UNITACTILE: GEOMETRY–CONNECTIVITY TRADE-OFFS IN AUTOMATED TACTILE SKIN DESIGN

Anonymous authors

ABSTRACT

Converting robot geometry into a tactile skin requires more than flattening a surface: the resulting sheets must also support electrode connectivity and physical connector footprints. We study this coupled problem for a fixed sandwich architecture comprising separate row and column flexible printed circuit boards. Unitactile combines prepared surface targets, a trusted geometry and circuit evaluator, a distortion-aware coverage objective, and numerical or program-generating search. Our main experiment evaluates a deterministic multi-sheet baseline on 21 prepared Allegro-hand targets under five segmentation policies, completing 104 of 105 conditions. The tested single-sheet policy produces no computationally valid sheets. Allowing up to eight sheets yields 18 valid sheets on 12 targets, with jointly reachable pads covering 8.3% of the original target area on valid sheets and 14.0% across all safely routed sheets, at 66 connectors. On the 20 targets completed under every policy, raw coverage peaks at four sheets, whereas valid-sheet coverage is higher at eight sheets. Splitting reduces distortion but leaves many regions without a sampled connector placement. On three development targets, an offline common-evaluator comparison of five deterministic flattening settings and saved language-model agent designs finds no computationally valid single sheet for any method and does not establish an advantage for specialization. The results characterize a geometry–connectivity trade-off under explicit hardware assumptions; they establish computational feasibility for selected sheets, not physical fit or fabrication readiness.

1 INTRODUCTION

Tactile skins must conform to the surfaces on which robots make contact. Robot hands contain flat panels, rounded links, concavities, and small fingertip domes, often within a single assembly. A flexible printed circuit board (FPCB) can bend, but a planar layout does not generally fit a doubly curved surface without in-plane deformation, cuts, or incomplete coverage. Automated sensor design must therefore decide which surface to cover, how to flatten it, and how to maintain an electrically useful sensing array after cutting.

These decisions are coupled. A seam can reduce distortion but sever a shared electrode bus; splitting can improve flattening but leave patches too small for their connectors. Consequently, a convincing UV map, a high pad count, and a connected circuit are individually insufficient evidence of a useful tactile skin.

Our baseline experiments examine this coupling under a fixed sensor architecture. Each sheet comprises two FPCBs with opposing row and column electrodes, connected through independent routing layers to separate flexible-flat-cable connectors. The electrode pitch, channel count, connector envelope, routing rules, and allowable geometric strain remain fixed. A pad contributes useful area only when it is reachable from both axes. These assumptions separate parameterization, connectivity and hardware failures. A subsequent extension replaces local connectors with backside flex-bus continuations toward a shared readout.

Language-model agents initially motivated a private evaluator and diagnostic feedback. Their geometric improvements did not yield completely valid single-sheet designs on the tested targets. We therefore also study multiple independently connected sheets, which changes hardware resources as well as geometry.

Our contributions are: a circuit-aware formulation separating geometric validity, electrical safety, completeness and physical qualification; a distortion-aware coverage objective and deterministic

Table 1: Frozen local-connector baseline hardware and evaluation settings. These are design assumptions, not experimentally validated material limits.

Quantity	Value
Grid pitch / square pad side	3.0 mm / 2.0 mm
Channels / connector pitch	32 per axis / 0.5 mm
Usable planar grid extent	95 mm square after edge clearance
FFC connector	Würth 687132183622
Rigid stiffener envelope	23.7×6.8 mm, inside sheet
Trace width / clearance / edge clearance	0.20 / 0.15 / 0.50 mm
Via / drill diameter	0.70 / 0.30 mm
Maximum principal stretch error	0.05
Search budget	6 layouts/sheet; up to 8 routes/board

multi-sheet baseline; a 21-target, five-policy study accounting for failures, discarded area and connector costs; and a common-evaluator comparison of deterministic and agent designs. The evidence identifies selected computationally feasible sheets and limitations of the tested search, without establishing optimal partitions, general single-sheet infeasibility or an advantage of one agent architecture.

2 Related Work

Custom tactile sensors. Olberding et al. (2013) study cuttable multi-touch sensors, emphasizing the relationship between cut shape and electrical connectivity. Flex-Glove automates personalized FPCB tactile-glove design from hand imagery (Murphy et al., 2025). These systems motivate treating sensing and fabrication constraints as part of layout generation. Our setting starts from prepared three-dimensional robot surfaces and measures both metric distortion and jointly reachable electrode area.

Parameterization and program search. Boundary First Flattening (BFF) offers boundary and cone control for conformal parameterization (Sawhney & Crane, 2018). SLIM optimizes locally injective mappings (Rabinovich et al., 2017), while OptCuts jointly optimizes cuts and parameterization under a distortion objective (Li et al., 2018). Such methods supply useful geometric building blocks, but their geometric energies do not directly enforce our electrode routing or connector constraints. Our partition baseline adapts normalized graph cuts (Shi & Malik, 2000), without claiming a new spectral method. Program-search systems such as FunSearch combine language-model proposals with executable evaluation (Romera-Paredes et al., 2024). We use that broad separation between proposal and verification for exploratory sensor-design agents; our small campaigns are not a comparison against FunSearch or a general evaluation of autonomous research systems.

3 Problem and Evaluation

3.1 Surface targets and the sensor architecture

Let $M = (V, F)$ be a prepared target mesh with vertex positions in millimeters and surface area A. The target is fixed before optimization. A design partitions a subset of its faces into connected regions $F_1, \ldots, F_m$, leaving a remainder F_0 unsensorized. Each region has a planar chart U_s, a rigid grid pose $g_s \in \mathrm{SE}(2)$, and separate row and column connector poses $c_s^r, c_s^c \in \mathrm{SE}(2)$. The source correspondence must preserve every face assigned to a sheet exactly once. In the original single-sheet setting, all target faces belong to one chart and internal seams can be introduced. The multi-sheet baseline instead uses disk regions without internal seams.

In the baseline protocol, the front copper of one board contains horizontal electrode connections and the other contains vertical electrode connections. Each board routes on its reverse copper layer through vias to its own connector. Intersections in a two-dimensional drawing can therefore belong to different physical layers; intersections between distinct nets on the same copper layer are disallowed. The connectors are FFC connectors leading to readout electronics, not multiplexers themselves. Mechanical coexistence of the two connector bodies in the assembled stack is not modeled.

Table 1 gives the fixed protocol. A full interior grid cell has pad fill factor $2^2/3^2 = 4/9$, a planar reference density rather than an upper bound on source coverage; we do not rescale coverage by it.

3.2 Trusted reconstruction and outcome classes

The evaluator receives chart arrays, face correspondence, and poses. It reconstructs the electrode grid and copper using fixed rules instead of accepting an agent's proposed connectivity or score. Geometry checks include source integrity, disk topology, a simple boundary without holes, non-degenerate and consistently oriented UV triangles, overlap checks, grid extent, and channel limits. Connector footprints must fit the outline with clearance. Routing checks assess connectivity to the prescribed pins and isolation between distinct nets.

Electrical *safety* and *completeness* are separate. A safely routed board pair can leave some pads unreachable without introducing a short. A pad is jointly reachable only when both its row and column reach the assigned connector pins. A sheet is *computationally valid* when the geometric checks and strain limit pass and every retained pad is jointly reachable. Unsafe board pairs receive no coverage credit. Safe but incomplete pairs retain diagnostic coverage for the pads that do connect. This distinction provides progress feedback while preventing partial layouts from being reported as complete sensors.

The evaluator also separates execution from design quality: a completed search without a usable layout has zero coverage, whereas a timed-out condition has missing metrics. A computationally valid sheet still lacks physical fitting or manufacturing qualification.

3.3 Distortion and coverage without a compression reward

For triangle t, let J_t be the Jacobian from a local orthonormal source frame to UV coordinates, with singular values $\sigma_{t,1}, \sigma_{t,2} > 0$. The hard geometric criterion uses

$$e_{\max} = \max_{t,i} |\sigma_{t,i} - 1| \le 0.05. \tag{1}$$

This is a proxy for local metric compatibility, not an elastic simulation or a measured FPCB failure threshold. To obtain a continuous score that treats stretch and reciprocal compression symmetrically, define

$$d_t = \max_i |\log \sigma_{t,i}|, \qquad \phi_t = \left(\frac{d_t}{\log 1.05}\right)^2, \qquad w_t = \frac{1}{1+\phi_t}. \tag{2}$$

The weight is one at isometry and one-half at stretch 1.05 or compression $1/1.05$. It is not exactly one-half at 0.95, so the continuous reference and the hard strain interval are intentionally distinguished. Nonpositive singular values are rejected, not clipped into a finite score.

Let a_t be source triangle area. We compute $E_{\text{all}} = \sum_t a_t \phi_t / A_s$ for sheet area A_s, and E_{tail}, the area-weighted mean over the worst 1% of source area, with fractional inclusion of the final triangle. For internal seam length L_s and slit-tip count n_s, the geometry factor is

$$Q_s = \frac{1}{1 + E_{\text{all}} + E_{\text{tail}}} \frac{1}{1 + L_s/\sqrt{A_s} + 0.1 n_s}. \tag{3}$$

The tail term discourages localized collapse from disappearing in an average. Seam length and tips penalize internal cutting in the single-sheet problem. Both are zero within the uncut regions of the multi-sheet baseline; inter-region boundaries are reported separately.

For a pad p intersecting UV triangle t, let b_{pt} be the planar intersection area and let $\tilde{a}_{pt} = b_{pt} a_t / a_t^{UV}$ be its corresponding source area. With $z_p = 1$ only for jointly reachable pads on a safe board pair, define

$$C_s^{\text{raw}} = \frac{\sum_{p,t} z_p \tilde{a}_{pt}}{A_s}, \quad C_s^{\text{cons}} = \frac{\sum_{p,t} z_p \min(\tilde{a}_{pt}, b_{pt})}{A_s}, \quad C_s^{q} = \frac{\sum_{p,t} z_p \min(\tilde{a}_{pt}, b_{pt}) w_t}{A_s}. \tag{4}$$

Pads are checked for overlap before accumulation. Crediting the smaller area prevents a compressed UV triangle from turning one small flat pad into arbitrarily large credited surface area. The per-sheet optimization score, called `metric_coverage_v1` in the implementation, is

$$R_s = Q_s (1 + C_s^q). \tag{5}$$

The additive one preserves a geometry signal before useful routing is found. It also means geometry dominates when distortion is severe; feedback therefore exposes the components rather than just the scalar. The score orders continuous design quality and does not itself certify validity. Controller stage offsets are described in Appendix B.

For a whole part, coverage uses the original target denominator, $C_M = \sum_s A_s C_s / A$. Discarded and unrouted faces remain in A. A valid-only version restricts the numerator to computationally valid sheets. We do *not* sum R_s across sheets: doing so would reward changing the number of regions without adequately pricing connectors or partition boundaries. The multi-sheet comparison instead reports coverage together with actual connector counts and geometry diagnostics.

4 Search Methods

4.1 Target preparation and geometric solvers

Inputs are pinned robot-part meshes assembled using their URDF transforms at one fixed pose. Preprocessing classifies external exposure using solid-membership and outside-air tests, excludes parts below a 5% exposed-area threshold, densifies eligible geometry, and removes thin features and small components. Target postprocessing produces a contiguous, hole-free manifold disk and rechecks exposure. Optimization begins from these frozen targets rather than changing the desired sensing area to improve a score. The final study uses the 21 prepared Allegro-right targets; it does not evaluate robustness to joint motion or target-selection errors.

The fixed-topology flattening implementation combines harmonic, least-squares conformal, and intrinsic initializations with ARAP/SLIM-style optimization and worst-stretch refinement. Development baselines additionally include BFF cone-count sweeps, OptCuts, and an incremental cut-growth heuristic. These optimize different geometric energies, so an energy bound or conformal map is not interpreted as satisfying Equation 1. Every submitted chart is remeasured by the same evaluator. The multi-sheet experiment uses only the fixed-topology solver, making segmentation policy its principal controlled variable.

4.2 Deterministic multi-sheet baseline

We form a dual graph with one vertex per face and edges between adjacent faces. For adjacent faces i, j with dihedral angle θ_{ij}, set

$$W_{ij} = \exp[-(\theta_{ij}/0.35)^2] + 10^{-6}. \tag{6}$$

Large creases have low affinity. A normalized-Laplacian Fiedler vector defines a threshold sweep; the minimum normalized cut is selected, and disconnected outputs become separate regions. Each region is screened for disk topology, strain and planar extent, with a minimum-area heuristic (Appendix B).

The automatic policy recursively splits failing regions until they pass, become too small, or reach depth 12. Capped policies permit at most $K \in \{1, 2, 4, 8\}$ retained regions, splitting the largest failure only when its connected components fit within the cap. At the cap, excessive-strain regions still receive diagnostic circuit evaluation. Accepted regions are not split to improve routing: the baseline is geometry-driven and does not feed connector failures back into partitioning.

4.3 Grid placement, connector placement, and routing

For the fixed-connector baseline, each region receives up to six grid poses. Connector anchors are sampled on a 1 mm lattice at 15° increments and tested with the actual stiffener-containment predicate, including its anchor offset and 0.5 mm edge clearance. Feasible poses are ranked and diversified; row and column poses may differ (Appendix B).

The router permits nonparallel back-layer paths and varies deterministic routing order and component priority over up to eight attempts per board. Safe row and column candidates are paired to maximize jointly reachable, distortion-weighted pad area. The highest-scoring safe layout is retained for each sheet. Because complete and incomplete layouts can have different retained pad sets, the score winner need not be the layout with the largest completeness indicator. Reported valid-sheet counts therefore describe the selected designs, not the existence of any valid proposal in the search history.

Table 2: Final study across all completed targets. Coverage percentages use original target area. "Raw" includes safely routed pads on invalid sheets; "valid" restricts coverage to valid sheets. Auto omits the palm timeout and therefore has a different denominator.

Policy	Done	Raw	Cons.	Weighted coverage (%)	Valid	Routed sheets	Valid sheets	Conn.
$K = 1$	21	5.8	3.5	1.0	0.0	8	0	16
$K \leq 2$	21	8.2	6.9	3.2	0.8	22	3	44
$K \leq 4$	21	15.0	14.4	11.7	5.1	36	12	72
$K \leq 8$	21	14.0	13.9	13.4	8.3	33	18	66
Auto	20	17.0	17.0	16.7	9.7	26	16	52

An empty connector search means no sampled placement was found for the tested grid proposals. It does not prove continuous infeasibility. Likewise, a routing failure under the fixed discretization and order budget is not a proof that no wiring exists. These finite searches provide reproducible baselines and interpretable failure records.

4.4 Exploratory language-model agents

The agent interface asks a language model to emit a Python construction program. The program runs in an isolated numerical environment and produces geometry and pose artifacts; the private evaluator builds and scores the circuit afterward. A direct agent receives one proposal call, a generic agent receives up to four feedback rounds, and two specialists alternate geometry and layout calls for two rounds, also totaling four calls. Later protocol revisions provide a public contract checker, read-only prior artifacts, and spatial feedback including pad reachability, failed nets, and extent violations. Evaluator implementation files remain outside the agent workspace.

The specialists decompose the design variables: the geometry role proposes the chart and cuts, while the layout role proposes grid and connector poses. Both depend on the same trusted evaluation. Equal call caps do not imply equal tokens, successful calls, or runtime. Our historical campaigns differ in interface revision and completion, so they are reported separately from the deterministic final experiment. The experiment does not include the previously planned co-scientist or evolutionary-agent variants.

5 Experiments

5.1 Protocol and analysis population

The final matrix comprises 21 prepared targets from an Allegro right hand crossed with five policies: $K = 1$, $K \leq 2$, $K \leq 4$, $K \leq 8$, and automatic segmentation. Their original target area totals approximately 85,378 mm^2, or 61.7% of the assembled mesh surface recorded by preprocessing. All conditions use the same hardware, target files, per-sheet proposal budget, and routing budget. They contain no language-model calls. We run one deterministic realization per condition with a 150-minute limit. Jobs use single-threaded numerical workers; concurrent scheduling differs between execution batches, so runtime is descriptive rather than a controlled speed comparison.

Of 105 official conditions, 104 complete and one, automatic segmentation of the palm, times out. Its partial outputs are excluded. All-policy paired comparisons use the same 20 completed targets, approximately 55,400 mm^2 in total. The palm alone accounts for about 35.1% of the full target area, making this distinction material. Coverage is area-weighted across targets, not an average of per-part percentages. Related links are not independent robot families; no statistical significance or cross-robot generalization is inferred from the 21-part inventory.

Two implementation failures were corrected during execution: an initial batch crashed while writing summaries, and a capped-path error affected $K \geq 3$. The final dataset combines unaffected Auto/$K = 1$/$K \leq 2$ jobs with rerun $K \leq 4$/$K \leq 8$ jobs; both source identities and the excluded runs are preserved (Appendix A).

5.2 Coverage improves, but raw and valid coverage disagree

Table 2 separates all safely routed coverage from coverage on computationally valid sheets. The tested $K = 1$ policy yields no valid sheets on any of the 21 targets. It nevertheless routes eight

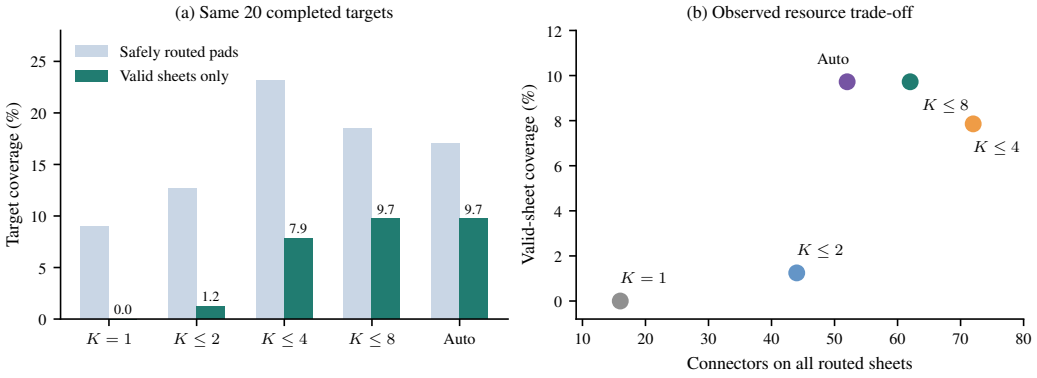


Figure 1: Paired comparison on the 20 targets completed under every policy. Left: raw coverage peaks at $K \leq 4$, but valid-sheet coverage is higher at $K \leq 8$ and Auto. Right: coverage and actual connector counts, including connectors on invalid routed sheets. These are observed policy outcomes, not an optimal Pareto frontier.

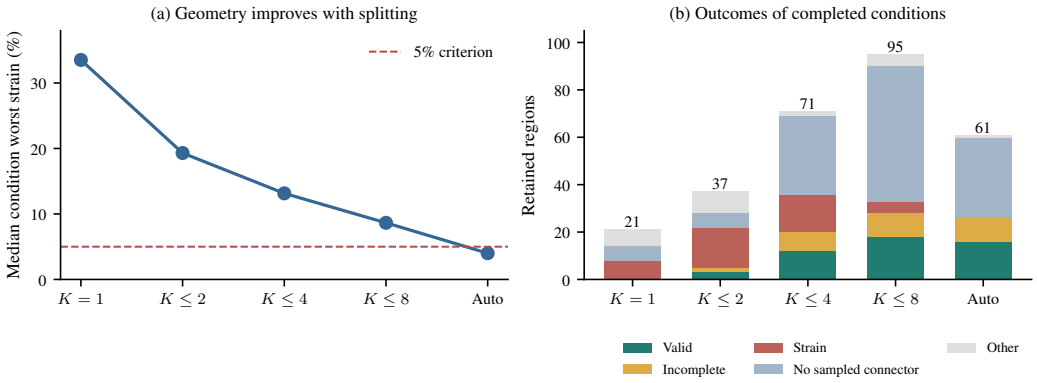


Figure 2: Geometry and region outcomes across completed conditions (21 targets per cap; 20 for Auto). Left: median of each condition's maximum assessed region strain, not a physical fit error. Right: mutually exclusive diagnostic outcome labels. "No sampled connector" describes a finite search failure, not impossibility. Strain labels take precedence over incompleteness when both occur.

sheets and reports 5.8% raw target coverage. All eight violate strain. The remaining regions include six with no sampled connector placement, six rejected by extent, and one without a safe evaluated layout. These outcomes show that connectivity alone does not make the single-sheet results acceptable; they do not prove that every possible single-sheet design fails.

Increasing the cap produces 3, 12, and 18 valid sheets for $K \leq 2, 4, 8$, respectively. Under $K \leq 8$, 12 of 21 parts have at least one valid sheet. Across all targets, 33 routed sheets require 66 connectors, and the valid subset contributes 8.3% coverage. This is partial coverage of the prepared targets, not full coverage of 12 parts or an installed hand. Counting only the 18 valid sheets would imply 36 connectors; we report all 66 associated with the selected routed designs to expose the resources spent on diagnostic, incomplete outputs too.

On the paired 20-target subset (Figure 1), raw coverage peaks at $K \leq 4$ (23.2%); valid-only coverage peaks at $K \leq 8$ and Auto (both 9.7%, with 62 and 52 connectors). These are different operating points, and the missing palm precludes full-hand equivalence. Conservative and weighted coverage also approach raw coverage as splitting reduces distortion (Table 2).

5.3 Splitting changes the dominant observed failure

The median condition maximum strain decreases from 33.5% at $K = 1$ to 19.3%, 13.1%, and 8.6% at the larger caps, and 4.0% under Auto (Figure 2). These statistics include assessed regions outside the final valid set. They are descriptive, and Auto has a different target population. Splitting

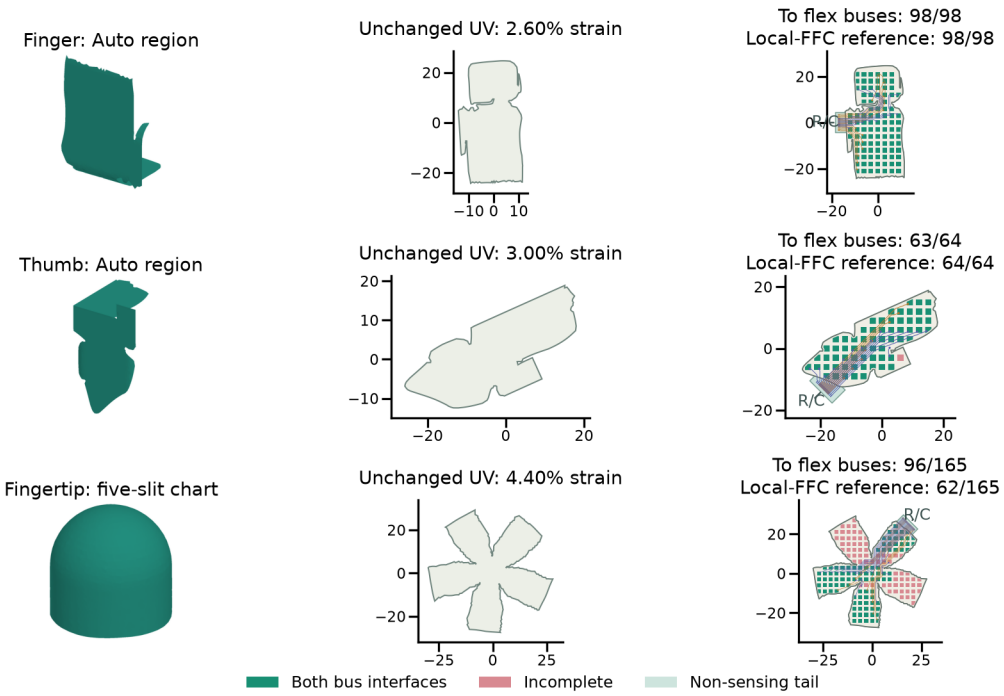


Figure 3: Local flex-bus extension on the saved finger, thumb and fingertip charts. Green pads reach both continuation interfaces; pink pads miss an axis. Orange/blue copper belongs to separate row/column FPCBs; R/C mark onward bus interfaces, not local connectors. Pale green is added non-sensing flex. Counts compare the new local-interface result with its local-FFC reference. No board body or stiffener is placed on a region. Source meshes show correspondence, not physical fit; routing to the hand base is not evaluated.

also increases partition boundary length: the corresponding medians are 0, 43, 102, 160, and 170 mm. These boundaries include interfaces to discarded regions; they are not a direct estimate of manufactured cut length or seam width.

The failure counts show the competing footprint constraint. Under $K \leq 8$, 57 of 95 retained regions produce no sampled connector placement; Auto has 34 such regions out of 61. Electrical incompleteness remains important: ten selected sheets in each of these policies have acceptable geometric assessment but incomplete connectivity. Thus improved flattening does not eliminate the circuit-design problem. The automatic policy also drops a mean 29.2% of target area across its completed parts, partly because of the minimum-area rule. This mean is per-part and should not be confused with area-weighted electrode coverage. Median completed-condition runtimes are 240, 543, 417, 424, and 619 s for $K = 1, 2, 4, 8$ and Auto; they are descriptive rather than equal-compute comparisons.

5.4 Representative regions and flex-bus continuation

The baseline includes complete local grids: 98/98 pads on one finger Auto region and 64/64 on a thumb region. This does not imply full-part coverage: Auto covers 22.7% of the finger target on valid sheets and only 5.4% of the thumb target. The saved five-slit fingertip meets the strain bound (4.40%) but connects only 62/165 pads to its local FFC interfaces.

Figure 3 tests a different architecture. We hold those UV maps, grids and pads fixed, remove local connector bodies and stiffeners, and reroute backside copper to continuation sections on non-sensing flex tails. These sections are unpopulated circuit boundaries for onward routing through the skin. A bounded search chooses tail position, width, routing order and independent row/column exits, maximizing jointly reached pads and then minimizing added substrate (Appendix G.2). Trace, clearance and via rules are retained; finite-width copper and intact via annuli determine reachability.

Table 3: Single-sheet method comparison under one evaluator (`metric_coverage_v1`, eight routes per board). Each cell gives the selected R (maximum over safely evaluated candidates) and safely evaluated/total candidates. A dash means no safely evaluated candidate, not zero; † marks an interrupted condition. No candidate in any row is computationally valid. Agent rows used 1–4 model calls (29k–229k tokens where recorded); deterministic rows used none. Single observations with unequal search budgets; not a ranking.

		Finger		Fingertip		Thumb	
Method	Placement	R	safe	R	safe	R	safe
Fixed topology	bank (18)	0.023	18/18	0.0032	18/18	–	0/18
Cut growth (fallback)	bank (18)	0.023	18/18	0.0032	18/18	–	0/18
BFF, 4 cones	bank (18)	0.071	18/18	0.0079	18/18	–	0/18
BFF, 8 cones	bank (18)	0.186	18/18	0.018	6/18	–	0/18
OptCuts, bound 4.05	bank (18)	0.023	18/18	0.0083	18/18	–	0/18
direct	agent	–†	0/0	0.0060	1/1	0.0074	1/1
generic feedback	agent	0.160	4/4	0.114	4/4	1.3×10^{-4}†	1/1
two specialists	agent	0.233	2/2	0.117	2/2	–	0/2
direct chart	bank (18)	–†	0/0	0.0060	18/18	–	0/18
generic feedback chart	bank (18)	–	0/72	0.0062	54/72	–†	0/18
two specialists chart	bank (18)	0.241	36/36	–	0/36	–	0/18

The selected layouts reach 98/98, 63/64 and 96/165 pads, respectively. The fingertip gains 34 reachable pads; the finger retains complete connectivity, while one thumb pad remains blocked. Removing local hardware therefore helps some layouts but does not guarantee complete routing. Added substrate is 39.15, 44.85 and 44.65 mm^2 per FPCB, excluded from sensing credit. These are local-interface results under changed hardware, not base-board connectivity or replacements for the frozen coverage and method tables.

The baseline palm also illustrates a partition limitation: its first split produces five components and exceeds caps up to four. At $K \leq 8$, two valid sheets cover 5.6% of its target; Auto times out. A cap constrains one heuristic trajectory, not all possible partitions.

5.5 Method comparisons on three development targets

On finger `link_0.0`, fingertip `link_3.0_tip`, and thumb `link_12.0`, we compare fixed-topology flattening, cut growth, four- and eight-cone BFF, OptCuts, and saved agent designs. None of the archived 30-second baseline charts meets the 5% strain bound. Eight-cone BFF is closest (9.4%, 22.1%, 6.1%), but its thumb chart exceeds the fixed extent; cut growth always falls back to the fixed map. Appendix E details these failures.

Offline evaluation under identical evaluator and hardware settings reproduces all 15 archived agent scores and six BFF results. Deterministic charts receive an 18-pose placement bank. Agent charts use their saved poses and, separately, the same bank; no new model calls are made.

Under the local-connector protocol, none of the 520 candidates is computationally valid (Table 3). Agent fingertip charts meet the strain bound but lose connectivity across their slits. Best agent scores exceed the deterministic results on finger and fingertip (0.233 versus 0.186; 0.117 versus 0.018), with unequal budgets and one replicate. The common bank improves the finger specialist chart to 0.241 but misses feasible poses for four other agent conditions. Thus chart quality and placement search interact; these observations do not establish an agent ranking. No baseline thumb chart passes the extent or connector check. Multi-sheet results use different hardware budgets and are not pooled with this comparison.

6 Discussion and Limitations

The local-connector results expose a geometry–hardware mismatch: partitioning reduces distortion, but fixed connector envelopes can exclude the resulting smaller regions. This helps explain non-monotone coverage. It is a frequent observed bottleneck, not the sole cause of failure.

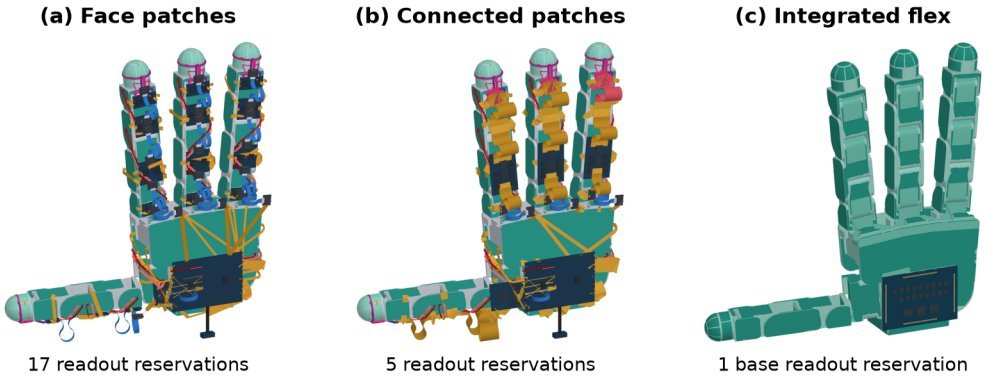


Figure 4: Whole-hand assembly concepts: **(a)** distributed face patches, **(b)** connected patches with one readout per finger plus a palm readout, and **(c)** the current integrated-flex carrier with one base readout. All retain 156 intended sensing targets; light-green carrier and joint webs are non-sensing. Board reservations decrease from 17 to five to one under different hardware assumptions. These are packaging proposals, not measured full-hand coverage or qualified interconnects.

Integrated flex and shared readout. Figure 4 compares face patches, connected patches and the current glove-like carrier. In the integrated design, paired row and column FPCBs extend along the parts, carrying shared backside returns through the skin toward a single base readout rather than hosting electronics at each patch. Close-fitting joint-web reservations replace the exposed harness envelopes. The 156 sensing targets are preserved; the added light-green carrier and joint webs are non-sensing.

Centralization exchanges finger-mounted electronics for longer analog paths and a larger base channel budget: 127 row and 140 column channels, exceeding the frozen experimental hardware budget. Address uniqueness is checked, and the local continuation trial verifies selected copper routes (Section 5.4). The cleaner whole-hand silhouette nevertheless remains a packaging proposal: global copper routing, the flat manufacturing pattern, attachment and joint compliance are unqualified. In particular, the independently reconstructed open and curled webs do not demonstrate that ordinary, non-stretching FPCB can follow the motion (Appendix G).

Figure 5 connects this architecture to a local design and its electrical organization. The illustrated Auto region of finger `link_2.0` has 2.60% worst strain and all 98 retained pads connected to the flex continuations. Their area covers 37.68% of the selected region, so complete grid connectivity does not imply complete surface coverage. At electrode scale, adjacent patches can tap one continuous backside return per shared row; distinct column addresses on the companion FPCB preserve unique row–column pairs. The row and column circuits remain on separate FPCBs. This suggests reducing patch-level connector overhead through shared returns and bridges, but the evaluated local circuit has not been routed into the proposed hand carrier. Joint optimization of sensing regions, bridges and readout locations remains a design direction.

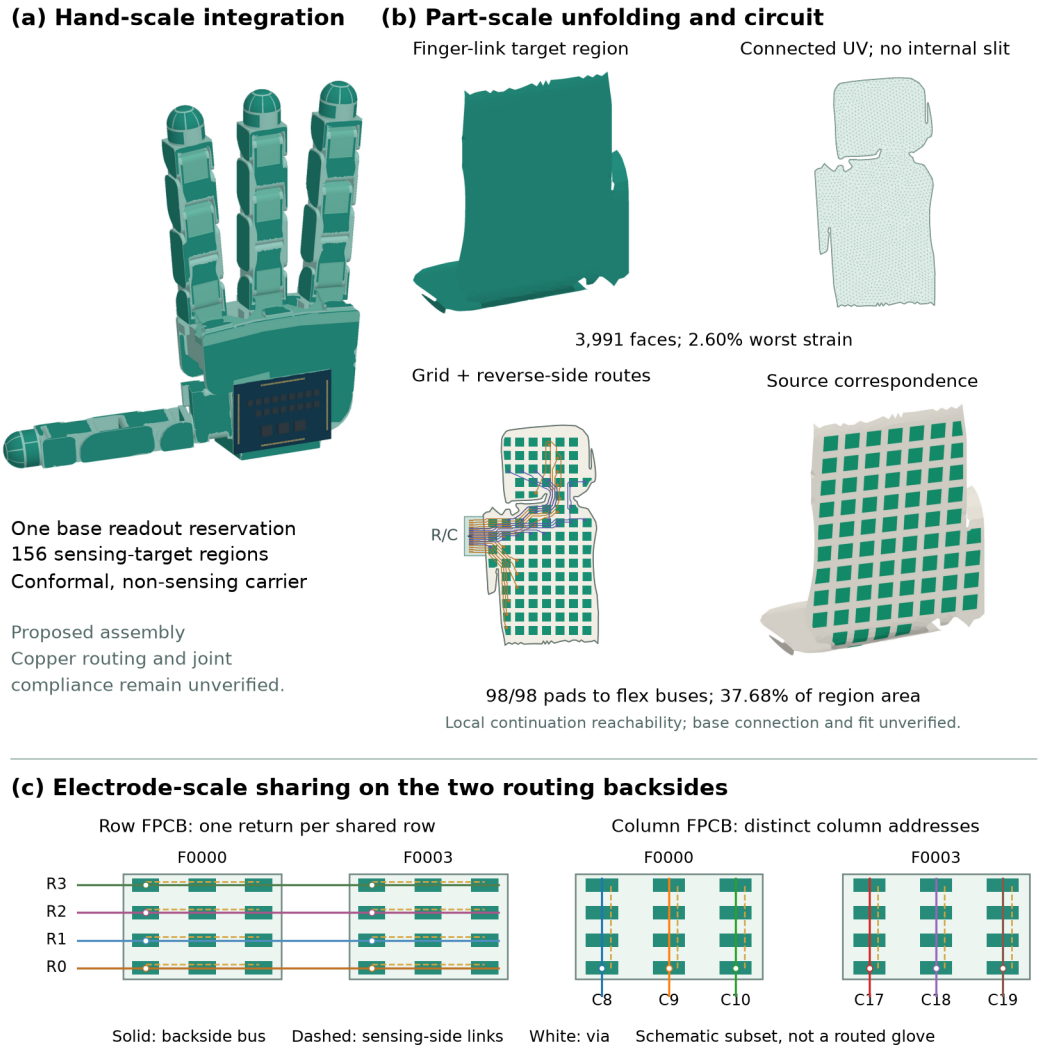


Figure 5: Sensorization from macro to micro. **(a)** Proposed glove carrier and base readout. **(b)** Finger-link region with no local connector body: 2.60% worst strain, 98/98 pads reaching flex continuations, and 37.68% region-area coverage. The source view shows correspondence, not physical fit; backside routes from separate boards are overlaid. **(c)** Shared row returns and distinct column returns on separate FPCBs. The local circuit has not been routed into the whole-hand proposal.

The study covers one hand and pose, related targets, one segmentation trajectory and finite placement searches. The assumed 5% strain bound is not calibrated to a fabricated stack. Neither local distortion nor source correspondence establishes physical wrapping, adhesion, collision clearance or connector access; dense fitting remains unqualified, and no design was fabricated. Agent comparisons have single replicates, censored conditions and unequal budgets. Testing specialization requires matched resources and a same-call generalist control.

7 Conclusion

Tactile-skin design couples geometry and connectivity. Across 104 completed conditions, the eight-sheet policy produced 18 computationally valid sheets on 12 targets, covering 8.3% of original target area. Splitting enables local feasibility but incurs connector overhead and leaves substantial uncovered area. No tested baseline or saved agent design yielded a valid single sheet on the three development targets under local connectors. Replacing them with flex continuations improves fingertip reachability, but whole-hand copper routing and physical validation remain open.

REFERENCES

Minchen Li, Danny M. Kaufman, Vladimir G. Kim, Justin Solomon, and Alla Sheffer. OptCuts: Joint optimization of surface cuts and parameterization. *ACM Transactions on Graphics*, 37(6), 2018. https://www.cs.ubc.ca/labs/imager/tr/2018/OptCuts/.

Devin Murphy, Yichen Li, Crystal Owens, Layla Stanton, Young Joong Lee, Paul Pu Liang, Yiyue Luo, Antonio Torralba, and Wojciech Matusik. Fits like a Flex-Glove: Automatic design of personalized FPCB-based tactile sensing gloves. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2025. doi: 10.1145/3706599.3720147.

Simon Olberding, Nan-Wei Gong, John Tiab, Joseph A. Paradiso, and Jürgen Steimle. A cuttable multi-touch sensor. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 2013. doi: 10.1145/2501988.2502048.

Michael Rabinovich, Roi Poranne, Daniele Panozzo, and Olga Sorkine-Hornung. Scalable locally injective mappings. *ACM Transactions on Graphics*, 36(2), 2017. https://igl.ethz.ch/projects/slim/SLIM2017.pdf.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, et al. Mathematical discoveries from program search with large language models. *Nature*, 625:468–475, 2024. doi: 10.1038/s41586-023-06924-6.

Rohan Sawhney and Keenan Crane. Boundary first flattening. *ACM Transactions on Graphics*, 37(1), 2018. doi: 10.1145/3132705.

Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. doi: 10.1109/34.868688.

A Reproducibility and Evidence Boundaries

The manuscript is based on the archived `unitactile-manuscript-evidence` bundle dated 26 October 2026. The numerical source for the final tables is `campaigns/final-multisheet/analysis/results.json`, checked against the official per-condition summaries. The manuscript's figure script performs only aggregation and rendering; it does not execute a solver, reroute a circuit, select a new incumbent, or invoke a model. A separate figure manifest records hashes of the input artifacts. Source snapshots and excluded executions are retained in the evidence bundle.

Table 4: Execution provenance. Prefixes identify archived source digests, not method variants.

Record	Use in manuscript	Source prefix
Initial matrix (r1)	Excluded: summary construction crash	See archive
Matrix r2	Auto, $K = 1$, $K \leq 2$	`d8ce468e85a4`
Matrix r2, caps 4 and 8	Excluded: capped-path crash	`d8ce468e85a4`
Matrix r3	$K \leq 4$, $K \leq 8$	`6001d87ec58c`
Packaged final source	Reproduction source; not a new run	`852271b5...`
Common-evaluator analysis	Table 3, Appendix E; offline re-evaluation	`852271b5...`

There are 105 official policy–target cells, not 105 successful designs. The automatic palm cell is a timeout and has null metrics. No partial score from this cell enters coverage aggregates. Completed searches with no routed layout are genuine zero-coverage outcomes. Figure 1 excludes the palm from *every* policy, while Table 2 includes all completed targets for each policy. This prevents silent substitution of a smaller Auto population into a full-hand comparison.

The final source corrects an execution bug in the capped path without changing the declared hardware or objective. This supports use of the corrected cap results but does not establish bitwise equivalence across all historical releases. The matrix uses a fixed eigensolver start vector; reproducibility across arbitrary libraries or platforms is not assumed. Full replication should use the recorded environment and target hashes, retaining new runs under new identities if software changes.

Prepared targets are derived from upstream robot meshes. Allegro geometry originates from the pinned `simlabrobotics/allegro_hand_ros` source; the bundle records provenance and checksums. Redistribution terms require review before public dataset release. We do not describe the archive as an unrestricted public benchmark. Other downloaded robot families, including LEAP and Allegro-left, were not included in the final prepared-target experiment.

B Metric and Routing Details

Placement implementation. For each region, up to six grid proposals combine offsets relative to an interior reference point with angles $0^\circ, \pm 30^\circ, 60^\circ$. Connector anchors are searched on an origin-aligned 1 mm lattice at 15° increments. Each candidate is tested with the evaluator's actual containment predicate: the transformed stiffener must lie inside the outline eroded by 0.5 mm. The anchor-to-stiffener-center offset is included. Feasible candidates are ranked by their orientation toward the interior and diversified by position and angle; selected row and column poses may differ.

Partition implementation. Large creases have low affinity and are inexpensive to separate. We compute a normalized-Laplacian Fiedler vector and sweep 41 thresholds spanning 20–80% cumulative face area, selecting the smallest normalized-cut objective. A fixed eigensolver initialization and sign convention make the implementation deterministic under its recorded environment. Disconnected outputs are separated into connected components before assessment.

A region is accepted geometrically if it is a disk, its fixed-topology flattening meets the strain bound, and its bounding square fits within 94.95 mm under a sampled rotation search. The 0.05 mm margin is numerical. The extent screen evaluates rotations every 3°; the subsequent proposed layout must still pass the actual grid-extent check. Near the strain threshold, the solver receives an additional 150 worst-stretch refinement iterations. Regions below $1.6(23.7 \cdot 6.8) = 257.856$ mm^2 are dropped by a size heuristic. This threshold is not a geometric proof that a connector cannot fit.

The *automatic* policy recursively splits failing regions until they pass, become too small, or reach depth 12. The *capped* policies permit at most $K \in \{1, 2, 4, 8\}$ retained regions. They split the

largest failing region first, with deterministic tie-breaking, only if the retained connected components fit within the cap. At the cap, regions exceeding strain are still evaluated to preserve diagnostic information; regions that cannot be flattened, are not disks, or exceed extent are rejected. Accepted regions are not split merely to improve routing. Thus this baseline is geometry-driven: it does not claim joint optimality or backpropagate connector failures into the partition.

The tail term in Equation 3 sorts triangle energies in descending order and accumulates exactly 1% of source area. If the last triangle crosses the quota, only the fraction needed to fill it contributes. Sorting is stable under ties. This makes the definition area-based rather than dependent on counting a fixed number of faces, although changing the triangulation can still change the represented geometry and Jacobians.

The conservative area is computed per triangle–pad intersection, not by taking the minimum of two already aggregated totals. In particular,

$$\sum_{p,t} \min(\tilde{a}_{pt}, b_{pt}) \quad \text{is generally different from} \quad \min\left(\sum_{p,t} \tilde{a}_{pt}, \sum_{p,t} b_{pt}\right).$$

The implementation uses the former. Overlapping electrode polygons are rejected before summing their areas. Both routing priorities and row/column board pairing use the resulting distortion-weighted pad values. If either board is unsafe, coverage for the pair is withheld. Thus the scalar cannot reward a short circuit simply because it connects more pads.

For a uniform isotropic scaling by α, $d = |\log \alpha|$ and $E_{\text{all}} = E_{\text{tail}} = (\log \alpha / \log 1.05)^2$. With no seam, $Q = [1 + 2(\log \alpha / \log 1.05)^2]^{-1}$. It tends to zero as $\alpha \to 0$ or $\alpha \to \infty$. This analytic behavior addresses the observed compression loophole in raw coverage, but does not prove resistance to every possible objective exploitation. The hard validity predicate remains necessary.

The agent controller uses stage information in addition to R. For safely evaluated layouts, its staged reward includes a $1 + R$ offset; geometric repair candidates retain separate geometry feedback and other failures have distinct statuses. A best-geometry artifact is retained alongside the optimization incumbent. Manuscript values denoted R always refer to Equation 5, not a stage-offset reward. Missing numerical values remain null. Exact ties are resolved using deterministic keys rather than a tolerance that could change ordering.

Grid indices $-16, \ldots, 15$ map to pins $1, \ldots, 32$ on each axis. The finite extent is $[-49.5, 46.5]^2$ mm before the 0.5 mm edge clearance. Connector containment uses the actual rotated stiffener rectangle. In local connector coordinates its bounds are $[-11.85, 11.85] \times [-4.8, 2.0]$ mm, so its center is offset from the anchor by $(0, -1.4)$ mm. Treating the anchor as the rectangle center changes the feasible-placement test.

The six grid proposals use offsets $(0, 0)$, $(1.5, 0)$, $(0, 1.5)$, $(0.75, 0.75)$, $(-0.75, 0.75)$, and $(1.5, 1.5)$ mm with angles $0, 0, 0, 30, -30, 60$ degrees. If an offset point is outside the outline, its center falls back to the interior reference point. Feasible connector poses are ranked and diversified; the implementation retains up to 12 candidate poses rather than exhaustively routing all sampled poses. Row and column connectors may overlap in projection because they belong to different boards, but this is not evidence that their physical bodies fit in a laminated assembly.

The reported condition-level "area-weighted RMS" diagnostic in the archived summaries averages per-region RMS values with region-area weights. It is not the square root of an average of all squared triangle errors. We avoid using it as a global RMS claim in the main analysis. Likewise, inter-region boundary length counts edges between different labels, including discarded-component labels. It is useful as partition complexity, not as a complete fabrication cost.

C Complete Per-Target Results

Table 5: All 21 targets. Each policy entry is valid-only coverage (%) / valid sheets / connectors on all routed sheets. A dash denotes the timeout, not zero. Target area is rounded to the nearest mm^2.

Part	Area	$K = 1$	$K \leq 2$	$K \leq 4$	$K \leq 8$	Auto
base_link	29,977	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0	5.6 / 2 / 4	–
link_1.0	5,181	0.0 / 0 / 0	0.0 / 0 / 0	17.3 / 2 / 6	17.3 / 2 / 6	17.3 / 2 / 6
link_5.0	5,175	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 8	2.5 / 1 / 8	2.5 / 1 / 6
link_9.0	5,132	0.0 / 0 / 0	0.0 / 0 / 0	13.2 / 1 / 6	21.4 / 2 / 6	21.4 / 2 / 4
link_12.0	4,659	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 6	5.4 / 1 / 8	5.4 / 1 / 8
link_14.0	4,305	0.0 / 0 / 0	0.0 / 0 / 2	11.2 / 1 / 8	16.6 / 2 / 6	16.6 / 2 / 6
link_2.0	3,681	0.0 / 0 / 0	0.0 / 0 / 2	22.7 / 2 / 8	22.7 / 2 / 4	22.7 / 2 / 4
link_10.0	3,679	0.0 / 0 / 0	0.0 / 0 / 2	14.9 / 2 / 8	14.9 / 2 / 4	14.9 / 2 / 4
link_6.0	3,600	0.0 / 0 / 0	0.0 / 0 / 2	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0
link_15.0	2,823	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 2
link_15.0_tip	1,676	0.0 / 0 / 2	0.0 / 0 / 4	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 0
link_3.0_tip	1,674	0.0 / 0 / 2	0.0 / 0 / 4	0.0 / 0 / 2	0.0 / 0 / 0	0.0 / 0 / 0
link_7.0_tip	1,674	0.0 / 0 / 2	0.0 / 0 / 4	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 0
link_11.0_tip	1,673	0.0 / 0 / 2	0.0 / 0 / 4	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0
link_0.0	1,620	0.0 / 0 / 0	14.1 / 1 / 4	14.1 / 1 / 4	14.1 / 1 / 4	14.1 / 1 / 4
link_8.0	1,619	0.0 / 0 / 0	14.3 / 1 / 4	14.3 / 1 / 4	14.3 / 1 / 4	14.3 / 1 / 4
link_4.0	1,572	0.0 / 0 / 0	14.7 / 1 / 4	14.7 / 1 / 2	14.7 / 1 / 2	14.7 / 1 / 2
link_13.0	1,529	0.0 / 0 / 2	0.0 / 0 / 2	14.8 / 1 / 4	14.8 / 1 / 4	14.8 / 1 / 2
link_3.0	1,386	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0
link_11.0	1,374	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0
link_7.0	1,369	0.0 / 0 / 2	0.0 / 0 / 2	0.0 / 0 / 0	0.0 / 0 / 0	0.0 / 0 / 0

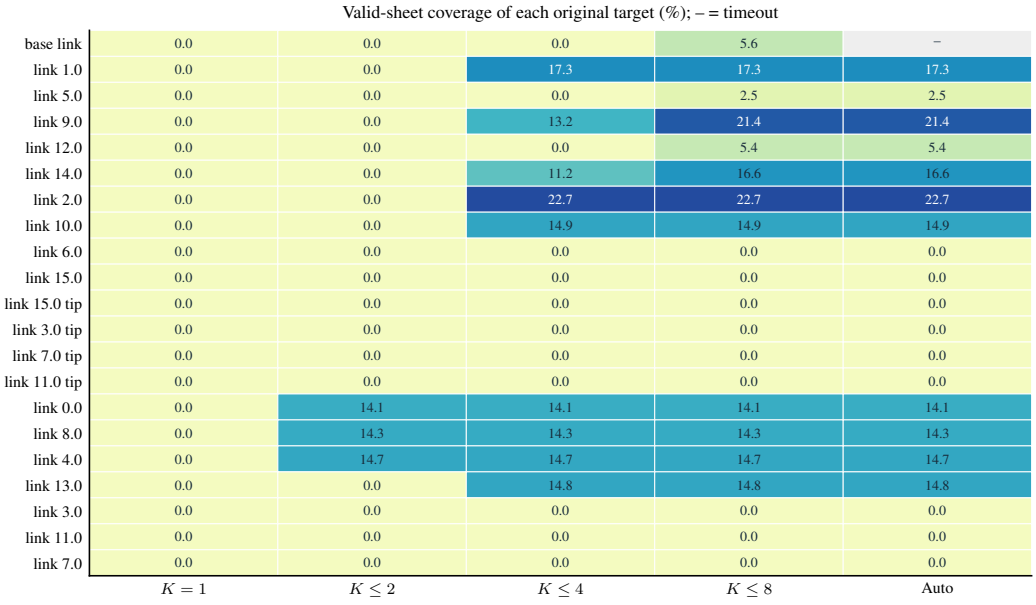


Figure 6: Valid-only target coverage underlying the table. Repeated or related link geometries remain separate physical targets, but do not constitute independent robot families. Color encodes a numerical result, not a significance estimate.

D Outcome and Area Accounting

Table 6: Region outcomes in completed conditions. These are selected-outcome labels with precedence, not independent failure indicators. Dropped fragments are excluded from region counts but remain in coverage denominators.

Outcome	$K = 1$	$K \leq 2$	$K \leq 4$	$K \leq 8$	Auto
Computationally valid	0	3	12	18	16
Electrically incomplete	0	2	8	10	10
Strain exceeds bound	8	17	16	5	0
No sampled connector placement	6	6	33	57	34
No safely evaluated layout	1	4	1	3	1
Rejected by extent	6	5	1	2	0
Total retained regions	21	37	71	95	61

The strain label takes precedence over electrical incompleteness when a capped region violates strain and has a selected safe layout. Therefore, a decline in the strain category does not measure a causal elimination of all geometric problems, and the incomplete category does not count every disconnected layout. Non-disk and flattening-failure rejections are zero among retained regions in these official completed conditions; this does not imply that preprocessing or discarded recursion branches never encountered such geometry.

Table 7: Mean fractions of target area by disposition, computed per part and then averaged across completed parts. These are sheet areas, not electrode coverage. Values may differ slightly from one when summed because of rounding.

Policy	Valid	Designed, invalid	No layout	Rejected	Dropped
$K = 1$	0.000	0.381	0.333	0.286	0.000
$K \leq 2$	0.059	0.480	0.271	0.184	0.007
$K \leq 4$	0.173	0.346	0.385	0.017	0.078
$K \leq 8$	0.206	0.222	0.452	0.020	0.100
Auto	0.209	0.184	0.316	0.000	0.292

All-policy paired raw coverage uses 20 targets and excludes the 29,977 mm^2 palm for every policy. The full $K \leq 8$ result includes the palm. A part with one valid sheet may still have most of its target unsensorized, and a valid sheet itself contains inactive space between electrodes. These distinct denominators explain why the fraction of parts with a valid sheet is much higher than the fraction of target area covered by pads on valid sheets.

The minimum-area threshold and on-sheet stiffener rule are both protocol choices. A connector body fitting in an isolated planar region is not sufficient for a cable to be inserted, routed across a joint, or retained during motion. Conversely, an excluded small region could become sensorizable using a tail to an off-target connector. Neither possibility is evaluated here.

E Method Comparison Details

Detailed development observations. To compare distinct methods beyond segmentation policy, we use three development targets, finger `link_0.0`, fingertip `link_3.0_tip`, and thumb `link_12.0`, keeping two questions separate. *Flattening:* archived 30-second runs of fixed-topology flattening, cut growth, BFF with four and eight cones, and OptCuts with distortion bound 4.05 (Appendix E) give no chart within the 5% bound. Eight-cone BFF has the lowest maximum stretch error on every target (9.4%, 22.1%, and 6.1%), but its thumb chart needs a 114.8 mm square, beyond the 95 mm extent. Cut growth found no valid net and fell back to the fixed map on all three targets. OptCuts meets its own energy bound with maximum stretch error 132–248%, and the fixed thumb map is nearly collapsed ($\sigma_{\min} = 0.002$).

Deterministic versus agent designs: we froze one evaluator configuration and a candidate manifest, then re-evaluated saved artifacts offline without model calls. Evaluator implementation hashes, metric constants, and hardware settings are identical across the earlier circuit study, the agent pilot, and the final matrix; the re-evaluation reproduces all 15 archived agent scores and six archived BFF

results exactly. Each deterministic chart receives a common bank of 18 grid and connector poses. Agent conditions are evaluated with their own saved poses, and, separately, their charts receive the same bank (*common placement applied to saved charts*).

No method yields a computationally valid sheet among 520 evaluated candidates (Table 3). Two agent fingertip designs meet the strain bound (4.4% and 4.5%) by introducing 118–120 mm of internal seams, but only 62 of 165 and 73 of 167 retained pads are jointly reachable. On the finger and fingertip, the best agent design scores above the best deterministic chart (0.233 versus 0.186; 0.117 versus 0.018), mainly through Q_{geom}. These are single observations with unequal budgets (at most four agent layouts per condition versus 18 per deterministic chart). Placement also interacts with chart shape. The common bank raises the finger specialist chart to $R = 0.241$, a deterministic post-processing result rather than an additional agent success, yet admits no pose for four agent conditions whose own poses were admitted. On the thumb, every deterministic chart fails the extent or connector check, and the agents produce one collapsed design and one with no reachable pad. Multi-sheet policies change the hardware budget instead: they produce one valid finger sheet (14.1% valid-only area coverage, four connectors), one valid thumb sheet at $K \leq 8$ (5.4%, eight connectors), and none on the fingertip. Later agent campaigns remain separate (Appendix F).

Protocol. The common-evaluator analysis was frozen before execution: evaluator source `852271b5...`, `metric_coverage_v1` with the constants of Section 3, 32 channels per axis, the 95 mm usable extent, the frozen connector footprint, eight routing attempts per board, and an explicit manifest of 44 evaluation units (28 charts with the common bank and 16 saved agent submissions). The deterministic settings are the archived 30-second Stage-2 charts; no chart was generated. The 18-pose bank combines the six fixed grid proposals of Appendix B, with connectors on the longest outline edges, and 12 chart-independent poses (six angles, two lattice offsets). Targets are verified array-identical to the frozen prepared targets, and every chart preserves the source triangles exactly. Within each analysis and condition, selection uses the maximum R among safely evaluated candidates, then Q_{geom}, then the earliest candidate. Rejected or unevaluated designs have no score. A safe design with no reachable pad has measured zero coverage. All 44 units completed, with no failures or timeouts.

Compatibility. Implementation hashes of the evaluator files, process configuration, pin maps, and metric constants are identical across the Stage-4 circuit follow-up, the agent pilot, the final matrix, and this analysis. Stage-4 placed BFF charts with the same 18-pose bank, and the agent pilot used agent-chosen poses with at most four submissions. Re-evaluation reproduced every archived value. The legacy Stage-2 circuit bank used a different score and two routing attempts, so it appears only as a validity count. Later agent campaigns used revised interfaces and are not additional replicates of the pilot. Equal nominal call caps do not establish matched successful compute.

Consistency with the final matrix. At $K = 1$ the matrix reports no sampled connector placement for the finger, whereas the common bank finds safe layouts on the archived fixed-topology finger chart. An offline diagnostic applying the matrix's stiffener lattice search to both charts explains the difference: the matrix chart is compressed (UV/source area 0.60) and admits no anchor under any of the six grid poses, whereas the archived chart (UV/source area 1.05) admits 501–526 anchors per pose. Single-sheet null outcomes in the matrix therefore describe the charts produced by its solver configuration, not the target.

Table 8: Deterministic flattening settings (archived 30-second runs). Maximum stretch error is $\max(\sigma_{\max} - 1, 1 - \sigma_{\min})$; compression is capped near one for a collapsed chart, so $\sigma_{\min}$ is also shown. "Square" is the smallest square containing the chart over all rotations (usable: 95 mm). "Selected solver" is the method actually used; cut growth fell back to the fixed map on every target. OptCuts satisfies its symmetric-Dirichlet bound of 4.05, which is not the 5% maximum-strain bound.

Target	Setting	Max err.	$\sigma_{\min}$	In 5% (% area)	Q_{geom}	Seam mm/tips	Square mm	Selected solver
finger	Fixed topology	0.506	0.494	60.1	0.0191	0.0/0	60.0	metric_L4
finger	Cut growth	0.506	0.494	60.1	0.0191	0.0/0	60.0	metric_L4 (fallback)
finger	BFF, 4 cones	0.196	0.804	91.3	0.0567	10.7/4	60.5	BFF
finger	BFF, 8 cones	0.094	0.906	95.7	0.1483	16.0/6	60.3	BFF
finger	OptCuts, bound 4.05	1.320	0.591	86.7	0.0191	0.0/0	59.5	OptCuts
fingertip	Fixed topology	0.481	0.519	0.0	0.0032	0.0/0	32.8	SLIM_ARAP
fingertip	Cut growth	0.481	0.519	0.0	0.0032	0.0/0	32.8	SLIM_ARAP (fallback)
fingertip	BFF, 4 cones	0.376	0.727	11.3	0.0074	51.0/2	56.6	BFF
fingertip	BFF, 8 cones	0.221	0.829	25.0	0.0158	78.1/1	58.3	BFF
fingertip	OptCuts, bound 4.05	1.566	0.870	19.0	0.0073	38.9/1	54.0	OptCuts
thumb	Fixed topology	0.998	0.002	0.0	0.0000	0.0/0	2.6	SLIM_ARAP
thumb	Cut growth	0.998	0.002	0.0	0.0000	0.0/0	2.6	SLIM_ARAP (fallback)
thumb	BFF, 4 cones	0.513	0.487	96.9	0.0137	19.2/3	115.1	BFF
thumb	BFF, 8 cones	0.061	0.947	100.0	0.3869	25.8/5	114.8	BFF
thumb	OptCuts, bound 4.05	2.478	0.554	85.8	0.0128	0.0/0	110.1	OptCuts

Table 9: Selected candidate per method and target under the common evaluator. Coverage values are fractions of the part target: weighted C_q, conservative, and raw jointly reachable area. Pairs: jointly reachable/retained pads. "A:" denotes an agent condition, and "+ bank" denotes common placement applied to its saved charts.

Target	Method	R	Q_{geom}	C_q %	Cons. %	Raw %	Strain	Pairs	Calls/tokens
finger	Fixed topology	0.023	0.0191	19.5	29.6	29.8	0.51	126/168	0
finger	Cut growth (fallback)	0.023	0.0191	19.5	29.6	29.8	0.51	126/168	0
finger	BFF, 4 cones	0.071	0.0567	25.1	29.5	30.2	0.20	125/159	0
finger	BFF, 8 cones	0.186	0.1483	25.4	28.4	28.8	0.09	120/162	0
finger	OptCuts, bound 4.05	0.023	0.0191	21.8	23.8	23.9	1.32	100/161	0
finger	A: direct	–	–	–	–	–	–	–/–	1 / unknown
finger	A: generic	0.160	0.1328	20.5	22.4	22.4	0.17	95/164	4 / 228,799
finger	A: specialists	0.233	0.1944	20.0	21.3	21.3	0.18	90/162	4 / 168,190
finger	A: direct + bank	–	–	–	–	–	–	–/–	1 / unknown
finger	A: generic + bank	–	–	–	–	–	–	–/–	4 / 228,799
finger	A: specialists + bank	0.241	0.1944	24.0	26.7	26.8	0.18	112/166	4 / 168,190
fingertip	Fixed topology	0.0032	0.0032	0.2	21.0	42.9	0.48	91/91	0
fingertip	Cut growth (fallback)	0.0032	0.0032	0.2	21.0	42.9	0.48	91/91	0
fingertip	BFF, 4 cones	0.0079	0.0074	7.5	34.8	41.0	0.38	171/176	0
fingertip	BFF, 8 cones	0.018	0.0158	12.3	34.2	37.3	0.22	159/169	0
fingertip	OptCuts, bound 4.05	0.0083	0.0073	12.8	39.5	41.3	1.57	176/176	0
fingertip	A: direct	0.0060	0.0060	0.4	27.9	42.3	0.60	121/121	1 / 28,824
fingertip	A: generic	0.114	0.1028	11.2	16.7	17.1	0.04	73/167	4 / 176,309
fingertip	A: specialists	0.117	0.1057	10.4	13.9	14.1	0.04	62/165	4 / 151,086
fingertip	A: direct + bank	0.0060	0.0060	0.4	28.1	42.5	0.60	120/120	1 / 28,824
fingertip	A: generic + bank	0.0062	0.0061	0.5	28.7	42.0	0.65	123/123	4 / 176,309
fingertip	A: specialists + bank	–	–	–	–	–	–	–/–	4 / 151,086
thumb	Fixed topology	–	–	–	–	–	–	–/–	0
thumb	Cut growth (fallback)	–	–	–	–	–	–	–/–	0
thumb	BFF, 4 cones	–	–	–	–	–	–	–/–	0
thumb	BFF, 8 cones	–	–	–	–	–	–	–/–	0
thumb	OptCuts, bound 4.05	–	–	–	–	–	–	–/–	0
thumb	A: direct	0.0074	0.0074	0.0	0.0	0.0	1.15	0/324	1 / 35,166
thumb	A: generic	1.3×10^{-4}	0.0001	0.1	23.6	33.8	80.91	772/772	2 / unknown
thumb	A: specialists	–	–	–	–	–	–	–/–	3 / 117,283
thumb	A: direct + bank	–	–	–	–	–	–	–/–	1 / 35,166
thumb	A: generic + bank	–	–	–	–	–	–	–/–	2 / unknown
thumb	A: specialists + bank	–	–	–	–	–	–	–/–	3 / 117,283

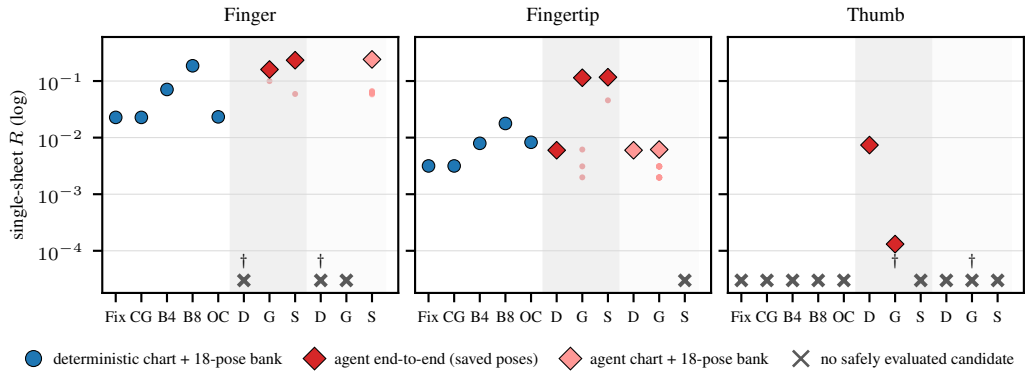


Figure 7: All safely evaluated candidates (small markers) and the selected candidate (large markers) per method. Fix, CG, B4, B8, and OC denote fixed topology, cut growth (fallback), BFF with four and eight cones, and OptCuts. D, G, and S denote direct, generic, and specialist agents: dark diamonds use their own saved poses, and light diamonds use the common bank on their charts. Crosses mark methods without a safely evaluated candidate (not zero); † marks interrupted conditions. Log scale; no candidate is computationally valid.

F HISTORICAL NUMERICAL AND AGENT CAMPAIGNS

Table 10: Development evidence, kept separate from the final matrix. Counts refer to the archived campaign units and are not a pooled success rate. No row supports a statistical architecture ranking.

Campaign	Scope	Recorded result and interpretation
Numerical single-sheet development	Three representative targets; fixed charts, BFF, OptCuts, cut heuristics	No valid full-part design. BFF eight-cone charts reached 9.4% worst strain on the finger and 6.1% on the thumb; the latter exceeded the fixed extent (Appendix E).
Agent pilot	Direct, generic, two specialists; single replicate	17 submissions, no valid design. Best reported finger specialist $R = 0.233$, with 90/162 pads reachable. Missing or rejected outputs remain distinct from evaluated failures.
Protocol block 1	Revised feedback and contract; three targets	Seven evaluated designs, no valid result; best reported $R = 0.073$ for a direct finger design. Several conditions interrupted or unstarted.
Protocol v1.1 follow-up	Finger specialists and thumb conditions	One evaluated design, $R = 0.223$, invalid. Thumb conditions lack completed evaluation; no zero-score substitution.
Early multi-sheet pilot	Three targets	Motivated the final matrix; used a different segmentation implementation and is not pooled.

The agent pilot used the archived `gpt-6-astra` model identifier with tools disabled during program generation. Programs executed afterward in a numerical sandbox. Direct construction did not receive a numerical baseline chart as its starting solution. Feedback-loop variants could revise their constructions using evaluator output. The intended call caps were one for direct, four for generic, and two geometry–layout pairs for specialists. These nominal caps are not realized matched resource allocations when calls fail or time out.

Readiness audits identified contract discoverability and information-flow issues in early interfaces. Subsequent releases added prior-artifact access and more spatial feedback. Those changes matter when interpreting the specialist hypothesis: an apparently unproductive geometry role could reflect unavailable layout information, failure to use available information, or an intrinsically difficult design. The archived campaigns do not isolate these explanations. We therefore include the agent results as unsuccessful or partial development evidence rather than claiming that specialization helps, that general-purpose agents solve the task, or that their failures establish computational hardness.

G From Local Sheets to a Connected Glove Architecture

Figure 5 connects three scales of the sensorization process. These are separate exploratory artifacts with different evidence boundaries: a whole-hand packaging proposal, an evaluated local finger-link circuit, and a schematic for sharing electrode returns. They do not constitute an end-to-end validated glove, and their results are not pooled with the frozen multi-sheet or method-comparison tables.

G.1 Face Patches and Connecting Patches

The assembly studies progress through three levels of integration. A *face patch* is a local sensing region; a *connecting patch* is a non-sensing substrate bridge carrying conductors between regions. Neither term denotes an individual mesh triangle. Throughout, opposing row and column electrodes occupy separate FPCBs, with return routing on each board's reverse side. Integrating the assembly does not collapse this sandwich into one electrical layer.

Face-patch assembly. The initial concept places small sensing islands on suitable part surfaces and connects their row and column tails to distributed readout reservations (Figure 4(a)). The illustrated target set contains 108 face regions and 48 fingertip-cap tiles. Local pieces can simplify fitting and allow replacement of a damaged region, but their integration requires tails, readout sites and connections across joints. The archived assembly uses 17 readout reservations and 32 joint-connector housings, together with printed branches and wire exceptions. These are different counts: a sensing patch is not necessarily assigned its own board or plug connector. The target-colored area represents intended placement, not electrically verified coverage.

Connected-patch assembly. This stage retains the sensing islands but joins them through non-sensing flex bridges, with shared backside buses where the address allocation permits. A bridge carries signals and does not earn sensing-area credit merely by covering the robot. In the continuous-finger trial, replacing internal jumper interfaces reduces joint-connector housings from 32 to eight, while retaining the 17 distributed readout reservations. A subsequent trial consolidates readout to one board per finger plus one palm board: five reservations in total. These changes are distinct—removing connectors does not automatically remove readout electronics. The resulting illustrations still include slack joint loops and unresolved wire exceptions. Connecting islands also couples their installation and repair, and concentrates mechanical demands in the bridges. A continuous two-board finger assembly means paired row and column flex blanks, not one shared copper sheet.

The current integrated-flex stage is described in the main discussion alongside Figures 4(c) and 5(a).

These stages explain architectural choices rather than a controlled performance ranking. Their hardware and interconnect assumptions differ, so board or connector reductions must not be interpreted as measured gains in tactile coverage or pooled with the fixed-hardware experiment scores.

G.2 Copper-Routed Flex Continuation Trial

The `local_flex_continuation_v1` extension changes the circuit boundary condition. A local board or connector no longer occupies each region: row and column returns terminate at labeled cross-sections of an unpopulated flex tail, for eventual continuation through the shared carrier. The implementation reuses the finite-width connectivity and isolation checks, but has a separate interface schema and does not instantiate a commercial connector, anchor, courtyard or stiffener. A terminal section is not evidence of a completed route to the hand-base electronics.

The three inputs are the unchanged Figure 3 charts and front-side pad sets. Eight outward radial directions and two tail widths are screened for a connected, hole-free planar support. A full-width strip joins the boundary to a continuation bank 3 mm outward. Initial bank pitch is 0.5 mm with 0.30 by 1.0 mm interface sections; traces retain 0.20 mm width, 0.15 mm clearance and 0.50 mm edge clearance. Via/drill diameters remain 0.70/0.30 mm. A 0.00001 mm substrate guard prevents numerical erosion from clipping unchanged pad boundaries; it is included in added-area accounting.

Each eligible tail receives two routing orders per axis. The best recorded directions by row, column or joint reachability are refined with eight attempts. The thumb additionally receives 16 attempts on its best direction and a 0.0625 mm routing-grid follow-up; the nominal grid is 0.125 mm. A 1.0 mm bank-pitch follow-up on the thumb and fingertip does not improve the selected joint count. Row and column candidates can use different exits, because they occupy separate FPCBs. We search their

pairings and replay the saved copper on the union support, maximizing joint pad count and breaking ties by added substrate area. This is finite search, not a proof of optimal routing or of impossibility for missed pads.

Table 11: Local continuation trial on fixed pad sets. The FFC column is the archived local-connector reference; the continuation column measures reachability to both bus interfaces. Added area is per FPCB, not sensing area. Whole-hand base reachability is unmeasured in all rows.

Region	Local FFC	Flex continuation	Added flex (mm^2)
Finger `link_2.0`, Auto 0	98/98	98/98	39.15
Thumb `link_12.0`, Auto 2	64/64	63/64	44.85
Fingertip, specialists round 2	62/165	96/165	44.65

Selected copper passes isolation, substrate containment, pin-label consistency and complete-annulus checks. Tests deliberately sever returns, remove vias, short nets, duplicate an assignment and place copper outside the support; these cases do not earn false connectivity. Pads remain in the denominator when unreachable, and neither extra tails nor logical whole-hand addresses earn sensing-area credit. The saved thumb archive lacks its per-triangle UV array, so we do not reconstruct a new map or report newly integrated source-area coverage for it. Source/outline correspondence and the original strain are inherited from the archive.

The comparison isolates a practical design change while allowing additional search and substrate; it is not a matched-compute algorithm ranking. The thumb loses one pad under the best tested continuation layout, demonstrating that eliminating the local footprint alone does not ensure a superior route. Global channel assignment, junction crossings, transitions between parts, motion, analog performance and the physical fit of the added tails remain unqualified. The whole-hand image continues to be a packaging proposal. Frozen local-connector scores and validity classes are unchanged. Run scripts, histories, selected copper and replay checks are under `experiments/flex_bus_escape` and `output/flex_bus_escape_20261028`.

G.3 Shared Addressing and Qualification

Connecting patches without a connector at every patch. Each sensing region retains opposing row and column electrodes on separate FPCBs. Their reverse sides carry the respective return traces. A shared row can continue through adjacent patch substrates and non-sensing bridges, with vias tapping the sensing-side conductor, rather than terminating at a separate connector for each patch. In the illustrated subset, two patches share rows R0–R3 while using columns C8–C10 and C17–C19, respectively. Distinct column labels avoid duplicate row–column addresses across those patches. This is a logical allocation rule, not a proof of crossing-free routing or ghost-free electrical readout. The yellow sensing-side links and solid backside buses in the schematic belong to different copper surfaces; row and column circuits belong to different FPCBs.

Whole-hand integration proposal. The new carrier preserves the 108 face-region and 48 cap-tile target meshes, removes the previous external harness and fixture objects, and reduces an intermediate five-readout arrangement to one 70×48 mm base-board reservation. Shelf packing preserves the four finger maps internally and assigns the palm and fingers to 127 global rows and 140 global columns. All 11,816 declared bounding-grid addresses are unique, and every target has a logical path through the robot's link tree to the base. These addresses are capacity reservations, not evaluated sensing pads. Eight nominal 16-channel row mux banks and nine column banks reserve 128/144-channel capacity. The illustration does not establish component selection, pin escape, analog performance, or detailed routing through the palm and joint bridges. This larger readout is a new hardware protocol, not the 32-channel-per-board protocol used in the scored experiments.

A locally connected finger-link region. Panel (b) now uses the connector-free continuation result on Auto region 0 of finger `link_2.0`, also shown in Figure 3. All 98 retained pads reach both bus interfaces. The unchanged region has 3,991 source faces and $1{,}025.44$ mm^2 of surface; worst strain is 2.60%. Mapping its retained pads through the archived UV chart gives 386.38 mm^2, or 37.68% of the region, unchanged from the complete local-FFC pad set. The added 39.15 mm^2 per-board tail is not mapped onto the source or credited as sensing area. This replaces the earlier archival `link_1.0` illustration so the local circuit and current assembly proposal use consistent connector-

free terminology. Neither full retained-grid reachability nor the source-mapped display establishes whole-link coverage or base-board connectivity.

Mechanical and evidential limits. The glove render contains thin rigid-link carrier shells and sixteen pose-dependent joint webs instead of projecting service loops. Webs are reconstructed independently in open and curled poses, with their ends projected onto parent and child surfaces; they are not a constant-length material simulation. The modeled close-fitting paths lengthen by up to 56.4% between these poses. This is a geometric accommodation demand, not predicted copper strain: ordinary bonded FPCB cannot be assumed to follow it. Sliding interfaces, compliant interconnects, attachment, fatigue, interference and a manufacturable planar blank remain unresolved. The new scene's absence of wires is therefore a proposed replacement architecture, not evidence that the legacy wiring failures have been electrically solved. Three integration tests check target preservation, unique address allocation and link-tree paths, and the saved motion calculations; they do not qualify fabrication.

The saved sources are `glove_base_readout_20261027/run_v1`, the flex-continuation result on region 0 of `link_2.0__auto`, and `backside_shared_bus_20261027/run_v3`. The figure script composes saved local-routing outputs and redraws the address-sharing schematic without rerouting during rendering; a separate manifest records source and output hashes.

H Physical Qualification and Responsible Interpretation

No physical sensor, force calibration, robot manipulation task, or durability measurement was performed in this project. The two-board stack is a design specification, not a tested device. In particular, the 5% mapping threshold does not establish that copper traces, vias, or solder joints tolerate that strain. Stiffener placement in a planar chart does not establish contact with a curved substrate or a collision-free connector assembly.

The existing inextensible-fitting path supports only small meshes and was not qualified for the dense targets in this study. Saved visual reports correctly mark the fit panels unavailable. Our source-mesh plots display where a region came from; they do not show a mechanically simulated wrap. A future physical study would need material and stack-up choices, assembly tolerances, dense fitting or actual fit tests, cable clearance, manufacturer review, and electrical and tactile characterization.

The manuscript consequently uses "computationally valid" only for the evaluator's implemented criteria. It avoids "manufacturable," "deployable," or "full-hand sensorization" as conclusions. The evidence establishes a reproducible computational design study under restrictive hardware, with explicit failures and a practical motivation for connector-aware surface partitioning.