

Robust variational segmentation of 3D bone CT data with thin cartilage interfaces

Tarun Gangwar^{a,*}, Jeff Calder^b, Takashi Takahashi^c, Joan E. Bechtold^d, Dominik Schillinger^a

^a*Department of Civil, Environmental, and Geo-Engineering, University of Minnesota, Twin Cities, USA*

^b*School of Mathematics, University of Minnesota, Twin Cities, USA*

^c*Department of Radiology, University of Minnesota, Twin Cities, USA*

^d*Department of Orthopedic Surgery, University of Minnesota, Twin Cities, USA*

Abstract

We present a two-stage variational approach for segmenting 3D bone CT data that performs robustly with respect to thin cartilage interfaces. In the first stage, we minimize a flux-augmented Chan-Vese model that accurately segments well-separated regions. In the second stage, we apply a new phase-field fracture inspired model that reliably eliminates spurious bridges across thin cartilage interfaces, resulting in an accurate segmentation topology, from which each bone object can be identified. Its mathematical formulation is based on the phase-field approach to variational fracture, which naturally blends with the variational approach to segmentation. We successfully test and validate our methodology for the segmentation of 3D femur and vertebra bones, which feature thin cartilage regions in the hip joint, the intervertebral disks, and synovial joints of the spinous processes. The major strength of the new methodology is its potential for full automation and seamless integration with downstream predictive bone simulation in a common finite element framework.

Keywords: Variational segmentation, 3D bone CT data, thin cartilage interfaces, flux-augmented Chan-Vese model, phase-field fracture mechanics, voxel finite elements, vertebra extraction, femur extraction

***Corresponding Author:** Tarun Gangwar, Department of Civil, Environmental, and Geo-Engineering, University of Minnesota, 500 Pillsbury Drive S.E., Minneapolis, MN 55455, USA; Phone: +1 612 624-0063; Fax: +1 612 626 7750;

Email addresses: gangw007@umn.edu (Tarun Gangwar), jcalder@umn.edu (Jeff Calder), takas005@umn.edu (Takashi Takahashi), bechtold@umn.edu (Joan E. Bechtold), dominik@umn.edu (Dominik Schillinger)

Contents

1	Introduction	3
2	Theoretical background and methods	4
2.1	A flux-augmented Chan-Vese functional for bone/tissue separation	4
2.1.1	The Chan-Vese energy functional	5
2.1.2	Allen-Cahn mean curvature flow and a local flux augmentation	5
2.1.3	Discretization based on voxel finite elements	7
2.1.4	A representative benchmark problem in 2D	8
2.2	The variational approach to brittle fracture and its phase-field approximation	9
2.2.1	Variational description of fracture	10
2.2.2	Phase-field approximation of the fracture surface and energy degradation	10
2.2.3	Variational formulation	11
2.3	A phase-field fracture inspired functional for eliminating fine-scale contacts .	12
2.3.1	A simplified phase-field fracture formulation based on a Laplace problem	12
2.3.2	Variational formulation and discretization with voxel finite elements .	13
2.3.3	A staggered solution algorithm	15
2.3.4	Choosing appropriate Dirichlet constraints	15
2.4	The complete methodology: overview and implementation	16
3	Results and discussion	19
3.1	Femur segmentation: separation from pelvis	19
3.1.1	Clinical CT data set from the University of Minnesota AHC	20
3.1.2	Automated variational segmentation	21
3.1.3	Validation by global overlap metrics	24
3.1.4	Validation by local HU averages at characteristic surface positions .	25
3.2	Vertebra segmentation: separation from spine	26
3.2.1	Test data from the OsiriX DICOM image library	28
3.2.2	Automated variational segmentation	28
3.3	Discussion of strengths and opportunities	29
4	References	31

1. Introduction

In the future, physicians will utilize patient-specific predictive simulations based on diagnostic imaging data to better treat individual patients. For many clinically relevant scenarios, physiology-based modeling and simulation methods exist today (see, e.g., [1, 2, 3, 4, 5, 6, 7, 8]). A major roadblock, however, is the intricate process of transferring diagnostic imaging information into explicit geometric models of patient-specific physiological systems. A fundamental component of this process is segmentation, that is, assigning a label to every voxel in a 3D image that belongs to a target physiological object [9]. Robust, accurate and fully automated segmentation of 3D imaging data, however, still constitutes one of the grand challenges in medical image analysis, despite significant research efforts over the last decades [9, 10, 11]. The segmentation methodologies most widely used in applications usually emphasize accuracy and robustness at the cost of automation. As a consequence, segmentation procedures in clinical practice and biomedical research still largely rely on the supervision and manual intervention of segmentation experts.

In the context of physiology-based modeling and simulation, using a variational approach to segmentation is particularly attractive, as both approaches can be based on a common variational framework, including the same numerical discretization method, to find solutions to the underlying partial differential equations. Variational segmentation methods are based on finding the minimizing function of an energy-type functional whose contour represents the boundaries of segmented objects to be identified [12]. Suitable energy functionals can flexibly incorporate edge based local information, region based global information or any other a-priori knowledge about the target segmentation region. Parametric active contour models [13, 14, 15] initialize a parametric contour surface and energy minimization moves the parametric representation towards the boundaries of the segmentation region. Geometric deformable models [16, 17] develop a segmentation by evolving implicit contour functions that are independent of any parametrization. Level set methods [18, 19] form the basic numerical framework for the solution of these models.

The presence of thin interfaces between different objects in practical medical image representations, usually coupled with limited resolution and fuzzy color information, is a major challenge for segmentation methodologies. In this paper, we consider this problem in the form of bone objects to be segmented from 3D CT scans that are separated by thin cartilage interfaces. An important example is the segmentation of the femur, which requires detecting the thin cartilage layer at the hip joint to establish a clean separation from the pelvis. Prior work on this problem has mostly focused on non-variational segmentation methods. For the femur bone, several authors have worked on improving thresholding techniques based on tensor-based filters [20, 21] or 3D correlation between bone objects [22]. Another approach is to integrate a-priori geometric information, for example the close similarity of the femoral head with a perfect sphere [23, 24]. Significant attention has also been devoted to statistical shape models that separate the femur and pelvis by determining the significant mode of variations from sets of training data [25, 26, 27].

In this article, we focus on developing a variational methodology for the segmentation of 3D bone CT data. Our goal is a methodology that is accurate, performs robustly with respect

to thin cartilage interfaces and can be seamlessly integrated with downstream predictive bone modeling and simulation methods in a common finite element framework. At the same time, our methodology should be able to operate independently from critical input or manual intervention of a computational analyst.

Our basic idea is to devise a two-stage segmentation procedure that leverages recent developments in variational image segmentation of bone structures with geometric deformable models. In the first stage, we use a flux-augmented two-phase Chan-Vese model [28, 29] to find a contour function that reliably segments the well-separated parts of all bone objects in the image. In the second stage, we devise a new phase-field fracture inspired method that reliably eliminates all spurious contact bridges across thin cartilage interfaces, resulting in an accurate segmentation topology, where the segmentation region of each individual bone object can be easily identified. The underlying conceptual idea for the new method is based on engineering intuition, establishing an analogy between fracture mechanics and the segmentation problem of eliminating thin spurious contacts. For its mathematical formulation, we generalize the phase-field approach to variational fracture based on the minimization of an energy functional [30, 31], which naturally fits to the minimization of the flux-augmented Chan-Vese functional in the first stage.

Our article is structured as follows: In Section 2, we present a detailed derivation of our two-stage variational segmentation methodology and its implementation in the context of voxel finite elements. In Section 3, we demonstrate accuracy, robustness and potential for full automation of the new segmentation methodology for the femur-pelvis problem and the separation of an individual vertebra from a vertebral column. In particular, we present an in-depth validation study for a 22-sample series of clinical 3D femur CT data obtained at the Academic Health Center of the University of Minnesota. We close our article by discussing strengths and weaknesses of our new methodology.

2. Theoretical background and methods

In this section, we introduce our approach to variational segmentation of bone CT data with thin cartilage interfaces. Our discussion follows the structure of the approach, which consists of two basic parts. In the first subsection, we review the derivation of an energy functional whose minimizer extracts bone geometry from imaging data. Its contour solution reliably separates bone and surrounding tissue, but small spurious contacts remain in the fine interface region between separate bone objects. In the second and third subsections, we motivate and derive in detail an additional phase-field fracture inspired functional whose minimizer is able to reliably detect and eliminate all contact regions from the contour solution. The fourth subsection provides a concise summary of all technical aspects.

2.1. A flux-augmented Chan-Vese functional for bone/tissue separation

We are given a 3D voxel (or a 2D pixel) representation of an image over a domain Ω , such that its local intensity is given as $I(x) : \Omega \rightarrow [0, 1]$. We will assume in the following that we can always generate approximations to the piecewise constant image I that have sufficient smoothness for the operations we need to perform on I . The first step of our methodology is based on an implicit geometric deformable model, which we briefly describe in the following.

2.1.1. The Chan-Vese energy functional

Implicit geometric deformable models are typically based on simplified Mumford-Shah invariants [32, 33] that enable a robust numerical treatment. One of the most successful is the Chan-Vese active contour model without edges [28]. It assumes that we can identify two regions $\Omega_1(C)$ and $\Omega_2(C)$ with approximatively piecewise constant intensities, separated by a deformable curve C . This assumption gives rise to the following two fitting terms:

$$F_1(C) + F_2(C) := \int_{\Omega_1} |I - c_1|^2 dx + \int_{\Omega_2} |I - c_2|^2 dx \quad (1)$$

where c_1 and c_2 are defined as the average of the image intensity I in $\Omega_1(C)$ and $\Omega_2(C)$, respectively. The curve C that minimizes the functional (1) represents the boundary of the object to be identified.

The unknown curve C can be represented as the zero isocontour of a deformable contour function $\phi(x)$. Adding a length penalty term and using the properties of the Heaviside function H , the functional (1) can be rewritten in terms of ϕ into the following Chan-Vese energy functional:

$$\begin{aligned} \mathcal{F}_{\text{CV}}(c_1, c_2, \phi) &:= \int_{\Omega} |\nabla H(\phi)| + \lambda \{ H(\phi)(I - c_1)^2 + (1 - H(\phi))(I - c_2)^2 \} d\Omega \quad (2) \\ \text{with } c_1 &= \frac{\int_{\Omega} H(\phi) I d\Omega}{\int_{\Omega} H(\phi) d\Omega} \quad \text{and} \quad c_2 = \frac{\int_{\Omega} (1 - H(\phi)) I d\Omega}{\int_{\Omega} (1 - H(\phi)) d\Omega} \end{aligned}$$

where $\lambda \geq 0$ is a weighting coefficient. The corresponding Euler-Lagrange equation for ϕ yields the following strong form gradient descent partial differential equation:

$$\frac{\partial \phi}{\partial t} = |\nabla \phi| \left\{ \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - \lambda \{ (I - c_1)^2 - (I - c_2)^2 \} \right\} \quad (3)$$

For the evolution of the contour function ϕ in (3), its curvature represented as

$$\kappa = \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right), \quad (4)$$

plays a key role. This term leads to Hamilton-Jacobi type equations, which is usually solved by upwind finite difference schemes [34].

2.1.2. Allen-Cahn mean curvature flow and a local flux augmentation

The Allen-Cahn energy functional and its corresponding partial differential equation [35, 36, 37] are expressed as follows:

$$\mathcal{F}_{\text{AC}}(u) = \int_{\Omega} \left(\frac{\epsilon^2}{2} |\nabla u|^2 + W(u) \right) d\Omega \quad (5)$$

$$\frac{\partial u}{\partial t} = \epsilon^2 \Delta u - W'(u) \quad (6)$$

If we choose a double-well potential $W(\xi) = \frac{1}{4}\xi^2(1 - \xi)^2$, the phase-field contour solution u is either zero or one, separated by a diffuse interface region, where the solution rapidly changes between zero and one. The parameter ϵ controls the characteristic length-scale of the diffuse interface region. The unknown curve C that defines the segmentation region is now represented as the isocontour $u = 0.5$.

Phase-field formulations based on the Allen-Cahn equation share many key features with deformable active contour models, based on the analogy between mean curvature flow expressed in terms of a contour gradient as in (4) and its regularization in terms of the Allen-Cahn equation. EVANS, SONER and SOUGANIDIS showed that in the limit as ϵ goes to zero, the Allen-Cahn formulation yields motion by mean curvature of the interface that separates the zero and one regions of the solution [38].

MERRIMAN, BENCE and OSHER [39] were among the first to represent the motion of mean curvature with the help of the Allen-Cahn phase-field formulation (5). Hence, the curvature term (4) in the Chan-Vese energy functional (3) can be replaced by the computationally more tractable terms of the Allen-Cahn energy functional (5). The result is a diffuse two-phase Chan-Vese model described by the following energy functional [40]:

$$\mathcal{F}_{CV}^\epsilon(u, c_1, c_2) := \int_{\Omega} \frac{\epsilon^2}{2} |\nabla u|^2 + W(u) + \lambda \{u^2(I - c_1)^2 + (1 - u)^2(I - c_2)^2\} dx \quad (7)$$

We emphasize that the phase-field formulation (7) leads to a variational form that can be robustly solved with finite element methods [41]. Its variation with respect to u gives the following strong form gradient descent partial differential equation:

$$\frac{\partial u}{\partial t} = \epsilon^2 \Delta u - W'(u) - 2\lambda \{u(I - c_1)^2 + (1 - u)(I - c_2)^2\} \quad (8)$$

The functional (7) represents a global measure of the characteristics of the segmentation region. Segmentation regions in bone CT scans, however, often have non-negligible image intensity boundaries, where the intensity gradient vector field is pointing inwards. To leverage this observation for image segmentation, CALDER, TAHMASEBI and MANSOURI [29] suggested a flux augmentation to (7) that incorporates a local measure of region boundaries in the image. This flux term reads:

$$\mathcal{F}_{flux} := - \oint \nabla I \cdot \mathbf{n} ds \quad (9)$$

where integration is performed over the isocontour $u = 0.5$ representing the segmentation curve C . The vector \mathbf{n} is the unit inward normal to C and the vector ∇I is the gradient of the image intensity function. Adding the flux augmentation (9) to the functional (7), the final energy functional for bone/tissue separation follows as:

$$\begin{aligned} \mathcal{F}_{flux}^\epsilon(u, c_1, c_2) := & \int_{\Omega} \frac{\epsilon^2}{2} |\nabla u|^2 + W(u) + \\ & + \lambda \{u^2(I - c_1)^2 + (1 - u)^2(I - c_2)^2\} d\Omega - \mu \oint \nabla I \cdot \mathbf{n} ds \end{aligned} \quad (10)$$

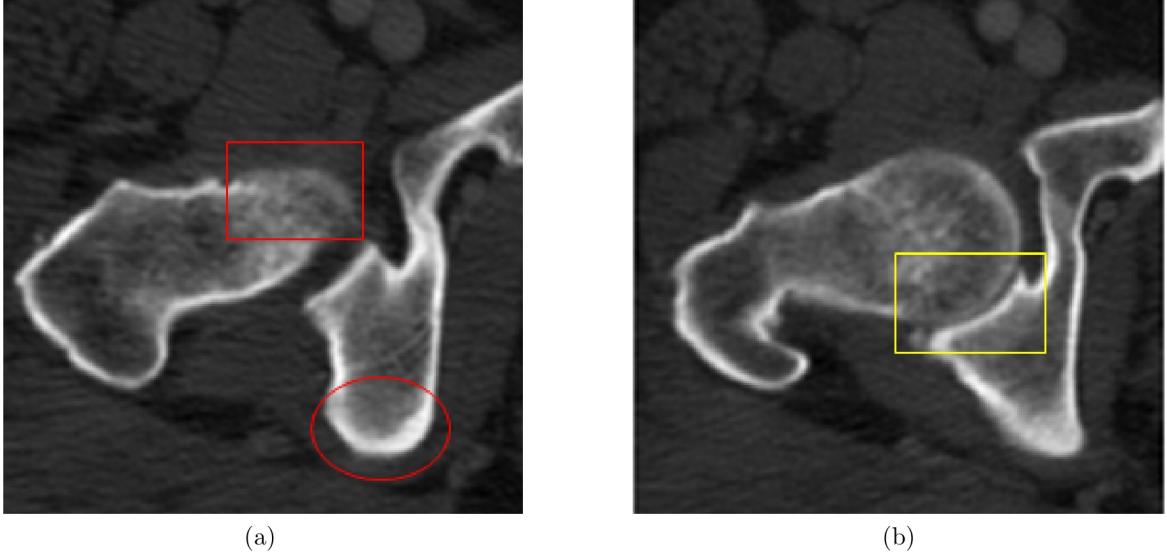


Figure 1: Two slices taken from a stack of CT images of the femur-pelvis configuration: (a) well separated regions (square and ellipse: strong cortical shell in contrast with interior cancellous region) (b) poor separation across thin cartilage region (square).

$$\text{with } c_1 = \frac{\int_{\Omega} u I \, d\Omega}{\int_{\Omega} u \, d\Omega} \quad \text{and} \quad c_2 = \frac{\int_{\Omega} (1-u) I \, d\Omega}{\int_{\Omega} (1-u) \, d\Omega}$$

where $\mu \geq 0$ is a weighting coefficient. The corresponding gradient descent equation reads:

$$\frac{\partial u}{\partial t} = \epsilon^2 \Delta u - W'(u) - 2\lambda \{u(I - c_1)^2 + (u - 1)(I - c_2)^2\} - \mu \Delta I \quad (11)$$

2.1.3. Discretization based on voxel finite elements

A salient feature of the formulation (10) and (11) is that it can be transferred into a variational format suitable for discretization with standard finite elements [41]. In the context of image analysis, voxel finite element methods [42, 43, 44] that associate each 3D voxel with one linear hexahedral element are particularly suitable, since their approximation power directly corresponds to the available image resolution. In addition, the discrete image representation directly provides the finite element mesh, so that voxel finite element methods do not require additional meshing procedures.

Multiplying the gradient descent equation (11) with a test function v , integrating over the image domain Ω and applying integration by parts to all terms that involve the Laplace operator results in the following variational formulation: Find $u_h \subset H^1(\Omega)$ such that

$$\begin{aligned} & \left(\frac{\partial u_h}{\partial t}, v_h \right) + \epsilon^2 \left(\nabla u_h, \nabla v_h \right) + \left(W'(u_h), v_h \right) + \\ & + 2\lambda \left(u_h \{(I - c_1)^2 + (I - c_2)^2\}, v_h \right) - 2\lambda \left((I - c_2)^2, v_h \right) + \end{aligned}$$

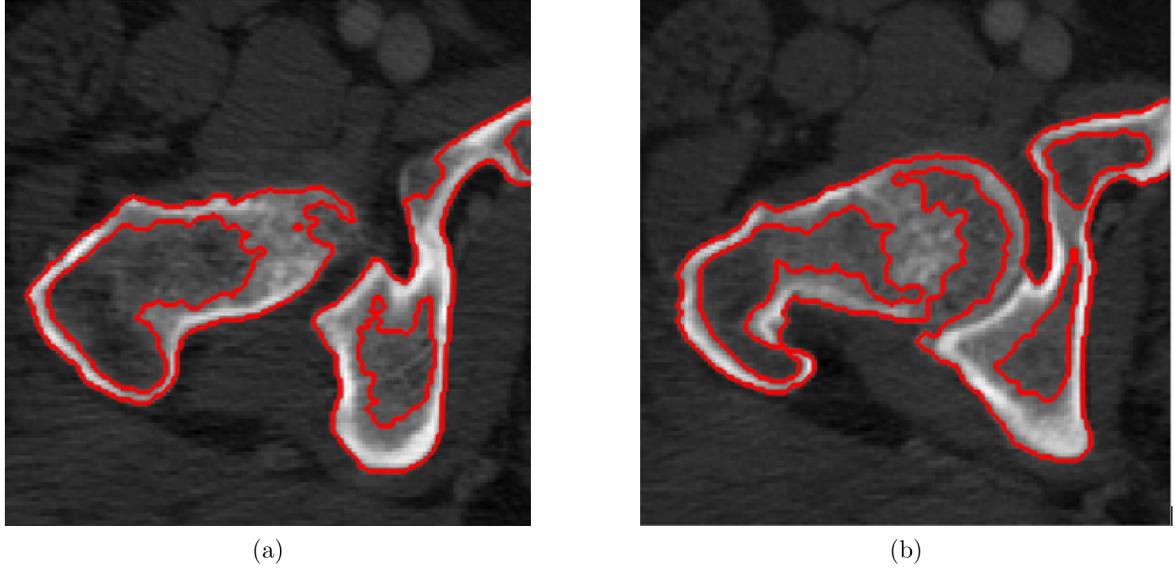


Figure 2: Diffuse Chan-Vese model (7) without flux augmentation: contour function plotted in red at $u = 0.5$ for the two slices of the 2D model problem.

$$+ \alpha(\Delta I, v_h) = 0, \quad \forall v_h \subset H^1(\Omega) \quad (12)$$

where we assume that functions $c_1(u)$ and $c_2(u)$ are known at any given point in time. The first and second derivatives in the term ΔI are evaluated by standard finite difference stencils on the center points of the structured voxel grid.

The variational form (12) is then discretized in time with a first order semi-implicit finite difference scheme [45] that reads: Find $u_h \subset H^1(\Omega)$ such that

$$\begin{aligned} \frac{1}{\Delta t} (u_h^{n+1} - u_h^n, v_h) + \epsilon^2 (\nabla u_h^{n+1}, \nabla v_h) + (W'(u_h^n), v_h) + \\ + 2\lambda (u_h^{n+1} \{(I - c_1^n)^2 + (I - c_2^n)^2\}, v_h) - 2\lambda ((I - c_2^n)^2, v_h) + \\ + \alpha(\Delta I, v_h) = 0, \quad \forall v_h \subset H^1(\Omega) \end{aligned} \quad (13)$$

where linear and nonlinear expressions are evaluated implicitly and explicitly, respectively.

2.1.4. A representative benchmark problem in 2D

We take a step back and consider two simple 2D model problems shown in Fig. 1. They consist of two slices from a CT scan of the femur and pelvis, exhibiting typical features of bone CT scans. The “coarse-scale” bone objects can be easily distinguished from the surrounding soft tissue due to the bright cortical shell region near the bone boundaries and the light gray textured cancellous part in the bone interior regions. In Fig. 1a, the 2D femur and pelvis regions are well separated. In Fig. 1b, however, the two bone regions are only

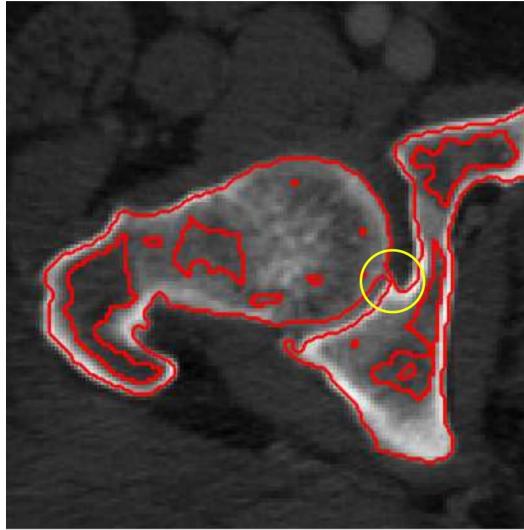


Figure 3: Diffuse Chan-Vese model (10) with flux augmentation: contour function plotted at $u = 0.5$ for the second slice of the 2D model problem. A spurious bridge between femur and pelvis remains (circle).

separated by a thin “fine-scale” sliver of cartilage in the hip joint whose thickness is resolved by just a few pixels.

We now use the two slices to demonstrate the performance of image segmentation based on the functionals discussed above. Each image has a resolution of 140×140 pixels. The weak form of each variational problem is discretized by two triangular finite elements per pixel and the finite difference time discretization is solved until the steady state is reached. Figure 2 illustrates the results obtained with the basic two-phase Chan-Vese model (7) by plotting the segmentation contour $u = 0.5$ in red over the image. We observe that this model works well for the slice where the two bone sections are well separated (see Fig. 2a). For the slice with the thin cartilage separation, however, it fails to resolve the two sections, integrating the femur and pelvis regions (see Fig. 2b).

Figure 3 illustrates the result obtained with the flux-augmented diffuse two-phase Chan-Vese model (10) for the second poorly resolved slice. We observe that the flux-augmented model is able to separate the two bone regions well. In particular, it detects the thin cartilage region, since it incorporates local information based on the flux augmentation (9). However, small bridges between the femur and pelvis sections still remain. They can be easily detected visually and removed manually by a computational analyst, but preclude this method from being applied autonomously.

2.2. The variational approach to brittle fracture and its phase-field approximation

Before deriving our new approach to resolve the issue of fine-scale contacts in the cartilage region, we briefly review the variational formulation of the phase-field model for quasi-static fracture. We largely follow the presentation given in SCHILLINGER, BORDEN and STOLARSKI [46] and the references therein.

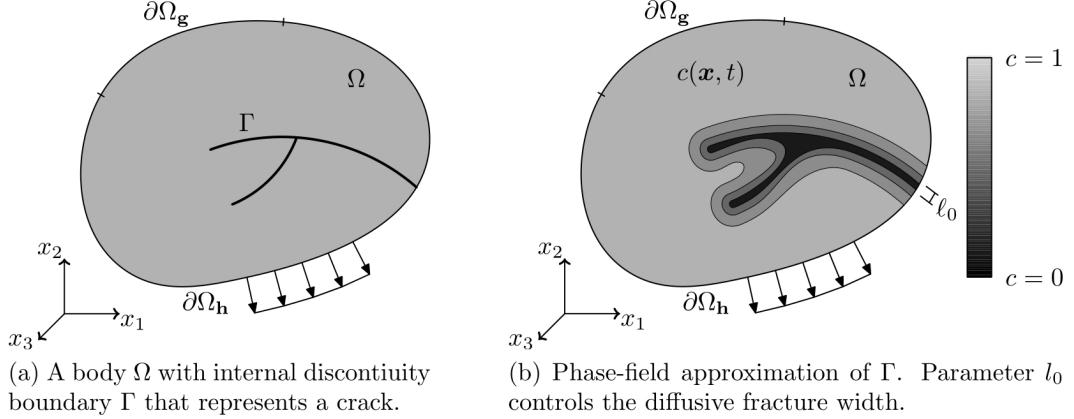


Figure 4: The concept of phase-field fracture (pictures from [49], courtesy of Michael Borden).

2.2.1. Variational description of fracture

We consider an arbitrary body $\Omega \subset \mathbb{R}^d$ (with $d \in \{1; 2; 3\}$) with external boundary $\partial\Omega$ and internal displacement discontinuity Γ , as illustrated in Fig. 4a. The symmetric strain tensor is given as:

$$\boldsymbol{\varepsilon} = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^T) \quad (14)$$

where $\mathbf{u}(x, t)$ denotes the displacement at each material point x . The elastic energy stored in the bulk of the solid per unit volume is described by the energy density function

$$\psi_0(\boldsymbol{\varepsilon}) = \frac{1}{2} \lambda_e \operatorname{tr}(\boldsymbol{\varepsilon})^2 + \mu_e \boldsymbol{\varepsilon} : \boldsymbol{\varepsilon} \quad (15)$$

where λ_e and μ_e are the Lamé constants of elasticity. The evolving internal discontinuity boundary $\Gamma(t)$ represents a set of discrete cracks. Following Griffith [47] and Irwin [48], the work required to create a unit area of fracture surface is equal to the critical energy release rate \mathcal{G}_c . The total internal potential of the body is given by the functional:

$$\mathcal{F}_{\text{pot}}(\boldsymbol{\varepsilon}, \Gamma) = \int_{\Omega} \psi_0(\boldsymbol{\varepsilon}) \, d\Omega + \int_{\Gamma} \mathcal{G}_c \, d\Gamma \quad (16)$$

The first term represents the energy stored in the bulk of the elastic solid, and the second term represents the work necessary to create the current fracture topology $\Gamma(t)$. Following FRANCFOR and MARIGO [30], the variational approach to fracture predicts the nucleation, propagation and interaction of cracks by finding a global minimizer of (16) for a given load.

2.2.2. Phase-field approximation of the fracture surface and energy degradation

To solve this variational problem numerically, BOURDIN, FRANCFOR and MARIGO [31] introduced a volumetric approximation to the surface integral:

$$\int_{\Gamma} \mathcal{G}_c \, d\Gamma \approx \int_{\Omega} \mathcal{G}_c \gamma_c \, d\Omega \quad (17)$$

This approximation uses a smooth scalar-valued phase-field, $c \in [0; 1]$, to represent the crack. It has a value of one away from the crack and a value of zero at the crack (see Fig. 4b). The phase-field approximation introduces a second-order crack density function:

$$\gamma_c = \frac{1}{4l_0} [(c - 1)^2 + 4l_0^2 |\nabla c|^2] \quad (18)$$

whose characteristic interface width is controlled by the length-scale parameter l_0 .

In the next step, we multiply the reference energy density function (15) associated with the undamaged elastic solid with the phase-field:

$$\tilde{\psi}(\boldsymbol{\varepsilon}, c) = [(1 - \kappa) c^2 + \kappa] \psi_0(\boldsymbol{\varepsilon}) \quad (19)$$

From a physical point of view, the phase-field locally penalizes the capability of the material to carry stress across cracks. The model parameter $\kappa \ll 1$ introduced by AMBROSIO and TORTORELLI [50] prevents the full degradation of the stored energy by maintaining a small artificial tensile energy density $\kappa\psi_0$ at the fully-broken state $c = 0$.

2.2.3. Variational formulation

The variational form of the phase-field fracture problem can be derived from an incremental variational principle that balances the rate of different energy terms. In the quasi-static case considered here, we have the rate of the stored energy, the rate of dissipated energy due to the work done by fracture, and the rate of the energy due to the work of external forces. Their balance leads to the following rate of energy functional

$$\dot{\mathcal{E}}_{\text{int}} + \dot{\mathcal{F}}_{\text{frac}} - \dot{\mathcal{P}}_{\text{ext}} = 0 \quad (20)$$

The rate of the stored energy can be computed from the stored energy function (19) as

$$\dot{\mathcal{E}}_{\text{int}} = \frac{d}{dt} \int \tilde{\psi} d\Omega = \int \boldsymbol{\sigma} : \dot{\boldsymbol{\varepsilon}} d\Omega + \int 2(1 - \kappa)\psi_0 c \dot{c} d\Omega \quad (21)$$

where the stress tensor can be computed from (19) as

$$\boldsymbol{\sigma} := \partial_{\boldsymbol{\varepsilon}} \tilde{\psi}(\boldsymbol{\varepsilon})_0 = [(1 - \kappa) c^2 + \kappa] (\lambda_e \langle \text{tr}(\boldsymbol{\varepsilon}) \rangle \boldsymbol{I} + 2\mu \boldsymbol{\varepsilon}) \quad (22)$$

The second term of (21) contains the reference energy density function ψ_0 and describes the local intensity of the deformation. This term can be interpreted as an energetic force that drives the crack evolution [51]. An intuitive, simple and effective idea to incorporate the irreversibility of the crack topology is to introduce a local history field:

$$\mathcal{H}(x, t) = \max_{s \in [0; t]} \psi_0(\boldsymbol{\varepsilon}(x, s)) \quad (23)$$

It records the maximum positive reference energy density that has occurred during quasi-static loading at a specific location x up to the current (pseudo-)time instant t . Replacing

ψ_0 by \mathcal{H} in (21) guarantees that the corresponding energy is dissipated from the system and cannot be returned.

The rate of the dissipated energy due to work done by fracture follows with the phase-field approximation of the fracture topology (17) as

$$\dot{\mathcal{F}}_{\text{frac}} = \frac{d}{dt} \int \mathcal{G}_c \gamma_c d\Omega = \int \frac{\mathcal{G}_c}{2l_0} (c - 1) \dot{c} d\Omega + \int 2\mathcal{G}_c l_0 \nabla c (\nabla \dot{c}) d\Omega \quad (24)$$

The rate of energy due to the work done by external forces is simply

$$\mathcal{P}_{\text{ext}} = \int \mathbf{b} \cdot \dot{\mathbf{u}} d\Omega + \int \mathbf{t} \cdot \dot{\mathbf{u}} d\partial\Omega \quad (25)$$

where \mathbf{b} and \mathbf{t} denote body forces and boundary tractions, respectively.

With the argument that the balance (20) must hold for arbitrary $\dot{\mathbf{u}}$ and \dot{c} , we can separate the rate of energy functional into a phase-field part and an elasticity part. Identifying $\dot{\mathbf{u}}$ and \dot{c} as test functions \mathbf{w} and q , we can write the weak form of the phase-field fracture problem: Find $c \in H^1(\Omega)$ and $\mathbf{u} \in H^1(\Omega)$ such that

$$\int_{\Omega} \left(\frac{4l_0(1 - \kappa)\psi_0}{\mathcal{G}_c} + 1 \right) c q d\Omega + \int_{\Omega} 4l_0^2 \nabla c \nabla q d\Omega = \int_{\Omega} q d\Omega \quad \forall q \in H^1(\Omega) \quad (26)$$

$$\int_{\Omega} \boldsymbol{\sigma} : \nabla \mathbf{w} d\Omega = \int \mathbf{b} \cdot \mathbf{w} d\Omega + \int_{\Omega} \mathbf{t} \cdot \mathbf{w} d\partial\Omega \quad \forall \mathbf{w} \in H^1(\Omega) \quad (27)$$

2.3. A phase-field fracture inspired functional for eliminating fine-scale contacts

Our central idea to automatically detect and remove fine-scale contacts due to thin cartilage regions from an existing contour function follows from engineering intuition. For a physical body, a simple way to separate contacts that are much smaller than the characteristic length scale of the body is to apply some forcing, which will break any contact bridges immediately due to their small size. To transfer this idea to the image segmentation problem at hand, it needs to be formulated as a mathematical model alongside an appropriate numerical discretization method. In our context, a phase-field fracture based approach offers significant advantages over other computational fracture methodologies. The phase-field fracture formulation can handle complex crack patterns in three dimensions, where most other computational fracture models have severe difficulties. In addition, it is intimately related to the variational approach to image segmentation, so that the two can be naturally combined. Both formulations share the concept of a diffuse phase-field solution controlled by a length-scale parameter that represents complex geometric features implicitly. In fact, their origin can be traced back to the same Mumford-Shah functional [32].

2.3.1. A simplified phase-field fracture formulation based on a Laplace problem

We return to the image domain Ω and the associated image intensity function I . Instead of the displacement vector function \mathbf{u} of the previous section, we consider a scalar function h , defined over Ω . If one prefers to work with a physical analogy in mind, the function h could be interpreted as a temperature distribution.

Following Section 2.1, we can assume that we know the contour solution u that minimizes the flux-augmented Chan-Vese functional (10). The isocontour $u = 0.5$ represents the boundaries of the correct segmentation regions, except for small spurious bridges in thin cartilage regions between individual bone objects. This contour function can therefore be used to define an indicator function $\alpha(u)$ as follows:

$$\alpha(u) = \begin{cases} 1.0, & \text{if } u \geq 0.5 \\ 0.0, & \text{if } u < 0.5 \end{cases} \quad (28)$$

which distinguishes the segmentation region $\alpha = 1$ from the rest of the domain Ω .

In analogy to the strain energy density function (15), we define the following energy density associated with the function h :

$$\psi_0(\nabla h) = \frac{1}{2} \alpha(u) (\nabla h)^2 \quad (29)$$

Adding the indicator function α enables us to extend the problem definition to the complete domain Ω in a fictitious domain sense [43]. We can naturally restrict the problem definition to the segmentation region again, when we discretize the weak form.

In analogy to the variational description of fracture (16) and its phase-field approximation (17), we define an equivalent potential energy functional for the modified scalar problem:

$$\mathcal{F}_{\text{pot}}(\nabla h, c) = \int_{\Omega} \tilde{\psi}_0(\nabla h, c) d\Omega + \int_{\Omega} \mathcal{G}_c \gamma_c d\Omega \quad (30)$$

where we use the same crack density function (18) as in the previous subsection. We keep the material parameter $\mathcal{G}_c > 0$ to tune the weighting between the bulk term and surface term. In (30), we penalize the energy density function by the phase-field c , following the concept of energy degradation (19):

$$\tilde{\psi}_0(\nabla h, c) = [(1 - \kappa) c^2 + \kappa] \frac{1}{2} \alpha(\nabla h)^2 \quad (31)$$

where the parameter $\kappa \ll 1$ again prevents full degradation. It is important to note that in contrast to the functionals discussed in Section 2.1, the minimizers of the potential energy functional (30) is constrained by appropriate Dirichlet boundary conditions for h .

2.3.2. Variational formulation and discretization with voxel finite elements

Following the concepts stated in Section 2.2.3, we can now derive the weak form of the modified coupled phase-field fracture problem: Find $c \in H^1(\Omega)$ and $h \in H^1(\Omega)$ such that

$$\int_{\Omega} \left(\frac{4 l_0 (1 - \kappa)}{\mathcal{G}_c} \mathcal{H} + 1 \right) c q d\Omega + \int_{\Omega} 4 l_0^2 \nabla c \cdot \nabla q d\Omega = \int_{\Omega} q d\Omega, \quad \forall q \in H^1(\Omega) \quad (32)$$

$$\int_{\Omega} [(1 - \kappa) c^2 + \kappa] \alpha \nabla h \cdot \nabla w d\Omega = 0, \quad \forall w \in \{H^1(\Omega), w|_{\partial\Omega_D} = 0\} \quad (33)$$

where we leave the definition of the Dirichlet boundary $\partial\tilde{\Omega}_D$ in (33) unspecified for the moment. In analogy to (23), we replace the reference energy density function ψ_0 in (32) by a history field:

$$\mathcal{H}(x, t) = \max_{s \in [0; t]} \psi_0(\nabla h(x, s)) \quad (34)$$

to dissipate fracture energy and prevent crack healing during quasi-static loading.

We can then discretize the variational problem (32) and (33) with voxel finite elements, using one hexahedral finite element per voxel. Since the part of the domain where the indicator function is zero is not of interest, all voxel elements with $\alpha(u) = 0$ in the center are removed from the discretization. We note that there exist other suitable voxel based schemes to discretize (32) and (33), such as the higher-order voxel finite cell method [44, 52, 53].

The remaining voxel finite elements form a modified discretization that covers all voxels associated with the segmented region, approximating the smooth contour solution u of the flux-augmented Chan-Vese functional. Based on this approximation, we define $\tilde{\Omega}$ as the domain covered by the voxel finite element discretization, enclosed by the set of boundary facets $\partial\tilde{\Omega}$, which can be split into a Neumann part $\partial\tilde{\Omega}_N$ and a Dirichlet part $\partial\tilde{\Omega}_D$. The discretized weak form of the coupled phase-field problem then reads: Find $c_h \subset H^1(\tilde{\Omega})$ and $h_h \subset H^1(\tilde{\Omega})$ such that

$$\left(\left(\frac{4l_0(1-\kappa)}{\mathcal{G}_c} \mathcal{H} + 1 \right) c_h, q_h \right) + \left(4l_0^2 \nabla c_h, \nabla q_h \right) - \left(1, q_h \right) = 0 \quad \forall q_h \subset H^1(\tilde{\Omega}) \quad (35)$$

$$\left([(1-\kappa) u^2 + \kappa] \nabla h_h, \nabla w_h \right) = 0 \quad \forall w_h \subset \left\{ H^1(\tilde{\Omega}), w|_{\partial\tilde{\Omega}_D} = 0 \right\} \quad (36)$$

When we go back to continuous functions c and h again, take into account that (35) and (36) hold for arbitrary q_h and w_h and apply integration by parts, we find the strong form of the modified phase-field fracture problem on $\tilde{\Omega}$:

$$\left\{ \begin{array}{ll} \left(\frac{4l_0(1-\kappa)}{\mathcal{G}_c} \mathcal{H} + 1 \right) c + 4l_0^2 \Delta c = 1 & \text{on } \tilde{\Omega} \\ \nabla c \cdot \mathbf{n} = 0 & \text{on } \partial\tilde{\Omega} \\ [(1-\kappa) c^2 + \kappa] \Delta h = 0 & \text{on } \tilde{\Omega} \\ h = g & \text{on } \partial\tilde{\Omega}_D \\ \nabla h \cdot \mathbf{n} = 0 & \text{on } \partial\tilde{\Omega}_N \end{array} \right. \quad (37)$$

The first two equations in (37) represent the partial differential equation for the evolution of the phase-field and its zero-flux boundary condition, respectively. We observe that the phase-field evolution is driven by the history field. If $\mathcal{H} = 0$, the phase-field solution is one everywhere. If $\mathcal{H} > 0$, the phase-field solution $c(x)$ needs to have non-zero second derivatives to satisfy the partial differential equation. A large local value of the history field forces the phase-field solution to have large local gradients, forming a diffuse crack whose width is controlled by the length-scale parameter l_0 [49].

The third equation in (37) is a scalar Laplace equation with a variable phase-field coefficient. It is complemented by zero-flux conditions at the Neumann boundary $\partial\tilde{\Omega}_N$ and non-zero Dirichlet conditions at the Dirichlet boundary $\partial\tilde{\Omega}_D$. The gradient ∇h of the Laplace solution determines the local intensity of the history function \mathcal{H} according to (29) and (34). With appropriate Dirichlet boundary conditions, the gradient field ∇h will naturally spike at the small contact bridges between different bone objects, since they are geometrically much finer than the bone objects themselves. This will initiate fracture at exactly those locations, removing the spurious contact bridges from the segmentation region.

Major advantages of the Laplace based formulation (37) with respect to the elasticity based phase-field fracture formulation shown in Section 2.2 is its smaller set of parameters as well as its reduced computational cost, since only two scalar functions need to be discretized instead of a scalar and a vector function.

2.3.3. A staggered solution algorithm

For the solution of the coupled system (35) and (36), we adjust the staggered solution algorithm based on operator splits [51, 54]. The idea is to consider the phase-field and Laplace parts independently, which leads to two well-defined sub-problems, and incrementally increase the Dirichlet boundary conditions from 0 to their full value g over pseudo-time t . In our case, we solve the phase-field part first to update the phase-field solution c and then solve the Laplace part that updates the Laplace solution h . Finally, we update the history field \mathcal{H} that records the maximum reference energy density at each quadrature point. The algorithm is summarized in the following box:

The field variables h_h^n and c_h^n , and the history field \mathcal{H}^n are known at the current time t^n . Update incremental prescribed Dirichlet boundary values g^{n+1} to the new time t^{n+1} .

1. *Update phase-field variable c_h (phase-field part):*

Solve linear system $K_{ij}^c c_j = F_i^c$ that evolves from (35) at frozen h_h .

2. *Update field variable h_h (Laplace part):*

Solve linear system $K_{ij}^h c_j = F_i^h$ using (36) at frozen phase-field c_h .

3. *Update history field:*

Determine maximum reference energy density during loading history

$$\mathcal{H}^{n+1}(h_h^{n+1}) = \begin{cases} \psi_0(h_h^{n+1}) & \text{if } \psi_0(h_h^{n+1}) > \mathcal{H}^n \\ \mathcal{H}^n & \text{otherwise} \end{cases}$$

Update time variable t^{n+1} to t^n and proceed to next step by restarting this procedure.

2.3.4. Choosing appropriate Dirichlet constraints

The success of our phase-field fracture inspired method critically depends on the appropriate choice of Dirichlet boundary constraints on the Laplace solution h . We assume that we know how many different bone objects we would like to segment from a given CT scan

and that we can identify a non-zero number of voxel facets on the boundary of each individual bone object. Based on this information, we can impose the following set of Dirichlet boundary conditions:

$$\begin{aligned}
h = 1 & \text{ on } \partial\tilde{\Omega}_{D,1} \\
h = 2 & \text{ on } \partial\tilde{\Omega}_{D,2} \\
& \vdots \quad \vdots \\
h = n & \text{ on } \partial\tilde{\Omega}_{D,n}
\end{aligned} \tag{38}$$

where $\partial\tilde{\Omega}_{D,i}$ denotes the boundary surface that consists of the known voxel boundary facets of the i^{th} bone object.

During the quasi-static loading procedure, we incrementally increase the boundary values at each Dirichlet boundary. Due to the zero-flux conditions at the remainder of the boundary, the solution h in the region of the i^{th} bone object will tend towards the constant value i . Before fractures occur, the instantaneous Laplace solution h in each bone region is only disturbed where small bridges between bone regions exist. This initiates sharp gradients in the solution h that drive the fracturing of fine-scale contacts. Once all bridges have been fully broken, we multiply the Laplace solution h with the phase-field solution c such that h is set to zero in the fully fractured contact regions. The resulting plateau function $c(x) \cdot h(x)$ consists of clear plateau regions with constant value i that indicate the i^{th} bone object. A simple overlay of the plateau function onto the original image I yields the final segmentation region for each individual bone object in the CT scan.

2.4. The complete methodology: overview and implementation

Our approach for segmenting CT bone objects in the presence of thin cartilage regions consists of two core components: (a) a model for bone/tissue separation based on a flux-augmented Chan-Vese functional, developed in Section 2.1; and (b) a model for automatic detection and removal of spurious fine-scale bridges based on a phase-field fracture inspired method, developed in Section 2.3. To improve robustness and automation of our approach, we apply the following simple voxel-wise manipulations in addition to the two core components. Before running the flux-augmented Chan-Vese model, we rescale the original image intensities to the zero-one range and set all tissue related information by a simple threshold to zero. Thresholding works well for this task, as the Hounsfield units (HU) of a CT scan yield a pronounced contrast between soft tissue such as blood muscles and hard tissue bone. After solving the flux-augmented Chan-Vese functional, we use voxel connectivity information to identify the biggest connected network of voxels covered by the nonzero solution field and remove any small regions unconnected to the main object. After obtaining the plateau function $c(x) \cdot h(x)$, we remove all voids within the bone interior by applying a dilation filter, a flood filling algorithm and an erosion filter. This ensures that the segmentation masks fully cover each bone object. We emphasize that these voxel-wise manipulations represent very common morphological operations in image processing, for which reliable and fully automatic algorithms are available (see, e.g., [55]). We summarize the corresponding algorithmic tasks in the following box:

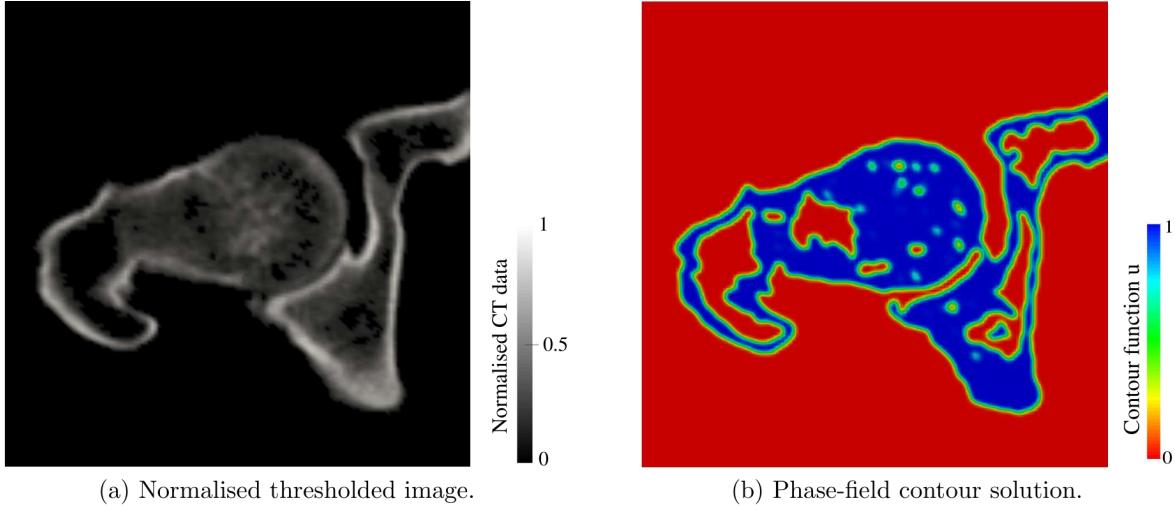
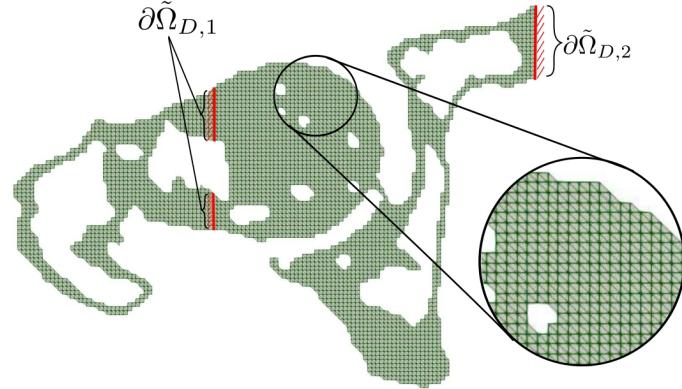


Figure 5: First component: bone/tissue separation based on a flux-augmented Chan-Vese model.

Input: Original CT imaging data, consisting of voxelized HU data.

1. Rescale the image intensity of the original CT data to the range $[0, 1]$ and perform global thresholding to set soft tissue information in the image domain to zero.
2. Find an implicit contour function u as the minimizer of the flux-augmented Chan-Vese functional (10), using voxel finite elements in space and first order semi-implicit finite differences in time. Use voxel connectivity to identify the biggest connected network of voxels covered by the nonzero solution and remove any small regions unconnected to the main object. [also refer to (12), (13)]
3. Remove all voxel finite elements in the region where $u < 0.5$ [also refer to (28)]. Identify suitable voxel facets to impose Dirichlet boundary values for the Laplace solution h at the boundary of each bone object. [also refer to (38)]
4. Find the coupled solution $\{c, h\}$ as the minimizer of the potential energy functional (30) under the Dirichlet constraints for h , using voxel finite elements and the staggered algorithm. [also refer to (35), (36), (38)]
5. Use the plateau function $c(x) \cdot h(x)$ on the voxel mesh to generate segmentation masks for each bone object. Apply a standard dilation filter, a standard filling algorithm and a standard erosion filter to fill potential voids in the bone interior.
6. Overlay the original CT data and the mask voxel-wise. Multiply the CT intensity with the one-zero mask of the desired bone object to obtain its segmentation, including the original HU data.

Output: Segmentation region, consisting of the HU voxel data of a bone object.



(a) Domain $\tilde{\Omega}$ covered by the voxel finite element mesh and Dirichlet constraints on the femur and pelvis parts.

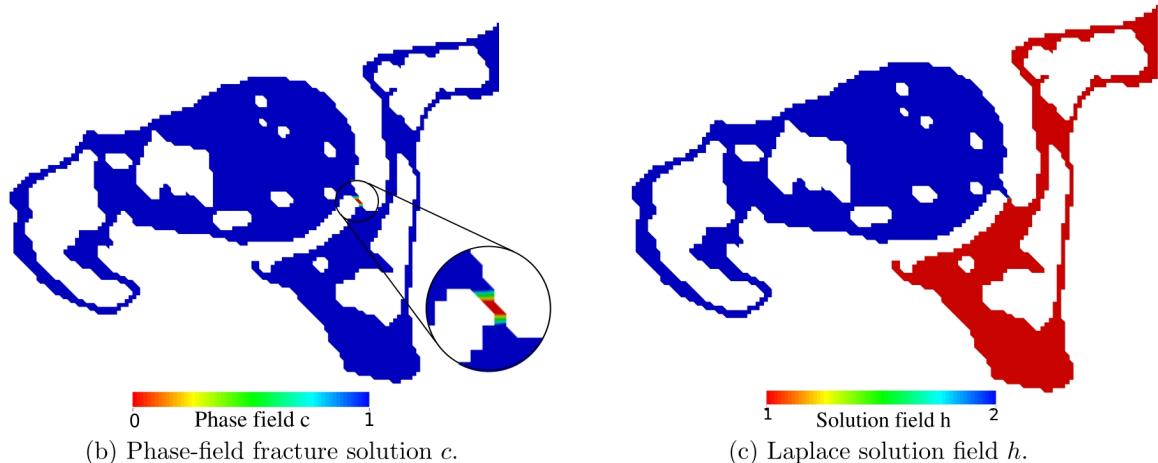
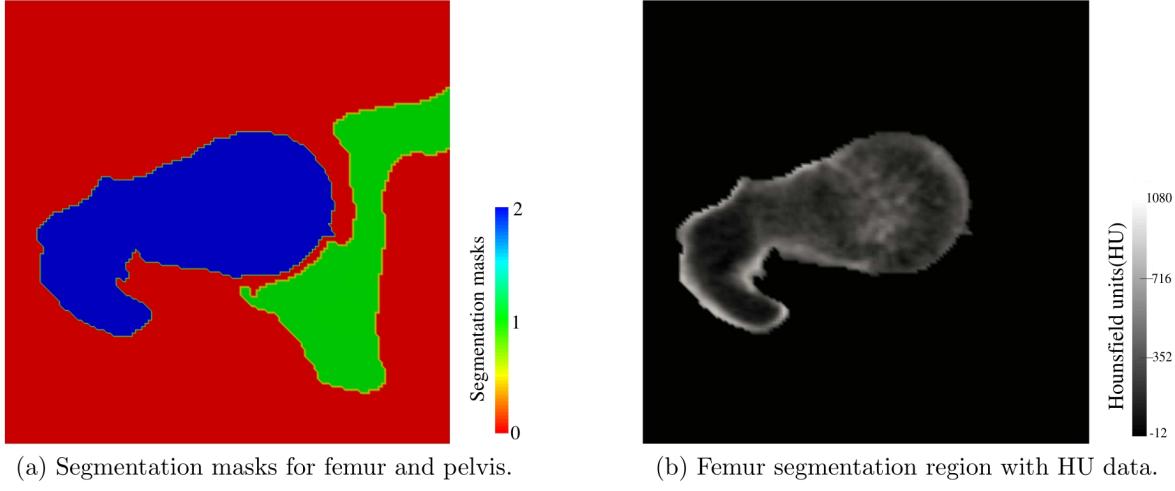


Figure 6: Second component: removal of spurious bridges in the cartilage region by the phase-field fracture inspired Laplace model.

The numerical treatment of the flux-augmented Chan-Vese and phase-field fracture inspired models largely relies on voxel finite element discretizations. Our implementation is based on the open source platform FEniCS¹. As FEniCS does not feature quadrilateral and hexahedral elements yet, we discretize each pixel in 2D by two triangular elements and each voxel in 3D by five tetrahedral elements.

We illustrate the effect of each step of our methodology with the representative 2D problem shown in Fig. 1b that features a thin sliver of cartilage between the femur and pelvis. Figure 5 illustrates the first component of our methodology, that is, the result of

¹FEniCS, an open-source finite element based software suite for solving partial differential equations; <https://fenicsproject.org/>



(a) Segmentation masks for femur and pelvis.

(b) Femur segmentation region with HU data.

Figure 7: Segmentation result.

thresholding and the contour result of the flux-augmented Chan-Vese model (steps 1 and 2). Figure 6 shows the voxel finite element discretization of the domain modified $\tilde{\Omega}$, including Dirichlet constraints, and the final result of the coupled phase-field and Laplace solution fields (steps 3 and 4). Figure 7 illustrates the segmentation results in terms of the segmentation masks and the femur segmentation region with HU data (steps 5 and 6). We observe that the phase-field fracture inspired model successfully detects and removes the spurious contact in the cartilage region, leading to two well-defined and accurate segmentation masks for the femur and the pelvis.

3. Results and discussion

In this section, we assess the performance of the variational segmentation methodology presented above, with a particular focus on the phase-field fracture inspired part. In the first subsection, we consider clinical 3D CT data of a series of femur that was obtained at the Academic Health Center of the University of Minnesota. We use different global and local overlap metrics to validate femur segmentations obtained with our two-stage variational approach against corresponding manual segmentations approved by an expert radiologist. In the second subsection, we demonstrate the generality of our two-stage segmentation method by segmenting 3D CT data of part of a vertebral column. In the third subsection, we discuss strengths and weaknesses of our two-stage variational segmentation method, with particular emphasis on accuracy, robustness and potential for full automation.

3.1. Femur segmentation: separation from pelvis

Segmenting the femur (thigh bone) from 3D CT data is an important task in clinical practice and biomedical research, for example to visually assess bone quality and multi-scale fractures or as a basis for running physiology-based finite element simulations. The main

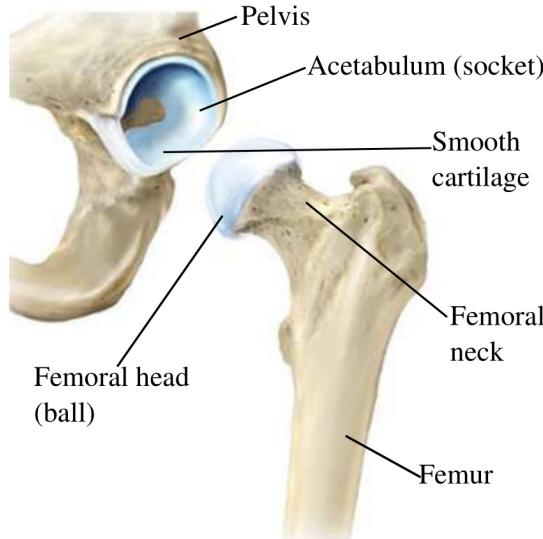


Figure 8: Anatomy of femur and pelvis joint.

challenge is the reliable and accurate separation of the femur from neighboring bones, in particular from the pelvic bone. The femoral head meets the pelvis at the acetabulum, forming the hip joint. Their surfaces are separated by a thin layer of slippery tissue called articular cartilage. Figure 8 illustrates the anatomical femur/pelvis configuration and the location of the articular cartilage.

3.1.1. Clinical CT data set from the University of Minnesota AHC

We obtain sets of clinical 3D CT imaging data of eleven patients from the Academic Health Center (AHC) of the University of Minnesota that were taken during regular clinical practice in the musculoskeletal and neuroradiology division. This results in twenty-two different femur bones (left and right bone of each patient) as input for our segmentation methodology. The data sets consist of quantitative CT scans in the form of DICOM files that provide the HU for a specific layer and pixel spacing. Ten data sets were obtained with a Siemens SOMATOM Definition Edge (Germany): 120 or 140 kVp, 300 or 350 mAs, 0.75 mm slice thickness; one data set was obtained with a clinical Siemens SOMATOM Sensation 64 CT (Germany): 140kVp, 350mAs, 0.75 mm slice thickness. For each CT data set, we crop the region of interest (ROI) that contains the femur and pelvis bone from the whole-body scan. Figure 9a illustrates the ROI taken from a the whole-body CT scans for one of the femurs with ROI pixel dimensions as $140 \times 140 \times 80$. The pixel spacing is 0.703125×0.703125 mm within each horizontal slice and slice thickness as 2.5 mm. We observe that in clinical CT scans such as the one shown in Fig. 9a, the articular cartilage region is usually resolved only by a few voxels, which makes detection and bone separation difficult.

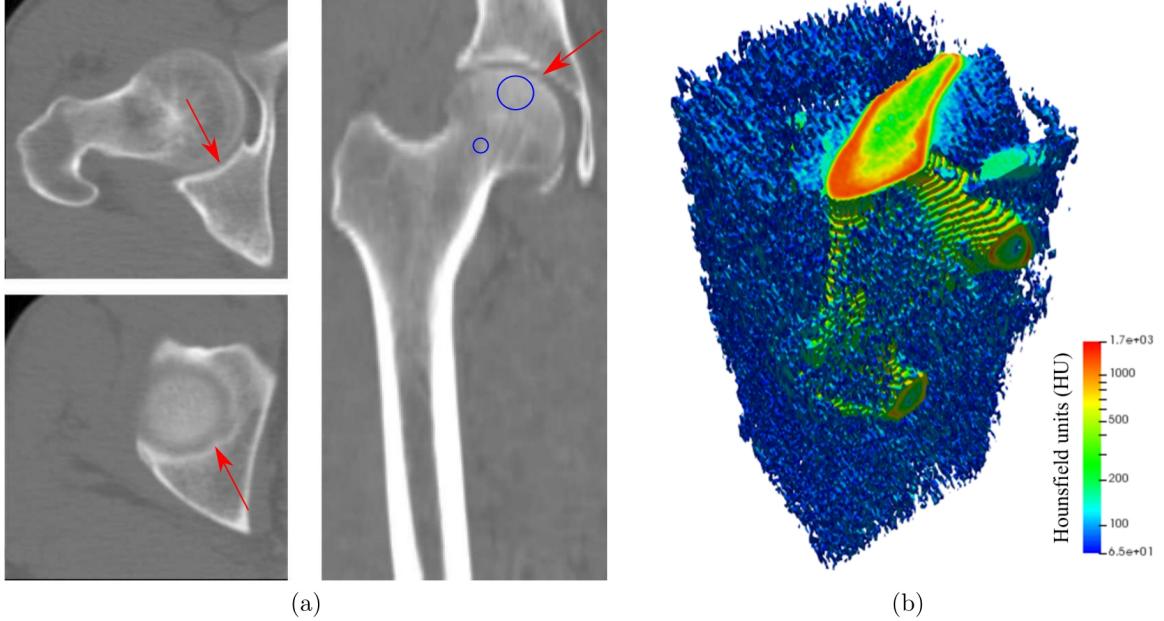


Figure 9: One of the femura of the clinical CT data set from the University of Minnesota AHC: (a) Cropped ROI (thin cartilage region between pelvis and femur is highlighted by red arrows); (b) 3D voxel representation (with light filter for better visualization).

3.1.2. Automated variational segmentation

In the next step, we run our two-stage variational segmentation approach for each data set, where we follow the methodology described in Section 2.4. In step 1, we normalize the ROI image to the zero-one range and subsequently apply a threshold of 0.05 to detect and level tissue regions to zero. In step 2, we find the contour function u by numerically minimizing the flux-augmented Chan-Vese functional. The choice of parameters plays a key role for the success of our methodology. To make the set of choices for this example transferrable to other examples, we non-dimensionalize all data sets such that each slice of the ROI covers a unit square. We generate a voxel finite element mesh in space, covering the complete resampled ROI, and solve the semi-implicit finite difference scheme with a time step of $\Delta t = 0.1$, until the difference between the current and previous solution falls below 10^{-10} in the 2-norm. The length-scale parameter ϵ of the Allen-Cahn part that controls the characteristic width of the diffuse boundary contour corresponds to the voxel finite element size, which is 0.005 in the resampled ROI. We empirically tested different values for λ and μ while keeping ϵ constant. Our conclusion is that the optimal ratio between λ and μ lies in the range of 0.5 and 2. For this example, we chose $\lambda = \mu = 20$. The resulting contour function is illustrated for one femur sample in Fig. 10. We observe that the flux-augmented Chan-Vese model provides an excellent representation of well-separated bone boundaries, but is not able to fully separate the femur in the cartilage region, where contact bridges between femur and pelvis remain.

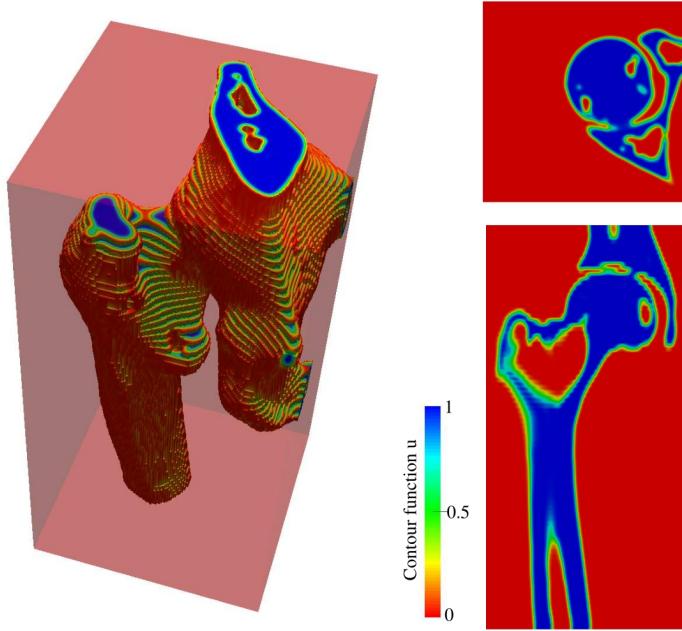
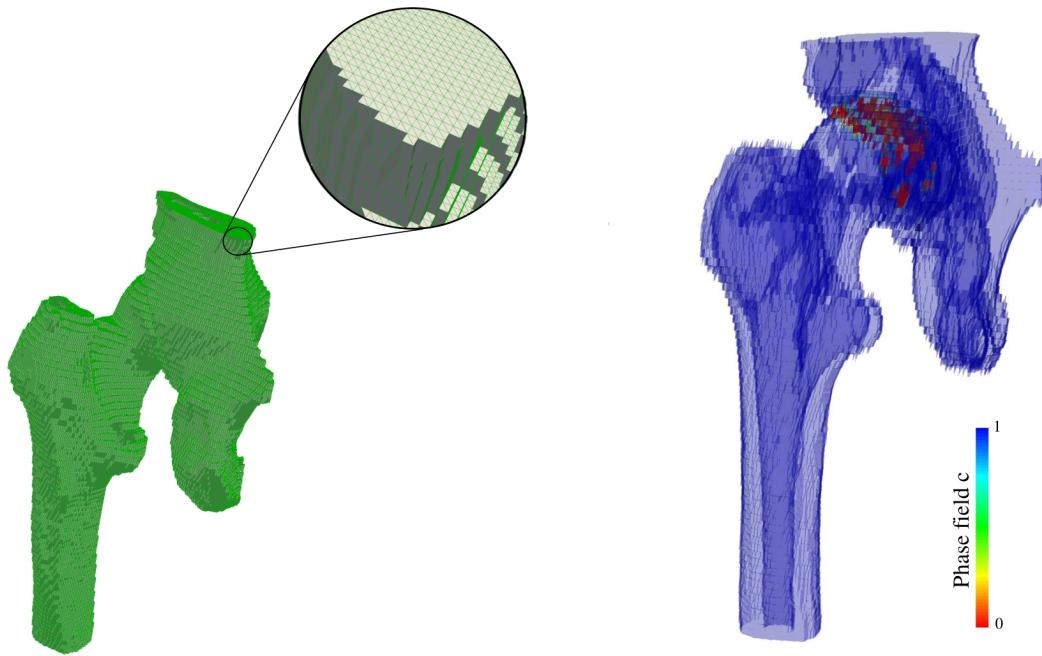


Figure 10: Femur example: contour function of the flux-augmented Chan-Vese model, featuring spurious contacts between femur and pelvis in the thin cartilage region.

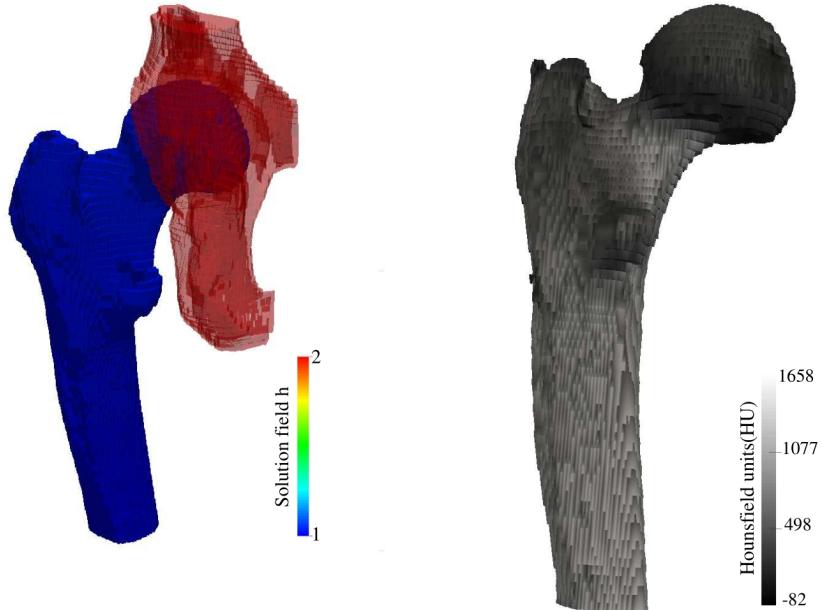
In step 3, we remove all voxel finite elements that are completely outside the contour $u < 0.5$. In this example, both individual bone objects touch the ROI boundary, where appropriate Dirichlet constraints for the Laplace solution h can be easily applied. For the femur, we set $h = 1$ on all voxel facets that touch the bottom plane of the ROI. For the pelvis, we set $h=2$ on all voxel facets that touch the top plane of the ROI. Figure 11a plots all remaining voxel finite elements, with a zoom on elements with facets on the top plane.

In step 4, we find the phase-field fracture function c coupled to the Laplace function h by numerically minimizing the simplified potential energy functional via the staggered algorithm on the voxel finite element mesh. The length-scale parameter for the phase-field c is the same as the length-scale parameter in the diffuse Chan-Vese model, such that $l_0 = \varepsilon = 0.005$. For numerical stability, we choose $\kappa = 10^{-4}$. Our numerical experiments with the staggered algorithm show that the solution is very robust with respect to the number of iterations. For the example data set shown in Figs. 9 to 11, we can find a suitable fracture topology with as few as three iterations to reliably eliminate all fine-scale contacts. For our examples, we apply a conservative number of ten iterations to compute the fracture topology, making the computational cost of the phase-field fracture inspired model comparable to the flux-augmented Chan-Vese model. We note that this is in marked contrast to the mechanical fracture formulation, where hundreds or thousands of iterations in the staggered algorithm are usually necessary to arrive at a solution accuracy of the fracture topology that is appropriate for physical interpretation. The resulting phase-field and Laplace solutions are plotted in Figs. 11b and 11c, respectively.



(a) Voxel finite element mesh.

(b) Phase-field solution c .



(c) Laplace solution h .

(d) Segmented femur with HU data.

Figure 11: Femur example: automated removal of contacts with the phase-field fracture inspired method.

In step 5, we first generate the plateau function $c(x) \cdot h(x)$ on the voxel mesh. For each bone object, we then apply a dilation filter, a flood filling algorithm and an erosion filter to fill all interior holes. The result is a segmentation mask for each bone object. In step 6, we first create a segmentation mask for the femur by writing out all voxels where $h = 1$. Overlaying and multiplying the mask with the original CT data yields the final segmentation of the HU data in the femur region. Figure 11d shows the final femur segmentation region with HU values.

3.1.3. Validation by global overlap metrics

To gain a better understanding of its accuracy, we validate the results obtained with our two-stage variational approach against corresponding femur segmentations that we know are of highest quality. Our gold standard are segmentation masks that were obtained manually from the same CT data sets and assessed for highest quality by Dr. Takashi Takahashi, an expert radiologist. For our validation, we overlay the pairs of segmented femur regions obtained manually and by the two-stage variational approach and classify each pair of voxels into four standard cardinalities. The correctly segmented bone voxels are true positive (TP). Non-bone voxels that are recognized as bone voxels by the two-stage variational approach are false positive (FP). Bone voxels that are missed by the two-stage variational approach are termed false negative (FN). Finally, non-bone voxels that are correctly segmented are true negative (TN). To assess the quality of the results of our variational approach, we consider three different overlap based statistical metrics: the Dice similarity coefficient (DSC), the sensitivity or true positive rate (TPR), and the specificity or true negative rate (TNR) [56]. The metrics are tabulated for all the datasets in Table 1.

The DSC, also called overlap index, measures the spatial overlap between two segmentations, in our case the manual segmentation A and the variational segmentation B, and is defined as:

$$\text{DSC}(A, B) = \frac{2(A \cap B)}{A + B} \quad (39)$$

where \cap denotes the intersection. The best possible case is $\text{DSC} = 1$, which corresponds to a perfect overlap. In terms of the cardinalities defined above, the DSC can be defined as:

$$\text{DSC} = \frac{2TP}{2TP + FP + FN} \quad (40)$$

In the context of segmentation, a DSC value greater than 0.8 is widely considered to be a good overlap between the results of a segmentation method and the ground truth [57]. We observe in Table 1 that all DSC values are significantly larger than 0.8, with average and standard deviation at (0.9339 ± 0.0287) , which indicates an excellent agreement of the results obtained with our two-stage methodology and the optimal segmentation result.

The sensitivity or TPR measures the proportion of hits (TP) to the total number of bone voxels in the ground truth, that is, the manual segmentation: $\text{TPR} = TP / (TP + FN)$. Similarly, the specificity or TNR measures the portion of negative voxels (background) in the ground truth segmentation that are also identified as negative by the automatic

Table 1: Validation metrics for all data sets (L stands for left, R for right femur)

Data set	Dice coefficient	Sensitivity	Specificity
1L	0.9445	0.9988	0.9843
1R	0.945	0.9996	0.9846
2L	0.9649	0.9979	0.993
2R	0.9631	0.9778	0.997
3L	0.9738	0.9872	0.9946
3R	0.9704	0.9712	0.9964
4L	0.8893	0.9655	0.9851
4R	0.9086	0.9976	0.9862
5L	0.9142	0.9912	0.98
5R	0.9155	0.9892	0.9831
6L	0.903	0.9954	0.9731
6R	0.8815	0.9906	0.9863
7L	0.9129	0.9928	0.9809
7R	0.916	0.9769	0.983
8L	0.9037	0.9732	0.9794
8R	0.9119	0.9871	0.9839
9L	0.9695	0.9713	0.9975
9R	0.9715	0.9771	0.9973
10L	0.9584	0.9608	0.9954
10R	0.959	0.9886	0.9928
11L	0.9402	0.9929	0.9921
11R	0.9297	0.999	0.9897

segmentation, that is, $TNR = TN/(TN + FP)$. We highlight that stand-alone values of TPR and TNR do not provide meaningful indications of the quality of the segmentation results with respect to the ground truth. The sensitivity can be equal to 1 for a very poor segmentation result, when the segmented region completely covers the ground truth data, but also contains many false positive hits. Similarly, the specificity can be equal to 1 for a very poor segmentation, when not a single voxel of the object could be detected. Hence, both metrics need to be checked simultaneously for a meaningful assessment of segmentation quality. We observe in Table 1 that for each data set, both the sensitivity and the specificity are far above 0.9 simultaneously, with an average and standard deviation of (0.9339 ± 0.0287) and (0.9855 ± 0.0115) , respectively. We can conclude that all metrics consistently indicate that our two-stage segmentation methodology is able to provide results that are in excellent agreement with the optimal manually obtained segmentation region.

3.1.4. Validation by local HU averages at characteristic surface positions

The overlap metrics presented in Table 1 provide a measure of global similarity between the segmentation region obtained with the two-stage variational approach and manual segmentation. However, they do not provide any information on the local similarity in the

femur region that is affected by the phase-field fracture inspired method and therefore of particular interest in our context. Moreover, proper segmentation and representation of HU values in local cartilage regions such as the femoral head are significant from the clinical perspective for the diagnosis of a number of bone diseases.

Therefore, we validate the local quality of HU representation at the bone surface with the following strategy. In the first step, an experienced radiologist (in our case Dr. Takashi Takahashi) measures, with the help of OsiriX², the average HU value locally at two characteristic positions on the bone surface. He obtains a corresponding set of HU measurements from all segmented femur regions obtained manually and with the two-stage variational approach. For the femoral head, the average HU value was measured in an identical spherical ROI with a radius of 10 pixels centered at the fovea capitis, a small concave depression within the femoral head. Similarly, average HU values are obtained for a spherical ROI with a radius of 2 pixels centered at the femoral neck. Both locations in one slice are shown by circles in Fig. 9a.

To assess the local similarity of the average HU measurements, we compute the normalized root mean square deviation (NRMSD) between the measured HU values in the two ROIs for the series of segmented femur samples:

$$\text{NRMSD} = \sqrt{\sum_{i=1}^n \frac{(\bar{H}U_{\text{var},i} - \bar{H}U_{\text{man},i})^2}{n}} / \bar{H}U_{\text{man},i} \quad (41)$$

where we normalize with the mean of the reference manual segmentation. For the femoral neck, the resulting NRMSD is 2.64%. For the femoral head, the resulting NRMSD is as low as 0.99%. We observe that the deviation between the variational segmentation regions and the reference is very small and at an extent that is far away from having an impact on clinically relevant HU based observations. Our local error analysis thus confirms that the variational segmentation approach captures local features in terms of HU averages in the bone with excellent accuracy. In particular, this also holds for bone regions of the femoral head that are in contact with the articular cartilage region, where we have applied the phase-field fracture inspired separation.

3.2. Vertebra segmentation: separation from spine

Clinical diagnosis and therapy of spine related diseases require knowledge of stress and strains in particular vertebra regions. An important task in clinical practice and biomedical applications is therefore to extract a segmentation of an individual vertebra from 3D CT data of the spine. The main challenge is the reliable separation of the target vertebra bone from the neighboring upper and lower vertebrae. Neighboring vertebral bodies do not touch each other directly, but are separated by an intervertebral disk that form a thin intervertebral fibrocartilage layer. In addition, some of the spinous processes of neighboring vertebrae form

²OsiriX, a medical image processing application for navigation and visualization of multimodality and multidimensional images; www.osirix-viewer.com

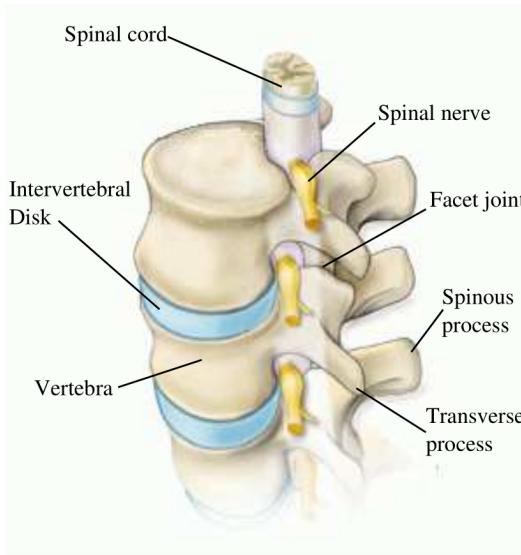


Figure 12: Anatomy of the vertebral column (spine).

synovial joints that consist of thin fibrous capsules of cartilage material. Figure 12 illustrates part of the vertebral column and the location of the intervertebral disks and fibrous capsules. In clinical CT scans such as the one shown in Fig. 7, these cartilage regions are usually resolved only by a few voxels, which makes detection and bone separation difficult.

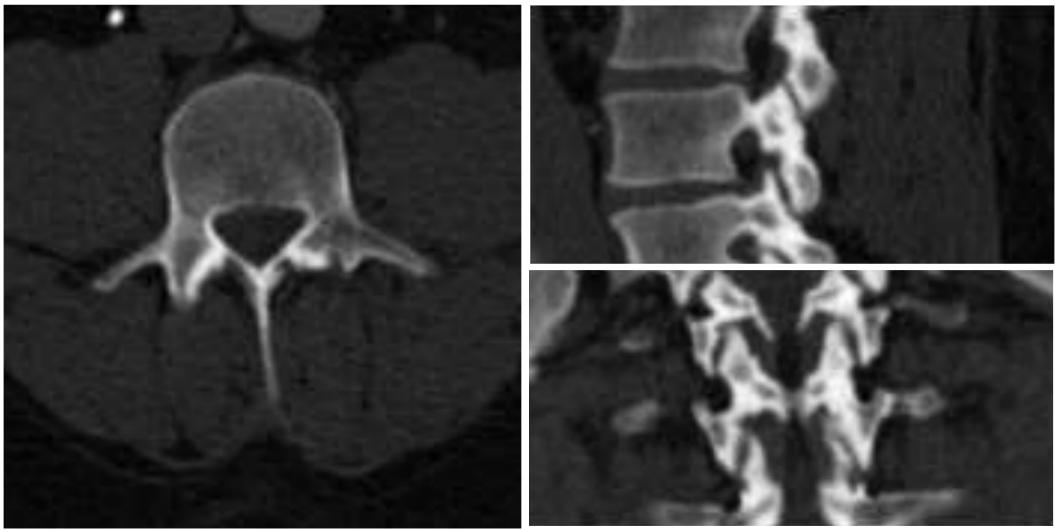


Figure 13: Vertebra example: cropped ROI, illustrating the complex structure of the target region.

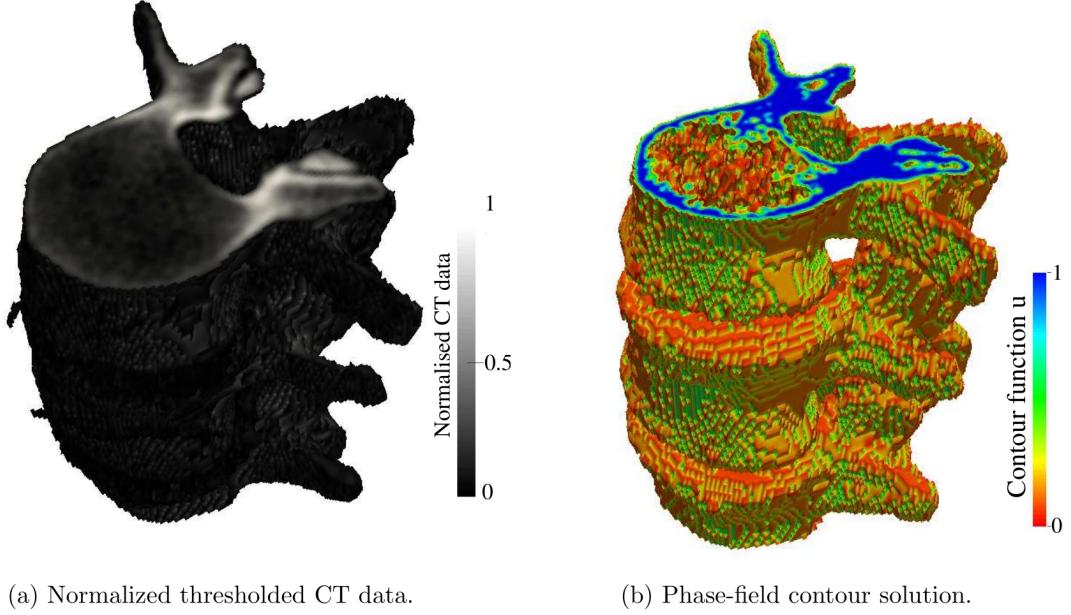


Figure 14: Vertebra example: contour function of the flux-augmented Chan-Vese model, featuring contacts between neighboring vertebrae.

3.2.1. Test data from the OsiriX DICOM image library

We download the test data set ‘‘OBELIX’’ from the OsiriX DICOM image library that is available online on the OsiriX webpage³. It consists of a whole body contrast CTA acquired on a 16 detector CT scanner in a normal study, providing quantitative CT scans in the form of a series of DICOM files. From the original whole body 3D data, we crop the region of interest shown in Fig. 13 that contains the target vertebra in the center and most of the neighboring upper and lower vertebrae. The ROI voxel dimensions are $180 \times 180 \times 91$. The pixel spacing is 0.7422×0.7422 mm within each horizontal slice and the distance between slices in axial direction is 1.0 mm.

3.2.2. Automated variational segmentation

We then apply our two-stage variational approach to segment the center vertebra from the ROI. We again strictly follow the methodology described in Section 2.4. We note that all parameters involved in our methodology are the same as in the femur study described above in Section 3.1.2. After rescaling the original HU values to $[0, 1]$ and applying a low-level threshold to level tissue information to zero (see Fig. 14a), we generate a voxel finite element mesh and find the minimizer of the flux-augmented Chan-Vese model, illustrated in Fig. 14b. For better visibility of the bone structures, the outside region where the contour field is zero is not plotted.

³<http://www.osirix-viewer.com/resources/dicom-image-library/> on January 18, 2017

We then use the phase-field fracture inspired method to eliminate all remaining connections of the center vertebra with its upper and lower neighbors. To this end, we first remove all voxel finite elements from the discretization that are completely outside the contour $u < 0.5$. The Dirichlet constraints on the Laplace solution h are applied as follows: At the middle plane of the ROI, we set $h = 2$, since we know that this plane lies completely within the target center vertebra. The lower and upper vertebrae touch the lower and upper boundaries of the voxel finite element discretization, where we apply $h = 1$ and $h = 3$, respectively. We then find the minimizer $\{c, h\}$ of the potential energy functional under the applied Dirichlet constraints via the staggered algorithm that is again run with ten iterations.

Figure 15a illustrates the phase-field solution c that represents the fracture topology in the critical cartilage regions. Figure 15b plots the final plateau solution $c(x) \cdot h(x)$, which is constant in each vertebra region according to our choice of Dirichlet constraints. Figures 15c and 15d show the segmentation mask and the final segmentation region with HU data for the targeted center vertebra.

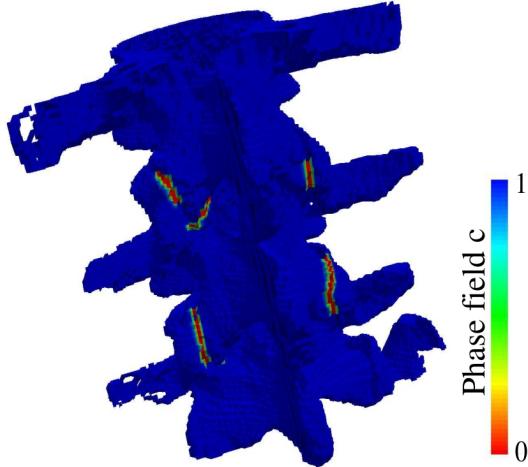
3.3. Discussion of strengths and opportunities

Even the most advanced segmentation approaches today typically compromise on at least one of the fundamental properties of accuracy, robustness and automation. An important goal of this research has therefore been to develop a variational segmentation methodology for complex 3D bone CT data that is accurate, robust and operates autonomously at the same time. A particular focus has been on improving the performance of existing variational methods in the presence of thin cartilage regions at practical image resolutions.

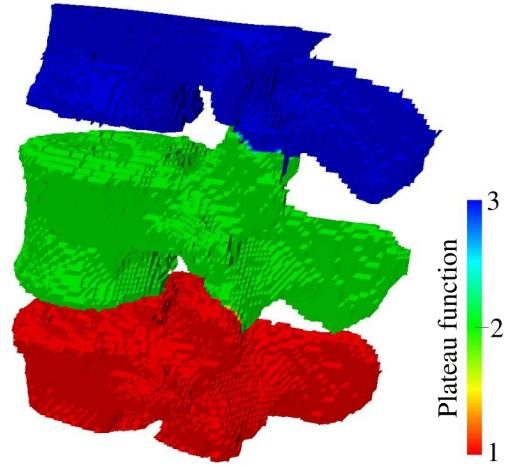
With respect to *accuracy*, we could show for a series of twenty-two clinically obtained femora that our two-stage segmentation methodology consistently provided excellent segmentation regions that were almost identical to segmentation regions obtained manually by an expert radiologist. The accuracy of our methodology largely relies on the flux-augmented Chan-Vese model employed in the first stage, which is responsible for identifying the bulk of the segmentation region. The Chan-Vese model has been fine-tuned over the last two decades and has been shown to yield excellent segmentation results in a variety of situations (see for example [29] and the references therein). We note that in our methodology, it could be easily exchanged by any other variational segmentation model, if desired.

With respect to *robustness*, the segmentation results presented in this paper showed that our two-stage methodology is algorithmically robust, in particular with respect to the choice of the computational parameters involved. Our approach is also computationally efficient and, for instance, could be operated on regular desktop computers that are readily available in hospitals or medical practices.

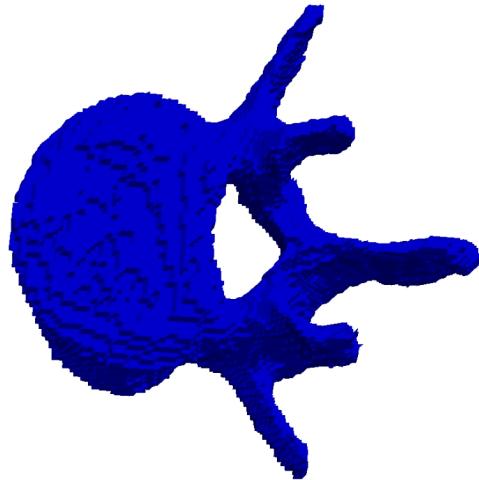
With respect to *automation*, our two-stage methodology constitutes a significant advance with respect to existing variational segmentation methods. On the one hand, we could show that if the flux-augmented Chan-Vese model is applied alone, the presence of thin cartilage interfaces at practical image resolutions lead to segmentation regions that feature spurious contacts between different individual bone objects. These bridges can be detected visually and removed manually by a computational analyst, but preclude autonomous operation. On



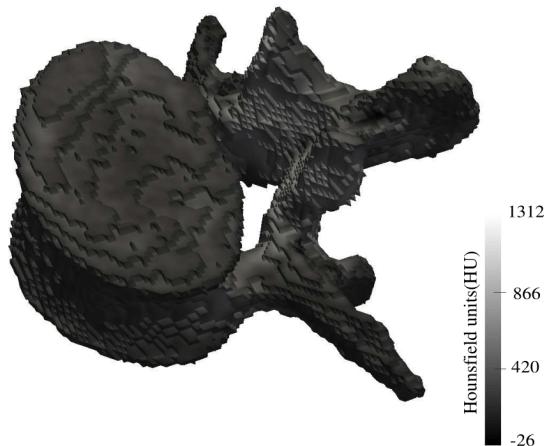
(a) Phase-field solution c .



(b) Plateau function in each vertebra.



(c) Mask for target vertebra.



(d) Segmented vertebra with HU data.

Figure 15: Vertebra example: automated removal of contacts with the phase-field fracture inspired method.

the other hand, we could show for several challenging 3D segmentation problems, including twenty-two femora and a vertebra, that the phase-field fracture inspired second stage of our methodology reliably removes all spurious contacts. We showed that our approach provides accurate segmentation topologies, from which segmentation regions of any bone object in the CT data can be easily extracted without further manual effort.

From a variational viewpoint, both stages of our methodology are based on the minimization of functionals which lead to variational problems that are closely related. In particular, the variational structure of both stages allows a discretization with the standard

finite element method. Both stages share the concept of a diffuse phase-field solution controlled by a length-scale parameter that represents complex geometric features implicitly. The solutions in both stages must be obtained separately, but can be computed on the same voxel finite element mesh. In the future, we plan to exploit and extend this property for seamlessly connecting variational segmentation procedures with downstream predictive bone simulation in a common finite element framework and on the same finite element mesh.

Acknowledgments. T. Gangwar is partially supported by a Sommerfeld Fellowship awarded by the Department of Civil, Environmental, and Geo-Engineering at the University of Minnesota, which is gratefully acknowledged. D. Schillinger gratefully acknowledges support from the National Science Foundation through the research grant CISE-1565997 and the NSF CAREER Award No. 1651577. We acknowledge the Academic Health Center (AHC) and the Institutional Review Board (IRB) of the University of Minnesota for providing clinical data and supporting its use in this study. The Minnesota Supercomputing Institute (MSI) of the University of Minnesota has provided computing resources that have contributed to the research results reported within this paper (<https://www.msi.umn.edu/>), which is also gratefully acknowledged.

4. References

- [1] L. Formaggia, A. Quarteroni, and A. Veneziani. Multiscale models of the vascular system. In *Cardiovascular Mathematics: Modeling and simulation of the circulatory system*, pages 395–446. Springer, 2009.
- [2] W.A. Wall, L. Wiechert, A. Comerford, and S. Rausch. Towards a comprehensive computational model for the respiratory system. *International Journal for Numerical Methods in Biomedical Engineering*, 26(7):807–827, 2010.
- [3] D. Ambrosi, G.A. Ateshian, E.M. Arruda, S.C. Cowin, J. Dumais, A. Goriely, G.A. Holzapfel, J.D. Humphrey, R. Kemkemer, E. Kuhl, et al. Perspectives on biological growth and remodeling. *Journal of the Mechanics and Physics of Solids*, 59(4):863–883, 2011.
- [4] I. Sazonov, S.Y. Yeo, R.L.T. Bevan, X. Xie, R. van Loon, and P. Nithiarasu. Modelling pipeline for subject-specific arterial blood flow – A review. *Int'l Journal for Numerical Methods in Biomedical Engineering*, 27(12):1868–1910, 2011.
- [5] D. Kamensky, M.-C. Hsu, D. Schillinger, J.A. Evans, A. Aggarwal, Y. Bazilevs, M.S. Sacks, and T.J.R. Hughes. An immersogeometric variational framework for fluid–structure interaction: application to bioprosthetic heart valves. *Computer Methods in Applied Mechanics and Engineering*, 284:1005–1053, 2015.
- [6] F. Auricchio, M. Conti, M. Ferraro, S. Morganti, A. Reali, and R.L. Taylor. Innovative and efficient stent flexibility simulations based on isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering*, 295:347–361, 2015.
- [7] A.L. Marsden and M. Esmaily-Moghadam. Multiscale modeling of cardiovascular flows for clinical decision support. *Applied Mechanics Reviews*, 67(3):030804, 2015.
- [8] R. Blanchard, C. Morin, A. Malandrino, A. Vella, Z. Sant, and C. Hellmich. Patient-specific fracture risk assessment of vertebrae: A multiscale approach coupling X-ray physics and continuum micromechanics. *International Journal for Numerical Methods in Biomedical Engineering*, doi:10.1002/cnm.2760, 2016.
- [9] M. Sonka, V. Hlavac, and R. Boyle. *Image processing, analysis, and machine vision*. Cengage Learning, 2014.

[10] T. Heimann and H.-P. Meinzer. Statistical shape models for 3D medical image segmentation: a review. *Medical Image Analysis*, 13(4):543–563, 2009.

[11] J.E. Iglesias and M.R. Sabuncu. Multi-atlas segmentation of biomedical images: a survey. *Medical Image Analysis*, 24(1):205–219, 2015.

[12] J.-M. Morel and S. Solimini. *Variational methods in image segmentation: with seven image processing experiments*, volume 12. Springer Science & Business Media, 2012.

[13] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[14] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2):91–108, 1996.

[15] L.D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Transactions on Pattern Analysis and machine intelligence*, 15(11):1131–1147, 1993.

[16] R. Malladi, J.A. Sethian, and B.C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE transactions on pattern analysis and machine intelligence*, 17(2):158–175, 1995.

[17] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997.

[18] S. Osher and J.A. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.

[19] J.A. Sethian. Fast marching methods. *SIAM Review*, 41(2):199–235, 1999.

[20] C.-F. Westin, S. Warfield, A. Bhalerao, L. Mui, J. Richolt, and R. Kikinis. Tensor controlled local structure enhancement of CT images for bone segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1205–1212. Springer, 1998.

[21] M. Krčah, G. Székely, and R. Blanc. Fully automatic and fast segmentation of the femur bone from 3D-CT images with no shape prior. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 2087–2090. IEEE, 2011.

[22] J. Zhang, C.-H. Yan, C.-K. Chui, and S.-H. Ong. Fast segmentation of bone in CT images using 3D adaptive thresholding. *Computers in Biology and Medicine*, 40(2):231–236, 2010.

[23] Y. Sato, K. Nakanishi, H. Tanaka, N. Sugano, T. Nishii, H. Nakamura, T. Ochi, and S. Tamura. A fully automated method for segmentation and thickness determination of hip joint cartilage from 3D MR data. In *International Congress Series*, volume 1230, pages 352–358. Elsevier, 2001.

[24] R.A. Zoroofi, Y. Sato, T. Sasama, T. Nishii, N. Sugano, K. Yonenobu, H. Yoshikawa, T. Ochi, and S. Tamura. Automated segmentation of acetabulum and femoral head from 3-D CT images. *IEEE Transactions on Information Technology in Biomedicine*, 7(4):329–343, 2003.

[25] H. Seim, D. Kainmueller, M. Heller, H. Lamecker, S. Zachow, and H.-C. Hege. Automatic segmentation of the pelvic bones from CT data based on a statistical shape model. *VCBM*, 2008:93–100, 2008.

[26] F. Yokota, T. Okada, M. Takao, N. Sugano, Y. Tada, and Y. Sato. Automated segmentation of the femur and pelvis from 3D CT data of diseased hip using hierarchical statistical shape model of joint structure. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 811–818. Springer, 2009.

[27] H. Lamecker, M. Seebaß, H.-C. Hege, and P. Deuflhard. A 3D statistical shape model of the pelvic bone for segmentation. In *Medical Imaging 2004*, pages 1341–1351. International Society for Optics and Photonics, 2004.

[28] T.F. Chan and L.A. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.

[29] J. Calder, A.M. Tahmasebi, and A.-R. Mansouri. A variational approach to bone segmentation in CT images. In *SPIE Medical Imaging*, pages 79620B–79620B. International Society for Optics and Photonics, 2011.

[30] G.A. Francfort and J.-J. Marigo. Revisiting brittle fracture as an energy minimization problem. *Journal of the Mechanics and Physics of Solids*, 46(8):1319–1342, 1998.

[31] B. Bourdin, G.A. Francfort, and J.-J. Marigo. The variational approach to fracture. *Journal of Elasticity*, 91(1-3):5–148, 2008.

[32] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42(5):577–685, 1989.

[33] L. Bar, T.F. Chan, G. Chung, M. Jung, N. Kiryati, R. Mohieddine, N. Sochen, and L.A. Vese. Mumford and Shah model and its applications to image segmentation and image restoration. In *Handbook of Mathematical Methods in Imaging*, pages 1095–1157. Springer, 2011.

[34] J.A. Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, volume 3. Cambridge University Press, 1999.

[35] A. Karma and W.-J. Rappel. Quantitative phase-field modeling of dendritic growth in two and three dimensions. *Physical Review E*, 57(4):4323, 1998.

[36] S.K.F. Stoter, P. Müller, L. Cicalese, M. Tuveri, D. Schillinger, and T.J.R. Hughes. A diffuse interface method for the Navier–Stokes/Darcy equations: Perfusion profile for a patient-specific human liver based on mri scans. *Computer Methods in Applied Mechanics and Engineering*, 321:70–102, 2017.

[37] L.H. Nguyen, S.K.F. Stoter, M. Ruess, M.A. Sanchez Uribe, and D. Schillinger. The diffuse Nitsche method: Dirichlet constraints on phase-field boundaries. *International Journal for Numerical Methods in Engineering*, page DOI:10.1002/nme.5628, 2017.

[38] L.C. Evans, H.M. Soner, and P.E. Souganidis. Phase transitions and generalized motion by mean curvature. *Communications on Pure and Applied Mathematics*, 45(9):1097–1123, 1992.

[39] B. Merriman, J.K. Bence, and S.J. Osher. Motion of multiple junctions: A level set approach. *Journal of Computational Physics*, 112(2):334–363, 1994.

[40] S. Esedoglu and Y.H.R. Tsai. Threshold dynamics for the piecewise constant Mumford–Shah functional. *Journal of Computational Physics*, 211(1):367–384, 2006.

[41] O.C. Zienkiewicz and R.L. Taylor. *The Finite Element Method – The Basis*, volume 1. Butterworth-Heinemann, 6th edition, 2005.

[42] J.H. Keyak, S.A. Rossi, K.A. Jones, and H.B. Skinner. Prediction of femoral fracture load using automated finite element modeling. *Journal of Biomechanics*, 31(2):125–133, 1997.

[43] D. Schillinger and M. Ruess. The Finite Cell Method: A review in the context of higher-order structural analysis of CAD and image-based geometric models. *Archives of Computational Methods in Engineering*, 22(3):391–455, 2015.

[44] L.H. Nguyen, S.K.F. Stoter, T. Baum, J.S. Kirschke, M. Ruess, Z. Yosibash, and D. Schillinger. Phase-field boundary conditions for the voxel finite cell method: surface-free stress analysis of CT-based bone structures. *International Journal for Numerical Methods in Biomedical Engineering*, doi:10.1002/cnm.2880, 2017.

[45] J. Shen and X. Yang. Numerical approximations of Allen-Cahn and Cahn-Hilliard equations. *Discrete and Continuous Dynamical Systems*, 28(4):1669–1691, 2010.

[46] D. Schillinger, M.J. Borden, and H.K. Stolarski. Isogeometric collocation for phase-field fracture models. *Computer Methods in Applied Mechanics and Engineering*, 284:583–610, 2015.

[47] A.A. Griffith. The phenomena of rupture and flow in solids. *Royal Society of London Philosophical Transactions Series A*, 221:163–198, 1921.

[48] G.R. Irwin. Fracture. In S. Flügge, editor, *Elasticity and Plasticity / Elastizität und Plastizität*, volume 3 / 6 of *Encyclopedia of Physics*, pages 551–590. 1958.

[49] M.J. Borden, C.V. Verhoosel, M.A. Scott, Hughes T.J.R., and C.M. Landis. A phase-field description of dynamic brittle fracture. *Computer Methods in Applied Mechanics and Engineering*, 217–220:77–95, 2012.

[50] L. Ambrosio and V.M. Tortorelli. Approximation of functional depending on jumps by elliptic functional via γ -convergence. *Communications on Pure and Applied Mathematics*, 43(8):999–1036, 1990.

[51] C. Miehe, M. Hofacker, and F. Welschinger. A phase field model for rate-independent crack propagation: Robust algorithmic implementation based on operator splits. *Computer Methods in Applied Mechanics and Engineering*, 199(45):2765–2778, 2010.

[52] A. Stavrev, L.H. Nguyen, R. Shen, V. Varduhn, M. Behr, S. Elgeti, and D. Schillinger. Geometrically accurate, efficient, and flexible quadrature techniques for the tetrahedral finite cell method. *Computer*

Methods in Applied Mechanics and Engineering, 310:646–673, 2016.

- [53] V. Varduhn, M.-C. Hsu, M. Ruess, and D. Schillinger. The tetrahedral finite cell method: Higher-order immersogeometric analysis on adaptive non-boundary-fitted meshes. *International Journal for Numerical Methods in Engineering*, 107:1054–1079, 2016.
- [54] B. Bourdin. Image segmentation with a finite element method. *ESAIM: Modélisation Mathématique et Analyse Numérique*, 33(2):229–244, 1999.
- [55] W. Birkfellner. *Applied Medical Image Processing: A Basic Course*. Taylor & Francis, 2014.
- [56] A.A. Taha and A. Hanbury. Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging*, 15(1):29, 2015.
- [57] K.H. Zou, S.K. Warfield, A. Bharatha, C.M.C. Tempany, M.R. Kaus, S.J. Haker, W.M. Wells, F.A. Jolesz, and R. Kikinis. Statistical validation of image segmentation quality based on a spatial overlap index. *Academic Radiology*, 11(2):178–189, 2004.