

Pieces-of-parts for supervoxel segmentation with global context: Application to DCE-MRI tumour delineation

Benjamin Irving^{a,*}, James M. Franklin^b, Bartłomiej W. Papież^a, Ewan M. Anderson^b, Ricky A. Sharma^c, Fergus V. Gleeson^b, Sir Michael Brady^c,
Julia A. Schnabel^{a,1}

^a*Institute of Biomedical Engineering, Department of Engineering Science, Old Road Campus Research Building, University of Oxford, Headington, Oxford, OX3 7DQ, UK*

^b*Department of Radiology, Oxford University Hospitals NHS Trust, Churchill Hospital, Oxford, OX3 7LE, UK*

^c*Department of Oncology, Old Road Campus Research Building, University of Oxford, Headington, Oxford, OX3 7DQ, UK*

Abstract

Rectal tumour segmentation in dynamic contrast-enhanced MRI (DCE-MRI) is a challenging task, and an automated and consistent method would be highly desirable to improve the modelling and prediction of patient outcomes from tissue contrast enhancement characteristics – particularly in routine clinical practice. A framework is developed to automate DCE-MRI tumour segmentation, by introducing: *perfusion-supervoxels* to over-segment and classify DCE-MRI volumes using the dynamic contrast enhancement characteristics; and the *pieces-of-parts* graphical model, which adds global (anatomic) constraints that further refine the supervoxel components that comprise the tumour. The framework was evaluated on 23 DCE-MRI scans of patients with rectal adenocarcinomas, and achieved a voxelwise area-under the receiver operating characteristic curve (AUC) of 0.97 compared to expert delineations. Creating a binary tumour segmentation, 21 of the 23 cases were segmented correctly with a median Dice similarity coefficient (DSC) of 0.63, which is close to the inter-rater variability of this challenging task. A

*Corresponding author

Email address: benjamin.irving@eng.ox.ac.uk (Benjamin Irving)

¹JAS is now at the Department of Biomedical Engineering, Division of Imaging Sciences and Biomedical Engineering, King's College London, St Thomas' Hospital, London, SE1 7EH, UK

second study is also included to demonstrate the method’s generalisability and achieved a DSC of 0.71. The framework achieves promising results for the underexplored area of rectal tumour segmentation in DCE-MRI, and the methods have potential to be applied to other DCE-MRI and supervoxel segmentation problems.

Keywords: Parts-based graphical models, supervoxel, classification, segmentation, DCE-MRI, rectal tumour

1. Introduction

Dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) shows promise for tumour characterisation and measuring response to therapy from tissue perfusion characteristics (Karahaliou et al., 2014; Gollub et al., 2012; 5 Goh et al., 2007). In DCE-MRI, a series of volumes are acquired during and after the injection of a Gadolinium chelate based paramagnetic contrast agent (CA), such as MultihanceTM, OmniscanTM, ProhanceTM, or DotaremTM. A bolus of intravenous CA circulates through the vascular system and perfuses into the extravascular-extracellular space (EES) of the tissue, resulting in 10 tissue enhancement, which depends on the tissue type and vascularisation.

Neovascularisation, the formation of new vessels, is a key process of tumour growth, with angiogenesis being a common mechanism in colorectal tumours. This results in highly vascularised tumours with vessels that are thin, fragile and tortuous (Goh et al., 2007). The vasculature is also chaotic 15 and leaky with regions of the tumour exhibiting hypoxia and necrosis, which leads to complex contrast enhancement patterns in DCE-MRI imaging of rectal tumours; Figure 1 shows an example of contrast enhancement in a rectal DCE-MRI scan. Quantitative and semi-quantitative measures such as pharmacokinetic modelling can be used to analyse these perfusion characteristics 20 in DCE-MRI (Tofts et al., 1999) and have shown the ability to assess the tumour (Goh et al., 2007). Heterogeneity in tumours can be further exploited to quantify tumour progression and response to therapy (Karahaliou et al., 2014; Gollub et al., 2012). Figure 2 shows example enhancement curves for a rectal tumour and illustrates the heterogeneity of the tumour and surrounding 25 region, while also demonstrating some consistency in tumour subregions.

To derive tumour-specific perfusion parameters, accurate delineation of the rectal tumour is required. It is difficult to visually distinguish tumour from surrounding structures on T1w sequences used for DCE-MRI; in clinical




Figure 1: A DCE-MRI axial slice at baseline, the 15th and 24th time point (Case 09).




Figure 2: Signal enhancement curves of the rectal tumour and surrounding tissue, which illustrates the heterogeneity in the tumour and surrounding region, as well as similarities in enhancement between local regions. *Green*, *red* and *blue* are examples of enhancement for regions of the tumour, while *cyan* shows the rectal wall and mucosa, *orange* shows neighbouring tissue, and *brown* the non-enhancing lumen.

practice T2w imaging is preferred (Brown et al., 2005), and can be registered
 30 to the DCE-MRI scan. DCE-MRI tumour delineation is further complicated
 due to the contrast-varying nature of the dynamic scan. However,
 even on T2w scans, manual segmentation of colorectal tumours is a difficult
 task, with substantial inter- and intra-rater variability (Franklin et al.,
 2014; Goh et al., 2008), due to: partial volume effects in the axial plane;
 35 complex anatomy in the lower rectum, making it difficult to delineate normal
 anatomical structures; wall thickening due to venous congestion; and
 mucinous tumours. There is also often local motion between the T2w and
 DCE-MRI scans, which can lead to poor registration of the volumes, and
 therefore a poor delineation of the tumour in DCE-MRI. An additional chal-
 40 lenge in this study, and medical imaging in general, is the limited number
 of cases available. Therefore, a requirement of any automated DCE-MRI

segmentation method is that it must be able to learn complex relationships efficiently from small datasets. This study is motivated by the rectal imaging trial (Churchill Hospital, Oxford, UK), which is an ongoing trial that aims
45 to use pre-therapy DCE-MRI to assess patient outcomes after therapy.

In this work, we propose an automated DCE-MRI segmentation approach that exploits features of the dynamic sequences to segment tumours, and provides a robust and consistent delineation for perfusion analysis. The key contributions of this method are: the introduction of *perfusion-supervoxels*
50 (supervoxels for dynamic contrast enhanced images, which are used to learn local characteristics of the tumour); and *pieces-of-parts*, a novel formulation of parts-based graphical models (PGM) (Felzenszwalb and Huttenlocher, 2005; Felzenszwalb et al., 2010) to directly improve supervoxel-based segmentations by including global anatomical constraints.

Figure 3 shows a flowchart of the framework, where: a 4D contrast-enhanced DCE-MRI volume is acquired, the region is over-segmented using perfusion-supervoxels and each supervoxel is classified using learnt neighbourhood features. Pieces-of-parts then provides global constraints to the
55 tumour segmentation using belief propagation on likely organ locations. Finally, the tumour segmentation is evaluated against an expert ground truth. This framework is applied to the novel and challenging task of automatic rectal tumour segmentation in DCE-MRI; the authors are not aware of any other method to automatically segment rectal tumours from DCE-MRI.

The initial steps in the development of this method were introduced in
65 Irving et al. (2014), which includes the formulation of supervoxels for perfusion images and development of a supervoxel classifier. This is extended to include the novel pieces-of-parts graphical model. The validation is also improved using a more robust ground truth from expert delineations on multiple modalities and, while, previously, a number of cases required manual
70 intervention, this method is evaluated in a fully automated fashion.

In the following section we place this work in the context of the state-of-the-art. Section 3 outlines the method which is divided into perfusion-supervoxels (Section 3.1), extraction of features (Section 3.2) and pieces-of-parts (Section 3.3) which relate to Steps 1-3 in Figure 3. The dataset, experimental set-up and results are described in Section 4. Finally, the discussion
75 is found in Section 5.




Figure 3: Flowchart of the method for a new DCE-MRI volume. Steps 1 - 3 show the key steps of the framework.

2. Background

Pharmacokinetic models are a common method of extracting quantitative measures of perfusion from the contrast enhancement curves of DCE-MRI volumes (Tofts et al., 1999). However, these measures depend on the selection of the compartment model, which in turn depends on the tissue characteristics, as well as the choice or measurement of an arterial input function (Irving et al., 2013). Tumour heterogeneity means that a single model is rarely appropriate for the whole tumour and surrounding region (Kallehauge et al., 2014). As an alternative to these model based methods, Hamy et al. (2014) apply principal component analysis (PCA) to the signal enhancement (SE) curves, and use this decomposition to distinguish motion from enhancement in a DCE-MRI based registration framework. In this study, we also take a data driven approach to avoid any assumptions about the underlying enhancement patterns in DCE-MRI; PCA is used to decompose the signal enhancement but is applied to supervoxel and feature extraction, unlike previously for registration (Hamy et al., 2014).

DCE-MRI tumour segmentation, and in particular colorectal tumour segmentation, remains relatively unexplored. Breast DCE-MRI segmentation methods are more common, and the fuzzy c-means clustering approach is an established method (Chen et al., 2006). A benefit of this method is that it is unsupervised; the method iteratively assigns a fuzzy label to each voxel from the sum-of-squared distance of each enhancement curve to the cluster centres. McClymont et al. (2014) propose an alternative approach using mean shift clustering and graph cuts. Such methods have been effective for breast tumour segmentation; but the much more complex anatomy of the lower abdomen makes unsupervised clustering less effective compared to a supervised method, as demonstrated in our previous comparison (Irving et al., 2014).

Alternatively, key approaches in computer vision include the use of a superpixel representation with supervised learning and graph-based analysis. Fulkerson et al. (2009) apply quick-shift clustering, and then objects of interest are detected using support vector machine classification of each superpixel, with conditional random fields (CRFs) to refine the segmentation. In medical imaging, Mahapatra et al. (2013) also use supervoxel over-segmentation for detection of Crohn’s disease in conventional 3D MRI, and Su et al. (2013) use superpixels with spectral clustering for brain tumour labelling. We also use a supervoxel and graph-based approach, but, unlike previous studies, we develop a method for 4D contrast imaging that uses: 1)

a novel perfusion supervoxel over-segmentation and classification approach
115 that exploits heterogeneous local neighbourhood characteristics; and 2) we
developed an approach that couples supervoxels to a parts-based graphical
model to efficiently include global anatomical relationships. Graph-based
methods provide a convenient basis for incorporating these global anatomical
relationships into our method.

120 2.1. Markov Random Fields and Parts-based graphical models

Markov Random Fields (MRFs), and more specifically CRFs, are an im-
portant approach for refining image segmentations, where the image is repre-
sented as a graph of connected nodes, and the nodes are generally image pix-
els – or superpixels (e.g. Fulkerson et al. (2009)). The Hammersley-Clifford
125 theorem guarantees (under certain assumptions) that the local interactions of
an MRF can attain the global optimum, but it uses global constraints implic-
itly and weakly. A number of methods have attempted to explicitly include
larger scale context information in the MRFs to capture both neighbourhood
and global interactions. Shotton et al. (2009) incorporate texture-layout po-
130 tentials into a CRF. This filter captures that layout of connected regions and
is useful for scene labelling. However, this does not capture long range rela-
tionships between sparsely labelled regions that is required to exploit organ
segmentations in large volumes. Alternatively, Kumar et al. (2005) present
a parts-based representation that includes global relationships between rigid
135 parts in a moving object using part location, shape and texture information.
Their method imposes strict constraints on the shape and relative location of
each part, which is less suitable when individual parts are highly variable, as
is the case for rectal tumours that vary both in shape and location relative
to the colorectal anatomy. More recently, Krähenbühl and Koltun (2011)
140 have introduced mean field inference on fully connected CRFs, which makes
fully connected CRFs a feasible approach for capturing long range relation-
ships in an image. This method shows improvements on previous methods
but the formulation of the pairwise term using contrast sensitive two-kernel
potentials cannot explicitly capture distance relationships between sparsely
145 labelled regions e.g. organs. The efficient inference method is also not suit-
able for a superpixel/supervoxel representation because the feature space
needs to be represented as a permutohedral lattice.

Alternatively, parts-based graphical models (PGMs), or pictorial graphi-
cal models, are an efficient method for landmark detection, and combine
150 local features of individual parts with the spatial relationships between these

parts (Fischler and Elschlager, 1973). Felzenszwalb and Huttenlocher (2005) introduce a Bayesian representation of the PGM by defining the graphical model as a graph $G = (\nu, \varepsilon)$ where ν is the set nodes and ε are the edges. The prior probability $P(L|I, \Theta)$ that the PGM has a certain configuration is
155 given by:

$$P(L|I, \Theta) \propto \left(\prod_{i=1}^n p(I|l_i, u_i) \prod_{(v_i, v_j) \in \varepsilon} p(l_i, l_j | c_{ij}) \right) \quad (1)$$

where $p(I|l_i, u_i)$ is the probability of an image I given that a part appearance parameters u_i is at location l_i and $p(l_i, l_j | c_{ij})$ is the likelihood that two parts are at locations l_i and l_j given the model parameters c_{ij} , and $\Theta = (u, c)$ are all model parameters. $l_i \in L$ are the locations of each part in the PGM.
160 Potesil et al. (2014) have applied PGMs to medical images for landmark localisation of anatomy, and have shown that the spatial distributions can be learnt from small datasets making it suitable for medical images, where large datasets of images are often not easily available.

A fully-connected CRF could be incorporated into our framework. However,
165 instead we introduce a pieces-of-parts approach – a modified PGM adapted for the task of image segmentation. Instead of each supervoxel forming a node of a fully-connected CRF, the centre points of all supervoxels act as candidates (potential locations) for the nodes (tumour, lumen and bladder) of the PGM. The marginal distribution of the tumour is then calculated
170 for these candidates using two-way belief propagation, and the parameters of the method are chosen in such a way that the marginal distribution provides an improved segmentation. This approach has a number of advantages (compared to fully connected CRFs) for segmentation problems with an expected spatial relationship between labelled regions of the image. First, only
175 a small number training cases are required as distances between parts can be pre-learned on a few cases, and inference is straightforward using two-way belief propagation on the three nodes of the graph. Most importantly, pieces-of-parts can explicitly encode distance relationships and while a fully-connected CRF such as Karahaliou et al. (2014) only use distance to weight
180 the importance of the connected nodes.

3. Methods

The key steps of the framework are discussed in Sections 3.1, 3.2 and 3.3 and are shown as steps 1, 2 and 3 of Figure 3. Perfusion-supervoxels are introduced in Section 3.1, which are used to over-segment the tumour and surrounding regions based on DCE-MRI perfusion characteristics. Features are extracted from each supervoxel and neighbourhood, and a classifier is trained to label supervoxels as tumour, lumen and bladder (Section 3.2). This supervoxel segmentation approach provides a powerful tool for DCE-MRI tumour segmentation but uses only local analysis (supervoxel and nearest neighbours) and does not take advantage of relationships between various anatomical structures in the region. In Section 3.3, the supervoxels are reformulated as pieces-of-parts in order to propagate the belief about each supervoxel via a parts-based graphical model to update the tumour probabilities based on the surrounding anatomy.

3.1. Perfusion-supervoxel extraction

Supervoxel or supervoxel segmentation methods are an effective method of reducing an image into a set of locally similar regions, which reduces the complexity and redundancy of the image, and provides a more natural set of subregions for analysis. Previously, these methods have been used for greyscale or colour images, and we extend them to perfusion images in order to over-segment the volume from locally similar enhancement patterns.

PCA is used to project a set of n-dimensional points into uncorrelated space using a linear transform and is commonly used for dimensionality reduction, and in our case is applied to the signal enhancement (SE) curves of the DCE-MRI. The transformed representation is given by $\mathbf{b} = \Phi^T(\mathbf{x} - \bar{\mathbf{x}})$, where $\mathbf{x}$ is a vector containing the SE curve values $x(t)$, $\mathbf{b}$ is the representation in uncorrelated space and Φ^T is the transposed matrix of eigenvectors of the covariance matrix of the SE curves. The original SE curve can be reconstructed from the mean curve and the weighted sum of each principal component by $\mathbf{x} \approx \bar{\mathbf{x}} + \Phi\mathbf{b}$. The standard deviation of each mode is given by $\sigma_i = \sqrt{\lambda_i}$ where λ_i are the eigenvalues of the covariance matrix. Figure 4(a) shows variation of the modes with one standard deviation from the mean for mode 1 ($\mu \pm \sigma_1$), and 3 standard deviations for modes 2 and 3 ($\mu \pm 3\sigma_i$). The enhancement curves constructed using the first one, two, and five modes




Figure 4: SLIC perfusion-supervoxel oversegmentation a) Variation of the first 3 modes of variation from the mean enhancement curve b) An axial cross section of the supervoxel oversegmentation shown at a single time point of the 4D DCE-MRI volume ($S = 350$ and $c = 0.11$). *Blue* shows the ground truth segmentation .

215 represents 97%, 99% and 99.9% of the variance, respectively².

Simple linear iterative clustering (SLIC) (Achanta et al., 2012) is a method for generating superpixels from colour or greyscale images using an adaptation of k -means clustering that uses a distance function with both colour and distance similarity terms. The method is initialised by k cluster centres that
 220 are sampled on a grid with separation $S = \sqrt{(N_v/k)}$, where N_v is the number of pixels/voxels. Pixels are then assigned to each cluster centre based on the distance function and by searching a $2S \times 2S$ neighbourhood. We have chosen SLIC because of its speed and memory efficiency when dealing with large images (Achanta et al., 2012). We extend SLIC to an arbitrary
 225 n -feature image so that 3D supervoxel regions can be extracted from the contrast enhancement patterns of 4D DCE-MRI volumes. The distance metric (D) is defined in terms of a feature distance (d_f) and a spatial distance (d_s):

²Each curve is first smoothed with a 1D Gaussian filter of $\sigma = 1$

$$d_f^2 = \frac{1}{n} \sum_k^n (b_{jk} - b_{ik})^2 \quad (2)$$

$$d_s^2 = \sum_k^3 (x_{jk} - x_{ik})^2 \quad (3)$$

As discussed earlier, principal components are used to represent the complex enhancement patterns, where b_{jk} is the k^{th} principal component of the j^{th} voxel, and (x_{j1}, x_{j2}, x_{j3}) denotes the location in 3D space of the j^{th} voxel in mm. The distance metric (D) is given by:

$$D = \sqrt{d_f + (d_s/r)} \quad (4)$$

where $r = \frac{S}{c}$ represents the average size / grid separation (S) and compactness (c) the influence of the spatial distance metric (d_s). In this study, principal components were extracted for each voxel of the volume from the SE of all voxels within the volume, $n = 3$ principal component features are used with the variation of the components are normalised between $[0, 1]$.

As discussed, a standard implementation of SLIC (Achanta et al., 2012) was adapted to use PCA features with the modified distance function (Eq. 4). As with the standard implementation, the method is initialised by placing cluster centres on a regular grid within the volume. Voxels are assigned the closest cluster centre using k-means, within a $2S \times 2S \times 2S$ distance of each cluster centre and with Eq. 4 as the distance metric. New cluster centres are computed and the process is repeated iteratively until the change in cluster centres falls below a tolerance. This method does not guarantee that the regions are contiguous, and in a post-processing step, small disconnected regions are reassigned to the closest cluster. Full implementation details of SLIC can be found in Achanta et al. (2012). Figure 4(b) shows a 2D axial slice of the 3D supervoxel over-segmentation of the tumour and surrounding region. In the following sections we call this method *perfusion-supervoxels*.

As a side note, rather than using PCA, an alternative to the outlined approach would be to apply Eq. 2 and 4 directly to the SE curves. However, extending the supervoxel representation to dynamic scans using PCA provides a compact representation of each time series, extracts meaningful variation while reducing noise, and most importantly, when normalised, the




Figure 5: Supervoxel representation with example enhancement curves for three supervoxel regions. Normalised PCA modes are used to extract supervoxels from enhancement curves to better represent curve shape. The magnitude of the curves can be thought of as being represented by a single PCA mode and, thus, the penalty is lower than if applied directly to the raw enhancement data. For example, a single blue curve may have the same magnitude as the curves of the red region but the shape is distinct.

255 supervoxel segmentation prioritises curve shape similarity over intensity similarity as shown in Figure 5. If this normalisation were not performed, the supervoxel over-segmentation would be strongly influenced by the first mode (i.e. the maximum signal), which does not capture the complex enhancement patterns.

260 3.2. Supervoxel-based classification from neighbourhood enhancement features

Having over-segmented the volume using perfusion-supervoxels, features are extracted from each supervoxel and the local neighbourhood. The tumour (as well as lumen and bladder) class probabilities are then assigned using a trained classifier.

265 Principal components of the SE curves are used as features. First, a single patient scan is used to extract representative principal components and these are then used to project the enhancement curves of all other cases. The mean and standard deviation of the first five PCA components (b_i) were used to create a feature vector (f_i) for each supervoxel (v_i). Five PCA components are used during classification, instead of three used during clustering, because
 270 a trained classifier is more robust to poor features than an unsupervised clustering method.

The supervoxel neighbourhood plays an important role in the detection




Figure 6: Supervoxel neighbourhood (2D illustration). *Orange* shows the tumour outline, with *red* representing the neighbourhood connectivity for a single supervoxel. Dashed lines illustrate the neighbours that are most representative of the orientations from which an analogue to magnitude of the gradient can be calculated for each feature (f_i).

of the tumour because it is likely to be surrounded by other heterogeneous
tumour regions; lumen, that may contain non-enhancing air or stool; and
thin rectal walls/mucosa (see Figure 2). Therefore, an additional rotationally
invariant set of features was developed to capture the enhancement of the
supervoxel neighbourhood.

For this purpose, the supervoxel neighbourhood connectivity can be rep-
resented by an adjacency graph $G = (\mathcal{V}, \mathcal{E})$ where $v_i \in \mathcal{V}$ are supervoxels,
and edges ($\mathcal{E}$) connect adjacent supervoxels (v_i, v_j). A unit vector $\hat{d}_{ij}$
is constructed between the centroid of a supervoxel of interest v_i and each
neighbour v_j to calculate the relative direction. A descriptor analogous to
magnitude of the gradient is introduced to capture neighbourhood variation
for each feature:

$$f_{\nabla i} = \sqrt{(f_{i,x+1} - f_{i,x-1})^2 + (f_{i,y+1} - f_{i,y-1})^2 + (f_{i,z+1} - f_{i,z-1})^2} \quad (5)$$

where $f_i \in \mathbf{f}_i$ is the i^{th} feature from six neighbouring supervoxels that
have a relative direction that is closest to that of the six orientations using
 $\hat{d}_{ij}$ and denoted by $[(x+1), (x-1), (y+1), (y-1), (z+1), (z-1)]$, as shown
in Figure 6.

The final supervoxel feature vector (f_i) is composed of 20 features: 10
(mean and variance of the 5 modes) from the supervoxel and 10 (magnitude
of the gradient of the 10 features) from the local neighbourhood. Each fea-
ture in f_i was normalised to $[0, 1]$ over the entire training set and f_i were
used to train a linear discriminant analysis classifier to assign a probability
of tumour, bladder and lumen to each supervoxel. We use the linear dis-

criminant analysis to limit the amount of parameter selection required for small training sets. These class probabilities were used as the unaries for the parts-based model.

3.3. *Pieces-of-parts to improve tumour delineation using spatial context*

300 The supervoxel classification method that is described in Sections 3.1 and 3.2 is used to detect and segment rectal tumours. These perfusion-supervoxels over-segment the tumour so that each tumour is accurately represented as a union of a number of supervoxels; and there are inevitably “false positives” (supervoxel components that are labelled as part of the colorectal tumour but are at some distance from the colorectum). The pieces-of-parts 305 method addresses these issues by using the classifier class probabilities for each supervoxel in conjunction with the anatomical relationships in the scan to improve the segmentation.

Previously, parts-based representations have been used for detecting points 310 of interest or, alternatively, a bounding box surrounding an object of interest (Felzenszwalb and Huttenlocher, 2005; Felzenszwalb et al., 2010; Potesil et al., 2014). We extend this approach to a segmentation problem by exploiting the relationship between long range parts, but using this information to improve segmentation at a local supervoxel level, which we call a *pieces-of-* 315 *parts* segmentation method. The tumour, lumen and bladder are defined as parts (nodes of the graph) while supervoxels belonging to each part are candidate locations of the parts. This allows the global spatial relationships to be incorporated into the supervoxel segmentation.

Eq. 1 can be reformulated into standard notation for a pairwise graph by 320 changing $\psi_p(x_p) = p(I, l_i, u_i)$ and $\psi_{p,q}(x_p, x_q) = p(l_i, l_j | c_{ij})$ (Murphy, 2012), which leads to:

$$p(x|v) = \frac{1}{Z(v)} \prod_{p \in \mathcal{V}} \psi_p(x_p) \prod_{(p,q) \in \mathcal{E}} \psi_{p,q}(x_p, x_q) \quad (6)$$

where $\psi_p(x_p)$ is the local evidence of node p at position x_p , $\psi_{p,q}(x_p, x_q)$ is the pairwise potential between nodes at x_p and x_q , and $Z(v)$ is the local normalisation constant. In the case of undirected edges the representation 325 is the same for both PGMs and pairwise MRFs (or CRFS). However, for PGMs, the pairwise potential represents relationships between labelled parts of an image. In our case, the nodes represent the parts of our tumour model, including tumour, lumen and bladder as shown in Figure 7, and the supervoxels are candidates that may belong to one of the parts.




Figure 7: Diagram illustrating the parts model where the tumour is defined as root and the lumen and bladder are children. Each part consists of pieces represented by supervoxels.

330 For inference, belief propagation (BP) is used for our *pieces-of-parts* formulation (Section 3.3.1). BP is used to update the tumour (root) probabilities (Section 3.3.2), and lumen and bladder (child) probabilities (Section 3.3.3) in the tree structure. Generation of the spatial prior distributions in this model are discussed in Section 3.3.4. The method makes no assumptions about the structure of each part involved other than the learnt spatial distributions, but in Section 3.3.5 some problem-specific modifications are included.

3.3.1. Formulation of belief propagation for pieces-of-parts

340 The graph is a tree-structured PGM except that all supervoxels act as candidates for all parts. The first step of the BP collects the evidence from all candidate supervoxels that they belong to a part. This evidence from the candidates of each part is used to form a message that is propagated to the root, which is used to calculate the marginal distribution (or belief) $bel_r(s_r)$ for each root candidate s_r , with location $l(s_r)$ as shown in Figure 8. 345 The evidence is then distributed to the child parts and used to update the hypothesis of each candidate belonging to each part. This is similar to the matching algorithm of (Felzenszwalb and Huttenlocher, 2005) but instead of just determining the most likely candidate we calculate the marginal distribution for all candidates. Therefore, the probability of each supervoxel belonging to the tumour is determined by: 1) the unary potential of s_r being tumour (based on the classifier from Section 3.2), and 2) the beliefs about the likelihood of every other supervoxel in the volume belonging to a particular part, given s_r is tumour. 350




Figure 8: The hierarchy of the supervoxel pictorial graphical model. The tumour root (r_1) has child nodes (c_1, c_2) representing the lumen and bladder and these parts are informed by the beliefs of each candidate supervoxel in the volume ($s_1, s_2, \dots, s_n$).

3.3.2. Collect evidence phase

355 In belief propagation, the belief of the root node, $bel_r(s_r)$, at supervoxel s_r , is given by the product of the local evidence, $\psi_r(s_r)$, and messages about the beliefs of the child nodes $m_{c \rightarrow r}^-(s_r)$ such that:

$$bel_r(s_r) = \frac{1}{Z_r} \psi_r(s_r) \prod_{c \in ch(r)} m_{c \rightarrow r}^-(s_r) \quad (7)$$

360 In our case, $ch(r) = \{\text{lumen, bladder}\}$, and s_r are not a sparse set of candidates but all supervoxels. The message $m_{c \rightarrow r}^-$ from each child candidate to the root is a weighted sum of the belief about each s_c being a piece of a part given its position of $l(s_c)$, the root at position $l(s_r)$, and a potential of $\psi_c(s_c)$ of belonging to class c :

$$m_{c \rightarrow r}^- = \sum_{s_c \in \hat{S}} \psi_{cr}(s_c, s_r) \psi_c(s_c) \quad (8)$$

365 where $\hat{S} = \{\forall s \in S \wedge c \neq r\}$ are all supervoxels, excluding the root candidate. $\psi_{cr}(s_c, s_r)$ is the probability that a root (tumour) candidate s_r has distance $\|l(s_r) - l(s_c)\|$ from a bladder or lumen candidate (see Fig. 9).

370 This step acts to collect evidence from all supervoxels based on the likelihood of being a component of child c and their location relative to the root supervoxel candidate (s_r). $bel_r(s_r)$ can be interpreted as a belief that each supervoxel $s_r \in S$ of belonging to tumour, and can be used to create a tumour segmentation using an appropriate threshold. Figure 10 shows the composition of the messages from the child parts $m_{c \rightarrow r}^-(s_r)$ and the local evidence $\psi_r(s_r)$ to form $bel_r(s_r)$.




Figure 10: Colormaps representing the potential of the bladder and lumen for each supervoxel (transparent is < 0.01). Belief propagation is used to form $\prod_{c \in ch(r)} m_{c \rightarrow r}^-(s_r)$ for each supervoxel, which is combined with the tumour unary $\psi_r(s_r)$ to create a pieces-of-parts belief for the tumour location. This is compared to the ground truth (*light blue*). Some bladder and lumen supervoxel probabilities might be incorrect but the majority vote leads to good tumour constraint.

385 Note that the lumen is in a similar location as the tumour but provides a strong constraint.

3.3.5. Problem specific constraints

Our pieces-of-parts method can in principle be applied to any supervoxel segmentation problem that has a number of spatially related regions. However, in the case of DCE-MRI rectal tumour segmentation, we only keep the largest connected region after segmentation in order to exclude some smaller misclassifications. The bladder part was also constrained to the upper wall using a learnt part prior.




Figure 11: Normal distributions of the relative spatial locations of the child components: the bladder (top) and lumen (bottom) for the xy axis (a) and yz axis (b)

4. Experiments and Results

395 This section describes the application of our method to a rectal DCE-MRI dataset. The dataset is described in Section 4.1, with the experimental set up of the evaluation outlined in Section 4.2 and Results presented in Section 4.3.

4.1. Materials

400 Our method was evaluated on a rectal DCE-MRI dataset acquired as part of the Rectal Imaging Trial, Churchill Hospital, Oxford, UK between 2007-2014. T2-weighted small field-of-view anatomical MRI scans and T1-weighted DCE-MRI scans were acquired for 23 patients with stage 3 or stage 4 rectal adenocarcinomas. A 1.5T GE scanner was used with a gradient echo, fat-suppressed sequence (LAVA) (TR=4.5 ms, TE=2.2 ms and flip angle 12°). Patients then underwent downstaging chemo-radiotherapy. 405 MultihanceTM contrast agent was used, and images were acquired at approximately 12 sec. intervals over 20 to 25 successive periods. The DCE-MRI voxel resolution was $0.78 \times 0.78 \times 2.0$ mm. T2w (resolution= $0.39 \times 0.39 \times 3.30$ mm, TR=14ms, TE=12ms, flip angle 40°) are acquired in the axial-oblique plane, perpendicular to the long axis of the tumour to reduce partial volume effects.

RHYTHM, a second smaller trial was also included to further demonstrate generalisability of the method. Scans were acquired in the same conditions as the first study at the Churchill Hospital, Oxford, and to date 415 includes 4 cases with DCE-MRI scans.

Manual delineations of the tumour were used as the “gold standard” from which to train and evaluate our approach. DCE-MRI rectal tumour delineation is very challenging, even for the trained observer, and instead tumours were annotated on high resolution T2w anatomical scans, registered to the DCE-MRI, and further corrected to remove potential misregistration and rater error on the T2w, as follows:

1. Rectal tumours were delineated on T2w scans by a trained radiologist.
2. T2w scans were aligned and registered to the mid-volume of the DCE-MRI to correct for minor abdominal motion between scans using (Heinrich et al., 2012).
3. The transformation was then applied to the delineations.
4. Delineations were further corrected on the DCE-MRI volume by the expert using the T2w MRI as a reference.

Fig. 12 illustrates Steps 1-3, and also highlights the inter-rater variability, which sets a limit on the evaluation of our method against the radiologist “gold standard”. Inter-rater variability was calculated by asking two additional experts to annotate 10 cases, and comparing the results to the original annotator. The Dice similarity coefficient (DSC) were 0.73 ± 0.13 and 0.77 ± 0.10 (Irving et al., 2014; Franklin et al., 2014), where $DSC \left(\frac{2|X \cap Y|}{|X| + |Y|} \right)$ provides a measure of the region overlap between two sets of labels.

Fig. 13 shows the manual correction in Step 4. Fig. 13(b) is an example where relabelling is particularly important because a region that had the appearance of tumour in the T2w scan, was found to be rectal wall from the DCE-MRI scan. The average DSC overlap between the registered and relabelled tumour volume for all cases was 0.91 ± 0.06 (with a minimum of 0.71 shown in Figure 13b). In addition to the expert tumour annotations, bladder and lumen labels were included for the pieces-of-parts training, as shown in Fig. 14. This section further highlights the challenges of manual delineation and the need for a automated and consistent approach.

4.2. Experimental setup

4.2.1. Preprocessing

Region-of-interest: Otsu thresholding is used to find a bounding box around the foreground of the MRI volume (Otsu, 1979), and hard thresholds are applied as a simple preprocessing step to reduce the size of the region that is analysed by the framework. 1/3 of the width was removed from each




Figure 12: Colorectal MRI scans (and cropped regions) with annotations to illustrate the inter-rater variability and registration method a) T2w axial-oblique MRI with tumour delineations from two experts (*blue* and *orange*) b) The corresponding DCE-MRI axial slice with registered annotations

side of the bounding boxing, $1/4$ of the height from the top and $1/8$ of the height from the bottom of the original region. Figure 1 has already shown an example of the reduced ROI.

Contrast-enhancement normalisation: Contrast injection time relative to the start of the image acquisition has some variation between scans and the steepest gradient in contrast enhancement of the entire volume (contrast in the arteries) was used to detect injection time. SE was calculated by $SE = \left(\frac{\text{signal} - \text{baseline}}{\text{baseline}} \right)$, where the baseline is the mean intensity of each voxel for the time points before contrast injection. The curves were further normalised by the 80th percentile of the maximum SE of the entire volume, which was designed to exclude the influence of blood vessels or late enhancement of the bladder in the normalisation. As there is also variation in the temporal resolution for some scans, linear interpolation was used to resample these scans to 12s intervals.

Motion correction: We use an existing non-rigid image registration framework based on the Diffeomorphic Demons modified with Normalized Gradient Fields (NGF) (Haber and Modersitzki, 2007; Hodneland et al., 2014); details of the current implementation can be found in (Papiez et al., 2014). The NGF is the contrast-invariant representation, which is used dur-




Figure 13: Two examples (a and b) showing the T2w delineation (left), transformed mask (red) and the corrected mask (blue) on the DCE-MRI volume (right). The second example shows a very large local change to the tumour shape between scans, which requires relabelling of the registered volume. Blue arrows show the region that was originally mislabelled.

ing estimation of the deformation vector field to handle intensity changes caused by contrast uptake, and has been previously shown to be suitable in DCE-MRI (Hodneland et al., 2014). The maximum number of iterations for each level is fixed to 25, and at each resolution level the standard deviation of the Gaussian smoothing was 2.8, 1.4 and 0.7 mm. These parameters were based on previous DCE-MRI datasets and are not changed for colorectal DCE-MRI. As expected, motion correction makes the segmentation more robust – particularly at the edges of the tumour. However, the performance was poorer for some cases using just the perfusion-supervoxel classifier compared to previously (Irving et al., 2014). This is because certain motion characteristics appear to have been learnt during classification. Therefore, global constraints, using pieces-of-parts, are even more important after motion correction because features from motion of the colon are excluded from the classification.

Training mask pre-processing: Expert delineations provide a voxel-wise ground truth. We create a processed ground truth for training, which only has a single label assigned to each supervoxel. Supervoxels at segmentation borders might only partially overlap the expert segmentation (see




Figure 14: Cross section and 3D rendering of the multilabel region of interest based on the T2w axial-oblique MRI. Green shows the bladder region, red the tumour, blue the lumen and magenta the lumen wall.

Figure 4b). Therefore, to generate a single label, supervoxels with $\geq 50\%$ tumour labels are assigned as tumour. These are in turn used to reclassify
 490 supervoxels with an overlap $\geq 10\%$ as either tumour or background based on their similarity to other tumour regions in the volume using a linear classifier and mean PCA features, and provides a single label to each supervoxel for training.

4.2.2. Post-processing

495 The supervoxel and pieces-of-parts methods both assign a tumour probability to each supervoxel. These probability maps are thresholded to create a tumour segmentation. After creating this segmentation, only the largest connected region is included, and smaller disconnected regions are removed.

500 4.2.3. Parameter selection

Given the limited data currently available for DCE-MRI, this method is designed to be effective and efficient for both large and small datasets. Therefore, we have attempted to limit the number of tunable parameters. The first four patients were used to find optimal parameters for the supervoxel method. The perfusion-supervoxel step requires two parameters (size
 505 S_n and compactness c). Note that the number of grid points (n_{grid}) and, therefore, S is derived from the supervoxel size by $n_{grid} = N_v/S_n$. The linear discriminant analysis is as a parameter free classifier, where the output can also be given as a probability with optimal separating threshold (T_s) of 0.5.
 510 T_s was verified on the four cases. Similarly, the *pieces-of-parts* classifier only requires a single threshold parameter T_p (equal weighting between the unary

perfusion-supervoxels:	$c = 0.05$
	$S_n = 350$
	$T_s = 0.5$
pieces-of-parts:	$T_p = 0.15$

Table 1: Algorithm parameters




Figure 15: DSC as a function of the threshold for cases used for parameter selection (pieces-of-parts (blue) and supervoxels (green)). A threshold is chosen for both the supervoxels (T_s) and pieces-of-parts (T_p).

and pairwise terms was used). T_p was also chosen from four cases, but includes a poorly performing case (Case 6), as shown in Figure 15. Parameters are shown in Table 1. These parameters are used for the key steps in the method. Other parameters, such as registration parameters, are discussed earlier. These parameters were also found to be fairly robust to the parameter choice. Supervoxel size (S_n), the mean number of voxels per supervoxel, in particular shows negligible change in accuracy in the range of 100 and 900 voxels.

4.3. Results

The method was evaluated using leave-one-patient-out cross-validation on the DCE-MRI rectal tumour dataset. The results are presented from both a detection perspective using a receiver operating characteristic (ROC) and a segmentation perspective using the Dice similarity coefficient (DSC).

Examples of the classifier output and final segmentations are shown in Figure 16. Case 09 is an example of when the *pieces-of-parts* method considerably improves the *perfusion-supervoxel* segmentation while Case 05 is an example of an already good segmentation that is barely improved. Case 06 shows an example of a case which is misclassified by the original *perfusion-supervoxel* algorithm but including *pieces-of-parts* leads to a good segmentation because of the spatial priors; at temporal resolutions of 12 sec., large

vessels may have similar characteristics to enhancing tumours, leading to the original misclassification. The DSC of this case is roughly equivalent to the median DSC and so is a representative example of the capabilities of the algorithm.




Figure 16: 2D cross sections of the 3D automatic segmentations of rectal DCE-MRI cases. Each row shows outputs of the framework and the final segmentation. (a) Unlabelled axial and sagittal cross sections through the DCE-MRI volume after contrast enhancement; (b) using perfusion-supervoxels, each supervoxel of an unseen case is assigned a probability of being tumour (shown by the color overlay where less than 0.01 is transparent) and then (c) improved using the pieces-of-parts graphical model. This is used to generate a segmentation by thresholding the probability maps (using thresholds T_s and T_p from Sec. 4.2.3) and removing small disconnected regions. The final segmentations are shown in d) where the shaded region is derived from the perfusion-supervoxels probability in (b) and the outline is derived from the final pieces-of-parts probability in (c), and compared to (e), which is the ground truth expert annotation.

4.3.1. Experiment 1: Tumour detection

The classifier returns a probability for each supervoxel, and treating this problem as a voxelwise detection problem, a ROC curve can be generated. Every voxel is labelled as tumour or background, and compared to the ground truth using sensitivity ($\frac{TP}{TP+FN}$) and specificity ($\frac{TN}{TN+FP}$), generated over a range of thresholds to obtain ROCs for each case and the mean ROC with 95% confidence intervals. As can be seen in Figure 17, this method achieves a high accuracy (AUC = 0.97), and improved on the first step of this method (AUC = 0.94). These scores show that the tumour can be accurately detected with minimal false positive regions.




Figure 17: Mean ROC curves with 95% confidence intervals for the perfusion-supervoxel classification (blue) and the entire pieces-of-parts method (red). Taking the mean for all cases at each threshold. Pieces-of-parts ROC curves are also shown for the individual cases (gray).

4.3.2. Experiment 2: Tumour segmentation

Considering the results from the segmentation perspective, we want to provide an alternative segmentation approach to manual delineation of the rectal tumour by experts. DSC was used to evaluate the accuracy of the segmentation. Unlike the AUC score, DSC does not account for the background being correctly labelled.

Figure 18 shows the DSC for each case³ using the perfusion-supervoxel

³Original case numbers from the trial are used in this study. However, not all patients

classifier (T_s threshold) and the pieces-of-parts extension (T_p threshold). The
 555 supervoxel approach without postprocessing (T_p threshold) is also included
 to illustrate the connection between perfusion-supervoxels and pieces-of-parts
 method. . While well segmented cases give similar results, poorly segmented
 or undetected cases are on average considerably improved using pieces-of-
 parts.




Figure 18: DSC compared to the manually annotated ground truth for each cases using the perfusion-supervoxel classification (light green), the perfusion-supervoxel classification without postprocessing (grey), and the final pieces-of-parts classification (blue). The two failed cases are 12 and 23, which failed due to the limited number of female training cases in the dataset. Parameter optimisation was performed on Cases 1, 3, 4, 5, 6 (see Section 4.2.3)

The median DSC for the entire dataset was 0.60 for the initial perfusion-
 560 supervoxel classification, and slightly improved to 0.63 for the pieces-of-
 parts. The number of correct detections (chosen as $DSC > 20\%$) is shown in
 Figure 19 a) which increases from 15 to 21 out of 23. Box plots of the DSC
 of the correctly detected cases are shown in b). This clearly shows that the
 pieces-of-parts achieves a similar segmentation accuracy while segmenting
 565 many more cases. An influencing factor is the post-processing (Sec. 4.2.2)

underwent imaging, leading to missing case numbers.




Figure 19: a) Number of cases correctly detected b) Boxplots DSC for all detected cases of each category (note that $n = 15$ for the perfusion-supervoxels and $n = 21$ for the pieces-of-parts method). The red bar shows the median value.

of the perfusion-supervoxel method, which improves the DSC at the expense of missed cases. In this method, only the largest connected region is kept, which can lead to misclassification (such as Case 06, Figure 16) even though a response was also generated for the tumour region.

570 A statistical comparison was made using the Wilcoxon signed-rank test for non-parametric comparison of paired samples. Assuming significance at $p = 0.05$, the pieces-of-parts method was significantly better than both the supervoxel method without post-processing ($p = 0.00016$) and the full supervoxel method ($p = 0.028$).

575 The framework presented in this paper is promising given the challenging nature of colorectal segmentation with an expert inter-rater variability of 0.73 ± 0.13 and 0.77 ± 0.10 , which is essentially an upper limit on the accuracy (Section 4.1) and is shown in Fig. 20. Furthermore, the experts both annotated on T2w images (making the experts more likely to agree),
580 while this method detects the tumour directly from the DCE-MRI volume. Nevertheless, the automated approach is close to the inter-rater variability.

4.3.3. Experiment 3: Male vs female anatomy

The framework we have presented is designed to be efficient for small training sets and has been found to be robust to patient variation. However,
585 the dataset includes both male and female patients, and because of anatomical differences it would be useful to develop separate gender-specific




Figure 20: Comparison of the automated DCE-MRI segmentation method to inter-rater variability on the T2w anatomical MRI

models. Female cases are 9, 12, 17, 19, 23, 26, 30 (Fig. 18), and include the poorest performing segmentations, including both failed cases (12 and 23, which are highlighted in the figure). Unfortunately, the available data contains only a small number of female patients (16 males and 7 females), making separate training challenging. Training the female cohort separately using an extended pieces-of-parts model that adds the uterus to the model shows potential, but ultimately more cases are required. Fig. 21 shows the DSC for female and male patients when trained together, and when trained using a separate model. While some improvement to the female cases are obtained by separate training, male cases are unaffected.

4.3.4. Experiment 4: Tumour segmentation for the RHYTHM trial

There are only a limited number of studies available that study rectal tumours using DCE-MRI and so, given the limited data, LOOCV was used to validate this study, where each case is unseen and the method is trained on the remaining cases. However, to further demonstrate the generalisability of this method, we segmented the four DCE-MRI scans that have recently been made available to us from a new independent trial. The method was applied blindly with all training and parameter choices based on the previous study. Figure 22 shows the results. The pieces-of-parts method achieved good results with a median DSC of 0.71. Case 1 is an example of a well defined tumour that is missed by the initial supervoxel method because a larger region is also detected as tumour, but pieces-of-parts provides the global context so that the tumour is still segmented accurately. It is interesting to note that the DSC is higher for this study with 0.71 than the previous study




Figure 21: DSC for gender specific models. (a) Examples of male (Case 03 and Case 04, left) and female (Case 12 and Case 26, right) DCE-MRI cross sections. Case 12 has uterine fibroids present in the scan that add further complexity to the region. (b) The box plots show results for both male and female using the general model and separate gender specific models. A larger cohort is required to fully evaluate gender specific models.

with 0.63. This appears to be due to less patient motion in the scans.




Figure 22: DSC for the RHYTHM trial and Case 1 showing a) the perfused tumour b) the manual delineation c) the perfusion-supervoxel probabilities and d) the final pieces-of-parts probabilities. Probability of zero is transparent.

4.3.5. Implementation

The framework was implemented in Python with numerically intensive steps optimised in C++ or Cython, and tested on a 6-core Intel Xeon system running Ubuntu 14.04 (x64). The mean time for a single case was 5.8 ± 4.0 minutes, where Table 2 shows the time and computational complexity for each step. These times are for a single processor, but multiprocessing was used when running a batch of cases, leading to an approximately five-fold increase in speed over the reported time. *Pieces-of-parts* barely adds any additional time to the perfusion-supervoxel classifier.

Step	Time (s)	Complexity
1. Supervoxel extraction	$160 \pm 86\text{s}$	$O(N_v)$
2. Feature extraction	$162 \pm 142\text{ s}$	$O(N_s S)$
3. Supervoxel classification	$3.9 \pm 1.2\text{s}$	$O(N_s)$
4. Pieces-of-parts:		
4.1 Training	$3.24 \pm 0.06\text{s}$	
4.2 Testing	$13.2 \pm 9.7\text{s}$	$O(N_s^2 p)$

Table 2: The mean time and approximate computational complexity of each step for a DCE-MRI volume. N_v are the number of voxels, N_s are the number of supervoxels ($N_s \ll N_v$), S is mean number of voxels in a supervoxel and p are the number of parts.

5. Discussion

We make two key contributions in this paper: *perfusion-supervoxels* and the *pieces-of-parts* method. This is applied to the challenging and previously unexplored area of automated rectal tumour segmentation from DCE-MRI scans. Accurate and consistent tumour segmentation is very important in DCE-MRI as it impacts any modelling and analysis that is performed, but is a challenging and time consuming task to perform manually due to the 4D contrast-varying nature of the image. The results are very promising with 21 of the 23 cases detected correctly, an AUC of 0.97 for the voxelwise labelling of tumour and background, and a median DSC score of 0.63. A DSC of 0.71 is obtained for a second trial to illustrate the generalisability of the method.

While the methods are applied to the problem of rectal tumour segmentation from DCE-MRI, there is a broad range of other potential applications. *Perfusion-supervoxels* are potentially applicable to any dynamic contrast enhanced sequence, while *pieces-of-parts* can be effective for adding

global constraints to superpixel/supervoxel over-segmentations. Pieces-of-parts method is suitable for images or volumes where there is an expected spatial relationship between the parts (regions of the image to be labelled) – such as organs. The method is also better suited to problems where the parts are small relative to the entire volume so as to provide meaningful spatial constraints, and, therefore, may not be optimal for scene labelling problems, where every voxel is assigned a label. In this case a fully connected CRF would be more appropriate. In future work it would be interesting to compare the trade-offs in more detail between pieces-of-parts and a fully connected CRF for variation in the type of scene. Pieces-of-parts also does not place any shape priors on the individual parts, which is beneficial for a tumour segmentation where shape is highly variable. For segmentation of well defined objects, a shape prior could be included into the segmentation.

This method applies normalisation and signal decomposition steps in the learning algorithm but does not explicitly convert the signal to Gadolinium concentration. We found that the noise introduced using a non-linear conversion to concentration did not improve segmentation accuracy.

As discussed previously, development of a ground truth segmentation on which to evaluate this method is challenging, due to the difficulty of delineating tumours in rectal DCE-MRI. This required delineations on T2w anatomical MRI scans, which were then registered to the DCE-MRI, and the delineations were further adjusted based on both aligned scans. This justifies the need for a automated and consistent approach to DCE-MRI rectal tumour segmentation. Even on the T2w anatomical images the mean DSC score was 0.73 ± 0.13 and 0.77 ± 0.10 between two additional experts and the main observer, which sets an upper limit on the evaluation of our method.

There are a number of potential extensions to this framework. A challenge is the performance on female patients due the limited female patient training data and anatomical differences. Developing gender specific models shows potential but a larger training dataset is required. This method is intended for automatic segmentation but would also be appropriate in a semi-automatic system to allow clinicians to adjust the threshold and weight parameters in order to quickly modify the segmentation.

670 Acknowledgements

This research is supported by the CRUK/EPSRC Oxford Cancer Imaging Centre.

References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., 2012. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Trans. Pattern Anal. Mach. Intell.* 34, 2274–2281.
- 675
- Brown, G., Daniels, I.R., Richardson, C., Revell, P., Peppercorn, D., Bourne, M., 2005. Techniques and trouble-shooting in high spatial resolution thin slice MRI for rectal cancer. *Brit. J. Radiol.* 78, 245–251.
- Chen, W., Giger, M.L., Bick, U., 2006. A fuzzy c-means (FCM)-based approach for computerized segmentation of breast lesions in dynamic contrast-enhanced MR images. *Academic Radiol.* 13, 63–72.
- 680
- Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D., 2010. Object detection with discriminatively trained part-based models. *IEEE Trans. Pattern Anal. Mach. Intell.* 32, 1627–1645.
- 685
- Felzenszwalb, P.F., Huttenlocher, D.P., 2005. Pictorial Structures for Object Recognition. *Int. J. Comput. Vision* 61, 55–79.
- Fischler, M.A., Elschlager, R.A., 1973. The representation and matching of pictorial structures. *IEEE Trans. Comput.* 22, 67–92.
- Franklin, J., Irving, B., Betts, M., Pureza, A., Brady, M., Schnabel, J., Gleeson, F., Anderson, E., 2014. Impact of interobserver variability on dcmri-derived pharmacokinetic parameters in patients with locally-advanced rectal cancer, in: *Proc. Radiological Society of North America 2014 Scientific Assembly and Annual Meeting (RSNA)*, Chicago IL.
- 690
- Fulkerson, B., Vedaldi, A., Soatto, S., 2009. Class segmentation and object localization with superpixel neighborhoods, in: *IEEE International Conference on Computer Vision (ICCV)*, pp. 670–677.
- 695

- Goh, V., Halligan, S., Gharpuray, A., Wellsted, D., Sundin, J., Bartram, C.I., 2008. Quantitative Assessment of Colorectal Cancer Tumor Vascular Parameters by Using Perfusion CT: Influence of Tumor Region of Interest. *Radiology* 247, 726–732.
- Goh, V., Padhani, A.R., Rasheed, S., 2007. Functional imaging of colorectal cancer angiogenesis. *Lancet Oncol.* 8, 245 – 255.
- Gollub, M., Gultekin, D., Akin, O., Do, R., Fuqua, J.L., I., Gonen, M., Kuk, D., Weiser, M., Saltz, L., Schrag, D., Goodman, K., Paty, P., Guillem, J., Nash, G., Temple, L., Shia, J., Schwartz, L., 2012. Dynamic contrast enhanced-MRI for the detection of pathological complete response to neoadjuvant chemotherapy for locally advanced rectal cancer. *Eur. Radiol.* 22, 821–831.
- Haber, E., Modersitzki, J., 2007. Intensity gradient based registration and fusion of multi-modal images. *Method. Inform. Med.* 46, 292–299.
- Hamy, V., Dikaïos, N., Punwani, S., Melbourne, A., Latifoltojar, A., Makanyanga, J., Chouhan, M., Helbren, E., Menys, A., Taylor, S., Atkinson, D., 2014. Respiratory motion correction in dynamic MRI using robust data decomposition registration - Application to DCE-MRI. *Med. Image Anal.* 18, 301–13.
- Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, M., Schnabel, J.A., 2012. MIND: modality independent neighbourhood descriptor for multi-modal deformable registration. *Med. Image Anal.* 16, 1423–35.
- Hodneland, E., Hanson, E.A., Lundervold, A., Modersitzki, J., Eikefjord, E., Munthe-kaas, A.Z., 2014. Segmentation-Driven Image Registration-Application to 4D DCE-MRI Recordings of the Moving Kidneys. *IEEE Trans. Image Process.* 23, 2392–2404.
- Irving, B., Cifor, A., Papiez, B., Franklin, J., Anderson, E.M., Brady, M., Schnabel, J.A., 2014. Automated Colorectal Tumour Segmentation in DCE-MRI Using Supervoxel Neighbourhood Contrast Characteristics, in: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (Eds.), *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*.

- 730 Springer International Publishing. volume 8673 of *Lecture Notes in Computer Science*, pp. 609–616.
- Irving, B., Tanner, L., Enescu, M., Bhushan, M., Hill, E., Franklin, J., Anderson, E., Sharma, R., Schnabel, J., Brady, M., 2013. Personalised Estimation of the Arterial Input Function for Improved Pharmacokinetic Modelling of Colorectal Cancer Using dceMRI, in: Yoshida, H., Warfield, S., Vannier, M. (Eds.), *Abdominal Imaging. Computation and Clinical Applications*. Springer Berlin Heidelberg. volume 8198 of *Lecture Notes in Computer Science*, pp. 126–135.
- 735 Kallehauge, J.F., Tanderup, K., Duan, C., Haack, S., Pedersen, E.M., Lindgaard, J.C., Fokdal, L.U., Mohamed, S.M.I., Nielsen, T., 2014. Tracer kinetic model selection for dynamic contrast-enhanced magnetic resonance imaging of locally advanced cervical cancer. *Acta Oncol.* 53, 1064–1072.
- Karahaliou, A., Vassiou, K., Arikidis, N., Skiadopoulos, S., Kanavou, T., Costaridou, L., 2014. Assessing heterogeneity of lesion enhancement kinetics in dynamic contrast-enhanced MRI for breast cancer diagnosis. *Brit. J. Radiol.* 83, 296–309.
- 745 Krähenbühl, P., Koltun, V., 2011. Efficient inference in fully connected crfs with gaussian edge potentials, in: Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. (Eds.), *Advances in Neural Information Processing Systems 24*. Curran Associates, Inc., pp. 109–117.
- 750 Kumar, M., Ton, P., Zisserman, A., 2005. Obj cut, in: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (ICVPR)*, pp. 18–25.
- Mahapatra, D., Schuffler, P., Tielbeek, J., Makanyanga, J., Stoker, J., Taylor, S., Vos, F., Buhmann, J., 2013. Automatic Detection and Segmentation of Crohn’s Disease Tissues From Abdominal MRI. *IEEE Trans. Med. Imag.* 32, 2332–2347.
- 755 McClymont, D., Mehnert, A., Trakic, A., Kennedy, D., Crozier, S., 2014. Fully automatic lesion segmentation in breast MRI using mean-shift and graph-cuts on a region adjacency graph. *J. Magn. Reson. Imaging* 39, 795–804.
- 760

- Murphy, K.P., 2012. Machine learning: a probabilistic perspective. MIT press.
- Otsu, N., 1979. A threshold selection method from gray-level histograms. 765 IEEE Trans. Syst., Man, Cybern., Syst. 9, 62–66.
- Papiez, B.W., Heinrich, M.P., Fehrenbach, J., Risser, L., Schnabel, J.A., 2014. An implicit sliding-motion preserving regularisation via bilateral filtering for deformable image registration. Med. Image Anal. 18, 1299–1311.
- 770 Potesil, V., Kadir, T., Brady, M., 2014. Learning New Parts for Landmark Localization in Whole-Body CT Scans. IEEE Trans.Med. Imag. 33, 836–48.
- Shotton, J., Winn, J., Rother, C., Criminisi, A., 2009. Textonboost for image understanding: Multi-class object recognition and segmentation by jointly 775 modeling texture, layout, and context. Int. J. Comput. Vision 81, 2–23.
- Su, P., Yang, J., Li, H., Chi, L., Xue, Z., Wong, S., 2013. Superpixel-Based Segmentation of Glioblastoma Multiforme from Multimodal MR Images, in: Shen, L., Liu, T., Yap, P.T., Huang, H., Shen, D., Westin, C.F. (Eds.), Multimodal Brain Image Analysis. Springer International Publishing. 780 volume 8159 of *Lecture Notes in Computer Science*, pp. 74–83.
- Tofts, P.S., Brix, G., Buckley, D.L., Evelhoch, J.L., Henderson, E., Knopp, M.V., Larsson, H.B.W., Lee, T.Y., Mayr, N.A., Parker, G.J.M., Others, 1999. Estimating kinetic parameters from dynamic contrast-enhanced T1-weighted MRI of a diffusable tracer: standardized quantities and symbols. 785 J. Magn. Reson. Imaging 10, 223–232.