A Steering Engine: Learning 3-D Anatomy Orientation Using Regression Forests

Fitsum A. Reda, Yiqiang Zhan, and Xiang Sean Zhou

Syngo Innovation, Siemens Healthcare, Malvern 19355, PA, USA
{fitsum.reda,yiqiang.zhan,xiang.zhou}@siemens.com

Abstract. Anatomical structures have intrinsic orientations along one or more directions, e.g., the tangent directions at points along a vessel central-line, the normal direction of an inter-vertebral disc or the base-to-apex direction of the heart. Although auto-detection of anatomy *orientation* is critical to various clinical applications, it is much less explored compared to its peer, "auto-detection of anatomy *location*". In this paper, we propose a novel and generic algorithm, named as "steering engine", to detect anatomy orientation in a robust and efficient way. Our work is distinguished by three main contributions. (1) Regression-based colatitude angle predictor: we use regression forests to model the highly non-linear mapping between appearance features and anatomy colatitude. (2) Iterative colatitude prediction scheme: we propose an algorithm that iteratively queries colatitude until longitude ambiguity is eliminated. (3) Rotation-invariant integral image: we design a spherical coordinates-based integral image from which Haar-like features of any orientation can be calculated efficiently. We validate our method on three diverse applications (different imaging modalities and organ systems), i.e., vertebral column tracing in CT, aorta tracing in CT and spinal cord tracing in MRI. In all applications (tested on a total of 400 scans), our method achieves a success rate above 90%. Experimental results suggest our method is fast, robust and accurate.

1 Introduction

"Where is the anatomical structure?", "What is the orientation of the anatomical structure?"- these are two fundamental questions frequently asked during medical image interpretation. Here, "anatomy orientation" denotes one or more intrinsic directions of an anatomical structure. These directions can be defined by local landmarks or global organ shapes, and are consistent across population. Some examples of anatomy orientations (see Fig. 1) include the tangent directions at points along a vessel central-line, the normal direction of an intervertebral disc or the base-to-apex direction of the heart. The ability to automatically detect anatomy orientation will not only benefit various clinical use cases, e.g., automatic MR scan planning [6], but will also pave a way to other high-level medical image analysis tasks. For example, predicting a local vessel orientation is a fundamental task for tracing algorithms [1], [2].

N. Navab et al. (Eds.): MICCAI 2015, Part III, LNCS 9351, pp. 612–619, 2015.
DOI: 10.1007/978-3-319-24574-4_73

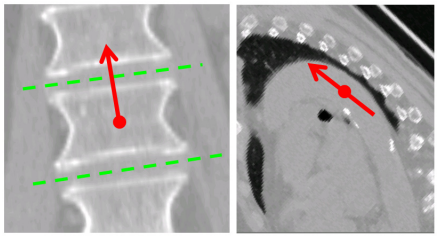

Fig. 1. Anatomy orientation (red arrow) of a vertebral column (left) and an aorta segment (right).

Early studies of anatomy orientation detection often exploit prior knowledge of specific anatomies. For example, template matching that relies on the Matched Filter Theorem or eigenanalysis of Hessian matrix has been widely used to determine local vessel orientation and perform tracing [1], [2]. These methods are dependent on strong priors of the anatomy under study, and do not generalize to other anatomies, such as the heart or the vertebral-column.

From another perspective, machine learning technologies have revolutionized the landscape of medical image analysis. In the areas of anatomy detection or segmentation, learning-based approaches have provided generic solution to different organs in different modalities [3]-[7]. However, learning technologies in detecting anatomy orientation remain under explored. One pioneer work is the marginal space learning (MSL) [8] which formulates the detection of anatomy orientation as a classification problem. This technique uses a classification model trained to detect a desired orientation. During testing, it uniformly samples the orientation space and performs an exhaustive test, which induces significant computational cost.

In contrast to classification, regression methods directly estimate the most probable hypothesis (anatomy orientation in our study) without exhaustively testing the hypothesis space. In fact, recent works have shown the power of regression-based approaches for medical image analysis, including landmark localization [4],[6], object detection [5],[7], and segmentation [3],[11]. Inspired by these works, we resort to regression paradigm and propose a method for efficient detection of anatomy orientation in a generic fashion (across modalities and organ systems), which is named as "steering engine". Our steering engine has three major novelties. (1) Regression-based colatitude angle predictor: we use regression forests to model the highly nonlinear mapping between appearance features and anatomy colatitude. (2) Iterative colatitude prediction scheme: we propose an algorithm that iteratively queries colatitude until longitude ambiguity is eliminated. (3) Rotation-invariant integral image: we design a spherical coordinates-based integral image from which Haar-like features of any orientation can be calculated very efficiently. With these three hallmarks, our method is able to detect anatomy orientations robustly and efficiently. In addition, thanks to the nature of machine learning techniques, our method is highly scalable to different organ systems and imaging modalities.

2 Methods

2.1 Problem Definition

Given a 3-D image patch $\wp$ (black cuboid in Fig 2a) containing an anatomy of interest (blue cuboid in Fig 2a), the goal of our steering engine is to determine the anatomy

orientation $\hat{\boldsymbol{v}}$ using appearance features. In this paper, we assume the anatomy is roughly located at the center of the patch. In practice, this can be achieved by existing anatomy detection algorithms or manual operators. The anatomy orientation $\hat{\boldsymbol{v}}$ is parameterized as (α, β), where $\alpha \in [0°, 90°]$ is the colatitude, which is the complement of the latitude or the angle measured from the positive z-axis (North Pole), and $\beta \in [0°, 360°)$ is the longitude, which is the counter-clockwise angle in the xy-plane measured from the positive x-axis (see Fig. 2b). This work does not discuss the prediction of the ϕ degree-of-freedom, which describes the anatomy's spinning around its own axis (see Fig 2b). In fact, after $\hat{\boldsymbol{v}}$ is determined, prediction of ϕ is a 2D orientation detection problem, which can also be solved by our steering engine. We plot the 3D orientation space (α, β) in a 2-D polar system, with α as the radius and β as the angle measured from the polar axis. As shown in Fig 2c, each black dot denotes a hypothesis in the orientation space. Fig 2d shows the corresponding 3D orientation space, in which the red dot denotes the North Pole.

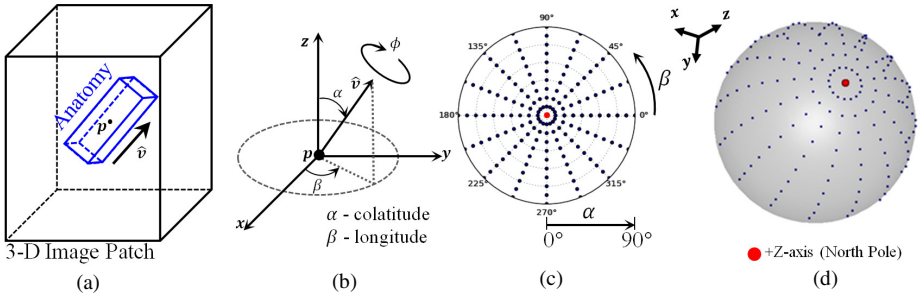

Fig. 2. Problem definition and orientation space.

2.2 Regression-Based Anatomy Colatitude Prediction

To achieve a generic solution, learning-based approaches are preferred. Recall one of the representative learning-based approaches, Marginal Space Learning (MSL) [8], in which orientation prediction is formulated as a classification problem as

$$(\alpha, \beta)_{\text{opt}} = \text{argmax}_{\alpha \in [0°, 90°], \beta \in [0°, 360°)} C(\wp_{\alpha,\beta}), \tag{1}$$

where $\wp_{\alpha,\beta}$ is the original patch $\wp$ rotated by (α, β), and $C(\cdot)$ is a pre-trained classifier that maps $\wp_{\alpha,\beta}$ to a score which indicates if the anatomy orientation within the patch is along the desired direction. As shown in Eq. 1, at run-time, each (α, β) hypothesis needs to be evaluated by $C(\cdot)$ and the hypothesis that maximizes the score is the predicted anatomy orientation. Assuming the orientation space is quantized into a 1° resolution, Eq. 1 needs to evaluate 32,400 (360×90) hypotheses, which induces high computational cost.

To achieve efficient orientation detection, we resort to regression paradigm. Mathematically, the orientation prediction is formulated as

$$(\alpha, \beta)_{\text{opt}} = \boldsymbol{R}(\wp). \tag{2}$$

At run-time, the pre-trained regressor $\boldsymbol{R}(\cdot)$ evaluates $\wp$ only *once* to estimate orientation. Compared to classification paradigm, which uses exhaustive test, the regression-based method has immense potential for run-time efficiency. The remaining challenge is how to learn the highly non-linear function that maps visual features of $\wp$ to anatomy orientation (α, β). In this study, we select random forests (RF) [10] as the regressor. Since RF comprises a set of non-linear decision trees with feature selection capability, it is able to model the highly non-linear mapping between visual features (observable variables) and anatomy orientation (latent variables). In addition, we simplify the regressand from the orientation vector (α, β) to the colatitude variable α. This strategy is designed for two reasons. First, a single-regressand mapping function is usually less complex to learn than a multi-regressand mapping function. Second, since the longitude β is a circular variable, the longitude statistics required by standard RF training cannot be calculated in a trivial way.

The colatitude predictor can then be described as

$$\alpha = \boldsymbol{RF}(\boldsymbol{\mathcal{F}}(\wp)), \tag{3}$$

where $\boldsymbol{RF}$ denotes a trained regression forest, and $\boldsymbol{\mathcal{F}}(\wp)$ denotes the visual features of $\wp$ [6], [12]. In this study, we use Haar-like features which comprise one or multiple cuboids (refer to [6] for more details). These features are automatically selected in the regression forest training stage. We train the regression forest as suggested in [10].

2.3 Iterative Colatitude Prediction and Patch Reorientation

Since Eq. 3 predicts only anatomy colatitude, it still cannot provide complete 3D anatomy orientation. To address this problem, we design an iterative algorithm exploiting a unique property of the 3D orientation space that all longitudes converge at the North Pole. In other words, as shown in Fig 2c, the smaller the colatitude the less the ambiguity of 3-D orientations with respect to different longitudes. For example, while orientations on the equator (colatitude = 90°) vary a lot with different longitudes, the orientation at the North Pole (colatitude = 0°) is not affected by longitude at all. Accordingly, our algorithm iteratively performs colatitude prediction and patch reorientation until the predicted colatitude from the reoriented patch is close to zero, i.e., until the orientation ambiguity from the unknown longitude is eliminated.

A schematic explanation of our algorithm is shown in Fig. 3. We denote a patch $\wp$ rotated by (α, β) as $\wp_{\alpha,\beta}$, and let $\wp_{\alpha_0,\beta_0\,(\alpha_0 = 0°, \beta_0 = 0°)}$ is the original patch. The green dot in Fig 3a denotes the ground truth of the anatomy orientation in $\wp_{\alpha_0,\beta_0}$. Our algorithm first predicts the anatomy colatitude $\alpha_1 = \boldsymbol{RF}\left(\boldsymbol{\mathcal{F}}\left(\wp_{\alpha_0,\beta_0}\right)\right)$, which narrows down the orientation hypothesis to the yellow circle in Fig 3a. Due to possible prediction errors, the narrowed hypothesis may not include the ground truth (whose colatitude is $\hat{\alpha}_0$). To further reduce the size of the hypothesis space along the longitude, we sample M uniformly spread hypothesis along the longitude (blue points on the yellow circle in Fig. 3a). Intuitively, the "best" hypothesis (the dark-blue point in Fig 3a) should be the one that is closest to the ground truth (green point), i.e., the hypothesis

with the minimum *relative* colatitude to the ground truth. Mathematically the optimal longitude β_1 is determined by,

$$\beta_{i+1} = \arg\min_{\{\beta_{ij}|j=1,\dots,M\}} \boldsymbol{RF}\left(\boldsymbol{\mathcal{F}}\left(\wp_{\alpha_{i+1},\beta_{ij}}\right)\right). \tag{4}$$

Then, $\wp_{\alpha_0,\beta_0}$(Fig. 3b) is re-oriented to $\wp_{\alpha_1,\beta_1}$ (Fig. 3d), and its orientation space is re-centered at α_1, β_1 (dark-blue point in Fig. 3a moved to orientation space center in Fig. 3c). Comparing Fig. 3b and 3d, the colatitude $\hat{\alpha}_0$ in $\wp_{\alpha_0,\beta_0}$has decreased to $\hat{\alpha}_1$ in $\wp_{\alpha_1,\beta_1}$. Hence, the longitude ambiguity (yellow circles in Fig. 3a and 3c) has greatly reduced. The colatitude prediction and patch reorientation steps are then iterated until the colatitude is less than a small threshold.

Although Eq. 4 looks similar to the classification-based approach (Eq. 1), it essentially has much higher run-time efficiency due to two reasons. First, the colatitude predictor shrinks the full hypothesis space to a circle at each iteration (yellow band in Fig 3a). Second, along with the iterations, even the sizes of the shrunk hypothesis reduce dramatically (yellow circle shrinks from Fig. 3a to Fig. 3c).

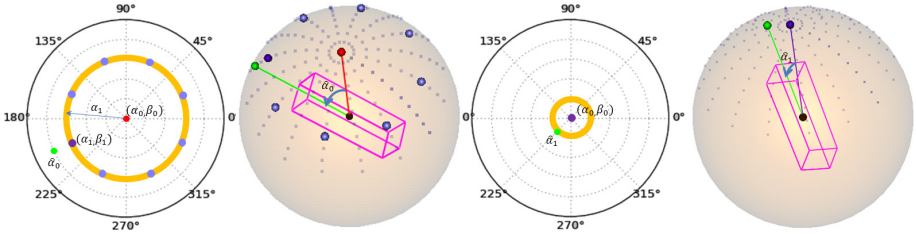

Fig. 3. Illustration of iterative colatitude prediction and patch reorientation.

2.4 Rotation-Invariant Integral Image for Fast Feature Calculation

The colatitude predictor uses Haar-like features, which can be efficiently computed using the integral image [9]. However, during the iterative colatitude prediction and patch re-orientation, the integral image needs to be recalculated, since it is not rotationally invariant. To speed-up this computationally expensive task, we propose a rotation-invariant integral image built on the spherical coordinate system. We illustrate our idea in 2D first. Given a 2-D image I (Fig. 4a), its integral image II (Fig. 4c) can be built on the polar-coordinate system as,

$$II(r,\theta) = \sum_{r'<r,\theta'<\theta} I(r'\cos\theta', r'\sin\theta'). \tag{5}$$

$II(r,\theta)$ is equal to the sum of the intensities within the red sector in Fig 4a. The basic element of Haar-like features, i.e., the sum of the intensity within a circular trapezoid (Fig. 4b) (approximation of a rectangle), can be quickly calculated using four entries in the integral image as,

$$A = II(r_2, \theta_2) + II(r_1, \theta_1) - II(r_2, \theta_1) - II(r_1, \theta_2) \ . \quad (6)$$

To calculate the same Haar element on a rotated image (rotated by $\Delta\theta$), instead of re-computing the integral image, we only change the coordinates of the four entries in the integral image as:

$$A' = II(r_2, \theta_2 + \Delta\theta) + II(r_1, \theta_1 + \Delta\theta) - II(r_2, \theta_1 + \Delta\theta) - II(r, \theta_2 + \Delta\theta) \quad (7)$$

Eq. 6 and Eq. 7 show the "rotation-invariant" property of the polar coordinate system-based integral image. The same integral image can be reused to calculate Haar-like features at any orientation. The same idea can be extended to 3-D using the spherical coordinate system.

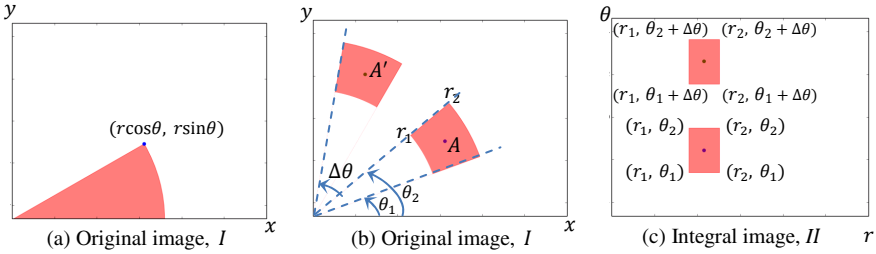

Fig. 4. Illustration of the rotation-invariant integral image in 2D.

3 Experimental Results

We evaluate our method on different anatomical structures, organ-systems and imaging modalities. The datasets have large variability due to different scanners (GE, Siemens and Philips), exposure time, field of view, or contrast level. Voxel resolutions range from 0.5 to 5 mm. To validate the generalization ability of our method, we use the same RF training parameters: number of trees (100), minimum sample size to split (20), maximum node size (10), number of features to examine per node (100).

Local Anatomy Orientation Detection: This experiment aims to evaluate the performance of our steering engine on local orientation detection on two anatomical structures, vertebrae and aorta segments (see Fig 1 for orientation definition). The input of the algorithm is a fixed-size volume of interest (VOI) containing the anatomical structure, which is roughly located at the center of the VOI. We compare the steering engine with (a) classification-based approach (as described in Eq. 1) and (b) Steering engine w/o RIII: steering engine without using the rotation-invariant integral image (RIII).

Table I shows the results on 280 vertebrae VOIs (the steering engine is trained from another 100 CTs). We achieved <4° average error, <8° error in ~95% VOIs with run-time speed of ~1.33 seconds (a single thread implementation on an Intel Xeon 2.3GHz CPU). The classification-based approach achieved inferior accuracy/robustness and is ~65 times slower. Without using RIII, the accuracy is similar but the speed is ~3 times slower. We conduct similar comparison on 220 aorta VOIs

(trained by another 50 CTs). Despite the lack of contrast in 50% of the VOIs, our steering engine is still accurate, robust and fast. The comparisons across the detection methods show similar pattern as the vertebrae use case.

Table 1. Quantitative comparison of orientation detection methods

Vertebrae	Accuracy (°)	Robustness (%)				Speed
		<1°	<2°	<4°	<8°	(seconds)
Classification	4.13±2.01	0.00%	11.36%	49.54%	97.27%	86.13±13.81
Steering engine w/o RIII	3.75±1.88	**5.00%**	19.09%	80.90%	97.73%	4.47±1.09
Steering engine	**3.74±2.17**	4.54%	**22.72%**	**85.91%**	95.90%	**1.33±0.42**
Aorta	Accuracy (°)	Robustness (%)				Speed
		<1°	<2°	<4°	<8°	(seconds)
Classification	5.09±2.45	0.19%	9.26%	39.67%	83.01%	29.29±5.25
Steering engine w/o RIII	**4.86±2.45**	2.41%	9.92%	33.97%	**84.98%**	0.88±0.23
Steering engine	4.97±2.75	**4.42%**	**22.12%**	**39.82%**	83.18%	**0.32±0.14**

Anatomical Structure Tracing: We further evaluate the steering engine on more clinically relevant applications - tracing anatomical structures. Given a starting location, we (a) identify anatomy orientation using the steering engine, (b) step to a new location along the orientation, and (c) iterate (a) and (b) until a maximum number of steps or a stopping location is reached. We tested it on tracing (a) the aorta in 100 CTs (with varying level of contrast, and some with no contrast), (b) the vertebral column in 100 CTs, and (c) the spinal cord in 200 DIXON-sequence MRIs, and achieved a success rate (based on evaluation of experienced professionals) of 95%, 96% and 90%, respectively. Fig 5 shows representative example tracing results. Tracing was successful despite the large anatomical shape variation in a scoliosis subject (Fig 5a), and a kyphosis subject (Fig 5b), the lack of contrast in the CT (Fig 5d), the noisy reconstruction CT kernel (Fig 5c) and large shape and intensity variation surrounding the spinal cord (Fig 5e and 5f).

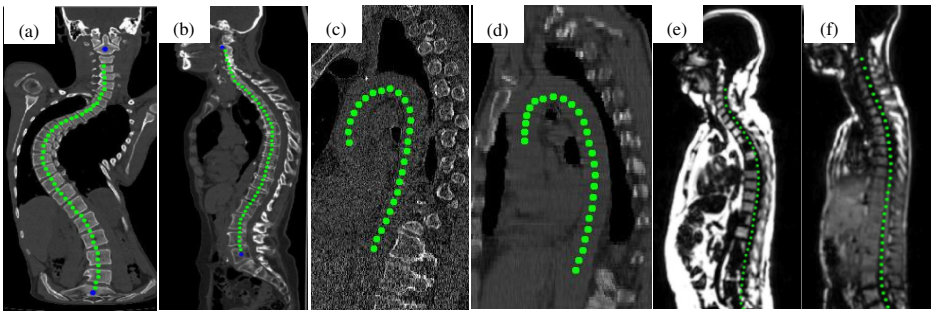

Fig. 5. Representative results of anatomy tracing in different use cases.

4 Conclusions

We presented a novel and generic steering engine to automatically detect anatomy orientation in a robust and efficient fashion. We employ regression forests to learn the highly non-linear relationship between anatomy visual features and anatomy colatitude. By exploiting a unique property of the 3D orientation space, we design an iterative algorithm to determine anatomy orientation by fully relying on colatitude prediction. In addition, we design a rotation-invariant integral image to efficiently calculate Haar-like features at any orientations. We tested our steering engine on both local anatomy orientation detection and anatomy tracing. In all three applications, including different imaging modalities (CT and MRI Dixon) and anatomies (vertebral column, spinal cord and aorta), experimental results show that our method is fast, robust and accurate.

References

1. Friman, O., Hindennach, M., Kühnel, C., Peitgen, H.: Multiple hypothesis template tracking of small 3D vessel structures. MedIA, 14(2010), 160–171 (2010)
2. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A.C.F., Delp, S.L. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
3. Lindner, C., Thiagarajah, S., Wilkinson, J.M., Consortium, T.A., Wallis, G.A., Cootes, T.F.: Fully Automatic Segmentation of the Proximal Femur Using Random Forest Regression Voting. IEEE Trans. Med. Imag. 32(8), 181–189 (2013)
4. Glocker, B., Feulner, J., Criminisi, A., Haynor, D.R., Konukoglu, E.: Automatic localization and identification of vertebrae in arbitrary field-of-view CT scans. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part III. LNCS, vol. 7512, pp. 590–598. Springer, Heidelberg (2012)
5. Yaqub, M., Kopuri, A., Rueda, S., Sullivan, P.B., McCormick, K., Noble, J.A.: A Constrained Regression Forests Solution to 3D Fetal Ultrasound Plane Localization for Longitudinal Analysis of Brain Growth and Maturation. In: Wu, G., Zhang, D., Zhou, L. (eds.) MLMI 2014. LNCS, vol. 8679, pp. 109–116. Springer, Heidelberg (2014)
6. Zhan, Y., Dewan, M., Harder, M., Krishnan, A., Zhou, X.S.: Robust automatic knee mr slice positioning through redundant and hierarchical anatomy detection. IEEE Trans. Med. Imag. 30(12), 2087–2100 (2011)
7. Criminisi, A., Robertson, D., Konukoglu, E., Shotton, J., Pathak, S., White, S., Siddiqui, K.: Regression Forests for Efficient Anatomy Detection and Localization in Computed Tomography Scans. MedIA 17(8), 1293–1303 (2013)
8. Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D.: Four-Chamber Heart Modeling and Automatic Segmentation for 3D Cardiac CT Volumes using Marginal Space Learning and Steerable Features. IEEE Trans. Med. Imag. 27(11), 1668–1681 (2008)
9. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: Proc. IEEE CVPR, pp. 1-511 (2001)
10. Breiman, L.: Random forests. Machine Learning 45(1), 5–32 (2001)
11. Gao, Y., Shen, D.: Learning distance transform for boundary detection and deformable segmentation in ct prostate images. In: Proc. MLMI, pp. 93–100 (2014)
12. Skibbe, H., Reisert, M., Schmidt, T., Brox, T., Ronneberger, O., Burkhardt, H.: Fast Rotation Invariant 3D Feature Computation Utilizing Efficient Local Neighborhood Operators. IEEE PAMI 34(8), 1563–1575 (2012)