Degree Project in Mathematics

Second cycle, 30 credits

Cut Finite Element Approximation of the Mean Curvature Vector

GUSTA[illegible]URSTAM

Cut Finite Element Approximation of the Mean Curvature Vector

GUSTAF BJURSTAM

Master's Programme, Applied and Computational Mathematics, 120 credits
Date: May 23, 2026

Supervisor: Sara Zahedi
Examiner: Sara Zahedi
School of Engineering Sciences
Swedish title: Approximation av medelkrökningsvektorn med skurna finita elementmetoder

Abstract

Accurate computation of the mean curvature vector is important for simulating certain interfacial phenomena. The Cut Finite Element Method, or CutFEM, provides a framework for such problems by embedding the surface in a stationary, unfitted background mesh, thereby avoiding remeshing for non-stationary problems. This thesis extends the CutFEM framework for mean curvature computation, originally established for piecewise linear elements, to multilinear spaces and investigates the theoretical limitations of higher-order polynomial approximations.

The first part of the thesis abstracts the existing piecewise linear analysis into a stabilization framework in arbitrary dimensions. The assumptions are formulated in terms of weak consistency, positivity, tangential gradient control, and trace control, leading to a first-order *a priori* error estimate.

The main theoretical contribution is the extension to multilinear elements on curved discrete surfaces. In the piecewise linear setting, the discrete surface Laplacian term vanishes locally after integration by parts. For multilinear elements this is no longer true, and an additional surface-Laplacian control estimate is required. This thesis shows that the proposed normal-based stabilizations provide this control and hence recover first-order convergence for the multilinear formulation.

For higher-order spaces, the analysis identifies a remaining obstruction. Stability can be proved, and an error estimate one order below the interpolation order is obtained for polynomial degree k. However, the present argument does not give an estimate of the same order as the interpolation error. The difficulty is that the discrete surface Laplacian contains tangential second-derivative terms that are not controlled by the considered normal-based stabilizations. Numerical experiments support the multilinear theory and show second-order convergence for quadratic elements in the tested two and three-dimensional geometries.

Sammanfattning

Noggrann beräkning av medelkrökningsvektorn är viktig vid simulering av vissa ytfenomen. Skurna finita elementmetoder (Cut Finite Element Method, CutFEM) ger ett ramverk för sådana problem genom att bädda in ytan i ett stationärt, icke-anpassat bakgrundsnät. På så vis undviks behovet av att generera nya nät för icke-stationära problem. Detta examensarbete utvidgar CutFEM-ramverket för beräkning av medelkrökningsvektorn, som ursprungligen utvecklades för styckvis linjära element, till multilinjära rum, och undersöker de teoretiska begränsningarna för polynomapproximationer av högre ordning.

Den första delen av arbetet abstraherar den befintliga analysen för styckvis linjära element till ett stabiliseringsramverk i godtycklig dimension. Antagandena formuleras i termer av svag konsistens, positivitet, kontroll av tangentiella gradienter och spårkontroll, vilket leder till en första ordningens *a priori*-feluppskattning.

Det huvudsakliga teoretiska bidraget är utvidgningen till multilinjära element på krökta diskreta ytor. I det styckvis linjära fallet försvinner den diskreta Laplace–Beltrami-termen lokalt efter partiell integration. För multilinjära element gäller detta inte längre, och ytterligare kontroll av Laplace–Beltrami-operatorn krävs. Arbetet visar att de föreslagna normalbaserade stabiliseringarna ger denna kontroll och därmed återfår första ordningens konvergens för den multilinjära formuleringen.

För rum av högre ordning identifierar analysen ett kvarstående hinder. Stabilitet kan visas, och en feluppskattning en ordning under interpolationsordningen erhålls för polynomgrad k. Däremot ger det nuvarande beviset inte en uppskattning av samma ordning som interpolationsfelet. Svårigheten är att den diskreta Laplace–Beltrami-operatorn innehåller termer med tangentiella andraderivator som inte kontrolleras av de normalbaserade stabiliseringar som behandlas här. Numeriska experiment stöder teorin för multilinjära element och visar andra ordningens konvergens för kvadratiska element för de testade två- och tredimensionella geometrierna.

Acknowledgements

I would like to express my sincere gratitude to Sara Zahedi for supervising this project, for suggesting the topic, and for introducing me to CutFEM. Her guidance and support throughout the work have been very valuable. I am also grateful to Sebastian Myrbäck for his help with the group's codebase and for patiently answering my questions during the implementation work. Finally, I would like to thank Mats G. Larson. His lectures at the Interfaces and Unfitted Discretization Methods programme at Institut Mittag-Leffler gave me a valuable introduction to CutFEM, and his input during the project helped shape the multilinear extension.

List of Symbols

Mathematical Relations and Spaces

$a \lesssim b$	Inequality $a \leq Cb$ for a generic positive constant C independent of the mesh size h.
$a \sim b$	Equivalence up to a generic constant, meaning $a \lesssim b$ and $b \lesssim a$.
$C^k(U, V)$	The space of k-times continuously differentiable functions from U to V.
$W_p^k(\Omega), H^k(\Omega)$	The Sobolev spaces of order k defined on Ω, $H^k(\Omega) = W_2^k(\Omega)$.
$\mathbb{P}_k(T)$	The space of polynomials of degree less than or equal to k on the element T.
$\mathbb{Q}_k(T)$	The space of tensor-product polynomials of degree less than or equal to k in each variable.
V_h, W_h	The scalar and vector-valued finite element spaces defined on the active volume.

Geometric Entities and Discretisation

Γ	The exact, continuous, and closed surface.
Γ_h	The approximate, discrete surface.
$U_{\delta_0}(\Gamma)$	The tubular neighbourhood of radius δ_0 surrounding Γ.
Ω_0	The bounding polytopal domain encompassing $U_{\delta_0}(\Gamma)$.
Ω_h	The active volume composed of the union of all intersected background elements.
$\mathcal{T}_{h,0}$	The quasi-uniform background mesh of Ω_0.
$\mathcal{T}_h$	The active background mesh, containing all elements $T \in \mathcal{T}_{h,0}$ intersected by Γ_h.
$\mathcal{K}_h$	The partition of the discrete surface, containing segments $K = T \cap \Gamma_h$.
$\mathcal{F}_h$	The set of interior faces shared by adjacent elements in $\mathcal{T}_h$.
$\mathcal{E}_h$	The set of interior interfaces (edges) shared by adjacent surface elements in $\mathcal{K}_h$.
h	The characteristic element size of the background mesh $\mathcal{T}_{h,0}$.

Operators and Mappings

$p(x)$	The closest point projection mapping from $U_{\delta_0}(\Gamma)$ onto Γ.
$\rho(x)$	The signed distance function to Γ.
v^e	The constant-normal extension of a function v from Γ to $U_{\delta_0}(\Gamma)$ via $v \circ p$.

v^l	The lifting of a function v from Γ_h to Γ along the exact normal.
$\pi_h, \mathcal{I}_h$	Abstract high-order interpolation operator and the standard Lagrange nodal interpolant.
P_Γ, P_{Γ_h}	Orthogonal projection matrices onto the tangent spaces of Γ and Γ_h.
B	The tangent space transformation matrix from $T_x(\Gamma_h)$ to $T_{p(x)}(\Gamma)$.
J	The surface Jacobian determinant associated with the transformation B.
$[\cdot]$	The jump operator across an interior face or edge.
$\nabla_\Gamma, \nabla_{\Gamma_h}$	Tangential gradients along Γ and Γ_h.
$\Delta_K, \Delta_{\Gamma_h}$	The discrete surface Laplace–Beltrami operator evaluated locally and globally.
$D^j_{\nu_h} v$	The j-th order directional derivative of v in the direction of the normal ν_h.

Variables and Bilinear Forms

ν, ν_h	The exterior unit normal vectors to Γ and Γ_h.
t_K, t_E	The outward-pointing unit co-normal vectors at the element boundaries.
H, H_h	The exact and discrete mean curvature vectors.
$a_h(\cdot,\cdot)$	The discrete inner product evaluated over Γ_h.
$s_h(\cdot,\cdot)$	The total combined stabilization form.
$s_{h,P}(\cdot,\cdot)$	The volume-based (patch) ghost penalty stabilization form.
$s_{h,F}(\cdot,\cdot)$	The face-based ghost penalty stabilization form.
$s_{h,N}(\cdot,\cdot)$	The volume-based normal gradient stabilization form.
$s_{h,n}(\cdot,\cdot)$	The surface-based normal derivative stabilization form.
τ_P, τ_N, τ_n	Non-negative penalty parameters associated with the respective stabilization forms.

Contents

1 Introduction

Computing the mean curvature vector on closed surfaces is required for simulating certain physical phenomena governed by interfacial laws. In many coupled bulk-surface systems, the interface acts as a distinct thermodynamic entity where mean curvature serves as the primary driving force. Examples include the geometric evolution of interfaces via Willmore flow [1], capillary forces in multiphase flows [2], [3], phase separation on manifolds [4], dendritic crystal growth [5], deformation of surfaces in solid mechanics and Gurtin-Murdoch theory [6].

Approximating these interfaces numerically presents specific geometric challenges [7]. While discrete differential geometry offers non-PDE approximations based directly on mesh geometry [8], finite element methods provide a robust alternative. This relies on the weak formulation for the Laplace–Beltrami operator established by Dziuk [9]. By letting this operator act on the position vector, mean curvature can be computed using only first-order tangential derivatives [10]. This formulation helps avoid the suboptimal geometric bounds associated with standard discrete polyhedral approximations [11].

Standard fitted finite element methods, such as Arbitrary Lagrangian–Eulerian (ALE) approaches, conform the computational mesh to the physical boundary. However, evolving interfaces or topological changes severely distort the mesh, necessitating computationally expensive remeshing and introducing projection errors [2], [12]. To circumvent this, unfitted techniques such as the Cut Finite Element Method (CutFEM) [13] embed the physical surface in a stationary background mesh [14]. Allowing the interface to cut arbitrarily through elements eliminates the need for remeshing.

Addressing both geometric representations, Hansbo, Larson and Zahedi [15] established a stabilized finite element approximation for the mean curvature vector on standard meshed and arbitrarily defined cut surfaces in $\mathbb{R}^3$. While the unstabilized formulation fails to converge for piecewise linear approximations, their stabilization framework recovers robust first-order convergence in the L^2-norm.

While this framework is established for linear tetrahedral elements, extending it to other discretizations, such as quadrilateral elements or higher-order simplicial elements, introduces specific mathematical challenges. Although the stabilization of higher-order simplicial CutFEM spaces has been investigated for general surface PDEs [16], applying these discretizations specifically to the mean curvature problem requires further analysis to evaluate their convergence behaviour. Previously, Frachon and Zahedi [17] have numerically investigated the convergence of higher order approximations of one geometry in $\mathbb{R}^2$.

1.1 Objectives

This thesis formalizes and extends the CutFEM mean curvature framework beyond the standard $\mathbb{P}_1$ setting. This work aims to:

- Abstract and generalize the existing $\mathbb{P}_1$ framework for arbitrary dimensions.
- Extend the formulation to multilinear ($\mathbb{Q}_1$) spaces on curved unfitted geometries in $\mathbb{R}^d$, and identify the additional stabilization needed to recover first-order

convergence.

- Consider high-order ($\mathbb{P}_k$) spaces, identify the terms preventing a direct order k error estimate in the present framework, and study the resulting numerical behaviour in $\mathbb{R}^2$ and $\mathbb{R}^3$.

The thesis is primarily methodological, and its direct contribution is to the analysis and numerical evaluation of finite element methods for surface PDEs. The social and ethical aspects are therefore indirect. More reliable curvature approximations may support simulations in interfacial mechanics, material modelling, and biological membranes, but such applications are not evaluated here. The work is instead assessed through accuracy, stability, and computational robustness.

1.2 Outline

The remainder of this thesis is structured as follows. Chapter 2 introduces the continuous geometric definitions and mathematical formulations. The baseline stabilization framework is abstracted in Chapter 3, establishing stability and first-order error estimates for $\mathbb{P}_1$ elements. Chapter 4 extends this framework to multilinear $\mathbb{Q}_1$ elements on curved discrete surfaces and proves first-order convergence under an additional surface-Laplacian control assumption. Chapter 5 considers high-order $\mathbb{P}_k$ extensions and identifies the surface-Laplacian term which prevents a direct order k estimate in the present framework. Numerical experiments are presented in Chapter 6 to validate the $\mathbb{Q}_1$ theory and evaluate the convergence behaviour of the $\mathbb{P}_2$ spaces. Finally, Chapter 7 summarizes the contributions and discusses remaining questions.

2 The Continuous Mean Curvature Vector

2.1 The Abstract Hypersurface

Before we can mathematically construct the mean curvature vector, we must establish its underlying geometric domain. We adopt the standard regularity assumptions necessary for classical geometric partial differential equations and mean curvature flow as presented by Deckelnick, Dziuk and Elliott [7].

Definition 2.1. A subset $\Gamma \subset \mathbb{R}^d$ is called a C^2 *hypersurface* if for each point $x_0 \in \Gamma$ there is a function u defined in a neighbourhood $U \subset \mathbb{R}^d$ of x_0 such that $u \in C^2(U, \mathbb{R})$,

$$U \cap \Gamma = \left\{ x \in U : u(x) = 0 \right\},$$

and

$$\nabla u(x) \neq 0, \quad \forall x \in U \cap \Gamma.$$

The *tangent space* $T_{x_0}\Gamma$ at a point $x_0 \in \Gamma$ is the $(d-1)$-dimensional linear subspace of $\mathbb{R}^d$ that is orthogonal to $\nabla u(x_0)$. It is clear that it is independent of the particular choice of u used to describe Γ.

The functions u give us a way to locally orient the surface; we must also establish the notion of global orientability before moving on to the curvature.

Definition 2.2. A C^2 hypersurface $\Gamma \subset \mathbb{R}^d$ is said to be *orientable* if there exists a neighbourhood U of Γ on which we can define $\nu \in C^1(U, \mathbb{R}^d)$ such that $\nu(x) \perp T_x\Gamma$ and $\left|\nu(x)\right| = 1$ for all $x \in \Gamma$.

Notice that due to the way we define the C^2 hypersurfaces, they implicitly do not have any boundaries. As such, we only need one simple qualifier in order to have sufficiently characterized the geometric domains we will be working with.

Definition 2.3. A C^2 hypersurface $\Gamma \subset \mathbb{R}^d$ is *closed* if it is compact.

In the following, Γ will always denote an orientable, closed C^2 hypersurface.

2.2 Tangential Calculus

In this section, we aim to define the differential operators necessary to formulate equations directly on Γ. By extending surface functions into a surrounding tubular neighbourhood, we can utilize standard Euclidean derivatives and project them orthogonally onto the local tangent space.

Definition 2.4 (Tangential Gradient)**.** Suppose $\Gamma \subset \mathbb{R}^d$ is a closed, orientable C^2 hypersurface with exterior unit normal ν, and let V be a neighbourhood of Γ. For a function $f \in C^1(V, \mathbb{R})$, we define the *tangential gradient* for $x \in \Gamma$ as

$$\nabla_\Gamma f(x) = \nabla f(x) - \Big(\nabla f(x) \cdot \nu(x)\Big)\nu(x), \tag{1}$$

which is the projection of the regular spatial gradient ∇f onto the tangent space $T_x\Gamma$.

In order to show that it makes sense to talk about tangential gradient for functions defined only on Γ, we must demonstrate that the operator is intrinsic to the surface and does not depend on what happens on $V \setminus \Gamma$. We can prove the following proposition:

Proposition 2.1. *The tangential gradient $\nabla_\Gamma f$ only depends on the values f assumes on Γ.*

Proof. Consider two functions $f_1, f_2 \in C^1(V, \mathbb{R})$ which agree on Γ. Then, their difference $g = f_1 - f_2$ is identically zero on Γ. Hence, the spatial gradient ∇g must be orthogonal to Γ. Therefore, $\nabla g(x) = (\nabla g(x) \cdot \nu(x))\nu(x)$, and we have that

$$\begin{aligned}\nabla_\Gamma g(x) &= \nabla g(x) - \Big(\nabla g(x) \cdot \nu(x)\Big)\nu(x) \\ &= \nabla g(x) - \nabla g(x) \\ &= 0.\end{aligned}$$

Since ∇_Γ clearly is a linear operator, we can conclude

$$\nabla_\Gamma f_1 = \nabla_\Gamma f_2.$$

□

Having established the tangential gradient for scalar functions, we naturally extend our calculus to vector fields. We define the tangential divergence by applying the tangential gradient component-wise.

Definition 2.5 (Tangential Divergence)**.** Suppose $\Gamma \subset \mathbb{R}^d$ is a closed, orientable C^2 hypersurface with exterior unit normal ν, and let V be a neighbourhood of Γ. For a vector field $f \in C^1(V, \mathbb{R}^d)$, we define the *tangential divergence* as

$$\nabla_\Gamma \cdot f = \sum_{i=1}^{d} (\nabla_\Gamma f_i)_i \,,$$

where $(\cdot)_i$ is the i-th component of $(\cdot)$.

2.3 The Laplace–Beltrami Operator and Mean Curvature

The Laplace–Beltrami operator is the tangential analogue to the Laplacian. It appears in many types of PDEs on surfaces, and can be used to model e.g. diffusion [18] or phase transitions [4] on surfaces.

Definition 2.6 (Laplace–Beltrami operator)**.** Let $\Gamma \subset \mathbb{R}^d$ be a closed orientable C^2 hypersurface, and let V be a neighbourhood of Γ on which the exterior unit normal ν can be defined. For a function $f \in C^2(V, \mathbb{R})$, the *Laplace–Beltrami operator* Δ_Γ produces the tangential divergence of the tangential gradient of f

$$\Delta_\Gamma f = \nabla_\Gamma \cdot \nabla_\Gamma f.$$

For a vector field $f \in C^2(V, \mathbb{R}^d)$, the operator is defined component-wise

$$(\Delta_\Gamma f)_i = \Delta_\Gamma f_i.$$

At each point $x \in \Gamma$, we can define a matrix $\Pi(x) \in \mathbb{R}^{d \times d}$ by

$$\Pi_{i,j}(x) = \big(\nabla_\Gamma \nu_j(x)\big)_i \,.$$

Since ν is defined in a neighbourhood of Γ, the matrix is well-defined. It is also straightforward to verify that the matrix is symmetric and that $\Pi\nu = 0$. The other $d-1$ eigenvalues and eigenvectors are called principal curvatures and principal directions. Notice that the $d-1$ principal directions form an orthonormal basis for $T_x\Gamma$. The scalar mean curvature is then easily found as $\bar{\kappa} = \frac{1}{d-1}\operatorname{tr}(\Pi)$. Since the trace is equal to the sum of the diagonal, we also see that $\bar{\kappa}(x) = \frac{1}{d-1}\nabla_\Gamma \cdot \nu(x)$.

In a slight abuse of the usual terminology, we shall now define the mean curvature vector.

Definition 2.7 (Mean curvature vector)**.** Let $\Gamma \subset \mathbb{R}^d$ be a closed orientable C^2 hypersurface, with the exterior unit normal ν. The *mean curvature vector* is defined as

$$H = (d-1)\bar{\kappa}(x)\nu(x).$$

We end this section with a strong differential equation which the mean curvature vector satisfies.

Proposition 2.2. *The mean curvature vector H satisfies*

$$H = -\Delta_\Gamma x_\Gamma, \tag{2}$$

where $x_\Gamma : \Gamma \to \Gamma$ is the identity map.

Proof. The statement (2) is of course equivalent to $H_i = -\nabla_\Gamma \cdot \nabla_\Gamma (x_\Gamma)_i$. From (1), we have that

$$\nabla_\Gamma (x_\Gamma)_i = e_i - \nu_i \nu,$$

where e_i is the i-th standard basis vector of $\mathbb{R}^d$. Since the projection onto $T_x\Gamma$ is linear, it is clear that ∇_Γ inherits the product rule from the spatial gradient, hence

$$\nabla_\Gamma \cdot \nabla_\Gamma (x_\Gamma)_i = -(\nu_i(\nabla_\Gamma \cdot \nu) + \nabla_\Gamma \nu_i \cdot \nu).$$

As noted previously, the tangential gradient of a C^1 function on Γ is in $T_x\Gamma$, so $\nabla_\Gamma \nu_i \cdot \nu = 0$. It is now easy to see that

$$\Delta_\Gamma x_\Gamma = -(\nabla_\Gamma \cdot \nu)\nu,$$

so using $\bar{\kappa} = \frac{1}{d-1}\nabla_\Gamma \cdot \nu$, we finally conclude

$$\begin{aligned} -\Delta_\Gamma x_\Gamma &= (d-1)\bar{\kappa}\nu \\ &= H. \end{aligned}$$

□

2.4 The Weak Formulation

While the strong differential equation $H = -\Delta_\Gamma x_\Gamma$ provides a complete continuous description of the mean curvature vector, it is not directly amenable to finite element discretization. Evaluating the Laplace–Beltrami operator requires the underlying surface geometry to possess well-defined second derivatives. As we will see in subsequent chapters, discrete surfaces constructed from piecewise polynomials are typically only C^0-continuous globally.

To bridge this gap and prepare for numerical approximation, we must recast the strong equation into a variational, or weak, formulation. This process will systematically reduce the geometric differentiability requirements.

We begin by establishing the standard functional analytic notation that will be used throughout the remainder of this text.

Definition 2.8 (L^2 Inner Product)**.** Let $L^2(\Gamma, \mathbb{R})$ denote the space of square-integrable scalar functions on Γ. For functions $u, v \in L^2(\Gamma, \mathbb{R})$, we define the *inner product* as

$$\langle u, v \rangle_\Gamma = \int_\Gamma uv \, ds,$$

where ds is the surface area measure on Γ. For vector fields $u, v \in L^2(\Gamma, \mathbb{R}^d)$, the inner product is naturally extended using the standard Euclidean dot product:

$$\langle u, v \rangle_\Gamma = \int_\Gamma u \cdot v \, ds.$$

With the inner product established, we must determine how to perform integration by parts on Γ. Because Γ is a closed surface, standard boundary terms will vanish, but the curvature of the surface introduces a new geometric term. The following lemma, adopted from Gilbarg and Trudinger [19] formalizes this.

Lemma 2.1. *For any vector field $w \in C^1(\Gamma, \mathbb{R}^d)$, the following identity holds:*

$$\int_\Gamma \nabla_\Gamma \cdot w \, ds = \int_\Gamma w \cdot H \, ds. \tag{3}$$

From here, we are able to derive the specific integration by parts formula for the Laplace–Beltrami operator that will be required to relax our strong differential equation.

Lemma 2.2. *For any $u \in C^2(\Gamma, \mathbb{R})$ and $v \in C^1(\Gamma, \mathbb{R})$,*

$$\int_\Gamma (\Delta_\Gamma u) v \, ds = -\int_\Gamma \nabla_\Gamma u \cdot \nabla_\Gamma v \, ds.$$

Proof. Set $w = v \nabla_\Gamma u$, and apply the product rule to see that

$$\nabla_\Gamma \cdot w = \nabla_\Gamma u \cdot \nabla_\Gamma v + (\Delta_\Gamma u) v.$$

Since $\nabla_\Gamma u \in T_x\Gamma$, we have $w \cdot H = 0$. Using Lemma 2.1 and the linearity of the integral yields the desired statement. □

We are now ready to introduce the weak problem, and show that the mean curvature vector H is the unique solution to it.

Definition 2.9 (Weak Problem)**.** Let $\Gamma \subset \mathbb{R}^d$ be a closed, orientable C^2 hypersurface with exterior unit normal ν. Define the bilinear form $a : L^2(\Gamma, \mathbb{R}^d) \times H^1(\Gamma, \mathbb{R}^d) \to \mathbb{R}$

by

$$a(u, v) = \langle u, v \rangle_\Gamma$$

and the linear form $L : H^1(\Gamma, \mathbb{R}^d) \to \mathbb{R}$ by

$$L(v) = \sum_{i=1}^{d} \left\langle \nabla_\Gamma \left(x_\Gamma\right)_i, \nabla_\Gamma v_i \right\rangle_\Gamma, \tag{4}$$

where $x_\Gamma : \Gamma \to \Gamma$ is the identity map. The *weak problem* for the mean curvature vector is then

$$\begin{cases} \text{Find } H \in L^2(\Gamma, \mathbb{R}^d) & \text{such that} \\ a(H, v) = L(v), & \forall v \in H^1(\Gamma, \mathbb{R}^d). \end{cases} \tag{5}$$

Remark. We will also write $\langle \nabla_\Gamma x_\Gamma, \nabla_\Gamma v \rangle_\Gamma$ at times for the linear form (4) and similar expressions.

Proposition 2.3. *The mean curvature vector H defined by* (2) *is the unique solution to the weak problem* (5)*.*

Proof. We first demonstrate existence. Let $H = -\Delta_\Gamma x_\Gamma$ as defined in the strong form (2). The identity map x_Γ is the restriction of the smooth coordinate map to Γ, and its tangential gradient is given by $\nabla_\Gamma (x_\Gamma)_i = e_i - \nu_i \nu$. Applying the tangential divergence to this expression requires differentiating the unit normal field ν. Because Γ is assumed to be a C^2 hypersurface, $\nu \in C^1(V, \mathbb{R}^d)$, which ensures that $\Delta_\Gamma x_\Gamma$ is well-defined and continuous. Since Γ is compact, this continuous vector field is bounded, implying $H \in L^2(\Gamma, \mathbb{R}^d)$. Taking the inner product of H with an arbitrary test function $v \in H^1(\Gamma, \mathbb{R}^d)$ and applying the integration by parts formula from Lemma 2.2 yields exactly the variational equation $a(H, v) = L(v)$.

To prove uniqueness, suppose there exist two solutions $H_1, H_2 \in L^2(\Gamma, \mathbb{R}^d)$ to the weak problem (5). By linearity, their difference $E = H_1 - H_2$ satisfies

$$\langle E, v \rangle_\Gamma = 0 \quad \forall v \in H^1(\Gamma, \mathbb{R}^d).$$

Because the Sobolev space $H^1(\Gamma, \mathbb{R}^d)$ is dense in $L^2(\Gamma, \mathbb{R}^d)$, the only function in $L^2(\Gamma, \mathbb{R}^d)$ that is orthogonal to the entirety of $H^1(\Gamma, \mathbb{R}^d)$ is the zero function. Therefore, $E = 0$ almost everywhere, proving that the solution is unique. □

Remark. The weak formulation (5) fundamentally relaxes the geometric regularity required to define the mean curvature vector. While the strong form (2) requires the surface Γ to be of class C^2 so that the classical Laplace–Beltrami operator can be evaluated, the linear form $L(v)$ only requires the computation of first-order tangential derivatives. Consequently, the weak problem remains well-defined for piecewise smooth or Lipschitz surfaces, where ∇_Γ is defined almost everywhere. This relaxation is the critical theoretical foundation for the finite element methods developed in the following chapters, as it permits the use of discrete surfaces constructed from piecewise polynomials, which lack global C^2 continuity.

3 Linear CutFEM for the Mean Curvature

In this section, we aim to describe how to compute a piecewise linear approximation of the mean curvature vector using Cut Finite Element Methods on a discrete surface. We follow the arguments in the seminal paper by Hansbo, Larson and Zahedi [15] and generalize them to $\mathbb{R}^d$ and abstract stabilization forms, bringing the analysis closer to the framework of the 2025 review article by Burman, Hansbo, Larson *et al.* [14]. Throughout the chapter we assume that Γ is regular enough such that $H \in H^2(\Gamma)$; this is for example the case if Γ is C^4.

3.1 The Discrete Surface

Since Γ is C^2, we have according to Lemma 14.16 in Gilbarg and Trudinger [20] that with $\delta_0 > 0$ such that

$$\sum_{i=1}^{d-1} \delta_0 \left|\kappa_i(x)\right| \leq C < 1, \qquad \forall x \in \Gamma \tag{6}$$

the closest point mapping $p : U_{\delta_0} \to \Gamma$ is well-defined, where

$$U_{\delta_0}(\Gamma) = \left\{ x \in \mathbb{R}^d : \min_{y \in \Gamma} \|x - y\| < \delta_0 \right\}.$$

Furthermore, we can define a signed distance function $\rho : U_{\delta_0} \to \mathbb{R}$ by $\rho(x) = \nu(p(x)) \cdot (x - p(x))$ and $\rho \in C^2(U_{\delta_0}(\Gamma))$. We consider a family of connected piecewise linear surfaces $\Gamma_h \subset U_{\delta_0}$ with $h \in (0, h_0]$ with exterior unit normals ν_h and the approximation properties

$$\left\|x - p(x)\right\|_{L^\infty(\Gamma_h)} \lesssim h^2, \tag{7}$$

$$\|\nu \circ p - \nu_h\|_{L^\infty(\Gamma_h)} \lesssim h. \tag{8}$$

Here, and throughout the rest of this text we use the notation $\alpha(h) \lesssim \beta(h)$ to mean that there is a constant $C \geq 0$ such that, for all h, $0 \leq \alpha(h) \leq C\beta(h)$. It follows from this definition that the relation is transitive (if $\alpha(h) \lesssim \beta(h)$ and $\beta(h) \lesssim \gamma(h)$, then $\alpha(h) \lesssim \gamma(h)$) and preserves addition (if $\alpha_1(h) \lesssim \beta_1(h)$ and $\alpha_2(h) \lesssim \beta_2(h)$, then $\alpha_1(h) + \alpha_2(h) \lesssim \beta_1(h) + \beta_2(h)$). Furthermore, we write $\alpha(h) \sim \beta(h)$ to mean both $\alpha(h) \lesssim \beta(h)$ and $\beta(h) \lesssim \alpha(h)$ hold.

In order to meaningfully compare functions defined on the exact surface Γ with those defined on Γ_h we will need to define two concepts, the first enables us to evaluate the exact curvature H on the discrete surface, and the second to evaluate the solution to (9) on the exact surface.

> **Definition 3.1** (Extension from Γ to $U_{\delta_0}(\Gamma)$)**.** Given a function v defined on Γ, we extend it to the whole tubular neighbourhood $U_{\delta_0}(\Gamma)$ by setting $v^e = v \circ p$.

> **Definition 3.2** (Lifting)**.** Given a function w defined on Γ_h, we define its *lifting* w^l on Γ by $w^l(x) = w(x + t\nu(x))$ where t minimizes $|t|$ under the condition $x + t\nu(x) \in \Gamma_h$.

Remark. The lifting can equivalently be defined by the push-forward $(w^l)^e = w^l \circ p = w$. Consequently, the restriction of p to Γ_h is a bijection between the exact surface Γ and the discrete surface Γ_h.

Remark. By definition, v^e is constant in the direction of ν, meaning that $\nabla_\Gamma v(x) = \nabla v^e(x)$ for $x \in \Gamma$.

3.2 Domain Discretization

Since $\Gamma \subset \mathbb{R}^d$ is compact, we can embed $U_{\delta_0}(\Gamma)$ in a polytopal domain $\Omega_0 \subset \mathbb{R}^d$. Let $\mathcal{T}_{h,0}$ be a quasi-uniform partition of Ω_0 into shape regular simplices T such that

$$\begin{aligned} \operatorname{Diam}(T) &\lesssim h, \quad \text{and,} \\ \frac{\operatorname{Diam}(T)}{\operatorname{diam}(T)} &\lesssim 1, \quad \forall T \in \mathcal{T}_{h,0}. \end{aligned}$$

Here, $\operatorname{Diam}(T) = \max_{x,y \in T} \|x - y\|$ is the diameter of the element T and $\operatorname{diam}(T)$ is the diameter of the largest possible hypersphere inscribed in T. Furthermore, assume that $T \cap \Gamma_h$ is either empty or a subset of a hyperplane for all $T \in \mathcal{T}_{h,0}$. We define the *active mesh* as

$$\mathcal{T}_h = \left\{ T \in \mathcal{T}_{h,0} : T \cap \Gamma_h \neq \emptyset \right\}.$$

Restricting h_0 such that $\bigcup_{T \in \mathcal{T}_h} T \subset U_{\delta_0}(\Gamma)$ for all $h \leq h_0$ ensures a well-defined closest point mapping on the entire active mesh $\mathcal{T}_h$. We also define the partitioning of the discrete surface as

$$\mathcal{K}_h = \left\{ T \cap \Gamma_h : T \in \mathcal{T}_h \right\},$$

and the active volume $\Omega_h = \bigcup_{T \in \mathcal{T}_h} T$; the final component of the mesh is the set of interior faces and edges

$$\begin{aligned} \mathcal{F}_h &= \{ \partial T_1 \cap \partial T_2 : T_1, T_2 \in \mathcal{T}_h \} \\ \mathcal{E}_h &= \{ \partial K_1 \cap \partial K_2 : K_1, K_2 \in \mathcal{K}_h \}. \end{aligned}$$

With the mesh defined, we are ready to introduce the finite element space

$$V_h = \left\{ v \in C^0(\Omega_h) : v|_T \in \mathbb{P}_1(T) \quad \forall T \in \mathcal{T}_h \right\},$$

where $v|_T$ is the restriction of v to T and $\mathbb{P}_k(T)$ is the set of degree k polynomials defined on T. In words, this means that V_h is the space of continuous element-wise linear functions on $\mathcal{T}_h$. As it will be of use to us later, let us also denote $W_h = [V_h]^d$.

Lemma 3.1. *Let $\mathcal{T}_h$ be the active background mesh and $\mathcal{K}_h$ be the associated discrete surface partition. Under the shape regularity assumptions for $\mathcal{T}_h$, the total $(d-2)$-dimensional measure of the element boundaries ∂K satisfies*

$$\sum_{K \in \mathcal{K}_h} \|1\|^2_{L^2(\partial K)} \lesssim h^{-1}.$$

Proof. The exact closed surface Γ has a finite, strictly positive $(d-1)$-dimensional measure. By the geometric approximation properties, the discrete surface Γ_h also has

a uniformly bounded $(d-1)$-dimensional measure. Because the background mesh $\mathcal{T}_{h,0}$ is quasi-uniform with elements of characteristic size h, the active volume Ω_h forms a band around Γ of thickness $\mathcal{O}(h)$. The d-dimensional volume of this band is therefore $\mathcal{O}(h)$. Since each shape-regular element T has a volume of $\mathcal{O}(h^d)$, the number of active elements scales as $\mathcal{O}(h^{-(d-1)})$.

For each $T \in \mathcal{T}_h$, the intersection $K = T \cap \Gamma_h$ is a planar $(d-1)$-polytope. The boundary ∂K is $\Gamma_h \cap \partial T$. Since T is shape-regular with diameter $\mathrm{Diam}(T) \lesssim h$, the maximum $(d-2)$-dimensional measure of any such planar intersection is bounded by $\mathcal{O}(h^{d-2})$.

Summing the maximum boundary measure over all active elements yields the global bound:

$$\sum_{K \in \mathcal{K}_h} \|1\|^2_{L^2(\partial K)} \lesssim h^{1-d} \max_{K \in \mathcal{K}_h} \|1\|^2_{L^2(\partial K)} \lesssim h^{-1}. \qquad \square$$

3.3 The Discrete Problem

With the discretized hypersurface and finite element space defined, we are ready to talk about the discrete counterpart to (5)

$$\begin{cases} \text{Find } H_h \in W_h & \text{such that} \\ A_h(H_h, v) = L_h(v), & \forall v \in W_h. \end{cases} \tag{9}$$

Here the forms are defined by

$$\begin{aligned} A_h(u,v) &= a_h(u,v) + s_h(u,v), \\ a_h(u,v) &= \langle u, v \rangle_{\Gamma_h}, \\ L_h(v) &= \left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} v \right\rangle_{\Gamma_h}, \end{aligned}$$

and s_h is a bilinear stabilization form.

To ensure the well-posedness and first-order convergence of the discrete problem, the stabilization form s_h cannot be chosen arbitrarily. To formalize the requirements, let us define an extended space $W_* = W_h + [H^2(\Omega_h)]^d$. This space is sufficiently regular to contain both our discrete finite element functions and the exact extended mean curvature vector H^e.

Before introducing the stabilization form, we require an interpolation operator to map the exact extended solution into our discrete space.

Definition 3.3. We assume the existence of a linear interpolation operator $\pi_h : W_* \to W_h$ that satisfies the following approximation bounds on the active volume mesh Ω_h for the exact extended mean curvature $H^e \in [H^2(\Omega_h)]^d$:

$$\|H^e - \pi_h H^e\|_{L^2(\Omega_h)} + h \left\|\nabla(H^e - \pi_h H^e)\right\|_{L^2(\Omega_h)} \lesssim h^2 \|H^e\|_{H^2(\Omega_h)}. \tag{10}$$

Furthermore, by applying the standard trace inequality for planar intersections (Hansbo, Hansbo and Larson [13, Lemma 4.2]) to the elements cut by the surface, we assume the interpolant satisfies the corresponding bound on the discrete surface Γ_h:

$$\|H^e - \pi_h H^e\|_{L^2(\Gamma_h)} + h \left\|\nabla_{\Gamma_h}(H^e - \pi_h H^e)\right\|_{L^2(\Gamma_h)} \lesssim h^2 \|H\|_{H^2(\Gamma)}. \tag{11}$$

We assume the existence of a bilinear stabilization form $s_h : W_* \times W_* \to \mathbb{R}$ that satisfies the following three abstract properties:

(A1) **Symmetry and Positivity:** The form $s_h(\cdot,\cdot)$ is symmetric and positive semi-definite on W_*. This naturally induces a stabilization semi-norm defined by $\|v\|_{s_h}^2 = s_h(v,v)$. Consequently, it satisfies the Cauchy–Schwarz inequality:

$$s_h(u,v) \leq \|u\|_{s_h} \|v\|_{s_h}, \quad \forall u, v \in W_*. \tag{A1}$$

(A2) **Weak Consistency and Interpolation:** The stabilization must not excessively penalize the true continuous solution nor its finite element interpolant. We assume that both the exact extended solution H^e and its interpolant $\pi_h H^e$ satisfy the bound:

$$\|H^e\|_{s_h} + \|\pi_h H^e\|_{s_h} \lesssim h \|H\|_{H^2(\Gamma)}. \tag{A2}$$

(A3) **Tangential Gradient and Trace Control:** The stabilization must provide robust control over both the tangential gradient and the traces on element boundaries. We assume the inverse inequality for all $v \in W_h$:

$$h^2 \left\|\nabla_{\Gamma_h} v\right\|_{L^2(\Gamma_h)}^2 + \sum_{K \in \mathcal{K}_h} h \|v\|_{L^2(\partial K)}^2 + \sum_{E \in \mathcal{E}_h} h \left\|[t_E \cdot \nabla_{\Gamma_h} v]\right\|_{L^2(E)}^2 \lesssim \|v\|_{L^2(\Gamma_h)}^2 + \|v\|_{s_h}^2. \tag{A3}$$

These assumptions abstract the stabilization requirements used in the CutFEM curvature analysis of Hansbo, Larson and Zahedi [15]. They separate the proof of the error estimate from the verification of any particular choice of s_h: once a concrete stabilization is shown to satisfy (A1)–(A3), the estimate below applies.

3.4 The *a priori* Error Estimate

The main result of the chapter is the following first-order error estimate for the discrete mean curvature vector. The remainder of the chapter is devoted to its proof and to verifying the assumptions above for concrete stabilization forms.

Theorem 3.2. *Let $H \in [H^2(\Gamma)]^d$ be the exact mean curvature vector of the C^2 continuous surface Γ, and let $H^e = H \circ p$ be its extension to the active volume Ω_h. Let $H_h \in W_h$ be the solution to the discrete formulation* (9)*. Assume that the discrete surface Γ_h satisfies the geometric approximation bounds* (7) *and* (8)*, the stabilization form s_h satisfies assumptions* (A1)–(A3)*, and the interpolation operator π_h satisfies the bounds of Definition 3.3. Then, the error satisfies the first-order bound*

$$\|H^e - H_h\|_{L^2(\Gamma_h)}^2 + \|H^e - H_h\|_{s_h}^2 \lesssim h^2. \tag{12}$$

3.5 Analysis of the Discrete Problem

We now turn to the analysis of the discrete problem. The aim is to prove Theorem 3.2, but we first collect the auxiliary estimates needed for the proof. The main difficulty is

that the exact mean curvature vector H is defined on Γ, whereas the discrete solution H_h is defined on Γ_h. We therefore begin with geometric estimates comparing quantities on the two surfaces, then establish stability of the discrete formulation, and finally derive the consistency bounds used in the proof of the *a priori* estimate.

3.5.1 Geometric estimates

Definition 3.4 (Tangent Space Transformation)**.** To map tangential gradients between the discrete and continuous surfaces, we define the transformation matrix $B_x : T_x(\Gamma_h) \to T_{p(x)}(\Gamma)$ as

$$B_x = P_\Gamma(I - \rho(x)\mathcal{H}(x))P_{\Gamma_h},$$

where P_Γ and P_{Γ_h} are the orthogonal projections onto the respective tangent planes, ρ is the signed distance function, and $\mathcal{H} = \nabla \otimes \nabla\rho$ is the Hessian of the distance function.

Applying the chain rule to the extension $v^e = v \circ p$, we obtain the fundamental relation $\nabla_{\Gamma_h} v^e = B^T \nabla_\Gamma v$. To bound the errors introduced by this change of variables, we rely on the following geometric estimates.

Lemma 3.3. *Let ρ be the signed distance function to Γ. For a sufficiently fine background mesh (i.e., for h_0 sufficiently small), the Hessian of the distance function, $\mathcal{H} = \nabla \otimes \nabla\rho$, is uniformly bounded on the active volume Ω_h:*

$$\|\mathcal{H}\|_{L^\infty(\Omega_h)} \lesssim 1.$$

Proof. According to Lemma 14.16 and Lemma 14.17 in Gilbarg and Trudinger [20], the Hessian $\mathcal{H}$ of ρ can at a point $x \in U_{\delta_0}(\Gamma)$ be written as

$$\mathcal{H}(x) = \sum_{i=1}^{d-1} \frac{\kappa_i(p(x))}{1 + \rho(x)\kappa_i(p(x))} a_i(p(x)) \otimes a_i(p(x)),$$

where κ_i are the principal curvatures and a_i the principal directions. Since $\left|\rho(x)\right| < \delta_0$, and using (6), we can see that

$$\begin{aligned}\left|\frac{\kappa_i(p(x))}{1 + \rho(x)\kappa_i(p(x))}\right| &= \frac{\left|\kappa_i(p(x))\delta_0\right|}{\left|\delta_0 + \rho\delta_0\kappa_i(p(x))\right|} \\ &\leq \frac{C}{\delta_0 - C\left|\rho\right|} \\ &\leq \frac{1}{\delta_0}\frac{C}{1 - C}.\end{aligned}$$

Hence, $\|\mathcal{H}\|_2$ is bounded on $U_{\delta_0}(\Gamma)$, noting that $\Omega_h \subset U_{\delta_0}(\Gamma)$; we are done. □

Definition 3.5. Let $J : \Gamma_h \to \mathbb{R}$ denote the surface Jacobian of the closest point mapping p, such that the change of variables between the surface measures satisfy $d\Gamma = J\, d\Gamma_h$. Algebraically, $J(x)$ is the absolute determinant of $B_x : T_x(\Gamma_h) \to T_{p(x)}(\Gamma)$.

Lemma 3.4. *If the surface approximation assumptions hold, the transformation matrix B and the surface Jacobian J satisfy the uniform bounds:*

$$\begin{aligned} \|B\|_{L^\infty(\Gamma_h)} \lesssim 1, \quad \left\|B^{-1}\right\|_{L^\infty(\Gamma)} &\lesssim 1, \\ \left\|P_\Gamma - BB^T\right\|_{L^\infty(\Gamma_h)} &\lesssim h^2, \\ \|1 - J\|_{L^\infty(\Gamma_h)} &\lesssim h^2. \end{aligned} \tag{13}$$

Remark. Because the discrete surface Γ_h is piecewise planar, the discrete normal ν_h is constant on the relative interior of each element $K \in \mathcal{K}_h$, but undefined on ∂K. Since the intersection skeleton $\bigcup_K \partial K$ has measure zero, it does not affect the essential supremum. Thus, in the following proof, all pointwise matrix bounds are evaluated for $x \in K \setminus \partial K$, which yields the L^∞ estimates.

Proof. The first inequality follows from

$$\left\|P_\Gamma(I - \rho\mathcal{H})P_{\Gamma_h}\right\|_2 \leq \|P_\Gamma\|_2 \|I - \rho\mathcal{H}\|_2 \left\|P_{\Gamma_h}\right\|_2,$$

the fact that the norms of the projections are 1, and $\|I - \rho\mathcal{H}\|_2 \leq 1 + |\rho| \frac{C}{1-C} \lesssim 1 + h^2 \lesssim 1$.

To bound the difference $P_\Gamma - BB^T$, we expand the product $B_x B_x^T$:

$$\begin{aligned} B_x B_x^T &= P_\Gamma(I - \rho\mathcal{H})P_{\Gamma_h}(I - \rho\mathcal{H})P_\Gamma \\ &= P_\Gamma(I - \rho\mathcal{H})(I - \nu_h \otimes \nu_h)(I - \rho\mathcal{H})P_\Gamma \\ &= P_\Gamma(I - \rho\mathcal{H})^2 P_\Gamma - \left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right) \otimes \left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right), \end{aligned}$$

where we have used $P_{\Gamma_h}^2 = P_{\Gamma_h}$. Expanding the first term yields

$$P_\Gamma(I - \rho\mathcal{H})^2 P_\Gamma = P_\Gamma - 2\rho P_\Gamma \mathcal{H} P_\Gamma + \rho^2 P_\Gamma \mathcal{H}^2 P_\Gamma,$$

and for the second term, we use $P_\Gamma \nu = 0$ so that

$$P_\Gamma(I - \rho\mathcal{H})\nu_h = P_\Gamma(\nu_h - \nu) - \rho P_\Gamma \mathcal{H}\nu_h.$$

Due to the approximation properties (7) and (8), we see that this vector has norm $\lesssim h$, so for the outer product we have

$$\left\|\left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right) \otimes \left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right)\right\|_2 \lesssim h^2.$$

Using the triangle inequality allows us to write

$$\begin{aligned} \left\|P_\Gamma - B_x B_x^T\right\|_2 &= \left\|2\rho P_\Gamma \mathcal{H} P_\Gamma - \rho^2 P_\Gamma \mathcal{H}^2 P_\Gamma + \left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right) \otimes \left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right)\right\|_2 \\ &\leq \|2\rho P_\Gamma \mathcal{H} P_\Gamma\|_2 + \left\|\rho^2 P_\Gamma \mathcal{H}^2 P_\Gamma\right\|_2 + \left\|\left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right) \otimes \left(P_\Gamma(I - \rho\mathcal{H})\nu_h\right)\right\|_2, \end{aligned}$$

leading to the conclusion

$$\begin{aligned} \left\|P_\Gamma - B_x B_x^T\right\|_2 &\lesssim h^2 + h^4 + h^2 \\ &\lesssim h^2. \end{aligned}$$

Taking the supremum over $K \setminus \partial K$ yields $\left\| P_\Gamma - BB^T \right\|_{L^\infty(\Gamma_h)} \lesssim h^2$.

To bound the spectral norm of B^{-1}, we use the identity

$$B_x^{-1} = B_x^T \left(B_x B_x^T \right)^{-1}.$$

Since B_x has full rank on these tangent spaces, the inverse is geometrically well-defined for all $h \leq h_0$. Furthermore, since P_Γ is the identity on $T_{p(x)}(\Gamma)$ and due to preceding arguments, we can write $B_x B_x^T = I - E$ with $\|E\|_2 \lesssim h^2$. For sufficiently small h where $\|E\|_2 < 1$, we can apply Theorem 7.3-1 in Kreyszig [21], yielding

$$\left\| (I-E)^{-1} \right\|_2 \leq \frac{1}{1 - \|E\|_2} \lesssim 1.$$

Now, since $\left\| B_x^T \right\|_2 = \|B_x\|_2 \lesssim 1$, taking the supremum yields $\left\| B^{-1} \right\|_{L^\infty(\Gamma)} \lesssim 1$.

Finally, we need to bound $\|1 - J\|_{L^\infty(\Gamma_h)}$. Since the determinant is continuous and $J = \sqrt{\det(B_x B_x^T)}$, this is straightforward using $\det(B_x B_x^T) = 1 + \mathcal{O}(\|E\|_2) = 1 + \mathcal{O}(h^2)$ and a first-order Taylor approximation of $\sqrt{1 + (\cdot)}$. □

The bounds on the transformation determinant imply that the measures of the two surfaces are equivalent up to a second-order perturbation, leading to the following norm equivalences.

Lemma 3.5. *For any $v_1 \in L^p(\Gamma), u_1 \in L^p(\Gamma_h)$ and $v_2 \in W_p^1(\Gamma), u_2 \in W_p^1(\Gamma_h)$, with $p \in [1, \infty)$, we have the equivalences*

$$\|v_1\|_{L^p(\Gamma)} \sim \left\| v_1^e \right\|_{L^p(\Gamma_h)}, \qquad \|u_1\|_{L^p(\Gamma_h)} \sim \left\| u_1^l \right\|_{L^p(\Gamma)},$$
$$\|\nabla_\Gamma v_2\|_{L^p(\Gamma)} \sim \left\| \nabla_{\Gamma_h} v_2^e \right\|_{L^p(\Gamma_h)}, \qquad \left\| \nabla_{\Gamma_h} u_2 \right\|_{L^p(\Gamma_h)} \sim \left\| \nabla_\Gamma u_2^l \right\|_{L^p(\Gamma)}.$$

Proof. By the change of variables $d\Gamma = J \, d\Gamma_h$, the L^p norm expands as

$$\|v_1\|_{L^p(\Gamma)}^p = \int_{\Gamma_h} \left| v_1^e \right|^p J \, d\Gamma_h.$$

From Lemma 3.4, the surface Jacobian and its inverse are uniformly bounded for the considered family of discrete surfaces, satisfying $\|J\|_{L^\infty(\Gamma_h)} \lesssim 1$ and $\left\| J^{-1} \right\|_{L^\infty(\Gamma_h)} \lesssim 1$. Applying these scalar bounds to the integrals yields $\|v_1\|_{L^p(\Gamma)}^p \lesssim \left\| v_1^e \right\|_{L^p(\Gamma_h)}^p$ and $\left\| v_1^e \right\|_{L^p(\Gamma_h)}^p \lesssim \|v_1\|_{L^p(\Gamma)}^p$, which establishes the function value equivalences.

For the tangential gradients, the chain rule yields $\nabla_{\Gamma_h} v_2^e = B^T \nabla_\Gamma v_2$. Integrating over the discrete surface and changing variables gives $\left\| \nabla_{\Gamma_h} v_2^e \right\|_{L^p(\Gamma_h)}^p = \int_\Gamma \left| B^T \nabla_\Gamma v_2 \right|^p J^{-1} \, d\Gamma$. Using the matrix norm bound $\left| B^T \nabla_\Gamma v_2 \right| \leq \|B\|_{L^\infty(\Gamma_h)} \left| \nabla_\Gamma v_2 \right|$ along with the uniform bound $\|B\|_{L^\infty(\Gamma_h)} \lesssim 1$ from Lemma 3.4, the integrand is bounded. Combined with the uniform bound on J^{-1}, this evaluates to $\left\| \nabla_{\Gamma_h} v_2^e \right\|_{L^p(\Gamma_h)}^p \lesssim \|\nabla_\Gamma v_2\|_{L^p(\Gamma)}^p$. The reverse inequality follows identically using the inverse relation $\nabla_\Gamma v_2 = B^{-T} \nabla_{\Gamma_h} v_2^e$ alongside the uniform L^∞ bounds for B^{-1} and J. □

We end this section with an error estimate for the discrete embedding.

Lemma 3.6. *If the surface approximation assumptions* (7) *and* (8) *hold, then*

$$\left\|x_\Gamma^e - x_{\Gamma_h}\right\|_{L^\infty(\Gamma_h)} + h\left\|\nabla_{\Gamma_h}(x_\Gamma^e - x_{\Gamma_h})\right\|_{L^\infty(\Gamma_h)} \lesssim h^2.$$

Proof. The first term can be bounded by noting that for $x \in \Gamma_h$ we have $x_\Gamma^e - x_{\Gamma_h} = p(x) - x$, so that $\left\|x_\Gamma^e - x_{\Gamma_h}\right\|_{L^\infty(\Gamma_h)} = \|\rho\| \lesssim h^2$. For the second term, we have the identities $\nabla_{\Gamma_h} x_{\Gamma_h} = P_{\Gamma_h}$ and $\nabla_{\Gamma_h} x_\Gamma^e = B_x^T P_\Gamma = B_x^T$, so

$$\begin{aligned}
\left\|\nabla_{\Gamma_h}(x_\Gamma^e - x_{\Gamma_h})\right\|_{L^\infty(\Gamma_h)} &= \left\|P_{\Gamma_h} - B_x^T\right\|_{L^\infty(\Gamma_h)} \\
&= \left\|P_{\Gamma_h}(I - P_\Gamma) + P_{\Gamma_h}\rho\mathcal{H}P_\Gamma\right\|_{L^\infty(\Gamma_h)} \\
&\leq \left\|P_{\Gamma_h}\nu \otimes \nu\right\|_{L^\infty(\Gamma_h)} + \|\rho\|_{L^\infty(\Gamma_h)}\left\|P_h\mathcal{H}P_\Gamma\right\|_{L^\infty(\Gamma_h)} \\
&\leq \left\|\left(P_{\Gamma_h}(\nu - \nu_h)\right) \otimes \nu\right\|_{L^\infty(\Gamma_h)} + \|\rho\|_{L^\infty(\Gamma_h)}\left\|P_h\mathcal{H}P_\Gamma\right\|_{L^\infty(\Gamma_h)} \\
&\lesssim h + h^2 \\
&\lesssim h.
\end{aligned}$$

Here, we have used Lemma 3.3 and the approximation assumptions (7) and (8). □

3.5.2 Stability

Theorem 3.7. *Let $H_h \in W_h$ be the solution to the discrete problem* (9)*. Assume that the discrete surface Γ_h satisfies the geometric approximation bounds* (7) *and* (8)*, and that the stabilization form s_h satisfies the tangential gradient control* (A3)*. Then the discrete solution satisfies the stability bound*

$$\|H_h\|_{L^2(\Gamma_h)}^2 + \|H_h\|_{s_h}^2 \lesssim 1. \tag{14}$$

Proof. Setting $v = H_h$ in the discrete form, we see that

$$\begin{aligned}
\|H_h\|_{L^2(\Gamma_h)}^2 + \|H_h\|_{s_h}^2 &= L_h(H_h) \\
&= \left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} H_h \right\rangle_{\Gamma_h} \\
&= \sum_{K \in \mathcal{K}_h} \left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} H_h \right\rangle_K
\end{aligned}$$

By applying partial integration to each $K \in \mathcal{K}_h$

$$\begin{aligned}
\left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} H_h \right\rangle_K &= \left\langle t_K \cdot \nabla_K x_{\Gamma_h}, H_h \right\rangle_{\partial K} - \left\langle \Delta_K x_{\Gamma_h}, H_h \right\rangle_K \\
&= \left\langle t_K \cdot \nabla_K x_{\Gamma_h}, H_h \right\rangle_{\partial K},
\end{aligned} \tag{15}$$

where t_K is the outward normal of ∂K which is orthogonal to ν_K, and the second term evaluates to zero since K is planar and x_{Γ_h} is linear. Since $\nabla_K x_{\Gamma_h}|_K$ is the identity on

$T_x(K)$, using Cauchy–Schwarz yields

$$\begin{aligned}\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} &= \sum_{E\in\mathcal{E}_h} \left\langle [t_E], H_h \right\rangle_E \\ &\leq \sum_{E\in\mathcal{E}_h} \left(h^{-1} \left\| [t_E] \right\|^2_{L^2(E)} \right)^{1/2} \left(h \|H_h\|^2_{L^2(E)} \right)^{1/2} \\ &\leq \left(\sum_{E\in\mathcal{E}_h} h^{-1} \left\| [t_E] \right\|^2_{L^2(E)} \right)^{1/2} \left(\sum_{E\in\mathcal{E}_h} h \|H_h\|^2_{L^2(E)} \right)^{1/2}\end{aligned}$$

The first factor can be bounded using (8). Recognizing that the jump in the co-normal t is the same as the jump in the normal ν_h (see the remark after the proof for justification), we evaluate the terms at the interface E. Because the exact normal is continuous and $\nu_{h,i}$ is polynomial on the interior of K_i, taking the limit to the boundary ensures the global approximation assumption holds pointwise on E:

$$\begin{aligned}\left\| [t_E] \right\|_{L^\infty(E)} &= \left\| \nu_{h,1} - \nu_{h,2} \right\|_{L^\infty(E)} \\ &\leq \left\| \nu_{h,1} - \nu \circ p \right\|_{L^\infty(E)} + \left\| \nu \circ p - \nu_{h,2} \right\|_{L^\infty(E)} \\ &\lesssim h.\end{aligned}$$

Using this pointwise bound, the sum over the edges becomes bounded

$$\sum_{E\in\mathcal{E}_h} h^{-1} \left\| [t_E] \right\|^2_{L^2(E)} \lesssim \sum_{E\in\mathcal{E}_h} h \|1\|^2_{L^2(E)} \lesssim 1,$$

due to Lemma 3.1.

Now, for the second factor, we must bound the L^2-norm of our discrete solution along the element interfaces. Because each interior edge $E \in \mathcal{E}_h$ is shared by exactly two adjacent elements in the discrete surface, we can bound the sum over the edges by the sum over the element boundaries:

$$\sum_{E\in\mathcal{E}_h} h \|H_h\|^2_{L^2(E)} \leq \sum_{K\in\mathcal{K}_h} h \|H_h\|^2_{L^2(\partial K)}.$$

Here, we can invoke (A3), such that

$$\sum_{K\in\mathcal{K}_h} h \|H_h\|^2_{L^2(\partial K)} \lesssim \|H_h\|^2_{L^2(\Gamma_h)} + s_h(H_h, H_h).$$

Substituting both bounded factors back into our Cauchy–Schwarz inequality yields:

$$\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \lesssim \left(\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \right)^{1/2}.$$

Dividing both sides by $\left(\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \right)^{1/2}$ gives

$$\left(\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \right)^{1/2} \lesssim 1.$$

Squaring this result yields the desired stability bound, concluding the proof. □

Remark (Equivalence of Tangent and Normal Jumps)**.** Let K_1 and K_2 be adjacent discrete surface elements sharing an interface E. The element normals $\nu_{h,i}$ and the outward co-normals t_{E,K_i} are all orthogonal to E. Consequently, these four vectors lie entirely within the two-dimensional plane normal to E.

Within this plane, the co-normal is obtained by a $\pi/2$ rotation of the element normal. Let R denote this orthogonal rotation matrix. Given a consistent orientation across the interface, the outward-pointing co-normals satisfy $t_{E,K_1} = R\nu_{h,1}$ and $t_{E,K_2} = -R\nu_{h,2}$.

The jump in the co-normal vectors is therefore

$$[t_E] = t_{K_1} + t_{K_2} = R(\nu_{h,1} - \nu_{h,2}).$$

Since R is an isometry, it preserves the Euclidean norm, yielding the identity

$$\left\|[t_E]\right\|_{L^2(E)} = \left\|\nu_{h,1} - \nu_{h,2}\right\|_{L^2(E)}.$$

3.5.3 Consistency estimates

It remains to quantify the consistency errors caused by replacing the exact surface Γ by the discrete surface Γ_h. These errors enter the discrete error equation through the L^2 inner product and the tangential gradient term. The following two lemmas bound these contributions.

Lemma 3.8. *For any test function $v \in W_h$, we have the inequality:*

$$\left|a_h(\pi_h H^e, v) - a(H, v^l)\right| \lesssim h^2 \|v\|_{L^2(\Gamma_h)}.$$

Proof. The statement follows from applying Lemma 3.4, (11), Cauchy–Schwarz, and Hölder's inequality:

$$\begin{aligned}
\left|a_h(\pi_h H^e, v) - a(H, v^l)\right| &= \left|\int_{\Gamma_h} \pi_h H^e \cdot v - H^e \cdot v J \, d\Gamma_h\right| \\
&= \left|\int_{\Gamma_h} (\pi_h H^e - H^e) \cdot v + (1 - J) H^e \cdot v \, d\Gamma_h\right| \\
&\leq \|\pi_h H^e - H^e\|_{L^2(\Gamma_h)} \|v\|_{L^2(\Gamma_h)} + \|1 - J\|_{L^\infty(\Gamma_h)} \|H^e\|_{L^2(\Gamma_h)} \|v\|_{L^2(\Gamma_h)} \\
&\lesssim h^2 \|v\|_{L^2(\Gamma_h)}.
\end{aligned}$$

□

Lemma 3.9. *Let $L^*(v) = \left\langle \nabla_{\Gamma_h} x^e_\Gamma, \nabla_{\Gamma_h} v \right\rangle_{\Gamma_h}$. For any $v \in W_h$, we have*

$$\left|L(v^l) - L^*(v)\right| \lesssim h^2 \left\|\nabla_{\Gamma_h} v\right\|_{L^2(\Gamma_h)}.$$

Proof. Using the identities $\nabla_{\Gamma_h}(x^e_\Gamma)_i = B^T P_\Gamma e_i = B^T e_i$ and $\nabla_{\Gamma_h} v_i = B^T \nabla_\Gamma v_i^l$, we map the integral to Γ via $d\Gamma_h = J^{-1} d\Gamma$:

$$L^*(v) = \sum_{i=1}^{d} \int_\Gamma e_i \cdot (B B^T \nabla_\Gamma v_i^l) J^{-1} \, d\Gamma.$$

Comparing this to $L(v^l) = \sum_{i=1}^d \int_\Gamma e_i \cdot (P_\Gamma \nabla_\Gamma v_i^l)\, d\Gamma$, the difference is bounded by applying the triangle inequality and Lemma 3.4:

$$\left|L(v^l) - L^*(v)\right| \leq \sum_{i=1}^d \int_\Gamma \left|e_i \cdot (P_\Gamma - J^{-1}BB^T)\nabla_\Gamma v_i^l\right| \, d\Gamma \lesssim h^2 \left\|\nabla_\Gamma v^l\right\|_{L^2(\Gamma)}.$$

Applying Lemma 3.5 yields the result. □

Proof of Theorem 3.2. We begin by using the triangle inequality to write

$$\begin{aligned}\|H^e - H_h\|_{L^2(\Gamma_h)}^2 + \|H^e - H_h\|_{s_h}^2 \lesssim \|H^e - \pi_h H^e\|_{L^2(\Gamma_h)}^2 + \|\pi_h H^e - H_h\|_{L^2(\Gamma_h)}^2 \\ + \|H^e - \pi_h H^e\|_{s_h}^2 + \|\pi_h H^e - H_h\|_{s_h}^2.\end{aligned}$$

By (11) and (A2) we have the first and third term $\lesssim h^2$. For the rest of the terms, set $v = \pi_h H^e - H_h$, and see that

$$\begin{aligned}\|v\|_{L^2(\Gamma_h)}^2 + \|v\|_{s_h}^2 &= a_h(\pi_h H^e, v) + s_h(\pi_h H^e, v) - L_h(v) \\ &= a_h(\pi_h H^e, v) - a(H, v^l) + s_h(\pi_h H^e, v) + L(v^l) - L_h(v).\end{aligned}$$

For the first term, use Lemma 3.8 to see

$$\left|a_h(\pi_h H^e, v) - a(H, v^l)\right| \lesssim h^2 \|v\|_{L^2(\Gamma_h)}.$$

The second term is bounded by using Cauchy–Schwarz and (A2) to see

$$\left|s_h(\pi_h H^e, v)\right| \lesssim h \|v\|_{s_h}.$$

We split what is left using the triangle inequality,

$$\left|L(v^l) - L_h(v)\right| \leq \left|L(v^l) - L^*(v)\right| + \left|L^*(v) - L_h(v)\right|.$$

By Lemma 3.9, the first term is $\lesssim h^2 \left\|\nabla_{\Gamma_h} v\right\|_{L^2(\Gamma_h)}$. By (A3) we get

$$\left|L(v^l) - L^*(v)\right| \lesssim h\sqrt{\|v\|_{L^2(\Gamma_h)}^2 + \|v\|_{s_h}^2}$$

We expand the second term as a sum over $\mathcal{K}_h$

$$\begin{aligned}L^*(v) - L_h(v) &= \left\langle \nabla_{\Gamma_h} x_\Gamma^e, \nabla_{\Gamma_h} v\right\rangle_{\Gamma_h} - \left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} v\right\rangle_{\Gamma_h} \\ &= \sum_{K \in \mathcal{K}_h} \left\langle \nabla_{\Gamma_h} x_\Gamma^e - \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} v\right\rangle_K\end{aligned}$$

Writing $w = x_\Gamma^e - x_{\Gamma_h}$, we perform partial integration on each $K \in \mathcal{K}_h$ and $i = 1, \ldots, d$

$$\begin{aligned}\left\langle \nabla_{\Gamma_h} w_i, \nabla_{\Gamma_h} v_i\right\rangle_K &= \left\langle w_i, t_K \cdot \nabla_{\Gamma_h} v_i\right\rangle_{\partial K} - \left\langle w_i, \Delta_K v_i\right\rangle_K \\ &= \left\langle w_i, t_K \cdot \nabla_{\Gamma_h} v_i\right\rangle_{\partial K},\end{aligned} \tag{16}$$

where the last term vanishes since K is planar and v is linear. Substituting this back and using Cauchy–Schwarz, we get

$$\begin{aligned}\left|L^*(v) - L_h(v)\right| &= \left|\sum_{E\in\mathcal{E}_h}\sum_{i=1}^{d}\left\langle w_i, [t_E\cdot\nabla_{\Gamma_h} v_i]\right\rangle\right| \\ &\le \sum_{E\in\mathcal{E}_h} \|w\|_{L^2(E)} \left\|[t_E\cdot\nabla_{\Gamma_h} v]\right\|_{L^2(E)} \\ &\lesssim \left(\sum_{K\in\mathcal{K}_h} h^{-1}\|w\|^2_{L^2(\partial K)}\right)^{1/2}\left(\sum_{E\in\mathcal{E}_h} h\left\|[t_E\cdot\nabla_{\Gamma_h} v]\right\|^2_{L^2(E)}\right)^{1/2}\end{aligned}$$

The first factor is $\lesssim h$ from Hölder's inequality, (7) and Lemma 3.1. The second factor is $\lesssim \sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}}$ from (A3). Looking back to (3.5.3) we see

$$\begin{aligned}\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h} &\lesssim h^2\|v\|_{L^2(\Gamma_h)} + h\|v\|_{s_h} + h\sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}} + h\sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}} \\ &\lesssim h\sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}}.\end{aligned}$$

Dividing both sides by $\sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}}$ and squaring the result finalizes the proof. □

3.6 Example Stabilization Forms

In this section, we provide concrete examples of stabilization forms that possess the three abstract properties assumed in the previous analysis, and as such inherit the results of Theorems 3.7 and 3.2.

Definition 3.6 (Face-Based Ghost Penalty)**.** Let $\mathcal{F}_h$ denote the set of all interior faces of the active volume mesh $\mathcal{T}_h$. For piecewise linear elements, the face-based ghost penalty introduced by Hansbo, Larson and Zahedi [15] penalizes the jump in the normal gradient across these internal faces. For $u, v \in W_*$, it is defined component-wise as

$$s_{h,F}(u,v) = \sum_{i=1}^{d}\sum_{F\in\mathcal{F}_h} \tau_F \left\langle [\nu_F\cdot\nabla u_i], [\nu_F\cdot\nabla v_i]\right\rangle_F, \tag{17}$$

where ν_F is a fixed unit normal to the face F, $[\cdot]$ denotes the jump operator across the face, and $\tau_F > 0$ is a chosen stabilization parameter.

To satisfy the interpolation bounds (Definition 3.3) and the weak consistency assumption (A2), we must also specify a suitable interpolation operator.

Definition 3.7 (Lagrange Interpolation)**.** For the finite element space W_h composed of continuous, piecewise linear vector functions ($\mathbb{P}_1$), we define the interpolation operator $\pi_h = \mathcal{I}_h$, where $\mathcal{I}_h : W_* \to W_h$ is the standard nodal Lagrange interpolation operator.

Proposition 3.1 (Suitability of the Ghost Penalty)**.** *The face-based ghost penalty* (17) *combined with the Lagrange interpolant* $\mathcal{I}_h$ *satisfies the abstract stabilization properties: symmetry and positivity, weak consistency* (A2)*, and tangential gradient control* (A3)*. Consequently, the discrete formulation* (9) *achieves the a priori error bound established in Theorem 3.2.*

Proof. The symmetry and positivity of $s_h(\cdot,\cdot)$ follow from its definition as a sum of inner products. The formal proofs establishing both the tangential gradient control and the weak consistency are detailed by Hansbo, Larson and Zahedi [15]. The gradient control is achieved by using the ghost penalty to extend the stability of the discrete formulation from the surface Γ_h to the entire active volume Ω_h.

For the weak consistency, we note that because the exact extended mean curvature $H^e \in [H^2(\Omega_h)]^d$ possesses a continuous gradient across the internal mesh faces, its jump $[\nu_F \cdot \nabla H^e]$ vanishes. Consequently, the stabilization form evaluated on the Lagrange interpolant reduces to a penalty on the gradient interpolation error, $[\nu_F \cdot \nabla(\mathcal{I}_h H^e - H^e)]$. As established in the aforementioned reference, bounding this face-based jump by the continuous volume norms yields the interpolation estimate:

$$s_h(\mathcal{I}_h H^e, \mathcal{I}_h H^e) \lesssim h \left\| H^e \right\|^2_{H^2(\Omega_h)} \lesssim h^2 \left\| H \right\|^2_{H^2(\Gamma)}.$$

Taking the square root gives $\left\| \mathcal{I}_h H^e \right\|_{s_h} \lesssim h$, which satisfies the abstract framework requirements. □

Remark. The face-based ghost penalty can also be applied to higher order finite element spaces. For a $\mathbb{P}_k$ space it would take the form

$$s_{h,F}(u,v) = \sum_{j=1}^{k} \sum_{i=1}^{d} \sum_{F \in \mathcal{F}_h} \tau_{F,j} h^{2(j-1)} \left\langle [D^j_{\nu_F} u_i], [D^j_{\nu_F} v_i] \right\rangle_F. \tag{18}$$

While the face-based ghost penalty (17) is standard, explicitly evaluating jumps of higher-order normal derivatives for high-order spaces becomes computationally expensive and complex to implement. An alternative that avoids explicit derivative evaluations is the volume-based (or patch) ghost penalty.

Definition 3.8 (Volume-Based Ghost Penalty)**.** Let $\omega_F = T_1 \cup T_2$ denote the macro-element patch consisting of the two adjacent elements sharing an interior face $F \in \mathcal{F}_h$. Let $u^*_{i,1}$ and $u^*_{i,2}$ denote the canonical polynomial extensions of the restrictions $u_i|_{T_1}$ and $u_i|_{T_2}$ to the entire patch ω_F, respectively. The volume-based ghost penalty is defined component-wise as

$$s_{h,P}(u,v) = \sum_{i=1}^{d} \sum_{F \in \mathcal{F}_h} \tau_P h^{-3} \left\langle [u_i], [v_i] \right\rangle_{\omega_F}, \tag{19}$$

where $[u_i] = u^*_{i,1} - u^*_{i,2}$ represents the jump between the extended polynomials evaluated over the volume of the patch, and $\tau_P > 0$ is a chosen stabilization parameter.

Proposition 3.2 (Equivalence of the Volume-Based Penalty)**.** *The volume-based ghost penalty* (19) *is functionally equivalent to the face-based stabilization regardless*

of the polynomial order. Consequently, it satisfies the abstract stabilization properties required for Theorem 3.2.

Proof. When introducing this penalty, Preuß [22] proved that

$$s_{h,P} \lesssim s_{h,F}$$

via a Taylor expansion. This inequality directly establishes the weak consistency properties (A2) of $s_{h,P}$.

The reverse inequality can be established by considering $v = u^*_{i,1} - u^*_{i,2}$ on a patch ω_F. Because $\left|D^j_{\nu_F} v\right| \leq \left|D^j v\right|$ we have

$$\sum_{j=1}^{k} \tau_{F,j} h^{2(j-1)} \left\| D^j_{\nu_F} v \right\|^2_{L^2(F)} \leq \sum_{j=1}^{k} \tau_{F,j} h^{2(j-1)} \left\| D^j v \right\|^2_{L^2(F)}.$$

From here, we can apply a discrete trace inequality, like Lemma 12.8 in Ern and Guermond [23], to see

$$\begin{aligned} h^{2(j-1)} \left\| D^j v \right\|^2_{L^2(F)} &\lesssim h^{2j-3} \left\| D^j v \right\|^2_{L^2(T_\alpha)} \\ &\lesssim h^{2j-3} \left\| D^j v \right\|^2_{L^2(\omega_F)}. \end{aligned}$$

Here T_α is either of the two elements contained in the patch ω_F. We now use Lemma 12.1 in Ern and Guermond [23] for each term so that

$$h^{2j-3} \left\| D^j v \right\|^2_{L^2(\omega_F)} \lesssim h^{-3} \left\| v \right\|^2_{L^2(\omega_F)}.$$

Summing over $F \in \mathcal{F}_h$ and $i = 1, \ldots, d$ then shows that

$$s_{h,F} \lesssim s_{h,P}.$$

This inequality establishes that $s_{h,P}$ possesses the control properties (A3). □

Because it is much simpler to implement while preserving first-order convergence, the volume-based penalty (19) will be utilized for the numerical examples presented later in this work.

The abstract framework established in Section 3.3 possesses a useful modularity. To satisfy the global stability requirements, specifically the tangential gradient control (A3), it is sufficient for only *one* of the stabilization components to provide the necessary geometric bound. Additional stabilization terms can then be added to control other properties of the discrete extension. As long as these additional terms are symmetric, positive semi-definite (so they do not diminish the existing gradient control), and weakly consistent, their inclusion will preserve the first-order convergence of the overall method.

To illustrate this, we note that it is possible to augment the patch ghost penalty with other physically motivated terms. For example, in the numerical treatment of surface

PDEs, a stabilization form that controls the variation of the discrete solution in the normal direction can be employed. This is for instance done by Larson and Zahedi [16]. This specific normal gradient volume stabilization was first introduced and analysed in the context of CutFEM by Burman, Hansbo, Larson *et al.* [24].

Definition 3.9 (Normal Gradient Penalty)**.** The normal gradient volume penalty is defined component-wise over the active volume Ω_h as

$$s_{h,N}(u,v) = \sum_{i=1}^{d} \tau_N h^{-1} \langle \nu_h \cdot \nabla u_i, \nu_h \cdot \nabla v_i \rangle_{\Omega_h}, \tag{20}$$

where $\tau_N > 0$ is a chosen stabilization parameter. The fully stabilized bilinear form can then be constructed as the sum $s_h(u,v) = s_{h,P}(u,v) + s_{h,N}(u,v)$, where $s_{h,P}$ is the form defined in Definition 3.8.

Remark. In the penalty formulation (20), the discrete normal ν_h must be evaluated throughout the active volume Ω_h. To maintain a fully discrete and computable formulation, we extend the discrete normal via the closest point projection in each $T \in \mathcal{T}_h$ onto $K \in \mathcal{K}_h$. Let $p_h : \Omega_h \to \Gamma_h$ denote this projection. For any $x \in \Omega_h$, the extended normal is defined as $\nu_h^E = \nu_h \circ p_h$.

This extension preserves the $\mathcal{O}(h)$ geometric approximation error over the entire volume. By adding and subtracting $\nu(p(p_h(x)))$ and applying the triangle inequality, we obtain:

$$\left|\nu_h(p_h(x)) - \nu(p(x))\right| \leq \left|\nu_h(p_h(x)) - \nu(p(p_h(x)))\right| + \left|\nu(p(p_h(x))) - \nu(p(x))\right|.$$

The first term is bounded by $\mathcal{O}(h)$ due to the surface approximation property (8). For the second term, the C^2 regularity of Γ implies that the extended exact normal $\nu \circ p$ is Lipschitz continuous in $U_{\delta_0}(\Gamma)$. Because the distance $\left\|p_h(x) - x\right\|$ for $x \in \Omega_h$ is bounded by the thickness of the active mesh ($\mathcal{O}(h)$), the second term is also bounded by $\mathcal{O}(h)$. Therefore, the volume bound is established:

$$\left\|\nu_h^E - \nu \circ p\right\|_{L^\infty(\Omega_h)} \lesssim h.$$

In the subsequent analysis, ν_h within volume integrals denotes this specific discrete extension.

Proposition 3.3 (Suitability of the Combined Stabilization)**.** *The combined stabilization form $s_h(u,v) = s_{h,P}(u,v) + s_{h,N}(u,v)$ satisfies all abstract properties laid out in Section 3.3.*

Proof. Because $s_{h,N}(v,v) \geq 0$ for all $v \in W_h$, the combined form preserves the tangential gradient control provided by the volume-based ghost penalty, since $s_h(v,v) \geq s_{h,P}(v,v)$. Symmetry and positivity also follow from the L^2 inner product.

The h^{-1} scaling in (20) geometrically balances the dimensional difference between the surface measure and the active volume integration, offsetting the thickness of the active volume Ω_h, which scales as h.

To prove weak consistency, we evaluate the penalty on the Lagrange interpolant. By adding and subtracting the extended continuous normal $\nu \circ p$, we can expand the

directional derivative as $\nu_h \cdot \nabla \mathcal{I}_h H^e = \nu_h \cdot \nabla(\mathcal{I}_h H^e - H^e) + (\nu_h - \nu \circ p) \cdot \nabla H^e + (\nu \circ p) \cdot \nabla H^e$. Because the exact extended mean curvature $H^e = H \circ p$ is constant along the normal direction, the last term vanishes.

The remaining terms consist of the interpolation error and a geometric approximation error. According to the interpolation bounds in Definition 3.3, the L^2 norm of the gradient error is bounded by $h \left\| H^e \right\|_{H^2(\Omega_h)}$. For the geometric error, using the normal approximation property (8), we obtain the bound $\left\| \nu_h - \nu \circ p \right\|_{L^\infty(\Omega_h)} \left\| \nabla H^e \right\|_{L^2(\Omega_h)} \lesssim h \left\| H^e \right\|_{H^2(\Omega_h)}$.

The thickness of the active volume scales with h. Consequently, the volume norm scales with the surface norm as $\left\| H^e \right\|^2_{H^2(\Omega_h)} \lesssim h \left\| H \right\|^2_{H^2(\Gamma)}$. Applying the h^{-1} scaling factor to the combined bounds, we obtain

$$s_{h,N}(\mathcal{I}_h H^e, \mathcal{I}_h H^e) \lesssim h^{-1} \left(h^2 \left\| H^e \right\|^2_{H^2(\Omega_h)} \right) \lesssim h^{-1} h^2 \left(h \left\| H \right\|^2_{H^2(\Gamma)} \right) = h^2 \left\| H \right\|^2_{H^2(\Gamma)} .$$

Since the added normal penalty is weakly consistent and positive semi-definite, it integrates into the abstract framework without compromising the stability bounds of the volume-based ghost penalty. □

4 Extension to $\mathbb{Q}_1$ elements

In this chapter, we show that the *a priori* error bound still holds for curved discrete surfaces, provided that suitable stabilization is employed. Specifically, we establish that the formulation remains valid for $\mathbb{Q}_1$ elements on orthogonal background meshes. In Chapter 3, the assumption of a $\mathbb{P}_1$ finite element space on a piecewise planar surface was relied upon in only a single instance: to conclude that $\langle w_i, \Delta_K v_i \rangle_K = 0$ in (16). Because all other steps proceed identically, our objective is to bound this non-vanishing term.

4.1 Domain Discretization

Since $\Gamma \subset \mathbb{R}^d$ is compact, we can cover its tubular neighbourhood $U_{\delta_0}(\Gamma)$ with a background mesh $\mathcal{T}_{h,0}$. Let $\mathcal{T}_{h,0}$ be a quasi-uniform partition of shape-regular orthogonal hyper-rectangles T such that $\mathrm{Diam}(T) \lesssim h$.

We assume the existence of a continuous discrete surface $\Gamma_h \subset U_{\delta_0}(\Gamma)$ that approximates the exact surface Γ and satisfies the geometric approximation bounds (7) and (8). Furthermore, we assume that for every active element $T \in \mathcal{T}_h$, the surface segment $K = T \cap \Gamma_h$ is C^2 and partitions T into exactly two distinct subdomains.

The active mesh $\mathcal{T}_h$, the surface partition $\mathcal{K}_h$, the active volume Ω_h, and the internal face and edge sets ($\mathcal{F}_h$ and $\mathcal{E}_h$) are defined identically to the sets in Chapter 3. With the background mesh established, we introduce the finite element space

$$V_h = \left\{ v \in C^0(\Omega_h) : v|_T \in \mathbb{Q}_1(T) \quad \forall T \in \mathcal{T}_h \right\},$$

where $v|_T$ is the restriction of v to T and $\mathbb{Q}_1(T)$ is the space of multilinear polynomials defined on the physical element T. We denote the corresponding vector space as $W_h = [V_h]^d$.

Because the discrete surface Γ_h is curved within the elements, we must generalize our trace bounds to evaluate the discrete solution along these non-planar interfaces.

Lemma 4.1 (Scaled Trace Inequality)**.** *Let $T \in \mathcal{T}_h$ be a shape-regular active element, and let $K = T \cap \Gamma_h \in \mathcal{K}_h$. Provided the mesh parameter h is sufficiently small ($h \leq h_0$) such that the discrete surface is sufficiently resolved, the following estimate holds for any $v \in H^1(T)$:*

$$\|v\|^2_{L^2(K)} \lesssim h^{-1} \|v\|^2_{L^2(T)} + h \|\nabla v\|^2_{L^2(T)}. \tag{21}$$

The following argument is a generalization of the flat-interface proof of Lemma 4.2 by Hansbo, Hansbo and Larson [13], exploiting the local mesh resolution to account for the surface curvature.

Proof. The surface segment K divides T into two cells. Let T^* denote the cell chosen such that its outward-pointing unit normal $n_{\partial T^*} = \nu_h$ on K. The boundary ∂T^* then consists of K and a set of flat facets $F \subset \partial T$.

Choose an arbitrary point $x_c \in K$ and define the constant vector field $w = h\nu_h(x_c)$. Applying the divergence theorem to the vector field $v^2 w$ over T^* yields

$$\int_{T^*} \nabla \cdot (v^2 w)\, dx = \int_K v^2 (w \cdot \nu_h)\, ds + \int_F v^2 (w \cdot n_{\partial T^*})\, ds.$$

Because $\nabla \cdot w = 0$, we have $\nabla \cdot (v^2 w) = 2v \nabla v \cdot w$. Rearranging the terms gives

$$\int_K v^2 (w \cdot \nu_h)\, ds = 2 \int_{T^*} v \nabla v \cdot w\, dx - \int_F v^2 (w \cdot n_{\partial T^*})\, ds. \tag{22}$$

Because Γ_h is sufficiently resolved by the mesh ($h \leq h_0$), the variation of the discrete normal ν_h within T is small. Consequently, the local constant approximation satisfies $\nu_h(x_c) \cdot \nu_h(x) \gtrsim 1$ for all $x \in K$, which implies $w \cdot \nu_h \gtrsim h$. Thus,

$$\int_K v^2 (w \cdot \nu_h)\, ds \gtrsim h \left\| v \right\|^2_{L^2(K)}.$$

For the volume integral, Cauchy–Schwarz and Young's inequalities yield

$$2 \int_{T^*} v \nabla v \cdot w\, dx \leq 2h \left\| v \right\|_{L^2(T)} \left\| \nabla v \right\|_{L^2(T)} \leq h \left(h^{-1} \left\| v \right\|^2_{L^2(T)} + h \left\| \nabla v \right\|^2_{L^2(T)} \right).$$

For the boundary integral over F, since $|w \cdot n_{\partial T^*}| \leq h$ and $F \subset \partial T$, we apply the standard polynomial trace inequality (Ern and Guermond [23, Lemma 12.15]) to obtain

$$\left| \int_F v^2 (w \cdot n_{\partial T^*})\, ds \right| \leq h \left\| v \right\|^2_{L^2(F)} \leq h \left\| v \right\|^2_{L^2(\partial T)} \lesssim h \left(h^{-1} \left\| v \right\|^2_{L^2(T)} + h \left\| \nabla v \right\|^2_{L^2(T)} \right).$$

Substituting these bounds into (22) and dividing by h yields the desired estimate. □

4.2 The Surface Laplacian Control Assumption

To bound the non-vanishing residual and recover the first-order convergence, we extend the abstract stabilization framework from Chapter 3 by adopting an additional assumption. For all $v \in W_h$, we assume the stabilization form provides control over the surface Laplacian:

(A4) **Surface Laplacian Control:** The stabilization must control the non-vanishing discrete curvature introduced by the multilinear element mapping. We assume the bound:

$$h^2 \left\| \Delta_{\Gamma_h} v \right\|^2_{L^2(\Gamma_h)} \lesssim \left\| v \right\|^2_{L^2(\Gamma_h)} + \left\| v \right\|^2_{s_h} \qquad \forall v \in W_h. \tag{A4}$$

With the abstract framework complete, we can immediately establish the first-order convergence of the discrete formulation.

Theorem 4.2. *Let $H \in [H^2(\Gamma)]^d$ be the exact mean curvature vector of the C^2 continuous surface Γ, and let $H^e = H \circ p$ be its extension to the active volume Ω_h. Let $H_h \in W_h$ be the solution to the discrete formulation* (9) *using the $\mathbb{Q}_1$ finite element space. Assume that the discrete surface Γ_h satisfies the geometric approximation bounds* (7) *and* (8)*. Assume further that the stabilization form s_h satisfies the abstract properties (A1) through (A4). Then, the error satisfies the first-order bound:*

$$\|H^e - H_h\|^2_{L^2(\Gamma_h)} + \|H^e - H_h\|^2_{s_h} \lesssim h^2. \tag{23}$$

Proof. The proof follows the linear case in Theorem 3.2. Setting $v = \pi_h H^e - H_h$, the argument diverges when evaluating the consistency error in (16). Writing $w = x^e_\Gamma - x_{\Gamma_h}$, partial integration over the components leaves the residual sum $\sum_{i=1}^d \sum_{K \in \mathcal{K}_h} \langle w_i, \Delta_K v_i \rangle_K$. Because the $\mathbb{Q}_1$ shape functions are not linear, the surface Laplacian does not vanish.

Summing over the vector components, we bound this residual directly using Cauchy–Schwarz and the geometric approximation bound (7):

$$\begin{aligned}\left| \sum_{K \in \mathcal{K}_h} \langle w, \Delta_K v \rangle_K \right| &= \left| \langle w, \Delta_{\Gamma_h} v \rangle_{\Gamma_h} \right| \\ &\leq \|w\|_{L^2(\Gamma_h)} \left\| \Delta_{\Gamma_h} v \right\|_{L^2(\Gamma_h)} \\ &\lesssim h \left(h \left\| \Delta_{\Gamma_h} v \right\|_{L^2(\Gamma_h)} \right).\end{aligned}$$

Applying the surface Laplacian control (A4) yields:

$$h \left(h \left\| \Delta_{\Gamma_h} v \right\|_{L^2(\Gamma_h)} \right) \lesssim h \sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}}.$$

This bound exactly replaces the vanishing term from the linear analysis. Substituting it into the error expansion yields the first-order estimate. □

With the *a priori* error estimate established under the abstract framework, the remainder of this chapter is dedicated to verifying that our specific choices of stabilization satisfy the requisite control properties (A1) through (A4).

4.3 Stabilization Forms and Volume Control

With the abstract error estimate established, we must define the specific stabilization forms used to enforce the requisite control properties (A1) through (A4). We consider two distinct normal-gradient penalties: the volume-based penalty $s_{h,N}$ established in Definition 3.9, and a surface-based alternative.

For surface PDEs, Zahedi [25] introduced a stabilization form evaluated directly on the discrete surface. We generalize this form to the vector-valued problem in $\mathbb{R}^d$.

Definition 4.1 (Surface-Based Normal Penalty)**.** Let $\tau_{i,n} > 0$ be stabilization parameters. For any $u, v \in W_h$, the surface-based normal penalty form is defined component-wise as

$$s_{h,n}(u,v) = \sum_{i=1}^{d} \sum_{j=1}^{d} \tau_{j,n} h^{2(j-1)} \left\langle D^j_{\nu_h} u_i, D^j_{\nu_h} v_i \right\rangle_{\Gamma_h}, \tag{24}$$

where $D^j_{\nu_h} v_i$ denotes the j-th order directional derivative of the i-th component of v in the direction of the discrete normal ν_h.

Before verifying that $s_{h,N}$ and $s_{h,n}$ satisfy the abstract assumptions, we establish a fundamental lemma that bounds the active volume norm by the discrete surface norm and the chosen stabilization penalty. Because the gradients of $\mathbb{Q}_1$ functions are not element-wise constant, this volume control is a necessary component for the subsequent proofs in $\mathbb{R}^d$.

Lemma 4.3 (Volume Control)**.** *For any $v \in V_h$, the active volume norm is bounded by the discrete surface norm and the combined stabilization form such that*

$$h^{-1} \|v\|^2_{L^2(\Omega_h)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + h^2 s_h(v,v), \tag{25}$$

where $s_h(v,v) = s_{h,P}(v,v) + s_{h,\perp}(v,v)$, and $s_{h,\perp} \in \{s_{h,N}, s_{h,n}\}$.

Proof. Let $T \in \mathcal{T}_h$ be an active element. Inspired by Demlow and Olshanskii [26, Proposition 4.2], we first establish the existence of an active element T' within a bounded distance from T, satisfying the lower bound $\|1\|^2_{L^2(K')} \geq Ch^{d-1}$ for $K' = T' \cap \Gamma_h$ and a uniform constant $C > 0$.

Let $x_0 \in T \cap \Gamma_h$. Because Γ_h is a closed continuous surface, we define a contiguous patch $S \subset \Gamma_h$ containing x_0, with diameter ch for a uniform constant $c > 0$, such that its $(d-1)$-dimensional measure satisfies $\|1\|^2_{L^2(S)} \sim h^{d-1}$. Let $\omega_S = \{T^* \in \mathcal{T}_h : T^* \cap S \neq \emptyset\}$ be the set of elements intersected by S.

Because the diameter of S is ch, the entire patch is contained within a ball $B_{ch}(x_0)$ of radius ch. By the shape regularity of the mesh, the minimum inscribed diameter of any element scales with h. Consequently, there is a uniform upper bound $N_{\max}$, independent of h, on the number of elements intersecting $B_{ch}(x_0)$. Thus, the cardinality of ω_S satisfies $|\omega_S| \leq N_{\max}$.

Since the total measure of S scales as h^{d-1}, we have

$$\sum_{T^* \in \omega_S} \|1\|^2_{L^2(T^* \cap \Gamma_h)} \gtrsim h^{d-1}.$$

By the pigeonhole principle, there exists at least one element $T' \in \omega_S$ satisfying

$$\|1\|^2_{L^2(T' \cap \Gamma_h)} \gtrsim \frac{1}{N_{\max}} h^{d-1}.$$

This establishes the required lower bound with the uniform constant $C = 1/N_{\max}$.

We construct a subdomain $T^* \subset T'$ by sweeping each point $x_0 \in K'$ inward along the normal direction ν_h up to a variable height $h^*(x_0)$. To ensure T^* remains entirely

contained within the element boundaries of T', we choose $h^*(x_0)$ such that it vanishes on the boundary $\partial K'$ and reaches a maximum height of ch (for a uniform constant $c > 0$) in the interior, effectively forming a topological cone with base K'. Because the curvature of the exact surface is bounded, the Jacobian determinant of this normal parametrization satisfies $J \sim 1$ for sufficiently small h. Consequently, the d-dimensional measure of this cone satisfies $\|1\|^2_{L^2(T^*)} \sim h \|1\|^2_{L^2(K')} \sim h^d$. Because v is a polynomial on $T' \supset T^*$, we have the norm equivalence $\|v\|^2_{L^2(T')} \sim \|v\|^2_{L^2(T^*)}$.

For any point $x \in T^*$, we write $x = x_0 + s\nu_h(x_0)$, where $x_0 \in K'$ and $s \in [0, h^*(x_0)]$. Because $v|_{T'}$ is a polynomial of degree d, its Taylor expansion along the normal yields:

$$v(x) = v(x_0) + \sum_{j=1}^{d} \frac{1}{j!} s^j D^j_{\nu_h} v(x_0).$$

This yields the pointwise bound

$$\left|v(x)\right|^2 \lesssim \left|v(x_0)\right|^2 + \sum_{j=1}^{d} h^{2j} \left|D^j_{\nu_h} v(x_0)\right|^2.$$

Integrating the 1D normal parameter over $s \in [0, h^*(x_0)]$ yields a factor of $h^*(x_0) \lesssim h$. Subsequent integration over the base $x_0 \in K'$ gives:

$$\|v\|^2_{L^2(T^*)} \lesssim h \|v\|^2_{L^2(K')} + \sum_{j=1}^{d} h^{2j+1} \left\|D^j_{\nu_h} v\right\|^2_{L^2(K')}.$$

Multiplying by h^{-1} and applying the norm equivalence between T^* and T' yields:

$$h^{-1} \|v\|^2_{L^2(T')} \lesssim \|v\|^2_{L^2(K')} + \sum_{j=1}^{d} h^{2j} \left\|D^j_{\nu_h} v\right\|^2_{L^2(K')}.$$

For two adjacent elements T_1, T_2, we have $\|v\|^2_{L^2(T_1)} \lesssim \|v\|^2_{L^2(T_2)} + \left\|[v]\right\|^2_{L^2(T_1 \cup T_2)}$ via the triangle inequality and polynomial norm equivalence. Chaining elements across the patch ω_S provides the bound

$$\begin{aligned}\|v\|^2_{L^2(T)} &\lesssim \|v\|^2_{L^2(T')} + \sum_{F \in \omega_S} \left\|[v]\right\|^2_{\omega_F} \\ &\lesssim h \left(\|v\|^2_{L^2(K')} + \sum_{j=1}^{d} h^{2j} \left\|D^j_{\nu_h} v\right\|^2_{L^2(K')} \right) + \sum_{F \in \omega_S} \left\|[v]\right\|^2_{\omega_F}.\end{aligned}$$

Summing this bound over all active elements $T \in \mathcal{T}_h$ yields

$$\|v\|^2_{L^2(\Omega_h)} \lesssim \sum_{K' \in \mathcal{K}_h} \left(h \|v\|^2_{L^2(K')} + \sum_{j=1}^{d} h^{2j+1} \left\|D^j_{\nu_h} v\right\|^2_{L^2(K')} \right) + \sum_{F \in \mathcal{F}_h} \left\|[v]\right\|^2_{L^2(\omega_F)}.$$

The first term evaluates to $h \|v\|^2_{L^2(\Gamma_h)}$. The final term is bounded by $h^3 s_{h,P}(v, v)$. We bound the derivative terms using $s_{h,\perp}$.

If $s_{h,\perp} = s_{h,n}$, the intermediate terms match the definition of the surface normal penalty (24), establishing the bound $h^3 s_{h,n}(v,v)$.

If $s_{h,\perp} = s_{h,N}$, we evaluate the directional derivatives by introducing the local constant normal approximation $\bar{\nu}$ satisfying $\|\nu_h - \bar{\nu}\|_{L^\infty(T')} \lesssim h$. Expanding the derivatives for each $j \in \{1, \dots, d\}$ yields the bounds

$$\left\| D^j_{\nu_h} v \right\|^2_{L^2(K')} \lesssim \left\| D^j_{\bar{\nu}} v \right\|^2_{L^2(K')} + h^2 \left\| D^j v \right\|^2_{L^2(K')}.$$

Because v is a polynomial on T' and $\bar{\nu}$ is constant, the directional derivatives $D^j_{\bar{\nu}} v$ are also polynomials. Applying the scaled trace inequality (Lemma 4.1) maps the surface norms to the volume T', and the standard inverse inequality (Ern and Guermond [23, Lemma 12.1]) bounds the higher-order spatial derivatives:

$$\begin{aligned}\sum_{j=1}^{d} h^{2j+1} \left\| D^j_{\nu_h} v \right\|^2_{L^2(K')} &\lesssim \sum_{j=1}^{d} \left(h^{2j+1} \left\| D^j_{\bar{\nu}} v \right\|^2_{L^2(K')} + h^{2j+3} \left\| D^j v \right\|^2_{L^2(K')} \right) \\ &\lesssim \sum_{j=1}^{d} \left(h^{2j} \left\| D^j_{\bar{\nu}} v \right\|^2_{L^2(T')} + h^{2j+2} \left\| D^j v \right\|^2_{L^2(T')} \right) + h^2 \|v\|^2_{L^2(T')}.\end{aligned}$$

For the constant-direction derivatives, we observe the identity $D^j_{\bar{\nu}} v = \bar{\nu} \cdot \nabla (D^{j-1}_{\bar{\nu}} v)$. Applying the inverse inequality on this relation yields

$$h^2 \left\| D^j_{\bar{\nu}} v \right\|^2_{L^2(T')} \lesssim \left\| D^{j-1}_{\bar{\nu}} v \right\|^2_{L^2(T')}.$$

Applying this argument recursively down to $j = 1$, together with Lemma 12.1 in Ern and Guermond [23], simplifies the dominant terms:

$$\begin{aligned}\sum_{j=1}^{d} h^{2j+1} \left\| D^j_{\nu_h} v \right\|^2_{L^2(K')} &\lesssim h^2 \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T')} + h^2 \|v\|^2_{L^2(T')} \\ &\lesssim h^2 \|\nu_h \cdot \nabla v\|^2_{L^2(T')} + h^4 \|\nabla v\|^2_{L^2(T')} + h^2 \|v\|^2_{L^2(T')} \\ &\lesssim h^2 \|\nu_h \cdot \nabla v\|^2_{L^2(T')} + h^2 \|v\|^2_{L^2(T')}.\end{aligned}$$

Summing over the active elements, the first term matches the definition of $s_{h,N}$ and establishes the bound $h^3 s_{h,N}(v,v)$. The second term sums to $h^2 \|v\|^2_{L^2(\Omega_h)}$. Thus, we conclude

$$\|v\|^2_{L^2(\Omega_h)} \lesssim h \|v\|^2_{L^2(\Gamma_h)} + h^3 s_{h,N}(v,v) + h^2 \|v\|^2_{L^2(\Omega_h)} + h^3 s_{h,P}(v,v).$$

The $h^2 \|v\|^2_{L^2(\Omega_h)}$ term is absorbed by the left-hand side.

In either configuration, dividing the sum by h yields the conclusion:

$$h^{-1} \|v\|^2_{L^2(\Omega_h)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + h^2 s_h(v,v). \qquad \square$$

4.4 Verification of Weak Consistency and Tangential Control

In this section, we verify that the combined stabilization $s_h = s_{h,P} + s_{h,\perp}$ satisfies the abstract control properties (A1) through (A3). By definition, (A1) is trivially satisfied. The weak consistency of the volume penalty $s_{h,N}$ was established in Proposition 3.3, so we proceed to establish the weak consistency of the surface penalty $s_{h,n}$.

Because $s_{h,n}$ explicitly penalizes higher-order normal derivatives, its analysis requires control over the second derivatives of the interpolant. Therefore, when utilizing the surface penalty, we strengthen the abstract interpolation assumption from Chapter 3 to include the extended volume approximation bound:

$$\|H^e - \pi_h H^e\|_{L^2(\Omega_h)} + h \left\|\nabla(H^e - \pi_h H^e)\right\|_{L^2(\Omega_h)} + h^2 \left\|D^2(H^e - \pi_h H^e)\right\|_{L^2(\Omega_h)} \lesssim h^2 \|H^e\|_{H^2(\Omega_h)} . \tag{26}$$

This directly guarantees the H^2-stability of the interpolant, which is necessary to bound the high-order volume traces in the subsequent proof.

Lemma 4.4. *Assume the interpolation operator π_h satisfies the extended volume approximation bound* (26)*. Assuming further that Γ is sufficiently regular such that $H^e \in [H^{d+1}(\Omega_h)]^d$, the surface penalty evaluated on the exact solution and its interpolant satisfies the bound*

$$s_{h,n}(H^e, H^e) + s_{h,n}(\pi_h H^e, \pi_h H^e) \lesssim h^2 \|H\|^2_{H^{d+1}(\Gamma)} . \tag{27}$$

The following argument is adapted from the more general framework presented by Larson and Zahedi [16].

Proof. By the triangle inequality, it is sufficient to bound $s_{h,n}(H^e, H^e)$ and the interpolation error $s_{h,n}(H^e - \pi_h H^e, H^e - \pi_h H^e)$ separately.

By definition, the extended solution $H^e = H \circ p$ is constant in the direction of the exact continuous normal ν. Hence, the normal derivative $D_\nu H_i^e = \nu \cdot \nabla H_i^e = 0$. Since ν is also constant along the normal direction, we have $D_\nu^j H_i^e = 0$ for all $j \geq 1$ and components $1 \leq i \leq d$.

This allows us to evaluate the discrete normal derivatives as the deviation from the exact normal derivatives. For any $j \in \{1, \dots, d\}$, the difference in the directional derivatives is bounded by the deviation of the normal vectors. We have the pointwise bound

$$\begin{aligned} \left| D^j_{\nu_h} H_i^e \right| &= \left| D^j_{\nu_h} H_i^e - D^j_{\nu} H_i^e \right| \\ &\lesssim |\nu_h - \nu| \left| D^j H_i^e \right| , \end{aligned}$$

where D^j denotes the full spatial derivative tensor of order j. Substituting this pointwise bound into the penalty term, summing over the components, and applying the geometric

approximation (8) alongside Lemma 3.5 yields

$$\begin{aligned} h^{2(j-1)} \sum_{i=1}^{d} \left\| D^j_{\nu_h} H^e_i \right\|^2_{L^2(\Gamma_h)} &\lesssim h^{2(j-1)} \left(h^2 \sum_{i=1}^{d} \left\| D^j H^e_i \right\|^2_{L^2(\Gamma_h)} \right) \\ &= h^{2j} \left\| D^j H^e \right\|^2_{L^2(\Gamma_h)} \\ &\sim h^{2j} \left\| (D^j H^e)^l \right\|^2_{L^2(\Gamma)} . \end{aligned}$$

Because $h^{2j} \leq h^2$, summing over j yields the bound $s_{h,n}(H^e, H^e) \lesssim h^2 \|H\|^2_{H^{d+1}(\Gamma)}$.

For the interpolant error, we define $w = H^e - \pi_h H^e$. Because the discrete normal vector has unit length ($|\nu_h| = 1$), the directional derivative is bounded pointwise by the full spatial derivative tensor, $\left| D^j_{\nu_h} w_i \right| \leq \left| D^j w_i \right|$. Applying Lemma 4.1 maps the surface penalty to the active volume:

$$\begin{aligned} h^{2j-2} \sum_{i=1}^{d} \left\| D^j_{\nu_h} w_i \right\|^2_{L^2(\Gamma_h)} &\lesssim h^{2j-2} \left\| D^j w \right\|^2_{L^2(\Gamma_h)} \\ &\lesssim h^{2j-3} \left\| D^j w \right\|^2_{L^2(\Omega_h)} + h^{2j-1} \left\| D^{j+1} w \right\|^2_{L^2(\Omega_h)} . \end{aligned}$$

For $j = 1$, we bound the resulting terms $h^{-1} \|Dw\|^2_{L^2(\Omega_h)}$ and $h \left\| D^2 w \right\|^2_{L^2(\Omega_h)}$ separately. From the extended interpolation bound (26), we directly obtain the estimates

$$\|Dw\|^2_{L^2(\Omega_h)} \lesssim h^2 \|H^e\|^2_{H^2(\Omega_h)} \quad \text{and} \quad \left\| D^2 w \right\|^2_{L^2(\Omega_h)} \lesssim \|H^e\|^2_{H^2(\Omega_h)} .$$

Substituting these bounds into the volume trace estimate gives

$$\begin{aligned} h^{-1} \|Dw\|^2_{L^2(\Omega_h)} + h \left\| D^2 w \right\|^2_{L^2(\Omega_h)} &\lesssim h \|H^e\|^2_{H^2(\Omega_h)} + h \|H^e\|^2_{H^2(\Omega_h)} \\ &\lesssim h \|H^e\|^2_{H^2(\Omega_h)} . \end{aligned}$$

For $j \geq 2$, we bound the higher-order terms using the triangle inequality $D^j w = D^j H^e - D^j \pi_h H^e$. Applying the triangle inequality and the interpolation bound (26) shows:

$$\left\| D^2 \pi_h H^e \right\|^2_{L^2(\Omega_h)} \lesssim \left\| D^2 H^e \right\|^2_{L^2(\Omega_h)} + \left\| D^2 w \right\|^2_{L^2(\Omega_h)} \lesssim \|H^e\|^2_{H^2(\Omega_h)} .$$

Using Lemma 12.1 in Ern and Guermond [23], the j-th derivative of the polynomial interpolant is mapped back to the second derivative element-wise:

$$\left\| D^j \pi_h H^e \right\|^2_{L^2(T)} \lesssim h^{4-2j} \left\| D^2 \pi_h H^e \right\|^2_{L^2(T)} .$$

Summing over the active elements and substituting the interpolant stability bound, the first term from the volume trace scales as:

$$\begin{aligned} h^{2j-3} \left\| D^j w \right\|^2_{L^2(\Omega_h)} &\lesssim h^{2j-3} \left\| D^j H^e \right\|^2_{L^2(\Omega_h)} + h^{2j-3} \left(h^{4-2j} \left\| D^2 \pi_h H^e \right\|^2_{L^2(\Omega_h)} \right) \\ &\lesssim h^{2j-3} \left\| D^j H^e \right\|^2_{L^2(\Omega_h)} + h \left\| D^2 \pi_h H^e \right\|^2_{L^2(\Omega_h)} \\ &\lesssim h \|H^e\|^2_{H^{d+1}(\Omega_h)} . \end{aligned}$$

An identical application of the inverse inequality resolves the $(j+1)$-th derivative term:

$$\begin{aligned} h^{2j-1}\left\|D^{j+1}w\right\|^2_{L^2(\Omega_h)} &\lesssim h^{2j-1}\left\|D^{j+1}H^e\right\|^2_{L^2(\Omega_h)} + h^{2j-1}\left(h^{4-2(j+1)}\left\|D^2\pi_h H^e\right\|^2_{L^2(\Omega_h)}\right) \\ &= h^{2j-1}\left\|D^{j+1}H^e\right\|^2_{L^2(\Omega_h)} + h\left\|D^2\pi_h H^e\right\|^2_{L^2(\Omega_h)} \\ &\lesssim h\left\|H^e\right\|^2_{H^{d+1}(\Omega_h)}. \end{aligned}$$

Because H^e is constant along the normal direction, the active volume integral scales geometrically with the surface integral:

$$\left\|H^e\right\|^2_{H^{d+1}(\Omega_h)} \lesssim h\left\|H\right\|^2_{H^{d+1}(\Gamma)}.$$

Substituting this geometric bound back into the combined volume trace bounds yields:

$$h^{2j-2}\sum_{i=1}^{d}\left\|D^j_{\nu_h}w_i\right\|^2_{L^2(\Gamma_h)} \lesssim h\left(h\left\|H\right\|^2_{H^{d+1}(\Gamma)}\right) = h^2\left\|H\right\|^2_{H^{d+1}(\Gamma)}.$$

Summing over all $1 \le j \le d$ establishes the bound $s_{h,n}(w,w) \lesssim h^2\left\|H\right\|^2_{H^{d+1}(\Gamma)}$, which concludes the proof. □

Remark. We note that the consistency bound (27) established for the surface normal penalty requires higher regularity (H^{d+1}) than the abstract consistency bound (A2) established in Chapter 3, which requires only H^2 regularity. However, because both bounds yield the requisite $\mathcal{O}(h^2)$ scaling, (27) fulfils the exact same functional role in the *a priori* error analysis without requiring modifications to the core abstract framework.

With (A2), or functional equivalents, established, we move on to verifying that (A3) holds for both stabilization combinations.

Proposition 4.1. *For $v \in W_h$, (A3) is satisfied:*

$$h^2\left\|\nabla_{\Gamma_h}v\right\|^2_{L^2(\Gamma_h)} + \sum_{K\in\mathcal{K}_h} h\left\|v\right\|^2_{L^2(\partial K)} + \sum_{E\in\mathcal{E}_h} h\left\|[t_E\cdot\nabla_{\Gamma_h}v]\right\|^2_{L^2(E)} \lesssim \left\|v\right\|^2_{L^2(\Gamma_h)} + s_h(v,v), \tag{28}$$

where $s_h(v,v) = s_{h,P}(v,v) + s_{h,\perp}(v,v)$.

Proof. We bound each term independently.

Term 1: Because the tangential gradient is defined via a projection, we have the pointwise bound $\left|\nabla_{\Gamma_h}v\right| \le |\nabla v|$. Applying Lemma 4.1 followed by Lemma 12.1 in Ern and Guermond [23] yields:

$$\left\|\nabla_{\Gamma_h}v\right\|^2_{L^2(K)} \le \left\|\nabla v\right\|^2_{L^2(K)} \lesssim h^{-1}\left\|\nabla v\right\|^2_{L^2(T)} \lesssim h^{-3}\left\|v\right\|^2_{L^2(T)}.$$

Multiplying by h^2 and summing over $K \in \mathcal{K}_h$ yields $h^2\left\|\nabla_{\Gamma_h}v\right\|^2_{L^2(\Gamma_h)} \lesssim h^{-1}\left\|v\right\|^2_{L^2(\Omega_h)}$. Applying Lemma 4.3 yields:

$$h^2\left\|\nabla_{\Gamma_h}v\right\|^2_{L^2(\Gamma_h)} \lesssim \left\|v\right\|^2_{L^2(\Gamma_h)} + h^2 s_h(v,v).$$

Term 2: The boundary ∂K consists of $(d-2)$-dimensional segments $E = F \cap \Gamma_h$ lying on the $(d-1)$-dimensional faces $F \in \partial T$. By restricting our analysis to the tangent space of the shape-regular face F, the cut E forms a continuous interface with a well-defined unit co-normal strictly within F. Evaluating the trace directly in this $(d-1)$-dimensional domain, we apply Lemma 4.1 alongside Lemma 12.1 in Ern and Guermond [23] to obtain the bound:

$$\|v\|^2_{L^2(E)} \lesssim h^{-1} \|v\|^2_{L^2(F)} + h \|\nabla_F v\|^2_{L^2(F)} \lesssim h^{-1} \|v\|^2_{L^2(F)}.$$

Applying Lemma 12.8 in Ern and Guermond [23] maps the norm from the face F to the volume T, yielding $\|v\|^2_{L^2(F)} \lesssim h^{-1} \|v\|^2_{L^2(T)}$. Chaining these inequalities establishes the boundary bound $\|v\|^2_{L^2(\partial K)} \lesssim h^{-2} \|v\|^2_{L^2(T)}$. Multiplying by h and summing over all discrete surface elements $K \in \mathcal{K}_h$ gives:

$$\sum_{K \in \mathcal{K}_h} h \|v\|^2_{L^2(\partial K)} \lesssim h^{-1} \|v\|^2_{L^2(\Omega_h)}.$$

Applying Lemma 4.3 bounds this final volumetric term by $\|v\|^2_{L^2(\Gamma_h)} + h^2 s_h(v,v)$.

Term 3: Let $E \in \mathcal{E}_h$ be the intersection of two adjacent surface elements K_1 and K_2, lying on the internal mesh face $F \in \mathcal{F}_h$ shared by volume elements T_1 and T_2. Defining the average co-normal vector $t^* = \frac{1}{2}(t_{E,K_1} - t_{E,K_2})$, the co-normal derivative jump decomposes as:

$$[t_E \cdot \nabla_{\Gamma_h} v] = (t_{E,K_1} - t^*) \cdot \nabla v_1 + (t_{E,K_2} + t^*) \cdot \nabla v_2 + t^* \cdot [\nabla v].$$

By the geometric approximation of the discrete surface (8), the deviation between the opposing co-normals satisfies $\left|t_{E,K_i} \pm t^*\right| \lesssim h$. Applying the triangle inequality yields:

$$h \left\|[t_E \cdot \nabla_{\Gamma_h} v]\right\|^2_{L^2(E)} \lesssim h^3 \|\nabla v_1\|^2_{L^2(E)} + h^3 \|\nabla v_2\|^2_{L^2(E)} + h \left\|t^* \cdot [\nabla v]\right\|^2_{L^2(E)}.$$

Evaluating the $(d-2)$-dimensional cut $E = F \cap \Gamma_h$ within the tangent space of F as in Term 2, we apply Lemma 4.1 and Lemma 12.1 in Ern and Guermond [23] to obtain the bound $\|\cdot\|^2_{L^2(E)} \lesssim h^{-1} \|\cdot\|^2_{L^2(F)}$. Furthermore, because the polynomial jump $[v]$ is identically zero across the continuous finite element face F, the gradient jump strictly aligns with the normal vector, satisfying $[\nabla v] = [\nu_F \cdot \nabla v]\nu_F$. Using these face identities and applying Lemma 12.8 in Ern and Guermond [23] maps the first two terms from the face to their respective volume elements:

$$h \left\|[t_E \cdot \nabla_{\Gamma_h} v]\right\|^2_{L^2(E)} \lesssim h \|\nabla v\|^2_{L^2(T_1)} + h \|\nabla v\|^2_{L^2(T_2)} + \left\|[\nu_F \cdot \nabla v]\right\|^2_{L^2(F)}.$$

Applying Lemma 12.1 in Ern and Guermond [23] to the first two terms and summing over all $E \in \mathcal{E}_h$ yields:

$$\sum_{E \in \mathcal{E}_h} h \left\|[t_E \cdot \nabla_{\Gamma_h} v]\right\|^2_{L^2(E)} \lesssim h^{-1} \|v\|^2_{L^2(\Omega_h)} + \sum_{F \in \mathcal{F}_h} \left\|[\nu_F \cdot \nabla v]\right\|^2_{L^2(F)}.$$

The final term matches the definition of the face penalty $s_{h,F}(v,v)$, which Proposition 3.2 bounds above by $s_{h,P}(v,v)$. Applying Lemma 4.3 to the volumetric term yields the final bound $\sum_{E \in \mathcal{E}_h} h \left\|[t_E \cdot \nabla_{\Gamma_h} v]\right\|^2_{L^2(E)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + s_h(v,v)$.

□

4.5 Verification of the Surface Laplacian Control

With assumptions (A1) through (A3) established, we only need to verify that (A4) holds for the method to satisfy our *a priori* bound from Theorem 4.2. To do this, we exploit a specific expansion of the Laplace–Beltrami operator on each surface element:

$$\Delta_K v = \Delta v - D^2_{\nu_h} v - H_K \cdot \nabla v, \tag{29}$$

where H_K is the mean curvature vector of the surface element K. In light of this expansion, establishing that $s_{h,n}$ satisfies (A4) becomes trivial on $\mathbb{Q}_1$ finite elements.

Proposition 4.2. *Assume the background mesh $\mathcal{T}_h$ consists of orthogonal hyper-rectangles in $\mathbb{R}^d$, and the continuous discrete surface Γ_h satisfies the geometric bound $\|H_K\|_{L^\infty(\Gamma_h)} \lesssim 1$. Then the surface-based normal penalty $s_{h,n}$ satisfies the surface Laplacian control bound (A4).*

Proof. On an orthogonal hyper-rectangle $T \in \mathcal{T}_h$, the $\mathbb{Q}_1$ shape functions are strictly multilinear, yielding the full volume Laplacian $\Delta v = 0$. Furthermore, because H_K is the mean curvature vector, it is aligned with the discrete normal vector, satisfying $H_K = \pm |H_K| \, \nu_h$. Substituting these identities into the expansion (29) yields:

$$\Delta_K v = -D^2_{\nu_h} v \mp |H_K| \, D_{\nu_h} v.$$

Applying the triangle inequality, squaring, and integrating over the discrete surface gives:

$$\left\| \Delta_{\Gamma_h} v \right\|^2_{L^2(\Gamma_h)} \lesssim \left\| D^2_{\nu_h} v \right\|^2_{L^2(\Gamma_h)} + \|H_K\|^2_{L^\infty(\Gamma_h)} \left\| D_{\nu_h} v \right\|^2_{L^2(\Gamma_h)}.$$

Multiplying by h^2 and applying the geometric assumption $\|H_K\|_{L^\infty(\Gamma_h)} \lesssim 1$ yields:

$$h^2 \left\| \Delta_{\Gamma_h} v \right\|^2_{L^2(\Gamma_h)} \lesssim h^2 \left\| D^2_{\nu_h} v \right\|^2_{L^2(\Gamma_h)} + h^2 \left\| D_{\nu_h} v \right\|^2_{L^2(\Gamma_h)}.$$

Because $h \leq 1$, the second term is bounded by $h^0 \left\| D_{\nu_h} v \right\|^2_{L^2(\Gamma_h)}$. These two terms correspond exactly to the $j = 2$ and $j = 1$ components of the surface penalty $s_{h,n}(v, v)$. Thus, we conclude:

$$h^2 \left\| \Delta_{\Gamma_h} v \right\|^2_{L^2(\Gamma_h)} \lesssim s_{h,n}(v, v) \leq \|v\|^2_{L^2(\Gamma_h)} + s_h(v, v). \qquad \square$$

We now also establish the suitability of $s_{h,N}$.

Proposition 4.3. *Assume the background mesh $\mathcal{T}_h$ consists of orthogonal hyper-rectangles in $\mathbb{R}^d$, and the continuous discrete surface Γ_h satisfies the geometric bound $\|H_K\|_{L^\infty(\Gamma_h)} \lesssim 1$. Then the volume-based normal gradient penalty $s_{h,N}$ satisfies the surface Laplacian control bound (A4).*

Proof. On an orthogonal hyper-rectangle $T \in \mathcal{T}_h$, the $\mathbb{Q}_1$ shape functions are strictly multilinear, yielding $\Delta v = 0$. Furthermore, the mean curvature vector H_K is aligned with the discrete normal, satisfying $H_K = \pm |H_K| \, \nu_h$. By the expansion (29), the discrete surface Laplacian simplifies to:

$$\Delta_K v = -D^2_{\nu_h} v \mp |H_K| \, D_{\nu_h} v.$$

Applying the triangle inequality, the geometric bound $\|H_K\|_{L^\infty(\Gamma_h)} \lesssim 1$, and multiplying by h^2 yields the element-wise bound:

$$h^2 \|\Delta_K v\|^2_{L^2(K)} \lesssim h^2 \left\|D^2_{\nu_h} v\right\|^2_{L^2(K)} + h^2 \left\|D_{\nu_h} v\right\|^2_{L^2(K)}.$$

To bound these normal derivatives using the volume penalty $s_{h,N}$, we introduce the local constant normal approximation $\bar{\nu}$ satisfying $\|\nu_h - \bar{\nu}\|_{L^\infty(T)} \lesssim h$. Expanding the directional derivatives via the triangle inequality for each j yields $\left\|D^j_{\nu_h} v\right\|^2_{L^2(K)} \lesssim \left\|D^j_{\bar{\nu}} v\right\|^2_{L^2(K)} + h^2 \left\|D^j v\right\|^2_{L^2(K)}$.

Applying Lemma 4.1 maps the constant-direction polynomials and the perturbations from the surface to the volume. We also apply Lemma 12.1 in Ern and Guermond [23] to the higher-order spatial derivatives:

$$\begin{aligned} h^2 \left\|D^2_{\nu_h} v\right\|^2_{L^2(K)} + h^2 \left\|D_{\nu_h} v\right\|^2_{L^2(K)} &\lesssim h \left\|D^2_{\bar{\nu}} v\right\|^2_{L^2(T)} + h \|D_{\bar{\nu}} v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)} + h \|v\|^2_{L^2(T)} \\ &\lesssim h^{-1} \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + h \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)} \\ &\lesssim h^{-1} \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)}. \end{aligned}$$

Mapping the constant normal back to the discrete normal via the triangle inequality bounds the element-wise surface Laplacian by:

$$h^2 \|\Delta_K v\|^2_{L^2(K)} \lesssim h^{-1} \|\nu_h \cdot \nabla v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)}.$$

Summing over the active elements $T \in \mathcal{T}_h$, the first term exactly matches the definition of the normal penalty $s_{h,N}(v, v)$. For the second term, we apply Lemma 4.3 to obtain:

$$h^2 \|\Delta_K v\|^2_{L^2(K)} \lesssim s_{h,N}(v, v) + \|v\|^2_{L^2(\Gamma_h)} + h^2 s_h(v, v).$$

Because $h^2 s_h(v, v) \leq s_h(v, v)$, we arrive at the final bound:

$$h^2 \left\|\Delta_{\Gamma_h} v\right\|^2_{L^2(\Gamma_h)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + s_h(v, v). \qquad \square$$

4.6 Concrete *A Priori* Error Estimates

With the requisite control properties verified for our chosen stabilization strategies, we establish the concrete *a priori* error estimates for the $\mathbb{Q}_1$ surface finite element method. This theorem replaces the abstract assumptions of Theorem 4.2 with the specific geometric and regularity requirements demanded by each penalty form on standard orthogonal grids.

Theorem 4.5. *Let $H_h \in W_h$ be the solution to the discrete formulation (9) on a background mesh of orthogonal hyper-rectangles satisfying the geometric approximation bounds (7) and (8). The following bounds hold:*

1. ***Volume Penalty:*** *Assume $s_h = s_{h,P} + s_{h,N}$. If the exact mean curvature satisfies $H \in [H^2(\Gamma)]^d$, the error satisfies*

$$\|H^e - H_h\|^2_{L^2(\Gamma_h)} + \|H^e - H_h\|^2_{s_h} \lesssim h^2. \tag{30}$$

2. ***Surface Penalty:*** *Assume $s_h = s_{h,P} + s_{h,n}$. If the exact mean curvature satisfies the higher regularity $H \in [H^{d+1}(\Gamma)]^d$ and the interpolation operator satisfies the extended volume bound (26), the error satisfies*

$$\|H^e - H_h\|^2_{L^2(\Gamma_h)} + \|H^e - H_h\|^2_{s_h} \lesssim h^2. \tag{31}$$

Proof. The result follows directly from the abstract estimate in Theorem 4.2 by verifying that the properties (A1) through (A4) hold for each stabilization configuration on orthogonal meshes.

Symmetry and positivity (A1) follow by definition for both methods. The tangential gradient control (A3) is satisfied for both configurations by Proposition 4.1.

For the volume penalty $s_{h,N}$, weak consistency (A2) is established in Proposition 3.3, and the surface Laplacian control (A4) is satisfied by Proposition 4.3.

For the surface penalty $s_{h,n}$, weak consistency (A2) is satisfied under the higher regularity and extended interpolation assumptions via Lemma 4.4. The surface Laplacian control (A4) is satisfied by Proposition 4.2.

With all abstract assumptions verified, the stated estimates hold. □

4.7 Non-Orthogonal Meshes in $\mathbb{R}^2$

While orthogonal background meshes are geometrically convenient, non-orthogonal discretizations are often advantageous. For materials exhibiting planar anisotropy, such as those governed by full permeability tensors or non-symmetric Cosserat stiffness matrices, imposing an orthogonal grid can introduce numerical artefacts like the grid orientation effect and artificial shear locking [27], [28]. Aligning the mesh with the oblique characteristic axes of these tensors helps accurately resolve directional stress and flux channelling. Consequently, it is useful to formulate our method for the affine equivalence class of parallelograms.

Let $\mathcal{T}_h$ be a shape-regular mesh of parallelograms in $\mathbb{R}^2$. We define the finite element space via an affine mapping from a reference element:

$$V_h = \left\{ v \in C^0(\Omega_h) : v|_T \circ F_T \in \mathbb{Q}_1(\hat{T}) \quad \forall T \in \mathcal{T}_h \right\},$$

where $F_T : \hat{T} \to T$ is the affine mapping from the reference unit square $\hat{T} = [0,1]^2$. Because F_T is affine, the composite function $v|_T$ is a second-order polynomial on T.

We now demonstrate that the normal penalty $s_{h,N}$ preserves the surface Laplacian control bound (A4) on these skewed elements.

Lemma 4.6. *Let $T \in \mathcal{T}_h$ be a parallelogram element, $K = T \cap \Gamma_h \in \mathcal{K}_h$, and $\bar{\nu}$ an arbitrary constant unit vector. For any $v \in V_h$, the spatial Hessian $D^2 v|_T$ satisfies the bound:*

$$\left\| D^2 v \right\|_{L^2(K)}^2 \lesssim h^{-3} \left\| \bar{\nu} \cdot \nabla v \right\|_{L^2(T)}^2 . \tag{32}$$

Proof. Because $v|_T$ is a second-order polynomial, its Hessian $D^2 v$ is spatially constant over T. We can therefore immediately map the surface norm to the volume by applying the geometric relation $h \left\| 1 \right\|_{L^2(K)}^2 \lesssim \left\| 1 \right\|_{L^2(T)}^2$:

$$\left\| D^2 v \right\|_{L^2(K)}^2 = \left| D^2 v \right|^2 \left\| 1 \right\|_{L^2(K)}^2 \lesssim h^{-1} \left| D^2 v \right|^2 \left\| 1 \right\|_{L^2(T)}^2 = h^{-1} \left\| D^2 v \right\|_{L^2(T)}^2 .$$

To bound this volume term, let A_T be the constant Jacobian matrix of F_T, and define the directional vectors spanning the parallelogram T as $g_1 = A_T e_1$ and $g_2 = A_T e_2$. Because $v|_T \circ F_T \in \mathbb{Q}_1(\hat{T})$, the mapped function is linear along the reference axes. Applying the chain rule guarantees the second directional derivatives along the spanning vectors vanish:

$$g_1 \cdot D^2 v \, g_1 = 0 \quad \text{and} \quad g_2 \cdot D^2 v \, g_2 = 0. \tag{33}$$

The space of 2×2 symmetric matrices is three-dimensional. Because g_1 and g_2 are linearly independent, the constraints (33) restrict $D^2 v$ to a one-dimensional subspace. Thus, $D^2 v = cM$ for some $c \in \mathbb{R}$ and a constant matrix M satisfying $\left\| M \right\|_2 = 1$.

We claim $M\bar{\nu} \neq 0$. If $M\bar{\nu} = 0$, symmetry dictates $M = \pm \bar{t} \otimes \bar{t}$, where $\bar{t}$ is orthogonal to $\bar{\nu}$. Substituting this into (33) requires $(\bar{t} \cdot g_1)^2 = (\bar{t} \cdot g_2)^2 = 0$. This implies $\bar{t}$ is orthogonal to both g_1 and g_2, which is impossible since they span $\mathbb{R}^2$.

The quantity $|M\bar{\nu}|$ is a continuous function of the relative angles between g_1, g_2, and $\bar{\nu}$. Due to mesh shape regularity, these angles reside within a compact set avoiding degenerate configurations. Because $|M\bar{\nu}|$ is strictly positive on this compact domain, it possesses a uniform lower bound, giving $|M\bar{\nu}|^{-1} \lesssim 1$.

The scalar magnitude of the entire Hessian satisfies the pointwise bound $\left| D^2 v \right| = |c| = \frac{\left| D^2 v \, \bar{\nu} \right|}{|M\bar{\nu}|} \lesssim \left| D^2 v \, \bar{\nu} \right|$. Because $\bar{\nu}$ is constant, we have the exact identity $D^2 v \, \bar{\nu} = \nabla(\bar{\nu} \cdot \nabla v)$. Substituting this pointwise bound into the volume integral yields:

$$\left\| D^2 v \right\|_{L^2(T)}^2 \lesssim \left\| \nabla(\bar{\nu} \cdot \nabla v) \right\|_{L^2(T)}^2 .$$

Applying Lemma 12.1 in Ern and Guermond [23] bounds this by $h^{-2} \left\| \bar{\nu} \cdot \nabla v \right\|_{L^2(T)}^2$. Chaining this with the initial volume mapping provides the desired bound. □

Proposition 4.4. *Assume the background mesh $\mathcal{T}_h$ consists of shape-regular parallelograms in $\mathbb{R}^2$, and the continuous discrete surface Γ_h satisfies the geometric bound $\left\| H_K \right\|_{L^\infty(\Gamma_h)} \lesssim 1$. Then the volume-based normal gradient penalty $s_{h,N}$ satisfies the surface Laplacian control bound (A4).*

Proof. Using (29), we write $\Delta_K v = (\Delta v - D^2_{\nu_h} v) - H_K \cdot \nabla v$. Because the volume Laplacian is the trace of the Hessian ($\Delta v = \operatorname{tr}(D^2 v)$) and the normal derivative is a projection ($D^2_{\nu_h} v = \nu_h \cdot D^2 v \, \nu_h$), the first bracket is bounded pointwise by $\left|D^2 v\right|$. Furthermore, because the mean curvature vector aligns with the discrete normal ($H_K = \pm |H_K| \, \nu_h$), the final term simplifies to $\pm |H_K| \, D_{\nu_h} v$.

Applying the triangle inequality, squaring, multiplying by h^2, and applying the geometric assumption $\|H_K\|_{L^\infty(K)} \lesssim 1$ yields the element-wise bound:

$$h^2 \|\Delta_K v\|^2_{L^2(K)} \lesssim h^2 \left\|D^2 v\right\|^2_{L^2(K)} + h^2 \left\|D_{\nu_h} v\right\|^2_{L^2(K)}.$$

We introduce the local constant normal $\bar{\nu}$ satisfying $\|\nu_h - \bar{\nu}\|_{L^\infty(T)} \lesssim h$. Applying Lemma 4.6 maps the full Hessian directly to the volume. For the remaining normal derivative, applying Lemma 4.1 maps the constant-direction polynomial and its perturbation from the surface to the volume. Applying Lemma 12.1 in Ern and Guermond [23] yields:

$$\begin{aligned} h^2 \left\|D^2 v\right\|^2_{L^2(K)} + h^2 \left\|D_{\nu_h} v\right\|^2_{L^2(K)} &\lesssim h^{-1} \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + \left(h \|D_{\bar{\nu}} v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)} + h \|v\|^2_{L^2(T)}\right) \\ &\lesssim h^{-1} \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + h \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)} \\ &\lesssim h^{-1} \|\bar{\nu} \cdot \nabla v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)}. \end{aligned}$$

Mapping the constant normal back to the discrete normal via the triangle inequality bounds the element-wise surface Laplacian by:

$$h^2 \|\Delta_K v\|^2_{L^2(K)} \lesssim h^{-1} \|\nu_h \cdot \nabla v\|^2_{L^2(T)} + h^{-1} \|v\|^2_{L^2(T)}.$$

Summing over the active elements $T \in \mathcal{T}_h$, the first term exactly matches the definition of the normal penalty $s_{h,N}(v, v)$. For the second term, we apply Lemma 4.3 to obtain:

$$h^2 \left\|\Delta_{\Gamma_h} v\right\|^2_{L^2(\Gamma_h)} \lesssim s_{h,N}(v, v) + \|v\|^2_{L^2(\Gamma_h)} + h^2 s_h(v, v).$$

Because $h^2 s_h(v, v) \leq s_h(v, v)$ and $s_{h,N}(v, v) \leq s_h(v, v)$, we arrive at the final bound:

$$h^2 \left\|\Delta_{\Gamma_h} v\right\|^2_{L^2(\Gamma_h)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + s_h(v, v). \qquad \square$$

5 Higher-Order Approximations ($\mathbb{P}_k$)

In Chapter 4, we established that the first-order *a priori* error bound holds on curved discrete surfaces, provided the stabilization form controls the discrete surface Laplacian (A4). While we successfully verified this control for $\mathbb{Q}_1$ elements, extending this framework to higher-order finite element spaces ($\mathbb{P}_k$ for $k \geq 2$) presents a fundamental theoretical challenge.

Specifically, the geometric properties of multilinear polynomials that allowed us to bound the full volume Laplacian no longer hold for higher-order spaces. Consequently, controlling the non-vanishing surface Laplacian without violating weak consistency remains an open problem within this abstract framework. Rather than providing a complete order k *a priori* error estimate, the objective of this chapter is twofold: to mathematically isolate the exact theoretical bottleneck preventing the closure of the order k bound, and to formulate a concrete empirical conjecture regarding the convergence rates, which will be subsequently tested in our numerical experiments.

5.1 The High-Order Discrete Surface

We maintain the assumption of a continuous, piecewise smooth discrete surface $\Gamma_h \subset U_{\delta_0}(\Gamma)$ equipped with exterior unit normals ν_h, as introduced in the previous chapter. To support an order k approximation, we elevate the geometric approximation properties. For a polynomial degree $k \geq 2$, we assume the discrete surface satisfies:

$$\left\|x - p(x)\right\|_{L^\infty(\Gamma_h)} \lesssim h^{k+1}, \tag{34}$$

$$\|\nu \circ p - \nu_h\|_{L^\infty(\Gamma_h)} \lesssim h^k. \tag{35}$$

Retaining the domain discretization consisting of the active background mesh $\mathcal{T}_h$ and the active volume Ω_h, we introduce the high-order finite element space:

$$V_h^k = \left\{ v \in C^0\left(\Omega_h\right) : v|_T \in \mathbb{P}_k(T) \quad \forall T \in \mathcal{T}_h \right\},$$

and denote the corresponding vector space as $W_h^k = [V_h^k]^d$. Because the error of the continuous-to-discrete mappings dictates the overall convergence rate of the method, we must elevate the bounds on our transformation matrices to match the high-order surface approximation.

Lemma 5.1 (Generalized Geometric Estimates)**.** *Assume the discrete surface Γ_h satisfies the high-order geometric approximation bounds* (34) *and* (35)*. Then the transformation matrix B and the surface Jacobian J satisfy the generalized uniform bounds:*

$$\begin{aligned} \|B\|_{L^\infty(\Gamma_h)} \lesssim 1, \quad \left\|B^{-1}\right\|_{L^\infty(\Gamma)} &\lesssim 1, \\ \left\|P_\Gamma - BB^T\right\|_{L^\infty(\Gamma_h)} &\lesssim h^{k+1}, \\ \|1 - J\|_{L^\infty(\Gamma_h)} &\lesssim h^{k+1}. \end{aligned} \tag{36}$$

Proof. The uniform boundedness of B and its inverse follows identically to the arguments in Lemma 3.4. For the perturbation bounds, we recall the expansion

$P_\Gamma - BB^T = 2\rho P_\Gamma \mathcal{H} P_\Gamma - \rho^2 P_\Gamma \mathcal{H}^2 P_\Gamma + E_\nu$, where E_ν is the outer product of the normal perturbation vectors, $E_\nu = \big(P_\Gamma(I - \rho\mathcal{H})\nu_h\big) \otimes \big(P_\Gamma(I - \rho\mathcal{H})\nu_h\big)$.

Substituting the high-order assumptions yields $\|\rho\|_{L^\infty(\Gamma_h)} \lesssim h^{k+1}$ and $\|E_\nu\|_2 \lesssim h^{2k}$. Because $k \geq 1$, the linear ρ term strictly dominates the higher-order normal perturbation, yielding $\left\| P_\Gamma - BB^T \right\|_{L^\infty(\Gamma_h)} \lesssim h^{k+1}$. The Jacobian estimate follows identically via the Taylor expansion of the determinant. □

Lemma 5.2. *Assume the discrete surface Γ_h satisfies the high-order geometric approximation bounds* (34) *and* (35)*. Then the discrete embedding error satisfies:*

$$\left\| x_\Gamma^e - x_{\Gamma_h} \right\|_{L^\infty(\Gamma_h)} + h \left\| \nabla_{\Gamma_h}(x_\Gamma^e - x_{\Gamma_h}) \right\|_{L^\infty(\Gamma_h)} \lesssim h^{k+1}. \tag{37}$$

Proof. The structure of this proof is identical to the proof of Lemma 3.6, but using Lemma 5.1. □

5.2 Generalization of the Stabilization Forms

To establish the well-posedness and convergence of the discrete problem on the high-order space W_h^k, we must generalize the abstract stabilization properties introduced in Chapter 3.

We assume the existence of a high-order interpolation operator $\pi_h : W_* \to W_h^k$ satisfying the standard interpolation bounds:

$$\|H^e - \pi_h H^e\|_{L^2(\Gamma_h)} + h \left\| \nabla_{\Gamma_h}(H^e - \pi_h H^e) \right\|_{L^2(\Gamma_h)} \lesssim h^{k+1} \|H\|_{H^{k+1}(\Gamma)}. \tag{38}$$

Consequently, for the exact extended solution H^e and its interpolant $\pi_h H^e$, the weak consistency assumption (A2) must be elevated to the polynomial degree k:

$$\|H^e\|_{s_h} + \|\pi_h H^e\|_{s_h} \lesssim h^k \|H\|_{H^{k+1}(\Gamma)}. \tag{A2*}$$

The remaining abstract properties—symmetry and positivity (A1), tangential gradient control (A3), and surface Laplacian control (A4)—remain structurally identical to those defined previously, but are now assumed to hold for all test functions $v \in W_h^k$.

In Chapter 4, we successfully satisfied all four assumptions for $\mathbb{Q}_1$ elements on curved discrete surfaces. For $\mathbb{P}_k$ elements ($k \geq 2$), however, a divergence occurs. While (A1) and (A3) are readily satisfied by standard interface penalties, formulating a computable stabilization form that simultaneously satisfies the high-order weak consistency (A2) and the surface Laplacian control (A4) remains an open mathematical problem.

To rigorously isolate this theoretical bottleneck, we will first establish the stability of the high-order formulation utilizing only the tangential gradient control (A3), before examining exactly how the failure of (A4) arrests the *a priori* consistency estimate.

5.3 Stability

With the generalized geometric estimates established, we can recover the stability of the high-order discrete formulation. While the proof structurally mirrors the linear case, it must explicitly account for the non-vanishing discrete curvature of the higher-order elements.

Theorem 5.3. *Let $H_h \in W_h^k$ be the solution to the discrete problem. Assume that the discrete surface Γ_h satisfies the high-order geometric approximation bounds* (34) *and* (35)*, and that the stabilization form s_h satisfies the tangential gradient control* (A3)*. Assume additionally that the discrete surface Γ_h satisfies the geometric bound $\left\|H_{\Gamma_h}\right\|_{L^\infty(\Gamma_h)} \lesssim 1$. Then, the discrete solution satisfies the stability bound:*

$$\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \lesssim 1. \tag{39}$$

Proof. Setting the test function $v = H_h$ in the discrete formulation, we obtain:

$$\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} = L_h(H_h) = \sum_{K\in\mathcal{K}_h} \left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} H_h \right\rangle_K .$$

Applying integration by parts on each curved discrete element $K \in \mathcal{K}_h$ yields:

$$\left\langle \nabla_{\Gamma_h} x_{\Gamma_h}, \nabla_{\Gamma_h} H_h \right\rangle_K = \left\langle t_K \cdot \nabla_K x_{\Gamma_h}, H_h \right\rangle_{\partial K} - \left\langle \Delta_K x_{\Gamma_h}, H_h \right\rangle_K .$$

Unlike the piecewise planar $\mathbb{P}_1$ formulation, the surface Laplacian $\Delta_K x_{\Gamma_h}$ does not vanish; it evaluates identically to the discrete mean curvature vector H_{Γ_h} of the approximate surface. By assumption, this is bounded.

Consequently, the summation of the interior terms can be bounded using the Cauchy–Schwarz inequality:

$$\left| \sum_{K\in\mathcal{K}_h} \left\langle H_{\Gamma_h}, H_h \right\rangle_K \right| \leq \left\|H_{\Gamma_h}\right\|_{L^2(\Gamma_h)} \|H_h\|_{L^2(\Gamma_h)} \lesssim \|H_h\|_{L^2(\Gamma_h)} .$$

For the boundary terms, we proceed exactly as in Theorem 3.7, evaluating the jumps across the element interfaces $E \in \mathcal{E}_h$:

$$\sum_{E\in\mathcal{E}_h} \left\langle [t_E \cdot \nabla_K x_{\Gamma_h}], H_h \right\rangle_E \leq \left(\sum_{E\in\mathcal{E}_h} h^{-1} \left\| [t_E] \right\|^2_{L^2(E)} \right)^{1/2} \left(\sum_{E\in\mathcal{E}_h} h \|H_h\|^2_{L^2(E)} \right)^{1/2} .$$

The normal jump is bounded pointwise by h^k due to the generalized approximation property (35), rendering the first factor $\lesssim h^{k-1} \lesssim 1$. The second factor is bounded by the tangential gradient control assumption (A3). Combining these bounds yields:

$$\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \lesssim \|H_h\|_{L^2(\Gamma_h)} + \left(\|H_h\|^2_{L^2(\Gamma_h)} + \|H_h\|^2_{s_h} \right)^{1/2} .$$

Applying Young's inequality to absorb the linear and square root terms into the left-hand side concludes the proof. □

5.4 An Order $k-1$ *a priori* Error Bound

While establishing an order k *a priori* estimate requires extended geometric control, the baseline tangential gradient control (A3) is sufficient to establish the convergence of the high-order formulation. Attempting to evaluate the consistency error via integration by parts reveals a theoretical bottleneck, forcing a fallback strategy that limits the proven convergence rate to order $k-1$.

Theorem 5.4. *Let $H \in [H^{k+1}(\Gamma)]^d$ be the exact mean curvature vector, and let $H^e = H \circ p$ be its extension. Let $H_h \in W_h^k$ be the solution to the discrete problem. Assume that the discrete surface Γ_h satisfies the high-order geometric approximation bounds* (34) *and* (35)*, the stabilization form s_h satisfies the tangential gradient control* (A3) *and the weak consistency bound* (A2*)*, and the interpolation operator π_h satisfies the standard bound* (38)*. Then, the error satisfies the bound:*

$$\|H^e - H_h\|^2_{L^2(\Gamma_h)} + \|H^e - H_h\|^2_{s_h} \lesssim h^{2(k-1)}. \tag{40}$$

Proof. Following the error splitting strategy from Chapter 3, we set $v = \pi_h H^e - H_h$. By the triangle inequality, (38), and (A2*), we have $\|H^e - \pi_h H^e\|^2_{L^2(\Gamma_h)} + \|H^e - \pi_h H^e\|^2_{s_h} \lesssim h^{2k}$.

Let $w = x^e_\Gamma - x_{\Gamma_h}$. Expanding the consistency error of the right-hand side over the active elements $K \in \mathcal{K}_h$ and applying integration by parts yields:

$$L^*(v) - L_h(v) = \sum_{K\in\mathcal{K}_h} \left\langle w, t_K \cdot \nabla_{\Gamma_h} v \right\rangle_{\partial K} - \sum_{K\in\mathcal{K}_h} \langle w, \Delta_K v\rangle_K . \tag{41}$$

For the interface term in (41), evaluating the jump across the internal edges $E \in \mathcal{E}_h$ and applying the Cauchy–Schwarz inequality yields:

$$\left| \sum_{E\in\mathcal{E}_h} \left\langle w, [t_E \cdot \nabla_{\Gamma_h} v] \right\rangle_E \right| \leq \left(\sum_{E\in\mathcal{E}_h} h^{-1} \|w\|^2_{L^2(E)} \right)^{1/2} \left(\sum_{E\in\mathcal{E}_h} h \left\| [t_E \cdot \nabla_{\Gamma_h} v] \right\|^2_{L^2(E)} \right)^{1/2} .$$

The first factor is $\lesssim h^k$ from (34) and Lemma 3.1. The second factor is $\lesssim \sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}}$ from (A3).

For the interior term in (41), the surface Laplacian $\Delta_K v$ does not vanish. Unless we can prove a bound like $h^2 \left\|\Delta_{\Gamma_h} v\right\|^2_{L^2(\Gamma_h)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}$, we are forced to backtrack. Instead, we apply Cauchy–Schwarz directly; using Lemma 5.2 and (A3), we obtain:

$$\begin{aligned} L^*(v) - L_h(v) &= \left\langle \nabla_{\Gamma_h} w, \nabla_{\Gamma_h} v \right\rangle_{\Gamma_h} \\ &\lesssim \left\|\nabla_{\Gamma_h} w\right\|_{L^2(\Gamma_h)} \left\|\nabla_{\Gamma_h} v\right\|_{L^2(\Gamma_h)} \\ &\lesssim h^k h^{-1} \sqrt{h^2 \left\|\nabla_{\Gamma_h} v\right\|^2_{L^2(\Gamma_h)}} \\ &\lesssim h^{k-1} \sqrt{\|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}}. \end{aligned}$$

From here the argument is as before. □

5.5 The Consistency Bottleneck and Empirical Conjecture

Theorem 5.4 explicitly demonstrates why the analytical machinery developed in Chapter 4 is insufficient for general $\mathbb{P}_k$ spaces. Equation (41) shows that achieving the order k estimate via integration by parts requires recovering the surface Laplacian control, recalling assumption (A4):

$$h^2 \left\|\Delta_{\Gamma_h} v\right\|^2_{L^2(\Gamma_h)} \lesssim \|v\|^2_{L^2(\Gamma_h)} + \|v\|^2_{s_h}, \quad \forall v \in W_h^k. \tag{A4}$$

If this assumption holds, applying the Cauchy–Schwarz inequality immediately yields the order k estimate.

In Chapter 4, we verified (A4) for $\mathbb{Q}_1$ elements by utilizing the expansion of the Laplace–Beltrami operator $\Delta_K v = \Delta v - D^2_{\nu_h} v - H_K \cdot \nabla v$. Because multilinear shape functions possess a vanishing (or constrained) volume Laplacian Δv, the normal penalty forms were sufficient to bound the remaining terms.

For general $\mathbb{P}_k$ elements ($k \geq 2$), this structural advantage disappears. The volume Laplacian Δv becomes a non-zero polynomial of degree $k-2$. The stabilization forms $s_{h,N}$ and $s_{h,n}$ penalize derivatives in the normal direction; they provide no control over the tangential second derivatives comprising the remainder of the full volume Laplacian. While one could theoretically devise a penalty that bounds the full spatial Hessian $D^2 v$ directly, doing so would excessively penalize the exact solution, violating the weak consistency constraint (A2*).

Consequently, establishing an order k *a priori* bound for curved surfaces remains an open problem. Despite this analytical limitation, numerical investigations reveal that the method achieves order k convergence in practice, provided appropriate stabilization is employed.

At this point, a technical distinction between the two supplementary normal penalties must be made. The volume-based normal penalty $s_{h,N}$ utilizes only first-order derivatives and naturally extends to arbitrary polynomial degrees. Conversely, the surface-based normal penalty $s_{h,n}$, as introduced in Definition 4.1, penalizes derivatives up to $j = d$. For $\mathbb{P}_k$ spaces, we must instead sum over the k first normal derivatives.

Due to previous results observed by Frachon and Zahedi [17] for a specific geometry in $\mathbb{R}^2$, we conclude this theoretical framework with the following empirical conjecture, which will be validated in the subsequent numerical experiments:

Conjecture 5.5. *Let the discrete surface Γ_h satisfy the high-order geometric approximation bounds* (34) *and* (35)*. If the discrete variational problem on W_h^k is stabilized using the volume-based ghost penalty $s_{h,P}$ combined with either the volume-based normal penalty $s_{h,N}$ or the appropriate k-th order generalization of the surface-based normal penalty $s_{h,n}$, the formulation achieves the convergence rate:*

$$\|H^e - H_h\|^2_{L^2(\Gamma_h)} + \|H^e - H_h\|^2_{s_h} \lesssim h^{2k}. \tag{42}$$

6 Numerical Experiments

In this section, we verify the theoretical *a priori* error bounds established in the preceding chapters through a series of numerical experiments. We begin by validating the standard linear formulation, proceed to the $\mathbb{Q}_1$ extension, and finally explore the generalization of the stabilization forms to higher-order finite element spaces empirically.

6.1 Construction of the Discrete Geometry

Throughout the theoretical analysis, the discrete surface Γ_h was abstractly assumed to satisfy specific geometric approximation properties. For the implementation, we concretely realize this surface using level set functions.

The exact surfaces are implicitly defined by

$$\Gamma = \left\{ x \in \mathbb{R}^d : \varphi(x) = 0 \right\},$$

where φ is the level set function. We embed Γ within a bounding hypercube Ω_0 aligned with the coordinate axes and construct a uniform background mesh $\mathcal{T}_{h,0}$ on Ω_0. By interpolating φ onto the finite element space associated with $\mathcal{T}_{h,0}$, we obtain an approximate level set function φ_h, which defines the discrete surface:

$$\Gamma_h = \left\{ x \in \mathbb{R}^d : \varphi_h(x) = 0 \right\}.$$

Implicitly defining geometries via level set functions is very convenient, as it provides immediate identities for the exterior unit normal and the exact mean curvature vector:

$$\nu = \frac{\nabla \varphi}{|\nabla \varphi|}, \qquad H = (\nabla \cdot \nu)\nu.$$

6.2 Model Problems

We evaluate the method on five benchmark geometries across $\mathbb{R}^2$ and $\mathbb{R}^3$.

6.2.1 Example geometries in $\mathbb{R}^3$

Figure 1 illustrates the two selected geometries in $\mathbb{R}^3$. The first is an elliptic torus defined by the level set function

$$\varphi(x, y, z) = \sqrt{a\left(\sqrt{x^2 + y^2} - R\right)^2 + (bz)^2} - 1,$$

utilizing the parameters $R = 1.0$, $a = 2.5$, and $b = 10/3$. The second geometry is the deco-cube benchmark introduced by Hansbo, Larson and Zahedi [15], given by

$$\varphi(x, y, z) = \phi(x, y, z)\phi(y, z, x)\phi(z, x, y) - \delta,$$

where $\delta = 0.02$, and the auxiliary function ϕ is defined as

$$\phi(u, v, w) = (u^2 + v^2 - c^2)^2 + (w^2 - 1)^2,$$

with $c = \sqrt{3/4}$.

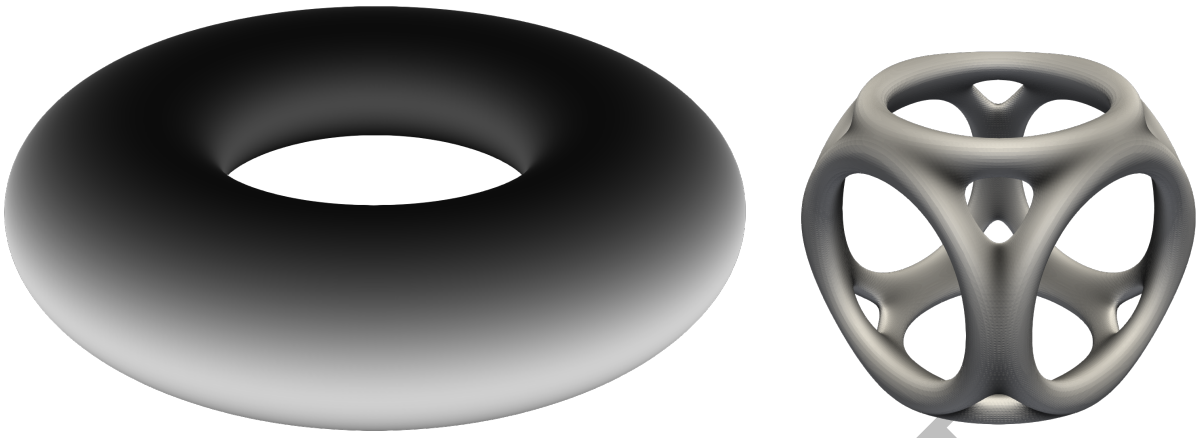

(a) Elliptic torus (b) Deco-cube

Figure 1: Example geometries in $\mathbb{R}^3$; the surface gradient denotes the variance in mean curvature.

6.2.2 Example geometries in $\mathbb{R}^2$

Figure 2 illustrates the three two-dimensional benchmark geometries. Their respective level set functions are defined as:

$$\varphi(x,y) = ((x-a)^2+y^2)((x+a)^2+y^2) - b^4 + cx \qquad \text{(Peanut)}$$

$$\varphi(x,y) = \sqrt{x^2+y^2} - \big(1 + 0.3\cos(\theta) + 0.2\sin(2\theta) + 0.1\cos(3\theta)\big) \qquad \text{(Amoeba)}$$

$$\varphi(x,y) = \sqrt{x^2+y^2} - \big(1 + 0.3\cos(7\theta)\big) \qquad \text{(7-Star)}$$

where θ is the principal argument of the complex coordinate $x+iy$. The parameters for the peanut geometry are fixed at $a = 1$, $b = 1.02$, and $c = 0.5$. The first two example geometries have rather modest curvature variations, and for each formulation below, only one will be presented in the figures as visually they are essentially indistinguishable.

6.3 Implementation Details

The numerical experiments were implemented in an in-house C++ code developed within the Department of Mathematics at the KTH Royal Institute of Technology. This framework is based on the open-source `FreeFEM++` library [29], which provides the core finite element space architecture. For reproducibility, an open-source version of the CutFEM infrastructure is available [30]. The resulting sparse linear systems were solved via the MUMPS direct solver [31], [32].

For all convergence studies, the background mesh is uniformly refined such that the characteristic element size h is reduced by a factor of $2^{-1/4}$ at each step. This scaling continuously shifts the relative position of the interface within the background elements, which demonstrates that the observed convergence and condition numbers are robust against arbitrary cut positions.

The empirical convergence rates are computed via a linear least-squares fit on a logarithmic scale. To isolate the asymptotic behaviour from pre-asymptotic variations,

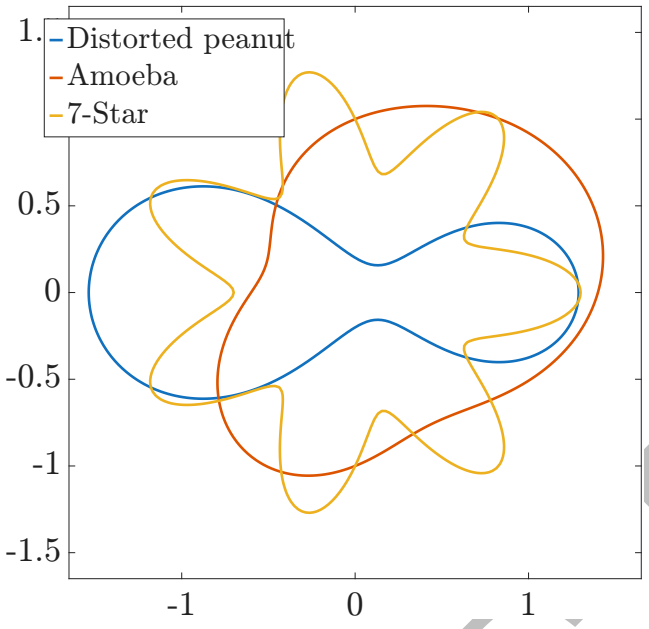

Figure 2: Example geometries in $\mathbb{R}^2$.

this fit is restricted to the data from the finest half of the mesh sequence.

6.4 The $\mathbb{P}_1$ Formulation

6.4.1 Convergence of Examples in $\mathbb{R}^3$

For the two examples in $\mathbb{R}^3$, we compare how the face-based ghost penalty (17) and the volume-based ghost penalty (19) perform in practice.

For the elliptic torus, the convergence plots can be seen in Figure 3. With both stabilization forms we observe first-order convergence, though the parameter sensitivity seems higher for the volume-based ghost penalty, with significant locking being observed with $\tau_P = 10^{-2}$ and a lower pre-asymptotic convergence rate for $\tau_P = 10^{-4}$. For well-chosen parameters τ_P and τ_F the final error is similar. We note that for sufficiently fine meshes, we are by Theorem 3.2 guaranteed to have an $\mathcal{O}(h)$ error bound for any $\tau > 0$, but since we cannot compute with infinitely small meshes, we may in practice not observe the asymptotic behaviour.

Because of memory constraints and the complex geometry of the deco-cube, it was not possible to run as many mesh refinements as for the elliptic torus. Hence, there are fewer points in the convergence plots in Figure 4. The plots show the same locking phenomena as observed for the elliptic torus when using $\tau_P = 10^{-2}$.

6.4.2 Convergence of Examples in $\mathbb{R}^2$

In this section, we investigate whether stabilization terms not needed for the *a priori* error bound have positive effects in practice. For this purpose, we compare the stabilization form from Definition 3.8 with the combined form stabilization form from Definition 3.9 on the example geometries in $\mathbb{R}^2$. When comparing different stabilization

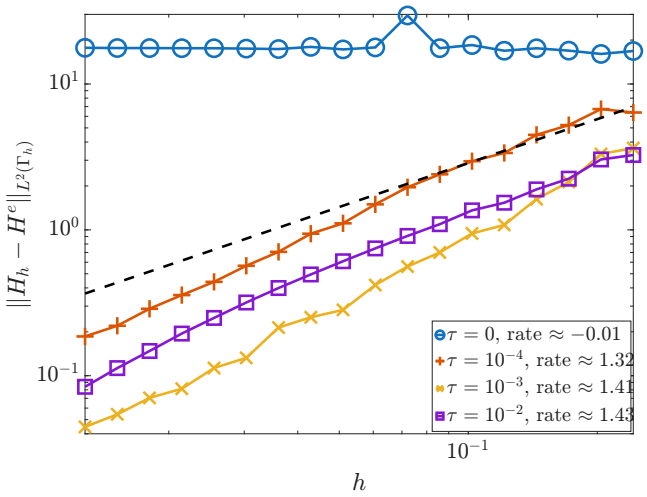

(a) Face-based ghost penalty

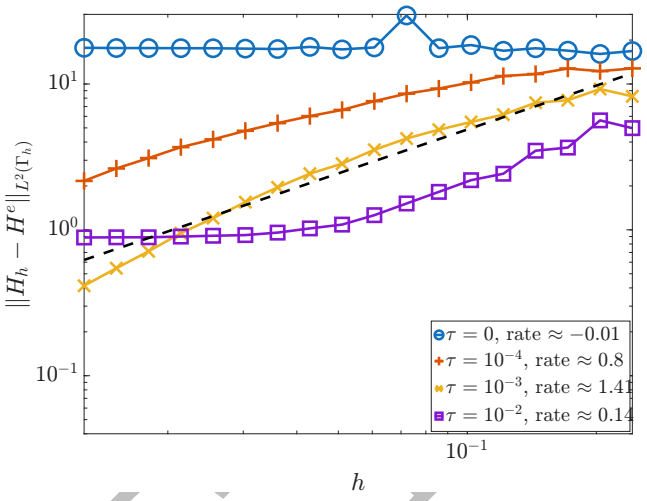

(b) Volume-based ghost penalty

Figure 3: Convergence plots for the elliptic torus. The legends indicate values of τ_F and τ_P respectively. The dashed reference line represents $\mathcal{O}(h)$.

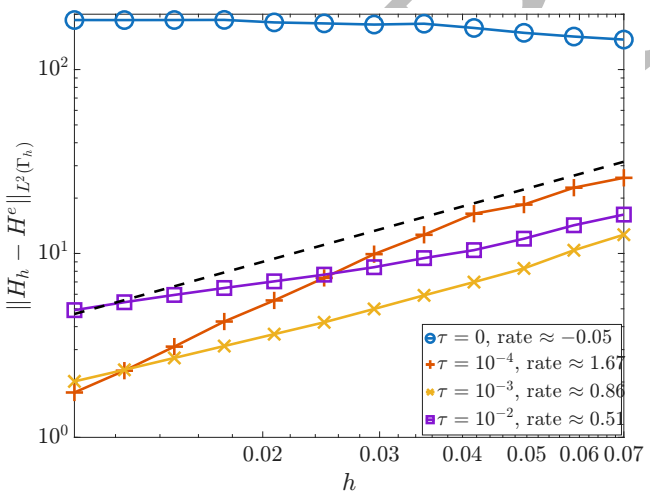

(a) Face-based ghost penalty

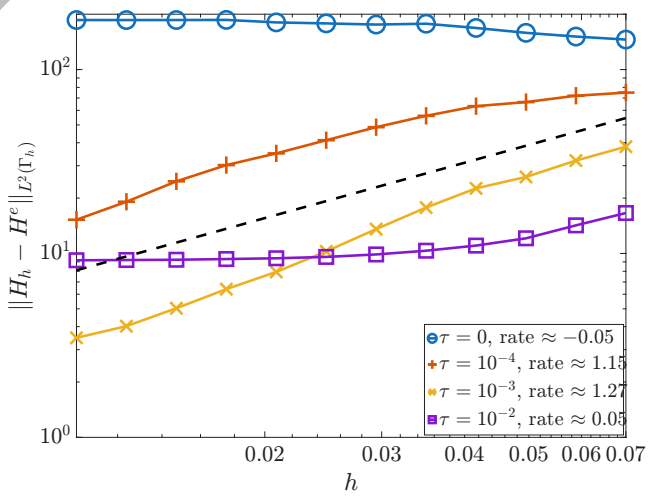

(b) Volume-based ghost penalty

Figure 4: Convergence plots for the deco-cube. The legends indicate values of τ_F and τ_P respectively. The dashed reference line represents $\mathcal{O}(h)$.

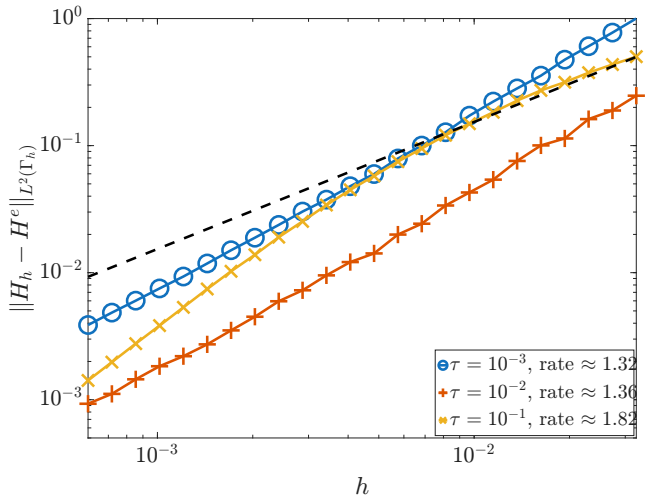

(a) Only $s_{h,P}$, the legend shows different values of τ_P.

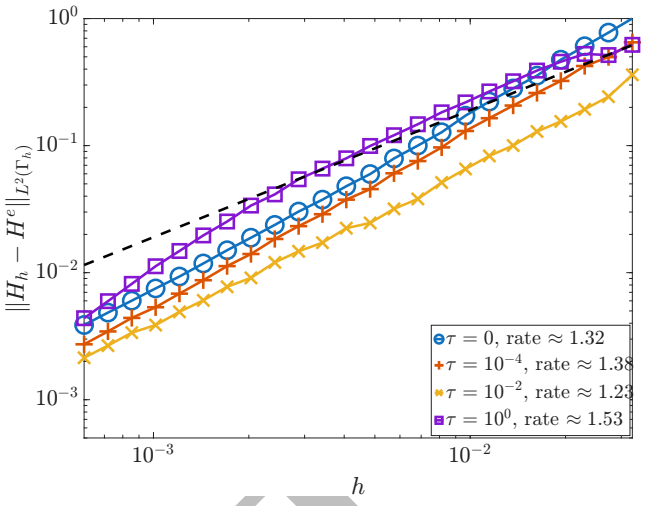

(b) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N. We take $\tau_P = 10^{-3}$ constant.

Figure 5: Convergence plots for the distorted peanut, using $\mathbb{P}_1$ elements. The dashed reference line represents $\mathcal{O}(h)$.

parameters we vary τ_P for the formulation with only $s_{h,P}$, and for the formulation with $s_{h,P} + s_{h,N}$ we keep $\tau_P = 10^{-3}$ constant and vary τ_N.

In Figure 5, we observe superlinear convergence for all parameter choices, a behaviour which is representative for both the modest-curvature geometries. Different stabilization parameters mostly affect the initial error. We observe that larger stabilization parameters generally improve the final convergence rate, though they sometimes converge slower for the initial meshes and have larger initial errors. This contrasts with the convergence for the 7-star, shown in Figure 6, whose curvature varies more; we generally see smaller errors with smaller stabilization parameters. It should be noted that the error reduction rate seems to increase for each refinement, which makes it difficult to draw any conclusion about what parameters would be best for finer meshes.

For the three-dimensional examples and the tetrahedral meshes, it was observed that the parameter choice of τ_P was very delicate. This was not observed in $\mathbb{R}^2$. Furthermore, in $\mathbb{R}^3$, the volume-based ghost penalty also produced notably higher error even when the convergence rate was the same; this was also not the case in $\mathbb{R}^2$. For the higher-order experiments below, which are restricted to $\mathbb{R}^2$, we thus choose to continue using the volume-based ghost penalty, as we want to avoid computing higher-order derivatives.

6.4.3 Condition Numbers for the Examples in $\mathbb{R}^2$

In this section we numerically investigate the scaling of the condition number of the resulting linear system of equations arising from (9). While in classic unstabilized finite element methods, a mass matrix, like the one arising from a_h, would have an $\mathcal{O}(1)$ condition number, in CutFEM, however, it would actually become unbounded due to the arbitrary cutting of the elements, meaning the measure of $K \in \mathcal{K}_h$ can be arbitrarily small.

The stabilization forms needed for convergence give us minimum bounds on the

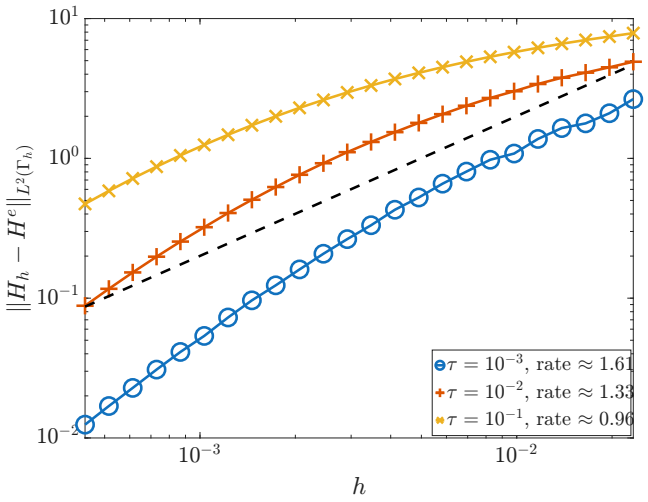

(a) Only $s_{h,P}$, the legend shows different values of τ_P.

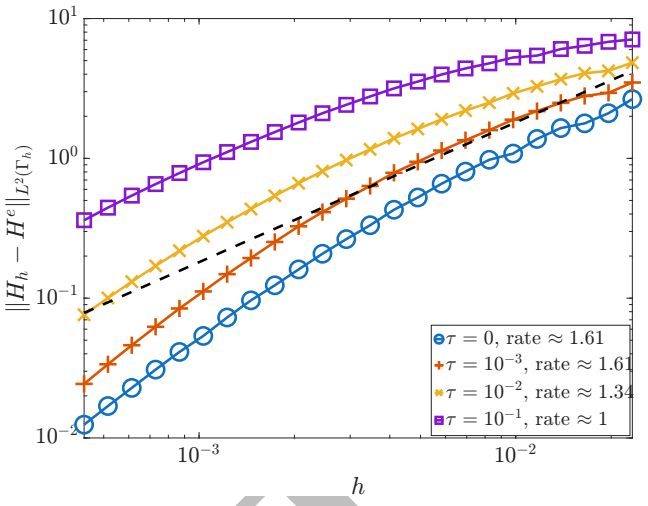

(b) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N. We take $\tau_P = 10^{-3}$ constant.

Figure 6: Convergence plots for the 7-star, using $\mathbb{P}_1$ elements. The dashed reference line represents $\mathcal{O}(h)$.

eigenvalues of the matrix by acting on the whole elements, which are shape regular, rather than only the cut subset, and hence restore well-behaved condition numbers. The stabilization forms behave similarly to second-order differential operators, so we expect the condition numbers to scale as $\mathcal{O}(h^{-2})$, a bound proven for stabilized surface methods by Larson and Zahedi [16].

When Larson and Zahedi [16] investigated stabilizing the mass matrix, they were able to pick a different scaling in terms of h for the stabilization form, this maintained the $\mathcal{O}(1)$ condition number bound. If we were to apply similar scaling here we would lose the properties needed for the approximation error.

Figure 7 shows condition numbers, estimated using `condest` in MATLAB, for the distorted peanut. The visual behaviour of the different geometries in $\mathbb{R}^2$ is indistinguishable, so this figure is representative for all the geometries. Notably, the value of τ_P largely does not affect the condition number, while introducing a very small τ_N can drop the condition number 2–3 orders of magnitude. For all combinations of parameters we observe that the condition number scales as $\mathcal{O}(h^{-2})$.

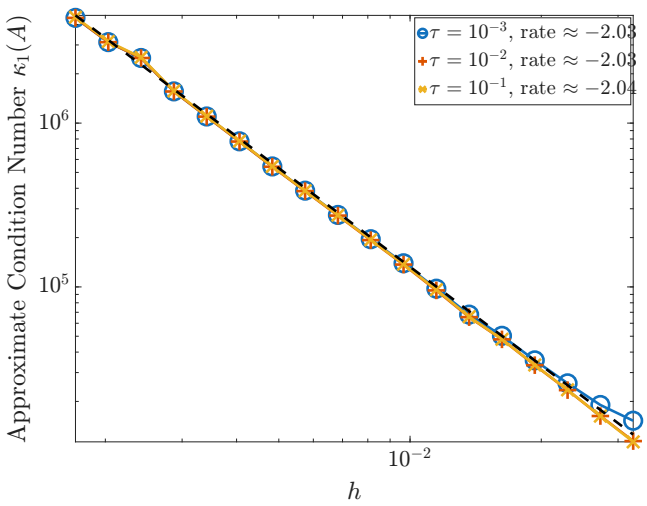

(a) Only $s_{h,P}$, the legend shows different values of τ_P.

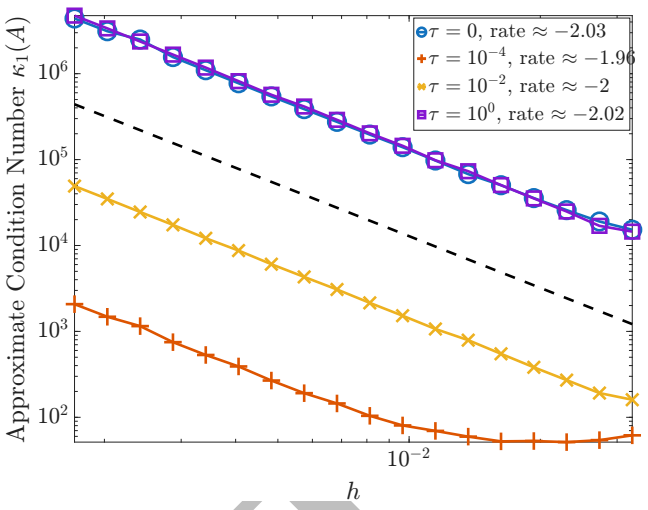

(b) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N. We take $\tau_P = 10^{-3}$ constant.

Figure 7: Condition numbers for the distorted peanut, using $\mathbb{P}_1$ elements. The dashed reference line represents $\mathcal{O}(h^{-2})$.

6.5 High-Order Integration on Curved Implicit Surfaces

In a numerical implementation of CutFEM for surface problems we need to be able to integrate on every $K \in \mathcal{K}_h$; when every K is completely flat it is easy to implement a quadrature rule with exact integration for polynomials of the required order. For a $\mathbb{P}_k, k \geq 2$ or a general $\mathbb{Q}_k$ approximation the element intersections K are curved and integration becomes much more complicated. A possible naive approach would be to simply tessellate K with smaller piecewise planar pieces, but this would for any fixed sub-tessellation resolution result in a quadrature error of $\mathcal{O}(h^2)$.

Saye has suggested two different approaches for different cell types, both of which are implemented in the CutFEM Library. The first approach was suggested in 2015, and allows for arbitrary order integration of smooth surfaces on hyper-rectangles [33]. The second approach was suggested in 2022 and allows for integration over surfaces defined by polynomials on simplicial cells [34]. Both approaches essentially work by treating the implicitly defined geometry as a height function and a dimension reduction strategy together with one-dimensional quadrature methods. These approaches also have the advantage of yielding only positive quadrature weights.

6.6 The $\mathbb{Q}_1$ formulation

6.6.1 Convergence of Examples in $\mathbb{R}^3$

Because we want to avoid computing the third-order normal derivatives needed for $s_{h,n}$ on tensor-product elements in $\mathbb{R}^3$, we only examine the performance of $s_h = s_{h,P} + s_{h,N}$ here. In Figure 8, the approximation errors for both example geometries are shown. When comparing these to Figures 3 and 4, it is evident that the error magnitudes are very similar at corresponding mesh sizes h. Surprisingly, the method achieves first-order convergence even when the normal-based stabilization is not present ($\tau_N = 0$).

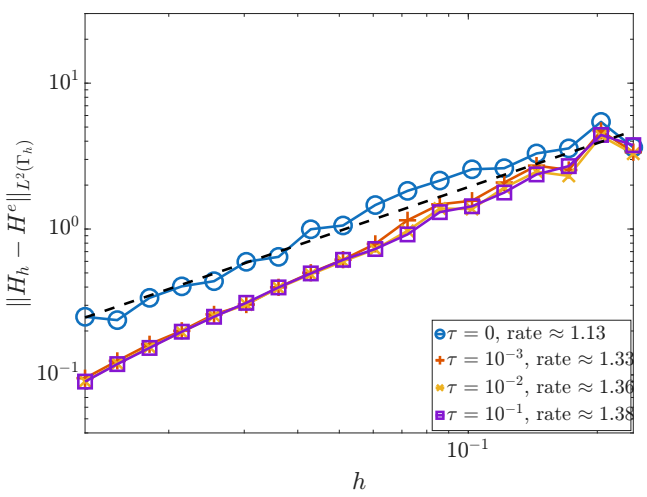

(a) Convergence for the elliptic torus using $s_h = s_{h,P} + s_{h,N}$, the legend shows different values of τ_N while $\tau_P = 10^{-2}$ is fixed.

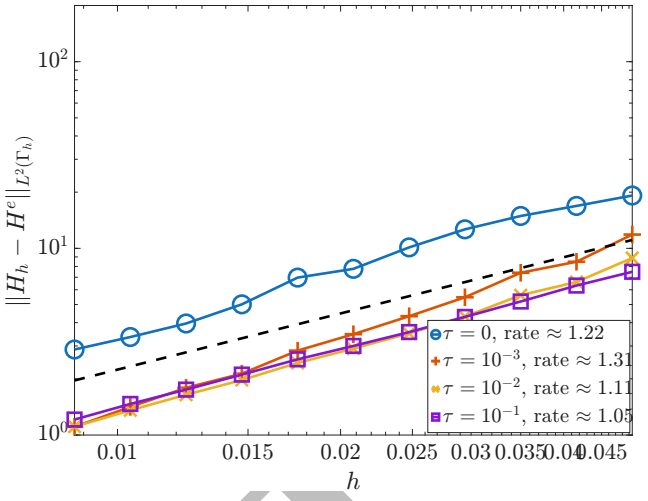

(b) Convergence for the deco-cube using $s_h = s_{h,P} + s_{h,N}$, the legend shows different values of τ_N while $\tau_P = 10^{-3}$ is fixed.

Figure 8: Convergence plots for the three-dimensional geometries, using $\mathbb{Q}_1$ elements. The dashed reference line represents $\mathcal{O}(h)$.

A closer inspection of the error curves does, however, reveal a subtle distinction. When $\tau_N = 0$, the local rate of error reduction exhibits slight variations between successive mesh refinements, although the overall trend remains monotonic. In contrast, in the presence of normal-based stabilization ($\tau_N > 0$), the reduction rate is noticeably more uniform, particularly on the finer meshes. This suggests that while normal control may not be required to achieve convergence in these specific three-dimensional examples, its presence dampens local geometric sensitivities and ensures a more predictable error reduction.

6.6.2 Construction of the discrete surface for $\mathbb{Q}_1$ Elements in $\mathbb{R}^2$

While the theoretical framework established in Chapter 4 accommodates curved discrete surfaces for $\mathbb{Q}_1$ elements, the standard implicit construction described in Section 6.1 does not naturally yield a simple geometric representation when φ_h is a piecewise $\mathbb{Q}_1$ function.

To streamline the numerical integration, we utilize a piecewise planar approximation of Γ for the quadrilateral elements in $\mathbb{R}^2$. One way to do this is by solving $\varphi_h = 0$ along each element face and then connecting those points with line segments. Because a piecewise planar interface is a valid geometric subcase of our continuous surface framework, this constructed Γ_h satisfies all of our geometric assumptions without violating the conditions stated in the *a priori* error analysis.

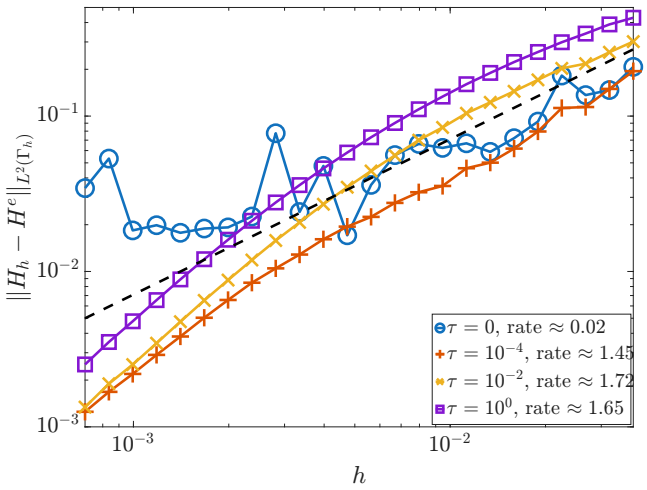

(a) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N.

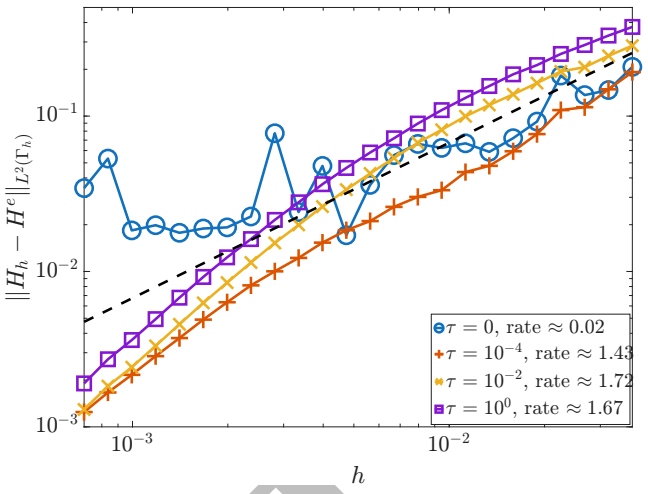

(b) $s_{h,P} + s_{h,n}$, the legend shows different values of τ_n.

Figure 9: Convergence plots for the amoeba, using $\mathbb{Q}_1$ elements. The dashed reference line represents $\mathcal{O}(h)$.

6.6.3 Convergence of Examples in $\mathbb{R}^2$

With the piecewise planar geometry established, we evaluate the discrete variational formulation using the combined stabilization forms

$$\begin{aligned} s_h^N &= s_{h,P} + s_{h,N}, \\ s_h^n &= s_{h,P} + s_{h,n}. \end{aligned} \tag{43}$$

As established in Chapter 4, both $s_{h,N}$ and $s_{h,n}$ provide the necessary control for $\mathbb{Q}_1$ elements on the rectangular mesh we will be using. In order to keep the range of parameters we need to sweep over manageable, we will keep $\tau_P = 10^{-1}$ constant.

In Figures 9 and 10, the resulting convergence plots are presented. As previously mentioned, the peanut geometry behaves very similarly to the amoeba. The graphs clearly demonstrate that the basic stabilization forms which work for $\mathbb{P}_1$ elements do not yield convergence for quadrilateral elements. We also observe that first-order convergence is achieved when $\tau_N, \tau_n > 0$. In general, both suggested stabilization forms perform similarly with no clear advantage to either alternative. It is also interesting to note that at the same h, the quadrilateral elements have an error that is approximately one order of magnitude larger than the triangular elements in the $\mathbb{P}_1$ formulation.

The requirement for normal-based stabilization, $\tau_N > 0$ or $\tau_n > 0$, and the order-of-magnitude error gap relative to $\mathbb{P}_1$ elements contrast with the three-dimensional $\mathbb{Q}_1$ formulation, which achieves $\mathbb{P}_1$ accuracy without normal-based stabilization. Replacing the curved zero-set of φ_h with a piecewise planar surface Γ_h introduces a secondary geometric approximation, inflating the error constant compared to the exact multilinear surface used in $\mathbb{R}^3$. While our theoretical analysis relied on normal-based stabilization, why the exact multilinear surfaces in our three-dimensional examples bypass this requirement remains unresolved.

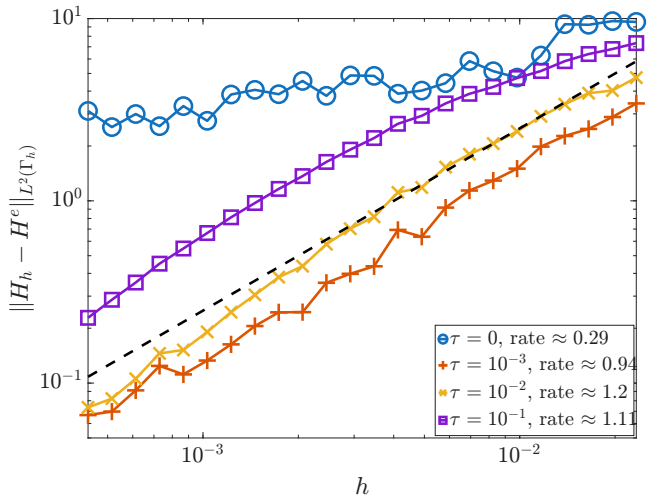

(a) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N.

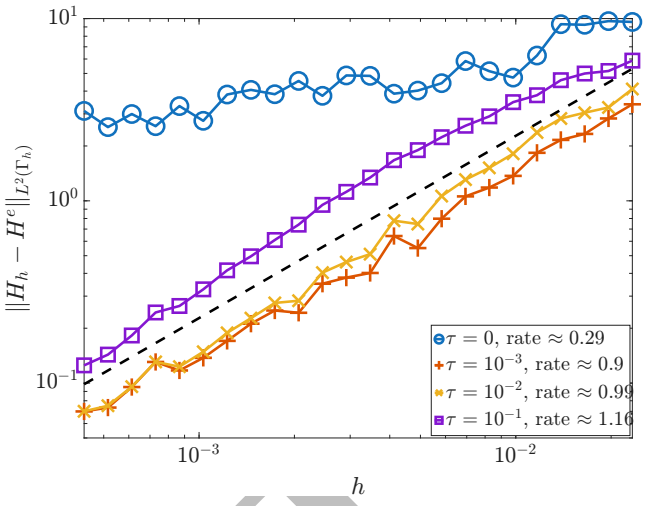

(b) $s_{h,P} + s_{h,n}$, the legend shows different values of τ_n.

Figure 10: Convergence plots for the 7-star, using $\mathbb{Q}_1$ elements. The dashed reference line represents $\mathcal{O}(h)$.

6.6.4 Condition Numbers for the Examples in $\mathbb{R}^2$

The condition numbers for the $\mathbb{Q}_1$ elements are similar to those of the $\mathbb{P}_1$ elements. As illustrated for the amoeba in Figure 11, the condition numbers scale approximately with h^{-2} and the final sizes are similar across all two-dimensional examples. In general there seems to be a sweet spot for the parameters τ_N and τ_n where the condition number is the lowest.

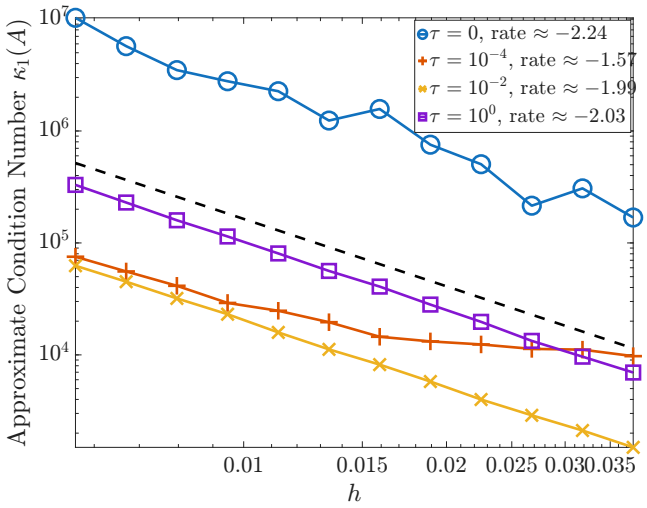

(a) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N.

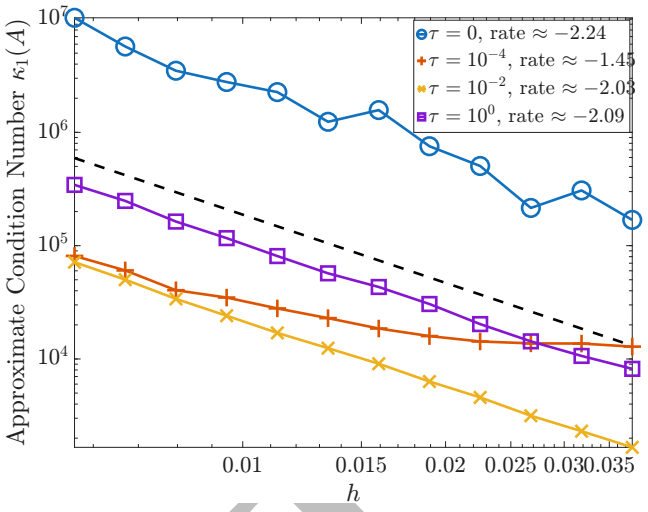

(b) $s_{h,P} + s_{h,n}$, the legend shows different values of τ_n.

Figure 11: Condition numbers for the amoeba, using $\mathbb{Q}_1$ elements. The dashed reference line represents $\mathcal{O}(h^{-2})$.

6.7 The $\mathbb{P}_2$ Formulation

6.7.1 Convergence of Examples in $\mathbb{R}^3$

Because the volume-based ghost penalty did not perform well in the $\mathbb{P}_1$ examples in $\mathbb{R}^3$, we here use the combined penalty $s_h = s_{h,F} + s_{h,N}$. The convergence plots are presented in Figure 12. Due to limited computational resources, the number of refinements had to be reduced, and as such it might be difficult to draw any strong conclusions, particularly for the deco-cube.

When normal-based stabilization is present, $\tau_N > 0$, we observe second-order convergence for both geometries. This is empirical validation of Conjecture 5.5 as established in the previous chapter. However, just like for the $\mathbb{R}^3$ examples using $\mathbb{Q}_1$ elements, we also observe second-order convergence without normal-based stabilization.

6.7.2 Convergence of Examples in $\mathbb{R}^2$

Here, we compare the two combined stabilization forms (43) again. Previously, Frachon and Zahedi [17] has used $s_h = s_{h,F} + s_{h,n}$ for an example geometry in $\mathbb{R}^2$, for both $\mathbb{P}_2$ and $\mathbb{P}_3$ formulations. For the parameter sweeps, we have chosen to keep $\tau_P = 10^{-3}$ constant.

The convergence plots for the two-dimensional examples can be seen in Figures 13 and 14, representing the general behaviour across the geometries. Consistent with the three-dimensional experiments above, we observe second-order convergence when the normal-based stabilization forms are present. This also aligns with what was observed by Frachon and Zahedi [17].

Furthermore, the two-dimensional geometries mirror the unexpected behaviour observed in $\mathbb{R}^3$; all three examples still exhibit second-order convergence even when normal-

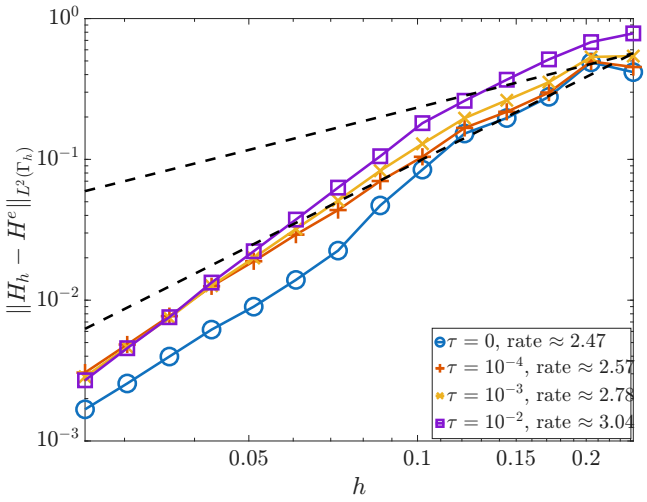

(a) Convergence for the elliptic torus using $s_h = s_{h,F} + s_{h,N}$, the legend shows different values of τ_N while $\tau_F = 10^{-4}$ is fixed.

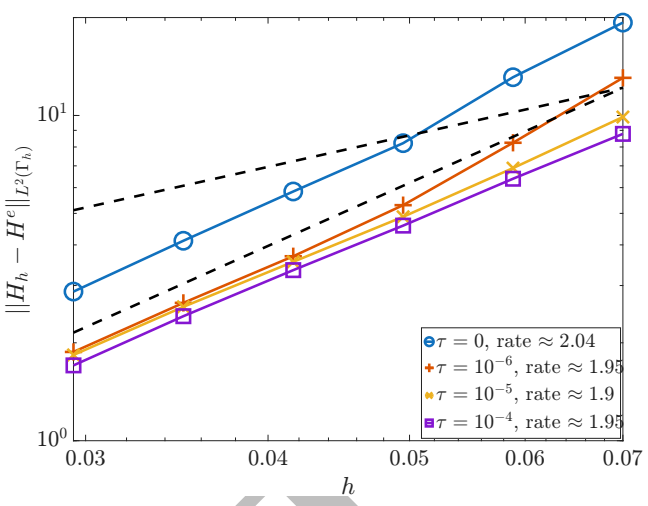

(b) Convergence for the deco-cube using $s_h = s_{h,F} + s_{h,N}$, the legend shows different values of τ_N while $\tau_F = 10^{-5}$ is fixed.

Figure 12: Convergence plots for the three-dimensional geometries, using $\mathbb{P}_2$ elements. The dashed reference lines represent $\mathcal{O}(h)$ and $\mathcal{O}(h^2)$.

based stabilization is absent. This contrasts with what we saw for the $\mathbb{Q}_1$ examples in $\mathbb{R}^2$, which required normal-based stabilization in theory and in practice to achieve convergence. Whether this convergence without normal-based stabilization generalizes to other geometries in $\mathbb{R}^d$, or if it is merely a pre-asymptotic artefact that would degrade for smaller h on these specific examples, is unknown. This behaviour could also be due to the particular way we construct Γ_h in these examples, via φ_h, since this is also the primary difference between the $\mathbb{R}^3$ and $\mathbb{R}^2$ examples for the $\mathbb{Q}_1$ experiments.

When comparing to the $\mathbb{P}_1$ formulation, it is interesting to note that here we usually do not observe any superquadratic convergence, while we consistently observed superlinear convergence in the $\mathbb{P}_1$ formulation. We also see that for the finest h we used, the error in the $\mathbb{P}_2$ formulation is about three orders of magnitude smaller.

6.7.3 Condition Numbers for the Examples in $\mathbb{R}^2$

While introducing the normal stabilizing forms did not significantly improve convergence compared to the pure volume ghost stabilization $s_{h,P}$, it drastically affects the condition numbers of the linear system. Figure 15 shows that when stabilizing only with $s_{h,P}$, the condition number grows as h^{-4}, a scaling observed consistently across all examples. To explain this degradation, we establish a lower bound on the condition number.

For a geometry in $\mathbb{R}^d$, the maximum eigenvalue of the stabilized system is dominated by the stabilization form, scaling as $\lambda_{\max} \sim h^{d-3}$. Using the Rayleigh quotient, we bound the minimum eigenvalue from above by evaluating the stabilized bilinear form $a_h + s_{h,P}$ with the approximate level set function $\varphi_h \in V_h^k$ and its nodal vector $\mathbf{\Phi}$:

$$\lambda_{\min} \leq \frac{a_h(\varphi_h, \varphi_h) + s_{h,P}(\varphi_h, \varphi_h)}{\mathbf{\Phi}^T \mathbf{\Phi}}.$$

According to Ern and Guermond [23, Lemma 11.7], the discrete vector norm satisfies

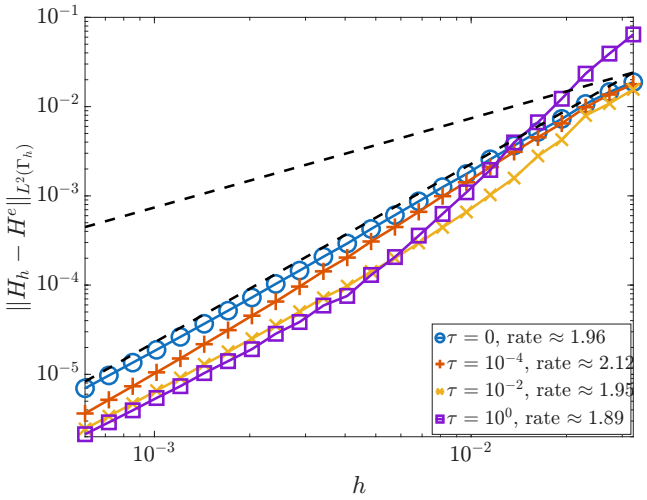

(a) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N.

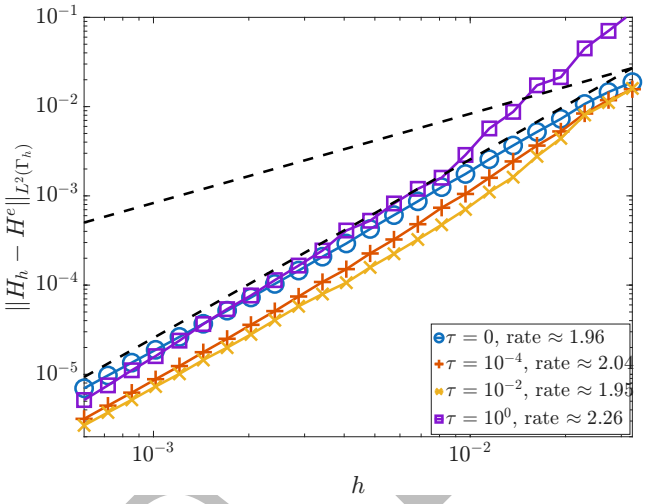

(b) $s_{h,P} + s_{h,n}$, the legend shows different values of τ_n.

Figure 13: Convergence plots for the distorted peanut, using $\mathbb{P}_2$ elements. The dashed reference lines represent $\mathcal{O}(h)$ and $\mathcal{O}(h^2)$.

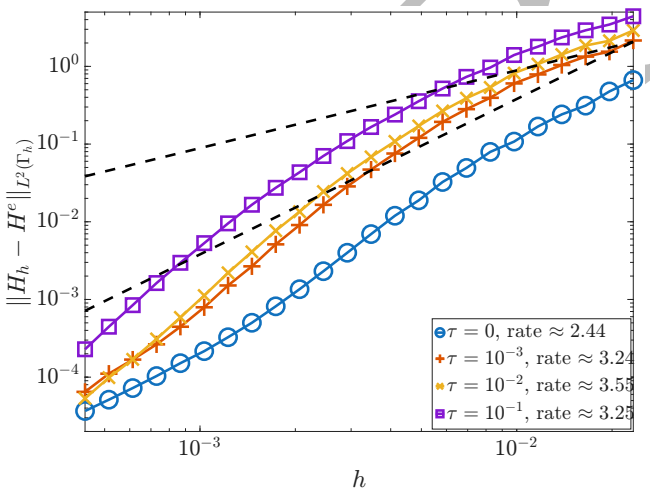

(a) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N.

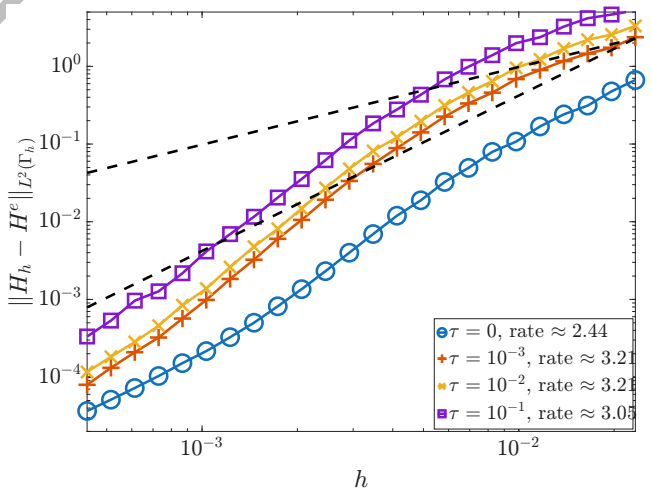

(b) $s_{h,P} + s_{h,n}$, the legend shows different values of τ_n.

Figure 14: Convergence plots for the 7-star, using $\mathbb{P}_2$ elements. The dashed reference lines represent $\mathcal{O}(h)$ and $\mathcal{O}(h^2)$.

$h^d \mathbf{\Phi}^T \mathbf{\Phi} \sim \|\varphi_h\|^2_{\Omega_h}$. Because φ_h interpolates an exact level set function with a non-vanishing gradient ($|\nabla\varphi| \sim 1$), it grows linearly away from the zero contour. Its magnitude scales as h over a macroscopic fraction of the active volume, yielding $\|\varphi_h\|^2_{\Omega_h} \sim h^3$ and the discrete scaling $\mathbf{\Phi}^T \mathbf{\Phi} \sim h^{3-d}$.

Since Γ_h is the zero contour of φ_h, the mass form vanishes ($a_h(\varphi_h, \varphi_h) = 0$), leaving only the ghost penalty contribution. For any two adjacent elements $T_1, T_2 \in \mathcal{T}_h$, let φ_1 and φ_2 denote the polynomial extensions of φ_h from T_1 and T_2 onto their shared patch $\omega = T_1 \cup T_2$. Let $q \in \mathbb{P}_k$ be the best polynomial approximation of the exact level set function φ on ω in the L^2 norm. Because φ_i and q are polynomials, discrete norm equivalence on shape-regular patches guarantees $\|\varphi_i - q\|_{L^2(\omega)} \lesssim \|\varphi_i - q\|_{L^2(T_i)}$.

Since φ_i coincides with the nodal interpolant $\mathcal{I}_h\varphi$ on T_i and $\mathcal{I}_h$ preserves polynomials ($\mathcal{I}_h q = q$), the error relative to φ is bounded by combining this local stability with standard approximation estimates. These estimates, obtained by applying the Bramble–Hilbert lemma [23, Lemma 11.9] on a reference patch and scaling back to ω, extract the required mesh dependence:

$$\begin{aligned}
\|\varphi_i - \varphi\|^2_{L^2(\omega)} &\lesssim \|\varphi_i - q\|^2_{L^2(\omega)} + \|q - \varphi\|^2_{L^2(\omega)} \\
&\lesssim \left\|\mathcal{I}_h(\varphi - q)\right\|^2_{L^2(T_i)} + \|q - \varphi\|^2_{L^2(\omega)} \\
&\lesssim h^{2(k+1)} \|\varphi\|^2_{H^{k+1}(\omega)}.
\end{aligned}$$

Adding and subtracting φ bounds the squared L^2 jump across the interface via the triangle inequality:

$$\begin{aligned}
\left\|[\varphi_h]\right\|^2_{L^2(\omega)} &= \|\varphi_1 - \varphi_2\|^2_{L^2(\omega)} \\
&\lesssim \|\varphi_1 - \varphi\|^2_{L^2(\omega)} + \|\varphi - \varphi_2\|^2_{L^2(\omega)} \\
&\lesssim h^{2(k+1)} \|\varphi\|^2_{H^{k+1}(\omega)}.
\end{aligned}$$

Summing these local bounds over the active volume Ω_h and applying the global h^{-3} scaling bounds the total algebraic energy:

$$\begin{aligned}
a_h(\varphi_h, \varphi_h) + s_{h,P}(\varphi_h, \varphi_h) &= s_{h,P}(\varphi_h, \varphi_h) \\
&\lesssim h^{-3} h^{2(k+1)} \|\varphi\|^2_{H^{k+1}(\Omega_h)} \\
&\sim h^{-3} h^{2k+2} h \\
&= h^{2k}.
\end{aligned}$$

Dividing this energy bound by the vector norm yields $\lambda_{\min} \lesssim h^{2k+d-3}$. Comparing this to $\lambda_{\max}$ eliminates the spatial dimension and establishes the condition number lower bound:

$$\begin{aligned}
\kappa &\geq \frac{\lambda_{\max}}{\lambda_{\min}} \\
&\gtrsim \frac{h^{d-3}}{h^{2k+d-3}} \\
&= h^{-2k}.
\end{aligned}$$

For $\mathbb{P}_1$ elements ($k = 1$), this agrees with the observed h^{-2} scaling, whereas for $\mathbb{P}_2$ elements ($k = 2$), it yields $\kappa \gtrsim h^{-4}$, predicting the observed degradation. Larson and

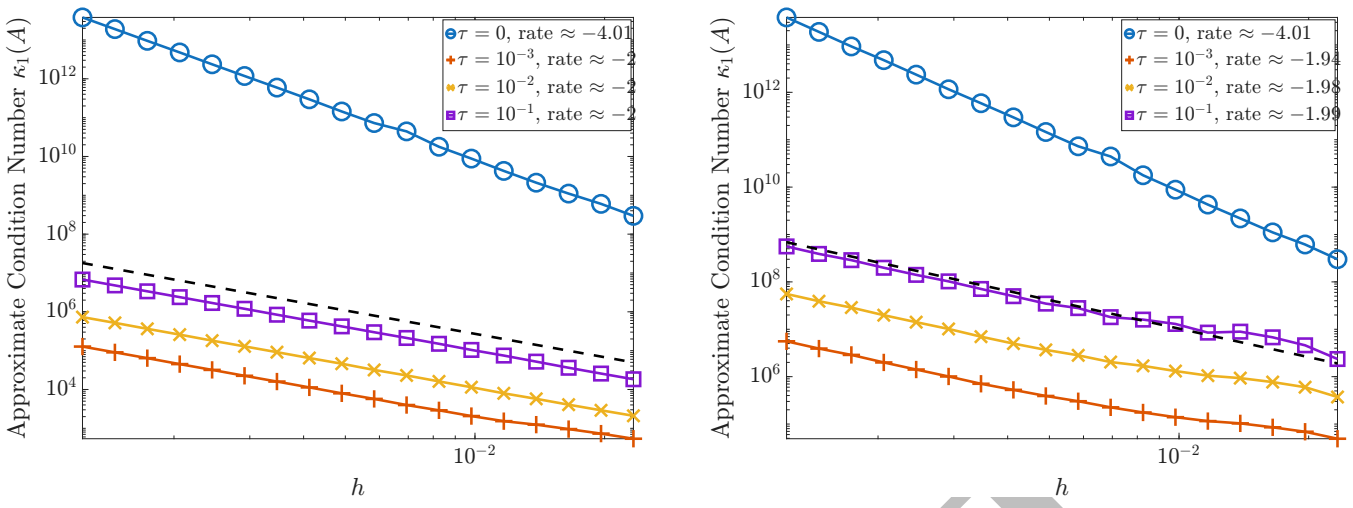

(a) $s_{h,P} + s_{h,N}$, the legend shows different values of τ_N.

(b) $s_{h,P} + s_{h,n}$, the legend shows different values of τ_n.

Figure 15: Condition numbers for the 7-star, using $\mathbb{P}_2$ elements. The dashed reference line represents $\mathcal{O}(h^{-2})$.

Zahedi [16] established that explicitly stabilizing the normal derivatives ($s_{h,N}$ or $s_{h,n}$) provides the necessary bounds over the entire finite element space to prevent this. These forms ensure $\lambda_{\min} \gtrsim h^{d-1}$, restoring the condition number scaling to h^{-2} as seen in Figure 15.

This improved scaling is important for both direct solver accuracy and iterative solver convergence. Empirically, small parameter values $\tau_n, \tau_N > 0$ produce the smallest condition numbers without altering the asymptotic scaling. Furthermore, when $\tau_n = \tau_N$, s_h^N yields condition numbers roughly one to three orders of magnitude smaller than s_h^n.

7 Conclusions and Outlook

The results of this thesis show that extending the CutFEM approximation of the mean curvature vector beyond piecewise linear elements is not only a matter of increasing the polynomial degree or changing the reference element. The stabilization has to compensate for different terms depending on the finite element space. In particular, the surface Laplacian of discrete test functions, which is harmless in the standard $\mathbb{P}_1$ setting, becomes central both for multilinear elements and for higher-order polynomial spaces.

For $\mathbb{P}_1$ elements, the existing first-order theory can be expressed in terms of abstract assumptions on the stabilization. These assumptions separate the roles of weak consistency, positivity, tangential gradient control, and trace control. Under these conditions, the method satisfies a first-order *a priori* error estimate. This formulation is useful because it makes clear which parts of the analysis depend on the particular ghost penalty and which parts only depend on the structural properties of the stabilization.

The $\mathbb{Q}_1$ case shows where this structure first changes. Since the discrete surface Laplacian no longer vanishes locally, the proof requires an additional control assumption. With this assumption, first-order convergence follows also for curved discrete surfaces. The proposed normal-based stabilizations provide this control on orthogonal hyper-rectangular meshes, and the volume-based normal penalty was also shown to do so on shape-regular parallelogram meshes in $\mathbb{R}^2$. Thus, first-order convergence for $\mathbb{Q}_1$ elements is recovered, but only after adding a stabilization mechanism that is not needed in the same way for $\mathbb{P}_1$ elements.

For higher-order $\mathbb{P}_k$ elements, the same issue becomes more restrictive. Stability can be proved under the tangential gradient control assumption, and the resulting error estimate is of order $k-1$. An order k estimate is not obtained by the present argument. The obstruction is that the discrete surface Laplacian contains tangential second-derivative contributions which are not controlled by the normal-based stabilizations considered here. Penalizing the full Hessian would remove this difficulty, but would also conflict with the weak consistency needed for an order k estimate. Whether such an estimate can be proved within a related framework therefore remains open.

The numerical results are consistent with the proven theory in the cases where the theory is complete, but they also indicate behaviour which is not yet fully explained. The $\mathbb{P}_1$ computations show first-order convergence, with some sensitivity to the stabilization parameter. For $\mathbb{Q}_1$ elements in $\mathbb{R}^2$, the experiments support the theoretical need for normal-based stabilization: the basic stabilizations are not sufficient, while first-order convergence is recovered once $s_{h,N}$ or $s_{h,n}$ is included. In the three-dimensional $\mathbb{Q}_1$ examples, however, first-order convergence was observed also without normal-based stabilization. This suggests that the specific construction of the discrete surface via the approximate level set function may possess additional properties that are unaccounted for in the present analysis.

The $\mathbb{P}_2$ experiments show a similar gap between what was conjectured and what was observed. With normal-based stabilization, the method exhibits second-order convergence in the tested $\mathbb{R}^2$ and $\mathbb{R}^3$ examples. However, the same rate was also observed when only the ghost penalty was used. Since the current analysis does not

control the corresponding surface Laplacian term, this behaviour should be interpreted as numerical evidence rather than as a theoretical conclusion. It remains unclear whether it persists for other geometries, higher degrees, or different surface approximations.

The conditioning results give a clearer practical conclusion. For $\mathbb{P}_2$ elements, the volume ghost penalty alone leads to condition numbers growing like h^{-4}, while adding normal-based stabilization gives h^{-2} scaling. Thus, even in cases where the convergence rate appears unaffected, the normal-based terms can be important for the algebraic properties of the method.

The main open problem is to determine whether an order k estimate can be proved for higher-order spaces on curved discrete surfaces. This requires either a sharper argument which avoids the missing surface Laplacian control, or a stabilization which supplies this control without losing weak consistency. It is also important to understand why convergence of the higher rate is observed in some cases without normal-based stabilization. Further numerical studies for $k > 2$, more general geometries, and less structured meshes would help determine whether this behaviour is robust. The methods can also be further developed for evolving interfaces and coupled bulk-surface problems.

8 References

[1] J. W. Barrett, H. Garcke and R. Nürnberg, 'Parametric approximation of Willmore flow and related geometric evolution equations,' *SIAM Journal on Scientific Computing*, vol. 31, no. 1, pp. 225–253, 2008. DOI: 10.1137/070700231.

[2] S. Gross and A. Reusken, *Numerical Methods for Two-Phase Incompressible Flows* (Springer Series in Computational Mathematics). Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, vol. 40. DOI: 10.1007/978-3-642-19686-7.

[3] C.-J. Heine, 'Computations of form and stability of rotating drops with finite elements,' *IMA Journal of Numerical Analysis*, vol. 26, no. 4, pp. 723–751, 2006. DOI: 10.1093/imanum/drl007.

[4] Q. Du, L. Ju and L. Tian, 'Finite element approximation of the Cahn–Hilliard equation on surfaces,' *Computer Methods in Applied Mechanics and Engineering*, vol. 200, no. 29–32, pp. 2458–2470, 2011. DOI: 10.1016/j.cma.2011.04.018.

[5] A. Schmidt, 'Computation of three dimensional dendrites with finite elements,' *Journal of Computational Physics*, vol. 125, no. 2, pp. 293–312, 1996. DOI: 10.1006/jcph.1996.0095.

[6] M. E. Gurtin, J. Weissmüller and F. Larché, 'A general theory of curved deformable interfaces in solids at equilibrium,' *Philosophical Magazine A*, vol. 78, no. 5, pp. 1093–1109, 1998, ISSN: 0141-8610, 1460-6992. DOI: 10.1080/01418619808239977.

[7] K. Deckelnick, G. Dziuk and C. M. Elliott, 'Computation of geometric partial differential equations and mean curvature flow,' *Acta Numerica*, vol. 14, pp. 139–232, 2005. DOI: 10.1017/S0962492904000224.

[8] M. Meyer, M. Desbrun, P. Schröder and A. H. Barr, 'Discrete differential-geometry operators for triangulated 2-manifolds,' in *Visualization and Mathematics III*, H.-C. Hege and K. Polthier, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 35–57. DOI: 10.1007/978-3-662-05105-4_2.

[9] G. Dziuk, 'Finite elements for the Beltrami operator on arbitrary surfaces,' in *Partial Differential Equations and Calculus of Variations*, S. Hildebrandt and R. Leis, Eds., vol. 1357, Berlin, Heidelberg: Springer Berlin Heidelberg, 1988, pp. 142–155. DOI: 10.1007/BFb0082865.

[10] G. Dziuk, 'An algorithm for evolutionary surfaces,' *Numerische Mathematik*, vol. 58, no. 1, pp. 603–611, 1990. DOI: 10.1007/BF01385643.

[11] K. Hildebrandt, K. Polthier and M. Wardetzky, 'On the convergence of metric and geometric properties of polyhedral surfaces,' *Geometriae Dedicata*, vol. 123, no. 1, pp. 89–112, 2006. DOI: 10.1007/s10711-006-9109-5.

[12] S. Gross and A. Reusken, 'Finite element discretization error analysis of a surface tension force in two-phase incompressible flows,' *SIAM Journal on Numerical Analysis*, vol. 45, no. 4, pp. 1679–1700, 2007. DOI: 10.1137/060667530.

[13] A. Hansbo, P. Hansbo and M. G. Larson, 'A finite element method on composite grids based on Nitsche's method,' *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 37, no. 3, pp. 495–514, 2003. DOI: 10.1051/m2an:2003039.

[14] E. Burman, P. Hansbo, M. G. Larson and S. Zahedi, 'Cut finite element methods,' *Acta Numerica*, vol. 34, pp. 1–121, 2025. DOI: 10.1017/S0962492925000017.

[15] P. Hansbo, M. G. Larson and S. Zahedi, 'Stabilized finite element approximation of the mean curvature vector on closed surfaces,' *SIAM Journal on Numerical Analysis*, vol. 53, no. 4, pp. 1806–1832, 2015. DOI: 10.1137/140982696.

[16] M. G. Larson and S. Zahedi, 'Stabilization of high order cut finite element methods on surfaces,' *IMA Journal of Numerical Analysis*, vol. 40, no. 3, pp. 1702–1745, 2020. DOI: 10.1093/imanum/drz021.

[17] T. Frachon and S. Zahedi, 'A cut finite element method for incompressible two-phase Navier–Stokes flows,' *Journal of Computational Physics*, vol. 384, pp. 77–98, 2019. DOI: 10.1016/j.jcp.2019.01.028.

[18] I. L. Novak, F. Gao, Y.-S. Choi, D. Resasco, J. C. Schaff and B. M. Slepchenko, 'Diffusion on a curved surface coupled to diffusion in the volume: Application to cell biology,' *Journal of Computational Physics*, vol. 226, no. 2, pp. 1271–1290, 2007. DOI: 10.1016/j.jcp.2007.05.025.

[19] D. Gilbarg and N. S. Trudinger, *Elliptic Partial Differential Equations of Second Order* (Grundlehren der mathematischen Wissenschaften v.224). Berlin, Heidelberg: Springer Berlin Heidelberg, 2014.

[20] D. Gilbarg and N. S. Trudinger, *Elliptic Partial Differential Equations of Second Order* (Classics in Mathematics 224). Berlin, Heidelberg: Springer Berlin Heidelberg, 2001. DOI: 10.1007/978-3-642-61798-0.

[21] E. Kreyszig, *Introductory Functional Analysis with Applications* (Wiley Classics Library), Wiley Classics Library edition. New York: Wiley, 1989.

[22] J. Preuß, 'Higher order unfitted isoparametric space-time FEM on moving domains,' M.S. thesis, University of Göttingen, Göttingen, Germany, 2018.

[23] A. Ern and J.-L. Guermond, *Finite Elements I: Approximation and Interpolation* (Texts in Applied Mathematics 72). Cham: Springer, 2021.

[24] E. Burman, P. Hansbo, M. G. Larson and A. Massing, 'Cut finite element methods for partial differential equations on embedded manifolds of arbitrary codimensions,' *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 52, no. 6, pp. 2247–2282, 2018. DOI: 10.1051/m2an/2018038.

[25] S. Zahedi, 'A space-time cut finite element method with quadrature in time,' in *Geometrically Unfitted Finite Element Methods and Applications*, S. P. A. Bordas, E. Burman, M. G. Larson and M. A. Olshanskii, Eds., vol. 121, Cham: Springer International Publishing, 2017, pp. 281–306. DOI: 10.1007/978-3-319-71431-8_9.

[26] A. Demlow and M. A. Olshanskii, 'An adaptive surface finite element method based on volume meshes,' *SIAM Journal on Numerical Analysis*, vol. 50, no. 3, pp. 1624–1647, 2012. DOI: 10.1137/110842235.

[27] T. Poulet, H. A. Sheldon, U. Kelka and P. Behnoudfar, 'Impact of permeability anisotropy misalignment on flow rates predicted by hydrogeological models,' *Hydrogeology Journal*, vol. 31, no. 8, pp. 2129–2137, 2023. DOI: 10.1007/s10040-023-02708-4.

[28] F. Löschner, J. A. Fernández-Fernández, S. R. Jeske, A. Longva and J. Bender, 'Micropolar elasticity in physically-based animation,' *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 6, no. 3, pp. 1–24, 2023. DOI: 10.1145/3606922.

[29] F. Hecht, 'New development in FreeFem++,' *Journal of numerical mathematics*, vol. 20, no. 3–4, pp. 251–266, 2012.

[30] T. Frachon and S. Myrbäck, *CutFEM-Library*, GitHub, 2026.

[31] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent and J. Koster, 'A fully asynchronous multifrontal solver using distributed dynamic scheduling,' *SIAM Journal on Matrix Analysis and Applications*, vol. 23, no. 1, pp. 15–41, 2001.

[32] P. R. Amestoy, A. Guermouche, J.-Y. L'Excellent and S. Pralet, 'MUMPS: A general purpose distributed memory sparse solver,' *Parallel Computing*, vol. 32, no. 3, pp. 261–281, 2006.

[33] R. I. Saye, 'High-order quadrature methods for implicitly defined surfaces and volumes in hyperrectangles,' *SIAM Journal on Scientific Computing*, vol. 37, no. 2, A993–A1019, 2015. DOI: 10.1137/140966290.

[34] R. I. Saye, 'High-order quadrature on multi-component domains implicitly defined by multivariate polynomials,' *Journal of Computational Physics*, vol. 448, p. 110 720, 2022. DOI: 10.1016/j.jcp.2021.110720.

TRITA – SCI-GRU
Stockholm, Sweden 2026

www.kth.se