Is a Little Slip Worth It?
Certified Fast Lifting for Learned Grasp Policies Under a Static-to-Kinetic Friction Drop

Clement Aidoo

CANLABS, Accra, Ghana clement.aidoo@canlabs.org

Code, data, protocol: github.com/CAN-Aidoo/lift-shield Archive: doi:10.5281/zenodo.23769371

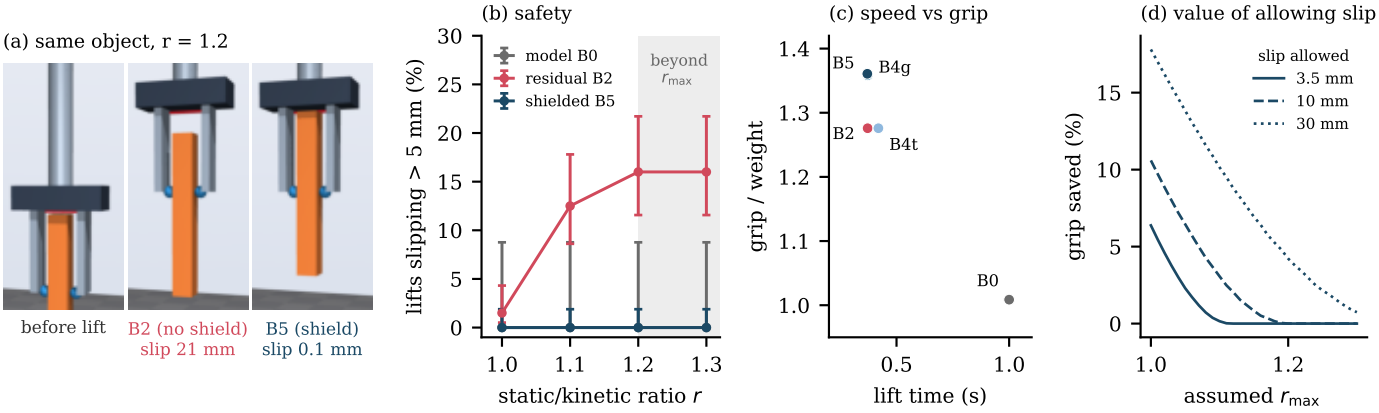


Figure 1. A learned lifting policy cuts lift time from 1.00 s to 0.37 s, and a physics shield is what keeps it safe. (a) The same object lifted by the unshielded residual policy (B2), which slips 21 mm in the fingers (the rod falls below the red mark on the palm), and by the shield-trained policy (B5), which does not. (b) Unshielded, the policy fails 12–16% of lifts once kinetic friction is lower than static ($r > 1$); shielded, it never fails, even beyond the shield's assumed $r_{max} = 1.2$ (40 test objects, 5 seeds, Wilson 95% intervals). (c) Every learned variant is 2.7× faster than the friction-model controller but grips harder. (d) Letting the shield allow some slip saves grip only if friction is assumed not to drop.

Abstract

A robot gripper lifting an unfamiliar object must decide how hard to squeeze and how fast to move. Squeezing less and moving faster are both attractive, and a few millimetres of slip in the fingers is often harmless. This suggests letting a learned policy trade a little slip for a lighter, faster lift, protected by a certificate that the slip stays small. So, is a little slip worth it? Mostly not. We study the question in MuJoCo simulation with explicit stick–slip friction, where friction drops from its static to its kinetic level once an object starts to slide. We derive a closed-form bound on slip during a minimum-jerk lift from a probed breakaway grip and an assumed bound r_{max} on the static-to-kinetic ratio, and use it as a shield that raises the grip or slows the lift until the bound is at most 3.5 mm. In a pre-registered study, a linear residual policy trained with evolution strategies cut lift time from 1.00 s to 0.37 s against a friction-model controller. Without the shield it learned to grip 6% below the model, exploiting the model's conservatism at short lift times, and 12–16% of its lifts slipped more than 5 mm once friction could drop; with the shield, at training or only at test time, none of 3,520 certified lifts did, at a cost of 6% more grip and no speed, and the results replicated on fresh objects. The certificate holds because the shield ends up forbidding slip altogether: when friction may drop by 10% or more, allowing even 10 mm of slip saves no grip. A controlled study that changes one design choice at a time (19 changes, 6 methods, frozen before it was run) shows that the shield's benefit survives heavier objects, softer pads, poor weight estimates and most training choices. It also finds two real weaknesses: the breakaway probe fails on slippery objects or with a noisy slip sensor, leaving the shield nothing to certify; and because the shield sits exactly at the onset boundary, it fails when friction is allowed to weaken during the creep that precedes slip. A margin on onset is the obvious fix.

1 Introduction

People lift objects with a grip only modestly above the force at which they would slip, scale that grip to the friction they feel, and correct it reflexively when the object starts to slip [1]. Robot grippers usually do the opposite. Because an object's weight and friction are unknown and a dropped object is costly, they squeeze hard and move slowly. Light grips matter for fragile, deformable or food items; fast lifts matter for throughput. Tactile slip detection and grip stabilisation [2, 3, 4, 5] make light grips safer by reacting when slip starts, but reaction comes after the fact, and a reflex cannot undo slip that has already happened.

A second idea is to plan the lift so that slip is bounded in advance. A friction model predicts the grip needed to hold an object during a known acceleration profile, and a learned policy

can adjust the model's choice where the model is conservative or wrong. This is *residual learning* [6, 7]. The risk is that a learned correction pushes the system exactly where the model is least reliable. In earlier work on in-hand pivoting [9] we found that this happens at *breakaway*: when a grasped object starts to slide, friction falls from its static to its kinetic level [12, 13], and a quasi-static model that ignores the drop under-predicts how far the object travels. Physics *shields* [22], which project a learned action onto a set certified by a friction model, prevented this in pivoting and in controlled sliding [10].

Lifting is the most common grasp task and the natural next test, and it raises a question the earlier tasks did not. If a few millimetres of slip are acceptable, a certificate need not forbid slip; it only needs to bound it. Can a policy then deliberately allow some slip and gain a lighter or faster lift, with a guarantee that the slip stays small? We answer this in simulation, in a study whose protocol was committed before any policy was trained (Fig. 1).

Our answer is a qualified no. A learned residual does lift 2.7× faster than a friction-model controller. Left unshielded, it also learns to grip 6% below the model, which is safe only as long as friction does not drop; with a drop, 12–16% of its lifts slip by 5 to 56 mm. A slip bound used as a shield removes every one of these failures at the cost of 6% more grip and no speed. But the shield's slip allowance turns out to be almost worthless: the bound behaves like a cliff, because once the object starts to slide the kinetic deficit persists for the rest of the acceleration phase. For any assumed drop of 10% or more, the only certified lifts are those in which slip never starts, and a 10 mm allowance buys no grip at all. **The practical lesson is that "a little slip" is cheap to tolerate but expensive to certify: unless friction is known not to drop, plan lifts in which slip does not start.**

Contributions.

- A ***closed-form slip bound for lifting*** (Proposition 1) that needs only a probed breakaway grip and an assumed bound $r_{\max}$ on μ_s/μ_k, used as a shield with two projection axes: raise the grip or slow the lift.
- A ***pre-registered simulation study*** (**Q1**–**Q4**, with a confirmation set run once) showing that an unshielded residual learns to exploit the quasi-static model's conservatism, that this fails once friction can drop, and that the shield removes every failure, with the certificate checked on 4,480 shielded lifts including lifts beyond $r_{\max}$.
- An ***analysis of why slip allowances are worth little***: the bound is a cliff in the assumed friction drop, quantified analytically, on an action-space safety map and in failure traces.
- A ***controlled one-change-at-a-time study*** (19 changes × 6 methods, frozen before it was run) of the environment, sensing, shield settings and training choices, which locates two failure modes, a probe that returns no information and friction that weakens before gross slip, and separates a simulator artefact from a real exposure with pre-registered diagnostics. All code, data and per-episode results are released.

2 Related Work

Grip force during lifting. Johansson and Westling showed that humans set grip force in proportion to load force, with a safety margin adapted to the friction of the contact, and that local slips trigger fast grip increases [1]. Robotic counterparts estimate friction or detect incipient slip with tactile sensors and raise the grip in response [4, 5, 2]; Romeo and Zollo survey slip-detection methods and sensors [3]. These approaches are reactive. We instead choose the grip and the motion *before* the lift, so that slip is bounded in advance, using only a probe of the breakaway grip; a reactive layer could sit on top.

Friction and stick–slip. The difference between static and kinetic friction has been measured for decades [12] and is central to models used in control, from Coulomb with a Stribeck drop to dynamic models such as LuGre [13, 14, 15]. Soft fingers add load-dependent friction and torsional effects [16, 17]. Physics engines approximate this behaviour: MuJoCo uses a convex, regularised contact model with a single friction coefficient [18], and Kolbert et al. found that simulated in-hand pivoting diverged from measurements [19]. Following our earlier work [10], we emulate the static-to-kinetic drop by switching the friction coefficient on the relative sliding speed, and we validate the resulting rig against the quasi-static model before training anything (Section 4).

Residual learning and safe learning. Residual policies learn corrections on top of a controller [6, 7]. Safe reinforcement learning restricts what a learner may do [20, 21]: shields override unsafe actions using a model [22]; control-barrier-function filters [23], reachability analysis [24] and predictive safety filters [25] generalise the idea. Our shield is a one-dimensional projection derived from contact mechanics. Its novelty is not the projection but what the certificate rests on: the static-to-kinetic transition during a commanded acceleration profile, which a single-μ Coulomb model, as used by rollout-based friction safety filters [26], cannot express.

Trajectories, learning and evaluation. We lift along a minimum-jerk profile [27], which fixes the peak acceleration for a given duration. Policies are linear and trained with antithetic evolution strategies [28, 29] and Adam [30]. Our evaluation follows recommendations for reliable RL reporting [31, 32]: multiple seeds, per-object paired tests and interval estimates, with the protocol committed before training [33]. Our audit of 2026 real-robot evaluations motivates reporting failure counts with intervals rather than success rates alone [11].

3 Certified Lifting

3.1 Task

A parallel gripper holds a rod of unknown mass $m \sim U[0.05, 0.2]$ kg and unknown static friction $\mu_s \sim U[0.4, 0.8]$ at its centre of mass, with gravity along the rod axis. Kinetic friction is $\mu_k = \mu_s/r$, with the ratio $r \geq 1$ also unknown. The gripper lifts the rod by $H = 10$ cm along a minimum-jerk trajectory of

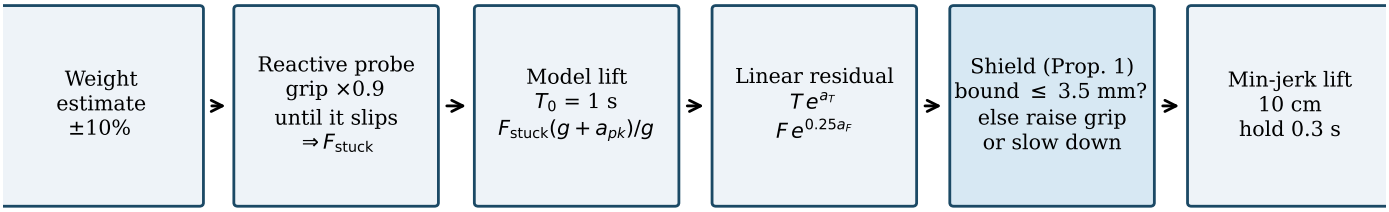


Figure 2. Episode pipeline. A wrist weight estimate and a reactive probe give the breakaway grip F_{stuck}; the friction-model controller proposes a lift; the residual policy rescales its duration and grip; the shield evaluates Proposition 1 and, if the slip bound exceeds 3.5 mm, raises the grip (grip axis) or lengthens the lift (time axis) until it does not, declining if neither suffices. Then the gripper lifts the object 10 cm along a minimum-jerk profile and holds for 0.3 s.

duration T [27],

$$z(t) = H\,(10s^3 - 15s^4 + 6s^5), \quad s = t/T, \qquad a_{pk} = \tfrac{10}{\sqrt{3}}\,\tfrac{H}{T^2} \approx 5.77\,\tfrac{H}{T^2}, \tag{1}$$

at a constant grip F per pad, then holds still for 0.3 s. The rod is supported only by friction. Two pads with Coulomb friction hold it while

$$m\,(g + a(t)) \;\le\; 2\mu_s F, \tag{2}$$

and once it slides its acceleration relative to the fingers is $g + a(t) - 2\mu_k F/m$. A lift *fails* if the rod slips more than $S_{\text{tol}} = 5$ mm relative to the fingers during the lift and hold, or is dropped. We want short lifts and low grips without failures.

3.2 Probing the breakaway grip

Before lifting, the gripper estimates the breakaway grip with the palm still. The grip starts at $\hat{m}g/(2 \cdot 0.3)$, the grip that would hold the estimated weight at $\mu = 0.3$, and is multiplied by 0.9 every 0.4 s (0.1 s ramp, 0.3 s hold) while a slip monitor with 0.5 mm noise watches. When the monitor reads more than 2.5 mm since the step began, the grip snaps to 2.5 times the grip that moved within 20 ms, so the probe never lets the rod run away. The weight estimate $\hat{m}$ is within $\pm 10\%$ of m (wrist force sensing). The probe returns F_{stuck}, the last grip at which the rod did not move. Since the rod held at F_{stuck}, (2) with $a = 0$ gives $2\mu_s F_{\text{stuck}} \ge mg$: F_{stuck} is an *upper* bound on the true breakaway grip $F_b = mg/(2\mu_s)$, and its step size makes it at most about 11% above it.

3.3 The friction-model controller

The quasi-static Coulomb lift uses a fixed duration $T_0 = 1$ s and the grip that just holds at peak acceleration if $r = 1$:

$$F_{\text{model}}(T) = F_{\text{stuck}}\,\frac{g + a_{pk}(T)}{g}. \tag{3}$$

It is slow and, because $F_{\text{stuck}} \ge F_b$, it carries a margin of up to 11%.

3.4 A slip bound for lifting

Proposition 1 (Lift slip bound)**.** *Let the palm follow an acceleration profile $a(t)$ from rest at constant grip F, let F_{stuck} be a grip at which the object was observed stuck with the palm at rest, and let $r \le r_{\max}$. Then slip cannot begin while $g + a(t) \le gF/F_{\text{stuck}}$. Let t_0 be the first instant with $g + a(t_0) > gF/F_{\text{stuck}}$ and define*

$$\bar{a}(t) = g + a(t) - \frac{gF}{r_{\max} F_{\text{stuck}}}, \qquad \dot{\bar{v}} = \bar{a} \text{ with } \bar{v} \ge 0 \text{ (clipped)},\ \bar{v}(t_0) = 0. \tag{4}$$

Then the slip of the object relative to the fingers is at most $\bar{s} = \int_{t_0} \bar{v}\,dt$.

Proof sketch (full proof in Appendix A). Because the object held at F_{stuck}, the static capacity per unit mass at grip F is $2\mu_s F/m = gF/F_b \ge gF/F_{\text{stuck}}$, so (2) holds while $g + a \le gF/F_{\text{stuck}}$. While sliding, the kinetic capacity is $2\mu_k F/m = gF/(rF_b) \ge gF/(r_{\max} F_{\text{stuck}})$, so the relative acceleration is at most $\bar{a}$; re-sticking can only lower it. A comparison argument gives $v(t) \le \bar{v}(t)$ and hence the bound. □

The bound needs no knowledge of m, μ_s or r beyond F_{stuck} and $r_{\max}$. We evaluate it over the whole commanded profile: the 0.6 s grip ramp and settle (during which the grip is at least F), the lift, and the 0.3 s hold. Figure 3 shows how it behaves.

3.5 The shield

A command (T, F) is *certified* if $\bar{s} \le s_{\text{lim}} = S_{\text{tol}} - 1.5$ mm, where 1.5 mm is a sensing margin (three standard deviations of the slip monitor). The *grip projection* raises F to the smallest certified grip (bisection, up to 30 N); the *time projection* lengthens T to the shortest certified duration (up to 2 s). If neither exists, the shield *declines* and the lift falls back to the model controller's command. The shield uses $r_{\max} = 1.2$ unless stated. It costs one bisection over a one-dimensional integral, under a millisecond in our implementation.

3.6 The residual policy

A linear policy maps six features of the probe to two actions through tanh: the log ratio of F_{stuck} to the estimated weight, the log ratio of the grip that moved to F_{stuck} (the probe's resolution), the estimated mass, the probe count, the grip scale and a bias. The actions rescale the model's command,

$$T = T_0\,e^{a_T} \in [0.25, 2]\,\text{s}, \qquad F = F_{\text{model}}(T)\,e^{0.25\,a_F}, \qquad a_T, a_F \in [-1, 1], \tag{5}$$

so zero weights reproduce the model controller and the grip action is relative to the model's grip *for the chosen duration*. The training reward is -10 for a failed lift and otherwise $-T/T_0 - \lambda F/(10mg)$ with $\lambda = 1$: one second of lift time is

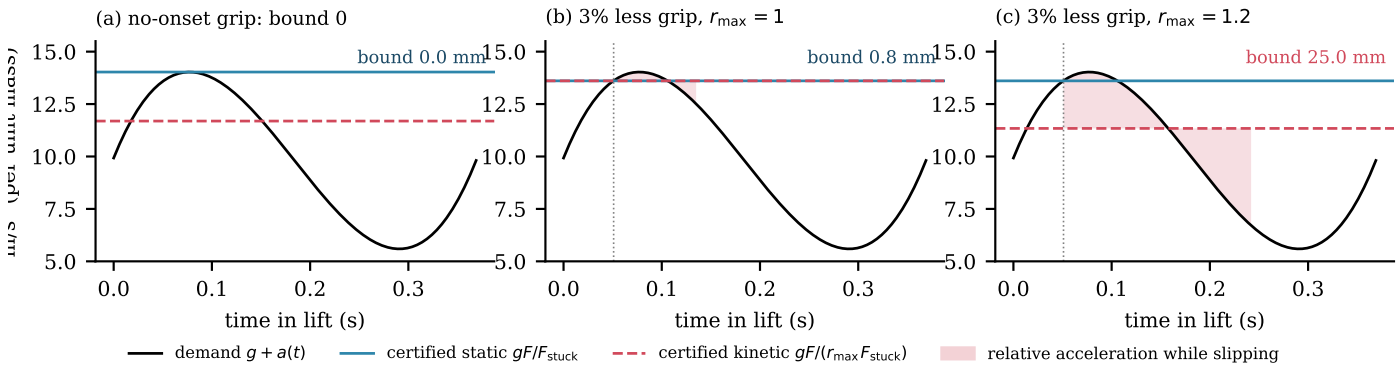


Figure 3. How Proposition 1 works (minimum-jerk lift, $T = 0.37$ s). (a) At the no-onset grip $F_{\text{model}}(T)$ the demand never exceeds the certified static capacity and the bound is zero. (b) With 3% less grip, slip may begin near peak acceleration; if friction does not drop ($r_{\max} = 1$) the excess is small and short, and the bound is 0.8 mm. (c) The same grip with a possible 20% drop: from the moment slip may begin, the certified kinetic capacity is far below the demand for most of the acceleration phase, and the bound jumps to 23 mm. This cliff is why slip allowances buy little (Section 6.4).

worth ten object-weights of grip. Training uses antithetic evolution strategies with rank-based fitness shaping [28, 29] and Adam [30] (Appendix C.2).

Remark 1 (Why the residual can beat the model)**.** *F_{model} is quasi-static and uses $F_{\text{stuck}} \geq F_b$. Both make it conservative: Table 1 shows that the grip needed during a fast lift is 2–11% below the quasi-static prediction, and F_{stuck} adds up to 11%. A residual can find this margin. It cannot know whether friction will drop once it uses it.*

4 Simulation Testbed

Rig. We simulate the task in MuJoCo 3 [18] at 1 ms steps with elliptic friction cones, an impedance ratio of 10 and 20 no-slip iterations. Two spherical pads (radius 6 mm, torsional friction) close on a $2\times2\times16$ cm rod; each pad is driven by a force-controlled slide joint. The palm moves on a vertical slide joint whose position and velocity are imposed at every step, so its acceleration follows the planned trajectory exactly. The rod is constrained to a vertical slide: a free rod rocked about the grasp axis after slip and then crept, which mixes in a pivoting motion we studied separately [8]; we study axial slip only. Appendix B lists the model.

Stick–slip. MuJoCo's contact model uses one friction coefficient. We emulate the drop from static to kinetic friction as in [10]: the pads' friction switches from μ_s to μ_s/r when the rod's speed relative to the fingers exceeds 2 cm/s, and back below 2 mm/s. We raised the switching thresholds during development, before any policy was trained, because slow elastic creep at lower thresholds triggered kinetic friction on rods that were not slipping.

Validation. Before training we checked the rig against the friction model (Table 1; details in Appendix B.3). The static breakaway grip matched $mg/(2\mu_s)$ within 1.1%. During lifts the grip at which the rod broke away was 2–11% below the quasi-static prediction $m(g+a_{pk})/(2\mu_s)$, more so for short lifts: the quasi-static model is conservative, increasingly so as lifts get faster. We did not isolate the cause; contact compliance and the short time the demand spends near its peak are candidates. In

Table 1. Rig validation (development objects only, run before training).

Check	Result
V1 static breakaway / $mg/(2\mu_s)$ ($n = 20$)	0.989 (0.986–1.014)
V2 lift breakaway / $m(g+a_{pk})/(2\mu_s)$:	
$T = 0.25$ / 0.35 / 0.5 / 1.0 s ($n = 8$ each)	0.890 / 0.930 / 0.952 / 0.983
V3 random lifts, $r \sim U[1, 1.2]$ ($n = 120$):	
slipped although bound was zero ($n = 67$)	0
largest slip above its bound (creep)	0.20 mm
held, non-zero bound: slip < bound	21 of 21
dropped (bound ≥ 0.56 m)	32
slip > 5 mm although bound ≤ 5 mm	0
probe slip, max; probes, median	5.3 mm; 8
V4 elastic creep, grip margin 1.1–2.0	0.09–0.19 mm

120 lifts with random grips and durations, no rod slipped whose bound was zero (beyond 0.2 mm of creep), every lift that kept the rod slipped less than its non-zero bound (the 32 rods that were dropped all had bounds above half a metre, i.e. no certificate at all), and no lift with a bound of at most 5 mm slipped more than 5 mm. Elastic creep in the pads, a stiff spring rather than friction slip, moved the rod 0.1–0.2 mm whatever the grip.

5 Experimental Setup

Methods. We compare six methods (Table 2). B0 is the friction-model controller and B0S adds the grip shield. B2 is a residual trained and tested without a shield. B4g and B4t apply the grip or time shield to B2's commands at test time only; B5 is a residual trained and tested with the grip shield in the loop. Residuals were trained with 5 seeds each on objects with $r \sim U[1.0, 1.2]$, $r_{\max} = 1.2$, for 1,560 episodes per run, starting from zero weights, i.e. from the model controller.

Evaluation sets. The *test set* has 40 new objects, each lifted at $r \in \{1.0, 1.1, 1.2, 1.3\}$; $r = 1.3$ lies beyond the shield's assumed $r_{\max} = 1.2$. Episode random seeds are shared across methods, so every method sees the same weight-estimate error and sensor noise on each object. The *confirmation set* has 40 further objects with $r \sim U[1.0, 1.3]$ per object and was run once, after the analysis code and the test-set results had been commit-

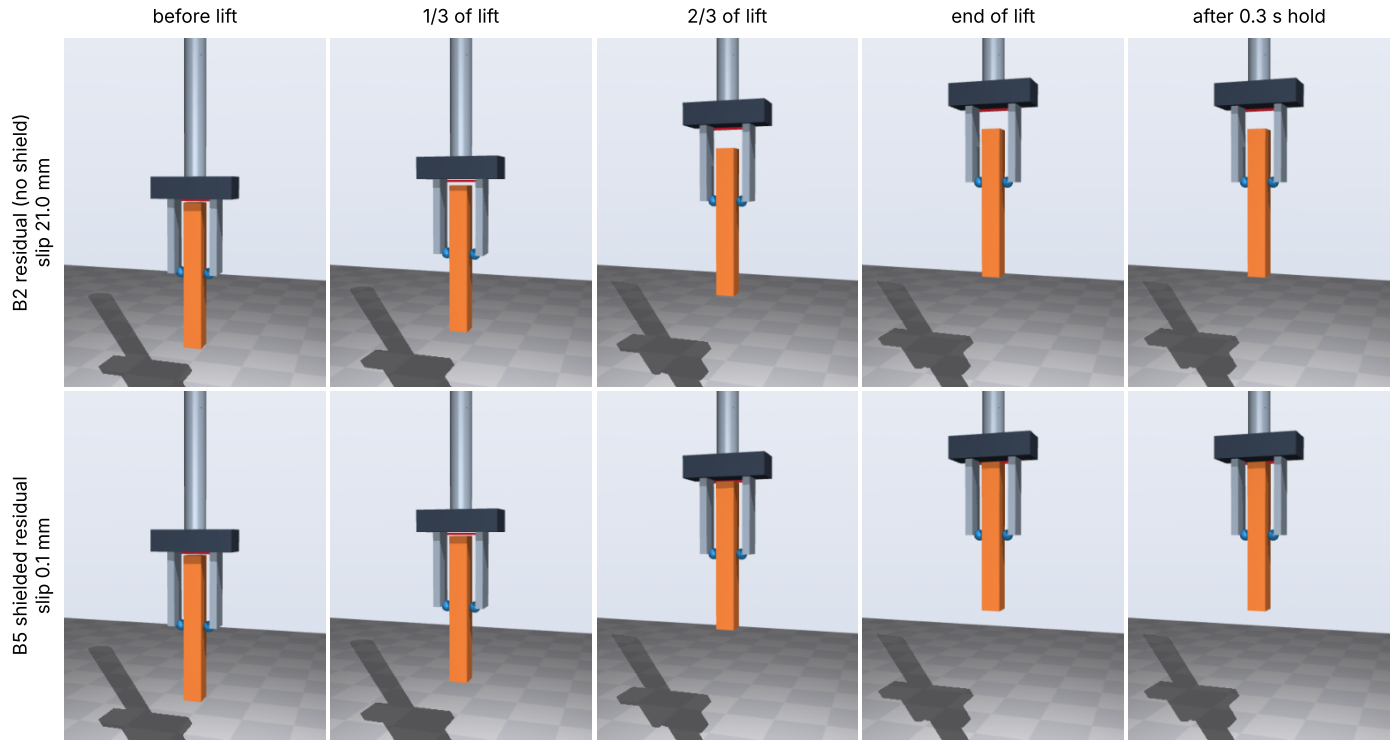


Figure 4. One lift, two policies (test object 2, 72 g, $\mu_s = 0.77$, $r = 1.2$; MuJoCo renders replayed from the recorded simulation states, with visual-only finger and palm geometry added). The red mark on the palm sits at the rod's initial top. The unshielded residual (top) grips 9% below the model's grip for its lift time; the rod breaks away near peak acceleration and ends 21 mm low. The shield-trained residual (bottom) lifts equally fast with a certified grip; the rod does not move relative to the fingers.

Table 2. Methods.

ID	Controller	Shield in training	Shield at test
B0	friction model	–	–
B0S	friction model	–	grip
B2	residual	no	no
B4g	residual (B2 weights)	no	grip
B4t	residual (B2 weights)	no	time
B5	residual	grip	grip

ted. For **Q4** we also evaluate B0S and B5 with the shield set to $r_{\max} \in \{1.0, 1.4\}$.

Questions. **Q1**: How much faster does a residual lift, and at what grip? **Q2**: Does it stay safe as friction drops? **Q3**: Should the shield raise the grip or slow the lift? **Q4**: Does the certificate hold, also when $r_{\max}$ is wrong, and what does a slip allowance buy?

Metrics and statistics. We report the failure count (slip > 5 mm or a drop) with Wilson 95% intervals [34], the mean lift time T, the mean grip relative to the object's weight F/mg, and the largest slip. Learned methods are averaged over the 5 seeds per object. We test lift time and grip against B0 with paired Wilcoxon signed-rank tests over objects [35], and per-object failures with exact McNemar tests [36], Holm-corrected within each question [37]; mean differences carry hierarchical bootstrap 95% intervals that resample seeds, then objects [38].

Table 3. Main results, test set ($r_{\max} = 1.2$). Lift time and grip are means over all lifts and do not depend on r; failures are lifts with slip > 5 mm out of 200 (40 for B0, B0S). No failure was a drop. Full statistics: Table 7.

	T (s)	F/mg	failures at $r =$ 1.0	1.1	1.2	1.3	max slip (mm)
B0 model	1.000	1.009	0	0	0	0	0.49
B0S model + shield	1.000	1.009	0	0	0	0	0.49
B2 residual	0.372	1.276	3	25	32	32	56.0
B4g B2 + grip shield	0.372	1.358	0	0	0	0	0.12
B4t B2 + time shield	0.420	1.276	0	0	0	0	0.14
B5 trained with shield	0.370	1.361	0	0	0	0	0.11

Pre-registration. The task, rig, questions, methods, sets, metrics and tests were committed to the repository before any policy was trained [33]. Three dated deviations are recorded there (Section 9). The controlled study, safety map, bound-tightness and trace analyses of Section 7 were designed after the confirmatory results were known; they were frozen in a protocol addendum before they were run and are reported as exploratory.

6 Results

6.1 Q1: How much faster, and at what grip?

Every residual variant lifts 2.7× faster than the model. B5 cut the lift time by 0.630 s (95% CI 0.629–0.630; Holm-corrected Wilcoxon $p < 10^{-10}$ at every r), from 1.00 s to 0.37 s (Table 3). All ten trained policies ended at the same corner of

the action space, $a_T \approx -1$ (Table 8), so the speed gain is set by our action range, not by the physics; Section 7 widens the range.

Faster lifts need more grip. Peak acceleration grows as $1/T^2$, so the model's own grip for a 0.37 s lift is 35% higher than for a 1 s lift. B5 gripped $0.352\,mg$ harder than B0 (CI 0.330–0.374). The unshielded residual gripped less, $+0.267\,mg$ (CI 0.237–0.304), because it set its grip on average 6% below the model's grip for its own lift time (range -11% to $+3\%$ across lifts; the five seeds chose $\bar{a}_F$ between -0.24 and -0.38).

The lift is a small part of an episode. The probe took a median of 4.1 s (7.5 probe steps, range 4–11), so the 0.63 s saved shortens a whole episode by about 10%. A faster probe is a larger opportunity than a faster lift.

6.2 Q2: Does the residual stay safe as friction drops?

No. B2's grip cut is nearly harmless without a friction drop: 3 of 200 lifts failed at $r = 1$, because the quasi-static model overpredicts the grip needed for short lifts by about 8% (Table 1) and F_{stuck} adds more margin. Once the rod does start to slide, kinetic friction is too low to stop it. Failures rose to 25, 32 and 32 of 200 at $r = 1.1$, 1.2 and 1.3, with slips up to 56 mm. Per object, B2 failed at least once on 14 objects where the model never failed at $r = 1.1$, and on 16 at $r = 1.2$ and 1.3, never the reverse (Holm-corrected McNemar $p = 0.001$ and $p = 0.0004$). The seeds cut differently: at $r = 1.2$ the five B2 policies failed 2, 6, 14, 2 and 8 times in 40 lifts. Every shielded method had zero failures at every r.

The failure has one mechanism. Figure 5 traces one failure. At 0.03 s, well before peak acceleration, the demand $g + a$ exceeds the static capacity and the rod starts to creep; at 0.09 s its speed relative to the fingers passes the 2 cm/s switching threshold, friction drops to its kinetic level, which lies well below the demand, and the rod slides at up to 0.2 m/s for 0.17 s, coming to rest 21 mm lower only once the palm decelerates. All 14 failures of the worst seed follow the same pattern (Fig. 12 in the appendix), with final slips of 9 to 25 mm. With a certified grip, the same lift shows only an elastic oscillation of ± 6 mm/s and at most 0.8 mm of transient displacement.

The result replicates. On the confirmation set, run once on 40 fresh objects with $r \sim U[1.0, 1.3]$, B2 failed 35 of 200 lifts (1, 16 and 18 in the bands $(1, 1.1]$, $(1.1, 1.2]$ and $(1.2, 1.3]$; 15 objects worse than the model, none better, $p = 6 \times 10^{-5}$), while B4g and B5 failed none at the same speed (0.371 s and 0.370 s) and the same extra grip ($+0.32\,mg$) as on the test set (Table 9).

6.3 Q3: Raise the grip, or slow the lift?

Applied to the same B2 commands, both projections intervened on 95.5% of lifts and removed every failure (Table 3). **The grip projection keeps the speed; the time projection keeps the grip.** Grip projection added 6.4% grip (1.358 vs $1.276\,mg$) at unchanged lift time; time projection kept B2's grip and slowed the lift by 13% (0.420 vs 0.372 s). Under our reward, which values one second of lift time like ten object-weights of grip, grip projection is better. Training with the shield in the loop (B5) gave the same result as adding it at test time (B4g): 0.370 vs 0.372 s and 1.361 vs $1.358\,mg$. The shield never declined: a certified command always existed within 30 N or 2 s.

6.4 Q4: Does the certificate hold, and what does a slip allowance buy?

The certificate held wherever its assumption did. Every certified lift at $r_{\max} = 1.2$ had a bound of exactly zero: the shield placed the grip where slip cannot start. Measured slip was 0.05–0.49 mm, which is the elastic creep of Table 1, not friction slip. Because any creep exceeds a zero bound, the pre-registered "exceedance" rate is 100% (deviation D1), but the largest excess is a third of the 1.5 mm margin and no certified lift came within 4.5 mm of the tolerance: 0 of 2,560 lifts at $r_{\max} = 1.2$, including 640 at $r = 1.3 > r_{\max}$, and 0 of 960 at $r_{\max} = 1.4$ (Table 10). Lifts beyond $r_{\max}$ were safe for a simple reason: slip never started, so the kinetic level never mattered.

Assuming friction does not drop is not safe. With $r_{\max} = 1.0$ the shield let B5 grip 6.4% less ($1.274\,mg$), and 60 of its 600 lifts at $r > 1$ slipped more than 5 mm (10, 25 and 25 at $r = 1.1$, 1.2, 1.3; worst excess over the bound 18 mm), against 0 of 200 at $r = 1$.

A slip allowance is worth little once friction may drop. Figure 1d and Fig. 3 explain why. Allowing 3.5 mm of slip saves 6.4% grip if friction does not drop, 4.6% if it may drop by 2%, 2.3% at 5%, 0.15% at 10% and nothing from 15%. Allowing 10 mm saves nothing from $r_{\max} = 1.2$. Once slip starts, the kinetic deficit $(g + a)(1 - 1/r_{\max})$ acts until the palm decelerates, so the certified slip is either zero or large. The ratio of the certified grip to the no-onset grip was identical for all 40 test objects, because it depends only on $r_{\max}$ and T (Appendix A.2).

7 Extended Study: What Survives When the Setup Changes?

The confirmatory study answers **Q1–Q4** for one configuration. A reader building on it needs to know which conclusions depend on our choices. We therefore ran a controlled study (overview in Fig. 13, Appendix E): starting from the test-set configuration, each of 19 rows changes *one* design choice and reruns all six methods on the same 40 objects, four values of r and the same episode seeds; four rows retrain the residuals. Four diagnostic rows (R0L, E7L, E9L, T2L) were added after some results were seen, each frozen before it was run. Figure 6 reports the change in failure rate, grip and episode time for every cell, with a paired bootstrap over objects marking changes whose 95% interval excludes zero (an unpaired bootstrap for E1 and E2, which draw different objects). Table 12 in the appendix describes each change and Table 13 gives absolute values. The study, and the analyses in Sections 7.4–7.5, were frozen in a protocol addendum before they were run and are exploratory.

7.1 Findings that hold

The unshielded residual is fragile, and the shield fixes it. B2's failure rate moved with almost every change: from 11.5% in the reference to 26% for heavier objects (E1), 30% with a

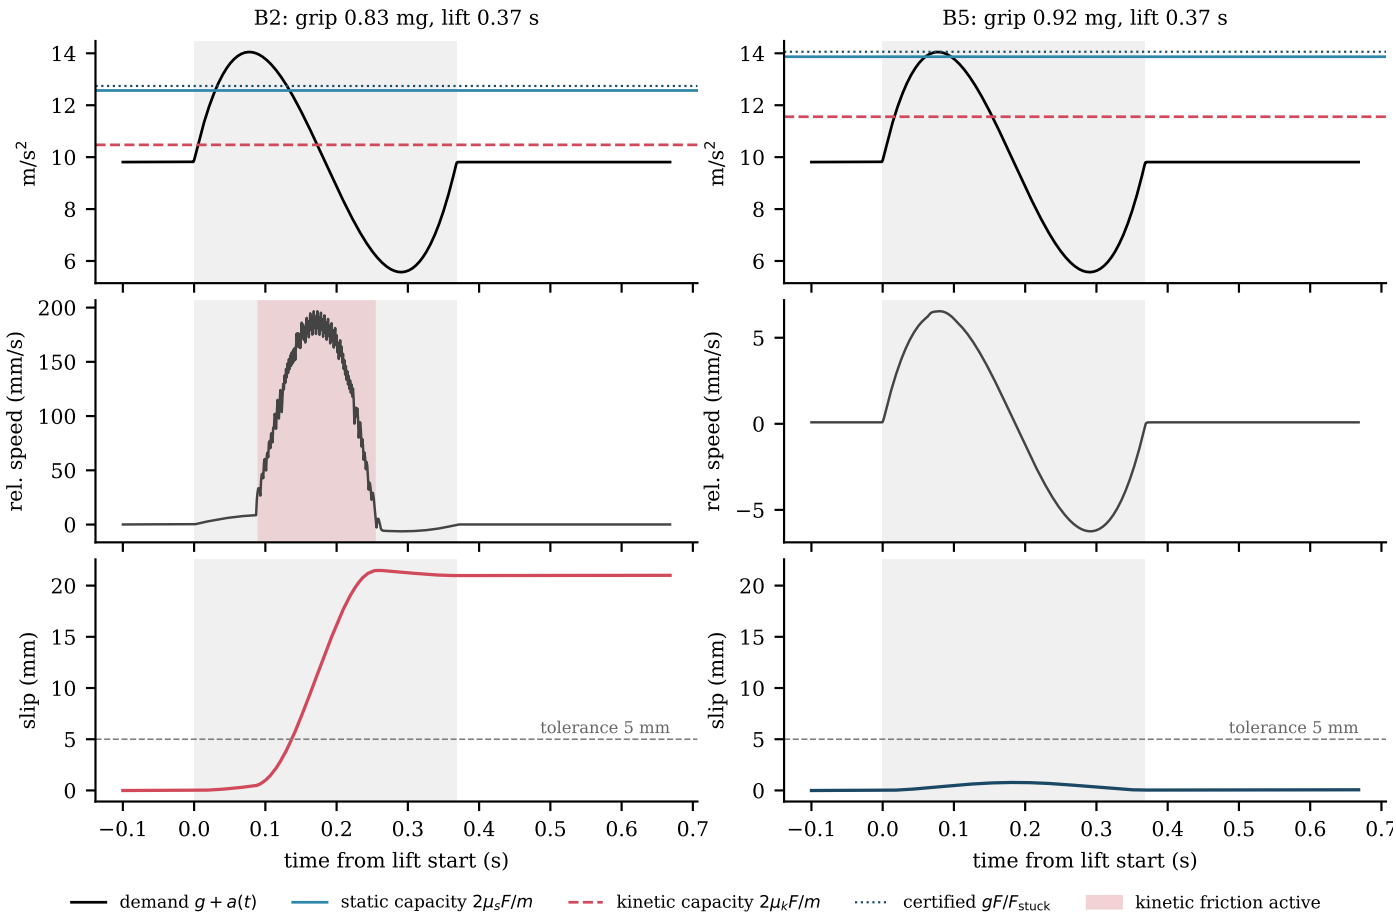


Figure 5. Anatomy of a failure (object 2, $r = 1.2$, seed 2). Top: demand $g + a(t)$ against the static, kinetic and certified capacities for the grip each policy chose. Middle: speed of the fingers relative to the rod, with kinetic friction active shaded. Bottom: slip. Left, B2: the demand crosses the static capacity at 0.03 s; the rod breaks away at 0.09 s, kinetic friction persists for 0.17 s, and slip ends at 21 mm. Right, B5: the certified grip keeps the demand below the static capacity throughout; the residual motion is elastic and returns to zero.

finer probe (E6), 15% with softer pads (E8) and 14% with a 2 mm tolerance (S4). In every one of these rows, and in the weight-error and probe-step rows, B4g, B4t and B5 had zero failures. A finer probe is an instructive case: it brings F_{stuck} closer to the true breakaway grip, which removes some of the hidden margin that B2 had learned to spend, so B2 fails more while the shielded methods do not notice.

The cliff holds. Setting the shield to $r_{\max} = 1.1$ (S2) lowered B5's grip by only 0.19 percentage points of the object's weight, and $r_{\max} = 1.4$ (S3) changed nothing: once any drop is assumed, the certified grip is the no-onset grip. Removing the 1.5 mm sensing margin (S6) or widening the tolerance to 10 mm (S5) did not change the shielded grip at all. Only $r_{\max} = 1.0$ (S1) bought grip, -7 to -9 points, and it cost 5–7.5% failures, as in **Q4**.

7.2 Findings that break

A probe that fails gives the shield nothing to certify. Our probe starts at the grip that would hold the estimated weight at $\mu = 0.3$. With slippery objects ($\mu_s \in [0.2, 0.4]$, E2) the first step already slips for many objects; with a noisier slip monitor (1.5 mm, E3) noise alone exceeds the 2.5 mm motion threshold on the first step. In both cases the probe reports $F_{\text{stuck}} = 30$ N, the maximum grip. The shield then certifies nothing: it declined on 68% (E2) and 100% (E3) of lifts and fell back to the slow model lift, at 21 to 31 times the object's weight in grip. This is a safe failure, but not a gentle one, and Fig. 6a shows why failure rates alone mislead here: B2's failure rate *fell* in these rows because it, too, was handed a crushing grip.

The certificate assumes friction stays static until gross slip, and that assumption can fail. Two rows broke the shield itself. Lifting 20 cm instead of 10 cm (E7) doubles the peak acceleration for the same duration; B5 then failed 38.6% of lifts and B4g 37.8%, every one with a certified bound of zero. Switching to kinetic friction at 5 mm/s instead of 2 cm/s (E9) gave 46.9% for both. Figure 7 shows the mechanism. The shield sets the grip at the edge of onset, where the demand reaches 95% of the true static capacity at peak acceleration. At that load MuJoCo's soft contact lets the rod creep relative to the pads, at up to 12 mm/s in E7 (6.5 mm/s in the reference). Our emulator decides when friction drops by comparing the rod's velocity from the previous step with the palm's com-

	(a) failed lifts (%), change						(b) grip (% of weight), change						(c) episode time (s), change					
	B0	B0S	B2	B4g	B4t	B5	B0	B0S	B2	B4g	B4t	B5	B0	B0S	B2	B4g	B4t	B5
Reference	0.0	0.0	11.5	0.0	0.0	0.0	101	101	128	136	128	136	6.03	6.03	5.40	5.40	5.45	5.40
E1 heavier objects 0.2-0.5 kg	0 (0.0)	0 (0.0)	**+14.5*** (26.0)	0 (0.0)	0 (0.0)	0 (0.0)	+1 (102)	+1 (102)	+2 (129)	+3 (139)	+2 (129)	+2 (138)	-0.05 (5.98)	-0.05 (5.98)	-0.05 (5.35)	-0.05 (5.35)	-0.04 (5.41)	-0.05 (5.35)
E2 slippery objects mu 0.2-0.4	0 (0.0)	0 (0.0)	**-8.1*** (3.4)	0 (0.0)	0 (0.0)	0 (0.0)	**+2008*** (2109)	**+2008*** (2109)	**+1920*** (2047)	**+1991*** (2127)	**+1995*** (2123)	**+1990*** (2126)	**-2.32*** (3.71)	**-2.32*** (3.71)	**-2.26*** (3.14)	**-1.89*** (3.51)	**-1.92*** (3.53)	**-1.89*** (3.51)
E3 slip-monitor noise 1.5 mm	0 (0.0)	0 (0.0)	**-11.5*** (0.0)	0 (0.0)	0 (0.0)	0 (0.0)	**+3007*** (3108)	**+3007*** (3108)	**+2867*** (2994)	**+2973*** (3108)	**+2981*** (3108)	**+2972*** (3108)	**-2.63*** (3.40)	**-2.63*** (3.40)	**-2.55*** (2.85)	**-2.00*** (3.40)	**-2.05*** (3.40)	**-2.00*** (3.40)
E4 weight error +-25%	0 (0.0)	0 (0.0)	-2.6 (8.9)	0 (0.0)	0 (0.0)	0 (0.0)	+1 (102)	+1 (102)	+1 (129)	+1 (137)	+1 (129)	+1 (137)	-0.02 (6.01)	-0.02 (6.01)	-0.02 (5.38)	-0.02 (5.38)	-0.02 (5.43)	-0.02 (5.38)
E5 coarser probe step 0.8	0 (0.0)	0 (0.0)	-4.9 (6.6)	0 (0.0)	0 (0.0)	0 (0.0)	**+5*** (106)	**+5*** (106)	**+7*** (135)	**+7*** (142)	**+7*** (135)	**+6*** (142)	**-1.26*** (4.77)	**-1.26*** (4.77)	**-1.26*** (4.14)	**-1.26*** (4.14)	**-1.26*** (4.19)	**-1.26*** (4.14)
E6 finer probe step 0.95	0 (0.0)	0 (0.0)	**+18.5*** (30.0)	0 (0.0)	0 (0.0)	0 (0.0)	**-3*** (98)	**-3*** (98)	**-5*** (123)	**-3*** (132)	**-5*** (123)	**-4*** (133)	**+2.56*** (8.59)	**+2.56*** (8.59)	**+2.56*** (7.96)	**+2.56*** (7.96)	**+2.57*** (8.02)	**+2.56*** (7.96)
E7 lift height 20 cm	0 (0.0)	0 (0.0)	**+46.4*** (57.9)	**+37.8*** (37.8)	**+4.1*** (4.1)	**+38.6*** (38.6)	**+6*** (106)	**+6*** (106)	**+38*** (166)	**+41*** (176)	**+38*** (166)	**+41*** (177)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	**-0.02*** (5.43)	0 (5.40)
E8 softer pads solref 0.006	0 (0.0)	0 (0.0)	+3.5 (15.0)	0 (0.0)	0 (0.0)	0 (0.0)	**-3*** (98)	**-3*** (98)	**-4*** (123)	**-4*** (132)	**-4*** (123)	**-4*** (132)	**+0.12*** (6.15)	**+0.12*** (6.15)	**+0.12*** (5.52)	**+0.12*** (5.52)	**+0.12*** (5.57)	**+0.12*** (5.52)
E9 earlier kinetic switch 5 mm/s	0 (0.0)	0 (0.0)	**+54.9*** (66.4)	**+46.9*** (46.9)	**+53.9*** (53.9)	**+46.9*** (46.9)	+0.10 (101)	+0.10 (101)	+0.13 (128)	+0.13 (136)	+0.13 (128)	+0.13 (136)	-0.005 (6.03)	-0.005 (6.03)	-0.005 (5.40)	-0.005 (5.40)	-0.01 (5.44)	-0.005 (5.40)
S1 shield r_max = 1.0	0 (0.0)	0 (0.0)	0 (11.5)	**+5.0*** (5.0)	**+5.0*** (5.0)	**+7.5*** (7.5)	0 (101)	0 (101)	0 (128)	**-7*** (129)	0 (128)	**-9*** (127)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	**-0.04*** (5.41)	0 (5.40)
S2 shield r_max = 1.1	0 (0.0)	0 (0.0)	0 (11.5)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	0 (128)	**-0.16*** (136)	0 (128)	**-0.19*** (136)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	**-0.000*** (5.45)	0 (5.40)
S3 shield r_max = 1.4	0 (0.0)	0 (0.0)	0 (11.5)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	0 (128)	0 (136)	0 (128)	0 (136)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	0 (5.45)	0 (5.40)
S4 slip tolerance 2 mm	0 (0.0)	0 (0.0)	**+2.9*** (14.4)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	0 (128)	0 (136)	0 (128)	0 (136)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	0 (5.45)	0 (5.40)
S5 slip tolerance 10 mm	0 (0.0)	0 (0.0)	**-3.0*** (8.5)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	0 (128)	0 (136)	0 (128)	0 (136)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	0 (5.45)	0 (5.40)
S6 no sensing margin	0 (0.0)	0 (0.0)	0 (11.5)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	0 (128)	0 (136)	0 (128)	0 (136)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	0 (5.45)	0 (5.40)
T1 trained on r = 1.0 only	0 (0.0)	0 (0.0)	**+12.9*** (24.4)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	**-2*** (125)	-0.02 (136)	**-2*** (125)	0 (136)	0 (6.03)	0 (6.03)	**-0.000*** (5.40)	**-0.000*** (5.40)	**+0.02*** (5.47)	0 (5.40)
T2 wider time action (scale 1.5)	0 (0.0)	0 (0.0)	**+30.8*** (42.3)	**+23.8*** (23.8)	**+13.8*** (13.8)	**+23.8*** (23.8)	0 (101)	0 (101)	**+47*** (175)	**+49*** (185)	**+47*** (175)	**+49*** (185)	0 (6.03)	0 (6.03)	**-0.12*** (5.28)	**-0.12*** (5.28)	**-0.15*** (5.30)	**-0.12*** (5.28)
T3 grip weight lambda = 10	0 (0.0)	0 (0.0)	**-6.6*** (4.9)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	**-2*** (125)	**-6*** (130)	**-2*** (125)	**-18*** (118)	0 (6.03)	0 (6.03)	**+0.03*** (5.43)	**+0.03*** (5.43)	**+0.01*** (5.46)	**+0.13*** (5.53)
T4 double training budget	0 (0.0)	0 (0.0)	**-11.0*** (0.5)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	**+4*** (131)	**+1*** (137)	**+4*** (131)	**+1*** (137)	0 (6.03)	0 (6.03)	**-0.003*** (5.40)	**-0.003*** (5.40)	**-0.02*** (5.42)	**-0.003*** (5.40)
R0L reference, lag-free switch	0 (0.0)	0 (0.0)	0 (11.5)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	0 (128)	0 (136)	0 (128)	0 (136)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	0 (5.45)	0 (5.40)
T2L wider time action, lag-free switch	0 (0.0)	0 (0.0)	**-9.1*** (2.4)	0 (0.0)	0 (0.0)	0 (0.0)	0 (101)	0 (101)	**+47*** (175)	**+49*** (185)	**+47*** (175)	**+49*** (185)	0 (6.03)	0 (6.03)	**-0.12*** (5.28)	**-0.12*** (5.28)	**-0.15*** (5.30)	**-0.12*** (5.28)
E7L lift height 20 cm, lag-free switch	0 (0.0)	0 (0.0)	**-2.9*** (8.6)	0 (0.0)	0 (0.0)	0 (0.0)	**+6*** (106)	**+6*** (106)	**+38*** (166)	**+41*** (176)	**+38*** (166)	**+41*** (177)	0 (6.03)	0 (6.03)	0 (5.40)	0 (5.40)	**-0.02*** (5.43)	0 (5.40)
E9L switch 5 mm/s, lag-free switch	0 (0.0)	0 (0.0)	**+54.4*** (65.9)	**+45.5*** (45.5)	**+23.4*** (23.4)	**+45.8*** (45.8)	+0.10 (101)	+0.10 (101)	+0.13 (128)	+0.13 (136)	+0.13 (128)	+0.13 (136)	-0.005 (6.03)	-0.005 (6.03)	-0.005 (5.40)	-0.005 (5.40)	-0.01 (5.44)	-0.005 (5.40)

Figure 6. Controlled study. Each cell is the change from the reference configuration (top row, absolute values), with the resulting absolute value in parentheses; * marks a 95% bootstrap interval that excludes zero. Red is worse (more failures, more grip, longer episodes). Failures count lifts that slip more than the row's tolerance or drop the rod, over 160 lifts (B0, B0S) or 800 lifts (learned methods, 5 seeds). Blocks, top to bottom: environment and sensing (E), shield settings (S), training (T), and diagnostics with the emulator's one-step lag removed (rows ending in L).

manded velocity for the current step; this one-step lag adds $a\,\Delta t$, about 8 mm/s at E7's peak acceleration, to the measured relative speed. Creep plus lag crossed the 20 mm/s threshold, friction dropped although Coulomb's law says the rod was stuck, and the rod slid. In the reference configuration the same quantity peaked at 10.8 mm/s, and no shielded lift ever switched to kinetic friction (0 of 2,560). The time projection was more robust in E7 (4.1% failures), because a longer lift lowers the peak demand instead of meeting it at the edge.

Removing the lag separates artefact from exposure. After seeing E7 and E9 we added diagnostic rows that repeat them with the lag removed, comparing rod and palm velocities from the same step (Addendum A6, frozen before running). In E7L the shielded failures disappeared: 0 of 2,560 shielded lifts failed, and the largest slip was 0.12 mm. The E7 failures were an artefact of our emulator. In E9L they did not: with friction allowed to drop at 5 mm/s, which the 6.5 mm/s creep at the onset edge (Fig. 5) exceeds without any lag, B4g, B4t and B5 still

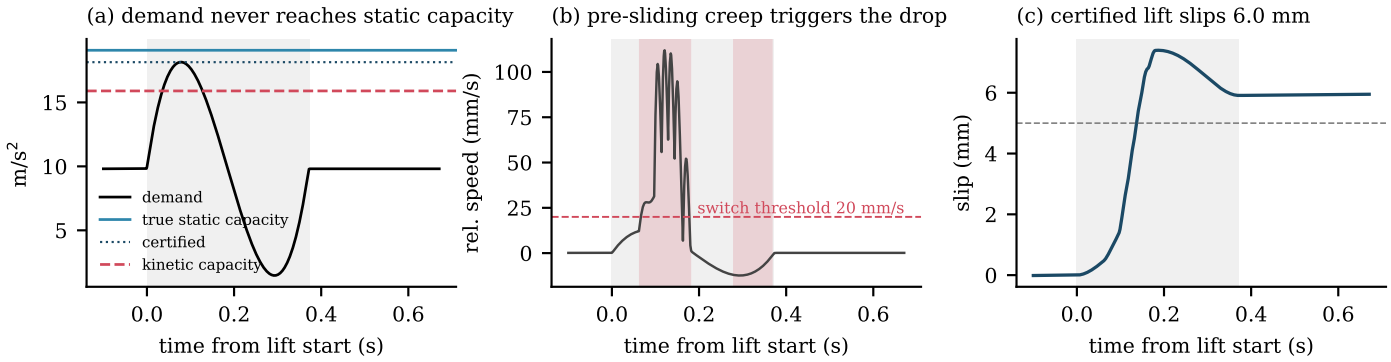


Figure 7. When the certificate fails (row E7, lift height 20 cm, B5, $r = 1.2$). (a) The demand peaks at 95% of the true static capacity, just below the certified level, so Coulomb friction should hold the rod. (b) The soft contact creeps at up to 12 mm/s before peak acceleration; the emulator's kinetic switch, which reads the relative speed with a one-step lag that adds about 8 mm/s here, fires at a true relative speed of 12 mm/s (kinetic friction shaded), and the rod slides. (c) The certified lift ends 6 mm low.

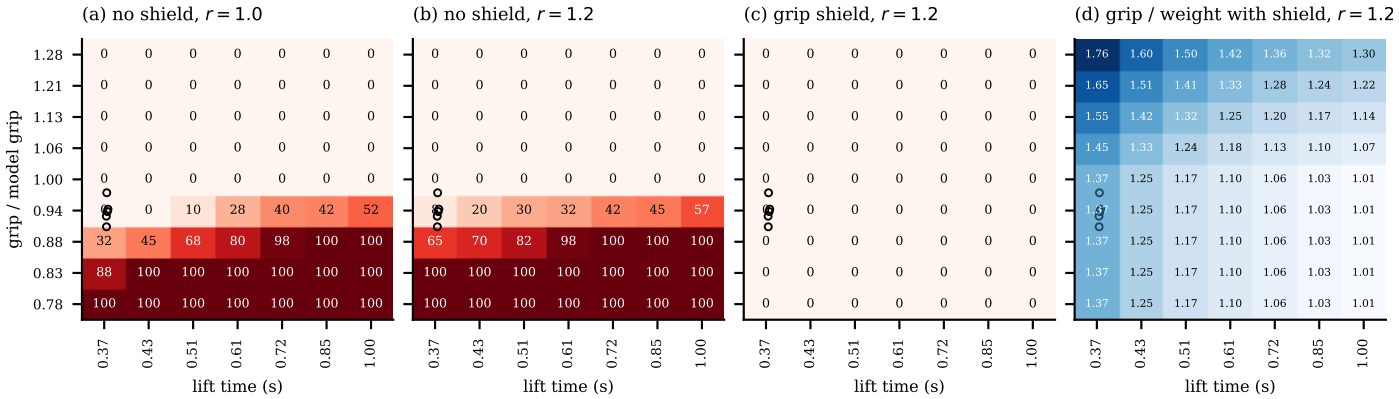


Figure 8. Action-space safety map. Fixed residual actions on the 40 test objects: lift time (columns) against grip relative to the model's grip for that lift time (rows). (a, b) Percentage of failed lifts without a shield at $r = 1.0$ and $r = 1.2$. (c) With the grip shield. (d) Mean grip relative to the object's weight after the shield. Circles mark the mean effective actions of the five trained B2 policies.

failed 364, 187 and 366 of 800 lifts. Repeating the reference configuration with the lag removed (row R0L) reproduced the main results exactly in every failure count (B2: 3, 25, 32 and 32 failures at $r = 1.0$–1.3; all shielded methods: 0), with 99% of episodes giving identical slip: the confirmatory results do not depend on the lag.

So E7 was an artefact, but E9 exposes something real. The shield places the grip exactly at the onset boundary, where the contact already creeps; if friction weakens during that creep, the certificate's premise (static friction until gross slip) fails. Real contacts also move before they slide: pre-sliding displacement is part of every dynamic friction model [14, 15], and whether friction weakens during it depends on the materials. **The practical fix is a margin on onset**, certifying $g + a \le (1 - \delta)\, gF/F_{\text{stuck}}$ for some $\delta > 0$, which our zero-margin shield did not have. We report this as found rather than adding the fix after seeing the result.

7.3 Training choices

Rows T1–T4 retrain B2 and B5 (five seeds each) and evaluate them in the reference environment (bottom block of Fig. 6).

Training without a friction drop makes the unshielded policy worse (T1). Trained on $r = 1$ only, B2 cut its grip further ($1.25\,mg$) and failed 24% of test lifts (195 of 800, against 92 in the reference); B5 and the test-time shields were unaffected.

Longer training made the unshielded policy safe on this test set (T4). With 40 instead of 20 iterations, B2 failed only 4 of 800 lifts (all from one seed), at 3.7 points more grip than the reference B2 and still 6 points less than B5. Part of B2's fragility in the main study is therefore a training-budget effect: with more episodes, the occasional -10 penalty for a slip at $r \in [1, 1.2]$ pushed the policy back from the edge. Nothing certifies this, and T1 shows how easily the lesson is unlearned when the training distribution omits the drop.

A grip-heavy reward separates training-time from test-time shielding (T3). With $\lambda = 10$, one object-weight of grip costs as much as a second of lift time. B5 then traded speed for grip, lifting in 0.50 s at $1.18\,mg$, while B4g, the shield applied only at test time to the fast B2 policy, lifted in 0.40 s at $1.30\,mg$. In the main study the two coincided because speed dominated the reward; when grip matters, putting the shield in the training loop lets the policy find commands that need little projection.

Wider action, faster lifts, and the lag again (T2). Allowed to lift in 0.25 s, every policy did so, at 1.75–$1.85\,mg$ of grip.

Certified lifts then failed at $r \geq 1.2$ (B5: 20 of 200 at $r = 1.2$, 170 of 200 at $r = 1.3$). The diagnostic row T2L, with the lag removed, had no shielded failures at any r, and B2's failures fell from 338 to 19: at 0.25 s the peak acceleration is 9.2 m/s^2 and the lag's $a\,\Delta t$ term crosses the switching threshold, as in E7.

7.4 An action-space safety map

Why does every policy end up at the fastest lift? We swept fixed residual actions over a 7×9 grid of lift times and grip scales on the 40 test objects (Fig. 8). Two facts explain the learned policies. First, **fast lifts tolerate grip cuts better**: at the model's grip minus 6% ($a_F = -0.25$), no lift failed at $T = 0.37$ s with $r = 1$, but 52% failed at $T = 1$ s. This is the quasi-static conservatism of Table 1 again: the faster the lift, the more the model over-predicts the grip it needs, so a policy rewarded for speed and light grip is pushed to the fastest corner, where the hidden margin is largest. The five B2 policies (circles) sit at that corner, 6–9% below the model's grip. Second, **the margin shrinks once friction may drop**: at $r = 1.2$ the same cut failed 8% of fast lifts and the safe region receded. With the grip shield, none of the 5,040 lifts in the grid failed, and the certified grip (panel d) depends only on the lift time, as the scale invariance of Appendix A.2 predicts.

7.5 How tight is the bound?

We lifted the 40 test objects with grips 0–15% below the no-onset grip, at three lift durations and three values of r, and evaluated Proposition 1 with $r_{\max}$ equal to the true r (Fig. 15). Three results stand out. **The bound is loose**: among the 1,309 lifts with a non-zero bound in which the rod stayed in the grasp, the median ratio of measured slip to bound was 0.023 (90th percentile 0.50). The probe's 0.9 step and the quasi-static model's conservatism together leave most under-gripped lifts with little or no slip. **A zero bound meant no slip**: the 360 lifts with a zero bound slipped at most 0.49 mm, all of it creep. **The bound can be exceeded during long slides**: 7 lifts, all at $T = 1$ s, $r = 1$ and 3% below the no-onset grip, slipped 26–44 mm against a bound of 25 mm. With $r = 1$ no friction switch is involved, so these slides expose a gap between the Coulomb model behind the bound and MuJoCo's regularised sliding friction [18]. They do not affect the shield, which never certifies bounds of this size, but they mark the limit of what the bound says about a simulated slide that lasts most of a second. The 491 lifts that dropped the rod all had bounds above 12 cm.

8 Discussion

Residuals find the model's margin, including the part that is not theirs to spend. The unshielded policy did exactly what its reward asked: it found that the friction model over-predicts the grip needed for a fast lift and spent that margin. Nothing in its training signal distinguished the margin that comes from quasi-static conservatism, which is real, from the margin that protects against a friction drop, which is needed only when the drop occurs. With r drawn from $[1, 1.2]$ during training, the cut paid off on most batches and failed on a few (Fig. 11). We expect the same pattern wherever a learned correction sits on a conservative model of a contact transition: the conservatism is visible in the average reward, the transition only in the tail.

Certify onset, not slip. The shield removed every failure and cost 6% grip. But it did so by forbidding slip, not by bounding it: every certified lift had a bound of zero. Section 6.4 and Appendix A.2 show why this is not an accident of our tolerances. With any friction drop, the kinetic deficit outlasts the moment of onset, so the certified slip jumps from zero to centimetres. For lifting, a useful certificate is therefore one about *onset*: $g + a(t) \leq gF/F_{\text{stuck}}$ throughout the profile. It is simple, it needs no $r_{\max}$ at all, and our results suggest it gives up almost nothing relative to the full slip bound once $r_{\max} \geq 1.1$. A slip allowance becomes valuable only when the friction drop is known to be small, for example on materials whose static and kinetic coefficients are measured to be close. The extended study adds one qualification: a certificate that sits exactly on the onset boundary inherits whatever happens there, and when friction weakened during pre-sliding creep (row E9) the certified lifts failed. An onset certificate should carry a margin, sized from measurements of how a contact creeps near its static limit.

Grip or time. When the shield must act, raising the grip kept the speed and slowing the lift kept the grip. Which is better depends on the application: for fragile objects the time projection is the natural choice, for throughput the grip projection. Both remove failures; neither declined in our study.

Where the time goes. The lift is a small part of an episode. Halving the probe time would save more than the learned lift does. Probes that use the lift itself, for example a slow initial acceleration that doubles as a breakaway test, are a promising direction.

9 Limitations

Simulation only. All results are from MuJoCo with an emulated stick–slip transition. Real contacts show rate-dependent friction, pre-sliding displacement, load-dependent soft-finger friction and wear [14, 16]; our switching thresholds were tuned in development, and the switch reads the relative speed with a one-step lag that we found only during the extended study. The lag did not change any confirmatory result (row R0L) but caused the certificate failures of row E7. The certificate assumes that F_{stuck} bounds the breakaway grip during the lift, which held here because the quasi-static model was conservative for fast lifts (Table 1); on hardware this must be measured. Domain randomisation [39, 40] over friction models would be a natural first step before hardware.

A probe with a blind spot. The probe assumes $\mu_s \geq 0.3$ and a slip monitor much quieter than its motion threshold; when either fails, it reports the maximum grip and the shield declines (rows E2, E3). Safe, but not gentle.

Narrow task. The rod is a rigid box constrained to slide axially, held at its centre of mass by two pads. Rotational slip, off-centre grasps, deformable objects and multi-finger hands are outside this study.

Policy class and budget. Policies are linear and were trained for 1,560 episodes; all of them saturated the duration action.

Training choices matter: twice the budget made the unshielded policy nearly safe on our test set, while training without a friction drop made it much less safe (Section 7.3). Richer policies might find other trade-offs.

Deviations. The pre-registered exceedance metric counts elastic creep (D1); the confirmation set was analysed pooled (D2); the slip-allowance analysis (D3) and the whole extended study are exploratory. Appendix F lists them.

10 Conclusion

We asked whether a learned lifting policy should be allowed to trade a little slip for a lighter, faster lift. A residual policy on a friction-model controller lifted 2.7 times faster; left unshielded, it also learned to spend a safety margin it could not see, and failed whenever friction dropped. A closed-form slip bound used as a shield removed every failure in simulation at 6% extra grip and no loss of speed, even beyond its assumed friction drop, and the results replicated on fresh objects. The slip allowance itself turned out to be worth little: once friction may drop by 10% or more, the only certifiable lifts are those in which slip never starts. Our recommendation for practice is to certify onset and to spend learning effort on faster probing. The next step is to measure the probe, the bound and the friction drop on hardware.

Broader Impact

This work aims to make robot grasping gentler and safer. Its main risk is over-trust: a certificate proved in simulation, or under an assumed friction drop, can fail on real objects, as our own $r_{\max} = 1.0$ results show. We report the conditions under which the certificate held and failed so that it is not applied beyond them.

Data Availability

Code, the frozen protocol with its addendum and deviations, trained weights, per-episode results and the scripts behind every number and figure: `https://github.com/CAN-Aidoo/lift-shield`, archived at doi:10.5281/zenodo.23769371.

Use of AI Tools

Code, analysis and drafting were carried out with the assistance of an AI model (Claude, Anthropic) under the author's direction. The author takes responsibility for the content.

References

[1] R. S. Johansson and G. Westling, "Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects," *Exp. Brain Res.*, vol. 56, no. 3, pp. 550–564, 1984.

[2] F. Veiga, B. Edin, and J. Peters, "Grip stabilization through independent finger tactile feedback control," *Sensors*, vol. 20, no. 6, art. 1748, 2020.

[3] R. A. Romeo and L. Zollo, "Methods and sensors for slip detection in robotics: a survey," *IEEE Access*, vol. 8, 2020.

[4] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, "Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor," in *Proc. IEEE-RAS Humanoids*, 2015.

[5] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson, "Measurement of shear and slip with a GelSight tactile sensor," in *Proc. IEEE ICRA*, 2015.

[6] T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Loskyll, J. A. Ojea, E. Solowjow, and S. Levine, "Residual reinforcement learning for robot control," in *Proc. IEEE ICRA*, 2019, arXiv:1812.03201.

[7] T. Silver, K. Allen, J. Tenenbaum, and L. Kaelbling, "Residual policy learning," arXiv:1812.06298, 2018.

[8] C. Aidoo, "Gravity pivoting from motor current: calibration-free angle estimation and the sensing budget it requires," engrXiv, doi:10.31224/8492, 2026.

[9] C. Aidoo, "Residual learning on a physics pivoting controller: faster search, and where the quasi-static model breaks," engrXiv, doi:10.31224/8542, 2026.

[10] C. Aidoo, "Physics guards for residual policies at breakaway: when learned corrections should defer to a friction model," engrXiv, doi:10.31224/8603, 2026.

[11] C. Aidoo, "How many trials back a force-aware manipulation claim? A statistical audit of 2026 real-robot evaluations," engrXiv, doi:10.31224/8576, 2026.

[12] E. Rabinowicz, "The nature of the static and kinetic coefficients of friction," *J. Appl. Phys.*, vol. 22, no. 11, pp. 1373–1379, 1951.

[13] B. Armstrong-Hélouvry, P. Dupont, and C. Canudas de Wit, "A survey of models, analysis tools and compensation methods for the control of machines with friction," *Automatica*, vol. 30, no. 7, pp. 1083–1138, 1994.

[14] C. Canudas de Wit, H. Olsson, K. J. Åström, and P. Lischinsky, "A new model for control of systems with friction," *IEEE Trans. Autom. Control*, vol. 40, no. 3, pp. 419–425, 1995.

[15] H. Olsson, K. J. Åström, C. Canudas de Wit, M. Gäfvert, and P. Lischinsky, "Friction models and friction compensation," *Eur. J. Control*, vol. 4, no. 3, pp. 176–195, 1998.

[16] N. Xydas and I. Kao, "Modeling of contact mechanics and friction limit surfaces for soft fingers in robotics, with experimental results," *Int. J. Robot. Res.*, vol. 18, no. 9, pp. 941–950, 1999.

[17] R. D. Howe and M. R. Cutkosky, "Practical force-motion models for sliding manipulation," *Int. J. Robot. Res.*, vol. 15, no. 6, pp. 557–572, 1996.

[18] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: a physics engine for model-based control," in *Proc. IEEE/RSJ IROS*, 2012, pp. 5026–5033.

[19] R. Kolbert, N. Chavan-Dafle, and A. Rodriguez, "Experimental validation of contact dynamics for in-hand manipulation," in *Proc. ISER*, 2016, arXiv:1702.07252.

[20] J. García and F. Fernández, "A comprehensive survey on safe reinforcement learning," *J. Mach. Learn. Res.*, vol. 16, pp. 1437–1480, 2015.

[21] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: from learning-based control to safe reinforcement learning," *Annu. Rev. Control Robot. Auton. Syst.*, 2022, arXiv:2108.06266.

[22] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe reinforcement learning via shielding," in *Proc. AAAI*, 2018, arXiv:1708.08611.

[23] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: theory and applications," in *Proc. ECC*, 2019, arXiv:1903.11199.

[24] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A general safety framework for learning-based control in uncertain robotic systems," *IEEE Trans. Autom. Control*, 2019, arXiv:1705.01292.

[25] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," *Automatica*, 2021, arXiv:1812.05506.

[26] A. Johansson, D. Lindmark, V. Wiberg, and M. Servin, "Safety filtering of robotic manipulation under environment uncertainty: a computational approach," arXiv:2509.12674, 2025.

[27] T. Flash and N. Hogan, "The coordination of arm movements: an experimentally confirmed mathematical model," *J. Neurosci.*, vol. 5, no. 7, pp. 1688–1703, 1985.

[28] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, "Evolution strategies as a scalable alternative to reinforcement learning," arXiv:1703.03864, 2017.

[29] H. Mania, A. Guy, and B. Recht, "Simple random search of static linear policies is competitive for reinforcement learning," in *Proc. NeurIPS*, 2018, arXiv:1803.07055.

[30] D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," in *Proc. ICLR*, 2015, arXiv:1412.6980.

[31] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," in *Proc. AAAI*, 2018, arXiv:1709.06560.

[32] R. Agarwal, M. Schwarzer, P. S. Castro, A. Courville, and M. G. Bellemare, "Deep reinforcement learning at the edge of the statistical precipice," in *Proc. NeurIPS*, 2021, arXiv:2108.13264.

[33] B. A. Nosek, C. R. Ebersole, A. C. DeHaven, and D. T. Mellor, "The preregistration revolution," *Proc. Natl. Acad. Sci. USA*, vol. 115, no. 11, pp. 2600–2606, 2018.

[34] E. B. Wilson, "Probable inference, the law of succession, and statistical inference," *J. Amer. Statist. Assoc.*, vol. 22, no. 158, pp. 209–212, 1927.

[35] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics Bull.*, vol. 1, no. 6, pp. 80–83, 1945.

[36] Q. McNemar, "Note on the sampling error of the difference between correlated proportions or percentages," *Psychometrika*, vol. 12, no. 2, pp. 153–157, 1947.

[37] S. Holm, "A simple sequentially rejective multiple test procedure," *Scand. J. Statist.*, vol. 6, no. 2, pp. 65–70, 1979.

[38] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*. Chapman & Hall, 1993.

[39] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *Proc. IEEE/RSJ IROS*, 2017, arXiv:1703.06907.

[40] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *Proc. IEEE ICRA*, 2018, arXiv:1710.06537.

Appendix Contents

A Proofs and Analysis 13
A.1 Proof of Proposition 1 13
A.2 Why a slip allowance is a cliff 13
B Simulation Details 14
B.1 Model and parameters 14
B.2 Development decisions 15
B.3 Rig validation 15
C Controller, Shield and Training 16
C.1 Probe, bound and shield code 16
C.2 Training details 18
D Additional Main Results 19
E Extended Study Details 22
F Protocol, Deviations, Reproducibility 24

List of Appendix Figures and Tables

Fig. 9 Certified grip vs. assumed drop 14
Fig. 10 Rig validation 16
Fig. 11 Training curves 18
Fig. 12 All failures of the worst seed 20
Table 5 Simulation parameters 15
Table 6 Training hyperparameters 18
Table 7 Test set, all methods and r 19
Table 8 Per-seed results 19
Table 9 Confirmation set 20
Table 10 Mismatched $r_{\max}$ 20
Table 11 Per-object results 21
Fig. 14 Controlled study: lift time 23
Table 12 Controlled changes 22
Table 13 Controlled study, absolute 24
Table 14 Training rewards per seed 25
Fig. 15 Bound against measured slip 24
Table 15 Compute 25

A Proofs and Analysis

A.1 Proof of Proposition 1

Setting. The rod (mass m) is held between two pads, each pressing with normal force F. Let $x(t) \geq 0$ be the downward displacement of the rod relative to the fingers and $v = \dot{x}$. The palm, and with it the fingers, moves up with acceleration $a(t)$. In the palm frame the rod feels gravity plus the inertial force, a downward demand of $m(g+a)$ per unit time, and a friction force from the two pads opposing relative motion. With Coulomb friction and a static-to-kinetic drop:

- while stuck ($v = 0$), the friction force can take any value up to $2\mu_s F$, so the rod stays stuck while $m(g+a) \leq 2\mu_s F$;
- while sliding down ($v > 0$), the friction force is $2\mu_k F$ upward, so $m\dot{v} = m(g+a) - 2\mu_k F$.

Step 1 (what the probe certifies). The probe observed the rod stuck at grip F_{stuck} with the palm at rest ($a = 0$) for 0.3 s. Hence $mg \leq 2\mu_s F_{\text{stuck}}$, i.e. $F_{\text{stuck}} \geq F_b := mg/(2\mu_s)$. For any grip F the static capacity per unit mass is $2\mu_s F/m = gF/F_b \geq gF/F_{\text{stuck}}$.

Step 2 (no slip before t_0). For $t < t_0$, $g + a(t) \leq gF/F_{\text{stuck}} \leq 2\mu_s F/m$, so the stick condition holds and $x(t) = v(t) = 0$.

Step 3 (relative acceleration after t_0). With $r = \mu_s/\mu_k \leq r_{\max}$, the kinetic capacity per unit mass is $2\mu_k F/m = gF/(rF_b) \geq gF/(r_{\max}F_{\text{stuck}})$. Whenever the rod slides, $\dot{v} = g + a - 2\mu_k F/m \leq \bar{a}$. Whenever it is stuck, $\dot{v} = 0$, and $v = 0 \leq \bar{v}$.

Step 4 (comparison). Let $\bar{v}$ solve $\dot{\bar{v}} = \bar{a}$ with $\bar{v}$ clipped at zero, from $\bar{v}(t_0) = 0$. Suppose $v(t) > \bar{v}(t)$ for some t, and let $t_1 = \sup\{s < t : v(s) \leq \bar{v}(s)\}$. On $(t_1, t]$ we have $v > \bar{v} \geq 0$, so the rod slides and $\dot{v} \leq \bar{a} \leq \dot{\bar{v}}$ wherever $\bar{v} > 0$; where $\bar{v} = 0$ is clipped, $\bar{a} \leq 0$ and again $\dot{v} \leq \bar{a} \leq 0$. Integrating from t_1, where $v(t_1) \leq \bar{v}(t_1)$ by continuity, gives $v(t) \leq \bar{v}(t)$, a contradiction. Hence $v \leq \bar{v}$ for all $t \geq t_0$, and $x(t) = \int_{t_0}^{t} v \leq \int_{t_0}^{t} \bar{v} = \bar{s}(t)$. □

What the proof does not cover. (i) Elastic compliance: the rod can move relative to the fingers without slipping by deforming the contact; we measured 0.1–0.5 mm and cover it with the 1.5 mm margin. (ii) Rate dependence and pre-sliding friction [14]: the proof assumes a static level that holds up to breakaway and a single kinetic level after it. (iii) Loss of grasp: once the rod's end passes the pads, there is no friction at all; this only happens for bounds of many centimetres, which the shield never certifies. (iv) The grip ramp: our grip ramps down to F over 0.3 s from a higher grip, so using F throughout the rest phase is conservative.

A.2 Why a slip allowance is a cliff

Scale invariance. Proposition 1 depends on the grip only through $\rho = F/F_{\text{stuck}}$: the onset condition is $g + a(t) > g\rho$ and $\bar{a} = g + a - g\rho/r_{\max}$. The certified grip is therefore $F^\star = F_{\text{stuck}}\,\rho^\star(T, r_{\max}, s_{\lim})$ and its ratio to the no-onset grip $F_{\text{model}} = F_{\text{stuck}}(g + a_{pk})/g$ is a function of T, $r_{\max}$ and $s_{\lim}$ only. This is why the ratio was identical for all 40 test objects.

The cliff. Write $\rho = (1-\varepsilon)(g + a_{pk})/g$ for a grip a fraction ε below the no-onset grip. Slip may begin a little before the acceleration peak, at the first t_0 with $a(t_0) > a_{pk} - \varepsilon(g + a_{pk})$. From then on the bound integrates

$$\bar{a}(t) = g + a(t) - \frac{(1-\varepsilon)(g+a_{pk})}{r_{\max}} = \underbrace{\big(a(t) - a_{pk}\big) + \varepsilon(g + a_{pk})}_{\text{as without a drop}} + \underbrace{(1-\varepsilon)(g+a_{pk})\Big(1 - \frac{1}{r_{\max}}\Big)}_{\text{kinetic deficit}}. \tag{6}$$

Without a drop ($r_{\max} = 1$) only the first term acts; it is positive only near the peak and integrates to a small slip that grows smoothly with ε. With a drop, the second term is a constant of order $(g + a_{pk})(1 - 1/r_{\max})$: for $T = 0.37$ s and $r_{\max} = 1.2$ it is about 2.3 m/s^2, larger than the first term's whole positive range for small ε. It keeps $\bar{v}$ growing until the minimum-jerk profile decelerates enough

to cancel it, roughly a third of the lift later, so even an infinitesimal ε certifies centimetres of slip. The certified grip therefore sits at $\varepsilon = 0$ unless the allowance s_{lim} is larger than this jump. Figure 9 shows the result for three allowances and three lift durations.

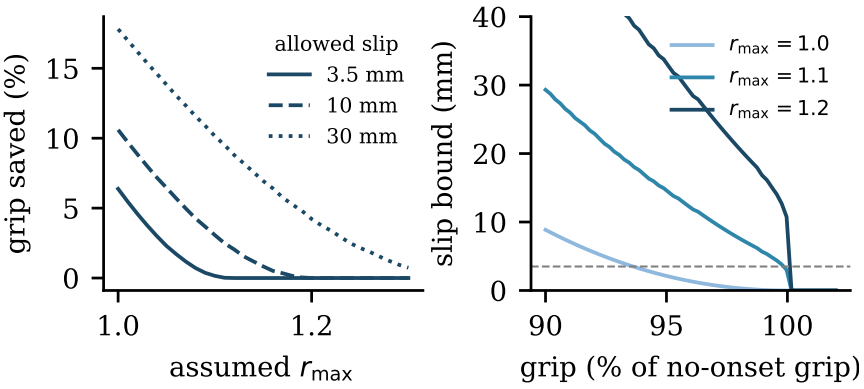


Figure 9. The cliff. Left: grip saved by allowing 3.5, 10 or 30 mm of slip at $T = 0.37$ s, as a function of the assumed drop $r_{\max}$. Right: Proposition 1's bound as a function of grip for three values of $r_{\max}$; with any drop, the bound jumps from zero to centimetres as soon as slip can begin.

Table 4. Grip saved by a slip allowance (% of the no-onset grip), computed from Proposition 1; the values are the same for every object.

T (s)	allowance (mm)	$r_{\max} = 1.0$	$r_{\max} = 1.02$	$r_{\max} = 1.05$	$r_{\max} = 1.1$	$r_{\max} = 1.15$	$r_{\max} = 1.2$	$r_{\max} = 1.3$
0.37	3.5	6.38	4.59	2.30	0.15	0.00	0.00	0.00
0.37	10	10.60	8.85	6.45	3.07	0.83	0.00	0.00
0.37	30	17.81	16.18	13.83	10.15	6.94	4.22	0.74
0.50	3.5	4.04	2.27	0.43	0.00	0.00	0.00	0.00
0.50	10	6.72	4.93	2.61	0.23	0.00	0.00	0.00
0.50	30	11.28	9.54	7.10	3.61	1.12	0.02	0.00
1.00	3.5	1.18	0.00	0.00	0.00	0.00	0.00	0.00
1.00	10	1.96	0.36	0.00	0.00	0.00	0.00	0.00
1.00	30	3.29	1.53	0.02	0.00	0.00	0.00	0.00

B Simulation Details

B.1 Model and parameters

Listing B.1 is the complete MuJoCo model; {m} and {mu} are filled per object. The palm body is massive (1000 kg), gravity-compensated and driven kinematically: before and after every physics step its slide joint's position and velocity are overwritten with the planned trajectory, so the pads move with the planned acceleration regardless of contact forces. Table 5 lists every parameter.

MuJoCo model (Listing 0)

```
<mujoco>
  <option timestep="0.001" integrator="implicitfast" cone="elliptic" impratio="10" noslip_iterations="20" gravity="0 0 -9.81"/>
  <worldbody>
    <body name="palm" pos="0 0 0.5" gravcomp="1"><joint name="pz" type="slide" axis="0 0 1" damping="0"/><geom name="palmg" type="box"
     size="0.002 0.002 0.002" mass="1000" contype="0" conaffinity="0"/>
      <body name="lf"><joint name="lj" type="slide" axis="1 0 0" damping="300" range="0 0.03" limited="true"/>
        <geom name="lpad" type="sphere" size="0.006" pos="-0.0262 0 0" mass="0.02" condim="4" friction="{mu} 0.005 0.0001" solref="0.003
       1" solimp="0.95 0.99 0.0005"/></body>
      <body name="rf"><joint name="rj" type="slide" axis="-1 0 0" damping="300" range="0 0.03" limited="true"/>
        <geom name="rpad" type="sphere" size="0.006" pos="0.0262 0 0" mass="0.02" condim="4" friction="{mu} 0.005 0.0001" solref="0.003
      1" solimp="0.95 0.99 0.0005"/></body>
    </body>
    <body name="rod" pos="0 0 0.5"><joint name="rz" type="slide" axis="0 0 1"/>
      <geom name="rodg" type="box" size="0.01 0.01 0.08" mass="{m}" condim="4" friction="{mu} 0.005 0.0001" solref="0.003 1" solimp="
     0.95 0.99 0.0005"/>
    </body>
  </worldbody>
  <actuator>
    <motor name="lgrip" joint="lj" gear="1" ctrlrange="0 40"/>
    <motor name="rgrip" joint="rj" gear="1" ctrlrange="0 40"/>
  </actuator>
</mujoco>
```

Stick–slip emulation: one physics step

```
def mj1(self):
    # the palm is driven kinematically: its position and velocity are imposed every step
    self.data.qpos[self._pq] = self.palm_z - 0.5; self.data.qvel[self._pv] = self.palm_v
    if self.ratio != 1.0 and self.armed:
        mujoco.mj_objectVelocity(self.model, self.data, mujoco.mjtObj.mjOBJ_BODY, self._rod, self._vel, 0)
        v = abs(self._vel[5] - self.palm_v)
        if not self.sliding and v > self.v_slip:
            self.sliding = True; self.n_breakaway += 1
            self.model.geom_friction[self._g, 0] = self.mu_k
        elif self.sliding and v < self.v_stick:
            self.sliding = False
            self.model.geom_friction[self._g, 0] = self.mu_s
    mujoco.mj_step(self.model, self.data)
    self.data.qpos[self._pq] = self.palm_z - 0.5; self.data.qvel[self._pv] = self.palm_v
```

Table 5. Simulation and task parameters.

Parameter	Value	Parameter	Value
time step	1 ms (implicit-fast)	friction cone	elliptic, impedance ratio 10
no-slip iterations	20	pad geometry	spheres, radius 6 mm, condim 4
contact solref / solimp	0.003, 1 / 0.95, 0.99, 0.0005	torsional / rolling friction	0.005 / 0.0001
rod	box $2\times2\times16$ cm on vertical slide	finger joints	slide, damping 300, force control
object mass m	$U[0.05, 0.2]$ kg	static friction μ_s	$U[0.4, 0.8]$
kinetic switch	$\lvert v_{\text{rel}}\rvert > 2$ cm/s; back below 2 mm/s	ratio r (test)	1.0, 1.1, 1.2, 1.3
lift height H	10 cm	lift profile	minimum jerk
grip ramp + settle	0.3 s + 0.3 s	hold after lift	0.3 s
slip tolerance S_{tol}	5 mm	sensing margin	1.5 mm
slip-monitor noise	0.5 mm (Gaussian)	weight estimate	$m \cdot U[0.9, 1.1]$
probe start	$\hat{m}g/(2 \cdot 0.3)$	probe step	$\times 0.9$ per 0.4 s
probe motion threshold	2.5 mm	probe re-clamp	$2.5\times$ moving grip in 20 ms
maximum grip	30 N	model lift time T_0	1 s
shield $r_{\max}$	1.2	lift time range	0.25–2 s

B.2 Development decisions

The rig went through several versions before the protocol was frozen. We list every change that affected behaviour, in order, because each one is a modelling choice a reader may want to revisit. None was made after a policy was trained.

1. **Driven palm.** A mocap (teleported) palm has no velocity, so contact friction never dragged the rod upward. We replaced it with a slide joint whose position and velocity are imposed.
2. **Gravity compensation.** The palm sagged under gravity within each step and the rod crept; we made the palm heavy (1000 kg) and gravity-compensated.
3. **Reactive probe.** With $r > 1$, holding a probe grip that turned out too low let the rod slide away. The probe now watches the slip and re-clamps at once.
4. **Re-clamp level.** Snapping to the maximum grip and then ramping down produced long, large ramps; the probe now snaps to 2.5 times the grip that moved, and the lift's grip ramp lasts 0.3 s. The bound's rest phase was lengthened to match (0.6 s).
5. **Switching thresholds.** Slow elastic creep at the original 5 mm/s threshold switched rods to kinetic friction although they were not slipping; we raised the thresholds to 2 cm/s and 2 mm/s.
6. **Axial slide.** A free rod rocked about the grasp axis after slip (about 0.24 rad/s) and crept afterwards, so the model controller dropped objects with a zero bound. We constrained the rod to a vertical slide. After this change the model controller held 40 of 40 development objects at $r \in [1, 1.2]$.

B.3 Rig validation

V1 measures the grip at which a rod first moves with the palm still, by bisection over grips, for 20 development objects at $r = 1$, against $mg/(2\mu_s)$. **V2** does the same during lifts of duration 0.25–1 s, against the quasi-static $m(g + a_{pk})/(2\mu_s)$. **V3** runs 120 complete episodes with a probe, $r \sim U[1, 1.2]$, a random duration $T \sim U[0.25, 1]$ s and a random grip between 0.6 and 1.4 times the quasi-static need, and compares the measured slip with the bound at $r_{\max} = 1.2$. **V4** measures creep in well-held lifts at four grip margins. Figure 10 shows the results. The bound is loose wherever it is non-zero, often by one to two orders of magnitude (panel b): with F_{stuck} above the true breakaway grip and the quasi-static model conservative for fast lifts, a rod that the bound allows to slip

usually does not slip at all. This looseness is one reason the shield ends up forbidding slip altogether; Section 7 measures it directly.

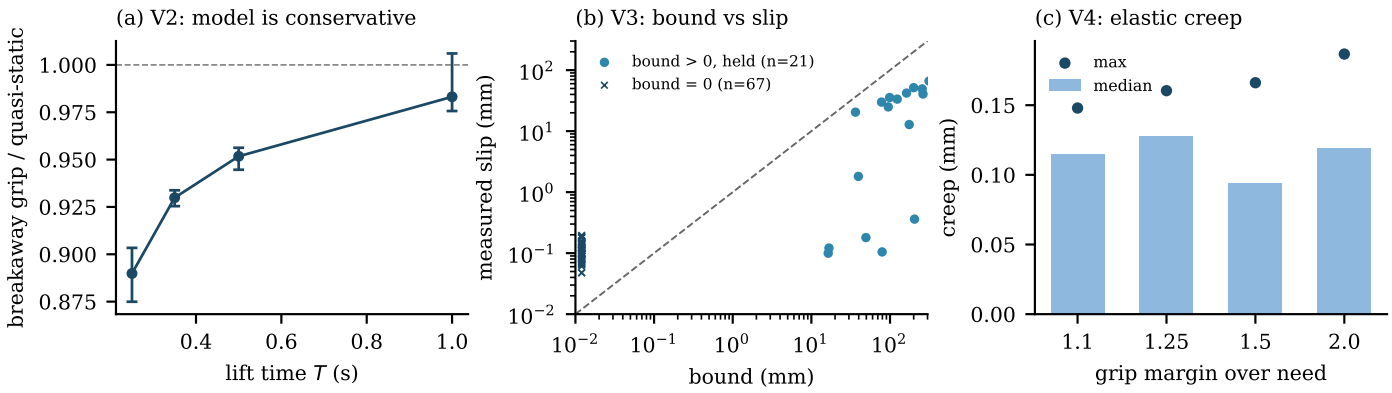


Figure 10. Rig validation. (a) Breakaway grip during a lift relative to the quasi-static prediction (mean and range over 8 objects). (b) Bound against measured slip in 120 random lifts (log axes; zero bounds drawn at the left edge; 32 dropped rods, all with bounds above 0.5 m, omitted). (c) Elastic creep in well-held lifts.

C Controller, Shield and Training

C.1 Probe, bound and shield code

The listings below are copied verbatim from `lift_env.py` by a script, so they match the code that produced every result.

Probe: reactive search for the breakaway grip

```
def probe(rig, m_est, rng=None, sig=None):
    """Search phase with the palm still: step the grip down by SEARCH from the mu = 0.3 grip until the rod moves.
    Each step ramps to the new grip and holds 0.3 s while a slip monitor watches; if slip since the step began
    exceeds MOVE_THR the grip snaps to 2.5x that grip within 20 ms (reactive re-clamp), so a probe never lets the rod run away.
    Returns F_stuck (last grip with no motion; FMAX if the first step moved), F_moved, number of probes."""
    sig = SIG if sig is None else sig
    F = min(FMAX, m_est * G / (2 * 0.3))
    F_stuck = FMAX; n = 0
    while True:
        n += 1; s0 = rig.slip(); F0 = rig.grip; moved = False
        steps = [F0 + (F - F0) * (k + 1) / 100 for k in range(100)] + [F] * 300
        for Fk in steps:
            rig.data.ctrl[:] = Fk; rig.palm_v = 0.0; rig.mj1()
            noise = rng.normal(0, sig) if rng is not None else 0.0
            if rig.slip() - s0 + noise > MOVE_THR:
                moved = True; break
        rig.t_sim += 0.4
        if moved or F < 0.05:
            Fs = min(FMAX, 2.5 * F)                  # snap to 2.5x the grip that moved (ample static margin)
            rig.grip = Fs
            for k in range(20):                      # 20 ms snap
                rig.data.ctrl[:] = Fs; rig.palm_v = 0.0; rig.mj1()
            rig.step(0.3, grip=Fs)
            return F_stuck, F, n
        rig.grip = F; F_stuck = F; F *= SEARCH
```

Proposition 1: the slip bound over the whole commanded profile

```
def lift_bound(F, T, F_stuck, r_max):
    """Proposition 3 for one lift at constant grip F: 0.6 s at rest (grip ramp 0.3 s + settle 0.3 s; the ramp runs
    from a higher grip, so using F there is conservative), the minimum-jerk lift, then 0.3 s at rest.
    Slip cannot start while g + a <= g F / F_stuck; from the first instant it could, integrate a_bar."""
    _, al = minjerk(T)
    a = np.concatenate([np.zeros(600), al, np.zeros(300)])
    cap = G * F / F_stuck                            # lower bound on static capacity (per unit mass)
    over = np.nonzero(G + a > cap)[0]
    if len(over) == 0:
        return 0.0
    abar = (G + a[over[0]:]) - G * F / (r_max * F_stuck)
    v = x = 0.0
    for ab in abar:
        v = max(0.0, v + ab * DT); x += v * DT
    return x
```

Shield projections: grip axis and time axis

```
def shield_grip(F, T, F_stuck, r_max, s_lim):
    """Grip projection: smallest grip >= F whose Proposition 3 bound is <= s_lim (None if FMAX is not enough)."""
    if lift_bound(F, T, F_stuck, r_max) <= s_lim:
        return F
    if lift_bound(FMAX, T, F_stuck, r_max) > s_lim:
        return None
    lo, hi = F, FMAX
    for _ in range(25):
        mid = 0.5 * (lo + hi)
        if lift_bound(mid, T, F_stuck, r_max) <= s_lim: hi = mid
        else: lo = mid
    return hi


def shield_time(F, T, F_stuck, r_max, s_lim):
    """Time projection: shortest duration >= T (<= T_MAX) whose bound is <= s_lim at grip F (None if none)."""
    if lift_bound(F, T, F_stuck, r_max) <= s_lim:
        return T
    if lift_bound(F, T_MAX, F_stuck, r_max) > s_lim:
        return None
    lo, hi = T, T_MAX
    for _ in range(25):
        mid = 0.5 * (lo + hi)
        if lift_bound(F, mid, F_stuck, r_max) <= s_lim: hi = mid
        else: lo = mid
    return hi
```

One episode, all methods

```
def model_grip(F_stuck, T):
    """Model controller: quasi-static Coulomb with r = 1 and F_stuck as the breakaway grip at rest."""
    return min(FMAX, F_stuck * (G + peak_acc(T)) / G)

def features(F_stuck, F_moved, m_est, n_probe):
    w = m_est * G
    return np.array([np.log(F_stuck / w), np.log(F_moved / F_stuck), m_est / 0.2, n_probe / 10.0,
                     np.log(F_stuck / FMAX) / 4.0, 1.0])

def run_episode(obj, rng, mode="model", policy=None, r_max=None, shield_axis="grip"):
    """mode: model | modelS | residual | residualS (shield in the loop). Returns a result dict."""
    r_max = RMAX if r_max is None else r_max
    rig = LiftRig(obj["m"], obj["mu"], ratio=obj.get("ratio", 1.0))
    m_est = obj["m"] * rng.uniform(1 - M_ERR, 1 + M_ERR)     # weight from wrist sensing, +-10%
    F_stuck, F_moved, n_probe = probe(rig, m_est, rng)
    t_probe = rig.t_sim
    T_m = T0; F_m = model_grip(F_stuck, T_m)
    if mode.startswith("residual"):
        a = np.clip(policy(features(F_stuck, F_moved, m_est, n_probe)), -1, 1)
        T_cmd = float(np.clip(T_m * np.exp(ST * a[1]), T_MIN, T_MAX))
        F_cmd = float(np.clip(model_grip(F_stuck, T_cmd) * np.exp(SF * a[0]), 0.1, FMAX))
    else:
        T_cmd, F_cmd = T_m, F_m
    s_lim = S_TOL - MARGIN
    bound0 = lift_bound(F_cmd, T_cmd, F_stuck, r_max)
    projected = declined = False
    if mode.endswith("S"):
        if shield_axis == "grip":
            Fs = shield_grip(F_cmd, T_cmd, F_stuck, r_max, s_lim)
            if Fs is None: declined = True
            else: projected = Fs > F_cmd + 1e-9; F_cmd = Fs
        else:
            Ts = shield_time(F_cmd, T_cmd, F_stuck, r_max, s_lim)
            if Ts is None: declined = True
            else: projected = Ts > T_cmd + 1e-9; T_cmd = Ts
    if declined:                                              # no certified lift: fall back to the model's lift
        T_cmd, F_cmd = T_m, max(F_m, F_cmd)
    bound = lift_bound(F_cmd, T_cmd, F_stuck, r_max)
    rig.set_grip(F_cmd); s0 = rig.slip(); rig.lift(F_cmd, T_cmd)
    slip = rig.slip() - s0
    held = rig.held()
    return dict(T=T_cmd, F=F_cmd, F_rel=F_cmd / (obj["m"] * G), slip=slip, held=held,
                ok=bool(held and slip <= S_TOL), probes=n_probe, t_probe=t_probe, F_stuck=F_stuck,
                bound=bound, bound_before=bound0, projected=projected, declined=declined,
                exceed=bool(slip > bound), n_breakaway=rig.n_breakaway)

def score(res, lam=None):
    lam = LAM if lam is None else lam
```

```
    """Training reward: failure (drop or slip > S_TOL) costs 10; otherwise faster and gentler is better.
    Lift time in units of T0, grip in units of the object's weight m g."""
    if not res["ok"]:
        return -10.0
    return -res["T"] / T0 - lam * res["F_rel"] / 10.0
```

C.2 Training details

Each iteration samples 6 training objects ($r \sim U[1, 1.2]$) and their episode seeds, draws 6 Gaussian perturbations ε_i of the 2×6 weight matrix, and evaluates the policy at $w \pm \sigma\varepsilon_i$ on all 6 objects, plus the current policy. Rewards are rank-transformed to $[-0.5, 0.5]$ and the gradient estimate $\sum_i (\text{rank}_i^+ - \text{rank}_i^-)\varepsilon_i/(6\sigma)$ is followed with Adam [28, 30]. Table 6 lists the settings, which are those of our earlier residual papers [9, 10], unchanged.

Table 6. Training hyperparameters (identical for B2 and B5).

antithetic pairs	6	objects per iteration	6
perturbation scale σ	0.1	Adam step	0.06 ($\beta = 0.9, 0.999$)
iterations	20	episodes per run	1,560
initial weights	0 (model controller)	seeds	0–4
duration scale	$T = T_0 e^{1.0 a_T}$	grip scale	$F = F_{\text{model}}(T) e^{0.25 a_F}$
failure reward	-10	success reward	$-T/T_0 - \lambda F/(10mg)$, $\lambda = 1$

Training loop (`train_es.py`)

```
def main():
    mode, seed = sys.argv[1], int(sys.argv[2])
    os.makedirs("weights", exist_ok=True)
    rng = np.random.default_rng(seed)
    w = np.zeros(2 * OBS_DIM); m = np.zeros_like(w); v = np.zeros_like(w); log = []; t0 = time.time()
    for it in range(ITERS):
        objs = [dict(sample_object(rng), ratio=float(rng.uniform(1.0, R_HI))) for _ in range(N_OBJ)]
        seeds = [int(x) for x in rng.integers(0, 2**31, N_OBJ)]
        eps = rng.normal(size=(POP_PAIRS, w.size))
        fp = np.array([evaluate(mode, w + SIGMA * e, objs, seeds) for e in eps])
        fm = np.array([evaluate(mode, w - SIGMA * e, objs, seeds) for e in eps])
        f0 = evaluate(mode, w, objs, seeds)
        allf = np.concatenate([fp, fm]); ranks = allf.argsort().argsort() / (allf.size - 1) - 0.5
        g = ((ranks[:POP_PAIRS] - ranks[POP_PAIRS:])[:, None] * eps).sum(0) / (POP_PAIRS * SIGMA)
        m = 0.9 * m + 0.1 * g; v = 0.999 * v + 0.001 * g * g
        w = w + LR * (m / (1 - 0.9 ** (it + 1))) / (np.sqrt(v / (1 - 0.999 ** (it + 1))) + 1e-8)
        log.append(dict(it=it, f_center=f0, f_mean=float(allf.mean()), episodes=(it + 1) * (2 * POP_PAIRS + 1) * N_OBJ,
                        wall=time.time() - t0))
        print(mode, seed, it, round(f0, 3), round(float(allf.mean()), 3), round(time.time() - t0), flush=True)
    np.save(f"weights/w_{mode}_s{seed}{TAG}.npy", w)
    json.dump(log, open(f"weights/log_{mode}_s{seed}{TAG}.json", "w"), indent=1)
```

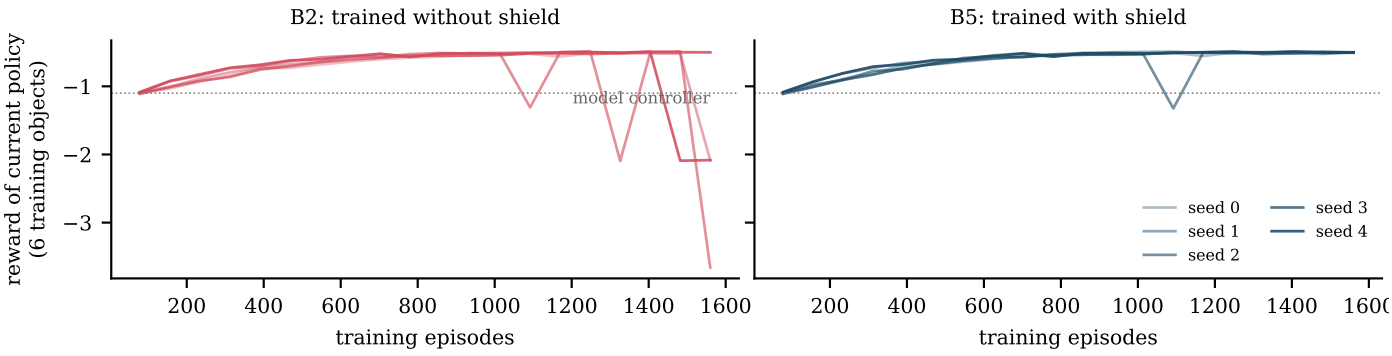


Figure 11. Training curves. Reward of the current (unperturbed) policy on each iteration's 6 training objects. Both methods improve on the model controller (dotted) within about 600 episodes. The unshielded runs dip whenever an iteration's batch contains a lift that slips: three of the five B2 runs ended on such a batch. The shielded runs never fail and converge to the same reward, about -0.50.

D Additional Main Results

Tables 7–11 give the complete test-set and confirmation results behind Section 6. Figure 12 shows every failure of the worst B2 seed.

Table 7. Test set, every method at every r ($r_{\max} = 1.2$). Failure counts with Wilson 95% intervals (%); ΔT and $\Delta F/mg$ against B0 with hierarchical bootstrap 95% intervals (seeds, then objects); "obj. worse" counts objects that failed at least once under the method but never under B0, with the Holm-corrected exact McNemar p (B0S and B4t are outside the pre-registered Holm families and show the raw p). All lift-time and grip differences have Holm-corrected Wilcoxon $p < 10^{-10}$.

Method	failures	Wilson CI (%)	ΔT (s)	$\Delta F/mg$	obj. worse	p	max slip (mm)
$r = 1.0$							
B0	0/40	[0.0, 8.8]	–	–	–	–	0.49
B0S	0/40	[0.0, 8.8]	+0.000	+0.000	0	1	0.49
B2	3/200	[0.5, 4.3]	-0.628[-0.629, -0.627]	+0.267[+0.237, +0.304]	2	1	8.59
B4g	0/200	[0.0, 1.9]	-0.628[-0.629, -0.627]	+0.350[+0.328, +0.372]	0	1	0.12
B4t	0/200	[0.0, 1.9]	-0.581[-0.597, -0.566]	+0.267[+0.237, +0.303]	0	1	0.14
B5	0/200	[0.0, 1.9]	-0.630[-0.630, -0.629]	+0.352[+0.330, +0.374]	0	1	0.11
$r = 1.1$							
B0	0/40	[0.0, 8.8]	–	–	–	–	0.49
B0S	0/40	[0.0, 8.8]	+0.000	+0.000	0	1	0.49
B2	25/200	[8.6, 17.8]	-0.628[-0.629, -0.627]	+0.267[+0.237, +0.304]	14	0.001	20.24
B4g	0/200	[0.0, 1.9]	-0.628[-0.629, -0.627]	+0.350[+0.328, +0.372]	0	1	0.12
B4t	0/200	[0.0, 1.9]	-0.581[-0.596, -0.566]	+0.267[+0.237, +0.303]	0	1	0.14
B5	0/200	[0.0, 1.9]	-0.630[-0.630, -0.629]	+0.352[+0.330, +0.375]	0	1	0.11
$r = 1.2$							
B0	0/40	[0.0, 8.8]	–	–	–	–	0.49
B0S	0/40	[0.0, 8.8]	+0.000	+0.000	0	1	0.49
B2	32/200	[11.6, 21.7]	-0.628[-0.629, -0.627]	+0.267[+0.237, +0.304]	16	0.0004	36.12
B4g	0/200	[0.0, 1.9]	-0.628[-0.629, -0.627]	+0.350[+0.328, +0.372]	0	1	0.12
B4t	0/200	[0.0, 1.9]	-0.581[-0.597, -0.566]	+0.267[+0.237, +0.303]	0	1	0.14
B5	0/200	[0.0, 1.9]	-0.630[-0.630, -0.629]	+0.352[+0.330, +0.374]	0	1	0.11
$r = 1.3$							
B0	0/40	[0.0, 8.8]	–	–	–	–	0.49
B0S	0/40	[0.0, 8.8]	+0.000	+0.000	0	1	0.49
B2	32/200	[11.6, 21.7]	-0.628[-0.629, -0.627]	+0.267[+0.238, +0.304]	16	0.0004	56.02
B4g	0/200	[0.0, 1.9]	-0.628[-0.629, -0.627]	+0.350[+0.328, +0.372]	0	1	0.12
B4t	0/200	[0.0, 1.9]	-0.581[-0.596, -0.566]	+0.267[+0.237, +0.303]	0	1	0.14
B5	0/200	[0.0, 1.9]	-0.630[-0.630, -0.629]	+0.352[+0.330, +0.374]	0	1	0.11

Table 8. Per-seed test results. $\bar{a}_T$, $\bar{a}_F$: effective mean duration and grip actions after any shield projection ($\bar{a}_F > 0$ means a grip above the model's for the chosen lift time). Failures out of 40 lifts at each r.

	seed	T (s)	F/mg	$\bar{a}_T$	$\bar{a}_F$	failures at $r =$ 1.0	1.1	1.2	1.3	max slip (mm)
B2	0	0.373	1.272	-0.99	-0.24	0	1	2	2	25.0
B2	1	0.372	1.276	-0.99	-0.26	1	4	6	6	43.9
B2	2	0.371	1.236	-0.99	-0.38	0	12	14	14	39.7
B2	3	0.371	1.331	-0.99	-0.11	0	2	2	2	24.6
B2	4	0.371	1.265	-0.99	-0.30	2	6	8	8	56.0
B4g	0	0.373	1.355	-0.99	-0.00	0	0	0	0	0.1
B4g	1	0.372	1.356	-0.99	-0.00	0	0	0	0	0.1
B4g	2	0.371	1.359	-0.99	-0.00	0	0	0	0	0.1
B4g	3	0.371	1.362	-0.99	+0.01	0	0	0	0	0.1
B4g	4	0.371	1.360	-0.99	-0.00	0	0	0	0	0.1
B4t	0	0.417	1.272	-0.88	-0.00	0	0	0	0	0.1
B4t	1	0.419	1.276	-0.87	-0.00	0	0	0	0	0.1
B4t	2	0.444	1.236	-0.81	-0.00	0	0	0	0	0.1
B4t	3	0.392	1.331	-0.94	+0.01	0	0	0	0	0.1
B4t	4	0.426	1.265	-0.86	-0.00	0	0	0	0	0.1
B5	0	0.370	1.361	-0.99	-0.00	0	0	0	0	0.1
B5	1	0.371	1.360	-0.99	-0.00	0	0	0	0	0.1
B5	2	0.370	1.361	-0.99	-0.00	0	0	0	0	0.1
B5	3	0.371	1.360	-0.99	-0.00	0	0	0	0	0.1
B5	4	0.370	1.361	-0.99	-0.00	0	0	0	0	0.1

Table 9. Confirmation set (40 fresh objects, $r \sim U[1.0, 1.3]$ per object, run once after the test results were committed). Failures with Wilson 95% intervals (%). B2's failures by r band: (1.0, 1.1]: 1, (1.1, 1.2]: 16, (1.2, 1.3]: 18.

	T (s)	F/mg	failures [Wilson %]	ΔT (s)	$\Delta F/mg$	obj. worse	max slip (mm)
B0	1.000	0.914	0/40 [0.0, 8.8]	–	–	–	0.55
B0S	1.000	0.914	0/40 [0.0, 8.8]	+0.000	+0.000	0 (p=1)	0.55
B2	0.371	1.154	35/200 [12.9, 23.4]	-0.629	+0.240	15 (p=6e-05)	49.69
B4g	0.371	1.232	0/200 [0.0, 1.9]	-0.629	+0.318	0 (p=1)	0.12
B5	0.370	1.234	0/200 [0.0, 1.9]	-0.630	+0.320	0 (p=1)	0.12

Table 10. Mismatched $r_{\max}$. Top: B0S and B5 (trained with $r_{\max} = 1.2$) evaluated with the shield set to $r_{\max} \in \{1.0, 1.2, 1.4\}$; failures out of 40 (B0S) or 200 (B5) lifts per r. Bottom: the certificate over all shielded methods (B0S, B4g, B4t, B5), split by whether the true r is within the assumed bound; "exceed" is the share of lifts whose slip exceeds the bound (dominated by elastic creep, deviation D1).

	$r_{\max}$	T (s)	F/mg	failures at r = 1.0	1.1	1.2	1.3	max slip (mm)
B0S	1.0	1.000	1.009	0	0	0	0	0.49
B0S	1.2	1.000	1.009	0	0	0	0	0.49
B0S	1.4	1.000	1.009	0	0	0	0	0.49
B5	1.0	0.370	1.274	0	10	25	25	21.53
B5	1.2	0.370	1.361	0	0	0	0	0.11
B5	1.4	0.370	1.361	0	0	0	0	0.11

$r_{\max}$	true r	lifts	exceed (%)	max excess (mm)	>5 mm
1.0	r<=	240	16.7	0.49	0
1.0	r>	720	25.7	18.03	60
1.2	r<=	1920	100.0	0.49	0
1.2	r>	640	100.0	0.49	0
1.4	r<=	960	100.0	0.49	0

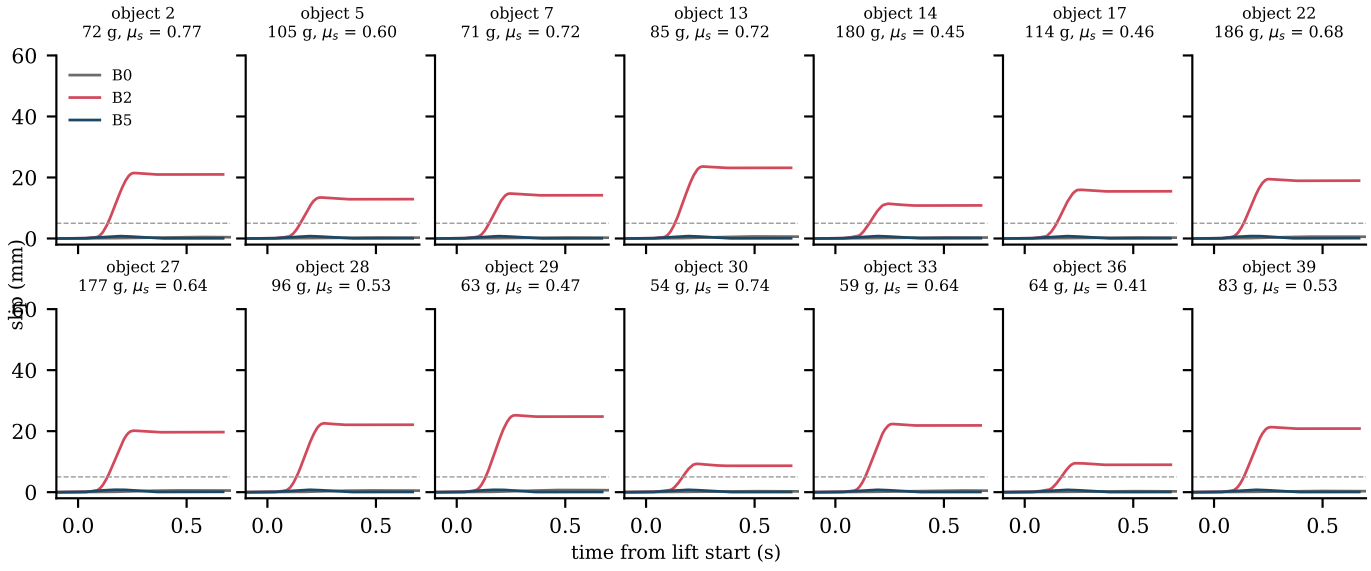


Figure 12. Every failure of the worst unshielded seed (B2 seed 2, $r = 1.2$, 14 of 40 test objects): slip against time for the model controller (B0), the unshielded residual (B2) and the shield-trained residual (B5, seed 2) on the same object and episode seed. All 14 failures share the mechanism of Fig. 5; final slips range from 9 to 25 mm.

Table 11. Per-object test results. m and μ_s are the object's mass and static friction; F_{stuck}/F_b is the probed breakaway grip relative to the true quasi-static breakaway grip $F_b = mg/(2\mu_s)$ (it is ≥ 1 by construction and at most $1/0.9$ up to sensing noise); probes and probe time are from the B0 episode at $r = 1$. Grip columns are F/mg (mean over r and seeds). B2 failures count lifts with slip > 5 mm over 5 seeds $\times$ 4 values of r.

Obj	m (g)	μ_s	F_{stuck}/F_b	probes	probe time (s)	B0 grip	B5 grip	B2 failures
0	69	0.60	1.075	7	3.9	0.95	1.28	0/20
1	140	0.41	1.051	4	2.7	1.35	1.82	0/20
2	72	0.77	0.986	11	5.5	0.68	0.92	11/20
3	61	0.45	1.075	5	3.1	1.26	1.69	0/20
4	192	0.65	1.054	8	4.3	0.86	1.16	0/20
5	105	0.60	1.012	8	4.3	0.89	1.20	3/20
6	149	0.51	1.077	7	3.9	1.12	1.51	0/20
7	71	0.72	1.001	11	5.5	0.74	1.00	5/20
8	151	0.60	1.083	8	4.3	0.95	1.28	0/20
9	173	0.62	1.041	8	4.3	0.89	1.20	0/20
10	197	0.48	1.029	6	3.5	1.13	1.53	3/20
11	133	0.59	1.060	8	4.3	0.95	1.28	0/20
12	103	0.64	1.091	9	4.7	0.91	1.23	0/20
13	85	0.72	0.983	11	5.5	0.72	0.98	12/20
14	180	0.45	1.012	6	3.5	1.19	1.60	6/20
15	120	0.51	1.074	6	3.5	1.11	1.50	0/20
16	62	0.76	1.028	11	5.5	0.72	0.97	0/20
17	114	0.46	1.012	5	3.1	1.17	1.57	3/20
18	151	0.48	1.025	6	3.5	1.13	1.52	0/20
19	185	0.49	1.025	7	3.9	1.11	1.51	3/20
20	55	0.48	1.059	6	3.5	1.17	1.57	0/20
21	102	0.59	1.037	8	4.3	0.93	1.26	0/20
22	186	0.68	0.990	10	5.1	0.77	1.04	11/20
23	101	0.41	1.072	5	3.1	1.40	1.88	0/20
24	74	0.80	1.081	11	5.5	0.72	0.97	0/20
25	119	0.68	1.043	10	5.1	0.82	1.10	0/20
26	58	0.41	1.036	5	3.1	1.33	1.78	0/20
27	177	0.64	0.991	9	4.7	0.83	1.12	10/20
28	96	0.53	0.990	8	4.3	1.00	1.34	7/20
29	63	0.47	0.986	6	3.5	1.11	1.50	6/20
30	54	0.74	1.011	10	5.1	0.73	0.98	2/20
31	120	0.45	1.066	5	3.1	1.25	1.68	0/20
32	161	0.48	1.086	6	3.5	1.20	1.62	0/20
33	59	0.64	0.985	9	4.7	0.82	1.10	3/20
34	184	0.41	1.087	4	2.7	1.40	1.89	0/20
35	171	0.48	1.057	6	3.5	1.18	1.59	0/20
36	64	0.41	1.019	5	3.1	1.33	1.78	2/20
37	94	0.69	1.076	10	5.1	0.82	1.11	0/20
38	124	0.74	1.043	11	5.5	0.74	1.01	0/20
39	83	0.53	0.996	7	3.9	1.00	1.35	5/20

E Extended Study Details

Design. Table 12 lists every controlled change. Rows E1–E9 and S1–S6 reuse the reference weights; T1–T4 retrain B2 and B5 with five seeds each under the changed setting and evaluate them in the reference environment (T2 also uses the wider duration scale at test). Every row uses the 40 test objects (for E1 and E2, the same seed with the changed ranges), the same four values of r and the same episode seeds as the reference. Significance in Fig. 6 uses a bootstrap over objects of the per-object means (10,000 resamples), paired for every row except E1 and E2. Table 13 gives absolute values; Fig. 14 adds the change in lift time.

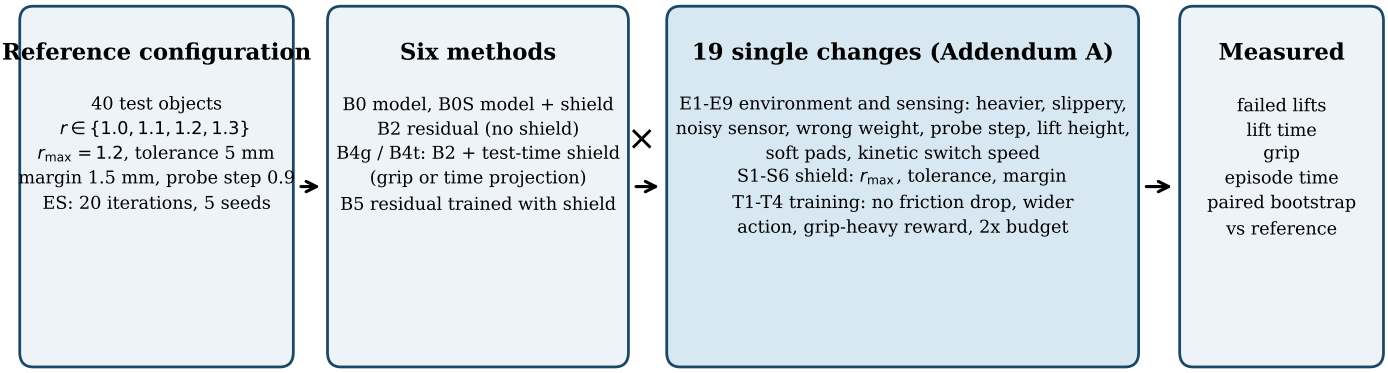


Figure 13. Controlled study overview. Each row of Fig. 6 changes a single design choice relative to the reference configuration and reruns every method on the same objects and episode seeds. Rows T1–T4 retrain B2 and B5 (5 seeds each); all other rows reuse the trained weights.

Table 12. Controlled changes. Each row modifies a single design choice relative to the reference configuration (test set: 40 objects, $r \in \{1.0, 1.1, 1.2, 1.3\}$, $r_{max} = 1.2$, tolerance 5 mm, margin 1.5 mm, probe step 0.9, slip-monitor noise 0.5 mm, weight error $\pm 10\%$, lift height 10 cm, pad stiffness solref 0.003, kinetic switch at 2 cm/s, training on $r \in [1, 1.2]$ for 20 iterations with $\lambda = 1$). Rows R0L, E7L, E9L and T2L are diagnostics added after E7, E9 and T2 were seen and frozen before they were run (Addendum A6).

Row	Change	What is modified	Retrained
Environment and sensing			
E1	heavier objects	$m \sim U[0.2, 0.5]$ kg (reference $U[0.05, 0.2]$)	no
E2	slippery objects	$\mu_s \sim U[0.2, 0.4]$ (reference $U[0.4, 0.8]$)	no
E3	noisy slip monitor	noise 1.5 mm (reference 0.5 mm)	no
E4	poor weight estimate	error $\pm 25\%$ (reference $\pm 10\%$)	no
E5	coarser probe	grip step $\times 0.8$ (reference $\times 0.9$)	no
E6	finer probe	grip step $\times 0.95$	no
E7	higher lift	$H = 20$ cm (reference 10 cm)	no
E8	softer pads	contact time constant 6 ms (reference 3 ms)	no
E9	earlier kinetic switch	switch at 5 mm/s (reference 2 cm/s)	no
E7L, E9L	diagnostics	E7, E9 with the one-step lag removed from the switch	no
R0L	diagnostic	reference with the lag removed	no
T2L	diagnostic	T2 with the lag removed	(T2 weights)
Shield settings			
S1, S2, S3	assumed drop	$r_{max} = 1.0$, 1.1, 1.4 (reference 1.2)	no
S4, S5	slip tolerance	2 mm, 10 mm (reference 5 mm)	no
S6	no margin	sensing margin 0 (reference 1.5 mm)	no
Training			
T1	no friction drop in training	train on $r = 1$ only	B2, B5
T2	wider duration action	$T = T_0 e^{1.5 a_T}$ (reaches 0.25 s)	B2, B5
T3	grip-heavy reward	$\lambda = 10$	B2, B5
T4	longer training	40 iterations (3,120 episodes)	B2, B5

Probe failures (E2, E3). With $\mu_s \in [0.2, 0.4]$ the probe's starting grip, chosen for $\mu = 0.3$, is already below the breakaway grip for many objects; with 1.5 mm monitor noise, the noise alone crosses the 2.5 mm threshold within the first 300 ms hold. In both cases the probe returns its default $F_{stuck} = 30$ N. The model controller then grips at 30 N, the residual's features fall far outside the training range, and the shield, whose bound is computed from F_{stuck}, finds no certified fast lift and declines. A probe that starts from a higher grip and treats a first-step motion as "unknown" rather than as a measurement would avoid this; we did not change it.

Change in lift time (ms)

	B0	B0S	B2	B4g	B4t	B5
Reference	1000	1000	372	372	420	370
E1 heavier objects 0.2-0.5 kg	0 (1000)	0 (1000)	**-2*** (369)	**-2*** (369)	**+12*** (431)	**-1*** (369)
E2 slippery objects mu 0.2-0.4	0 (1000)	0 (1000)	**+56*** (428)	**+429*** (801)	**+395*** (815)	**+430*** (800)
E3 slip-monitor noise 1.5 mm	0 (1000)	0 (1000)	**+82*** (454)	**+628*** (1000)	**+580*** (1000)	**+630*** (1000)
E4 weight error +-25%	0 (1000)	0 (1000)	+0.18 (372)	+0.18 (372)	+0.40 (420)	+0.15 (371)
E5 coarser probe step 0.8	0 (1000)	0 (1000)	**+2*** (373)	**+2*** (373)	**-5*** (415)	**+1*** (372)
E6 finer probe step 0.95	0 (1000)	0 (1000)	**-2*** (370)	**-2*** (370)	**+7*** (427)	**-1*** (369)
E7 lift height 20 cm	0 (1000)	0 (1000)	0 (372)	0 (372)	**-19*** (400)	0 (370)
E8 softer pads solref 0.006	0 (1000)	0 (1000)	**-0.23*** (372)	**-0.23*** (372)	**+0.31*** (420)	**-0.18*** (370)
E9 earlier kinetic switch 5 mm/s	0 (1000)	0 (1000)	+0.01 (372)	+0.01 (372)	-0.03 (420)	+0.00 (370)
S1 shield r_max = 1.0	0 (1000)	0 (1000)	0 (372)	0 (372)	**-39*** (381)	0 (370)
S2 shield r_max = 1.1	0 (1000)	0 (1000)	0 (372)	0 (372)	**-0.04*** (420)	0 (370)
S3 shield r_max = 1.4	0 (1000)	0 (1000)	0 (372)	0 (372)	0 (420)	0 (370)
S4 slip tolerance 2 mm	0 (1000)	0 (1000)	0 (372)	0 (372)	0 (420)	0 (370)
S5 slip tolerance 10 mm	0 (1000)	0 (1000)	0 (372)	0 (372)	0 (420)	0 (370)
S6 no sensing margin	0 (1000)	0 (1000)	0 (372)	0 (372)	0 (420)	0 (370)
T1 trained on r = 1.0 only	0 (1000)	0 (1000)	**-0.20*** (372)	**-0.20*** (372)	**+17*** (437)	0 (370)
T2 wider time action (scale 1.5)	0 (1000)	0 (1000)	**-122*** (250)	**-122*** (250)	**-153*** (267)	**-120*** (250)
T3 grip weight lambda = 10	0 (1000)	0 (1000)	**+28*** (400)	**+28*** (400)	**+14*** (434)	**+127*** (497)
T4 double training budget	0 (1000)	0 (1000)	**-3*** (369)	**-3*** (369)	**-25*** (395)	**-3*** (368)
R0L reference, lag-free switch	0 (1000)	0 (1000)	0 (372)	0 (372)	0 (420)	0 (370)
T2L wider time action, lag-free switch	0 (1000)	0 (1000)	**-122*** (250)	**-122*** (250)	**-153*** (267)	**-120*** (250)
E7L lift height 20 cm, lag-free switch	0 (1000)	0 (1000)	0 (372)	0 (372)	**-19*** (400)	0 (370)
E9L switch 5 mm/s, lag-free switch	0 (1000)	0 (1000)	+0.01 (372)	+0.01 (372)	-0.03 (420)	+0.00 (370)

Figure 14. Controlled study: change in lift time (ms), same layout as Fig. 6. Only the shield rows that decline (E2, E3: the model's 1 s lift is used), the time projection and the retrained wider-action policies (T2) move the lift time.

Table 13. Controlled study, absolute values. For every row and method: failed lifts (slip above the row's tolerance, or a drop) out of all lifts (40 objects × 4 values of r, × 5 seeds for learned methods), mean lift time T (ms) and mean grip F/mg (%). Row descriptions are in Table 12.

	B0			B0S			B2			B4g			B4t			B5		
	fail	T	F	fail	T	F	fail	T	F	fail	T	F	fail	T	F	fail	T	F
REF	0/160	1000	101	0/160	1000	101	92/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
E1	0/160	1000	102	0/160	1000	102	208/800	369	129	0/800	369	139	0/800	431	129	0/800	369	138
E2	0/160	1000	2109	0/160	1000	2109	27/800	428	2047	0/800	801	2127	0/800	815	2123	0/800	800	2126
E3	0/160	1000	3108	0/160	1000	3108	0/800	454	2994	0/800	1000	3108	0/800	1000	3108	0/800	1000	3108
E4	0/160	1000	102	0/160	1000	102	71/800	372	129	0/800	372	137	0/800	420	129	0/800	371	137
E5	0/160	1000	106	0/160	1000	106	53/800	373	135	0/800	373	142	0/800	415	135	0/800	372	142
E6	0/160	1000	98	0/160	1000	98	240/800	370	123	0/800	370	132	0/800	427	123	0/800	369	133
E7	0/160	1000	106	0/160	1000	106	463/800	372	166	302/800	372	176	33/800	400	166	309/800	370	177
E8	0/160	1000	98	0/160	1000	98	120/800	372	123	0/800	372	132	0/800	420	123	0/800	370	132
E9	0/160	1000	101	0/160	1000	101	531/800	372	128	375/800	372	136	431/800	420	128	375/800	370	136
S1	0/160	1000	101	0/160	1000	101	92/800	372	128	40/800	372	129	40/800	381	128	60/800	370	127
S2	0/160	1000	101	0/160	1000	101	92/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
S3	0/160	1000	101	0/160	1000	101	92/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
S4	0/160	1000	101	0/160	1000	101	115/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
S5	0/160	1000	101	0/160	1000	101	68/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
S6	0/160	1000	101	0/160	1000	101	92/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
R0L	0/160	1000	101	0/160	1000	101	92/800	372	128	0/800	372	136	0/800	420	128	0/800	370	136
T2L	0/160	1000	101	0/160	1000	101	19/800	250	175	0/800	250	185	0/800	267	175	0/800	250	185
E7L	0/160	1000	106	0/160	1000	106	69/800	372	166	0/800	372	176	0/800	400	166	0/800	370	177
E9L	0/160	1000	101	0/160	1000	101	527/800	372	128	364/800	372	136	187/800	420	128	366/800	370	136
T1	0/160	1000	101	0/160	1000	101	195/800	372	125	0/800	372	136	0/800	437	125	0/800	370	136
T2	0/160	1000	101	0/160	1000	101	338/800	250	175	190/800	250	185	110/800	267	175	190/800	250	185
T3	0/160	1000	101	0/160	1000	101	39/800	400	125	0/800	400	130	0/800	434	125	0/800	497	118
T4	0/160	1000	101	0/160	1000	101	4/800	369	131	0/800	369	137	0/800	395	131	0/800	368	137

Bound-tightness sweep. 2,160 lifts: grips of 85, 88, 91, 94, 97 and 100% of the no-onset grip, lift times 0.37, 0.5 and 1 s, $r \in \{1.0, 1.1, 1.2\}$, 40 objects, bound evaluated with $r_{\max} = r$. Of these, 491 dropped the rod (all with bounds above 12 cm), 360 had a zero bound (largest slip 0.49 mm) and 1,309 had a non-zero bound and kept the rod (median slip/bound 0.023; 7 exceedances, Section 7.5).

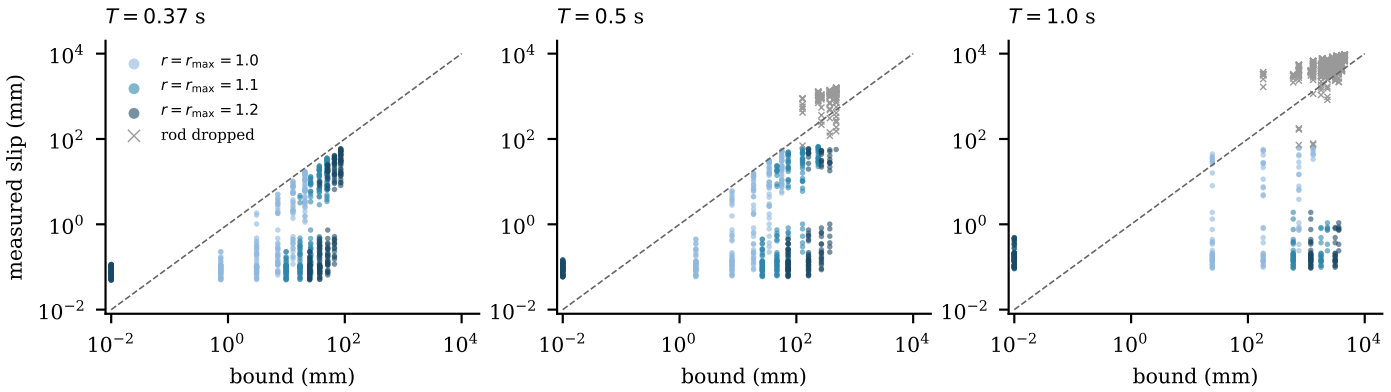


Figure 15. Bound against measured slip for deliberately under-gripped lifts (grip 85–100% of the no-onset grip, 40 objects; log axes, zeros drawn at the left and bottom edges; dashed line: slip equal to the bound). Points below the line are covered by the bound. Crosses are lifts that dropped the rod, where the slip is the rod's fall.

Safety map. 10,080 lifts: 63 fixed actions × 40 objects × $r \in \{1.0, 1.2\}$ × with and without the grip shield (Fig. 8).

F Protocol, Deviations and Reproducibility

Timeline. The repository history records the order of work: (1) rig, validation and the frozen protocol; (2) training, test-set evaluation and the analysis code, committed before the confirmation run; (3) the confirmation run, deviations D1–D3, the exploratory slip-allowance analysis and the first paper draft; (4) Addendum A, frozen before any of the extended study was run; (5) the extended study and this paper.

Pre-registered questions and outcomes (excerpt of `PROTOCOL.md`, frozen before training)

Q1 (speed and grip). Does a residual lift faster than the model, and how much of that survives the shield? Outcomes: lift time T, relative grip F/mg, success. Test B2, B4g, B5 vs B0 per r.

Table 14. Final training reward per seed (reward of the last unperturbed policy on that iteration's 6 training objects; failures cost −10, so values below −1 mean the batch contained a failed lift). Rewards under T3 are on a different scale ($\lambda = 10$).

Training	method	seed 0	seed 1	seed 2	seed 3	seed 4
reference	B2	-0.50	-2.08	-3.66	-0.50	-2.08
reference	B5	-0.51	-0.50	-0.50	-0.50	-0.51
T1 no drop	B2	-0.50	-0.50	-3.66	-0.50	-0.50
T1 no drop	B5	-0.51	-0.50	-0.50	-0.50	-0.51
T2 wider action	B2	-0.43	-3.61	-0.42	-6.81	-2.02
T2 wider action	B5	-0.43	-2.01	-0.43	-0.43	-0.43
T3 $\lambda = 10$	B2	-1.65	-3.05	-1.59	-3.02	-1.66
T3 $\lambda = 10$	B5	-1.67	-1.61	-1.63	-1.61	-1.67
T4 40 iterations	B2	-0.50	-0.48	-0.49	-0.48	-0.49
T4 40 iterations	B5	-0.50	-0.48	-0.49	-0.48	-0.49

Q2 (safety vs r). Failure rate (drop or slip > 5 mm) per method and r; is the unshielded residual's failure rate higher at larger r?

Q3 (which axis). Grip projection (B4g) vs time projection (B4t) on the same commands: T, F/mg, success, fraction projected and declined.

Q4 (certificate). For every shielded episode: fraction with slip $>$ bound, max exceedance, and number with slip $> S_{\text{tol}}$ among certified lifts, at $r \le r_{\max}$ and at $r > r_{\max}$; effect of $r_{\max} = 1.0/1.2/1.4$.

Statistics. Learned methods summarised per object as the mean over 5 seeds; paired Wilcoxon signed-rank tests vs B0 over the 40 objects for T and F/mg at each r; Holm correction within each question family; failure rates with Wilson 95% intervals; exact McNemar tests on per-object any-failure; hierarchical bootstrap 95% CIs (seeds, then objects; 10,000 draws). Only measured results are reported.

Deviations. All are dated in the repository.

- **D1.** The pre-registered exceedance metric (slip $>$ bound) counts elastic creep: every certified lift had a bound of exactly zero and 0.05–0.49 mm of creep. We report it as registered and add the size of the excess and the count of lifts over the tolerance.
- **D2.** The confirmation set draws r per object from $U[1.0, 1.3]$, so per-r tests do not apply; it is analysed pooled over objects with the same tests, plus failures by r band.
- **D3.** The slip-allowance analysis (Fig. 9, Table 4) was not pre-registered.
- **Addendum A.** The extended study (Section 7, Appendix E) was designed after the confirmatory results and frozen before it was run. Before running it we refactored module constants used as default arguments so they are read at call time; re-running 40 randomly chosen test episodes reproduced the committed results exactly.

Software and compute. Python 3.13, MuJoCo 3.15.0, NumPy 2.5, SciPy 1.18, pandas 3.0, statsmodels 0.15, Matplotlib 3.11. All experiments ran on two cores of a cloud CPU (Intel Xeon, 2.1 GHz); no GPU was used. Table 15 lists the episodes and wall-clock time.

Table 15. Compute. Episodes simulated per experiment and single-core CPU hours at the measured 0.075 s per episode (each episode simulates about 6 s of contact-rich dynamics at 1 ms). Everything ran on two CPU cores; no GPU.

Experiment	episodes	CPU hours
Rig validation (V1–V4)	–	< 0.1
Main training: B2, B5 × 5 seeds	15,600	0.3
Test set	5,440	0.1
Confirmation set	680	< 0.1
Training rows T1–T3 (30 runs)	46,800	1.0
Training row T4 (10 runs, 40 iterations)	31,200	0.7
Controlled rows E1–E9, S1–S6	52,800	1.1
Controlled rows T1–T4	14,080	0.3
Diagnostic rows R0L, E7L, E9L, T2L	14,080	0.3
Safety map	10,080	0.2
Bound tightness	2,160	< 0.1
Total	192,920	4.0

Reproduce every number

```
python validate.py                                                    # rig validation (Table 1)
./train_all.sh                                                          # B2, B5:  5 seeds each
python evaluate.py test && python analyze.py test
python evaluate.py confirm && python analyze_confirm.py
python explore_allowance.py                                                                  # D3
./train_extra.sh && python controlled.py ALL_ENV && python controlled.py ALL_TRAIN
python extended.py tightness && python extended.py safetymap && python extended.py traces
python controlled_analysis.py && python figs_ext.py ...  && python make_tables.py ...
```