UPPSALA
UNIVERSITET

UPTEC X 22001
Examensarbete 30 hp
Juni 2022

Development and evaluation of an inter-subject image registration method for body composition analysis for three slice CT images

Hugo Dahlberg

Civilingenjörsprogrammet i molekylär bioteknik

UPPSALA
UNIVERSITET

Development and evaluation of an inter-subject image registration method for body composition analysis for three slice CT images

Hugo Dahlberg

Abstract

Over 30 000 liver, abdomen, and thigh slices have been acquired by computed tomography for the SCAPIS and IGT study. To utilise the full potential of the large cohort and enable statistical pixel-wise body composition analysis and visualisation of associations with other biomarkers, a point-to-point correspondence between the scans is needed.

For this purpose, an inter-subject image registration pipeline that combines the low-level information from CT images with high-level information from segmentation masks have been developed. It uses tissue-specific regularisation and processes images efficiently.

The pipeline was used to deform 4603 CT scans of each slice into a respective common reference space in less than 30 hours. All but the ribs in the liver slices and the intra abdominal region of the abdomen were generally registered correctly. Vector and intensity magnitude errors indicating inverse consistency were on average less than 2.5 mm and 40 Hounsfield units respectively.

The method may serve as a starting point for statistical pixel-wise body composition analysis, its association with non-imaging data, as well as saliency mapping analysis of the three-slice CT scans from the large SCAPIS and IGT cohorts.

Teknisk-naturvetenskapliga fakulteten

Uppsala universitet, Utgivningsort Uppsala/Visby

Handledare: Joel Kullberg Ämnesgranskare: Robin Strand

Examinator: Siv Andersson

Sammanfattning

Hjärt- och lungsjukdomar är den största dödsorsaken världen över. Trots det, så kan nuvarande strategier för att förebygga och diagnostisera sjukdomarna vara utdaterade. Det har medfört att en studie, SCAPIS, ska karaktärisera 30 000 vuxna män och kvinnor för att se om det går att förbättra strategierna. Datortomografibilder, ett slags röntgenbilder har tagits på alla 30 000 testpersoner tillsammans med hälsorelaterad information.

För att möjliggöra nya typer av analyser där man pixelvist jämför alla testpersonernas kroppssammansättning med dess hälsorelaterade information, så krävs det att alla bilder har en punkt-till-punkt överensstämmelse. Exempelvis att punkterna som beskriver levern i bilden på en testperson ska motsvara punkterna som beskriver levern i bilden på en annan testperson. En metod som används för att skapa punkt-till-punkt överensstämmelser kallas bildregistrering.

I detta arbete har jag utvecklat och utvärderat en bildregistreringsmetod för datortomografibilderna i SCAPIS. Den har testats på 4603 bilder och visar på relativt god kvalité. Metoden kan även användas till datortomografibilder som tagits fram på identiskt vis för en studie som undersöker personer med nedsatt glukosupptagningsförmåga.

Den utvecklade bildregistreringsmetoden kan ses som ett försteg som möjliggör nya avancerade analyser på pixelnivå som jämför kroppsammansättning och annan hälsorelaterad information. Exempelvis hur storleken på levern korrelerar med diabetes.

Table of contents

Abbreviations

BMI	Body Mass Index
COPD	Chronic Obstructive Pulmonary Disease
CT	Computed Tomography
CVD	CardioVascular Disease
FOV	Field Of View
HU	Hounsfield Units
IGT	Impaired Glucose Tolerance
IME	Intensity Magnitude Error
ISAT	Inside Subcutaneous Adipose Tissue
PET	Positron Emission Tomography
ROI	Region Of Interest
SAT	Subcutaneous Adipose Tissue
SCAPIS	Swedish CArdioPulmonary bioImage Study
SSD	Sum of Squared Distances
VME	Vector Magnitude Error

1 Introduction

Cardiopulmonary diseases are a major cause of death worldwide, but current strategies for prevention and diagnosis could be outdated (Bergström *et al.* 2015). In order to improve these strategies, the Swedish CArdioPulmonary bioImage Study (SCAPIS) was initialised. It is a nation-wide study that uses advanced imaging technologies together with biomarkers and epidemiological analyses to characterise 30 000 adults. One of the principal aims of SCAPIS is to use imaging methods to examine fat deposits together with information obtained by ”omics” technologies to improve the understanding of the role of obesity and diabetes in cardiovascular disease (CVD) and chronic obstructive pulmonary disease (COPD). As a step to reach this goal, the PET/MR research group at Uppsala University has started a SCAPIS innovation project ”Advanced Body Composition Analysis in SCAPIS using Deep Learning and Image Registration”. In parallel with SCAPIS, a mirror cohort of 1923 subjects have been selected with different forms of glucose dysregulation, aimed at a deeper understanding of impaired glucose tolerance (IGT). Both studies have identical three slice (liver, abdomen and thighs) computed tomography (CT) scans.

Image registration is one of the most important technologies in imaging and especially in medical imaging (Modersitzki *et al.* 2015). It is a widely used tool that can be employed to find a spatial transformation aligning a set of images and creating a common reference space. The common reference space can be used to fuse images acquired by different modalities, investigate anatomical and structural changes in longitudinal studies, and do statistical voxel-wise body composition analysis between subjects (Sotiras *et al.* 2013). In an article from 2017 (Strand *et al.* 2017), a three step whole-body image registration method for MRI images is described. It utilises relative elasticity constraints of different tissue and shows successful deformations to the common reference space. The method enables analyses of relationships to non-imaging data in a new type of holistic manner. Similarly, Jönsson *et al.* (2021) have developed a step-by-step method for image registration of whole-body PET-CT images. It is an almost fully automatic pipeline, which enables efficient image processing and voxel-wise studies. However, to the best of our knowledge no image registration method have been developed for the three-slice-CT scans from the SCAPIS and IGT cohorts.

The aim of this project was to enable pixel-wise statistical analysis of the three-slice-CT scans in the SCAPIS and IGT cohorts and visualisation of associations between body composition and other biomarkers.

The research objectives are:

- Develop and optimise an inter-subject image registration pipeline for the three-slice-CT scans in the SCAPIS and IGT cohorts.
- Evaluate the performance of the registration method based on its inverse consistency, diffeomorphic property, computational time, and group-wise statistics.

2 Background

This section is divided into three parts. The first part describes CT. The second part gives a background of what image registration is. The last part describes how to evaluate an image registration method.

2.1 Computed Tomography (CT)

A common way to acquire medical images is to use a CT scanner. A CT scanner uses an X-ray source and a detector that are positioned on opposite sides of the patient, mounted on a rotational gantry that can spin at high speeds (Ginat & Gupta 2014). After the X-ray source have completed a full rotation around the body a cross-sectional image (slice) is created by mathematical reconstruction of the measured X-ray intensities that the detector receives at different positions around the patient. A tissue-specific attenuation coefficient can be derived from the detected intensity. The coefficients are normalised with respect to the attenuation coefficient for water, which results in an expression of Hounsfield Units (HU) (Brooks & Chiro 1976). This means that the intensities are standardised between different machines and importantly for analysis applications - different tissues corresponds to different intervals of HU.

2.2 Image registration

As aforementioned, image registration is one of the most important technologies in imaging and especially in medical imaging (Modersitzki *et al.* 2015). Figure 1 presents an example of registering one source image to a target image. The target image is the reference space and the source images are all of the images in the dataset that should be registered. The objective of a registration process is to find an optimal transformation that can transform the source image to the reference space, so that if the transformation

is applied to the source image, it should be as similar as possible to the target image. To obtain a common reference space for a full dataset, all of the images in the dataset needs to be registered to the target image which means that an optimal transformation for each image needs to be found.

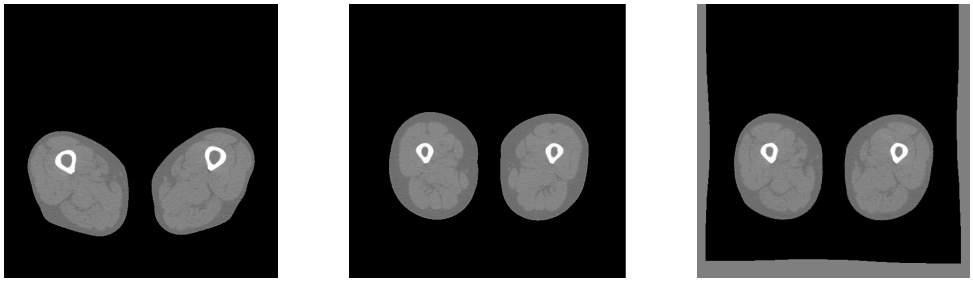

(a) Source (b) Target (c) Result

Figure 1: Example of a source image (a) registered to a target image (b). It is a successful registration since the resulting transformed image (c) is similar to the target image.

In order to automatically find an optimal transformation from each source image to the target image, the image registration method needs to consist of three main components; a transformation model, an objective function and an optimisation method (Ekström 2020). The transformation model can be global and restricted with few degrees of freedom which could prevent the method to find meaningful transformations. For complex applications a less restricted model is needed which increases the search space and makes it more difficult and time-consuming to find a solution (Holden 2008). The least restricted model allows each point in the source image to be displaced arbitrarily by generating a dense displacement field that consists of a displacement vector for each point in the source image. This enables local displacements in the image which is needed in for example inter-subject image registration.

The second component of the image registration method is the objective function. It can either be based on landmarks or on image intensities (Ekström 2020). A simple metric based on image intensities is sum of squared differences (SSD). It assumes that there is a direct correspondence between the intensities in the two images which makes it only suitable for some modalities, such as CT. Usually, the objective function includes regularisation, especially if it is used on a transformation model with no restriction. A common explicit regularisation method is to penalise first-order derivatives in order to impose smoothness.

The optimisation method is the third component of the image registration method. It depicts how to find the transformation that optimally align the two images based on the objective function (Ekström 2020). Gradient descent is a common algorithm for contin-

uous optimisation method and graph cut optimisation is a popular discrete optimisation method. A problem that often occurs during the optimisation is that a local minimum is found instead of the global minimum. One common way to avoid these local minima is to use a hierarchical multi-resolution strategy. It means that the registration is performed in a hierarchical matter with increasing resolution, e.g. increasing data complexity, until the registration reaches the final resolution.

Except for the main components of a registration method, the choice of reference image can have a great impact on the resulting method (Ekström 2020). A good reference should represent the mean image, so it is as similar as possible to all of the other images in the dataset. If metadata is available one possible way to find a good reference image is by choosing the subject with the mean body mass index (BMI) (Sjöholm *et al.* 2019) and by visual assessment. Another way is to compute a synthetic reference space (Pilia *et al.* 2019).

2.3 Evaluation of image registration

There is no ground truth or unique solution to an image registration method, making it hard to evaluate. However, a common way is to evaluate a combination of the image similarity and the quality of the transformations (Song 2017). One straightforward method to evaluate the image similarity is by visual assessment. To evaluate the quality of the transformations it is common to check if they are inverse consistent and diffeomorphic. For an image registration method to be inverse consistent, a transformation in the reversed direction from the original transform needs to be equal to the inverse transform. However, since there is no unique solution it is rare that the transform in the reversed direction is equal to the inverse transform. Instead, two measures of how far the transformation is from being inverse consistent, called vector magnitude error (VME) and intensity magnitude error (IME), are commonly used (Jönsson *et al.* 2021; Ekström 2020). A diffeomorphic transformation is differentiable and its inverse is differentiable (Song 2017). If the Jacobian determinant, from here on referred to as the Jacobian, of the transformation is positive in all points then it is diffeomorphic. Furthermore, the Jacobian in a point describes the local tissue change of that point. A Jacobian above one implies a local expansion and a Jacbian between zero and one implies a local contraction. However, a point with a negative Jacobian infers that the transformation inverts or folds at that point which implies a non-diffeomorphic property. A metric to evaluate an image registration methods' ability to produce diffeomorphic transformations is the number of foldings it produces, which can be derived from the computed Jacobian of the transformation.

3 Material & methods

This section is divided into three parts. The first part describes the images and how they were acquired. The second part describes the image registration method. The last part describes how the image registration method was evaluated.

3.1 Data

CT scans from two mirror cohorts were used. The first cohort, from SCAPIS, comprised of axial liver, abdomen and thigh slices from 2680 women and men aged between 50 and 64 years. The second cohort, from the IGT study, consists of the same slices from 1923 subjects with different forms of glucose dysregulation. The scans were acquired with a specific SCAPIS protocol, most images have a field of view (FOV) of 500 mm and they were in DICOM format. The image size were 512x512x1 with a slice thickness of 5 mm. Slices are visualised from the feet in a proximal direction. Therefore, the left in the image corresponds to the right in subjects' perspective and vice versa. Some details specific for each slice is described in the subsections below.

3.1.1 Liver images

The SCAPIS protocol states that the liver slice should be reconstructed from a thorax inspiration scan. It should consist of spleen, both liver lobes, and some lung. The scan should maximise the area of the liver and the area of the spleen should be at least two square centimeters. Kidney and heart should be avoided. A very small subset of the liver slices can be seen in figure 2. There are many differences between the slices of the subjects, including the number of ribs, if and where there is visible air, and the size and shape of the body. Even though the shape and size may differ between the subjects they all contain liver, spleen, vertebra, SAT, and a body shape.

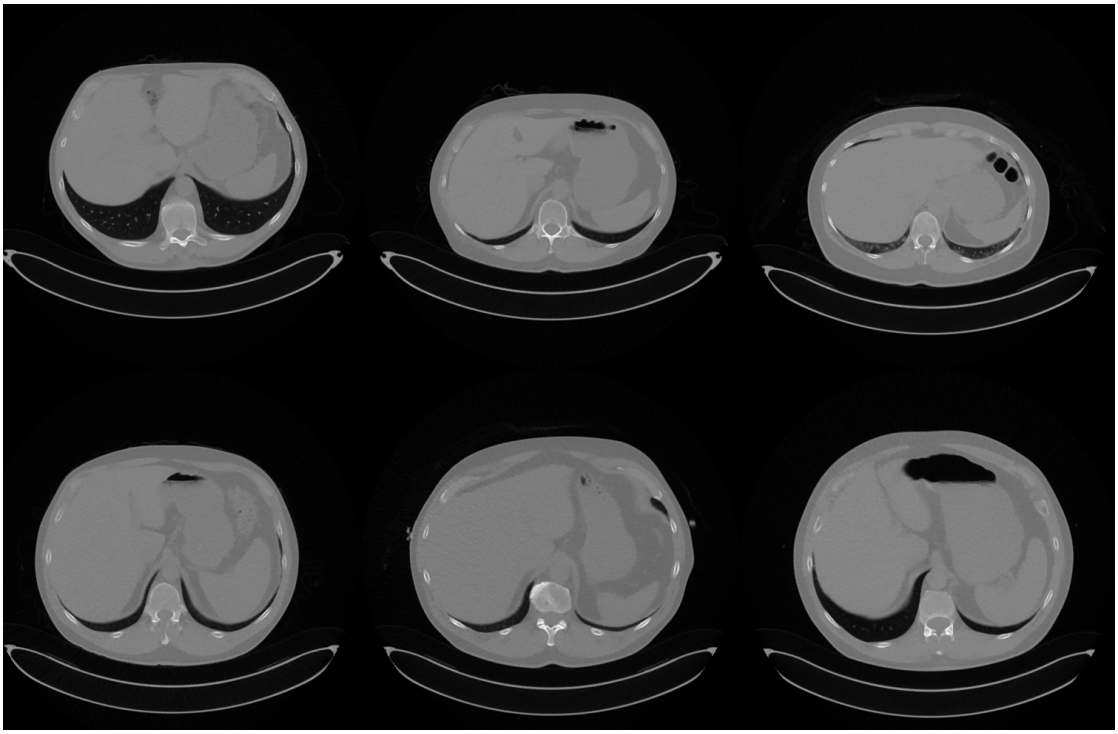

Figure 2: A small subset of the liver slices

3.1.2 Abdomen images

For abdomen slices, the SCAPIS protocol states that it should be taken in the middle of L4. Crista should be avoided and if possible, liver, spleen, and kidney should also be avoided. The scan should be acquired during maximum expiration. Figure 3 shows a very small subset of the abdomen slices and it is possible to see large differences between the subjects. Because of the inconsistent nature of the intestines, both within and between subjects, the axial slices contain different parts of the intestines and they are located differently between the subjects. The abdomen in the lower corner to the left of figure 3 does not have any SAT at some locations, which most other subjects have around all of the body. Even though the size and shape can differ between the subjects they all contain skeletal muscles, vertebra, and a body shape.

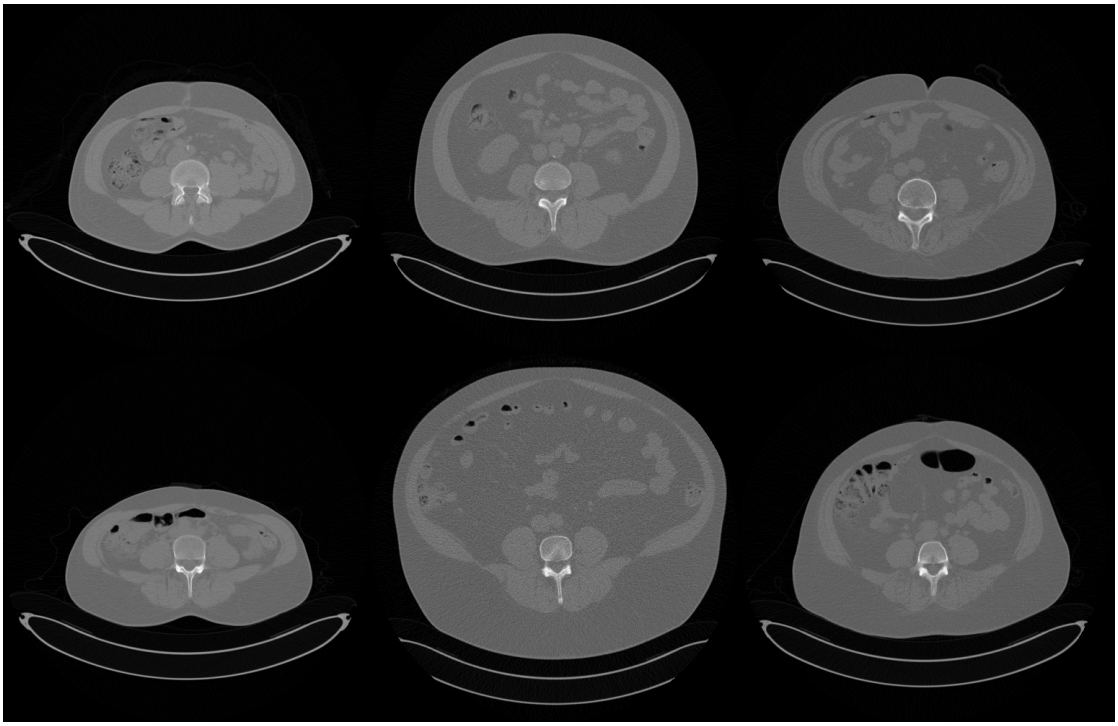

Figure 3: A small subset of the abdomen slices

3.1.3 Thigh images

The SCAPIS protocol states that the thigh slices should be taken in the middle of the acetabulum and the synovial surface of the right knee. Figure 4 shows a small subset of the thigh slices and it is possible to see the simplicity compared to the other two slices. A big difference between the subjects is the amount of SAT, they have any. However, the bone and muscle tissue vary less between the subjects.

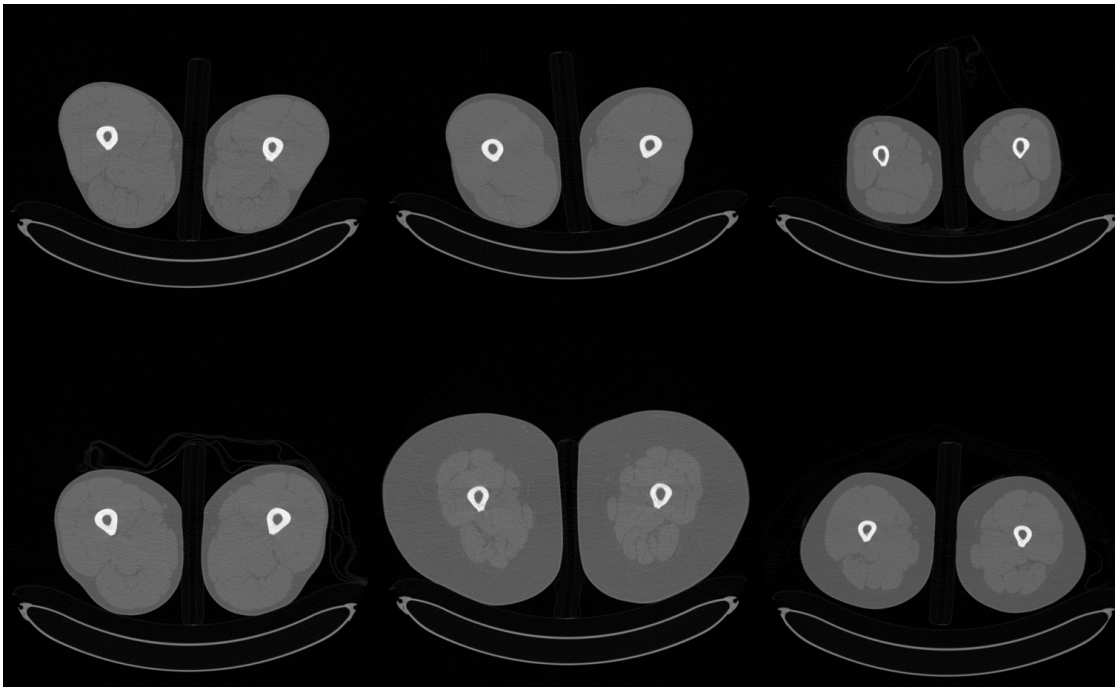

Figure 4: A small subset of the thigh slices

3.2 Image registration

The registration pipeline was implemented in Python and a schematic flowchart can be seen in figure 5. Before running the pipeline a reference space needs to be chosen. In this work, the reference space for each slice was chosen by visual assessment of the images in the dataset. The next step is to preprocess both the reference image and the source image. When they are preprocessed, the source image is registered to the reference image and the pipeline outputs a transform that can be applied to align the images. The preprocessing and registration step is repeated for every source image in the dataset. Ultimately, the transforms are evaluated both individually and all together. Most of the pipeline is the same between the slices, however some preprocessing and parameters differs which is further described in the next subsections.

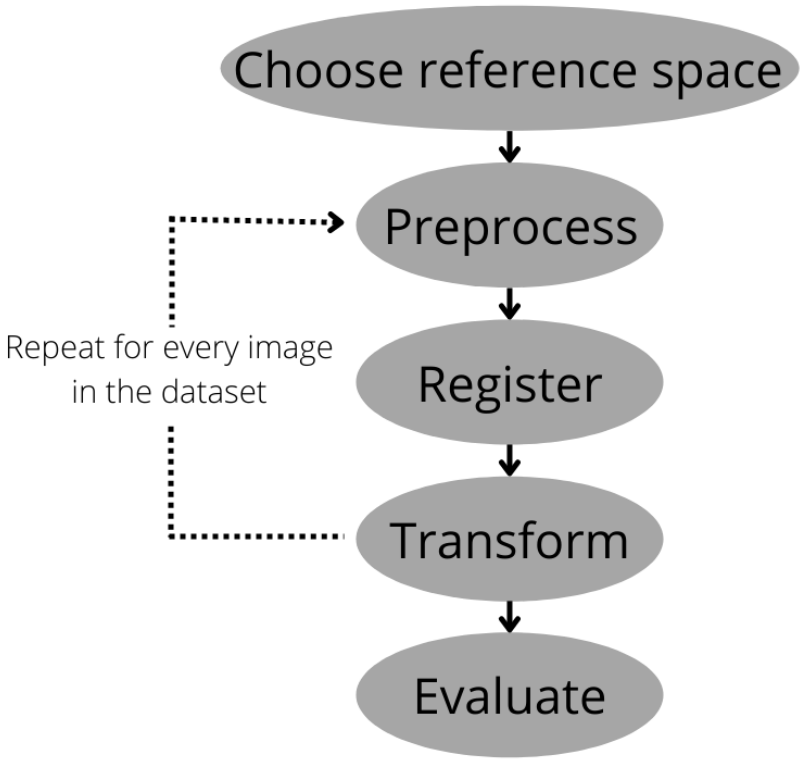

Figure 5: A schematic flowchart of the registration pipeline

3.2.1 Preprocessing

The first preprocessing step is to set the origin to a default [0,0,0], so all images have the same origin and can be registered to the same space. Intensities above 1024 HU was set to 1024 HU and images with at least 50 pixels above 2000 HU were noted to be checked for metallic implants.

The protocol to obtain the CT scans states that they should have a FOV of 500 mm (Bergström *et al.* 2015). However, there are a few images with different FOV and they are resampled before any further preprocessing step so all images have a FOV of 500 mm.

Each CT image was used to generate binary masks used for preprocessing and/or input during registration. Generally, masks for input during registration were generated for organs and tissues that existed in all images for each type of slice. The region of interests (ROI) of the binary masks had an intensity of one and the rest of the image had an

intensity value of zero. A body mask was generated for all slices by thresholding (above -200 HU), labeling and choosing the largest object (for thighs the two largest objects were chosen). The body mask was used to remove the CT table and other non-body objects from all the CT images.

For the liver registration both a liver mask and a spleen mask was generated from the output of two separate Ghost-Net/U-Net deep learning models trained to segment the respective organ on a subset of the slices and developed by the PET/MR research group. A vertebra mask was also generated by a combination of thresholding (above 200 HU (Broder 2011)) and morphological operations.

For the abdomen registration a muscle mask was generated from the output of a Ghost-Net/U-Net deep learning model trained on a subset of the slices to segment the skeletal muscles. This was used in combination with the algorithm developed by Jönsson *et al.* (2021) to generate an inside subcutaneous adipose tissue (ISAT) mask. Furthermore, a vertebra mask was generated in the same way as for the liver registration.

A bone mask and a lean tissue mask was generated for the thigh registration. The bone mask was produced by thresholding (above 200 HU (Broder 2011)) and then it was filled so the bone marrow were included in order to decrease the number of local minima for the registration. Similarly, the lean tissue mask was produced by thresholding (between -29 HU and 150 HU (Mitsiopoulos *et al.* 1998)), labeling and then extracting the largest object to exclude the bone marrow from the mask.

3.2.2 Registration

Images were registered in one step with a multi-channel input of the preprocessed CT-image together with the generated masks for each slice, figure 6 shows an example of the preprocessing and the multi-channel input to the registration of the abdomen images. SSD was used as objective function for all input images. However, the CT image had a higher weight than the masks, making the registration more dependent on the CT image than the masks. The exact image weights and all parameters used for the registration can be found in table 2 & 3, in the Appendix. The registration method utilised the least restricted transformation model to allow local displacements and generate a displacement field. A tissue-specific regularisation weight map was generated based on HU-thresholds. A lower regularisation weight was applied to tissues with high elasticity (e.g. fat) and a higher weight to tissues with low elasticity (e.g. bone). To optimise the objective function a fast graph-cut based approach (Ekström *et al.* 2020) was used. To avoid local minima a hierarchical multi-resolution strategy was implemented. It started on the lowest resolution possible and terminated on level 1, meaning that the final resolution is the original resolution downsampled by a factor of two in each dimension.

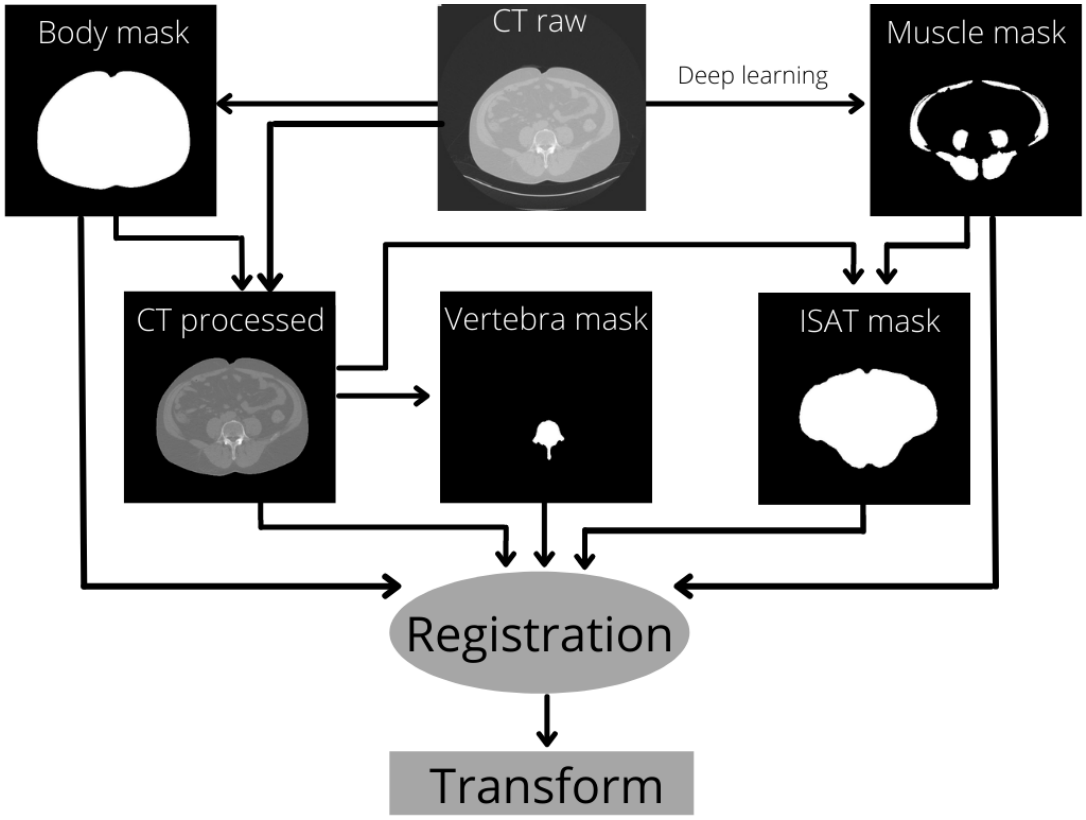

Figure 6: Flowchart of abdomen registration

3.3 Evaluation of the image registration

The performance of the registration method was evaluated based on the resulting transformed images, the quality of the transforms and computation time. During the development of the registration method each transformed image was visually assessed by its similarity to the reference image. Pixel-wise mean, standard deviation, and inter-quartile range were computed of all registered images in the three slices and evaluated by visual assessment. The quality of the transforms were evaluated by inverse consistency and their diffeomorphic property. To evaluate the inverse consistency, all registrations were performed in reversed direction and an average of VME and IME were computed over the segmented body region for the three slices. The Jacobians of the resulting transforms were computed. To evaluate the diffeomorphic property the number of points where the Jacobian was less than zero were counted, i.e. the number of folds. The natural logarithm of the Jacobian were computed to easily interpret local tissue volume change between the subjects. Computation time of the registration method was also measured, including

the time for all preprocessing steps except the step to generate masks with a deep learning model. An Intel Xeon Silver processor with 4216 CPU were used on a computer with 32 GB installed RAM and a SSD hard-drive. To evaluate the performance of the developed registration pipeline the evaluation metrics used to evaluate each individual registration were averaged over all registrations performed on the dataset.

4 Results

In this section the evaluation of the image registration method is described. The overall performance metrics for the quality of the transformations for each slice can be seen in table 1. It is possible to see that the thigh registration has performed best but the liver registration was the fastest. The abdomen registration took the most time and the transforms' quality is not as good as for the other slices, especially the IME is significantly higher. The rest of the result section is divided into three parts, one for each slice in the dataset.

*Table 1: Performance metrics for the three registration methods. Data presented as mean (standard deviation). VME = vector magnitude error, IME = intensity magnitude error, HU = Hounsfield Units. *Only mean was measured.*

	Liver registration	**Abdomen registration**	**Thigh registration**
VME (mm)	2.44 (1.33)	2.51 (1.67)	1.23 (0.66)
IME (HU)	36.81 (12.75)	47.44 (10.95)	30.57 (8.04)
Number of folds	0.22 (5.05)	0.49 (7.49)	0.12 (4.21)
Registration time (s)*	6.1	7.9	7.5

4.1 Liver registration

Figure 7 shows a resulting collage of pixel-wise statistics of the liver registration. The mean of the registered images is sharp enough to distinguish the lungs, liver, spleen, vertebra, and SAT. However, the ribs are not visible and the anterior left side in the subjects perspective is darker than its surroundings and relatively blurry. High intensity standard deviation was mostly located in the lungs and at the borders to air and vertebrae. The liver, spleen and SAT have low intensity standard deviation.

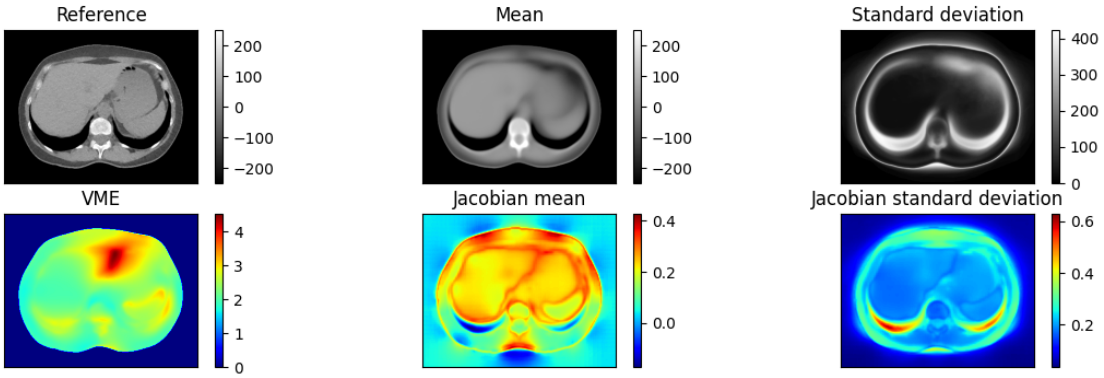

Figure 7: Resulting collage for the liver registration of 4603 CT scans. (Reference) shows the preprocessed reference image. (Mean) shows the mean intensity of all registered images. (Standard deviation) shows the intensity standard deviation of all registered images. (VME) visualise the mean VME for all registrations. (Jacobian mean) visualise the mean of the logarithm of the Jacobian in every point. (Jacobian standard deviation) show the standard deviation of the logarithm of the Jacobian in every point. For visibility purposes the intensity scale of the reference and mean images have been changed from -1024 HU to 1024 HU to -250 HU to 250 HU. VME = Vector magnitude error.

The visualisation of the average VME for the registration shows where the registration is consistent. A higher VME means a greater error and that the registration is further away from being inverse consistent and a lower VME means the opposite. From figure 7 it is possible to see that the registration is the least inverse consistent at the upper right edge of the liver. Other parts with high VME are the border of the spleen and the lungs. Low VME can be found in the middle of the liver, the upper part of the vertebrae and the SAT. The Jacobian for every generated displacement field were computed. For easier interpretation, the logarithm of the Jacobian in each point were calculated and the average for all points can be seen in figure 7. A negative value corresponds to a local contraction and a positive value corresponds to a local expansion. As the mean Jacobian is mostly green, yellow, and red, its average is positive meaning that the reference image generally is smaller than the other images in the dataset. However, the mean Jacobian is negative in the lungs meaning that the reference image generally has a lung larger than the source images. The lungs also shows a high Jacobian standard deviation which means that the size of the lungs vary between the subjects in the dataset. Generally the adipose tissue has a higher standard deviation than the lean tissue and bone. Meaning that the size of adipose tissue varies between the subjects more than what the size of the lean tissue and bone varies between the subjects in the dataset.

4.2 Abdomen registration

The mean of the registered abdomen images is sharp at the skeletal muscles, vertebrae and SAT (see fig. 8). However, the intra abdominal region is blurry. The same region has a high intensity standard deviation together with borders to air and bone. A thicker white border can be found the top and bottom of the body.

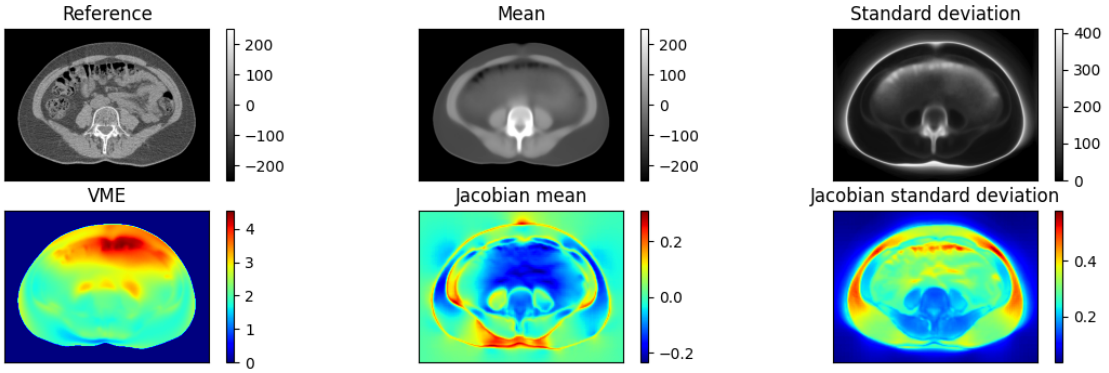

Figure 8: Resulting collage for the abdomen registration of 4603 CT scans. (Reference) shows the preprocessed reference image. (Mean) shows the mean intensity of all registered images. (Standard deviation) shows the intensity standard deviation of all registered images. (VME) visualise the mean VME for all registrations. (Jacobian mean) visualise the mean of the logarithm of the Jacobian in every point. (Jacobian standard deviation) show the standard deviation of the logarithm of the Jacobian in every point. For visibility purposes the intensity scale of the reference and mean images have been changed from -1024 HU to 1024 HU to -250 HU to 250 HU. VME = Vector magnitude error.

The visualisation of the average VME shows that the registration of the most anterior skeletal muscles are the least consistent and that the registration of the most posterior SAT are most consistent. The resulting average logarithm of the Jacobian of the abdomen registration implies that the lateral SAT of the reference image generally is larger than for the rest of the dataset, since it has a negative value. The same applies to the anterior skeletal muscles and the middle of the abdomen. The positive values at the lateral skeletal muscles indicates that the reference image generally has smaller lateral skeletal muscles than the rest of the dataset. There are also positive values above the anterior side and at the posterior side of the body. Furthermore, it is possible to see that the size of the SAT varies greatly between the subjects in the dataset as it has a high Jacobian standard deviation. The intra abdominal region and intestinal gases also have a high Jacobian standard deviation. However, the skeletal muscles and vertebrae have a lower value, which implies that their size do not vary as much between the subjects in the dataset.

4.3 Thigh registration

Figure 9 shows a resulting collage of pixel-wise statistics of the thigh registration. The average pixel-wise intensity of the registered thigh images results in a sharp image with lean tissue, adipose tissue and bone separated. There are only high intensity standard deviation at the borders between tissue types for the images.

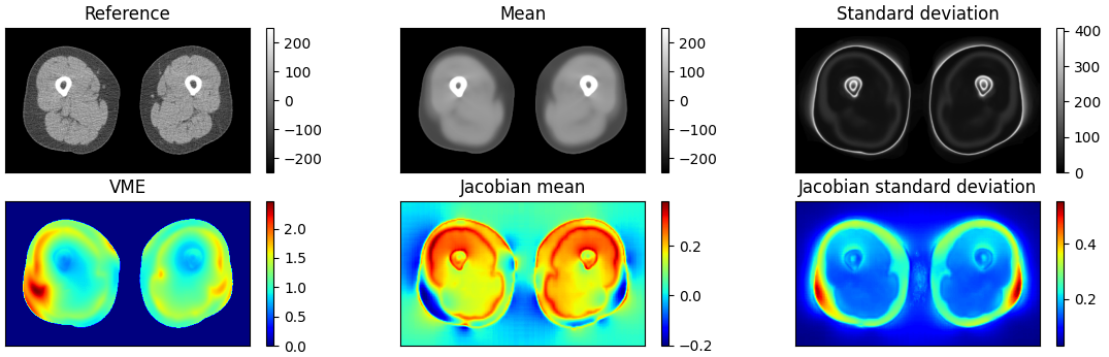

Figure 9: Resulting collage for the thigh registration of 4603 CT scans. (Reference) shows the preprocessed reference image. (Mean) shows the mean intensity of all registered images. (Standard deviation) shows the intensity standard deviation of all registered images. (VME) visualise the mean VME for all registrations. (Jacobian mean) visualise the mean of the logarithm of the Jacobian in every point. (Jacobian standard deviation) show the standard deviation of the logarithm of the Jacobian in every point. For visibility purposes the intensity scale of the reference and mean images have been changed from -1024 HU to 1024 HU to -250 HU to 250 HU. VME = Vector magnitude error.

The average VME of every registration of thighs in the dataset is low around the bone, which indicates that the registration is consistent around the bone. Note that the scale differs from the liver and abdomen registration. The registration is the least consistent at the lateral SAT, as well as the border between the adipose tissue and the lean tissue. The logarithm of the Jacobian mean of the generated displacement fields is negative in the lateral SAT meaning that the reference image generally has more adipose tissue there than the rest of the dataset. The reference image generally has less lean tissue than the rest of the thigh slices since the mean Jacobian has a positive value at the lean tissue. From the Jacobian standard deviation it is possible to see that the lean tissue does not vary between the thigh slices as much as the adipose tissue. The lateral SAT varies the most between the images in the dataset since that is where the Jacobian standard deviation has the greatest value.

5 Discussion

To utilise the full potential of the large SCAPIS and IGT cohorts an inter-subject image registration pipeline have been developed and evaluated. It enables pixel-wise statistical body composition analysis, associations with other biomarkers, and saliency mapping analysis.

During development of the registration method a several-step approach with every step adapted to tissue elasticity was tested similar to the methods in (Strand *et al.* 2017; Jönsson *et al.* 2021). It increased the computation time but showed no increase in performance, which is why the one-step approach was favored. However, to adapt the registration to the tissue elasticity a tissue-specific regularisation weight map were generated, which increased performance. Moreover, the main characteristic of the developed image registration method is the multi-channel input. The deformation algorithm inputs a CT image as well as several generated masks of structures that exists in the respective slices between all subjects, which drives the algorithm to find point-to-point correspondence on a high level between the same structures and avoid local minima.

The resulting performance metrics shows that the quality of the transforms are relatively good (see table 1). The average VME is less or similar compared to similar methods (Strand *et al.* 2017; Jönsson *et al.* 2021). However, those methods register whole-body volumes with different spacing from the method described in this thesis, making it difficult to assess the value of the VME. The results also show that the thigh registration generated best quality of the transforms of the three registration pipelines. It is expected as thigh slices are more similar to each other and not as complex as the other two slices which makes the registration task easier. Although the liver registration have similar VME as the abdomen registration the IME of the abdomen is significantly higher. An explanation to this could be that the liver scans all have relatively large areas with the same HU so there are relatively few borders with big intensity difference. Whereas the intra abdominal region of the abdomen scans varies greatly in HU, mostly because of the intestinal gases. Because of this, every small error in the intra abdominal region of the abdomen registration can generate a big intensity difference and consequently a high IME, while on average a greater error needs to occur to generate a big intensity difference in the liver registration. Hence, the IME is lower for the liver registration than for the abdomen registration.

One possible explanation why the liver registration is faster than the other two pipelines could be that the measured time only includes the time for generating the body mask and vertebrae mask, and not the two deep learning generated masks spleen and liver. The time measured for the abdomen and thigh registrations include the time for three

generated masks instead of the liver registrations two, which could be the difference. However, the preprocessing steps are not the computational heavy steps in the pipeline so it should not make that big of a difference. The rest of the discussion is divided into four sections, one for the each slice and one for limitations.

5.1 Liver registration

The sharp mean image of the registered livers (see fig. 7 with visible major organs together with low standard deviation at the liver, spleen, and SAT indicate successful inter-subject registration to the majority of the reference space. However, due to between-subject variability and structural differences some parts of the image is not perfectly registered. The mean of the registered liver image is darker, in its anterior left side from the subjects perspective, than the surrounding area. This is because of air in some images (see fig. 2), that can not be physically contracted into its removal of existence. As the number of visible ribs differ between the images it is impossible for the registration method to register them correctly and therefore they are not visible in the mean image.

The high intensity standard deviation in the borders between air and other tissue as well as bone and other tissue can be explained by the big difference in intensity value at the borders. Air having an intensity of around -1000 HU and most other tissue having an intensity of -190 HU to 200 HU, makes a small spatial error in the registration of a source image of the border between air and tissue to generate a high intensity error. The same reasoning applies to the border between bone (200 HU - 1024 HU) and other tissue types. However, there is high intensity standard deviation in all of the lungs and not just the borders, which originates from the fact that some source images does not have any visible lung and it is impossible for the registration to expand an organ that is not originally visible. This makes the resulting registered image to have other tissue than lungs at the location of the lung in the reference image, which contributes to a high intensity standard deviation at the location. The high intensity standard deviation in the anterior left side from the subjects perspective is caused for the same reason.

The overall average VME is low which indicates that the registration is consistent and the quality of the produced transforms are good. The Jacobian mean image shows that the reference image generally is smaller than the rest of the images in the dataset. This means that it most probably is not the mean or median image of the dataset, which usually is preferred when choosing a reference space. Hence, a possible improvement to the pipeline would be to choose a different reference image or compute a synthetic one with the proposed method by Pilia *et al.* (2019). The Jacobian standard deviation shows that

the different tissues vary in size between the subjects at different rates. Lungs vary in size the most, adipose tissue size varies relatively much, lean tissue size varies relatively little, and bone/vertebrae size varies almost nothing between the subjects in the dataset. This is expected since it corresponds to the elasticity of the tissues, where lungs/air are the most elastic and bone is not elastic at all.

5.2 Abdomen registration

Similarly to the liver registration, the visualisation of the mean of the registered abdomens is sharp in most areas and the intensity standard deviation is low in most areas (see fig. 8). This indicates successful registration to the abdomen reference space in most areas, such as SAT, skeletal muscles and vertebrae. However, the intra abdominal region in the mean image is blurry, meaning that the registration is unsuccessful in this area. This is somewhat expected because of the details and between-subject variability (examples seen in fig. 3) in the area. The great amount of details causes many borders between tissues which makes it easier for the registration algorithm to find a local minimum instead of the global minimum and therefore misalign the point-to-point correspondence between the source image and the reference image. As the content of the intra abdominal region is mobile, it is near impossible to acquire the same content of the region in every subject in solely a slice of the abdomen. A difference in content means that there exist no point-to-point correspondence to obtain between the reference image and the source image, making it impossible to register. Hence, it is difficult or impossible to register the intra abdominal region of all abdomen slices, which is why the mean image of the abdomen is blurry in the area and it shows relatively high intensity standard deviation.

Why it is a higher intensity standard deviation at the anterior side of the intra abdominal region is because intestinal gases are located there, and as mentioned before a misalignment between air and other tissue causes high intensity standard deviation. The intestinal gases rises to the anterior side because the subject is laying on its back during the scan.

Some slices are taken over the subjects' umbilicus which is something the registration method does not manage to transform and therefore the anterior of the mean intensity image has a small pit/cavity and there is a high intensity standard deviation at the same location.

A possible explanation to the relatively high average VME in the most anterior skeletal muscles could be that some subjects lack anterior skeletal muscles partly (see fig. 3) or entirely. Besides the non-existent possibility to find a point-to-point correspondence,

the lack of anterior skeletal muscles can cause the automatic generation of ISAT masks to warp. Consequently, when the warped ISAT mask is used in the multi-channel input of the registration pipeline it will perform worse. However, generally the VME is low for the registration which implies good quality of the produced transforms.

As there are mostly average logarithm Jacobian values around zero and the mean of all values visually approximates to zero, it is possible to say that the reference image is relatively close to being the mean image of the dataset. Hence, it is a relatively good reference space.

Similarly to the Jacobian standard deviation of the liver registration the Jacobian standard deviation of the abdomen registration shows that the size of air and adipose tissue varies the most between the subjects in the dataset. The bone size varies barely at all and the lean tissue varies to some degree. It is as expected since adipose tissue is the most elastic tissue and bone is the least elastic.

5.3 Thigh registration

The sharp mean image of the registered thigh images together with the low intensity standard deviation implies that the registered images are similar to the reference image which equals a successful registration. The high intensity standard deviation found at the borders between tissue is expected, as discussed earlier about the liver registration.

The relatively high average VME at the lateral SAT and the border between the lean tissue and the adipose tissue probably originates from the reference image having relatively much SAT, whereas some subjects have essentially no SAT. This makes it harder for the registration to be consistent in these areas, and areas with the highest VME is where the SAT is the thickest in the reference image. However, the registration is generally consistent and it is the most consistent around the bone. This could be explained by the fact that the bone is located at roughly the same place between subjects and therefore none of the bone pixels is transformed far. If pixels are not transformed far in either direction of the registration, than there is no risk to obtain a high VME.

The average of the logarithm of the Jacobian of the thigh registrations visualise that the reference image has more lateral SAT and less lean tissue than the average thigh slice. As the already discussed Jacobian standard deviations of the liver and abdomen registration, the thigh registration shows similar expected results. The tissue types varies in size between the subjects as their relative elasticity.

5.4 Limitations

Time have been a limiting factor to this master thesis project. With more time, synthetic reference spaces could have been created as described by Pilia *et al.* (2019) which would probably have improved the performance of the registrations. Another limitation to this project was the absence of available metadata. With metadata, several reference spaces for each slice could have been created, for example registration of all male subjects to a male reference space and all female subjects to a female reference space, as recommended by Ekström (2020). This could also have improved the performance of the registrations, especially for the thigh registration since there generally are clear anatomical differences between female and male thighs. However, a drawback with dividing the registrations into several reference spaces is that the pixel-wise statistically analysis will also be divided into the different groups. But since most analysis are divided on the subjects' sex, it should not be a major drawback.

6 Conclusion

In this work, an inter-subject, single-step image registration pipeline for the three-slice CT scans from the SCAPIS and IGT cohorts have been developed and evaluated. It combines the low-level information from CT images with high-level information from segmentation masks. The pipeline was used to deform 4603 CT scans of each slice into a respective common reference space in less than 30 hours. The method may serve as a starting point for statistical pixel-wise body composition analysis, its association with non-imaging data, as well as saliency mapping analysis of the three-slice CT scans.

7 Ethical aspects

Radiation exposure is a major ethical dilemma when using CT scans as it causes a small potential risk for cancer (Bach *et al.* 2012). Although it is a small risk it is hard to quantify as the harm from radiation can take one or two decades before it appears. However, SCAPIS is expected to result in many more early detection of cancers than those caused by radiation (Bergström *et al.* 2015). In addition, the scientific potential of the large number of CT scans is huge and can in the future lead to better strategies to prevent and

treat different diseases. In this work, the CT scans have already been performed and the purpose is to utilise the full potential of the large cohort. Therefore there is no real ethical doubts for this work to be performed.

8 Ethical approval

Ethical approval has been obtained for development and application of advanced body composition analysis from CT images (Dnr 2021-05856-01).

9 Acknowledgements

This project would not have been possible without the continuous support and the many ideas from my supervisor Joel Kullberg, thank you. I would also like to thank my subject reader Robin Strand for all feedback and inputs to my project. Next, a big thanks should go to the PET/MR research for an inspiring working environment and support to my project. Lastly, I would like to thank my examiner Siv Andersson and the course coordinator Lena Henriksson for support during the course.

References

Bach PB, Mirkin JN, Oliver TK, Azzoli CG, Berry D, Brawley OW, Byers T, Colditz GA, Gould MK, Jett JR, Sabichi AL, Smith-Bindman R, Wood DE, Qaseem A, Detterbeck FC. 2012. Benefits and Harms of CT Screening for Lung Cancer: A Systematic Review. JAMA : the journal of the American Medical Association 307: 2418–2429.

Bergström G, Berglund G, Blomberg A, Brandberg J, Engström G, Engvall J, Eriksson M, de Faire U, Flinck A, Hansson MG, Hedblad B, Hjelmgren O, Janson C, Jernberg T, Johnsson □, Johansson L, Lind L, Löfdahl C, Melander O, Östgren CJ, Persson A, Persson M, Sandström A, Schmidt C, Söderberg S, Sundström J, Toren K, Waldenström A, Wedel H, Vikgren J, Fagerberg B, Rosengren A. 2015. The Swedish CArdioPulmonary BioImage Study: objectives and design. Journal of Internal Medicine 278: 645–659.

Broder J. 2011. Chapter 9 - imaging of nontraumatic abdominal conditions. Broder J, editor, Diagnostic Imaging for the Emergency Physician, W.B. Saunders, Saint Louis, 445–577.

Brooks RA, Chiro GD. 1976. Principles of computer assisted tomography (CAT) in radiographic and radioisotopic imaging. Physics in Medicine and Biology 21: 689–732.

Ekström S. 2020. Efficient GPU-based Image Registration : for Detailed Large-Scale Whole-body Analysis. Ph.D. thesis, Uppsala University, Radiology.

Ekström S, Malmberg F, Ahlström H, Kullberg J, Strand R. 2020. Fast graph-cut based optimization for practical dense deformable registration of volume images. Computerized Medical Imaging and Graphics 84: 101745.

Ginat DT, Gupta R. 2014. Advances in Computed Tomography Imaging Technology. Annual Review of Biomedical Engineering 16: 431–453. _eprint: https://doi.org/10.1146/annurev-bioeng-121813-113601.

Holden M. 2008. A review of geometric transformations for nonrigid body registration. IEEE Transactions on Medical Imaging 27: 111–128.

Jönsson H, Ekström S, Ahlström H, Strand R, Kullberg J. 2021. An image registration method for voxel-wise studies of whole-body oncological pet-ct.

Mitsiopoulos N, Baumgartner RN, Heymsfield SB, Lyons W, Gallagher D, Ross R. 1998. Cadaver validation of skeletal muscle measurement by magnetic resonance

imaging and computerized tomography. Journal of Applied Physiology 85: 115–122. Publisher: American Physiological Society.

Modersitzki J, Heldmann S, Papenberg N. 2015. Nonlinear Registration Via Displacement Fields. Toga AW, editor, Brain Mapping, Academic Press, Waltham, 307–314.

Pilia M, Kullberg J, Ahlström H, Malmberg F, Ekström S, Strand R. 2019. Average volume reference space for large scale registration of whole-body magnetic resonance images. PLOS ONE 14: e0222700. Publisher: Public Library of Science.

Sjöholm T, Ekström S, Strand R, Ahlström H, Lind L, Malmberg F, Kullberg J. 2019. A whole-body FDG PET/MR atlas for multiparametric voxel-based analysis. Scientific Reports 9: 6158.

Song JH. 2017. Methods for evaluating image registration. Ph.D. thesis, University of Iowa.

Sotiras A, Davatzikos C, Paragios N. 2013. Deformable medical image registration: A survey. IEEE Transactions on Medical Imaging 32: 1153–1190.

Strand R, Malmberg F, Johansson L, Lind L, Sundbom M, Ahlström H, Kullberg J. 2017. A concept for holistic whole body MRI data analysis, Imiomics. PLOS ONE 12: e0169966. Publisher: Public Library of Science.

A APPENDIX

Table 2: Image weights

	Liver registration	Abdomen registration	Thigh registration
CT image	0.5	0.6	0.5
Body mask	0.15	0.1	0.2
Vertebrae mask	0.15	0.1	-
Bone mask	-	-	0.2
Lean tissue mask	-	-	0.1
Liver mask	0.1	-	-
Spleen mask	0.1	-	-
Skeletal muscle mask	-	0.1	-
ISAT mask	-	0.1	-

Table 3: Settings for the registration

Parameter	Value
Step size	0.5
Update rule	Additive
Block size	[16, 16, 16]
Regularisation weight map	[0.08, 0.1, 0.15, 0.2]
Pyramid levels	8
Pyramid stop level	1