


INSEF PROJECT

Origami-Folded Sensor Geometries for Cylindrical Tracking Detectors

A Geant4 comparison of Miura-ori, Yoshimura, and Kresling fold patterns against a flat-tiled barrel control — evaluated across 500,000 simulated 5 GeV/c π^+ tracks per configuration.

Kanishk Sharma

Air Force Bal Bharti School





Abstract & Novelty

In a field which is quickly running out of firsts, this is a genuinely novel study.

Cylindrical tracking detectors are conventionally built from flat silicon tiles, forcing a trade-off between dead zones and material-adding overlaps. This project presents the first systematic Geant4 comparison of three cylindrically deployable origami fold families — Miura-ori, Yoshimura, and Kresling — against a matched flat-tiled barrel control, using a differentiated rigid-silicon / flexible-Kapton material stack.

Across 500,000 primary events per configuration, all three fold geometries increase geometric hit acceptance relative to the flat barrel, at the cost of a proportional increase in radiation-length traversal — forming a genuine three-way Pareto trade-off rather than a single optimal design.



WHAT'S NEW HERE

THE NOVELTY

- First systematic comparison**
No prior work evaluates origami-folded shells against a matched flat-tile control with full particle transport.
- Physically realistic stack**
Rigid silicon facets on a continuous flexible Kapton backing — not one idealized shell.

WHY YOU CAN TRUST IT

- Independent physics check**
Multiple-scattering output validated against Highland and Rutherford, not just self-consistency.
- Statistical robustness**
5 replicated runs, Welch's t-test with Holm–Bonferroni, sweeps over facet count, angle, vertex position.

Why Flat Tiles Fail on Cylinders

Silicon is brittle and cannot bend through a sharp crease and so tiling a cylinder with flat modules leaves gaps or forces overlaps that add material.

ALICE ITS3

bends ultra-thin (40–50 μm) CMOS silicon into seamless half-cylinders — no joints, but the silicon itself must be extremely thin.

LHCb Upstream Tracker


pairs rigid tiles with a flexible Kapton substrate — but only for electrical routing, not for changing shape.

CMS Phase-2 Outer Tracker


tilts flat modules along the beam axis to close gaps — the modules stay flat and planar.

No prior work bends the tiling pattern itself using origami creases. That is the gap this project addresses.


FOUR MATCHED-ENVELOPE GEOMETRIES




Barrel (control)



Miura-ori



Yoshimura



Kresling

Two Terms People Misread

Two words carry almost the entire paper's argument, and both are easy to mishear from everyday English. Neither means what it sounds like.

ACCEPTANCE

$$A = \text{hits} / \text{primaries fired}$$

Fire 500,000 pions out from the center in every direction. Acceptance is just the fraction that leave at least one hit in the silicon. It's a coverage number, not a quality number — it says nothing about how precisely a hit is measured, only whether a track was seen at all.

In this paper: barrel 83.2% vs. Kresling 85.5% — a +2.36 pp gap, meaning about 1 in 43 more pions get seen at all.

MATERIAL BUDGET (X / X_0)

It is **not money**. It is the amount of material a particles faces.

A radiation length (X_0) is how far a particle travels, on average, before the material scrambles its path (via scattering) or its energy. X/X_0 is a track's path length in units of that scale. More material crossed → more deflection → the worse a tracker can reconstruct where a track came from.

In this paper: $\approx 0.4\%$ of one X_0 at 300 μm Si — a thin sliver. "Budget" means how much of that sliver each geometry spends, not a price tag.

The Main Problem

Folding a flat sensor onto a cylinder can catch more tracks (increasing acceptance) — but every fold also adds a facet a track can cross edge-on, or a seam where sensor overlaps sensor, which adds path length (increasing material budget i.e. X_0). You cannot raise one without asking what happened to the other.

Pseudorapidity (η)


GEOMETRY OF THE BEAM

A stand-in for angle from the beam axis, chosen because particle rates from an isotropic source come out roughly flat in η — not in θ . Equal slices of η collect roughly equal statistics, which is why physicists bin in η instead of θ .

Why not just use θ ? $\theta = 90^\circ$ (sideways) is over-sampled by an isotropic source; θ near $0^\circ/180^\circ$ (forward) is under-sampled. η rescales the axis so this bias disappears.

IN THIS PAPER

$|\eta| \leq 0.80$ marks the central belt of the barrel ($\theta \in [48.4^\circ, 131.6^\circ]$). **94–99%** of every fold's acceptance gain comes from outside that belt — the barrel's grazing-angle ends, not its middle.



Geometry & Material Budget

All four envelopes share the same outer radius and length, but their facets don't trace the same surface area as the smooth cylinder ($A_{\text{cyl}} = 301.59 \text{ cm}^2$) — that area ratio is the geometric root of every material-cost difference.

Geometry	Segmentation	Fold parameter	A/A_{cyl}	Si mass (g)
Barrel (fine)	6 × 141 cols	—	0.9999	21.07
Miura-ori	6 × 13 cols	$\alpha = 45^\circ$	0.9722	20.49
Kresling	6 layers × 6 sides	$\theta = 30^\circ/\text{layer}$	1.0360	21.83
Yoshimura	6 rings × 8 facets	$\theta = 30^\circ$	1.1102	23.39

WHY THE MATERIAL COST DIFFERS

Miura-ori	inscribed facets sit slightly inside the envelope (97.2% of A_{cyl}) — less silicon than the barrel itself.
Yoshimura	diamond facets bulge outward, adding +11.0% surface area — more area intercepts more tracks and adds material per track.
Kresling	chiral triangulated facets add a moderate +3.6% surface area — intermediate material cost between the other two.

Knobs & Locks: Constraints and Controls

The barrel exists to answer one question: how much of a fold's behavior comes from folding, versus from simply being a cylinder at that radius? Without a matched control, a fold's under-coverage or radiation length has no reference point.

HELD IDENTICAL ACROSS ALL FOUR GEOMETRIES

Radius / length	R = 40 mm, L = 120 mm
Material stack	300 μm Si, 50 μm Kapton
Particle beam	500,000 × 5 GeV/c π ⁺
Vertex spread	3 mm 3-D Gaussian
Scoring	primaries only (parentID = 0)

THE ONE VARIABLE THAT CHANGES

The presence and shape of the fold lines.

Every acceptance gain and every radiation-length increase reported is therefore attributable to fold topology alone.

A tile-count sensitivity check confirmed this matters: coarse faceting (24 cols) vs. fine faceting (141 cols) shifted acceptance by +0.51 pp on its own — so the fine-tiled barrel, not an arbitrary tiling, was fixed as the baseline throughout.

Simulation Pipeline

01

Parametric geometry

4 cylindrical envelopes — barrel control plus Miura-ori, Yoshimura, Kresling — generated in Python.

02

GDML export

Watertight tessellated solids: rigid Si facets (per-facet normal offset) plus continuous Kapton backing (per-vertex area-weighted offset).

03

Geant4 simulation

500,000 primary 5 GeV/c π^+ tracks per geometry, isotropic 4π source, 3 mm 3-D Gaussian vertex smear.

04


Step-level scoring

Hits, 3-D path length, X/X_0 , local incidence angle — primaries only (parentID = 0).

05

Statistics

5 replicated runs per geometry, Welch's t-test, Holm–Bonferroni correction, Cohen's h effect size.



Simulated 5 GeV/c π^+ event — primaries (blue), hadronic secondaries (red). Only primaries are scored for acceptance.

Headline Performance

KRESLING

+2.36 pp

acceptance gain over barrel · +5.69% X₀

YOSHIMURA

+1.30 pp

acceptance gain over barrel · +8.01% X₀

MIURA-ORI

+0.49 pp


acceptance gain over barrel · +1.88% X₀

Comparison	ΔA (pp)	95% CI (pp)	Cohen's h
Barrel vs. Kresling	+2.362	[+2.326, +2.398]	+0.065
Barrel vs. Yoshimura	+1.304	[+1.249, +1.358]	+0.035
Barrel vs. Miura-ori	+0.493	[+0.452, +0.535]	+0.013

All six pairwise comparisons significant at $p < 10^{-7}$ (Welch's t-test, Holm–Bonferroni corrected); run-to-run SD ≤ 0.05 pp, an order of magnitude below the inter-geometry gaps.


Headline Performance — The Pareto Plots

Three-Objective Trade-off Matrix (Acceptance, X_0 , Relative Under-Coverage)



Kresling maximizes acceptance; Miura-ori minimizes material and coverage gaps; Yoshimura is dominated by Kresling on every axis.

Reading The Trade-off: No Free Lunch



HOW TO READ THIS

Up and left is better.

More acceptance for less material — no geometry sits there. That's the point.

The dashed line is the frontier.

Barrel, Miura-ori, and Kresling each improve on at least one axis over the last — none is beaten on both.

Yoshimura is the exception.



It sits strictly worse than Kresling on both acceptance and material — the only geometry here that's actually dominated.

So the choice depends on what you need.

Acceptance-limited work picks Kresling. Material-limited, resolution-sensitive work picks Miura-ori.



SECTION 4.5 · TABLE 5 · FIGURES 10-13

Spatial Structure Of Coverage — Miura-ori





Smooth, quasi-uniform hit density ($R \approx 1$) across the full azimuthal and axial range.

Spatial Structure Of Coverage — Yoshimura



Sharp periodic diamond bands — overlap seams alternate with under-dense valleys.


Spatial Structure Of Coverage — Kresling




Helical bands with localized material peaks at fold vertices.

Under-coverage Is Not Blindness

Relative under-coverage fraction vs. threshold, per geometry



Relative under-coverage response, 2D (mean $\pm$ SEM across runs)



Relative under-coverage fraction swept across detection threshold R ($0 \rightarrow 1$) for all three fold geometries.

True dead zones are (θ, z) bins with zero recorded hits — equivalent to threshold $R \rightarrow 0$.


For all three folds, $f_{\text{zero}} = 0.000\%$ — no origami geometry contains any truly blind region.

Non-zero under-coverage first appears only at $R \approx 0.43$ (Miura-ori), 0.31 (Kresling), 0.26 (Yoshimura) — all well above zero.

What f_{under} counts are relative efficiency dips near fold seams, which reconstruction algorithms can compensate for — unlike a physical inter-tile gap.


Twisting the Knobs: Facet-count (N) & Fold-angle (θ) Sweeps

Fold-Count (N) Sweep — Acceptance and X_0 vs. N



Acceptance gains asymptotically converge toward the flat-barrel limit as facet count N increases — Kresling: +2.31 pp at N=6 $\rightarrow$ $\sim$ 0 pp at N=32.

Fold-Angle (θ) Sweep - Acceptance and X_0 vs. θ



Sweeping the fold angle: Miura-ori and Kresling stay nearly flat (<0.15 pp variation); Yoshimura grows steeply — acceptance 0.845 $\rightarrow$ 0.942 while X_0 triples.

Catching Degenerate Configurations

Two points in the parameter sweeps aren't just weak — they're geometrically trivial. Rather than average them in silently, they're marked excluded, and shown here.


Yoshimura, $N = 6$, $\theta = 30^\circ$

Rotational periodicity cancels exactly — the fold collapses into a plain tapered hexagonal frustum. No diamond pattern remains.





Kresling, $N = 12$, $\theta_{\text{twist}} = 30^\circ$

The twist angle cancels across facets — the structure collapses into an untwisted, faceted cylinder. No chirality remains.

This is the same discipline as the scattering-angle refit on the previous slides: when a data point doesn't behave as expected, the right response is to check whether the model or method is producing an artifact — not to include it uncritically, and not to drop it without saying so.



Testing Sensitivity: Vertex Smear





Performance is entirely stable for Gaussian vertex smear σ up to 10 mm ($|\Delta A| < 0.2$ pp) — well beyond the nominal 3 mm used throughout.

At $\sigma = 20$ mm, the beam spot becomes comparable to the 40 mm detector radius, causing substantial off-center track leakage and a sharp acceptance drop for every geometry.

How Far Is Too Far? Testing Vertex Displacement

Radial vertex displacement shows the same pattern: stable acceptance and path length until displacement approaches the detector radius.

Experiment 3: outcome vs. realized per-event vertex displacement (pooled across all sigma points)



A Puzzling Result

The first check of multiple scattering against the Highland formula gave a number with no obvious explanation attached to it:

$$\theta_{0, \text{sim}} / \theta_{0, \text{Highland}} = 0.537 \pm 0.008$$

across all 20 momentum–geometry configurations tested

Two explanations were on the table.


A bug in the pipeline

— some step silently under-scattering tracks.

The wrong statistical model

— the fitting method doesn't match the physics of an isotropic 4π source. This was it.





Particles are fired isotropically over 4π solid angle, so the recorded scattering angle is a 3-D space angle — it cannot be Gaussian. On a log-log plot, the simulated distribution tracks Highland's prediction at small angle, then flattens into a power-law tail extending more than two decades past where Highland has already dropped by orders of magnitude. An unconstrained Gaussian fit is pulled wide by that tail.



SECTION 2.6

The shortfall traced to the fitting method — not to a failed simulation run.

Multiple-Scattering Fits



- Refit as a **Rayleigh core + Rutherford tail** mixture: the core matches Highland to within $\sim 7\%$ ($\sigma/\theta_0, \text{Highland} = 0.930 \pm 0.003$) across all 4 geometries and 5 momenta.
- The tail's fitted power index, $\alpha = 2.999 \pm 0.046$, is indistinguishable from the textbook Rutherford value ($\alpha = 3$) — and it dominates the raw variance (81–90% of $\Sigma \theta^2$) despite being only 3.3–3.9% of events.

A Rayleigh Core and a Rutherford Tail

The distribution was refit as an explicit two-component mixture — a Rayleigh core for the many-small-deflections regime Highland's formula describes, plus a Rutherford power-law tail for the rare single-scatter events it was never designed to capture.

CORE / HIGHLAND RATIO

0.930 ± 0.003

a stable 7% shortfall — not the 2× gap the naive fit implied

TAIL POWER INDEX α

2.999 ± 0.046

Indistinguishable from the textbook Rutherford value, $\alpha = 3$

TAIL SHARE OF $\Sigma\theta^2$


86.5%

from a tail that is only 3.3–3.9% of events


So what does this actually mean:

- The core agrees with theory — Highland's formula matches the simulated core to within its own stated accuracy.
- The tail is real physics — a Rutherford cross-section $d\sigma/d\Omega \propto \theta^{-4}$ converts to $dN/d\theta \propto \theta^{-3}$ via the solid-angle Jacobian, exactly the fitted slope.
- The shape is scale-invariant — rescaling by each momentum's own Highland θ_0 collapses all five momenta onto one curve.

A Universal Scale-invariant Shape



Barrel reference, rescaled by each momentum's own Highland ϑ_0



Kresling, all momenta — the same core+tail fit tracks Highland at every momentum

Rescaling the barrel's scatter-angle axis by each momentum's own Highland θ_0 collapses all five simulated momenta onto a single curve. The same shape recurs for the folded geometries too — across all 20 geometry $\times$ momentum configurations, the fitted σ/θ_0 , Highland stays within a narrow 0.924–0.935 band. The $\sim 7\%$ core offset and the power-law tail are structural, scale-invariant properties of the geometry and the physics — not momentum-dependent artifacts.

Areal Decomposition

ONE NUMBER, TWO REGIONS

A single acceptance number hides where it comes from. Areal decomposition splits the sphere of possible track directions into regions, scores each region on its own, then adds the pieces back up — each weighted by how much of the sphere it covers.

Areal weight

The fraction of all track directions a region covers (its solid angle). The source fires isotropically, so a region's weight is just its share of the sphere.

Fiducial range · $|\eta| \leq 0.80$

The central band, θ from 48.4° to 131.6° — 66.4% of all directions. “Fiducial” means the zone you trust as a fair test: tracks arrive nearly head-on and every geometry is near-hermetic (no gaps).

Peripheral range · $|\eta| > 0.80$

The other 33.6%: directions that skim toward the barrel's open ends, where a flat barrel's finite length runs out before the track does.

ΔA and “pp”


ΔA = a fold's acceptance minus the flat barrel's. pp = percentage points: 50.02% $\rightarrow$ 56.89% is +6.87 pp, a plain subtraction — not a “6.87% increase.”

Share of gain

How much of a fold's overall ΔA was earned in one region. Peripheral share: 94.4% Miura-ori, 98.9% Yoshimura, 97.5% Kresling.

WHERE THE GAIN LIVES

barrel cross-section, shaded by where gain is earned



← 94–99% of every fold's gain is earned here →

fiducial · 66.4%

peripheral · 33.6%

PUTTING IT BACK TOGETHER

Barrel $0.664 \times 99.98\% + 0.336 \times 50.02\% \approx 83.2\%$

Kresling $0.664 \times 99.98\% + 0.336 \times 56.89\% \approx 85.5\%$

Only the peripheral term changes — so 97.5% of Kresling's overall +2.36 pp gain is earned there.

Analyzing the Gains: Areal Decomposition

Every primary track was classified by true generated angle into a central fiducial band ($|\eta| \leq 0.80$) or a peripheral band ($|\eta| > 0.80$) and scored separately, to test whether folding helps everywhere or only at the edges.

Geometry	Global ΔA	Fiducial A	Peripheral A	Peripheral ΔA	Share of gain
Barrel (fine)	—	99.98%	50.02%	—	—
Miura-ori	+0.49 pp	99.98%	51.41%	+1.39 pp	94.4%
Yoshimura	+1.30 pp	99.98%	53.86%	+3.84 pp	98.9%
Kresling	+2.36 pp	99.98%	56.89%	+6.87 pp	97.5%


94–99%

of every fold's acceptance gain comes from the peripheral band alone. Inside the central fiducial region, all four geometries are statistically identical — folding adds nothing there, and costs nothing there either. The entire effect lives at the barrel's grazing-angle ends, where inclined facets keep catching tracks ($\theta_{\text{local}} \gtrsim 65^\circ$) that a flat barrel geometrically cannot.


Two regimes, one boundary

- **Central band** ($|\eta| \leq 0.80, \theta \in [48.4^\circ, 131.6^\circ]$) — encloses 66.4% of the isotropic 4π solid angle. All geometries near-hermetic here.
- **Peripheral band** ($|\eta| > 0.80$) — where a flat barrel's finite length physically runs out before the track does; folded end-facets are angled inward and keep intercepting it.


Peripheral-band acceptance



Thin-Silicon Variant & ITS3 Benchmark



(a) Silicon-only, matched physical scope



(b) Si + Kapton vs. full engineered ITS3 layer

At 50 μm silicon (matching ALICE ITS3's 40–50 μm bent sensors), the same ranking holds: Kresling +2.37 pp, Yoshimura +1.26 pp, Miura-ori +0.59 pp.

Silicon-only X/X_0 sits +0.010 to +0.016 pp above ITS3's continuously bent baseline — flat facets can't fully match a curved surface at non-normal incidence.

Total material traversal scales down by $\sim 4.8\times$ when silicon drops from 300 μm to 50 μm , since the Kapton backing now makes up a larger share of the stack.

The full Si+Kapton stack falls below ITS3's engineered layer budget — but that isn't apples-to-apples, since ITS3's figure includes structural glue, cooling and supports this study excludes.

Discussion & Conclusion

DISCUSSION

- Within the central barrel ($|\eta| \leq 0.80$) all four geometries hit 99.98% acceptance — folding changes nothing there.
- 94–99% of every fold's acceptance gain comes from peripheral, grazing-angle tracks a flat barrel geometrically cannot catch past $\theta_{\text{local}} \approx 65^\circ$.
- Folding trades material for reach — it doesn't out-perform a flat facet head-on, it simply keeps catching tracks the flat barrel has already lost.

CONCLUSION

Origami folding is a genuine, quantifiable new design axis for cylindrical tracking detectors: it reliably buys geometric acceptance at a cost in material, at any matched silicon thickness. Which fold to use is a choice, not a default.

Choose Kresling

when maximizing acceptance is the priority — highest gain, intermediate material cost.

Choose Miura-ori

when momentum resolution matters most — minimal material, near-zero under-coverage.

Avoid Yoshimura

Pareto-dominated by Kresling on every axis: less acceptance, more material, more under-coverage.

From Simulation Result To Engineering Problem

The Pareto trade-off is a real geometric result. Turning it into a usable detector component is a separate, harder problem.

01

What this study shows

Origami folding reliably increases geometric acceptance at any matched silicon thickness, at a proportional radiation-length cost — a genuine three-way Pareto trade-off.

02

What it doesn't yet show

Whether a rigid-silicon-on-Kapton facet survives real fold-vertex stress, thermal cycling, and magnetic loads without fatigue or delamination at the crease.

03

The concrete next step

Fabricate a flexible silicon-on-Kapton test module and run structural/thermal FEA on fold-vertex stress, ahead of any beam test.

Claiming a fabrication-ready design would overstate what a Geant4 acceptance study can establish. The honest claim is narrower and stronger — folding is a genuine, quantifiable new design axis, worth the engineering investment to test physically.

An Independent Idea, Tested Against Outside Scrutiny

This project had no institutional backing, no lab, and no supervising research group — the fold-geometry idea, the simulation pipeline, and the analysis were developed independently. Four physicists across three institutions gave feedback at different stages, unprompted by each other.

ON THE INITIAL, BARE-BONES IDEA

Prof. Biplob Bhattacharjee

Professor, Centre of High Energy Physics, IISc - early feedback on the initial concept

Prof. Sudhir Vempati

Chair, Centre of High Energy Physics, IISc — early feedback on the initial concept

ON THE MANUSCRIPT, WITH RESULTS IN HAND

Prof. Mei Dongming

University of South Dakota — detailed feedback on the completed manuscript

Prof. Daniel Marlow

Princeton University — feedback on the completed manuscript

None of the above endorse this project's conclusions, and none were involved in its design or execution — the idea, the methods, and any mistakes in this work are the author's alone. The reason to name them is narrower: at two different stages, independent researchers looked at the reasoning, did not find a reason to stop it and went on to call it interesting

Limitations, Next Steps & References

L I M I T A T I O N S

- No structural or thermal modeling of fold-vertex creases — stress, cooling, gravitational/B-field loads are untested.
- Material budget covers Si + Kapton only; supports, cabling, glue and readout electronics are excluded.
- No physical prototype has been built or beam-tested — this is a simulation-only study.

N E X T S T E P S

- Fabricate a flexible silicon-on-Kapton folded test module.
- Add a complete engineering material budget.
- Run structural/thermal FEA on fold-vertex stress concentrations.
- Beam-test a physical prototype against the simulation.

R E F E R E N C E S

[1] ALICE Collab., ALICE ITS3: a bent stitched MAPS-based vertex detector, JINST 19 (2024) C02048.

[2] ALICE Collab., ITS3: a truly cylindrical inner tracker for ALICE, PoS Vertex2023 (2023) 015.

[3] LHCb Collab., LHCb upgrade: Upstream Tracker, PoS Vertex2014 (2014) 009.

[4] CMS Collab., The phase-2 upgrade of the CMS tracker, CERN-LHCC-2017-009 (2017).

[5] K. Miura, Method of packaging & deployment of large membranes in space, ISAS Rep. 618 (1985).

[6] Y. Yoshimura, On the buckling of a circular cylindrical shell under axial compression, NACA TM 1390 (1955).

[7] B. Kresling, Natural twist buckling in shells, IASS-IACM (2008).

[8] GEANT4 Collab., GEANT4 — a simulation toolkit, NIM A 506 (2003) 250.

[9] V. L. Highland, Some practical remarks on multiple scattering, NIM 129 (1975) 497.

[10] G. R. Lynch & O. I. Dahl, Approximations to multiple Coulomb scattering, NIM B 58 (1991) 6.

Check out the preprint at Zenodo (<https://doi.org/10.5281/zenodo.22775533>)
github.com/Alphacodde/Origami-Particle-Detectors-Geometries