

Simulating Thin Shells by Bicubic Hermite Elements

Xingyu Ni ^{a,d,1}, Xuwen Chen ^{b,d,1}, Chengyu ^{c,d}, Bin Wang ^{e,*}, Baoquan Chen ^{b,d}

^a School of CS, Peking University, China

^b School of IST, Peking University, China

^c School of EECS, Peking University, China

^d State Key Laboratory of General Artificial Intelligence, Peking University, China

^e State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI), China

ARTICLE INFO

Keywords:

Thin-shell simulation
High-order finite elements
Hermite interpolation

ABSTRACT

In this study, we present the bicubic Hermite element method (BHEM), a new computational framework devised for the elastodynamic simulation of thin-shell structures. The BHEM is constructed based on quadrilateral Hermite patches, which serve as a unified representation for shell geometry, simulation, collision avoidance, as well as rendering. Compared with the commonly utilized linear FEM, the BHEM offers higher-order solution spaces, enabling the capture of more intricate and smoother geometries while employing significantly fewer finite elements. In comparison to other high-order methods, the BHEM achieves conforming C^1 continuity for Kirchhoff–Love (KL) shells with minimal complexity. Furthermore, by leveraging the subdivision and convex hull properties of Hermite patches, we develop an efficient algorithm for ray-patch intersections, facilitating collision handling in simulations and ray tracing in rendering. This eliminates the need for laborious remodeling of the pre-existing surface as the conventional approaches do. We substantiate our claims with comprehensive experiments, which demonstrate the high accuracy and versatility of the proposed method.

1. Introduction

The mechanical properties of thin-shell structures are commonly described using the Kirchhoff–Love (KL) theory [1], which neglects transverse shear effects and focuses on the in-plane stretching and lateral bending of the shell’s midsurface. Under the KL assumption, the energy density function derived from the elastic strain incorporates second-order derivatives of displacements (as elaborated in Section 3). This requirement necessitates H^2 regularity of the geometric representation of midsurface, which implies C^1 continuity, to ensure a well-defined analysis of the elastic energy.

To effectively address the second-order thin-shell energy, many endeavors have been developed, tested, and honed. In the realm of linear finite-element analysis, which is highly favored within the field of computer graphics, a customary solution is to reformulate the bending energy as a specialized function of the dihedral angle based on the discrete differential geometry principle [2–4]. Nevertheless, this discretized edge-based energy allows bending motion along the common edges of elements only and eventually fails to converge towards the shape operator of the smooth surface at the element interfaces, irrespective of the chosen discretization scheme and mesh resolution [5]. Isogeometric Analysis (IGA) [6,7] is an appealing alternative to attaining high-order continuity solution space. It utilizes various basis

functions emanating from computer-aided geometric design (CAGD) in finite element analysis for their global smoothness. One of the earliest depictions of this approach was presented by Cirak et al. [8,9], who devised a finite element formulation grounded in Loop subdivision surfaces for Kirchhoff–Love thin-shell simulation. Loop subdivision scheme can easily represent smooth surfaces of arbitrary topology with polygonal mesh data structures. By employing the same convergent shape function for displacement field interpolation, the subdivision finite element scheme [8,9] requires only nodal displacement degrees of freedom while retaining C^1 continuity across elements, which is necessary for thin-shell simulation. The use of more general non-uniform rational B-splines (NURBS) basis functions in the finite element context was also proposed by Hughes et al. [6] in 2005. NURBS patch is more memory-friendly than subdivision surface, but it needs additional constraints to maintain conforming C^1 continuity within multi-patch models [10]. Since the smooth surface does not pass through the coarse control mesh nodes for both NURBS and subdivision FEM scheme, applying boundary conditions and resolving contacts on control nodes need complicated treatment. Hermite elements stand out for their inherent guarantee of C^1 continuity among patches by incorporating shared derivatives as degrees of freedom. This particular type of element has

* Corresponding author.

E-mail addresses: binwangbuaa@gmail.com (B. Wang), baoquan.chen@gmail.com (B. Chen).

¹ Xingyu Ni and Xuwen Chen contributed equally to this paper.

been demonstrated to be advantageous when solving Kirchhoff plate problems [11–14]. Even though the curved Hermite surface exactly goes through the nodal position, existing studies still heavily resort to surface triangulation for collision handling and rendering, diminishing the value of high-order methods.

In this paper, in order to fully leverage the advantages provided by the high-order method, while concurrently minimizing the modifications to the FEM simulation pipeline, we develop a novel framework for thin-shell simulation. The most crucial aspect of our framework lies in the utilization of a unified bicubic-Hermite-patches-based representation, which serves as the foundation of midsurface geometric modeling, dynamics simulation, and ray-tracing rendering purposes. Specifically, we present three main contributions to achieving this goal.

First, we develop a C^1 -continuous finite element solution for dynamic simulation of Kirchhoff–Love thin shells, dubbed BHEM (Bicubic Hermite Element Method). We listed the derivation process of the governing equations and its weak form (Section 4) of KL thin shells from the first principles of continuum mechanics. In the Hermit polynomial space, the discretized form of the governing equations, along with the gradient and the Hessian matrix (see the supplementary material) of the hyperelastic energy, are provided to facilitate a seamless implicit solve (Section 6.1). Furthermore, the BHEM also incorporates a tailored Hessian matrix which entails much less computational overhead, to cater to applications with restricted time constraints.

Further consolidating the integration between geometry representation and simulation, and bypassing mesh-based collision detection, our second main contribution is a BH-surface intersection detection algorithm that serves for both CCD and rendering tasks (Section 5). In our method, the bicubic Hermite surfaces are transformed into their equivalent Bézier form. Based on the convex hull property of the Bézier surface, a bounding-volume-hierarchy (BVH) tree is constructed through dynamic subdivision of the surface and then pruned using Newton’s method to enhance the computation efficiency. Our method ensures the discovery of the first intersection point in both static and dynamic settings. It can also be extended to other spline-based surface intersection detection, such as rational Bézier patches and NURBS, which we believe is a critical missing piece of current IGA research.

Last but not least, we conduct a broad array of experiments meticulously designed to showcase the fidelity and efficiency of the BHEM framework. We first test our method under typical quasi-static settings. The simulation results are highly consistent with the theoretical solutions. Compared with linear FEMs, BHEM can capture rich geometric features with much fewer DoFs, better convergence speed, and less time cost. Furthermore, we demonstrate the applicability of the framework by applying it to a variety of graphics scenarios with complex collisions and diverse boundary conditions. We emphasize the smoothness and highlight the characteristic wrinkles and folds of cloth geometries during dynamic simulations in all these experiments.

2. Related work

2.1. Thin-shell simulation

Physics-based thin shell simulation, such as cloth [15–20], paper [21–24], and skin [25,26], is a long-standing topic in computer graphics owing to their visually appealing geometry and desirable mechanical properties. Following the seminal work of Terzopoulos et al. [27], a series of discrete constitutive models [3–5,28] have been developed for the elastic strain energy by applying geometric operators over the piecewise-linear surface, but with the bending energy being a non-integrable function of the dihedral angle. Researchers have taken many measures to circumvent the issue, such as subdivision FEM [8,29–31] and discontinuous Galerkin (DG) FEM [32], but both with numerical issues due to their unavoidable special treatments to approximate or meet the continuity requirement of a Kirchhoff–Love shell.

Another branch of the solution follows the route of isogeometric analysis, in which the basis functions that express geometric shapes are utilized in FE analysis for the physical field interpolation [10,33]. After the first attempt [34] for using NURBS surface to represent a deformable thin object, NURBS has been widely exploited in cloth motion to describe the characteristic wrinkles or folds [10] and in volumetric object simulation [35]. However, due to its complex math formulation, the NURBS surfaces have involved derivation and high computational cost in graphics applications.

Studies have explored Hermite-interpolated surfaces on triangular elements [36,37], quadrangular elements [38–40], and hybrid elements [41] in engineering and numerical analysis. Triangular Hermite elements use higher-order polynomials as basis functions so as to achieve C^1 continuity, leading to more computational cost. By contrast, quadrangular Hermite elements can have simpler but accurate expressions, though fall short in geometric flexibility. In graphics, Hermite patches were originally adopted for freeform surface modeling [42,43], but they lack applications in deformable objects. In this work, we try to establish a complete framework that enables dynamic simulation and rendering of Hermite-interpolated surfaces.

2.2. Rendering of parametric surfaces

The calculation of ray tracing a parametric surface is to find the first intersection point between a ray and a surface, which is equivalent to solving a system of nonlinear equations obtained by combining the parametric equations of the two geometric primitives. A commonly adopted approach to ray-trace a high-order parametric surface is to triangulate the surface and render the approximated mesh. But to reserve the special advantages of high-order representation, it is beneficial to explore an efficient and robust rendering method that solves the problem directly on the parametric surface.

Existing studies have developed various numerical methods for this problem. Resultant elimination [44,45] transforms the problem into finding the minimum of a uni-variant polynomial, which works but requires heavy numerical computation with complex situational discussion, thus lack of robustness. Newton’s method [46–51] is a powerful tool that can deliver an accurate solution with a fast convergence speed, but it depends severely on the initial conditions and has no guarantee for a correct solution since it only solves the problem locally. Also, the system would become ill-conditioned if a ray intersects the surface tangentially.

To solve for a global solution robustly, researchers have introduced subdivision methods [52–55] that recursively prune the parameter domain until convergence. Bézier clipping [52,56,57] is a classical subdivision method that utilizes the convex-hull property of Bézier patches. It has been widely used for ray-tracing any surfaces that can be converted into rational Bézier patches, including B-splines [57] and NURBS [58]. Meanwhile, subdivision-based methods often expend more time and space cost than Newton’s methods. Therefore, a popular idea is to first conduct subdivision methods to construct a BVH for culling and getting a rough solution as an initial guess and then use Newton’s descent to get a final answer [47–49,51]. Simple geometric primitives, such as Chebyshev boxing/sphere [51,59], Convex hulls [49,60], and axis aligned bounding box [47,48] are the popular choices for BVH representations.

2.3. Collision detection of parametric surfaces

Continuous collision detection (CCD) for high-order meshes or parametric surfaces is mostly conducted by detecting collisions between two approximated triangle meshes [61], due to the maturity of traditional CCD techniques for linear FEM. As can be imagined, this would introduce inevitable false positives and negatives where the original surface sinks or extrudes from the linear mesh.

Conducting DCD/CCD directly on parametric surfaces would result in solving a 4D/5D nonlinear system with a complex solution manifold, which is usually reduced to point pairs in implementation [62,63]. Similar to the ray-tracing problem, solutions for CCD are always established on either subdivision methods [62–64] to recursively solve for a global optimum or Newton's methods [10,63,65] to directly solve for a local optimum. Researchers need to get an appropriate initial guess for Newton's method and cull unnecessary checks via assisting strategies including tessellation [10], BVH [66], spatial partition [67], etc. Recently, another promising approach proposed by Zhang et al. [68,69] modeled the CCD problem between polynomial surfaces as a sum-of-squares programming (SOSP) so bypassed linearizing high-order surfaces. This method includes an auxiliary hyperparameter in SOSP that directly determines the certificate of the exact solution, which requires a trade-off between efficiency and effectiveness.

3. Thin-shell theory

To make the paper self-contained, in this section, we will briefly review the thin-shell theory based upon the Kirchhoff–Love assumption, which generally follows the work of Cirak et al. [8] but is not restricted to linearized deformation.

Conventions and notations. The Einstein summation convention is followed, where an index variable appears twice in a single term implying summation of that term over all the values of the index. We further assume that the value of an index denoted by a lowercase Latin letter (e.g., i, j , and k) ranges over the set $\{1, 2, 3\}$, while that denoted by a lowercase Greek letter (e.g., α, β , and γ) ranges over the set $\{1, 2\}$. Besides, indices appearing after commas imply partial derivatives.

3.1. Geometries

We begin by considering a *midsurface* $\Omega \subset \mathbb{R}^3$. As a surface, Ω is parameterized with curvilinear coordinates (ξ^1, ξ^2) . The possible values of these coordinates form a parameter space ω , and the parameterization is then given by a mapping $\mathbf{x} : \omega \rightarrow \Omega$ such that the following properties hold:

- At each point of Ω , the two partial derivatives $\partial \mathbf{x} / \partial \xi^1$ and $\partial \mathbf{x} / \partial \xi^2$ exist and are linearly independent;
- As a function $\omega \rightarrow \mathbb{R}^3$, $\mathbf{x}(\xi^1, \xi^2)$, as well as its first- and second-order derivatives, is square-integrable.

The former property allows the definition of the unit normal vector $\mathbf{a}_3 = \mathbf{a}_1 \times \mathbf{a}_2 / \|\mathbf{a}_1 \times \mathbf{a}_2\|$, in which \mathbf{a}_α denotes $\partial \mathbf{x} / \partial \xi^\alpha$, while the latter property is required by analysis of elastic energy.

According to the Kirchhoff–Love assumption, the midsurface is extruded by a constant distance $h/2$ both along and opposite to the surface normal direction, which forms the volume of a h -thick shell. With ω^h defined as $\omega \times [-h/2, +h/2]$, the extrusion of Ω is described by a function $\mathbf{r} : \omega^h \rightarrow \Omega^h$ as follows:

$$\mathbf{r}(\xi^1, \xi^2, \xi^3) = \mathbf{x}(\xi^1, \xi^2) + \xi^3 \mathbf{a}_3(\xi^1, \xi^2), \quad -\frac{h}{2} \leq \xi^3 \leq +\frac{h}{2}, \quad (1)$$

which lays the foundation of a thin-shell geometry.

To analyze shell deformation, the geometric difference between deformed and undeformed (reference) configurations needs evaluation. We assume that each point of ω^h always maps to the same material point, and use symbols with overbars to indicate quantities in the reference configuration. A function $\bar{\mathbf{r}} : \omega^h \rightarrow \bar{\Omega}^h$ can be similarly defined by

$$\bar{\mathbf{r}}(\xi^1, \xi^2, \xi^3) = \bar{\mathbf{x}}(\xi^1, \xi^2) + \xi^3 \bar{\mathbf{a}}_3(\xi^1, \xi^2), \quad -\frac{h}{2} \leq \xi^3 \leq +\frac{h}{2}. \quad (2)$$

3.2. Strains

Given the parametric description of shell geometry, we acquire the tangent basis vectors of Ω^h as follows:

$$\mathbf{g}_i = \frac{\partial \mathbf{r}}{\partial \xi^i} = \begin{cases} \mathbf{a}_\alpha + \xi^3 \mathbf{a}_{3,\alpha}, & i = \alpha < 3, \quad (\text{a}) \\ \mathbf{a}_3, & i = 3. \quad (\text{b}) \end{cases} \quad (3)$$

Dot products of these vectors result in covariant components of the metric tensor. To be specific, $g_{ij} = \mathbf{g}_i \cdot \mathbf{g}_j$ is formulated by

$$g_{ij} = \begin{cases} a_{\alpha\beta} - 2b_{\alpha\beta}\xi^3 + c_{\alpha\beta}(\xi^3)^2, & i = \alpha < 3 \wedge j = \beta < 3, \quad (\text{a}) \\ 1, & i = j = 3, \quad (\text{b}) \\ 0, & \text{otherwise}, \quad (\text{c}) \end{cases} \quad (4)$$

where $a_{\alpha\beta} = \mathbf{a}_\alpha \cdot \mathbf{a}_\beta$, $b_{\alpha\beta} = \mathbf{a}_{\alpha,\beta} \cdot \mathbf{a}_3$, and $c_{\alpha\beta} = \mathbf{a}_{3,\alpha} \cdot \mathbf{a}_{3,\beta}$ correspond to the 1, 2, and 3 fundamental forms of S , respectively.

With $\{\bar{\mathbf{g}}_i\}$ being the basis of the tensor space, the *Green–Lagrange strain* is defined as half the difference between the metric tensors in the deformed and undeformed configurations:

$$E_{ij} = \frac{1}{2}(g_{ij} - \bar{g}_{ij}). \quad (5)$$

As will be readily seen, E_{ij} can be expanded as

$$\begin{cases} E_{\alpha\beta} = A_{\alpha\beta} - 2B_{\alpha\beta}\xi^3 + C_{\alpha\beta}(\xi^3)^2, & (\text{a}) \\ E_{3\alpha} = E_{\alpha 3} = 0, & (\text{b}) \\ E_{33} = 0, & (\text{c}) \end{cases} \quad (6)$$

followed by definitions $A_{\alpha\beta} = (a_{\alpha\beta} - \bar{a}_{\alpha\beta})/2$, $B_{\alpha\beta} = (b_{\alpha\beta} - \bar{b}_{\alpha\beta})/2$, and $C_{\alpha\beta} = (c_{\alpha\beta} - \bar{c}_{\alpha\beta})/2$. Note that the 0-order term $A_{\alpha\beta}$ in Eq. (6a) represents the *membrane strain*, while the other terms $B_{\alpha\beta}\xi^3 + C_{\alpha\beta}(\xi^3)^2$ characterize the *curvature strain*. With the linear stain assumed in the thickness direction, the highest-order term of Eq. (6a), specifically $C_{\alpha\beta}(\xi^3)^2$, is omitted in the following calculation.

3.3. Energies

As suggested by previous studies [8,70], the formula of the elastic strain energy V_e is derived from the Green–Lagrange strain based upon a St. Venant–Kirchhoff constitutive model, which is given in the form of areal density \bar{v}_e as

$$\bar{v}_e = \frac{dV_e}{d\bar{\Omega}} = \left(A_{\alpha\beta} A_{\gamma\delta} h + \frac{1}{3} B_{\alpha\beta} B_{\gamma\delta} h^3 \right) \bar{H}^{\alpha\beta\gamma\delta}, \quad (7)$$

in which $\bar{H}^{\alpha\beta\gamma\delta}$ is defined by

$$\bar{H}^{\alpha\beta\gamma\delta} = \frac{\lambda}{2} \bar{a}^{\alpha\beta} \bar{a}^{\gamma\delta} + \mu \bar{a}^{\beta\gamma} \bar{a}^{\alpha\delta}. \quad (8)$$

Here, $\bar{a}^{\alpha\beta}$ is a contravariant tensor component, which can be calculated by the matrix inversion $(\bar{a}^{\alpha\beta})_{2 \times 2} = (\bar{a}_{\alpha\beta})_{2 \times 2}^{-1}$. λ and μ , known as the first and the second *Lamé parameters*, are deduced from Young's modulus Y and Poisson's ratio ν as $\lambda = Y\nu/(1-\nu^2)$ and $\mu = Y/2(1+\nu)$, respectively.

The motion of a thin shell is also influenced by the kinetic energy T . With $\bar{\rho}$ denoting mass density, similar to V_e , T is also written in the form of areal density \bar{T} as

$$\bar{T} = \frac{dT}{d\bar{\Omega}} = \frac{1}{2} \bar{\rho} h (\dot{\mathbf{x}} \cdot \dot{\mathbf{x}}), \quad (9)$$

which is derived in the material space by integrating energy density along the thickness and ignoring high-order infinitesimals.

3.4. Equations of motion

Finally, the equations of motion for a thin shell can be constructed by analyzing the interchange of energy. Since the volumetric shell is replaced by its midsurface, we should also reduce any force that applies to the shell into surface force by integration. Suppose that the areal density of external force applied on $\bar{\Omega}$ is $\bar{\mathbf{f}}$, and the linear density on $\bar{\Gamma} = \partial \bar{\Omega}$ is $\bar{\mathbf{t}}$, D'Alembert's principle states that for any virtual deformation $\delta \mathbf{x}$,

$$\iint_{\bar{\Omega}} (\delta \bar{T} + \delta \bar{v}_e) d\bar{\Omega} = \iint_{\bar{\Omega}} \bar{\mathbf{f}} \cdot \delta \mathbf{x} d\bar{\Omega} + \int_{\bar{\Gamma}} \bar{\mathbf{t}} \cdot \delta \mathbf{x} d\bar{\Gamma}, \quad (10)$$

in which $\delta\bar{\mathcal{T}} = \bar{\rho}h\dot{\mathbf{x}} \cdot \delta\mathbf{x}$ can be interpreted as the virtual work done by the inertia force, and the variation of $\bar{\mathcal{V}}_e$ can be expanded as

$$\delta\bar{\mathcal{V}}_e = \left(hA_{\gamma\delta} \delta a_{\alpha\beta} + \frac{1}{3}h^3 B_{\gamma\delta} \delta b_{\alpha\beta} \right) \bar{H}^{\alpha\beta\gamma\delta}, \quad (11)$$

owing to the exchange symmetry of indices.

It is noteworthy that the surface area element $d\bar{\Omega}$ can be expressed in the parameter space by

$$d\bar{\Omega} = \sqrt{\bar{a}} d\xi^1 d\xi^2, \quad (12)$$

$$\bar{a} = (\bar{\mathbf{a}}_1 \times \bar{\mathbf{a}}_2)^2 = \bar{a}_{11}\bar{a}_{22} - \bar{a}_{12}\bar{a}_{21}. \quad (13)$$

4. The bicubic Hermite element method

4.1. Geometric discretization

To numerically analyze the statics and dynamics of a thin shell, we divide its midsurface into a collection of bicubic Hermite patches that share information via nodes, where the bicubic Hermite interpolation is used to naturally maintain the C^1 smoothness and square-integrability (Section 3.1) in surface reconstruction.

As shown in Fig. 5, each patch corresponds to an axis-aligned rectangle in the parameter space, and each rectangle side is entirely shared by two adjacent patches. In the context of such discretization, with the function value and its derivatives (i.e., \mathbf{x} , $\partial\mathbf{x}/\partial\xi^1$, $\partial\mathbf{x}/\partial\xi^2$, and $\partial^2\mathbf{x}/\partial\xi^1\partial\xi^2$) stored at nodes, any point on a single patch satisfying $\xi^1 \in [\xi_{\min}^1, \xi_{\max}^1]$ and $\xi^2 \in [\xi_{\min}^2, \xi_{\max}^2]$ can be expressed with the piecewise bicubic Hermite interpolation (see the supplementary material), such that the midsurface can be parameterized by

$$\mathbf{x}(\xi^1, \xi^2) = \sum_{I=1}^N \Phi^I(\xi^1, \xi^2) \mathbf{q}_I, \quad (14)$$

in which $\mathbf{q}_I \in \mathbb{R}^3$ ($I = 1, 2, 3, \dots, N$), treated as *generalized coordinates*, denotes a value or derivative that is stored at nodes, and each \mathbf{q}_I corresponds to a shape function $\Phi^I(\xi^1, \xi^2)$. These shape functions have compact supports, so the summand $\Phi^I \mathbf{q}_I$ takes nonzero values only if \mathbf{q}_I is stored at the 4 nodes of the patch that (ξ^1, ξ^2) lies in. Typically, the number of such \mathbf{q}_I is 16.

The bicubic Hermite patches can be readily used to discretize surfaces that are homeomorphic to a rectangle by embedding them in a Cartesian grid and taking grid cells as patches (Fig. 5). More complex surfaces, e.g., cylinders and torus, can also be divided into BH patches with the help of continuity boundary conditions (Section 4.2), which is illustrated in Figs. 6(a) and 6(b).

4.2. Boundary conditions

Continuity boundary conditions. For BHEM simulations, it is significant to incorporate the continuity boundary conditions in order to complete geometric discretization. As an example, in the discretization of a cylinder (Fig. 6(a)), a node with an azimuthal coordinate of 2π is identical to that with an azimuthal coordinate of 0. Whenever we encounter such a situation, only one of the overlapping nodes needs maintenance. Any calculation that involves these nodes is referred to as the maintained one.

Positional constraints. The BHEM prescribes two categories of positional constraints, namely Dirichlet and Neumann boundary conditions, which respectively constrain $\mathbf{x}(\xi^1, \xi^2)$ and its first-order derivatives at specific points. When at nodes, these constraints can be simply realized by removing their corresponding generalized coordinates from the solution variables and taking their influence back to the governing equations as a given term. Based on this, a constraint imposed on an element's whole boundary curve can be realized through the combination of Dirichlet constraints and Neumann constraints along the boundary tangent direction for every node on the boundary. Furthermore, to apply the constraints at arbitrary points of the midsurface, Lagrange multipliers should be introduced, as done in conventional finite element methods.

4.3. Governing equations

Now we consider discretizing Eq. (10) using the variational method based on the BH patches. Taking $\delta\bar{\mathcal{V}}_e = \partial\bar{\mathcal{V}}_e/\partial\mathbf{q}_I \cdot \delta\mathbf{q}_I$, substituting Eq. (14) into Eq. (10) yields

$$\sum_{I=1}^N \left(\iint_{\bar{\Omega}} \left(\Phi^I \bar{\mathbf{f}}^* - \frac{\partial\bar{\mathcal{V}}_e}{\partial\mathbf{q}_I} \right) d\bar{\Omega} + \int_{\bar{\Gamma}} \Phi^I \bar{\mathbf{t}} d\bar{\Gamma} \right) \cdot \delta\mathbf{q}_I = 0, \quad (15)$$

$$\bar{\mathbf{f}}^* = \bar{\mathbf{f}} - \bar{\rho}h \sum_{J=1}^{4N} \Phi^J \bar{\mathbf{q}}_J, \quad (16)$$

which indicates that Eq. (10) is always true for any virtual deformation interpolated through Eq. (14). Due to the arbitrariness of $\delta\mathbf{q}_I$, every coefficient of $\delta\mathbf{q}_I$, i.e., the terms in the outermost parentheses of Eq. (15), must equal to zero. Thus, for any I ($1 \leq I \leq N$), the following equations hold:

$$\iint_{\bar{\Omega}} \Phi^I \bar{\mathbf{f}} d\bar{\Omega} + \int_{\bar{\Gamma}} \Phi^I \bar{\mathbf{t}} d\bar{\Gamma} - \iint_{\bar{\Omega}} \frac{\partial\bar{\mathcal{V}}_e}{\partial\mathbf{q}_I} d\bar{\Omega} - \sum_{J=1}^N M^{IJ} \bar{\mathbf{q}}_J = \mathbf{0}, \quad (17)$$

where each coefficient of the mass matrix is defined as

$$M^{IJ} = \iint_{\bar{\Omega}} \bar{\rho}h \Phi^I \Phi^J d\bar{\Omega} = \iint_{\bar{\Omega}} \bar{\rho}h \Phi^I \Phi^J \sqrt{\bar{a}} d\xi^1 d\xi^2. \quad (18)$$

These constitute the governing equations of a dynamic BHEM shell. For static analysis, the last term in the left-hand side of Eq. (17) is omitted.

Properties of the mass matrix. Due to the compact support of the shape functions Φ , a coefficient M^{IJ} is nonzero only when the affected scopes of Φ^I and Φ^J overlap, i.e., the corresponding generalized coordinates of Φ^I and Φ^J belong to the nodes of the same patch. This implies the sparsity of the mass matrix. Furthermore, with Φ^I being a polynomial of no more than degree 3, $\Phi^I \Phi^J$ reaches a sixth-order at most in each parametric dimension. Thus we conclude that M^{IJ} can be calculated precisely without much effort, given that $\sqrt{\bar{a}}$ is also a polynomial concerning the position.

External forces. We briefly introduce how to calculate the external force term in Eq. (17), taking $\iint_{\bar{\Omega}} \Phi^I \bar{\mathbf{f}} d\bar{\Omega}$ as an example. For an areal force, the integration cannot be avoided. Typical examples are gravity force (given by $\bar{\mathbf{f}}_{\text{grav}} = \bar{\rho}h\mathbf{g}$ with \mathbf{g} standing for the gravity acceleration) and pressure force (given by $\bar{\mathbf{f}}_{\text{press}} = p\mathbf{n}$ with p denoting the magnitude of pressure). On the other hand, a point force $\bar{\mathbf{f}}_{\text{pt}}$ exerted at an arbitrary point can be reformulated as an areal force multiplied by a Dirac δ function. Given the point $\mathbf{x}_{\text{pt}}(\xi_{\text{pt}}^1, \xi_{\text{pt}}^2)$, the closed integral form of $\bar{\mathbf{f}}_{\text{pt}}$ can be calculated by

$$\iint_{\bar{\Omega}} \Phi^I \bar{\mathbf{f}}_{\text{pt}} \delta(\mathbf{x}_{\text{pt}}) d\bar{\Omega} = \sum_I \Phi^I(\xi_{\text{pt}}^1, \xi_{\text{pt}}^2) \bar{\mathbf{f}}_{\text{pt}}, \quad (19)$$

which implies that a point force only influences the 16 generalized coordinates of its nearest 4 nodes.

5. Ray-patch intersection detection

We can equivalently transform a BH surface into a Bézier surface by regarding the cubic polynomials as linear combinations of *Bernstein basis polynomials* of degree 3:

$$\mathbf{x}(\xi^1, \xi^2) = \sum_{i=0}^3 \sum_{j=0}^3 B^i(\xi^1) B^j(\xi^2) \mathbf{p}_{ij}, \quad (20)$$

where $B^i(x)$ is defined as $\binom{3}{i} x^i (1-x)^{3-i}$ and $\mathbf{p}_{ij} \in \mathbb{R}^3$ denote the control points. A Bézier form provides us with a strong property that a surface lies completely within the convex hull of its control points, and thus also completely within the bounding box of them in any given Cartesian coordinate system, which lays the foundation of our ray-surface intersection detection algorithms.

The surface intersection is an important geometric operation in CAGD. We mainly focus on finding the intersection point between a ray and a surface in this paper. According to whether the surface moves during ray propagation, two kinds of intersections are detected.

5.1. Static ray–patch intersection

Given a ray defined as

$$\mathbf{x}_{\text{ray}}(\tau) = \mathbf{x}_0 + \tau \mathbf{d}, \quad (21)$$

where \mathbf{d} represents the direction and τ denotes the (pseudo) time, the goal of intersection tests is to find (ξ^1, ξ^2) and minimum $\tau > 0$ such that $\mathbf{x}_{\text{ray}}(\tau) = \mathbf{x}(\xi^1, \xi^2)$ holds. The convex-hull property of Bézier surfaces implies that for any ξ^1, ξ^2 , and τ satisfying this intersection equation, the following inequalities hold:

$$\min_{i,j \in \{0,1,2,3\}} \{\mathbf{p}_{ij} \cdot \hat{\mathbf{e}}_k\} \leq (\mathbf{x}_0 + \tau \mathbf{d}) \cdot \hat{\mathbf{e}}_k \leq \max_{i,j \in \{0,1,2,3\}} \{\mathbf{p}_{ij} \cdot \hat{\mathbf{e}}_k\}. \quad (22)$$

Here, $\hat{\mathbf{e}}_k$ ($k \in \{1, 2, 3\}$) denotes the k th vector of the standard basis. These inequalities correspond to an intersection test between the ray and an axis-aligned bounding box (AABB). By solving Eq. (22) for τ , we can obtain an interval $[\tau_{\min}, \tau_{\max}]$ or determine that no $\tau > 0$ satisfies all the inequalities, indicating that the ray does not intersect the current surface.

When the former case is encountered after solving Eq. (22), a divide-and-conquer strategy is employed. The Bézier surface is split into four smaller Bézier surfaces using De Casteljau's algorithm [71] at specific parameter coordinates. This recursive process allows for ray-AABB tests to be performed, enabling the elimination of surfaces that cannot be intersected by the ray. This continues until either all surfaces are discarded or an arbitrarily small interval of τ is obtained.

The above subdivision-based algorithm is trivial but effective and robust, converging linearly to the accurate intersection point. To minimize the number of surfaces requiring collision checks, the selection of the next surface is improved by using a min-heap of candidate surfaces. Every time a new surface is generated, it is inserted into the heap with τ_{\min} as the key. When moving to another surface, the top of the heap is chosen. The first surface that satisfies the termination condition provides an accurate approximation of the minimum τ among all the intersection points.

Finally, Newton's method is employed to determine where to subdivide surfaces by directly solving $\mathbf{x}(\xi^1, \xi^2) = \mathbf{x}_{\text{ray}}(\tau)$. The central parameter coordinates of a surface serve as the initial guess for Newton's method during the subdivision process. Within a fixed number of iterations, if Newton's method converges to $\xi^1 \in [u_{\min}, u_{\max}]$ and $\xi^2 \in [v_{\min}, v_{\max}]$, we obtain an intersection point and use the corresponding τ to update the current optimal solution τ_{opt} if τ is smaller and positive. Note that this intersection point may not be the closest one. We remove it from the search space by abandoning the neighborhood of (ξ^1, ξ^2) in the parameter space and subdividing the other region as illustrated in Fig. 9(b). In the case Newton's method fails to converge, the surface is split from its midpoint as shown in (Fig. 9(a)). The entire algorithm is summarized in Alg. 1.

5.2. Dynamic ray–patch intersection

Assuming constant velocity within a single time step Δt , a time-varying Bézier surface is given by

$$\mathbf{x}(\xi^1, \xi^2, t) = \sum_{i=0}^3 \sum_{j=0}^3 B_i^3(\xi^1) B_j^3(\xi^2) (\mathbf{p}_{ij} + t \dot{\mathbf{p}}_{ij}), \quad (23)$$

where \mathbf{p}_{ij} and $\dot{\mathbf{p}}_{ij}$ are the initial position and velocity of the (i, j) -th control point, and $t \in [0, \Delta t]$ indicates the time. The trajectory of a moving point can be depicted as a parameterized ray as

$$\mathbf{x}_p(t) = \mathbf{x}_{p0} + t \dot{\mathbf{x}}_p, \quad (24)$$

where \mathbf{x}_{p0} and $\dot{\mathbf{x}}_p$ denote the initial position and velocity.

According to the convex-hull property, a necessary condition for the intersection to occur at the moment t is that the point lies inside the

ALGORITHM 1: Static Intersecting Detection

Input: The ray \mathbf{x}_{ray} and the surface.

Output: The nearest intersection point (ξ^1, ξ^2) or **null** indicating no intersection.

```

 $\tau_{\text{opt}} \leftarrow \infty;$ 
Construct a min heap;
if IntersectsAABB( $\mathbf{x}_{\text{ray}}$ , the whole surface) then Push the
  whole surface into the heap with its  $\tau_{\min}$  as the key;
else return null;
while the heap is not empty do
  The current surface  $\leftarrow$  the top of the heap;
  if  $\tau_{\min} > \tau_{\text{opt}}$  then break;
  if the surface is small enough then return the middle
    parameter coordinates of the surface ;
  Find an intersection point by Newton's method;
  if converged then
    Update  $\tau_{\text{opt}}$  and  $(\xi_{\text{opt}}^1, \xi_{\text{opt}}^2)$ ;
    Subdivide the surface by the scheme in Fig. 9(b);
  else
    Subdivide the surface by the scheme in Fig. 9(a);
  end
  foreach subdivided surface do
    if IntersectsAABB( $\mathbf{x}_{\text{ray}}$ , the subdivided surface) then
      Push the surface into the heap with its  $\tau_{\min}$  as the key;
    end
  end
end
if  $\tau_{\text{opt}} \neq +\infty$  then return  $(\xi_{\text{opt}}^1, \xi_{\text{opt}}^2)$ ;
else return null;

```

axis-aligned bounding box of the surface at that moment, which can be written as

$$\min_{i,j} \{(\mathbf{p}_{ij} + t \dot{\mathbf{p}}_{ij}) \cdot \hat{\mathbf{e}}_k\} \leq (\mathbf{x}_{p0} + t \dot{\mathbf{x}}_p) \cdot \hat{\mathbf{e}}_k \leq \max_{i,j} \{(\mathbf{p}_{ij} + t \dot{\mathbf{p}}_{ij}) \cdot \hat{\mathbf{e}}_k\}, \quad (25)$$

with $i, j \in \{0, 1, 2, 3\}$. By solving Eq. (25) for t similar to the approach described in Section 5.1 for τ , we can employ a similar subdivision-based algorithm.

For each inequality on the left-hand side, the result is equivalent to the union of the intervals solved from the inequalities

$$(\mathbf{p}_{ij} + t \dot{\mathbf{p}}_{ij}) \cdot \hat{\mathbf{e}}_k \leq (\mathbf{x}_{p0} + t \dot{\mathbf{x}}_p) \cdot \hat{\mathbf{e}}_k, \quad (26)$$

where i and j ranges in $\{0, 1, 2, 3\}$. The same interpretation holds for the right-hand side of Eq. (25). The solution of Eq. (25) is the intersection of the two unions, which can be found using segment trees or greedy algorithms. If no feasible solution exists within $[0, \Delta t]$, no intersection is detected during the time step. Otherwise, we can utilize the divide-and-conquer framework described in Section 5.1 to recursively detect possible intersections. Similarly, Newton's method can be employed for acceleration.

6. Implementation

6.1. Implicit BHEM solver

Having established spatially discretized equations of motion (Section 4.3), we here discuss how to discretize Eq. (17) in time.

As a preparation, we assemble the mass matrix $\mathbf{M} \in \mathbb{R}^{3N \times 3N}$ by a 16-point Gauss–Legendre quadrature method (see the supplementary material) according to the component definition in Eq. (18). The generalized coordinates and the corresponding generalized forces

exerted on them are also numerically integrated and stacked into a $3N$ -dimensional q and F , respectively. Then Eq. (17) is reformulated as

$$\mathbf{M}\ddot{q} = F, \quad (27)$$

which is further discretized by the implicit Euler scheme in time:

$$\begin{cases} \dot{q}^{n+1} = \dot{q}^n + \Delta t \mathbf{M}^{-1} F(q^{n+1}, \dot{q}^{n+1}, t^{n+1}), & \text{(a)} \\ q^{n+1} = q^n + \Delta t \dot{q}^{n+1}. & \text{(b)} \end{cases} \quad (28)$$

By convention, we adopt Newton's method to solve the above equations iteratively, where quadratic and cubic line searchers referring to the widely used library *ArcSim* [21,72] are integrated. During each Newton step, the linear system is solved by a direct sparse LDLT Cholesky factorizations [73].

Besides, for the aforementioned positional constraints at arbitrary points, we use the augmented Lagrangian method [74] to solve the optimization problem, which acquires a better convergence rate than a pure method of Lagrange multipliers.

The Hessian matrix. In a Newton's method, the analytical form of $\partial F / \partial q_I$ ($I = 1, 2, 3, \dots, N$) is required. The most tricky part of the derivatives is the contribution of the elastic force, namely the Hessian matrix of V_e w.r.t. the generalized coordinates. We have carefully derived the elastic energy's first- and second-order derivatives for the BHEM solver and provided the concrete formulations in the supplementary material. In some situations, an inexact Hessian matrix may reduce the time cost of the solver's convergence, due to the high overhead of assembling the exact one. This alternative, which we called the *pseudo* Hessian matrix, is also given in the appendix. Moreover, when facing with ill-conditioned Hessian matrix, a diagonal regularizer is added to make the matrix positive definite.

6.2. Collision handling

We utilize the point-to-surface intersection detection scheme for CCD through sampling strategies adjusted to specific scenarios. For collision detection between shells and colliders with simple geometry shapes, such as spheres, cylinders, or planes, the midsurface is uniformly sampled. As for the collision detection with colliders possessing complex shapes, we sample the collider surface, or directly take the surface vertices as the sample points, if the collider has a triangular mesh. For the thin-shell self-collision, we iteratively sample each Hermite patch and perform CCD on sampling points against all the other surfaces. Penetrations occurring within a single surface can be handled by subdividing the current surface with the scheme illustrated in Fig. 9(b) and treating the 4 subdivided regions recursively.

After we have examined all the collision primitive pairs and got the list of earliest simultaneous collision pairs, we roll back to the collision moment. This prohibits penetration throughout the simulation. Then we follow the impulse-based method [75] to compute the simultaneous collision responses according to the conservation of momentum. Using a zero restitution coefficient [2], we model each collision as the momentum change at the collision point, which is formulated as a linear constraint of generalized velocities. Then we handle the multiple collisions in one batch by solving the linear system of constraints for the impulses in a least-square sense and update velocities. When friction exists, we compute the applied frictional impulses on the tangent direction of collision points according to the Column cone and solve for the normal and tangent updates together.

To avoid the resolved collision being detected at the beginning of the next round of CCD, we additionally update the displacement by pushing out each collision point a subtle distance opposite the normal direction. This is also modeled as a constraint linear to the generalized positions. Then the position constraints are resolved in a similar way after the velocity updates.

6.3. Rendering

In a ray-tracing rendering framework, the most expensive part is to test whether, where, and when a ray intersects the object. Traditional subdivision-based methods own their advantage in robustness but often fall in efficiency. We implement our ray-tracing algorithm in the rendering system of *pbrt-v4* [76] using our algorithm of static ray-patch intersection detection. Our enhancements to the subdivision-based algorithm have significantly accelerated performance, achieving an approximate 10-fold increase in speed compared to theoretically quadratic-convergent algorithms, such as Bézier clipping, as shown in Table 1 and Fig. 10. Newton's method tends to converge efficiently when applied to patches of low curvature. For example, the convergence percentages for each of the subdivisions in the three scenes depicted in Fig. 10 are 99%, 54%, and 34%, respectively. Additionally, Newton's method is adept at identifying the nearest intersection point. Consequently, the specially designed subdivision scheme can effectively prune much of the division tree.

7. Experimental results

We design a wide range of validation tests and simulation experiments to evaluate the validity, fidelity, and effectiveness of our framework from various aspects. All of the experiments are run on a 3.50 GHz 13th Gen Intel® Core™ i5-13600KF desktop with 32 GB RAM. The parameters and statistics are reported in Table 2.

In comparison with linear FEMs, we include two different bending formulas: the hinge energy proposed by Grinspun et al. [3] and the mid-edge bending (the discrete Koiter's shell energy) [22,70] implemented in *LibShell* [77]. It should be noted that the latter, along with the energy model employed in BHEM, directly originates from the continuous equations of the first and second fundamental forms. For the sake of fairness, we have standardized the linear search algorithm in each method.

7.1. Validation

Wrinkled sheets. We first validate the accuracy and convergence of our BHEM method with a standard stretched sheet experiment, which was first proposed by the pioneering work of Cerda et al. [78] and later used in both physical engineering [79] and computer graphics [18] community as well. In this experiment, a rectangular thin sheet is pulled apart from its two ends. Due to the high Poisson ratio, the sheet compresses in the perpendicular (vertical) and generates horizontal wrinkles.

We simulate this problem with the same physical parameter settings as in the work of Chen et al. [18], which are 0.25 m \times 0.1 m for size, 0.1 mm for thickness, Poisson ratio $\nu = 0.5$ and Young's modulus $Y = 1$ MPa. According to the conclusion of Chen et al. [18], the triangular finite-element-based method needs a high resolution of up to tens of thousands of vertices to generate correct wrinkle patterns in this example. As shown in Fig. 12, the BHEM starts to give apparent wrinkles at a resolution of merely 10^2 patches (1452 DoFs in total). As the resolution increases over 15^2 , the difference in the shape of wrinkles is already imperceptible.

We further verify this observation through a quantitative experiment that measures the peak amplitude of wrinkles produced with increased resolution. According to the physical experiment results reported by Wang et al. [79], the sheet is expected to produce wrinkles with a peak amplitude of 0.35 mm. The curve illustrated in Fig. 16 shows that our BHEM can stably yield 0.34 mm peak amplitude when the simulation resolution is greater than 30^2 patches (11.5k DoFs in total). While the peak amplitude can only reach 0.31 mm for the linear FEM simulated on a triangular mesh with 130k vertices (390k DoFs in total).



Fig. 1. A sheet of square cloth (30×30 patches) drapes on a ball, exhibiting rich wrinkle patterns in the process of reaching a steady state.



Fig. 2. A sheet of square cloth (30×30 patches) drapes on an armadillo model, where the sharp edges of the model are reflected by the cloth deformation.




Fig. 3. A sheet of cloth (30×30 patches) falls on a needle array and then gets pulled away from aside. The bulges on the cloth surface are pushed by the needle tips. This sharp geometry deformation caused by the tiny contact demonstrates well the fine resolution of the patch interpolation.



Fig. 4. The folding process of an oriental paper parasol (280×5 patches), simulated by jointly controlling nodal positions and their first-order derivatives. Bending and subtle wrinkles along the ribs can be observed.




Fig. 5. A surface that is homeomorphic to a rectangle is embedded into a Cartesian grid, of which grid cells are taken as patches. Each BH patch corresponds to an axis-aligned rectangle in the parameter space.

In Fig. 13, we also demonstrate the efficiency of BHEM by comparing its convergence curves with those of randomly triangulated thin shells by aforementioned linear FEMs at various resolutions. Generally speaking, linear FEM solvers require fewer iterations but significantly more time to achieve a converged solution. This is probably because the BH system mixes DoFs of different orders, which, in the meanwhile, allows the BH surface to present similar high-frequency visual effects with fewer DoFs.

Draped cloth. We conduct a set of comparative experiments to further demonstrate the superiority of our geometric discretization format. In




Fig. 6. Some complex surfaces can be divided into BH patches with the help of continuity boundary conditions.




Fig. 7. A hollow cylindrical shell (20×20 patches) deforms severely under the gradually increased compression. Initially, the cylinder is mounted on the ground; as the wood plank gets pushed down, the deformation of the cylinder manifests locally at first; at some point, the unshaping takes a sudden change; further compression produces severe deformation and high degree of self contact.

this scenario, a square piece of cloth, with its four corners moved inward a bit and clamped, drapes from a flattened configuration under the influence of gravity. In Fig. 14, we demonstrate the results obtained with our BHEM and *LibShell* at varying discretization resolutions respectively. The comparison shows that BHEM is capable of producing more detailed wrinkles with a comparable number of DoFs (11.5k DoFs for 30^2 -patch BH surface, and 12k DoFs for 4k-vertex FE surface). Additionally, at low discretization resolutions, BHEM (432 DoFs for 5^2 patches) tends to yield an over-smoothed surface, whereas *LibShell* (1.5k DoFs for 500 vertices) generates artificial wrinkles, as illustrated in the first column. At high discretization resolutions, although the quality produced by *LibShell* improved, however, BHEM can achieve similar quality with only 11.5k DoFs (30^2 BH patches), thereby requiring significantly fewer iterations and less time to converge, as depicted in Fig. 15. The convergence speed of different methods is provided in Fig. 15, including results from the linear FEM with hinge-based bending.

Parametric surface rendering. Our proposed ray-surface intersection algorithm would be regarded as a substantial enrichment for the current off-the-shell ray-tracing rendering engine. By simply substituting the current ray-polygon intersection detection module in *pbrt-v4* [76] with our algorithm, *pbrt-v4* can realize parametric surface ray-tracing rendering with good visual effects, as shown in the rightmost columns of Figs. 8 and 11. The other columns are the rendering results for the polygon meshes generated from the parametric surface subdivision. From left to right, the subdivision level gradually increases. A parametric surface naturally processes continuous normal vectors on its surface. Discontinuities in surface normal vectors can result in visual artifacts, such as abnormal reflections on the teapot and cluttered caustics on the ground.

2D cantilever beam and 3D lateral buckling. We further validate our method by comparing two experiments with their theoretical solutions [80]. The first is the cantilever beam experiment, where a beam is fixed at one end and bent under gravity. Its master curve uniquely determines the aspect ratio H/W of the cantilever beam under equilibrium as a function of the dimensionless parameter $\Gamma^* = 12(1 - \nu^2)\rho g L^3 / Y h^2$. The second is the lateral buckling experiment with the plate lying vertically in the (x, z) plane. We let the plate hang and sag



Fig. 8. Parametric surface rendering. A Utah teapot represented in bicubic patches is rendered by our scheme, shown in (e). The teapots rendered by traditional methods with each patch triangulated into 4/16/64/256 facets are shown in (a)/(b)/(c)/(d), where artifacts can be seen at the mouth and silhouettes of the teapots, especially for coarse mesh decimation.




Fig. 9. Ray surface intersection subdivide schemes. Left: splitting at midpoint; Right: abandoning the neighborhood of intersection point and splitting the rest region.




Fig. 10. Three test cases with only 1 patch for Plane and 10×10 patches for Cylinder and Drape each. The scenes all have the same setting of lights and 16 samples per pixel.




Fig. 11. Parametric surface rendering. The caustic lighting effect amplifies the imperfection of the surface, such as the un-smooth surface normal distribution. Fine-grained caustic rays clutter the ground for the scenarios with low-resolution mesh objects.




Fig. 12. Wrinkled sheets. Sheets of different resolutions buckle under uniaxial stretching. (a) There appears only one artificial bump on a single patch. (b) No apparent wrinkle on a shell composed of 5×5 patches. (c) A 10×10 -patch sheet manages to produce several shallow wrinkles. (d) Clearer wrinkles are produced by a 15×15 -patch sheet. (e) Though wrinkles are finer, the difference is subtle compared with the previous. (f) The wrinkle pattern converges when it is finer than 30×30 patches.

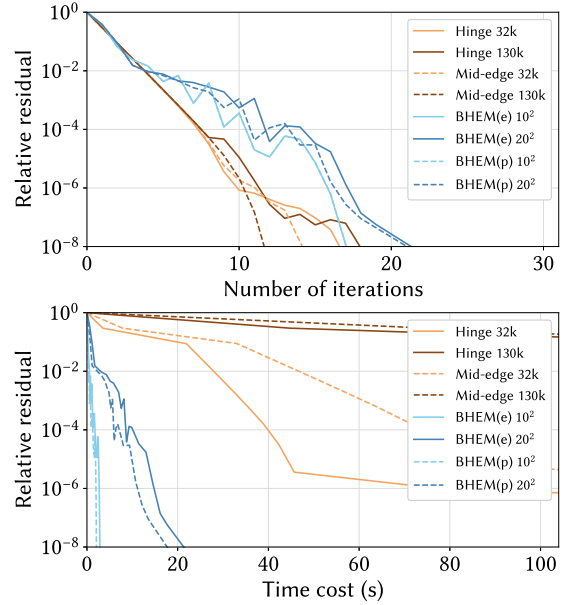


Fig. 13. Performance curves of wrinkled sheets with respect to different resolutions. Linear FEM solvers employing hinge-based bending and mid-edge bending are denoted as *hinge* and *mid-edge*, respectively. We use *BHEM(p)* and *BHEM(e)* to distinguish between BHEM solvers that utilize pseudo-Hessian and exact Hessian matrices, respectively. Note that the light-blue solid and dashed lines overlap in the upper subplot. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

under its weight, waiting for a bifurcation to occur. Its master curve determines when the plate buckles in the third direction as Γ^* increases under a given aspect ratio w/L . We conduct our experiments using methods and physical parameters provided by Romero et al. [80]. The beam in the cantilever test is discretized in 2×100 patches. And the plate is discretized in 20 patches per meter. Our results perfectly match the master curves as shown in Figs. 18 and 19.

7.2. Experiments

Locking. As a high-order method, the BHEM formulation alleviates locking issues remarkably. It does not need any special treatment, e.g., dynamic remeshing, to get plausible effects, even with very few degrees of freedom. As shown in the top row of Fig. 20, a squared piece of cloth, which is composed of 10^2 patches, is initially pinned at its two diagonal corners. As the pinning points slide along the boundary, the cloth can naturally fold down along any direction as illustrated in the bottom row.

Twisting cloth. Our method allows precise and intuitive control over the first-order derivatives of an arbitrary point on the surface. In this example, the center point of a squared piece of cloth is fixed. Its two first-order partial derivatives (indicated by the red arrows in Fig. 17 are rotated with a constant speed in the horizontal plane. This prescript motion results in a persistent wrinkling perpendicular to the first

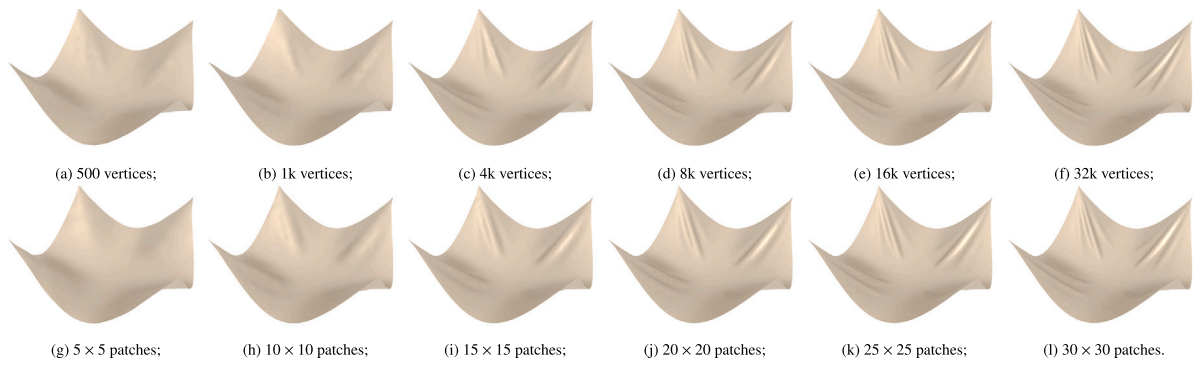


Fig. 14. Draped cloth. The top row and bottom row show the results obtained by *LibShell* and BHEM, respectively. The resolution of the discretized surfaces increases from left to right. We choose the examples with the highest resolutions and similar plausible wrinkles and test the performance of the solvers under these resolutions. The performance curves are shown in Fig. 15.

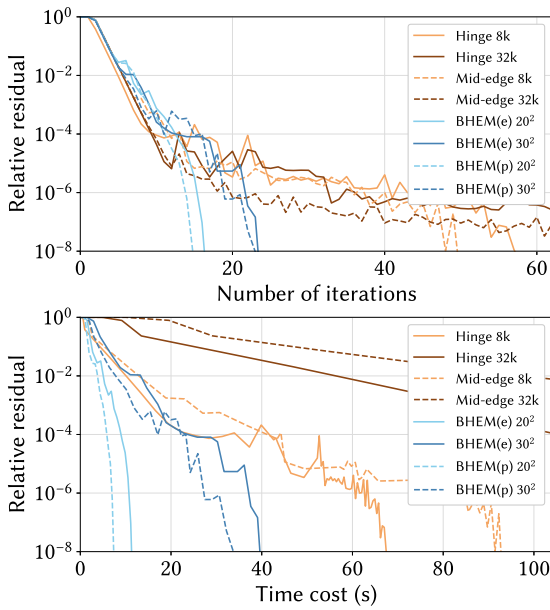


Fig. 15. Performance curves of draped cloth with respect to different resolutions. Same abbreviations are used as in Fig. 13.

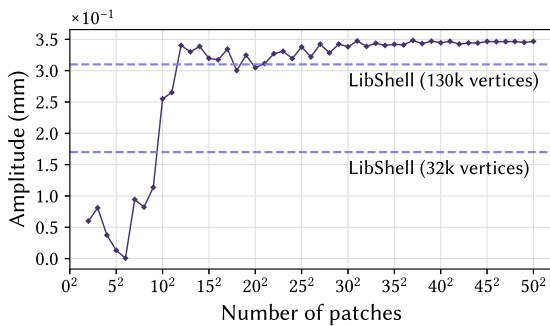


Fig. 16. Statistic of the wrinkle's peak value. The simulated amplitude increases with the refinement of the sheet and converges rapidly to a physical real value (3.5×10^{-1} mm). Results of *LibShell* with 32k and 130k vertices are plotted in horizontal dashed lines.

derivative directions and eventually invokes the buckling swirl around the center. In a pure displacement-based method, a similar result can only be achieved through prescript at least three nodes' motion.




Fig. 17. Twisting cloth (30×30 patches). A buckling swirl around the center as the first-order derivatives (indicated by the red arrows) of the center point rotate horizontally with a constant speed.

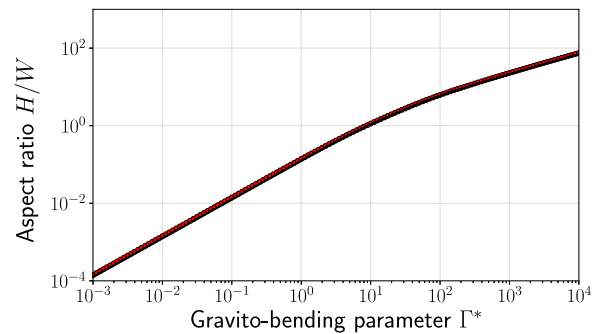


Fig. 18. Comparisons with theoretical solutions on the cantilever test. We simulate 140 Γ^* values and superimpose the data (red dots) onto the master curve (black line). Our results perfectly match the master curve.




Fig. 19. Comparisons with theoretical solutions on the lateral buckling test. The master curve separates two areas colored to indicate whether the plate has buckled in 3D (turquoise) or lies in 2D (orange). Our results perfectly match the theoretical solution. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



Fig. 20. A thin sheet (10×10 patches) with its two diagonal ends fixed bends by gravity. After reaching a steady state, the fixed points slide smoothly along the opposite edges without any locking artifacts from discretization. We emphasize that only 10×10 patches are used here.




Fig. 21. Hollow cylindrical shells (30×30 patches) with different thickness buckle under the gradually increased axial compression. Inside each shell, there is a rigid cylindrical mandrel. Its radius differs from that of the shell by 10%.

Cloth draping on objects. Both resting and dynamic contact can be faithfully detected and resolved within our framework. Two square sheets of cloth, both of which are composed of 30^2 patches, drape on a parametric sphere and a triangular meshed armadillo are demonstrated in Figs. 1 and 2 respectively. Rich deformation details due to the interference from external objects can be observed in the results. Please see the supplemental video for more visual evidence.

Cloth sliding over needles. In this example, we drop a sheet of cloth on a needle array and then pull it away from aside. The subtle bulges on the cloth surface, which are pushed out by the needle tips, can be clearly observed in Fig. 3. In the supplemental video, we can easily notice that when pulling the cloth over the needle array, the cloth exhibits natural choppy movement around the needle tips due to the contact interaction. These two points verify that our BHEM formulation and collision detection algorithm can handle sharp geometry features robustly.

Folding an oriental paper parasol. By jointly controlling node positions and their first-order derivatives, we can mimic the folding process of an oriental paper parasol driven by the motion of its rib as displayed in Fig. 4.

Cylindrical shell buckling. In Fig. 7, a hollow cylindrical shell severely buckles under the gradually increased compression is simulated with our method. When there is an inner mandrel, local buckles can no longer grow unrestrained because their radial displacements are arrested by the mandrel. Consequently, as illustrated in Fig. 21, each buckle folds along its central ridge, orthogonal to both the pressure and the axis of the cylinder. Under compression, more buckles emerge, forming a densely packed, diamond-like pattern across the entire cylinder. Given that the bending stiffness of a shell is proportional to the cube of its thickness, thicker shells display more regular, rapid, and stable buckling development, as shown in the accompanying video.

Table 1

Rendering performance. The three algorithms from left to right are: our subdivision-based algorithm that splits only at the midpoint, our algorithm with Newton's method, and Bézier clipping.

	Subdivision	Subdivision (Opt.)	Bézier clipping
Plane	58.3 s	4.5 s	6.0 s
Cylinder	91.6 s	9.7 s	14.2 s
Drape	86.0 s	14.8 s	33.9 s

Furthermore, our experiment setup ensures a constant ratio between the mandrel radius and the shell thickness, thereby keeping the annular gap unchanged. As depicted in Fig. 21, the size of the buckle remains consistent regardless of variations in thickness, which aligns well with the conclusion presented in Seffen et al. [81].

8. Conclusions and discussions

In this study, we propose a new computational framework designed for elastodynamic simulation of thin-shell structures. The central piece of our framework is a high-order finite element formulation equipped with an implicit Euler solver. This formulation, based on bicubic Hermite interpolation, naturally ensures conforming C^1 continuity. Capitalizing on the advancements in surface modeling and rendering, we have crafted an intersection detection paradigm that is custom-tailored for bicubic Hermite surfaces. This unified approach empowers us to achieve high-fidelity CCD and rendering without resorting to any form of auxiliary tessellation mesh. The significance of research on BHEM spans various applications: (1) facilitating a seamless flow of information throughout the design-to-production workflow (CAD, CAE, and CAM), where geometric consistency with IGA is needed, (2) enhancing inverse optimization and control, where fewer DoFs (compact representation) aid in fast convergence, (3) enabling locking-free simulation with rigorous strain limiting, which necessitates sufficient DoFs of each vertex, and (4) advancing development of conforming contact solver, where accurately capturing contact regions is essential.

Rotation dependence. The basis functions for tensor-product bicubic Hermite splines are $(\xi^1)^i(\xi^2)^j$ ($0 \leq i, j \leq 3$), which include a complete set of rotation-free basis functions, specifically $(\xi^1)^i(\xi^2)^j$ ($0 \leq i + j \leq 3$). Despite that the geometric representation of BHE is affected by the directions of parameterization, the differences only exist in several high-order bases.

Comparisons with relevant methods. The BHEM achieves an impressive balance between computational efficiency and geometric continuity, compared to relevant methods. For instance, subdivision surfaces also ensure C^1 continuity but do not interpolate control points, resulting in difficulties of enforcing boundary conditions and detecting collisions. Specifically, subdivision surfaces utilizing Loop's scheme [8] employ quartic shape functions, which demand more computation than ours. Meanwhile, those using Catmull–Clark's scheme [30] generate bicubic B-splines like ours, but their values and derivatives of the shape functions cannot be evaluated analytically. In the case of DG-FEM, although it employs lower-order (quadratic) shape functions, it necessitates an additional penalty between each pair of adjacent elements, which is expected to quickly offset the extra computational cost of evaluating higher-order polynomials as mesh size increases. Moreover, for the same number of DoFs, DG-FEM has been observed to be less accurate than the FEMs used in CG [32].

Limitations and future work. Some challenges remain unresolved and await deeper investigation. First, in order to maintain C^1 -continuous everywhere, our current framework is limited to working with surfaces that are homeomorphic to squares, cylinders, or tori. For more complex shapes (e.g., spheres), special treatments, such as integrating hinge-based bending energy, are still necessary around C^0 seams and

Table 2

Here we list the parameters and the time consumed in each example, including the degrees of freedom, the total number of nonzero elements in the BHEM system matrix, the total number of sampling points on the external collider, the properties of the cloth (Young's Modulus Y [Pa], Poisson's ratio ν and thickness h [mm], mass density ρ [10^3 kg/m³]), Rayleigh damping coefficient α , time step Δt [ms], time consumed in one integration step t_{int} [s], in one CCD step t_{obj} [s], and whether we detect self-collision (SC, '-' means that CCD has not been performed in this example).

Figure	Example	#DoFs	#NNZs	#SPs	Y	ν	h	ρ	α	Δt	t_{int}	t_{obj}	SC
12	Wrinkled ^a	–	–	–	1×10^6	0.5	0.1	0.93	–	–	63.0	–	–
14	Draped ^a	–	–	–	1×10^5	0.3	1	0.93	0	10	6.01	–	–
17	Twisting	1.15×10^4	1.19×10^6	–	1×10^4	0.3	1	0.2	50	2	0.20	–	–
20	Locking ^a	–	–	–	1×10^4	0.3	1	0.2	2	2	3.38	4.37	Yes
1	Ball	1.15×10^4	1.17×10^6	–	1×10^5	0.3	0.1	0.2	2.5	2	4.80	0.0023	No
2	Armadillo	1.15×10^4	1.19×10^6	2.8k	8.21×10^5	0.243	0.32	0.4726	2.5	1	2.90	1.84	No
3	Needles	1.15×10^4	1.19×10^6	3.6k	1×10^4	0.3	0.1	0.2	5	2	5.05	26.66	Yes
4	Umbrella	2.02×10^4	1.76×10^6	–	1×10^6	0.5	0.1	0.93	5	2	12.47	–	–
7	Can	5.3×10^3	4.97×10^5	–	1×10^8	0.3	2	2.7	5	2	4.55	–	–
21	Cans ^a	1.15×10^4	1.13×10^6	–	1×10^9	0.47	–	1.4	5	0.04	5.29	0.0024	No

^a The statistics is obtained by 30×30 -patch shells for *Wrinkled* and *Draped*, 20×20 -patch shells for *Locking*, and 1 mm-thick shell for *Cans*.

points. Second, while Hermite interpolation excels in depicting complex characteristics like wrinkles or folds in cloth motion, it encounters difficulties when addressing plastic deformation and impact dynamics. The high-order formulation of Hermite interpolation exacerbates the inherent non-linearity and non-smoothness of these tasks, which need sophisticated modeling and analysis. Besides, this paper only introduces an algorithm of ray-patch intersection detection, which alone is insufficient for achieving non-penetration in CCD. To overcome this limitation, we intend to expand our algorithm to include patch–patch intersection detection. In addition, as our current collision-resolving scheme struggles with intricate contact scenarios like tying ribbons into a reef knot, we view the integration of Incremental Potential Contact (IPC) as a natural and significant follow-up step.

CRediT authorship contribution statement

Xingyu Ni: Conceptualization, Formal analysis, Investigation, Methodology, Software, Validation, Writing – original draft, Writing – review & editing. **Xuwen Chen:** Data curation, Methodology, Software, Validation, Writing – original draft, Writing – review & editing. **Cheng yu:** Resources, Software, Visualization, Writing – review & editing. **Bin Wang:** Conceptualization, Formal analysis, Project administration, Supervision, Writing – original draft, Writing – review & editing. **Baoquan Chen:** Funding acquisition, Project administration, Supervision.

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

No data was used for the research described in the article.

Acknowledgement

This work was supported in part by National Key R&D Program of China (2022ZD0160801) and Shenzhen Collaborative Innovation Program (CJGJZD2021048092601003). The Houdini Education license is credited for the video generations.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.cad.2024.103734>.

References

- [1] Ciarlet PG. Theory of shells. Mathematical elasticity, vol. 3, Devon, England, UK: North Holland; 2000.
- [2] Bridson R, Fedkiw R, Anderson J. Robust treatment of collisions, contact and friction for cloth animation. ACM Trans Graph 2002;21(3):594–603.
- [3] Grinspun E, Hirani AN, Desbrun M, Schröder P. Discrete shells. In: Proceedings of the 2003 ACM SIGGRAPH/eurographics symposium on computer animation. Goslar, DEU: Eurographics Association; 2003, p. 62–7.
- [4] Bergou M, Wardetzky M, Harmon D, Zorin D, Grinspun E. Discrete quadratic curvature energies. In: ACM SIGGRAPH 2006 courses. New York, NY, USA: Association for Computing Machinery; 2006, p. 20–9.
- [5] Gingold Y, Secord A, Han JY, Grinspun E, Zorin D. A discrete model for inelastic deformation of thin shells. In: ACM SIGGRAPH/eurographics symposium on computer animation. Goslar, DEU: Eurographics Association; 2004, p. 1–12.
- [6] Hughes T, Cottrell J, Bazilevs Y. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. Comput Methods Appl Mech Eng 2005;194(39):4135–95.
- [7] Cottrell JA, Hughes TJR, Bazilevs Y. Isogeometric analysis: Toward integration of CAD and FEA. 1st ed. Hoboken, NJ, USA: Wiley Publishing; 2009.
- [8] Cirak F, Ortiz M, Schröder P. Subdivision surfaces: a new paradigm for thin-shell finite-element analysis. Internat J Numer Methods Engr 2000;47(12):2039–72.
- [9] Cirak F, Ortiz M. Fully C1-conforming subdivision elements for finite deformation thin-shell analysis. Internat J Numer Methods Engr 2001;51(7):813–33.
- [10] Lu J, Zheng C. Dynamic cloth simulation by isogeometric analysis. Comput Methods Appl Mech Eng 2014;268:475–93.
- [11] Phusakulkajorn W. Finite element based solutions of thin-shell problems with a small strain [Master's thesis], The University of Manchester; 2013.
- [12] Greco L, Cuomo M. An implicit G1-conforming bi-cubic interpolation for the analysis of smooth and folded Kirchhoff–Love shell assemblies. Comput Methods Appl Mech Eng 2021;373:113476.
- [13] Greco L, Cuomo M, Contrafatto L. A quadrilateral G1-conforming finite element for the Kirchhoff plate model. Comput Methods Appl Mech Eng 2019;346:913–51.
- [14] Zienkiewicz O, Taylor R, David F. The finite element method for solid and structural mechanics. Oxford: Butterworth-Heinemann; 2005.
- [15] Baraff D, Witkin A. Large steps in cloth simulation. In: Proceedings of the 25th annual conference on computer graphics and interactive techniques. New York, NY, USA: Association for Computing Machinery; 1998, p. 43–54.
- [16] Choi K-J, Ko H-S. Stable but responsive cloth. ACM Trans Graph 2002;21(3):604–11.
- [17] Tamstorf R, Jones T, McCormick SF. Smoothed aggregation multigrid for cloth simulation. ACM Trans Graph 2015;34(6).
- [18] Chen Z, Chen H-Y, Kaufman DM, Skouras M, Vouga E. Fine wrinkling on coarsely meshed thin shells. ACM Trans Graph 2021;40(5).
- [19] Wang H. GPU-based simulation of cloth wrinkles at submillimeter levels. ACM Trans Graph 2021;40(4).
- [20] Zhang JE, Dumas J, Fei YR, Jacobson A, James DL, Kaufman DM. Progressive simulation for cloth quasistatics. ACM Trans Graph 2022;41(6).
- [21] Narain R, Pfaff T, O'Brien JF. Folding and crumpling adaptive sheets. ACM Trans Graph 2013;32(4).
- [22] Chen H-Y, Sastry A, van Rees WM, Vouga E. Physical simulation of environmentally induced thin shell deformation. ACM Trans Graph 2018;37(4).
- [23] Burgoon R, Wood ZJ, Grinspun E. Discrete shells origami. In: Jackson DJ, editor. 21st international conference on computers and their applications, CATA-2006, Seattle, Washington, USA, March 23–25, 2006, proceedings. ISCA; 2006, p. 180–7.
- [24] Pfaff T, Narain R, de Joya JM, O'Brien JF. Adaptive tearing and cracking of thin sheets. ACM Trans Graph 2014;33(4).

- [25] Rémillard O, Kry PG. Embedded thin shells for wrinkle simulation. *ACM Trans Graph* 2013;32(4).
- [26] Li P, Kry PG. Multi-layer skin simulation with adaptive constraints. In: *Proceedings of the 7th international conference on motion in games*. New York, NY, USA: Association for Computing Machinery; 2014, p. 171–6.
- [27] Terzopoulos D, Platt J, Barr A, Fleischer K. Elastically deformable models. In: *Proceedings of the 14th annual conference on computer graphics and interactive techniques*. New York, NY, USA: Association for Computing Machinery; 1987, p. 205–14.
- [28] Bridson R, Marino S, Fedkiw R. Simulation of clothing with folds and wrinkles. In: *Proceedings of the 2003 ACM SIGGRAPH/eurographics symposium on computer animation*. Goslar, DEU: Eurographics Association; 2003, p. 28–36.
- [29] Thomaszewski B, Wacker M, Straßer W. A consistent bending model for cloth simulation with corotational subdivision finite elements. In: *Proceedings of the 2006 ACM SIGGRAPH/eurographics symposium on computer animation*. Goslar, DEU: Eurographics Association; 2006, p. 107–16.
- [30] Clyde D, Teran J, Tamstorf R. Modeling and data-driven parameter estimation for woven fabrics. In: *Proceedings of the ACM SIGGRAPH / eurographics symposium on computer animation*. New York, NY, USA: Association for Computing Machinery; 2017, p. 1–11.
- [31] Guo Q, Han X, Fu C, Gast T, Tamstorf R, Teran J. A material point method for thin shells with frictional contact. *ACM Trans Graph* 2018;37(4).
- [32] Kaufmann P, Martin S, Botsch M. Implementation of discontinuous Galerkin Kirchhoff-love shells. *CIT technical reports series* 622, 2009.
- [33] Wawrzinek A, Hildebrandt K, Polthier K. Koiter's thin shells on catmull-clark limit surfaces. In: Eisert P, Hornegger J, Polthier K, editors. *Vision, modeling, and visualization* (2011). Berlin, Germany: The Eurographics Association; 2011, p. 113–20.
- [34] Qin H, Terzopoulos D. Triangular NURBS and their dynamic generalizations. *Comput Aided Geom Design* 1997;14(4):325–47.
- [35] Trusty T, Chen H, Levin DW. The shape matching element method: Direct animation of curved surface models. *ACM Trans Graph* 2021;40(4).
- [36] Clough R, Tocher J. Finite element stiffness matrices for analysis of plates in bending. In: *Proceedings of the conference on matrix methods in structural mechanics*. 1965, p. 515–45.
- [37] Bell K. A refined triangular plate bending finite element. *Internat J Numer Methods Engrg* 1969;1(1):101–22.
- [38] Bogner F. The generation of interelement compatible stiffness and mass matrices by the use of interpolation formulae. In: *Proceedings of the conference on matrix methods in structural mechanics*. 1965, p. 397–444.
- [39] Petera J, Pittman JFT. Isoparametric Hermite elements. *Internat J Numer Methods Engrg* 1994;37:3489–519.
- [40] Beheshti A. Novel quadrilateral elements based on explicit Hermite polynomials for bending of Kirchhoff-Love plates. *Comput Mech* 2018;62:1199–211.
- [41] Šolín P, Segeth K. Hierarchic higher-order hermite elements on hybrid triangular/quadrilateral meshes. *Math Comput Simulation* 2007;76(1):198–204, *Mathematical Modelling and Computational Methods in Applied Sciences and Engineering*.
- [42] Farin G. *Curves and surfaces for CAGD: A practical guide*. 5th ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 2001.
- [43] Salomon D. *Curves and surfaces for computer graphics*. Berlin, Heidelberg: Springer-Verlag; 2005.
- [44] Kajiji JT. Ray tracing parametric patches. *SIGGRAPH Comput Graph* 1982;16(3):245–54.
- [45] Manocha D. Solving systems of polynomial equations. *IEEE Comput Graph Appl* 1994;14(2):46–55.
- [46] Benthin C. *Realtime ray tracing on current CPU architectures* [Ph.D. thesis], Saarbrücken, Germany: Saarland University; 2006, p. 1–194.
- [47] Abert O, Geimer M, Muller S. Direct and fast ray tracing of NURBS surfaces. In: *Symposium on interactive ray tracing*. Salt Lake City, UT, USA: IEEE; 2006, p. 161–8.
- [48] Geimer M, Abert O. Interactive ray tracing of trimmed bicubic bézier surfaces without triangulation. In: *International conference in central europe on computer graphics and visualization*. Bory, Plzen, Czech Republic: University of West Bohemia; 2005, p. 71–8.
- [49] Martin W, Cohen E, Fish R, Shirley P. Practical ray tracing of trimmed NURBS surface. *J Graph Tools* 2000;5.
- [50] Toth DL. On ray tracing parametric surfaces. *SIGGRAPH Comput Graph* 1985;19(3):171–9.
- [51] Barth W, Stuerzlinger W. Efficient ray tracing for bezier and B-spline surfaces. *Comput Graph* 1993;17:423–30.
- [52] Nishita T, Sederberg TW, Kakimoto M. Ray tracing trimmed rational surface patches. In: *Proceedings of the 17th annual conference on computer graphics and interactive techniques*. New York, NY, USA: Association for Computing Machinery; 1990, p. 337–45.
- [53] Campagna S, Slusallek P, Seidel H. Ray tracing of spline surfaces: Bézier clipping, Chebyshev boxing, and bounding volume hierarchy - a critical comparison with new results. *Vis Comput* 1997;13(6):265–82.
- [54] Rogers DF. *Procedural elements for computer graphics*. USA: McGraw-Hill, Inc.; 1984.
- [55] Woodward C. Ray tracing parametric surfaces by subdivision in viewing plane. In: *Theory and practice of geometric modeling*. Berlin, Heidelberg: Springer-Verlag; 1989, p. 273–87.
- [56] Sederberg T, Nishita T. Curve intersection using Bézier clipping. *Comput Aided Des* 1990;22(9):538–49.
- [57] Tejima T, Fujita M, Matsuoka T. Direct ray tracing of full-featured subdivision surfaces with bezier clipping. *J Comput Graph Tech (JCGT)* 2015;4(1):69–83.
- [58] Efremov A, Havran V, Seidel H-P. Robust and numerically stable Bézier clipping method for ray tracing NURBS surfaces. In: *Proceedings of the 21st spring conference on computer graphics*. New York, NY, USA: Association for Computing Machinery; 2005, p. 127–35.
- [59] Slusallek P, Seidel H-P. Vision - An architecture for global illumination calculations. *IEEE Trans Vis Comput Graphics* 1995;1(1):77–96.
- [60] Pabst H, Springer J, Schollmeyer A, Lenhardt R, Lessig C, Froehlich B. Ray casting of trimmed NURBS surfaces on the GPU. In: *Symposium on interactive ray tracing*. Salt Lake City, UT, USA: IEEE; 2006, p. 151–60.
- [61] Ferguson Z, Jain P, Zorin D, Schneider T, Panozzo D. High-order incremental potential contact for elastodynamic simulation on curved meshes. In: *ACM SIGGRAPH 2023 conference proceedings*. New York, NY, USA: Association for Computing Machinery; 2023, p. 1–11.
- [62] Von Herzen B, Barr AH, Zatz HR. Geometric collisions for time-dependent parametric surfaces. *SIGGRAPH Comput Graph* 1990;24(4):39–48.
- [63] Snyder JM, Woodbury AR, Fleischer K, Currin B, Barr AH. Interval methods for multi-point collisions between time-dependent curved surfaces. In: *Proceedings of the 20th annual conference on computer graphics and interactive techniques*. New York, NY, USA: Association for Computing Machinery; 1993, p. 321–34.
- [64] Hughes M, DiMattia C, Lin M, Manocha D. Efficient and accurate interference detection for polynomial deformation. In: *Proceedings computer animation '96*. 1996, p. 155–66.
- [65] Lu J. Isogeometric contact analysis: Geometric basis and formulation for frictionless contact. *Comput Methods Appl Mech Engrg* 2011;200(5):726–41.
- [66] Krishnan S, Gopi M, Lin M, Manocha D, Pattekar A. Rapid and accurate contact determination between spline models using ShellTrees. *Comput Graph Forum* 1998;17(3):315–26.
- [67] Teschner M, Kimmerle S, Heidelberger B, Zachmann G, Raghupathi L, Fuhrmann A, et al. Collision detection for deformable objects. *Comput Graph Forum* 2005;24(1):61–81.
- [68] Marschner Z, Zhang P, Palmer D, Solomon J. Sum-of-squares geometry processing. *ACM Trans Graph* 2021;40(6).
- [69] Zhang P, Marschner Z, Solomon J, Tamstorf R. Sum-of-squares collision detection for curved shapes and paths. In: *ACM SIGGRAPH 2023 conference proceedings*. New York, NY, USA: Association for Computing Machinery; 2023, p. 1–11.
- [70] Weischedel C. *A discrete geometric view on shear-deformable shell models* [Ph.D. thesis], Göttingen, Niedersachsen, Germany: Georg-August-Universität Göttingen; 2012.
- [71] Boehm W, Müller A. On de Casteljau's algorithm. *Comput Aided Geom Design* 1999;16(7):587–605.
- [72] Narain R, Samii A, O'Brien JF. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans Graph* 2012;31(6).
- [73] Guennebaud G, Jacob B, et al. *Eigen v3*. 2010, <http://eigen.tuxfamily.org>.
- [74] Hestenes MR. Multiplier and gradient methods. *J Optim Theory Appl* 1969;4:303–20.
- [75] Harmon D, Vouga E, Tamstorf R, Grinspun E. Robust treatment of simultaneous collisions. *ACM Trans Graph* 2008;27(3):1–4.
- [76] Pharr M, Jakob W, Humphreys G. *Physically based rendering, fourth edition: From theory to implementation*. 4th ed. Cambridge, MA, USA: The MIT Press; 2023.
- [77] Vouga E, Jourdan D. *LibShell*. 2021, <https://github.com/evouga/libshell>.
- [78] Cerda E, Mahadevan L. Geometry and physics of wrinkling. *Phys Rev Lett* 2003;90:074302.
- [79] Wang T, Fu C, Xu F, Huo Y, Potier-Ferry M. On the wrinkling and restabilization of highly stretched sheets. *Internat J Engrg Sci* 2019;136:1–16.
- [80] Romero V, Ly M, Rasheed AH, Charrondière R, Lazarus A, Neukirch S, et al. Physical validation of simulators in computer graphics: A new framework dedicated to slender elastic structures and frictional contact. *ACM Trans Graph* 2021;40(4).
- [81] Seffen KA, Stott SV. Surface texturing through cylinder buckling. *J Appl Mech* 2014;81(6):061001.