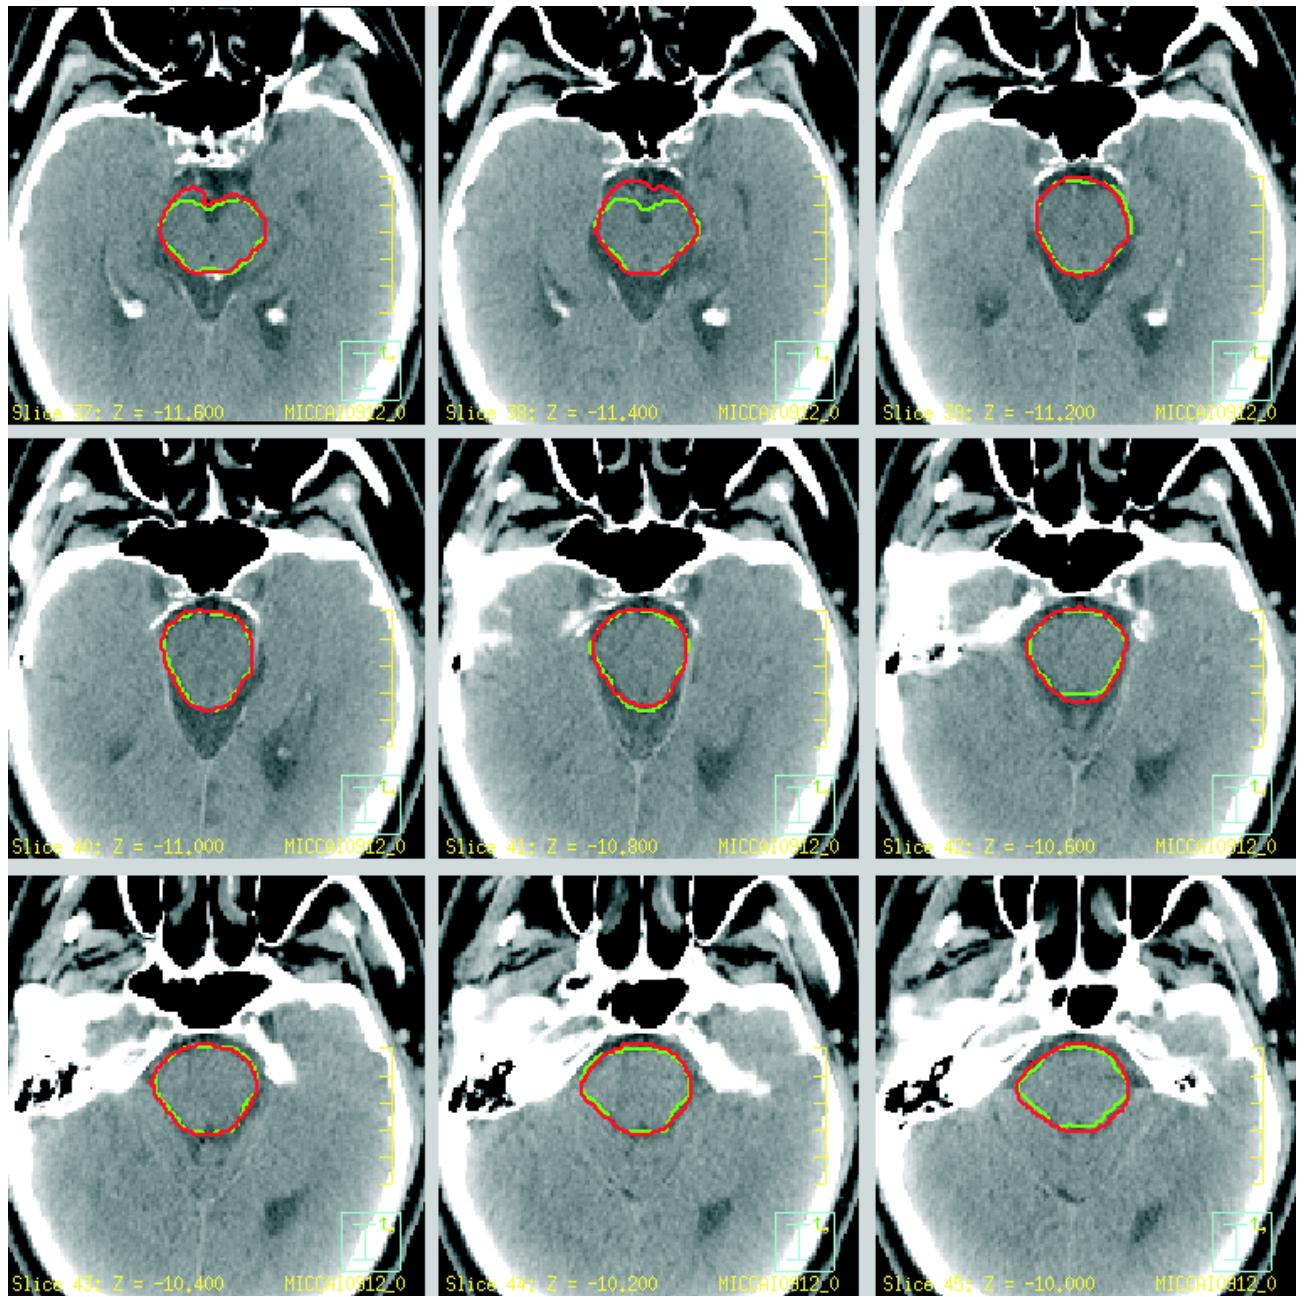


Figure 5: Snapshots of the brainstem segmentation result for one test subject. Red curves indicate the auto-segmentation result and green curves are the corresponding manual segmentation.

Dataset No.	Mean HD	Median HD	No. of slices (HD > 3 mm)
11	10.81	2.76	39 (18)
12	7.01	2.18	40 (14)
13	6.18	2.76	35 (14)
14	13.20	3.09	34 (18)
15	10.00	1.97	37 (12)
16	5.17	2.47	35 (12)
17	3.87	2.93	43 (20)
18	8.29	3.54	37 (21)

Table 1: Hausdorff distance (HD) statistics for mandible segmentation in the testing datasets.

Dataset No.	Average slice OV	Median slice OV	Total volume OV
11	83.5 %	89.3 %	88.3 %
12	87.4 %	89.8 %	90.6 %
13	88.4 %	92.3 %	91.5 %
14	81.4 %	86.3 %	89.4 %
15	87.2 %	91.4 %	91.3 %
16	88.1 %	88.8 %	91.9 %
17	88.9 %	90.4 %	91.4 %
18	84.1 %	88.7 %	88.2 %

Table 2: Overlap (OV) statistics for mandible segmentation in testing datasets.

Dataset No.	Mean HD	Median HD	No. of slices (HD > 3 mm)
11	3.73	3.52	28 (22)
12	3.31	2.93	29 (13)
13	-	-	-
14	3.67	3.52	30 (18)
15	2.48	2.49	18 (5)
16	3.05	2.93	27 (13)
17	3.71	2.93	27 (12)
18	3.90	3.34	29 (21)

Table 3: Hausdorff distance (HD) statistics for brainstem segmentation in the testing datasets.

The last column in Table 1 and Table 3 shows the number of slices with a Hausdorff distance greater than 3 mm for the mandible and the brainstem auto-segmentation results respectively. These numbers directly reflect the proportion of manual contouring time that can be saved when the proposed auto-segmentation method is applied, since a Hausdorff distance less than 3 mm is often considered as acceptable by clinicians [16]. In total, it was found that 171 out of 300 ($\sim 57\%$) mandible slices have a Hausdorff distance less than 3 mm and 84 out of 188 ($\sim 45\%$) brainstem slices are less than 3 mm. Although manual editing is considered necessary for the remaining slices, extra time-saving is expected since editing normally takes

Dataset No.	Average slice OV	Median slice OV	Total volume OV
11	86.3 %	86.8 %	87.3 %
12	90.7 %	90.7 %	92.1 %
13	-	-	-
14	87.0 %	87.4 %	89.6 %
15	90.8 %	90.9 %	91.0 %
16	89.1 %	90.6 %	90.9 %
17	89.2 %	89.6 %	89.7 %
18	85.7 %	85.6 %	85.3 %

Table 4: Overlap (OV) statistics for brainstem segmentation in testing datasets.

much less time than drawing contours from scratch.

4 Conclusion

In this work, we have developed an atlas-based method for fully automatic segmentation of H&N CT images that employs a hierarchical atlas registration and a multi-atlas fusion strategy. From the experimental results generated for the MICCAI H&N Auto-segmentation Challenge workshop, it was shown that the method gave very accurate segmentation of the mandible and the brainstem on real clinical data, as indicated by a mean volume Dice coefficient above 0.89 for both structures. We also showed that GPUs offer a cost-effective way to improve the computational performance and make multi-atlas segmentation feasible in practice. Future work will study the dependency of segmentation accuracy with the number of atlases used and investigate efficient and effective atlas selection methods.

References

- [1] Teng, C., Shapiro, L., Kalet, I.: Head and neck lymph node region delineation using a hybrid image registration method. In: Proc. ISBI 2006. (2006) 462–465 [1](#)
- [2] Isambert, A., Grégoire, V., Bidault, F., et al.: Atlas-based automatic segmentation (ABAS) of head and neck (H&N) structures in conformal radiotherapy: Atlas development and preliminary results in clinical context. In: ESTRO 2007. (2007) [1](#)
- [3] Zhang, T., Chi, Y., Meldolesi, E., Yan, D.: Automatic delineation of on-line head-and-neck computed tomography images: Towards on-line adaptive radiotherapy. *Int. J. Radiation Oncology Biol. Phys.* **68**(2) (2007) 522–530 [1](#)
- [4] Commowick, O., Grégoire, V., Malandain, G.: Atlas-based delineation of lymph node levels in head and neck CT images. *Radiother. Oncol.* (2008) [1](#), [2.1](#)
- [5] Han, X., Hoogeman, M., Levendag, P., Hibbard, L., Teguh, D., Voet, P., Cowen, A., Wolf, T.: Atlas-based auto-segmentation of head and neck ct images. In: MICCAI 2008, Part II. LNCS 5242, Springer-Verlag, Heidelberg (2008) 434–441 [1](#), [2.1](#), [2.3](#), [2.4](#)

- [6] Gorthi, S., Duay, V., Houhou, N., Cuadra, M., Schick, U., Becker, M., Allal, A., Thiran, J.P.: Segmentation of head and neck lymph node regions for radiotherapy planning, using active contour based atlas registration. *IEEE J. Selected Topics in Signal Proc.* **3** (2009) 135–147 [1](#)
- [7] Rohlfing, T., Brandt, R., Menzel, R., Russakoff, D., Maurer, Jr., C.: Quo vadis, atlas-based segmentation? In: *The Handbook of Medical Image Analysis*. Kluwer Academic / Plenum Publishers, New York, NY (2005) 435–486 [2.1](#)
- [8] Warfield, S., Zou, K., Wells, W.: Simultaneous truth and performance level estimation (STAPLE): An algorithm for the validation of image segmentation. *IEEE Trans. Med. Imag.* **23**(7) (2004) 903–921 [2.1](#)
- [9] Viola, P., Wells, W.: Alignment by maximization of mutual information. *Int. J. Comput. Vision* **24**(2) (1997) 137–154 [2.2](#), [2.4](#)
- [10] Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A fast and log-Euclidean polyaffine framework for locally rigid or affine registration. In: *Int. Workshop Biomed. Image Registration. LNCS 4057*, Springer, Heidelberg (2006) 120–7 [2.3](#), [2.3](#)
- [11] Suárez, E., Santana, J., Rovaris, E., Westin, C., Ruiz-Alzola, J.: Fast entropy-based nonrigid registration. In: *EUROCAST 2003. LNCS 2809*, Springer, Heidelberg (2003) 607–615 [2.3](#)
- [12] Stefanescu, R., Pennec, X., Ayache, N.: Grid powered nonlinear image registration with locally adaptive regularization. *Med. Image Anal.* **8** (2004) 325–342 [2.4](#)
- [13] Han, X., Hibbard, L., Willcut, V.: GPU-accelerated, gradient-free MI deformable registration for atlas-based MR brain image segmentation. In: *MMBIA 2009: IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*. (2009) [2.4](#), [2.4](#)
- [14] Rogelj, P., Kovačič, S., Gee, J.: Point similarity measures for non-rigid registration of multi-modal data. *Comput. Vis. Imag. Under.* **92** (2003) 112–140 [2.4](#)
- [15] NVIDIA: *NVIDIA CUDA Compute Unified Device Architecture Programming Guide: Version 2.0*. NVIDIA Corporation (2008) [2.4](#)
- [16] Pekar, V.: <http://www.grand-challenge2009.ca/Evaluation.pdf>. [3](#), [3](#)