

A Flexible Mold for Facade Panel Fabrication

FLORIAN RIST*, KAUST, KSA
ZHECHENG WANG*[†], UofT, CA
DAVIDE PELLIS, ISTI-CNR, IT
MARCO PALMA, TU Wien, AT
DAOMING LIU, KAUST, KSA
EITAN GRINSPUN, UofT, CA
DOMINIK L. MICHELS, KAUST, KSA




Fig. 1. A facade composed of nine panels fabricated with our flexible mold, arranged as a 3×3 array. The facade surface (left) demonstrates the challenge of segmentation (second and third from the left) where seven out of nine panels exhibit both positive and negative Gaussian curvature, each with distinct curvature ranges, including one nearly developable panel. The fabricated and assembled facade (right) is shown, with panels mounted on a metal frame. The rendered virtual facade (left and second from the left) shows a remarkable visual similarity with the fabricated facade, capturing even the minutest details, such as panel gaps. This exemplifies the scalability and precision of our solution, advancing towards a reusable, zero-waste, and cost-efficient mold for transformative architectural applications.

Architectural surface panelling often requires fabricating molds for panels, a process that can be cost-inefficient and material-wasteful when using traditional methods such as CNC milling. In this paper, we introduce a novel solution to generating molds for efficiently fabricating architectural panels. At the core of our method is a machine that utilizes a deflatable membrane as a *flexible mold*. By adjusting the deflation level and boundary element positions, the membrane can be reconfigured into various shapes, allowing for mass customization with significantly lower overhead costs. We devise an efficient algorithm that works in sync with our *flexible mold* machine

that optimizes the placement of *customizable* boundary element positions, ensuring the fabricated panel matches the geometry of a given input shape: (1) Using a quadratic Weingarten surface arising from a natural assumption on the membrane's stress, we can approximate the initial placement of the boundary element from the input shape's geometry; (2) we solve the inverse problem with a simulator-in-the-loop optimizer by searching for the optimal placement of boundary curves with sensitivity analysis. We validate our approach by fabricating baseline panels and a facade with a wide range of curvature profiles, providing a detailed numerical analysis on simulation and fabrication, demonstrating significant advantages in cost and flexibility.

*F. Rist and Z. Wang contributed equally to this work.

[†]Z. Wang has worked on the project at KAUST and UofT.

Authors' addresses: Florian Rist, KAUST, Thuwal 23955, KSA; Zhecheng Wang, UofT, Toronto, ON M5S, CA; Davide Pellis, ISTI-CNR, Pisa, IT; Marco Palma, TU Wien, Vienna, AT; Daoming Liu, KAUST, Thuwal 23955, KSA; Eitan Grinspun, UofT, Toronto, ON M5S, CA; Dominik L. Michels, KAUST, Thuwal 23955, KSA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

0730-0301/2024/12-ART231

<https://doi.org/10.1145/3687906>

CCS Concepts: • **Computing methodologies** → **Modeling and simulation**.

Additional Key Words and Phrases: Computational Fabrication, Freeform Architecture, Geometric Initialization, Inverse Design, Mass Customization, Panel Fabrication, Flexible Mold, Thin Shell Simulation, Weingarten Surfaces.

ACM Reference Format:

Florian Rist, Zhecheng Wang, Davide Pellis, Marco Palma, Daoming Liu, Eitan Grinspun, and Dominik L. Michels. 2024. A Flexible Mold for Facade Panel Fabrication. *ACM Trans. Graph.* 43, 6, Article 231 (December 2024), 15 pages. <https://doi.org/10.1145/3687906>

1 INTRODUCTION

Contemporary architectural streams, such as ‘blobitecture’, often showcase striking freeform facades, aiming to emulate organic shapes. Iconic structures like Foster’s Sage Gateshead music center (2004), Hadid’s Heydar Aliyev Center (2007), the Dongdaemun Design Plaza (2009), and Snøhetta’s King Abdulaziz Centre for World Culture (Ithra) (2017) stand as prominent examples.

Cladding these complex facades is a challenging task, as it requires the production of panels of different shapes. Indeed, the fabrication techniques used to bring such panels into life generally involve the production of molds. Molds for custom panels are typically produced by computerized numerical control (CNC) machining; one mold per panel type that – in the worst case – may only be used for a single panel. Efficient solutions for mold production exist only for special cases: For instance, fiber-reinforced concrete panels can be built on molds from styrofoam, which – for certain shapes – can be fabricated effectively through hot wire cutting. However, general freeform shapes usually require CNC machining. While some parts of the molds are potentially used at a later stage to protect the panels during shipping, the process still remains not environmentally friendly.

One strategy to minimize the quantity of molds required involves employing a single machined mold to form multiple panels. This is generally achieved by clustering and shape optimization techniques (see Pellis et al. [2020] and associated literature).

An alternative solution involves creating a mold with a flexible shape. Flexible molds offer significant advantages over static ones. First, they provide superior flexibility by easily adapting to produce a wider range of panel shapes and sizes, particularly beneficial for unique or complex geometries. Secondly, this reduces manufacturing time and resource consumption by eliminating the need for production of multiple static molds for each panel variation. State-of-the-art flexible molds employ either a dense grid of actuated pins (e.g., see *Dynapixel*¹), or a slightly coarser pin grid covered by elastic material (e.g., see *Adapa*²). The latter typically yields a smoother surface finish. However, pin-based molds require a high number of supporting elements and entail complex mechanical components, rendering fabrication resolution dependent on the number of pins.

In this study, we propose a novel method for designing a flexible mold based on a simpler principle. Rather than mechanically supporting an entire surface with numerous elements, we demonstrate that supporting only a carefully chosen boundary curve suffices. While adjusting a full surface support in existing machines necessitates straightforward surface evaluation, our mold relies on a computational fabrication technique, requiring an efficient algorithm to solve a challenging inverse design problem.

In particular, our specific technical contributions include: (1) the mechanical design, and (2-3) the algorithms developed for it. These contributions are as follows: (1) We devise a novel machine for efficient mass customization of panels which is based on a deflatable membrane functioning as the mold. By the adjustment of the deflation level and the movement of boundary elements restricting

the membrane, the mold can be fully (re)configured. (2) We present an efficient optimization scheme based on a thin shell model which computes the deflation level and the boundary positions in order to produce panels of a given input shape. (3) We utilize special piecewise quadratic Weingarten surfaces to approximate the input shape’s geometry in order to obtain a proper initialization for the aforementioned optimization scheme.

2 RELATED WORK

The development of more effective and efficient approaches to paneling in computational architecture and in computational fabrication have been two prominent ongoing research topics within the academic community. While this spans a breadth of work that we cannot conclusively discuss here, we provide references to recent developments in computational architecture focusing on the use of pneumatic systems in the construction industry, on the paneling problem in architectural geometry, and related work in geometry processing and differential geometry. As we are utilizing a thin shell simulator for our optimization process, we also provide pointers to references of related work on simulating these systems.

2.1 Computational Architecture

The use of pneumatic systems in architecture and structural design comes in a variety of forms, but typically on a much larger scale and not for high-precision fabrication of panels. This is primarily due to the inherent challenges in achieving the precise control required for small-scale shapes, as opposed to the relative ease of inflating large-scale structures where fine details are less critical. Notable examples include pneumatic formwork systems for the production of concrete shells (for a comprehensive survey, see Kromoser and Huber [2016]), pneumatic forming of hardened concrete [Kromoser and Kollegger 2015] and the erection of gridshells with pneumatic formwork [Quinn and Gengnagel 2014]. A thorough overview of topics in this area can be found in Quinn and Gengnagel et al. [2008]. A significant challenge in this domain is the simulation-based design, particularly for high-precision applications. In this context, we highlight a recent, notable inverse design solution for surface-based inflatables [Panetta et al. 2021].

The fundamental motivation for a pneumatic system was to create a cost-effective and environmentally friendly solution for the fabrication of architectural panels. This is in line with the broader trend in computational fabrication research towards sustainable design [Suzuki et al. 2023]. A typical approach to cost reduction in freeform architecture is rationalization, i.e., the approximation of the original design by a simpler one which can be built at lower cost [Fischer 2012]. A large portion of research in architectural geometry is essentially dealing with the rationalization problem [Pottmann et al. 2015]. Eigensatz et al. [2010] presented a general solution to the paneling problem that has been successfully used in practice. It reduces the number of panels that require a custom mold to be built and also includes the reuse of the same mold for different panels. However, this approach necessitates a controlled deviation from the original geometry, often resulting in a less smooth surface due to visible kinks between adjacent panels. In contrast, our objective in this paper is to fabricate the exact panel shape using the “same”

¹<https://www.compositesworld.com/articles/dynapixel-automated-reconfigurable-molds>

²<https://adapamoulds.com/>




Fig. 2. High-level overview (flowchart from left to right) of the different steps within the presented approach: A reference panel is taken from a design surface (left) and approximated by a quadratic Weingarten surface (second from left) in order to estimate the support curves; this estimation is used as the initialization of an inverse simulation loop which optimizes the support curves (middle); the mold is then reconfigured according to these support curves (second from right), and the membrane is finally deflated functioning as the mold (right).

mold. This task requires a inherently more physics-aware design than the rationalization problem, which is often driven only by geometric considerations.

Paneling techniques also depend on the materials to be used: If a material such as sheet metal is isometrically bendable after being formed over a mold, only very few molds (spherical or pseudospherical) are necessary [Jiang et al. 2021]. One may even apply bending to brittle material like glass [Datsiou 2017] supported by an adapted fabrication-aware computational design system [Gavriil et al. 2020]. To achieve a smaller deviation from the original geometry while achieving easier reconfiguration, we aim to use a stretchable membrane that can easily be reused.

From a methodological perspective, the most closely related work on paneling has been published by Pellis et al. [2021; 2020] who rationalize the design surface by a so-called Weingarten surface. These surfaces are characterized by a relation between their principal curvatures and had previously received interest in differential geometry. Rogers and Schief [2003] detected a remarkable relation to integrable systems in those Weingarten surfaces that are characterized by a linear relation between Gaussian and mean curvature; they deduce them as models for equilibrium shapes of membranes under normal loading. However, the underlying assumption on the curvature-stress relations is not ideal for our application, as we are working with a real-world membrane with nonlinear stress-strain behavior. We use a piecewise quadratic relation between principal curvatures for the geometric initialization of the inverse design problem, whose computation can be based on an appropriate adaptation of the method that has been presented by Pellis et al. [2021]. A simple consequence of using this model is the correct orientation and avoidance of side change of the mean curvature normals of the membrane surface, which has previously been used [Konaković-Luković et al. 2018].

2.2 Simulating Thin Shells

The simulation of thin shells has been addressed within different scientific communities ranging from mathematical modeling and computational physics to computer engineering. We refer to the excellent textbook on elasticity and geometry by Audoly and Pomeau [2010] for a comprehensive overview. Moreover, Weischedel [2012] has provided a discrete geometric view on shear-deformable shell models. In this project, we adopt the Koiter’s model in the




Fig. 3. A coarse segmentation of the top of the *Lilium tower* (design by Zaha Hadid Architects) into panels yields a surprisingly good fit by panels that can be produced with our mold. The red panels on the left are replaced by panels very close to the quadratic Weingarten surface estimated ones on the right. The coherent reflection pattern visualizes the preservation of curvature properties and smoothness of the panels estimated with Weingarten surfaces.

form derived by Weischedel [2012] and Van Rees [2017]. Using such a thin shell model over the “classical” finite element (FEM) approach [Clough and Johnson 1968] comes with several advantages as the thin shell model is computationally more efficient as it uses fewer elements to represent the membrane compared to the FEM and appropriately takes into account the membrane’s bending behavior and thickness. This allows for faster analysis and simulation of the membrane’s behavior under different loading conditions and shows its capabilities, particularly when predicting the membrane’s response to external loads and deformations. Furthermore, a thin shell model can provide a simplified and more intuitive representation of the membrane, which can facilitate the understanding and interpretation of the results.

The visual computing community has substantially contributed to the development of sophisticated methods for the accurate and efficient simulation of thin shells (e.g., see Grinspun et al. [2006; 2003] and references therein) and several problems which occur in this context such as handling collisions and interpenetrations (e.g., see English et al. [2012]). Related problems such as simulating the dynamics of cloth [Bridson et al. 2002; Wang et al. 2018], the folding and crumpling of sheets [Narain et al. 2013], or the behavior of fiber meshes [Michels et al. 2015] are active fields of research. As the numerical handling of the underlying differential equations can

be quite challenging, a diverse portfolio of integration methods for elastodynamic systems has been devised [Chen et al. 2020; Michels et al. 2017, 2014].

Finally, we would like to mention that in recent years, differentiable simulations for cloth and thin shells have been devised and studied in the literature [Jourdan et al. 2022; Li et al. 2022; Liang et al. 2019]. Notably, the work by Zhang et al. [2022] presents a spatial-temporal progressive simulation framework for thin shells quasi-statics, which is particularly relevant for our implementation of a quasi-static formulation.

3 OVERVIEW

Panel manufacturing usually involves the casting of material on a base, commonly known as a mold, to form the desired panel shape. We aim to develop an inverse design approach contributing to the solution of the panel design and the mold reuse problems by devising a flexible mold. The pipeline overview of our approach is shown in Figure 2.

3.1 Flexible Mold

Traditional panel manufacturing involves the use of static molds. Therefore, the first step to designing a *flexible mold* is to design a resilient yet flexible surface that can withstand the weight of a panel and can be reconfigured into different shapes to meet the needs of production. Naturally, we then would like to find a flexible material as our surface and design a mechanism to support such a surface with a strong force.

In our design, we use a pressure force of constant pressure p to deflate a deformable membrane $S \subset \mathbb{R}^3$ with flexible *support frames* placed underneath to implement the *flexible mold*. With the Kirchhoff-Love assumption, we could represent the membrane S as $\bar{\Omega} \times [-h/2, h/2] \subset \mathbb{R}^3$, where $\bar{\Omega}$ is the rest/undeformed midsurface of the membrane and h is the thickness of the membrane. With a deformation map function $\varphi : \bar{\Omega} \rightarrow \Omega$ mapping the undeformed midsurface $\bar{\Omega}$ into the deformed Ω , such membrane S could be repurposed into different shapes, where its deformed midsurface Ω shape could be reconfigured to match the *target panel* $\mathcal{T} \subset \mathbb{R}^3$. In this paper, when we refer to *membrane*, we are referring to the deformed midsurface Ω of S .

Our machine prototype has an elastic membrane mounted on an air-tight metal box case, where the negative pressure p is applied by creating a low-pressure chamber inside the metal case (see Figure 4). The *support frames* consist of 4-piece support frame with bullnose edges, and are placed under the membrane in the metal case, where the entirety of the support frames stays below the rest state of the membrane. Figure 4 shows the cross-section of our machine during the deflation process, where the support frames placed under the membrane Ω helped it to form the ideal shape to match target panel \mathcal{T} . During fabrication, the elastic membrane deforms to match the shape of the target panel \mathcal{T} . After each panel is produced, the membrane returns to its original shape by equalizing the pressure with ambient levels ($p \rightarrow 0$), which gives the system its name, “flexible mold”. The flexibility comes not only from the elastic nature of the membrane but also from the replaceable support frames, which

allow for different panel shapes, supporting “flexible” production needs.

3.2 Shape Optimization for Fabrication

From the fabrication pipeline of the *Flexible Mold*, we can see that the shape of the deformable membrane Ω is mainly determined by the placement of the *support frame* and the constant pressure p . Since *support frames* only make contact with the membrane Ω at the top bullnose edges, we can model the support frames with a set of 4-piece pipe surfaces of radius r , where the centerlines of the 4-piece set are denoted by $\gamma = \{\gamma_1, \gamma_2, \gamma_3, \gamma_4\}$ (see Figure 4). Intuitively, given a constant pressure p , the placement of pipe surface centerlines γ determines the static equilibrium shape of the membrane Ω and therefore determines an integrated L^2 distance between the target panel \mathcal{T} and the membrane Ω . Following the formulation in Iterative Closest Points (ICP) for rigid registration [Pottmann et al. 2006], we adapt the *linearized* integrated L^2 distance as our optimization objective function \mathcal{D} :

$$\mathcal{D}(\mathcal{T}, \Omega) = \int_{y \in \mathcal{T}} ((y - \text{proj}_{\Omega}(y))n)^2 dA, \quad (1)$$

where $\text{proj}_{\Omega}(y)$ is the closet-point projection operator of y onto Ω , and n is the unit normal vector of the tangent plane at $\text{proj}_{\Omega}(y) \in \Omega$. While metrics based on geometric measurements such as the Chamfer and Hausdorff distances could be a good indicator of set distance evaluation, they are not ideal for optimization purposes due to their non-differentiability. The definition of distance (Eq. (1)) is often referred to as the “point-to-plane” distance, a locally smoothed derivation from the L^∞ distance (Hausdorff distance):

$$\mathcal{H}(\mathcal{T}, \Omega) = \max \left\{ \sup_{y \in \mathcal{T}} \inf_{x \in \Omega} \|y - x\|, \sup_{x \in \Omega} \inf_{y \in \mathcal{T}} \|y - x\| \right\}. \quad (2)$$

While this distance definition is purely geometric, the fabrication process of reaching a static equilibrium state for the membrane Ω is a physical process, giving rise to the need for a physics-aware optimization algorithm to optimize for the placement of pipe surface centerlines γ .

We first model the target panel \mathcal{T} to be a *rigid body* with *rigid pose* (\mathcal{R}, t) , where $\mathcal{R} \in \text{SO}(3) \subset \mathbb{R}^{3 \times 3}$ is a rotation matrix and $t \in \mathbb{R}^3$ is the translation vector. We solve this problem with a two-fold optimization scheme: A geometric initialization algorithm to find the initial placement of pipe surface centerlines $\bar{\gamma}$ and the initial rigid pose $(\bar{\mathcal{R}}, \bar{t})$ of the target panel \mathcal{T} (Section 4.2); a physics-aware optimization algorithm to validate $\bar{\gamma}$ and improve the design of pipe surface centerlines γ to optimize the shape of Ω to better align with the target panel \mathcal{T} of rigid pose (\mathcal{R}, t) (see Section 4.3).

4 METHODOLOGY

In this section, we describe the mechanical design and the fabrication process of our machine (see Figure 5) for the efficient mass customization of panels, the design of the geometric initialization process, and the simulation-based optimization.




Fig. 4. Our machine’s cross-sectional view shows the arrangement of the support under the membrane in the area that acts as the mold. The outer aluminum case ensures air tightness, while the deflatable silicone membrane serves as the mold surface. The smooth contact between the slightly rough bullnose edges of the wooden frame and the membrane results in a very high friction contact with improved strain distribution.

4.1 Fabrication Process with Flexible Mold

The machine is composed of a membrane and an aluminum alloy box case of the same size, where the membrane is mounted on top of the box before every fabrication session. From the top view, we could define the machine to be of the same size as undeformed membrane midsurface $\bar{\Omega}$ in width and length. We then fabricate lightweight replaceable *support frames* made of medium density fiber (MDF) boards for fast prototyping. For every new panel, we pre-fabricate the 4-pieces support frame set accordingly referred to as the *customizable support frames*.

The fabrication process involves deflating the membrane by applying a negative constant pressure p to the inside of the machine. We deflate the membrane until it reaches an *equilibrium* state, in which the forces in this simple system can be modeled as a *static equilibrium equation*:

$$\underbrace{f(x)}_{\text{total}} = \underbrace{f_e(x)}_{\text{internal}} - \underbrace{(f_p(x) + mg)}_{\text{external}} = 0, \quad \text{s.t. } d_\gamma(x) \geq R = r + h/2, \quad (3)$$

where internal force being the membrane elastic force f_e , external forces include pressure force f_p , gravitational force mg , and the non-penetration constraint d_γ is defined as the distance between the membrane Ω and pipe surfaces centerlines γ must be greater than a fillet pipe surface radius R with membrane radius r and thickness h (see Figure 7).

4.2 Geometric Initialization with Weingarten Surfaces

Inspired by previous work on shape approximation in the field of architectural geometry [Pellis et al. 2021], we use Weingarten surface with simplified physical assumptions to approximate an extended surface containing the target panel, then extract a 4-piece curve set $\bar{\gamma}$ from the Weingarten surface as the initial placement of pipe surface centerlines. During this step, $(\bar{\mathcal{R}}, \bar{t})$ is defined relative to the placement of $\bar{\gamma}$.

To initialize the shape of the support curves, we use a purely geometric approach based on a simplified mechanical model of the elastic membrane of our machine based on Eq. (3).

As first, we consider the sheet mechanically as a shell membrane resisting tensile stresses only, loaded by a constant pressure p in the direction of the surface normal \hat{n} . Let us consider a principal curvature parametrization of the deformed membrane surface, with unit basis $\{e_1, e_2\}$ and with unit normal $\hat{n} = e_1 \times e_2$. Let κ_1 and κ_2 be the principal curvatures along the directions e_1, e_2 respectively, and let N_1, N_2 be the corresponding membrane stresses, with positive values indicating tension. The equilibrium in the normal direction \hat{n} requires

$$N_1 \kappa_1 + N_2 \kappa_2 + p = 0. \quad (4)$$

We apply the following ansatz on the membrane stresses:

$$N_1 = \lambda |\kappa_1| + \mu, \quad N_2 = \lambda |\kappa_2| + \mu \quad (5)$$

with $\lambda > 0, \mu \geq 0$ to ensure tension. This assumption is based on the specific constraints of the membrane in our machine: A higher curvature of the membrane entails a higher stretching along that direction, and consequently a higher tensile stress. In particular, in Equation 5 we assumed a linear relation between the absolute curvature and the tensile stress. We allowed also a homogeneous pre-stretching of the membrane through the coefficient μ . Actual values of curvatures κ_i and corresponding normal stresses N_i found in simulations are shown in Figure 6.

By plugging Eq. (5) into Eq. (4), we get the following condition for the principal curvatures of the surface:

$$\lambda(\kappa_1 |\kappa_1| + \kappa_2 |\kappa_2|) + \mu(\kappa_1 + \kappa_2) + p = 0 \quad \text{with } \lambda > 0, \mu \geq 0. \quad (6)$$

A surface where the principal curvatures are in functional relation belongs to the class of Weingarten surfaces. The connection between the equilibrium of shell membranes under normal loading and Weingarten surfaces has been already highlighted in the work of Rogers and Schief [2003]. In particular, they show that a shell membrane where principal stresses follow the linear relation $N_1 = \lambda \kappa_2 + \mu$ and $N_2 = \lambda \kappa_1 + \mu$ is in equilibrium under constant pressure p if it




Fig. 5. Step by step overview of the panel fabrication process with our machine. From top to bottom, left to right: (1) Alignment of the laser distance sensor to a referenced point of interest for later pressure control. (2) Placement of the support frame inside the machine, centered relative to its base. (3) Positioning of the silicone membrane and its metal frame on the machine walls. (4) Pressure sealing of the machine with fixtures and inflation of the membrane. (5) Deflation of the membrane until the target pressure value is reached. (6) Option 1, manual fabrication. Panel production by casting of alabaster gypsum on the membrane, with the help of a removable foam frame. (7) Option 2, robotic fabrication. Panel production by extrusion 3D printing of gypsum with a 6 axis robotic arm. (8) Closeup of the robotic 3D printing process: Materialization of the designed toolpath on the deflated silicone membrane.

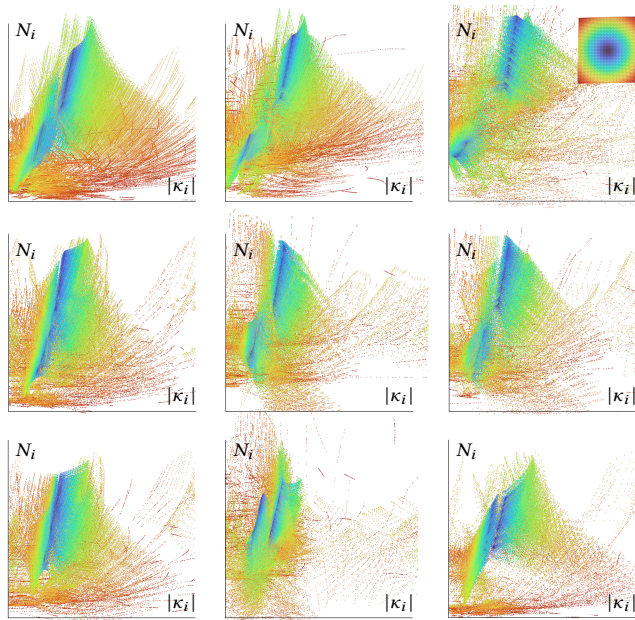


Fig. 6. Relation between principal curvatures κ_i and corresponding normal stresses N_i with respect to the 9 panels of the facade in Figure 1 and numbering the array. Each diagram refers to 50 forward simulations of the membrane for given centerlines γ of one of facade panel, subject to different pressures ranging linearly from -0.01 MPa to -0.02 MPa. For each mesh face enclosed by boundary curves, points $(|\kappa_1|, N_1)$ and $(|\kappa_2|, N_2)$ are represented. The membrane stress tensor, used to compute normal stresses, is derived from simulation data. The color gradient indicates the distance of the face centroids from the centroid of the centerlines γ . Points closer to the curves centroid tend to cluster into two straight segments, corresponding to the two principal curvatures. We observe that the linear relation of Eq. (5) tends to hold better for those points.




Fig. 7. Derivation of pipe surfaces centerlines. Each pipe surface centerline $\bar{\gamma}_{i=1,4}$ is given by the directrix of a fillet pipe surface of radius R between the Weingarten surface and a section plane. The fillet pipe surface is given by the envelope of spheres with constant radius R , tangent to both the Weingarten surface and the section plane.





Fig. 8. Different initialization qualities between a linear Weingarten surface as in Rogers and Schief [2003] (middle) and a quadratic Weingarten surface (right). The target panel is shown on the left. The quadratic Weingarten surface is colored with the same error range as the linear Weingarten surface, which is defined between 0 and the maximum point-to-mesh distance between the linear Weingarten surface and the simulated thin shell.

takes the shape of a Weingarten surface with curvature relation $2\lambda K + \mu H + p = 0$, where H and K are respectively the mean and the Gaussian curvature given by $H = 1/2(\kappa_1 + \kappa_2)$ and $K = \kappa_1\kappa_2$. They also show that, for this kind of surface, if principal stresses are aligned with principal curvature directions e_1, e_2 , the equilibrium in the tangent plane of the membrane is automatically satisfied since it delivers the Gauss-Mainardi-Codazzi equations of the surface. Due to the specific geometric and mechanical characteristics of our pressurized membrane, we found it necessary to formulate the alternative geometric hypothesis on membrane stresses of Equation 5.



However, note that our stress ansatz does not entail tangent equilibrium as well. Nevertheless, since our shell membrane model is intended as an initialization for a more accurate thin shell simulation, it turned out to be sufficient to consider equilibrium in the normal direction alone (see also Figure 8). If we consider a negative pressure p (i.e., a force of magnitude p acting in the opposite direction of the surface normal \hat{n}), Equation 6

gives a curvature relation represented in the plane κ_1 - κ_2 by branches of hyperbola in the quadrants II and IV, with asymptotes given by their bisecting line, and a portion of a circle in the quadrant I, with center in its bisecting line. This curve lies entirely above the asymptotic line $\kappa_1 + \kappa_2 = 0$. This property highlights the fact that with this membrane model, if pressure p is negative, equilibrium can be reached only for surfaces with a positive mean curvature H . The opposite is true if pressure p is positive. In our setup, a membrane pre-stretching can be obtained by positioning the centerlines γ above the boundary of the membrane by a specific height h_0 . The process then entails inflating the membrane above the height h_0 and gradually deflating it. As first approximation, the pre-stretching can be derived from the extension achieved by the membrane upon contact with the centerlines γ during deflation. In this configuration, for given coefficients λ and μ , we can assume $\mu \approx \lambda\kappa_1 \approx \lambda\kappa_2$. Assuming a spherical inflated shape for the membrane, we can set $\kappa_1 = \kappa_2 = 1/\varrho$, where ϱ denotes the sphere radius. Consequently, we can infer that $\varrho = \lambda/\mu$. The height h_0 can be readily determined from the sphere radius ϱ , knowing the span of the spherical membrane, which we can consider as the smaller membrane dimension d , as follows:

$$h_0 = \lambda/\mu - \sqrt{(\lambda/\mu)^2 - (d/2)^2}. \quad (7)$$

As shown in Figure 6, we can make the assumption that a Weingarten surface fulfilling Eq. (6) is a good approximation of a region on the membrane surface Ω matching the target panel \mathcal{T} . The Weingarten surface is then cut with four planes at a prescribed distance from the panel boundaries, as shown in Figure 9. From the four intersection curves, we derive the initialized pipe surfaces centerlines $\bar{\gamma}$, taking into account the thickness h of the membrane and the radius r of the bullnose edges of the support frames (see Figure 7). Utilizing the coefficients λ and μ from the Weingarten


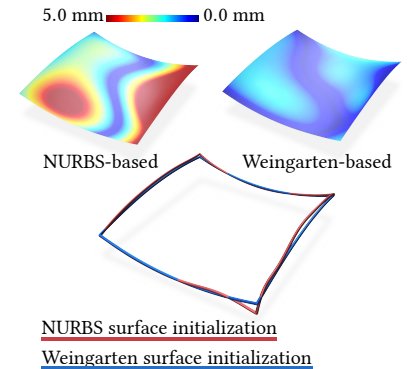


Fig. 9. Initialization pipeline. (a) A target panel is given as a B-spline surface (blue). The surface is extended beyond the panel boundaries. (b) The surface is optimized towards a Weingarten surface with curvature relation given by Equation 6. The color scale shows the distance from the target panel. The initial guess of the pipe surface centerlines γ is obtained by intersecting the optimized surface with four sides of a box (red curves). (c) The curvature diagrams before (yellow) and after (magenta) optimization are shown.

surface, we can determine a suitable height h_0 to position the centroid of the curves $\bar{\gamma}$ using Eq. (7). Then, these curves $\bar{\gamma}$ undergo rotation around its centroid to minimize the height of their bounding box. Lastly, we derive the initial rigid pose $(\bar{\mathcal{R}}, \bar{t})$ of the target panel \mathcal{T} by aligning the panel to the pipe surface centerlines $\bar{\gamma}$.

One might argue that a naive initialization could be sufficient for surface extension. While it is true that simple methods, such as fitting a NURBS surface and applying an outward offset, may work for panels with simple, positively curved geometries, these methods quickly produce visually unappealing and physically implausible



results when applied to negatively curved panels.

To illustrate this, we compare the surface extension results using a NURBS-based approach with those obtained using Weingarten surface extension. The NURBS extension tends to create distortions and inaccuracies in areas of negative curvature, leading to significant deviations from the desired geometric properties. On the other hand, the Weingarten surface extension provides a more accurate and visually coherent result, especially for complex geometries. Furthermore, only with the Weingarten surface approach can we achieve a reasonable estimation of the prestress distribution and the initial height h_0 . This method takes into account the underlying curvature of the surface, allowing for a more precise optimization process that minimizes error and better captures the physical behavior of the material.

4.3 Inverse Design with Physical Simulation

Our geometric initialization gives an initial estimation of the initial geometry of $\bar{\gamma}$, the estimated height h_0 of boundary-panel pair $(\bar{\gamma}, \mathcal{T})$, and initial rigid pose of the panel $(\bar{\mathcal{R}}, \bar{t})$, the geometric-only approach does not account for physical parameters such as Young's modulus E and Poisson's ratio ν , resulting in differences between

fabricated panels and desired panels. This raises the need for further physical optimization to reduce the mean and maximum error on the panel – by incorporating a simulation-based optimization, we can utilize the accurate geometry-aware prior while also achieving a physically feasible design.

It might seem intuitive to formulate the problem as a pipe surfaces centerlines optimization problem, where we minimize the distance between the panel \mathcal{T} and the surface Ω by tweaking pipe surfaces centerlines γ design.

However, we would like to reparametrize the problem from a pipe surface centerlines γ optimization problem to a membrane Ω shape optimization problem. There are several practical reasons behind this as described in the following.

No-Slip Contact. We treated the pipe surface centerlines used in fabrication to have an extremely high friction coefficient, putting our physical setup outside the typical working range of frictional contact simulation in modern computer graphics literature (e.g., IPC [Li et al. 2020] and C-IPC [Li et al. 2021]). Note that this modeling is based on observations in physical experiments, not simplified assumptions, as this model matches the no tangential sliding and no penetration behavior of the mold in the fabrication process.

Panel Alignment Complexity. Despite the Ω can be interpreted as a function of γ , we are only interested in how panel \mathcal{T} is aligned to Ω . Using barrier formulation adds an extra layer of complexity due to the choice of thickness of the barrier (a common conservative option would be around 1mm), as the distance energy varies with millimeter level offset.

Penalty/Barrier Force Issues. While modeling differentiable contact through a penalty/barrier force is common, in our setup, accumulating gradients from the panel-membrane proximity energy to boundary curves proved numerically unstable. This instability persisted even with an implicit simulator unless using impractically small time steps.

Thus, instead of optimizing γ , we aim to minimize the distance (Eq. (1)) between the membrane Ω and panel \mathcal{T} subject to $x \in \Omega$ and panel rigid pose (\mathcal{R}, t) where $\mathcal{R} \in \text{SO}(3) \subset \mathbb{R}^{3 \times 3}$ is a rotation matrix and $t \in \mathbb{R}^3$ is the translation vector. Consequently, this motivates us to optimize the rigid pose of \mathcal{T} to align with Ω and change the shape of Ω to match \mathcal{T} simultaneously. Since Ω is a deformable surface governed by elastic energy, our shape optimization problem deviates from the traditional rigid alignment (registration) problem formulation. Here, let us revisit Eq. (1) from the perspective of a rigid panel \mathcal{T} , and a deformable surface Ω , essentially splitting into two sub-problems in ARAP style [Sorkine and Alexa 2007].

Panel Rigid Alignment. Assuming deformable membrane surface Ω is in stasis, the rigid alignment problem is formulated as an optimization of panel rigid pose (\mathcal{R}, t) to minimize the proximity objective \mathcal{D} (Eq. (1)):

$$\begin{aligned} \min_{\mathcal{R} \in \text{SO}(3), t \in \mathbb{R}^3} \int_{y \in \mathcal{T}} \|\mathcal{R}y + t - \text{proj}_{\Omega}(y)\|^2 dA, \\ \text{s.t. } \theta(\mathcal{R})_y = 0, t_{x,z} = 0, \end{aligned} \quad (8)$$

where θ is a $\text{SO}(3) \rightarrow \mathbb{R}^3$ function that maps a rotation matrix \mathcal{R} to a Eulerian angles vector in \mathbb{R}^3 . In this paper, we assume the coordinate system to be y -up. Therefore, the constraints imply we should not apply the change on the y -axis rotation and x, z -axis translation of the rigid pose – limiting the rigid pose of \mathcal{T} to 3 degrees of freedom to ensure it will never transform out of the bounding box \mathcal{B} defined by pipe surfaces centerlines γ .

Membrane Surface Shape Optimization. Considering the rigid pose of \mathcal{T} is fixed, we can optimize the shape of Ω to minimize the distance in Eq. (1) subject to elastic energy. Similar to Eq. (8), we consider the left inverse of $\text{proj}_{\Omega}(\cdot)$, $\text{proj}_{\Omega}^{-1}(\cdot)$, as a function that maps a point $y \in \text{proj}_{\Omega}(y)$ to a point $x \in \mathcal{T}$. Therefore, we can formulate the shape optimization problem as:

$$\min_{x \in \Omega} \int_{x \in \text{proj}_{\Omega}(\mathcal{T})} \|\mathcal{R} \text{proj}_{\Omega}^{-1}(x) + t - x\|^2 dA, \quad \text{s.t. } f = 0. \quad (9)$$

Note that f is the system force applied on Ω , so $f = 0$ is a nodal force equilibrium state (Eq. (3)).

5 IMPLEMENTATION

In this section, we provide details on the implementation of our flexible mold system. We first describe the physical design of the flexible mold in Section 5.1. We then detail the geometric initialization in Section 5.2, the physical simulation in Section 5.3, and the local-global optimization in Section 5.4. Finally, we discuss the finalizing design in Section 5.4.3.

5.1 Mechanical Design of the Flexible Mold

In the physical proof of concept implementation of our flexible mold (Section 3.1 and Section 4.1), we used a 3mm thick silicone membrane with a Shore A hardness of 40 and a density of 1190 kg/m³. The membrane is clamped in a two-part machined aluminum frame with inner dimensions of 660 mm × 880 mm × 3 mm (*width × depth × thickness*). This frame is placed on an almost airtight box constructed from aluminum profiles. Dowel pins ensure the frame sits at the correct location. Inside this box, we place the CNC-machined MDF board support frames. Once the air pressure inside the box is lowered, the membrane deflates and contacts the MDF frame. The MDF board's slightly rough surface at the rounded edge ensures the membrane does not slip across the edge during deflation. The air pressure inside the box is lowered below ambient pressure by two pneumatic Venturi vacuum generators. One of the generators is always on (generating a base airflow), and a fast-acting magnetic valve on the high-pressure side controls the second one. We realize a closed-loop controller by measuring the deflation using a laser triangulation distance sensor (Micro-Epsilon ILD1420-200). Note that although we monitor internal pressure difference p within the machine interior, it is not used as a control variable. Instead, we regularize deflation based on the laser sensor's distance measurement rather than pressure, mitigating any error caused by material parameter measurements. For example, running the optimization with a 10% variation in Young's modulus results in a nearly identical shape but significantly different pressures. Controlling the deflation by measuring the membrane shape implicitly eliminates this source of error.

The support frames are formed by CNC-machined MDF boards. The important top edge is shaped using a 5 mm radius sunk bead cutter. The panels in our experiments are made of both manually applied and robotically 3D printed alabaster gypsum. At an architectural scale, materials like concrete or fiber-reinforced laminate would be used.

For verification, we used a laser scanner to digitize the shape of the deflated membrane and compare it to simulation results and the panel geometry. We use a Metris MCA with an MMD200 laser scan head. The set-up features a volumetric accuracy of about $100\mu\text{m}$.

Membrane Mechanical Properties Measurement. In theory, our setup is designed to work with any incompressible membrane material with a Young's modulus E and Poisson's ratio ν that can be modelled as a thin shell. However, the physical properties of the membrane material will affect the deflation process and the final shape of the panel. We collect physical properties including thickness 3 mm and density 1190 kg/m^3 from the specification sheet the vendor provided. We calibrate the mechanical properties of the silicone membrane used in our experiments under the tensile testing machine, conclude the Young's modulus $E = 1.2\text{ MPa}$, and Poisson's ratio $\nu = 0.5$. We cross-validated these estimations on actual machine membrane scans overlaid with scans of the deflated membrane on the machine using the fitted Young's modulus.

5.2 Geometric Initialization

Motivated by Section 4.2, we use the method of [Pellis et al. 2021] to approximate a panel with a Weingarten surface. This approach allows us to optimize a given B-spline surface towards a Weingarten surface by enforcing a specified relation between the mean curvature H and the Gaussian curvature K at sample points. The optimization is based on the Levenberg-Marquardt algorithm, where constraints and target functions are expressed by a system of quadratic polynomial equations solved iteratively. Starting from the implementation of [Pellis et al. 2021], we add the variables λ , μ , p , and the auxiliary variables $\bar{\kappa}_1 = \kappa_1|\kappa_1|$ and $\bar{\kappa}_2 = \kappa_2|\kappa_2|$ at each sample point, with corresponding auxiliary constraints

$$2H^2 - K + 2H\sqrt{H_*^2 - K_*} - \text{sign}(\bar{\kappa}_1^*)\bar{\kappa}_1 = 0,$$

$$2H^2 - K - 2H\sqrt{H_*^2 - K_*} - \text{sign}(\bar{\kappa}_2^*)\bar{\kappa}_2 = 0,$$

where the subscript $*$ denotes a variable taken from the previous iteration, and where we used the relation $\kappa_{1,2} = H \pm \sqrt{H^2 - K}$. The target function of (6) is then enforced as $\lambda(\bar{\kappa}_1 + \bar{\kappa}_2) + 2\mu H + p = 0$ at each sample point. As a convention, we orient the surface normal towards the open side of the membrane; the vacuum will produce than a constant negative pressure p . To ensure positive coefficients λ , μ , and a negative p , we add constraints $c - |c_*| = 0$, with $c = \lambda, \mu, -p$.

For the optimization, we used 4-th degree B-spline surfaces with eight control points in each parameter direction. The algorithm has been implemented in Python and tested on a Intel® Core™ i7-10750H CPU with 32 GB RAM. Computation times for the examples presented in this paper ranged from 4.47 s to 26.03 s.

5.3 Forward Physical Simulation

For optimization, we would like to simulate until the membrane reaches equilibrium (i.e. $f = 0$) where the internal force and external force are balanced in Eq. (3). In this section, we provide details on our implementation of the forward quasi-static simulation (Section 4.3) of the deflation process to reach equilibrium. Our simulation model is tightly connected to our machine design, and all modeling decisions were *motivated by observations from the physical experiments*.

In this paper, we build our model based on Koiter's model [Koiter 1966] in the form derived by Weischedel [2012]. We discretize our membrane Ω into a triangular mesh $\{V, E, F\}$ where $V \in \mathbb{R}^{|V| \times 3}$ is the set of vertices, $E \in \mathbb{R}^{|E| \times 2}$ is the set of edges, and $F \in \mathbb{R}^{|F| \times 3}$ is the set of faces. We denote the set of vertices on the boundary of Ω as $\partial\Omega$. We implement our flexible mold simulator based on the thin shell simulator `LibShell` from [Chen et al. 2018]. In all experiments, we utilize the *StVK* (St. Venant-Kirchhoff) material model and the midedge average formulation for bending energy discretization. We opted not to use a more advanced model because, although Neo-Hookean performs better under compression, our application is primarily dominated by stretching. Additionally, *StVK*'s simplicity and its sparser Hessian facilitate sensitivity analysis.

5.3.1 Simulator Setup. We set up the membrane following the physical parameter detailed in Section 5.1. While it is evident that pressure force f_p will play a crucial role in reaching the equilibrium and defining the final shape of membrane Ω , it is unclear how we should choose the pressure constant p to simulate the deflation process.

Pressure Force. We first define the discrete pressure force f_p in Eq. (3) following [Skouras et al. 2014] as:

$$f_p = \frac{p}{3} \sum_{x_i \in \Omega} \sum_{e \in \mathcal{N}(x_i)} A_e \hat{n}_e, \quad (10)$$

where p is the pressure constant, x_i is the i -th vertex in Ω , $\mathcal{N}(x_i)$ is the set of faces that share x_i , A_e is the area of face e , and \hat{n}_e is the unit normal vector of face e .

Pressure Constant. To initialize the pressure constant p in Eq. (10), we utilize two simple strategies to estimate p in a similar fashion be described later in the *Global Step* in Section 5.4. We minimize the L^2 distance between the panel \mathcal{T} and the membrane Ω in Eq. (1) to find the optimal pressure constant p^* :

$$p^* = \arg \min_p \mathcal{D}(p). \quad (11)$$

Similar to Section 5.4, we run one forward quasi-static simulation (Section 5.3.2), and then use sensitivity analysis to relate Δp to Δx . See the following sections for details.


5.3.2 Quasi-Static Simulation. We model the deflation process as a quasi-static simulation following the formulation of Teran et al. [2005] where we assume the membrane deflates slowly enough that the system is always in equilibrium, omit the inertial term in the equation of motion, and iteratively solve the following equation

until convergence. Considering one iteration of the quasi-static simulation, updating the vertex position x from x^k to x^{k+1} , we obtain

$$\begin{aligned} x^{k+1} &= x^k - \alpha \left(\nabla f(p^k)^{-1} \right) f(p^k), \\ \text{s.t. } d_\gamma(x^{k+1}) &\geq R, \quad f = 0, \end{aligned} \quad (12)$$

where ∇f is a $\mathbb{R}^{3|V| \times 3|V|}$ force Jacobian matrix, $\alpha \in \mathbb{R}$ is the descent rate, and $d_\gamma(\cdot)$ is the distance to pipe surface γ which constrained to be greater than constant radius $R = r + h/2$ in Figure 7 throughout the simulation. Note that for every equilibrium state, x^k will be subject to a pressure force with pressure p^k . We slowly increase pressure $p^{1:k}$ until it reaches the optimal pressure p^* in Eq. (11) where distance \mathcal{D} (see Eq. (1)) is minimized. We use the *Newton-Raphson* method with backtracking line search to solve the equation above, and the linear system is solved using the *Conjugate Gradient* (CG) method. We terminate the simulation when the squared norm of displacement $\|\delta x\|_2^2$ is below a threshold ϵ .

Collision Detection And Response. We experimented with penalty force based frictionless contact and found the simulation to be *viable* only with an extremely high penalty coefficient and small descent rate α , given the other physical parameters (i.e. thickness h and pressure constant p) in the system. Despite our attempt to use an implicit time integrator to stabilize the simulation, the simulation is still unstable. Furthermore, a frictionless contact model would be against our intention to let the mold withhold the weight of a panel, as any weight will make the membrane slide along γ . To take into account the physical design and the numerical challenge, instead of a barrier [Li et al. 2020] or penalty formulation, our implementation is essentially a geometry level step-and-project: When we advance from one equilibrium state x^k subject to external pressure constant p^k to the next equilibrium state x^{k+1} subject to external pressure constant p^{k+1} , we take the step, then detect the distance d_γ between the support curves segments γ and membrane vertices x , and make sure those vertices no longer slide along tangent plane to move perpendicular to tangent plane. In summary, we will enforce Dirichlet boundary conditions on vertices in contact with γ in the previous to the next state to simply model the non-penetration non-sliding contact constraint in Eq. (12). Our machine prototype is designed to have a high friction coefficient on the pipe surface, making the membrane virtually stick to the pipe surface upon contact. Therefore, we can assume there is *zero* penetration along the contact normal direction and *zero* sliding along the contact tangent plane. This formulation, combined with our curve fitting extraction, was later proven to be highly accurate for our problem: In the inline figure below, we visualize the difference between two membranes forward simulated under the same physical parameters but one with the original input γ and the other with the extracted support curve γ' . The colormap highlights that there is almost 0 deviation in vertices positions, with the vertices in contact having a numerical epsilon difference.



5.4 Local-Global Optimization

The local-global optimization, as described by Sorkine [2007] involves two key steps. In the local step, the rigid pose of the panel \mathcal{T} is optimized to minimize the proximity objective \mathcal{D} (Eq. (1)). In the global step, the shape of the membrane Ω is optimized to further minimize the same proximity objective \mathcal{D} (Eq. (1)) subject to total force f equilibrium condition.

We start with discretizing the target panel \mathcal{T} with a triangular mesh $\{\mathcal{V}, \mathcal{E}, \mathcal{F}\}$.

5.4.1 Local Step. Assume the membrane Ω stays static after reaching an equilibrium state, we formulate the local step as a \mathcal{T} rigid pose optimization problem with Iterative Closest Point (ICP) with the point-to-plane objective [Chen and Medioni 1992] discretizing Eq. (8):

$$\begin{aligned} \min_{\mathcal{R} \in \text{SO}(3), t \in \mathbb{R}^3} \quad & \sum_{i=1}^{|\mathcal{V}|} \left((\mathcal{R}y_i + t - x_{y_i}) \cdot \hat{n}_{y_i} \right)^2, \\ \text{s.t. } \quad & \theta(\mathcal{R})_y = 0, \quad t_{x,z} = 0, \end{aligned} \quad (13)$$

where x_{y_i} is the closest point on Ω and \hat{n}_{y_i} is the unit normal defined on Ω at x_{y_i} . We enforce the rotation/translation constraint with *projected gradient descent* where we project the updated rotation \mathcal{R} to the closest admissible rotation \mathcal{R}' in $\text{SO}(3)$, and set the updated translation t to zero along the x and z axis.

5.4.2 Global Step. We make the ansatz that for an infinitesimal change in the design of the pipe surfaces centerlines γ , the set of indices of the vertices in contact remain the same and only the vertex positions change. We denote the set of **design** vertices in contact after one forward simulation reaching force *equilibrium* in Eq. (3) as $x_d \subset x \in \Omega$, and the set of **free** vertices in Ω as $x_f = x \setminus x_d$. Therefore, we could model the optimal change in pipe surfaces centerlines design $\Delta\gamma \approx \Delta x_d = \frac{\partial \mathcal{D}}{\partial x_d}$ as a sensitivity analysis problem [Umetani et al. 2011]:

$$\Delta x_d = \frac{\partial \mathcal{D}}{\partial x_r} \frac{\partial x_r}{\partial x_d}. \quad (14)$$

Upon finishing one run of forward simulation (Section 5.3), we could improve design vertices x_d by updating the design vertices x_d with the optimal change Δx_d in Eq. (14) without running more forward simulation. It apparent that $\frac{\partial \mathcal{D}}{\partial x_r}$ is the gradient of the proximity objective \mathcal{D} in Eq. (1) with respect to the free vertices x_f . To derive $\frac{\partial x_r}{\partial x_d}$, we need to compute the sensitivity matrix $\mathcal{S} = \frac{\partial x_r}{\partial x_d}$ from the force Jacobian matrix ∂f in Eq. (12).

Incremental Update Equation. We could relate the change of a set of design vertices to the change of all other free vertices $x_f = x \setminus x_d$ through the *static equilibrium equation* (Eq. (3)) by the following *incremental update equation* in *Sensitivity Contour* [Umetani et al. 2011]:

$$\Delta x_d = (\nabla_{x_d} f)^{-1} \nabla_{x_f} f \Delta x_f = \mathcal{S} \Delta x_f, \quad (15)$$

where \mathcal{S} is the $\mathbb{R}^{|x_d| \times |x_f|}$ *sensitivity matrix* mapping Δx_f to Δx_d , changes in free vertices to design handles. We solve this linear system in Eq. (15) with an LDLT direct solver.

In experiments, we found the gradient Δx_d to be highly localized at every optimization step, resulting in spiky and abrupt changes.

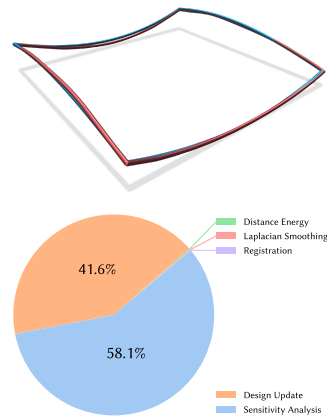
These changes Δx_d can no longer be interpreted as a smooth change of the pipe surfaces centerlines $\Delta \gamma$. To mitigate this issue, we apply a Laplacian smoothing filter [Nicolet et al. 2021] on Δx_d before updating the design vertices x_d :

$$x_d^{k+1} = x_d^k - \eta(I + \lambda_s L)^{-1} \Delta x_d, \quad (16)$$

where η is the descent rate, M is the mass matrix, L is the Laplacian matrix, and λ_s is the Laplacian smoothing hyperparameter. We set the λ_s to be the averaged edge length of a triangular mesh of Ω in all experiments.

After applying the Laplacian smoothing to update the design vertices x_d as described in Eq. (16), we relax the updated shell configuration with a quasi-static simulation step as described in Eq. (5.3). During this process, the pressure remains fixed at the optimal value p^* and the rigid pose of the target panel \mathcal{T} is held constant.

5.4.3 Finalizing Design. We iterate the local-global optimization until convergence. Following a similar procedure of extracting polylines from the Weingarten surface in Section 4.2, we extract the final pipe surfaces centerlines γ from the final membrane Ω after optimization. We then fabricate the panel \mathcal{T} with the final pipe surfaces centerlines γ using the fabrication pipeline as described in Section 3.1. The inline figure shows the extracted optimal curve for panel TP3 Section 6.1 (red) and the Weingarten surface initialized curve (blue). For results in the next sections, we observe a runtime of 10 s to 34 s per optimization step, with a high-resolution forward simulation of 10 minutes. On average, this sums up to around 30 to 40 minutes optimization time for one panel on a machine with a 32 cores AMD Ryzen 9 7950X and 64 GB RAM. A runtime breakdown of the physical optimization is provided in the pie chart illustrating the distribution across various computational steps as described in Section 5.4. The majority of the time is consumed by Sensitivity Analysis (Eq. (15)) and Design Update (relaxation in Section 5.4.2, which involves constructing and solving large linear systems that require numerous iterations). In contrast, the other steps are smaller optimizations that depend only on the geometric state and thus take less time.



6 RESULTS

In this section, we present the results of our computational fabrication process on 4 testing baseline panels and a facade composed of 9 panels, varying in curvature properties. In computational architectural design, we often categorize panels based on their curvature properties into three categories: Panels with positive Gaussian curvature ($K > 0$), panels with negative Gaussian curvature ($K < 0$), and developable panels with zero Gaussian curvature ($K = 0$).

6.1 Baselines: Range And Accuracy

Here, we choose four baseline test panels to evaluate the potential range and accuracy of panels that can be fabricated with our flexible mold. Following the real-world expectations, we select the following test panels (TP) of roughly the same size and shape, but with different curvature properties: Panel TP1 with a positive Gaussian curvature K ranging from 2.72 to 4.89; panel TP2 with a positive Gaussian curvature K greater than TP1, where K is in the range of 4.33 to 20.71; panel TP3 is saddle-shaped, exhibiting negative Gaussian curvature, with K ranging from -12.12 to -0.80 ; panel TP4 featuring a sign change in Gaussian curvature, with K ranging from -3.78 to 3.76.

We then obtain the simulation/fabrication results for these panels, shown in columns 1-4 of Figure 11. Here, we provide discussions and insights into the results of the panels.

Range of Curvature. Panel TP2 ($K \gg 0$), TP3 ($K < 0$), and TP4 ($K > 0$ and $K < 0$) exhibit unusual curvature properties that are challenging to fabricate with traditional molds. Moreover, those three panels have a higher Gaussian curvature value than most typical architectural structures, with panel TP4 notably having a sign change in Gaussian curvature. The panel TP2 has a ridiculously high Gaussian curvature value in the positive domain of up to 20.71, compared to our fabricated facade with a maximum Gaussian curvature of 3.49. Similarly, our TP3 has a negative Gaussian curvature of -12.12 , which is also lower than that of the facade with a minimum Gaussian curvature of -4.41 . Given the context,

One might consider panel TP4 as a potential failure case given the higher error in the initial simulation, shown as red on panel visualization, however, it has an explanation: The target panel exhibits a side change of the mean curvature normal. The Weingarten surface approximation removes those properties in the geometric optimization by a Weingarten surface, but subsequent simulation-based optimization towards the original panel just shifts the error to the boundary.

Accuracy. In Figure 11, we show surfaces obtained from the closest point projecting panels onto mold midsurface of the geometric initialization and simulation-based optimization for the four test panels, where the geometric initialization results is obtained with forward simulation with pressure optimization as the initial state for the physical optimization. In Figure 11, we show 5 frames sampled logarithmically from the optimization process. As more optimization steps are taken, the mold midsurface gradually converges to the target panel shape. Naturally, the error concentrated on boundary or corners, as the boundary error for all panels is higher than the mean error (MAE/MAPE in Table 1) throughout the optimization for all panels. We observe a great reduction in mean distance error and max distance error for TP2 and TP4 as the best fit Weingarten surface has a high initial error as high as 6mm, due to the initial pose of the target panel being far from the best fit Weingarten surface. Then the error rapidly decreases as the local registration step in the optimization process aligns the target panel to the simulated mold midsurface, eventually bring the mean error down to around 1mm, achieving a 83% decrease in mean distance error for both panels. Panel TP1 shows little decrease in mean error, as TP1 was designed

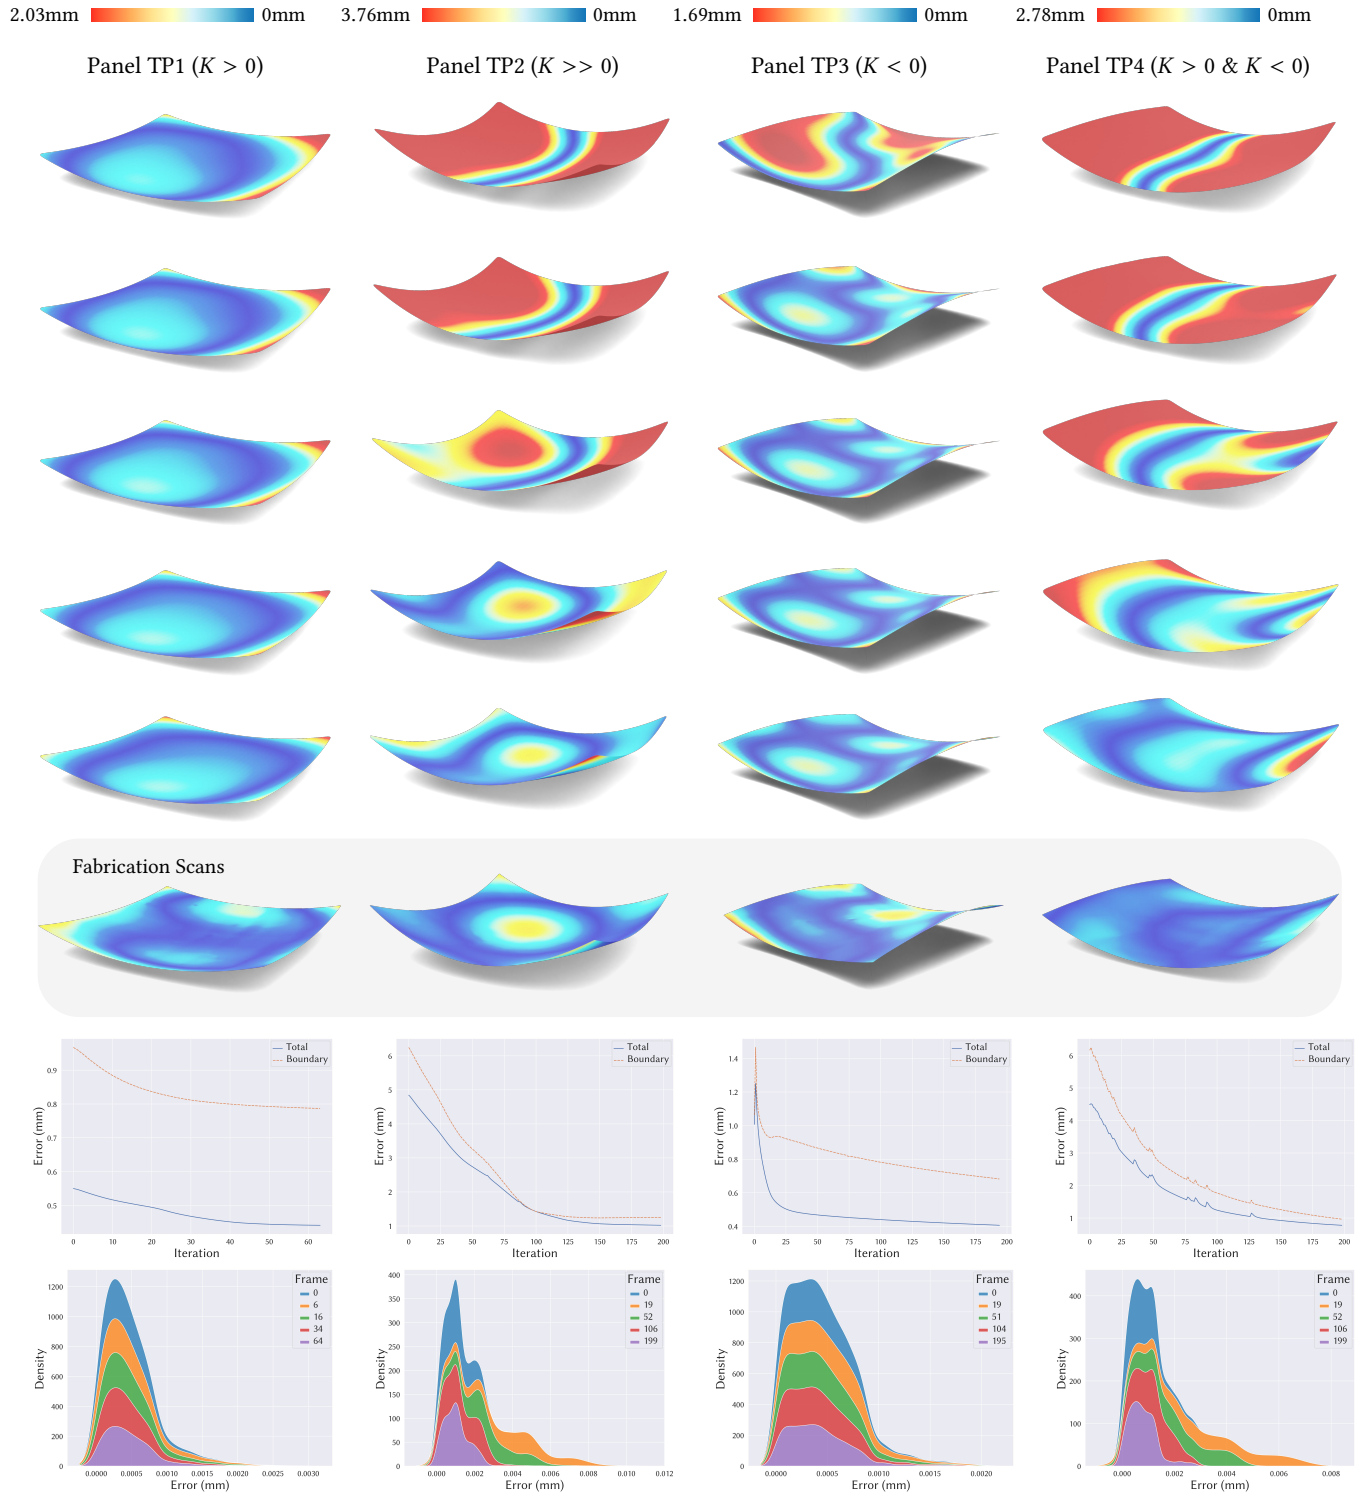


Fig. 11. Results on the four test panels TP1-TP4 (columns 1-4). Colormap encodes distance to the target panel: (1) Row 1: Forward simulation (Section 5.3) with geometric initialization (Section 4.2). (2) Rows 2-5: Simulation-based optimization results at steps $N/10$, $N/4$, $N/2$ and N (total number of optimization steps). The displayed shape is the closest point projection of the target panel onto the mold midsurface. (3) Row 6: Laser scans of the mold surface using the optimized curve γ (Section 5.4). (4) Row 7: Shows the integrated L^2 distance vs. number of optimization steps. The blue curve indicates the mean error across the panel, and the orange curve indicates the mean error at the panel boundary. (5) Row 8: Stacked KDE plots of the distance distribution to the target panel, showing that as optimization converges, both the upper bound and the variance of the L^2 distance distribution decreases.




Fig. 12. Negatively curved panel TP3 and the laser scan of the membrane with two vertical planes through areas of “large” deviation.

to align well with the geometry of a deflated membrane surface, making the initial error sufficiently low for panel fabrication. We run the optimization process on TP1 to showcase that the optimization can converge in a few steps (< 30) for panels that are already close to the mold midsurface, which serves as a good sanity check for the optimization process. The negatively curved TP3 achieves an impressive deduction in mean error from 1.05mm to 0.39mm, a 63% decrease in mean distance error. In the final state of the optimization, we observe that all panels have a relative error under 1% with respect to the target panel in Table 1.

To give a further basis for judging the quality of results, we show two planar sections through panel TP3 and the corresponding laser scan in Figure 12. The sections show a high agreement between the actual scan of the membrane surface and the target panel shape with negligible deviation on the boundary of the panel. We also provide illustrations of the membrane’s cross sections and fitted target panel shapes for all test panels in supplementary material.

6.2 Facade Design: Scalability And Robustness

In Figure 1 and Figure 14, we showcase a facade composed of panels fabricated with our flexible mold. The facade is specifically designed to be significantly more challenging than real-world examples, as our facade is engineered to include three challenging geometric properties to realize in fabrication: (1) positive and negative Gaussian curvature, (2) wide range of Gaussian curvature in both positive and negative end, (3) almost developable panels. A part of the facade is segmented into 9 panels, arranged in a 3×3 array. We give the panels numbering based on its position in the array, from top left to bottom left as P1, P2, P3, and top left to top right as P1, P4, P6, with P5 being the center panel.

Out of the 9 panels, 7 panels have mixed curvature with a distinct range of Gaussian curvature distribution, with one panel being almost developable.

Facade Surface. Compared to baseline test panels, in facade panelling, we would like to build a coherent surface that is smooth and continuous, while every panel should be aware of its neighbours’ shape. In previous work on panelling surfaces (see [Eigensatz et al. 2010]), to avoid kinks between adjacent panels, one would need to optimize adjacent panels simultaneously. In our experiment, we take the challenge of optimizing each panel individually and assemble a facade from independent panels. The results show that the panels are well aligned with each other, see Figure 1.

Almost Developable Surface. By almost developable surface, we meant some panels have Gaussian curvature $|K| < \epsilon$ that is significantly lower than the mean Gaussian curvature of the facade. The



Fig. 13. From left to right, manually fabricated gypsum panels TP1 ($K > 0$), TP2 ($K \gg 0$), TP4 ($K < 0$ and $K > 0$), and TP3 ($K < 0$).

inline figure is the Gauss image of the 9 panels in the facade, where the ordering remains consistent with Figure 1.

The area of the image of the Gauss map (Gauss image) is equivalent to the surface integral of the Gaussian curvature. The Gauss image of panel P8 is marked with blue for its interesting property, being the smallest in the Gauss image area (0.021) and is the most “curve”-like one, implying one curvature direction is significantly higher than the other.

Almost developable surfaces were not considered in the baseline test panels, as they were not a part of the initial design objective. Despite the challenging nature of those panels, our prestress from elevation scheme successfully aided us in fabricating P8. Panel P8, having three adjacent panels, was installed in the middle of the rightmost column.



7 DISCUSSION

We conclude with discussing scalability and versatility as well as limitations and future work.

Scaling. Given the size of our machine, we produced panels at the decimeter level, which is at model scale. In actual architectural design and manufacturing, one needs a larger machine to achieve the required panel sizes. Although different material properties and air pressure settings will affect the convergence and accuracy of the thin shell model, we expect to be able to produce panels with curvatures occurring in architectural applications. Due to practical constraints, we were unable to construct a large-scale machine for fabricating full-sized concrete panels. Instead, we developed precise scale models at approximately a 1 to 10 ratio. We anticipate that scaling up by a factor of 10 will not present significant challenges. Theoretically, the membrane’s thickness and material properties can be adjusted to accommodate larger panels. The membrane’s thickness and stiffness (Young’s Modulus) can be increased proportionally to the panel size to maintain the same panel weight to pressure ratio.

Weight of Panel. In our study, we manually fabricated 13 panels and verified their accuracy using laser scans of the membrane. The mean error was found to be 0.13 mm, with a maximum error of

Table 1. Quantitative evaluation of the panels TP1, TP2, TP3, and TP4: Mean absolute error (MAE) and mean absolute percentage error (MAPE) are measured in total and with respect to x -, y -, and z -directions. The maximum absolute and maximum relative errors are also measured with respect to each spatial direction. The upper table refers to the results obtained from our optimization process while the lower table refers to the scans of the deflated membrane. For fair comparison between panels, let the divider of the mean absolute percentage error (MAPE) being the corresponding dimension length of the panel bounding box for x -MAPE, y -MAPE, and the spatial diagonal length of the panel bounding box for the rest MAPE.

Panel	Total MAE (MAPE)	x -MAE (x -MAPE)	y -MAE (y -MAPE)	z -MAE (z -MAPE)	Total MAX (RMAX)	x -MAX (x -RMAX)	y -MAX (y -RMAX)	z -MAX (z -RMAX)
TP1	0.41 mm (0.13%)	0.04 mm (0.01%)	0.06 mm (0.02%)	0.40 mm (0.13%)	2.03 mm (0.64%)	0.39 mm (0.12%)	0.50 mm (0.16%)	1.93 mm (0.61%)
TP2	0.99 mm (0.30%)	0.14 mm (0.04%)	0.24 mm (0.07%)	0.93 mm (0.29%)	3.76 mm (1.16%)	0.85 mm (0.26%)	2.11 mm (0.65%)	3.00 mm (0.92%)
TP3	0.39 mm (0.12%)	0.03 mm (0.01%)	0.06 mm (0.02%)	0.38 mm (0.12%)	1.69 mm (0.52%)	0.22 mm (0.07%)	0.57 mm (0.18%)	1.59 mm (0.49%)
TP4	0.75 mm (0.23%)	0.03 mm (0.01%)	0.16 mm (0.05%)	0.72 mm (0.22%)	2.78 mm (0.86%)	0.19 mm (0.06%)	0.92 mm (0.29%)	2.62 mm (0.82%)

Panel	Total MAE (MAPE)	x -MAE (x -MAPE)	y -MAE (y -MAPE)	z -MAE (z -MAPE)	Total MAX (RMAX)	x -MAX (x -RMAX)	y -MAX (y -RMAX)	z -MAX (z -RMAX)
TP1	0.37 mm (0.10%)	0.04 mm (0.01%)	0.06 mm (0.02%)	0.36 mm (0.10%)	1.33 mm (0.37%)	0.30 mm (0.08%)	0.46 mm (0.13%)	1.25 mm (0.35%)
TP2	0.87 mm (0.27%)	0.13 mm (0.04%)	0.20 mm (0.06%)	0.82 mm (0.26%)	2.72 mm (0.85%)	0.97 mm (0.30%)	1.32 mm (0.41%)	2.69 mm (0.84%)
TP3	0.41 mm (0.13%)	0.03 mm (0.01%)	0.07 mm (0.02%)	0.40 mm (0.13%)	1.68 mm (0.53%)	0.21 mm (0.07%)	0.55 mm (0.17%)	1.59 mm (0.50%)
TP4	0.39 mm (0.12%)	0.01 mm (0.00%)	0.10 mm (0.03%)	0.38 mm (0.12%)	1.26 mm (0.39%)	0.13 mm (0.04%)	0.53 mm (0.17%)	1.15 mm (0.36%)

0.82 mm for a 205×205 mm² panel. This verification process indicates that the weight of the panels has a negligible influence on the overall accuracy. The weight per area of our panels is 70 N/m², which is relatively small compared to the applied pressure range (−600 Pa to −2000 Pa) and the membrane stress. This results in a weight-to-pressure ratio of approximately 1 to 9 to 1 to 28. Our control system ensures constant maximum deflation, which helps to mitigate some of the weight’s impact. At full scale, using 5 cm thick concrete and an appropriate membrane, the weight-to-pressure ratio could improve significantly, reaching up to 1 to 80. This demonstrates that the weight of the panels is not a critical factor affecting the performance and accuracy of the fabricated panels.

Materials of Panels and 3D Printing. In our proof of concept, at model scale, we used gypsum. At a larger scale, panels from a wide range of materials could be fabricated as long as processing temperatures are compatible with the flexible membrane. Examples include, but are not limited to, the thermoforming of thermoplastic material like PMMA (acrylic glass) or PC (polycarbonate), cast materials like concrete and ceramics, and fiber/resin systems like laminates from glass fiber mates and epoxy resin. Especially interesting is the combination with additive manufacturing technologies. As we demonstrate, 3D printing on the flexible mold results in an accurate, perfectly smooth outer panel surface with precise boundaries. At the same time, the 3D printing process would facilitate control of panel thickness and even the integration of additional features, for example, to mount the panel. In recent years, 3D concrete printing has matured to allow large-scale architectural applications. In combination with our work, it would enable low-cost fabrication of smooth panels.

Durability. For all fabrications in our paper, we used two membranes, but the exchange was only cautionary. We reproduced TP2 settings for after seven weeks reaching the same error numbers.

Full Reconfigurability. The small-scale machine developed to verify simulation and optimization results requires a new CNC-machined MDF support frame for each panel shape, making it not *fully reconfigurable*. While this machining process is faster and more resource-efficient than milling a full mold, a fully reconfigurable mold could eliminate this step. Such a mold would allow surface shape changes with the push of a button, achievable through existing technology.

The wooden frame can be replaced by a small number of digitally controlled pistons connected by a flexible member. The technology used would be very similar to the one used by Adapta Adaptive Molds for their adaptable molds but much more cost-efficient; while Adapta uses a dense 2D grid of pistons (540 for their smallest machine), we would only use a small number of pistons along the edges (we estimate 28 to 44 for a similarly sized machine).

Shape Space. Analyzing the space of all possible target panels \mathcal{T} that can be fabricated by deflating the membrane \mathcal{S} in contact with a support frame centerline γ is interesting. Our method approximates the target panel’s shape by using a subset of the membrane’s shape space, assuming similarity when varying γ . Facade panels are typically not highly curved due to their scale. In fact, the 3×3 model we present is considered highly curved by industry standards. Panels with consistent mean curvature normals suit our mold well, as uniform curvature ensures smooth deformation, reduces tension imbalances, and allows for accurate reproduction of curved surfaces. However, how different support frame centerlines γ affect the final shape space remains unclear. Recently, Huang et al. [2024] worked on a differentiable solver for time-dependent deformation problems with contact, which could provide a promising direction for exploring the influence of γ on membrane shape optimization.

ACKNOWLEDGMENTS

The authors are grateful to Helmut Pottmann for his valuable contributions to conceptual discussions and to Xiaole Li for measuring the membrane’s material parameters. Moreover, the authors would like to express their appreciation to Arnav Verma, Bingjian Huang, Kinjal Parikh, and Yuta Noma for their assistance in proofreading. This work has been supported by KAUST and TU Wien baseline funds, and an NSERC Discovery Grant (RGPIN-2021-03733). The reviewers’ valuable comments are gratefully acknowledged.

REFERENCES

- B. Audoly and Y. Pomeau. 2010. *Elasticity and geometry: from hair curls to the non-linear response of shells*. Oxford university press.
- R. Bridson, R. Fedkiw, and J. Anderson. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. 594–603.
- H.-Y. Chen, A. Sastry, W. M. van Rees, and E. Vouga. 2018. Physical simulation of environmentally induced thin shell deformation. *ACM Trans. Graph.* 37, 4 (2018), 1–13.



Fig. 14. Photos of the facade composed of 3×3 panels fabricated with our flexible mold. Please zoom in for a closer look at the surface.

- Y. Chen and G. Medioni. 1992. Object modelling by registration of multiple range images. *Image and vision computing* 10, 3 (1992), 145–155.
- Y. J. E. Chen, S. H. Sheen, U. M. Ascher, and D. K. Pai. 2020. SIERE: A Hybrid Semi-Implicit Exponential Integrator for Efficiently Simulating Stiff Deformable Objects. *ACM Trans. Graph.* 40, 1, Article 3 (aug 2020), 12 pages.
- R. W. Clough and C. P. Johnson. 1968. A finite element approximation for the analysis of thin shells. *International Journal of Solids and Structures* 4, 1 (1968), 43–60.
- K. C. Datsiou. 2017. *Design and performance of cold bent glass*. Ph.D. Dissertation. University of Cambridge. <https://doi.org/10.17863/CAM.15628>
- M. Eigensatz, M. Kilian, A. Schiftner, N. Mitra, H. Pottmann, and M. Pauly. 2010. Paneling Architectural Freeform Surfaces. *ACM Trans. Graphics* 29, 4 (2010), #45:1–10. Proc. SIGGRAPH.
- R. E. English, M. Lentine, and R. Fedkiw. 2012. Interpenetration free simulation of thin shell rigid bodies. *IEEE transactions on visualization and computer graphics* 19, 6 (2012), 991–1004.
- T. Fischer. 2012. Geometry Rationalization for Non-Standard Architecture. *Architecture Science* 5 (01 2012), 25–47.
- K. Gavriil, R. Guseinov, J. Perez, D. Pellis, P. Henderson, F. Rist, H. Pottmann, and B. Bickel. 2020. Computational design of cold bent glass facades. *ACM Trans. Graphics* 39, 6 (2020), 208:1–208:16.
- E. Grinspun, Y. Gingold, J. Reisman, and D. Zorin. 2006. Computing discrete shape operators on general meshes. In *Computer Graphics Forum*, Vol. 25. Wiley Online Library, 547–556.
- E. Grinspun, A. N. Hirani, M. Desbrun, and P. Schröder. 2003. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Citeseer, 62–67.
- Z. Huang, D. C. Tozoni, A. Gjoka, Z. Ferguson, T. Schneider, D. Panozzo, and D. Zorin. 2024. Differentiable solver for time-dependent deformation problems with contact. *ACM Trans. Graph.* 43, 3, Article 31 (may 2024), 30 pages. <https://doi.org/10.1145/3657648>
- C. Jiang, H. Wang, V. C. Inza, F. Dellinger, F. Rist, J. Wallner, and H. Pottmann. 2021. Using isometries for computational design and fabrication. *ACM Trans. Graphics* 40, 4 (2021), 42:1–42:12.
- D. Jourdan, M. Skouras, E. Vouga, and A. Bousseau. 2022. Computational Design of Self-Actuated Surfaces by Printing Plastic Ribbons on Stretched Fabric. *Computer Graphics Forum* 41, 2 (2022), 493–506. <https://doi.org/10.1111/cgf.14489>
- W. T. Koiter. 1966. On the nonlinear theory of thin elastic shells. *Proc. Koninkl. Ned. Akad. van Wetenschappen, Series B* 69 (1966), 1–54.
- M. Konaković-Luković, J. Panetta, K. Crane, and M. Pauly. 2018. Rapid deployment of curved surfaces via programmable auxetics. *ACM Trans. Graph.* 37, 4, Article 106 (2018), 13 pages.
- B. Kromoser and P. Huber. 2016. Pneumatic formwork systems in structural engineering. *Advances in Materials Science and Engineering*, Article 4724036 (2016).
- B. Kromoser and J. Kollegger. 2015. Pneumatic forming of hardened concrete – building shells in the 21st century. *Structural Concrete* 16, 2 (2015), 161–171.
- M. Li, Z. Ferguson, T. Schneider, T. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D. M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Trans. Graph. (SIGGRAPH)* 39, 4, Article 49 (2020).
- M. Li, D. M. Kaufman, and C. Jiang. 2021. Codimensional Incremental Potential Contact. *ACM Trans. Graph.* 40, 4, Article 170 (2021).
- Y. Li, T. Du, K. Wu, J. Xu, and W. Matusik. 2022. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Trans. Graph.* 42, 1 (2022), 1–20.
- J. Liang, M. Lin, and V. Koltun. 2019. Differentiable cloth simulation for inverse problems. *Advances in Neural Information Processing Systems* 32 (2019).
- D. L. Michels, V. T. Luan, and M. Tokman. 2017. A Stiffly Accurate Integrator for Elastodynamic Problems. *ACM Trans. Graph.* 36, 4, Article 116 (July 2017), 14 pages. <https://doi.org/10.1145/3072959.3073706>
- D. L. Michels, J. P. T. Mueller, and G. A. Sobottka. 2015. A Physically based Approach to the Accurate Simulation of Stiff Fibers and Stiff Fiber Meshes. *Computers & Graphics* 53 (2015), 136–146.
- D. L. Michels, G. A. Sobottka, and A. G. Weber. 2014. Exponential Integrators for Stiff Elastodynamic Problems. *ACM Trans. Graph.* 33, 1, Article 7 (feb 2014), 20 pages. <https://doi.org/10.1145/2508462>
- R. Narain, T. Pfaff, and J. F. O'Brien. 2013. Folding and Crumpling Adaptive Sheets. *ACM Trans. Graph.* 32, 4 (July 2013), 51:1–8. <http://graphics.berkeley.edu/papers/Narain-FCa-2013-07/> Proceedings of ACM SIGGRAPH 2013, Anaheim.
- B. Nicolet, A. Jacobson, and W. Jakob. 2021. Large Steps in Inverse Rendering of Geometry. *ACM Trans. Graph.* 40, 6 (Dec. 2021). <https://doi.org/10.1145/3478513.3480501>
- E. Oñate and B. Kröplin. 2008. Textile Composites and Inflatable Structures II. (2008).
- J. Panetta, F. Isvoranu, T. Chen, E. Siefert, B. Roman, and M. Pauly. 2021. Computational inverse design of surface-based inflatables. *ACM Trans. Graph.* 40, 4, Article 40 (2021), 14 pages.
- D. Pellis, M. Kilian, H. Pottmann, and M. Pauly. 2021. Computational Design of Weingarten Surfaces. *ACM Trans. Graph.* 40, 4, Article 114 (jul 2021), 11 pages. <https://doi.org/10.1145/3450626.3459939>
- D. Pellis, M. Kilian, H. Wang, C. Jiang, C. Müller, and H. Pottmann. 2020. Architectural freeform surfaces designed for cost-effective paneling mold re-use. In *Advances in Architectural Geometry*.
- H. Pottmann, M. Eigensatz, A. Vaxman, and J. Wallner. 2015. Architectural Geometry. *Computers and Graphics* 47 (2015), 145–164. <https://doi.org/10.1016/j.cag.2014.11.002>
- H. Pottmann, Q.-X. Huang, Y.-L. Yang, and S.-M. Hu. 2006. Geometry and convergence analysis of algorithms for registration of 3D shapes. *Int. J. Computer Vision* 67, 3 (2006), 277–296.
- G. Quinn and C. Gengnagel. 2014. A review of elastic grid shells, their erection methods and the potential use of pneumatic formwork. *WIT Transactions on the Built Environment* 136. <https://doi.org/10.2495/MAR140111>
- C. Rogers and W. Schief. 2003. On the equilibrium of shell membranes under normal loading. Hidden integrability. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 459, 2038 (2003), 2449–2462.
- M. Skouras, B. Thomaszewski, P. Kaufmann, A. Garg, B. Bickel, E. Grinspun, and M. Gross. 2014. Designing inflatable structures. *ACM Trans. Graph.* 33, 4 (2014), 1–10.
- O. Sorkine and M. Alexa. 2007. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, Vol. 4. Citeseer, 109–116.
- S. Suzuki, A. Martin, Y. Ren, T.-Y. Chen, S. Parascho, and M. Pauly. 2023. BamX: Rethinking Deployability in Architecture through Weaving. *Advances in Architectural Geometry* 2023 (2023).
- J. Teran, E. Sifakis, G. Irving, and R. Fedkiw. 2005. Robust quasistatic finite elements and flesh simulation. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '05)*. Association for Computing Machinery, New York, NY, USA, 181–190. <https://doi.org/10.1145/1073368.1073394>
- N. Umetani, D. M. Kaufman, T. Igarashi, and E. Grinspun. 2011. Sensitive couture for interactive garment modeling and editing. *ACM Trans. Graph.* 30, 4 (2011), 90.
- W. M. Van Rees, E. Vouga, and L. Mahadevan. 2017. Growth patterns for shape-shifting elastic bilayers. *Proceedings of the National Academy of Sciences* 114, 44 (2017), 11597–11602.
- Z. Wang, L. Wu, M. Fratarcangeli, M. Tang, and H. Wang. 2018. Parallel Multigrad for Nonlinear Cloth Simulation. *Computer Graphics Forum* (2018). <https://doi.org/10.1111/cgf.13554>
- C. Weischedel. 2012. A discrete geometric view on shear-deformable shell models. (2012).
- J. E. Zhang, J. Dumas, Y. R. Fei, A. Jacobson, D. L. James, and D. M. Kaufman. 2022. Progressive Simulation for Cloth Quasistatics. *ACM Trans. Graph.* 41, 6, Article 218 (nov 2022), 16 pages.