Feasible Is Not Faithful: Task Preservation in Feasibility Restoration

Jeonghyeon Kim and Shinsuk Park

***Abstract*—Restoring feasibility can sacrifice task coordinates: tracking a corrected reference, regulating force and keeping contact do not show that the coordinates a task prescribes were preserved. We analyze this choice for static affine corrections and test it on contact with unknown surfaces. A correction preserves task-selected coordinates Lx when its directions lie in $\ker L$, and it can reach every constraint when the constraint is reachable through that subspace. For any other correction direction B with AB invertible, the preserved coordinates move by $G(B)\rho$, $G(B) = -LB(AB)^{-1}$, where ρ is the mismatch removed. Intent-preserving constraint realization (IPCR) implements the condition with joint sensing only, estimating no surface, normal or contact event; drawing and box transport share one setting. On hardware, the nearest-point (NP) shift of a drawn circle on an 18° board reached 1.04–1.06 of a prediction fixed before the runs (21.2–21.6 mm, against 0.21 mm for IPCR from the encoders), and on an 11° board video of the ink confirmed the NP–IPCR offset independently of the encoders. In registered simulation, structures given the true normal and an ideal force sensor moved the figure 10–36 times as far as IPCR at 0.90–1.09 of their predicted shift; a crossed design showed that with a task-aware target, absorbing the residual along the normal still leaks the predicted amount; and fixed-axis hybrid control, which meets the condition, preserved placement as well. Blind predictions on new boards failed in four of twelve cells; after an audit fixed two input constants, a separate prospective set on unseen boards passed.**

***Index Terms*—Force control, compliance and impedance control, contact modeling, motion control.**

I. Introduction

REFERENCES planned from nominal geometry, demonstrations or policies trained elsewhere become infeasible on the actual surface, and restoring feasibility means choosing which task coordinates to sacrifice; controllers make this choice implicitly. Suppose a robot draws a circle at a position fixed in an external frame, as when marking drill positions or dispensing along a seam, on a board tilted by 11°. Nearest-point realization on our arm, even given the fitted normal, moved the circle 16.6–17.1 mm from its planned place, about half its radius (Fig. 1). The strokes finish and the force is regulated, but tracking and force errors are measured against the corrected point and do not reveal where the figure landed. In simulation at the same speed both rules kept contact on 100 % of the figure, while the figure moved 0.25–0.30 mm against 8.1–12.7 mm (Sec. V-E). Feasibility leaves a task decision open: which coordinates may change to restore contact?

The decision arises in CBF safety filters, constrained diffusion policies, MPC tracking infeasible references, and compliant controllers and learned policies [4]–[7], [29]. Our results concern static affine corrections, such as a minimum-norm projection onto one active plane, not filters whose active constraints switch or MPC that optimizes over time; contact with an unknown surface is where we test them. Preserving task intent is an explicit goal of recent contact control [1], [2]. Hybrid control, oblique projection and constraint-based specification allocate directions or outputs to keep [10], [11], [15], [25], and adaptive control learns the normal for surface-defined goals [3], [9], for which nearest-point (NP) realization is the right rule. Drawing systems that yield along the normal [26] or absorb scanned-surface deviations of up to 5 cm along the pen [13] report shape or line quality, not where the figure landed. What remains unstated is a task-relative characterization that links preservation, reachability and the displacement caused by a correction that changes prescribed coordinates. Better geometry cannot supply it: an exact normal can still direct motion into coordinates the task requires fixed.

We base the choice on the task. The task names *preserved coordinates* Lx and a *yield subspace* $\mathcal{D} = \ker L$: drawing may yield in depth but must keep the figure's frontal position; transport may yield sideways but must keep progress, height and possibly yaw. IPCR (Fig. 2) applies commands, corrections and force readings through $\mathcal{D}$, so estimation errors affect depth, yaw and timing but not commanded Lx; one proprioceptive setting serves drawing and transport without a surface, normal or contact-event estimate. Direction allocation is classical; our contribution characterizes which task coordinates feasibility restoration preserves, when preservation is feasible, and how far another correction direction moves them, with predictions tested across execution structures. A prediction fixed in advance held on hardware: at 18°, NP's shift reached 1.04–1.06 of the prediction from a plane measured before the runs.

Contributions.

1) A characterization of static affine corrections: preservation iff $LB = 0$ for spanning correction amounts (Theorem 1) and reachability iff $\operatorname{rank}(AD) = k$ (Theorem 2). For an alternative direction B with AB invertible, the closed-form displacement $G(B)\rho$, with nearest-point drift and residual leak as instances (Proposition 1, Corollary 1), and a projection metric that recovers the preserving target.
2) A proprioceptive implementation. Continuous contact weighting, force-regulated depth speed and permitted yaw use no surface or contact-event estimate; drawing and transport share settings.
3) Evaluation on contact. Hardware: drawing on two board tilts (26 runs), one with the NP prediction fixed before execution, an encoder-independent ink check, and box transport with fixed and yielding yaw. Registered simulation: structures from the literature with ideal sensing, a 2×2 target-absorption ablation that separates the

Preprint, November 2026. Supplementary material follows the references.

J. Kim and S. Park are with the School of Mechanical Engineering, Korea University, Seoul, Republic of Korea. Corresponding author: Shinsuk Park (drsspark@korea.ac.kr).

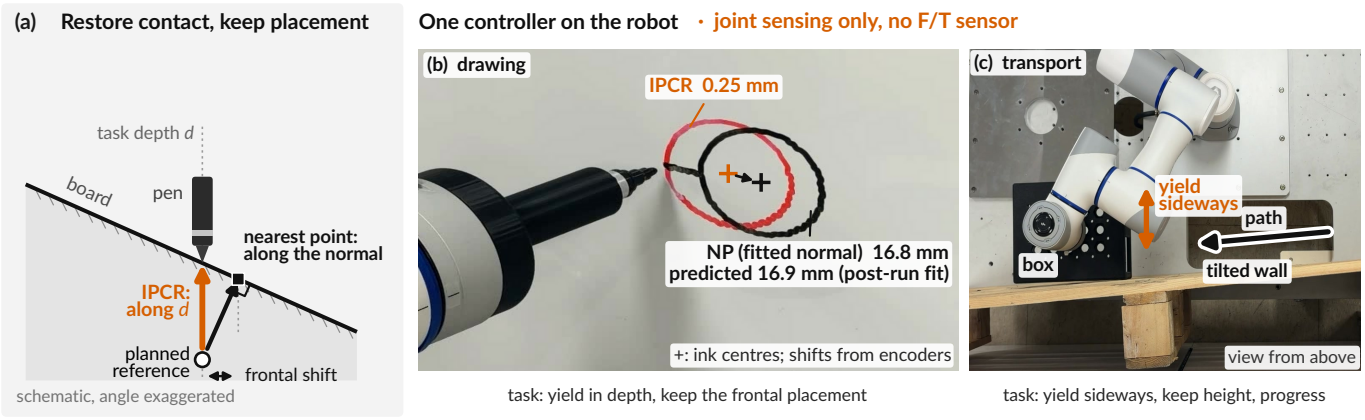

Fig. 1. Restoring contact without moving the task. (a) Schematic top view, board angle exaggerated: from a planned reference behind the board, the nearest point (NP) corrects along the normal and shifts the figure in the task's frontal coordinates; IPCR corrects only along the task depth $d \in \ker L$ and keeps them. (b) The robot, joint sensing only, on a board tilted by about $11°$: the circle drawn by IPCR (red), then by NP given the fitted normal (black; the short stroke on the left is its tip sliding from the planned start); encoder-based shifts 0.25 and 16.8 mm, the latter against 16.9 mm predicted from the plane fitted after the runs (the prior prediction at $18°$ is in Fig. 3); crosses mark the fitted ink centers, the planned center is not registered in the photo. (c) Box transport along a tilted wall with the same controller and setting, seen from above: the task lets the box yield sideways, keeps height and progress, and sets whether yaw may yield.

two mechanisms, curved boards and transport, and new boards where failed blind predictions led to an input audit and a separate prospective test.

II. Related Work

Hybrid control and oblique projection. Raibert–Craig's selection matrix S selects position-controlled directions in a constraint frame; Mason's natural/artificial constraints allocate motion and force [11], [15]–[17]. Our task fixes L and a basis D of $\mathcal{D}$: a fixed task frame shares this allocation, whereas a normal-defined frame gives $B = n$, leaking $G(n)\rho$ into external Lx. Lipkin–Duffy's metric selects the correction complement B in $P = I - B(AB)^{-1}A$ [10], [18]; Proposition 1 quantifies its cost in Lx. Subsequent decoupling and compliance formulations refine these classical choices [12], [19]–[22]. iTaSC realizes $x_D^\star$ by constraining Lx and freeing $\mathcal{D}$ [25]. Allocation, the projector and the rank test are classical foundations. Our contribution is their task-relative characterization: preservation iff $\operatorname{range} B \subseteq \ker L$ for spanning mismatches, reachability iff $\operatorname{rank}(AD) = k$ over expected surfaces, the closed-form cost $G(B)\rho$ and principal-angle drift, and registered prospective tests, including failures.

Unknown geometry. A wrong normal misplaces and distorts the control target [8], [9]. Slope identification from joint signals alone corrects along the normal [3], and constraint estimation remains central to hybrid control [23], [24]. These works define the task on the surface, where a better normal helps. A task fixed off the surface needs no normal for its command; it needs the directions in which it may yield.

Least-squares realization. A control-barrier-function program returns the min-norm correction along the constraint gradient [4], as do constrained diffusion policies [7]. Operational-space barrier functions weight it in task space [6], yet on one plane they still correct along a weighted normal.

Drawing and learned contact policies. Song et al. keep the planned in-plane position and comply along the surface normal after touching each control point [26], or along the pen from a modeled surface [13]. This pairs a target step with an absorption step, but whether that compliance moves the in-plane position under a persistent residual was not evaluated; shape, not placement, was reported. Weingartshofer et al. map patterns onto a known model in two ways: parallel projection, run with pure motion control and the pen held along the projection axis, and conformal mapping, run with motion or hybrid force control along the local normal. They report the pen slipping and leaving the planned path where the projection leaves a large angle to the surface [27]. Their parallel projection gives the target of Theorem 2; it was not combined with force-regulated contact, and the surface was known. Learned policies predict an interaction frame [28], force or stiffness modes [29] or a virtual target [30], or resolve forces in an estimated contact frame [2]; model-based admittance uses motor-current force estimates [14], and a proprioceptive reflex is trained toward a Euclidean projection [1]. None states which coordinates a correction may change or what changing the wrong ones costs.

III. Intent-Preserving Constraint Realization

Problem. A reference path $r(s) \in \mathbb{R}^n$, parameterized by progress s, is infeasible for an affine constraint $\mathcal{C} = \{x : Ax = b\}$, $A \in \mathbb{R}^{k\times n}$ of full row rank. The task fixes L and the yield subspace $\mathcal{D} = \operatorname{range}(D) = \ker L$, so $LD = 0$, and seeks motion that meets $\mathcal{C}$ while keeping Lx close to $Lr(s)$. The theory concerns any correction within this static affine model; our contact implementation receives neither A nor b and also regulates a contact-force component.

A. Preservation and Realization

Theorem 1 (Preservation)**.** *Let $x = r + B(y)\alpha$, where B may depend on the measurements y, and let $\mathcal{S}_y$ be the set of coefficients α attainable for that y. Then $Lx = Lr$ for*

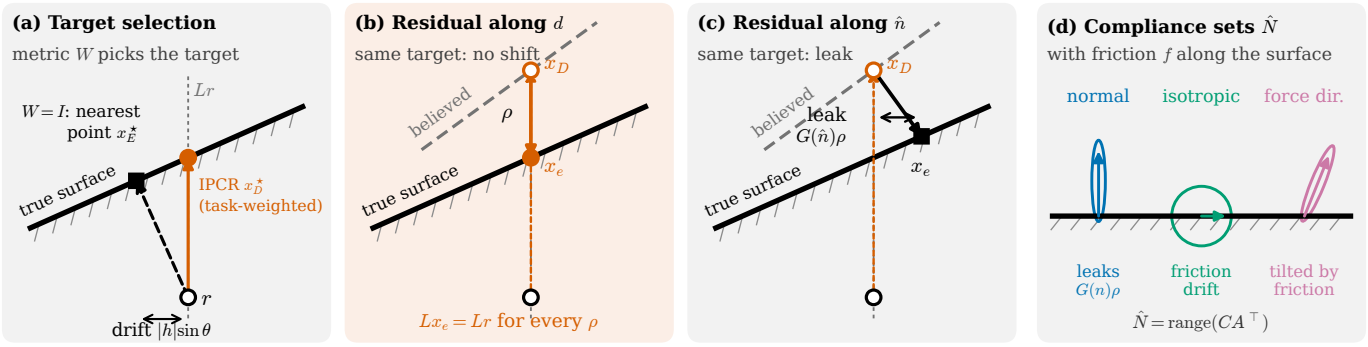

Fig. 2. Correction directions determine placement. (a) A projection with metric $W = I$ gives NP, $x_E^\star$, which shifts preserved coordinates by $|h|\sin\theta$ (Corollary 1); a task-weighted W gives $x_D^\star$, a correction along d (Theorem 2). (b, c) Residual absorption along d preserves $Lx = Lr$; along $\hat{n}$ it adds $G(\hat{n})\rho$ (Proposition 1). (d) Compliance-dependent absorption (Eq. (4)).

every $\alpha \in \mathcal{S}_y$ iff $LB(y)\alpha = 0$ on $\mathcal{S}_y$. If $\mathcal{S}_y$ spans $\mathbb{R}^m$, this is equivalent to $\operatorname{range} B(y) \subseteq \ker L$.

Proof. $L(x - r) = LB\alpha$, so the first claim is immediate; on a spanning set, linearity gives $LB = 0$. □

In static contact, $Ax = b$ with AB nonsingular gives $\alpha = -(AB)^{-1}\rho$ with $\rho = Ar - b$, so the span condition holds when the mismatches consistent with the same y span $\mathbb{R}^k$.

Theorem 2 (Realization)**.** *Let the columns of D form a basis of* $\ker L$*. For every b there is an x with $Ax = b$ and $Lx = Lr$ iff* $\operatorname{rank}(AD) = k$*. It is unique iff AD is square and nonsingular, and then*

$$x_D^\star = r + D(AD)^{-1}(b - Ar). \tag{1}$$

Proof. $Lx = Lr$ iff $x = r + D\alpha$; the constraint becomes $AD\alpha = b - Ar$. A solution exists for every right side iff AD is onto, and is unique iff its null space is trivial. □

The projector $I - D(AD)^{-1}A$ is standard; here D comes from the task. For one contact, $\mathcal{D} = \operatorname{span}(d)$ with unit d and unit normal n requires $n^\top d \neq 0$: the correction is amplified by $1/|n^\top d|$, and the depth must change at $|n^\top \dot{r}|/|n^\top d|$, at most $v\tan\theta$ for a path at constant depth. This is a geometric condition; unilateral contact, friction and drive limits add conditions on the equilibrium that it does not replace. Solutions with $\dim\mathcal{D} > k$ exist but are not unique: U-yaw transport has $k = 1$ and two yield directions, and the yaw law selects among them. Every command $x_{\mathrm{cmd}} = r(s) + D\alpha$ preserves $Lr(s)$ despite errors in the estimate setting α and s; Theorem 2 characterizes contact reachability, while convergence depends on the feedback.

B. What Another Direction Costs

Proposition 1 (Transfer)**.** *Let $x = z + B\nu$ meet $Ax = b$ with AB nonsingular, and $\rho_z = Az - b$. Then*

$$L(x - r) = L(z - r) + G(B)\rho_z, \qquad G(B) = -LB(AB)^{-1}. \tag{2}$$

With $Lz = Lr$ this vanishes for every ρ_z iff $LB = 0$.

Proof. $AB\nu = -\rho_z$; apply L. As ρ_z ranges over $\mathbb{R}^k$ so does ν. □

With $z = r$, the law predicts full-correction drift; with $z = x_D$ on a believed surface, it predicts absorption of the remaining mismatch. Preservation without relying on cancellation therefore requires both target placement and residual absorption to respect $\ker L$. Let NP be $x_E^\star = r + A^\top(AA^\top)^{-1}(b - Ar)$, i.e. $B = A^\top$.

Corollary 1 (Nearest-point drift)**.** *With AD nonsingular, $\Delta = x_E^\star - x_D^\star \in \ker A$ and $\|\Delta\|^2 = \sum_i (u_i^\top \bar{e})^2 \tan^2\theta_i$, where θ_i are the principal angles between* $\operatorname{range}(A^\top)$ *and $\mathcal{D}$, Q_A and Q_D orthonormal bases with $Q_A^\top Q_D = U\cos\Theta V^\top$, u_i the columns of U and $\bar{e} = Q_A^\top(x_E^\star - r)$. For one plane violated by the normal depth h, $\|\Delta\| = |h|\tan\theta$ and, for orthonormal L, $\|L(x_E^\star - r)\| = \|G(A^\top)\rho\| = |h|\sin\theta$.*

Proof. $x_D^\star - r = Q_D(Q_A^\top Q_D)^{-1}\bar{e}$ has component $Q_A\bar{e} = x_E^\star - r$ in $\operatorname{range}(Q_A)$, so $\|x_D^\star - r\|^2 = \|\bar{e}\|^2 + \|\Delta\|^2 = \sum_i (u_i^\top \bar{e})^2/\cos^2\theta_i$. The last identity is (2) with $LL^\top = I$, $\|Ln\| = \sin\theta$. □

The drift belongs to the rule, not to estimation error: it exists with the exact constraint and grows with θ. For a single static projection $\min_{Ax=b}\|x - r\|_W^2$, as in a safety filter with one active plane, $W = I$ gives the nearest point and its drift, whereas $W = wL^\top L + DD^\top$ approaches $x_D^\star$ as $w \to \infty$ when AD is nonsingular (Supplementary S6, Sec. VI); this fixes the target, and residual absorption must also respect $\ker L$. NP is the right correction for a task with $Ln = 0$, i.e. $\mathcal{D} = \operatorname{span}(n)$.

A robot misses and slips. If the executed $x = r + B\nu + \varepsilon$ has contact error $\eta = Ax - b$ and other tracking errors ε, then

$$L(x - r) = G(B)(\rho - \eta) + L\big[I - B(AB)^{-1}A\big]\varepsilon. \tag{3}$$

With $LB = 0$, only the tracking term $L\varepsilon$ remains; bounding it requires position feedback. Admittance, filters and NMPC also make B state-dependent. We therefore use (2) as a static prediction and (3) to read deviations from it.

C. Compliance, Information and Construction

Where B comes from. A servo with fixed quasi-static compliance $C \succeq 0$ about its target x_t, infinitely stiff on $\ker C$,

settles on the true constraint at $x = x_t + C(A^\top\lambda + f)$, $Ax = b$, where f collects tangential contact wrenches such as friction. Eliminating λ gives, with $\rho = Ax_t - b$ and $H_C = C - CA^\top(ACA^\top)^{-1}AC$,

$$L(x - x_t) = G(CA^\top)\rho + LH_C f. \tag{4}$$

Thus $B = CA^\top$ is built from the *true* constraint. (i) If $\operatorname{range} C = \mathcal{D}$ and an equilibrium satisfying unilateral contact and the friction law exists, it is the IPCR point of the true constraint for any normal estimate and friction; conversely, for one frictionless contact with a fixed C and nonzero reaction, preservation for all normals in an open set forces $LC = 0$ (if C is chosen from n this fails). (ii) Rank-k compliance such as $c\hat{n}\hat{n}^\top$ has $H_C = 0$ and leaks $G(\hat{n})\rho$. (iii) Isotropic compliance adds cLf, and compliance along the sensed force leaks $\mu|h|$ while sliding. For finite stiffness, $\|L(x - x_t)\| \le \|LC\|\,\|A^\top\lambda + f\|$.

What the command must know. With L fixed and $\operatorname{range} B \subseteq \ker L$, the commanded Lx does not depend on the surface; the amount α may. If the task itself moves with the surface, as a figure meant to sit on it ($Ln = 0$, $\mathcal{D} = \operatorname{span}(n)$), a controller must obtain whatever distinguishes surfaces that demand different commands, by an estimate, a direct measurement or a mechanical constraint. Using an estimate $\hat{n}$ then costs $\alpha L(\hat{n} - n)$ for the true L. Normal estimation is task-dependent: it can determine surface-relative commands or support contact and force regulation without entering the preserved coordinates.

Construction. (1) Write L from the task and take D spanning $\ker L$ (or the part the task frees). (2) Check $\operatorname{rank}(AD) = k$ over the surfaces expected, and the amplification and speed it costs. (3) Make every correction, force reading and rotation act through D, the contact force entering as $D^\top w$. (4) Close the feedback on α and s; its convergence, contact and drive limits must be shown for the feedback itself.

IV. One Controller for Unknown Surfaces and Objects

A. One Direction for Absorption and Force

For one contact ($k = 1$), $\mathcal{D} = \operatorname{span}(d)$. Hybrid controllers usually yield and regulate force along the surface normal n, true or estimated. Theorem 1 permits measurement-dependent correction amounts within $\ker L$, allowing implementation without a normal. With d_y the direction that absorbs depth and d_f the direction of the regulated force reading:

$$\begin{array}{ll} \text{NP:} & d_y = n, \quad d_f = n \quad \text{(true normal, given)}, \\ \text{IPCR:} & d_y = d, \quad d_f = d. \end{array} \tag{5}$$

Without friction, on a board at angle θ that reading is the normal force times $\cos\theta$; friction along d adds to it. The boards of the experiments reach $23.2°$. The admittance and drive loop determine the pressing and sliding components without a separate Cartesian wrench command.

B. Controller

On the Dobot CR3A, torque-level NMPC (20 Hz) tracks the target through the drives' joint-velocity loop. Depth laws act along d without estimating contact events or surface geometry. The target preserves Lr; the NMPC depth weight drops to a floor. Depth follows $v = k_f(F^\star - \hat{f})$, with $\hat{f}$ from joint-torque residuals ($k_f = 0.3\,\text{mm}\,\text{s}^{-1}\text{N}^{-1}$, retreat capped at 8 mm/s, $F^\star = 12$ N for the pen and 8 N for the box), without passing the reference depth. During tool travel, retreat rejects readings attributable to joint friction. Free-space approach uses the path's inward speed plus up to 5 mm/s. A continuous contact weight $w \in [0, 1]$ rises with increasing readings during inward motion, inward path lag, or readings exceeding friction alone; it falls with inward motion without a reading or an outward path. w scales the pressing terms; yielding is always on. The path clock waits when the tool is 100 mm behind and slows with the force opposing the path (14–22 N). Attitude is held except in U-yaw, whose yaw rate follows the estimated contact moment beyond a 0.3 N m dead band, widened at rest; yaw is limited by path travel. The force estimate reads 0.4–0.8 of the true force in simulation and has no sensor check on hardware.

TABLE I
Direction Handed to the Controller in Each Arm (Ideal Wrist Sensor Where Stated)

Arm	Yield direction	Target
IPCR	d	along d
NP-true	true n	along n
IPCR target, n absorption	n	along d
NP target, d absorption	d	along n
Soft task priority	$(1-\alpha)d + \alpha n$	along it
Iso., force-dir., motion-stiff	true n	ideal-sensor offset
HYB-C, HYB-S (fixed axis)	d	along d, PD on L

C. What the Task and Each Arm Specify

The task enters only through $(\mathcal{D}, L)$, an attitude specification and its inputs (force target, path, tool length); all other controller settings are identical across drawing and transport. Drawing yields in depth and keeps the frontal position. Transport yields sideways and keeps progress and height; U-fixed also keeps yaw, U-yaw lets it yield to the contact moment. Comparison arms share the control stack, with correction directions and position feedback specified in Table I.

V. Experiments

A. Setup and Protocol

The CR3A is simulated in Gazebo and ROS 2 with a patched DART engine (1 ms steps) and a plant identified on the physical arm. Every reference comes from nominal geometry and is infeasible: in leakage and drawing it commands 65 mm behind a nominal vertical plane, and in transport the box would pass through the board. Simulator forces and geometry serve only evaluation, except that the leakage test and one force diagnosis give IPCR the true normal and some comparison arms receive ideal sensing. Arms, metrics, decision rules and conditions were registered before each set of runs; invalid runs (simulator failures and supervisor faults not caused by the arm's motion) were repeated under a stated rule, and all failures are reported. "Unseen" denotes boards excluded from controller development; boards reused in later sets are

TABLE II
HARDWARE DRAWING (PER-RUN RANGES, ONE PASS, 12 N)

	Circle	Square	Star	Circle, 18°
Runs (IPCR / NP)	4 / 4	3 / 3	3 / 3	3 / 3
IPCR shift (mm)	0.23–0.27	0.28–0.30	0.33–0.34	0.21
NP shift (mm)	16.6–17.1	16.9–17.4	16.7–17.2	21.2–21.6
NP / predicted	1.00–1.01	1.00–1.01	1.00–1.01	1.04–1.06†

Columns 1–3: board at 11°, normal fitted after the runs. †Prediction fixed before the runs (registered).

identified as such. The physical CR3A runs the same controller setting as the simulation, with the NMPC on a host computer. A 6 N target lifted the pen on hardware but not in simulation (Sec. V-C). The simulated 1.1 N/mm spring represents arm and board compliance observed through the encoders; the board occupies the hardware drawing location, where the pen stiffness is 14.7 N/mm. On hardware, the board normal serves only for NP-true and the shift prediction: at 11° it is fitted afterwards to the pen-tip path of an IPCR run (fits 10.9° and 11.5°, measured 11°), at 18° measured with a tape before the runs; metrics come from joint encoders.

B. Drawing on Tilted Boards

Does contact preserve the figure's frontal placement? The simulated pen, aligned with the tool axis with its tip 134 mm from the flange, draws a circle of 0.20 m perimeter once. The metric is the frontal error $L(x_{\text{tip}} - r)$ over contact samples: its mean (shift) and RMS about the mean (shape error). Contact means a true force above 1 N, averaged over 50 ms; coverage is the share of the drawing phase in contact, reported because shift says nothing about the part not drawn.

C. Hardware: Drawing

At about 11°, IPCR and NP-true alternated single-pass circles, squares and stars at 5 mm/s (4+4, 3+3 and 3+3 runs). NP-true received the fitted normal. Runs that hit the board during set-up and one stopped by contact with a person were excluded. This set had no preregistered decision rule. A second circle set at 18° was registered after a simulation gate, with the taped plane committed before the first run (IPCR within 0.5 mm; NP-true 0.85–1.15 of the prediction; 3/3 each).

The 20 valid runs at 11° showed the placement difference across all three figures (Table II). Shape RMS was 0.50–0.60 mm (one circle run 0.90) against 0.79–1.28 mm. At 18° (6 valid runs) both criteria passed, NP-true within 0.7° of the predicted direction; its excess matches the tip's 3.5–4.9 mm offset from the taped plane. An independent video ink check measured an NP–IPCR displacement of 16.87–17.09 mm, 1.00–1.02 of the encoder value (Supplementary S2); it confirms the large offset, while IPCR's sub-millimeter placement rests on the encoders. The force estimate peaked at 24–34 N and was below 1 N for 0–6 % of the drawing phase at 11° and 0–13 % at 18°; at 6 N the share was 15–20 % and the pen lifted.

D. Hardware: Box Transport

A 250 × 200 × 150 mm box followed a straight path past parallel (within 0.6°) and 4–5° walls, with yaw held (U-fixed)

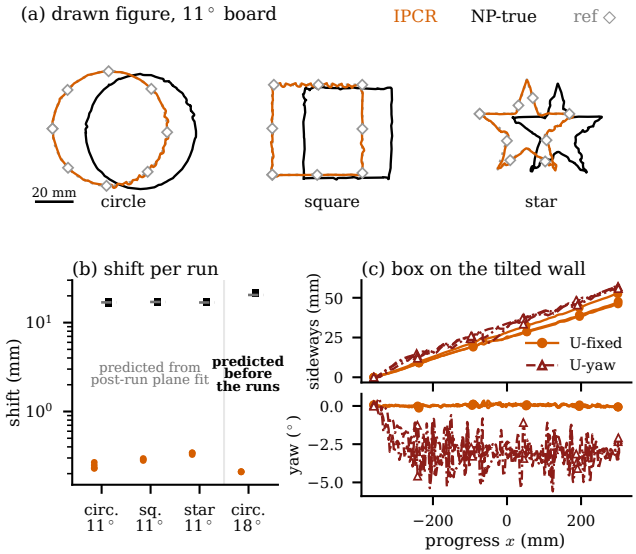

Fig. 3. Hardware results (U41M setting). (a) Drawn figures on the 11° board, IPCR against NP-true, with the reference (dotted, ◇). (b) Shift per run from the encoders; dashes: NP-true's prediction, from the plane fitted after the runs at 11° and from the plane fixed before the runs at 18°. (c) Box transport: front of the box and yaw for U-fixed and U-yaw on the tilted wall.

or yielding (U-yaw), three runs each after a free-space run of each task.

On the tilted wall, U-fixed preserved yaw within 0.4° while yielding 45–52 mm sideways over 0.64 m, with 99 % contact (force estimate above 1 N) and mean force 7.8–8.0 N against its 8 N target. U-yaw turned toward the wall (4.6–4.8°) to −3.0 to −3.1° on average with a standard deviation of 1.1° and excursions to −5.6°, with contact for 84–93 %. Its sideways yield, 56–57 mm over the same travel, was not smaller than U-fixed's on these walls. Progress error stayed at most 4.6 mm in every run and the attitude of U-fixed within 0.76°. First-touch force estimates peaked at 32–46 N, exceeding the preset 35 N threshold (simulation: 16–20 N); the source of the excess remains unresolved. On the parallel wall both tasks kept contact (91–100 %).

E. Against Execution Structures from the Literature

Does $G(B)\rho$ predict placement across execution structures with ideal sensing? At 5 mm/s, the same stack and pen tested each structure with the true normal and an ideal force sensor (Supplementary S1): NP-true, the single-plane minimum-norm case of safety filtering [4], [6], [7]; isotropic compliance used by learned position policies [29]; force-direction compliance [28], [30]; motion-stiff compliance [14]; and soft task-priority arms at $\alpha = \frac{1}{3}, \frac{2}{3}$. IPCR has neither the true normal nor a force sensor. Boards were tilted by 8.5°, 15.6° and 23.2° (frictionless) and 15.6° with $\mu = 0.6$, five runs per cell (140 runs). The registered criteria required every rival's shift to exceed IPCR's (Welch interval above zero), NP-true and the isotropic and force-direction arms to reach 0.85–1.15 of their per-run predictions, the soft arms to fall between IPCR and NP-true, and IPCR to stay within 0.5 mm.

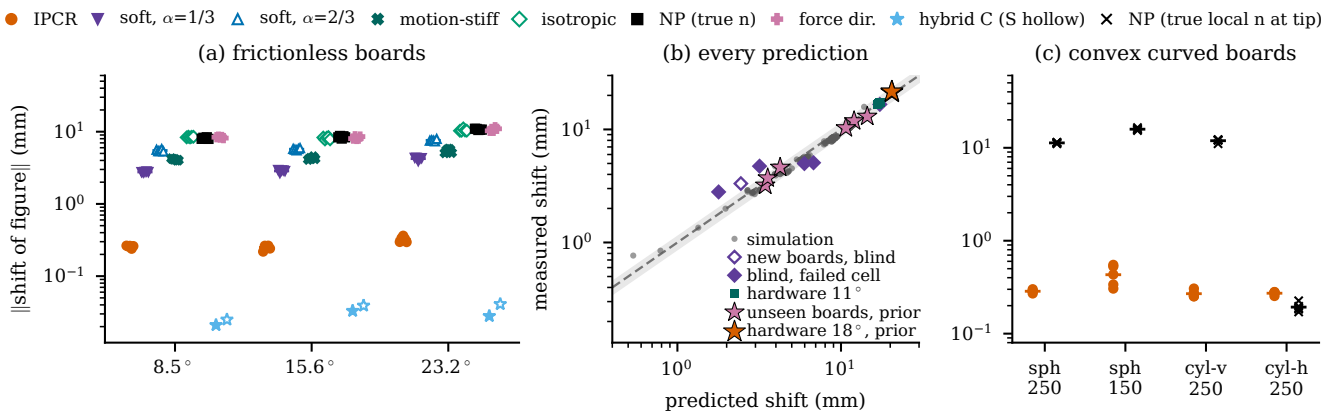

Fig. 4. Comparison at 5 mm/s; every rival has the true normal and an ideal force sensor, IPCR neither. (a) Shift of the figure per arm on the frictionless boards (5 runs per cell; bars: means; log scale); stars: fixed-axis hybrid controllers (cell means, HYB-S hollow), which meet the condition and were expected to stay near zero. (b) Every prediction above 0.4 mm against its measurement (dashed: identity; band: ±15 %); grey: simulation (rivals per run, target–absorption and offset cells, curved boards); diamonds: the blind first new-board set (filled: the four failed cells; its d/n band was 0.6–1.4); stars: predictions fixed before the runs. Zero predictions (IPCR, hybrid, d/d) are in (a) and Table III. (c) Convex curved boards (radius in mm; 5 runs per cell); NP-local takes the true normal at the tip and moves the figure except on the horizontal cylinder, where NP predicts no mean shift.

All runs were valid and all criteria passed (Fig. 4a). IPCR shifted the figure by 0.25–0.34 mm; rivals shifted it 10–36 times as far (board means: NP-true, isotropic and force-direction 28–36, motion-stiff 15–17, soft arms 10–23), at 0.90–1.09 of prediction (motion-stiff: 0.89–0.95, without a registered magnitude criterion). Soft arms followed the predicted order on every board, and every arm achieved at least 0.90 contact coverage. On the friction board force-direction compliance distorted the figure (shape RMS 22.8 mm against IPCR's 1.2 mm) and isotropic compliance 4.6 times IPCR's; NP-true's shift grew by 18 %. At 10 mm/s on four boards up to 23.2° (24 runs) IPCR kept full coverage and moved the figure by 0.25–0.36 mm against NP-true's 8.2–11.8 mm; at 5 mm/s on two further boards (15.0°, 18.8°; 12 runs) the values were 0.28 and 0.30 mm against 9.9 and 12.7 mm, again with full coverage for both arms.

Fixed-axis hybrid control. Does Theorem 1 predict preservation across architectures satisfying the same direction condition? Two fixed-axis hybrid controllers with $D = d$ ran on the same boards, seeds, force estimate and drive (52 runs, registered): task-space PD on the preserved axes with IPCR's own depth law (HYB-C), and a textbook guarded approach that latches force control along d once the reading exceeds F_{on} (8 N, chosen beforehand on development boards) and releases below 3 N (HYB-S). Both met the registered equivalence with IPCR (90 % Welch interval of the shift difference within ±0.5 mm on all six board–speed cells) and lay closer to the reference, at 0.02–0.09 and 0.03–0.10 mm. IPCR's remaining 0.25 mm is a phase-dependent tracking lead $L\varepsilon$, removed by their position feedback on L (Supplementary S3). HYB-S latched before true contact in 10 of 26 runs and recovered each time; its median first-contact peak was 13.3 N against 7.8 N for HYB-C and IPCR at 5 mm/s.

Reachability and speed. Theorem 2 requires depth speed up to $v \tan\theta$: 4.3 mm/s at 10 mm/s on the 23.2° board; NP instead changes frontal motion. IPCR supplies it through force-regulated depth speed without a normal. An earlier controller that took this rate from a normal estimate kept contact on half of the 23.2° board at 10 mm/s (Supplementary S7).

F. Target Versus Execution on the Full Controller

An earlier minimal-stack test is in Supplementary S5. We crossed the target rule (d or NP) with the absorption direction (d or n), all with the true normal, and injected a residual by offsetting the believed plane by $\delta = 10$ mm along n, a model of a surface error that contact does not correct. The prediction for the leak Λ (offset minus no offset) is zero when target and absorption both use d or both use n, and $\mp\delta L n$ when exactly one uses n. The design ran on the 15.6° and 23.2° boards with four repeats (64 runs), and a size series with δ = 2, 5 and 20 mm added 36 runs (registered).

All runs were valid, and all eight cells met the transfer-law criteria (Table III). The target alone did not preserve intent under the residual: the IPCR target with normal absorption leaked 2.90 and 4.02 mm, 1.08 and 1.02 of the prediction, against at most 0.09 mm for IPCR. This arm is the structure of one drawing system that keeps the in-plane position and complies along the normal [26]. The NP target with absorption along d leaked with the opposite sign (1.05 and 1.20 of the prediction), and both matched arms stayed at or below 0.28 mm, all within 1° of the predicted direction. Across δ = 2–20 mm the normal-absorption leak grew with slope 1.03 and 1.00 against the prediction while IPCR's stayed at or below 0.18 mm.

G. Predictions on New and Curved Boards

New planar boards. Does $G(B)\rho$ transfer without fitting? On three new boards (10.3°, 14.2°, 18.5°; 36 valid runs), blind magnitude predictions failed in four of twelve cells (NP-true: 0.75–0.84; normal absorption: 1.49–1.57 of prediction), although directions agreed within 4.8° and IPCR and d/d shifts stayed at or below 0.26 mm. A post hoc audit identified two

TABLE III
LEAK $|\Lambda|$ (MM) UNDER A 10 MM BELIEVED-PLANE OFFSET, MEANS OF FOUR RUNS PER CONDITION

Target / absorption	Predicted	15.6°	23.2°
d / d (IPCR)	0	0.09	0.05
d / n	2.69, 3.94	2.90	4.02
NP / d	2.69, 3.94	2.81	4.74
NP / n	0	0.08	0.28

residual-input constants accounting for the failed magnitudes (Supplementary S4). With both fixed from earlier boards, a second registered set on three unseen boards (13.4°, 13.9°, 16.6°; 45 runs, all valid) met every criterion: NP-true reached 0.90–0.99 and normal absorption 0.93–1.09 of the prediction (registered band 0.85–1.15), within 4.5° in direction, while IPCR stayed at or below 0.32 mm. A $\pm 3^\circ$ bias on the normal at 8.5° changed the NP figure by 2.8 mm, 0.94–0.95 of the prediction, and left IPCR's at 0.23–0.28 mm, since IPCR's target does not use the normal for the preserved coordinates.

Curved boards. Does command preservation extend beyond planes? Theorem 1 still applies on $r+\mathcal{D}$; NP's radial projection can also change figure size. On spheres of radius 250 and 150 mm and cylinders of radius 250 mm (40 runs, all valid; Fig. 4c), IPCR kept place and size: shift 0.27–0.43 mm, scale (mean centroid distance, drawn over reference) 0.999. NP-local, an oracle with the true normal at the tip, moved the figure by 11.3–15.9 mm on the spheres and the vertical cylinder (predicted 11.0–13.8 mm) with its scale within 0.02 of prediction; on the horizontal cylinder, where NP predicts no mean shift, it moved by 0.19 mm and changed the scale as predicted (1.124 against 1.120).

H. Same Controller, Different Task

The arm carries an unmodeled $250 \times 200 \times 150$ mm box 200 mm along a straight path, and an unknown board intrudes 20 mm into it. U-fixed keeps progress, height and attitude and yields sideways; U-yaw also places yaw in $\mathcal{D}$. The grid spans boards at 8.5°, 20° and 30°, friction 0 and 0.3 and payloads of 0.5–2.5 kg, three repeats each (108 runs).

Across 108 valid runs, five of six task-specification criteria passed (Fig. 5). U-fixed kept yaw within 1° in 54/54 runs and yielded a median 21.2 mm against the geometric 20 mm. U-yaw turned toward the board (median -3.9°) and yielded less than U-fixed in 16/18 conditions (median 16.1 mm), but turned past -3° in only 42/54 runs, short of the registered 80 %. Progress and height stayed within 4.3 mm in every run, roll and pitch within 1° in 107/108, and no yield exceeded 35 mm. Only the task specification differed within each pair.

Predicted progress drift. Does Corollary 1 predict transport displacement? Twenty-one further registered, valid runs (15 NP, six U-fixed), with nine U-fixed runs reused from the grid, compared U-fixed with true-normal NP on frictionless boards, testing the unfitted prediction $h_d \sin\theta\cos\theta$ for intrusion $h_d = 20$ mm. IPCR's maximum progress error stayed at 1.8–3.7 mm (Fig. 5d). NP exceeded it by 1.8, 4.8 and 6.0 mm at 8.5°, 20° and 30° (predicted 2.9, 6.4 and 8.7 mm), within the registered $\pm 35\,\%$ on two of three boards; the 8.5° magnitude failed.

VI. DISCUSSION AND LIMITATIONS

What the condition explains. The condition makes placement a consequence of the corrections a task permits. Fixed-axis hybrid control preserved position at least as well as IPCR (Sec. V-E), confirming the condition in a second architecture; NP instead preserves a surface task with $\mathcal{D} = \mathrm{span}(n)$. It also measures the metric recipe and its limit: in the ablation, all arms with the true plane and no injected offset, changing only the target from the nearest point ($W = I$) to $x_D^\star$ (the task-weighted limit) cut the shift from 8.5–10.8 to 1.3–1.9 mm under normal absorption, as in height-sensing drawing, and absorbing along d as well brought it to 0.26–0.32 mm; an offset added the predicted amount (Sec. V-F). Within the static affine model, a designer can thus predict task displacement before execution from A, L, B and an expected mismatch; measured shifts were 0.90–1.09 of prediction across rival structures and 1.04–1.06 on hardware with a prior plane (Fig. 4b).

Contact force. IPCR provides no force bound. In simulation the mean 10 ms peak per cell was 21.7–26.7 N for IPCR at 5 and 10 mm/s and 17.8–26.9 N for the rivals, against 30 N, a reporting threshold rather than a standard limit [31]. IPCR's peaks were no lower than NP-true's, so intent preservation did not suppress them.

Limits of the comparison. We tested execution structures under ideal sensing, not complete learned policies or safety filters. Failed magnitude predictions are reported in Secs. V-G and V-H. IPCR's residual tracking lead is removed by hybrid position feedback on L; its phase dependence remains unexplained.

Friction and slip. IPCR's tool meets the surface at up to 24°, and pens slip at large approach angles on real surfaces [27]; our friction boards reached only 15.6°.

Scope. The plant was identified on the physical arm; hardware tests cover two board tilts and two wall conditions, while box NP was evaluated in simulation. Local surface attitude is not tracked; preservation requires correct $(\mathcal{D}, L)$ and applies to the specified coordinates.

VII. CONCLUSION

Feasibility restoration does not by itself establish task preservation. Within a static affine model, a task's permitted corrections determine which coordinates are preserved and which constraints can be reached, and every other admissible direction has a closed-form cost $G(B)\rho$. The controller preserved hardware drawing placement using joint sensing and transported a box with held or yielding yaw; rivals with ideal sensing still displaced the figure. The displacement law serves as a diagnostic for correction structures, and a task-weighted metric recovers the preserving target in a static projection. Force bounds and contact convergence remain unproved. Contact realization must specify permitted task changes as well as surface geometry.

REFERENCES

[1] Q. Nguyen, Y. Kim, and J. Lee, "Learning reflexive behavior for contact-rich manipulation," 2026, *arXiv:2610.02811*.
[2] J. Hong, D. Song, S. Kim, and M. Seo, "Contact-aware imitation learning through contact factorization," 2026, *arXiv:2610.09533*.

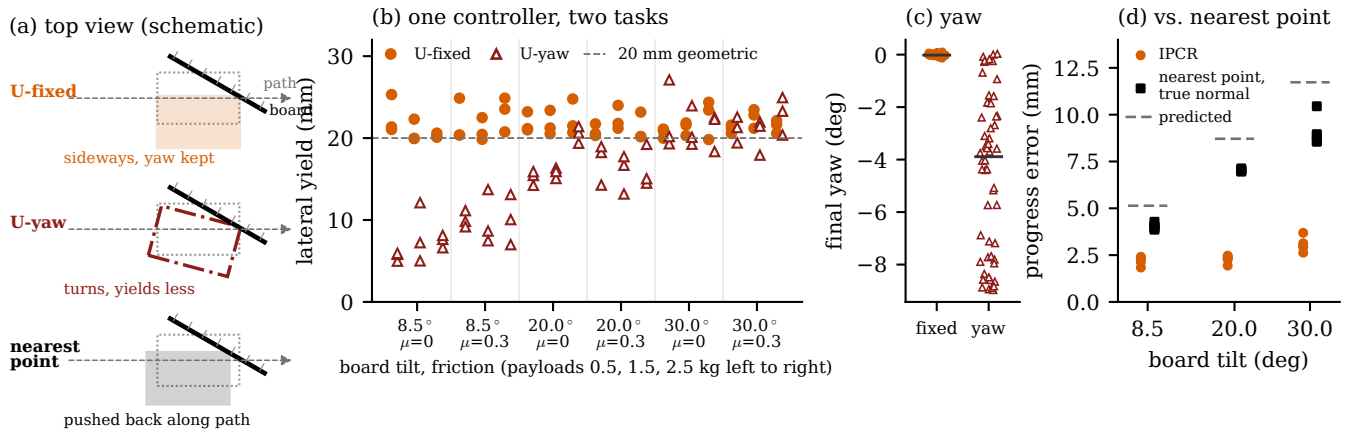

Fig. 5. Box transport with one controller. (a) Top view, schematic: U-fixed yields sideways, U-yaw turns, and NP pushes the box back along its path. (b) Lateral yield per condition (108 runs). (c) Yaw at the end of contact. (d) Maximum progress error on the frictionless boards (5 runs per arm and board), IPCR (U-fixed) versus NP with the true normal; dashed: IPCR's level plus $h_d \sin\theta \cos\theta$, predicted without fitting.

[3] Z. Doulgeri and Y. Karayiannidis, "Force/position regulation for a robot in compliant contact using adaptive surface slope identification," *IEEE Trans. Autom. Control*, vol. 53, no. 9, pp. 2116–2122, 2008.

[4] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. Eur. Control Conf.*, 2019, pp. 3420–3431.

[5] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-Jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Syst. Mag.*, vol. 43, no. 5, pp. 137–177, 2023.

[6] D. Morton and M. Pavone, "Safe, task-consistent manipulation with operational space control barrier functions," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2025.

[7] R. Römer, A. von Rohr, and A. P. Schoellig, "Diffusion predictive control with constraints," in *Proc. Learn. Dyn. Control Conf.*, PMLR vol. 283, 2025, pp. 791–803.

[8] Y. Karayiannidis and Z. Doulgeri, "Blind force/position control on unknown planar surfaces," *IET Control Theory Appl.*, vol. 3, no. 5, pp. 595–603, 2009.

[9] Y. Karayiannidis and Z. Doulgeri, "Robot contact tasks in the presence of control target distortions," *Robot. Auton. Syst.*, vol. 58, pp. 596–606, 2010.

[10] H. Lipkin and J. Duffy, "Hybrid twist and wrench control for a robotic manipulator," *ASME J. Mech., Trans., Autom. Des.*, vol. 110, no. 2, pp. 138–144, 1988.

[11] M. T. Mason, "Compliance and force control for computer controlled manipulators," *IEEE Trans. Syst., Man, Cybern.*, vol. 11, no. 6, pp. 418–432, 1981.

[12] J. M. Schimmels and M. A. Peshkin, "Force-assembly with friction," *IEEE Trans. Robot. Autom.*, vol. 10, no. 4, pp. 465–479, 1994.

[13] D. Song, J. Park, and Y. J. Kim, "SSK: Robotic pen-art system for large, nonplanar canvas," *IEEE Trans. Robot.*, vol. 39, no. 4, pp. 3106–3119, 2023.

[14] H. Shi, S. Hu, Y. Hou, W. Wang, K. Liu, and S. Song, "Minimalist compliance control," 2026, *arXiv:2603.00913*.

[15] M. H. Raibert and J. J. Craig, "Hybrid position/force control of manipulators," *ASME J. Dyn. Syst., Meas., Control*, vol. 103, no. 2, pp. 126–133, 1981.

[16] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE J. Robot. Autom.*, vol. 3, no. 1, pp. 43–53, 1987.

[17] H. Bruyninckx and J. De Schutter, "Specification of force-controlled actions in the 'task frame formalism'—a synthesis," *IEEE Trans. Robot. Autom.*, vol. 12, no. 4, pp. 581–589, 1996.

[18] J. Duffy, "The fallacy of modern hybrid control theory that is based on 'orthogonal complements' of twist and wrench spaces," *J. Robot. Syst.*, vol. 7, no. 2, pp. 139–144, 1990.

[19] W. D. Fisher and M. S. Mujtaba, "Hybrid position/force control: A correct formulation," *Int. J. Robot. Res.*, vol. 11, no. 4, pp. 299–311, 1992.

[20] R. Featherstone, S. S. Thiebaut, and O. Khatib, "A general contact model for dynamically-decoupled force/motion control," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1999, pp. 3281–3286.

[21] J. De Schutter, D. Torfs, H. Bruyninckx, and S. Dutré, "Invariant hybrid force/position control of a velocity controlled robot with compliant end effector using modal decoupling," *Int. J. Robot. Res.*, vol. 16, no. 3, pp. 340–356, 1997.

[22] Z. Doulgeri, N. Fahantidis, and A. Konstantinidis, "On the decoupling of position and force controllers in constrained robotic tasks," *J. Robot. Syst.*, vol. 15, no. 6, pp. 323–340, 1998.

[23] T. Yoshikawa and A. Sudou, "Dynamic hybrid position/force control of robot manipulators—On-line estimation of unknown constraint," *IEEE Trans. Robot. Autom.*, vol. 9, no. 2, pp. 220–226, 1993.

[24] E. Nasiri and L. Wang, "Hybrid force motion control with estimated surface normal for manufacturing applications," in *Proc. Int. Conf. Ubiquitous Robots*, 2024, pp. 125–132.

[25] J. De Schutter *et al.*, "Constraint-based task specification and estimation for sensor-based robot systems in the presence of geometric uncertainty," *Int. J. Robot. Res.*, vol. 26, no. 5, pp. 433–455, 2007.

[26] D. Song, T. Lee, and Y. J. Kim, "Artistic pen drawing on an arbitrary surface using an impedance-controlled robot," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 4085–4090.

[27] T. Weingartshofer, C. Hartl-Nesic, and A. Kugi, "Automatic and flexible robotic drawing on complex surfaces with an industrial robot," *IEEE Trans. Control Syst. Technol.*, vol. 32, no. 5, pp. 1602–1615, 2024.

[28] H. Fang *et al.*, "Force Policy: Learning hybrid force-position control policy under interaction frame for contact-rich manipulation," in *Proc. Robot.: Sci. Syst.*, 2026.

[29] Y. Yang *et al.*, "TDC: Sim-to-real transferable directional compliance for contact-rich manipulation," 2026, *arXiv:2602.14174*.

[30] Y. Hou, Z. Liu, C. Chi, E. Cousineau, N. Kuppuswamy, S. Feng, B. Burchfiel, and S. Song, "Adaptive compliance policy: Learning approximate compliance for diffusion guided control," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2025, pp. 4829–4836.

[31] J. A. Falco, J. A. Marvel, R. J. Norcross, and K. Van Wyk, "Benchmarking robot force control capabilities: Experimental results," Nat. Inst. Standards Technol., Gaithersburg, MD, USA, NISTIR 8097, 2016.

Supplementary Material: Feasible Is Not Faithful

Jeonghyeon Kim and Shinsuk Park

All runs, registration documents and analysis scripts are in the project repository; file names are given in brackets. Part I (S1–S6) concerns the controller and setting of the paper. Part II (S7–S15) records the development history on earlier controller versions; its numbers do not describe the paper's controller and are kept for the audit trail.

Part I: The Paper's Controller

S1. Comparison Arms

TABLE S1
Comparison Arms and the Methods They Represent (reference numbers as in the main text)

Arm	Target / absorption	Represents
IPCR	d / d	this work
NP-est	NP / $\hat{n}$ (own estimate)	adaptive and estimated-normal hybrid control (Karayiannidis–Doulgeri; Yoshikawa–Sudou; Nasiri–Wang); geometry-aligned compliance modes (TDC)
NP-true, NP-local	NP / true (local) n	the same without estimation error; min-norm safety filters (CBF, OSCBF, DPCC)
Force-direction	NP / sensed force (ideal F/T)	force-direction compliance (Force Policy; ACP)
Isotropic	NP / $\hat{n}$ + tangential (ideal F/T)	position policies with isotropic compliance (the default described by TDC)
Motion-stiff	NP / soft except along the commanded motion (true n, ideal F/T)	Minimalist Compliance
Soft task-priority	both along $(1-\alpha)d + \alpha\hat{n}$	weighted-metric realizations (Lipkin–Duffy)
IPCR target, $\hat{n}$ absorption	d / $\hat{n}$	pen drawing (Song et al.); the target–absorption ablation
HYB-C, HYB-S	d / d, PD on L	fixed-axis hybrid force/position control (Raibert–Craig); HYB-S with a guarded force latch

S2. Hardware Drawing: Metric and the Second Tilt

Metric. The tip path comes from the joint encoders and the identified kinematics; the ink point is where the tool axis meets the board plane, which removes the pen's compression along its axis. The shift is the mean of $L(x_{\mathrm{ink}} - r)$ over the drawing pass, and shape RMS is the spread about that mean [hw_eval.py, INK_PROJECT=1]. The board plane serves only this projection, NP-true and the prediction. *Ink check* [PREREG_INKVID; inkvid.json]. To test the encoder metric independently, the last frames of the 11° NP circle videos (1920×1080) were segmented by color; a least-squares ellipse was fitted to each ink circle, and the affine map that turns the red (IPCR) ellipse into a circle of the reference radius (31.8 mm) rectified the board. The black (NP) circle then had semi-axes of 30.7–32.3 mm, and its center lay 16.87, 16.92 and 17.09 mm from the red one, against 16.59, 16.90 and 16.85 mm from the encoders (ratios 1.017, 1.001, 1.014; registered band 0.90–1.10). The first NP run was drawn with a red pen and could not be separated by color, so the registered 4/4 was not met; the second video was analysed before the registration. The scale comes from the drawn red circle, so only the relative NP–IPCR placement is tested.

Second tilt [PREREG_HW_TILT2; tilt2_hw.json]. Registered after a simulation gate (6/6 runs passed on the robot-geometry board). The board stood about 18° from the frontal plane; its plane (distance 0.62 m from the base axis, yaw 18°, vertical) was measured with a tape and committed before the first run (sha256 1516112f…). Deviations, all decided before the runs: the plane was taped instead of fitted to touched points; the torque scale was 1.0 as in the 11° set (the handoff's 0.5 was an error of the development side); three set-up trials with the board lifted at its back were not logged. Criteria: IPCR at most 0.5 mm in 3/3 runs; NP-true within 0.85–1.15 of the prediction and 10° in direction in 3/3. Both passed (Table S2). The tip ran 3.5–4.9 mm off the taped plane along its normal; with the plane level taken from each run's own path, as for the 11° set, the ratios are 0.985–0.991 (post hoc).

TABLE S2
SECOND TILT, PER RUN (I N I N I N; MM)

Run	Shift	Prior prediction	Measured / prior	Shape RMS
IPCR 1, 2, 3	0.210, 0.209, 0.211	0	–	0.46, 0.48, 0.49
NP-true 1, 2, 3	21.52, 21.23, 21.59	20.51, 20.48, 20.42	1.049, 1.037, 1.057	2.91, 2.64, 2.89

Direction error of NP-true 0.5–0.7°. Peak force estimate 24.5–29.7 N; no stop criterion was met.

S3. IPCR'S REMAINING 0.25 MM

[DIAG_IPCR_RESID; registered before the analysis, existing runs only.] On the same boards and seeds IPCR moved the figure by 0.264 mm (median; 0.219–0.421) and the fixed-axis hybrid HYB-C by 0.028 mm (0.007–0.062). Each tick's frontal error was split into a reference part and an execution part, and along and across the path tangent. For IPCR the reference part was zero and 0.86–1.06 of the shift lay along the tangent (hybrid: 0.37); its correlation with the true force was -0.14. The first registered hypothesis, a tracking error along the path, therefore holds, with one difference from its wording: the error is a lead, not a lag, and depends on phase (on one run $+0.30$ to $+0.56$ mm over the right half of the circle and -0.05 to -0.12 mm over the left, at an unchanged speed of 4.3–4.7 mm/s), so it does not cancel over a lap. This is the term $L\varepsilon$ of the execution equation, not a leak through $\mathrm{range}\, B$; the hybrid's position feedback on L removes it. The cause of the phase dependence was not examined.

S4. FIRST NEW-BOARD SET AND ITS AUDIT

Blind set. Arms I, NP-true, d/d and d/n (absorption along the normal, 10 mm offset) ran on three new boards ub1–ub3 (36 runs, all valid), with predictions frozen before the runs. I and d/d stayed at or below 0.26 mm and every leak direction was within $4.8°$, but four of twelve magnitude cells failed the registered band (Table S3). The verdict stands. *Audit (post hoc).* The failures were traced to the residual ρ supplied to the prediction, not to $G(B)$. (i) Contact level: in the 51 IPCR runs on frictionless boards the tip sat below the modeled plane by $-0.30 - 0.147F$ mm (F the mean true force, 13.6–20.9 N; correlation -0.82, with tilt 0.14), i.e. an effective pen compliance 16 % above nominal or a force reading 14 % low; the run data cannot tell these apart. Adding the same board's IPCR contact level (-2.6 to -3.0 mm) to the NP prediction gives the corrected column. (ii) Baseline absorption: the blind d/n prediction omitted the 4.9 mm by which the d-directed target lies inside the surface without an offset (from the $\delta = 0$ runs of the target–absorption set, 4.95 and 4.80 mm on the two boards). Both inputs come from other runs of the same simulator, so the corrected ratios are a diagnosis, not a test. *Prospective test.* With the two constants frozen (-2.86 mm, 4.87 mm), a second registered set on three unseen boards ub4–ub6 (45 runs, all valid) passed every criterion (Table S3); the blind predictor would have given d/n ratios of 1.38–1.62 on these boards.

TABLE S3
MEASURED / PREDICTED SHIFT ON THE NEW BOARDS (BAND 0.85–1.15; BLIND d/n: 0.6–1.4)

Board	NP-true	NP-true, corrected	d/n	d/n, corrected
ub1 (blind)	**0.835**	0.913	**1.567**	1.054
ub2 (blind)	0.964	1.001	1.359	0.911
ub3 (blind)	**0.747**	0.870	**1.489**	1.001
	NP-true, prior		d/n, prior	d/n, blind rule
ub4	0.988		0.930	1.383
ub5	0.963		1.043	1.552
ub6	0.901		1.087	1.617

Bold: failed. The d/n band on ub1–ub3 was 0.6–1.4, so ub2 (1.359) passed. Corrected: post hoc (see text). Prior: constants frozen before the runs. Sources: u41m_confirmatory.json (F), f_audit.json, fpro_confirmatory.json.

S5. THE EXECUTION LEAK IN ISOLATION

An earlier minimal stack, given the true normal and without geometric holding, rotated the absorption direction $\hat{N}(\alpha) = \mathrm{normalize}((1-\alpha)d + \alpha\hat{n})$, $\alpha \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$, under a ± 15 mm anchor offset on three walls tilted by $12.8°$–$16.6°$ (72 runs, registered).

Every criterion passed (Fig. 1): the change of the stroke-mean preserved error followed $G(\hat{N})\rho$ with slope 0.93 and $r = 0.93$, and fell from 3.0–4.3 mm at $\alpha = 1$ to 0.12–0.69 mm at $\alpha = 0$.

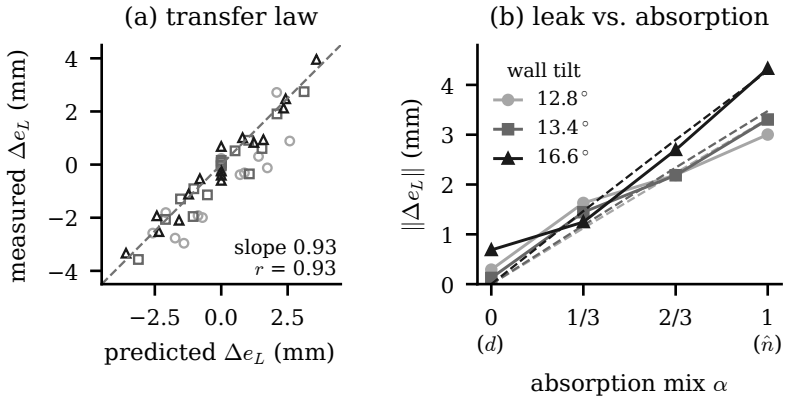

Fig. 1. Leakage law on the robot controller (72 runs, 3 walls). (a) Measured versus predicted change of the preserved coordinates for a ± 15 mm residual; dashed: identity. (b) Leak magnitude versus the absorption mix α (0: absorption along the yield direction d; 1: along the estimated normal); dashed: $\|G(\hat{N})\rho\|$.

S6. A Reading for MPC

Reading for MPC. A predictive controller meeting an infeasible reference usually minimizes $\|x-r\|_W^2$ subject to the contact constraint. With $W = I$ this returns the nearest feasible point, so the drift of Corollary 1 is the default of an unweighted tracking cost; with $W = wL^\top L + DD^\top$ the solution tends to $x_D^\star$ as $w \to \infty$ (on three tilts of 5.8–17.7°, $|x - x_D^\star| < 10^{-4}$ mm at $w = 10^6$). Where a residual is absorbed is decided by the servo, not by the cost (Proposition 1). Our NMPC only tracks the realized target.

Part II: Earlier Controller Versions (Development Record)

These sections concern controller versions that preceded the paper's setting (a separate normal-estimated force axis, contact commitment, auxiliary attitude terms, and earlier evaluation windows). They are kept because the main text cites S7 and because the failures they record shaped the final design.

S7. Depth-Rate Analyses

Speed tests (registered; first placement of the 23.2° board, which includes the approach contact of S12). At half speed coverage rose to 0.95 and 0.96 at 15.6° and 18.8° but stayed at 0.47 at 23.2°, where the estimate stayed at its prior in two of three runs, so the test failed by its rule. A prediction from the measured coverage curve for quarter speed at 23.2° (coverage of at least 0.75) passed at 0.78 (0.50, 0.85, 1.00, first stroke only).

Post hoc on stored runs [speed_limit.py, speed_limit2.py]. Criteria for the 10 mm/s contact-quality comparison (coverage, force s.d. and tangential goal error of NP-true no worse than IPCR's by more than 20 %, while its shift is at least five times larger) were written before computing and passed [silent_failure.json]: coverage 0.93, 0.79, 0.49 against 1.00, 0.98, 0.93; force s.d. 6.2, 6.0, 8.9 against 3.8, 4.1, 5.8 N; tangential goal error 2.3, 3.1, 3.8 against 2.3, 3.1, 3.3 mm. Coverage against the true depth rate s: IPCR 0.90 ($|s| < 1$ mm/s), 0.81, 0.66, 0.32, 0.15 ($s > 4$ mm/s); NP-true 0.95–0.99. Leave-one-board-out prediction of coverage from the rate curve: 5 of 6 boards within 0.15, missing 23.2° (0.86 predicted, 0.49 measured). The first and second one-dimensional models missed in turn: the second failed on the held-out half-speed 23.2° cell (1.0 predicted, 0.47 measured); the third, written after seeing that cell and using the mismatch rate, reached 6 of 9 cells. None is used. Registered half-speed set [PREREG_SPEED; one invalid run repeated]: overall fail. Registered matched comparison [PREREG_MATCH; match_confirmatory.json; IPCR_SIM_T_MAX = 130 s]: pass. Three runs faulted with a stale controller command at 94–106 s and one failed to start; each was repeated once, and the NP-true 15.6° repeat faulted again. These faults were not caused by the arm's motion, but they occurred only in the longer runs and are worth watching. The half- and quarter-speed PREREG_SPEED runs ended at 84 s and omit part of the scheduled drawing; the PREREG_MATCH runs used 130 s and completed both strokes.

S8. Attribution of Contact-Force Transients

Threshold. Peaks are reported against 30 N, about twice the mean drawing force (9–15 N). This is a reporting threshold, not a standardized limit for this contact. Following the force-control benchmark of Falco et al. (NISTIR 8097), transient peaks are reported apart from the mean force. The 10 ms window is used because 1 ms forces of the stiff simulated contact depend on solver settings.

Incidence. On the unseen boards the 10 ms peak exceeded 30 N in 27/29 IPCR runs (median 42, maximum 73 N) and 18/25 NP-true runs (median 34, maximum 76 N); the isotropic and force-direction medians were 40 and 43 N. IPCR's peaks were no lower than NP-true's.

Mechanism seen in traces. In the main drawing set, while the true force rose through 10–30 N the joint-torque estimate read 3–13 N; the controller released its contact commitment, and the approach continued until it re-engaged.

Registered attribution [PREREG_FGUARD; 29 valid runs; realization rule unchanged; new flags default off]. Three sources were separated. Each intervention removed most of the time above 30 N in its regime, but none kept peaks reliably below it.

- First touch: the uncommitted approach kept pushing during the contact-belief latency. A guarded approach that withdraws its lead as the estimated force rises cut the first-touch time above 30 N from 130–970 to 0–70 ms per run; 3 of 14 runs still exceeded, against 10 of 15 (criterion F1 failed).
- Sliding: after belief dropouts the joint-torque estimate was uncorrelated with the true force. Detection and force from the simulator kept 6 of 10 runs on the $8.5°$ and $15.6°$ boards at or below 30 N, against 2 of 10 (median time above 30 N: 0 versus 310 ms); remaining peaks were 30–35 N (criterion F2 failed).
- Normal acceptance: at $23.2°$ the committed force axis stayed on the task prior. Giving IPCR the true normal cut the median time above 30 N from 3.4 to 0.73 s, with peaks of 39–49 N (criterion F3 passed).

Drawing shift stayed at or below 0.8 mm in every arm (F4 passed). Peaks of 30–35 N remain unexplained.

S9. Residual Shift of an Earlier Version

IPCR's shift on the unseen boards was 0.25–1.64 mm (median below 1 mm). An offline decomposition of the 29 valid runs [residual_decomp.py] gave:

- a zero contribution from the tool attitude, because the flange target is recomputed from the current pen axis (the pen tilted by up to $9°$, which alone would move the tip by 21 mm);
- the remainder in the flange's tracking of its goal.

A contact-force model of the residual fitted in a pilot did not replicate (criterion P3).

In those runs, the shift grew as contact covered less of the circle (Spearman -0.55). A registered replication [PREREG_RESID] on 29 independent runs reversed this association ($+0.54$, criterion R1 failed; R3 failed). Restricting fully covered traces to the contact pattern of interrupted runs left their shift at 0.51 versus 0.49 mm (R4 failed). These tests did not support incomplete contact as the cause. Over the 58 runs the shift stayed at or below 1.75 mm; its cause remains open.

S10. Pilots and Protocol Deviations

- *First drawing set:* before the verdict, an evaluation error that compared the tip with the flange target was found and corrected.
- *Normal-bias test (unseen boards):* invalid, because its injection hook was inactive; this was found only at the verdict. It was repeated with the hook verified (curved-board set) and passed.
- *Horizontal cylinder:* the face point lies on the figure's center line, so the nearest point predicts no mean shift there (0.0 mm predicted, 0.4–3.2 mm measured). This design flaw was noticed at the verdict and made the confirmatory claim fail on that board.
- *NP-local oracle:* it evaluated the local normal at the commitment anchor rather than at the tip, which put it $4°$ (cylinders) to 8–$14°$ (spheres) off. Projecting along the realized direction reproduced shift and scale in 13 of 15 runs (e.g. scale 0.755 measured, 0.730 predicted). A corrected oracle, taking the normal at the current tip, was rerun [PREREG_ABL, C1–C2; 20 valid runs]. Both criteria passed: scale 1.188, 1.140, 1.058, 1.169 against 1.195, 1.161, 1.069, 1.186, and shifts 15.5, 21.9, 16.5 mm, 28, 47 and $24\times$ IPCR's mean from the earlier set; the horizontal cylinder shifted 0.65 mm.
- *Ablation launch stamp:* the frozen analysis excluded runs stamped before 18:00, but the confirmatory runs began at 17:58. It was run with the stamp set to 17:58, before results were seen; the smoke runs (17:54–17:55) stay excluded and no run was dropped.
- *Friction equivalence (P4a):* friction lowered the NP-true leak by 10–15 %; the 90 % CI of the change (-2.9 to -0.7 mm) left the ± 1.9 mm margin.
- *Leak injection on the frozen controller:* in a six-run pilot, its geometric depth hold diverted most of the injected anchor offset, leaving a 1–2 mm effect against 0.9 mm at $\alpha = 0$. The registered target–absorption ablation uses a believed-plane offset that reaches both target rules.

S11. Offset Analysis and All Predictions

20 mm set. A second set at 20 mm (three further boards, 18 runs) failed one of three cells; its predictor held an empirical 1.4 mm term found post hoc (below), so it is exploratory, with shifts reported uncorrected: 2.39, 4.79 and 9.79 mm for normal absorption against at most 0.49 mm for IPCR, within $6°$ of the predicted direction.

Status. Everything in this section is exploratory. The 1.4 mm term below was found after looking at the PREREG_ABL and PREREG_BLIND runs and was then written into the PREREG_BLIND2 predictor; that set failed and is reported below with uncorrected shifts. Measured-residual decomposition [blind_decompose.json]: measured over $|\rho_{\log} Ln|$ for the six blind cells

was 0.27, 0.85, 0.55 (10 mm; 11.65°, 18.8°, 12.99°) and 0.77, 0.76, 0.98 (20 mm; 9.99°, 14.96°, 21.14°); per-run values at 21.14° 0.87, 0.94, 1.09.

A post hoc analysis of the stored runs found that the realized correction falls short of the normal by a tangential offset of about 1.4 mm on most boards, independent of tilt and growing slowly with the residual, so that the leak is close to $\max(0, \rho\|Ln\| - e_0)$. Over 13 conditions, the RMS prediction error was 0.70 mm using the injected offset, 1.05 mm using the logged residual alone and 0.33 mm using the logged residual with e_0 taken from the other boards. We could not identify the cause: pen tilt, a force-proportional servo share, joint stiction and the general tracking error were examined and did not explain it. A second registered set, at $\delta = 20$ mm on three further new boards (9.99°, 14.96°, 21.14°; 18 runs, all valid), fixed the form with δ for ρ and $e_0 = 1.4$ mm before the runs. IPCR stayed at or below 0.49 mm. Normal absorption leaked 2.39, 4.79 and 9.79 mm against predicted 2.07, 3.76 and 5.81 mm, within 6° in direction and about 10 to $29\times$ IPCR. The last cell exceeded its limit of 1.3 (ratio 1.68), so the set failed by its registered rule. There, the three runs realized residuals of 17, 38 and 29 mm and leaked 5.3, 12.7 and 11.4 mm, while IPCR's residual stayed near 18 mm: normal absorption's size followed ρ, and ρ itself was not predicted.

Fig. 2 collects every mean-level prediction with its measurement, split by what the prediction used. Those using no logged absorption direction or residual had a median measured-to-predicted ratio of 0.95, with 28 of 35 within 40 % and 31 of 34 zero predictions at or below 1 mm. Those using a logged direction or residual were no closer on nonzero predictions (0.89, 32 of 42 within 40 %), and all 21 of their zero predictions stayed at or below 1 mm. This summary is post hoc and its conditions share boards and hardware, so it is descriptive. The misses in this group are the four ablation cells, three cells of the registered prediction sets, the 8.5° transport progress and two role-reversal cells at 1.3–1.4 mm. Some NP predictions still took their plane level and contact window from the logged run, so this group is not wholly independent of execution data.

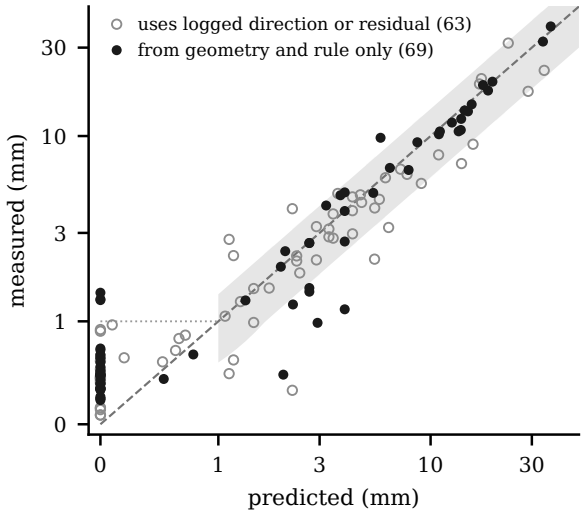

Fig. 2. Mean-level predictions against measurements across all experiments (one marker per condition; symmetric-log axes, linear below 1 mm). Dashed: identity. Shaded: within 40 %.

S12. Verdict on the Unseen Boards

Earlier set and why it was replaced. The table below is the verdict of the first comparison (PREREG_EVIDENCE, 2026-11-04, 170 runs), in which IPCR moved the figure by 0.47, 0.82 and 0.89 mm and NP-true by 10.2, 12.4 and 32.3 mm on the frictionless boards. An audit then found that the isotropic, force-direction and soft arms had received the controller's estimated normal (12–23° off on the steeper boards), not the true one, and that on the 23.2° board the common start pose placed the pen tip 5.2 mm past the face, so 50 of its 150 runs touched the board during the approach joint move. The next development sets reran every rival with the true normal on a base setting with a 0.5 s pen-axis filter and no commitment during the approach (PREREG_REBASE2, 120 runs; its 23.2° cells share the collision and are not used) and on the 23.2° board moved 40 mm back (PREREG_CF24D40, 35 runs; four rival runs touched it briefly in an approach overshoot, 25–49 N, without changing any verdict). Rival shifts on the moved board are compared only with their own predictions, since its commanded depth is 40 mm smaller. *Role reversal (this set).* For a task defined on the surface the roles must swap. In board-tangent coordinates, NP-true (IPCR with $\mathcal{D} = \mathrm{span}(n)$) kept the figure within 0.72–1.43 mm, and IPCR with $\mathcal{D} = \mathrm{span}(d)$ moved it by 10.6–39.1 mm: the appropriate rule is the one that uses the task's yield subspace.

TABLE S4
PRE-REGISTERED VERDICT ON THE UNSEEN BOARDS (167 VALID OF 170)

Claim	Verdict
P1 (confirmatory): IPCR shift $\leq$ NP-true/5 on all 3 boards	pass (15–36×)
P2: per-tick slope $\in [0.8, 1.2]$, $r \geq 0.9$ in $\geq 80\,\%$ of runs	fail (26 %)
P3: pilot model of IPCR residual replicates	fail
P4a: NP-true friction-immune ($\pm 15\,\%$)	fail (-10 to $-15\,\%$)
P4b: isotropic friction leak follows physics model	fail
P4c: force-direction distorts on friction ($\geq 3\times$ IPCR)	pass (4–10×)
P5: IPCR $<$ soft $\alpha{=}\frac{1}{3} < \frac{2}{3} <$ NP on $\geq 2/3$ boards	pass (2/3)
P6: normal-bias response	invalid (hook off; passed in a later set)

S13. ABLATION DETAILS

Approach contact on the 23.2° *board* [AUDIT_ABL_CF24_COLLISION, post hoc]. On this board the common start pose places the pen tip 5.2 mm past the face, so 16 of the 64 runs touched it during the approach joint move with more than 5 N (at most 44.2 N); all 61 runs on the 15.6° board stayed at 0 N. Recomputed from runs at most 5 N, the normal-absorption leak with the hold on was 4.06 mm (1.03 of prediction; registered 3.94) and IPCR's 0.56 mm (0.60); with at most 1 N only one or two runs per cell remain (3.04 and 0.53 mm). The conclusion is unchanged, but the 23.2° values rest partly on these runs; the claim rests on the 15.6° board and the hardware set.

Registered verdict [PREREG_ABL; abl_confirmatory.json]: A3 (IPCR leak) passed; A1 (hold off) and A2 (hold on) failed, in four cells of this earlier experiment, distinct from the main paper's ablation. Table S5 lists the shift and residual behind each leak. With no injection, the target effect (NP minus d, absorption along d, hold on) was 11.8 and 29.8 mm against 13.7 and 33.4 mm predicted for the same boards.

Follow-ups [PREREG_VAL; val_confirmatory.json; 68 valid runs, one timeout repeated]. Both registered sets passed. Own normal, offset 10 mm, four repeats: IPCR 0.09 and 0.16 mm, normal absorption 0.72 and 1.07 mm at 8.5° and 15.6°; the 8.5° prediction (0.64 mm) fell below 1 mm, so that cell tested only $|\Lambda| \leq 1$ mm, and these predictions use the logged absorption direction. Dose, true normal, three repeats with a baseline shared across offsets: slope 0.93 and 0.87 over $\delta = 2, 5, 10, 20$ mm. These sets did not retest the NP-target cells, whose misses remain unexplained.

Predictions fixed before the runs [PREREG_BLIND; PREREG_BLIND2]. Both registered sets failed by their all-cell rule: 10 of 12 and 5 of 6 cells passed. Offline analyses of the offset e_0 are in share_law.py, deadband_check.py and band_check.py; six tracking-band measures failed their written criterion.

Post hoc. The residual change realized by the offset varied (5.0–12.3 mm in the mixed cells) rather than equaling the injected 10 mm, which accounts for part of the spread among those cells. The standard error of Λ from run-to-run spread was 0.3–0.7 mm for d-target cells and up to 2.0 mm for NP-target cells at 23.2°.

TABLE S5
MEAN SHIFT MAGNITUDE AND REALIZED RESIDUAL ρ PER CELL (MM; FOUR RUNS PER VALUE), OFFSET $0 \to 10$ MM

Hold, target / absorption	Shift 15.6°	ρ 15.6°	Shift 23.2°	ρ 23.2°
on, d / d	0.44 → 0.10	2.2 → 11.2	0.69 → 0.09	5.1 → 12.1
on, d / n	0.28 → 2.37	1.7 → 13.3	0.38 → 4.18	1.8 → 14.1
on, NP / d	12.27 → 14.91	0.9 → 11.1	29.09 → 30.25	-3.6 → 1.3
on, NP / n	12.56 → 12.22	1.3 → 10.0	30.73 → 32.00	-3.9 → 3.9
off, d / d	0.34 → 0.25	1.8 → 16.0	0.49 → 0.10	5.1 → 12.7
off, d / n	1.10 → 2.55	4.8 → 12.4	0.16 → 2.86	1.0 → 8.3
off, NP / d	12.19 → 13.69	1.1 → 8.0	27.26 → 32.21	-6.4 → 4.2
off, NP / n	12.38 → 12.80	1.1 → 11.6	31.10 → 30.64	-4.0 → 3.3

S14. CONTROLLER OVERVIEW AND FURTHER RESULTS

Auxiliary terms and configurations. Three terms outside the realization rule act on attitude and effort only. The *moment reflex* turns the stored contact attitude toward a believed tilting moment while the believed lever $|M_t|/|F|$ is one a contact on the held object can produce (limit 0.05 m; 0 switches it off). The *face-alignment turn* re-orients the tool face to the estimated normal in a press–release–turn–press sequence. The *slide feed-forward* adds the predicted Coulomb friction $\hat{\mu} f_n$ along the reference's tangential direction on the committed plane. None enters the target $x_D^\star$, $(\mathcal{D}, L)$ or the absorption direction. On the hardware the full-gain slide feed-forward made the pen oscillate at 1.2–1.3 Hz on a flat board; switching it off alone removed the oscillation.

- *Comparison set at 10 mm/s.* IPCR's per-run shift was 0.25–1.64 mm; NP-true was 22, 15 and 36 times larger by board. Shape errors about the mean were 2.3–6.0 against 2.2–5.7 mm. Soft task-priority arms lay between IPCR and NP on two of three boards. Isotropic compliance added 1 mm/N tangential; force-direction absorbed depth along the sensed force.

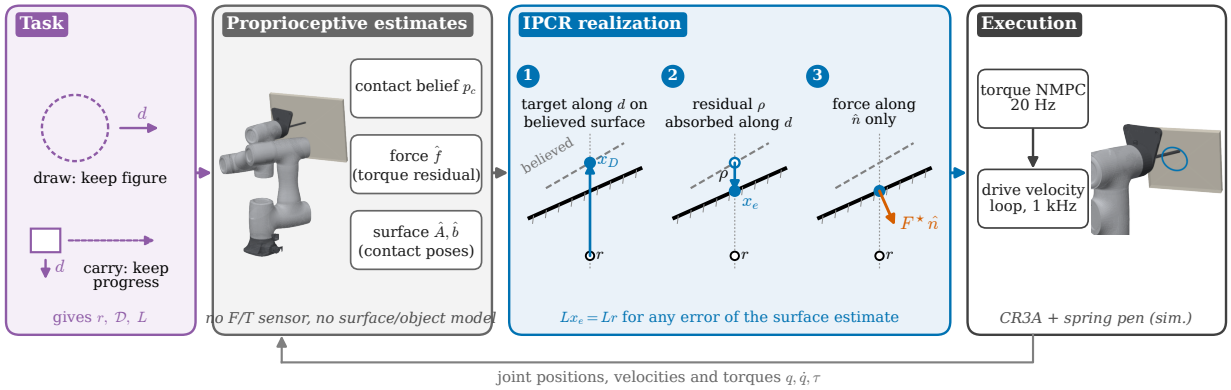

Fig. 3. Controller overview. The task supplies r and $(\mathcal{D}, L)$; joint signals supply the contact belief, force estimate and believed surface. IPCR (1) places the target on the believed surface along d, (2) absorbs the residual along d and (3) regulates force along $\hat{n}$ only. A torque NMPC and the drives' velocity loop execute it, with one setting for drawing and transport.

TABLE S6
AUXILIARY-TERM CONFIGURATION PER EXPERIMENT SET (ALL ARMS OF A SET SHARE IT)

Set	Moment reflex	Face alignment	Slide FF
Comparison at 10 mm/s, ablation, predictions, curved, transport	on (0.05 m)	on	on
Same-speed comparison (earlier version)	off	off	off
Hardware	off	off	off

- *Ablation misses.* Hold off, $15.6°$: mixed cells 0.54 and 0.56 of the prediction (limit 0.6), in the right direction. Hold on, $23.2°$: NP/d ratio 0.29; NP/n 1.32 mm against a 1 mm limit; post hoc standard errors of 1–2 mm. With the hold on the injected residual stayed at 10–14 mm, except 1.4–3.9 mm for the NP target at $23.2°$.
- *Offset follow-ups.* Normal absorption at 2, 5, 20 mm: 0.44, 1.30, 4.93 mm at $15.6°$; 0.68, 1.97, 6.56 mm at $23.2°$. Own normal: $\|L\hat{N}\| = 0.064$ and 0.107 against 0.148 and 0.269; 0.99 of the prediction at $15.6°$.
- *Matched comparison, per run* [match3_confirmatory.json, hardware setting]. IPCR shift per run 0.42–0.51, 0.36–0.76, 0.23–0.94, 0.50–0.58 mm ($8.5°$, $15.0°$, $15.6°$, $18.8°$); NP-true 10.45–10.67, 14.34–14.54, 13.59–13.78, 18.09–18.17 mm. The same cells with the auxiliary terms on [match_confirmatory.json] gave 0.55/10.4, 0.84/13.7, 0.65/12.9, 0.21/17.6 mm (cell means, IPCR/NP-true).
- *Curved boards.* Original oracle: 19, 32 and 20 mm on the spheres and vertical cylinder, 30–68× IPCR; size 0.81–1.02 against predicted 1.07–1.20, normal 4–14° off. Corrected oracle: 24–47× IPCR's mean in the earlier set.
- *Transport.* U-fixed median yaw $-0.3°$, yield 19.0–21.0 mm. Roll and pitch within $0.83°$; progress and height within 5 mm in 107 of 108 runs (one: 6.2 mm). NP's sideways yield 19.7, 18.8, 13.6 mm against predicted 19.6, 17.7, 15.0 mm. High friction resists U-yaw's turning (yield 7–16 mm at $30°$, still below U-fixed).
- *First drawing set, boards seen in development.* IPCR's median shift 0.31 mm on tilted boards (rigid pen: 0.20–0.36 mm); NP 4.0 and 7.8 mm at $8.5°$ and $16.6°$, 0.92–1.08 of prediction.
- *IPCR target with normal absorption, offline.* On logged goals and normals of curved-board runs, where $\hat{n}$ missed the local normal by a median of 5°–31°, this structure predicts a figure within 0.2 mm of IPCR's (one run per board).
- *Soft tools.* If the belief follows the kinematic tip, a tool's compression $s = F(n^\top z)/k$ along its axis z moves the believed surface with it, so a soft tool alone does not create a persistent residual.

S15. EARLIER HARDWARE METRIC WITHOUT SIMULATOR GROUND TRUTH

On hardware, contact samples are those with the pen tip at or behind the plane fitted to touched points. On the 30 simulated runs on the three frictionless unseen boards, this encoder-only metric reproduced the force-based one to 0.08 mm on average for IPCR (maximum 0.29 mm) and 0.02 mm for NP-true (maximum 0.16 mm). With the plane shifted by ± 0.5 mm the agreement stayed within 0.11 mm on average, and a $1°$ error of the measured normal changed IPCR's shift by at most 0.23 mm (NP prediction: 2 %). Transport metrics (yield, progress error, end yaw) are computed from joint encoders alone and matched the simulation evaluator exactly. The controller's contact belief is not used to select samples, because it covered only 5–30 % of the true contact samples.