



NOVA University of Newcastle Research Online

nova.newcastle.edu.au

Chandra, Shekhar S.; Dowling, Jason A.; Greer, Peter B.; Martin, Jarad; Wratten, Chris; Pichler, Peter; Fripp, Jurgen; Crozier, Stuart "Fast automated segmentation of multiple objects via spatially weighted shape learning". Published in Physics in Medicine and Biology Vol. 61, Issue 22, p. 8070-8084 (2016)

Available from: <http://dx.doi.org/10.1088/0031-9155/61/22/8070>

Accessed from: <http://hdl.handle.net/1959.13/1345358>

Fast Automated Segmentation of Multiple Objects via Spatially Weighted Shape Learning

Shekhar S. Chandra¹, Jason A. Dowling², Peter B. Greer^{3,4}, Jarad Martin^{3,4}, Chris Wratten^{3,4}, Peter Pichler⁴, Jurgen Fripp², Stuart Crozier¹

E-mail: shekhar.chandra@uq.edu.au

¹ School of Information Technology and Electrical Engineering, The University of Queensland, Australia.

² Australian e-Health Research Centre, CSIRO, Australia.

³ Calvary Mater Newcastle Hospital, Australia.

⁴ University of Newcastle, Australia.

Abstract. Active Shape Models (ASMs) have proved successful in automatic segmentation by using shape and appearance priors in a number of areas such as prostate segmentation, where accurate contouring is important in treatment planning for prostate cancer. The ASM approach however, is heavily reliant on a good initialisation for achieving high segmentation quality. This initialisation often requires algorithms with high computational complexity, such as three dimensional (3D) image registration. In this work, we present a fast, self-initialised ASM approach that simultaneously fits multiple objects hierarchically controlled by spatially weighted shape learning. Prominent objects are targeted initially and spatial weights are progressively adjusted so that the next (more difficult, less visible) object is simultaneously initialised using a series of weighted shape models. The scheme was validated and compared to a multi-atlas approach on 3D Magnetic Resonance (MR) images of 38 cancer patients and had the same (mean, median, inter-rater) Dice's similarity coefficients of (0.80, 0.81, 0.85), while having no registration error and a computational time of 12-15 minutes, nearly an order of magnitude faster than the multi-atlas approach.

Keywords: Active shape models, MRI, Prostate, Segmentation

1. Introduction

Automatic segmentation of three dimensional (3D) Magnetic Resonance (MR) images has become important because of its increasing use in treatment planning of prostate cancer due to its dramatically improved soft-tissue contrast (Karlsson et al. 2009; Nyholm and Jonsson 2014). Manual contouring, especially from such high detail MR images, are time consuming and prone to inter-rater variability. In our hospital, contouring these organs from MR images requires approximately 3 hours per patient (most of this being the MR bone contouring).

Recently, Active Shape Models (ASMs) (Cootes et al. 1995) have proved successful in automatically segmenting 3D MR images for a number of anatomical structures in humans including the (individual) bones (Fripp et al. 2010; Schmid, Kim and Magnenat-Thalmann 2011; Xia et al. 2013; Chandra et al. 2014), muscles (Engstrom et al. 2011; HajGhanbari et al. 2011), prostate (Martin, Troccaz and Daanen 2010; Chandra et al. 2012), liver (Heimann, Wolf and Meinzer 2006) and the heart (Zheng et al. 2008; Ecabert et al. 2008). They have even been utilised in the shape analysis of structures such as the hippocampus of the brain in

Alzheimer’s disease (Shen et al. 2012) and inter-vertebral discs with degenerative-type changes of the lumbar spine (Neubert et al. 2013).

An ASM consists of a triangulated surface S (having vertices v_i with $0 \leq i < N$, where N is the total number of vertices), a statistical shape model that encapsulates the object shape information and a gray-level model to represent the object (gray-level intensity) appearance in proximity to the object as one dimensional (1D) profiles per vertex of the surface. These models are pre-computed offline from a set of training images that typically have the objects of interest manually contoured. During the testing phase, the surface is then deformed to match local object appearance similar to the gray-level model while only allowing the global shape closest to that of the learned shape subspace. The reader is directed to a number of works that describe this process in greater detail (Martin, Troccaz and Daanen 2010; Schmid, Kim and Magnenat-Thalmann 2011; Chandra et al. 2012).

However, ASMs require a good initialisation to improve convergence and avoid being caught in local minima. For example in prostate segmentation work, Martin, Troccaz and Daanen (2010) utilised an ASM approach with spatial constraints that was initialised by a single atlas, i.e. where a template image or atlas with a known segmentation and surface is propagated to the test case using 3D image registration. After the success of using multiple atlases in segmenting prostate MR images on their own (Klein et al. 2008), Chandra et al. (2012) utilised multiple atlases to initialise their ASM and improved ASM convergence using a selection of atlases with most similar appearance for the gray level profile template matching. These initialisation methods, although very accurate, are themselves of high computational complexity, requiring rigid and non-rigid registrations of a large number of detailed 3D images, taking up a significant portion of the total computation time. A thorough review of prostate segmentation work can be found in Ghose et al. (2012) or more recently in Litjens et al. (2014).

Limited work has been performed in self-initialising deformable models or multi-object segmentation in MR imaging due to the challenges of the high dimensionality of the image data. Most methods involve some pre-segmentation of the image either by landmarks (Smith, Najarian and Ward 2009) or pixel classifications (Cosío 2008). Multi-object segmentation methods have been developed for computerised tomography (CT) images (Yokota et al. 2009; Glocker et al. 2012) or photography (Shotton et al. 2007). Although methods exist for MR images, these tend to be either semi-automatic (Gao et al. 2012) or ideally suited to terrain-like objects unless shape models are used initially (while having similar run times to multi-atlas approaches) (Song et al. 2013). A promising approach are hierarchical deformable models, where the segmentation is done at multiple stages, with each stage initialising the next. The initialisation for most of these methods still utilise some form of registration (Okada et al. 2008) or sub-sampling/grouping (Shang, Su and Dössel 2006; Ma et al. 2010).

In this work, we present a self-initialised multi-object ASM approach that segments the target object in a hierarchical manner using weighted shape learning. This shape weighted hierarchical approach progressively initialises the subsequent (more difficult to segment) object by segmenting a more discernible object first, such as larger objects that possess clearer and more contrasting appearance like the bone, and exploiting any connectedness between these objects to segment subsequent objects. The approach is free of image registrations and is of low computational complexity compared to the 3D images involved. as it optimises a two dimensional (2D) surface embedded into 3D space using the image volume sampled along 1D profiles. In this way, the proposed method propagates towards the target object without needing any other forms of initialisation. This hierarchical targeting is achieved by adjusting the global accuracy of the individual object segmentations via (continuous) spatial weights in the shape model, allowing particular objects to be segmented accurately, while the next (more difficult, less clearly visible) object is simultaneously initialised for the next step. When these

steps are pipelined together, the segmentation approach achieves a robust initialisation and fast segmentation convergence of the ultimate object of interest.

The proposed method is validated by segmenting the pelvic structures (body outline, bone, bladder, rectum and prostate) with the eventual goal to segment the prostate accurately in 3D MR images. The proposed method is shown to be nearly an order of magnitude faster, while having similar accuracy, than the multi-atlas approach. In the next section, we present the proposed multi-object approach and subsequent sections outline the results and discuss these results.

2. Method

The objective of our work is to use a single combined deformable model, also referred to as a ASM (Cootes et al. 1995), for multi-object segmentation with spatially weighted shape learning. A series of weighted shape models are used to hierarchically activate/deactivate objects, as well as partially fit/initialise and adjust their corresponding segmentation accuracies via the spatially weighted Principal Component Analysis (PCA) (Thomaz et al. 2010). This weighted hierarchical model encapsulates the inter-shape relationships in between the multiple objects within the dataset. Each of the above aspects, as well as the segmentation pipeline will be described in subsections that follow.

2.1. Training Surfaces

To segment each newly acquired case using deformable models, we require a shape model trained offline from manual segmentations. For each of the M objects to be segmented in new scans, we assume that there are s training cases (each having the M objects in the same patient space) manually segmented by an expert to be used for training our algorithm. For example in our case, 38 patients (so that $s = 38$) receiving external beam radiation therapy (EBRT) at the Calvary Mater Newcastle Hospital, Australia had their body outline, bone, bladder, rectum and prostate (so that $M = 5$) manually segmented from their T2 3D MR images (Dowling et al. 2015). The set of s training surfaces per object can then be created from these manual segmentations by any one of the well established shape correspondence methods such as the shape-image technique (Davies, Twining and Taylor 2008), a surface registration method (Hufnagel et al. 2008), particle-based systems (Datar et al. 2011) or minimum descriptor lengths (Davies et al. 2010).

In this work, we use the correspondence approach also presented by (Chandra et al. 2014) (appendix in their work) to obtain the training surfaces. To do this, a template surface for each object was initially obtained via marching cubes (Lorensen and Cline 1987) from a single case. It was then smoothed (Taubin, Zhang and Golub 1996) and decimated (Garland and Heckbert 1997) to reduce the number of vertices and hence the computational complexity. To better account for the field of view (FOV) differences (of the superior and inferior areas) in the body outline and bone, an average atlas was created from the labels instead and a surface created in a similar manner to the other organs from the labelled atlas for these two objects. Figure 1 shows the resulting surfaces, their FOV and sampling.

The template surfaces are then fit to each of the s manual segmentation labelled images using the surface to image correspondence approach to obtain training surfaces for each object. This results in a series of s surfaces for each object that is essentially a non-rigidly deformed template surface and thus has point-to-point correspondence between surfaces of the same object. The point-to-point correspondence is important to ensure similar vertices across the population of s cases represent the same part of the anatomy of each object. The template

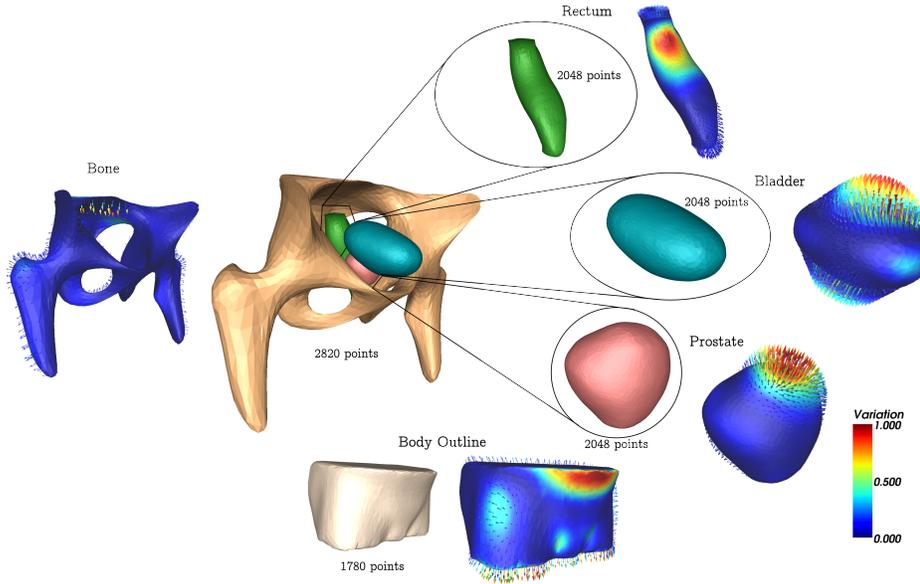


Figure 1. The template surfaces, their sampling and resulting PCA-based shape modelling (showing the primary mode of variation) obtained from the shape creation pipeline by Chandra et al. 2014 for each object in this study. Body outline is kept separate to allow easy visualisation of the other (internal) organs.

surface that was selected initially to generate the model was then replaced by the resulting mean shape from these surfaces and the surface to image registration approach re-run to obtain the final training surfaces. The result is s cases, their MR image with M object surfaces accurately represented within them ready for shape modelling.

2.2. Spatially Weighted Learning

Once each of the s cases have all object surfaces accurately represented, the surfaces representing each object (and the template surfaces) are combined into a single surface for each case. Figure 1 shows the template surfaces, point sampling and standard PCA-based shape models (described shortly) used in this work. The next step is to weight each object depending on the segmentation stage using the spatially weighted PCA.

The PCA (Pearson 1901) provides an eigen-decomposition of a rectangular ($3N \times s$) data matrix $\mathbf{X} = \mathbf{x}_j$, where $0 \leq j < s$ with \mathbf{x} being the (3D) position vector of each surface vertex. The decomposition is usually computed via the singular value decomposition (SVD) of the covariance matrix $\mathbf{C} = (1/s) \mathbf{X}'^T \mathbf{X}'$, where $\boldsymbol{\mu}$ is the mean shape and $\mathbf{X}' \leftarrow \mathbf{X} - \boldsymbol{\mu}$, i.e. each column of \mathbf{X}' consists of the variation of the points from the mean. The eigen-decomposition produces a dimensionality reduction in the form of a set of eigenvectors $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N] = \mathbf{u}_i$ and eigenvalues λ_i that effectively represent the major modes, which model the global shape variation of the surfaces.

The spatial weighted PCA effectively provides regions of interest (ROI) shape modelling of the surfaces by essentially pre-multiplying the rows of \mathbf{X}' , i.e. the positional variations, with weights

$$\mathbf{w}_i = [\sqrt{w_1}, \sqrt{w_2}, \dots, \sqrt{w_N}], \quad (1)$$

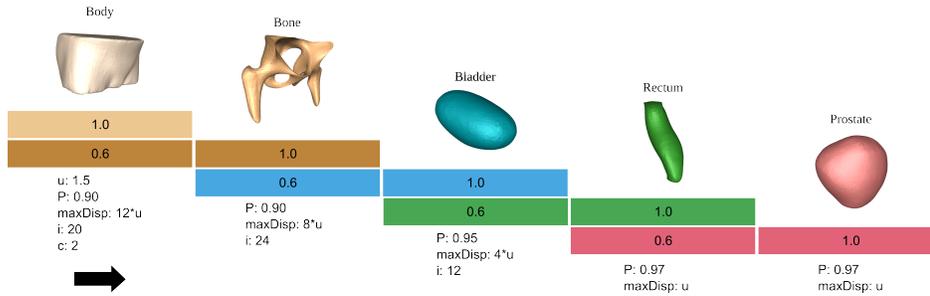


Figure 2. The segmentation weighting and parameters used in this work, where profile length $x = 12$, u is the profile spacing, P is the shape model precision, i is the number of iterations and c is the capture range of the deformable model. Each stage and its corresponding weights are shown as a unique colour (see online version for colours). The overlapping spatial weights at 1.0 and 0.6 (see matching colour bars), which effectively a weighted version of the models in figure 1. This simultaneously fits the current object accurately in the pipeline (with a weight of 1.0), while initialising the next object (with a weight of 0.6). The other objects (shown with absent colour bars) are “disabled” by using a value close to zero (such as 1×10^{-6}). This shortens the overall pipeline than it would be with other unweighted (hierarchical) methods. The weights are implemented as spatial weights using the spatially weighted PCA (Thomaz et al. 2010). Note that the objects in the figure are not to scale.

directly before the eigen-decomposition. This manipulates the variations from the mean shape, thus also the global shape representation energy of the model, to those that are highly weighted. The points for which the shape variations are down-weighted have their variations revert to the mean shape as the weights approach zero and are effectively ignored. By setting these weights, either manually or automatically via an algorithm, we can focus the accuracy of the shape model to certain spatial areas deemed more important to the problem. To construct the multi-object approach, we must engineer a series of pipelined weights for segmenting and initialising objects.

2.3. Segmentation Scheme

The main purpose of the proposed segmentation scheme is to use spatial weights in order to adjust the segmentation accuracy of an ASM approach to facilitate a weighted hierarchical scheme for self-initialised multi-object segmentation.

Assuming that there are spatial relationships between the position of the multiple objects in question, a series of weights are produced to partially fit more difficult objects, while segmenting more easily discernible objects first. In our case, the body outline is the most easily distinguishable object in MR images, and since the bones are present within the body with relative positions that can be learned using training data, we can fit the body outline and partially fit the bones. A similar argument is true for the bladder with respect to the bones. Once we have accurately segmented the body and suitably initialised the bone via a simultaneous partial fit, we can then fit the bones and partially fit the bladder. This process is repeated for the bladder and rectum and finally the rectum and prostate.

The choice of weights for the hierarchical pipeline can be found empirically or by automatic tuning processes, however a weight of 1.0 for fully fitting an object and 0.5-0.75 for partially fitting the subsequent object seems to suffice in the cases explored by the authors. Incorporating this weighting type into the segmentation scheme, one gets a pipeline that looks like figure 2. The overlapping of the spatial weights ensures that the target object, here a prostate, is reached with as few iterations as possible, decreasing the overall processing time.

After the weights have been decided, the spatially weighted PCA is used to create a weighted shape model for each stage (the blocks or steps in figure 2). The remaining aspects of the segmentation method hierarchically utilises spatial weights as highlighted above resulting in a computationally efficient segmentation pipeline using template matching along 1D profiles.

2.4. Template Matching

Deformable models when applied to imaging data involving organs such as the prostate traditionally utilise a gray level modelling approach similar to that of Martin, Troccaz and Daanen 2010 or Chandra et al. 2012 due to a lack of strong edges. The gray level model essentially consists of 1D profiles of x elements extracted along each vertex of the s training surfaces (that accurately represent the manually segmented cases) thereby approximately encapsulating the surrounding appearance of the organ. Thus, each of the N vertices of the deformable model have s training profiles \mathbf{p}_j .

During the fit procedure of the deformable model, new profiles \mathbf{q}_i of length cx are extracted from the case to be segmented and template matched to the training profiles \mathbf{p}_j for each vertex, where $c \geq 1$ and is effectively a capture range. The template matching is done using normalised cross-correlation (NCC) and the position of the vertices of the surface adjusted to where there is greatest correlation γ defined as

$$\gamma = \frac{\mathbf{p}_j \cdot \mathbf{q}_i}{\|\mathbf{p}_j\| \|\mathbf{q}_i\|}, \quad (2)$$

where $0 \leq \gamma \leq 1$, so that negative correlations are ignored, $\|\cdot\|$ is the magnitude and \cdot represents the dot product. Once the vertices are deformed, the surface is smoothed (Taubin, Zhang and Golub 1996) to account for neighbourhood deformations and the shape model applied. The remaining parameters for the proposed weighted deformable model approach is shown in figure 2.

To improve the computational performance of the template matching, we detect disabled objects as those vertices with near zero weights via a threshold to ignore these points in the matching. Having established all components of a multi-object weighted deformable model, it was applied to a cohort of male subjects receiving treatment for prostate cancer.

2.5. Imaging Data

The imaging data utilised for this work was originally used for substitute CT generation for MR imaging (alone) based EBRT planning (Dowling et al. 2015). A total of 38 patients (aged between 58-78 years) all diagnosed with tumours had two MR sequences utilised in this work acquired with a Siemens Skyra 3T scanner and positioned by two radiation therapists. The planning MR was a 3D, T2 weighted 1.6 mm isotropic Sampling Perfection with Application optimized Contrasts using different flip angle Evolution (SPACE) sequence with FOV to cover the entire pelvis and bladder. The prostate delineation sequence was a 2D axial T2 weighted sequence with FOV approximately 200x200 mm² acquired in the same session. The parameters for each of these scans are summarised in Dowling et al. 2015.

Three raters (two radiation oncologists and one radiation therapist) each segmented the prostate on the (small FOV) T2 weighted MR image. All 38 patients had their body outline, bone, bladder and rectum manually segmented from their large FOV T2 MR images. To ensure a prostate model was also available for the large FOV segmentation, the small FOV prostate manual segmentations were propagated to the large FOV images for each case. The prostate contours of the most experienced expert was used for training and validation in this work. The proposed deformable model segmentation approach was applied to the large FOV MR images

first then to the small FOV MR images once the prostate was suitably initialised. Only the prostate was segmented in the small FOV MR images using the a shape model built from the contours of the most experienced rater.

The surfaces were deformed along vertex normals using template matching (see section 2.4). To reduce computational complexity, the matching was done with a gray level model trained using a hand selected subset of cases that had well behaved appearance and shape characteristics. For the body and bones, the segmentation was also done at two image resolutions using an image pyramid, while the remaining objects were segmented at full resolution. The final prostate segmentation was obtained by also segmenting the second smaller FOV T2 weighted MR image.

2.6. Validation

To validate the proposed multi-object weighted deformable model approach, the automated segmentations \mathcal{A} were validated against the manual segmentations \mathcal{M} of the 38 cases using Dice’s similarity coefficient (DSC) (Dice 1945) as a measure of voxel overlap between the segmentations defined as

$$\text{DSC}(\mathcal{A}, \mathcal{M}) = \frac{2 |\mathcal{A} \cap \mathcal{M}|}{|\mathcal{A}| + |\mathcal{M}|} \quad (3)$$

Since the voxel based validation is a global measure of segmentation accuracy, we also validate our results using the mean absolute surface distance (MASD) that calculates the average absolute Euclidean distance between \mathcal{A} and \mathcal{M} surfaces. The proposed method was then compared to another two methods on the same dataset.

3. Results

Table 1. The DSC and surface distance errors between the proposed weighted multi-object method, a standard unweighted deformable model approach and a multi-atlas method. The t -tests show the statistical differences (if any) to the multi-atlas approach.

Method	Body Median DSC	Bone Median DSC	Bladder Median DSC	Rectum Median DSC	Prostate Median DSC	Prostate Median MASD (mm)	Prostate Median Hausdorff (mm)	Prostate p -value versus Multi-Atlas
Weighted	0.94	0.81	0.87	0.79	0.81	2.08	9.6	0.44
Unweighted	0.94	0.81	0.83	0.66	0.70	3.20	12.9	1×10^{-6}
Multi-Atlas	1.00	0.92	0.86	0.85	0.82	2.04	13.3	-

The proposed multi-object weighted deformable model approach was compared using DSC (Dice 1945) to the multi-atlas approach as reported by Dowling et al. 2015 on the same dataset and a standard (unweighted) deformable model approach that does not utilise spatial weights, i.e. with similar steps to figure 2, but with no partial fitting and all weights for all objects are set to unity. The results for each approach is shown in table 1. A visualisation of the proposed approach (presented as a movie) when applied to a case in the dataset is provided as supplementary material.

The automatic prostate segmentations of the proposed method was found to have a (mean DSC, median DSC, Hausdorff distance, median MASD) of (0.79, 0.81, 9.6 ± 2.9 mm, 2.0 ± 0.9 mm). The approximate maximum (large FOV, small FOV) time required for the

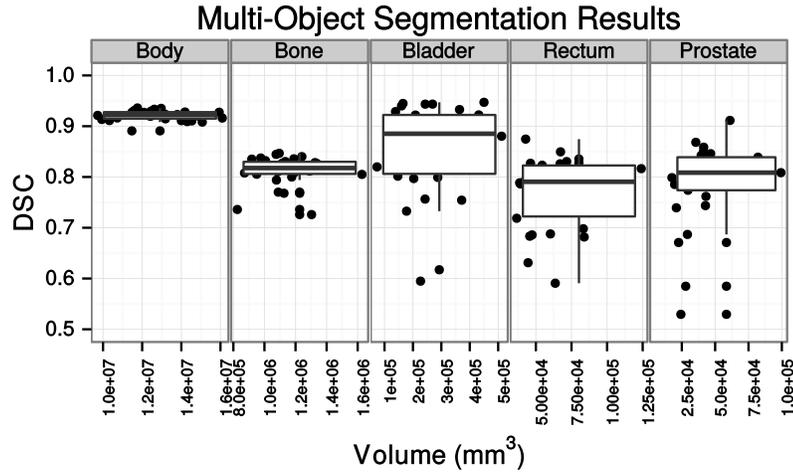


Figure 3. The Dice's similarity coefficient (DSC) box plot with respect to the resulting volume of the proposed method per object.

entire prostate segmentation result to be obtained was (12 minutes, 3 minutes), leading to a total maximum computation time of round 15 minutes. The individual DSC (as box plots) for each case and their respective volumes for the proposed approach are shown in figure 3 for all objects. The optimal spatial weights for initialising the subsequent object, while simultaneously fitting the current object (such as those weights shown in figure 2), was found to be between 0.6-0.8 across all objects. Figure 4 shows the study of these spatial weights with respect to the resulting DSC of the subsequent object for the different objects in this work. The DSCs for the weights in figure 4 were computed using brute force on a subset of 12 cases by keeping the targeted object fixed at maximum accuracy (i.e. a weight of 1.0).

To ensure that the actual error of the segmentations are reasonable for applications such as treatment planning for EBRT, the MASD errors of the prostate, bone and rectum were computed and are shown in figures 5 and 6. Following the work of Klein et al. 2008, figure 5 shows the overall prostate MASD errors across all cases as a Mollweide or homolographic projection (Snyder 1987) (assuming a genus zero embedding of the prostate surfaces) to highlight the problematic areas for the segmentation method.

4. Discussion

The validation of the proposed multi-object approach showed that the prostate segmentations were of a similar DSC as the multi-atlas approach, showing no significant difference ($p > 0.4$), and inter-rater scores as reported by Dowling et al. 2015 on the same data. Although similar results are obtained, the computation times are vastly different between the two methods with the entire proposed pipeline per case computation time being nearly an order of magnitude faster (approximately 12-15 minutes) when compared to a (multi-threaded) multi-atlas approach without cloud infrastructure (approximately taking 120 minutes in total). Figure 4 also showed that the partial weights used for the simultaneous initialisation of the subsequent object is important and that a weight in between 0.5-0.8 should suffice for most use cases.

The main locations of the prostate surface errors shown in figure 5 seem to relate to the medial and lateral regions marked in light blue and green on the region labelling of

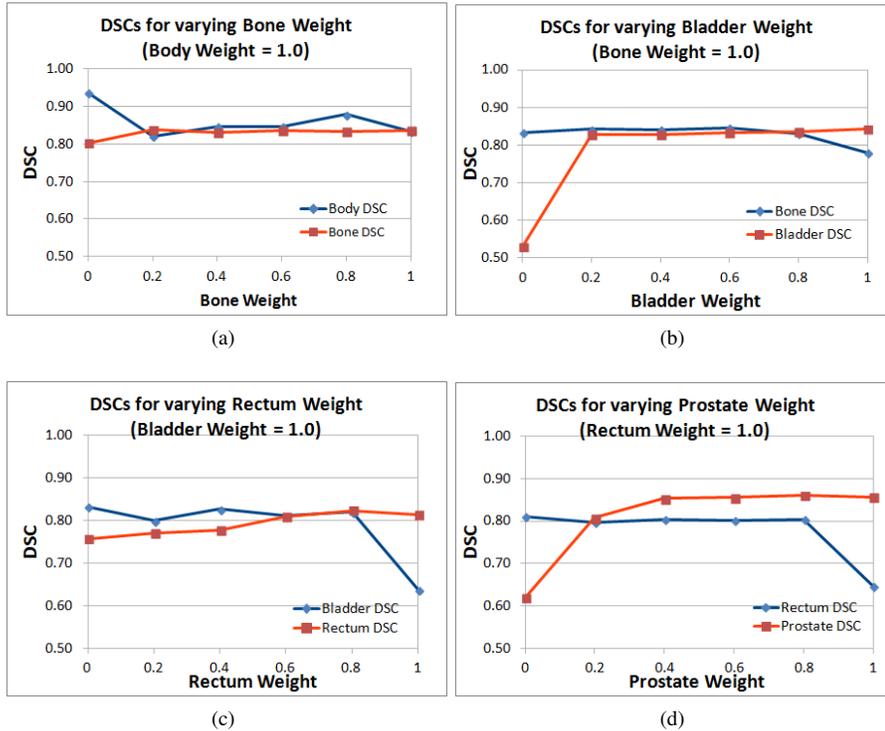


Figure 4. The spatial weight values of the subsequent object while accurately segmenting the current target object (weight of 1.0) as a function of resulting DSC. The stages of the pipeline in figure 2 are shown as body to bone (a), bone to bladder (b), bladder to rectum (c) and rectum to prostate (d).

figure 5(b). This appears to correspond to local minima due to other tendons and soft tissue edges in proximity to the prostate that are similar to the prostate edges. Spatially constrained segmentation, such as those presented in Martin, Troccaz and Daanen 2010 or a more sub-focused model (per region in the prostate as opposed to just the multiple objects) to restrict deformation in this region as a potential remedy for these errors.

There were four low performing cases with the proposed method. Three of these cases were found to have had a transurethral prostate resection (TURP) in the past, which may have led to unusual bladder-prostate interfaces in the small FOV T2 MR images (Hricak and Carrington 1992). These appearances are shown in figure 7. Since the proposed method utilises gray level models, these cases introduce errors for normal bladder-prostate appearances as they essentially allow normal cases to fit sections of the prostate within the bladder. A possible solution to these cases is to use a separate gray level model for TURP patients, since the patient history is known for the training cases and the new case to be segmented. The remaining low performing case had an unusual bladder interfaces in the large FOV to be discussed shortly, leading to a sub-optimal initialisation for the small FOV image resulting in a low DSC.

There were no distinct advantages of the proposed method with respect to the (relatively) rigid structures in the dataset, namely the body outline and bones, when compared to the unweighted approach (as seen in table 1 and figure 4). The multi-atlas approach out-performed the proposed method and unweighted method in both these two objects. The main reasons for

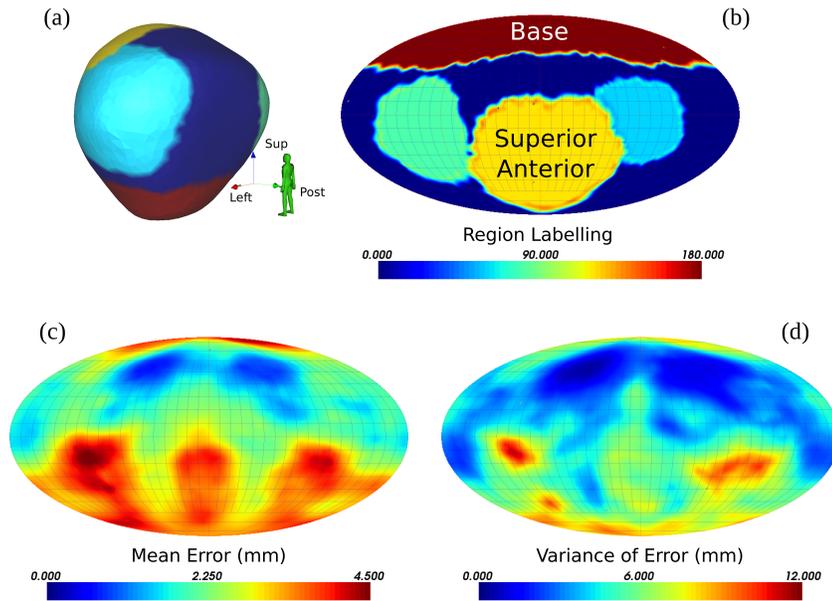


Figure 5. Surface distance errors with respect to the manual segmentations of the automatic prostate segmentation using the proposed multi-object approach. (a) and (b) shows the prostate region labels. (c) and (d) show the mean and variance of the distance errors. The distances and regions were parameterised from prostate shapes to corresponding spheres and mapped to 2D using Mollweide projection (Snyder 1987).

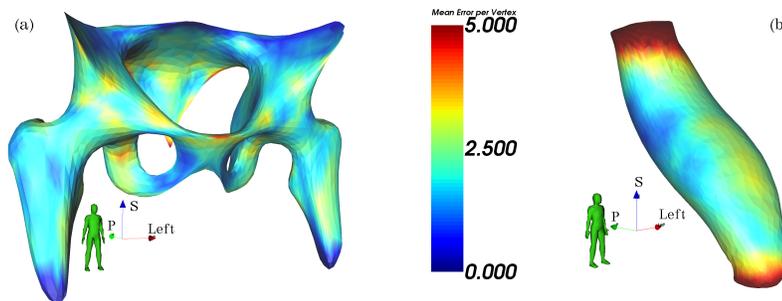


Figure 6. Surface distance errors with respect to the manual segmentations of the (a) bone and (b) rectum automatic segmentations (not to scale) using the proposed multi-object approach. The areas in proximity to the prostate useful for tasks such as treatment planning for EBRT have relatively small errors.

this discrepancies were found to be the FOV differences within the dataset and the template surfaces. Although these surfaces were created to account as best as possible for these differences, the average (or larger than average) consensus were used to generate the template surfaces and thus the differences were still large. It is likely that larger and more accurate (in terms of sampling) template surfaces could remedy this problem, though at a potential cost to computation times. Methods for dealing with FOV differences have already been devised for the bone and could also be incorporated, however with an increased computation time (Schmid, Kim and Magnat-Thalmann 2011; Chandra et al. 2014).

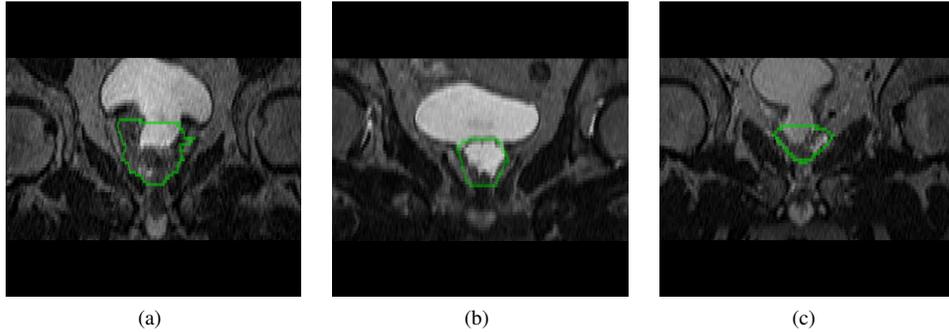


Figure 7. A set of unusual bladder-prostate interfaces present in the dataset who have had a TURP. (a)-(c) each show the three cases that perform with low DSC as a slice of the small FOV MR image with green contours outlining the manual segmentations of the prostate. Notice the bladder like appearance within the prostate, which is problematic when using a gray level model to fit deformable models in these cases.

The unweighted approach also performed as well as the proposed approach for the rigid objects, namely the body and bone. This is because the objects are suitably discernible and large enough to be segmented without much initialisation. It is possible that more centred scans, where the bone is properly centred, the body outline segmentation step could be omitted altogether without much loss in prostate segmentation accuracy in applications where dosimetry is not needed. Herein also lies one limitation of the proposed approach; it is dependent on having multiple objects that can be segmented automatically present in the dataset, some of whom must be easily discernible for an accurate self-initialisation. The unweighted approach however, does not perform well for more highly shape variable (non-rigid) objects such as the bladder, rectum and prostate, achieving low DSC values as shown in table 1.

The low bone segmentation DSC of the proposed approach, when compared to the multi-atlas method, could be due to the fact that the labels for the pelvis and femur are not separated. This results in a deformable model that has low sensitivity in the joint area because it cannot take into account the flexions of the joint. Chandra et al. (2014) have developed methods to account for these issues by using large CT-based template surfaces with individual bone components. This could also be added to improve bone segmentations, but again at the cost of computation time, which would be of the order of 6-8 minutes just for the bone. Also note that figure 6(a) shows that the surface distance errors for the bone segmentations are not that high near the central pelvic area, an area of important in applications such as prostate EBRT.

The rectum segmentation DSC scores were slightly lower than the multi-atlas approach, but surface distance analysis showed that the errors of the overall areas closest to the prostate are small (around 2 mm, see figure 6(b)). This is acceptable for applications such prostate EBRT treatment planning, where the prostate margins are specified to reduce rectal toxicity. The greatest areas of disagreement between the automated and manual rectum results were the sigmoid junction and inferior cut-off regions (see figure 6(b)), which was defined as the end point for the manual segmentations based on the position of the prostate. Note that this area was also subject to larger on average inter-observer variability. The disagreement was likely because it is difficult to profile match for the partial fitting of a tube like structure such as the rectum and where the “end” of the structure might be. A large and more complete rectum surface that covers more of the anatomy would likely perform better.

To achieve the computation times of the proposed approach, the gray level models were

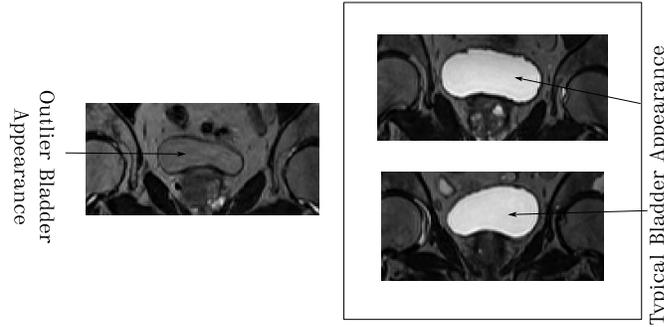


Figure 8. Outlier in bladder appearance within the current dataset leading to low DSC scores for the bladder and prostate.

restricted to subset of the training cases that were chosen by hand. A set of eight cases were chosen to meet the general appearance of most cases. For the bladder however, two particular cases had to be omitted due to highly irregular appearance as shown in figure 8. The gray level model for this case could not account for the drastic difference in internal and external intensity differences. Performance of the segmentation scheme could be improved by (automatically) choosing a subset from a much larger training set of cases using methods such as pre-clustering (Langerak et al. 2013). These cases were among the lowest DSC scores of the proposed method.

Further work needs to be done in improving the accuracy and computation times of the proposed approach. More analysis of the weights, especially at a sub-object level could lead to a more accurate method. An automated method for selecting and clustering cases based on object profiles for use in gray level models could result in improved accuracy and speed. A recent method for improving the capacity of shape models to efficiently represent multi-object shape using multi-resolution approach (Cerroloza et al. 2015), as well as using graphical processing unit (GPU) implementations, could also be useful in these endeavours.

Conclusion

This work presented a fast, self-initialised multi-object ASM approach using applied to pelvic structures in 3D MR images. The self-initialisation is achieved via a hierarchical multi-object pipeline controlled using spatially weighted shape learning to simultaneously fit the current (more easily discernible) object and initialise the subsequent (more difficult) object. The proposed method achieved the same segmentation accuracy for the prostate as a multi-atlas approach, but with a computation time of nearly an order of magnitude faster (approximately 12-15 minutes compared to 120 minutes) of such an approach without cloud infrastructure (see table 1 and figure 5). With recent enhancements for multi-resolution shape models, it should be possible to engineer further improvements in performance of the proposed approach in terms of computation time.

Acknowledgments

This work was supported by Cancer Council New South Wales research grant RG11-05, the Prostate Cancer Foundation of Australia (Movember Young Investigator Grant YI2011) and Cure Cancer Australia.

References

- Cerrolaza, Juan J., Mauricio Reyes, Ronald M. Summers, Miguel Ángel González-Ballester and Marius George Linguraru (2015). “Automatic multi-resolution shape modeling of multi-organ structures”. In: *Medical Image Analysis* 25.1, pp. 11–21. ISSN: 1361-8415. DOI: [10.1016/j.media.2015.04.003](https://doi.org/10.1016/j.media.2015.04.003).
- Chandra, Shekhar S., Ying Xia, Craig Engstrom, Stuart Crozier, Raphael Schwarz and Jurgen Fripp (Apr. 2014). “Focused shape models for hip joint segmentation in 3D magnetic resonance images”. In: *Medical Image Analysis* 18.3, pp. 567–578. ISSN: 1361-8415. DOI: [10.1016/j.media.2014.02.002](https://doi.org/10.1016/j.media.2014.02.002).
- Chandra, S.S., J.A. Dowling, Kai-Kai Shen, P. Raniga, J.P.W. Pluim, P.B. Greer, O. Salvado and J. Fripp (2012). “Patient Specific Prostate Segmentation in 3-D Magnetic Resonance Images”. In: *Medical Imaging, IEEE Transactions on* 31.10, pp. 1955–1964. ISSN: 0278-0062. DOI: [10.1109/TMI.2012.2211377](https://doi.org/10.1109/TMI.2012.2211377).
- Cootes, T. F., C. J. Taylor, D. H. Cooper and J. Graham (1995). “Active Shape Models-Their Training and Application”. In: *Computer Vision and Image Understanding* 61.1, pp. 38–59. ISSN: 1077-3142. DOI: [10.1006/cviu.1995.1004](https://doi.org/10.1006/cviu.1995.1004).
- Cosío, F. Arámbula (Aug. 2008). “Automatic initialization of an active shape model of the prostate”. In: *Medical Image Analysis* 12.4, pp. 469–483. ISSN: 1361-8415. DOI: [16/j.media.2008.02.001](https://doi.org/10.1016/j.media.2008.02.001).
- Datar, Manasi, Yaniv Gur, Beatriz Paniagua, Martin Styner and Ross Whitaker (2011). “Geometric Correspondence for Ensembles of Nonregular Shapes”. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI* 14.Pt 2, pp. 368–375.
- Davies, R. H., C. J. Twining, T. F. Cootes and C. J. Taylor (2010). “Building 3-D Statistical Shape Models by Direct Optimization”. In: *Medical Imaging, IEEE Transactions on* 29.4, pp. 961–981. ISSN: 0278-0062. DOI: [10.1109/TMI.2009.2035048](https://doi.org/10.1109/TMI.2009.2035048).
- Davies, Rhodri H., Carole J. Twining and Chris Taylor (2008). “Groupwise surface correspondence by optimization: Representation and regularization”. In: *Medical Image Analysis* 12.6, pp. 787–796. ISSN: 1361-8415. DOI: [DOI:10.1016/j.media.2008.03.009](https://doi.org/10.1016/j.media.2008.03.009).
- Dice, Lee R (1945). “Measures of the Amount of Ecologic Association Between Species”. In: *Ecology* 26.3, pp. 297–302. ISSN: 00129658.
- Dowling, Jason A. et al. (2015). “Automatic substitute CT generation and contouring for MRI-alone external beam radiation therapy from standard MRI sequences”. English. In: *International Journal of Radiation Oncology Biology Physics* 93.5, pp. 1144–1153. ISSN: 0360-3016. DOI: [10.1016/j.ijrobp.2015.08.045](https://doi.org/10.1016/j.ijrobp.2015.08.045).
- Ecabert, O., J. Peters, H. Schramm, C. Lorenz, J. Von Berg, M.J. Walker, M. Vembar, M.E. Olszewski, K. Subramanian, G. Lavi and J. Weese (2008). “Automatic Model-Based Segmentation of the Heart in CT Images”. In: *Medical Imaging, IEEE Transactions on* 27.9, pp. 1189–1201. ISSN: 0278-0062. DOI: [10.1109/TMI.2008.918330](https://doi.org/10.1109/TMI.2008.918330).
- Engstrom, Craig M., Jurgen Fripp, Valer Jurcak, Duncan G. Walker, Olivier Salvado and Stuart Crozier (June 2011). “Segmentation of the quadratus lumborum muscle using statistical shape modeling”. en. In: *Journal of Magnetic Resonance Imaging* 33.6, pp. 1422–1429. ISSN: 1522-2586. DOI: [10.1002/jmri.22188](https://doi.org/10.1002/jmri.22188).
- Fripp, J., S. Crozier, S.K. Warfield and S. Ourselin (Jan. 2010). “Automatic Segmentation and Quantitative Analysis of the Articular Cartilages from Magnetic Resonance Images of the Knee”. In: *Medical Imaging, IEEE Transactions on* 29.1, pp. 55–64. DOI: [10.1109/TMI.2009.2024743](https://doi.org/10.1109/TMI.2009.2024743).
- Gao, Yi, Ron Kikinis, Sylvain Bouix, Martha Shenton and Allen Tannenbaum (Aug. 2012). “A 3D interactive multi-object segmentation tool using local robust statistics driven active contours”. In: *Medical Image Analysis* 16.6, pp. 1216–1227. ISSN: 1361-8415. DOI: [10.1016/j.media.2012.06.002](https://doi.org/10.1016/j.media.2012.06.002).
- Garland, Michael and Paul S. Heckbert (1997). “Surface simplification using quadric error metrics”. In: *Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH)*. New York,

- NY, USA: ACM Press/Addison-Wesley Publishing Co., pp. 209–216. ISBN: 0-89791-896-7. DOI: [10.1145/258734.258849](https://doi.org/10.1145/258734.258849).
- Ghose, Soumya, Arnau Oliver, Robert Martí, Xavier Lladó, Joan C. Vilanova, Jordi Freixenet, Jhimli Mitra, Désiré Sidibé and Fabrice Meriaudeau (Oct. 2012). “A survey of prostate segmentation methodologies in ultrasound, magnetic resonance and computed tomography images”. English. In: *Computer Methods and Programs in Biomedicine* 108.1, pp. 262–287. ISSN: 0169-2607, 1872-7565. DOI: [10.1016/j.cmpb.2012.04.006](https://doi.org/10.1016/j.cmpb.2012.04.006).
- Glocker, Ben, Olivier Pauly, Ender Konukoglu and Antonio Criminisi (Oct. 2012). “Joint Classification-Regression Forests for Spatially Structured Multi-object Segmentation”. en. In: *Computer Vision – ECCV 2012*. Ed. by Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato and Cordelia Schmid. Lecture Notes in Computer Science 7575. Springer Berlin Heidelberg, pp. 870–881. ISBN: 978-3-642-33764-2, 978-3-642-33765-9. DOI: [10.1007/978-3-642-33765-9_62](https://doi.org/10.1007/978-3-642-33765-9_62).
- HajGhanbari, Bahareh, Ghassan Hamarneh, Neda Changizi, Aaron D. Ward and W. Darlene Reid (Feb. 2011). “MRI-Based 3D Shape Analysis of Thigh Muscles: Patients with Chronic Obstructive Pulmonary Disease Versus Healthy Adults”. In: *Academic Radiology* 18.2, pp. 155–166. ISSN: 1076-6332. DOI: [10.1016/j.acra.2010.09.008](https://doi.org/10.1016/j.acra.2010.09.008).
- Heimann, Tobias, Ivo Wolf and Hans-Peter Meinzer (2006). “Active Shape Models for a Fully Automated 3D Segmentation of the Liver - An Evaluation on Clinical Data”. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*. Ed. by Rasmus Larsen, Mads Nielsen and Jon Sporring. Lecture Notes in Computer Science 4191. Springer Berlin Heidelberg, pp. 41–48. ISBN: 978-3-540-44727-6, 978-3-540-44728-3. DOI: [10.1007/11866763_6](https://doi.org/10.1007/11866763_6).
- Hricak, Hedvig and Bernadette M. Carrington (Feb. 1992). *MRI of the Pelvis: A Text Atlas*. English. Norwalk, Conn. : London, England: Appleton & Lange. ISBN: 9780838565278.
- Hufnagel, Heike, Xavier Pennec, Jan Ehrhardt, Nicholas Ayache and Heinz Handels (2008). “Generation of a statistical shape model with probabilistic point correspondences and the expectation maximization- iterative closest point algorithm”. In: *International Journal of Computer Assisted Radiology and Surgery* 2.5, pp. 265–273. ISSN: 1861-6410. DOI: [10.1007/s11548-007-0138-9](https://doi.org/10.1007/s11548-007-0138-9).
- Karlsson, Mikael, Magnus G. Karlsson, Tufve Nyholm, Christopher Amies and Björn Zackrisson (June 2009). “Dedicated magnetic resonance imaging in the radiotherapy clinic”. eng. In: *International Journal of Radiation Oncology, Biology, Physics* 74.2, pp. 644–651. ISSN: 1879-355X. DOI: [10.1016/j.ijrobp.2009.01.065](https://doi.org/10.1016/j.ijrobp.2009.01.065).
- Klein, Stefan, Uulke A. van der Heide, Irene M. Lips, Marco van Vulpen, Marius Staring and Josien P. W. Pluim (2008). “Automatic segmentation of the prostate in 3D MR images by atlas matching using localized mutual information”. In: *Medical Physics* 35.4, pp. 1407–1417. DOI: [10.1118/1.2842076](https://doi.org/10.1118/1.2842076).
- Langerak, Thomas R., Floris F. Berendsen, Uulke A. Van der Heide, Alexis N. T. J. Kotte and Josien P. W. Pluim (Aug. 2013). “Multiatlas-based segmentation with preregistration atlas selection”. In: *Medical Physics* 40.9, p. 091701. ISSN: 0094-2405. DOI: [10.1118/1.4816654](https://doi.org/10.1118/1.4816654).
- Litjens, Geert et al. (Feb. 2014). “Evaluation of prostate segmentation algorithms for MRI: The PROMISE12 challenge”. In: *Medical Image Analysis* 18.2, pp. 359–373. ISSN: 1361-8415. DOI: [10.1016/j.media.2013.12.002](https://doi.org/10.1016/j.media.2013.12.002).
- Lorensen, William E and Harvey E Cline (1987). “Marching cubes: A high resolution 3D surface construction algorithm”. In: *Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH)*. Vol. 21. 4. ACM Press/Addison-Wesley Publishing Co., pp. 163–169. DOI: [10.1145/37402.37422](https://doi.org/10.1145/37402.37422).
- Ma, Jun, Le Lu, Yiqiang Zhan, Xiang Zhou, Marcos Salganicoff and Arun Krishnan (2010). “Hierarchical Segmentation and Identification of Thoracic Vertebra Using Learning-Based Edge Detection and Coarse-to-Fine Deformable Model”. en. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2010*. Ed. by Tianzi Jiang, Nassir Navab, Josien P. W. Pluim and Max A. Viergever. Lecture Notes in Computer Science 6361. Springer Berlin Heidelberg, pp. 19–27. ISBN: 978-3-642-15704-2, 978-3-642-15705-9. DOI: [10.1007/978-3-642-15705-9_3](https://doi.org/10.1007/978-3-642-15705-9_3).

- Martin, Sebastien, Jocelyne Troccaz and Vincent Daanen (2010). “Automated segmentation of the prostate in 3D MR images using a probabilistic atlas and a spatially constrained deformable model”. In: *Medical Physics* 37.4, pp. 1579–1590. DOI: [10.1118/1.3315367](https://doi.org/10.1118/1.3315367).
- Neubert, A, J Fripp, C Engstrom, D Walker, M-A Weber, R Schwarz and S Crozier (Nov. 2013). “Three-dimensional morphological and signal intensity features for detection of intervertebral disc degeneration from magnetic resonance images”. eng. In: *Journal of the American Medical Informatics Association: JAMIA* 20.6, pp. 1082–1090. ISSN: 1527-974X. DOI: [10.1136/amiajnl-2012-001547](https://doi.org/10.1136/amiajnl-2012-001547).
- Nyholm, Tufve and Joakim Jonsson (July 2014). “Counterpoint: Opportunities and Challenges of a Magnetic Resonance Imaging–Only Radiotherapy Work Flow”. English. In: *Seminars in Radiation Oncology* 24.3, pp. 175–180. ISSN: 1053-4296. DOI: [10.1016/j.semradonc.2014.02.005](https://doi.org/10.1016/j.semradonc.2014.02.005).
- Okada, Toshiyuki, Keita Yokota, Masatoshi Hori, Masahiko Nakamoto, Hironobu Nakamura and Yoshinobu Sato (2008). “Construction of Hierarchical Multi-Organ Statistical Atlases and Their Application to Multi-Organ Segmentation from CT Images”. en. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2008*. Ed. by Dimitris Metaxas, Leon Axel, Gabor Fichtinger and Gabor Szekely. Lecture Notes in Computer Science 5241. Springer Berlin Heidelberg, pp. 502–509. ISBN: 978-3-540-85987-1, 978-3-540-85988-8. DOI: [10.1007/978-3-540-85988-8_60](https://doi.org/10.1007/978-3-540-85988-8_60).
- Pearson, K. (1901). “On lines and planes of closest fit to systems of points in space”. In: *Philosophical Magazine* 2.6, pp. 559–572.
- Schmid, Jérôme, Jinman Kim and Nadia Magnenat-Thalmann (2011). “Robust statistical shape models for MRI bone segmentation in presence of small field of view”. In: *Medical Image Analysis* 15.1, pp. 155–168. ISSN: 1361-8415. DOI: [DOI:10.1016/j.media.2010.09.001](https://doi.org/10.1016/j.media.2010.09.001).
- Shang, Yan, Guangda Su and Olaf Dössel (2006). “Hierarchical 3D Shape Model for Segmentation of 4D MR Cardiac Images”. en. In: *Medical Imaging and Augmented Reality*. Ed. by Guang-Zhong Yang, TianZi Jiang, Dinggang Shen, Lixu Gu and Jie Yang. Lecture Notes in Computer Science 4091. Springer Berlin Heidelberg, pp. 333–340. ISBN: 978-3-540-37220-2, 978-3-540-37221-9. DOI: [10.1007/11812715_42](https://doi.org/10.1007/11812715_42).
- Shen, Kai-kai, Jurgen Fripp, Fabrice Mériaudeau, Gaël Chételat, Olivier Salvado and Pierrick Bourgeat (2012). “Detecting global and local hippocampal shape changes in Alzheimer’s disease using statistical shape models”. In: *NeuroImage* 59.3, pp. 2155–2166. ISSN: 1053-8119. DOI: [10.1016/j.neuroimage.2011.10.014](https://doi.org/10.1016/j.neuroimage.2011.10.014).
- Shotton, Jamie, John Winn, Carsten Rother and Antonio Criminisi (Dec. 2007). “TextonBoost for Image Understanding: Multi-Class Object Recognition and Segmentation by Jointly Modeling Texture, Layout, and Context”. en. In: *International Journal of Computer Vision* 81.1, pp. 2–23. ISSN: 0920-5691, 1573-1405. DOI: [10.1007/s11263-007-0109-1](https://doi.org/10.1007/s11263-007-0109-1).
- Smith, Rebecca, Kayvan Najarian and Kevin Ward (Nov. 2009). “A hierarchical method based on active shape models and directed Hough transform for segmentation of noisy biomedical images; application in segmentation of pelvic X-ray images”. en. In: *BMC Medical Informatics and Decision Making* 9.1, pp. 1–11. ISSN: 1472-6947. DOI: [10.1186/1472-6947-9-S1-S2](https://doi.org/10.1186/1472-6947-9-S1-S2).
- Snyder, J.P. (1987). “Map projections - A working manual”. In: *U.S. Geological Survey Professional Paper* 1395, pp. 249–252.
- Song, Qi, Junjie Bai, M.K. Garvin, M. Sonka, J.M. Buatti and Xiaodong Wu (Feb. 2013). “Optimal Multiple Surface Segmentation With Shape and Context Priors”. In: *IEEE Transactions on Medical Imaging* 32.2, pp. 376–386. ISSN: 0278-0062. DOI: [10.1109/TMI.2012.2227120](https://doi.org/10.1109/TMI.2012.2227120).
- Taubin, Gabriel, Tong Zhang and Gene H Golub (1996). “Optimal Surface Smoothing as Filter Design”. In: *Proceedings of the 4th European Conference on Computer. ECCV*. Springer-Verlag, pp. 283–292. ISBN: 3-540-61122-3.
- Thomaz, C.E., G.A. Giralddi, J.F.P. da Costa and D.F. Gillies (2010). *A Simple and Efficient Supervised Method for Spatially Weighted PCA in Face Image Analysis*. Technical Report. Department of Computing, Imperial College London.

- Xia, Ying, Jurgen Fripp, Shekhar S. Chandra, Raphael Schwarz, Craig Engstrom and Stuart Crozier (Oct. 2013). "Automated bone segmentation from large field of view 3D MR images of the hip joint". en. In: *Physics in Medicine and Biology* 58.20, p. 7375. ISSN: 0031-9155. DOI: [10.1088/0031-9155/58/20/7375](https://doi.org/10.1088/0031-9155/58/20/7375).
- Yokota, Futoshi, Toshiyuki Okada, Masaki Takao, Nobuhiko Sugano, Yukio Tada and Yoshinobu Sato (2009). "Automated Segmentation of the Femur and Pelvis from 3D CT Data of Diseased Hip Using Hierarchical Statistical Shape Model of Joint Structure". en. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2009*. Ed. by Guang-Zhong Yang, David Hawkes, Daniel Rueckert, Alison Noble and Chris Taylor. Lecture Notes in Computer Science 5762. Springer Berlin Heidelberg, pp. 811–818. ISBN: 978-3-642-04270-6, 978-3-642-04271-3. DOI: [10.1007/978-3-642-04271-3_98](https://doi.org/10.1007/978-3-642-04271-3_98).
- Zheng, Yefeng, A. Barbu, B. Georgescu, M. Scheuering and D. Comaniciu (2008). "Four-Chamber Heart Modeling and Automatic Segmentation for 3-D Cardiac CT Volumes Using Marginal Space Learning and Steerable Features". In: *IEEE Transactions on Medical Imaging* 27.11, pp. 1668–1681. ISSN: 0278-0062. DOI: [10.1109/TMI.2008.2004421](https://doi.org/10.1109/TMI.2008.2004421).