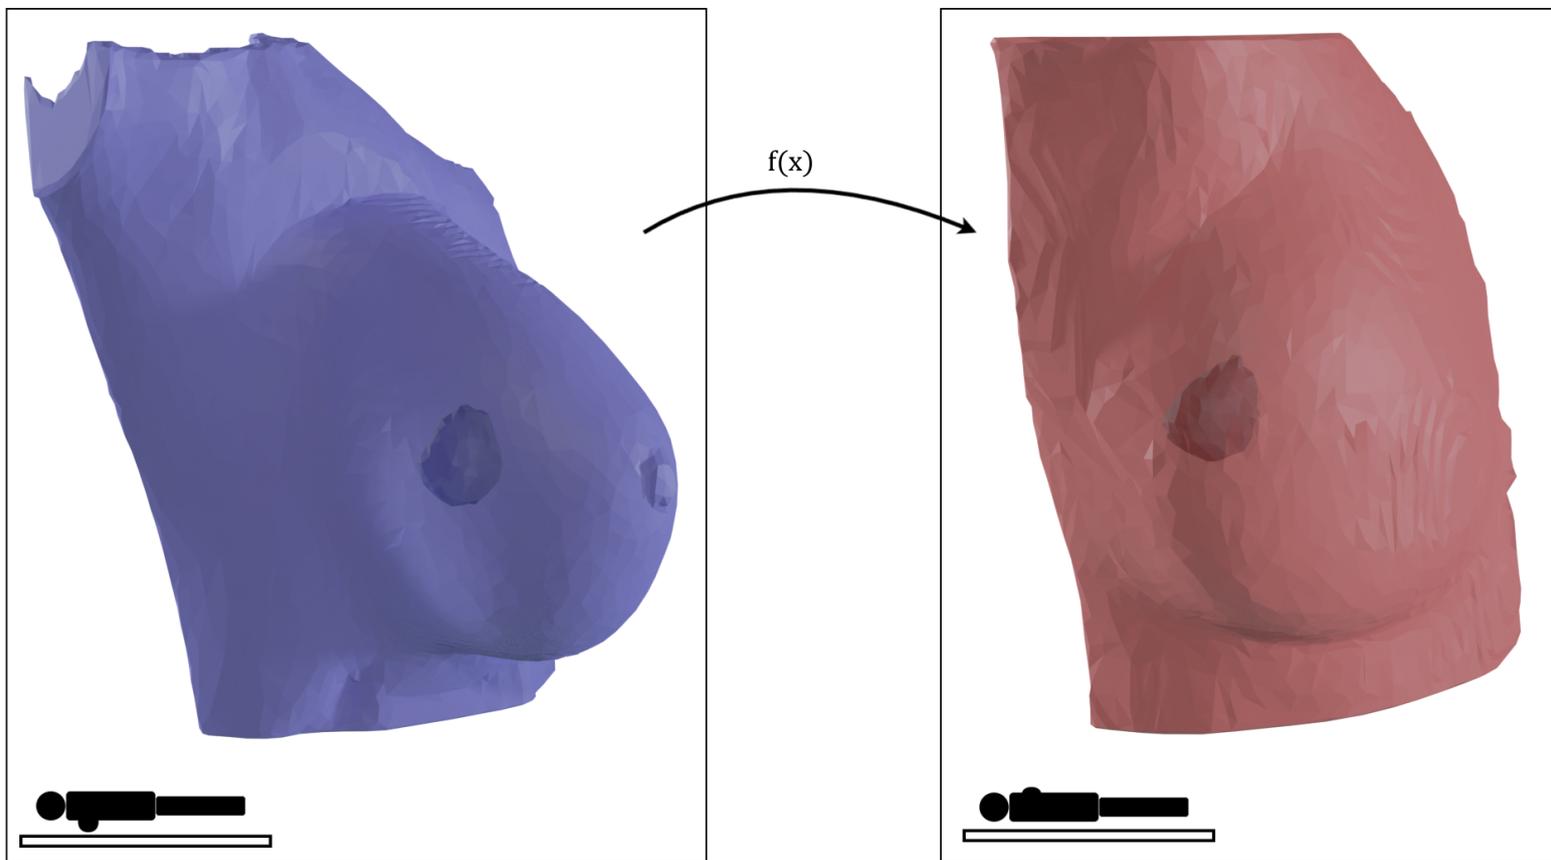


Prone-to-supine Breast Tumor Registration Using Landmarks

H.M. Yeh



Prone-to-supine Breast Tumor Registration Using Landmarks

by

H.M. Yeh

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday April 20, 2023 at 10:00 AM.

Student number:	4386027	
Thesis committee:	Prof. Dr. E. Eisemann,	TU Delft
	Dr. K. Hildebrandt,	TU Delft, supervisor
	Dr. P. Ambrosini,	Erasmus MC, supervisor
	Dr. Xucong Zhang	TU Delft, independent member

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

During the preoperative planning for breast-conserving surgery, the surgeon makes use of an MRI scan of the breast cancer patient in the prone position to accurately locate the tumour. However, surgery is performed with the patient in the supine position. The surgeon needs to mentally translate the location of the tumour from the prone position towards the supine position. The usage of mixed reality systems in the form of the Microsoft HoloLens could visually aid the surgeon by projecting the tumour in the supine position onto the patient. This requires us to obtain the tumour in the supine position by dealing with the prone-to-supine breast image registration problem.

We propose to maximize the overlap between the prone and supine surfaces of the breast through the use of soft constraints to guide the deformation. An initial subspace FEM simulation pulls the prone and supine surfaces towards each other at specified landmark correspondences between the surfaces. The resulting prone surface is projected onto the supine surface and these surface vertices are restricted to tangential movement across the surface of the breast. The tissues at the interior of the breast are deformed accordingly.

Our experiments with synthesized data indicate the robustness of our method under various scenarios, except for the rigid misalignment between the prone and supine meshes. This problem of misalignment also affects the experiments with real-world data, which results in distances greater than 12 mm between the centroids of the resulting deformed tumour and the ground-truth tumour.

Contents

1	Introduction	1
2	Related Works	3
2.1	Deformable Simulation	3
2.2	Subspace Simulation	4
2.3	Prone to supine breast image registration	5
3	Background	7
3.1	Finite Element Method	7
3.2	Newton's method	8
3.3	Constraints	8
3.4	Subspace Simulation	9
4	Method	10
4.1	Overview of the method.	10
4.2	Preprocessing.	10
4.3	Landmark guided deformation	12
4.4	Surface guided deformation	13
5	Experiments	15
5.1	Implementation	15
5.2	Experiments with synthetic data	15
5.2.1	Experiment Setup	15
5.2.2	Practical Settings.	16
5.2.3	Rest Shape	18
5.2.4	Rotated Gravity	20
5.2.5	Misalignment	21
5.2.6	Material Parameters	24
5.2.7	Landmark Correspondence Weight	25
5.2.8	Landmarks.	26
5.3	Phantom Breast Experiments	26
5.4	Real-World Experiments	27
6	Discussion	31
7	Conclusion	33

1

Introduction

Breast cancer is a prevalent illness affecting women worldwide. The treatment plan for breast cancer commonly consists of: chemotherapy to reduce the size of the tumours; surgery to excise the tumours from the healthy tissue; radiotherapy to eradicate remaining tumour cells and reduce the risk of recurrence. Breast cancer surgery is divided into two different types, namely mastectomy and lumpectomy. Mastectomy removes the entire breast, whereas lumpectomy attempts to preserve the shape of the breast by only removing the tumours and a margin of healthy tissue. If the margin was not appropriately chosen, then the patient might have to undergo additional surgery. Therefore, to accurately locate the tumours for the preoperative planning of lumpectomy, it is important that the surgeon uses appropriate imaging modalities such as e.g., mammography, Positron Emission Tomography-Computed Tomography (PET-CT) and Magnetic resonance imaging (MRI). MRI scan using breast coils with the patient lying flat on the stomach (prone position) offers the most accurate image to localize the tumour (C.-B. Wang et al., 2020). However, the surgery is performed with the patient lying on the back (supine position). There are significant differences between the prone and supine positions due to large deformations of the breast. This makes it a challenging problem for the surgeon, who needs to mentally translate the MRI from the prone position to the supine position.

The surrounding context regarding this Master's project is focused on visually aiding the surgeon during lumpectomy with the use of mixed reality systems. For example, Perkins et al. proposed an augmented reality method using the HoloLens mixed-reality glasses (Microsoft, USA) (Perkins et al., 2017). The surgeon wearing the glasses could see the MRI projection of the breast and the tumours onto the actual breast of the patient in real-time. In their study, they rigidly align the MRI scan acquired in the supine position with the patient in surgery using qr-code and MR-visible fiducial markers attached to the patient's skin around the breast. Deformation of the breast was not taken into account. It is also not standard protocol for a hospital to acquire the MRI scan in the supine position instead of the prone position with the breast coils. In this Master's project, we would like to focus on the breast deformation and alignment between a scan acquired in the prone position and the patient during surgery in the supine position. There are three practical settings to help tackle this problem: the first setting has an additional 3D surface scan of the breast in the supine position to perform the registration; the second setting has landmark correspondences between the prone and supine scans; the final setting has both the surface scan as well as the landmark correspondences. The 3D surface scan of the breast in the supine position could be obtained from the HoloLens depth map sensor. However, the acquisition of the data would require administrative paperwork, patient consent and timely acquisition schedule. In order to only focus on the deformable registration, we use already available PET-CT scans of the breast acquired both in the prone and supine position during breast chemotherapy (H. Wang & Mao, 2020). The results of our method are evaluated by comparing the Euclidean distance between the centroids of the deformed tumour and the ground-truth tumour from the PET-CT scan in the supine position.

The problem of prone-to-supine image registration of the breast is commonly tackled by physically simulating the breast in the prone position using the Finite Element Method (FEM). The simulation simplifies the real-world forces acting upon the breast to the gravitational forces and elastic forces of the breast. This simplification disregards unknown constraints, such as the position of the arm. The material model and its parameters define the elastic behaviour of the breast during the simulation, which should properly reflect the characteristics of the breast in the reality. The elasticity of the breast differs from patient to patient and can also change over time (Hipwell et al., 2016). The structure of the breast consists of different types of tis-

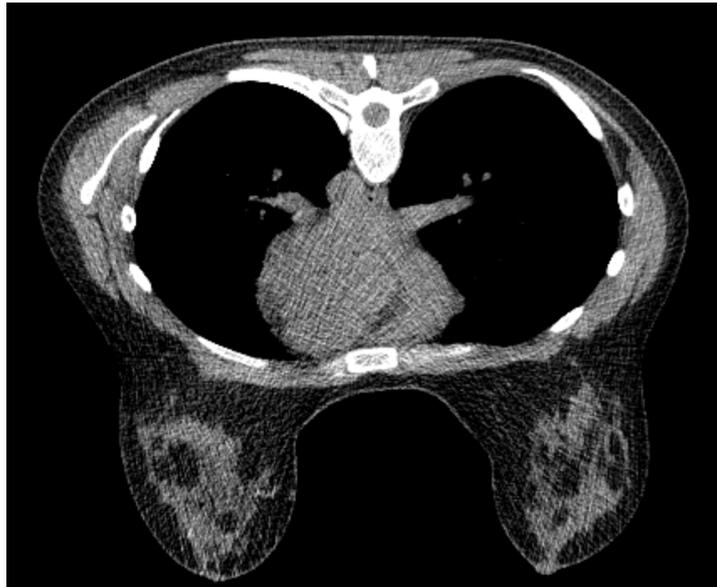


Figure 1.1: The PET-CT scan of patient 1 in the prone position, where the luminance values indicate the different tissue types.

sues (see Fig. 1.1). However, the segmentation of those different tissue regions is a difficult problem and will not be considered in this thesis. Therefore the material of the breast is assumed to be homogeneous. Besides that, the rest shape of the breast is unknown for the FEM simulation. The rest shape is the shape of the breast without any forces affecting it. All of these assumptions and challenges make it difficult to solve the prone-to-supine image registration.

In our method, we propose to maximize the overlap between the breast skin surface mesh in the prone and supine positions to obtain a sufficiently accurate localization of the tumour for clinical visualization on the Microsoft HoloLens. The method consists of a landmark guided deformation step and a surface guided deformation step. The landmark guided deformation uses landmark correspondences between the prone and supine meshes to guide the regions of the prone surface closer towards the supine surface. This is used as an initialization for the surface guided deformation, where the surface of the prone mesh is projected onto the surface of the supine mesh. Afterwards, the projected surface is aligned and the interior of the breast is deformed accordingly.

The report will discuss and compare the related works in Chapter 2. In Chapter 3, we will provide the background information to understand our approach and, our method will be then explained in Chapter 4. The setup and results of the experiments are discussed in Chapter 5. Finally, Chapter 6 will discuss the insights and guidelines gained from the Master's project.

2

Related Works

2.1. Deformable Simulation

The physical simulation of deformable objects is commonly performed with the Finite Element Method (FEM) (Sifakis & Barbic, 2012). It splits the deformable object into a finite number of elements. Each finite element uses a material model to compute its elastic energy, which is afterwards combined using a weighted sum to obtain the total elastic energy of the entire object. The quasi-static simulation of the object is solved by minimizing the total elastic energy subject to constraints. Additionally, inertial effects can be introduced to simulate the dynamics of an object using a numerical time integration scheme (e.g., implicit Euler). The time integration scheme results into a system of equations, which needs to be solved to obtain the solution for each timestep. In the particular case of the implicit Euler scheme, the system of equations can be reformulated as an optimization problem (Gast et al., 2015). An accurate solution to the optimization problem can be obtained with the application of Newton’s method. It uses the Hessian and gradient to produce a descent direction to search for a solution. The stability of the simulation requires the Hessian matrix to be positive semi-definite (PSD). The Hessian matrix is computed as the sum of the local Hessian matrices for each finite element. If all elementary Hessian matrices are positive-definite, then the constructed Hessian matrix is also positive-definite. Consequently, the Hessian matrix can become indefinite due to summing an indefinite elementary Hessian from an inverted or degenerate finite element.

For the stability of the simulation, an indefinite Hessian matrix can be replaced with a positive definite matrix close to the true Hessian matrix. A popular method is to project the local Hessian matrix for each finite element to become PSD (Irving et al., 2004; Teran et al., 2005). It modifies the deformation gradient by inverting its smallest singular value. The modified deformation gradient is used to compute a positive definite Hessian matrix. The material model used in the FEM simulation does need to be able to deal with flattening or inversion of the finite element. With the definition of a new set of invariants and the closed-form formulas for the eigensystems of arbitrary isotropic material models (Smith et al., 2019), a PSD Hessian matrix can be analytically constructed for any isotropic material model. The approach to fix the definiteness of the Hessian matrix has been applied to a number of different applications. For the application to simulate contact between deformable objects, Incremental Potential Contact (M. Li et al., 2020) introduces local barrier functions to create repulsion forces between contacting finite elements. The physical simulation for cutting deformable objects can model the cut by changing the topology of the object (J. Wu et al., 2015). An alternative approach is to use the extended FEM (Koschier et al., 2017) to model the cut using discontinuous enrichment functions without affecting the topology of the object. Another application is the simulation of elastoplastic objects (X. Li et al., 2022; B. Wang et al., 2021), where the deformation gradient is the product of the plastic and elastic deformation gradients. The plastic deformation changes the rest configuration, whereas the elastic deformation produces elastic forces to recover the rest configuration.

The material model defines the behaviour of the deformable object to any type of deformation by designing a function, which relates the stress and strain in an appropriate manner for the chosen application (Kim & Eberle, 2020). A popular material model in the field of Computer Graphics for its performance is the linear corotated material (Kugelstadt et al., 2018). The function consists of a term to preserve the lengths of the edges between vertices, and a linearized term to preserve the volume of the finite element. This results in an efficient approximation for small deformations. However, for larger deformations the volume can grow

or shrink substantially. The Projective Dynamics framework (Bouaziz et al., 2014) defines its Projective Dynamics material as a function with a nonlinear constraint projection and a quadratic distance measure. The formulation of the material model lends the local-global solver of Projective Dynamics its performance. The existing hyperelastic materials (e.g., Neo-Hookean, St. Venant–Kirchhoff and Mooney–Rivlin) can be applied to the deformable object for more realistic deformation. These materials are commonly used in the engineering literature. The Neo-Hookean material has been redesigned to become stable and able to deal with invertible elements (Smith et al., 2018). This stable version of the Neo-Hookean material consists of a stretching term, invertible nonlinear volume preservation term and a term for stability of the rest configuration.

A different approach to design a material is to allow for interactive editing of the material model using the principal stretches (Xu, Sin, et al., 2015). This method formulates the interactive material model according to the Valanis–Landel hypothesis into three elastic energy functions expressed in the principal stretches. The first derivative of these functions can then be edited to design the deformation behaviour of the material.

A material model can define parameters to weight the importance of specific terms in the function (e.g., stretching and volume preservation), which can heavily affect the deformation behaviour of the object. In order to move closer to the deformation behaviour in the real-world situation, the material parameters can be optimized based on given target forces and positions (Xu, Li, et al., 2015). The method uses a model reduction of the material distribution space for interactive performance. A different approach to capture the realistic deformation behaviour uses an alternating optimization between the rest configuration and the material parameters (B. Wang et al., 2015). The optimization of the material parameters uses manually specified control points to linearly blend the material parameters.

2.2. Subspace Simulation

Model reduction (subspace simulation) projects the high-dimensional space of the simulation into a properly chosen low-dimensional subspace (Sifakis & Barbic, 2012). The projection should ensure that the subspace simulation sufficiently approximates the full space simulation. There is thus a need to construct an appropriate subspace basis for the projection.

The global deformation of the object can be taken into account for the construction of the subspace basis. Linear Modal Analysis (LMA) (Sifakis & Barbic, 2012) constructs the subspace basis from the low-frequency vibration modes. These linear modes show the deformation with the lowest increase in elastic strain energy. LMA provides an appropriate basis for small deformations from the rest configuration. In the case of large deformation, a subspace basis can be constructed with the linear vibration modes and their modal derivatives (Barbič & James, 2005). The modal derivatives show deformations that co-appear in a nonlinear system when the system is excited in the direction of the linear modes. Principal Component Analysis is used on the linear vibration modes and their modal derivatives in order to reduce the dimensionality of the subspace.

With a localized approach to the subspace construction, the subspace simulation can handle more spatially localized deformations. The linear blend skinning space (Brandt et al., 2018) consists of the affine transformations of a number of sampled points, which uses a radial basis function to indicate its local influence to the rest of the points. This subspace basis is used in our method to deform localized landmark regions. A different approach to add local deformations to the subspace basis starts with the global subspace basis for modal derivatives, and during the simulation it updates the basis by introducing additional basis functions for local deformations (Harmon & Zorin, 2013). These basis functions are computed for the regions of the object that are experiencing external forces. It does not work well if the region under load increases in size. Domain decomposition has been used in combination with subspace simulations (Kim & James, 2011; X. Wu et al., 2015). It partitions the object into multiple domains with its own (global) subspace basis. The partitioned domains are coupled with the use of elastic coupling forces. However, this can still lead to discontinuities among multiple domains due to separate simulation in their own subspace.

Data-driven methods for the subspace simulation require the generation of appropriate training data from the full space simulation. Cubature approximation (An et al., 2008) is used to approximate the reduced internal forces and stiffness matrix. The cubature weights are optimized using nonnegative least squares on the training data. A common subspace basis can be obtained with the application of mass-PCA on the training data (Sifakis & Barbic, 2012). An extension to this subspace basis attaches the nonlinear autoencoder after the mass-PCA subspace in order to introduce nonlinearity to the subspace basis (Fulton et al., 2019). However, the autoencoder subspace can cause global motion from locally applied forces.

2.3. Prone to supine breast image registration

The tumour of a breast cancer patient in the supine position needs to be accurately located in order plan the surgery. The MRI scan of the breast in the prone position gives the most accurate location for the tumour. The registration problem is thus concerned with transforming the breast in the prone position to the supine position. An overview of the different methods discussed in this section is shown in Table 2.1. The initial approach commonly used to deal with the problem of prone-to-supine breast image registration is to use a FEM simulation of the breast in the prone position (Babarenda Gamage et al., 2019; Bessa et al., 2020; Danch-Wierzchowska et al., 2020; Eiben et al., 2016; Hipwell et al., 2016). The FEM simulation has gravitational forces acting on the deformable breast, while constraining the vertices of the chest wall to be fixed. There are a number of unknowns in the FEM simulation, which are approximated using different methods in the related works.

A deformable object in a FEM simulation always tries to recover back to its rest configuration, where there are no forces acting on the object. However, the MRI scan of the breast in the prone position does not correspond to the actual rest configuration of the breast due to the gravitational forces. An iterative gravity unloading method has been proposed (Bessa et al., 2020; Danch-Wierzchowska et al., 2020; Eiben et al., 2016; Hipwell et al., 2016) to approximate the rest configuration of the breast by simulating the prone mesh with gravity acting in the inverse direction. The resulting rest configuration is then simulated with gravity acting back towards the prone position, and the difference between the simulated and actual prone mesh is iteratively used to update the approximation of the rest configuration. This method is sensitive for the choice of material parameters for the simulations. On the other hand, our method uses the prone configuration of the breast as its rest shape. The guidance of landmark correspondences counteracts the elastic forces that pull the breast back to its prone configuration.

The parameters for the material model to physically simulate the behaviour of the breast are unknown and can also differ from patient to patient. A simple method to optimize the material parameters is to use a constant multiplicative factor based on an image similarity measure (Eiben et al., 2016). A classification into different size classes can help with identifying the optimal material parameters into specific ranges (Danch-Wierzchowska et al., 2020). The material parameters can also be optimized to become spatially varying. With the knowledge that stiffness parameters are higher at the posterior region of the torso than the breast tissue, a Laplacian problem can be solved to allow for a smooth transition between these different stiffness parameters (Babarenda Gamage et al., 2019). For the FEM simulations in our method, we have chosen to keep the material model and its parameters fixed. The method puts more emphasis on the soft constraints to guide the transformation from prone position to the supine position.

Even with the aforementioned improvements to the physical simulation of the breast, it cannot fully capture the deformation of the breast from the prone position to the supine position. The related works have instead looked into the application of additional methods onto the result of the FEM simulation. The image similarity of the prone and supine scans is used to derive forces to deal with any differences between both scans, which could not be modeled using only the known physical forces (Eiben et al., 2016). On a similar note, our method uses constraint forces to deal with the differences between the prone and supine scans. Another method to tackle the problem is to apply Free Form Deformation (FFD) with B-splines onto the result of the FEM simulation. FFD uses a distance measure between the result of the simulation and the supine scan to determine the final deformation of the prone mesh (Bessa et al., 2020). The application of FFD can result into unrealistic deformation of the breast. In comparison, our method projects the prone surface onto the supine surface and uses the physical simulation for the interior of the breast. For the automatic construction of tetrahedral meshes for individual patients, a Convolutional Neural Network (CNN) is trained to segment the different tissue types (skin, ribs, lungs) from the prone scan (Babarenda Gamage et al., 2019). The segmented image data is then processed with the use of Statistical Shape Modelling (SSM) in order to construct personalized tetrahedral meshes for each patient. This method allows for a large variation of breast shapes in the FEM simulations. However, our focus is placed on the registration task, and therefore does not make use of the automatic construction of personalized tetrahedral meshes. A geometric approach to the problem uses Laplacian editing in order to constrain a set of landmark correspondences onto the positions in the supine position (Alfano et al., 2019). It uses nonrigid iterative closest point (ICP) algorithm to project the entire surface of the prone scan onto the surface of the supine scan. The resulting transformation gives a dense point to point matching of the two surfaces, and this is used to perform FFD with B-splines to deform the interior of the breast. This approach is similar to our method in the use of landmark correspondences to pull the prone and supine meshes closer together, and performing a surface projection onto the supine mesh. However, our method uses FEM simulations to physically simulate the behaviour of the breast and uses the

Method	Imaging Data	Single or Both breasts	Results
FEM + FFD B-splines (Bessa et al., 2020)	MRI + 3D scan	Single	TRE markers 18.50 ± 3.88 mm
FEM + material optimization + image derived forces (Eiben et al., 2016)	MRI + CT	Single	TRE markers CT 7.3 mm, MRI 6.4 mm
Laplacian editing + nonrigid ICP + FFD B-splines (Alfano et al., 2019)	MRI + CT	Both	Tumour Centroid Distance 10.40 mm
CNN and SSM for FE model + FEM + Laplacian material optimization (Babarenda Gamage et al., 2019)	MRI	Single	Nipple 13mm, Internal tissues $7\text{mm} \pm 2\text{mm}$

Table 2.1: Summary of the different prone-to-supine breast image registration methods. The evaluation for the methods use the Target Registration Error (TRE) of markers placed on the surface of the breast to indicate its performance. The tumour centroid distance is the distance between the centroid of the deformed tumour and the ground-truth tumour in the supine position. The last row in the results indicates the displacement of the nipple and its internal tissues.

landmarks and surface projection to guide the optimization of the breast.

3

Background

3.1. Finite Element Method

The Finite Element Method (FEM) is commonly used for the physical simulation of a deformable object. FEM discretizes the deformable object into a finite number of elements (usually triangles for 2D and tetrahedrons for 3D). Each finite element stores the position of the vertices in the rest configuration X . By deforming the element, the positions of the vertices are transformed from the rest configuration X to a deformed configuration x following the deformation function:

$$\phi(X) = FX + b = x, \quad (3.1)$$

where F is the deformation gradient matrix and b is the rigid translation of the finite element (see Fig. 3.1). The deformation gradient F stores rotation and deformation (stretching and shearing) information of the finite element.

A material model function $\psi(F)$ uses the deformation gradient F to determine the amount (and direction) of deformation and relates it with a corresponding scalar amount of elastic energy. An example of a simple material model function is a linear spring: $\psi_{linear\ spring}(L) = \frac{k}{2}(L - l)^2$, where L is the length of the spring in the rest configuration, l is the length of the spring in the deformed configuration, and k is the stiffness of the spring. This material model shows a linear relationship between the amount of deformation and the elastic energy of the spring.

A common choice for a material model to represent the deformable breast in the prone-to-supine breast image registration is the NeoHookean material (Babarenda Gamage et al., 2019; Eiben et al., 2016; Hipwell et al., 2016). The formulation of the original NeoHookean material introduces a singularity, which explodes the FEM simulation during squashing. The stable NeoHookean material (Smith et al., 2018) offers a different formulation that does not contain the singularity, and allows degenerate or inverted elements to recover to a valid configuration. The material model is defined using the invariants shown in Equations 3.2, 3.3, and

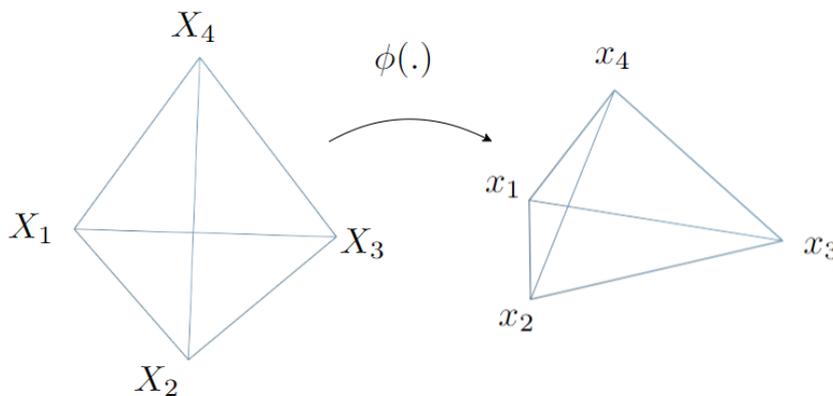


Figure 3.1: The rest configuration X (left) and deformed x (right) configuration of a tetrahedron

3.4. The invariants use the deformation gradient F to indicate a measure for the amount of deformation by removing the rotational part of the deformation gradient F . A simpler formulation of the invariants uses the singular values $\sigma_1, \sigma_2, \sigma_3$ of the right Cauchy–Green strain tensor $F^T F$.

$$I_1 = \text{tr}(F^T F) = \sigma_1 + \sigma_2 + \sigma_3 \quad (3.2)$$

$$I_2 = \text{tr}(F^T F F^T F) = \sigma_1^2 + \sigma_2^2 + \sigma_3^2 \quad (3.3)$$

$$I_3 = \det(F^T F) = \sigma_1 \sigma_2 \sigma_3 \quad (3.4)$$

The stable NeoHookean material is defined as:

$$\psi = \frac{\mu}{2}(I_2 - 3) - \mu(I_3 - 1) + \frac{\lambda}{2}(I_3 - 1)^2 \quad (3.5)$$

It consists of three weighted terms. The edge length term $\frac{\mu}{2}(I_2 - 3)$ pulls the edges of a tetrahedron closer until it collapses into a point. The volume term $\frac{\lambda}{2}(I_3 - 1)^2$ preserves the volume of the tetrahedron, and it is invertible and well-defined everywhere. The combination of the previous terms moves the elastic energy of the rest configuration away from 0 (i.e., $\psi(X) \neq 0$). The elastic energy of 0 coincides with a volume of $\det(F) < 1$. The last term $-\mu(I_3 - 1)$ resolves this by inflating the tetrahedron in such a way that the tetrahedron deflates into the rest configuration at $\det(F) = 1$.

The total elastic energy for a deformable object is computed as the sum of the elastic energy for each finite element weighted by its contribution to the whole object (areas for 2D and volumes for 3D). The equation for a tetrahedral mesh is shown in Equation 3.6. The elastic forces are computed as $f_{elastic} = -\nabla E(x)$.

$$E(x) = \sum_{i \in \text{tetrahedrons}} \text{volume}_i \cdot \psi(F_i) \quad (3.6)$$

For quasi-static problems, Equation 3.6 can be minimized including any additional energy terms (e.g., gravitational energy or penalty functions), which will result in the object attempting to recover back to its rest configuration while satisfying any additional constraints. Newton's method is used to solve the resulting equation, and is described in further detail in Section 3.2.

To simulate the dynamics for a deformable object, time needs to be discretized into different time samples and integrated. An example for numerical time integration is the implicit Euler method, for which the resulting update equations are given in Equation 3.7.

$$\begin{aligned} x_{n+1} &= x_n + h v_{n+1} \\ v_{n+1} &= v_n + h M^{-1} (f_{elastic}(x_{n+1}) + f_{ext}), \end{aligned} \quad (3.7)$$

where h is the simulation step size, v is the velocity, M is the mass matrix and f_{ext} are any external forces.

3.2. Newton's method

The quasi-static FEM simulation can be formulated as an optimization problem with the energy function $E(x) = E_{elastic} + E_{gravity} + E_{constraints}$. Newton's method can be used to iteratively solve such optimization problems. In each iteration, the gradient g and the Hessian H of the energy function E are used to compute the Newton descent direction $d = -H^{-1}g$. If the symmetric Hessian H is positive definite, then d is guaranteed to be a descent direction. The choice of using the stable NeoHookean material allows us to compute a positive definite Hessian matrix from the analytical formulas of its eigendecomposition (Smith et al., 2018). Any negative eigenvalues of the Hessian are clamped to either 0 or a small positive number. A line search is commonly performed along the descent direction d to find a suitable step length α that sufficiently reduces the energy function E , which improves the convergence of the method. The resulting step length α and descent direction d are combined to update the positions of the vertices $x = x + \alpha d$. These steps are repeated until convergence of the method or a maximal number of iterations has been reached.

3.3. Constraints

The FEM simulation can be constrained to model any intended behaviour of the deformable object. The types of constraints are categorized as either hard or soft constraints (Nocedal & Wright, 1999), and are solved in different ways.

Hard constraints must always be satisfied to obtain a feasible solution. An example of a hard constraint is fixing the position of a vertex. To find a feasible solution for this specific type of hard constraint, the quasi-static FEM simulation can simply remove the degrees of freedom for the fixed vertices. A more general method to satisfy different types of hard constraints is to solve the Newton–Karush–Kuhn–Tucker system (Nocedal & Wright, 1999) described in Equation 3.8.

$$\begin{bmatrix} H & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} -d \\ \lambda \end{bmatrix} = \begin{bmatrix} g \\ b \end{bmatrix}, \quad (3.8)$$

where A is the selection matrix for the constraints, b is the value for the constraints, and λ is the Lagrange multipliers.

Soft constraints can be violated to obtain a feasible solution, but it incurs a penalty to the energy function corresponding to the severity of the violation. Therefore soft constraints commonly need to define their constraint penalty function. An example of a constraint penalty function $c(x)$ to fix the position of a vertex is $c(x) = \frac{w}{2} \|x - x_{target}\|^2$. This constraint penalty function acts as a spring-type energy to pull the current position of the vertex closer towards its initial position. The weight of a constraint penalty function w controls whether the optimization of the energy function allows for violating the constraint to incur the weighted penalty. In the case that the weight w approaches infinity, then the soft constraint will act as a hard constraint.

3.4. Subspace Simulation

The system of equations of the quasi-static FEM simulation can be simplified by applying the model reduction technique (Sifakis & Barbic, 2012). It approximates the solution of the original high-dimensional space by projecting from the original space of the deformable object onto a properly chosen lower-dimensional subspace. The subspace matrix U maps the subspace onto the original space (prolongation operator), and U^T maps the original space onto the subspace (restriction operator). This means that $U^T U = I$ should hold for the chosen subspace matrix U . The vertex positions in the original space can then be defined as $u = Ux$.

An initial choice for a subspace matrix U can be obtained from the linear vibration modes (Linear Modal Analysis). The linear vibration modes are obtained from solving the generalized eigenvalue problem in Equation 3.9. After ordering the eigenvectors e based on ascending eigenvalues λ , then the first k eigenvectors can be used to construct the subspace matrix $U = (e_1 | \dots | e_k)$

$$He_i = \lambda_i Me_i, \quad (3.9)$$

where H is the Hessian matrix in the rest configuration, M is the diagonal mass matrix, λ is the i -th eigenvalue, and e is the i -th eigenvector.

A different choice for the construction of the subspace matrix U is the linear skinning subspace (Brandt et al., 2018; Nasikun & Hildebrandt, 2022), which is relatively inexpensive to construct in comparison to the linear vibration mode subspace. It samples k vertices of the deformable object using farthest point sampling. A Radial Basis Function (RBF) is centered around each sampled vertex with a specified falloff radius, and assigns a weight w to each vertex. The assigned weights w are then used with the positions of the vertices x to construct the linear skinning subspace U in Equations 3.10 and 3.11.

$$U_j = \begin{bmatrix} w_0^j & x_0^x \cdot w_0^j & x_0^y \cdot w_0^j & x_0^z \cdot w_0^j \\ w_1^j & x_1^x \cdot w_1^j & x_1^y \cdot w_1^j & x_1^z \cdot w_1^j \\ \dots & & & \end{bmatrix}, \quad (3.10)$$

where w_i^j is the weight of vertex j for RBF i , and x_i^j is coordinate number $j \in (x, y, z)$ of vertex i .

$$U = (U_1 | \dots | U_k) \quad (3.11)$$

4

Method

4.1. Overview of the method

Our method optimizes a nonrigid mapping of the breast in the prone position (prone mesh) onto the breast in the supine position (supine mesh) in order to maximize the overlap between the surface of both meshes. The resulting nonrigid mapping should offer a good approximation for the mapping of the tumour from prone to supine.

The MRI or PET-CT scan in the prone position is semi-automatically segmented to extract the triangle meshes of the surfaces of the breast, the tumour, and the remaining tissue as shown in Figure 4.1. Additionally, we manually segment the single breast containing the tumour from these meshes. As a consequence of substituting a 3D surface scan of the breast in the supine position with the available PET-CT scan in the supine position, the same steps need to be taken to extract the triangle meshes for the supine breast and tumour. Cancer cells are generally also drained to the sentinel lymph nodes in the armpit. This results in tumours being located in the armpit. Since our method does not include simulating the deformable tissue of the armpit, those tumour parts are disregarded in the application of our method.

In Figure 4.2, an overview of the method is illustrated. The method takes as input the triangle mesh of the breast in the prone position, the triangle mesh of the breast skin surface in the supine position, a list of vertices representing the chest wall for the prone mesh, and a correspondence of manually annotated landmark vertices between both surfaces. For the preprocessing of our method, the prone mesh is tetrahedralized for use in the FEM simulations. The tumour is implicitly mapped using barycentric coordinates onto the tetrahedrons of the tetrahedralized prone mesh. The prone mesh and supine surface mesh are rigidly aligned using manual alignment and/or Iterative Closest Points method. After the rigid alignment, the landmark guided deformation step pulls the landmark regions closer together in order to maximize the overlap between the prone and supine surfaces. This serves as an initialization for the surface guided deformation step, where the surface of the deformed prone mesh is projected onto the supine surface. Afterwards, the interior of the prone mesh is optimized in a FEM simulation while keeping the projected surface vertices fixed.

4.2. Preprocessing

Tetrahedralization Tetrahedralization converts the triangle mesh of the breast in the prone position into a tetrahedral mesh. The tetrahedral mesh can then be used for the FEM simulation of the breast. The quality of the tetrahedralization (e.g., equally sized tetrahedrons) can affect the accuracy of the FEM simulation. The input for our method uses a list of vertex indices to define the landmark correspondences as well as the fixed points of the chest wall. In order to ensure that these vertex indices stay consistent, the tetrahedralization should construct tetrahedrons using existing vertices and add vertices when necessary.

Tumour barycentric mapping The application of the FEM simulation onto the tetrahedral mesh of the breast results in the deformation of the breast. However, this does not include the deformation of the tumour. The triangle meshes of the breast and the tumour could be explicitly tetrahedralized together. The quality of this tetrahedral mesh would be reduced, and could impact the accuracy of the FEM simulation. Since the constructed tetrahedral mesh would have smaller tetrahedrons located at the tumour and bigger tetrahedrons elsewhere. The material of the breast is chosen to be modeled to be homogeneous in order to

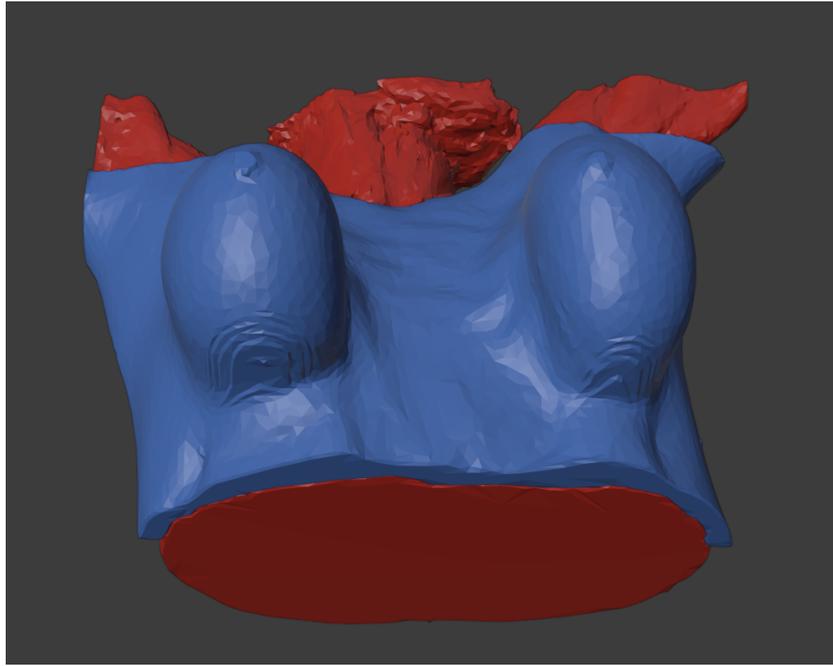


Figure 4.1: Segmentation of a breast CT scan. The breast and its surrounding tissues (blue) is separated from the rest of the chest (red).

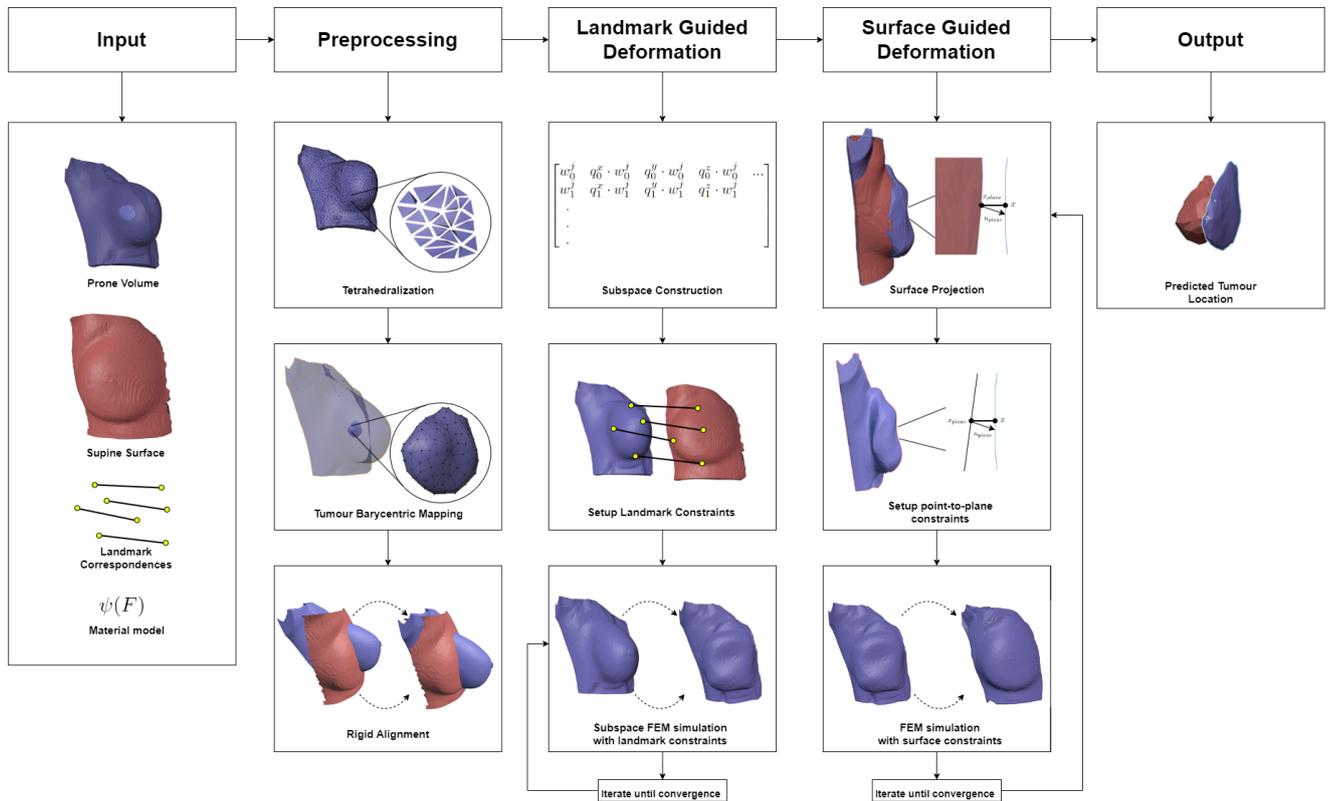


Figure 4.2: An overview of our method

avoid the problem of tissue segmentation. The entire breast will have the same material and parameters applied to it. Therefore we have chosen to implicitly map the geometry of the tumour onto the tetrahedrons of the tetrahedral mesh using barycentric mapping. The implicit mapping can recover the deformation of the tumour from the result of the FEM simulation.

Rigid Alignment The subsequent steps of our method require a good rigid alignment between the prone mesh and supine surface mesh. The rigid alignment is needed in order to find the appropriate positions of the corresponding landmark vertices. Besides that, the rigid chest wall restricts the deformation of the tumour. This could be problematic if the chest wall of the prone mesh and the chest wall of the patient in the supine position do not line up due to bad rigid alignment. It could become impossible for the tumour of the prone mesh to deform to the actual location of the tumour in the supine position. The rigid alignment is performed either manually and/or with Iterative Closest Points.

4.3. Landmark guided deformation

The shape of the prone and supine surfaces of the breast are very different from each other. The landmark guided deformation pulls the prone surface closer towards the supine surface in order to provide an initialization for the surface guided deformation (see Section 4.4). The sparse set of landmark correspondences pulls the regions of the prone mesh surrounding each landmark closer towards the corresponding landmark of the supine mesh. The steps for the landmark guided deformation are shown in Algorithm 1.

Subspace construction The linear skinning subspace is used to apply model reduction to the FEM simulation (see Sec. 3.4). The application of the subspace is used to suppress high-frequency deformations. Any deformation applied to a single vertex will also affect the neighbouring vertices.

The subspace chosen for our method is the linear skinning subspace (Brandt et al., 2018) (see Subsection 3.4). In contrast to the construction of the subspace in Brandt et al., we have chosen to use Euclidean distances instead of geodesic distances to perform farthest point sampling due to the convex geometry of the breast. The sampling process takes the vertices from the landmark correspondences as the initial sampled points. As a result, the subspace FEM simulation becomes deterministic. The radius for the subspace is computed with the adapted heuristic in Equation 4.1 (Nasikun & Hildebrandt, 2022).

$$r = \sqrt{\frac{\sigma \cdot volume}{number\ of\ samples \cdot \pi}}, \quad (4.1)$$

where the area of the surface in the original formulation is changed for the volume of the mesh. An alternative choice for the subspace could be the subspace constructed from the linear vibration modes. However, this requires us to solve an expensive generalized eigenvalue problem.

Setup landmark constraints The landmark constraints need to pull the regions of the prone mesh surrounding each landmark closer towards the corresponding landmark of the supine surface mesh. The pulling forces are modeled using soft penalty energy constraints in the following form:

$$c(x) = \frac{w}{2} \|x - x_{target}\|^2, \quad (4.2)$$

where w is the weight of the constraint, x is the position of the vertex and x_{target} is the target position. The defined constraint pulls only the landmark vertex of the prone mesh towards the corresponding landmark vertex of the supine mesh in the full-space FEM simulation. The result of using the landmark constraints is shown in Figure 4.3a. It mostly affects the landmark vertices instead of the surrounding region of vertices in order to avoid high energy states (Braess, 2007).

A solution for pulling landmark regions is to rigidly transform the neighbouring vertices along with the landmark vertex as a rigid frame. A radial basis function assigns a weight between 0 and 1 to each neighbouring vertex depending on its distance from the landmark vertex. This can then be used to setup the same type of constraint energy as described in Equation 4.2 for each vertex inside a landmark region. This approach requires the user to determine the number of neighbouring rings that should be inside each rigid frame. It can also become problematic when landmark regions start to overlap, since each vertex inside the overlapping region will have multiple spring constraints to pull the vertex to different target positions. Therefore we have opted to instead use the landmark constraints in the constructed subspace. The subspace acts as a filter to avoid high-frequency deformations resulting from the point constraints. Figure 4.3b shows the result of applying the landmark constraints inside the subspace.

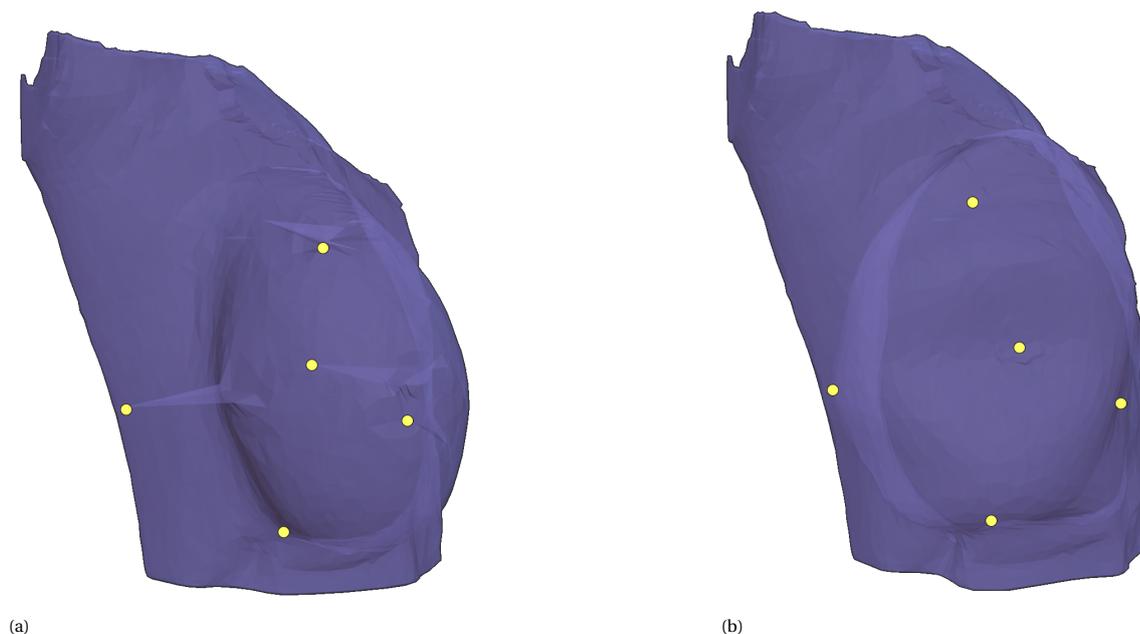


Figure 4.3: The landmark guided deformation method applied on patient 1 of the real-world experiments (a) with and (b) without the use of a subspace. Without the use of a subspace, the landmark constraint mainly affects the landmark vertex and not the region surrounding the vertex. This problem is resolved with the use of a subspace.

Subspace FEM simulation with landmark constraints The landmark FEM simulation needs to pull the landmark regions of the prone mesh and supine surface mesh closer together in order to serve as an initialization for the surface guided deformation. The energy function consists of the elastic energy, gravitational energy and the landmark constraints of the breast. This function is reduced with the application of the linear skinning subspace. The hard constraints for the fixed vertices at the chest wall are resolved by setting the corresponding Degrees Of Freedom (DOFs) in the subspace to 0. The FEM simulation iteratively uses Newton's method to compute the descent direction d and linesearch to find the optimal step length α . The subspace matrix U is used to project the descent direction d back to the full-space and update the vertex positions of the prone mesh.

Algorithm 1 Landmark guided deformation

Input: prone breast, chest wall vertices, landmarks

Output: prone breast

Construct subspace matrix U

Setup landmark constraints (Eqn. 4.2)

for each $v \in$ chest wall vertices **do**

Set subspace U DOFs to 0

end for

while $\|g\|^2 < 1e-3$ and the maximum number of iterations is not reached **do**

$d \leftarrow -(U^T H U)^{-1} U^T g$

$d \leftarrow U d$

$\alpha \leftarrow \text{Linesearch}(d)$

$q \leftarrow q + \alpha d$

end while

4.4. Surface guided deformation

The shape of the surface deformed prone mesh from the landmark registration should be closer towards the actual surface of the supine mesh, but there can still be differences between the surfaces. To improve the overlap between the surfaces, the surface of the prone mesh is projected onto the surface of the supine

surface mesh. Then the interior of the breast is optimized while keeping the surface vertices of the projected prone mesh restricted on the projected surface. The steps for the surface guided deformation are shown in Algorithm 2.

Surface projection The simplest way to maximize the overlap of both meshes is to project the surface of the prone mesh onto the surface of the supine mesh. The vertices in the interior and at the chest wall are disregarded for the surface projection. Each surface vertex of the prone mesh is projected onto the corresponding closest point on the surface of the supine mesh. The surface normal at the location of the projection point is also stored. This step can be sped up by using a mesh data structure, such as an Axis-Aligned Bounding Box.

Setup point-to-plane constraints The point-to-plane constraints in Equation 4.3 constrain the surface vertices of the prone mesh x to their projected closest points x_{plane} . The surface vertices are allowed to slide across the tangential plane of the projected point using the stored surface normal n_{plane} . This should prevent the emergence of high energy states from constraining the entire surface to specified points. On top of that, the soft fixed point constraint in Equation 4.2 is used to constrain the surface vertices x onto their projected closest points x_{plane} . The combination of both soft constraints should allow for some tangential movement, but it should keep the surface vertices from moving too far from the projected closest points.

$$c(x) = \frac{w}{2} ((x - x_{plane}) \cdot n_{plane})^2 \quad (4.3)$$

FEM simulation with surface constraints The interior FEM simulation keeps the projected surface fixed and optimizes the interior of the breast. The prone and supine surfaces should be very similar to each other, and this should result in a better approximation for the mapping of the tumour from prone to supine. The energy function consists of the elastic energy, gravitational energy and the point-to-plane and fixed point constraints of the breast. The hard constraints for fixing the vertices of the chest wall are resolved in the full-space by solving the Newton-KKT system to obtain feasible solutions. The result of this system is a Newton descent direction d , which is used in a linesearch to find an optimal step length α . The combination of the descent direction d and the step length α is used to update the vertex positions of the prone mesh. If the method has not converged, then it will return to the surface projection step.

Algorithm 2 Surface guided deformation

Input: prone breast, chest wall vertices, supine breast surface

Output: prone breast

```

while  $\|g\|^2 < 1e-3$  and the maximum number of iterations is not reached do
  for each  $x \in$  surface prone mesh do
     $x_{plane}, n_{plane} \leftarrow$  Project  $x$  onto surface of supine mesh
    Setup point constraint  $(x, x_{plane})$  (Eqn. 4.2)
    Setup point-to-plane constraint  $(x_{plane}, n_{plane})$  (Eqn.4.3)
  end for
   $d \leftarrow$  Newton-KKT(H, g) (Eqn. 3.8)
   $\alpha \leftarrow$  Linesearch( $d$ )
   $q \leftarrow q + \alpha d$ 
end while

```

5

Experiments

5.1. Implementation

The triangle meshes of the breast and tumour are initially provided as a result of a semi-automatic segmentation, based on region growing using manually annotated seeds for the breast, tumour and chest. The segmentation and triangle mesh reconstruction was done with the MeVisLab software. The tetrahedralization of the triangle mesh is performed with TetGen (Hang, 2015). The parameters used for TetGen keep the original triangle mesh, and insert edges and vertices to construct the tetrahedral mesh. For the quality of the tetrahedral mesh, the maximum allowable radius-edge ratio is set to 1.414. The FEM simulation is implemented in C++ (<https://github.com/hmyeh/BreastTumourRegistration>), and uses the libigl (Jacobson, Panozzo, et al., 2018) and Eigen (Guennebaud, Jacob, et al., 2010) libraries to solve the system of equations. The material models are implemented according to the Dynamic Deformables course (Kim & Eberle, 2020), which defines the analytic formulas for the gradient and Hessian for arbitrary isotropic materials. The analytic formulas do require the computation of simple derivatives. These are automatically computed using the autodiff library (Leal, 2018). Thus any isotropic material model can be used in the simulation by adding the elastic energy function $\psi(F)$.

For a more detailed explanation about FEM simulations, the reader is referred to the SIGGRAPH course of FEM simulation of 3D deformable solids (Sifakis & Barbic, 2012). Afterwards, the Dynamic Deformables course (Kim & Eberle, 2020) allows for a deeper dive into the topic of material models and the practicalities of implementing your own FEM simulation.

5.2. Experiments with synthetic data

The application of the method on the real-world scans of patients could be subject to various problems due to the simplification of the real-world conditions. This makes it difficult to identify any issues with the method. To this end, different controllable test scenarios are setup using a symmetric synthesized mesh to represent the breast. Each test scenario is used to determine the influence of a particular variable on the result of the method. It allows us to identify the strengths and limitations of the method. An overview of the experiments is given in Table 5.1. In the following sections, we will explain the setup of the experiments and discuss the results.

5.2.1. Experiment Setup

The rest configuration of the synthesized breast is represented as half an icosphere in Figure 5.1b. The half icosphere is symmetric and does not contain mesh singularities. It is scaled (72x149x149 mm) to be proportional to the breast in the prone configuration of patient 1 in the real-world experiments (see Subsection 5.4). The prone and supine configurations of the synthesized breast are generated by physically simulating the gravitational forces acting in the upwards or downwards direction respectively. The literature on deformable breast models (Griesenauer et al., 2017; Hipwell et al., 2016) commonly segments the breast into adipose, fibroglandular and tumour tissues. Each tissue type is assigned a Neo-Hookean material and the corresponding material parameters are optimized with the following value ranges: $E = 50 - 12000$ Pa, $\nu = 0.45 - 0.49$ and $\rho = 1000$ kg/m³. However, our method disregards the different tissue types and assigns a single homogeneous

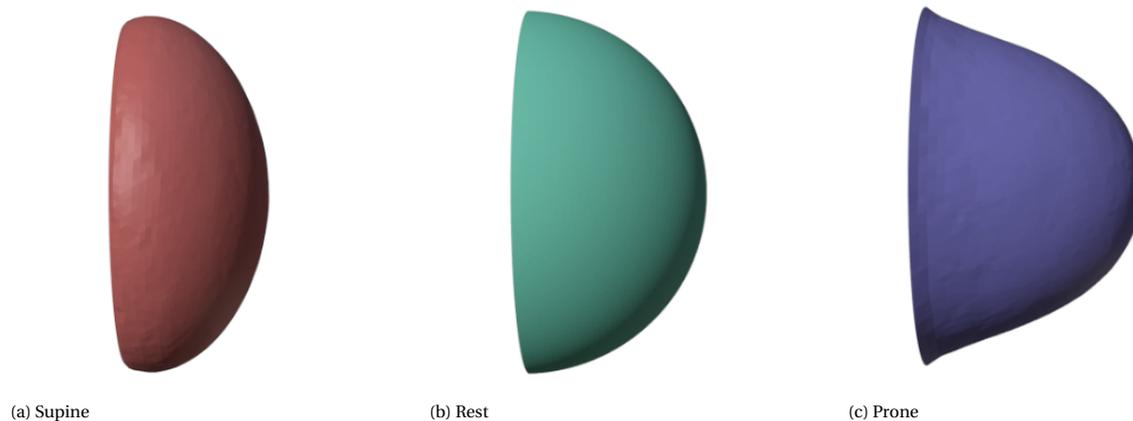


Figure 5.1: The supine, rest and prone configuration of the half icosphere mesh for the synthesized experiments

material to the entire breast. We choose to use the stable NeoHookean material with a fixed $\nu = 0.49$ and $\rho = 1000 \text{ kg/m}^3$. The choice for ν is to satisfy the incompressibility property of the breast as best as possible. The elasticity of the breast is coupled to the Young's modulus E , for which we have sampled the values [200, 400, 600, 800, 1000, 2000, 3000, 4000, 5000, 6000, 7000, 8000, 9000, 10000, 11000, 12000] Pa from the value range of E in the literature. The values for E have been sampled to leave out any degenerate prone-supine configurations. The extremes of the value range does show unrealistic prone-supine configurations of the synthesized breast in Figure 5.3.

The tumour is represented as a simple sphere placed inside the synthesized breast at different locations: top, middle, bottom, right, left (see Fig. 5.2a). This allows us to determine whether the location of the tumour affects the results of our method in the various experiments.

For the application of our method in the experiments, the following constants and variables are kept fixed unless otherwise specified: gravitational constant of 9.81 m/s^2 , gravitational direction vector downwards, and a landmark constraint weight of $1e4$. The input for our method consists of the material for the FEM simulations, the landmark correspondences between the prone and supine scans of the breast as well as the 3D surface scan of the breast in the supine position. The NeoHookean material model and its parameters ($E = 3400 \text{ Pa}$, $\nu = 0.49$ and $\rho = 1000 \text{ kg/m}^3$) are taken from the literature (Hipwell et al., 2016). The parameters are fixed for the experiments with exception of the material parameters experiment in Subsection 5.2.6. For the landmark correspondences, we have manually chosen five vertices spread across the synthesized breast shown in Figure 5.2b. The spread out placement of the landmarks should ensure that each region of the breast is affected by the landmark guided deformation of our method. The surface scan of the breast in the supine position takes the surface of the corresponding generated supine configuration of the synthesized breast. The experiments evaluate the results of our method with the Euclidean distance between the centroid of the resulting deformed tumour and the centroid of the tumour from the generated supine configuration.

5.2.2. Practical Settings

For the practical application of the method in a clinical setting, there are three relevant settings to consider in the context of the mixed reality project. In the first setting, a 3D surface scan of the breast in the supine position can be obtained from either the Microsoft HoloLens during the surgery or the PET-CT scan during the chemotherapy. The second setting requires the surgeon to choose appropriate landmark correspondences in order to pull the regions of the prone scan closer towards the breast in the supine position. It does not include the 3D surface scan of the breast in the supine position. The third setting has access to both the 3D surface scan of the breast in the supine position as well as the chosen landmark correspondences. The additional data accessible for each setting limits the use of our method. This means that the first setting can only use the surface guided deformation of our method (see Sec. 4.4). The second setting can only use the landmark guided deformation of our method (see Sec. 4.3), whereas the third setting allows for the application of the entire method. However, the first setting with the surface guided deformation also has the option to use a gravity FEM simulation without the landmark constraints. In this experiment, we want to compare and assess the accuracy of the tumour location for each of the different settings. Ideally, it would be beneficial to reduce the burden of the patient and the surgeon by avoiding the PET-CT scan acquisition and/or the

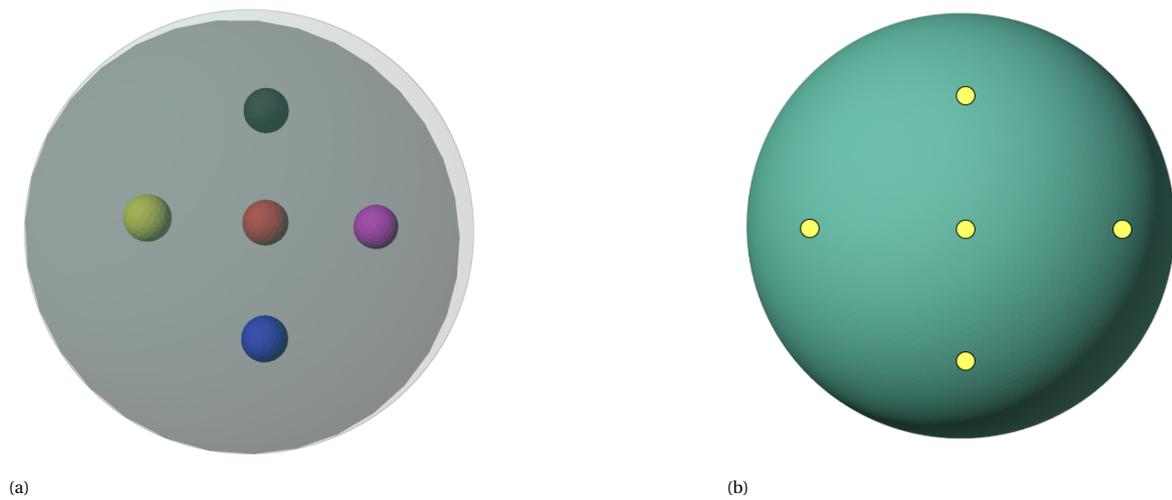


Figure 5.2: The settings for the synthesized data experiments. (a) The locations of the tumour inside the symmetric synthesized breast (top, left, bottom, right, middle). (b) The five landmark locations on the rest configuration for the experiments.

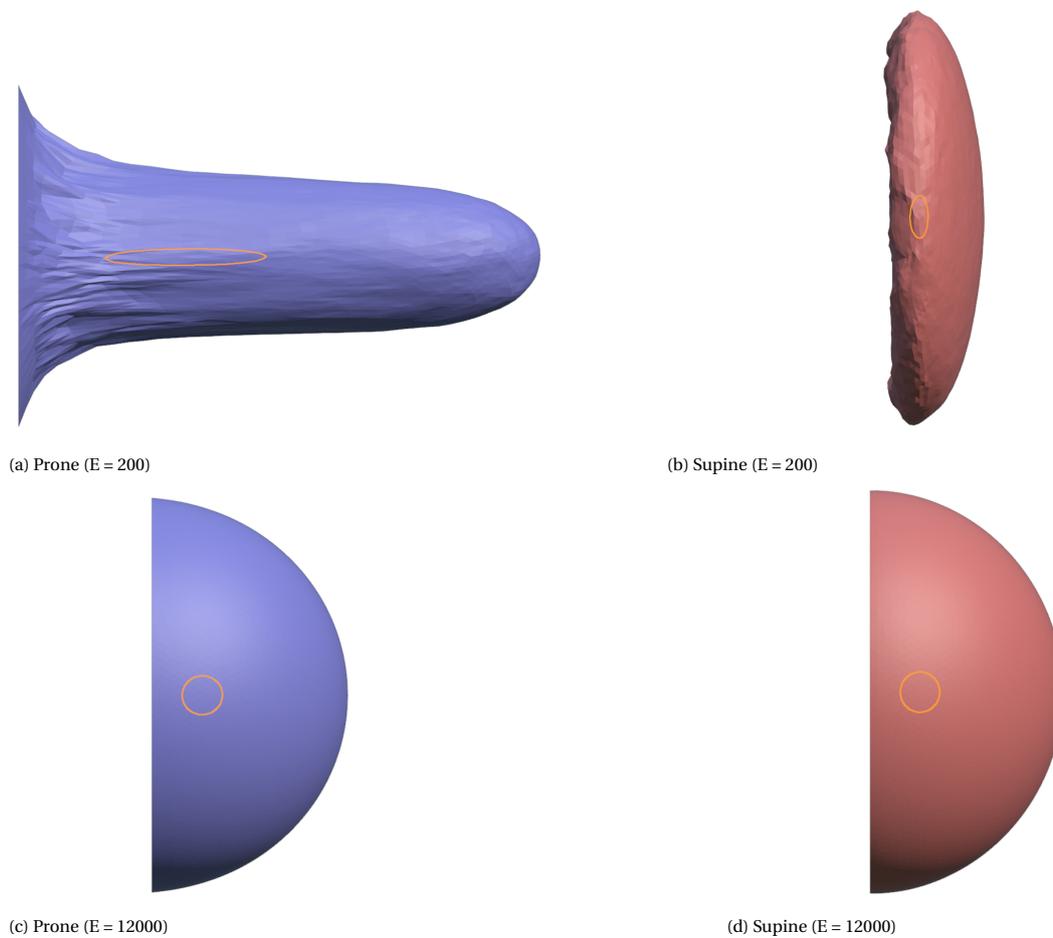


Figure 5.3: The extreme cases for the generated prone-supine configurations for the synthesized experiments

Name (Section)	Input	Variable	Variable Value Range
Practical Settings (5.2.2)	prone mesh, supine surface	choice of method	landmark guided deformation, surface guided deformation, full method
Rest Shape (5.2.3)	prone mesh, rest shape mesh, supine surface	choice of rest shape	prone/rest shape
Rotated Gravity (5.2.4)	prone mesh, supine surface	gravitational direction vector	rotation 0 – 15°
Misalignment (5.2.5)	prone mesh, supine surface	rigid alignment prone-supine meshes	rotation 0 – 15°, translation -10 – 10 mm
Material Parameters (5.2.6)	prone mesh, supine surface	Young’s modulus E	200 – 12000 Pa
Landmark Correspondence Weight (5.2.7)	prone mesh, supine surface	Landmark weight	1e2 – 1e5

Table 5.1: An overview of the synthetic data experiments. The input includes a stable Neo-Hookean material model with $E = 3400$ Pa, $\nu = 0.49$ and $\rho = 1000$ kg/m³, and the five landmarks from Figure 5.2b across the experiments. Each experiment is evaluated with the Euclidean distance between the tumour centroids of the computed result and the ground-truth supine tumour.

manual annotation of the landmark correspondences. For this experiment, our method takes as input the different prone-supine configurations, the different tumour locations, the five landmark correspondences and the fixed material settings (see Sec. 5.2.1).

Discussion In Figure 5.4, a boxplot visualizes the distribution of the tumour centroid distance across the different prone-supine configurations and tumour locations for the various applicable steps of our method. The rigid registration method shows the initial high distance distribution between the tumour in the prone and supine configurations (0.5 - 19.8 mm). The color of the boxplot bar indicates which methods are applicable for each of the different practical settings. The methods applicable to the first practical setting with the supine surface scan are the surface deformation (0.3 - 5.7 mm) and the gravity simulation + surface deformation (0.0 - 5.4 mm). From the corresponding plots, we can see that they have similar distance distributions that have been reduced in comparison to the rigid registration. The gravity simulation without landmarks does help the surface guided deformation of our method, since the gravitational forces push the surface of the prone configuration closer towards the supine configuration. The second setting with the landmarks applies the landmark deformation method, which shows a considerably higher distance distribution from 0.0 - 9.0 mm. The soft constraints used in the landmark deformation can only pull the prone regions closer to the supine regions, but the elastic forces can oppose the constraint forces in order to recover to the rest configuration and preserve the volume of the breast. On the other hand, the surface guided deformation restricts the location of the tumour inside the supine volume of the breast by projecting the surface of the prone mesh onto the supine mesh. The third setting contains both the landmarks and the supine surface scan, and is thus able to apply the full method onto the prone-supine configurations, which results in the smallest distance distribution from 0.0 - 1.0 mm.

5.2.3. Rest Shape

A deformable object undergoing deformation always tries to recover to its rest configuration. The FEM simulation needs the rest configuration as input. However in the real-world experiments, the actual rest configuration of the breast is unknown. The scans of the breast in the prone and supine positions are subject to deformation due to gravitational forces. The related works propose to approximate a rest configuration of the breast (see Section 2). In this experiment, we determine whether not knowing the actual rest configuration is detrimental to our method. This is evaluated by comparing the results of our method when using the prone configuration as the rest configuration against using the actual rest configuration shown in Figure 5.1b. Thus the input to our method becomes the different prone-supine and rest-supine configurations, the different tumour locations, the five landmark correspondences and the fixed material settings (see Sec. 5.2.1).

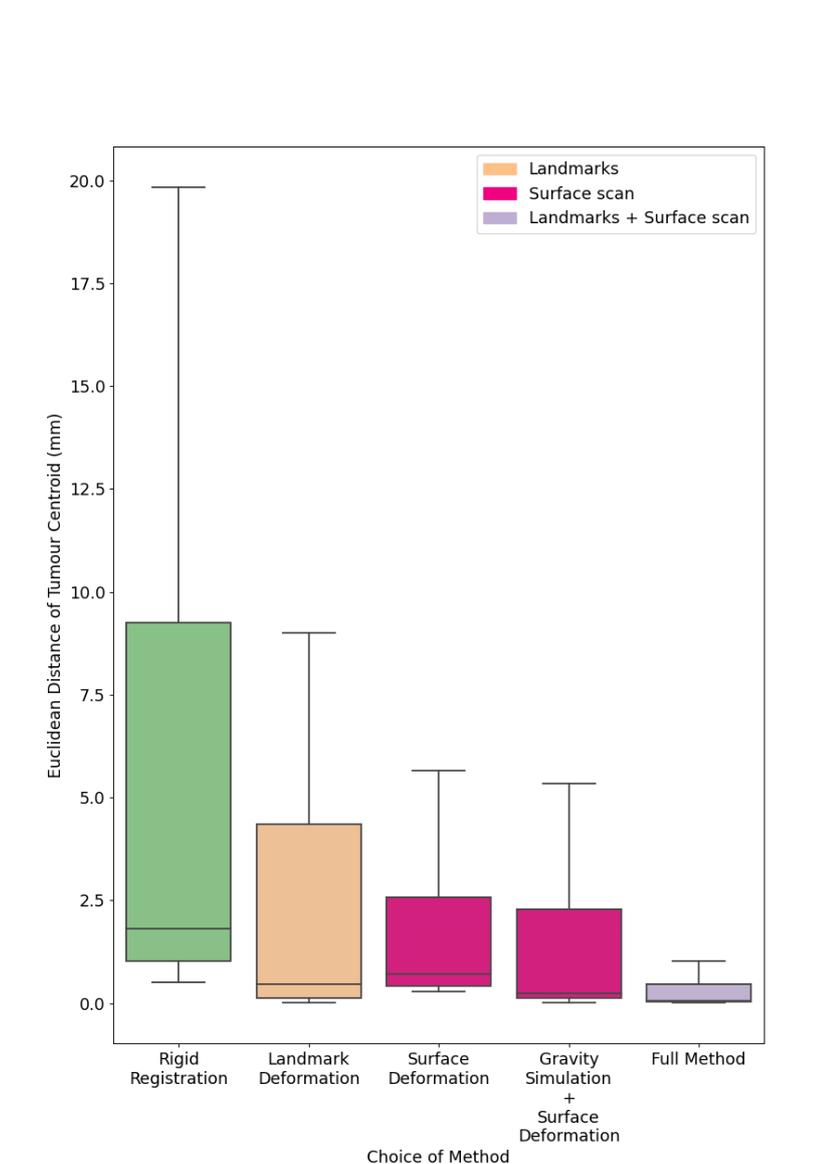


Figure 5.4: The boxplots of the tumour centroid distances for the three practical settings. Each plot indicates the distribution of the tumour centroid distances across the prone-supine configurations and tumour locations for the applied method. The plot splits the distribution into four quartiles with the median at the center of the plot, and leaves out any outliers. The color of the plot indicates which setting corresponds to the method with the exception of the rigid registration method.

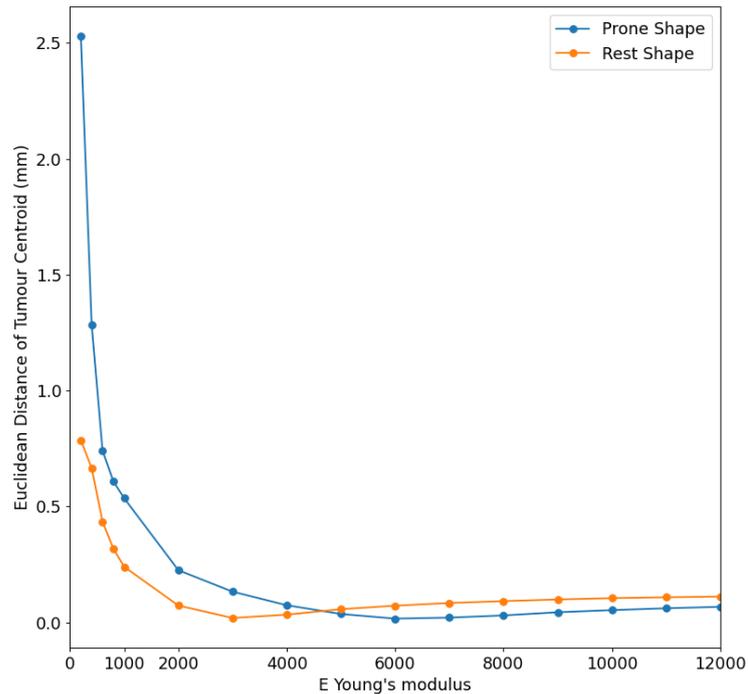


Figure 5.5: The results of the rest shape experiment. The plots average the results across the different tumour locations.

Discussion Figure 5.5 plots the results of the experiments with different values for the Young's modulus E for the generated prone-supine / rest-supine configurations on the x-axis and the tumour centroid distance from the ground-truth on the y-axis. The Young's modulus E is coupled to the elasticity of the breast, where higher values of E indicate a stiffer material and lower values of E indicate a more elastic material. The plots of the prone and rest shape indicate small differences between the averaged results for most values of E with tumour centroid distances below 1.0 mm. The prone shape keeps performing worse for increasingly lower values of E in comparison to using the actual rest shape. The surface of the actual rest configuration is always closer to the surface of the supine configuration, which makes it easier to obtain a good initialization for the surface guided deformation of our method. The extreme case of $E = 200$ Pa does result in a high tumour centroid distance of 2.5 mm, but the deformation of the corresponding prone configuration is unrealistically soft as can be seen in Figure 5.3a. Therefore it is not very problematic for our method to use the prone configuration as the rest configuration for the FEM simulations.

5.2.4. Rotated Gravity

The precise direction of gravity affecting the prone and supine scans of the breast is unknown. It could be that gravity is affecting the breast at a slight angle (see Fig. 5.6b), whereas our method assumes that the direction of gravity is in a straight downwards direction (see Fig. 5.6a). This experiment is used to determine whether rotating the direction of gravity has a significant influence on the result of our method. The rotations on the direction of gravity are performed on the X axis. Since applying a rotation on the Z axis to the direction of gravity does not change its direction, and due to the symmetry of the synthesized breast it does not matter whether it is rotated on the X or Y axis. The angle of rotations should be small. (Griesenauer et al., 2017) estimates the stiffness of the breast by rotating the torso by 15° . It assumes that the gravitational forces have sufficiently deformed the breast from the supine to the rotated torso position in order to compute the stiffness. Therefore, we choose to keep our rotations within that range of 15° to keep the amount of deformation limited. The range for the angles of rotation are sampled as: $[0, 2, 4, 6, 8, 10, 15]^\circ$. The input to our method is then the rotated gravity direction vector, the prone-supine configurations, the different tumour locations, the five landmark correspondences and the fixed material settings (see Sec. 5.2.1).

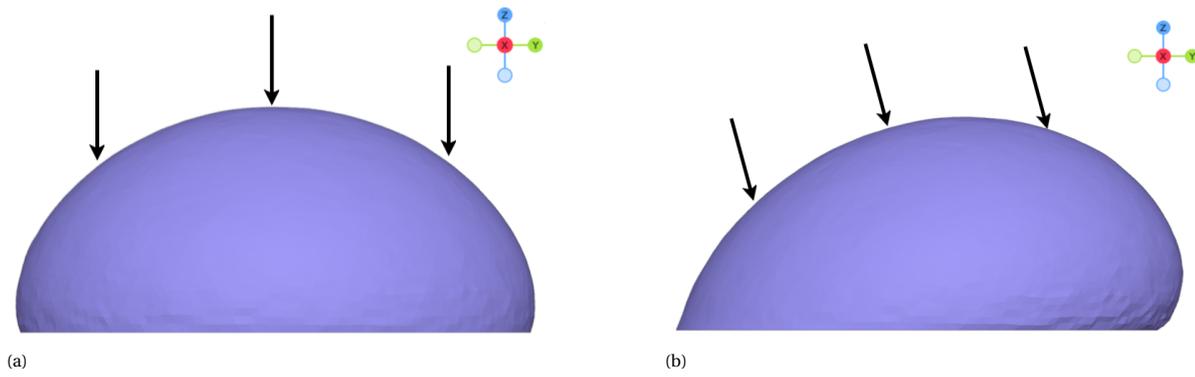


Figure 5.6: The rotated gravity experiments change the gravitational direction to determine its influence on our method. (a) shows the assumed situation with gravity acting straight downwards. (b) shows the situation when the gravitational direction has been rotated by 15 degrees.

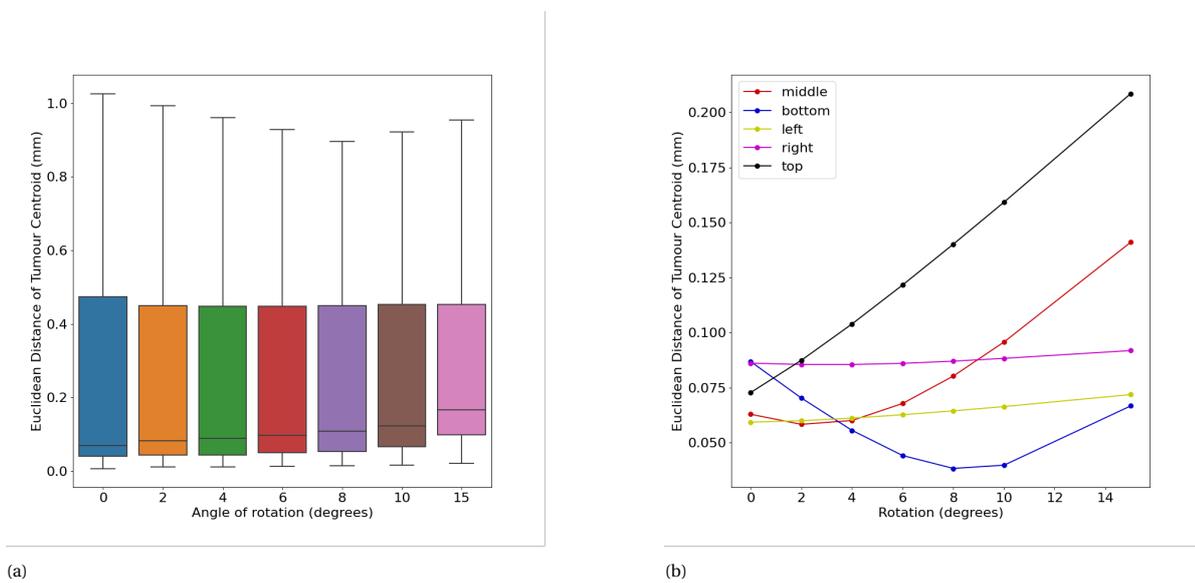


Figure 5.7: The results of the rotated gravity experiment. (a) shows a boxplot for the tumour centroid distance across the 16 prone-supine configurations and five tumour locations. (b) shows distance plots for each tumour location for the prone-supine configuration with $E = 4000$.

Discussion The boxplot in Figure 5.7a shows the distributions of the resulting distance between the tumour centroids for the different angles of rotations applied on the downwards gravity vector across the five tumour locations and different prone-supine configurations. The median of the distance distributions slightly increases with larger angles, but the overall spread of the distance distributions stays around 1.0 mm for the different rotated gravity directions. This indicates that the result of our method is not significantly influenced by the direction of gravity. The guidance of the landmark constraints and the surface projection of our method reduce the impact of a rotated gravitational direction. In Figure 5.7b, the distance metric is plotted for each tumour location of the prone-supine configuration with $E = 4000$ Pa. The tumour at the bottom location improves its results with a decreasing tumour centroid distance for smaller angles of rotation. The direction of gravity is rotated in such a way that the tumour at the bottom location is pushed downwards towards its ground-truth position, whereas the other tumours can be pushed away from their corresponding ground-truth positions. However with larger angles of rotation, the distance between the deformed tumour at the bottom location and the ground-truth tumour increases due to the significant deformation of the breast.

5.2.5. Misalignment

Our method uses landmark correspondences between the prone and supine meshes to perform the landmark guided deformation (see Subsection 4.3). This requires an accurate initial rigid alignment between the prone

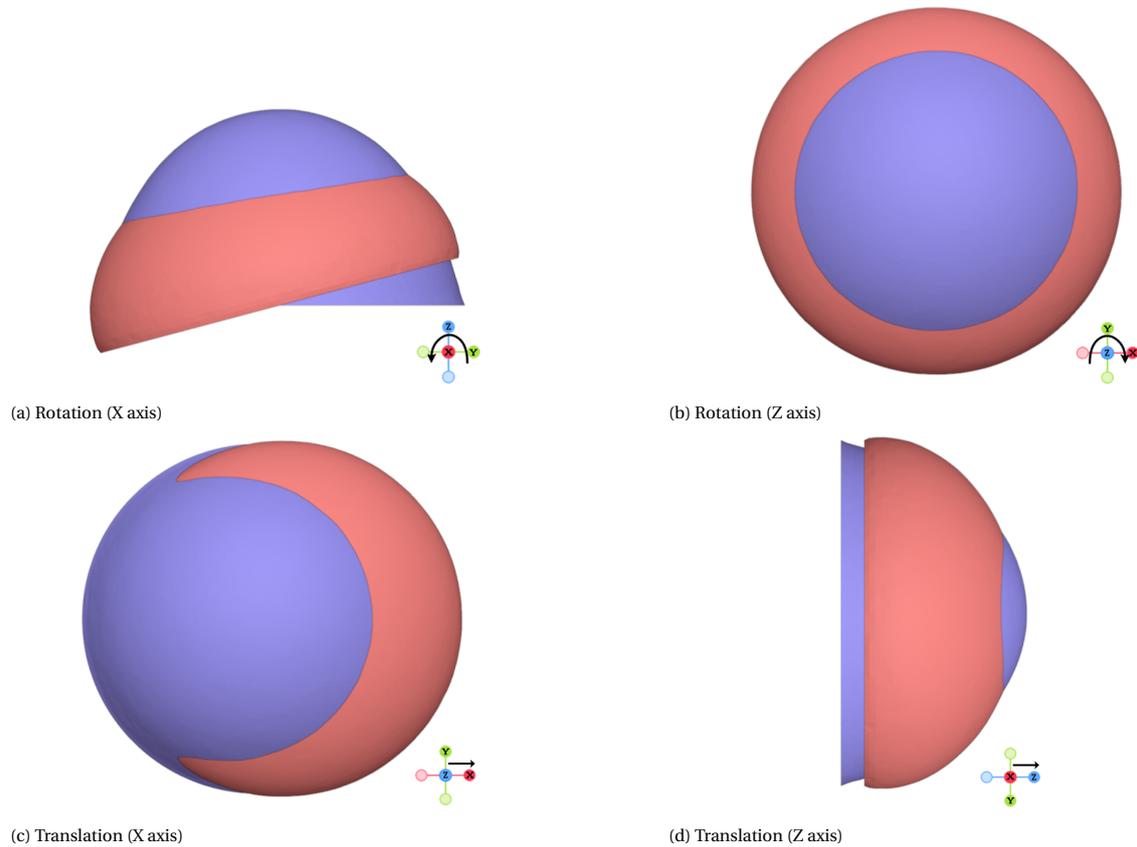
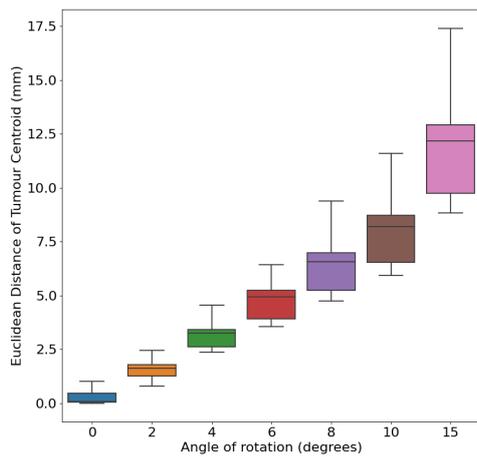


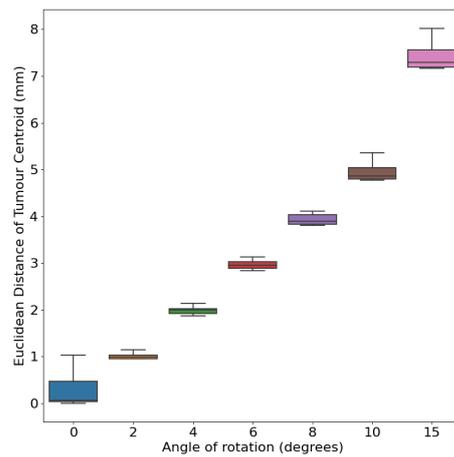
Figure 5.8: The misalignment experiment rotates or translates the supine mesh in order to determine how our method would be affected by misalignment of the prone and supine meshes.

and supine meshes, for which we have chosen to use the chest walls of the both configurations to perform the alignment. However in the real-world patient data (see Sec. 5.4), the chest walls in the prone and supine scans show significant deformation from each other. This can therefore still result in misalignment between the prone and supine meshes. This experiment is used to determine how significant any misalignment between the prone and supine meshes is to our method. The misalignment in translation and rotation are considered separately for the experiment. A visualization of the rotations and translations performed on the supine meshes is shown in Figure 5.8. Any rotation or translation on the X or Y axis should obtain similar results due to the symmetry of the synthesized breast. Therefore rotations or translations are performed on the X and Z axis. The same range of angles for the rotation will be used as the rotated gravity experiment (see Subsec. 5.2.4): $[0, 2, 4, 6, 8, 10, 15]$ $^\circ$. The amount of translation should be kept small and performed in the positive and negative directions, which is particularly important for translations on the Z axis. Translations on the Z axis in the negative direction can result in the ground-truth tumour being positioned behind the fixed chest wall of the prone mesh (see Fig 5.8), whereas in the positive direction the supine tumour should be inside the volume of the prone mesh. The sampled range of translations is chosen to be: $[-10, -8, -6, -4, -2, 0, 2, 4, 6, 8, 10]$ mm.

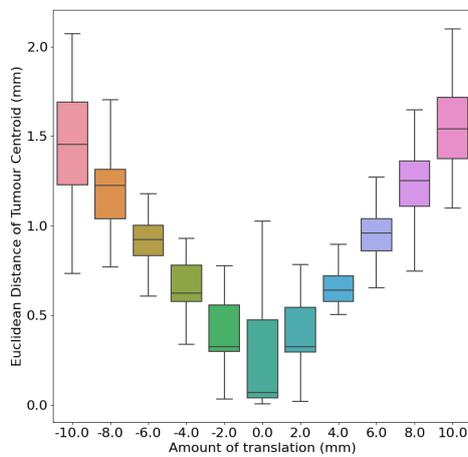
Discussion The results for the misalignment between prone and supine configurations have been separated into rotation (Fig. 5.9a and 5.9b) and translation (Fig. 5.9c and 5.9d). The boxplots in these figures visualizes the distance distributions for the chosen rotations or translations of the supine mesh over the different tumour locations and prone-supine configurations. In comparison to the previous experiments, these plots showcase significantly higher distances for the results of our method. The boxplots regarding the rotational misalignment show that rotations on the X/Y axis incur worse results than rotations on the Z axis. The spread of the distance distribution for rotations on the X/Y axis increases with higher angles of rotation, whereas the distance distribution for rotations on the Z axis stays small in comparison. The rotations on the Z axis are turning the synthesized breast with the locations of the tumours while keeping the alignment of the chest wall (see Fig. 5.8b). The surface guided deformation then ensures that the tumour is located within



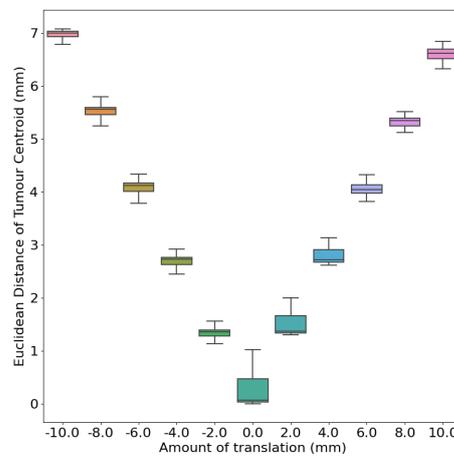
(a) Rotation (X/Y axis)



(b) Rotation (Z axis)



(c) Translation (X/Y axis)



(d) Translation (Z axis)

Figure 5.9: The results for the misalignment experiment for rotation and translation around the different axes.

the volume of the breast in the supine configuration. On the other hand, the rotations on the X/Y axis breaks the alignment of the chest walls (see Fig. 5.8a). It also results in the rotation of the ground-truth tumour locations, which can become impossible to reach for our method due to the fixed chest wall. Another problem with the rotations around the X/Y axis is that the surface guided deformation can incorrectly project vertices onto the same closest point.

For the boxplots for the translational misalignment, the translations on the X/Y axis do not have a significant effect on the result of our method. However, the translations on the Z axis does significantly increase the distance distribution. The landmark constraints partially counteract the misalignment in translation on the X/Y axis. The translations on the Z axis move both the tumour and the chest wall either forwards or backwards (see Fig. 5.9d), which can make it impossible for the result of our method to reach the ground-truth due to the fixed chest wall.

5.2.6. Material Parameters

The choice of material parameters defines the elastic behaviour of the deformable breast, which would influence the result of our method due to the use of FEM simulations. In the real-world experiments, the actual material parameters for a patient are unknown. The related works make use of different material optimization techniques to identify optimal material parameters for a patient (see Section 2). The application of our method during the experiments has been using fixed material parameters instead, and relying on the landmark constraints and surface projection to reduce the influence of the unknown material parameters. Thus with this experiment, we want to determine the influence of the material parameters on the result of our method. We are particularly interested in the choice of the Young's modulus E , which is coupled to the elasticity of the material. The Poisson's ratio $\nu = 0.49$ and density $\rho = 1000 \text{ kg/m}^3$ are kept fixed. The sampled range of values for E is chosen to be the same range used to generate the prone-supine configurations in Subsection 5.2.1: [200, 400, 600, 800, 1000, 2000, 3000, 4000, 5000, 6000, 7000, 8000, 9000, 10000, 11000, 12000] Pa. Our method takes as input the generated prone-supine configurations, the different tumour locations and the five landmark correspondences (see Sec. 5.2.1).

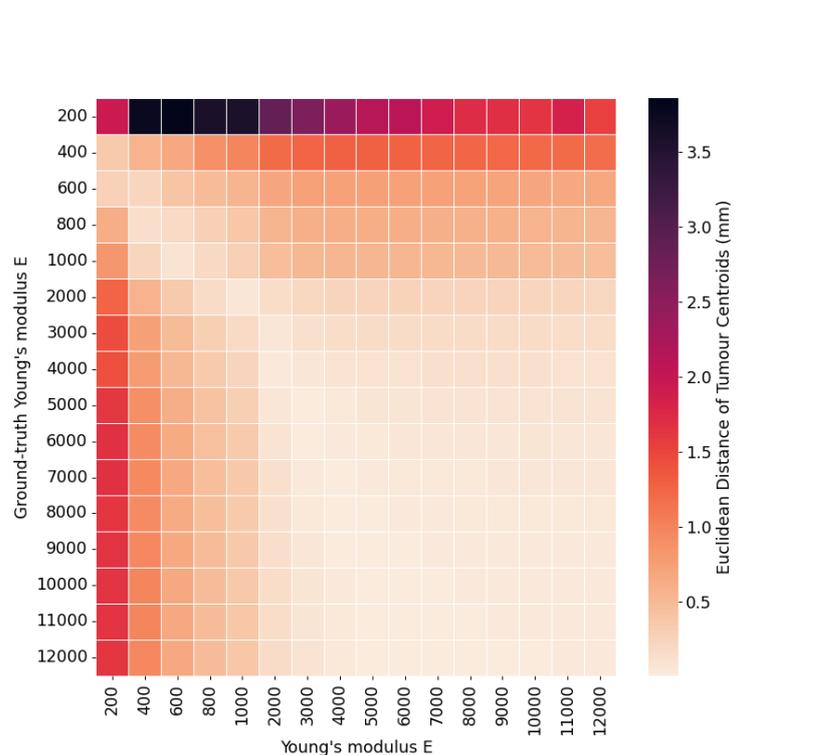


Figure 5.10: The heatmap visualizes the results of the material parameters experiment, where the Y axis shows the Young's modulus used to generate the prone-supine configuration and the X axis is the Young's modulus used by our method for the FEM simulations.

Discussion The heatmap in Figure 5.10 visualizes on the Y axis the ground-truth E that was used to generate the prone and supine configurations for the experiment, and on the X axis the choice for E for applying

our method onto the corresponding generated prone and supine configurations. From the heatmap, it can be seen that applying the method on a more elastic breast with $E < 1000$ performs better with lower values for E . On the other hand, the performance of the method does not change significantly for different values for E when applied to breasts with stiffer tissues. An exception to this is $E = 200$ Pa, which makes the material of the breast unrealistically elastic as can be seen in Figure 5.3. Overall, the method is less dependent on the choice of material parameter E due to the usage of soft constraints.

5.2.7. Landmark Correspondence Weight

The landmark correspondence constraints are defined as soft penalty energy functions with a specified weight (see Eq. 4.2). The weight indicates how strongly the constraint pulls the landmark regions closer together. On the other hand, the elastic forces will work against the landmark constraint to keep the deformable object in its rest shape. This means that a high weight will not guarantee that the landmark correspondence points are perfectly overlapping. It can also lead to a stiff system of equations in the FEM simulations. The landmark constraints can be re-defined as hard constraints in order to ensure that the landmark correspondence points always overlap. With this experiment, we want to determine an appropriate weight for the landmark constraint. The weights for the landmark constraints are chosen to be: $[1e2, 1e3, 1e4, 1e5]$. The input for our method is the generated prone-supine configurations, the different tumour locations, the five landmark correspondences and the fixed material settings (see Sec. 5.2.1).

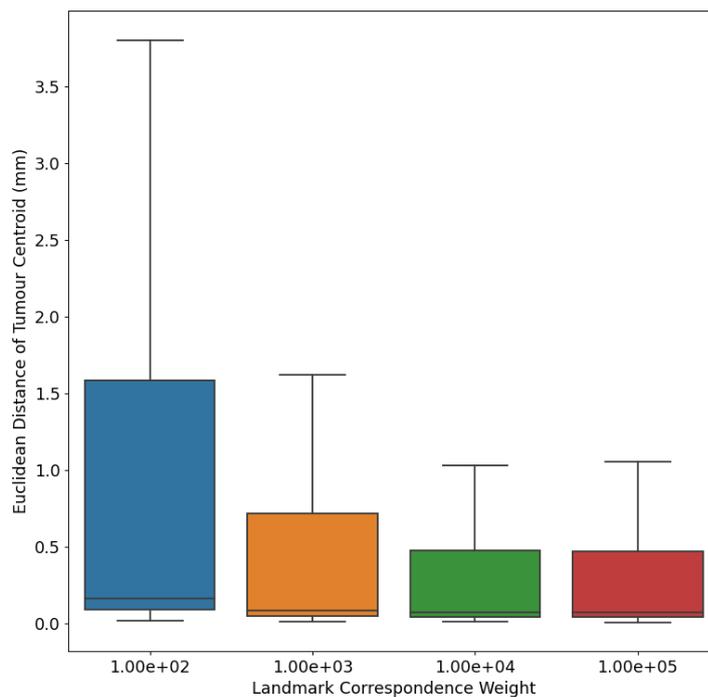


Figure 5.11: The results of the landmark correspondence weight experiment.

Discussion The boxplot in Figure 5.11 visualizes the distance distribution across each prone-supine configuration and each tumour location categorized for each landmark correspondence weight. It can be seen that a higher landmark weight results in a more condensed and lower distance distribution, except for the similar performing landmark weights $1e4$ and $1e5$. Therefore we choose to use the best performing landmark weight $1e4$ for our method. The higher landmark weight $1e5$ does not improve the performance of the method, and would only introduce stronger elastic forces to recover to the rest configuration by pulling on the landmark regions.

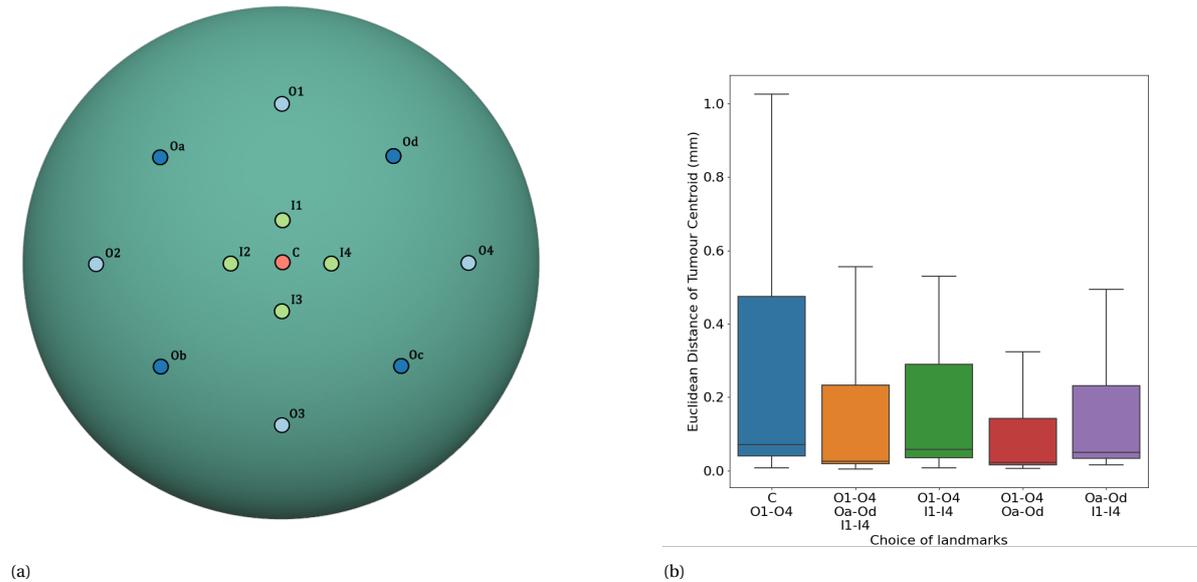


Figure 5.12: The landmarks experiment looks into using different sets of landmarks. (a) The different sets of landmarks O1-O4, Oa-Od, and I1-I4 placed on the rest shape. (b) The results for the landmarks experiment.

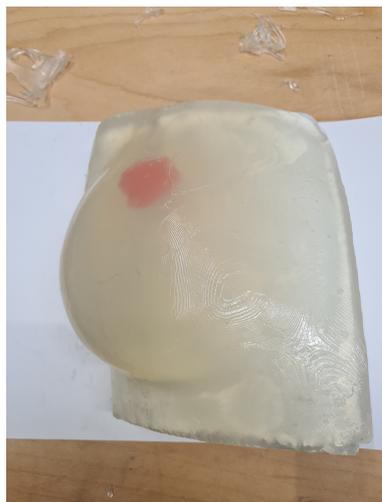
5.2.8. Landmarks

The usage of landmarks to guide the deformation of the breast makes up for not approximating the unknown variables, such as the rest shape and the material parameters. Each landmark pulls the surrounding region of the prone mesh closer towards the surface of the supine mesh, where the linear skinning subspace uses a heuristic to determine the size of the affected region. It is important to properly distribute the landmarks across the surface of the breast in order to maximize the overlap of the prone and supine surfaces. Otherwise the unaffected regions of the prone surface could be too far apart from the supine surface for the surface projection step of our method. With this experiment we want to take a preliminary look into appropriate choices for the positioning of the landmarks. The related works have made use of surface markers on top of the breast for the alignment of multimodal scans of the breast (Bessa et al., 2020). In Figure 5.12a we have used these surface markers to position the different sets of landmarks on the rest shape of the synthesized breast. The five landmark positions from Figure 5.2b have been annotated for this experiment as C, O1, O2, O3, O4. To evaluate the choice of landmarks, we have chosen to compare the following groups of landmarks: $\{C, O1-O4\}$, $\{O1-O4, Oa-Od, I1-I4\}$, $\{O1-O4, I1-I4\}$, $\{O1-O4, Oa-Od\}$, $\{Oa-Od, I1-I4\}$. The input for our method is the generated prone-supine configurations, the different tumour locations, the different setups for the landmark correspondences and the fixed material settings (see Sec. 5.2.1).

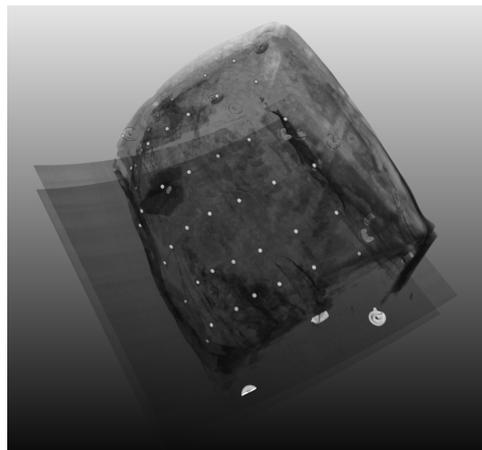
Discussion In Figure 5.12b, the boxplot shows the distance distributions across each prone-supine configuration and each tumour location categorized for the different groups of landmarks. The landmarks used throughout the synthesized data experiments $\{C, O1-O4\}$ sets a baseline for the other groups of landmarks with a maximum distance of 1.0 mm. Each of the other landmark groups performs better than the baseline, where the landmark group $\{O1-O4, Oa-Od\}$ achieves the lowest maximum distance of 0.3 mm. The resulting tumour centroid distances in this experiment are less than 1 mm. Therefore the choice of landmarks does not seem as impactful on the result of our method given that the landmarks are distributed across the surface of the breast.

5.3. Phantom Breast Experiments

The experiments with synthesized data allowed for an exploration of the different strengths and limitations of the method using idealized scenarios. To step closer towards a practical application of the method, a phantom breast has been CT scanned in the prone and supine positions with a resolution of $957 \times 512 \times 512$ voxels and an isotropic voxel size of 0.6 mm. The material used to fabricate the phantom breast is soft polyvinyl chloride plastisol (Lure Flex soft, Lure Factors, UK), which has been tuned to deform similar to a real breast. However, the phantom breast used for the scans is quite stiff in comparison to the real-world data exper-



(a) Picture of the Phantom Breast



(b) Phantom Breast Scan with markers



(c) Prone mesh



(d) Supine pushed mesh

Figure 5.13: The setup for the phantom breast experiments.

iments. Therefore the supine configuration has been additionally deformed by placing a weight on top of the breast during the CT scan. As can be seen from Figure 5.13c and 5.13d, the surface of the breast barely changes and the tumour has been slightly translated between the prone and supine meshes. The prone and supine configurations of the phantom breasts are rigidly aligned using a transformation matrix obtained from point-based registration of ten fiducial markers placed on the rigid torso of the phantom breast during the CT scans. For the landmark correspondences, the CT scans have also 42 markers placed on the surface of the phantom breast. These markers are visualized in Figure 5.13b. The FEM simulations use the Stable Neo-hookean material model with the following material parameters: $E = 3400 \text{ Pa}$, $\nu = 0.49$ and $\rho = 1000.0 \text{ kg/m}^3$.

The results show that the original stiff prone and supine configurations start with a low tumour centroid distance of 2.71mm, and the application of the method barely improves upon it. The pushed version of the prone and supine configurations show an improved relative performance between the baseline and the application of the method.

5.4. Real-World Experiments

For the real-world experiments, anonymized data has been retrospectively obtained from four patients who underwent breast-conserving surgery with radiotherapy treatment. The patients have been CT scanned in the prone and supine positions at the Erasmus Medical Center, Rotterdam, the Netherlands. Each CT scan has been preprocessed with MeVisLab in order to segment the tumors, breast and skin, and thorax using region growing segmentation with manual annotations. These are then extracted into 3D triangle meshes. To keep the chest wall of the breast rigid and fixed, the corresponding vertices have been manually selected

	Tumour Centroid Distance	Tumour Centroid Distance (Rigid Registration)
Prone - Supine	2.59	2.71
Prone - Supine Pushed	2.80	10.81

Table 5.2: The results of the phantom breast experiments with measurements in mm.

	Tumour Centroid Distance	Tumour Diameter (Supine)	Tumour Centroid Distance (Rigid Registration)
Patient 1	17.3	41.0	68.8
Patient 2	12.6	26.5	19.3
Patient 3	23.4	39.0	56.9
Patient 4	22.7	36.0	68.1

Table 5.3: The results of the real-world experiments with measurements in mm.

from the triangle mesh and kept fixed during the FEM simulations with hard constraints. The CT scans do not include landmark correspondences for the prone-supine configurations. Thus we choose to manually select five approximately corresponding vertices between the prone and supine configurations to act as the landmark correspondences. The material for the breast is chosen to be the stable Neo-Hookean material with the following material parameters: $E = 3400$ Pa, $\nu = 0.49$, $\rho = 1000$ kg/m³.

The prone and supine configurations with an outline for the tumours are shown in Figure 5.14. For patient 1 and 4, there are additional tumours located near the axillary gland. These are disregarded for the method and not shown in the figure. However, patient 3 has a tumour in the supine position partially outside of the mesh of the breast due to problems with the segmentation of the breast and chest wall (see Fig. 5.15a). This makes it impossible for the resulting tumour of our method to completely overlap with the ground-truth tumour for patient 3 due to the fixed chest wall. Figure 5.15b illustrates another problem with the real-world data, which is that the chest wall significantly differs between the prone and supine configurations. This is the case for all patients. As a result, it is difficult to find an accurate rigid alignment.

The results in Table 5.3 show that the method is able to reduce the tumour centroid distance significantly from the initial distance. However, these distances are still significantly large, especially in comparison to the results of the idealized scenario in the synthetic data experiments in Section 5.2. The most problematic variable in the synthesized data experiments is the misalignment between the prone and supine configurations. The poor rigid alignment does show up for the data of each patient due to the significant deformation between prone and supine configurations (see Fig. 5.15b). This would indicate that the biggest bottleneck would be the rigid alignment. On top of that, the tumours are located close to the chest wall. The combination of these problems can make it impossible for the tumour in the prone position to deform to the correct location in the supine position. Another problem that comes up is that the tumour gets flattened during the FEM simulations, whereas the ground-truth tumour in the supine position is quite stiff and keeps similar shape as the tumour in the prone position. This is due to our choice for applying a homogeneous material in combination with the amount of deformation that is needed to push the prone breast towards the shape of the supine breast.

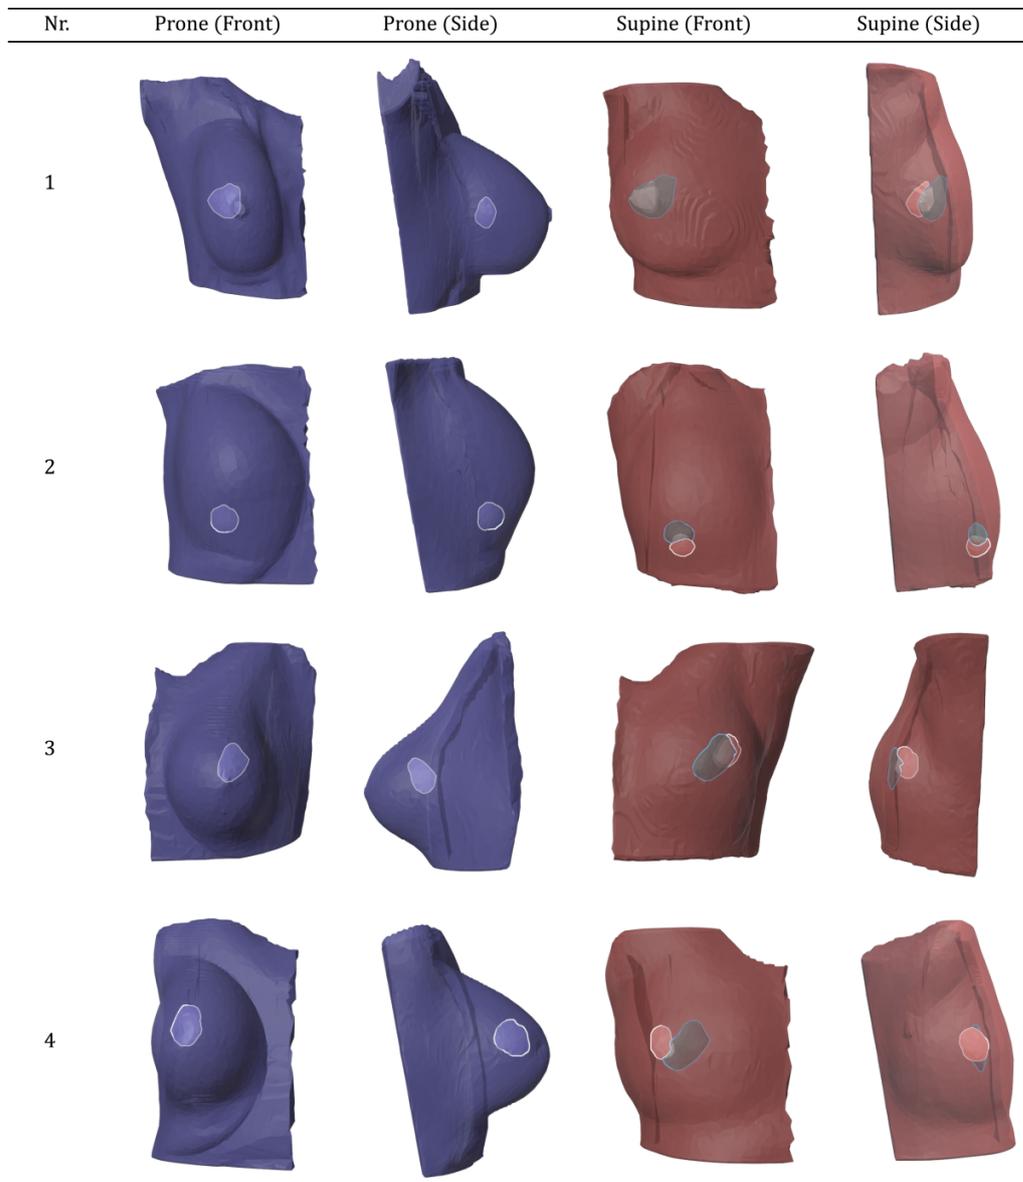
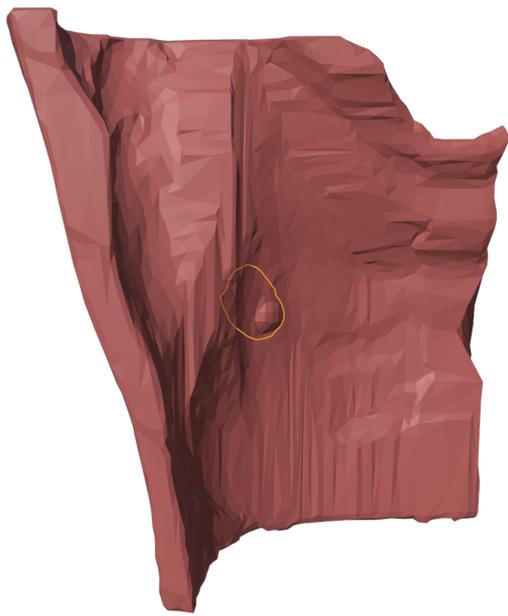
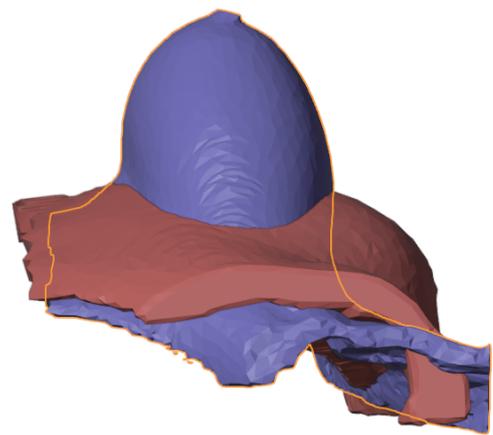


Figure 5.14: The prone and supine configurations of the patients in the real-world experiments. The resulting tumours of our method are indicated with a blue outline, and the ground-truth tumours are shown with a white outline.



(a) Patient 3 Supine (Back)



(b) Patient 1 Prone-Supine alignment

Figure 5.15: Examples of problems with the real-world data.

6

Discussion

Our method deforms the volumetric mesh of the breast in the prone position towards the supine position with the use of landmark constraints in a subspace simulation, surface projection onto the surface mesh in the supine position and optimization of the interior tissues of the breast. The result of the method is a nonrigid prone-to-supine mapping for the tumour, which should give a good approximation for the location and deformation of the tumour in the supine position. The synthesized data experiments have tested the method for different variables under idealized conditions using a half icosphere to model the rest shape of the breast. The prone and supine configurations are synthesized from a specified range of Young's moduli with a gravitational FEM simulation. This specified range of Young's moduli does include extremely unrealistic cases as can be seen in Figure 5.3. The generation of the prone-supine configurations does not make use of varying the Poisson's ratio, whereas the related works (Sec. 2.3) use a range of Poisson's ratio to optimize the material parameters. Besides that, the tumours generated from these prone-supine configurations are deformed with the same material as the rest of the breast. For the more elastic configurations, this can lead to very unrealistic stretchy tumours for the prone-supine configurations. The tumours from the real-world experiments are quite stiff and show small amounts of deformation from the prone towards the supine configuration. The three different practical settings relevant for the clinical application of the method have been evaluated on the synthesized data by applying the parts of the method, which are applicable according to the data available for the corresponding setting. The setting with the landmarks and the setting with the supine surface both result in significantly large tumour centroid distances under these idealized conditions. These settings are therefore unsuitable for clinical application. On the other hand, the setting with both landmarks and the supine surface obtains a significantly lower distance by applying the full method.

With regards to the other tested variables, the method showed a difference up to 2.5 mm between using the prone shape as the rest shape in comparison to using the actual rest shape. This tumour centroid distance increases with prone-supine configurations synthesized from increasingly lower Young's moduli. Similarly, the Young's modulus used in the FEM simulations of the method is more sensitive with elastic prone-supine configurations with a low Young's modulus, which results in distances up to 3.5 mm. This is reduced to 1.7 mm for stiffer prone-supine configurations. The rotation of the gravitational direction evaluates to similar tumour centroid distance distributions across the tested range of angles. These variables do not significantly affect the result of the method even though they heavily influence the deformable object in FEM simulations. The method reduces the impact of these variables with the number of constraints applied during the FEM simulations, which are used to guide the deformation of the deformable object. The landmarks play an important role in guiding the deformation with their constraint forces. The method showed tumour centroid distances below 1.0 mm for the tested groups of landmarks, which are properly distributed across the surface of the breast. Therefore the usage of these groups of landmarks does not significantly affect the results of our method. A more thorough investigation into the positioning and distribution of the landmarks could be done in future works. The most problematic variable in the experiments is the misalignment between the prone mesh and the supine surface mesh. Any misalignment with regards to the rotation of the meshes results in tumour centroid distances up to 17.5 mm on the symmetric X/Y axis and 8.0 mm on the Z axis. As for misalignment in translation between the meshes is mostly relevant for the Z axis where it can go up to 7.1 mm, whereas the X/Y axis has an upper limit of 2.1 mm. The reason for the significantly large distances comes as a result of our reliance on landmark and surface constraints, which are very sensitive to the rigid alignment

of both meshes. Since transforming one of the meshes would also transform the corresponding landmark or surface projection constraints as well.

The phantom breast experiments uses CT scans of a fabricated breast in order to model a more realistic situation than the synthesized data experiments. It ensures a good rigid alignment between the prone and supine CT scans with the use of fiducial markers. The landmark correspondences are also obtained from markers placed during the CT scans. These experiments have two settings due to the stiffness of the phantom breast: the prone-supine CT scans and the prone-supine pushed CT scans. The supine pushed CT scan is obtained from additionally deforming the phantom breast in the supine position with an additional force in the gravitational direction. The application of our method barely improves the location of the tumour from an initial distance of 2.71 mm to 2.59 mm. The improvement is more significant for prone-supine pushed CT scans from 10.81 mm to 2.80 mm. We expected the results of the method to be closer to the ground-truth location due to the stiffness of the phantom breast. Since the choice for Young's modulus and rest shape has less impact on the results of the synthesized data experiments.

The real-world experiments contain PET-CT scans of four patients in the prone and supine positions without any marker placements. The rigid alignment between the prone and supine CT scans has been performed manually. This is a simplification of the relevant clinical setting with the surface scan in the supine position instead of the volumetric supine CT scan. The landmark correspondences have been chosen at visually similar locations by a layman. The results of these experiments are significantly worse than any of the previous experiments with tumour centroid distances from 12.6 mm to 24.5 mm. It emphasizes the problem pointed out by the synthesized data experiments of the misalignment between the prone and supine meshes. The prone and supine meshes are poorly aligned even with the simplification of the supine CT scan. The chest wall shows significant deformation between the prone and corresponding supine configurations. This could partially be attributed to the segmentation and meshing process of the CT scan. However, a nonrigid registration of the torso would need to be performed in order to solve the problem of the misalignment for the prone and supine CT scans. Additionally, the tumour is located near the chest wall. In the case of patient 3, the tumour in the supine position is even partially outside of the chest wall. Another problem is that the tumour deformed by the method has also been flattened significantly in comparison to the ground-truth tumour in the supine position. The cause is the choice to apply a homogeneous material model to the entire breast. The tumour is flattened due to the large compressional forces, which are pushing onto the surface of the breast in order to match the prone and supine surfaces. A solution would be to assign a stiffer material model to the tumour. In order to assign a different material model to the tumour, the implicit mapping would need to be able to deal with the tetrahedrons with partial tumour or breast tissue. The tumour could also be explicitly meshed in the tetrahedral mesh, which would reduce the quality of the tetrahedralization. However, this does for simple assignment of a different material model to the breast and tumour tissues. The related works offer a different solution to this problem by optimizing the material parameters across the breast.

The clinical application of the method is to visually aid the surgeon during lumpectomy with mixed reality systems. The method deforms the scans of the breast in the prone position and deforms the breast and tumour into the supine position. Before the proposed method could be applied in a clinical setting, the misalignment and flattening issues in the real-world experiments need to be addressed. The resulting static meshes would be overlaid onto the patient in the supine position using the Microsoft HoloLens. The static supine mesh in the mixed reality system does not take into account any additional deformation of the breast and torso due to the breathing of the patient. A FEM simulation with a linear material model could be applied to compensate for the breathing, since the amount of deformation should be quite small.

7

Conclusion

Mixed reality systems could visually aid the surgeon by overlaying the tumour from the supine position on top of the patient. Our method uses as input a volumetric scan of the breast in the prone configuration, a surface scan of the breast in the supine configuration, and landmark correspondences in order to deform the tumour in the prone position towards the supine position. In the preprocessing stage, the prone volumetric scan is tetrahedralized and the tumour is implicitly mapped. The prone and supine meshes are rigidly aligned manually or with ICP. A subspace FEM simulation with landmark constraints pulls regions in the prone configuration closer towards similar regions in the supine configuration. The resulting surface of the prone configuration is projected onto the surface of the supine surface scan. Then a FEM simulation with surface constraints is used to optimize the interior of the breast.

The synthetic data experiments showed that our method needs the landmarks and the supine surface scan to obtain good results. It also showed that the method is not significantly affected by the choice of rest shape, direction of gravity or the choice of material for the breast. The biggest problem is the sensitivity to the misalignment between the prone and supine configurations. The real-world experiments resulted in large tumour centroid distances, which makes it unsuitable for the clinical application of the mixed reality systems. The major problems with the method are the difficult rigid alignment between the prone and supine configurations and the flattening of the tumour. The rigid alignment affects the important landmark and surface constraints used in the method. The choice for assigning a homogeneous material onto the entire breast results in large compressional forces flattening the tumour.

Bibliography

- Alfano, E, Fisac, J. O., García-Sevilla, M., Conde, M. H., Zamora, O. B., Lizarraga, S., Santos, A., Pascau, J., & Carbayo, M. L. (2019). Prone to supine surface based registration workflow for breast tumor localization in surgical planning. *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, 1150–1153.
- An, S. S., Kim, T., & James, D. L. (2008). Optimizing cubature for efficient integration of subspace deformations. *ACM transactions on graphics (TOG)*, 27(5), 1–10.
- Babarenda Gamage, T. P., Malcolm, D. T., Maso Talou, G., Míra, A., Doyle, A., Nielsen, P. M., & Nash, M. P. (2019). An automated computational biomechanics workflow for improving breast cancer diagnosis and treatment. *Interface Focus*, 9(4), 20190034.
- Barbič, J., & James, D. L. (2005). Real-time subspace integration for st. venant-kirchhoff deformable models. *ACM transactions on graphics (TOG)*, 24(3), 982–990.
- Bessa, S., Gouveia, P. F., Carvalho, P. H., Rodrigues, C., Silva, N. L., Cardoso, E., Cardoso, J. S., Oliveira, H. P., & Cardoso, M. J. (2020). 3d digital breast cancer models with multimodal fusion algorithms. *The Breast*, 49, 281–290.
- Bouaziz, S., Martin, S., Liu, T., Kavan, L., & Pauly, M. (2014). Projective dynamics: Fusing constraint projections for fast simulation. *ACM transactions on graphics (TOG)*, 33(4), 1–11.
- Braess, D. (2007). *Finite elements: Theory, fast solvers, and applications in solid mechanics*. Cambridge University Press.
- Brandt, C., Eisemann, E., & Hildebrandt, K. (2018). Hyper-reduced projective dynamics. *ACM Transactions on Graphics (TOG)*, 37(4), 1–13.
- Danch-Wierzchowska, M., Borys, D., & Swierniak, A. (2020). Fem-based mri deformation algorithm for breast deformation analysis. *Biocybernetics and Biomedical Engineering*, 40(3), 1304–1313.
- Eiben, B., Vavourakis, V., Hipwell, J. H., Kabus, S., Buelow, T., Lorenz, C., Mertzaniidou, T., Reis, S., Williams, N. R., Keshtgar, M., et al. (2016). Symmetric biomechanically guided prone-to-supine breast image registration. *Annals of biomedical engineering*, 44(1), 154–173.
- Fulton, L., Modi, V., Duvenaud, D., Levin, D. I., & Jacobson, A. (2019). Latent-space dynamics for reduced deformable simulation. *Computer graphics forum*, 38(2), 379–391.
- Gast, T. F., Schroeder, C., Stomakhin, A., Jiang, C., & Teran, J. M. (2015). Optimization integrator for large time steps. *IEEE transactions on visualization and computer graphics*, 21(10), 1103–1115.
- Griesenauer, R. H., Weis, J. A., Arlinghaus, L. R., Meszoely, I. M., & Miga, M. I. (2017). Breast tissue stiffness estimation for surgical guidance using gravity-induced excitation. *Physics in Medicine & Biology*, 62(12), 4756.
- Guennebaud, G., Jacob, B. et al. (2010). Eigen v3.
- Hang, S. (2015). Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw*, 41(2), 11.
- Harmon, D., & Zorin, D. (2013). Subspace integration with local deformations. *ACM Transactions on Graphics (TOG)*, 32(4), 1–10.
- Hipwell, J. H., Vavourakis, V., Han, L., Mertzaniidou, T., Eiben, B., & Hawkes, D. J. (2016). A review of biomechanically informed breast image registration. *Physics in Medicine & Biology*, 61(2), R1.
- Irving, G., Teran, J., & Fedkiw, R. (2004). Invertible finite elements for robust simulation of large deformation. *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 131–140.
- Jacobson, A., Panozzo, D. et al. (2018). libigl: A simple C++ geometry processing library [https://libigl.github.io/].
- Kim, T., & Eberle, D. (2020). Dynamic deformables: Implementation and production practicalities. *Acm siggraph 2020 courses* (pp. 1–182).
- Kim, T., & James, D. L. (2011). Physics-based character skinning using multi-domain subspace deformations. *Proceedings of the 2011 ACM SIGGRAPH/eurographics symposium on computer animation*, 63–72.
- Koschier, D., Bender, J., & Thuerey, N. (2017). Robust extended finite elements for complex cutting of deformables. *ACM Transactions on Graphics (TOG)*, 36(4), 1–13.
- Kugelstadt, T., Koschier, D., & Bender, J. (2018). Fast corotated fem using operator splitting. *Computer Graphics Forum*, 37(8), 149–160.

- Leal, A. M. M. (2018). Autodiff, a modern, fast and expressive C++ library for automatic differentiation. <https://autodiff.github.io>
- Li, M., Ferguson, Z., Schneider, T., Langlois, T. R., Zorin, D., Panozzo, D., Jiang, C., & Kaufman, D. M. (2020). Incremental potential contact: Intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.*, 39(4), 49.
- Li, X., Li, M., & Jiang, C. (2022). Energetically consistent inelasticity for optimization time integration. *ACM Transactions on Graphics (TOG)*, 41(4), 1–16.
- Nasikun, A., & Hildebrandt, K. (2022). The hierarchical subspace iteration method for laplace–beltrami eigenproblems. *ACM Transactions on Graphics (TOG)*, 41(2), 1–14.
- Nocedal, J., & Wright, S. J. (1999). *Numerical optimization*. Springer.
- Perkins, S. L., Lin, M. A., Srinivasan, S., Wheeler, A. J., Hargreaves, B. A., & Daniel, B. L. (2017). A mixed-reality system for breast surgical planning. *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*, 269–274.
- Sifakis, E., & Barbic, J. (2012). Fem simulation of 3d deformable solids: A practitioner’s guide to theory, discretization and model reduction. *Acm siggraph 2012 courses* (pp. 1–50).
- Smith, B., Goes, F. D., & Kim, T. (2018). Stable neo-hookean flesh simulation. *ACM Transactions on Graphics (TOG)*, 37(2), 1–15.
- Smith, B., Goes, F. D., & Kim, T. (2019). Analytic eigensystems for isotropic distortion energies. *ACM Transactions on Graphics (TOG)*, 38(1), 1–15.
- Teran, J., Sifakis, E., Irving, G., & Fedkiw, R. (2005). Robust quasistatic finite elements and flesh simulation. *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 181–190.
- Wang, B., Wu, L., Yin, K., Ascher, U. M., Liu, L., & Huang, H. (2015). Deformation capture and modeling of soft objects. *ACM Trans. Graph.*, 34(4), 94–1.
- Wang, B., Matcuk, G., & Barbič, J. (2021). Modeling of personalized anatomy using plastic strains. *ACM Transactions on Graphics (TOG)*, 40(2), 1–21.
- Wang, C.-B., Lee, S., Kim, T., Hong, D., Kim, G. B., Yoon, G. Y., Kim, H. H., Kim, N., & Ko, B. (2020). Breast tumor movements analysis using mri scans in prone and supine positions. *Scientific reports*, 10(1), 1–9.
- Wang, H., & Mao, X. (2020). Evaluation of the efficacy of neoadjuvant chemotherapy for breast cancer. *Drug Design, Development and Therapy*, 14, 2423.
- Wu, J., Westermann, R., & Dick, C. (2015). A survey of physically based simulation of cuts in deformable bodies. *Computer Graphics Forum*, 34(6), 161–187.
- Wu, X., Mukherjee, R., & Wang, H. (2015). A unified approach for subspace simulation of deformable bodies in multiple domains. *ACM Transactions on Graphics (TOG)*, 34(6), 1–9.
- Xu, H., Li, Y., Chen, Y., & Barbič, J. (2015). Interactive material design using model reduction. *ACM Transactions on Graphics (TOG)*, 34(2), 1–14.
- Xu, H., Sin, F., Zhu, Y., & Barbič, J. (2015). Nonlinear material design using principal stretches. *ACM Transactions on Graphics (TOG)*, 34(4), 1–11.