



Interpolation-based immersed finite element and isogeometric analysis

Jennifer E. Fromm^a, Nils Wunsch^b, Ru Xiang^a, Han Zhao^a, Kurt Maute^b, John A. Evans^b, David Kamensky^{a,*}

^a University of California San Diego, San Diego, CA 92093, USA

^b University of Colorado Boulder, Boulder, CO 80309, USA

Received 13 September 2022; received in revised form 9 December 2022; accepted 5 January 2023

Available online 12 January 2023

Abstract

We introduce a new paradigm for immersed finite element and isogeometric methods based on interpolating function spaces from an unfitted background mesh into Lagrange finite element spaces defined on a foreground mesh that captures the domain geometry but is otherwise subject to minimal constraints on element quality or connectivity. This is a generalization of the concept of Lagrange extraction from the isogeometric analysis literature and also related to certain variants of the finite cell and material point methods, as well as the concept of free-form deformation from design. Crucially, the interpolation may be approximate without sacrificing high-order convergence rates, which distinguishes the present method from existing high-order finite cell, CutFEM, and immersogeometric approaches. The interpolation paradigm also permits non-invasive reuse of existing finite element software for immersed analysis. We analyze the properties of the interpolation-based immersed paradigm for a model problem and implement it on top of the open-source FEniCS finite element software, to apply it to a variety of problems in fluid, solid, and structural mechanics where we demonstrate high-order accuracy and applicability to practical geometries like trimmed spline patches.

© 2023 Elsevier B.V. All rights reserved.

Keywords: Immersed finite element method; Immersogeometric analysis; Trimmed isogeometric analysis; CutIGA; Lagrange extraction

1. Introduction

The finite element (FE) method [1] is a widely-used approach for approximating solutions to partial differential equations (PDEs). In summary, it approximates solutions to variational formulations of PDEs in function spaces defined piece-wise over a collection of disjoint elements, called a mesh, whose closures cover (an approximation of) the PDE domain. Typically, each element is diffeomorphic to one of a few standard reference elements, like a unit right triangle in 2D or a tetrahedron in 3D, and the boundary of the (approximate) PDE domain coincides with a collection of element boundary parts, e.g., edges of elements in 2D and faces of elements in 3D. An example of such a “boundary-fitted” mesh is shown in Fig. 1(a). However, the approximation power and regularity of FE function spaces depend on the sizes, shapes, and connectivity of the elements, placing nontrivial geometric and

* Corresponding author.

E-mail address: dmkamensky@eng.ucsd.edu (D. Kamensky).




Fig. 1. (a) Domain discretized using a body-fitted mesh; (b) Domain discretized using a tensor-product background mesh with a background-fitted foreground mesh; (c) Domain discretized using a tensor-grid background mesh with a background-unfitted foreground mesh.

topological constraints on the mesh generation process. These constraints are often quite challenging to satisfy, especially for the complicated PDE domains endemic to engineering analysis. Thus, mesh generation remains a pervasive bottleneck in practical FE analysis [2].

These meshing constraints become *much* easier to satisfy if one removes the requirement that the PDE boundary be a collection of element boundary parts, instead allowing it to cut arbitrarily through element interiors. We refer to such meshes as “unfitted”, and call numerical methods using unfitted meshes “immersed boundary methods” — or simply “immersed methods” — in an extrapolation of terminology originally applied to early work by Peskin [3]. For instance, it is easy to construct highly-regular function spaces with excellent approximation properties on structured tensor-product grids of elements, and it would be convenient to simply immerse an arbitrary complex PDE domain geometry into such a grid without needing its boundary to coincide with element boundary parts.

The earliest immersed methods amounted, essentially, to worsening the geometric approximation of the PDE domain, leading to low-order accuracy and the (still widespread) belief that immersed methods necessarily represent a trade-off between order of accuracy and ease of meshing. However, over the past few decades, higher-order immersed methods have emerged based on the insight that one only needs an accurate representation of the domain geometry for numerical quadrature of variational forms from the FE method and not necessarily for constructing the approximate function spaces. These quadrature-based immersed methods serve as our point of departure in the present work.

Quadrature-based immersed methods typically integrate weak forms by constructing a fitted mesh of the PDE domain and applying Gaussian quadrature within each element. For instance, in Fig. 1(b), the fitted mesh of blue triangles would be used for quadrature, while the unfitted black structured mesh would be used for defining function spaces. However, quadrature imposes *far* fewer geometric and topological constraints than construction of function spaces. Further, even brute-force high-resolution/low-order quadrature meshes (such as those in early iterations of the finite cell method [4]) will not add degrees of freedom to the discrete function space and may remain efficient in applications whose cost is dominated by algebraic solution of the discretized PDE. Thus, an easy-to-generate, “foreground” mesh for quadrature can be superposed on an unfitted “background” mesh for defining function spaces to obtain high-order-accurate immersed methods. Some examples of accurate quadrature-based immersed methods from the literature include the finite cell approach mentioned earlier [5–7] and various cut finite element (CutFEM) [8] and level-set [9] schemes. One can also view the original material point method (MPM) [10] as a quadrature-based immersed method using a low-order quadrature rule for the mass measure of a solid body.

A historically-separate line of research aimed at circumventing the FE meshing problem is isogeometric analysis (IGA) [11,12] where function spaces of smooth splines from computer-aided design (CAD) are directly used to approximate PDE solutions. However, these spline spaces often have *more* topological constraints on their construction than classical FE spaces, and industrial CAD models often represent complex geometries by “trimming” [13] unfitted rectangular meshes called patches to obtain the desired shapes. Thus, IGA is hardly an alternative to immersed methods; on the contrary, it has only intensified interest in them! The symbiotic union of isogeometric and immersed methods is referred to as “immersogeometric analysis” (IMGA) [14] and the case of defining smooth spline spaces on a background mesh is sometimes also called CutIGA [15] by analogy to CutFEM.

An optimistic reading of the published results on immersed finite element and isogeometric methods might lead one to conclude that the FE mesh-generation problem is nearly solved. However, the majority of work on

immersed methods has relied on specialized research codes whose code architectures differ greatly from existing finite element or isogeometric solvers. Thus, adoption of immersed methods in practice has been very limited. A notable attempt at general-purpose implementation of quadrature-based immersed methods in an existing FE code was the MultiMesh [16] functionality developed within the FEniCS [17] software framework. However, the invasive nature of the implementation made it difficult to complete and maintain, and MultiMesh functionality was not retained in the FEniCSx redesign of the software.

A similar barrier to entry once existed for IGA, but a major advance in the design of IGA software was the concept of Bézier extraction [18,19] where smooth spline basis functions are represented as linear combinations of less-regular Bernstein basis functions. This was later reformulated as *interpolation* of IGA basis functions using Lagrange FE bases, referred to as “Lagrange extraction” [20] and allowing for direct reuse of FE software implementing Lagrange bases for IGA (e.g., [21,22]). **The central idea of the present work is to build on the concept of Lagrange extraction by interpolating basis functions defined on an unfitted background mesh into a Lagrange FE space defined on a fitted foreground mesh where quadrature is performed.** In this setting, the background basis does not need to be a smooth spline space; it could also be a classical FE space defined on an unstructured mesh. This not only permits reuse of existing FE software, but also suggests some efficient approximations, leading to a novel class of immersed methods.

An earlier related idea is the projection-based adaptive Gaussian quadrature (PAGQ) implementation of the finite cell method [23] where shape functions are defined on quadrature sub-cells. However, we now push the concept further by considering the possibility of *approximate* Lagrange extraction. While extraction-based IGA and PAGQ exactly reproduce basis functions through interpolation, we explore the possibility of allowing this interpolation to be approximate. In situations like Fig. 1(b), where the foreground mesh is fitted to the intersection of each background element with the domain, it may be possible to capture all monomials of the background basis functions exactly using a sufficiently high-order foreground Lagrange space. However, one may wish to reduce the order of the foreground space to reduce the cost of assembling algebraic equations. Such an approximation is especially useful if the foreground elements are non-affine mappings of reference elements to capture the domain geometry with higher-order accuracy. Further, one may consider more radical approximations, like Fig. 1(c), where the foreground mesh is unrelated to the background mesh beyond having a proportional quasi-uniform element size. Notably, Gaussian quadrature on such a foreground mesh would *not* be accurate for the background basis functions (which are non-smooth within foreground elements), but it *is* accurate for their interpolations into a foreground Lagrange function space. Such an interpolation-based immersed method can be viewed as a generalization of the convected particle domain integration (CPDI) variant of the MPM [24] proposed earlier for simulating large-strain solid mechanics, where our foreground elements are analogous to CPDI “particles” that carry their own shape functions. It is also conceptually analogous to the method of free-form deformation (FFD) [25] from the graphics and design literature, but uses the background FFD block’s degrees of freedom to parameterize an approximate PDE solution instead of a geometrical deformation.

The remainder of this article is structured as follows. Section 2 provides a concise list of notation and terminology for reference while reading subsequent sections (and may be skipped on an initial reading). Section 3 introduces the concept of interpolation-based immersed methods using a scalar elliptic model PDE along with a sketch of how a basic error analysis of the method might proceed. Section 4 then describes how to implement interpolation-based immersed analysis for arbitrary PDE systems in a way that efficiently reuses existing FE software. We then explore the accuracy of the method for a variety of PDEs through numerical experiments in Section 5 and apply it to analysis of shell structures modeled geometrically by trimmed spline surfaces in Section 6. Section 7 draws conclusions and outlines future work.

2. Notation and terminology

The use of multiple meshes with different function spaces can lead to ambiguity in terms like “element size” and “basis function”, so we must be careful to distinguish between background, foreground, and interpolated entities through precise terminology and notation. We shall define notations more comprehensively as they arise in context later, but we also summarize them here for convenient reference:

- h : background mesh element size
- η : foreground mesh element size

- N_i : background basis functions
- \widehat{N}_i : interpolated background basis functions
- ϕ_i : foreground basis functions
- k : polynomial degree of the background FE space
- κ : polynomial degree of the foreground FE space
- \widehat{k} : limiting degree $\min\{k, \kappa\}$ of the interpolated background FE space
- \mathcal{P} : interpolation into the foreground FE space
- \mathcal{Q} : interpolation into the background FE space
- $\widehat{\mathcal{Q}}$: interpolation $\mathcal{P} \circ \mathcal{Q}$ into the interpolated background FE space

We also define terms to distinguish the mesh immersions shown in Fig. 1(b) and (c):

- **Background-fitted foreground mesh:** A foreground mesh which is fitted not just to the domain geometry’s boundary but also to interior facets of the background mesh in the sense that the closure of the intersection of a background element with the PDE domain is the closure of a union of foreground elements, as shown in Fig. 1(b).
- **Background-unfitted foreground mesh:** A foreground mesh such as the one shown in Fig. 1(c) whose elements are fitted to the domain geometry but not to interior facets of the background mesh.

We may sometimes use the terms “fitted” and “unfitted” without specifying the prefix “boundary-” or “background-”, but only when the omitted prefix should be clear from context.

3. Immersed finite element analysis for a scalar elliptic model problem

This section defines our proposed interpolation-based immersed method for a model scalar elliptic problem. However, we will apply it to several more complicated problems in subsequent sections. We define our model PDE in Section 3.1, then define a variational formulation suitable for immersed methods in Section 3.2. Section 3.3 reviews the formulation and challenges of quadrature-based immersed methods, and Section 3.4 defines the construction of a function space used for interpolation-based immersed analysis. In Section 3.5 we outline an *a priori* error estimate for the case of a background-fitted foreground mesh.

3.1. The model problem

We choose as our model PDE the Poisson equation with Dirichlet boundary conditions. Let $\Omega \subset \mathbb{R}^d$ be a spatial domain of interest. For simplicity of exposition, we assume that Ω is polyhedral and can be decomposed into affine simplicial elements exactly. The strong form of the problem is as follows: Find $u : \Omega \rightarrow \mathbb{R}$ such that

$$\begin{aligned} -\Delta u &= f && \text{in } \Omega, \\ u &= g && \text{on } \partial\Omega, \end{aligned} \tag{1}$$

where $f : \Omega \rightarrow \mathbb{R}$ is a given source term and $g : \partial\Omega \rightarrow \mathbb{R}$ is given Dirichlet boundary data. The weak form we use as a starting point for formulating numerical methods is: Find $u \in \mathcal{V}_g$ such that

$$\int_{\Omega} \nabla u \cdot \nabla v \, d\Omega = \int_{\Omega} f v \, d\Omega \quad \forall v \in \mathcal{V}_0, \tag{2}$$

where

$$\mathcal{V}_g := \{v \in H^1(\Omega) : v|_{\partial\Omega} = g\}, \tag{3}$$

$$\mathcal{V}_0 := \{v \in H^1(\Omega) : v|_{\partial\Omega} = 0\}, \tag{4}$$

in which restriction is understood in the sense of trace, and f and g are in at least $L^2(\Omega)$ and $H^{1/2}(\partial\Omega)$, respectively.

3.2. Discretization of the model problem

The most common discretization of the weak problem (2) is a conforming one employing finite dimensional subsets of \mathcal{V}_g and \mathcal{V}_0 . However, the construction of these subsets typically relies on a boundary-fitted mesh to

conform to the boundary conditions. Many immersed methods therefore use nonconforming discretizations which enforce the Dirichlet boundary condition in a weak sense. In particular, we employ Nitsche’s method [26] for weak enforcement of Dirichlet boundary conditions, leading to the discrete problem: Find $u^h \in \mathcal{V}^h$ such that, $\forall v^h \in \mathcal{V}^h$,

$$\begin{aligned} \int_{\Omega} \nabla u^h \cdot \nabla v^h d\Omega - \int_{\partial\Omega} (\nabla u^h \cdot \mathbf{n}) v^h d\Gamma \mp \int_{\partial\Omega} (\nabla v^h \cdot \mathbf{n}) u^h d\Gamma + \int_{\partial\Omega} \frac{C_{\text{pen}}}{h} u^h v^h d\Gamma \\ = \int_{\Omega} f v^h d\Omega \mp \int_{\partial\Omega} (\nabla v^h \cdot \mathbf{n}) g d\Gamma + \int_{\partial\Omega} \frac{C_{\text{pen}}}{h} g v^h d\Gamma, \end{aligned} \tag{5}$$

where the discrete space \mathcal{V}^h is a finite-dimensional subspace of $H^1(\Omega)$ with no extra restrictions on $\partial\Omega$, h is a length scale indicating the refinement level of the discretizations (e.g., the element diameter in a standard boundary-fitted discretization), \mathbf{n} is the outward-pointing unit normal vector to $\partial\Omega$, and $C_{\text{pen}} \geq 0$ is a dimensionless constant. The “ \mp ” in (5) must be selected consistently on both sides of the equation, and toggles between the symmetric (“ $-$ ”) and non-symmetric (“ $+$ ”) variants of the Nitsche method. The original symmetric variant proposed by [26] remains well-posed for C_{pen} sufficiently large, with the lower bound emanating from trace-inverse estimates (but often chosen on an *ad hoc* basis in practice). The non-symmetric variant is well-posed for C_{pen} arbitrarily small (or zero), which has led some authors to advocate its use in immersed discretizations [27] to sidestep technical concerns over the appropriate values for C_{pen} and h . The non-symmetric variant retains optimal accuracy in the H^1 norm, which is often more relevant to applications like stress analysis. However, the symmetric variant is optimal in both H^1 and L^2 norms, and is currently more widely used in practice.

3.3. Quadrature-based immersed discretization

In quadrature-based immersed methods, such as those reviewed in Section 1, the discrete space \mathcal{V}^h is defined by simply restricting functions in an FE (or isogeometric) space defined on a background mesh to the PDE domain Ω :

$$\mathcal{V}^h = \text{span} \{N_i|_{\Omega}\}_{i=1}^n, \tag{6}$$

where $\{N_i\}_{i=1}^n$ are the basis functions of the background-mesh FE space whose supports intersect Ω . The principal challenge in implementing such a method is to accurately compute integrals of the forms

$$\int_{\Omega} (\dots) d\Omega = \sum_{E=1}^{n_{\text{el}}} \int_{\Omega_E \cap \Omega} (\dots) d\Omega \quad \text{and} \quad \int_{\partial\Omega} (\dots) d\Gamma = \sum_{E=1}^{n_{\text{el}}} \int_{\partial\Omega \cap \overline{\Omega_E}} (\dots) d\Gamma, \tag{7}$$

where Ω_E is the E th element of the background mesh and n_{el} is the total number of background mesh elements. One can use standard Gaussian quadrature on Ω_E when $\Omega_E \subset \Omega$ or on boundary parts of Ω_E when the background mesh is fitted to $\partial\Omega$. Various schemes have been proposed for integrating nontrivial intersections, but we focus in this paper on those which decompose $\Omega_E \cap \Omega$ into a collection of quadrature elements, as shown in Fig. 1(b), and we refer to the mesh formed by these quadrature elements as a background-fitted foreground mesh. With such a mesh, one can then decompose integrals like

$$\int_{\Omega} (\dots) d\Omega = \sum_{e=1}^{v_{\text{el}}} \int_{\omega_e} (\dots) d\Omega \quad \text{and} \quad \int_{\partial\Omega} (\dots) d\Gamma = \sum_{e=1}^{v_{\text{el}}} \int_{\partial\Omega \cap \partial\omega_e} (\dots) d\Gamma, \tag{8}$$

where $\{\omega_e\}_{e=1}^{v_{\text{el}}}$ are the elements of the foreground mesh and standard Gaussian quadrature rules can be accurately used on each one of its boundary parts.

While mathematically-appealing, this presents a major challenge from a software development standpoint as now the mesh over which quadrature rules are defined is a separate data structure from the mesh on which the basis functions are constructed. Implementation of such methods therefore requires either custom software to be written from scratch or complicated invasive modifications to be made in mature FE or IGA software designed for standard boundary-fitted analysis. The primary motivation of the present work is to circumvent these difficulties while retaining the high-order accuracy of quadrature-based immersed methods.

3.4. Interpolation-based immersed discretization

The main idea of the present work is to use a discrete space that consists of interpolations of background-mesh basis functions on a Lagrange FE space defined over the foreground mesh:

$$\mathcal{V}^h = \text{span} \{ \widehat{N}_i \}_{i=1}^n, \tag{9}$$

where

$$\widehat{N}_i = \mathcal{P}N_i := \sum_{j=1}^v N_i(\mathbf{x}_j)\phi_j, \tag{10}$$

in which \mathcal{P} is the operator that interpolates a function into the Lagrange FE space with basis $\{\phi_j\}_{j=1}^v$ and \mathbf{x}_j is the nodal point associated with foreground Lagrange basis function ϕ_j , i.e.,

$$\phi_i(\mathbf{x}_j) = \delta_{ij}. \tag{11}$$

The space \mathcal{V}^h is a subset of the Lagrange FE space defined on the foreground mesh, so its functions can be accurately integrated using Gaussian quadrature rules on each foreground element. However, unlike in quadrature-based immersed methods, this remains true even with a background-unfitted foreground mesh, like the one shown in Fig. 1(c), where, for an arbitrary background element Ω_E , $\overline{\Omega_E} \cap \overline{\Omega}$ is not necessarily equal to the closure of a union of foreground elements.

For either background-fitted or background-unfitted foreground meshes, the basis $\{\widehat{N}_i\}_{i=1}^n$ exhibits the following properties:

1. If $\{N_i\}_{i=1}^n$ forms a partition of unity, then $\{\widehat{N}_i\}_{i=1}^n$ does as well.
2. If the background function space spans polynomials of degree k and the foreground function space spans polynomials of degree κ , then $\{\widehat{N}_i\}_{i=1}^n$ spans polynomials of degree $\widehat{k} := \min\{k, \kappa\}$.
3. If both $\{N_i\}_{i=1}^n$ and $\{\phi_i\}_{i=1}^v$ have local support, then $\{\widehat{N}_i\}_{i=1}^n$ also has local support.
4. The continuity of $\{\widehat{N}_i\}_{i=1}^n$ is at least that of $\{\phi_i\}_{i=1}^v$. Note that this may be lower than the continuity of $\{N_i\}_{i=1}^n$.

Moreover, for the specific case of a background-fitted foreground mesh, the interpolated background basis $\{\widehat{N}_i\}_{i=1}^n$ is equivalent to the original background basis $\{N_i\}_{i=1}^n$ for a sufficiently high foreground polynomial degree κ . If the background function space consists of piecewise polynomials of degree k defined over a simplicial background mesh, the original background basis is recovered with a Lagrange interpolation of degree $\kappa \geq k$ onto a simplicial background-fitted foreground mesh. If the background function space consists of piecewise tensor-product polynomials of degree k defined over a tensor-product background mesh, the original background basis is recovered with a Lagrange interpolation of degree $\kappa \geq k^d$ onto a simplicial background-fitted foreground mesh.

To better convey the concept of interpolated background basis function, several interpolations of an example background basis function are displayed in Fig. 2. Fig. 2(a) shows the geometric domain and background mesh in red and black respectively, while Fig. 2(b) shows the example background basis function. The background basis function is a bi-quadratic ($k = 2$) B-spline basis function of maximal C^{k-1} continuity defined over the background mesh. Fig. 2(c) and (d) show representative background-fitted and background-unfitted foreground meshes. Fig. 2(e) and (f) show linear ($\kappa = 1$) Lagrange interpolations of the example background basis function onto the background-fitted and background-unfitted foreground meshes. These two interpolated background basis functions share the same shape as the background basis function, but they are both inaccurate approximations of the background basis function. Alternatively, Fig. 2(g) and (h) show quadratic ($\kappa = 2$) Lagrange interpolations of the background basis function onto the background-fitted and background-unfitted foreground meshes. These two interpolated background basis functions are much more accurate approximations of the background basis function than the interpolated background basis functions displayed in Fig. 2(e) and (f).

3.5. Best approximation estimates for the background-fitted case

Having presented our new paradigm for immersed finite element and isogeometric analysis, we now derive estimates for the H^1 -norm best approximation error

$$\inf_{v^h \in \text{span}\{\widehat{N}_i\}_{i=1}^n} \|u - v^h\|_{H^1(\Omega)} \tag{12}$$

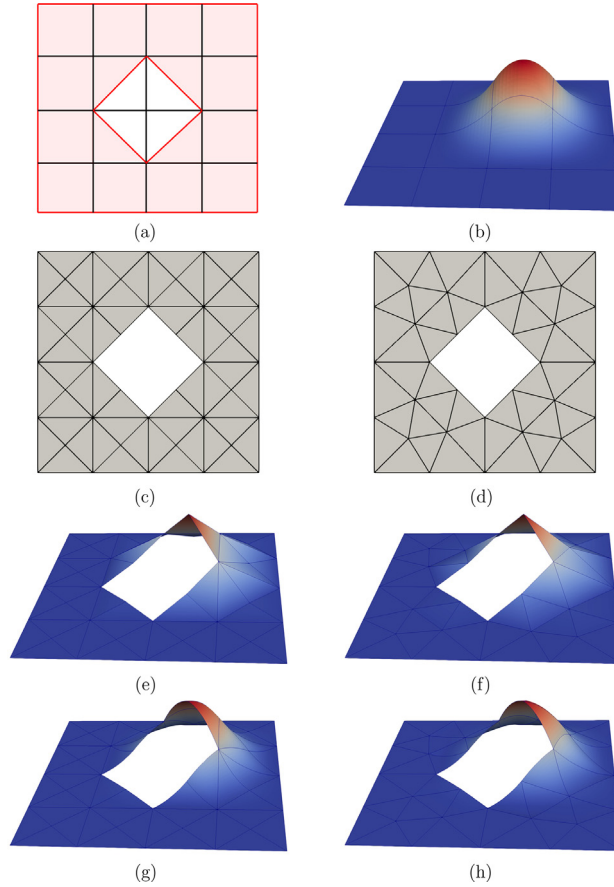


Fig. 2. (a) Geometric domain (red) and background mesh (black). (b) Background basis function. (c) Background-fitted foreground mesh. (d) Background-unfitted foreground mesh. (e) Linear Lagrange interpolation of background basis function on background-fitted mesh. (f) Linear Lagrange interpolation of background basis function on background-unfitted mesh. (g) Quadratic Lagrange interpolation of background basis function on background-fitted mesh. (h) Quadratic Lagrange interpolation of background basis function on background-unfitted mesh. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

associated with the interpolated background FE space where u is the exact solution of our model problem. Best approximation estimates of the form given by (12) can be combined with coercivity and continuity results to arrive at *a priori* error estimates for a given immersed finite element method (as done, e.g., for quadrature-based CutFEM with ghost penalty stabilization in [8]). In our mathematical analysis, we must make two assumptions regarding the foreground mesh:

Assumption 1. The foreground mesh is background-fitted.

Assumption 2. The foreground mesh is free of hanging nodes.

Assumption 1 ensures that the background basis functions are smooth within every element of the foreground mesh, while Assumption 2 ensures that the interpolated background FE space is H^1 -conforming and thus the H^1 -norm best approximation error is well-defined. Note that an immediate consequence of Assumption 1 is that the foreground mesh size η is less than or equal to the background mesh size h . We must make two further assumptions regarding the exact solution u and the background FE space:

Assumption 3. The exact solution u lies in the space $H^{\widehat{k}+1}(\Omega)$.

Assumption 4. There exists an interpolation operator $\mathcal{Q} : H^1(\Omega) \rightarrow \text{span}\{N_i\}_{i=1}^n$ onto the background FE space such that

$$\|v - \mathcal{Q}v\|_{H^1(\Omega)} \leq C_1 h^{\hat{k}} \|v\|_{H^{\hat{k}+1}(\Omega)} \tag{13}$$

and

$$\sum_{e=1}^{v_{el}} \|\mathcal{Q}v\|_{H^{\hat{k}+1}(\omega^e)}^2 \leq C_2 \|v\|_{H^{\hat{k}+1}(\Omega)}^2 \tag{14}$$

for all $v \in H^{\hat{k}+1}(\Omega)$ where $C_1 > 0$ and $C_2 > 0$ depend on the background FE space polynomial degree k and the shape regularity of the background mesh but not the background mesh element size h .

Assumption 3 holds provided the problem geometry and parameters are sufficiently smooth, while **Assumption 4** holds when the background FE space consists of classical finite element functions or smooth spline functions provided there exists a continuous extension operator $E : H^{k+1}(\Omega) \rightarrow H^{k+1}(\mathbb{R}^d)$ (cf. [8, (9)]).

Now let us define $\tilde{\mathcal{Q}} = \mathcal{P} \circ \mathcal{Q}$. Then we can bound the H^1 -norm best approximation error associated with the interpolated background FE space by the H^1 -norm error associated with the interpolation operator $\tilde{\mathcal{Q}}$:

$$\inf_{v^h \in \text{span}\{\tilde{N}_i\}_{i=1}^n} \|u - v^h\|_{H^1(\Omega)} \leq \|u - \tilde{\mathcal{Q}}u\|_{H^1(\Omega)}. \tag{15}$$

By the triangle inequality,

$$\|u - \tilde{\mathcal{Q}}u\|_{H^1(\Omega)} \leq \|u - \mathcal{Q}u\|_{H^1(\Omega)} + \|\mathcal{Q}u - \mathcal{P}\mathcal{Q}u\|_{H^1(\Omega)}. \tag{16}$$

Invoking **Assumption 3**, we can immediately bound the first term on the right hand side of the above equation:

$$\|u - \mathcal{Q}u\|_{H^1(\Omega)} \leq C_1 h^{\hat{k}} \|u\|_{H^{\hat{k}+1}(\Omega)}. \tag{17}$$

To bound the second term, we first recognize that

$$\|\mathcal{Q}u - \mathcal{P}\mathcal{Q}u\|_{H^1(\Omega)}^2 = \sum_{e=1}^{v_{el}} \|\mathcal{Q}u - \mathcal{P}\mathcal{Q}u\|_{H^1(\omega^e)}^2. \tag{18}$$

By **Assumption 1**, we have that $\mathcal{Q}u|_{\omega^e}$ is smooth for every foreground mesh element ω^e . Thus we have that

$$\|\mathcal{Q}u - \mathcal{P}\mathcal{Q}u\|_{H^1(\omega^e)}^2 = C_3 \eta^{2\hat{k}} \|\mathcal{Q}u\|_{H^{\hat{k}+1}(\omega^e)}^2 \tag{19}$$

for every foreground element ω^e where C_3 is a constant that depends on the foreground FE space polynomial degree κ and the shape regularity of the foreground mesh but not the foreground mesh element size η . We can immediately combine (18), (19), **Assumptions 1**, and **4** to arrive at

$$\|\mathcal{Q}u - \mathcal{P}\mathcal{Q}u\|_{H^1(\Omega)} = \sqrt{C_2 C_3} h^{\hat{k}} \|u\|_{H^{\hat{k}+1}(\Omega)}. \tag{20}$$

Combining (15), (16), (17), and (18), we finally arrive at the following estimate for the H^1 -norm best approximation error associated with the interpolated background FE space:

$$\inf_{v^h \in \text{span}\{\tilde{N}_i\}_{i=1}^n} \|u - v^h\|_{H^1(\Omega)} \leq C_{\text{interp}} h^{\hat{k}} \|u\|_{H^{\hat{k}+1}(\Omega)} \tag{21}$$

where $C_{\text{interp}} = C_1 + \sqrt{C_2 C_3}$. It is immediately evident that the above estimate is optimal with respect to both the background mesh size h and the limiting degree \hat{k} of the interpolated background space. Note that the convergence rate does not improve by decreasing the foreground mesh size η faster than the background mesh size h or increasing the foreground polynomial degree κ above that of the background polynomial degree k .¹

In the above analysis, we made the rather stringent assumption that the foreground mesh be background-fitted. However, our later numerical experiments suggest that best approximation estimates of the form given by (21) also hold when the foreground mesh is background-unfitted. This is perhaps not too unexpected as the interpolated background FE space is complete up to polynomial degree \hat{k} even when the foreground mesh is background-unfitted,

¹ If there is geometry error in the foreground mesh, then there can be some benefit to over-refining η near boundaries (cf. Section 5.3), but geometry error is outside the scope of the analysis in this section.

but our analysis does not extend to this case since we require that the background basis functions be smooth over each element of the foreground mesh.

We also made the assumption in our analysis that the foreground mesh be free of hanging nodes. This was to ensure that the interpolated background FE space was H^1 -conforming. That being said, our later numerical experiments suggest we may not need an interpolated background FE space that is H^1 -conforming provided that the original background FE space is H^1 -conforming. In particular, we are able to attain optimal convergence rates for the biharmonic problem using an interpolated background FE space that is only H^1 -conforming when the original background FE space is H^2 -conforming. For the analysis of non-conforming methods, we note that a similar interpolation estimate to (21) could be derived element-wise, using a weakened version of Assumption 2, where the mesh only needs to be free of hanging nodes within each background element. However, we have focused here on the globally $H^1(\Omega)$ -conforming case, for a more direct analogy to prior convergence analysis of quadrature-based immersed methods.

Finally, the best approximation estimate given by (21) depends on the shape regularity of the foreground mesh through the interpolation constant C_{interp} , but our later numerical experiments suggest that high quality results are attained with interpolation-based immersed finite element and isogeometric analysis even when the foreground mesh is of poor quality.

4. Reusing FE software for interpolation-based immersed methods

As each basis function \widehat{N}_i of \mathcal{V}^h in our proposed method is a linear combination of basis functions defined on the foreground mesh used for element-by-element assembly, one may implement interpolation-based immersed methods using existing FE software in a similar manner as one would implement IGA via Lagrange extraction [20]. The cited paper emphasizes an implementation based on modifying element-level shape function routines which could also be followed for immersed methods. However, the method we shall outline here is based on global linear algebra operations applied to the full, assembled system of equations, similar to what is done in [21,22] for IGA. The latter approach requires fewer modifications to the existing FE code and is referred to as “non-invasive” by [22].

4.1. Non-invasive implementation using extraction matrices

Consider a scalar variational problem of the form: Find $u^h \in \mathcal{V}^h$ such that, $\forall v^h \in \mathcal{V}^h$,

$$a(u^h, v^h) = L(v^h) \quad , \tag{22}$$

where a and L are a bilinear and linear form, i.e., the left- and right-hand sides of the Nitsche formulation (5) for the Poisson model problem. $\mathcal{V}^h = \text{span}\{\widehat{N}_i\}$ is the solution space for the interpolation-based immersed method. The state variable field is discretized by the interpolated basis functions as

$$u^h = \sum_{i=1}^n d_i \widehat{N}_i \quad , \tag{23}$$

where $\{d_i\}_{i=1}^n$ are unknown coefficients. Then we need to solve the linear system

$$\mathbf{Kd} = \mathbf{F} \quad , \tag{24}$$

where the global stiffness matrix and forcing vector are

$$K_{ij} = a(\widehat{N}_j, \widehat{N}_i) \quad \text{and} \quad F_i = L(\widehat{N}_i) \quad . \tag{25}$$

In the interpolation-based immersed method,

$$\widehat{N}_j = \sum_{i=1}^v M_{ij} \phi_i \quad \text{with} \quad M_{ij} = N_j(\mathbf{x}_i) \quad . \tag{26}$$

M_{ij} projects the background functions onto the space of the foreground elements and can be obtained by evaluating the background basis N_j at the nodal points of the foreground mesh \mathbf{x}_i . Using the (bi-) linearity of a and L , we can re-write (25) as

$$K_{ij} = \sum_{k,\ell} M_{ki} a(\phi_\ell, \phi_k) M_{\ell j} \quad \text{and} \quad F_i = \sum_k M_{ki} L(\phi_k) \quad . \tag{27}$$

In matrix form, this reads as

$$\mathbf{K} = \mathbf{M}^T \mathbf{A} \mathbf{M} \quad \text{and} \quad \mathbf{F} = \mathbf{M}^T \mathbf{B} , \quad (28)$$

where

$$A_{ij} = a(\phi_j, \phi_i) \quad , \quad B_i = L(\phi_i) , \quad (29)$$

We will refer to \mathbf{M} as the “extraction matrix”. With some additional straightforward bookkeeping, this method extends to linear PDE systems with multiple scalar solution fields, as presented for extraction-based IGA in [21]. It can also be applied to each step of Newton iteration to solve a nonlinear problem, where a is the Gateaux derivative of a residual form $-L$.

Crucially, the matrix \mathbf{A} and vector \mathbf{B} can be assembled using a standard boundary-fitted FE code in a loop over elements in the foreground mesh. Thus, given a foreground mesh and an extraction matrix \mathbf{M} , one can directly reuse an existing FE code for interpolation-based immersed analysis without making any modifications to the assembly procedures. The post processing features of existing FE codes can also be reused by leveraging a foreground FE representation of the solution:

$$u^h = \sum_{j=1}^v c_j \phi_j , \quad (30)$$

where the foreground FE coefficients $\{c_i\}$ are obtained from the solution \mathbf{d} of (24) by

$$\mathbf{c} = \mathbf{M} \mathbf{d} . \quad (31)$$

4.2. Example implementation using FEniCS

For the numerical results in this paper, we apply the method outlined in Section 4.1 to the existing FE automation toolchain FEniCS [17]. FEniCS allows arbitrary linear and bilinear forms to be specified over a variety of classical FE function spaces using the Python-based Unified Form Language (UFL) [28]. UFL forms are compiled into efficient element-level routines [29] which are used by the FE library DOLFIN [30] to assemble matrices and vectors. In the notation of Section 4.1, UFL would be used to specify the forms a and L , and DOLFIN would be used to assemble the matrix \mathbf{A} and vector \mathbf{B} , applying the compiled element-level routines to elements $\{\omega_e\}_{e=1}^{v_{el}}$ of the foreground mesh. DOLFIN’s application–programmer interface (API) provides access to the linear algebra objects \mathbf{A} and \mathbf{B} as data structures in the parallel linear algebra library PETSc [31–33], which is used to form \mathbf{K} and \mathbf{F} via (28), solve the linear system (24) for \mathbf{d} , and compute the foreground FE coefficients \mathbf{c} via (31). DOLFIN’s output and visualization routines can then be applied to this foreground FE representation of the solution.

This non-invasive reuse of FEniCS provides great flexibility for implementing formulations of a variety of different PDE systems as we illustrate in the sequel. What remains is to generate the foreground- and background meshes alongside the extraction operators \mathbf{M} . The tools and methods used to do so are outlined in the following subsection.

4.3. Mesh generation

This section describes the techniques used in the numerical examples of Sections 5 and 6 to generate both background-fitted and -unfitted foreground meshes.

4.3.1. Background-fitted foreground meshes

To generate background-fitted foreground meshes and their extraction matrices \mathbf{M} , the XFEM Tool Kit of the research code presented in [34] was used. Within this toolkit, the foreground elements are created by repeated subdivision of a structured background grid. First, each background cell is subdivided into a set of triangular or tetrahedral elements as is shown in Fig. 3(b). Then, all edges of the elements are checked for intersections with the geometry. Using a set of subdivision templates, elements with intersected edges are then again subdivided into straight-edged triangular and tetrahedral elements such that new element vertices are created at the intersection locations. The new element facets (i.e. edges in 2D and faces in 3D) are a first-order geometric approximation of the domain boundary as shown in Fig. 3(c).




Fig. 3. Generation of background fitted foreground meshes. (a) Single background element with description of the boundary shown by a red curve; (b) regular subdivision of the background element into triangles; (c) subdivision of triangles based on intersections of the interface with edges; interface approximated by red lines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

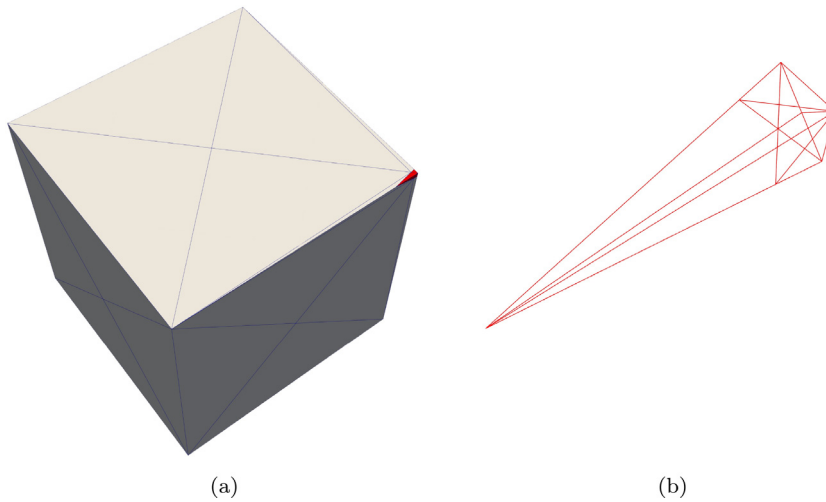


Fig. 4. (a) A tessellation of a hexahedral background element from the 3D mesh shown in Fig. 5, highlighting the sliver element in the rightmost corner. (b) Only the red section consisting of small tetrahedral elements intersects the PDE domain Ω . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The XFEM tool kit constructs the extraction matrices \mathbf{M} such that basis functions are enriched following a Heaviside enrichment strategy [34]. Enrichment is needed when the support of a basis contains two or more geometrically disconnected domains. Note that for the numerical examples studied in this paper enrichment is not required to ensure optimal convergence.

While the procedure outlined enables the robust background-fitted triangulation of the material domain, it may lead to elements of arbitrarily small size and arbitrarily large aspect ratio. As an example, the 2D foreground meshes used in Section 6.2 contain elements with aspect ratios as high as ~ 300 and elements with volumes as small as $\sim 10^{-4}$ of the biggest elements. The poor mesh quality is exacerbated in 3D. The 3D foreground meshes used for the example in Section 5.1 contain elements that differ by a factor of $\sim 10^{27}$ in volume, and aspect ratios as high as $\sim 10^{29}$. An example for the poor elements resulting from this mesh generation process in 3D is shown in Fig. 4.

4.3.2. Background-unfitted foreground meshes

Mesh cutting tools, like the XFEM tool kit discussed in Section 4.3.1, generate background-fitted meshes and do not require an additional mesh generator to create foreground meshes. However, background-unfitted analysis allows substantially greater flexibility in the generation of foreground meshes; essentially any existing mesh generator can be used, with the only necessary information from the background mesh being an approximate element size. We apply two different approaches within the examples shown in this paper, with the particular details largely motivated by convenience. The background-unfitted foreground meshes used in Section 5.5 were generated using FEniCS’s built in structured mesh functions, as were the corresponding background meshes. For these examples,

the background function spaces were C^0 Lagrange FE spaces, and the extraction operators were generated directly in FEniCS, using its `PETScDMCollection` functionality. The background-unfitted foreground meshes used in Sections 6.1 and 6.2 use an isogeometric background basis. These meshes were created with FEniCS's `mshr` module, which generates unstructured simplicial meshes from geometric descriptions. The extraction operators were generated using the FEniCS-based isogeometric analysis library `tIGAr` [21] by simply overriding methods that normally generate structured meshes for exact Lagrange extraction.

5. Numerical results

We now explore the properties of interpolation-based immersed analysis numerically using a variety of benchmark problems.

5.1. Comparing quadrature-based and interpolation-based methods

The first question we seek to answer is: How do interpolation-based immersed methods with background-fitted foreground meshes compare with the more thoroughly-studied class of quadrature-based immersed methods? (We defer the comparison of background-fitted and -unfitted foreground meshes to Section 5.5.) In particular, we use the Poisson equation as a model problem and consider a background function space of tensor-product B-splines of degree k (at maximal continuity) and a foreground Lagrange FE space of equal polynomial degree $\kappa = k$ on a simplicial mesh. The background function space contains monomials of degree greater than k which cannot be represented exactly by the foreground function space, so interpolation-based methods differ from quadrature-based ones.

To carry out this test, we construct an instance of the Poisson problem using the method of manufactured solutions, where we choose the exact solutions

$$u(\mathbf{x}) = \sin(\pi(x_1^2 + x_2^2)) \cos(\pi(x_1 - x_2)) \quad (32)$$

and

$$u(\mathbf{x}) = \sin(\pi(x_1^2 + x_2^2 + x_3^2)) \cos(\pi(x_1 + x_2 + x_3)) \quad (33)$$

in 2D and 3D respectively, and define the source term as $f := -\Delta u$. For the 2D problem, Ω is taken to be a unit square rotated by 45° about the origin (i.e., $\Omega = \{\mathbf{x} \in \mathbb{R}^2 : \|\mathbf{x}\|_{\ell^1} < 1/2\}$), and the manufactured solution u is applied as Dirichlet data on the boundary of this rotated square. Ω is then immersed into a structured background mesh of an axis-aligned bi-unit square on which the background B-spline spaces are constructed. A suite of progressively-refined pairs of background and background-fitted foreground meshes are used with background element size $h = 2^{-(R+1)}$, where $R \in \{0, \dots, 6\}$ denotes the refinement level and “ h ” is the square root of background element area.

For the 3D case, we take Ω to be an analogous rotated unit cube (rotated by 45° about the x_3 axis and then by 45° about the x_2 axis) and define a sequence of structured hexahedral background meshes with element size $h = 2^{-(R+1)}$, where $R \in \{0, \dots, 4\}$ is the refinement level and h is the cube root of the element volume. The foreground meshes in 3D are background-fitted tetrahedral meshes.

Dirichlet boundary conditions are enforced weakly using the non-symmetric Nitsche formulation (5), with $C_{\text{pen}} = 0$. A representative foreground mesh and numerical solution are shown in Fig. 5.

Figs. 6 and 7 compare the convergence of interpolation-based and quadrature-based immersed calculations.² Both types of immersed method are seen to converge optimally for $k = 1$ and $k = 2$ in both H_0^1 and L^2 norms. We also include convergence results of standard FE analysis using the Lagrange FE space defined on the foreground mesh. This also converges optimally with respect to h , but is less accurate per degree of freedom due to excess foreground mesh refinement in cut background elements and less-efficient approximation of smooth functions than the maximally-smooth B-splines of the background function space.

² The interpolation-based calculations use the FEniCS-based implementation described in Section 4, whereas the quadrature-based calculations use a different research code with the requisite functionality which is outlined in [34].

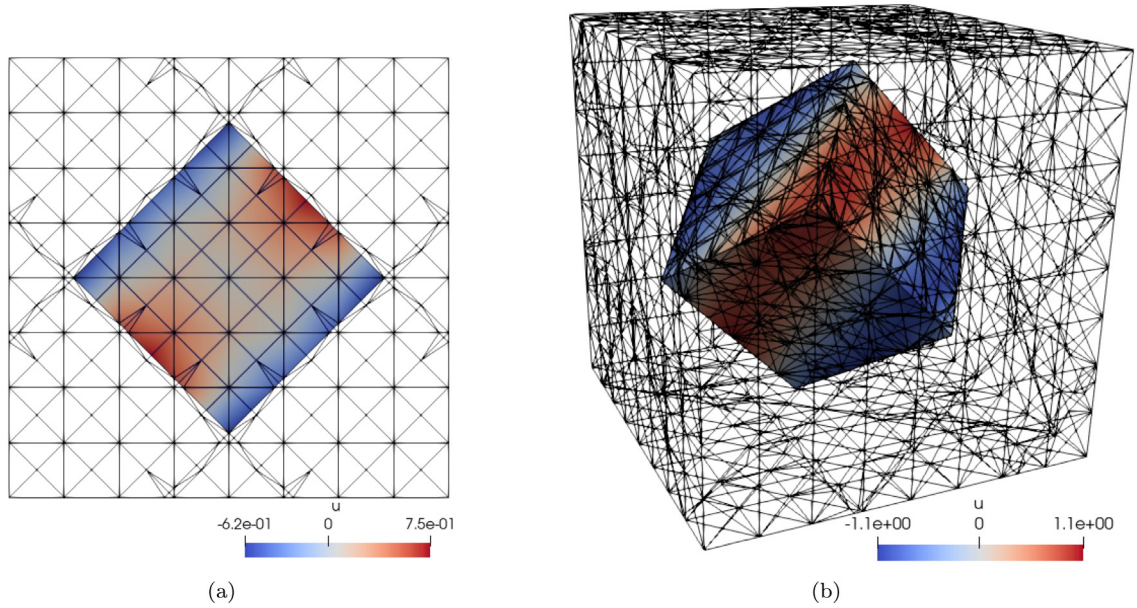


Fig. 5. Numerical solution to the Poisson problem using interpolation-based immersed analysis, visualized on the foreground mesh of refinement level $R = 2$. Tests were performed on both the 2D square (a) and 3D rotated cube (b).

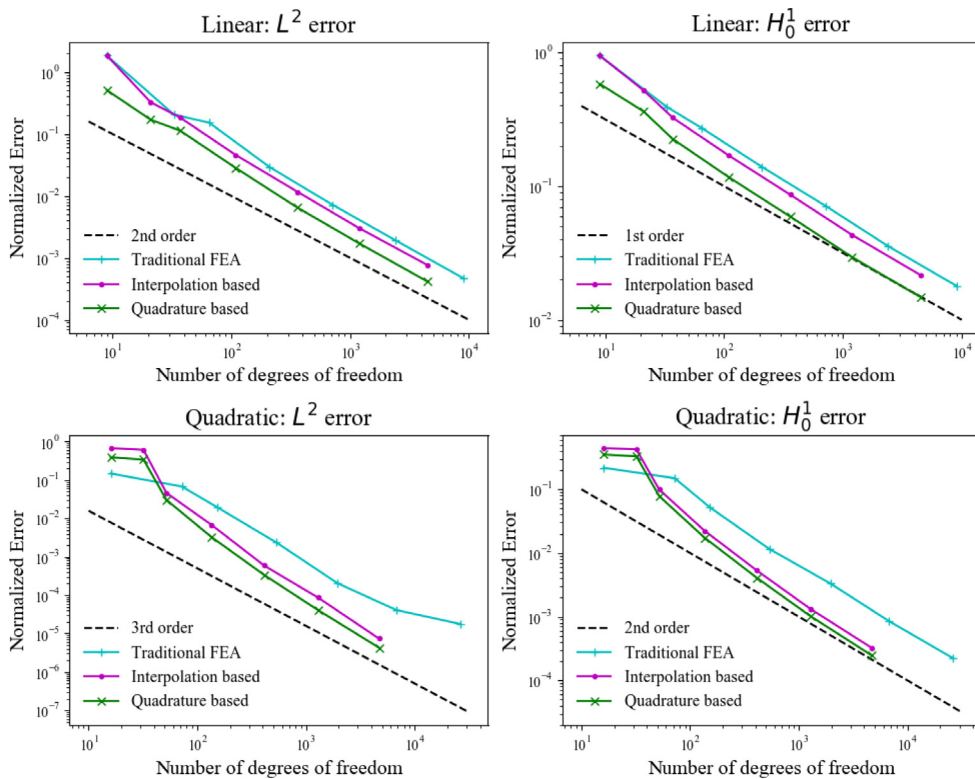


Fig. 6. Convergence data for Poisson's problem in 2D, comparing the results of traditional FEA on the foreground mesh with the Lagrange FE space, the interpolation based immersed method developed in this paper, and the quadrature based immersed method described in Section 3.3.

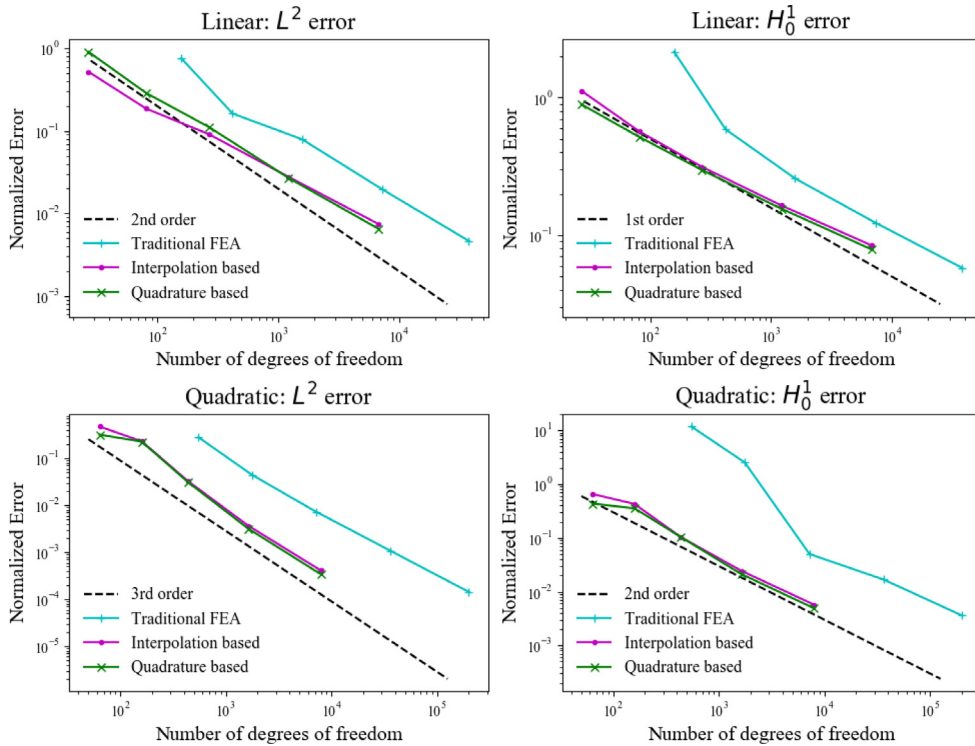


Fig. 7. Convergence data for Poisson’s problem in 3D, comparing the results of traditional FEA on the foreground mesh with the Lagrange FE space, the interpolation based immersed method developed in this paper, and the quadrature based immersed method described Section 3.3.

5.2. Non-conforming background spaces

The numerical analysis sketched in Section 3.5 made the assumption of a conforming discrete space, i.e., that the discrete space \mathcal{V}^h of interpolated background functions was a subset of the infinite dimensional space used to define the continuous variational problem (viz., H^1 in the case of the Poisson equation). However, one may wish to relax that requirement in practice to eliminate topological constraints on the foreground mesh and/or interpolate smooth background functions inexactly with less-regular foreground FE functions.

As a model problem to illustrate non-conforming interpolation-based immersed analysis, we consider the biharmonic problem: Find u such that

$$\Delta^2 u = f, \tag{34}$$

with boundary conditions

$$u = \sigma \text{ on } \partial\Omega, \tag{35}$$

$$\nabla u \cdot \mathbf{n} = \nabla \sigma \cdot \mathbf{n} \text{ on } \partial\Omega, \tag{36}$$

where $f : \Omega \rightarrow \mathbb{R}$ is a given source term and $\sigma : \Omega \rightarrow \mathbb{R}$ is an auxiliary function used to define Dirichlet boundary data for u and $\nabla u \cdot \mathbf{n}$ on $\partial\Omega$.³ Unlike the nonsymmetric Nitsches method used in the Poisson problem in (5), we

³ One could instead define separate functions to prescribe boundary data for u and its normal derivative, but an equivalent σ must exist for the problem to be well-posed and the present perspective is more convenient for applying the method of manufactured solutions.

discretize this problem using a symmetric Nitsche-like method: Find $u^h \in \mathcal{V}^h$ such that, $\forall v^h \in \mathcal{V}^h$,

$$\begin{aligned} & \int_{\Omega} \Delta u^h \Delta v^h d\Omega + \int_{\partial\Omega} \nabla \Delta u^h \cdot \mathbf{n} v^h - \Delta u^h \nabla v^h \cdot \mathbf{n} d\Gamma \\ & + \int_{\partial\Omega} (\nabla \Delta v^h) \cdot \mathbf{n} (u^h - \sigma) - \Delta v (\nabla u^h \cdot \mathbf{n} - \nabla \sigma \cdot \mathbf{n}) d\Gamma \\ & + \int_{\partial\Omega} \frac{\alpha}{h^3} (u^h - \sigma) v^h + \frac{\beta}{h} (\nabla u^h \cdot \mathbf{n} - \nabla \sigma \cdot \mathbf{n}) \nabla v^h \cdot \mathbf{n} d\Gamma = \int_{\Omega} f v^h d\Omega, \end{aligned} \tag{37}$$

where $\alpha > 0$ and $\beta > 0$ are independent of h and ϵ , but must be sufficiently large to ensure stability of the formulation.⁴ In the computations of this paper, we use $\alpha = \beta = 5$. In the case of immersed methods, the integrals $\int_{\Omega} (\dots) d\Omega$ and $\int_{\partial\Omega} (\dots) d\Gamma$ should be understood as being split up over elements of the foreground mesh, as in (7). This distinction is important when non-conforming spaces are used as some spatial derivatives may have distributional contributions at interior mesh facets that are ignored, so (37) technically involves some abuse of notation.

As with the Poisson problem in Section 5.1, we use the method of manufactured solutions, choosing exact solutions

$$u(\mathbf{x}) = \cos(0.05\pi x_1 + 0.1) \cos(0.05\pi x_2 + 0.1) \tag{38}$$

and

$$u(\mathbf{x}) = \cos(\pi x_1 + 0.5) \cos(\pi x_2 + 0.5) \cos(\pi x_3 + 0.5) \tag{39}$$

in 2D and 3D, respectively, and setting the problem data to $f := \Delta^2 u$ and $\sigma := u$.

We reuse the rotated square and cube domains and mesh sequences from Section 5.1. We select the *background* function spaces to be H^2 -conforming. In particular, we use quadratic B-splines of maximal continuity. However, the *foreground* function spaces are taken to be quadratic Lagrange FE spaces, which, on simplicial meshes, do not represent all multivariate monomials of the background B-spline spaces. Further, interpolations of background functions on this Lagrange FE space are only C^0 , containing jumps in their first derivatives (so they can be at most $H^{3/2-\epsilon}$ -conforming, for $\epsilon > 0$ arbitrarily-small). Interpolation-based immersed analysis will therefore be non-conforming in this case. 2D and 3D convergence results are shown in Fig. 8. We get optimal rates of 1 and 2 for the H^2 and H^1 norms respectively, where “ H^2 norm” is understood in a broken sense (i.e., the square root of a sum over squared foreground-element H^2 norms) in this non-conforming setting. We get only a rate of 2 for the L^2 norm, but this is expected, since an Aubin–Nitsche-type duality argument to obtain an optimal L^2 error bound breaks down for quadratic B-spline discretizations of the biharmonic problem [35].

Remark 1. The 3D foreground meshes used in this section include some elements where $\eta \ll h$, such as the example shown in Fig. 4. The harmful effects of small background element–domain intersections on linear system conditioning have been studied previously in the context of quadrature-based immersed methods, along with a variety of possible remedies, but we leave the adaptation of such techniques to the interpolation-based setting to future work, as discussed further in Section 7. However, small foreground elements, where $\eta \rightarrow 0$, present a distinct challenge unique to interpolation-based methods, due to the blow-up of foreground basis function derivatives.⁵ The resulting large numbers in the matrix \mathbf{A} defined by (29) should, in exact arithmetic, be canceled by small entries in the extraction matrix \mathbf{M} when forming \mathbf{K} by (28). However, this cancellation is not exact due to finite precision. For the fourth-order biharmonic problem in 3D, we find that this can prevent convergence even with direct solvers. In the 3D biharmonic examples from the present work, we circumvent the issue by removing near-zero volume elements from the assembly of \mathbf{A} , skipping integration over foreground elements (and their associated boundary facets) if their volume falls below 0.001% of the largest foreground element. The resulting geometric perturbation of the domain has no observable effect on solutions, as illustrated by the convergence results of Fig. 8. One can imagine pathological domain–mesh intersections where such a strategy would lead to significant errors, but these do not arise in our tests, and the solution of fourth-order problems on topologically-3D domains is itself relatively uncommon in applications.

⁴ For the method to be usable with piecewise-quadratic function spaces, α must be strictly positive, even if a non-symmetric variant is used, as $\nabla \Delta v^h \equiv 0$ and only the penalty term is left to enforce $u^h \approx \sigma$.

⁵ This blow-up does not manifest in quadrature-based methods since they directly evaluate background basis function derivatives.




Fig. 8. Convergence data for the biharmonic problem.

5.3. Geometric approximations of domain geometry

The convergence tests of Sections 5.1 and 5.2 deliberately used polyhedral PDE domains so that a foreground mesh of affine simplex elements could fit them exactly, allowing us to isolate errors due to interpolation of the background basis functions. We now consider a problem with a curved boundary where there is some geometry error between Ω and the foreground mesh.

In particular, we choose the classic linear elasticity benchmark of an infinite 2D plate with a circular hole subjected to biaxial tension. We consider the plane strain variant of the problem. An exact solution for the stress field is readily available in elasticity textbooks, e.g., [36, Section 7.7.5]. We truncate the problem to a finite square domain whose side length is four times the diameter of the hole and apply the traction from the exact solution to the infinite-domain problem on the boundary. This finite-domain problem is further truncated by restricting it to the upper-right quadrant of the x_1 - x_2 plane and applying symmetry boundary conditions on the axes, which eliminates rigid-body modes from the displacement solution of the original pure Neumann problem. We formulate linear elasticity as a problem for displacement, viz., Navier’s equations, which also provide an opportunity to demonstrate interpolation-based immersed analysis for a system of PDEs: Find a displacement field $\mathbf{u} : \Omega \rightarrow \mathbb{R}^2$ such that

$$-\nabla \cdot \boldsymbol{\sigma}(\mathbf{u}) = \mathbf{0} \tag{40}$$

subject to boundary conditions

$$\mathbf{u} \cdot \mathbf{n} = 0 \quad \text{on } \Gamma_{\text{sym}} , \tag{41}$$

$$(\mathbf{I} - \mathbf{n} \otimes \mathbf{n})(\boldsymbol{\sigma} \cdot \mathbf{n}) = \mathbf{0} \quad \text{on } \Gamma_{\text{sym}} , \tag{42}$$

$$\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{t} \quad \text{on } \Gamma_t , \tag{43}$$

where $\Gamma_{\text{sym}} \subset \partial\Omega$ is the union of the symmetry planes where sliding boundary conditions are applied, $\Gamma_t = \partial\Omega \setminus \Gamma_{\text{sym}}$ is the Neumann boundary on which traction is prescribed, \mathbf{t} is the traction on Γ_t taken from the exact solution, and the Cauchy stress $\boldsymbol{\sigma}$ is

$$\boldsymbol{\sigma}(\mathbf{u}) = 2\mu\boldsymbol{\varepsilon}(\mathbf{u}) + \lambda \text{tr}(\boldsymbol{\varepsilon}(\mathbf{u}))\mathbf{I} \tag{44}$$

in terms of the symmetric gradient operator $\boldsymbol{\varepsilon}(\cdot)$ (which gives strain when applied to the displacement) and Lamé parameters μ and λ , which we set from Young’s modulus $E = 200 \times 10^9$ and Poisson’s ratio $\nu = 0.3$ by standard formulas.

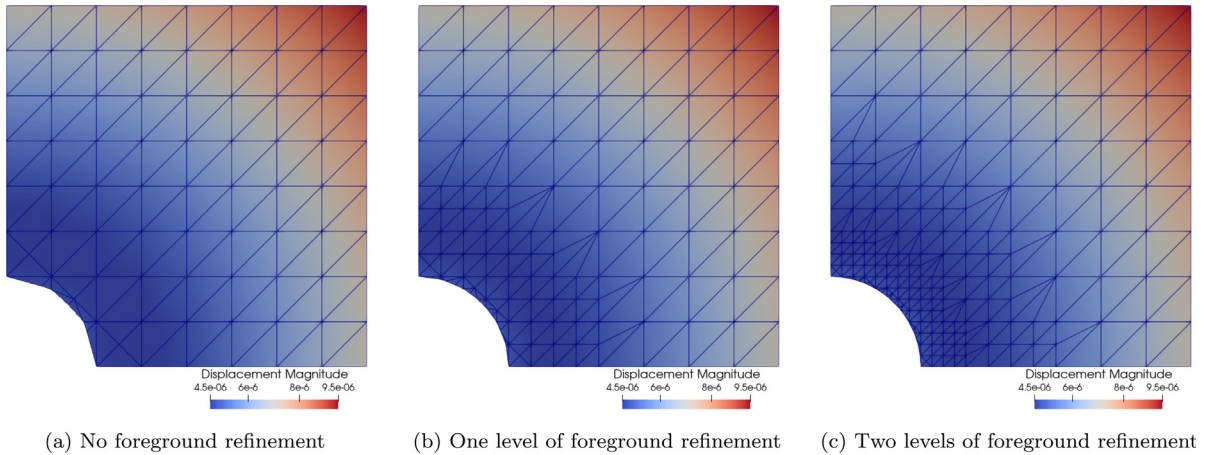


Fig. 9. Plots showing the magnitude of displacement. To resolve curved surfaces, the foreground mesh is locally refined, while the same uniform background mesh is used.

We discretize Navier’s equations with slip boundary conditions using a symmetric Nitsche-like formulation analogous to (5): Find $\mathbf{u}^h \in \mathcal{V}^h$ such that, $\forall \mathbf{v}^h \in \mathcal{V}^h$,

$$\begin{aligned}
 & \int_{\Omega} \boldsymbol{\sigma}(\mathbf{u}^h) : \nabla \mathbf{v}^h \, d\Omega \\
 & - \int_{\Gamma_{\text{sym}}} (\mathbf{u}^h \cdot \mathbf{n}) \mathbf{n} \cdot \boldsymbol{\sigma}(\mathbf{v}^h) \cdot \mathbf{n} \, d\Gamma - \int_{\Gamma_{\text{sym}}} (\mathbf{v}^h \cdot \mathbf{n}) \mathbf{n} \cdot \boldsymbol{\sigma}(\mathbf{u}^h) \cdot \mathbf{n} \, d\Gamma \\
 & + \int_{\Gamma_{\text{sym}}} \frac{\beta \mu}{h} (\mathbf{u}^h \cdot \mathbf{n}) \mathbf{n} \cdot \mathbf{v}^h \, d\Gamma = \int_{\Gamma_t} \mathbf{t} \cdot \mathbf{v}^h \, d\Gamma, \tag{45}
 \end{aligned}$$

where $\beta > 0$ is a penalty parameter associated with Nitsche’s method (in this work we set $\beta = 10$) and integrals should be understood in the sense of (7) (which is technically an abuse of notation in the present case because it ignores the geometry error in the foreground mesh). Similarly, the outward-facing unit normal \mathbf{n} depends on the geometry of the foreground mesh, but this dependence has been suppressed in the notation.

To study the behavior of interpolation-based immersed methods, we construct a sequence of uniform, structured background meshes of quadrilateral elements, covering the truncated plate and hole. We consider several corresponding sequences of foreground meshes, with different levels of local refinement around the curved boundary of the hole. The three different foreground mesh refinement levels are shown for an 8×8 background mesh in Fig. 9, along with representative displacement solutions. In all cases, the foreground mesh elements are affine triangles. For the background function spaces, we select B-spline spaces of degrees $k = 1$ and $k = 2$, with maximal continuity for each, and, for the foreground function spaces, we use Lagrange FE spaces of corresponding polynomial degrees $k = \kappa$.

Fig. 10 shows the convergence of the L^2 norm of error in the Cauchy stress for different polynomial degrees and levels of foreground refinement. We see that, even without foreground refinement the case of $k = 1$ attains the optimal convergence rate. For $k = 2$, geometry error dominates in the case without foreground refinement and limits the rate of convergence. However, even with just one level of foreground refinement near the curved boundary, we see an optimal effective convergence rate. For affine foreground elements and a fixed level of local foreground refinement, we still expect the geometry error to dominate in the ultimate asymptotic limit (as occurs also with the even lower-order stair-step approximations of curved boundaries in finite cell adaptive quadrature), but this phenomenon is unlikely to be encountered in practical problems where the model error would dominate well before this asymptotic regime. We hypothesize that the use of high-order isoparametric elements instead of local foreground refinement would provide true optimal convergence, but high-order foreground mesh generation is currently outside the capabilities of the software framework used for this paper.




Fig. 10. Frobenius norm of stress error tensor plotted against the background element size shown with both linear basis functions and quadratic basis functions. In the case of quadratic basis functions, the foreground mesh was refined as shown in Fig. 9.

5.4. Nonlinear and unsteady problems

We now consider the unsteady incompressible Navier–Stokes equations as an example of a more complicated nonlinear PDE system. This problem also illustrates the use of interpolation-based immersed finite element analysis in conjunction with stabilized finite element formulations that use mesh-dependent variational forms, originally designed with the standard body-fitted FE setting in mind.

In particular, we consider the PDE system: Find velocity and pressure fields $(\mathbf{u}, p) : \Omega \times (0, T) \rightarrow \mathbb{R}^d \times \mathbb{R}$ such that

$$\begin{aligned} \rho (\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u}) &= -\nabla p + \mu \Delta \mathbf{u} \\ \nabla \cdot \mathbf{u} &= 0 \end{aligned} \tag{46}$$

subject to the boundary condition

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial \Omega, \tag{47}$$

where d is the spatial dimension of Ω , $(0, T)$ is the time interval on which the problem is posed, ρ is mass density, μ is dynamic viscosity, and \mathbf{g} is given Dirichlet boundary data. We discretize this in space using a stabilized variational multiscale (VMS) formulation [37] with non-symmetric Nitsche-like weak enforcement of the Dirichlet boundary condition [38], leading to the semi-discrete problem: Find $(\mathbf{u}^h, p^h) \in \mathcal{V}^h \times \mathcal{Q}^h$ such that, $\forall (\mathbf{v}^h, q^h) \in \mathcal{V}^h \times \mathcal{Q}^h$,

$$\begin{aligned} &\int_{\Omega} (\rho (\partial_t \mathbf{u}^h + \mathbf{u}^h \cdot \nabla \mathbf{u}^h) \cdot \mathbf{v}^h + \boldsymbol{\sigma}(\mathbf{u}^h, p^h) : \nabla \mathbf{v}^h + \nabla \cdot \mathbf{u}^h q^h) d\Omega \\ &+ \int_{\partial \Omega} (-(\boldsymbol{\sigma}(\mathbf{u}^h, p^h) \mathbf{n}) \cdot \mathbf{v}^h + (\boldsymbol{\sigma}(\mathbf{v}^h, q^h) \cdot \mathbf{n}) \cdot (\mathbf{u}^h - \mathbf{g})) d\Gamma \\ &\quad - \int_{\partial \Omega} (\rho \min\{\mathbf{u}^h \cdot \mathbf{n}, 0\} (\mathbf{u}^h - \mathbf{g}) \cdot \mathbf{v}^h) d\Gamma \\ &+ \int_{\Omega} \tau_M \left((\mathbf{u}^h \cdot \nabla \mathbf{v}^h + \frac{1}{\rho} \nabla q^h) \cdot \mathbf{r}_M - \mathbf{v}^h \cdot \mathbf{r}_M \cdot \nabla \mathbf{u}^h \right) d\Omega \\ &\quad - \int_{\Omega} \frac{\tau_M^2}{\rho} (\nabla \mathbf{v}^h) : (\mathbf{r}_M \otimes \mathbf{r}_M) d\Omega \\ &\quad + \int_{\Omega} \tau_C (r_C \nabla \cdot \mathbf{v}^h) d\Omega = 0 \end{aligned} \tag{48}$$

where

$$\boldsymbol{\sigma}(\mathbf{u}, p) = 2\mu \boldsymbol{\varepsilon}(\mathbf{u}) - p \mathbf{I} \tag{49}$$

gives the functional form of the Cauchy stress in terms of the velocity and pressure fields,

$$\mathbf{r}_M = \rho(\partial_t \mathbf{u}^h + \mathbf{u}^h \cdot \nabla \mathbf{u}^h) - \nabla \cdot \boldsymbol{\sigma}(\mathbf{u}^h, p^h) \tag{50}$$

is the residual of the strong momentum balance equation,

$$r_C = \rho \nabla \cdot \mathbf{u}^h \tag{51}$$

is the residual of the strong continuity equation, and τ_M and τ_C are the momentum and continuity stabilization parameters, respectively. A number of definitions for these stabilization parameters have been studied in the literature. We follow [39] in defining

$$\tau_M = \left(\mathbf{u}^h \cdot \mathbf{G} \cdot \mathbf{u}^h + C_1 v^2 \mathbf{G} : \mathbf{G} + \frac{C_t}{\Delta t^2} \right)^{-1/2}, \tag{52}$$

and

$$\tau_C = \frac{1}{\tau_M \text{tr}(\mathbf{G})}, \tag{53}$$

with $v = \mu/\rho$ the kinematic viscosity, Δt the time step size of the temporal discretization, $C_1 > 0$ and $C_t > 0$ are dimensionless constants (set to $C_1 = 60$ and $C_t = 4$ in this work), and \mathbf{G} is an anisotropic mesh size tensor, which simplifies to

$$\mathbf{G} = 4h^{-2} \mathbf{I} \tag{54}$$

in the case of a uniform isotropic mesh with element size h . In the computations of this paper, we take this element size to be from the uniform background mesh. We discretize the system of ordinary differential equations emanating from the semi-discrete problem (48) using the implicit midpoint rule in time, with a time step Δt proportional to the background element size h .

As a representative instance of the Navier–Stokes problem, we consider the 2D Taylor–Green vortex. The statement and spatially-periodic exact solution of this problem are well-known. In numerical studies, it is often restricted to the square $(0, \pi)^2$ with symmetry boundary conditions (e.g., [40, Section 9.10.1]). However, in the present work, we take Ω to be the rotated unit square used as the domain of the Poisson problem in Section 5.1, apply the exact solution as the boundary data \mathbf{g} , and use the same sequence of background and foreground mesh pairs defined in Section 5.1. The stabilized formulation (48) permits the stable use of equal-order discretizations, i.e.,

$$\mathcal{V}^h = (\mathcal{V}^h)^d \quad \text{and} \quad \mathcal{Q}^h = \mathcal{V}^h \tag{55}$$

for a single scalar function space \mathcal{V}^h . Our numerical tests in this section use the same linear and quadratic choices for \mathcal{V}^h as used for the scalar Poisson problem in Section 5.1.

The convergence of velocity and pressure as $h \sim \Delta t \rightarrow 0$ (resulting in a CFL number of approximately 1) is shown in Fig. 11. We see effectively optimal rates of velocity convergence with respect to the polynomial degree of the spatial discretizations in the L_2 and H^1 norms. In principle, using $h \sim \Delta t$ and the implicit midpoint rule should limit convergence to second-order in the L_2 norm,⁶ but the quadratic spatial discretizations appear to be third-order in velocity L^2 error within the range of resolutions tested. Pressure H^1 error convergence is also nearly optimal, but with a drop in convergence rate on the finest meshes when using quadratic spatial discretizations. This sub-optimal H^1 pressure convergence is consistent with the analysis of related stabilized methods for Stokes and Oseen flow, e.g., [42,43], where error is bounded in a norm on the velocity–pressure product space in which the pressure gradient is scaled by element size, weakening its contribution.

5.5. Background-unfitted foreground meshes

The analysis of Section 3.5 and preceding numerical examples in Sections 5.1–5.4 all used background-fitted foreground meshes. Background-fitted foreground meshes are a necessity for quadrature-based immersed methods,

⁶ The assumption of $h \sim \Delta t$ is tacit in the stabilized formulation used; refining Δt faster than h would weaken stabilization in the asymptotic regime. See [41] for additional discussion.




Fig. 11. Convergence data for the velocity and pressure fields of the Taylor–Green vortex problem using both linear and quadratic discretizations. Kinematic viscosity was set to 1/100 yielding a Reynolds number of 100, and error is evaluated at time $T = 1$.

since the quadrature rules defined on foreground elements require smoothness of the integrand for accuracy. However, interpolations of background basis functions on a foreground mesh are, by construction, smooth within each foreground element and can be accurately integrated with foreground element-by-element quadrature rules. Of course, the solution space for interpolation-based immersed methods consists of interpolations of background basis functions that are non-smooth within each foreground element, so the analysis of Section 3.5 is not applicable. However, the numerical results of the current section suggest that optimal accuracy is maintained with background-unfitted foreground meshes.

In the preceding examples, the background function space was taken to be a B-spline space of maximal continuity. However, as a stress test of background-unfitted interpolation-based immersion, we instead test it with C^0 Lagrange background function spaces defined on simplicial meshes. As test problems, we consider the Poisson equation with a manufactured solution (cf. Section 5.1) and the Taylor–Green vortex (cf. Section 5.4). For both of these test problems, the domain is taken to be a rotated square, and the sequence of pairs of background and background-unfitted foreground meshes are constructed similarly to the representative pair shown in Fig. 12 but at varying levels of refinement. In all tests, the foreground element size η is kept proportional to the background element size h , and the polynomial degrees k and κ of the background and foreground are equal. For the Taylor–Green vortex problem, the same time implicit midpoint rule with time step Δt proportional to the background element size h was used as in Section 5.4.

The convergence results for both the Poisson and Navier–Stokes examples are shown in Figs. 13 and 14 respectively. We see optimal convergence rates of L^2 and H^1 error (of velocity in the case of Navier–Stokes) for linear and quadratic discretizations.

6. Application to IMGA of trimmed shells

A major success of IGA has been its application to thin shell analysis, originating in the work of Kiendl and collaborators [44–46]. Smooth discrete function spaces from IGA make it possible to apply the Bubnov–Galerkin method directly to Kirchhoff–Love theory where the displacement satisfies a fourth-order PDE system. This provides much better per-degree-of-freedom accuracy than classical FE discretizations of shells, sometimes resolving important qualitative solution features with orders of magnitude less computation [47]. However, the original isogeometric Kirchhoff–Love shell formulation of [44] is restricted to boundary-fitted spline representations of the midsurface geometry. With suitably-flexible spline technologies, such as T-splines [48], it is practical to design such “analysis-suitable” geometries for industrial applications as demonstrated for thin shell analysis of composite wind turbine blades [49], prosthetic heart valve leaflets [50], and automobile parts [51]. However, the vast majority of industrial CAD software defines surface geometries in an immersed way by cutting them out of unfitted parametric coordinate charts along arbitrary trim curves [13]. Thus, immersogeometric formulations are needed to directly analyze CAD models of shell structures. Quadrature-based immersed methods for thin shells have previously been used in a number of studies [52–54]. In this section of the present work, we apply interpolation-based IMGA to Kirchhoff–Love analysis of shells whose geometries are defined by subsets of a B-spline patch’s parameter space.




Fig. 12. A representative example of a background mesh and a background-unfitted foreground mesh.

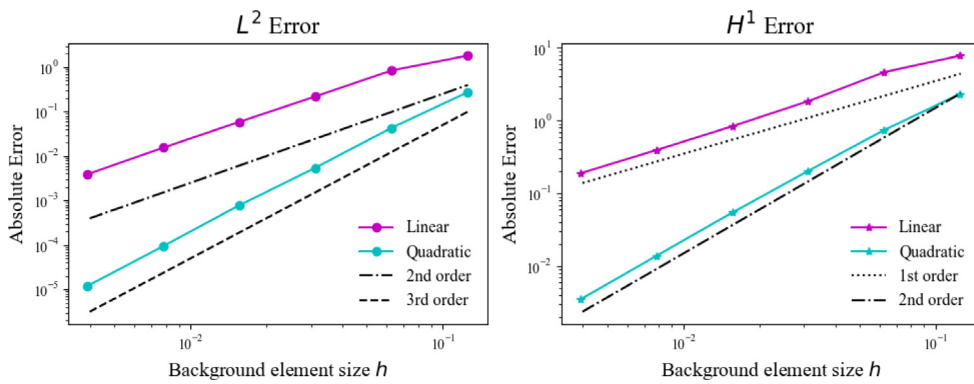


Fig. 13. Convergence data for the Poisson problem, generated using extraction onto unfitted background meshes (Fig. 12).




Fig. 14. Convergence data for the Taylor–Green vortex problem generated using extraction onto unfitted meshes (Fig. 12), with Reynolds number 100, after elapsed time $T = 1$.

The specific variational formulation that we use in this section is that of a geometrically-nonlinear Kirchhoff–Love shell of uniform thickness h_{th} whose material behavior follows a homogeneous isotropic St. Venant–Kirchhoff model defined by a Young’s modulus E and Poisson ratio ν . A complete mathematical statement of this problem is given by [45], and its translation into FEniCS UFL is available online in the library ShNAPr [55] developed to support the work in [56]. We therefore omit the extensive kinematic definitions for brevity as they are neither novel nor essential to following the present study. We mention briefly, though, that we implement pinned boundary conditions on the edges of trimmed shells using a penalty formulation based on [57]. This consists of adding

$$+\frac{1}{2} \int_{\mathcal{L}_{pin}} \frac{\alpha E h_{th}}{h} |\mathbf{u}^h - \mathbf{g}|^2 d\mathcal{L} \tag{56}$$

to the shell structure’s elastic energy functional where \mathcal{L}_{pin} is the portion of the trim curve (in the reference configuration) on which pinned boundary conditions are to be applied, $\alpha > 0$ is a dimensionless constant, h is the background-mesh element size, \mathbf{u}^h is the unknown discrete displacement field of the shell structure’s midsurface, and \mathbf{g} is the prescribed midsurface displacement on the pinned boundaries. For quadratic spline discretizations, this “naive” penalty method for pinned boundaries is essentially as accurate as a full Nitsche-type formulation (cf. [58]) because the associated consistency term involves third-order derivatives. The use of (56) in conjunction with an analogous rotational penalty has been demonstrated to be accurate for coupling non-matching shell structures in practical problem settings [57,59–61].

The examples of this section are limited to trimmed single-patch shell geometries. However, we see no significant conceptual barrier preventing the application of the above-cited patch coupling methodologies to collections of multiple trimmed patches, modeling shell geometries with sharp crease-like features. Given a foreground FE mesh of the physical part of each patch, the software framework of [61] could be applied more-or-less directly, with minimal modifications, since it is already based on Lagrange extraction.

6.1. Verification with boundary-fitted IGA

We first test interpolation-based IMGA of thin shells by trimming out a rectangular domain which can easily be modeled with an untrimmed B-spline to provide a reference solution for verification purposes. In particular, for each Cartesian component of the shell midsurface displacement, we reuse the $k = 2$ sequence of scalar B-spline background and Lagrange-interpolated foreground spaces from the 2D Poisson and biharmonic tests of Sections 5.1 and 5.2, where the parametric-space domain is a rotated unit square trimmed out of an axis-aligned bi-unit square, shown in Fig. 15. Additionally, we use a new background-unfitted $k = 2$ sequence of scalar B-spline background and Lagrange-interpolated foreground spaces, constructed as described in Section 4.3.2. In this test, we take the mapping from the B-spline parameter space to the physical-space reference configuration of the shell’s midsurface to be given by

$$\mathbf{X}(\boldsymbol{\xi}) = \xi_1 \mathbf{e}_1 + \xi_2 \mathbf{e}_2, \tag{57}$$

where $\boldsymbol{\xi} = (\xi_1, \xi_2)$ is a point in the B-spline parameter space, \mathbf{X} is the corresponding 3D point on the reference configuration midsurface, and $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is the standard orthonormal basis for a Cartesian coordinate system of 3D physical space into which the reference configuration is embedded. The construction of an analogous sequence of untrimmed reference B-spline discretizations is straightforward due to the square geometry.

We apply homogeneous (i.e., $\mathbf{g} = \mathbf{0}$) pinned boundary conditions to the entire boundary of the unit square domain, using penalty coefficient $\alpha = 10^5$, and assume material parameters of $E = 4.8 \times 10^5$ and $\nu = 0.38$, and thickness $h_{th} = 0.1$. We then apply a uniform force density per unit reference area of $90\mathbf{e}_3$, leading to an “inflated” equilibrium configuration, of which a representative numerical approximation using interpolation-based IMGA is shown in Figs. 15(a) and 15(b). We also extract the out-of-plane displacement component at the center of the plate using IMGA at several levels of refinement and plot it as a function of number of degrees of freedom in Fig. 15(c). This displacement converges to the reference value obtained from a highly-refined computation using standard untrimmed IGA. (This is a geometrically-nonlinear problem, for which there is no analytical Kirchhoff–Love solution available.)

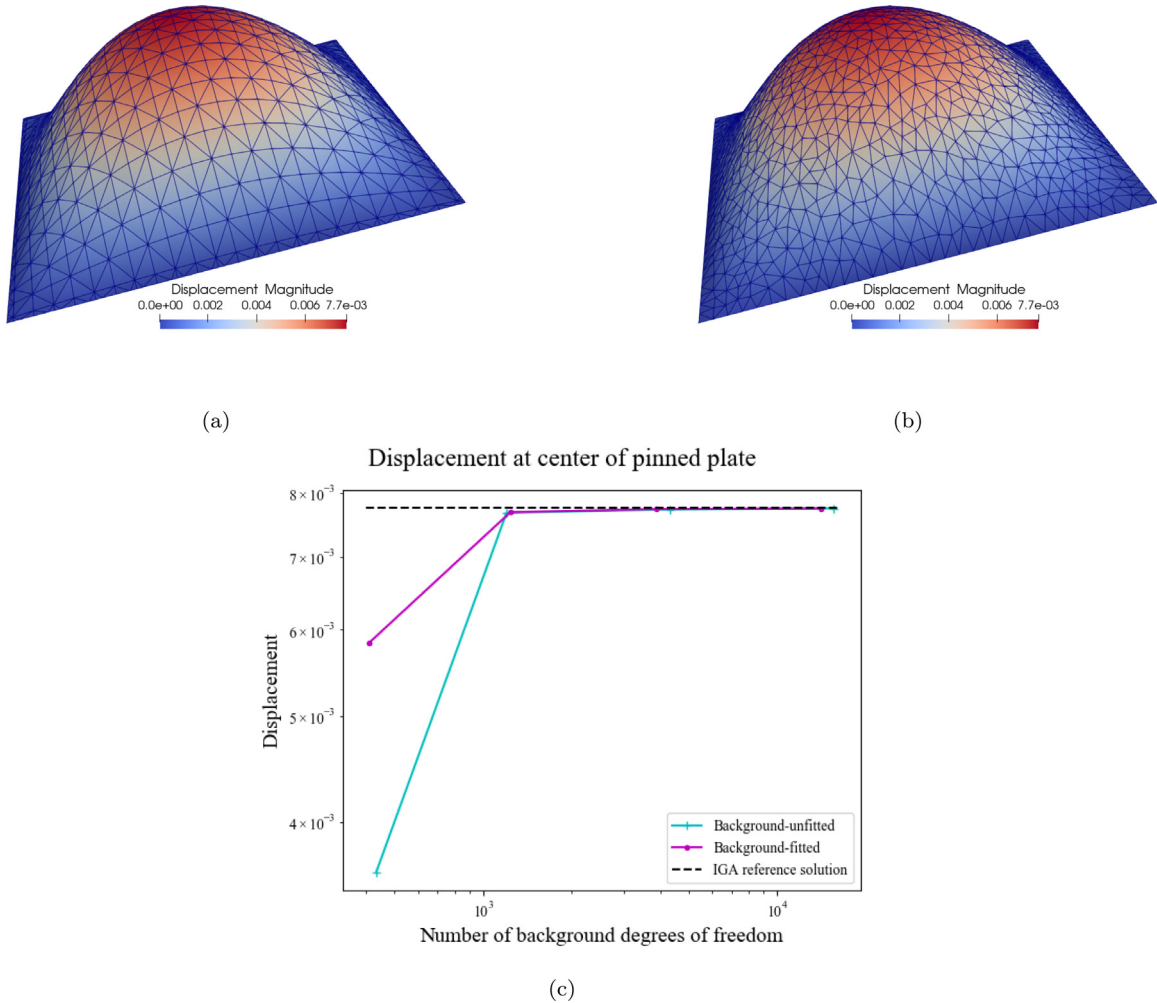



Fig. 15. Examples of the background-fitted (a) and background-unfitted (b) discretizations used to verify the IMGGA formulation. (c) Displacement of the center of the plate is plotted against the number of degrees of freedom in each mesh.

6.2. Background-fitted and -unfitted IMGGA for a complicated trimmed geometry


As a final example, we consider a more complicated trimmed spline geometry, for which a boundary-fitted IGA model would require unstructured spline technologies that are not currently part of standard design workflows. Background-fitted and -unfitted meshes of the trimmed geometry are shown in the spline parameter space in Figs. 16(a) and 16(b). The parametric space is then mapped to a curved physical-space reference configuration by

$$X(\xi) = \xi_1 e_1 + \xi_2 e_2 + A(1 - \xi_1^2) e_3, \tag{58}$$

as visualized in Figs. 16(c) and 16(d) where $A = 1/2$. The physical problem that we solve on this geometry is a static Kirchhoff–Love shell analysis with $E = 3 \times 10^4$, $\nu = 0.3$, and $h_{th} = 0.03$, pinned boundary conditions applied to the outer left and right sides of the parametric domain, and a pressure follower load with force density per unit reference midsurface area of $f = 90n$ where n is the unit normal vector to the deformed midsurface. Representative numerical approximations of the resulting equilibrium configuration are shown in Figs. 16(e) and 16(f).




(a) Background-fitted mesh




(b) Background-unfitted mesh




(c) Background-fitted reference configuration



(d) Background-unfitted reference configuration



(e) Background-fitted deformed configuration



(f) Background-unfitted deformed configuration

Fig. 16. Background-fitted and -unfitted meshes (in B-spline parameter space) and reference and deformed configurations (in physical space) of a trimmed shell structure subjected to a follower load.

Remark 2. The static equilibrium configurations shown in Figs. 16(e) and 16(f) were arrived at by computing the steady-state solution of a dynamic Kirchhoff–Love shell problem with strong mass damping and numerical dissipation from under-resolved backward-Euler time stepping. This pseudo-time continuation method was necessary to avoid bifurcations that arise on some meshes when attempting to apply static load stepping.




Fig. 17. Convergence of background-fitted and -unfitted analyses of a trimmed shell subject to a follower load.

To study convergence with respect to mesh refinement quantitatively, we extract the displacement magnitude at the parametric point $\xi = (0, -1/4)$ (corresponding to the bottom of the inner circle) as a quantity of interest. This displacement is shown as a function of number of degrees of freedom for both the background-fitted and -unfitted discretizations in Fig. 17. We can see that both types of discretizations converge to a consistent value as the number of degrees of freedom increases, where we estimate the converged value of 0.0078 based on a highly-refined background-unfitted discretization.

7. Conclusion

This paper introduces a new paradigm for defining and implementing high-order-accurate immersed-boundary methods based on interpolating basis functions defined on a boundary-unfitted background mesh into a function space defined on a foreground mesh that is fitted to the domain boundary but subject to fewer topological or geometrical constraints than a standard FE mesh. Compared with existing variational immersed methods that use foreground meshes to define quadrature rules for functions on background meshes, interpolation-based immersed finite element analysis provides more flexibility with respect to the construction of the foreground mesh and more straightforward reuse of existing FE software. The interpolation-based methodology is capable of exactly reproducing quadrature-based immersed methods under certain conditions, but it also enables efficient approximations that reduce computational cost with no significant loss of accuracy.

One major challenge of quadrature-based immersed methods that remains in the interpolation-based paradigm is the problem that the conditioning of algebraic equation systems like (24) can become arbitrarily-bad in certain situations where a background mesh element's intersection with the PDE domain is very small [62]. In the present work, we have not applied any special treatment of these situations aside from the *ad hoc* volume based cell-filtering discussed in Remark 1 for the 3D biharmonic tests, although that was motivated by floating-point precision issues, not linear system conditioning.⁷ In practice, we observe that iterative linear solvers remain effective for the examples considered in this paper, but we recognize that pathological conditioning like that identified in [62] can still arise in interpolation-based immersed methods. We hypothesize that techniques from the literature on quadrature-based immersed methods can be adapted *mutatis mutandis* to the interpolation-based setting with similar benefits, such as basis function agglomeration [63], ghost penalty formulations [64], basis function removal [15], and specialized preconditioners [65]. However, verifying this hypothesis is left to future work.

While Section 3.5 briefly sketches a path toward analyzing the approximation properties of the interpolated solution space, our verification efforts are primarily empirical in the present work, relying on numerical experiments

⁷ The technique of Remark 1 may nonetheless incidentally improve system conditioning, by effectively removing basis functions whose supports have small intersections with the domain, similar to the approach analyzed by [15].

with a variety of PDE systems. These experiments indicate that interpolation-based immersed methods are optimally-accurate, even in conditions that are far outside the scope of our initial analysis. It is therefore an interesting topic for future research to perform a more comprehensive *a priori* analysis of interpolation-based immersed methods, especially in the case of foreground meshes that are not fitted to elements of the background mesh.

While this paper has largely focused on comparing interpolation-based immersed methods with their prevailing quadrature-based counterparts, a skeptical reviewer might also raise the following question⁸: If interpolation-based immersed methods require construction of a foreground mesh and execution of an assembly loop over its elements, why not just directly solve the system (29) instead of transforming it to (24)? The preceding sections contain several different answers to this question, which we now compile into an itemized summary of the advantages of interpolation-based immersion over classical body-fitted FE analysis using the foreground mesh:

- The matrix \mathbf{A} defined by (29) would be singular in many problems of interest, e.g., the fourth-order biharmonic (Section 5.2) and Kirchhoff–Love (Section 6) problems.
- Even for second-order problems, one might wish to use a non-conforming approximation of H^1 , e.g., a foreground mesh containing hanging nodes, whose construction is subject to fewer topological constraints. This is outside the scope of Section 3.5’s preliminary analysis, but the empirical success of non-conforming H^2 approximations (as mentioned above) is encouraging.
- Even if the domain Ω is easy to mesh in a conforming way, one may still wish to approximate PDE solutions with functions such as B-splines that are only convenient to construct on a tensor-product grid.
- The numerical results of Section 5.1 demonstrate that interpolation-based immersed methods consistently require fewer degrees of freedom for the same accuracy when compared to classical body-fitted FE analysis using the foreground mesh (in cases where the latter is well-posed at all).

The extent (or even existence) of true computational cost savings (in wall-clock time) due to the last point would depend strongly on choices of algebraic solver algorithms and their parameters, the problem-dependent ratio of assembly time to solution time, and implementation details of the assembly process, which this paper does not explore in detail. In any case, we believe the other advantages of (interpolation-based) immersed methods are of greater significance, since computational cost is often of secondary importance when compared to the labor cost of model preparation in modern engineering analysis workflows.

The non-invasive implementation of interpolation-based immersed methods discussed in Section 4 is of great practical significance, as implementation cost has been a major barrier to widespread adoption and application of immersed methods proposed in the academic literature. We have made the FEniCS-based prototype used for experiments in this paper available online [66]. However, we plan to develop a more comprehensive open-source tool chain for non-invasive implementation of interpolation-based immersed methods in the near future.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Partial source code and data files have been made publicly available on GitHub, as cited in the manuscript.

Acknowledgments

J. E. Fromm was supported by National Science Foundation, United States award number 1650112. R. Xiang and D. Kamensky were partially supported by National Science Foundation, United States award number 2103939. N. Wunsch, K. Maute, and J. A. Evans were supported by National Science Foundation, United States award number 2104106. H. Zhao was partially supported by National Aeronautics and Space Administration, United States award number 80NSSC21M0070.

⁸ This question is not unique to the interpolation-based immersed method proposed here; we have seen similar questions raised about quadrature-based methods relying on foreground quadrature meshes.

References

- [1] T.J.R. Hughes, *The Finite Element Method. Linear Static and Dynamic Finite Element Analysis*, Prentice-Hall, Englewood Cliffs, New Jersey, 1987.
- [2] M.F. Hardwick, R.L. Clay, P.T. Boggs, E.J. Walsh, A.R. Larzelere, A. Altshuler, *DART System Analysis*, Technical Report SAND2005-4647, Sandia National Laboratories, 2005.
- [3] C.S. Peskin, Flow patterns around heart valves: A numerical method, *J. Comput. Phys.* 10 (2) (1972) 252–271.
- [4] J. Parvizian, A. Düster, E. Rank, Finite cell method, *Comput. Mech.* 41 (1) (2007) 121–133.
- [5] D. Schillinger, M. Ruess, The finite cell method: A review in the context of higher-order structural analysis of CAD and image-based geometric models, *Arch. Comput. Methods Eng.* 22 (3) (2015) 391–455.
- [6] S. Hubrich, A. Düster, Numerical integration for nonlinear problems of the finite cell method using an adaptive scheme based on moment fitting, *Comput. Math. Appl.* 77 (7) (2019) 1983–1997.
- [7] S.C. Divi, C.V. Verhoosel, F. Auricchio, A. Reali, E.H. van Brummelen, Error-estimate-based adaptive integration for immersed isogeometric analysis, *Comput. Math. Appl.* 80 (11) (2020) 2481–2516, *High-Order Finite Element and Isogeometric Methods 2019*.
- [8] E. Burman, S. Claus, P. Hansbo, M.G. Larson, A. Massing, CutFEM: Discretizing geometry and partial differential equations, *Internat. J. Numer. Methods Engrg.* 104 (7) (2015) 472–501.
- [9] T.-P. Fries, S. Omerović, Higher-order accurate integration of implicit geometries, *Internat. J. Numer. Methods Engrg.* 106 (5) (2016) 323–371.
- [10] D. Sulsky, S.-J. Zhou, H.L. Schreyer, Application of a particle-in-cell method to solid mechanics, *Comput. Phys. Comm.* 87 (1–2) (1995) 236–252.
- [11] T.J.R. Hughes, J.A. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry, and mesh refinement, *Comput. Methods Appl. Mech. Engrg.* 194 (2005) 4135–4195.
- [12] J.A. Cottrell, T.J.R. Hughes, Y. Bazilevs, *Isogeometric Analysis: Toward Integration of CAD and FEA*, Wiley, Chichester, 2009.
- [13] B. Marussig, T.J.R. Hughes, A review of trimming in isogeometric analysis: Challenges, data exchange and simulation aspects, *Arch. Comput. Methods Eng.* 25 (4) (2018) 1059–1127.
- [14] D. Kamensky, M.-C. Hsu, D. Schillinger, J.A. Evans, A. Aggarwal, Y. Bazilevs, M.S. Sacks, T.J.R. Hughes, An immersogeometric variational framework for fluid–structure interaction: Application to bioprosthetic heart valves, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 1005–1053.
- [15] D. Elfverson, M.G. Larson, K. Larsson, CutIGA with basis function removal, *Adv. Model. Simul. Eng. Sci.* 5 (1) (2018) 6.
- [16] A. Johansson, B. Kehlet, M.G. Larson, A. Logg, Multimesh finite element methods: Solving PDEs on multiple intersecting meshes, *Comput. Methods Appl. Mech. Engrg.* 343 (2019) 672–689.
- [17] A. Logg, K.-A. Mardal, G.N. Wells, et al., in: A. Logg, K.-A. Mardal, G.N. Wells (Eds.), *Automated Solution of Differential Equations By the Finite Element Method*, Springer, 2012.
- [18] M.J. Borden, M.A. Scott, J.A. Evans, T.J.R. Hughes, Isogeometric finite element data structures based on Bézier extraction of NURBS, *Internat. J. Numer. Methods Engrg.* 87 (2011) 15–47.
- [19] M.A. Scott, M.J. Borden, C.V. Verhoosel, T.W. Sederberg, T.J.R. Hughes, Isogeometric finite element data structures based on Bézier extraction of T-splines, *Internat. J. Numer. Methods Engrg.* 88 (2011) 126–156.
- [20] D. Schillinger, P.K. Ruthala, L.H. Nguyen, Lagrange extraction and projection for NURBS basis functions: A direct link between isogeometric and standard nodal finite element formulations, *Internat. J. Numer. Methods Engrg.* 108 (6) (2016) 515–534, nme.5216.
- [21] D. Kamensky, Y. Bazilevs, tIGAR: Automating isogeometric analysis with FEniCS, *Comput. Methods Appl. Mech. Engrg.* 344 (2019) 477–498.
- [22] M. Tirvaudey, R. Bouclier, J.-C. Passieux, L. Chamoin, Non-invasive implementation of nonlinear isogeometric analysis in an industrial FE software, *Eng. Comput.* 37 (1) (2020) 237–261.
- [23] B. Liu, D. Tan, A Nitsche stabilized finite element method for embedded interfaces: Application to fluid-structure interaction and rigid-body contact, *J. Comput. Phys.* 413 (2020) 109461.
- [24] A. Sadeghirad, R.M. Brannon, J. Burghardt, A convected particle domain interpolation technique to extend applicability of the material point method for problems involving massive deformations, *Internat. J. Numer. Methods Engrg.* 86 (12) (2011) 1435–1456.
- [25] T.W. Sederberg, S.R. Parry, Free-form deformation of solid geometric models, *SIGGRAPH Comput. Graph.* 20 (4) (1986) 151–160.
- [26] J. Nitsche, Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind, *Abh. Aus Dem Math. Semin. Univ. Hamburg* 36 (1971) 9–15.
- [27] D. Schillinger, I. Harari, M.-C. Hsu, D. Kamensky, S.K.F. Stoter, Y. Yu, Y. Zhao, The non-symmetric Nitsche method for the parameter-free imposition of weak boundary and coupling conditions in immersed finite elements, *Comput. Methods Appl. Mech. Engrg.* 309 (2016) 625–652.
- [28] M.S. Alnæs, A. Logg, K.B. Ølgaard, M.E. Rognes, G.N. Wells, Unified form language: A domain-specific language for weak formulations of partial differential equations, *ACM Trans. Math. Software* 40 (2) (2014) 9:1–9:37.
- [29] R.C. Kirby, A. Logg, A compiler for variational forms, *ACM Trans. Math. Software* 32 (3) (2006) 417–444.
- [30] A. Logg, G.N. Wells, DOLFIN: Automated finite element computing, *ACM Trans. Math. Software* 37 (2) (2010) 20:1–20:28.
- [31] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, K. Rupp, B.F. Smith, S. Zampini, H. Zhang, PETSc web page, 2015, <http://www.mcs.anl.gov/petsc>.
- [32] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, K. Rupp, B.F. Smith, S. Zampini, H. Zhang, PETSc Users Manual, Technical Report ANL-95/11 - Revision 3.6, Argonne National Laboratory, 2015.

- [33] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A.M. Bruaset, H.P. Langtangen (Eds.), *Modern Software Tools in Scientific Computing*, Birkhäuser Press, 1997, pp. 163–202.
- [34] L. Noel, M. Schmidt, K. Doble, J.A. Evans, K. Maute, XIGA: An eXtended IsoGeometric Analysis approach for multi-material problems, *Comput. Mech.* (2022).
- [35] G. Strang, G.J. Fix, *An Analysis of the Finite Element Method*, Prentice-Hall, Englewood Cliffs, NJ, 1973.
- [36] M.V. Lubarda, V.A. Lubarda, *Intermediate Solid Mechanics*, Cambridge University Press, 2020.
- [37] Y. Bazilevs, V.M. Calo, J.A. Cottrell, T.J.R. Hughes, A. Reali, G. Scovazzi, Variational multiscale residual-based turbulence modeling for large eddy simulation of incompressible flows, *Comput. Methods Appl. Mech. Engrg.* 197 (1) (2007) 173–201.
- [38] Y. Bazilevs, T.J.R. Hughes, Weak imposition of Dirichlet boundary conditions in fluid mechanics, *Comput. & Fluids* 36 (2007) 12–26.
- [39] Y. Bazilevs, V.M. Calo, T.J.R. Hughes, Y. Zhang, Isogeometric fluid–structure interaction: Theory, algorithms, and computations, *Comput. Mech.* 43 (2008) 3–37.
- [40] J.A. Evans, *Divergence-free B-spline Discretizations for Viscous Incompressible Flows* (Ph.D. thesis), University of Texas at Austin, Austin, Texas, United States, 2011.
- [41] M.C. Hsu, Y. Bazilevs, V.M. Calo, T.E. Tezduyar, T.J.R. Hughes, Improving stability of stabilized and multiscale formulations in flow simulations at small time steps, *Comput. Methods Appl. Mech. Engrg.* 199 (13) (2010) 828–840, *Turbulence Modeling for Large Eddy Simulations*.
- [42] T.J.R. Hughes, L.P. Franca, M. Balestra, A new finite element formulation for computational fluid dynamics: V. Circumventing the Babuška–Brezzi condition: A stable Petrov–Galerkin formulation of the Stokes problem accommodating equal-order interpolations, *Comput. Methods Appl. Mech. Engrg.* 59 (1) (1986) 85–99.
- [43] L.P. Franca, S.L. Frey, Stabilized finite element methods: II. The incompressible Navier–Stokes equations, *Comput. Methods Appl. Mech. Engrg.* 99 (2) (1992) 209–233.
- [44] J. Kiendl, K.-U. Bletzinger, J. Linhard, R. Wüchner, Isogeometric shell analysis with Kirchhoff–Love elements, *Comput. Methods Appl. Mech. Engrg.* 198 (2009) 3902–3914.
- [45] J. Kiendl, *Isogeometric Analysis and Shape Optimal Design of Shell Structures* (Ph.D. thesis), Lehrstuhl für Statik, Technische Universität München, 2011.
- [46] J. Kiendl, M.-C. Hsu, M.C.H. Wu, A. Reali, Isogeometric Kirchhoff–Love shell formulations for general hyperelastic materials, *Comput. Methods Appl. Mech. Engrg.* 291 (2015) 280–303.
- [47] S. Morganti, F. Auricchio, D.J. Benson, F.I. Gambarin, S. Hartmann, T.J.R. Hughes, A. Reali, Patient-specific isogeometric structural analysis of aortic valve closure, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 508–520.
- [48] T. Sederberg, J. Zheng, A. Bakenov, A. Nasri, T-splines and T-NURCCS, *ACM Trans. Graph.* 22 (3) (2003) 477–484.
- [49] Y. Bazilevs, M.-C. Hsu, M.A. Scott, Isogeometric fluid–structure interaction analysis with emphasis on non-matching discretizations, and with application to wind turbines, *Comput. Methods Appl. Mech. Engrg.* 249–252 (2012) 28–41.
- [50] M.-C. Hsu, D. Kamensky, F. Xu, J. Kiendl, C. Wang, M.C.H. Wu, J. Mineroff, A. Reali, Y. Bazilevs, M.S. Sacks, Dynamic and fluid–structure interaction simulations of bioprosthetic heart valves using parametric design with T-splines and Fung-type material models, *Comput. Mech.* (2015) 1–15.
- [51] H. Casquero, X. Wei, D. Toshniwal, A. Li, T.J.R. Hughes, J. Kiendl, Y.J. Zhang, Seamless integration of design and Kirchhoff–Love shell analysis using analysis-suitable unstructured T-splines, *Comput. Methods Appl. Mech. Engrg.* 360 (2020) 112765.
- [52] Y. Guo, J. Heller, T.J.R. Hughes, M. Ruess, D. Schillinger, Variationally consistent isogeometric analysis of trimmed thin shells at finite deformations, based on the STEP exchange format, *Comput. Methods Appl. Mech. Engrg.* 336 (2018) 39–79.
- [53] L. Coradello, D. D’Angella, M. Carraturo, J. Kiendl, S. Kollmannsberger, E. Rank, A. Reali, Hierarchically refined isogeometric analysis of trimmed shells, *Comput. Mech.* 66 (2) (2020) 431–447.
- [54] L. Coradello, J. Kiendl, A. Buffa, Coupling of non-conforming trimmed isogeometric Kirchhoff–Love shells via a projected super-penalty approach, *Comput. Methods Appl. Mech. Engrg.* 387 (2021) 114187.
- [55] <https://github.com/david-kamensky/ShNAPr>, *ShNAPr source code*.
- [56] D. Kamensky, Open-source immersogeometric analysis of fluid–structure interaction using FEniCS and tIGAr, *Comput. Math. Appl.* 81 (2021) 634–648, *Development and Application of Open-source Software for Problems with Numerical PDEs*.
- [57] A.J. Herrema, E.L. Johnson, D. Proserpio, M.C.H. Wu, J. Kiendl, M.-C. Hsu, Penalty coupling of non-matching isogeometric Kirchhoff–Love shell patches with application to composite wind turbine blades, *Comput. Methods Appl. Mech. Engrg.* 346 (2019) 810–840.
- [58] J. Benzaken, J.A. Evans, S.F. McCormick, R. Tamstorf, Nitsche’s method for linear Kirchhoff–Love shells: Formulation, error analysis, and verification, *Comput. Methods Appl. Mech. Engrg.* 374 (2021) 113544.
- [59] E.L. Johnson, M.-C. Hsu, Isogeometric analysis of ice accretion on wind turbine blades, *Comput. Mech.* 66 (2) (2020) 311–322.
- [60] N. Liu, E.L. Johnson, M.R. Rajanna, J. Lua, N. Phan, M.-C. Hsu, Blended isogeometric Kirchhoff–Love and continuum shells, *Comput. Methods Appl. Mech. Engrg.* 385 (2021) 114005.
- [61] H. Zhao, X. Liu, A.H. Fletcher, R. Xiang, J.T. Hwang, D. Kamensky, An open-source framework for coupling non-matching isogeometric shells with application to aerospace structures, *Comput. Math. Appl.* 111 (2022) 109–123.
- [62] F. de Prenter, C.V. Verhoosel, G.J. van Zwieten, E.H. van Brummelen, Condition number analysis and preconditioning of the finite cell method, *Comput. Methods Appl. Mech. Engrg.* 316 (2017) 297–327, *Special Issue on Isogeometric Analysis: Progress and Challenges*.
- [63] F. Kummer, Extended discontinuous Galerkin methods for two-phase flows: The spatial discretization, *Internat. J. Numer. Methods Engrg.* 109 (2) (2017) 259–289.
- [64] E. Burman, Ghost penalty, *C. R. Math.* 348 (21) (2010) 1217–1220.

- [65] J.N. Jomo, F. de Prenter, M. Elhaddad, D. D'Angella, C.V. Verhoosel, S. Kollmannsberger, J.S. Kirschke, V. Nübel, E.H. van Brummelen, E. Rank, Robust and parallel scalable iterative solutions for large-scale finite cell analyses, *Finite Elem. Anal. Des.* 163 (2019) 14–30.
- [66] <https://github.com/jefromm/interpolation-based-immersed-fea>, Source code for FEniCS-based solver.