Technical Documentation

Automatic Knee Cartilage Segmentation using Quantitative MRI

TSBB11

Rasmus Dencker Justus Johansson-Lindkvist Simon Karlsson Maria Kastberg Emma Riedl

December 13, 2017

v.1.0



Abstract

This study aims to investigate different methods for knee cartilage segmentation from quantitative magnetic resonance images. Images were acquired using a 3 Tesla machine together with a knee coil. The data sets were in 3D with a slice thickness of three millimeters. Image information from the three weighting scans T1, T2 and PD were used and the complete algorithm operates on one image slice at a time, resulting in a specific segmentation for each slice. The algorithm is based on several segmentation methods. Initially, a bone segmentation is done by morphological and watershed segmentation. This segmentation acts as a base for a slice specific bone contour mask. This mask operates together with a cartilage mask, and describes which pixels are considered as possible cartilage pixels in each slice. The classification of each pixel, as cartilage or non cartilage, is done by a k-means classification. The classification is based on pretrained centroids, generated from the clustering on data from several data sets. The setup used in the final algorithm showed a recall of 0.8394, a precision of 0.2835 and a flscore of 0.4238. The resulting algorithm is concluded to be good enough for further development but not yet ready for clinical use. The different parts of the method are implemented in a way which makes it easy to exchange them if better alternatives are found.

Contents

1	Intr	roduction 2		
	1.1	Problem description	2	
	1.2	Objectives	2	
2	Met	chod	3	
	2.1	Watershed segmentation and morphological operations	3	
		2.1.1 Approach \ldots	4	
		2.1.1.1 Bone segmentation \ldots \ldots \ldots \ldots \ldots \ldots \ldots	4	
		2.1.1.2 Cartilage segmentation $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$	9	
	2.2	Otsu thresholding	11	
		2.2.1 Implementation	11	
	2.3	Clustering	12	
		2.3.1 K-means	12	
		2.3.2 DBSCAN	13	
	2.4	Geometric methods	14	
		2.4.1 Bone contour dilation mask	14	
		2.4.2 Positioned cartilage mask	14	
	2.5	Active contours - Snakes	15	
	2.6	The complete cartilage segmentation algorithm $\ldots \ldots \ldots \ldots \ldots \ldots$	15	
3	Exp	periments	16	
	3.1	Evaluation metrics	16	
	3.2	Watershed segmentation and morphological operations	16	
		3.2.1 Parameters	17	
		3.2.2 Qualitative results	18	
	3.3	Morphological operations - Cartilage	20	
		3.3.1 Results	20	
	3.4	Otsu thresholding	21	
		3.4.1 Results	22	
	3.5	Clustering	23	
		3.5.1 Results	23	
	3.6	Geometric methods	24	
		3.6.1 Bone contour dilation mask	25	
		3.6.1.1 Setup	25	
		3.6.1.2 Results	25	

		3.6.2	Positioned cartilage mask	25	
			3.6.2.1 Setup	26	
			3.6.2.2 Results	26	
	3.7	Active	contours - Snakes	27	
		3.7.1	Results	27	
	3.8	The co	omplete cartilage segmentation algorithm	28	
		3.8.1	Quantitative results	29	
		3.8.2	Qualitative results	30	
		3.8.3	Performance	32	
4	Disc	cussion	L	32	
	4.1	Water	shed segmentation and morphological operations	33	
		4.1.1	Problems	33	
	4.2	Otsu t	hresholding	34	
	4.3	Cluste	ring	34	
	4.4	Geome	etric methods	35	
		4.4.1	Bone contour dilation mask	35	
		4.4.2	Positioned cartilage mask	36	
	4.5	Active	contours - Snakes	36	
	4.6	The co	omplete cartilage segmentation algorithm	37	
	4.7	Future	e actions	38	
		4.7.1	Bone mask	38	
		4.7.2	Better centroids	39	
		4.7.3	DBSCAN	39	
		4.7.4	Image registration	40	
		4.7.5	Testing new combinations	40	
5	Con	clusio	n	40	
R	References 42				

Terminology

- **Magnetic Resonance Imaging** Uses the magnetic properties of the hydrogen atoms in the body of organisms (in this case humans) to generate several 2D slices by applying different magnetic fields
- $\mathbf{TR}\,$ Time of repetition
- $\mathbf{TE}\,$ Time of echo
- **PD-weighted** Proton density weighted. Enhances areas and tissues containing high amounts of protons. Created through using a long TR and short TE
- ${\bf T1\text{-}weighted}\,$ Enhances muscles and fat. Created through using a short TR and short TE
- T2-weighted Enhances fluids and fat. Created through using a long TR and long TE
- R1-weighted A R1-weighted image is the inverse of the T1-weighted image
- R2-weighted A R2-weighted image is the inverse of the T2-weighted image

1 Introduction

Magnetic Resonance Imaging, more often known as MRI, are medical images generated through the usage of the magnetic property of the atomic cores within the human body. Depending on how different magnetic fields are generated over the human body, different weightings of the resulting images may be created. The most common ones are T1, T2 and proton density weightings. MRI generates several image slices over the observed object/human body part, which may be interpolated into a 3D-volume. *Quantitative MRI* contains more information than normal MRI images. Meaningful physical or chemical variables that can be measured in physical units and compared between tissue regions and different patients as well as between different time occasions can be obtained (Pierpaoli 2010).

Cartilage is protecting bone tissues from damaging each other within the body. After a knee injury it is interesting to monitor possible changes of the cartilage within the knee, but manually observing images is expensive and time consuming. Due to this, it is of interest to find a method that automatically can locate the cartilage in MRI images.

Automatic Knee Cartilage Segmentation using quantitative MRI is a project on behalf of the company **SyntheticMR**. This project is a part of a research project that SyntheticMR is working on together with the local hospital in Linköping. In the project course *TSBB11 - Images and graphics* the aim is to research and develop methods for automatic segmentation of the cartilage in MRI images. The client has no previous working method that solves this problem.

1.1 Problem description

The project will investigate different methods capable of automatically segmenting cartilage in knee MRI images. This is problematic due to the blurry appearance and unclear contour of the cartilage, as well as the similarity in pixel values between cartilage and muscle tissues in MRI.

1.2 Objectives

Several methods will be investigated and tested, as well as a combination of these methods to try to reach an accuracy as high as possible.

The client's wishes are that the group shall examine and compare different methods and variations of steps to reach the goal of cartilage segmentation. The goal is not to achieve an

application but to research the segmentation step by step to evaluate different algorithms and their performance.

2 Method

There are currently several methods for image segmentation. Some examples are seeded region growing segmentation (Abdelsamea 2014) and histogram based segmentation (Delon *et al.* 2017). The methods chosen for the cartilage segmentation, such as k-means clustering and watershed segmentation, were chosen for their strong abilities on different kinds of segmentation. In the following sections, several segmentation methods are described in detail, with the last section describing how to combine them for a complete segmentation.

2.1 Watershed segmentation and morphological operations

The watershed segmentation method segment objects from the background and from each other by different river systems, see Figure 1. One can imagine that bright areas (higher pixel values) are heights and dark areas (lower pixel values) are lower valleys. The method finds the lines that run along the tops of ridges for a gray-scale image. The key is to change the image into another image whose catchment basins are the objects desired to identify. When the segmentation is done labels are assigned to each segment, which gives opportunities for different kinds of identification.



Figure 1: The watershed river system. The figure to the left illustrates an example of variations of the intensity, and the right how the watershed segmentation works. The pink and yellow area illustrates the catchment basins which are the objects desired to segment, where the local maxima defines the watershed ridge line, that separates the two objects from each other.

Oversegmentation (too many regions) is common when using the watershed transform. It is dependent on well defined structures and smoothness. The method used for the purpose of bone segmentation is based on a method called "Marker-controlled watershed segmentation" (S.Beucher and F.Meyer 1990). The watershed transform is used in a combination with several morphological operations where foreground objects and background locations are marked to simplify the watershed segmentation.

2.1.1 Approach

The morphological segmentation was used in two different ways; segmenting the bones to produce a bone mask and segmenting the cartilage. For the bone segmentation the marker-controlled watershed segmentation was used, which follows this basic procedure:

1. The gradient magnitude of the image is calculated and used as a segmentation function. This is an image whose dark regions are the objects to be segmented.

2. Foreground objects are marked and different kinds of morphological techniques will be used to "clean up" the image and create flat maximas inside each object.

3. Background objects are marked (pixels that are not part of any object). This is done by using thresholding operations.

4. The segmentation function from step one is modified so that it only has minimas at the foreground and background marker locations.

5. The watershed transform is computed.

2.1.1.1 Bone segmentation

The bone segmentation can be divided into two parts; tibia/femur segmentation and patella segmentation. The parameter values differs for the two parts and the patella segmentation is done after the tibia/femur segmentation.

In the list below each step of the bone segmentation is described:

- The image is changed into a gray-scale image and the contrast is modified by enhancing the local contrast to enhance the difference between the bones and the surrounding tissues. To enhance the local contrast, the minimum intensity of the edges that is to be left intact is needed, as well as the amount of enhancement that is desired.
- The gradient magnitude is calculated for usage in a later step as a segmentation

function. The result shows that the gradient is high at the edges of the objects and low inside.

- A morphological structuring element is created. The flat binary shape is used in a few different ways and the shape and size of it will give different results.
- A variety of procedures could be applied here to find the foreground marker. The operation used in this case is *opening-by-reconstruction*. A new image is reconstructed with the bones clearly highlighted. The difference between the original image and the reconstructed image can be seen in Figure 2.





(a) The PD-weighted MRI image of the knee (b) T without any processing. struct



Figure 2: The original image 2a and the morphological reconstructed image 2b. The reconstructed image has the bones highlighted in comparison to the other tissues in the knee.

• The intensity is adjusted to highlight the image.



Figure 3: The intensity is adjusted for image 2b. The bones are now even more highlighted and therefore easier to separate from the rest of the image.

- The image in Figure 3 is then converted to a binary image with a threshold value chosen with Otsu's method. The threshold value differs between the tibia/femur and the patellas.
- The binary images is then put through a shrinking operation to free the bones from surrounding tissue. After that the holes are filled in.



(a) Thresholded binary image



(b) Shrunken binary image

Figure 4: A threshold value is applied to remove everything with a pixel value beneath the threshold. The image is then converted into a binary image (4a). The objects of the binary image is then shrunken to free the bones from the surrounding tissue. If the bones is still connected to other tissue the watershed segmentation will result in labels containing more than just bones (4b).

- The distance transform is calculated from the binary image. The watershed transformation is then applied on the new image.
- A new gradient magnitude is calculated by using the watershed transformed image, the binary image and the first gradient magnitude with an imposed minima function.



Figure 5: The gradient magnitude is calculated using the watershed transformation. The image shows the resulting gradient magnitude in black on top of the original image.

• The watershed transformation is applied on the new gradient magnitude image and the bones receives labels as seen in Figure 6b.



(a) Watershed transformation



(b) Colored labels from watershed transformation



Bone identification The watershed segmentation part gives each segment a label and what remains is to identify which labels that contain bones. Three approaches were used for this purpose: location, mean intensity and size. Depending on which bone part that was intended to be identified (femur, tibia or patella), different parameters were used. These steps are combined to produce a mask of the bones for each slice.

• Location

At first the centroids were calculated for each label. These centroids were useful in the purpose of locating the bones. Since the tibia and femur bone is almost always located in the center of the knee, bounderies were put in the x-direction of where these two are likely to be located. Regions that are not within these limits will be excluded.

In the case of the patella, the boundaries in x-direction were decided to be limited to be between where the first non-zero value was found in the x-direction, starting from the left (in the original image) and the starting point (looking from the left) of the found femur or tibia. The method is dependent on knowing the direction of the knee.

The limitations of both the tibia, femur bones and the patella can be seen in Figure 7 marked as green lines.



(a) Tibia and femur bone



(b) patella

Figure 7: Boundaries for identification of tibia and femur bone (a) tibia and the patella (b).

• Area

If only the position dependent property is used, it is likely that some regions with little or no interest will be found (see for example Figure 7b) and therefore the property of area is added. This is useful to get rid of smaller or bigger regions that does not belong to the bones.

• Mean Intensity

As an extra property, the mean intensity were calculated for each label and added as an extra restriction.

These steps are combined to produce a mask of the bones for each slice.

2.1.1.2 Cartilage segmentation

For the cartilage segmentation only the R2-weighted images are used as the cartilage is easily discerned and bright in the image.

The cartilage segmentation is made by:

- First making the image into a gray-scaled one.
- Using contrast-limited adaptive histogram equalization to enhance the contrast.



(a) Original R2-weighted image



(b) Adaptive histogram equivation

Figure 8: In the R2-weighted image the cartilage is a bright edge around the bones. A adaptive histogram equalization is used to enhance the contrast to make the bright parts clearer (8b).

- Top-hat filtering the image.
- Adjusting the intensity.





Figure 9: A morphological top-hat filtering is done (9a). The intensity of filtered image is adjusted which highlights the brighter parts that contains the cartilage (9b).

• Using Otsu's thresholding method in the purpose of making a binary image.



Figure 10: The image is submitted to a treasholding that removes the darker parts of the image. The resulting image is converted into a binary image.

• Then binary image is shrunk a bit and the small areas are removed.





(a) Small areas removed and the contours shrunken.



Figure 11: The final cartilage mask (11a). The resulting mask is then outlined in green to visualize the result (11b).

2.2 Otsu thresholding

Otsu thresholding is a method which is based on the assumption that there are two kinds of pixels in an image; those that belong to the background and those that belong to the foreground. Otsu thresholding then estimate the optimal threshold to separate these two kinds of pixels. Input to the Otsu thresholding method is a grayscale or binary image, while the output is a binary image. (Patil *et al.* 2016)

2.2.1 Implementation

The implemented Otsu thresholding resulted in something different from a normal Otsu thresholding. Instead of trying to separate the two kinds of pixels a more classical thresholding was achieved, separating higher valued pixels from lower. In this case a separation removing the bones from the rest of the image was desired since in PD-weighted images bones have higher values than other tissues. To do this each row was normalized to try to even out the image. After this the normal Otsu thresholding method was used, increasing the contrasts in the image, making it binary, followed by connecting components. Thereafter the image was labeled and one or more of the labels were chosen as the ones of interest. Both the binary mask and the image with the mask applied onto it was the final result.

2.3 Clustering

Clustering is used to place data points into groups, where points within the same group, or cluster, are more similar to each other in some measurable way compared to other groups.

2.3.1 K-means

K-means clustering is a simple unsupervised machine learning algorithm. It can classify data by identifying a given number of groups, or clusters, of data points. As Matteucci (2017) explains, each classification label is represented with a centroid and each data point is associated with the closest centroid. This means the positions of the centroids are of most importance and this is also what changes every iteration of the algorithm. The position update is based on an objective function that minimizes the error distance between each data point and its corresponding centroid. The algorithm continues until none of the centroids move.

Since a higher dimensional space allows for better separation of the data, all three channels (R1, R2 and PD) where used when implementing the algorithm. This feature space can be seen in Figure 12 below. This feature space is normally normalized but could be weighted if one dimension contains more information, or in this case a possible better separation of the cartilage from other tissues. Weighting would be stretching out one of the feature axis, e.g. by multiplying it with a scalar. The feature space could also be expanded with additional dimensions which could give more information and allow for a better separation of the data points representing the cartilage.



(a) Original feature space

(b) Labelled feature space, with 15 labels

Figure 12: Original and labelled feature space, where each axis describe what value the pixel has in each channel (R1, R2 PD).

It should be noted that using the three channels in this manner only works if each pixel in one of the weighted images represents the same pixel in the images of the other weightings. It should also be noted that k-means is not necessarily an algorithm that you pretrain and later use to classify new data. This is however a possibility that can be reviewed.

Since the cartilage segmentation is supposed to be done on quantitative MR images, the k-means algorithm produces similar centroid sets on different images. Of course there are differences between patients and image slices but there are enough similarities, thanks to the absolute scale of the quantitative MRI, to use the same centroids for different images. Avoiding the centroid generation on every image in the final cartilage segmentation algorithm can also be seen as an optimization.

To use a set of centroids as a classifier, one could take the euclidean distance as a measure on which centroid is closest to each data point. To avoid the computationally heavy operation of the square root the square distance can be used instead. Each centroid represents a label and each pixel in an image lies closest to one of the centroids in the three dimensional space. The closest centroid determines the label for the pixel. In this way, centroids can be pretrained using many images, of different patients and slices. Also, different number of centroids in different centroid sets can be pretrained and used in a final segmentation method.

2.3.2 DBSCAN

Density-based spatial clustering of applications with noise (DBSCAN) is another clustering method that instead of positioning centroids in the feature space finds clusters based on density, i.e. it groups together points that have many nearby neighbours. Points that do not end up in a cluster are labelled as noise. As described in the original paper by Ester (1996), how the clusters are calculated, are determined by two parameters given to the algorithm. The first parameter, *epsilon*, determines how large the neighborhood is which will be searched for nearby points. The second parameter is the minimum number of points needed to define a cluster. The algorithm starts by selecting an arbitrary point, and then checking if there are points within the range *epsilon* from this point. If there are the minimum required points within this range, these points are defined as a cluster. A new search for nearby points will start from each point in the cluster and the cluster grows recursively. When no more points can be found, the algorithm takes a new arbitrary point and continues until all points are labelled. A visualization tool for the algorithm has been created by Harris (2015) and can be of help to understand the algorithm.

2.4 Geometric methods

Cartilage is always placed next to, and in between, bones. This gives the opportunity to use this fact to narrow down the area where the cartilage may be found. Two different methods have been defined for this; one that locate a wide area close to the bone edges, and one that uses a mask created from known cartilage patterns from other cases. Both these methods are not made to find only the cartilage, but to narrow down the area where the cartilage is searched to create a good base for other cartilage finding methods. The two methods are used together with each other, the bone contour dilation mask followed by the positioned cartilage mask.

2.4.1 Bone contour dilation mask

The contour dilation mask is a segmentation with the goal of creating a mask surrounding the cartilage and nearby tissue. The segmentation is based on the result of the bone segmentation. It uses the mask from the bone segmentation to create two new masks, a large one that is a dilated bone segmentation and a smaller one which is an eroded bone segmentation. Thereafter, the smaller mask is subtracted from the larger one, leading to a mask following the edges of the bones, both containing the inside of the bone edge and the tissues surrounding the bone. This is made large enough to make sure that all cartilage is within this area, even if that means that even more uninteresting tissue might be within the segmented area as well. This is not supposed to find only the cartilage, but to narrow down the interesting area where cartilage can be found to create a base for other methods.

The downside with this method is that it relies on a good bone segmentation. If the bone segmentation is bad, that will result in a mask which highly possibly will not contain the cartilage.

2.4.2 Positioned cartilage mask

Since medical knowledge provides information on where the cartilage is located in the knee, a mask can be derived containing all possible cartilage pixels. This can be generalized to work on all patients and slices while still removing some areas, where cartilage can be falsely segmented, when used together with the bone contour dilation mask. For example the lower part of the segmented tibia never contains any cartilage, but is within the bone contour dilation mask.

Even if the mask is big enough to cover all cartilage on any patient slice it still has to be positioned to the correct location in the knee of the specific image set, since this varies between MRI examinations. There are most likely many ways of doing this automatically but one way is to use the previously defined, and patient specific, bone contour dilation mask. Since the knee does not move noticeably during the image acquisition it is enough to position the mask according to one of the bone contour dilation masks. The best choice of mask is one made from a bone segmentation that has revealed the tibia, femur and patella in a good way. Then the cartilage mask can be position simply by locating the position where the correlation between the masks is highest, i.e. where the overlap between the masks is the biggest. To be noted is that this implementation relies on the bone contour mask which, as mentioned in turn relies on the bone segmentation. Other methods, not dependent on the bone segmentation, could make use of the cartilage mask in a way that results in a more stable overall cartilage segmentation solution.

A suggested method that does not necessarily rely on the bone segmentation could be Image Registration. As Zitová and Flusser (2003) explains, this method could be implemented in both linear and nonlinear variations and will allow any transformation including translation, rotation, shearing and scaling. This could be used to attain a motion transformation for the cartilage mask, specific to each data set, to achieve optimal positioning.

2.5 Active contours - Snakes

Active contours, also called snakes, are methods which searches for contours and edges over a number of iterations, creating base contours which shrink or grow depending on different velocities or energies. Active contours may be a good choice in case that the shape of the desired segmentation is somewhat known. The strength of active contours is that it is good at finding the contours in an otherwise noisy image. (Kass *et al.* 1988)

The Chan-Vese algorithm was implemented during this project. Chan-Vese is a method which belong to the active contour methods called *levels*. The reason for this is that level based segmentation is stronger at clusters and unclear edges than normal snakes (Getreuer 2012). Since the contour of the cartilage is very unclear, Chan-Vese is a good choice.

2.6 The complete cartilage segmentation algorithm

The different methods explained above can be combined in several ways to achieve different results. Some have already been mentioned, e.g. using the bone segmentation from the watershed method to derive a slice specific bone contour mask.

The Otsu thresholding can be used to remove a lot of unwanted pixels in an image and can be used as a preprocessing step in different situations, perhaps to optimize another method by reducing the input data. The geometric methods could be used to determine were in the image it is likely for cartilage to exist, and used to remove the other areas which could also reduce the input data. Both the watershed and k-means methods could be used to segment both cartilage and other parts of the knee than those explained above. This could lead to other combinations than those mentioned in this report. See Section 3.8 for the proposed method combination of this report.

3 Experiments

The experiments section will describe the parameters and implementation choices made along with the result for each method. If the method or method combination was used for direct cartilage detection the quantitative and qualitative result will be included as well. All the experiments were conducted on images acquired from a 3 Tesla machine together with a knee coil. The data sets were in 3D with a slice thickness of three millimeters.

3.1 Evaluation metrics

The evaluation metrics used were **recall**, **precision and f1score**. **Recall** is the fraction of correctly labelled cartilage pixels out of all true cartilage pixels. **Precision** is the fraction of correctly labelled cartilage pixels out of all pixels labelled as cartilage. **f1score** is the harmonic average of precision and recall.

3.2 Watershed segmentation and morphological operations

The steps mentioned before results in a mask for the bones in each slice as can be seen in Figure 13. This mask can be used for geometrical algorithms to find the cartilage.



Figure 13: An example of the resulting bone mask. The femur and tibia identified first. The red star represent the edge of the mask containing only the femur and tibia and is the reference point used to identify the patella.

3.2.1 Parameters

An important part of the setup is the choice of parameters. In the bone segmentation there are a few different parameters that can be modified and that differs between the different parts. The parameters that are used to modify the method is the shape and size of the morphological structuring element, the level that determines which values that should become black and which should become white in the binary image and the size that should be shrunken and dilated. All the changeable parameters is part of the morphological segmentation.

The parameters that were used in the final algorithm are:

Part	Morphological structuring	Binary threshold	Erode	Dilate
	element	level		
Tibia/	Rectangle, [15 20]	0.55	ones	ones
femur			(15, 15)	(4,4)
Patella	Rectangle, [10 15]	0.65	ones	ones
			(10, 10)	(4,4)

 Table 1: Parameter values used for the bone segmentation.

For the identification part the parameters that is not calculated but set is for the intensity and the area. The mean intensity should be between 0.5 and 0.7 and the area for tibia/femur between 5000 and 55000 and for the patella between 300 and 9000.

Many parameter values were tried and tested for the final algorithm. The chosen values were the ones that gave the best results looking at the different data sets at hand. There is no numerical proof of this as it would have demanded correctly labeled sets of the bone for all of the data sets, which was not at hand. All the test were evaluated by appearance alone.

3.2.2 Qualitative results

When executing the full function of the bone segmentation, including the parameters for tibia/femur and patella bones, as well as the different methods of identification, the results in Figures 14, 15, 16 and 17 are achieved.

The result of the bone mask function differs between patients and between slices depending on the clarity of the image and how well defined the bones are in the image. In Figure 14 the corresponding bone mask contains all three bones correctly labeled, but in Figures 15, 16 and 17 the bone mask are faulty or incomplete.





(a) PD weighted image



Figure 14: The original image and the resulting bone mask.



Figure 15: The original image and the resulting bone mask. The tibia in the mask is not completely separated from surrounding tissue.



(a) PD weighted image (b) Bone mask

Figure 16: The original image and the resulting bone mask. The bone mask is missing the femur.





Figure 17: The original image and the resulting bone mask. The algorithm gives a false label for the patella.

3.3 Morphological operations - Cartilage

The cartilage segmentation method using morphological operations was not used in the final complete algorithm but is an alternative in a future solution.

Different morphological operations were applied to the R2-weighted images, as mentioned in the Methods section. All of these operations needs input values that is up to the user to set. In this case the values and the Matlab functions were:

- adapthisteq 'NumTiles',[10 10]
- *imtophat strel*('disk',15)
- *imerode strel*('disk', 1)
- bwareaopen 100

3.3.1 Results

The resulting mask does contain more than only the cartilage and is not a complete solution to the problem by itself. In Figure 18 the resulting mask for the above mentioned input values can be seen. The mask looks like it contains most if not all the cartilage as well as much of the muscles.



(a) Resulting cartilage mask.



(b) Cartilage mask outlined in green on the original image.

Figure 18: The resulting mask as a binary image and visualized in green on top of the original image.

3.4 Otsu thresholding

Otsu thresholding is usually done through separating the foreground from the background. But after trying this, it fairly soon showed that this method did not deliver what was desired. Due to this a kind of normalization was done over the rows of the input image, where each pixel value was divided with the sum of its row, to the power of 1.5. This specific normalization was found through testing. After that the contrasts within the image was increased through the function *imadjust* followed by binarizing the new image with a global image threshold, using the function *imbinarize*. The function *bwconncomp* finds connected points within an image, where the connectivity was set to 4, meaning that each pixel will look for connections in the x- and y-direction, not diagonally. Thereafter, the image was labeled based on the connectivity and the label which corresponds to the connectivity of the pixels. This does most often divide the image into two labels, one containing the bones and muscles (and some more tissue) and the other one the rest of the tissues. Therefore, the one not containing the bones or muscle was chosen and both the mask and the mask applied on the input image was returned as the result. The only input to the created Otsu thresholding function is a gray scaled image. Since Otsu was not able to find only the cartilage, the goal was to make it a preprocessing step for other methods, removing parts of the image that do not contain any cartilage.

3.4.1 Results

First a normal Otsu thresholding was implemented, which is presented in the figure below. As can be seen, the labeled image is quite noisy and does not create any label that contain all cartilage. One label is fairly close (the one colored orange) but that one also contains muscle tissue, which has intensity values that are almost the same as cartilage, which might make future separation of the two tissues harder.



Figure 19: To the left, the input image. To the right, the labeled image. Created by separating the foreground from the background through Otsu thresholding.

Secondly, the implementation was changed into something that more applied to narrow down the area where cartilage can be found. As seen in the figure below, the cyan colored label remove most bone and muscle tissue while still containing all of the cartilage. This can then be used as a preprocessing step for other algorithms and methods, since a lot of incorrect tissue has been removed. And even though the black background (original image) is within the cyan label, this part will also be removed when the mask is applied on the input image, since that black area has zero intensity.



Figure 20: To the left, the input image. To the right, the labeled image. Created by normalizing each row followed by using a global threshold.

3.5 Clustering

Many different pretrained centroid sets were examined. Varying numbers of centroids on different images, both with and without a preprocessing step using Otsu thresholding. Also pretraining on different amounts of training data was tested. The centroid set that seemed most reliable on different image sets and slices was the five centroids created using ten images from three different patients, using different slices near the center of the knee and preprocessed using Otsu thresholding.

The centroid set was then used for each image slice, with channels R1, R2 and PD, in an implemented function that used the square distance to label each pixel. Later the cartilage is picked out using the label color of the cartilage pixels.

3.5.1 Results

The results of the k-means are based on the results of the complete cartilage segmentation algorithm described in Section 3.8 below. Different parameters for the k-means classifier in the complete cartilage algorithm was evaluated. One comparison using different number of symaps to pretrain the k-means centroids was examined. The result is shown below in Table 2.

# training symaps	Recall	Precision	F1score
10	0.6452	0.3133	0.4218
20	0.6019	0.3038	0.4038

Table 2: Mean recall and mean precision for k-means pretrained with different number of symaps.

The number of labels in the k-means classifier, and different number of labels that represented cartilage was also evaluated. Again the result was evaluated by running the complete cartilage segmentation algorithm, with differently trained centroid sets. The result can be seen in Table 3 below.

# labels	# cartilage labels	Recall	Precision	F1score
5	1	0.5641	0.2119	0.3081
5	2	0.8995	0.2673	0.4121
4	1	0.4873	0.5136	0.5001
8	2	0.8082	0.3223	0.4608
12	1	0.3951	0.5459	0.4584
12	2	0.6878	0.4576	0.5496

Table 3: Mean recall and mean precision for k-means pretrained with different number of labels and different number of labels chosen to represent cartilage.

The k-means classifier was also evaluated with different weightings of the feature space. The result can be seen in Table 4 below.

R1 weight	R2 weight	PD weight	Recall	Precision	F1score
1	1	1	0.5603	0.4995	0.5282
3	1	1	0.5776	0.2914	0.3874
1	3	1	0.5638	0.4431	0.4962
1	1	3	0.8777	0.3179	0.4667
1	2	3	0.4508	0.5154	0.4809

 Table 4: Mean recall and mean precision for k-means pretrained with different weighting of feature space.

3.6 Geometric methods

Two different geometric algorithms were created. Both of these do not find the cartilage by itself, but are preprocessing steps for further algorithms.

3.6.1 Bone contour dilation mask

The contour dilation mask is not used for direct cartilage segmentation, but for simplifying the cartilage segmentation for other methods through making the area of interest smaller.

3.6.1.1 Setup

The function has five in-parameters which controls how much the contour of the inimage should be dilated as well as if pixels within the mask close to the top or bottom should be removed. For the final results, outwards dilation is set to 60 pixels with the Matlab-function *imdilate* and erosion is set to 10 pixels with the Matlab-function *imerode*. Neither top or bottom pixels are removed. Figure 23 below shows the three steps of creating the contour mask, the eroded bone contour, the dilated bone contour and then the final mask which is the difference between the two masks.

Eroded contour







Figure 21: The different parts of the creation of the contour dilation mask.

3.6.1.2 Results

As can be seen in the figure above the final mask is covering the areas close to the three bone segments. Since cartilage is located between and close to bones, the final mask gave desired results.

3.6.2 Positioned cartilage mask

Provided by SyntheticMR were two image sets containing the correct cartilage segmentation. The location of the cartilage over all images in both sets was put together into a mask. Due to the low amount of image sets with correct cartilage segmentation, this mask was not generic enough. Due to this the mask was dilated to make sure that the cartilage of new image sets would be within the mask. The mask is shown in Figure 22 below.



Figure 22: The dilated cartilage mask.

3.6.2.1 Setup

To position the cartilage mask a function was implemented. This function takes the cartilage mask together with the bone contour mask of slice 20 of the specific image set and slides the masks over one another, storing the location of the highest correlation. A parameter to the function is the step length, i.e. number of pixels, that the cartilage mask is incremented in each iteration of the algorithm. This is optimally chosen as small as possible and was set to three.

3.6.2.2 Results

The results of different sizes of the cartilage mask there were, as expected, more mislabelled cartilage pixels when using a bigger mask. It was also obvious that when the bone segmentation failed for slice 20, the cartilage mask was positioned poorly, resulting in a bad cartilage segmentation using the complete cartilage algorithm.

3.7 Active contours - Snakes

The Chan-Vese method created by Wu (2009) was used as a framework and a simpler run-file was created to easily be able to load and use images from the project. The Chan-Vese method takes in several parameters. The most important ones are starting shapes, number of iterations and initial weight, where the weight adjust the speed of contour reshape per iteration.

3.7.1 Results

Several tests was performed with the Chan-Vese method. For visualization, a run with many small starting shapes (octagons) was used. This way, many contours, and possibly non-excising, contours should be found. As seen in the figure below, the area where cartilage normally is found have no clear contour around it, and when the Chan-Vese divide the contour into segments, the cartilage is segmented with muscle tissue and more. For this test, an initial weight of 0.1, since lower initial weights more often detect weaker contours, and 1000 iterations was chosen. The algorithm stopped after 208 iterations since no snake was changing over the last iteration.



Figure 23: Top left - Input image. Top right - Starting shapes. Bottom left - Final contours. Bottom right - Final result.

3.8 The complete cartilage segmentation algorithm

The combination of methods consists of watershed transform with morphological operations, the geometric methods and Otsu thresholding together with k-means clustering. First the centroids for the k-means clustering were pretrained using several of the given images. Then the bone segmentation through the watershed algorithm was applied to the image that would be labelled. The bone segmentation result in a bone mask that was used to create the contour dilation mask. The contour dilation mask was then combined with the cartilage mask to create the final mask which was adjacent to the bones and that contained some bone, the cartilage and other nearby tissues. This final mask was applied to the original image, which results in an image were only the areas close to bone are kept. To finally extract the cartilage, the pretrained centroids were used to label this stripped down image.

This method is heavily dependent on the bone segmentation working correctly, since everything not considered to be close to bone is removed. If the bone segmentation misses a bone the nearby cartilage could be removed by the bone mask, and if some tissue is mislabelled as bone then the surrounding tissue will be searched for cartilage which could increase the number of false positives.



Figure 24: Diagram of method steps. Input data consist of 3D matrix containing slices with PD-weightings, and a 4D matrix with all channels. The combination of methods consists of watershed transform with morphological operations to create a bone contour mask, and together with the geometric methods, Otsu thresholding and k-means clustering a final cartilage label can be picked out.

3.8.1 Quantitative results

The complete algorithm was evaluated using the parameters that yielded the highest recall and precision from the evaluation of the k-means classifier, i.e. pretrained with 10 symaps, 5 labels and a weight of three on the PD-channel. The mean result for all slices can be seen for one of the test patients in the Table 5 below.

# cartilage labels	Recall	Precision	F1score
1	0.6511	0.2776	0.3892
2	0.9044	0.2540	0.3966

Table 5: Mean recall and mean precision when evaluating the complete cartilage segmentation algorithm.

The recall, precision and f1score per slice can be seen in figure 25 and 26. In the former, one centroid was used to label cartilage and in the latter, two centroids were used.



Figure 25: Result per slice (x-axis), with 1 labels representing cartilage.



Figure 26: Result per slice (x-axis), with 2 labels representing cartilage.

Note that the first and last couple of slices contain little to no cartilage, and are not as interesting from a clinical perspective. The mean recall, precision and f1score between slices 8-22 were 0.6497, 0.3186 and 0.4275 using one label as cartilage, and 0.8394, 0.2835 and 0.4238 when using two labels as cartilage. The recall is lower here since the recall is 1 in the first slices since there are no cartilage there, which misleadingly inflates the recall and precision.

3.8.2 Qualitative results

The segmentation was also plotted in the original image together with a coloring indicating if the labelling was done correctly. This visualization can be seen in Figures 27-30.







green - correct red - missed cartilage pink - mislabelled as cartilage

Figure 27: Segmentation of slice 11 with 2 labels representing cartilage.







green - correct red - missed cartilage pink - mislabelled as cartilage

Figure 28: Segmentation of slice 25 with 2 labels representing cartilage.







green - correct red - missed cartilage pink - mislabelled as cartilage

Figure 29: Segmentation of slice 11 with 1 labels representing cartilage.



green - correct red - missed cartilage pink - mislabelled as cartilage

Figure 30: Segmentation of slice 25 with 1 labels representing cartilage.

3.8.3 Performance

Running the complete algorithm on a 576x576 symap image data set of 26 slices takes 2 minutes and 43 seconds. This using the computer and Matlab specifications in Table 6 below.

Computer	MacBook Pro
Operating system	macOS High Sierra, version 10.13.1
Processor	2,3 GHz Intel Core i5
Memory	8 GB 2133 MHz LPDDR3
Matlab version	R2017b

Table 6: Computer and Matlab specifications.

The bone segmentation takes 23 seconds, the following creation of all the masks takes another 8 seconds and the cartilage segmentation on all slices takes up the remaining 2 minutes and 12 seconds.

4 Discussion

In the discussion section each method and its experiments will be discussed separately. In addition to the discussion of the experiments and results there will also be discussion and ideas of what can be improved and what the next step would have been in the project in a future aspect.

4.1 Watershed segmentation and morphological operations

The most important thing to take in account about the bone segmentation, is that it is the preprocessing part (morphological operations) that matters in the end. That is why the method is called Marker-Controlled Watershed Segmentation. If the watershed algorithm would have been applied directly on any of the images without any processing, it would have resulted in a over segmented image that would not have given any useful information at all. This part turned out to be the most important when looking at the complete cartilage segmentation algorithm, as a mistake in the bone mask leads to missed or miss-labeled cartilage.

The method for segmenting the cartilage directly using morphological segmentation was never evaluated as a part of the complete algorithm so it is difficult to say exactly how well it worked on finding the cartilage. This method could possibly be used as a step in a complete solution but the fact that it find so much of the muscles might be a problem.

4.1.1 Problems

There are a few difficulties to consider when looking at the bone segmentation.

- The different shapes and sizes of the bones between slices.
- The size of the patella in comparison to the other bones.
- Smudged edges for the bones.

In the preprocessing steps the morphological structuring element is the part that decides which sections that will be labeled and the size of these sections. So in the middle slices where the femur and tibia is big the structuring element should also be big to make sure that each bone is segmented and labeled as one part and not multiple smaller parts. But this structuring element will not work for the earlier and later slices where these bones are much smaller. Using different structuring elements for the different slices was considered to try to match the size of the structuring element with that of the femur and tibia but decided not to. There was no reliable way to determine which slices contained bigger bones and might therefore be improved by using another structuring element. Changing the structuring elements could just as easily give the opposite result. The change from smaller to bigger bones is gradual and is not necessarily at the same slice from patient to patient. There is also a difference if the images is of the right or left leg, so this would have to be taken in to consideration as well. This also means that a separate part for the segmentation of the patella is necessary as the size of the patella is considerably smaller than the size of the femur and tibia. The shape is also very unlike that of the other bones. One of the main reasons for missing bones is that the edges of some bones are partly indiscernible from surrounding tissue and therefore hard to segment correctly. Figures 15 and 16 in the experiment part shows two different examples of how the bone segmentation can be unsuccessful. In Figure 15b the tibia has not been correctly separated from the surrounding tissue. It might be possible to fix this problem with other parameter values but that might also result in other problems in other slices. In Figure 17b the whole femur is missing in the mask. The original image (Figure 17a) shows that part of the tibia's edge is melting into the tissue to the left of the bone. This means that the segmentation won't label just the bone but a much bigger area as one part, and that part doesn't fit the criteria to be in the bone mask.

4.2 Otsu thresholding

The implementated Otsu thresholding method, which appeared to be more like a mix between normal thresholding and Otsu thresholding, got results which was not capable of segmenting the cartilage itself, but only removing some tissues which was not cartilage. Several attempts was done, trying to find optimal values for cartilage segmentation, but it fairly soon showed that this was a hard thing to achieve. Either muscle tissue and cartilage was grouped together above the threshold, or cartilage was grouped with the background tissue, due to its fade contour and fairly weak signal. This was a bad result, since separating cartilage from muscle tissue is hard, due to their close intensity in MRI, and when the cartilage got mixed into the background, the Otsu thresholding method was more or less useless. So after modifying it, the result became better. It was still not able to segment the cartilage, but it was now capable of narrowing down the area where it could be found. Due to this, it is a good processing step for further segmentation, but it's not good on its own. There are most probably better methods to use for cartilage segmentation than Otsu thresholding.

4.3 Clustering

When pretraining the centroids one might assume it is always better to use as many images as possible. Both from different patients and slices. From the experiments, this was not the case. A few things to consider should however be:

- Pretrain on different patients, but only the slices that are of most importance.
- Try a different number of images.
- Scale the different axes (R1, R2 and PD) differently depending on which contains the most information.

- More dimensions to the feature space could enhance the final performance, e.g. a geometric dimension such as distance to a certain, patient invariant, point in the image.
- A different optimal number of labels might occur on different patients and slices.

Differently trained centroids naturally results in different labelling. One of the parameters during the training is the number of images the centroids are trained on. Two centroids were trained with the same parameters but different amount of training images and then applied in the segmentation through the complete cartilage algorithm. One of the centroids was trained with 10 images, and the other with 20 images. The images were from different patients and slices. The latter centroid set had 10 additional images, i.e. 10 of the images were the same for both centroids. The mean recall and mean precision when evaluating the result from the complete cartilage algorithm with the different centroids were slightly higher for the centroid that was trained with the least images. This centroid set had a mean recall of 0.65 and a mean precision of 0.31 while the centroids trained with 20 images had 0.60 and 0.30 respectively, as mentioned in in Section 3.5.1. This shows that training a centroid set with more images does not necessarily improve the results. Similarly the results varied when using different weightings of the feature axes, where the recall increased significantly when a larger weight was put on the PD-channel of the image.

4.4 Geometric methods

Two different geometric algorithms were created. Both of these do not find the cartilage itself, but are preprocessing steps for further algorithms. They both have their advantages and disadvantages, which will be described further below.

4.4.1 Bone contour dilation mask

The goal of the bone contour dilation mask was to decrease the area of the image in which the cartilage could be found. Since cartilage is located close to bones, removing areas within and far away from bones can decrease the area of interest, while assuring no cartilage is lost. The mask that was created does fulfill this goal since it is big enough to make sure no cartilage is lost, but contains a lot less pixels than the full image. The downside is that it does very heavily rely on a good bone segmentation. If the bone segmentation is incorrect, the mask created will be around the contour of the bone segmentation, not around the correct bone contour, which will risk that cartilage gets lost outside of the mask.

The fact that cartilage in the knee appears in contact with bone, mainly in between adjacent bones, is something that should be taken advantage of. It is a well known fact and fairly easy to consider when implementing a knee cartilage segmentation algorithm. There are probably other possible implementations than the one used here, but the bone contour dilation did work well, especially when combined with the positioned cartilage mask described below.

4.4.2 Positioned cartilage mask

The bone contour mask described above will allow cartilage to be found around the entire contour of the bones. Since cartilage only appears in some areas of the mask, another mask was generated to remove bone contour pixels that cannot contain cartilage. The positioned cartilage mask does this quite well. To be noted however is that to make sure it does not miss any cartilage pixel on any patient, it is rather large. The result contains therefore a bit more mislabelled cartilage pixels than if the mask would have been a bit more slim and selective. The cartilage mask size could be adjusted according to what is most important in the current situation. In the case of this report, a higher recall was prioritized and therefore a big cartilage mask was used.

The positioning of the cartilage mask by sliding over the dilated bone contour is a simple implementation. It does not rotate to fit a data set of a slightly rotated knee in any way which could be a smaller problem. A bigger problem is that the mask is positioned according to the dilated bone contour mask of slice 20, independent of which patient. Therefore it heavily relies on a good bone segmentation of slice 20. This problem could however be solved in many intuitively simple ways. For example by matching the bone segmentations to a standardized bone template and base the cartilage positioning to the dilated bone contour mask related to the "best" bone segmentation.

Using Image Registration with a reference MRI image at a certain depth in the knee, with the cartilage mask manually positioned to reference, a new data set could use a matching slice to get a motion transformation for the cartilage mask to position, and in many other ways transform it, according to the new data set. This would allow the cartilage mask to initially be quite small and still act as a general mask that can be used for all possible data sets. This would result in a better precision but would not affect the recall, when compared to the suggested method above.

4.5 Active contours - Snakes

Active contours have been tested for the cartilage segmentation during this project, and it fairly quickly showed that it does not work well at finding the cartilage. Even though active contours are good at finding contours where noise is present, the issue is that there are so many contours, stronger than the one for cartilage, which are found instead of the desired one. Contours from bone parts, muscles, nerves, cavities, etc. are all prioritized over the contours of the cartilage. Due to this, it is very hard to use active contours for cartilage segmentation.

Active contours may still be of interest though. With the correct preprocessing, it might be possible to more clearly define the contours of the cartilage and then use active contours. This has not been investigated within this project, but active contours are something that might be worth keeping in mind. Also, testing more active contours based on the level methods is recommended, since they're overall better at detecting weak contours.

4.6 The complete cartilage segmentation algorithm

Since the pixel values of cartilage tissue aren't unique in any MRI image weighting, a geometrical distinction is required. This is possible since the cartilage appears in connection to bone tissue, between the tibia and femur and behind the patella. This is the motivation for the bone segmentation part of the complete algorithm. This information combined with another mask, a cartilage mask, creates a rather small space of the total image where the cartilage is allowed to be found. This small space does however still contain other tissues than cartilage and more segmentation steps are required. The k-means algorithm is a good choice since it does a pixel wise classification and does not require certain structures to be present.

The result using the highest performing setup from the k-means evaluation showed that it was possible to reach a high recall, 0.9, however the precision was significantly lower. The segmentation can find almost all pixels labelled as cartilage, however it also has a lot of mislabelled pixels. This can also be seen in Figures 27-30, where pink pixels represent pixels mislabelled as cartilage. In the experiment where two labels represented cartilage, the number of mislabelled pixels naturally increased which can be seen when comparing Figure 27 and 29. The former contains more pink compared to the latter, however it also finds a lot more of the cartilage. The result also clearly depends on the result of the bone segmentation, which can be seen in Figure 26 and 25 where the numbers drops to zero at slice 25. When looking at slice 25 in Figure 28 and 30 all cartilage is missed, which is due to the lower bone not being segmented correctly.

For different ways to further improve the results, see Section 4.7.

4.7 Future actions

During the project, multiple new ideas on how the algorithm could be improved developed that the group unfortunately did not have time to pursue. As it's possible that the client might want to continue to develop an algorithm for automatic cartilage segmentation, this section will focus on what might be done or tested to improve the current algorithm that is proposed in this report.

4.7.1 Bone mask

There are multiple things that could possibly improve the bone mask function. A few possible ideas are

- Using information from previous and following slices to reduce the chance for missing bones.
- Use a 3D representation of the knee instead of using slice by slice.
- Separate the segmentation of the femur and the tibia to be able to use each of the bones characteristic features.

With the current method, it is not entirely uncommon that a bone is missing in one slice but is present in the slices before and after. This means that it would be possible to in these cases fill out the bone mask with information from the previous and following bone masks. The fit would not be as good as if the bone was correctly segmented but would give a better mask than if the bone was completely absent. This problem could also be avoided by instead of looking at the slices in 2D and one by one, represent the slices in 3D as a whole. This approach could possibly be more precise but it would take a lot of work and would possibly mean a completely different segmentation method overall. It was not looked into if these methods could be implemented in 3D. When looking at the implemented method and what could be done to improve it, is that it is built for images in 2D and it might not be possible to rebuild it for 3D.

Separating the segmentation of the femur and tibia should be relatively easy to add to the current method. By doing the segmentation in several steps the chosen structuring element could be more specific to the size and shape of each bone and therefore result in better bone masks.

4.7.2 Better centroids

During the experiments, it was shown that different parameters when training the centroid set effected the final performance. A next step could be optimizing these parameters to further improve the results of the classifier. The centroid sets were effected by the number of images used during the training, the number of centroids that would be defined in a set, the scaling of the feature space and the number of labels chosen to contain cartilage. These could all be optimized which could improve the final performance. One possibility that could be explored is adding a certainty to each label, which could give the pixels a high or low certainty of being cartilage. This could then be used for pixels which are uncertain whether they contain cartilage or not. If a pixel with uncertainty is close to several pixels that have a high or low certainty of containing cartilage it could be used to decide if the pixel with uncertainty contain cartilage or not. One way to measure certainty could be proximity to the nearest centroid. Neighbouring pixels can belong to different centroids since that calculation is done in the feature space, and a mislabelled pixel could be corrected if it has low certainty while neighbouring pixels have high certainty.

4.7.3 DBSCAN

One of the next steps would be to test out the DBSCAN-algorithm. The algorithm could replace the labelling done by the k-means algorithm, or it could be used before k-means to determine the optimal number of clusters which then could be fed to the k-means algorithm. Currently the number of clusters for k-means is set by the user when pretraining the centroids, which could be automatized by DBSCAN. The problem that stopped DBSCAN to be fully implemented was that it calculates the pair-wise distance between all points in the image, which requires a large amount of memory to handle. There are different approaches to solve this problem. One would be to only check distance to nearby data points. This is not as easy as just checking the distance to adjacent pixels, since the distance is measured in the feature space. Pixels which were adjacent in the image are not necessarily adjacent in the feature space.

Another solution is to decrease the number of pixels by using superpixels. Superpixels are pixels that have been grouped together by the simple linear iterative clustering (SLIC) algorithm, and could be used to represent a cluster of pixels. The DBSCAN algorithm could be applied onto these superpixels, labelling each superpixel which could then be used to apply the same label to all pixels within the superpixel.

4.7.4 Image registration

Since the pixel values of cartilage resembles the ones of other tissues, mainly muscle tissue, any cartilage segmentation method could result in mislabelled pixels far from the general cartilage locations. A cartilage mask solves this problem. Also, a mask that is in some way optimal in size and position can further reduce the number of mislabelled pixel and in turn enhance the overall precision of the segmentation algorithm.

The mask and positioning method used in this report will result in a high recall but relatively bad precision and a possible next step, to improve the algorithm, is Image Registration. This by creating a better cartilage mask, based on more data than the data used in the implementation described in this report, together with an image registration algorithm capable of registering new data to a chosen reference.

4.7.5 Testing new combinations

The complete algorithm is designed so that introducing new steps or replacing steps would be quite easy.

A possible combination to test is to use the cartilage mask from the morphological segmentation together with a bone mask and a geometrical function. Because of the muscles that are currently present in the cartilage mask, an additional step of muscle removal might be needed to reduce false positives that are marked as cartilage. Removing the muscles from the mask might be done with some kind of thresholding, for example Otsu's thresholding method.

5 Conclusion

The proposed method does segment the cartilage to some extent, but it needs to be improved in several aspects and is not yet ready for clinical use. The method can be seen as a base for further development.

Algorithms such as watershed, k-means and geometrical algorithms have shown to have capabilities of either preprocessing or finding the cartilage, while others such as Otsu's thresholding and Chan-Vese have not shown as useful. It has also shown that the combination of several methods lead to a better result than staying to one single method.

The final algorithm has been built in a way where different steps could be relatively easily replaced. If a better option for segmentation is found, the k-means classifier can be exchanged, and similarly the other parts can be exchanged as well. While the resulting recall is high using optimal settings, the precision is low which is the main problem with the final algorithm. Suggestions for how the different methods can be improved has been provided, together with possible replacements for some methods that could be investigated, all which could be used in future research to further improve the final segmentation accuracy.

References

- Abdelsamea, Mohammed. M. (2014). An Automatic Seeded Region Growing for 2D Biomedical Image Segmentation. URL: https://arxiv.org/ftp/arxiv/papers/1412/1412. 3958.pdf (visited on 2017-11-23).
- Delon et al., Julie (2017). Color Image Segmentation Using Acceptable Histogram Segmentation. URL: http://www.math-info.univ-paris5.fr/~desolneux/papers/DDLP_ IbPRIA_05.pdf (visited on 2017-11-23).
- Ester, M. et al. (1996). "A Density-Based Algorithm for Discovering Clusters". In: KDD'96 Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, pp. 226–231.
- Getreuer, Pascal (2012). Chan-Vese Segmentation. URL: http://www.ipol.im/pub/art/2012/g-cv/article.pdf (visited on 2017-11-30).
- Harris, Naftali (2015). Visualizing DBSCAN Clustering. URL: https://www.naftaliharris. com/blog/visualizing-dbscan-clustering/ (visited on 2017-11-13).
- Kass *et al.*, Micheal (1988). *Snakes: Active Contour Models*. URL: http://web.cs.ucla. edu/~dt/papers/ijcv88.pdf (visited on 2017-11-30).
- Matteucci, Matteo (2017). *K-Means Clustering*. URL: https://home.deib.polimi.it/ matteucc/Clustering/tutorial_html/kmeans.html (visited on 2017-11-07).
- Patil et al., Amruta B. (2016). OTSU Thresholding Method for Flower Image Segmentation. URL: http://www.ijceronline.com/papers/Vol6_issue5/A0650106.pdf (visited on 2017-11-30).
- Pierpaoli, Carlo (2010). Quantitative Brain MRI. URL: https://insights.ovid.com/ pubmed?pmid=21613871 (visited on 2017-12-07).
- S.Beucher and F.Meyer (1990). *Marker-Controlled Watershed Segmentation*. URL: http://cmm.ensmp.fr/~beucher/publi/SB_watershed.pdf (visited on 2017-11-21).
- Wu, Yue (2009). MATLAB CODES FOR ACTIVE CONTOURS. URL: https://sites. google.com/site/rexstribeofimageprocessing/chan-vese-active-contours/ matlabcodesforactivecontours (visited on 2017-12-05).
- Zitová, B. and J. Flusser (2003). "Image registration methods: a survey". In: *Image Vision Computing 977-1000* 11.21.