DÉPARTEMENT GÉNIE CIVIL ET CONSTRUCTION

Projet de Fin d'Études

Multiscale Differentiable Physics for the Inverse Design of Kirigami Sheets

Julien Kloers

Promotion 2025

Organisme d'accueil Form-Finding Lab, Princeton University
Princeton, NJ, USA

Encadrement Prof. Sigrid Adriaenssens, Princeton University (*Tuteur professionnel*)
Dr. Rafael Pastrana, Princeton University
Dr. Isabel Moreira de Oliveira, Princeton University
Dr. Robin Oval, École des Ponts (*Tuteur académique*)

Avril–Août 2026

Abstract

Kirigami turns a flat sheet into a mechanism by cutting it. Inverse design determines the cuts that open it into a prescribed shape under a chosen actuation. The uniform sheet constrains the answer at two scales. At the pattern scale, the sheet must open as a mechanism, which strongly restricts where the cuts fall. At the hinge scale, each hinge is a thin ligament of that same material: it buckles out of plane, yields and accumulates damage as the sheet opens. This work therefore designs the cut pattern *and* the hinges in one optimisation, differentiating the objective through the equilibrium path from flat sheet to deployed shape. First, the cut pattern is parameterised to guarantee deployability, within a family of quadrangular tessellations of periodic topology and non-uniform geometry, without penalty or repair. Second, the sheet is modelled as rigid panels joined by compliant hinges, its deployed shape minimising the total potential, the energy stored in the hinges less the work of the loads. A neural surrogate supplies that stored energy, learned once from elasto-plastic finite-element solves of one ligament, about eight orders of magnitude faster. Third, hinge geometry joins the cut pattern in one design vector, set by one gradient. Matching a target outline alone leaves the problem under-constrained, since a continuous set of cut patterns reaches the same shape; what a design can then be tuned against is the distribution of internal forces the actuation develops. Given a target and an actuation, the method then finds where to cut the sheet so that it opens onto that target, while limiting hinge damage and the compression and shear that drive it. One design was cut from polyester; its measured deployed shape departs from the prediction by 0.70% of the sheet length.

Keywords: kirigami, inverse design, differentiable physics, deployable structures, elasto-plastic hinges

Contents

List of Figures

List of Tables

Glossary

Terms appear in the order this work introduces them, with the symbols attached to each.

panel P_k One of the n rigid quadrilateral faces the cut pattern produces, held rigid by Assumption A1: centroid $\boldsymbol{c}_k$, centroid-to-corner vectors $\boldsymbol{d}_{kj}$.

hinge $h \in \mathcal{H}$ The compliant link between two panels. The term is general and says nothing about how the link is realised.

topology The connectivity of the panels and hinges.

geometry The shape and size of each panel and hinge.

void ϕ The opening that forms at a cut as the sheet deploys.

unit cell The four panels that meet around one void.

closed state The flat state in which every void is collapsed onto its cut and the panels tile the sheet exactly; the sheet is cut in it.

cut pattern $\boldsymbol{\kappa}$ The arrangement of cuts decomposing the sheet into panels, fixing the flat reference geometry and a connectivity that deformation does not change. Its parameters are $\boldsymbol{\kappa} = (\boldsymbol{r}, \boldsymbol{\tau})$ (2).

deployable A pattern is deployable if and only if its rigid panels form a mechanism with exactly one internal degree of freedom.

ligament The realisation of a hinge in a sheet cut from one uniform material: the thin strip of material left between two cuts. Every hinge in this work is a ligament, and the two words are not interchangeable, one naming the link and the other the material that makes it.

surrogate $\boldsymbol{\chi}$, χ_D The function fitted once, offline, to the three-dimensional solves and evaluated in their place inside the design loop (14).

damage D The fraction of the material's ductility consumed at a hinge (12): the equivalent plastic strain averaged over the window, over the fracture strain ε_f.

inner problem The equilibrium of the sheet at a fixed design: the minimisation of the total potential Π that returns the deployed state $\boldsymbol{q}^\star$ under the prescribed actuation (10).

deployed state $\boldsymbol{q}^\star$ The configuration the sheet takes up at equilibrium under a prescribed actuation: the minimiser $\boldsymbol{q}^\star(\zeta; \boldsymbol{\theta})$ of (10).

outer problem The design search: the minimisation of the design objective $\mathcal{L}$ over the design vector $\boldsymbol{\theta}$ in the admissible set Θ, evaluated at the state the inner problem returns (16).

hinge design vector $\boldsymbol{g}_h$ The vector $\boldsymbol{g}_h = (w_{\mathrm{lig},h}, \alpha_h(\boldsymbol{\kappa}), \rho_f, w_c, r_{\mathrm{win}})$ specifying one hinge: ligament width, cut angle, fillet ratio, kerf and window radius (7).

window Ω, r_{win} The three-dimensional domain on which the fine-scale solve is posed: a half-disc of radius r_{win} facing away from the cut.

chart $\boldsymbol{z}$, $\boldsymbol{\ell}$, $\boldsymbol{\omega}$ A smooth map from unconstrained coordinates onto the admissible design set, so that every value the optimiser can reach is a fabricable, deployable pattern (5).

condensation The reduction of a resolved fine-scale problem to one effective law the sheet-level model can evaluate.

design vector $\boldsymbol{\theta}$ Everything the optimiser sets, $\boldsymbol{\theta} = (\boldsymbol{\kappa}, \{\boldsymbol{g}_h\})$ of (1), ranging over the admissible set Θ the charts define.

void aspect ratio r One number per interior cut, placing its two endpoints at the fractions r and $1 - r$ between the two neighbouring cuts (3).

boundary position τ One number per non-corner cut on the sheet outline, locating it along its edge as a fraction (4). The field is $\boldsymbol{\tau}$, a component of $\boldsymbol{\kappa}$.

configuration $\boldsymbol{q}$, $\mathcal{C}$ A rigid motion of every panel away from the closed state, for each panel k a translation $\boldsymbol{u}_k$ and a rotation ϕ_k about its centroid $\boldsymbol{c}_k$; the $3n$ numbers $\boldsymbol{q} = (\boldsymbol{u}_k, \phi_k)_{k=1}^{n}$, of which $\boldsymbol{q} = \mathbf{0}$ is the closed state itself. The admissible set $\mathcal{C}$ is the configurations meeting the values the supports prescribe (10).

open state The terminal state of a deployment.

main cut L_{main} Of the two cuts meeting at the corner of a hinge, the one that ends there, bounding the hinge with its filleted tip O.

secondary cut L_{sec} The other, which runs on through that corner.

cut frame $\hat{\boldsymbol{e}}_{\text{sec}}$, $\hat{\boldsymbol{e}}_{\text{sec}}^{\perp}$ Per hinge: the unit tangent of its secondary cut in the flat sheet, with the counter-clockwise normal, signed so the hinge material lies on the non-positive side. Carried by the deployment: at $\boldsymbol{q}$ it is $R(\bar{\phi})\,\hat{\boldsymbol{e}}_{\text{sec}}$, with $\bar{\phi}$ the mean panel rotation that (8) removes.

ligament width w_{lig} The material left between the tip of the main cut and the flank of the secondary cut.

cut angle α_h The angle $\alpha_h \in (0, \pi)$ subtended at the cut tip between the secondary and the main cut, measured into the material.

kerf w_c The width of material the cutting tool removes, w_c of (7).

fillet ratio ρ_f The ratio of the cut-tip fillet radius over the ligament width.

hinge deformation $\boldsymbol{\varepsilon} = (a, s, \vartheta)$, $\boldsymbol{\varepsilon}_{ih}$ The three deformation measures defining the mechanical state of one hinge (8): the axial opening a, the shear s and the relative rotation ϑ of the two panels.

stored energy $\mathcal{W}$ The energy one hinge stores as a function of its deformation and its hinge design vector (9). Summed over $\mathcal{H}$ it gives the elastic energy U.

total potential Π The elastic energy less the work of the applied loads (10). Actuation enters as panel forces $\boldsymbol{F}_k$ and moments M_k scaled by the load factor ζ.

linear hinge law $\mathcal{W}^{\text{lin}}$ The standard energetic model of rotating-square kirigami: three uncoupled quadratic springs of stiffness k_a, k_s and k_ϑ (11).

training domain η_a, η_s, ϑ The box of hinge deformations the database covers.

1 Introduction

Kirigami (切り紙) joins *kiru* (切る), to cut, with *kami* (紙), paper. It is the cut counterpart of origami (折り紙), from *oru* (折る), to fold. The craft itself is difficult to date. Papermaking reached Japan from the Asian mainland early in the seventh century, and China has its own long tradition of cut paper in *jianzhi* (剪紙), but the inference from the arrival of paper to the beginning of a paper craft does not hold: for the closely related case of folding, the record cannot be traced back more than a few hundred years, and Japanese paper folding is dated instead to ceremonial wrappers of the fourteenth century [1]. During the Edo period the craft divided into an ornamental branch and an industrial one. In *monkiri* (紋切り) a sheet is folded once and cut once to yield a family crest. In *Ise-katagami* (伊勢型紙) stencils carry dyed patterns onto kimono cloth. Whatever is cut from such a stencil, the remainder must hold together as one sheet. European counterparts, *Scherenschnitte*, *wycinanki*, *papel picado* and silhouette work, evolved independently, and merged only in the 1860s and 1870s [1].

Cutting a flat sheet along a prescribed pattern turns it into a *compliant mechanism*, one that moves by deformation of its own material. The tessellation is an assembly of nearly rigid panels joined by hinges, and it deploys along kinematic modes fixed by the pattern [2]. A hinge is the compliant link between two panels. Inspired by the craft, an extensive family of mechanical metamaterials is built on this principle, and fabricated by several routes: three-dimensional printing, moulding, assembly and cutting [2]. Konaković-Luković *et al.* deploy a planar linkage of rigid triangles on rotational pins from a flat state to a doubly curved surface, and offer it for deployable architecture at large scale [3]. Their prototypes are cut as separate triangles and connected by metal rings. In a sheet that is cut, the tool removes material in plane only, so panels and hinges share the thickness of the sheet material. Applications include shape morphing under remote actuation [4], modular assemblies used as bodies for soft robots [5], deployable aerospace structures and kinetic facades [6], and load-carrying building components deployed at the scale of metres [7, 8]. A hinge is then a remainder, the part the tool did not cut, rather than a component. Its width is the only control the fabrication route allows over how strongly it resists deployment.

The mechanical response of a kirigami metamaterial is governed by its topology, the connectivity of the panels and hinges, and by its geometry, the shape and size of each. The pair can be represented as a planar graph. A few topologies have been studied extensively. Parallel rows of straight slits give the linear-cut sheet, whose panels are strips joined end to end [9, 10]. In rotating squares, quadrangular panels meet at their corners and turn in alternating directions as the voids between them open [11]. A topology is normally fixed at the outset, so the graph is given and the embedding is the design variable: panel size and aspect, slit length, hinge width. When the geometry is periodic, the response of the sheet follows from the kinematics of a single unit cell [12]. When only the topology is periodic, the connectivity repeats while the geometry of each cell may differ, so the sheet has a fixed, finite parameterisation without being uniform.

In a flat kirigami sheet, deployment increases the area the sheet encloses, while the material area of the panels remains constant. The size of that increase follows from where the cuts fall. In the closed state every void is collapsed onto its cut, and the panels tile the sheet exactly. Cutting kirigami from a single sheet of uniform material constrains the design at two scales. The first is geometric. A cut pattern must leave the sheet connected and must not self-intersect. A pattern is deployable if its rigid panels form a mechanism with exactly one internal degree of freedom. Deployability is a set of equality constraints on the cut positions, so the deployable patterns form a lower-dimensional variety inside the space of all patterns. For rigid planar quadrilateral tessellations with no interior hole and at least three panels in each direction, Dang *et al.* prove that a pattern is rigidly deployable if and only if every void of cuts is a parallelogram [13]. The second scale is the hinge. A hinge cut from sheet material is a strip of finite width, and the mechanics that resist deployment occur inside that strip. Reported kirigami metamaterials are predominantly at the centimetre scale and below [2, 14, 15]. During the deployment of a cut sheet, and at that scale, the ligaments buckle out of plane, the panels rotate out of plane with them, and the ligaments yield [14]. Bending a thin ligament out of plane costs less than deforming it in plane. For a strip of thickness t and width w the two flexural stiffnesses are in the ratio $(t/w)^2$, so a slender ligament reaches the out-of-plane mode first [10, 14]. The same instability is a route to three-dimensional structures cut from flat sheet [16]. These phenomena are three-dimensional, materially non-linear and irreversible. They are confined to the ligament and the sheet material around it, a region an order of magnitude below the size of the sheet.

Forward simulation maps design parameters to an equilibrium state. *Inverse design* selects those parameters so that the resulting design meets a prescribed requirement. For cut sheets it has so far been posed mainly on the geometry alone, and requirements beyond the geometry remain underexplored. Choi *et al.* invert a closed and compact tessellation of the plane onto a prescribed target [17], and describe a later compact construction of their own as material-independent [18]. Dudte *et al.* generate compact patterns by recursion on the voids [19]. Segall *et al.* widen that line from one tiling family to the tilings themselves. They give a condition under which a planar tiling can be cut into a rotationally deployable hinged pattern, and then minimise rigidity, planarity, shape and fairness terms over the coordinates of the flat tiling and of the deployed configuration at once [20]. Their flat state is gapless, as it is here. They record that material properties, actuation forces and the stability of the deployed structure lie outside their formulation. Mechanics enters elsewhere, but not on the cut pattern of a flat sheet. A bar-and-hinge network represents the mechanics of a cut sheet at reduced cost in forward simulation, with an analytically derived hinge stiffness and a linear elastic material [21]. A network mapping the geometry of a hinged-quadrilateral cell to a non-linear stress–strain response accelerates an evolution strategy over that geometry, for a moulded metamaterial rather than a cut sheet [22]. A conditional generative model, fine-tuned by reinforcement learning, produces cut patterns for prescribed silhouettes, on a forward simulator that omits forces and material non-linearity [15]. The determination of a shape as the equilibrium configuration of a structure under prescribed action is called form finding in the literature on shells and spatial structures [23]. What remains is to form find the cut pattern itself, so that a flat sheet reaches a prescribed shape under a prescribed actuation with the mechanics of the hinge in the loop. Requiring equilibrium makes the deployed shape the output of a solver, so a gradient on the cut pattern requires a derivative of the solver.

Differentiable physics denotes the implementation of a physical simulator whose output is differentiable with respect to its inputs [24]. The derivative is obtained by automatic differentiation of the solver, or by the implicit function theorem applied to its solution [25]. It permits end-to-end gradient-based optimisation of the quantities the simulator takes as input. Its principal uses are the inference of model parameters from observations and the replacement of solver components by learned surrogates [26]. A third is the inverse design of a system against a prescribed response [25]. Inverse problems on bar and cable systems date to the force-density method [27]. They have since been posed as gradient-based optimisation through a differentiable form of that solver [28]. A network trained to supply the parameters of that solver allows shape approximation at interactive rates [29]. Two applications to kirigami exist. Bordiga *et al.* differentiate a dynamic simulator of a hinged-panel sheet and tune the panel geometry so that the sheet focuses an elastic pulse on a target region, or shields that region from it [30]. Wang *et al.* differentiate an analytical model of the kinematics and the stored energy of a magnetically actuated sheet. They co-design the cuts and the magnetisation for a target shape [4]. A cut sheet has design parameters for every cut and every hinge, so their number grows with the sheet. Differentiating the solution returns the whole gradient at a cost set by one solve, and not by the number of parameters. What is missing is a single formulation that holds three properties at once: a parameterisation on which every reachable design is fabricable and deployable, an equilibrium solve that reaches a prescribed shape, and a hinge law that represents buckling, yielding and damage without a finite-element solve inside the design loop. That combination is what this work supplies.

This work treats flat-sheet kirigami at metre scale, in which the mechanism is cut from a single sheet of uniform material, so a hinge cannot be specified independently of the sheet it is cut from. At that size the sheet is cut flat, shipped flat and deployed as one piece, so the cost of fabrication follows a cutting path rather than a count of assembled parts. The sheet material and the actuation are given. Two choices remain: where the cuts fall, and how wide the ligament left between them is. The problem is inverse: the objective is the discrepancy between the deployed geometry and a prescribed target, evaluated at equilibrium under prescribed actuation.

From a computational point of view, the formulation is therefore nested. The inner problem is the equilibrium of the sheet: it holds the design fixed and returns the deployed state $\boldsymbol{q}^\star$ that minimises the total potential energy Π. The outer problem is the design search: it sets the cut pattern and every hinge, collected in $\boldsymbol{\theta}$, from an admissible set Θ. Its objective is $\mathcal{L}$, evaluated at the state the inner problem returns,

$$\boldsymbol{\theta}^\star = \arg\min_{\boldsymbol{\theta}\in\Theta} \ \mathcal{L}\big(\boldsymbol{q}^\star(\boldsymbol{\theta})\big), \qquad \boldsymbol{q}^\star(\boldsymbol{\theta}) = \arg\min_{\boldsymbol{q}} \ \Pi(\boldsymbol{q};\, \boldsymbol{\theta}).$$

The outer objective cannot be evaluated until the inner problem is solved.

The first scale is answered by construction. The topology is the rigid planar quadrilateral tessellation. The flat sheet is generated so that the condition of Dang *et al.* [13] holds by construction. Every admissible parameter value therefore yields a connected, non-self-intersecting and rigidly deployable pattern. Designing the pattern is then the choice of an embedding of a fixed graph. Penalising is the alternative [20], and outside the admissible set the flat geometry is undefined rather than poor, so a penalty cannot be measured there.

The second scale is answered by an energy. The energy Π is that of a sheet modelled as rigid panels joined by compliant hinges. Each hinge has a scalar stored-energy function $\mathcal{W}(\boldsymbol{\varepsilon}; \boldsymbol{g})$ [30]. Its first argument $\boldsymbol{\varepsilon}$ is the relative panel motion the assembly imposes, three rigid-body-invariant measures. Its second is the hinge design vector $\boldsymbol{g}$. That function is itself the value of a third minimisation, over a window about the hinge, in three-dimensional elasto-plastic finite elements. It is fitted to those solutions once, so the inner problem of the sheet reduces to a differentiable network evaluation. Its gradient is the internal generalised force and its Hessian is the tangent stiffness.

The parameterisation, the simulator and the hinge energy are differentiated together, so one gradient reaches the cut pattern and every hinge design at once (Figure 1). This work makes four contributions.

The cut pattern is parameterised so that fabrication and deployability hold by construction, on a criterion that is necessary as well as sufficient. Measured against penalised alternatives of the same task, the constructed chart is the one whose objective is differentiable where it is optimised (Sections 2.1 and 3.1). Within that chart, a condensation of the hinge brings out-of-plane buckling, yielding and ductile damage into a rigid-panel simulator: the hinge scale is resolved once, offline, and each subsequent hinge evaluation has fixed cost (Sections 2.4 and 4.2). Solving the problem those two make well posed then shows that the actuation, and not the target alone, selects the cut pattern. One sheet held identically develops distinct internal force distributions under different actuations, and designing against the actuation reduces the boundary error of a pattern designed from kinematics alone (Sections 3.1 and 3.2). The chain is finally tested against the physical sheet. One design is cut from polyester sheet, deployed and measured against its prediction, and one ligament is photographed in its buckled state and compared with the finite-element solution the surrogate was fitted to (Section 4).

Section 2 sets out the parameterisation, the hinge energy and the objective. The Glossary at the front of the document collects every term and symbol the work defines. Section 3 measures what the construction forgoes against unconstrained formulations, whether a target determines a design, what the actuation fixes and what designing against it recovers, all under the linear-spring hinge law. Section 4 identifies a polyester and measures its ligament, fits the condensed hinge energy to that material, replaces the linear-spring law by it, and cuts, deploys and measures one design. Section 5 establishes when out-of-plane deformation localises in the ligaments and the error the state-sufficiency assumption introduces. It verifies the gradient the design loop uses.

2 Methods

Cutting a kirigami sheet for its physical deployment requires two decisions, and they are taken at two scales. At the scale of the sheet, the position of the cuts sets the reachable shapes. At the scale of one hinge, the amount of material left uncut sets

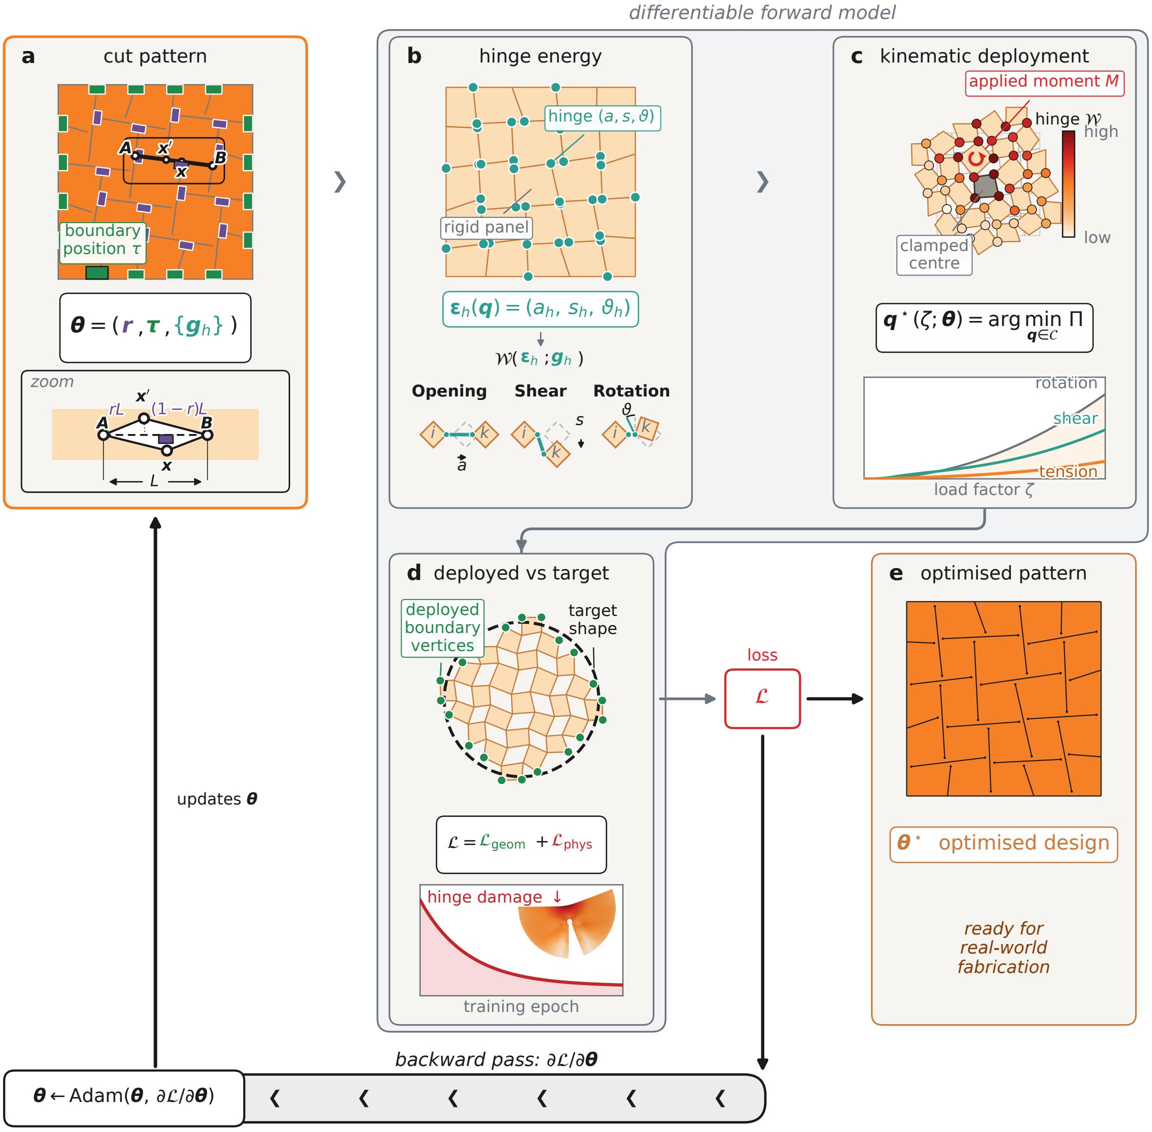

Fig. 1 The differentiable inverse-design pipeline. **(a)** The flat cut pattern set by the design parameters $\boldsymbol{\theta}$, one colour for each of $\boldsymbol{r}$, $\boldsymbol{\tau}$ and the hinge designs $\{\boldsymbol{g}_h\}$, drawn as idealised segments without kerf or fillets. The zoom shows one interior cut opened: $\boldsymbol{A}$ and $\boldsymbol{B}$ are the facing endpoints of its two neighbours and $\boldsymbol{x}$, $\boldsymbol{x}'$ its own, the dashed line is the closed-state segment $[\boldsymbol{A}, \boldsymbol{B}]$ of length L, and the void has sides $(1-r)L$ and rL (Section 2.1). **(b)** The hinge deformation $\boldsymbol{\varepsilon}(\boldsymbol{q}) = (a_h, s_h, \vartheta_h)$ of (8), on two adjacent panels, and the stored energy $\mathcal{W}(\boldsymbol{\varepsilon}; \boldsymbol{g})$ evaluated on them. **(c)** The deployed configuration of (10), panels coloured by their stored hinge energy, with that energy supplied by the learned surrogate (14). **(d)** The deployed boundary vertices against the target, and the objective (16) they enter. **(e)** The optimised pattern $\boldsymbol{\theta}^{\star}$, realised with kerf and cut-tip fillets as it would be cut. The mismatch at (d) is differentiated back to $\boldsymbol{\theta}$ through (20).

the deployment force and the margin against tearing. The two decisions are the two blocks of the design vector

$$\boldsymbol{\theta} = \big(\boldsymbol{\kappa},\ \{\boldsymbol{g}_h\}\big). \tag{1}$$

The first block is the cut pattern. It is an exact parameterisation of the cuts and panels, where a hinge is reduced to a point. The second block is one hinge design vector $\boldsymbol{g}_h$ per hinge to specify the material left at a hinge, and that material is what resists deployment. This work designs both blocks at once: the sheet is modelled as an assembly of rigid panels joined by slender compliant hinges (A1). The reduction holds when the panels are much stiffer than the hinges. The sheet material, the outline, the topology and the actuation are fixed before the design begins. The formulation rests on five assumptions:

(A1) The panels are rigid. All deformation localises in the hinges.

(A2) The panels remain coplanar throughout deployment. The condition on panel size under which this holds is derived in Section 5.1.

(A3) Hinge contributions are constitutively local. Given $\boldsymbol{\varepsilon}$ and $\boldsymbol{g}$ the stored energy of a hinge is determined, and the energy of the assembly is the sum over hinges.

(A4) Stored energy and damage are single-valued functions of the current $\boldsymbol{\varepsilon}$ and $\boldsymbol{g}$, not of the deformation path. Every hinge is driven along a straight ray from the undeformed state, where this holds exactly. Section 5.2 measures the error it introduces on the paths a deployed sheet follows.

(A5) Deployment is quasi-static.

Under Assumptions A1 and A2 the kinematics of a deployable pattern form a one-parameter family. Actuation is applied at selected panels. The sheet is in equilibrium where the total potential, the stored hinge energy less the work of the applied loads, is stationary. That equilibrium fixes the deployed configuration, the load-displacement response, and the damage accumulated in the hinges.

2.1 The cut pattern

A cut pattern is the set of slits cut into the flat sheet. It removes no material. Each slit separates the material on its two sides everywhere but at its endpoints. The regions the slits enclose are the panels, and the endpoints where two panels stay attached are the hinges. Cutting therefore fixes the flat shape of every panel and the connectivity between them. No deformation alters that connectivity, and it alone fixes how the assembly can move. The topology is fixed throughout this work: quadrilateral panels on an $M \times N$ planar grid, adjacent panels attached at their shared corner. Taken as rigid, the panels have $3MN$ degrees of freedom in the plane, and each corner hinge removes two of them. Three of those that remain are rigid-body motions of the whole sheet. The pattern is deployable with exactly one degree of freedom.

Every cut is a slit with two endpoints, and the sheet has $P = 2(M{+}1)(N{+}1)$ of them. Collect their coordinates in $\boldsymbol{X} \in \mathbb{R}^{P\times 2}$. Deployability is a system of polynomial equality constraints on $\boldsymbol{X}$, so the deployable patterns are a proper real algebraic subset $\mathcal{D} \subsetneq \mathbb{R}^{2P}$, of Lebesgue measure zero. An optimiser working in $\boldsymbol{X}$ must therefore impose $\mathcal{D}$ from outside, by projecting each iterate back onto it or by penalising the constraint residual. The alternative is to drop the free coordinates and parameterise $\mathcal{D}$ itself, so that every iterate deploys and no constraint enters the objective.

The rigid deployable planar quadrangular kirigami (RDPQK) construction of Dang *et al.* [13] is posed in the closed state: the cuts are slits of zero width, the panels tile the plane without gaps, and this is the state the sheet is cut in. It places the cut endpoints so that the four cuts around every interior vertex are collinear in pairs, which makes every void a parallelogram once the sheet opens, and that is necessary as well as sufficient for rigid deployability (Appendix A). $\boldsymbol{\kappa}$ places the cuts, and the panels follow from the cuts by one linear solve. The void aspect ratios and the boundary positions constitute the cut pattern.

$$\boldsymbol{\kappa} = \big(\boldsymbol{r}, \boldsymbol{\tau}\big), \tag{2}$$

An $M \times N$ sheet has $(M-1)(N-1)$ aspect ratios and $2(M+N-2)$ boundary positions. Cuts are indexed on an $(M+1)(N+1)$ grid. The cut C_{ij} is horizontal where $i+j$ is even and vertical where it is odd, with endpoints $\boldsymbol{x}_{ij}$ and $\boldsymbol{x}'_{ij}$. An interior cut C_{ij} has one neighbouring cut on either side, along i when it is horizontal and along j when it is vertical. Let $\boldsymbol{A}_{ij}$ and $\boldsymbol{B}_{ij}$ be the two endpoints those neighbours face it with (Fig. 1a): $(\boldsymbol{x}_{i-1,j}, \boldsymbol{x}'_{i+1,j})$ when $i+j$ is even and $(\boldsymbol{x}'_{i,j-1}, \boldsymbol{x}_{i,j+1})$ when it is odd, at a separation $L_{ij} = \|\boldsymbol{A}_{ij} - \boldsymbol{B}_{ij}\|$. One void aspect ratio $r_{ij} \in (0,1)$ places both endpoints of C_{ij} on the segment $[\boldsymbol{A}_{ij}, \boldsymbol{B}_{ij}]$,

$$\boldsymbol{x}_{ij} = r_{ij}\,\boldsymbol{A}_{ij} + (1-r_{ij})\,\boldsymbol{B}_{ij}, \qquad \boldsymbol{x}'_{ij} = (1-r_{ij})\,\boldsymbol{A}_{ij} + r_{ij}\,\boldsymbol{B}_{ij}. \tag{3}$$

The two endpoints lie symmetrically about the midpoint of $[\boldsymbol{A}_{ij}, \boldsymbol{B}_{ij}]$, a distance $|1-2r_{ij}|L_{ij}$ apart. The void that opens at C_{ij} is a parallelogram of sides $r_{ij}L_{ij}$ and $(1-r_{ij})L_{ij}$ at every opening angle. The four panels meeting around C_{ij} are its *unit cell*, and disjoint unit cells tile the sheet.

A cut on the grid border has both endpoints at one point, and that point lies on the sheet outline. One boundary position $\tau \in (0,1)$ places it along its edge, between that edge's endpoints $\boldsymbol{o}_0$ and $\boldsymbol{o}_1$,

$$\boldsymbol{x}_{ij} = \boldsymbol{x}'_{ij} = (1-\tau)\,\boldsymbol{o}_0 + \tau\,\boldsymbol{o}_1. \tag{4}$$

The four cuts at the grid corners are held at the corners of the outline, and the positions on one edge keep their order. The outline is held fixed, so τ only slides its cut along one edge.

Both conditions are linear in the cut endpoints, so the P equations, (3) at an interior cut and (4) on the outline, assemble into one sparse system. The aspect ratios enter the matrix as weights and the boundary positions the right-hand side as data. One direct solve, once per design, returns the whole flat sheet from $\boldsymbol{\kappa}$. The panels follow immediately from those endpoints. The panels across a cut part as the sheet opens; those meeting at its endpoints stay attached, and those attachments define the hinge set $\mathcal{H}$. On the outline, panels stay attached only where a cut runs along the edge. The $n = MN$ panels are numbered P_k, each with a centroid $\boldsymbol{c}_k$ and centroid-to-corner vectors $\boldsymbol{d}_{kj}$: the flat reference geometry.

The optimiser steps in unconstrained coordinates $(\boldsymbol{z}, \boldsymbol{\ell})$, mapped onto the two blocks by the charts

$$\boldsymbol{r} = \sigma(\boldsymbol{z}), \qquad \tau_{e,s} = \sum_{m \le s} \big[\mathrm{softmax}(\boldsymbol{\ell}_e)\big]_m, \tag{5}$$

with σ the logistic function and $\tau_{e,s}$ the sth boundary position on outline edge e, read off k_e+1 logits $\boldsymbol{\ell}_e$. Softmax weights are positive, so the partial sums increase strictly in $(0, 1)$: no step length reaches an inadmissible design, no projection enters the gradient path, and the flat pattern is a smooth, non-self-intersecting function of $(\boldsymbol{z}, \boldsymbol{\ell})$, as in Tutte's embedding [31].

Once a cut pattern is defined, kinematic deployment follows. A configuration $\boldsymbol{q}$ is a rigid motion of each panel about its centroid, a translation $\boldsymbol{u}_k$ and a rotation ϕ_k, with the flat sheet at $\boldsymbol{q} = \boldsymbol{0}$, so corner j of panel k moves by $\Delta \boldsymbol{x}_{kj} = \boldsymbol{u}_k + (R(\phi_k) - I)\, \boldsymbol{d}_{kj}$ with $R(\phi)$ the rotation matrix. Kinematic deployment is a single path through those configurations, and the shape at each opening angle follows from $\boldsymbol{\kappa}$ in closed form. Figure 2 shows three states on that path. A deploying state has the voids partly open, the enclosed area growing while the panel area is conserved, and the open state is the terminal one, at the opening the design is driven to. The single internal degree of freedom fixes an *ideal deployment*: at opening angle ϑ every panel rotates by $\pm\vartheta/2$ about its own centroid. It bounds what the sheet can reach, and the target of Section 2.5 is read off it. Let a panel diagonal make an angle ψ with a given axis. Opening every hinge by a relative rotation ϑ turns each panel by $\vartheta/2$ and rotates that diagonal towards the axis, so the panel pitch along it grows by

$$\lambda(\vartheta) = \frac{\cos(\psi - \vartheta/2)}{\cos\psi}, \qquad \lambda_{\max} = \frac{1}{\cos\psi} \quad \text{at} \quad \vartheta = 2\psi. \tag{6}$$

The maximum is reached where two panel centres and their shared hinge become collinear, and the ratio falls beyond it.[1]

The pattern is a mathematical description of the flat sheet: no material and no load appear in it. A simulation of the physical actuation is built on it directly: the panels are independent rigid bodies (A1), coupled only through the hinge set $\mathcal{H}$. It remains to give those hinges a mechanical behaviour, and Section 2.2 begins with the geometry that behaviour is defined on.

2.2 Geometry of a hinge

The RDPQK construction treats a hinge as a mathematical point, a corner shared by the two panels it joins. The fabricated hinge is a strip of finite width, and the mechanics occur there. Every hinge is bounded by the tip of a main cut and the flank of a secondary cut, the two cuts that meet at a shared corner. The secondary cut gives it a direction. Each hinge therefore has its own local frame, the cut frame $(\hat{\boldsymbol{e}}_{\mathrm{sec}}, \hat{\boldsymbol{e}}_{\mathrm{sec}}^{\perp})$:

[1] The law follows from the closure of one hinged pair. Both panels are rigid and their shared corner stays shared. For a pair mirror-symmetric about the perpendicular bisector of its two centroids the arms $\boldsymbol{d}_{kj}$ meeting at that corner share a length d, the deployed pitch is $2d\cos(\psi - \vartheta/2)$ and the flat pitch is $2d\cos\psi$, whose ratio is (6) at every opening angle.

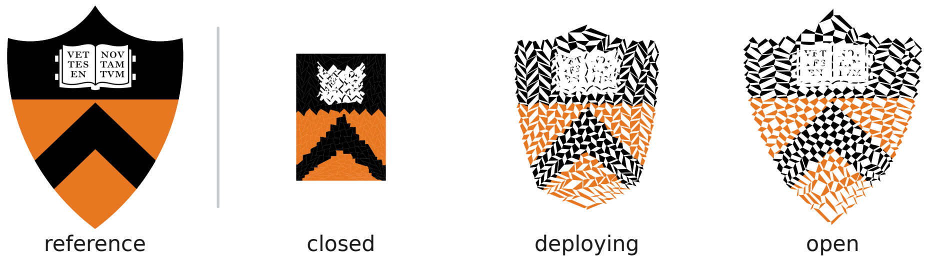

Fig. 2 The three states of one sheet. *Reference*: the Princeton University seal, whose silhouette is the design target of Section 3.1, set apart by the rule. *Closed*, *deploying*, *open*: the same sheet at opening angles $\vartheta = 0$, $\frac{1}{2}\vartheta^\star$ and $\vartheta^\star = 90°$. The panels are rigid and the hinges are pins, so each state is a closed-form configuration and not an equilibrium under load.

the unit vector along its secondary cut with the counter-clockwise normal, oriented so that the material lies on the non-positive side of the normal. The cut pattern fixes it in the flat state and the deployment carries it.

The hinge design vector $\boldsymbol{g}_h$ specifies the geometry of the remaining material at the end of a cut (Fig. 3b).

$$\boldsymbol{g}_h = \Big(\underbrace{w_{\mathrm{lig},h}}_{\text{learnable}},\ \underbrace{\alpha_h(\boldsymbol{\kappa})}_{\text{set by the cut pattern}},\ \underbrace{\rho_f,\ w_c,\ r_{\mathrm{win}}}_{\text{held fixed here}}\Big), \tag{7}$$

The hinge design vector $\boldsymbol{g}_h$ is the only channel by which a change of design reaches the hinge. Its components are defined at the scale of the hinge (Fig. 3b). It is fixed when the sheet is cut and stays constant during deployment. The ligament width w_{lig} is the material left between the main and the secondary cut. The cut angle $\alpha_h(\boldsymbol{\kappa})$ is the angle subtended at the cut tip between the secondary and the main cut, measured into the material. It is measured from $\boldsymbol{\kappa}$ as $\alpha_h = \operatorname{atan2}\big(-\hat{\boldsymbol{m}}\cdot\hat{\boldsymbol{e}}_{\mathrm{sec}}^{\perp}, \hat{\boldsymbol{m}}\cdot\hat{\boldsymbol{e}}_{\mathrm{sec}}\big) \in (0,\pi)$, with $\hat{\boldsymbol{m}}$ the unit vector from the cut tip along the main cut into the material. The kerf w_c is the width the tool removes in the main cut, and the fillet ratio ρ_f fixes a cut-tip fillet radius as a multiple of w_{lig}. The fifth component is the window radius r_{win}. It is not a cut feature but the extent of sheet material assigned to each hinge, and Section 2.4 defines it. Of the five, only the ligament width is set directly by the optimiser, charted from an unconstrained $\boldsymbol{\omega}$ as $\boldsymbol{w}_{\mathrm{lig}} = w_- + (w_+ - w_-)\,\sigma(\boldsymbol{\omega})$, with $[w_-, w_+]$ the interval of admissible ligament widths. Including the hinge design vectors in $\boldsymbol{\theta}$ brings the hinge design inside the optimisation loop, alongside the cut pattern, while the deployment stays parameterised by coordinates that are valid at every value they take. At every step of the outer problem, the optimiser moves the cuts and redesigns every hinge.

All deformation localises in the hinges (A1). What deforms a hinge is the relative motion of the two panels it joins, measured in that hinge's own cut frame; measured in a fixed frame it would also change under a rigid rotation of the whole sheet, which deforms nothing. That motion is the hinge deformation $\boldsymbol{\varepsilon}_h(\boldsymbol{q})$, the mechanical state of one hinge, given by three deformation measures (Fig. 1b). A configuration reaches the stored-energy function of Section 2.3 through it and through nothing else. Let hinge

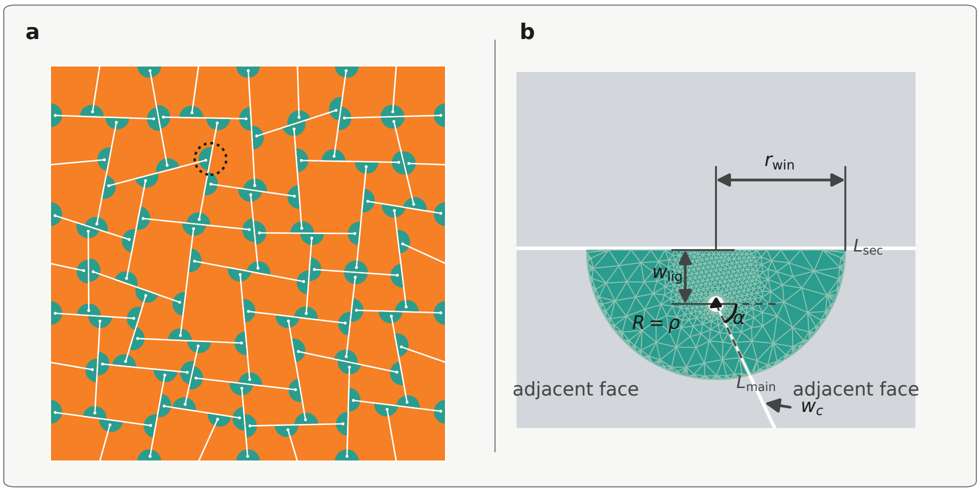

Fig. 3 Hinge geometry. **(a)** A closed tessellation, with the hinge window drawn around every hinge and one circled. **(b)** That hinge, drawn away from 90° so the two cuts are distinguishable, with the five components of the hinge design vector $\boldsymbol{g}_h$ (7): the ligament width w_{lig} between the main cut L_{main} and the secondary cut L_{sec}, the cut angle α at the tip, the kerf w_c and the tip fillet of radius $\rho = \rho_f w_{\text{lig}}$. Teal is the hinge window of radius r_{win}, bounded by the adjacent panel faces, on which the response of Section 2.4 is solved. Panel (b) is a schematic and is not to scale.

h join panels P_k and P_l. Its two panel corners separate by $\Delta\boldsymbol{u} = \Delta\boldsymbol{x}_l(\boldsymbol{q}_l) - \Delta\boldsymbol{x}_k(\boldsymbol{q}_k)$ under the rigid-body map of Section 2.1, which already contains the rotation of each corner about its own centroid. Let $\bar{\phi} = \frac{1}{2}(\phi_k + \phi_l)$ be the mean panel rotation and $\vartheta_h = \phi_l - \phi_k$ the relative rotation. Removing the mean rotation leaves $\boldsymbol{d}_h = R(-\bar{\phi})\left(\boldsymbol{e}_h^0 + \Delta\boldsymbol{u}\right) - \boldsymbol{e}_h^0$. The scalar measures are read off it,

$$\boldsymbol{\varepsilon}_h(\boldsymbol{q}) = \left(a_h,\ s_h,\ \vartheta_h\right), \tag{8}$$

with reference bond vector $\boldsymbol{e}_h^0 = \boldsymbol{x}_l^0 - \boldsymbol{x}_k^0$. The first two are the components of $\boldsymbol{d}_h$ in the cut frame. The axial opening $a_h = \boldsymbol{d}_h \cdot \hat{\boldsymbol{e}}_{\text{sec},h}$ is its projection along the secondary cut, and the shear $s_h = \boldsymbol{d}_h \cdot \hat{\boldsymbol{e}}_{\text{sec},h}^{\perp}$ its projection transverse to that cut. Both are lengths, and ϑ_h is an angle.

Each hinge now has both arguments its mechanical behaviour requires: a design $\boldsymbol{g}_h$, fixed once the sheet is cut, and a state $\boldsymbol{\varepsilon}_h(\boldsymbol{q})$, which changes as the sheet deploys. Section 2.3 states what a function of the two must satisfy, and puts the sheet in equilibrium under an actuation.

2.3 Energy and equilibrium of the sheet

Each hinge has a stored-energy function $\mathcal{W}(\boldsymbol{\varepsilon}; \boldsymbol{g})$ of its deformation and its hinge design vector. The elastic energy of the kirigami tessellation is the sum of independent hinge contributions (A3). Any frame-invariant $\mathcal{W}$ of $\boldsymbol{\varepsilon} = (a, s, \vartheta)$, parameterised by $\boldsymbol{g}_h$, is admissible, if it meets three requirements:

$$U(\boldsymbol{q};\,\boldsymbol{\theta}) = \sum_{h\in\mathcal{H}} \mathcal{W}\big(\boldsymbol{\varepsilon}_h(\boldsymbol{q});\ \boldsymbol{g}_h\big),$$
$$\mathcal{W} \geq 0, \qquad \mathcal{W}(\mathbf{0};\boldsymbol{g}) = 0, \quad \nabla_{\boldsymbol{\varepsilon}}\mathcal{W}(\mathbf{0};\boldsymbol{g}) = \mathbf{0}, \qquad \mathcal{W} \in C^2 \text{ in } (\boldsymbol{\varepsilon}, \boldsymbol{g}). \tag{9}$$

Non-negativity keeps the assembly energy U of (9) bounded below, so the deployed state minimises a coercive potential. Vanishing value and gradient at $\boldsymbol{\varepsilon} = \mathbf{0}$ make the flat sheet a stress-free equilibrium. Twice-differentiability supplies the internal generalised force work-conjugate to (a, s, ϑ), $\nabla_{\boldsymbol{\varepsilon}}\mathcal{W} = (F_a, F_s, M_\vartheta)$, which the deployment solver balances against the applied loads, as well as the tangent stiffness $\nabla^2_{\boldsymbol{\varepsilon}}\mathcal{W}$, which assembles into the Hessian H that the adjoint of the equilibrium solve inverts in (20). Without it the deployed state can be computed but not differentiated. The third is $\partial\mathcal{W}/\partial\boldsymbol{g}$, by which a sheet-level objective reaches the hinge designs. Section 2.4 defines $\mathcal{W}$ as the response of the representative volume element, and the requirements above are what any evaluation of it must satisfy.

A configuration $\boldsymbol{q}$ of the kirigami tessellation has three degrees of freedom per panel, the two components of $\boldsymbol{u}_k$ and the rotation ϕ_k. The supports prescribe some of them and leave the rest free. The admissible set $\mathcal{C}$ is the configurations meeting the prescribed values. Actuation enters as a force $\boldsymbol{F}_k$ or moment M_k applied to a free panel.

The actuation described above and a choice of $\mathcal{W}$ together determine the deployed shape (Fig. 1c). At load factor ζ the sheet minimises its total potential energy over the admissible configurations $\mathcal{C}$. This defines the inner problem:

$$\boldsymbol{q}^\star(\zeta;\,\boldsymbol{\theta}) = \arg\min_{\boldsymbol{q}\,\in\,\mathcal{C}}\ \Pi(\boldsymbol{q};\,\boldsymbol{\theta},\zeta),$$
$$\Pi(\boldsymbol{q};\,\boldsymbol{\theta},\zeta) = \underbrace{\sum_{h\in\mathcal{H}} \mathcal{W}\big(\boldsymbol{\varepsilon}_h(\boldsymbol{q});\,\boldsymbol{g}_h\big)}_{\text{stored in the hinges}} \;-\; \underbrace{\zeta\sum_{k=1}^{n}\big(\boldsymbol{F}_k\!\cdot\!\boldsymbol{u}_k + M_k\phi_k\big)}_{\text{work of the applied loads}}. \tag{10}$$

The potential Π is the elastic energy U of (9) less the work of the external loads. The panels are rigid (A1), so the first sum is the whole of the stored energy. Stationarity of Π states that the generalised forces the hinges develop balance the applied ones. The minimisation is performed with limited-memory Broyden–Fletcher–Goldfarb–Shanno, a gradient-based quasi-Newton method that approximates the inverse Hessian from a short history of past steps rather than forming it [32]. The load factor ζ is advanced in small increments, so that the solver follows the quasi-static path (A5).

The linear hinge law and its range of validity

The standard energetic model of rotating-square kirigami, the linear hinge law, gives each hinge three uncoupled quadratic springs, one per deformation measure: k_a resists opening along the ligament, k_s resists shearing across it, and k_ϑ resists relative rotation of the two panels [30, 33]:

$$\mathcal{W}_h^{\text{lin}} = \tfrac{1}{2}k_a\, a_h^2 + \tfrac{1}{2}k_s\, s_h^2 + \tfrac{1}{2}k_\vartheta\, \vartheta_h^2. \tag{11}$$

This is the second-order truncation of $\mathcal{W}$ about the closed state, $\nabla^2_{\boldsymbol{\varepsilon}}\mathcal{W}(\mathbf{0};\boldsymbol{g})$, diagonalised by neglecting the coupling between opening, shear and rotation. Interpenetration is prevented by a contact barrier on each void angle ϕ between adjacent panel edges, identically zero for $\phi \geq \phi_{\text{cut}}$ and divergent as $\phi \to \phi_{\min}$ [30]. The three constants are the diagonal of $\nabla^2_{\boldsymbol{\varepsilon}}\mathcal{W}(\mathbf{0};\boldsymbol{g})$, condensed from the element of Section 2.4 by driving one deformation measure at a time.

2.4 Non-linear hinge response

Equation (11) accounts for neither out-of-plane buckling of the ligament, which for a sheet cut from thin flat material follows almost immediately upon opening [14], nor plasticity of the material. Both are non-linear, so $\mathcal{W}$ is computed here by a finite-element solve of one hinge, licensed by (A3) and modelled in three dimensions as a solid meshed from the hinge design vector $\boldsymbol{g}_h$ of (7).

The domain of that solve, the window, is a half-disc of radius r_{win} below the secondary cut, bounded by the two adjacent panel faces and extruded through the sheet thickness (Fig. 3). Its radius must cover the deformed ligament and the decay of that deformation into the panels, beyond which the panels are rigid (A1) and store no energy. Both scales vary with the actuation, the sheet material and the hinge design, so r_{win} is fixed once at a value covering the range considered here.

The hinge is meshed with quadratic 15-node pentahedral elements, refined about the cut tip, and solved in CalculiX 2.23 [34]. The analysis is geometrically non-linear and elasto-plastic, under a finite-strain J_2 law with isotropic hardening tabulated from the tensile response of the sheet material (Section 4.1), so out-of-plane buckling, yielding and hardening are all resolved.

The two bounding panel faces are displaced by the hinge deformation $\boldsymbol{\varepsilon}$ of (8). They are moved as rigid bodies (A1) and held coplanar (A2), under the condition derived in Section 5.1. One solve ramps $\boldsymbol{\varepsilon}$ along a straight ray from the undeformed state, and both the solve and the deployment it feeds are quasi-static (A5).

Each converged increment returns the stored energy $\int_\Omega \psi\, \mathrm{d}V$ over the window, the work-conjugate forces (F_a, F_s, M_ϑ) and the plastic dissipation. It is therefore one sample of $(\boldsymbol{\varepsilon}, \boldsymbol{g}) \mapsto (\mathcal{W}, \nabla_{\boldsymbol{\varepsilon}}\mathcal{W}, D)$. The damage D measures how much of the material's ductility has been consumed, zero in a pristine hinge and one at fracture,

$$D = \frac{\langle \varepsilon_p \rangle_\Omega}{\varepsilon_f}, \qquad \langle \varepsilon_p \rangle_\Omega = \frac{1}{|\Omega|}\int_\Omega \varepsilon_p\, \mathrm{d}V, \tag{12}$$

where ε_p is the equivalent plastic strain, ε_f the measured fracture strain and Ω the whole hinge window. Reading energy and damage as functions of the state alone is Assumption A4: the material is plastic, so both are path functionals in principle, and Section 5.2 measures the resulting error.

The solve resolves every non-linearity of the fine scale, at a cost too high for the design loop: each hinge path is a sequence of incremental Newton solves, and with no adjoint exposed by the solver, a gradient with respect to $\boldsymbol{g}_h$ could only be taken by finite differences, at one further path per component and per hinge. The condensed hinge energy below replaces it by a fitted interpolant.

The condensed hinge energy

The hinge scale enters the design loop through a surrogate. The construction is that of computational homogenisation [35]: a fine-scale boundary-value problem is solved on a representative volume element, and its response is reduced to an effective law the scale of the cut pattern evaluates. That representative volume element is the hinge window, its solves are run in advance over a family of hinge design vectors. Thus $\mathcal{W}$ represents out-of-plane buckling, yielding and a ductile damage variable.

Sweeping hinge deformations and ligament geometries produces a database of solves, all on one window geometry, the rounded half-disc defined above. Each solve kept for the fit drives one geometry $(w_{\mathrm{lig}}, \alpha)$ along a straight ray from the undeformed state to one hinge deformation (a, s, ϑ), retaining the intermediate frames, so one solve contributes many samples. Diverged solves are removed before fitting. The rows retained are split by hinge geometry into a training and validation set. Two properties of the database bear on any fit built from it. Rays from the origin do not cross, so the condensed energy is single valued on the states it is fitted to. The states the database covers are bounded, and delimit the training domain. A barrier $\mathcal{B}$ penalises states outside it.

The surrogate predicts two quantities from one set of features: the stored energy $\mathcal{W}$ and the damage variable D. Each is returned by its own output head on a shared input. The generalised forces are the gradient $\nabla_{\boldsymbol{\varepsilon}}\mathcal{W} = (F_a, F_s, M_\vartheta)$ of the first, taken by differentiating the fitted energy, so the hinge returns a force field derived from a potential rather than three independently fitted components. Those features are

$$\boldsymbol{\xi}(\boldsymbol{\varepsilon}, \boldsymbol{g}) = [\,\eta_a,\ \eta_s,\ \vartheta,\ \log w_{\mathrm{lig}},\ \alpha\,], \qquad \eta_a = \frac{a}{w_{\mathrm{lig}}}, \quad \eta_s = \frac{s}{w_{\mathrm{lig}}}, \tag{13}$$

the two translational measures entering through the dimensionless ratios η_a and η_s. The training domain is delimited in these coordinates. Let $\boldsymbol{\chi}$ embed those features in a vector space and χ_D return a scalar. The surrogate is

$$\begin{aligned} \mathcal{W}(\boldsymbol{\varepsilon}; \boldsymbol{g}) &= \mathcal{W}_s \left\| \boldsymbol{\chi}\big(\boldsymbol{\xi}(\boldsymbol{\varepsilon}, \boldsymbol{g})\big) - \boldsymbol{\chi}\big(\boldsymbol{\xi}(\mathbf{0}, \boldsymbol{g})\big) \right\|_2^2, \\ \hat{D}(\boldsymbol{\varepsilon}; \boldsymbol{g}) &= D_s \,\mathrm{softplus}\big(\chi_D\big(\boldsymbol{\xi}(\boldsymbol{\varepsilon}, \boldsymbol{g})\big)\big). \end{aligned} \tag{14}$$

The stored energy is the squared displacement of the embedding from its position at zero deformation and the same geometry. That form satisfies the three requirements of (9) identically, without any penalty term and for any weights. A squared norm is non-negative, it vanishes when its two arguments coincide, and its gradient vanishes there with it. The second head predicts the damage variable D of (12), which (17) penalises. The scales $\mathcal{W}_s$ and D_s are the root-mean-square training energy and damage.

The embedding $\boldsymbol{\chi}$ is a multilayer perceptron with tanh activations and a linear output layer: three hidden layers of 256 units and a 128-dimensional output, 166 016 parameters. The damage head χ_D is a separate two-layer network of 256 units. The five features of (13) are standardised by the training-set mean and standard deviation. Every sample includes the energy and its gradient, so the loss matches both,

$$\mathcal{L}_{\text{data}} = \gamma\, \frac{\|\mathcal{W} - \mathcal{W}^*\|^2}{\mathcal{W}_s^2} + \frac{1-\gamma}{3} \sum_{c\in\{a,s,\vartheta\}} \frac{\|F_c - F_c^*\|^2}{\sigma_{F,c}^2} + w_D\, \frac{\|\hat{D} - D^*\|^2}{D_s^2}, \tag{15}$$

in which starred quantities denote finite-element targets and F_c are the components of $\nabla_{\boldsymbol{\varepsilon}}\mathcal{W}$. The weight γ balances energy against force, $\sigma_{F,c}$ is the root-mean-square of component c over the training split, w_D is the damage weight and D_s the damage scale. A fit accurate in energy and inaccurate in force produces the wrong equilibrium [36]. The force term is normalised per component. Pooled, the rotational component would absorb almost all of it, and $\partial\mathcal{W}/\partial a$ and $\partial\mathcal{W}/\partial s$ would be effectively unsupervised. The output dimension sets the rank of the squared form in (14), and it is sized by the force rather than by the energy. A network with seventy times fewer parameters is half a point worse in energy, and its force error is a quarter larger and four times as dispersed over five training seeds. Training minimises (15) with Adam at a learning rate of 10^{-3} and a batch of 256, weighting energy and force equally. The checkpoint reported throughout is the best-validation epoch.

The charts of Section 2.1 parametrize rigidly deployable cut patterns. The potential of Section 2.3 brings a tessellation to a deployed equilibrium, and the surrogate above supplies the hinge energy that potential requires, representing out-of-plane buckling, yielding and damage in a network evaluation rather than a finite-element solve. Every one of these maps is differentiable in the design vector, so the deployed state is a differentiable function of the cut pattern and the hinge widths.

2.5 Gradient-based inverse design

The design is chosen by solving an outer minimisation problem over the design vector, in which Θ is the admissible set of (5), so every iterate the optimiser reaches is a connected, non-self-intersecting and rigidly deployable pattern. The *design objective* splits into a geometric and a physical part (Fig. 1d),

$$\begin{gathered}\boldsymbol{\theta}^\star = \arg\min_{\boldsymbol{\theta}\in\Theta}\ \mathcal{L}(\boldsymbol{\theta}),\\ \mathcal{L}(\boldsymbol{\theta}) = \mathcal{L}_{\text{geom}}\big(\mathcal{P}(\boldsymbol{q}_N^\star(\boldsymbol{\theta}));\, \mathcal{T}\big) + \mathcal{L}_{\text{phys}}\big[\{\boldsymbol{\varepsilon}_h(\boldsymbol{q}_i^\star(\boldsymbol{\theta}))\}_{i,\, h\in\mathcal{H}}\big].\end{gathered} \tag{16}$$

The first term measures the discrepancy between the deployed geometry and a prescribed target. The second measures the severity of the hinge deformations that realise it. Here $\boldsymbol{q}_i^\star(\boldsymbol{\theta}) = \boldsymbol{q}^\star(\zeta_i;\boldsymbol{\theta})$ is the equilibrium of the inner problem (10) at load step i, of which N is the last. The map $\mathcal{P}$ returns the vertices on the sheet outline, so $\mathcal{P}(\boldsymbol{q}_N^\star(\boldsymbol{\theta}))$ is the deployed point set the shape term is evaluated on. The set $\mathcal{T}$ is the target shape, represented as a point set prescribed before the run. The shape term is evaluated on the deployed state alone, whereas $\mathcal{L}_{\text{phys}}$ is evaluated on the whole path. $\mathcal{L}_{\text{geom}}$ is the bidirectional Chamfer distance [37],

$$\mathcal{L}_{\text{geom}}(\mathcal{P};\,\mathcal{T}) = \frac{1}{|\mathcal{P}|}\sum_{\boldsymbol{p}\in\mathcal{P}}\min_{\boldsymbol{t}\in\mathcal{T}}\|\boldsymbol{p}-\boldsymbol{t}\|^2 + w_{\text{cov}}\,\frac{1}{|\mathcal{T}|}\sum_{\boldsymbol{t}\in\mathcal{T}}\min_{\boldsymbol{p}\in\mathcal{P}}\|\boldsymbol{t}-\boldsymbol{p}\|^2,$$

whose first sum measures deployed material lying off the target (precision) and whose second target regions left uncovered (coverage). Neither is sufficient alone. Precision is minimised by a sheet that deploys onto a small, well-placed part of the target and leaves the rest of it empty; coverage by one that spreads over the whole target and past its boundary. The weight w_{cov} sets the balance between the two.

The physical part is a weighted sum $\mathcal{L}_{\text{phys}} = \sum_j w_j\,\Phi_j$ of functionals of the deformation path $\boldsymbol{\varepsilon}_{ih} \equiv \boldsymbol{\varepsilon}_h(\boldsymbol{q}_i^\star)$, which the incremental solve supplies at no additional cost. Plastic strain only accumulates. Damage read from the hinge deformation alone would run backwards as a hinge closes again and return less than the hinge has taken, so the damage functional reads the path maximum rather than the endpoint. Path extrema are taken smoothly. Write $\max^\tau$ and $\min^\tau$ for the softmax-weighted mean along the path at temperature τ, which stays within the range of the values it summarises. Three functionals enter, one on the damage variable and two on the translational components of $\boldsymbol{\varepsilon}_{ih}$, the latter normalised by w_{lig} as in (13),

$$\Phi_{\text{damage}}\big[\{\boldsymbol{\varepsilon}_{ih}\}\big] = \frac{1}{|\mathcal{H}|}\sum_h\Big(\max\nolimits_i^\tau\,\hat{D}_{ih}/D_s\Big)^2,\qquad \hat{D}_{ih} = \hat{D}(\boldsymbol{\varepsilon}_{ih};\,\boldsymbol{g}_h), \tag{17}$$

$$\Phi_{\text{comp}}\big[\{\boldsymbol{\varepsilon}_{ih}\}\big] = \frac{1}{|\mathcal{H}|}\sum_h\max\big(0,\ -\min\nolimits_i^\tau\,\eta_{a,ih}\big), \tag{18}$$

$$\Phi_{\text{shear}}\big[\{\boldsymbol{\varepsilon}_{ih}\}\big] = \frac{1}{|\mathcal{H}|}\sum_h\max\nolimits_i^\tau|\eta_{s,ih}|. \tag{19}$$

The first penalises predicted damage, $\hat{D}$ being the predicted value of the measure D of (12). Division by D_s, the root-mean-square D over the training split, applies the normalisation of (15) and fixes $\Phi_{\text{damage}} = O(1)$ there. It is the material's own failure measure, and no other failure quantity enters the objective. Compression is penalised because it drives the path towards that edge and because it bears on Assumption A2. Depending on the material, shearing can have a similar effect and also cause tearing of the hinge.

Computing gradients via backpropagation

The objective above is minimised by gradient descent, and its gradient has to pass through an equilibrium that the design itself determines. Algorithm 1 sets out one design step and Fig. 1 shows the same chain; the iterates are the unconstrained coordinates of (5), which parameterise the cut pattern and the hinge designs of (1).

The design gradient runs back through the whole pipeline. Every link in that chain is differentiable by construction. This covers the charts (5), the collinearity solve, and

Algorithm 1 One design step of the differentiable pipeline.

Require: unconstrained coordinates $(\boldsymbol{z}, \boldsymbol{\ell}, \boldsymbol{\omega})$, target $\mathcal{T}$, fixed process constants (ρ_f, w_c), window radius r_{win}

1: **for** $k = 0, 1, \ldots$ **do**

2: $\quad (\boldsymbol{r}, \boldsymbol{\tau}, \boldsymbol{w}_{\text{lig}}) \leftarrow \text{charts}(\boldsymbol{z}, \boldsymbol{\ell}, \boldsymbol{\omega})$ ▷ Eq. (5)

3: $\quad$ panel vertices $\leftarrow$ solve the collinearity system in $\boldsymbol{\kappa}$ ▷ Sec. 2.1

4: $\quad \boldsymbol{g}_h \leftarrow (w_{\text{lig},h},\ \alpha_h,\ \rho_f,\ w_c,\ r_{\text{win}})$ for every hinge ▷ Eq. (7)

5: $\quad \boldsymbol{q}^\star \leftarrow \arg\min_{\boldsymbol{q}} \Pi(\boldsymbol{q}; \boldsymbol{\theta})$, hinge energy from the surrogate ▷ Eqs. (10), (14)

6: $\quad \mathcal{L} \leftarrow \mathcal{L}(\boldsymbol{q}^\star, \mathcal{T})$ ▷ Eq. (16)

7: $\quad \nabla_{\boldsymbol{\theta}} \mathcal{L} \leftarrow$ reverse-mode sweep of lines 2–6, the solve differentiated by (20)

8: $\quad (\boldsymbol{z}, \boldsymbol{\ell}, \boldsymbol{\omega}) \leftarrow \text{Adam}\big((\boldsymbol{z}, \boldsymbol{\ell}, \boldsymbol{\omega}),\ \nabla_{\boldsymbol{\theta}} \mathcal{L}\big)$

9: **end for**

the hinge deformations of Section 2.2. It covers the hinge energy of Section 2.4 in its geometric argument $\boldsymbol{g}$ as well as its kinematic one $\boldsymbol{\varepsilon}$, so $\partial \mathcal{W}/\partial \boldsymbol{g}$ propagates the objective to every hinge design. The deployed state $\boldsymbol{q}^\star$ is an argument of the minimum. Differentiating the solver iterates is exact but costs memory linear in the iteration count. The implicit function theorem, applied to the stationarity condition of (10), yields the same derivative from the converged state alone,

$$\frac{\partial \boldsymbol{q}^\star}{\partial \boldsymbol{\theta}} = -\, H^{-1} \, \frac{\partial^2 \Pi}{\partial \boldsymbol{q} \, \partial \boldsymbol{\theta}}, \qquad H = \nabla^2_{\boldsymbol{q}} \, \Pi, \tag{20}$$

exposed as a custom vector–Jacobian product, line 7 of Algorithm 1. That line is what separates a differentiable physics loop from a forward one. The backward pass costs the same whatever number of iterations the solve of line 5 took. The composition is differentiable end to end (Fig. 1).

The objective $\mathcal{L}$ is minimised with Adam [38], at one value-and-gradient evaluation per iteration. The gradient is obtained by reverse-mode differentiation of the forward chain $\boldsymbol{\theta} \to \Pi \to \boldsymbol{q}^\star \to \mathcal{L}$, so one backward pass returns $\partial \mathcal{L}/\partial \boldsymbol{r}$, $\partial \mathcal{L}/\partial \boldsymbol{\tau}$ and $\partial \mathcal{L}/\partial \boldsymbol{w}_{\text{lig}}$ together. The implementation is JAX 0.6.2 in double precision [39], with JAX MD [40] assembling the hinge energies.

3 Numerical results

3.1 Validity by construction

Four approaches to one geometric inverse design problem, run at an equal computing budget, measure the effect of the constraint enforcement strategy on the outcome (Fig. 4). A square 7×7 closed sheet is deployed to a circle of radius 0.81 times the sheet side, a diameter of 1.61 times it, while the flat sheet remains a single cuttable piece. Nothing here is at a physical scale, so every length in this section and the next is a fraction of that side. In this problem, deployment is purely kinematic, and only the geometric part outer problem (16) is solved.

A tessellation is valid when three conditions hold at every step of the deployment: every hinge keeps its two panels joined, no two panels overlap, and together the panels cover the whole closed sheet. Three of the approaches optimise the panel-corner positions directly, two coordinates for each of the four corners of every panel, or 392 here, so any validity they hold has to be imposed on those coordinates, as it is in the constrained programme of Choi *et al.* [17] and in the joint optimisation of Segall *et al.*, which treats the coordinates of the flat tiling and of its deployed configuration together and holds rigidity and planarity by penalty [20]. Free panel vertices (Fig. 4a) impose none of the three. An alternating projection (b) closes the hinge gaps, snapping each pair of bonded corners back together after every optimiser step, while overlap is penalised. A tiling penalty (c) is added to both, penalising uncovered material. The fourth (d) optimises the coordinates of the construction of Section 2.1 [13], 92 here, counted as raw scalars like the 392 above and not reduced to the 60 of them that are independent, and is valid by construction. The budget is equalised at 150 optimisation epochs for every chart, from the same starting geometry, tessellation and target, each at the better of two learning rates. The comparison is at fixed budget, not at convergence.

Free panel vertices with nothing enforced leave the sheet torn, at a worst hinge gap of 7.8% of the sheet side, and overlapping. Adding projection and a global non-overlap penalty removes both defects and the sheet then does not partition its material. A tiling penalty halves the residual void without closing it, and degrades the target match (Fig. 4). The constructed parameterisation satisfies all three requirements simultaneously. Its residual void and overlap are zero, its hinge gap closes to machine precision, and it attains the best target match of the four at 0.47% of the target radius, without a projection, a penalty or a repair step.

Figure 2, on page 14, is the constructed parameterisation applied to a larger tessellation. The closed state is an $1 \times \sqrt{2}$ rectangle tiled by $16 \times 22 = 352$ rigid panels. The target is the silhouette of the Princeton University seal, 1.77 times the rectangle's height. The deployed outline attains an intersection-over-union of 0.967, at a mean distance to the target of 0.7% of the target height. The flat sheet is gapless to a relative error below 10^{-15} and void-free by construction.

The construction reaches the target. It does not follow that it reaches it by one design. If many cut patterns match the target to the same precision, then at that precision the problem is under-constrained, and the choice among them is free for a criterion the geometry does not measure.

Structure of the solution set

The construction of Section 2.1 was run two hundred times on the same 7×7 tessellation, each run from an independent random cut pattern, against one fixed circle of radius 0.76 times the sheet side, at the same budget. Every run reached the same accuracy on purely deployable sheets (Figure 5c).

The converged patterns are not one single design. Below 4% of relative outline error, they separate into distinct families (Figure 5a) that deploy at different opening angles, 41.5° and 54.7°, and shape their panels differently. Two designs from different families lie 9.40 times further apart, on average, than two designs from the same

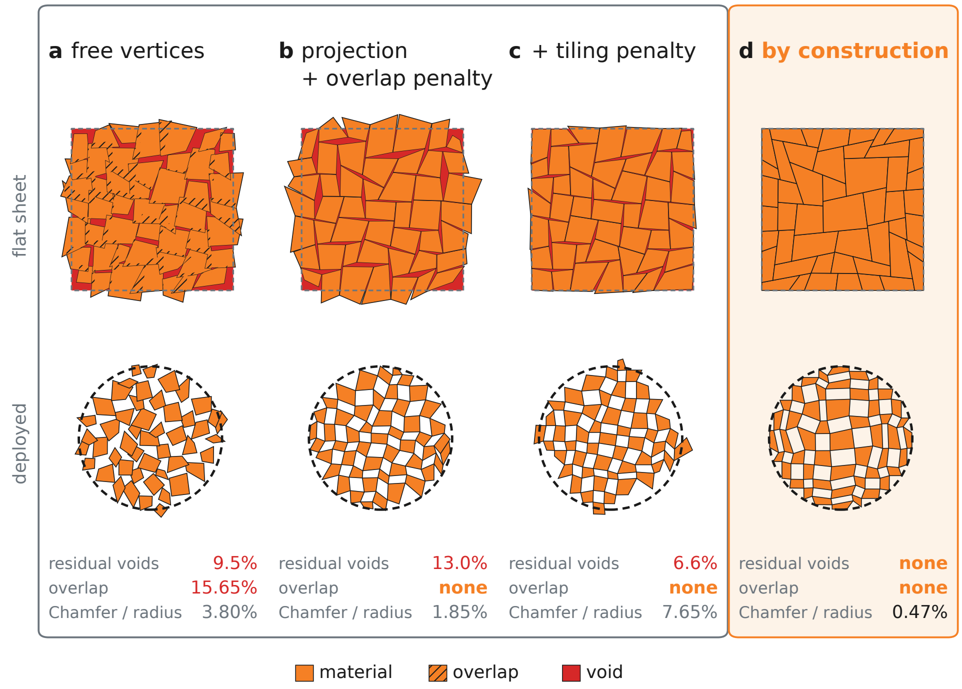

Fig. 4 Four approaches to one task, at equal budget. A square sheet of 49 panels is deployed to the dashed circle, of radius 0.81 times the sheet side. **(a)** Free panel vertices, with nothing enforced. **(b)** An alternating projection onto closed hinges, with a global non-overlap penalty. **(c)** The same with a tiling penalty. **(d)** The constructed parameterisation. All four columns share the tessellation, target, hinge model, epoch budget and starting geometry.

family. The objective is also flat in a large fraction of its directions. At a converged design the Hessian of the shape term has 37 negligible eigenvalues out of 93, one per design variable, the 92 of the cut pattern together with the opening angle (Figure 5b). A design moves along those directions with negligible change in the target match, so their number is the count of independent further requirements this problem can accommodate. This geometric target is therefore under-constrained. At a given level of target shape precision, a continuous set of cut patterns reaching the same shape remains.

A kirigami cut sheet is deployed by an actuation applied to it, so the force distribution is the quantity a design can be tuned against and the outline is not. Imposing a requirement on it uses the capacity counted above, and specifies a problem the target alone does not determine. What distinguishes the designs of Section 3.2 is the distribution over the sheet of the generalised forces the hinges develop at equilibrium.

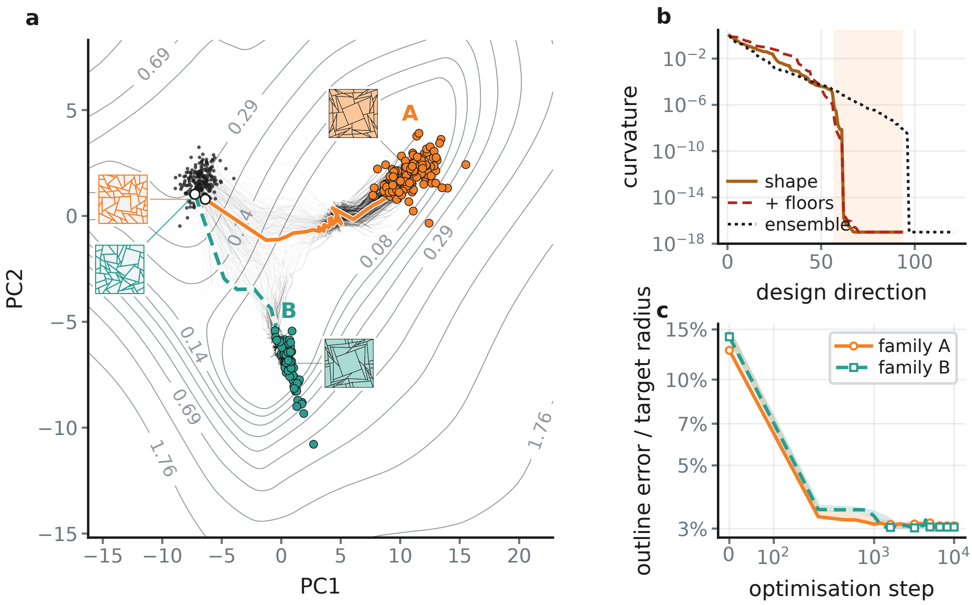

Fig. 5 Structure of the geometric solution set. Two hundred optimisations of one target, each from its own random starting pattern, with no physics. **(a)** The designs in the plane of the two leading principal components of the starting and converged patterns together. Black dots are the starting patterns, orange and teal the two families of converged designs. Grey contours are the objective evaluated on this plane. **(b)** Curvature of the objective at a converged design, by design direction. The shaded band marks the directions in which the target is flat. **(c)** Outline error against optimisation step for the two highlighted runs, over their family medians.

3.2 Internal forces under prescribed actuation

The designs above are kinematic: no force acts in them. Restoring the equilibrium inner problem of Section 2.3 and differentiating the objective through its solution makes the actuation part of the design problem. The hinge law is the linear-spring reference model (11), with the calibrated constants of Section 4.2 based on the material law of a PET sheet.

The sheet is a 15×10 closed tessellation of 150 panels joined by 275 hinges on a square panel pitch. It measures $6096 \times 4064\,\mathrm{mm}$ flat. The prescribed deployment is below the ideal per-axis deployment ratio $\lambda_{\max} = \sqrt{2}$ (6), at $\vartheta = 70°$, with $\lambda = 1.40$ on each axis and the sheet covers 1.96 times its flat area. Every run starts from the uniform rotating-squares sheet [11], at $r = \frac{1}{2}$ on every cut. The supports are a pin on one corner panel and a roller on the panel at the far end of the same edge. Case I is a concentrated axial load on the single panel nearest the sheet centre. Case II adds a transverse load at the far corner. Case III instead adds a couple on the loaded panel. Case IV is an eccentric axial load spread over three panels of the free edge. Case V is a self-equilibrated shear couple, opposed transverse tractions on the two halves of that

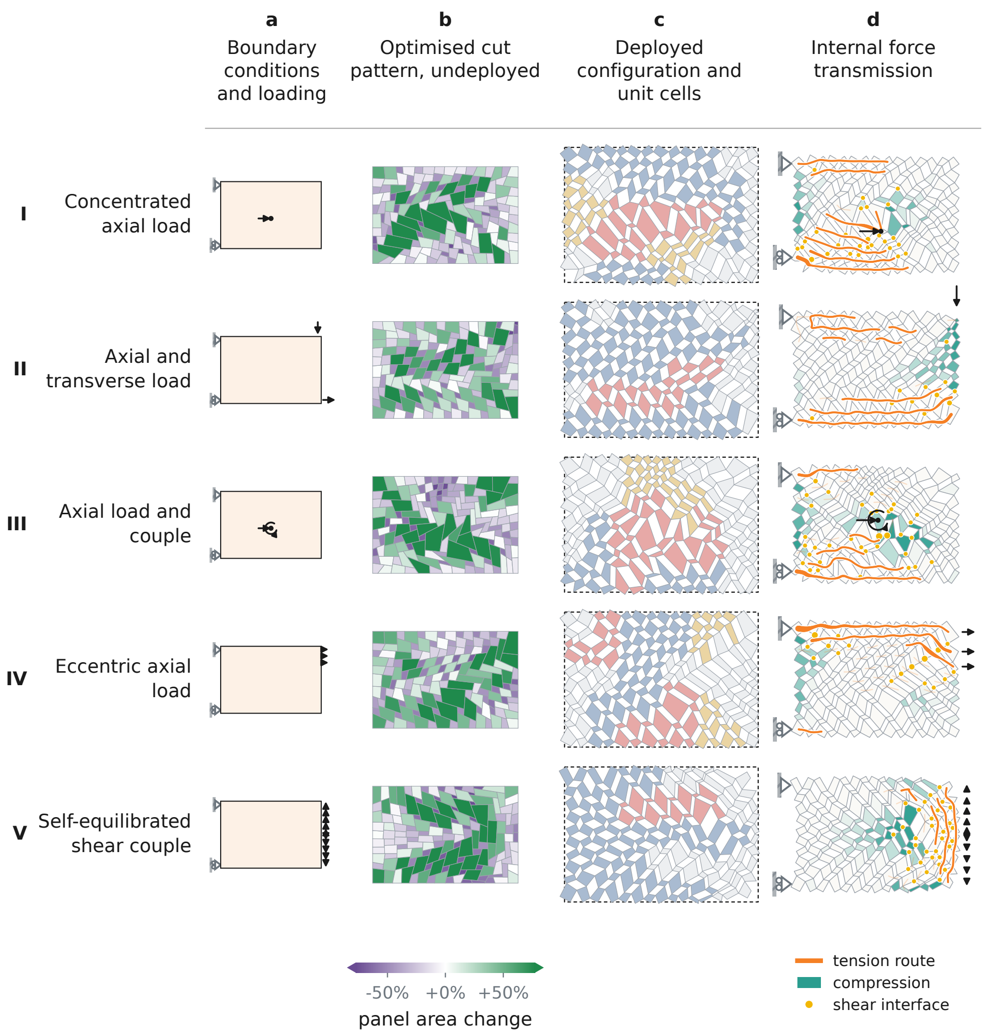

Fig. 6 Inverse design of one sheet under five prescribed mechanical actions. Each row is the same sheet, identically supported, under a different actuation. **(a)** The boundary-value problem. **(b)** The optimised cut pattern, flat and undeployed, panels coloured by their area relative to the uniform sheet. **(c)** The deployed configuration against the target, dashed. The sheet is tiled by 35 disjoint unit cells, each the four panels around one void. A unit cell is classified by that void's two side lengths, which the cut pattern fixes and deployment does not change. Pooled over all five designs these fall into three families. **(d)** The internal force transmission at that deployment, as tension routes, compressed panels and sheared interfaces, the last marked as points. Panels (b) and (c) are drawn at one scale, so the growth on deployment is the difference in size on the page. Panel (d) is drawn slightly smaller.

edge whose resultant vanishes. For each case the magnitude is tuned as the minimum required to match the same target rectangle. The couple of the third case supplies about 60% of the drive. The weights of the hinge-deformation terms in (16) are set per case so that each enters with a gradient of the same norm as the shape term. The weight that achieves that balance spans a factor 3.5 across the five cases.

All five designs converge and remain admissible, and the total panel overlap stays below 3×10^{-6} of the sheet area. The root-mean-square distance from the deployed boundary to the target falls to 1.2% of the target width (Fig. 6c). Write $\|\cdot\|$ for the root-mean-square norm over the $(M-1)(N-1)$ aspect ratios of (3) and $\boldsymbol{r}^0$ for the uniform sheet the five runs start from. Each design then lies farther from every other than from that common start, $\left\|\boldsymbol{r}-\boldsymbol{r}'\right\| > \max\left\{\|\boldsymbol{r}-\boldsymbol{r}^0\|,\ \|\boldsymbol{r}'-\boldsymbol{r}^0\|\right\}$, on all ten pairs $(\boldsymbol{r}, \boldsymbol{r}')$ of the five. Each actuation produces a different cut pattern (Fig. 6b).

The constants of (11) are a first-order fit, condensed from the same finite-element solves as the hinge energy of Section 2.4, on the sheet material those solves use. Section 4.2 evaluates that fit; here (11) is simply the working law. Each interface is resolved on its own corotated cut axis, the frame of a, s and ϑ in (8), into a normal force f_n and a tangential force f_t. The rotational spring of (11) stores almost all of the hinge energy in the four cases driven by a net pull, with axial and shear straining together under 1.5% (Table 1). Where the load cannot be carried by rotation alone, it is carried as tension, compression or shear. The non-rotational share rises monotonically with the force the interface carries, at a rank correlation of +0.99 in every case, and is largest under the self-equilibrated couple (Table 1). A strongly loaded panel is 1.9 to 2.8 times more likely than chance to lie beside another, which sets the tension routes and the compressed band already described. Tension forms straight routes (Fig. 6d). Panels rotate into alignment, and the load reaches the supports along nearly direct routes, the arc-over-chord ratio lying between 1.02 and 1.03 in all five cases. The longest route links 15 panels. The tension routes therefore behave partially as a cable net. Tension accounts for 40 to 53% of the transmitted force and shear 36 to 44%, compression carrying the remainder, whose share per case is the last column of Table 1. No tension route reaches the region between the pin and the roller. The only route across that region is compressive, so a membrane would transmit no force there [41]. An auxetic rotating-square mechanism [11] instead produces compression in this slack region [2]. The teal band in panel (d) of Fig. 6 is present in the three cases driven by a net axial pull. Compression is the most localised of the three channels. Where a couple must be transmitted alongside a net pull, the panels are larger. Case III transmits its moment through the largest panels of the five: the tenth of panels carrying the most moment average 1.61 times the median panel area, against 1.00 under the self-equilibrated couple. Those enlarged panels are the green ones in row III of Fig. 6(b). The larger panel lengthens the lever arm between its hinges, so a pair of forces transmits the moment.

The sheet's unit cells fall into three families, of mean void size 725×297, 491×360 and 411×226 mm (Fig. 6c). Each cell is a disjoint block of four panels, classified by the two side lengths of the void it encloses. The three families are the groups of a k-means partition of the logarithm of those two lengths, fitted once over the cells of all five designs so that a family denotes the same shape in every sheet. Where a unit cell

Case	Drive (N)	ϑ energy (%)	Non-rotational (%)	Compression (%)
I	495	98.6	4.3	12
II	325	99.2	2.6	11
III	211	99.5	1.6	16
IV	232	99.2	2.6	10
V	1 723	89.6	22.4	16

Table 1 Internal force transmission of the five designs of Fig. 6, deployed. ϑ energy is the share of (11) stored in rotation; non-rotational, its complement over the quarter of interfaces carrying the most force; compression, the share of transmitted force in $f_n < 0$.

family lies and where the load travels are related, but the relation changes with the actuation. Whether the recurring local geometry follows from the global aperiodic cut pattern or organises it remains open.

One sheet, held identically and matched to a common envelope, admits five distinct internal force distributions. Differentiating the objective through equilibrium reaches them. The cut pattern is selected by its mechanics and not by the outline alone.

Accuracy and load capacity of the cut pattern

The same sheet is loaded until its deployment degenerates, so that what the hinge-deformation terms of (16) return can be read as a load carried and not only as a routing. The sheet, the supports and the target rectangle are those of the five cases above. A constant opening pull of 247.7 N is held on the mid panel of the free edge. A transverse load is applied to the mid panel of the top border. Their ratio ρ is swept, each value an independent incremental solve.

Four strategies are compared, from the same start and against the same target. The uniform sheet degenerates at $\rho = 1.66$. A pattern matched to the target through the rigid-panel kinematics alone, with no equilibrium in its objective, reaches $\rho = 1.80$. Two further patterns are designed at $\rho = 1.0$ with the objective differentiated through equilibrium. Under the shape term alone the deployment degenerates at $\rho = 1.29$. With Φ_{comp} (18) and Φ_{shear} (19) added, the deployment holds to $\rho = 2.95$. In this scenario, matching the target through equilibrium halves the boundary error against the purely geometric pattern. The two hinge-deformation terms raise the capacity by a factor of 2.3 over the shape term alone, and by 64% over the geometric pattern.

The physics of a cut sheet is therefore tunable. One parameterisation, held to one target, reaches a distinct internal force distribution under each prescribed actuation, and a distinct load capacity under each objective.

Every design to this point is reached under one hinge law, the linear spring of (11), whose rotational secant is fitted at 20°. These designs open their hinges to a mean relative rotation of 55.8 to 68.8°. The construction is rigid-deployable throughout [13], so over that range it is the hinge law and not the kinematics that limits the model. Section 4 replaces it by the condensed hinge energy of Section 2.4, on the material that energy was fitted to, and takes one design under the same parameterisation through to a sheet that is cut, deployed and measured.

4 Design and validation on polyester sheet

The pipeline of Section 2 designs a cut pattern under equilibrium, and Section 3 shows that it tunes the internal forces and the load capacity of the deployed sheet. A sheet that is actually cut is governed by more than the model of Section 2 represents. Its ligaments buckle out of plane and their response softens [10, 14], and, according to the law of the material they are cut from, they yield and tear once they are narrow. The material must enter the design rather than parameterise it. This work treats sheets at metre scale, where gravity on a panel outweighs the moment the folded ligament exerts on it (Section 5.1). The panels then stay coplanar, which is Assumption A2, and the deformation is confined to the ligaments. This section runs the pipeline in full, on one adequate material.

4.1 The sheet material and its ligament

The material law is measured here from the material the kirigami design is cut from, and it enters the chain at two points: tabulated, it calibrates the finite-element solve of Section 2.4, and through its fracture strain it fixes the scale of the damage variable (12).

The sheet is a biaxially oriented thermoplastic polyester (PET) film 0.5 mm thick, chosen for its plastic behaviour, and its post-yield response is measured here. Strips of the ASTM D882 geometry [42], cut in both sheet directions, were pulled at $1\,\mathrm{mm\,min}^{-1}$ on an Instron 34SC-5. The yield stress is direction-independent within scatter, 48.4 against 49.5 MPa, so the film is treated as in-plane isotropic. Under load the film cold-draws through a propagating neck with the natural draw ratio $\Lambda = 3.28$, the area ratio of the necked section. The true stress is measured at $\Lambda\,\sigma_{\mathrm{plateau}} = 110.2\,\mathrm{MPa}$ and the plastic strain at 1.15.

The solve of Section 2.4 runs on that measurement directly. The film is elastic at $E = 3.0\,\mathrm{GPa}$, nominal for biaxially oriented film, with an assumed Poisson ratio of 0.40, and it yields at $\sigma_y = 48.8 \pm 1.6\,\mathrm{MPa}$ over ten specimens at the 0.2% offset. Hardening is tabulated as a nine-point curve in true stress and true plastic strain. It starts at the upper-yield peak of 44.0 MPa, passes through the 33.6 MPa draw plateau and the anchor above, and is extended to (200 MPa, 1.784), the fracture strain ε_f measured by running one specimen to break.

Hinge specimens were cut from the same sheet, at the hinge design vector $\boldsymbol{g} = (w_{\mathrm{lig}}, \alpha, \rho_f, w_c, r_{\mathrm{win}}) = (18\,\mathrm{mm},\ 90°,\ 0.16,\ 0.2\,\mathrm{mm},\ 100\,\mathrm{mm})$ of (7), and tested in fold and in opening.

These measurements enter the model in two ways. The material law parametrizes the finite element solve of Section 2.4. A hinge specimen sets the magnitude of a small numerical out-of-plane deviation to initiate the buckling branch at 0.5 mm, one sheet thickness. At that amplitude the simulated $u_z(\vartheta)$ tracks the measurement to 50°, with no material tuning. Both saturate, at 25.0 mm simulated against 25.5 mm measured. The measurement matches the one-parameter law $u_z(\vartheta) = u_0\sqrt{\sin\vartheta}$, with $u_0 = 25.7\,\mathrm{mm}$ and $R^2 = 0.97$.

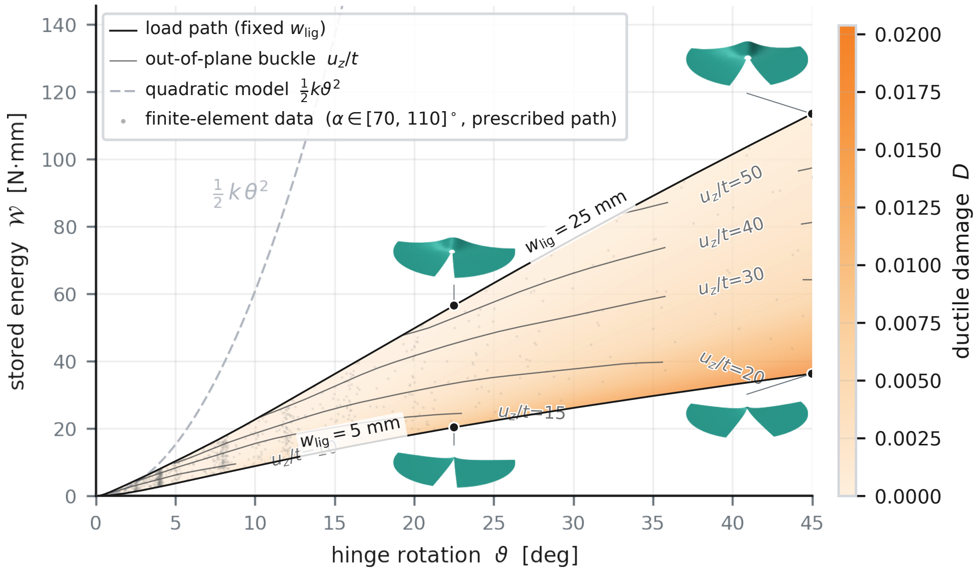

Fig. 7 The response of a single PET ligament in pure relative rotation ($a = s = 0$), learned by the surrogate (14) across the sampled range of ligament widths w_{lig}. The two bounding curves are $w_{\text{lig}} = 5$ and 25 mm. The scattered points are the 877 finite-element samples nearest this state. The grey contours are the buckle amplitude u_z/t read from the finite-element database. The renders are the CalculiX hinges of those two widths at 22.5 and 45°, at scale. Ductile damage (fill) is drawn only to report its range.

4.2 The condensed hinge energy of a polyester ligament

With the material law fixed, the hinge solve of Section 2.4 can be run on this film. The surrogate is fitted to a sweep of that solve over ligament width and cut angle. Each solve drives one geometry along a straight ray from the closed state and returns a $(\mathcal{W}, D)$ sample at every converged increment, of order ten along the path. The rays are independent and do not intersect, so a sample determines the path that reached it and the condensed energy is single valued on the set. Of order 10^3 such solves give a dataset of 12000 samples. A further set of solves, in which only the fold is imposed and the translations are left to the solver, is held out of the fit and used in Section 5.2 to test what that assumption costs. The states covered constitute the training domain: $\eta_a \in [-0.140, 0.233]$, $|\eta_s| \leq 0.220$ and $\vartheta \in [0, 71.9°]$.

Figure 7 shows that fitted energy in pure relative rotation, $a = s = 0$, where rotation is the only deformation. At any fixed geometry the energy is monotone in rotation. The buckling of the ligament shows as a falling slope. Past 40° the median deflection reaches 53.8 sheet thicknesses at ligament widths of 21 to 25 mm against 20.9 at 5 to 9, so the wider ligament deflects further and, in doing so, relieves membrane strain: at $\vartheta = 45°$ it stores 113.7 N mm against 36.4 and accumulates seventeen times less damage. Across the whole slice the surrogate stays below 0.03 and the database maximum is 0.065, against unity at fracture. A polyester hinge of this thickness does not tear

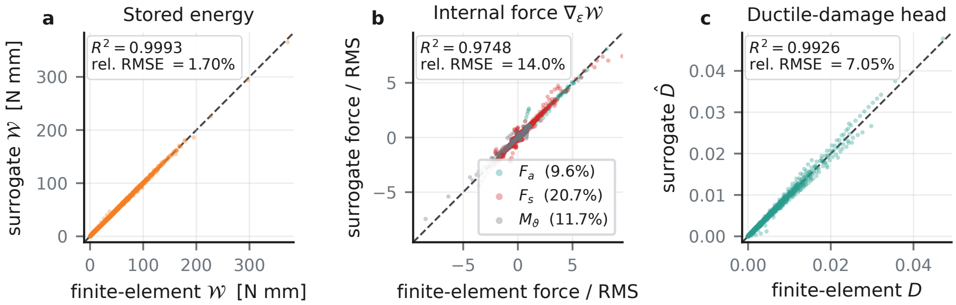

Fig. 8 Accuracy of the condensed hinge-energy surrogate against the finite-element solutions. **(a)** Stored energy. **(b)** Internal force. Each component is divided by its own root-mean-square; the legend gives per-component relative root-mean-square errors (rel. RMSE). **(c)** The ductile-damage head; no solve reached fracture, so $\hat{D}$ is reported against its own scale and not against a tear line.

by folding. The dashed parabola of Figure 7 is the linear hinge law (11), calibrated at zero rotation on the widest ligament. It diverges from the learned response as soon as that ligament buckles.

Figure 8 evaluates the fit on the independent validation set. The energy is reproduced to a relative root-mean-square error of 1.70%, median 0.97% per state; the forces, being its gradient, to 14.0%; and the damage head to 7.1%. No bin of rotation or of ligament width exceeds 3.6%. One converged window of Section 2.4 took an order of 10^3 s and returned a whole load path; one surrogate evaluation returns $\mathcal{W}$ at one state in 6 μs per hinge, a few hundred at a time. With the surrogate, the hinge energy and damage are evaluated about eight orders of magnitude faster. The learned energy landscape $\mathcal{W}(\boldsymbol{\varepsilon}; \boldsymbol{g})$ is therefore a differentiable function of the hinge deformation and of the ligament width, evaluated in place of the solve. Section 4.3 puts it in the design loop.

4.3 Inverse design with the condensed hinge energy

The condensed hinge energy is now put in the design loop, in place of the linear spring that carried every design of Section 3. It makes the deformation path of each hinge available to the objective, and the question is whether acting on that path changes the design that comes out. One problem is specified by the flat sheet and its topology, the grip, the target shape, the actuation and the material. The loop returns the design vector $\boldsymbol{\theta} = (\boldsymbol{\kappa}, \{\boldsymbol{g}_h\})$ of (1): the cut pattern coordinates and one ligament width per hinge.

The sheet is the 3×3 tessellation of nine panels and twelve hinges, cut from a single 4 ft × 8 ft sheet of that film at a 406.4 mm panel pitch. Neighbouring hinge windows do not overlap and the sheet is designed at full size. Every hinge starts at a nominal $w_{\text{lig}} = 18$ mm. The tessellation is clamped in all three degrees of freedom at one edge-centre panel and pulled collinearly at the opposite one (P_3 and P_5 in Figure 10). The

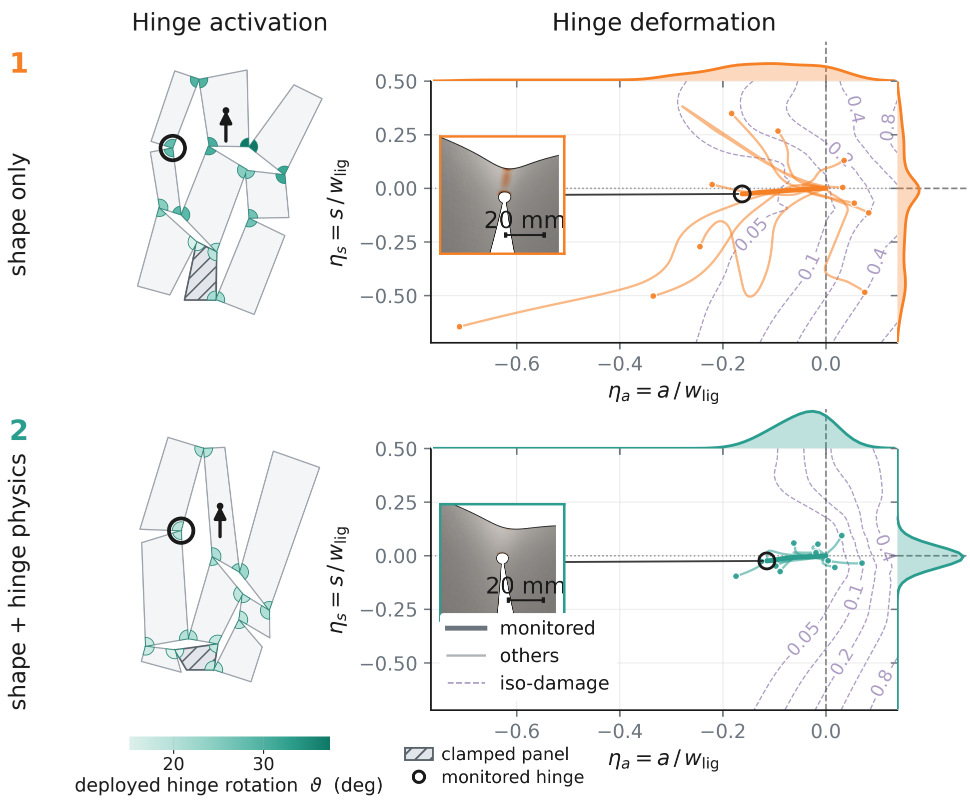

Fig. 9 The inverse problem of (Section 2.5) solved with and without the hinge-path objectives. **Row 1** uses the shape term alone; **row 2** adds Φ_{damage} (17) and Φ_{comp} (18). **Left:** the deployed sheet, showing at every hinge the window of Fig. 3 at its true radius, shaded by that hinge's deployed rotation on one scale common to the rows. The monitored hinge is ringed. **Right:** the (η_a, η_s) path of every hinge of that design, the monitored one emphasised, deployed states marked, against iso-contours of the surrogate's damage prediction in multiples of D_s, sliced at each row's own monitored geometry. The marginals are densities over path length. Insets are finite-element replays of that hinge along its own measured path, seen in plan, painted by accumulated plastic strain on one scale common to both rows.

chiral pattern swings about that centreline, so the target is oriented to the swing at 15.4°. The design load of 27.85 N is set by bisection until the worst hinge reaches 35°.

The design follows from one solve. The gradient of Section 2.5 descends through the equilibrium of Section 2.3, with the condensed hinge energy evaluated at every hinge at every step, and it returns the cut pattern coordinates and the twelve ligament widths together. That design is compared against a weaker one, obtained by removing Φ_{damage} and Φ_{comp} and leaving the shape term alone, *ceteris paribus*.

Figure 9 shows the regularised design in the second row and the purely geometric one in the first. The left column shows the panels each cut pattern produces, and they differ. The original sheet and the topology are the same in both, so the total panel area is the same. In the physical design, three panels grow, the largest by a factor of

3.2. The clamped panel and all three of its neighbours shrink, up to a fifth of their area. The panel proportions are the geometry through which the objective reaches the hinge, not a lever acting hinge by hinge. The interface forces shift with the panels. Under the same actuation the tension share of the transmitted force rises from 43 to 53%, and the compression share falls from 16.9 to 10.6%. Eight of the twelve interfaces are in tension against seven. The most compressed interface halves, from 13.9 to 6.9 N. The peak interface force is nearly unchanged, 25.8 against 23.9 N, while the mean falls from 10.2 to 7.2 N.

The right column plots every hinge in the (η_a, η_s) plane of (13). The two coordinates are the opening and the shear of the hinge, each divided by the ligament width, and both resolved in their cut frame. All twelve ligaments enter axial compression under either objective: what changes is the depth and the spread. Under the shape term alone the paths run out to $\eta_a = -0.711$ and past $|\eta_s| = 0.5$, and seven of them leave the training domain of the surrogate along the way. With the two functionals added the paths collapse into a cluster about the origin. The share of path samples inside the training domain rises from 60.2 to 95.9%. The population minimum rises fourfold, to -0.175, and the marginals narrow with it. The mean predicted damage at the path maximum falls by a factor of five, from 0.175 to 0.033 D_s. The insets of Figure 9 are finite-element renders of the deployed monitored ligament. In the shape-only case, the ligament folds into a sharp crease and reaches a peak accumulated plastic strain of 0.037. In the regularised one it bends smoothly and reaches 0.017. The two designs are read at unequal deployment, 28.0° of ligament rotation against 19.1°, so part of the reduction is less rotation rather than better routing.

Overall, the ligaments average 17.4 ± 2.4 mm under the shape term alone against 21.6 ± 0.7 mm with the two functionals added, so the regularised design's are wider and far more uniform.

4.4 Deployment and measurement of a cut sheet

The design of the preceding subsection was cut from the film characterised in Section 4.1, deployed and measured.

The sheet cut is the regularised design of Section 4.3. Cut pattern and cutting are by hand from the lattice vertices. Twelve fillet holes are drilled first, each sized to its own ligament, and each cut then runs from hole to hole, so the uncut material between two holes is the hinge. Ligament widths range from 19.9 to 22.3 mm, according to the per-hinge values the optimiser selected in Section 4.3. Panel P_3 is held in place with 15 kg weights. Panel P_5 is pierced at its centroid and pulled with a string at a speed of about $1\,\mathrm{cm} \cdot \mathrm{s}^{-1}$. The comparison is at matched opening and not at matched load.

Figure 10 shows the deployment of the sheet at three stages in (d) to (f), against the numerical predictions. The deployed photograph (a) is matched to the load–extension curve, which places it at 92.5% of the design load. Measurement is photogrammetric [43]. Tape marks every panel edge and every hinge, and a target rectangle is marked on the floor. Three photographs record the closed, part deployed and deployed states, each rectified into the millimetre frame of the design. The closed sheet is the metric standard, being a rectangle of known size cut from a squared sheet.

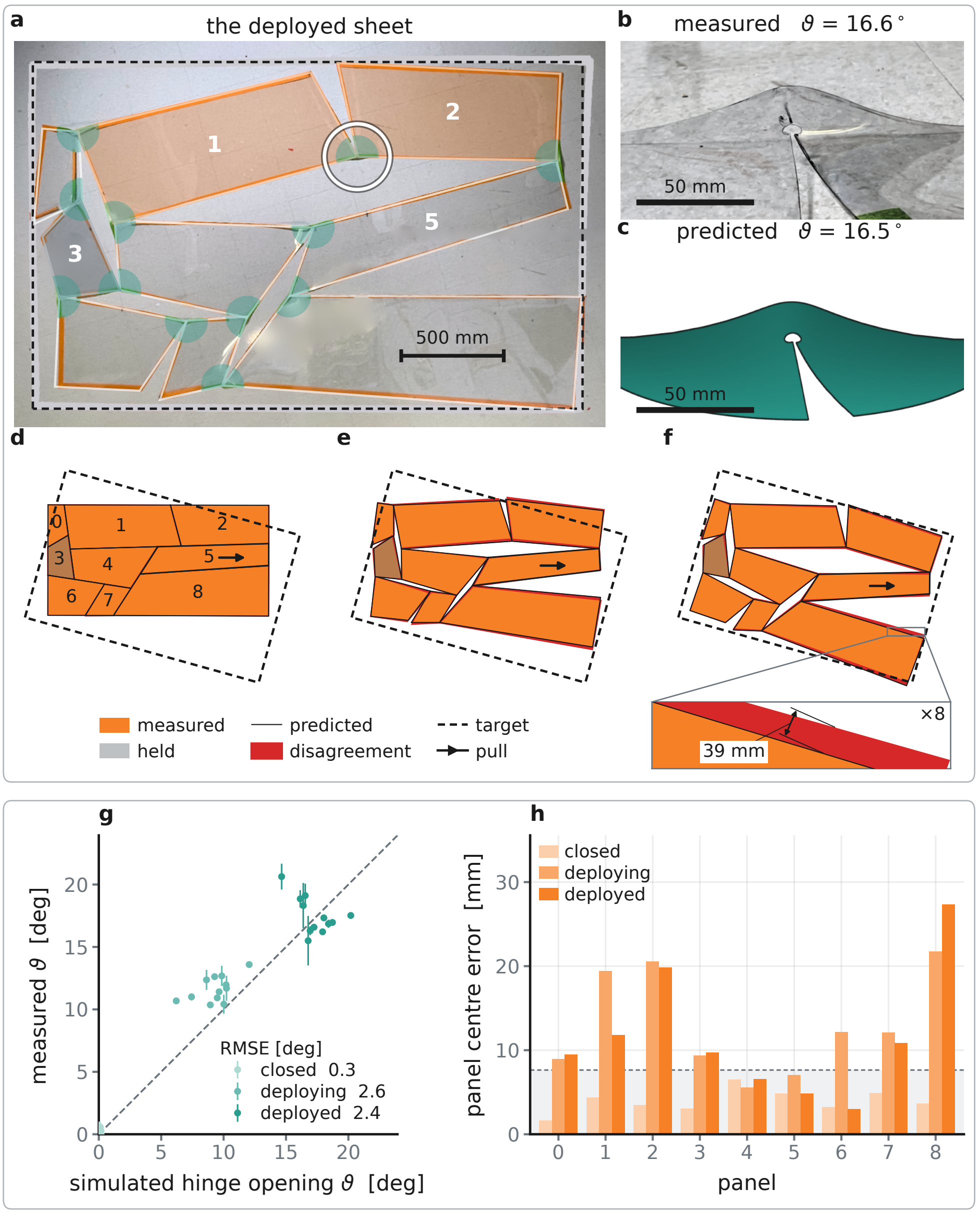

Fig. 10 A cut sheet deployed by hand, and one of its ligaments leaving the plane. **(a)** The deployed sheet photographed in plan and squared to the taped target, teal the hinge windows and the ring the ligament shown **(b)** measured and **(c)** predicted, at one shared camera and scale fixed by the relief hole both contain. **(d)** Closed, **(e)** part deployed, **(f)** deployed, the prediction drawn with the rigid-body offset between it and the measurement removed, and one region of (f) magnified beside the key. **(g)** Hinge opening measured against predicted. **(h)** Panel-centre error against the measurement floor. Panel (a) is retouched, its empty corners, a reflection, a weight and a cable reconstructed by diffusion; the homography was fitted unmodified.

Measured and predicted panel corners in (f) agree to 17.1 ± 3.6 mm root mean square, 0.70% of the 2438 mm sheet length. The rigid-body motion of the sheet is removed: the held panel P_3 slipped 26.6 mm during the experiment. Per panel centre the residual in (h) is 13.6 mm root mean square, against a measurement floor of 7.6 mm set by the closed sheet, which the model reproduces by construction. Six of the nine centres stand above that floor, from 3.0 mm at P_6 to 27.4 mm at P_8, and in the closed state none does, the largest there being 6.5 mm. The twelve hinge openings are in (g). The mean opening is 17.5° measured against 17.3° predicted, and the twelve span 15.5 to 20.6° measured against 14.7 to 20.2° predicted. Per hinge the two differ by 2.4° root mean square about a mean difference of 0.2°. One hinge of the deployed sheet was photographed alone, and panels (b) and (c) of Figure 10 compare it with a finite element prediction. The ligament joining panels 1 and 2 has buckled out of plane, reaching $u_z = 23.7$ mm, which is 47 times the 0.5 mm material thickness. It decays to under 1 mm at the rim of the window, so r_{win} does not truncate the mode. The experiment supports the locality the reduction is built on. The deformation appeared where the model places it.

Three effects lie outside the model, two of them out of plane. The sheet material was supplied rolled and the sheet kept a residual curvature, which a planar kinematics cannot represent. The pulled panel bent about its own centroid, where the load is applied at a point, while Assumption A1 holds every panel rigid. The sheet was deployed lying on the floor and every panel slid on it, whereas the model has no ground: the friction opposing that sliding is absent from the equilibrium, and the 26.6 mm slip of the weighted panel is the one measurement of it.

5 Discussion

5.1 Localisation of out-of-plane buckling in the ligaments

Assumption A2 asserts that the panels stay coplanar while the ligaments deform out of plane. In a sheet lying flat on a horizontal surface, gravity acts along the panel normals and can hold the panels coplanar. The condition under which that holds is a competition between gravity on the panels and the bending stiffness of the ligament. The sheets deployed in Section 4 are horizontal. $Y = Et$ is the stretching modulus and $B = Et^3/[12(1-\nu^2)]$ the bending stiffness of the ligament. The local fold angle [2] of the crease that forms in the ligament is θ. In the thin-ligament regime $w_{\text{lig}} \gg t$, the tension-induced instability need not tilt the panels out of plane. The panels may be held coplanar by their own weight, of areal load $p = \rho g t$. Buckling then localises in the ligaments [14].

The relative motion of two adjacent panels shares only the main-cut tip O. A planar rigid motion that leaves a point fixed is a rotation about that point. That O is fixed follows energetically. Translating it by u strains the junction, at a cost $\sim Yu^2$. Holding it fixed and folding costs only bending, $\sim B\theta^2$. With $u \sim \theta\, w_{\text{lig}}$ the former exceeds the latter by $12(1-\nu^2)(w_{\text{lig}}/t)^2 \gg 1$. Symmetry about the fold plane suppresses the remaining out-of-fold rotations. A developable fold concentrates an angle deficit $\sim 2\theta$.

[2] It is related to, but distinct from, the hinge rotation ϑ of (8).

In the interior that deficit would form a disclination requiring stretching, at a cost $\sim Y\theta^2 w_{\text{lig}}^2$. At the free edge it is absorbed by the opening of the cut at no stretching cost, through the cut-opening angle $2\theta_o$. The bulge is therefore a developable cone anchored at O [44], and its reach is what the hinge window of Section 2.4 must contain. With w_{lig} the only available transverse length, its width is $\ell_f \simeq w_{\text{lig}}$. The idealised cone's curvature discontinuities are smoothed over a bending boundary layer, whose width follows from balancing bending against stretching, $\ell_B = [B/(Y\theta^2)]^{1/2}$ [45]. The fold angle at the onset of buckling is $\theta \simeq \sqrt{2\varepsilon_c} = t/w_{\text{lig}}$, with $\varepsilon_c \simeq \frac{1}{2}(t/w_{\text{lig}})^2$ the critical strain [14]. This gives $\ell_B = w_{\text{lig}}/\sqrt{12(1-\nu^2)} \approx 0.3\, w_{\text{lig}}$. The fold is therefore a single smooth mound of finite curvature rather than a sharp crease.

Mode selection follows from a moment balance about the hinge axis of one panel of edge length L. Gravity resists tilting with $\mathcal{M}_{\text{grav}} \simeq \frac{1}{2}pL^3$, while the folded ligament drives it with $\mathcal{M}_{\text{el}} = EI_O\,\theta/\ell_f \simeq Et^4/(12\,w_{\text{lig}})$, where $I_O = w_{\text{lig}}t^3/12$. Their ratio is

$$\Gamma \equiv \frac{\mathcal{M}_{\text{grav}}}{\mathcal{M}_{\text{el}}} = \frac{6\,p\,L^3\,w_{\text{lig}}}{E\,t^4} = \frac{1}{2(1-\nu^2)}\left(\frac{L}{\ell_{\text{eg}}}\right)^3 \frac{w_{\text{lig}}}{t}, \tag{21}$$

with ℓ_{eg} the elasto-gravitational length,

$$\ell_{\text{eg}} = \left(\frac{B}{p}\right)^{1/3}, \qquad B = \frac{Et^3}{12(1-\nu^2)}, \qquad p = \rho g t,$$

equal to 0.18 m for the present sheet. The panels remain coplanar, and buckling localises in the ligaments, for $\Gamma > 1$. The condition $\Gamma = 1$ defines a critical panel edge length

$$L^\star = \left(\frac{E\,t^3}{6\,\rho\,g\,w_{\text{lig}}}\right)^{1/3}, \tag{22}$$

in which the Poisson ratio cancels. Values for several thin sheet materials are given in Table 2. The prefactor is of order unity and follows from the choices made above, so $L^\star$ should be read as a scale. For a PET sheet at $w_{\text{lig}} = 18$ mm with panels of pitch 406.4 mm, (21) gives $\Gamma \approx 3\times10^2$, and $L^\star$ is a sixth of the panel pitch (Table 2). At that pitch the condition $L > L^\star$ holds by a wide margin for paper, PET and polypropylene, and marginally for aluminium. Mild steel at 1.5 mm does not, at $\Gamma = 0.79$. The critical length therefore sets a lower bound on the panel size a given sheet material admits. Table 2 places $L^\star$ between 3 and 44 cm across those sheet materials, and it depends on the ligament width only as $w_{\text{lig}}^{-1/3}$. A pattern whose panels are smaller than $L^\star$ has $\Gamma < 1$. Its panels then tilt out of plane with the ligaments rather than staying flat, so the out-of-plane displacement is no longer confined to the hinge, although the straining still is.

5.2 Path dependence of the condensed energy

The surrogate reads the stored energy and the damage off the hinge deformation alone. The route that produced that deformation is not an input (A4). The claim can fail, because the ligament yields as it deforms: work done in one order is not work done

Material	E (GPa)	ρ (kg m^{-3})	t (mm)	$L^\star$ (cm)
Paper	3.0	800	0.2	3.0
PET (this work)	3.0	1 390	0.5	6.3
Polypropylene	1.5	910	0.8	9.3
Aluminium	70	2 700	1.0	29.0
Mild steel	210	7 850	1.5	44.0

Table 2 Critical panel edge length $L^\star$, Eq. (22), at the ligament width of this work, $w_{\mathrm{lig}} = 18\,\mathrm{mm}$, and $g = 9.81\,\mathrm{m\,s^{-2}}$.

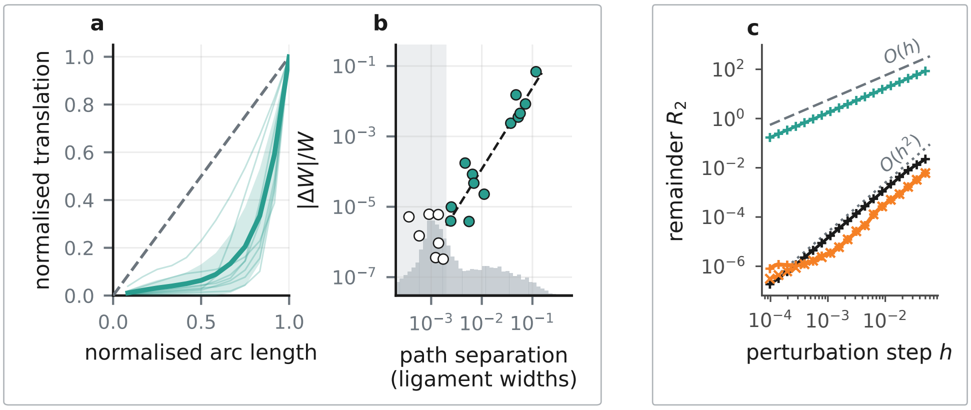

Fig. 11 **(a)** Worst eight of the twenty matched pairs of hinge routes. Each is plotted against the fraction of its journey and normalised by its own final translation (thin), with their mean and one standard deviation (thick, band). **(b)** Relative error in stored energy against the widest gap between the two legs of a pair, all twenty pairs, on logarithmic axes. Open marks over the shaded band are pairs whose two legs coincide to within the solver's resolution; the dashed line is the least-squares fit to the thirteen outside it. The histogram behind the marks is the same gap measured over every hinge route recorded from a deployed sheet. **(c)** The Taylor remainder R_2 of (23) against the perturbation step, averaged over eight random unit directions, for the constructed chart under the learned hinge law (orange) and under the linear hinge law of (11) (black), and for the free-vertex chart with the non-overlap penalty (teal); the dashed orange curve is the learned hinge law with the forward solve tightened. Grey lines are $O(h)$ and $O(h^2)$ references.

in another. The fitted database cannot expose the failure, since every route in it runs along a straight ray from the undeformed state and rays from the origin meet only there (Section 2.4).

The test was therefore built in matched pairs. Twenty routes were recorded from deployed sheets, and each was run twice through the finite-element model: once along the recorded route, once along the straight ray to the same finishing deformation. The two runs of a pair end identically and only the route between differs. Two quantities are read from each pair. The first is the separation between the legs, taken as the widest distance between them at equal fractions of the journey and measured in ligament widths. The second is the difference in stored energy and in accumulated damage. All twenty pairs are shown in Figure 11b. Panel (a) divides each route by its own final translation, so that the rays of all of them collapse onto one diagonal. Most deployed hinges ($\frac{12}{20}$) turn with negligible translation. The eight that do translate differ from their rays in the order they deform, not in direction. At the half way point the median recorded hinge has turned through 64% of its fold and covered 6% of its translation, where the ray has covered half of each, and every recorded route runs below the diagonal (Figure 11a): the translation lags the fold and arrives late.

Panel (b) measures what that lag costs. Each pair gives one point: its separation on the horizontal axis, the relative error in stored energy on the vertical. Seven pairs separate by less than 0.002 ligament widths, which is the resolution of the forward solve; their two legs are the same route to within that resolution, so the error measured across them is numerical rather than physical. Those seven are drawn open over the shaded band and left out of the fit. Over the thirteen filled marks outside it the error grows as the square of the separation, the fitted exponent of the dashed line being 2.3. A quadratic law carries no linear term, so bending a route away from its ray at fixed endpoints changes the stored energy only at second order in the separation. The histogram behind the marks gives the same separation over every hinge route of a deployed sheet, and its mass lies two decades to the left of the pairs that carry a measurable error. At the separation of a median deployed hinge the energy is in error by 0.001%; at the widest pair solved, the largest separation the training domain admits, by 6.9%, and the damage by 14.3%.

The panel-size condition of Section 5.1 and the path-separation bound above are what bound the model the optimiser descends. Whether the vector it descends is the gradient of that model is a separate question, and it is settled numerically.

5.3 Gradient verification of the design loop

The design objective is minimised through an equilibrium solve differentiated implicitly by (20). Its Taylor remainder tests whether the vector returned is the gradient of the objective implemented. Fix a design $\boldsymbol{\theta}$ and a unit direction $\boldsymbol{v}$, and write

$$R_2(h) = \left|\mathcal{L}(\boldsymbol{\theta} + h\boldsymbol{v}) - \mathcal{L}(\boldsymbol{\theta}) - h\,\langle \nabla_{\boldsymbol{\theta}}\mathcal{L}, \boldsymbol{v}\rangle\right|. \tag{23}$$

If $\nabla_{\boldsymbol{\theta}}\mathcal{L}$ is that gradient then $R_2 = O(h^2)$. The order is dimensionless, so charts of different dimension are compared on it directly, each in the parameterisation it is optimised in. Figure 11c reports the remainder over the charts of Section 3.1. Under the linear hinge law the constructed chart attains order 1.98 (black); the free-vertex chart with the validity projection alone attains 1.99, measured but not drawn. That chart needs a non-overlap penalty to stay embedded, and with the penalty active the order falls to 1.00 at every step tested (teal), the remainder never falling below 0.17. The design objective this work optimises attains order 2.05 (orange) over $1.1\times10^{-3} \le h \le 1.3\times10^{-2}$. Below that band the remainder is the difference of two objective values that agree to solver tolerance, so it measures that tolerance rather than the expansion. Tightening the forward solve lowers that floor (dashed orange); raising the iteration cap does not move it. The three charts share one adjoint, one solver, one energy and one load, and differ only in the parameterisation and in how validity is enforced, so only the penalty can account for the loss of order.

Constructing the chart, rather than penalising an unconstrained one, is therefore what keeps the design problem differentiable at the optimum the optimiser reaches. The learned hinge law attains order two as well, so condensing the elasto-plastic solve into a fitted function is not, on this evidence, a source of gradient error.

6 Conclusion

A kirigami is inverse-designed and cut from a flat sheet so that it opens onto a prescribed target under a chosen actuation. A cut pattern and the geometry of every hinge are selected by gradient descent through one differentiable chain: from a material law identified on tensile specimens, through a condensed hinge energy that carries out-of-plane buckling, yielding and ductile damage, to the equilibrium of the assembled sheet. At the pattern scale, every parameter the optimiser can reach is already a valid deployable sheet (Section 2.1); at the hinge scale, the mechanics of the ligament enters as a condensed energy, one network evaluation per hinge, fitted offline to elasto-plastic solves of a single ligament and about eight orders of magnitude faster than the solve it stands in for (Section 2.4). The consequence is control over the deployed sheet. In the loading scenario of Section 3.2, one actuation, one target, one set of supports and one starting geometry, matching the target through equilibrium rather than through the kinematics alone halves the boundary error of the deployed pattern, and penalising the hinge deformations along the deployment path raises the load the sheet carries by 64% over the same geometric pattern. To test the chain end to end, one design was cut from polyester sheet and deployed by hand. Its measured deployed shape departs from the prediction by 0.70% of the sheet length (Section 4.4).

The use that would exercise this chain fully is a kirigami sheet of structural metal: a deployable element cut flat from one sheet and opened on site. Such a sheet does not sit inside Assumption A2 the way the polyester one does. What holds the panels of a deployed sheet coplanar is gravity acting against the moment the folded ligament exerts on them, a balance measured by Γ (21), equal to $(L/L^{\star})^3$ with $L^{\star}$ the critical panel edge length (22). A structural metal is one to two orders stiffer than polyester, and its greater density does not make that up: at the 406.4 mm panel pitch used here Γ falls from 3×10^2 for the polyester sheet to 0.79 for mild steel (Table 2). The panels of such a sheet tilt out of plane with their ligaments (Section 5.1), so the out-of-plane displacement is no longer confined to the hinges, although the straining still is. What lies outside what is established here is the panel's own out-of-plane response, and it would enter as a three-dimensional parameterisation of the relative displacements.

The architecture that was developed is not specific to kirigami: the tessellation enters as a graph whose nodes are panels and whose edges are hinges, each edge carrying an energy learned from an inner problem, the assembly energy being their sum (9). Machine-learned interatomic potentials [40, 46], computational homogenisation [35] and learned closures inside conventional solvers [26] share that structure.

References

[1] Hatori, K.: History of origami in the East and the West before interfusion. In: Wang-Iverson, P., Lang, R.J., Yim, M. (eds.) Origami 5: Fifth International Meeting of Origami Science, Mathematics, and Education, pp. 3–12. CRC Press, Boca Raton, FL (2011)

[2] Dang, X., Paulino, G.H.: Kirigami engineering: The interplay between geometry and mechanics. Applied Mechanics Reviews **77**(5), 050801 (2025) https://doi.

org/10.1115/1.4068659

[3] Konaković-Luković, M., Panetta, J., Crane, K., Pauly, M.: Rapid deployment of curved surfaces via programmable auxetics. ACM Transactions on Graphics **37**(4), 106 (2018) https://doi.org/10.1145/3197517.3201373

[4] Wang, L., Chang, Y., Wu, S., Zhao, R.R., Chen, W.: Physics-aware differentiable design of magnetically actuated kirigami for shape morphing. Nature Communications **14**, 8516 (2023) https://doi.org/10.1038/s41467-023-44303-x

[5] Zhao, T., Dang, X., Manos, K., Zang, S., Mandal, J., Chen, M., Paulino, G.H.: Modular chiral origami metamaterials. Nature **640**(8060), 931–940 (2025) https://doi.org/10.1038/s41586-025-08851-0

[6] Li, S., Daqaq, M.F.: Origami-/kirigami-inspired structures: from fundamentals to applications. Philosophical Transactions of the Royal Society A **382**(2283), 20240018 (2024) https://doi.org/10.1098/rsta.2024.0018

[7] de Oliveira, I.M., Baker, E., Adriaenssens, S.: The effect of geometric tiling parameters on the stiffness of a rotational kirigami system. In: Proceedings of the IASS Symposium 2024: Redefining the Art of Structural Design, Zurich, Switzerland (2024). Paper 91

[8] Zhu, Y., Filipov, E.T.: Large-scale modular and uniformly thick origami-inspired adaptable and load-carrying structures. Nature Communications **15**, 2353 (2024) https://doi.org/10.1038/s41467-024-46667-0

[9] Shyu, T.C., Damasceno, P.F., Dodd, P.M., Lamoureux, A., Xu, L., Shlian, M., Shtein, M., Glotzer, S.C., Kotov, N.A.: A kirigami approach to engineering elasticity in nanocomposites through patterned defects. Nature Materials **14**(8), 785–789 (2015) https://doi.org/10.1038/nmat4327

[10] Isobe, M., Okumura, K.: Initial rigid response and softening transition of highly stretchable kirigami sheet materials. Scientific Reports **6**, 24758 (2016) https://doi.org/10.1038/srep24758

[11] Grima, J.N., Evans, K.E.: Auxetic behavior from rotating squares. Journal of Materials Science Letters **19**(17), 1563–1565 (2000) https://doi.org/10.1023/A:1006781224002

[12] Zheng, Y., Niloy, I., Tobasco, I., Celli, P., Plucinsky, P.: Modelling planar kirigami metamaterials as generalized elastic continua. Proceedings of the Royal Society A **479**(2272), 20220665 (2023) https://doi.org/10.1098/rspa.2022.0665

[13] Dang, X., Feng, F., Duan, H., Wang, J.: Theorem for the design of deployable kirigami tessellations with different topologies. Physical Review E **104**(5), 055006 (2021) https://doi.org/10.1103/PhysRevE.104.055006

[14] Rafsanjani, A., Bertoldi, K.: Buckling-induced kirigami. Physical Review Letters **118**(8), 084301 (2017) https://doi.org/10.1103/PhysRevLett.118.084301

[15] Yazdani, M., Shalileh, S., Shahriari, D.: Reinforcement learning for inverse silhouette design and rapid laser cutting of compact parallelogram quad kirigami prototypes. Materials & Design **268**, 116545 (2026) https://doi.org/10.1016/j.matdes.2026.116545

[16] de Oliveira, I.M., Sosa, E.M., Baker, E., Adriaenssens, S.: Experimental and numerical investigation of a rotational kirigami system. Thin-Walled Structures **192**, 111123 (2023) https://doi.org/10.1016/j.tws.2023.111123

[17] Choi, G.P.T., Dudte, L.H., Mahadevan, L.: Programming shape using kirigami tessellations. Nature Materials **18**(9), 999–1004 (2019) https://doi.org/10.1038/s41563-019-0452-y

[18] Choi, G.P.T., Dudte, L.H., Mahadevan, L.: Compact reconfigurable kirigami. Physical Review Research **3**(4), 043030 (2021) https://doi.org/10.1103/PhysRevResearch.3.043030

[19] Dudte, L.H., Choi, G.P.T., Becker, K.P., Mahadevan, L.: An additive framework for kirigami design. Nature Computational Science **3**(5), 443–454 (2023) https://doi.org/10.1038/s43588-023-00448-9

[20] Segall, A., Ren, J., Padilla, M., Sorkine-Hornung, O.: Reconfigurable hinged kirigami tessellations. In: Proceedings of the SIGGRAPH Asia 2025 Conference Papers, pp. 1–11. ACM, Hong Kong (2025). https://doi.org/10.1145/3757377.3763895

[21] Khawale, R.P., Mehdizadeh, E., Brigham, J., Filipov, E.T.: Rapid kirigami simulation using the bar and hinge approach. arXiv:2608.13770 (2026). https://arxiv.org/abs/2608.13770

[22] Deng, B., Zareei, A., Ding, X., Weaver, J.C., Rycroft, C.H., Bertoldi, K.: Inverse design of mechanical metamaterials with target nonlinear response via a neural accelerated evolution strategy. Advanced Materials **34**(41), 2206238 (2022) https://doi.org/10.1002/adma.202206238

[23] Adriaenssens, S., Block, P., Veenendaal, D., Williams, C. (eds.): Shell Structures for Architecture: Form Finding and Optimization. Routledge, London (2014). https://doi.org/10.4324/9781315849270

[24] Newbury, R., Collins, J., He, K., Pan, J., Posner, I., Howard, D., Cosgun, A.: A review of differentiable simulators. IEEE Access **12**, 97581–97604 (2024) https://doi.org/10.1109/ACCESS.2024.3425448

[25] Xue, T., Liao, S., Gan, Z., Park, C., Xie, X., Liu, W.K., Cao, J.: JAX-FEM:

A differentiable GPU-accelerated 3D finite element solver for automatic inverse design and mechanistic data science. Computer Physics Communications **291**, 108802 (2023) https://doi.org/10.1016/j.cpc.2023.108802

[26] Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P., Hoyer, S.: Machine learning-accelerated computational fluid dynamics. Proceedings of the National Academy of Sciences **118**(21), 2101784118 (2021) https://doi.org/10.1073/pnas.2101784118

[27] Schek, H.-J.: The force density method for form finding and computation of general networks. Computer Methods in Applied Mechanics and Engineering **3**(1), 115–134 (1974) https://doi.org/10.1016/0045-7825(74)90045-0

[28] Pastrana, R., Oktay, D., Adams, R.P., Adriaenssens, S.: JAX FDM: A differentiable solver for inverse form-finding. In: Differentiable Almost Everything Workshop of the 40th International Conference on Machine Learning (2023). arXiv:2307.12407

[29] Pastrana, R., Medina, E., de Oliveira, I.M., Adriaenssens, S., Adams, R.P.: Real-time design of architectural structures with differentiable mechanics and neural networks. In: The Thirteenth International Conference on Learning Representations (ICLR) (2025). arXiv:2409.02606

[30] Bordiga, G., Medina, E., Jafarzadeh, S., Bösch, C., Adams, R.P., Tournat, V., Bertoldi, K.: Automated discovery of reprogrammable nonlinear dynamic metamaterials. Nature Materials **23**(11), 1486–1494 (2024) https://doi.org/10.1038/s41563-024-02008-6

[31] Tutte, W.T.: How to draw a graph. Proceedings of the London Mathematical Society **s3-13**(1), 743–767 (1963) https://doi.org/10.1112/plms/s3-13.1.743

[32] Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. Mathematical Programming **45**(1–3), 503–528 (1989) https://doi.org/10.1007/BF01589116

[33] Coulais, C., Kettenis, C., Hecke, M.: A characteristic length scale causes anomalous size effects and boundary programmability in mechanical metamaterials. Nature Physics **14**(1), 40–44 (2018) https://doi.org/10.1038/nphys4269

[34] Dhondt, G.: The Finite Element Method for Three-Dimensional Thermomechanical Applications. John Wiley & Sons, Chichester (2004). https://doi.org/10.1002/0470021217 . Software version 2.23

[35] Le, B.A., Yvonnet, J., He, Q.-C.: Computational homogenization of nonlinear elastic materials using neural networks. International Journal for Numerical Methods in Engineering **104**(12), 1061–1084 (2015) https://doi.org/10.1002/nme.4953

[36] Vlassis, N.N., Sun, W.: Sobolev training of thermodynamic-informed neural networks for interpretable elasto-plasticity models with level set hardening. Computer Methods in Applied Mechanics and Engineering **377**, 113695 (2021) https://doi.org/10.1016/j.cma.2021.113695

[37] Fan, H., Su, H., Guibas, L.: A point set generation network for 3D object reconstruction from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2463–2471 (2017). https://doi.org/10.1109/CVPR.2017.264

[38] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations (ICLR) (2015)

[39] Bradbury, J., Frostig, R., Hawkins, P., Johnson, M.J., Katariya, Y., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., Zhang, Q.: JAX: composable transformations of Python+NumPy programs. Software, version 0.6.2 (2018). http://github.com/jax-ml/jax

[40] Schoenholz, S.S., Cubuk, E.D.: JAX, M.D. A framework for differentiable physics. Journal of Statistical Mechanics: Theory and Experiment **2021**(12), 124016 (2021) https://doi.org/10.1088/1742-5468/ac3ae9

[41] Pipkin, A.C.: The relaxed energy density for isotropic elastic membranes. IMA Journal of Applied Mathematics **36**(1), 85–99 (1986) https://doi.org/10.1093/imamat/36.1.85

[42] ASTM International: Standard Test Method for Tensile Properties of Thin Plastic Sheeting. ASTM D882, West Conshohocken, PA (2018). https://doi.org/10.1520/D0882-18

[43] Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision, 2nd edn. Cambridge University Press, Cambridge (2004). https://doi.org/10.1017/CBO9780511811685

[44] Cerda, E., Mahadevan, L.: Conical surfaces and crescent singularities in crumpled sheets. Physical Review Letters **80**(11), 2358–2361 (1998) https://doi.org/10.1103/PhysRevLett.80.2358

[45] Liang, T., Witten, T.A.: Crescent singularities in crumpled sheets. Physical Review E **71**(1), 016612 (2005) https://doi.org/10.1103/PhysRevE.71.016612

[46] Behler, J., Parrinello, M.: Generalized neural-network representation of high-dimensional potential-energy surfaces. Physical Review Letters **98**(14), 146401 (2007) https://doi.org/10.1103/PhysRevLett.98.146401

[47] Dang, X., Feng, F., Duan, H., Wang, J.: Theorem on the compatibility of spherical kirigami tessellations. Physical Review Letters **128**(3), 035501 (2022) https://

doi.org/10.1103/PhysRevLett.128.035501

Appendix A Deployability of the admissible set

The flat cut pattern of Section 2.1 is constructed after Dang *et al.* [13], whose criterion degenerates a compatibility theorem for spherical kirigami to the plane [47]. This appendix proves that every design the chart (5) admits is rigidly deployable, and that no deployable pattern of the same topology and outline is excluded.

The void of an interior cut is enclosed by the four hinges of its unit cell: a quadrilateral with a hinge at each corner, flat in the closed sheet and opening as the panels turn. Its corner angles are the void angles the contact barrier of Section 2.3 keeps positive. Deployability is a condition on void shape.

THE PARALLELOGRAM CRITERION (Dang *et al.* [13]). *A quadrilateral kirigami sheet of $M \times N$ panels with $M, N \geq 3$ and no interior hole deploys if and only if every void is a parallelogram, that is, if the four hinge-to-hinge distances around each void are equal in opposite pairs.*

It constrains side lengths alone, so it can be read off the closed sheet. The restriction $M, N \geq 3$ is not a technicality: a sheet two panels wide deploys whatever the shape of its voids.

Dang *et al.* prove it with a transfer map, which the collinearity system of Section 2.1 inherits. It suffices to treat three panels by three, whose four voids form one cycle around the central panel; any larger sheet without interior holes is covered by such blocks. In the closed sheet a void is a flat segment, so its side lengths satisfy $a+b = c+d$ and its two interior corners keep the same total distance to the two extreme ones, travelling on one confocal ellipse as the void opens. A quadrilateral with fixed side lengths has a single freedom, its opening. Two voids meeting at a panel corner share a rigid panel angle, so the opening of one fixes the opening of the next through a map $\mathcal{F}_i$, and the sheet deploys exactly when composing the four around the cycle returns an opening unchanged over an interval rather than at one configuration:

$$\mathcal{F}_1 \circ \mathcal{F}_2 \circ \mathcal{F}_3 \circ \mathcal{F}_4(\cos\beta_1) \equiv \cos\beta_1. \tag{A1}$$

Each $\mathcal{F}_i$ is the identity when its void is a parallelogram, and has $\mathcal{F}_i' > 0$ with $\mathcal{F}_i'' > 0$ when it is not, so the composition departs from the identity as soon as one void ceases to be a parallelogram: (A1) may then hold at isolated configurations, but a motion is what it requires, and the sheet is rigid.

The collinearity system of Section 2.1 is the criterion written out in the cut endpoints. The aspect ratio r_{ij} is by definition the fraction $a/(a+b)$ of the two side lengths of the parallelogram void at C_{ij}, and eliminating the void in favour of the cut endpoints turns that definition into (3). Every $\boldsymbol{r} \in (0,1)^{(M-1)(N-1)}$ and every ordered $\boldsymbol{\tau}$ produce parallelogram voids.

The converse is read off the same definition. In a sheet of this topology and outline that deploys, the criterion makes every void a parallelogram, r_{ij} is read off its two side lengths and the boundary vertices off the sheet; those values return the sheet through the collinearity system, whose solution is unique, so the chart reaches it. It does not reach a sheet of another outline, since (5) moves each boundary vertex along one edge and holds the corners while [13] leaves it free in the plane, nor one with an interior hole, for which the criterion is stated on a deployed state.