Registration by Optimisation over Differentiable Simulation: Topology-adaptive Deformable Registration with MPM

Rasoul Sharifian[1,2,a] Navid Rabbani[1,b] Adrien Bartoli[1,2,c]

[1]DRCI, CHU Clermont-Ferrand, France [2]SurgAR, Surgical Augmented Reality, Clermont-Ferrand, France

[a]rasoul.sharifian.cs@gmail.com [b]navid_rabbani@yahoo.com [c]adrien.bartoli@gmail.com

Abstract

Deformable 3D object registration is a central computer vision problem. Existing methods perform reasonably well under moderate deformations but poorly under topology changes. This is because, regardless of the optimisation-based or learning-based computational approach they use, these methods rely on a fixed 3D model topology. We propose a novel computational approach accommodating the use of advanced physics simulation engines with inherent topological adaptivity. This approach, called Registration by Optimisation over Differentiable Simulation (RODS), can be generically used with any existing differentiable simulation engine. In RODS, registration is formulated as an inverse physics problem and solved using numerical optimisation, where the objective is to estimate the initial conditions that drive the simulation to reach the desired registration state. We implement this approach in a method called RODS-MPM using gradient descent over the mesh-free Material Point Method (MPM). By internally representing objects as collections of material particles without explicit connectivity, RODS-MPM inherently supports large deformations and topology changes. We use the proposed method for surgical guidance, where incisions are modelled as particle separation, without the need for remeshing or manual connectivity updates, enabling a unified treatment of both minor and extreme deformations in the presence of topology changes. Extensive experiments, including toy examples, an ex-vivo dataset with Target Registration Error (TRE) evaluation, and patient data, show that RODS-MPM outperforms existing methods, demonstrating robust registration performance under topology-changing scenarios. Code is available at https://github.com/Rasoul-Sharifian/RODS-MPM.

1. Introduction

Deformable registration is a fundamental problem in computer vision that aims to establish dense spatial correspondences between images or geometric objects. Unlike rigid registration, deformable registration must account for complex deformations arising from articulated motion or material elasticity. Accurate modelling of these transformations is critical in applications ranging from Augmented Reality (AR) [6, 12, 33], dynamic 3D reconstruction [27], and shape analysis [7, 24, 36]. Despite significant progress, achieving robust and physically plausible registration remains a challenging task, particularly in the presence of large deformations [14, 30].

Existing approaches to deformable registration broadly fall into two main categories: optimisation-based and learning-based. Optimisation-based approaches [6, 24] formulate registration as the estimation of a deformation field that minimises a data fidelity objective subject to regularisation constraints. These regularisation terms enforce desirable deformation properties such as smoothness, volume preservation, or as-rigid-as-possible behaviour, and may incorporate simplified biomechanical models based on hyperelastic constitutive laws. While effective in well-constrained scenarios, these methods produce deformation fields that are physically implausible. Learning-based approaches [1, 31, 32] instead learn deformation fields from annotated correspondences or complex simulation engines. These methods can achieve fast inference and capture complex deformation patterns, but their performance heavily depends on the training data distribution and may generalise poorly to real scenarios.

Despite their differences, both categories share a fundamental limitation: the deformation space is not inherently constrained by the physical evolution of the object. As a result, estimated deformations may violate physical consistency. In addition, most existing methods rely on a mesh-based object representation and are topology-preserving. In other words, they do not handle *topological changes*, which occur when the transformation is discontinuous; physically, this corresponds to a loss of proximity between material points that were originally neighbours.

A motivation to study deformable registration is its application to minimally invasive surgery (MIS) [1, 6, 31],

where a camera and specialised surgical instruments are inserted into the body through small incisions. Although MIS has become the preferred approach for many surgical procedures due to its significant clinical benefits, including lower blood loss and faster recovery [38], it also introduces significant technical challenges. Surgeons must operate with limited visual and haptic feedback, while critical internal structures, including vessels and tumours, remain concealed beneath the organ surface. AR has emerged as a promising solution to this limitation by overlaying patient-specific 3D models reconstructed from preoperative CT or MR imaging directly onto intraoperative surgical video. This enables enhanced visualisation of hidden anatomical structures and improves surgical guidance. A fundamental requirement for AR-assisted surgery is deformable registration, which estimates and maintains the spatial correspondence between the preoperative 3D model and the intraoperative images. Accurate registration ensures that virtual anatomical information is correctly aligned with the physical anatomy throughout the surgical procedure.

However, deformable registration of soft tissues is highly challenging due to the complex mechanical behaviour of biological organs. Soft tissues undergo large nonlinear deformations driven by external manipulation, physiological processes such as respiration, and interactions with surrounding organs. Common surgical actions including incision, resection, and stitching constitute topological changes, as they directly alter the proximity of the material points in the anatomical structures. Mesh-based representations inherently struggle to accommodate such changes. Only a few studies [14, 30] have attempted to address topological changes, typically by introducing complex and often suboptimal strategies that require manual mesh modifications or ad hoc handling of connectivities, resulting in numerical instability and reduced reproducibility. As a result, topology-adaptive deformable registration has remained largely unexplored due to its inherent difficulty and the limitations of mesh-based frameworks. The recent advent of meshless methods that do not rely on fixed connectivity makes it timely to take a new look at this important problem.

We propose a novel computational approach termed Registration by Optimisation over Differentiable Simulation (RODS). Instead of directly estimating deformation fields, RODS formulates deformable registration as an inverse physics problem. Specifically, registration is achieved by optimising the initial state of a differentiable physics simulator such that its forward evolution produces a configuration that aligns with intraoperative observations. This formulation inherently constrains the deformation space to physically plausible states defined by the simulator dynamics, eliminating the need for hand-crafted regularisation terms and ensuring physically consistent deformations. Importantly, RODS is a general framework that can be implemented using any differentiable simulation engine.

We accommodate topology changes using a particle-based meshless representation. Unlike mesh-based methods, particle representations [16, 19, 41] do not rely on fixed connectivity and therefore naturally support separation, merging, and discontinuities. Topology changes emerge as a natural consequence of particle motion, rather than requiring explicit connectivity modification or remeshing. Among meshless simulation methods, the Material Point Method (MPM) [19, 28, 41] provides an ideal computational framework for this formulation. MPM represents objects as collections of Lagrangian particles carrying material properties, while computing interactions on a Eulerian grid. This hybrid formulation combines the flexibility of particle methods with the numerical stability of grid-based computation, without requiring explicit mesh connectivity. MPM has been extensively studied [17, 22] in physics simulation and graphics for modelling extreme deformations.

We instantiate the proposed framework in a method termed RODS-MPM, which integrates gradient-based optimisation with a differentiable MPM simulator. By representing anatomical structures as material particles, RODS-MPM naturally supports large deformations and topology changes such as incision and tissue separation, without remeshing or manual intervention. This enables a unified and physically grounded registration framework capable of handling both moderate deformations and topology-changing surgical events. We evaluate the proposed method on deformable registration tasks in MIS involving surgical incisions. Our experiments include 2D illustrative examples, quantitative in-silico assessments, and a qualitative clinical experiment. The main contributions of this work are threefold:

1. RODS, a novel approach that formulates deformable registration as an inverse differentiable physics problem, constraining solutions to physically valid deformations.
2. RODS-MPM, a topology-adaptive registration method that instantiates RODS with an MPM differentiable simulator engine, and naturally handles topology changes without remeshing or manual connectivity updates.
3. The application of RODS-MPM to MIS, with extensive experimental validation demonstrating accurate and physically consistent registration of organ 3D models in spite of topological changes.

2. Related Work

Topology changes under different shape representations. The ability of a simulation method to accommodate topology changes is closely linked to its underlying approach for shape representation. Shape representations can broadly be categorised based on whether they explicitly encode connectivity between structural elements or rely on implicit connectivity. Explicit connectivity representations,

such as surface meshes or volumetric tetrahedral meshes, define fixed relationships between elements through predefined connectivity [45]. This structure provides precise geometric and physical modelling capabilities but does not naturally accommodates topology changes.

In contrast, implicit connectivity representations do not rely on predefined structural relationships. For instance, the level-set representation has an implicit connectivity, which uses the zero level of a scalar field, generally represented on a grid [29, 37]. In this formulation, topology changes such as merging or splitting arise naturally as the level-set surface evolves. This representation has been widely used for image segmentation [5, 23] and attempted for 2D image registration [43]. However, it has not been combined with detailed volumetric deformation models [28]. Particle-based methods provide another implicit representation. Objects are represented as collections of material points that interact through governing physical laws rather than fixed connectivity. In such formulations, topology changes emerge naturally from particle dynamics without requiring explicit connectivity modification. Early particle-based methods such as Particle-In-Cell (PIC) [16] and Fluid-Implicit Particle (FLIP) [4], and more recently the Material Point Method (MPM) [28, 41], thus inherently support topology changes. In these methods, separation and fragmentation arise directly from the underlying physical simulation, making particle-based representations particularly well suited for modelling surgical scenarios involving incisions, tearing, or tissue fragmentation.

Registration in the presence of topology change. Most existing deformable registration methods rely on explicit connectivity representations, limiting their ability to handle topology changes [6, 40]. Only a limited number of works have attempted to address this problem by incorporating explicit topology-handling strategies. Specifically, [14] proposes an image-based method for handling surgical incisions, where topology changes are handled by modifying the mesh connectivity. Similarly, [30] introduces a method to manage cutting through local mesh refinements. While effective in controlled settings, these methods introduce additional complexity, requiring careful mesh-level management and refinements techniques to maintain mesh validity. They are generally not robust, as they depend on detailed prior knowledge about the nature of the topology change.

In contrast, particle-based implicit representations naturally accommodate topology changes through their connectivity-free formulation. However, they have primarily been used for forward physical modelling and animation [17, 22], and their potential for solving deformable registration problems remains largely unexplored.

Differentiable simulation. Differentiable simulators extend classical physics engines by enabling gradient computation through the simulation process, thereby supporting gradient-based optimisation in control, system identification, design optimisation, and policy learning [26]. Existing approaches differ mainly in how gradients are obtained—via analytical derivations, symbolic differentiation, adjoint methods, or automatic differentiation (AD) [2, 15, 34]—and in how they handle contacts, discontinuities, and numerical integration [9, 10].

Early work focused on differentiating rigid-body dynamics with smooth contact models to ensure stable gradients [9], while more recent methods incorporate implicit time integration and differentiable collision handling to improve robustness [15, 20]. Frameworks such as DiffTaichi [18] and Warp [22] demonstrate high-performance differentiable programming for particle-based, rigid, and soft-body systems, whereas gradSim [25] targets robotics and visuomotor learning by providing end-to-end differentiability from physics to rendering.

Collectively, these simulators highlight key trade-offs between physical fidelity, computational efficiency, and gradient accuracy, particularly in the presence of stiff dynamics and non-smooth events such as contact and topology changes [42]. We select DiffTaichi [18] as computational backbone of our topology-adaptive registration framework. Its flexible data-oriented design facilitates the implementation of custom constitutive models and topology-evolving formulations. These properties make it particularly suitable for our setting, where large deformations, proximity changes, and topology variations must be handled within a unified, end-to-end differentiable optimisation pipeline.

3. Methodology

3.1. Registration by Optimisation over Differentiable Simulation

RODS is a computational approach for deformable registration, built around the core component of a differentiable simulator. Let $\mathcal{S}_T$ denote the simulation operator over the time horizon T, which propagates the system state $\mathbf{z}_t$ from time step 0 to time step T, such that $\mathbf{z}_T = \mathcal{S}_T(\mathbf{z}_0)$. The state variables are split into $\mathbf{z}_t = [\mathbf{p}_t, \mathbf{q}_t]^\intercal$. Vector $\mathbf{p}_t$ encodes the object geometry; thus, $\mathbf{p}_0$ and $\mathbf{p}_T$ correspond to the source object and the candidate registered object, respectively. In contrast, $\mathbf{q}_t$ encode all state variables unknown at $t = 0$, primarily the object dynamics, which are optimised by RODS during the registration process.

The objective of RODS is to find the registered object as the optimal simulated configuration $\mathbf{p}_T^*$ that matches a target configuration $\mathcal{T}$ using a cost function $\mathcal{L}(\mathbf{p}_T, \mathcal{T})$. The principle of RODS is to optimise the initial state $\mathbf{q}_0$, obtaining the final state from simulation as $[\mathbf{p}_T^*, \mathbf{q}_T^*]^\intercal = \mathcal{S}_T(\mathbf{p}_0, \mathbf{q}_0^*)$. This is formulated as the following optimisation problem:

$$\mathbf{q}_0^* = \arg\min_{\mathbf{q}_0} \mathcal{L}(\mathbf{p}_T, \mathcal{T}), \quad [\mathbf{p}_T, \mathbf{q}_T]^\intercal = \mathcal{S}_T(\mathbf{p}_0, \mathbf{q}_0). \quad (1)$$

Concretely, in each optimisation iteration, the initial state $[\mathbf{p}_0, \mathbf{q}_0]^\intercal$ is first propagated forward in time to $t = T$ by the simulator. Since the simulator is differentiable with respect to the state variables, gradients of the cost with respect to the initial state can be computed directly. The initial state is then updated from the gradient, with α as the learning rate:

$$\mathbf{q}_0^{n+1} = \mathbf{q}_0^n - \alpha \frac{\partial \mathcal{L}}{\partial \mathbf{q}_0}, \tag{2}$$

This iterates until the termination condition is satisfied.

As the registered model, $\mathbf{p}_T^*$, is the output of a physics simulator, it ensures that the RODS registration output remains physically plausible; in other words the optimisation search space is bound to physically plausible deformations.

3.2. MPM Forward Model for Simulation

Continuum mechanics and elastic materials. Continuum mechanics models the deformation of a solid body through a mapping between a source configuration and its deformed state. Let $\mathbf{X} \in \Omega_0 \subset \mathbb{R}^3$ denote a material point in the source configuration and $\mathbf{x} \in \Omega_t \subset \mathbb{R}^3$ its position in the deformed configuration at time t. The deformation is described by the mapping $\mathbf{x} = \boldsymbol{\varphi}(\mathbf{X}, t)$. The local deformation is characterised by the deformation gradient $\mathbf{F} = \nabla_{\mathbf{X}} \boldsymbol{\varphi}(\mathbf{X}, t)$, which encodes rotation, stretching, and compression of material elements. For hyperelastic materials, stresses are derived from a strain energy density function $\psi(\mathbf{F})$. The corresponding Cauchy stress tensor is given by $\boldsymbol{\sigma} = \frac{1}{J} \frac{\partial \psi}{\partial \mathbf{F}} \mathbf{F}^\intercal$, with $J = \det(\mathbf{F})$.

The dynamics of the material are governed by the conservation laws of continuum mechanics:

$$\rho \frac{D\mathbf{v}}{Dt} = \nabla \cdot \boldsymbol{\sigma} + \mathbf{f}, \quad \frac{D\rho}{Dt} + \rho \nabla \cdot \mathbf{v} = 0, \tag{3}$$

where ρ is the density, $\mathbf{v}$ is the velocity field, and $\mathbf{f}$ represents external forces.

We model soft tissue using a Neo-Hookean hyperelastic material, whose strain energy density function is:

$$\psi(\mathbf{F}) = \frac{\mu}{2} \left(\mathbf{F}^T \mathbf{F} - 3 \right) - \mu \ln J + \frac{\lambda}{2} (\ln J)^2. \tag{4}$$

The Lamé parameters μ and λ are computed from the Young's modulus E and Poisson's ratio ν.

Material Point Method. MPM solves the governing equations (3) numerically using a hybrid Eulerian Lagrangian simulation technique, widely used for modelling deformable materials. In MPM, the continuum is discretised as a set of N material particles. Each particle p represents a small volume of material and carries properties including mass m^p, volume V^p, position $\mathbf{x}_t^p$, velocity $\mathbf{v}_t^p$, and deformation gradient $\mathbf{F}_t^p$.

MPM evolves the system through a particle–grid transfer scheme. At each time step, particle quantities are first transferred to a background grid (particle-to-grid, P2G). The momentum at a grid node i is computed as:

$$m_t^i \mathbf{v}_t^i = \sum_p N(\mathbf{x}^i - \mathbf{x}_t^p) \Bigg[m^p \mathbf{v}_t^p + \left(m^p \mathbf{C}_t^p - \frac{4 \Delta t V^p}{(\Delta x)^2} \frac{\partial \psi}{\partial \mathbf{F}^p} \mathbf{F}_t^{p\intercal} \right) (\mathbf{x}^i - \mathbf{x}_t^p) \Bigg] + \mathbf{f}_t^i, \tag{5}$$

where $N(\cdot)$ denotes the B-spline interpolation kernel, Δx is the grid spacing, Δt the simulation time step, and $\mathbf{f}_t^i$ external forces applied on the grid. The updated grid velocities are then interpolated back to the particles (grid-to-particle, G2P) to update their positions as:

$$\mathbf{v}_{t+1}^k = \sum_i N(\mathbf{x}^i - \mathbf{x}_t^k) \mathbf{v}^i, \quad \mathbf{x}_{t+1}^p = \mathbf{x}_t^p + \Delta t \mathbf{v}_{t+1}^p. \tag{6}$$

Finally, to update the states, the local velocity gradient and deformation gradient are updated as:

$$\mathbf{C}_{t+1}^p = \frac{4}{(\Delta x)^2} \sum_i N(\mathbf{x}^i - \mathbf{x}_t^p) \mathbf{v}_t^i (\mathbf{x}^i - \mathbf{x}_t^p)^\intercal \tag{7}$$

$$\mathbf{F}_{t+1}^p = \left(\mathbf{I} + \Delta t \sum_i \mathbf{v}_t^i \nabla N(\mathbf{x}^i - \mathbf{x}_t^p)^\intercal \right) \mathbf{F}_t^p. \tag{8}$$

The background grid acts as a temporary computational structure that is reconstructed at every simulation step, allowing efficient computation of spatial derivatives while naturally supporting large deformations and topology changes. Modern MPM implementations such as Diff-Taichi [17] and Warp [22] are differentiable. Consequently, gradients of simulation outputs with respect to particle states can be computed through automatic differentiation across all simulation steps, enabling the use of MPM as a differentiable forward model within RODS.

3.3. RODS-MPM

We instantiate the general RODS approach using a differentiable MPM simulator; the forward operator $\mathcal{S}_T(\mathbf{p}_0, \mathbf{q}_0)$ is thus MPM simulation over T time steps.

Object representation and landmarks. The source object is represented using N material particles. Following the RODS formulation, the optimisable system states are selected as initial particle velocities, $\mathbf{q}_0 = [\mathbf{v}_0^i]_{i=1}^N$, while the rest of states which encodes the object geometry are the initial particle positions $\mathbf{p}_0 = [\mathbf{x}_0^i]_{i=1}^N$. A set of source landmarks $\mathcal{M}_s = \{\mathbf{x}_0^k\}_{k \in \mathcal{K}}$, with corresponding target landmarks $\mathcal{M}_t = \{\mathbf{y}^k\}_{k \in \mathcal{K}}$, is selected, where $\mathcal{K} \subset \{1, \ldots, N\}$ denotes the set of landmark indices with $|\mathcal{K}| \ll N$. A target landmark $\mathbf{y}^k$ is either a point, a line, a curve, or a region.

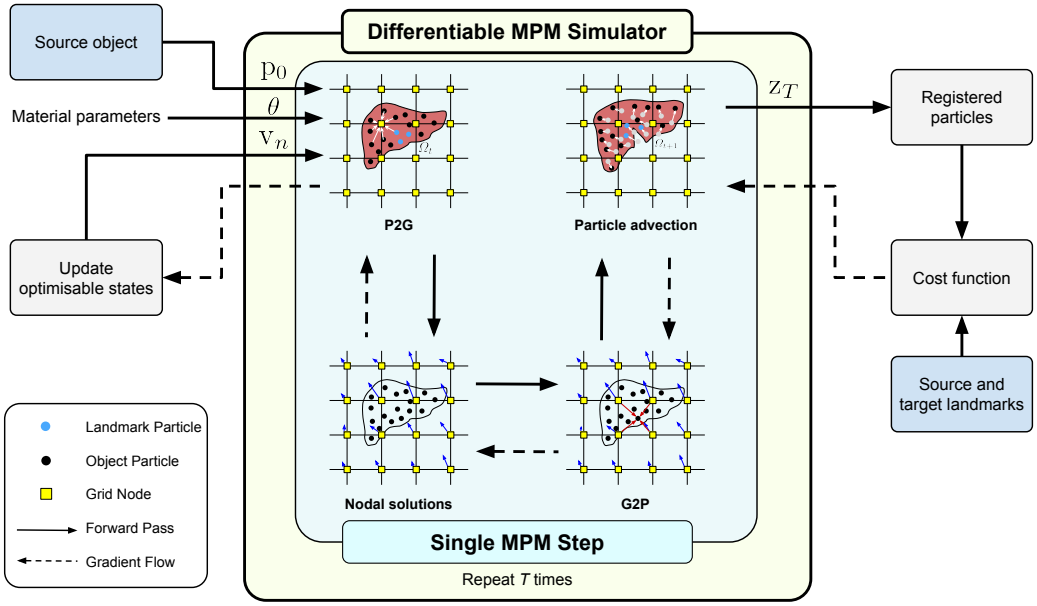


Figure 1. Overview of the proposed RODS-MPM. The source configuration is evolved using a differentiable MPM solver. A landmark-based cost is computed with respect to the target configuration, and gradients are backpropagated to optimise the initial particle velocities.

Registration formulation and optimisation. In the RODS formulation, registration is achieved by optimising the initial particle state such that the simulated landmark positions match the target landmarks. By optimising the initial particle velocities, the registration problem (1) becomes:

$$\begin{aligned} \mathbf{v}_0^* &= \arg\min_{\mathbf{v}_0} \frac{1}{|\mathcal{K}|} \sum_{k\in\mathcal{K}} \mathrm{dist}\left(\mathbf{x}_T^k, \mathbf{y}^k\right) \\ \text{s.t.} \quad & [\mathbf{x}_T^k, \mathbf{v}_T^k]^\top = \mathcal{S}_T(\mathbf{x}_0, \mathbf{v}_0), \end{aligned} \tag{9}$$

where $\mathrm{dist}(\cdot)$ denotes the distance between a particle and its target landmark. Following [35], we use the distance between the particle and the closest point on the target landmark, which is known to be differentiable [13] and can therefore be efficiently optimised using gradient-based methods. The optimised velocities should be interpreted as proxy control variables rather than estimates of the true physical motion. The objective is to obtain a physically plausible deformation that satisfies the landmark alignment objective under the constraints imposed by the simulator. Instead of explicitly estimating unknown forces and boundary conditions, their cumulative effect is implicitly encoded in the optimised velocity field.

Since the MPM simulator is fully differentiable, gradients of the cost with respect to the initial particle velocities are computed by automatic differentiation across all simulation steps and used to update the initial state using gradient descent. The optimisation drives the simulated object toward alignment with the target configuration and ensure that the deformation remains physically consistent with the underlying material model. We refer to this differentiable physics-based registration method, illustrated in Fig. 1, as RODS-MPM; the algorithmic workflow is in Algorithm 1.

Algorithm 1 RODS-MPM

Require: Initial particle positions $\mathbf{x}_0 = \{\mathbf{x}_0^i\}_{i=1}^N$; material parameters $\boldsymbol{\theta}$; landmark indices $\mathcal{K}$; target landmarks $\mathcal{M}_t = \{\mathbf{y}^k\}_{k\in\mathcal{K}}$; differentiable MPM simulator $\mathcal{S}_T$; distance function $\mathrm{dist}(\cdot)$; learning rate α; iterations N_{it}

Ensure: Optimised initial velocities $\mathbf{v}_0^\star$ and registered configuration $\mathbf{x}_T^\star$

Initialise velocities $\mathbf{v}_0 \leftarrow \mathbf{0}$; initialise internal states (e.g., $\mathbf{F}_0^i \leftarrow \mathbf{I}$)
for $n = 1$ to N_{it} **do**
  Forward simulation $(\mathbf{x}_T, \cdot) \leftarrow \mathcal{S}_T(\mathbf{x}_0, \mathbf{v}_0)$
  Compute cost

$$\mathcal{L} = \frac{1}{|\mathcal{K}|} \sum_{k\in\mathcal{K}} \mathrm{dist}(\mathbf{x}_T^k, \mathbf{y}^k)$$

  Backpropagate gradient $\mathbf{g} \leftarrow \nabla_{\mathbf{v}_0}\mathcal{L}$
  Update velocities $\mathbf{v}_0 \leftarrow \mathbf{v}_0 - \alpha\,\mathbf{g}$
end for
return $\mathbf{v}_0^\star \leftarrow \mathbf{v}_0$ and $\mathbf{x}_T^\star \leftarrow \mathbf{x}_T$

4. Experiments and Results

We evaluate the proposed RODS-MPM method in 3 settings. First, illustrative 2D examples, demonstrating the ca-

pability of the method; second, *in-silico* experiments with realistic liver deformations and surgical incisions, enabling quantitative evaluation; third, clinical case. All experiments are implemented in DiffTaichi [17].

4.1. Illustrative Examples in 2D

We use simple 2D geometric examples designed to illustrate RODS-MPM on both topology-preserving and topology-changing deformations. Figure 2 presents four scenarios in which a square domain, uniformly sampled with particles, undergoes deformations guided by different point and curve landmarks. The scenarios illustrate representative behaviours including a) large elastic deformation, b) curve-landmark-driven registration, c) topological change (rupture), and d) combined large deformation and topology change. Depending on the task, the registration target is defined either through sparse point landmarks (a, b), through a distance function to a target curve (c), or through a combination of both (d). Within the RODS-MPM framework, topology changes arise naturally from the spatial configuration of these constraints. We observe that, in spite of landmark sparsity, the registered object naturally follows the deformation behaviour of hyperelastic materials. In particular, in Fig. 2(a), guided by only two point landmarks, RODS-MPM stretches the object horizontally while compressing its height, reproducing the characteristic response of elastic materials. In addition, landmark placements that induce separation behaviour, implicitly drive topology-changing deformations without requiring explicit modification of connectivity, as seen in Fig. 2(c) and (d).

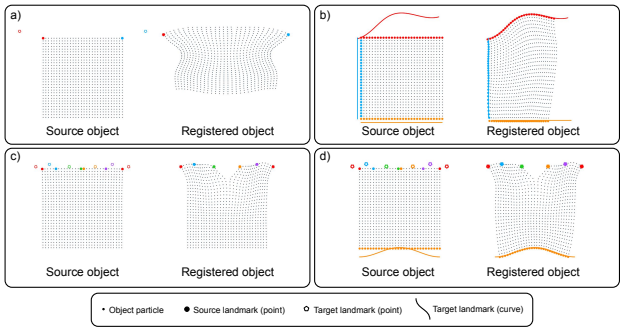


Figure 2. Illustrative examples showing the behaviour of RODS-MPM under different geometric constraints. The mesh-free particle representation allows topology changes to emerge naturally without explicit connectivity modification or remeshing.

4.2. In-Silico Experiments

Baseline methods. We compare RODS-MPM against four representative baselines: 1) ARAP [40] (geometry-driven regularisation widely used in non-rigid registration), 2) Deformable Registration with Biomechanical Constraints (DRBC) [6] (minimises data fidelity regularised by a Neo-Hookean internal strain energy), 3) LDDMM [3], and 4) the resection-aware registration method [14].

Dataset generation. We generate a synthetic dataset of liver deformations with simulated surgical incisions. Patient-specific 3D liver models reconstructed from preoperative imaging data [35] are used as the source anatomy, over which biomechanically realistic deformations with topology changes are simulated. We consider two clinically relevant scenarios [8]: *(i) non-anatomical resection*, typically performed for localised tumour removal with a 1 cm oncologic margin, and *(ii) anatomical resection*, corresponding to the removal of an entire liver Couinaud segment along anatomical boundaries. The dataset comprises 8 simulated resections, 4 of each scenarios.

The simulation pipeline consists of two stages. First, a surgical incision is created by subtracting a CAD model of a surgical blade from the volumetric liver mesh, introducing a geometric discontinuity in the mesh. Second, mechanical forces are applied to the two opposing surfaces of the incision using FEBio [21] to simulate tissue separation during surgery. FEBio (Finite Elements for Biomechanics) is a software package for finite element analysis specifically designed for biomechanics and bioengineering.

The liver tissue is modelled as a nearly incompressible Neo-Hookean hyperelastic material with standard liver parameters (Young's modulus $E = 4.1$ kPa and Poisson's ratio $\nu = 0.49$). To simulate tissue separation along the incision, equal opposite forces are applied to the two sides of the cut surface, inducing progressive deformation and opening of the tissue. The simulation is run for $T = 200$ time steps. The initial configuration is the source model and the final one is the target geometry. The two sides of the incision boundary are annotated using polylines in both the source and target geometries. These polylines represent sparse surgical landmarks that guide the registration process. Clinically, they would be automatically detected [14].

Implementation details. The source volumetric liver model is uniformly sampled with a spacing of 0.6 mm to generate material particles, resulting in approximately 6.8k particles. Material behaviour is modelled using a hyperelastic constitutive model parameterised by $\theta_1 = 4.1$ kPa and $\theta_2 = 0.49$. Source landmark particles are defined as the material particles closest to the annotated incision points in the source configuration.

The registration target is represented as a polyline. The cost function measures the Euclidean distance between each source landmark particle and the target polyline. Gradients of this cost with respect to the initial particle velocities are computed through automatic differentiation across the differentiable MPM simulator. The optimisation is performed for 500 iterations using gradient descent with an initial learning rate of 25 and a decay factor of 0.5. The result-

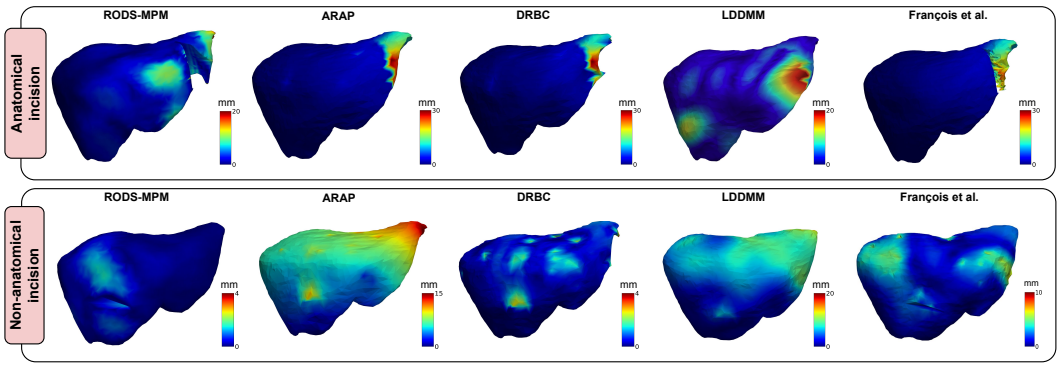


Figure 3. Spatial error maps of the registration results, measured as Euclidean distances between registered nodes and ground-truth model.

ing deformed particle set is finally converted into a triangulated surface mesh using the Alpha Shapes algorithm [11].

Results. We report the Target Registration Error (TRE) measured at mesh nodes. Concretely, we measure the node-wise Euclidean distance between the registered model and the ground-truth model. Qualitative results are shown in Fig. 3, where the estimated deformations are compared against the FEM-simulated ground truth.

The results show that RODS-MPM successfully handles topology changes in both anatomical and non-anatomical resection scenarios. In contrast, ARAP, DRBC, and LDDMM struggle to accommodate incision-induced discontinuities. Because these methods rely on explicit mesh connectivity, they attempt to satisfy the landmark constraints by stretching the mesh across the incision, leading to unrealistic deformations and large local errors. The method of [14] also captures the topology change but introduces larger errors in regions away from the incision. By contrast, RODS-MPM maintains lower errors across the entire organ due to its particle-based representation, which naturally supports material separation and topological changes.

To analyse the spatial distribution of registration errors, Fig. 4 plots the mean and standard deviation of the error as a function of the distance to the nearest incision landmark. In the non-anatomical resection scenario, RODS-MPM achieves consistently lower errors across all distances from the incision. In the anatomical resection scenario,

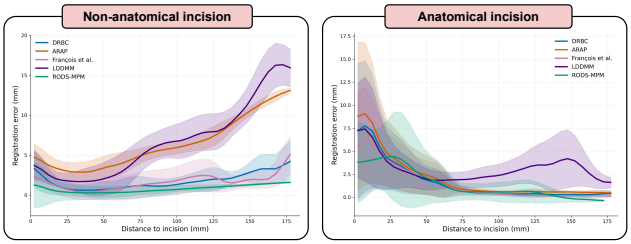


Figure 4. Quantitative evaluation of registration error as a function of the distance to the closest incision landmark (mean and std).

Table 1. TRE (mean ± standard deviation, mm).

	# landmark points			Gaussian noise σ (mm)		
	20	50	250	2.5	5.0	10.0
ARAP	3.12±2.54	2.05±2.43	1.95±2.08	2.82±2.92	3.04±3.12	3.41±3.48
DRBC	2.45±2.10	1.67±2.16	1.61±1.06	1.98±1.92	2.18±2.01	2.31±2.05
François et al.	2.31±2.18	1.80±1.86	1.72±1.14	1.93±1.87	2.08±1.99	2.24±1.91
LDDMM	4.50±3.23	3.46±2.86	2.83±2.26	3.73±3.07	4.24±3.52	4.54±3.89
Ours	**1.75±1.95**	**1.47±1.60**	**1.03±0.82**	**1.65±1.79**	**1.82±1.96**	**2.03±2.24**

the improvement is particularly stronger near the incision boundary, where topology changes occur. Farther away from the incision region, all methods produce comparable results, reflecting the smoother deformation regime away from the topological discontinuity. While [14] outperforms the other baselines near the incision, RODS-MPM achieves the lowest errors overall, including in the vicinity of the topology change.

Robustness. To evaluate the impact of landmark quantity and quality on the proposed method, we conducted additional experiments measuring the TRE under varying numbers of landmark points (20, 50, 250) and under different levels of Gaussian noise added to the landmark positions ($\sigma = 2.5, 5.0, 10.0\text{mm}$), as reported in table 1. The results demonstrate the robustness of the method to both sparse and noisy landmark configurations.

Sensitivity analysis. We evaluated the sensitivity of RODS-MPM to key simulation parameters, namely the Young's modulus E, grid resolution n_{grid}, simulation horizon T, and the number of particles. As shown in figure 5, RODS-MPM is largely insensitive to moderate variations in both material and numerical parameters. TRE remains stable across a wide range of Young's modulus values, indicating robustness to imperfect material calibration. Increasing the grid resolution, simulation horizon, and particle count improves accuracy only up to a convergence point, beyond which the gains become negligible. These results suggest that the RODS-MPM achieves stable performance over a broad range of parameter choices.

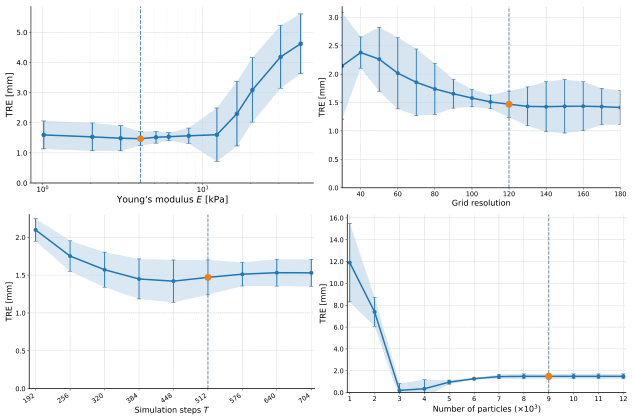


Figure 5. Sensitivity to MPM parameters. Orange circles indicate the experimental settings used in the experiments.

4.3. Clinical Experiment

We evaluate RODS-MPM retrospectively on stereoscopic images from a liver surgery case (collected from partner hospital under ethical clearance, IRB approval **anonymous**). The intraoperative data is a stereoscopic image sequence and the preoperative 3D liver model is reconstructed from an MRI scan through the segmentation and meshing pipeline in MITK [44]. The stereoscopic camera is calibrated prior to the procedure using a checkerboard pattern and OpenCV. The captured stereo images are rectified using the calibrated intrinsic and extrinsic camera parameters; depth is then recovered by computing disparity with the stereo matching models [39], allowing the reconstruction of a dense intraoperative 3D point-cloud of the liver.

The task is to register the preoperative 3D model to the current stereo image pair. In a real-time surgical registration system, the preoperative 3D model is already registered to prior image pairs, whilst the system processes image pairs sequentially. In order to emulate these conditions, we select two stereo image pairs from the video (Fig. 6). The first stereo image pair represents the past. It is taken before the incision occurs, but shortly after the surgeon marks the intended incision curve on the organ surface. The second stereo image pair represents the present. It is taken after the incision has been made. These deliberate choices demonstrate RODS-MPM's ability to handle incisions. The preoperative 3D model is manually registered to the reconstructed intraoperative surface using identifiable anatomical landmarks. The incision is also annotated in the left image of the stereo pair and transferred to the preoperative model.

We now have a preoperative 3D model comprising the incision boundary, serving as source landmarks for registration, and a stereo image pair showing the incised organ, where the anterior surface incision boundary is marked, serving as target landmarks for registration. Importantly, the preoperative model corresponds to the organ prior to

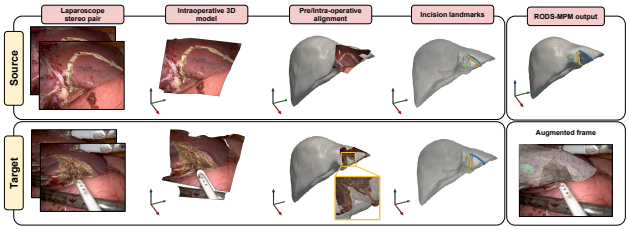


Figure 6. Registration of a preoperative liver model to intraoperative stereo targets using RODS-MPM. The reconstructed intraoperative geometry and annotated anterior surface incision guide the deformation; the model is then augmented on the surgical image.

incision, resulting in significant deformation and topological changes between the model and the organ configuration observed in the present image pair. The next step is deformable registration. We apply RODS-MPM whose inputs are the preoperative 3D model and the incision landmarks. The registered particles are then used to estimate the deformation field by simple barycentric coordinate interpolation, which is subsequently used to register the internal structures, such as the tumour. Since the output of RODS-MPM is a particle set, the resulting geometry is converted to a triangulated surface mesh using the Alpha Shapes algorithm [11] for display. The registered liver and the tumour it contains are projected onto the stereo image and visualised as an augmented reality overlay on the target frame.

5. Conclusion

We have proposed RODS, a computational registration approach that formulates deformable registration as an inverse problem over a differentiable physics simulator. By optimising the simulator's initial states, RODS restricts the solution space to physically plausible trajectories, reducing reliance on hand-crafted regularisation while improving realism under large deformations. We instantiated RODS as RODS-MPM using a differentiable Material Point Method solver. This formulation naturally accommodates incision-induced topology changes without remeshing or explicit connectivity updates, enabling a unified treatment of both continuous deformation and tissue separation. Experiments demonstrate that RODS-MPM achieves robust registration, particularly near incision regions where classical methods introduce unrealistic stretching. These results highlight the potential of differentiable, topology-adaptive simulators as a computational framework for surgical registration, enabling a unified treatment of both hyperelastic deformation and topology-changing events. The price of high accuracy is the higher computational cost required by the multiple forward simulations. Future work will integrate faster simulation engines, test the joint estimation of physical parameters, and integrate richer image targets.

References

[1] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019. 1

[2] Atilim Günes Baydin, Barak A. Pearlmutter, Alexey A. Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, 18:1–43, 2018. 3

[3] M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005. 6

[4] Jeremiah U Brackbill, Douglas B Kothe, and Hans M Ruppel. Flip: a low-dissipation, particle-in-cell method for fluid flow. *Computer Physics Communications*, 48(1):25–38, 1988. 3

[5] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International journal of computer vision*, 22(1):61–79, 1997. 3

[6] Toby Collins, Daniel Pizarro, Simone Gasparini, Nicolas Bourdel, Pauline Chauvet, Michel Canis, Lilian Calvet, and Adrien Bartoli. Augmented reality guided laparoscopic surgery of the uterus. *TMI*, 40(1):371–380, 2020. 1, 3, 6

[7] Timothy F Cootes, Christopher J Taylor, David H Cooper, and Jim Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1): 38–59, 1995. 1

[8] Claude Couinaud. *Le foie; études anatomiques et chirurgicales*. Masson, 1957. 6

[9] Jonas Degrave, Michiel Hermans, and Joni Dambre. A differentiable physics engine for deep learning in robotics. *Frontiers in Neurorobotics*, 13:6, 2019. 3

[10] Tao Du, Josiah P. Hughes, Sean Wah, and Wojciech Matusik. DiffPD: Differentiable projective dynamics. In *ACM SIGGRAPH Asia*, 2022. 3

[11] Herbert Edelsbrunner and Ernst P Mücke. Three-dimensional alpha shapes. *ACM Transactions On Graphics (TOG)*, 13(1):43–72, 1994. 7, 8

[12] Yamid Espinel, Lilian Calvet, Karim Botros, Emmanuel Buc, Christophe Tilmant, and Adrien Bartoli. Using multiple images and contours for deformable 3d-2d registration of a preoperative ct in laparoscopic liver surgery. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 657–666. Springer, 2021. 1

[13] Andrew W Fitzgibbon. Robust registration of 2d and 3d point sets. *Image and vision computing*, 21(13-14):1145–1153, 2003. 5

[14] Tom François, Lilian Calvet, Callyane Sève-d'Erceville, Nicolas Bourdel, and Adrien Bartoli. Image-based incision detection for topological intraoperative 3d model update in augmented reality assisted laparoscopic surgery. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 647–656. Springer, 2021. 1, 2, 3, 6, 7

[15] Matthias Geilinger, David Hahn, Bernhard Thomaszewski, and Stelian Coros. ADD: Analytically differentiable dynamics for multi-body systems with frictional contact. In *ACM SIGGRAPH Asia*, 2020. 3

[16] Francis H Harlow. The particle-in-cell method for numerical solution of problems in fluid dynamics. Technical report, Los Alamos Scientific Lab., N. Mex., 1962. 2, 3

[17] Yuanming Hu, Tzu-Mao Li, Luke Anderson, Jonathan Ragan-Kelley, and Frédo Durand. Taichi: a language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics (TOG)*, 38(6): 1–16, 2019. 2, 3, 4, 6

[18] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable programming for physical simulation. *International Conference on Learning Representations (ICLR)*, 2020. 3

[19] Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. The affine particle-in-cell method. *ACM Transactions on Graphics (TOG)*, 34(4):1–10, 2015. 2

[20] Yuanming Li, Mridul Aanjaneya, Yuanming Hu, and Eftychios Sifakis. A moving least squares material point method with displacement discontinuity and two-way rigid body coupling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 37(4):150:1–150:14, 2018. 3

[21] Steve A Maas, Benjamin J Ellis, Gerard A Ateshian, and Jeffrey A Weiss. Febio: finite elements for biomechanics. 2012. 6

[22] Miles Macklin. Warp: A high-performance python framework for gpu simulation and graphics. `https://github.com/nvidia/warp`, 2022. NVIDIA GPU Technology Conference (GTC). 2, 3, 4

[23] Ravi Malladi, James A Sethian, and Baba C Vemuri. Shape modeling with front propagation: A level set approach. *IEEE transactions on pattern analysis and machine intelligence*, 17(2):158–175, 2002. 3

[24] Jan Modersitzki. *Numerical methods for image registration*. OUP Oxford, 2003. 1

[25] J. K. Murthy, Miles Macklin, Florian Golemo, Vikram Voleti, Ludovic Petrini, Martin Weiss, Brennan Considine, Jonathan Parent-Lévesque, Junjie Xie, Stefano Ermon, et al. gradSim: Differentiable simulation for system identification and visuomotor control. In *International Conference on Learning Representations (ICLR)*, 2021. 3

[26] Rhys Newbury, Jack Collins, et al. A review of differentiable simulators. *IEEE Access*, 2024. Early Access. 3

[27] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015. 1

[28] Vinh Phu Nguyen, Alban De Vaucorbeil, and Stephane Bordas. The material point method. *Cham: Springer International Publishing*, 2023. 2, 3

[29] Stanley Osher and James A Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of computational physics*, 79 (1):12–49, 1988. 3

[30] Christoph J Paulus, Nazim Haouchine, Seong-Ho Kong, Renato Vianna Soares, David Cazier, and Stéphane Cotin. Handling topological changes during elastic registration: application to augmented reality in laparoscopic surgery. *International journal of computer assisted radiology and surgery*, 12(3):461–470, 2017. 1, 2, 3

[31] Micha Pfeiffer, Carina Riediger, Jürgen Weitz, and Stefanie Speidel. Learning soft tissue behavior of organs for surgical navigation with convolutional neural networks. *International journal of computer assisted radiology and surgery*, 14(7): 1147–1155, 2019. 1

[32] Micha Pfeiffer, Carina Riediger, Stefan Leger, Jens-Peter Kühn, Danilo Seppelt, Ralf-Thorsten Hoffmann, Jürgen Weitz, and Stefanie Speidel. Non-rigid volume to surface registration using a data-driven biomechanical model. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 724–734. Springer, 2020. 1

[33] Julien Pilet, Vincent Lepetit, and Pascal Fua. Fast non-rigid surface detection, registration and realistic augmentation. *IJCV*, 76(2):109–122, 2008. 1

[34] René-E. Plessix. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167(2):495–503, 2006. 3

[35] Navid Rabbani, Lilian Calvet, Yamid Espinel, Bertrand Le Roy, Mathieu Ribeiro, Emmanuel Buc, and Adrien Bartoli. A methodology and clinical dataset with ground-truth to evaluate registration accuracy quantitatively in computer-assisted laparoscopic liver resection. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, 10(4):441–450, 2022. 5, 6

[36] Daniel Rueckert, Alejandro F Frangi, and Julia A Schnabel. Automatic construction of 3D statistical deformation models of the brain using nonrigid registration. *IEEE transactions on medical imaging*, 22(8):1014–1025, 2003. 1

[37] James A Sethian et al. *Level set methods and fast marching methods*. Cambridge Cambridge UP, 1999. 3

[38] Abdul Razak Shaikh, Ahmed Khan Sangrasi, and Gulshan Ara Shaikh. Clinical outcomes of laparoscopic versus open appendectomy. *JSLS: Journal of the Society of Laparoendoscopic Surgeons*, 13(4):574, 2009. 2

[39] Rasoul Sharifian, Navid Rabbani, Yongcong Zhang, and Adrien Bartoli. Camera augmentation: enabling uncalibrated stereo matching of minimally invasive surgery images by training from the wealth of public synthetic image datasets. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–9, 2026. 8

[40] Olga Sorkine, Marc Alexa, et al. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, pages 109–116, 2007. 3, 6

[41] Deborah Sulsky, Zhen Chen, and Howard L Schreyer. A particle method for history-dependent materials. *Computer methods in applied mechanics and engineering*, 118(1-2): 179–196, 1994. 2, 3

[42] Marc A Toussaint, Kelsey Rebecca Allen, Kevin A Smith, and Joshua B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. 2018. 3

[43] Baba C Vemuri, J Ye, Yunmei Chen, and Christiana Morison Leonard. Image registration via level-set motion: Applications to atlas-based segmentation. *Medical image analysis*, 7 (1):1–20, 2003. 3

[44] Ivo Wolf, Marcus Vetter, Ingmar Wegner, Marco Nolden, Thomas Bottger, Mark Hastenteufel, Max Schobinger, Tobias Kunert, and Hans-Peter Meinzer. The medical imaging interaction toolkit (mitk): a toolkit facilitating the creation of interactive software by extending vtk and itk. In *Medical imaging 2004: Visualization, image-guided procedures, and display*, pages 16–27. SPIE, 2004. 8

[45] Afra J Zomorodian. *Topology for computing*. Cambridge university press, 2005. 3