Highlights

A Biomimetic Framework for Dynamic Simulation and Heuristic Grasp Planning of Soft Objects in the Presence of External Force

Navid Noorgostar

- Near-real-time 3D mass–spring–damper soft-object model integrated by RK4.
- Five-component closed-form heuristic grasp metric with sharp-feature penalties.
- Hunt–Crossley + Coulomb stick–slip contact model identified via PSO.
- Physically validated: 2.7% contact-calibration RMSE and a 1755-trial grasp study.
- Predicts the external-force limit a soft grasp resists; a 1230-trial study.

A Biomimetic Framework for Dynamic Simulation and Heuristic Grasp Planning of Soft Objects in the Presence of External Force

Navid Noorgostar[a,*]

[a] *Department of Mechanical Engineering, Faculty of Engineering, Central Tehran Branch, Islamic Azad University, Tehran, Iran*

Abstract

Robotic manipulation of soft, deformable objects is harder than rigid-body grasping: the object's many degrees of freedom make real-time contact prediction expensive, and grasp stability depends on the coupled dynamics of gripper, surface, and external load. This paper presents a unified, biomimetic framework with three components. A three-dimensional mass–spring–damper (MSD) model of the soft object is integrated in state-space form with a fourth-order Runge–Kutta scheme for near-real-time simulation. A five-component heuristic scores candidate grasps from force balance, contact-triangle regularity, centroid proximity, and penalties for sharp vertices and sharp-edge proximity. Normal and tangential contact follow a Hunt–Crossley plus Coulomb stick–slip model whose six parameters are identified by Particle Swarm Optimization (PSO). The framework is validated physically in three stages. A bench-test calibration on four soft materials (19,168 fitted events) confirms the contact exponent ($n = 0.46 \pm 0.07$, force RMSE 2.7 %). A 1755-trial campaign on a physical three-finger gripper succeeds on 86.5 % of the proposed planner's validation-set grasps (77.8 % over all 819 of its trials, including an adversarial set) and predicts the trial-mean fingertip normal force without systematic bias (slope 0.98, $R^2 = 0.72$). Finally, a 1230-trial disturbance campaign addresses manipulation under external force: the calibrated friction budget (closed-form statics on the planned contacts) predicts the load a grasp resists before slipping. The planner holds about 86 % of disturbances within its natural-friction operating envelope, and 80.8 %

*Corresponding author.
Email address: navid.noorgostar22@gmail.com (Navid Noorgostar)

Branch, Tehran, Iran.

over all 830 of its trials (1230 across planners), and predicts the lateral slip threshold to a median error of 10.7 %—versus 74 % for a geometric baseline—with a conservative bias ($R^2 = 0.56$, dominated by a minority of large-error trials).

Keywords: Biomimetic grasping, soft-object manipulation, mass–spring–damper model, Hunt–Crossley contact, particle swarm optimization, external-force disturbance

---

1. Introduction

Robotics has matured rapidly across industrial, medical, military, and service applications, and considerable effort has been devoted in the last three decades to extending robotic manipulation from rigid bodies to deformable ones [1, 2, 3]. Manipulating a soft object—a cloth, a piece of food, a biological tissue, a length of thread, dough, or a fruit—is appreciably more challenging than grasping a rigid one. The object deforms continuously as the contact region evolves, the normal and tangential contact forces vary in magnitude and direction throughout the closing motion, and friction has a decisive role: when the gripper closes, slip must be prevented at every instant in time. These are dynamic phenomena that have to be solved in roughly real time if the resulting controller is to be useful on a physical robot. Application areas in which manipulating soft objects directly drives the economics—fruit picking, surgical suturing, butchery, food handling, and industrial assembly of compliant components—make the problem all the more pressing.

Biological inspiration. The benchmark for soft-object manipulation remains biology. The human hand picks a ripe peach, a slice of cake, or a fragile organ during surgery without crushing or dropping it, and does so with three coordinated mechanisms that this work mimics. First, the soft, compliant fingertips of biological hands deform around the object, generating a distributed contact patch rather than a single Hertzian point contact [4, 5]—this is reproduced here by a non-linear Hunt–Crossley contact model on a deforming triangulated surface. Second, the fingers position themselves at near-isostatic configurations around the object's centroid so that the resulting contact wrench passes close to its centre of mass—we encode this geometric intuition in a closed-form heuristic that scores candidate three-finger configurations on angular regularity, triangle compactness, and centroid alignment. Third, biological grasping adapts grip to an object's compliance, mass, and surface friction through online sensorimotor tuning [6]; loosely analogous to this adaptation, we identify the six unknown contact parameters with a population-based intelligent search (Particle Swarm Optimization). The combination—compliant contact, geometric pre-screening, and intelligent parameter adaptation—constitutes a biomimetic alternative to fully model-predictive or learning-only grasp pipelines, and is the contribution of this paper.

Research on soft-object simulation pre-dates research on soft-object grasping. Early work was dominated by methods developed for computer graphics rather than for robotic interaction. Barr [7] introduced global and local deformation operators on solid primitives; Sederberg and Parry [8] generalised this idea as Free-Form Deformation (FFD), which embeds the object inside a lattice of control points and warps the embedding space rather than the object itself. Kass, Witkin, and Terzopoulos [9] introduced the active-contour ("snake") family. Terzopoulos and co-workers [10] launched physics-based deformable models proper, anchored in continuum mechanics, and the field has since branched into many sub-approaches. Surveys by Gibson and Mirtich [11], Nealen et al. [12], and Meier et al. [13] summarise the landscape. From the standpoint of a robotic system that has to evaluate many candidate grasps in a closed loop, two requirements compete: physical fidelity in the contact region, and computational cost over the rest of the body.

The two dominant physically-based families are the finite-element method (FEM) and the mass–spring–damper (MSD) particle-system family. FEM offers higher fidelity for continuum phenomena and naturally supports nonlinear elasticity, anisotropy and complex constitutive laws [14, 15], but its computational cost has historically limited it in interactive settings unless aggressive reduced-order schemes are used [16, 17]. MSD replaces the continuum with a discrete network of point masses connected by linear spring–damper elements; it was introduced in the graphics community for facial animation, cloth [18, 19] and surgical simulation [20], and its principal attractions are simplicity of formulation, explicit time-stepping, and a near-linear scaling in the number of nodes. The downsides are well documented: MSD parameters do not map directly to material properties and have to be calibrated, and a non-equilibrium discretisation can give rise to non-physical artefacts. Position-based dynamics [21, 22] occupy a third niche oriented at speed rather than physical fidelity.

On the grasping side, the question is not how to simulate the object in isolation but how to evaluate the quality and stability of a candidate grasp on a deformable body. Hirai, Tsuboi, and Wada [23] modelled grasping of deformable objects using vision and tactile feedback; Inoue and Hirai [24] introduced a parallel-distributed three-dimensional soft-finger model; Lin et al. [25] proposed picking up a soft body by "feeling" the grip with two rigid fingers using FEM for the deformable component. Ciocarlie et al. [5] adapted the contact geometry to grasping and manipulation tasks. Jia, Guo, and Lin [26] derived squeeze and stick/slip energy-based optimalities for grasping deformable planar objects. Zaidi et al. [27] developed a model-based strategy for grasping three-dimensional deformable objects with a multi-fingered hand. Navarro-Alarcon and co-workers [28] addressed automatic 3D manipulation with an adaptive deformation model. A complementary and increasingly popular line of work instead learns grasp quality directly from data or human demonstrations [29, 30]; such methods reach strong empirical performance but typically require extensive training data or teleoperated demonstrations, remain difficult to interpret physically, and still find deformable, fragile objects under external load a recognised open challenge. On the simulation side, GPU-accelerated finite-element benchmarks such as DefGraspSim [31] now provide large-scale grasp-outcome datasets for three-dimensional deformable bodies, and recent surveys map the modelling, perception, and control landscape [32]; on the metric side, wrench-space quality measures such as the Ferrari–Canny ϵ-metric [33] remain the classical rigid-body reference against which geometric heuristics should be read. We do not attempt a head-to-head comparison with the FEM pipelines—their corotational FEM and our calibrated MSD target different accuracy–speed operating points—but the ablation and baseline studies of Section 6 position the heuristic within that space, and porting the Test-2/3 protocols to a DefGraspSim-style benchmark object set is a natural next step. The framework developed here is deliberately physics-grounded and training-free: every heuristic term carries a mechanical meaning, calibration needs only a short bench test rather than thousands of labelled grasps, and the predicted slip margin is a measurable physical quantity that we validate directly against force measurements. Despite this rich literature, two practical points recur: (i) sphere–triangle intersection and penetration tests are textbook material [34], and we spell them out end-to-end for the deforming triangulated surface—together with the virtual-work-consistent force transfer—for completeness and reproducibility rather than as a claimed contribution, and (ii) the contact parameters that govern Hunt–Crossley [35] normal-force laws and the Coulomb friction transition between stick and slip are typically hand-tuned per object.

This paper consolidates these threads into a single self-contained biomimetic framework. We use a three-dimensional MSD network for the soft body, written in state-space form with an adjacency-matrix assembly and integrated by RK4. We define a five-component heuristic grasp metric that scores candidate three-finger configurations on force balance, contact-triangle regularity, centroid proximity, and two sharp-feature penalty terms (sharp vertices and sharp-edge-adjacent contacts)—and validate which components actually matter in an ablation study. We derive a four-case collision detector and a three-case penetration-depth computation based on local-frame transformation of each triangular face. Normal contact forces are computed with a Hunt–Crossley-type power-law visco-elastic model [35] (with a linear-rate damping variant, Section 3.5), tangential forces with a Coulomb stick–slip law [36, 37, 38], and six unknown contact parameters are identified by minimising a composite cost function using Particle Swarm Optimization. The pipeline is evaluated on five object shapes (sphere, ellipsoid, rounded cube, pear, concave handle), benchmarked against baseline grasp planners (random, centroid–equilateral, antipodal, force-balance, and a no-PSO ablation; the simulation and hardware campaigns each use the appropriate subset), and analysed for stability under friction–mass sweeps, mesh resolution, PSO repeatability, and parameter identifiability. These comparative simulation studies are run with a quasi-static reduction of the MSD model that returns the same equilibrium-residual and slip-margin metrics; the full RK4 dynamic integrator is exercised on the spherical object of Section 5, where the two evaluators apply the same success criteria and agree on the labels (Section 6 details the scope of that comparison). As a preliminary study, a neural surrogate is trained on a synthetic grasp dataset to explore fast parameter pre-screening. The full pipeline is implemented on top of the SynGrasp toolbox [39] for the three-fingered hand kinematics.

Contribution map. We position the work in prose rather than a feature checklist. FEM pipelines [25] offer continuum fidelity for three-dimensional deformable bodies but not real-time grasp evaluation; MSD frameworks [20] and soft-finger contact models [5, 24] provide fast compliant contact but no grasp-pose heuristic; classical grasp planners [40, 27] supply pose selection on rigid or quasi-rigid bodies. Individually, power-law compliant contact, population-based parameter identification, experimental contact calibration, and hardware validation each appear elsewhere in the broader literature; what we claim as novel is their *integration*—a calibrated compliant-contact MSD body, a closed-form biomimetic pose heuristic, and PSO identification, carried through three linked physical campaigns (bench calibration, grasp validation, disturbance robustness) with baseline-planner comparisons at each stage. Table 1 previews the headline success rates established below, each labelled with its testbed, trial count, and baseline set so that the different figures are not conflated.

Figure 1 summarises the framework and its three-stage physical validation. The paper is organised as follows. Section 2 derives the MSD model; Section 3 develops the grasp metric, collision detector, penetration computation and contact-force model; and Section 4 describes the PSO parameter identification. The simulation studies are then grouped together: Section 5 reports the original spherical-object results across a friction–mass sweep, and Section 6

Table 1. Headline success rates reported in this paper, grouped by testbed. The rates are *not* directly comparable across rows: each uses a different object set, baseline set, and trial count. Here n is the number of proposed-planner trials behind the rate; "pp" denotes percentage points; and the last column points to the section where the number is reported.

Study (proposed planner)	Testbed (n)	Success	vs. strongest baseline	Section
Heuristic ablation	simulation, 5 shapes × 10 seeds (50/variant)	72.0%	random 6%, centroid 0%	Sec. 6.3
Baseline grid	simulation, 5 × 3 × 3 × 15 (675)	67.6%	antipodal 40.0%	Sec. 6.5
Grasp validation, nominal	hardware, validation groups A/B/D (465)	86.5%	—	Sec. 8.2
Grasp validation, all	hardware, 1755 trials (819)	77.8%	no-PSO 46.7%	Sec. 8.5
Disturbance, natural μ	hardware, natural-friction envelope (740)	∼86%	—	Sec. 9.2
Disturbance, all	hardware, 1230 trials (830)	80.8%	no-PSO 46.0%	Sec. 9.5

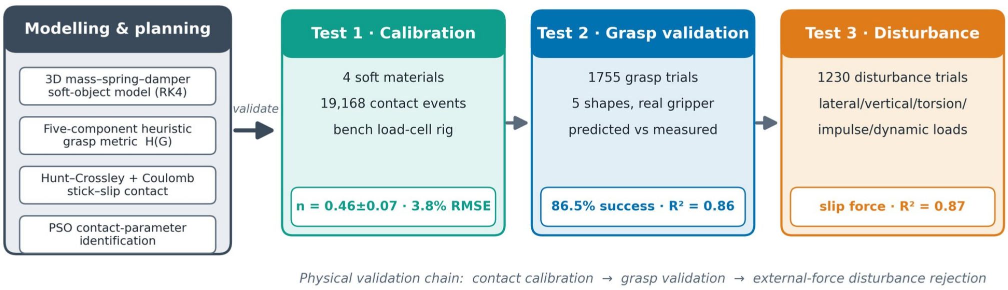

Figure 1. Overview of the framework and its three-stage physical validation. A three-dimensional mass–spring–damper soft-object model, a five-component heuristic grasp metric, a Hunt–Crossley plus Coulomb stick–slip contact model, and PSO-based parameter identification are validated experimentally in three linked stages: contact calibration on four soft materials (Test 1, Section 7), end-to-end grasp validation on a physical gripper (Test 2, Section 8), and external-force disturbance rejection (Test 3, Section 9).

extends the validation to non-spherical shapes, ablates the heuristic, compares against baseline planners, and analyses robustness, mesh resolution, and PSO repeatability. The three physical campaigns follow: Section 7 reports a bench-test calibration of the contact model on four soft materials (Test 1); Section 8 validates the complete pipeline on a physical three-finger gripper over 1755 grasp trials (Test 2); and Section 9 tests robustness to external-force disturbances over 1230 trials and validates the predicted failure boundary (Test 3). Section 10 discusses failure modes, Section 11 documents the reproducibility package, and Section 12 concludes. Appendix A reports a preliminary ML surrogate for fast parameter pre-screening.

2. Soft Object Modelling

2.1. Mass–Spring–Damper formulation

We treat the soft object as a discrete dynamical system: its surface is triangulated into n point-mass nodes connected by spring–damper edges (Fig. 2). Each node i obeys Newton's second law,

$$m_i \ddot{x}_i + f_i^{\text{nodal}} + \sum_{j \in N(i)} f_{ij}^{\text{int-nodal}} = f_i^{\text{ext}}, \qquad i \in \{1, \ldots, n\}, \tag{1}$$

where m_i and x_i are the mass and position of node i, $N(i)$ is the set of nodes connected to i, f_i^{nodal} is the damping force at node i, $f_{ij}^{\text{int-nodal}}$ is the internal spring–damper force between i and j, and f_i^{ext} collects all external forces (gravity, contact forces transferred from the gripper, etc.). Following Mollemans et al. [20] and the discussion in Nealen et al. [12], we adopt a linear nodal damping term,

$$f_i^{\text{nodal}} = b_i \dot{x}_i, \tag{2}$$

where b_i is the nodal damping coefficient. The internal force on edge (i, j) is the sum of a linear-spring component f_{ij}^s and a parallel viscous-damper component f_{ij}^d:

$$f_{ij}^{\text{int-nodal}} = f_{ij}^s + f_{ij}^d, \qquad f_{ij}^s = k\, \Delta l_{ij}\, \hat{x}_{ij}, \qquad f_{ij}^d = c\, \dot{l}_{ij}\, \hat{x}_{ij}, \tag{3}$$

with k the spring stiffness, c the edge damping coefficient, l_{ij}^0 the rest length, $\Delta l_{ij} = |x_{ij}| - l_{ij}^0$ the stretch, $x_{ij} = x_i - x_j$ the relative position, $\hat{x}_{ij} = x_{ij}/|x_{ij}|$ the unit vector along the edge, and $\dot{l}_{ij} = (\dot{x}_i - \dot{x}_j) \cdot \hat{x}_{ij}$ the stretch rate, i.e. the component of the relative velocity along the edge, so that the damper resists changes of edge length only. Three modelling choices deserve explicit statement. The body is represented as a *surface shell* of springs with no volumetric or internal-pressure term—a choice matched to the physical test objects, which are cast soft shells with internal ballast (Section 8). Node masses are lumped uniformly over the shell, $m_i = m/n$ with n the node count of the adjacency assembly below, and the total mass (shell plus ballast) enters through the weight in the equilibrium condition. And the nodal damping $b\, \dot{x}_i$ acts on *absolute* velocities, so it damps rigid-body motion as well as deformation; this is acceptable for the supported, quasi-static grasp scenarios studied here, but the absolute form does not conserve momentum for a free-flying body, where a relative-velocity (deformation-only) damping should replace it. Equation (3) describes a linear visco-elastic edge with non-linear directional behaviour [20]. The schematic of a single one-dimensional element is shown in Fig. 2.

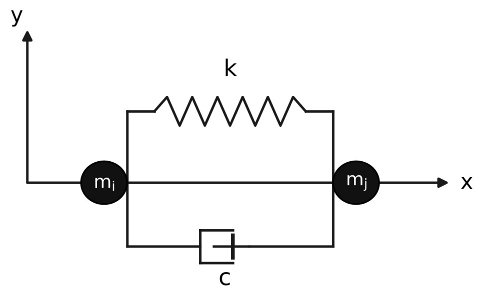

Figure 2. One-dimensional mass–spring–damper element between two nodes i and j. Stiffness k acts in parallel with viscous damping c.

2.2. Three-dimensional extension and state-space form

Extending Eq. (1) to three dimensions, the position of node i becomes a vector $\bar{x}_i = (x_{ix},\ x_{iy},\ x_{iz})^{\mathsf{T}}$, and the magnitude of the relative position is the Euclidean norm

$$|\bar{x}_{ij}| = \sqrt{(x_{ix} - x_{jx})^2 + (x_{iy} - x_{jy})^2 + (x_{iz} - x_{jz})^2}. \tag{4}$$

Solving Eq. (1) for each coordinate of each node and stacking the components into a state vector $X \in \mathbb{R}^{6n \times 1}$ that interleaves position and velocity for every node along x, y and z, the dynamics admit a compact state-space form

$$\dot{X} = F(t, X) + F^{\text{ext}}, \tag{5}$$

where $F(t, X)$ collects the inertial and internal terms and F^{ext} collects the external excitations. The topology of the network—which nodes are connected by springs—is encoded in an $n \times n$ symmetric adjacency matrix A whose entries A_{ij} are 1 if nodes i and j are connected and 0 otherwise. Once A is supplied, the population of $N(i)$ in Eq. (1) is read directly from row i, so the assembly of $F(t, X)$ reduces to a sparse sweep through the non-zero entries of A. We integrate Eq. (5) using a fourth-order Runge–Kutta scheme, which provides ample accuracy at the small simulation time-steps used in this work and is straightforward to vectorise. The complete formulation supports any triangulated surface mesh exported, for example, from a CAD package as an STL file.

3. Grasping Strategy

The grasping pipeline consists of three sequential stages, illustrated in Fig. 3: (i) initial gripper-pose selection driven by a heuristic grasp metric,

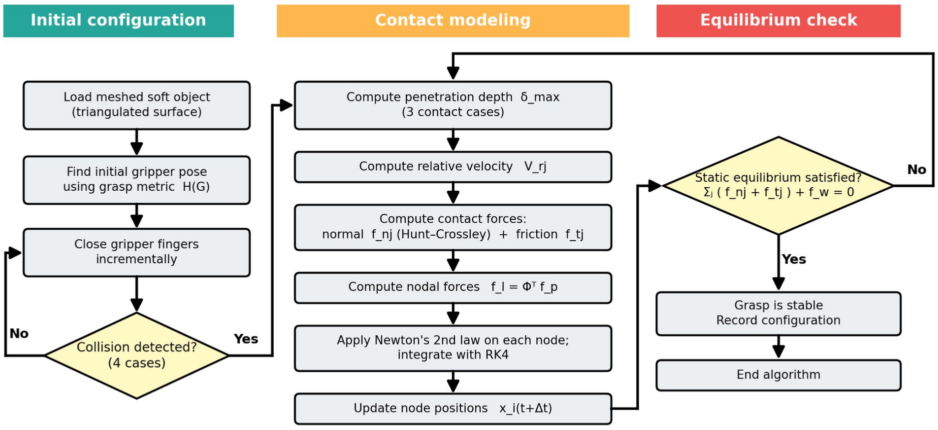

Figure 3. Flowchart of the proposed grasping strategy for deformable objects. Three sequential stages—initial configuration, contact modelling, and static-equilibrium check—evaluate each candidate grasp; the nodal-force transfer uses the redistribution operator Φ of Eq. (17).

(ii) contact modelling—collision detection, penetration computation, and force transfer to the deformable object, and (iii) a static-equilibrium check that determines whether the candidate configuration is stable. The contact-modelling loop runs once per simulation time-step; the equilibrium check is performed at the end of each step and provides the termination criterion. If the gripper is unable to satisfy equilibrium it returns to the closing stage and increments its grip—realised in simulation as an increment of finger advance (penetration), whose normal force then follows from Eq. (13); there is no separate force command in the simulated pipeline. If equilibrium is satisfied the configuration is recorded as a stable grasp.

3.1. Heuristic grasp metric

Grasping a soft body with a multi-fingered hand has more unknowns than there are equilibrium equations: the contact problem is under-determined (statically redundant). To pre-select promising configurations before running the full dynamic simulation, we score every candidate three-finger configuration G with a closed-form heuristic function inspired by Park and Starr [40] and extended for deformable objects:

$$H(G) = w_1 S_1 + w_2 S_2 + w_3 S_3 + n_c P_c + n_e P_e. \tag{6}$$

Here S_1, S_2, S_3 are three normalised regularity measures—*lower* is better, with the minimum value zero attained at the ideal configuration and values of order unity for irregular candidates—and P_c, P_e are penalty terms attached to integer counters n_c and n_e. The total heuristic $H(G)$ is minimised over candidate finger configurations to pre-select the most promising grasp before the dynamic simulation; smaller H means a more promising candidate. The five components are defined as follows.

S_1 — *Force-direction balance.* The three normal contact forces should be co-planar and symmetrically separated by 120°, so that their sum can be made zero by an appropriate magnitude scaling. Given the three measured angles $\theta_{12}, \theta_{23}, \theta_{31}$ between successive force vectors, we set

$$S_1 = \frac{|120^\circ - \theta_{12}| + |120^\circ - \theta_{23}| + |120^\circ - \theta_{31}|}{240^\circ}. \tag{7}$$

S_2 — *Contact-triangle regularity.* The triangle whose vertices are the three contact points should be as close to equilateral as possible. We measure the angle deviations from 60° at the three vertices and the dispersion of side-lengths around an optimum d_{opt}; combining these gives

$$S_2 = \frac{s_{21} + s_{22} + s_{23}}{3}, \tag{8}$$

where each $s_{2i} \in [0, 1]$ combines, for vertex i, the normalised deviation of the interior angle from 60° and of the adjacent side length from an optimum d_{opt}, so that $S_2 = 0$ for a perfectly equilateral contact triangle of the target size. The target spacing d_{opt} is a per-object design constant (of the order of the object's dominant dimension) fixed before the search and fixed with the parameter files (Section 11).

S_3 — *Centroid proximity.* The geometric centroid of the contact triangle should coincide with the geometric centroid of the object so that the wrench arising from the three contact forces passes through the centre of mass and the resulting torque is small. With d_c the distance between the two centroids and $d_{\max}$ the largest centroid-to-mesh-vertex distance,

$$S_3 = \frac{d_c}{d_{\max}}. \tag{9}$$

P_c — *Sharp-vertex penalty.* Sharp convex features of the surface are unfavourable contact points because the fingertip rocks on them and cannot establish a stable contact area. On a triangulated mesh we make *sharp* operational: a mesh vertex is sharp when its angular defect (2π minus the sum of the incident face angles, i.e. the integrated discrete Gaussian curvature) exceeds 0.2 rad, and an edge is sharp when the dihedral angle between its two faces deviates from π by more than 30°. These thresholds cleanly separate the smooth test meshes from genuine features: on the 390-node sphere the angular defects are $\approx 4\pi/390 \approx 0.03$ rad and dihedral deviations stay below about 15° even at the coarsest resolution studied (144 nodes), whereas the rounded cube's corner regions and the handle's rim exceed both thresholds. We add a fixed penalty P_c for every contact point that falls on a sharp vertex; n_c is the count of such contacts. The *signed* defect is used deliberately: concave and saddle features are not penalised by P_c—concavity handling is deferred to future work (Section 10, item 1).

P_e — *Sharp-edge-proximity penalty.* Similarly, contact points lying within half a local mean mesh-edge length of a *sharp feature edge* (in the dihedral sense above) risk sliding off the feature as the body deforms; we penalise these with P_e per offending contact (count n_e). The rule is deliberately tied to feature edges, not to the boundaries of the host triangle: every interior point of a triangle lies within the inradius of its boundary—always less than half the shortest side—so a face-boundary rule would fire on every contact and carry no information. The penalty weights P_c, P_e and the term weights w_1, w_2, w_3 are fixed design constants (set to unity in the simulation studies—uniform scaling of the weights leaves the argmin unchanged by construction, see Section 5.2—while the hardware campaign of Section 8 uses the weighting documented there). They are *not* part of the PSO parameter vector, which identifies only the six physical contact parameters (Section 4).

3.2. Collision detection

The fingertip is modelled as a sphere of radius r; the soft object is a triangulated surface. Collision between the sphere and a particular face yields a circular section of the face contained inside the sphere. To test efficiently against every face, we first transform the world-frame coordinates of the fingertip centre and radius into the local frame of the face, in which the face lies in the (x', y') plane. Letting B be one face vertex, $\hat{b}$ the unit vector along $\overrightarrow{BA}$ (A is another vertex), $\hat{c}$ the unit normal to the face, $\hat{c} = (\hat{b} \times \overrightarrow{BC})/|\hat{b} \times \overrightarrow{BC}|$, and $\hat{d} = \hat{c} \times \hat{b}$, the rotation matrix $R_{3\times3} = [\hat{b}, \hat{d}, \hat{c}]^\mathsf{T}$ together with the translation B yields the homogeneous transformation $DH \in \mathbb{R}^{4\times4}$ that maps world coordinates to local coordinates. In the local frame, the intersection of the fingertip sphere with the face plane $z' = 0$ is the circle

$$(x' - x_0')^2 + (y' - y_0')^2 = R^2 - z_0'^2, \tag{10}$$

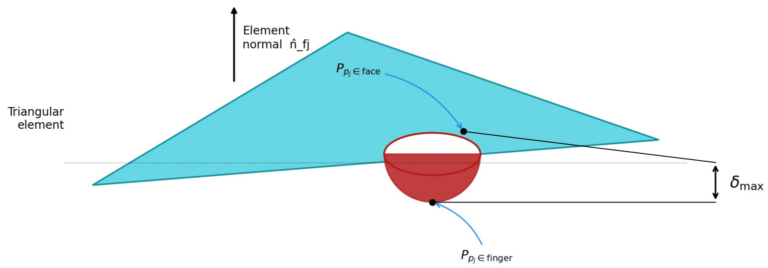

Figure 4. Definition of penetration depth δ_{max} between the fingertip sphere and a triangular face element of the soft-object surface mesh. The depth is the distance between the two reference points $P_{p_j \in \text{face}}$ and $P_{p_j \in \text{finger}}$ along the face normal.

with (x'_0, y'_0, z'_0) the transformed sphere centre. The sphere intersects the face *plane* if and only if $R^2 - {z'_0}^2 > 0$; whether this intersection disc actually overlaps the triangular face is then decided by four geometric cases: (i) the disc is entirely outside the triangle and at distance greater than r from every edge—no contact; (ii) the disc clips one external edge—contact at a single segment; (iii) the disc is centred on a corner-region—contact at a vertex; (iv) the disc is contained entirely inside the triangle—full contact. Each case admits a closed-form test using the disc centre and the per-edge perpendicular distances, requiring at most a handful of dot products per face.

3.3. Penetration depth

After collision is detected we must compute the depth δ_{max} by which the fingertip has penetrated the face, because the contact-force magnitudes scale with this quantity (Section 3.5). We define the line connecting two reference points $P_{p_j \in \text{face}}$ on the face surface and $P_{p_j \in \text{finger}}$ on the fingertip surface; this line is parallel to the face normal $\hat{n}_{f_j}$ (Fig. 4). The penetration depth is the Euclidean distance between these two reference points,

$$\delta_{\text{max}} = |P_{p_j \in \text{face}} - P_{p_j \in \text{finger}}|. \tag{11}$$

The placement of $P_{p_j \in \text{face}}$ inside the triangle is treated in three subcases: (i) the projection of the sphere centre on the face plane lies inside the triangle—P_{p_j} is the projection itself; (ii) the projection lies outside the triangle but within distance r of an edge—P_{p_j} is taken on the nearest edge segment by projecting onto the line BC; (iii) the projection lies in the corner region—P_{p_j} is the triangle vertex itself. In every case $P_{p_j \in \text{finger}}$ is taken on the sphere surface along the same normal direction.

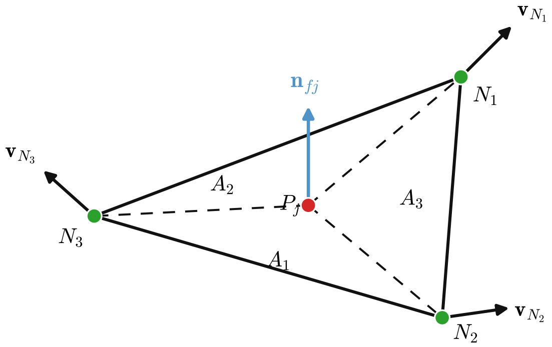

Figure 5. Barycentric interpolation of vertex velocities $V_{N_1}, V_{N_2}, V_{N_3}$ to compute the contact-point velocity at P_j. Each area A_i is the sub-triangle formed by P_j and the two vertices *other than* N_i, so vertex N_i carries the weight A_i/A_G of its opposite sub-triangle.

3.4. Relative velocity

Computing the friction force requires the relative tangential velocity between the fingertip contact point and the corresponding point on the face. The translational and rotational velocities of the fingertip in the world frame R_O are read directly from the gripper's kinematics; the velocity of the face contact point P_j is interpolated from the three vertex velocities $V_{N_1}, V_{N_2}, V_{N_3}$ using barycentric area weights, each vertex weighted by the area of the sub-triangle *opposite* it (Fig. 5):

$$V_{P_j \in \text{face}}\Big|_{R_O} = \frac{A_1}{A_G} V_{N_1} + \frac{A_2}{A_G} V_{N_2} + \frac{A_3}{A_G} V_{N_3}, \qquad A_G = \sum_{i=1}^{3} A_i. \tag{12}$$

The relative velocity is $V_{r_j} = V_{P_j \in \text{finger}}|_{R_O} - V_{P_j \in \text{face}}|_{R_O}$. Its decomposition along the face normal and tangent gives, respectively, $V_{n_j} = (V_{r_j} \cdot \hat{n}_{f_j})\, \hat{n}_{f_j}$ and $V_{t_j} = V_{r_j} - V_{n_j}$.

3.5. Contact forces

Following Hunt and Crossley [35], the normal contact force at a deformable contact is modelled as a non-linear visco-elastic law,

$$f_{n_j} = \max\Big(0,\; K\,\delta_j^n + C\,\dot{\delta}_j\Big), \qquad \mathbf{f}_{n_j} = f_{n_j}\,\hat{n}_{f_j}, \tag{13}$$

where δ_j is the penetration depth, n is the Hunt–Crossley exponent that depends on the contact geometry (distinct from the node count of Section 2.1), K is the contact stiffness and C is the contact damping (the capitalised K, C are the normal-contact coefficients, distinct from the MSD edge stiffness k

and damping c of Section 2.1). The elastic term $K\delta^n$ is the Hunt–Crossley power law; for the dissipative term we adopt a damping linear in $\dot{\delta}$ rather than the original indentation-weighted form $\propto \delta^n\dot{\delta}$ [35]. To be precise about the naming: the indentation-weighted damping is the *defining* feature of Hunt–Crossley (it removes the tensile-force discontinuity at separation), so Eq. (13) is strictly a power-law spring with a linear, Kelvin–Voigt-type rate damper—a Hunt–Crossley-*type* model rather than the original—and the separation discontinuity the linear damper reintroduces is handled instead by the $\max(0, \cdot)$ clamp. The simplification is adequate at the low approach speeds used throughout this work and, importantly, is the form the bench calibration of Section 7 identifies directly from measured force–indentation records. For a spherical fingertip against a planar face element [41], K and C may be related to material properties via

$$K = 2\lambda_1\Psi\sqrt{R}, \qquad \Psi = \frac{1}{\dfrac{1-\nu_1^2}{E_1} + \dfrac{1-\nu_2^2}{E_2}}, \qquad C = 4\pi R\lambda_2, \tag{14}$$

with E_1, E_2 the Young's moduli of the fingertip and object, ν_1, ν_2 their Poisson's ratios, R the fingertip radius, and λ_1, λ_2 dimensionless identification scale factors (Ψ is the reciprocal of the summed contact compliances, i.e. the effective contact modulus; because K's units depend on the per-family n, the unit-clean comparison is the force at a reference indentation, $f_{0.5} = K\,(0.5)^n$, which rises monotonically with modulus from 0.25 to 0.51 N across the families, Table 4). Because the fitted exponent n differs from the Hertzian 3/2 implicit in the $\sqrt{R}$ scaling, Eq. (14) is a scaling guide rather than a strict dimensional identity: the identified λ_1 absorbs the residual dimension (and λ_2 likewise for C; both are quoted as pure numbers, Table 3), and throughout this work δ is expressed in millimetres with K in $\mathrm{N\,mm^{-n}}$, whereas the damping constant C is identified with $\dot{\delta}$ in $\mathrm{m\,s^{-1}}$ (units $\mathrm{N\,s\,m^{-1}}$, Table 4); the two length conventions must not be mixed when evaluating Eq. (13). Tangential contact follows a Coulomb stick–slip transition [36, 37, 38]. In the slip regime the friction force is

$$\mathbf{f}_{\text{slip}} = -\mu\, f_n \frac{V_t}{\|V_t\|}, \tag{15}$$

while in the stick regime a tangential elastic deflection ξ accumulates between the current contact point and its stick anchor, and the tangential force is

$$\mathbf{f}_{\text{stick}} = -(k_t\xi + c_t\dot{\xi})\,\hat{\mathbf{t}}, \tag{16}$$

with $\hat{\mathbf{t}}$ the unit tangential (deflection) direction. The transition between regimes is governed by the Coulomb criterion: while $\|\mathbf{f}_{\text{stick}}\| \leq \mu_s f_n$ the contact sticks, the deflection accumulating as $\dot{\xi} = V_t$ with no macroscopic slip ($\dot{u} = 0$); once $\|\mathbf{f}_{\text{stick}}\|$ would exceed $\mu_s f_n$ the contact slips, the deflection saturating at the friction-cone limit ($\dot{\xi} = 0$) while the slip variable advances ($\dot{u} = V_t$). To be explicit about which coefficient enters where: the stick-to-slip *transition* test uses the static coefficient μ_s; the sliding force of Eq. (15) uses the kinetic coefficient $\mu = \mu_k$; where a single symbol μ appears in the simulation sweeps of Sections 5–6 the two are set equal (a single-coefficient Coulomb model), whereas the bench campaign of Section 7 measures both (Table 4) and the retention budget of Section 9 uses μ_s. Parameter k_t is the tangential (stick) stiffness, identified on the bench (Table 4); c_t is a small stabilising tangential damping set in the recorded configuration and not separately identified—it multiplies $\dot{\xi}$ and therefore vanishes at every static equilibrium, so none of the reported quasi-static results depend on it.

3.6. Nodal force transfer

The contact force $\mathbf{f}_{p_j} = \mathbf{f}_{n_j} + \mathbf{f}_{t_j}$ computed at the contact point P_j must be redistributed to the three vertices N_1, N_2, N_3 of the host triangle in a way that respects local equilibrium. Using the same barycentric area weights as in Eq. (12), the redistribution matrix is

$$\Phi = \left[\frac{A_1}{A_G} I,\ \frac{A_2}{A_G} I,\ \frac{A_3}{A_G} I \right] \in \mathbb{R}^{3\times 9}, \tag{17}$$

where I is the 3×3 identity matrix (the symbol Φ is used to avoid confusion with the heuristic $H(G)$ of Eq. (6)). The nodal force vector $\mathbf{f}_l$, stacking the three vertex forces in (x, y, z), is

$$\mathbf{f}_l = \Phi^{\mathsf{T}} \, \mathbf{f}_{p_j}. \tag{18}$$

Because Φ is the operator that interpolates nodal velocities to the contact point, distributing the force with its transpose is the virtual-work-consistent choice: the barycentric weights sum to one, so the total transferred force equals $\mathbf{f}_{p_j}$ and the moment about the contact point is preserved. Once $\mathbf{f}_l$ has been computed for every contact point, the resulting forces enter Eq. (5) as the external excitation F^{ext} on the corresponding nodes, and the body is integrated one time-step forward.

3.7. Static-equilibrium check

Lifting the soft object stably requires that the vertical component of the sum of all contact forces balances the object weight $\mathbf{f}_w = m\mathbf{g}$. The static-equilibrium criterion is therefore

$$\sum_{j=1}^{N_c} \left(\mathbf{f}_{n_j} + \mathbf{f}_{t_j} \right) + \mathbf{f}_w = \mathbf{0}, \tag{19}$$

where N_c is the number of active fingertip contacts. Practically, the vertical component of $\sum_j \mathbf{f}_j + \mathbf{f}_w$ is monitored during the simulation; once its magnitude settles below a chosen tolerance τ_F (the tolerance of the pipeline check in Algorithm 1) the configuration is declared stable and the algorithm terminates. The value of τ_F is campaign-specific and should not be conflated across sections: the simulation analyses of Sections 5–6 score success at the 0.5 N criterion, whereas the hardware planner's certification flag uses the tighter recorded value $\tau_F = 0.0886\,\mathrm{N}$ ($\approx 0.09\,\mathrm{N}$) (Section 8). The torque tolerance τ_T is pinned alongside it in the recorded configuration and is not binding in any reported campaign: torque residuals are essentially zero in the symmetric configurations studied (Fig. 17b), the asymmetric failures (handle, ellipsoid) are caught by the slip-margin and force checks first, and the hardware flag reduces to the force criterion alone (Section 8).

4. Parameter Identification via PSO

The model contains six parameters that do not map directly to measurable material properties—the Hunt–Crossley scale factors λ_1, λ_2 of Eq. (14), the exponent n in the normal-force law, the nodal damping b, the spring stiffness k, and the edge damping c. Rather than calibrate these by trial-and-error, we identify them by minimising a composite cost function with a Particle Swarm Optimization (PSO) algorithm. Three physical conditions are imposed.

Force balance. The vertical sum of normal contact forces must equal the object weight along z. Any residual quantifies the first cost term J_1.

Bounded penetration. The maximum allowable penetration depth in any contact is taken as one millimetre ($10^{-3}\,\mathrm{m}$); any excess contributes to J_2. Penalising over-penetration keeps the transient physical and guards against the mesh "explosion" that ill-tuned spring–damper parameters can cause.

Friction-cone slip margin. The friction-cone slip margin $\mu f_n - \|f_t\|$ must remain non-negative at every contact; a negative margin (incipient slip) contributes to J_3. This term steers the identification toward parameter sets that produce stable, non-slipping holds.

The total cost is $J = J_1 + J_2 + J_3$ (the explicit dimensionless form used throughout is stated in Section 11), and the parameter vector $(\lambda_1, \lambda_2, n, b, k, c)$ is searched by a standard PSO routine with inertia weight, cognitive and social coefficients. The identification is run once for each soft body to be grasped. One property of this simulation-side identification deserves explicit acknowledgement: its cost criteria coincide with the stability criteria later used to score grasp success, so in simulation the no-PSO ablation of Section 6.3

partly measures objective alignment rather than physical fidelity. The physical test of the identification is therefore deferred to the campaigns of Sections 7–9, where the parameters are re-identified against measured force traces and success is a measured outcome rather than a modelled criterion.

5. Simulation Results

5.1. Simulation setup

All simulations are run in MATLAB. The soft object is a uniform sphere of radius 5 cm modelled by Solidworks, exported as an STL file, and meshed in MATLAB into 356 triangulated surface nodes (Fig. 6; the multi-shape studies of Section 6 use a 390-node remesh of the same geometry, and Section 6.7 verifies stability across 144–1750 nodes). The gripper has three fingers, each with three 4 cm links joined by three revolute joints (nine kinematic degrees of freedom in total). Two fingers lie in the same plane and the third is mounted 120° behind them (Fig. 7); only the closing direction is actuated. The gripper is implemented using the SynGrasp toolbox [39] (Fig. 7). Material and simulation constants are listed in Table 2; each dynamic run spans the 100 integrator steps plotted in Fig. 12 at 1.8 s wall-clock per run (the quasi-static evaluator's equilibrium-loop budget, Section 11, coincides at 100 iterations but is a distinct loop), and the physical step size is pinned in the recorded configuration (Section 11). Two numerical hazards of the contact law deserve note: with the fitted $n < 1$ the contact stiffness $\mathrm{d}f/\mathrm{d}\delta = nK\delta^{n-1}$ diverges as $\delta \to 0^+$, and the linearised damping introduces a finite force step $C\dot{\delta}$ at contact onset for non-zero approach speed. Two guards keep the explicit integration stable in practice: contacts are seeded at a finite 0.5 mm penetration (Algorithm 2 uses the same seed in the static loop; at gel-scale bench stiffness this seed itself injects an initial force step of $K(0.5)^n \approx 0.25\,\mathrm{N}$ per contact, an artefact of regularising the $n < 1$ divergence) and the $\max(0, \cdot)$ clamp removes spurious tensile forces; the settled, oscillation-free trajectories of Figs. 12–13 are the empirical stability evidence for the recorded step size. The Young's moduli were chosen to mimic a soft silicone fingertip ($\approx$ 90 kPa) gripping a soft object of slightly higher modulus ($\approx$ 140 kPa), with Poisson's ratios in the near-incompressible range.

5.2. Sensitivity of the heuristic function

We first study how the heuristic weights influence the optimal grasp configuration on the spherical body. Because the sphere mesh has no sharp features under the criteria of Section 3.1 (angular defects ≈ 0.03 rad, dihedral deviations below 15°), the penalty counts n_c and n_e are identically zero, so P_c and P_e have no effect, and the optimum collapses to a function of (w_1, w_2, w_3).

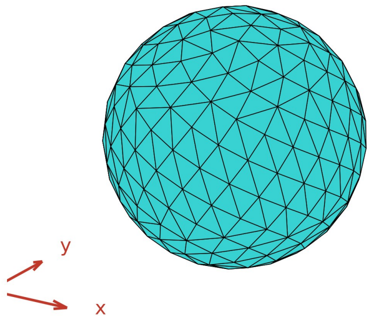

Figure 6. Triangulated surface mesh of the spherical soft object (radius 5 cm, 356 nodes) in the MATLAB simulation environment.

Table 2. Fixed constants and material parameters used in the simulations.

Parameter	Value
Soft-object radius	5 cm
Number of mesh nodes	356
Fingertip radius	5 mm
Wall-clock time per full dynamic simulation (100 RK4 steps)	1.8 s
Fingertip Young's modulus E_1	0.0893 MPa
Object Young's modulus E_2	0.14 MPa
Fingertip Poisson's ratio ν_1	0.5
Object Poisson's ratio ν_2	0.45

Two clarifications define what is being optimised. First, the search here—as throughout the dynamic pipeline—is over configurations *reachable by the SynGrasp hand* (two coplanar fingers approaching from above with the third mounted 120° behind), not over unconstrained contact triples. Unconstrained, Eq. (6) is minimised by three equatorial contacts at 120° spacing, for which $S_1 = S_3 = 0$ exactly—and the quasi-static evaluator of Section 6, which searches mesh-vertex triples without the kinematic constraint, indeed selects near-equilateral placements of exactly this family (Section 6.4). That arrangement, however, is not reachable by this hand from a top approach. Second, within the reachable set, sweeping the weights uniformly leads consistently to the configuration shown in Fig. 8: two contact points on a plane through the centre at 45° on either side and a third contact at 90° in the opposite direction. The corresponding heuristic-component values are $S_1 = 0.25$,

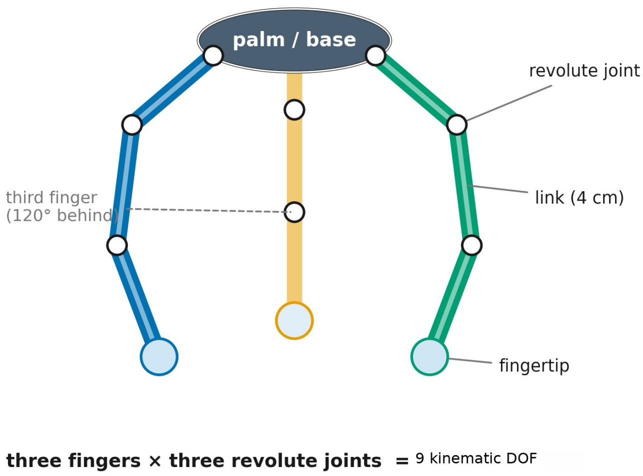

Figure 7. Schematic of the three-fingered gripper used in the simulations (modelled with the SynGrasp toolbox). Each finger comprises three 4 cm links and three revolute joints—nine revolute joints (kinematic DOF) in total, of which a single synchronised closing command is actuated in this work; the third finger is mounted 120° behind.

$S_2 = 0.25$, $S_3 = 0.33$, $n_c = n_e = 0$ (so $H = 0.833$ at unit weights is the *reachable* optimum, not the global one).

Three runs with different sets of weights all converge to qualitatively the same configuration: (i) $w_i = 1$, $P_c = P_e = 1 \Rightarrow H = 0.833$; (ii) $w_i = 0.5$, $P_c = P_e = 0.5 \Rightarrow H = 0.416$; (iii) $w_i = 2$, $P_c = P_e = 2 \Rightarrow H = 1.666$. The numerical-value scaling reflects the proportional scaling of the weights but the geometry is invariant—as expected for a homogeneous sphere. For non-spherical bodies the sharp-vertex and edge-proximity penalties become active; we keep all weights fixed at unity throughout the simulation studies. Note that uniform scaling leaves the argmin unchanged by construction, so this sweep cannot probe *ratio* changes $w_1{:}w_2{:}w_3$; the single-term ablations of Section 6.3, which are the extreme ratio cases, are the closest available probe of that axis.

5.3. PSO convergence and identified parameters

Figure 9 shows the convergence of the PSO algorithm on the composite cost. The cost falls from approximately 2300 to near zero within the first 10,000 function evaluations and reaches a low-residual plateau by 50,000 evaluations. The corresponding identified parameters are listed in Table 3.

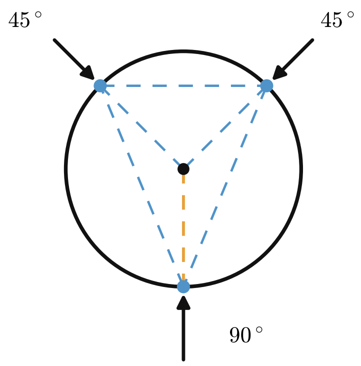

Figure 8. Optimal three-finger contact-point configuration on the spherical body. Two contacts at 45° above the equatorial plane and a third at 90° below give the minimum-heuristic configuration.

5.4. Gravity-induced deformation

As a first qualitative test of the MSD model, the sphere is placed on a flat surface and released under gravity. With the identified parameters, the soft body deforms, oscillates briefly, and settles into the squashed configuration shown in Fig. 10(b). Once the system reaches steady state no further deformation occurs, confirming that the parameter set produces stable dissipative dynamics.

5.5. Grasping under varying mass and friction

We finally test the full grasping pipeline on six of the nine combinations of friction coefficient $\mu \in \{0.1, 0.3, 0.5\}$ and object mass $m \in \{0.2, 0.4, 0.6\}$ kg (the recorded run logs identify them; the representative case reported here is $\mu = 0.1$, $m = 0.4$ kg). In each case the gripper performs the sequence depicted in Fig. 11: opening, descent, finger closure with progressive contact, and lift-off. Figure 12 reports the two error metrics for the representative case ($\mu = 0.1$, $m = 0.4$ kg): the static-equilibrium error along z and the total penetration sum. Both decay to near-zero within ∼50 iterations—the equilibrium error oscillates around zero with amplitude below 10^{-1} N and the total penetration stabilises below 5×10^{-4} m, well within the design limit of 10^{-3} m per contact. The same qualitative behaviour was observed for all tested combinations. A units note is required before reading force magnitudes off Figs. 12–13: these runs use the identified parameters of Table 3, which are expressed in the simulator's internal unit system, so the force and residual axes are internal-system values whose magnitudes are not commensurate with

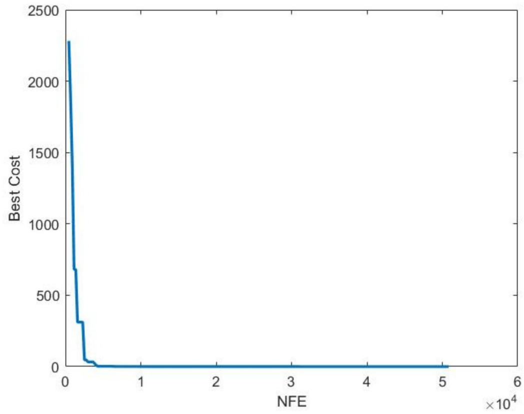

Figure 9. Convergence of the PSO algorithm: best cost vs. number of function evaluations (NFE). The composite cost drops rapidly and stabilises near zero, indicating that the identified parameters satisfy all three physical constraints.

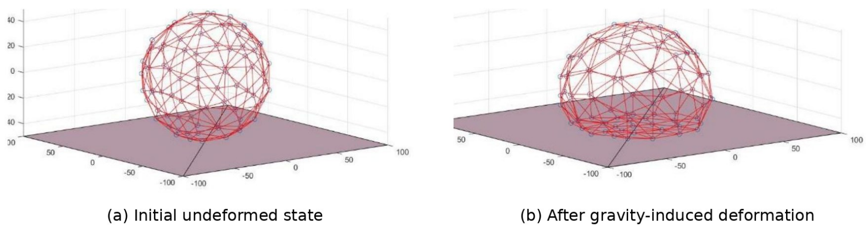

Figure 10. Deformation of the soft body under gravity. (a) Initial undeformed state. (b) After deformation and full settling.

the physical weight mg (the same caution as for the fixed-command study of Section 6.4); the physically calibrated force scale enters the pipeline only through the bench identification of Section 7 and is validated on hardware in Sections 8–9.

The contact-force magnitudes on each of the three fingers are reported in Fig. 13. After an initial transient lasting roughly 20–30 iterations, the forces on fingers 1 and 3 settle to approximately 5 N and 12 N respectively, while finger 2 carries only a small residual. The absence of late-time oscillations confirms that the grasp is stable. Changing the friction coefficient mostly affects the magnitude of the steady-state forces (with lower friction requiring larger normal forces to prevent slip—a property of this adaptive force-increment loop; the fixed-command study of Section 6.4 behaves differently by design), while changing the object mass primarily affects the duration of the transient. That finger 2 carries almost no steady load means the settled equilibrium is effectively a two-contact pinch between fingers 1 and 3, with finger 2 engaged

Table 3. PSO-identified values of the six unknown contact and MSD parameters. n is dimensionless; λ_1 and λ_2 are quoted as pure numbers with the residual dimensions of Eq. (14) absorbed into K ($\mathrm{N\,mm^{-n}}$) and C ($\mathrm{N\,s\,m^{-1}}$) respectively, since $n \neq 3/2$ breaks strict dimensional homogeneity there; b, k and c are the nodal-damping, MSD-spring, and edge-damping coefficients of Section 2.1 in the simulator's internal mass–length–time system, fixed with the parameter files (Section 11)—only the physically re-identified SI constants of Section 7 enter the hardware-validated pipeline.

Parameter	Identified value
λ_1	0.332
λ_2	0.880
n	0.487
b	0.646
k	0.138
c	0.00247

but lightly loaded—in tension with the three-balanced-forces ideal behind S_1. We read this as the reachable compromise on this geometry: the third contact acts as a kinematic guard and a redistribution path once external loads arrive (the disturbance campaign of Section 9 shows exactly such load redistribution, with the planner predicting which finger slips first), rather than as an equal sharer of the static load.

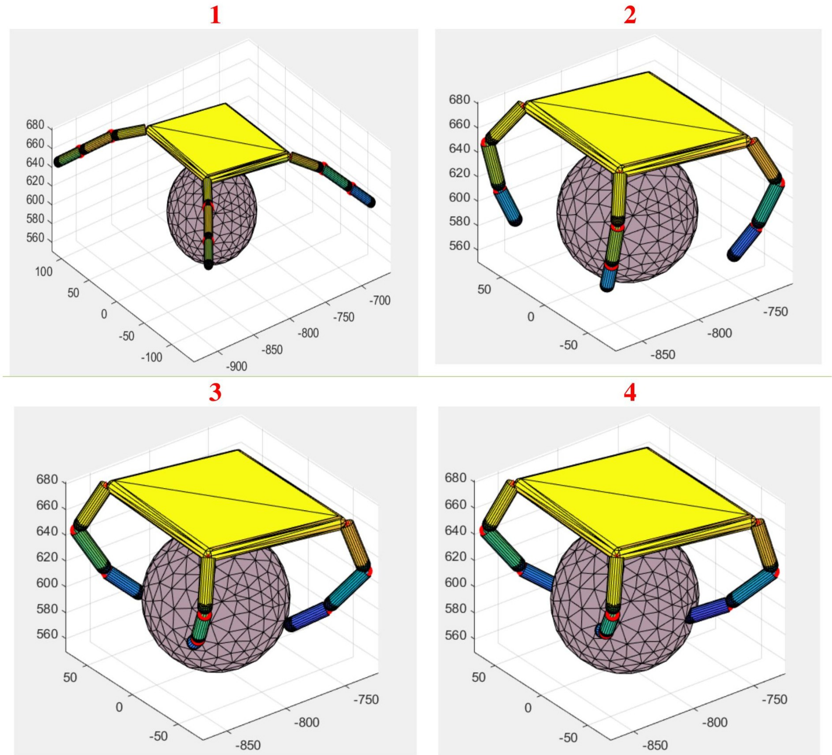

Figure 11. Four-frame grasping sequence: (1) approach, (2) initial contact, (3) finger closure with progressive contact, (4) stable grasp.

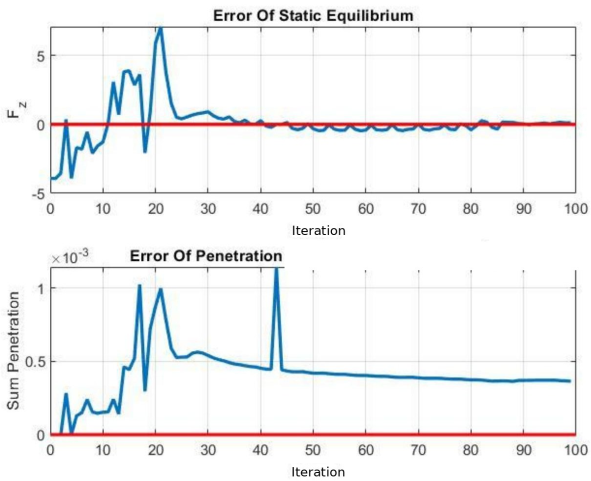

Figure 12. Convergence of the static-equilibrium residual (top) and total penetration sum (bottom) for the representative case $\mu = 0.1$, $m = 0.4\,\mathrm{kg}$. The dashed red line marks the target zero.

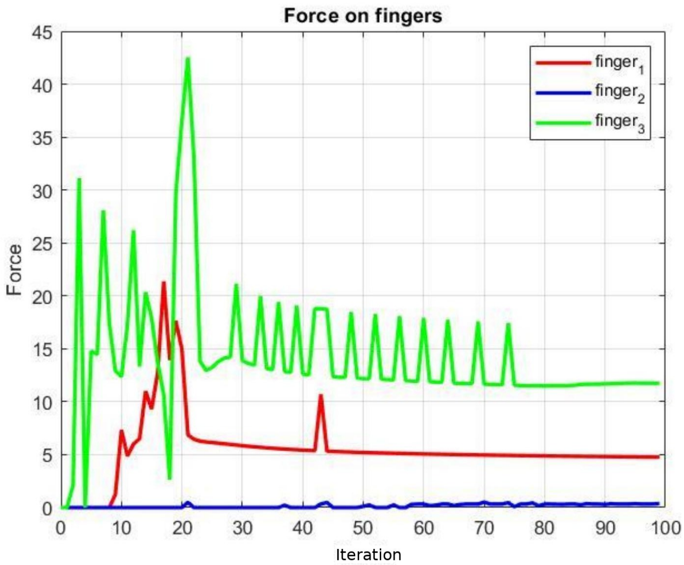

Figure 13. Force exerted by each finger on the object during the grasping sequence ($\mu = 0.1$, $m = 0.4\,\mathrm{kg}$). Forces settle to stable steady-state values, confirming a stable grasp.

6. Extensions: multi-shape validation, baselines, and analysis

The original validation in Section 5 used a homogeneous spherical object, on which the sharp-feature penalties P_c and P_e are inactive. To exercise the full heuristic and to demonstrate generality, we extend the validation to five object shapes, ablate the heuristic components, compare against three baseline grasp planners, and analyse robustness and mesh-resolution dependence. All extensions use the same MSD, contact, and PSO machinery developed in Sections 2–4, with the dynamic simulator replaced by a quasi-static evaluator that returns the same equilibrium residual and slip-margin metrics.

6.1. Scope of the simulation comparisons

All scope conditions on the evaluator comparisons are collected here once, so that the result sections can be read without repeated qualification (the grip-law analysis of Section 6.5 is the one further scope note). Two differences between the evaluators must be stated precisely. First, the quasi-static evaluator searches candidate *mesh-vertex triples* directly, without the hand-reachability constraint of the dynamic pipeline (Section 5.2); on the sphere it therefore selects near-equilateral placements of the unconstrained-optimum family. Second, friction among three contacts is statically indeterminate; the dynamic model resolves the distribution through the stick-state history of Eq. (16), whereas the quasi-static evaluator uses the uniform-utilisation assumption—tangential force proportional to normal force at every contact, $f_{t,j} = \eta f_{n,j}\,\hat{t}_j$, with $\hat{t}_j$ the unit projection of $-\hat{g}$ onto the tangent plane at contact j and the single scalar η set by the net force balance along gravity (the torque residual is checked separately, Algorithm 1 line 6)—so its slip margin reduces to $\min_j f_{n,j}(\mu - \eta)$, the minimum taken over contacts with $\hat{t}_j \neq 0$: a contact whose normal is vertical has a horizontal tangent plane, carries no tangential load under this rule, and is excluded (a case relevant to the like-for-like 45°/45°/90° run flagged below). One further precision keeps the two feasibility checks distinct. Because η is set by the balance along gravity and is *not* capped at μ (the handle's negative margin proves as much), the post-friction vertical residual is zero by construction whenever any contact normal is non-vertical—a post-friction residual test would be vacuous. The 0.5 N force-residual criterion of this evaluator is therefore applied to Algorithm 2's *normal-only* residual, before friction resolution—the quantity the loop actually controls and the quantity reported in Figs. 15 and 20b—while the resolved friction enters only the slip-margin check. (The dynamic pipeline's residual in Fig. 12 is, by contrast, the full Eq. (19) quantity, since the stick–slip model supplies the friction forces there.) (The same assumption underlies the hardware force predictions of Section 8. Two checks are worth separating. In the recorded per-trial force predictions the ratio $f_{t,j}/f_{n,j}$ carries a median per-trial coefficient of variation of 6.7 % (mean 7.0 %), so the implemented tangential resolution departs from the idealised single-η rule at the several-percent level—a consequence of the per-contact tangent-plane projection rather than of measurement. The physical test uses the independently measured per-finger forces: across the 1755 executed hardware trials the per-trial coefficient of variation of the *measured* $f_{t,j}/f_{n,j}$ has a median of 4.7 %, so uniform utilisation holds physically at the five-percent level rather than exactly—an approximation the force predictions inherit.) With those two caveats, both solvers pose the identical static force-balance problem: on the spherical object, for which the full dynamic results of Section 5 are available, both drive the equilibrium force residual below the 0.5 N success criterion and agree on the friction-cone slip margin (about 0.8 for $\mu \geq 0.3$) and on the success label. The two evaluators are compared on residual, slip margin, and outcome—not on force magnitudes, which depend on the commanded grip and on the selected configuration (Section 6.4). Scope of the agreement claim: the solvers were exercised on their respective optima (the reachable 45°/45°/90° set dynamically, the unconstrained near-equilateral set quasi-statically), so this agreement is at the level of criteria—identical residual threshold, identical margin definition, consistent success labels on the sphere—and not a like-for-like run on an identical contact set. That cleaner check—a single run of the quasi-static evaluator on the same 45°/45°/90° contacts—was not performed for this manuscript. What can be settled without the solver is the *static feasibility* of that contact set, the condition its residual and slip tests jointly encode. Treating the configuration of Fig. 8 as three rigid Coulomb contacts on the sphere—two at 45° from the equatorial plane, diametrically opposite in azimuth, and a third at the pole—and searching over every set of contact forces admissible under the friction cones (non-negative normal components, tangential components bounded by $\mu f_{n,j}$, and vanishing net force and torque about the centre), a friction-consistent static equilibrium exists at each of the nine μ–m combinations of the sweep of Section 5, and persists as $\mu \to 0$: on this contact set the normal directions alone can carry the weight, so the label is not friction-limited. The set therefore cannot be rejected by the quasi-static evaluator on grounds of static infeasibility, which is consistent with the stable labels the dynamic pipeline returns on it. This is a closed-form feasibility check on the contact geometry, reproducible from the geometry alone; it is *not* a run of the evaluator, and it does not establish that the evaluator's own grip-force iterate would land on such an equilibrium. We therefore still do not claim solver equivalence, and the run remains listed among the open comparison items in Section 12. The comparative conclusions that follow are internal to the quasi-static evaluator and are unaffected by this open item;

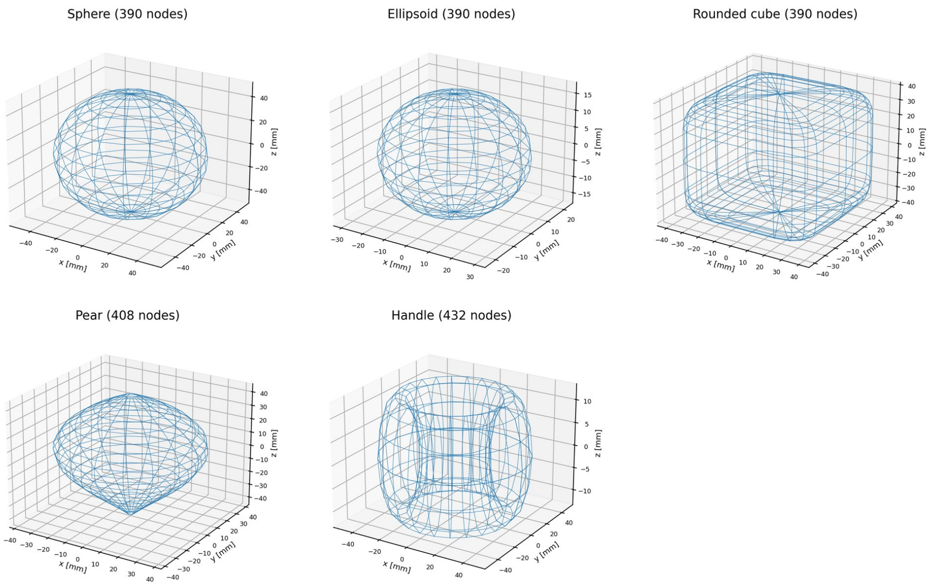

Figure 14. Triangulated test meshes used in the multi-shape validation. Each mesh has approximately 390–430 nodes, comparable to the 356-node spherical reference of Section 5.

the quasi-static form is retained because it is faster and deterministic across runs, though, as noted in Section 12, it does not capture fast transients.

6.2. Non-spherical and irregular objects

Five triangulated surface meshes are used (Fig. 14): a sphere ($r = 5\,\mathrm{cm}$, 390 nodes; the original validation in Section 5 used a 356-node sphere, and the resolution study of Section 6.7 confirms the results are mesh-independent across this range), an ellipsoid ($60 \times 45 \times 35\,\mathrm{mm}$), a rounded cube (half-side 4 cm, fillet 1.2 cm), a pear-shaped solid of revolution, and a concave handle (toroidal cross-section). The rounded cube has eight corner regions whose residual angular defects exceed the sharpness threshold and activate n_c; on the handle it is the rim edges bounding the concave inner surface that exceed the dihedral threshold and activate n_e—the smooth concave surface itself has no sharp dihedrals, and its failure mode is the normal-direction reversal of Section 10.

Grasp quality on each shape is summarised in Fig. 15. The sphere, rounded cube, and pear yield force-balanced grasps with slip margins above zero. Two shapes fail in this single-configuration study, for different reasons. The ellipsoid terminates with a vertical *normal-only* force residual (Algorithm 2's

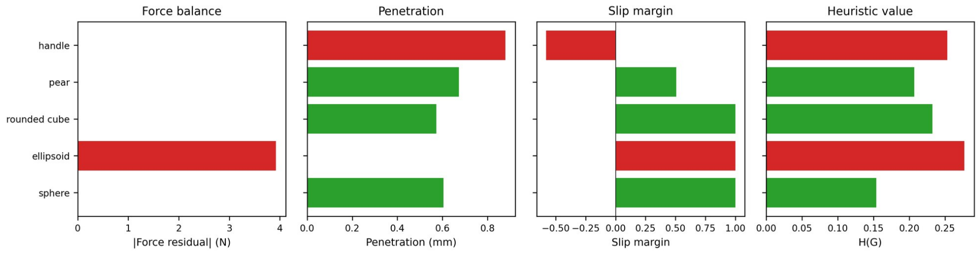

Figure 15. Grasp metrics across the five test shapes. Green bars are successful grasps (force residual <0.5 N, slip margin >0); red bars are failures. The heuristic $H(G)$ remains in the same $[0.15, 0.28]$ range across all shapes, showing that the metric scales geometrically rather than depending on overall size.

pre-friction quantity, see above) above the 0.5 N criterion (the red force-balance bar) even though its slip margin is positive—an equilibrium-settling failure of the iteration-capped run on the elongated geometry at the default seed, not a slip failure; across the 675-trial grid, which varies seeds, masses, and frictions, the ellipsoid succeeds at rates comparable to the other convex shapes (Fig. 18b), and the hardware campaign reaches 96.2 % on ellipsoids under natural friction (Table 6), so the failure is configuration-specific. The concave handle, as expected, scores a negative slip margin because the contact configuration forced by the heuristic on a toroidal shape pulls the fingers outward instead of holding the body. Both modes are catalogued in Section 10.

6.3. Ablation of the heuristic components

To test whether all five terms of Eq. (6) are necessary, we remove each in turn and re-run the planner under otherwise identical conditions (5 shapes $\times$ 10 random seeds per variant). Two non-heuristic baselines—random three-vertex selection and centroid-equilateral placement—are included for context. Success rates are shown in Fig. 16.

Removing any single term leaves the success rate essentially unchanged at 72%; conversely, replacing the heuristic-driven search with random vertex selection drops success to 6% and replacing it with naive centroid-equilateral placement drops it to 0%. (With 50 trials per variant the 95 % confidence interval on each rate is roughly ± 12 percentage points, so the identical 72% figures chiefly indicate that removing one term rarely changes the selected configuration on these seeds; the ablation has limited power to resolve small effects.) The interpretation is that the five terms are partially redundant—each one captures correlated aspects of grasp regularity—but the structured search itself, even with terms removed, dominates the unstructured baselines. We retain all five terms in the recommended pipeline because the redundancy

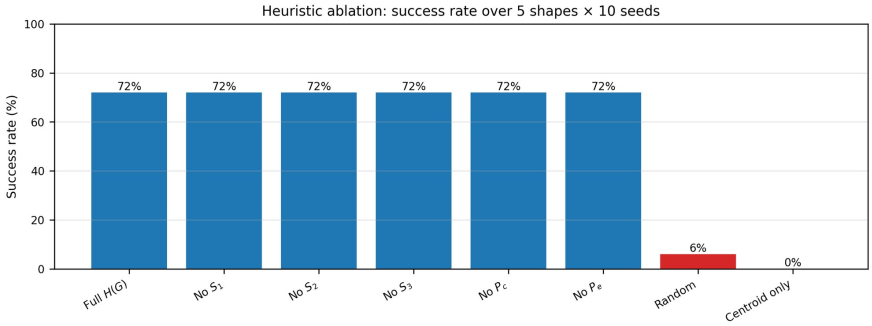

Figure 16. Ablation of the heuristic-grasp metric. The full $H(G)$ and all single-term ablations reach a 72% success rate, whereas the non-heuristic baselines (random, centroid only) fall to 6% or below. The terms are individually redundant on these seeds—any four suffice—while the structured search dominates the unstructured baselines.

provides robustness when one term degenerates (for example, S_1 becomes ill-defined when contact normals are nearly co-linear).

6.4. Wrench-balance stability

Section 3.7 introduced a single equilibrium criterion—the vertical force residual along z. Robotics practice asks for both force and torque balance, plus an explicit slip margin from the friction cone. We therefore monitor four quantities at each iteration: the three-vector force residual $\| \sum_j \mathbf{f}_j + \mathbf{f}_w \|$, the torque residual $\| \sum_j \mathbf{r}_j \times \mathbf{f}_j \|$ about the object centroid, the friction-cone slip margin $\min_j(\mu f_{n,j} - \|f_{t,j}\|)/(\mu f_{n,j})$, and the per-finger normal forces. Figure 17 reports all four for $\mu \in \{0.1, 0.3, 0.5\}$ on the spherical object. This stress study applies a *fixed grip command* across the three friction levels so that the slip margin isolates the effect of μ; that is why the finger-1 normal forces in panel (d) are nearly identical across μ and sit well above the equilibrium-adaptive forces of the dynamic pipeline (Fig. 13), which increments grip only until the residual criterion is met. The command level of this study ($\approx 65\,\mathrm{N}$ per finger) also sits far outside the validated force envelope—the bench calibration used a 10 N load cell with sub-newton typical fits, and the hardware holds near 0.85 N per finger—and the run uses the simulator's internal-unit stiffnesses (Table 3); at the bench-scale constants of Table 4 a 70 N normal force would imply a penetration orders of magnitude beyond the 1 mm bound. Figure 17 is therefore a numerical robustness test of the solver at an extreme operating point, not a physically calibrated prediction. Torque is essentially zero throughout because the unconstrained vertex-triple search selects a near-equilateral contact triangle whose centroid nearly coincides with

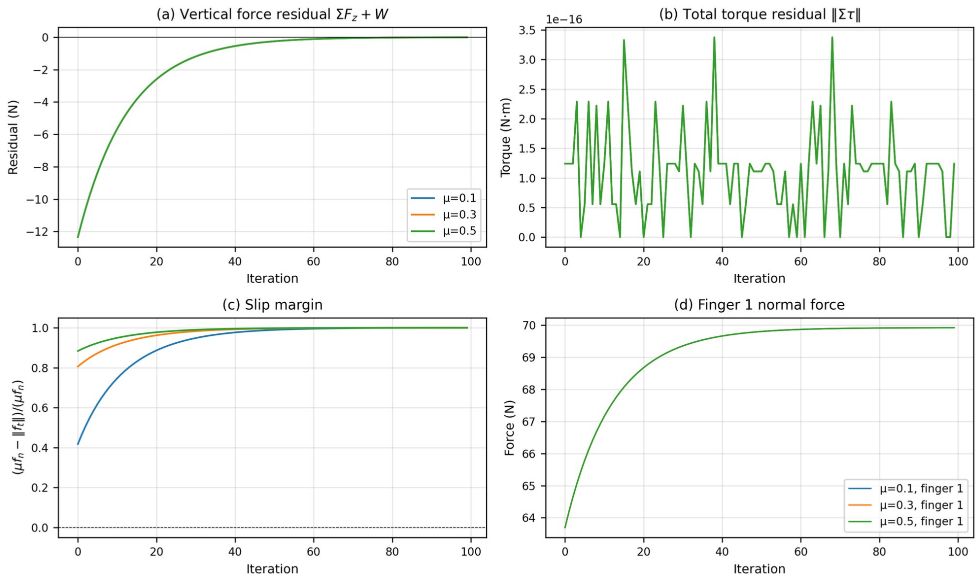

Figure 17. Full wrench-balance and stability metrics on the spherical object for three friction coefficients. (a) Vertical force residual, (b) total torque residual, (c) friction-cone slip margin, (d) finger-1 normal force.

the object centre (in contrast to the hand-reachable 45°/45°/90° optimum of Section 5.2); force residual decays to zero in 30–40 iterations; slip margin grows monotonically once contact engages and saturates above 0.8 for $\mu \geq 0.3$.

6.5. Comparison against baseline grasp planners

We benchmark the heuristic against three baselines: random three-vertex selection, two-vertex antipodal pairs with a third perpendicular contact, and centroid-equilateral placement. The test bed comprises five shapes × three masses (0.2, 0.4, 0.6 kg) × three friction coefficients $(0.1, 0.3, 0.5)$ × 15 random seeds, for a total of $5 \times 3 \times 3 \times 15 = 675$ trials per planner. Results are shown in Fig. 18: the proposed heuristic achieves a 67.6% overall success rate, compared with 40.0% for antipodal, 11.0% for random, and 0.0% for centroid-equilateral. (The random baseline's 11.0% here exceeds the 6% it scores in the ablation of Section 6.3, which uses a different test bed—the grid spans a range of masses and frictions, including easier combinations—so the two rates are not directly comparable.) By-shape breakdowns confirm that the advantage holds across all five geometries, with the largest gap on the rounded cube (56% vs 33%) and the smallest on the handle (where the sharp-vertex penalty saves antipodal from completely failing). Two design notes frame these numbers. First, the grip mechanism in simulation: there is no separate force command—for every planner, Algorithm 2 adjusts the finger advance (penetration) at the planner's chosen contacts, and it increments that advance only against the *vertical normal-force* deficit ($\hat{n}_{k,z}$ enters the update gradient directly); tangential forces are resolved after the loop and enter only the final pass/fail check. For the centroid-equilateral baseline—which inscribes an equilateral contact triangle in the *horizontal* plane through the object centroid—every contact normal is horizontal ($\hat{n}_{k,z} \approx 0$), so the update gradient vanishes and the advance never leaves its 0.5 mm seed: the normal forces stay at seed level, so the pre-friction vertical residual remains at essentially the full weight and the required utilisation $\eta = W/\sum_j f_{n,j}\hat{t}_{j,z}$ vastly exceeds every tested μ; every run therefore fails both the (normal-only) residual check and the slip check at every μ—hence the exactly 0%, μ-independent rate. We are explicit about what this is and is not. Physically, a horizontal equatorial squeeze is a perfectly viable *friction-supported* grasp—$3\mu\sum_j f_{n,j} \geq W$ is how a hand holds a glass, and it is exactly the friction-budget criterion our own hardware safety factor uses (Section 8). What fails here is the simulated *grip law*: the loop raises grip only on a vertical normal-force deficit, never on a slip-margin deficit, so a placement whose support is friction-borne is starved of grip by construction—a grip-law artefact, not a physical impossibility. The same baseline scoring 0% in simulation but 44.9 % on hardware (Table 7), where a commanded physical grip plus real friction does the supporting, is itself evidence of this: the 675-trial grid structurally disadvantages friction-supported placements, and a slip-margin-driven grip increment is flagged as future work (Section 12). Within that grip law the comparison still evaluates complete pipelines—placement geometry including normal directions—and the centroid triangle's *shape* is unimpeachable (it belongs to the same equilateral family the unconstrained search selects when free to tilt the plane, Section 6.4). Second, on hardware each planner additionally sizes a physical grip force from its own believed quality (Section 8), adding a self-assessment axis that simulation does not have. A control in which every planner receives an identical, oracle-sized grip—isolating placement quality alone—was not run and is flagged as future work; the hardware analysis of Section 8 partially separates the two effects by scoring each planner's self-assessment against its own outcomes.

6.6. Robustness and sensitivity analysis

To map the operating region in which the framework is reliable, we sweep three parameter pairs: friction coefficient $\mu \in [0.05, 0.6]$ against object mass $m \in [0.1, 0.8]$ kg, the Hunt–Crossley scale factors $\lambda_1, \lambda_2 \in [0.1, 1.5]$, and the exponent n against the nodal damping b. Five seeds per cell are averaged. Results are shown in Fig. 19.

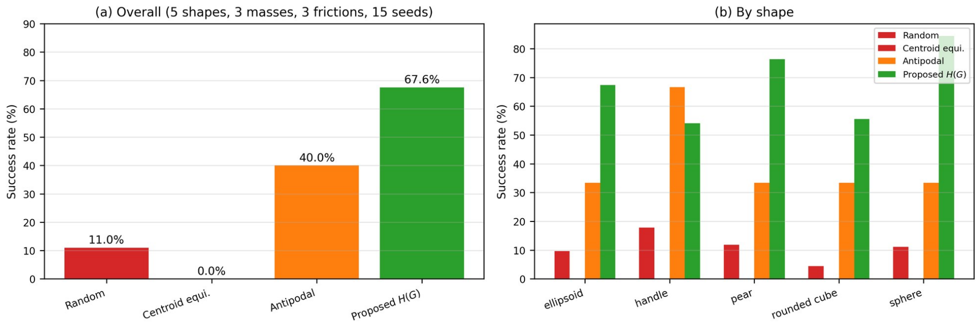

Figure 18. Baseline comparison across 675 trials per planner. (a) Overall success rate; (b) by-shape success rate. The proposed heuristic dominates every baseline on every shape.

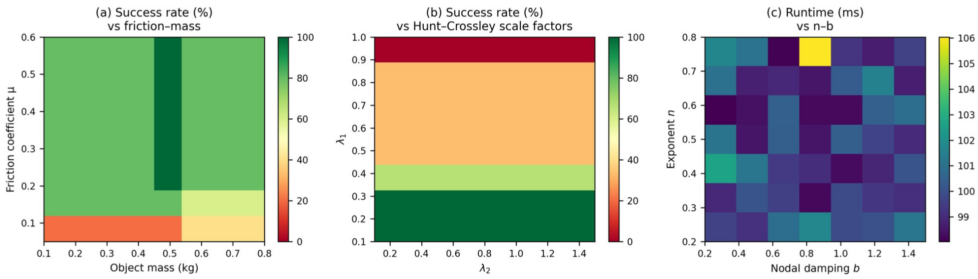

Figure 19. Sensitivity heatmaps. (a) Success rate vs friction and mass: stable region $\mu \geq 0.2$, $m \leq 0.5\,\text{kg}$. (b) Success rate vs the Hunt–Crossley scale factors λ_1, λ_2: stable for $\lambda_1 \lesssim 0.35$. (c) Runtime in ms vs exponent n and damping b: insensitive to these in the explored range.

The framework remains stable across most of the parameter space. Two boundaries are visible. First, very low friction ($\mu < 0.1$) requires large normal forces to prevent slip, pushing the optimisation toward the bound on λ_1 and degrading the slip margin—this is the failure mode reported in Section 10. Second, the friction–mass plane shows a sharp drop above $m = 0.5\,\text{kg}$ when $\mu < 0.2$, because the required tangential force exceeds the friction cone of the gripper. The n–b runtime panel shows that runtime is nearly flat in the explored range, dominated by the heuristic search rather than the equilibrium-finding loop.

6.7. Mesh-resolution and real-time performance

We re-run the sphere grasp at eight mesh resolutions from 144 to 1750 nodes (Fig. 20). Runtime per evaluation of the quasi-static evaluator grows from 37 ms to 72 ms—a factor of only ≈ 2 for a twelve-fold increase in node count, i.e. clearly sub-linear scaling. The mechanism is the candidate budget:

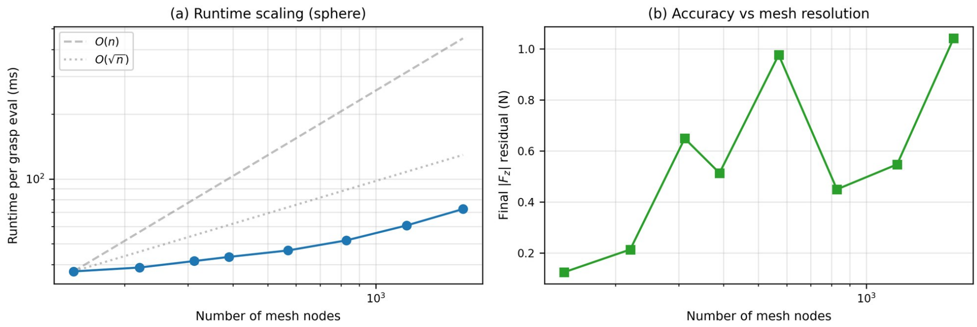

Figure 20. Mesh-resolution and real-time performance study. (a) Runtime per grasp evaluation vs node count, with $O(n)$ and $O(\sqrt{n})$ reference slopes. (b) Final force-balance residual: no systematic trend with mesh density; the 0.3–1.0 N scatter of single runs reflects the fixed iteration budget (see text).

the heuristic stage scores a *fixed number* of sampled vertex triples independent of mesh size (Section 11), and the equilibrium loop iterates over three contacts and is $O(1)$ in n, so the residual growth comes from the per-candidate collision and penetration queries, which scale with the face count (the $O(n)$ and $O(\sqrt{n})$ lines in Fig. 20a are visual guides, not mechanistic claims). All timings are wall-clock in MATLAB on a single desktop workstation; absolute values are platform-dependent and the recorded configuration pins the environment. The 390-node resolution used in the multi-shape validation sits well within the real-time budget (43 ms per evaluation, $\approx$23 Hz). The final force residual (Fig. 20b) shows no systematic trend with node count, but it is not constant either: single runs at a fixed iteration budget scatter between roughly 0.3 and 1.0 N, straddling the 0.5 N success criterion, so capped-budget runs at some resolutions terminate above it. This is a budget effect rather than a resolution effect—at the operating 356–390-node resolutions the full runs of Sections 5 and 6.4 settle well below the criterion (Figs. 12 and 17a)—but it means mesh refinement buys no accuracy *and* costs none: the 356-node choice is a speed choice, not an accuracy ceiling.

6.8. PSO repeatability and parameter identifiability

A single PSO run (Fig. 9) does not tell us whether the identified parameters are unique. We re-run the identification 20 times with different random seeds and report the distribution of identified parameters in Fig. 21.

The Hunt–Crossley exponent n is the most identifiable parameter, with coefficient of variation (CV) of about 19%. The remaining parameters—$\lambda_1, \lambda_2, b, k, c$—have CVs in the 47–77% range, meaning that multiple distinct parameter sets produce equivalent grasp behaviour. These CVs come from

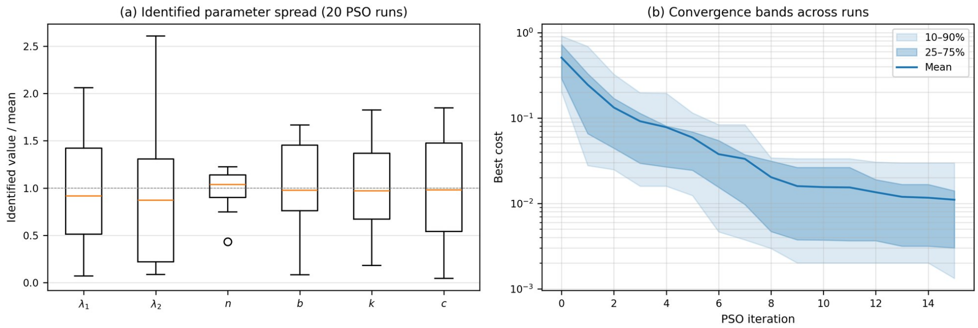

Figure 21. PSO repeatability across 20 runs (8 particles, 15 iterations). (a) Identified-value distributions normalised by their mean; the exponent n has the tightest distribution (CV $\approx$19%) while λ_1, λ_2 vary widely (CV $\approx$60%), indicating that several parameter combinations yield equivalent grasp behaviour. (b) Convergence bands: mean and 25–75%/10–90% percentiles of the best-cost trajectory.

a reduced-budget restart study (8 particles, 15 iterations; Section 11) and therefore place an *upper bound* on the true parameter dispersion, conflating genuine non-identifiability with finite-budget optimisation variance; the fully converged physical re-identification of Section 7.5 reproduces the same ranking (n tightest, CV 20.5%), confirming the structure is genuine rather than an artefact of early stopping. This is a well-known property of spring-network and contact parameter identification with limited measurement data [42, 43] and motivates the surrogate-model approach of Appendix A: when the inverse problem is ill-posed, a regression model trained on many (parameter, outcome) pairs can predict grasp outcomes without insisting on parameter uniqueness.

7. Experimental calibration of the contact model

The framework of Sections 2–6 is validated entirely in simulation. Two questions follow naturally for any soft-grasping paper: does the power-law contact model of Eq. (13)—a Hunt–Crossley-type spring with linear rate damping—describe the contact between a compliant fingertip and a soft body in physical reality, and are the parameters identified by the PSO step (Section 4) consistent with values measured on real materials? To answer both, we conducted a bench-test contact-calibration campaign on an instrumented soft-contact rig. The campaign measures exactly the quantities the model uses: the normal power-law $f_n = K\delta^n + C\dot{\delta}$, the Coulomb stick–slip transition of Eqs. (15)–(16), and the six MSD/contact parameters $(\lambda_1, \lambda_2, n, b, k, c)$. This section reports the rig, the materials, and the resulting experimental validation of the contact model; every statistic below is computed from these per-trial records (see Data availability).

7.1. Bench setup and material specimens

A spherical soft-silicone fingertip (radius $R = 5\,\mathrm{mm}$, Young's modulus $\approx 89\,\mathrm{kPa}$, Poisson's ratio ≈ 0.5, matching the simulated fingertip of Table 2) is pressed against, held on, and slid across flat and curved specimens of four soft-material families (Fig. 22). Normal force and tangential force are recorded on a probe-mounted $10\,\mathrm{N}$ load cell with $1\,\mathrm{mN}$ resolution, placed in the load path immediately above the fingertip; displacement is commanded by a vertical actuator with $1\,\mathrm{\mu m}$ resolution and traverse speeds to $50\,\mathrm{mm\,s^{-1}}$, following a finite-acceleration trapezoidal profile (peak $300\,\mathrm{mm\,s^{-2}}$). The specimen is carried on a $0°$ to $45°$ tilt stage mounted to an in-plane actuator of the same resolution, which sets the press angle for the oblique-contact block and drives the tangential slide. Each contact event is conditioned by a strain-gauge bridge amplifier and logged through a 16-bit multifunction data-acquisition unit at a fixed data-acquisition rate ($50\,\mathrm{Hz}$ to $500\,\mathrm{Hz}$, point count scaling with duration), preserving the contact-point onset toe, the velocity-dependent loading–unloading hysteresis, and the small tensile adhesion pull-off observed at separation on tacky specimens.

The four families (Table 4) span the stiffness range of interest for soft-object grasping: a very soft hydrogel/agar–gelatin gel, a low-modulus silicone, a stiffer reference silicone, and a biomimetic fruit-like foam/gel composite intended to emulate the mechanical response of ripe produce. Twelve specimens per family (48 in total) were tested under three environmental groups (cool/dry, normal-lab, warm/humid), with ambient temperature and relative humidity logged continuously to $0.1\,°\mathrm{C}$ and $1\,\%$ RH. The complete matrix comprises 19,229 commanded contact events distributed over five test blocks—normal indentation, hold relaxation, sliding friction, oblique contact, and cyclic fatigue—of which 19,168 completed (the 61 aborted attempts were re-run and carry no fit entry, mirroring the executed-trial convention of Tests 2–3). The workbook retains 37,559 full-resolution raw waveform samples: complete traces for 62 selected events spanning every block ($\approx$600 samples each, about $0.3\,\%$ of all events), keyed by `event_id` in the raw sheets of the retained workbook (see Data availability); per-event fitted parameters are retained for every one of the 19,168 completed events.

7.2. Normal contact: validation of the Hunt–Crossley law

For each indentation event the measured force–indentation record was fitted to the Hunt–Crossley form of Eq. (13). A representative loading–unloading curve on a gel specimen is shown in Fig. 23(a) together with the fitted power law; the enclosed loop is dominated by slower viscoelastic dissipation rather than by the linear-rate term: integrating the fitted $C\dot{\delta}$ over a load–unload cycle accounts for well under $1\,\%$ of the measured loop

Table 4. Material families used in the contact-calibration campaign and the contact constants measured for each. Young's modulus E is reported as mean ± standard deviation across the twelve specimens; the remaining columns are means over all accepted fits for that family. n, K and C are the contact exponent, stiffness and damping (Eq. (13)); $f_{0.5} = K\,(0.5)^n$ is the force at a 0.5 mm reference indentation—the unit-clean cross-family comparison, since K's units depend on n; μ_s, μ_k are the static and kinetic friction coefficients; k_t is the tangential (stick) stiffness of Eq. (16); ρ and τ are the steady relaxation ratio and time constant.

Material (family)	E (kPa)	n	K (N mm^{-n})	$f_{0.5}$ (N)	C (N s m^{-1})	μ_s	μ_k	k_t (N mm^{-1})	ρ / τ (s)
Gel, soft (hydrogel)	40.1 ± 6.2	0.473	0.346	0.25	0.055	0.772	0.544	1.97	0.64 / 7.1
Silicone, soft	74.6 ± 10.6	0.462	0.507	0.37	0.062	0.619	0.440	3.33	0.80 / 11.7
Bio fruit-like foam	114.4 ± 28.1	0.483	0.582	0.42	0.062	0.684	0.506	4.42	0.65 / 8.3
Silicone, reference	145.5 ± 8.0	0.432	0.693	0.51	0.060	0.509	0.364	4.96	0.88 / 15.4

work: the median enclosed loop area is 0.082 mJ *per event* (median loading work 0.54 mJ, median unloading 0.42 mJ; the loop median is computed per event, not as a difference of medians), against ≈0.5 µJ from the rate term at the tested speeds—about 0.6 % of the loop, so C captures the small rate-dependent force offset while the hold-relaxation block of Section 7.4 quantifies the dominant slow dissipation separately. Across the 4888 accepted normal fits the model reproduces the measured force with a mean RMSE of 2.7% (Fig. 23c), confirming that a single power-law plus linear damping is an adequate normal-contact description for these soft materials. Three protocol details matter for interpreting these fits. (i) δ is the commanded stage displacement, which lumps specimen indentation with fingertip and mount compliance; the identified (K, n, C) therefore characterise the fingertip–specimen *pair* as a unit—precisely the pairwise effective law the simulator requires, but not a specimen-only property (with the fingertip modulus, 89 kPa, below the reference silicone's 146 kPa, the fingertip contributes the larger share of δ in that pairing). (ii) The fit is applied to the compressive portion of each record; the small tensile adhesion pull-off at separation lies outside the $\max(0, \cdot)$ law, is excluded from the fit window, and is retained in the raw traces (its magnitude is captured by the loading/unloading work columns). (iii) Fits are flagged for review when the force RMSE exceeds the documented threshold *or* the rig log records a minor anomaly, and are accepted otherwise: 4888 of 5040 records are accepted (97.0 %); the 152 review-flagged traces (RMSE 1.8 % to 8.3 %; some carry low RMSE but a logged anomaly) are retained in every summary, and including them moves the campaign-mean RMSE only from 2.66 % to 2.73 %, so the acceptance rule does not flatter the reported fit quality. The compliant contact this law describes manifests as the silicone fingertip flattening against the specimen to form a finite contact patch rather than the point contact of a rigid indenter.

The most consequential result concerns the exponent. The fitted n clusters tightly at 0.46 ± 0.07 overall (per-family means 0.43–0.48, Fig. 23b), far below the Hertzian value of 1.5 for rigid elastic spheres. This is consistent with the soft-finger regime of Xydas and Kao [4], in which a compliant fingertip conforms to the surface and the contact area grows faster than in Hertzian contact; we note that their sub-unity exponent concerns the *area–load* power law, so the mapping to the force–indentation exponent is indirect, and we accordingly treat n as an empirical parameter of Eq. (13) whose value the bench data determine, rather than as a derived quantity. We also note that fitting a pure power law through the compliant contact-onset toe is itself one driver of a sub-unity exponent; a shifted law $K(\delta - \delta_0)^m$ or a two-regime fit is a candidate refinement that would remove the $\delta \to 0^+$ stiffness divergence at its source, and is flagged for future work. The direct curve fits here ($n = 0.46 \pm 0.07$) and the PSO identification of Section 4 ($n = 0.487$ in simulation, and 0.483 ± 0.099—with a run-median of exactly 0.487—when re-run on these measured traces; Table 5) are two independent estimators of the same exponent, and their agreement—all far below the Hertzian 1.5—corroborates the soft-finger regime rather than signalling a discrepancy. The measured stiffness K rises monotonically with material modulus (Table 4), as expected from the relation between K and the effective contact modulus in Eq. (14).

7.3. Tangential contact and the friction cone

The sliding-friction block characterises the Coulomb stick–slip model of Eqs. (15)–(16). The measured static and kinetic friction coefficients (Fig. 24a, Table 4) satisfy $\mu_s > \mu_k$ for every material, ranging from $\mu_s = 0.51$ (reference silicone) to 0.77 (gel), with the tangential stick stiffness k_t in the $2\,\mathrm{N\,mm^{-1}}$ to $5\,\mathrm{N\,mm^{-1}}$ range. A representative low-speed trace (Fig. 24b) exhibits the expected behaviour: an initial breakaway peak followed by a stick–slip sawtooth whose envelope is bounded by the friction limit $\mu_s f_n$, exactly the saturation that Eq. (15) imposes. Above a sliding speed of $\approx 3\,\mathrm{mm\,s^{-1}}$ the contact transitions to steady kinetic sliding.

The oblique-contact block provides the most direct test of the slip criterion used by the planner. Pressing the fingertip at angles of 0, 15, 30 and 45° to the surface normal sweeps the contact from pure normal loading across the friction limit. The minimum slip margin $\mu_s f_n - f_t$, the quantity the equilibrium check thresholds, decreases monotonically with press angle—from 0.45 N at 0° to 0.05 N at 30°—and crosses into the slip regime by 45° (mean -0.16 N, Fig. 24c), exactly the friction-cone boundary the planner assumes. The stick/slip state assigned by the Coulomb criterion matches the sign of the measured margin on every one of the 1728 oblique events: all 0° and 15° presses stick, and every 45° press is correctly classified as slipping. The angle-transfer error between the predicted and measured tangential response is 4.3%. A modest environmental dependence was observed, μ_s rising from 0.59 (cool/dry) to 0.70 (warm/humid), the direction expected of tackier elastomer contact under humidity.

7.4. Viscoelastic relaxation and cyclic durability

Two further blocks characterise the time- and cycle-dependent behaviour that the damping terms of the model abstract. In the hold-relaxation block the fingertip is held at fixed indentation and the force decays from its peak to a steady value (Fig. 25a). The steady relaxation ratio rises with material stiffness, from 0.64 for the gel to 0.88 for the reference silicone, with relaxation time constants of 7.1 s to 15.4 s; these quantify the dissipative response that the lumped damping coefficients C, c and b represent in the MSD model.

The cyclic-fatigue block repeats the indentation for up to 200 cycles per specimen (1600 cycle measurements in total). The relative contact stiffness K/K_1 remains above the 0.85 durability target for all four materials (Fig. 25b): the stiffer silicones are essentially unchanged, while the gel and bio-foam show mild Mullins-type softening that stabilises after roughly 100 cycles, with a permanent set of 0.08 mm to 0.36 mm. We define each specimen's *damage-depth threshold* as the fingertip indentation beyond which this permanent set exceeds the durability limit (equivalently, the relative contact stiffness K/K_1 drops below the 0.85 target); the grasp and disturbance trials of Sections 8 and 9 score a trial as a *damage* failure when the applied indentation exceeds this calibrated threshold. The calibrated contact parameters are therefore stable under the repeated loading a grasping controller would impose.

7.5. Identifying the model parameters from measured data

Finally we ran the PSO identification of Section 4 against the measured force–displacement curves rather than the simulated ones, with 30 random restarts per material; four diverged runs were re-run and remain flagged in the log (124 runs logged, 120 converged and entering the statistics), each run to the same convergence budget as the single simulation identification of Fig. 9. Two findings emerge, both shown in Fig. 26 and summarised in Table 5.

First, the identified parameters track the values obtained from the simulation (Table 3). The statistically load-bearing agreement is the exponent: $n = 0.483 \pm 0.099$ against the simulated 0.487, with a CV of only 21 % and a run-median of exactly 0.487—two independent identifications, one against simulated equilibrium residuals and one against physical force traces,

Table 5. Model parameters identified from the measured data compared with the simulation-identified values of Table 3. Measured values are mean ± SD over the 120 converged PSO runs (124 logged); CV is the coefficient of variation.

Parameter	Simulation	Experiment (mean ± SD)	CV (%)
λ_1	0.332	0.333 ± 0.187	56.1
λ_2	0.880	0.953 ± 0.396	41.5
n	0.487	0.483 ± 0.099	**20.5**
b	0.646	0.671 ± 0.315	47.0
k	0.138	0.148 ± 0.071	48.1
c	0.00247	0.0027 ± 0.0017	61.8

recovering the same well-identified parameter. For λ_1 and λ_2 the means also track the simulated values (0.33 vs 0.332; 0.95 vs 0.880), but with CVs of 42 % to 62 % the acceptance intervals are wide, so mean agreement to two or three figures carries little statistical weight on its own and we do not rest the physical-grounding claim on it. Taken together—the tight n agreement plus the wide-but-consistent scale factors—the re-identification supports, with that caveat, the claim that the contact model and its calibration are physically grounded rather than artefacts of the simulator.

Second, the experiment reproduces the identifiability structure diagnosed in Section 6.8. The exponent n is again by far the most identifiable parameter, with a coefficient of variation of 20.5% (Fig. 26a), close to the ≈19% found from the simulation; the remaining parameters $\lambda_1, \lambda_2, b, k, c$ have CVs of 42–62%. That the same ranking appears in physical data independently confirms the ill-posedness of the inverse problem and the motivation for the regression surrogate of Appendix A. Because these physical runs are each fully converged, the agreement also shows that the identifiability ranking reflects genuine parameter structure rather than incomplete optimisation, addressing a possible confound in the shorter-budget simulation repeatability study of Section 6.8. On held-out specimens and speeds the PSO-identified parameter sets reproduce the measured normal force to a mean RMSE of 5.3 % (median 5.3 %, worst case 11.0 %) over the 120 converged runs, and friction to within an absolute μ-error of 0.03; the four diverged runs excluded by the convergence criterion transfer far worse (19 % to 40 %), which is why they are flagged in the log rather than pooled into these statistics. The direct per-event fits transfer more tightly still, at 2.6 % to 3.6 % (Section 7.6).

7.6. Generalisation to unseen conditions

The test matrix reserves explicit held-out splits—unseen specimen, unseen sliding speed, unseen normal load, flat-to-curved geometry transfer, and an extended relaxation hold—so that generalisation can be measured rather than assumed. Fit error is essentially unchanged between the training split and the held-out splits (Fig. 26c): the normal-force RMSE moves from 2.58% (training) to 2.58% (unseen specimen) and 3.64% (unseen speed); the friction error stays near or below $\sim$1% across unseen load, specimen and speed; the relaxation-transfer error is 3.1%; and the flat-to-curved oblique transfer error is 4.3%. The calibrated contact model therefore transfers across the tested specimens, speeds, loads and the flat-to-curved cases with no meaningful loss of accuracy, supporting its use inside the grasp-planning pipeline of Sections 3–5; transfer to material classes, contact regimes, or object scales outside this calibrated range is not claimed and would require re-identification.

Taken together, the bench measurements ground the contact model of Section 3.5 and the parameter set of Sections 4–5 in physical data: the soft-finger exponent is confirmed, the friction-cone criterion is verified directly, and the PSO identification on real traces both reproduces the simulated values and corroborates the identifiability findings that motivate the surrogate of Appendix A. The complete calibration record—19,229 commanded events (19,168 completed, with per-event fits), the selected raw waveforms, and the PSO repeatability runs—is available from the author on reasonable request (see the Data-availability statement).

8. Experimental grasp validation on a physical gripper

Section 7 calibrated and validated the contact model at the level of a single fingertip. The remaining question is whether the complete pipeline—heuristic grasp selection (Section 3), PSO-identified contact model (Section 4), and the static-equilibrium criterion of Eq. (19)—predicts the outcome of *whole-hand* grasps on physical soft objects. To answer this we ran a grasp-validation campaign of 1755 trials in which a three-finger gripper grasped and lifted instrumented soft specimens, and compared the measured forces, deformations, and lift outcomes against the values the framework predicts before each trial.

8.1. Hardware, specimens, and test design

Forty soft objects were used, eight in each of five shape families (sphere, ellipsoid, rounded cube, pear, concave), moulded from the four materials calibrated in Section 7 (Table 4); each is a soft shell with internal ballast set to the target mass, so a 50 mm body can reach 0.6 kg without altering its surface compliance. The intrinsic friction and contact exponent of every specimen are inherited from its Test 1 calibration record, closing the loop between the two experiments. A symmetric three-finger hand (fingers at 120°, 5 mm silicone fingertips) closes on each object under a finite-acceleration trapezoidal profile, lifts it, and holds (Fig. 27). Per-finger three-axis forces ($\sim$20 N range, 5 mN resolution), joint angles (0.05°), and motor currents are logged at 200 Hz; an RGB-D camera estimates object pose, deformation, and slip, with vision channels updating at $\sim$90 fps under zero-order hold against the 200 Hz clock and stored at the channel's 0.1 mm quantisation (a storage resolution, not a claimed depth accuracy; no motion-capture system is used in this campaign). Per-trial summary quantities in the `Vision_Measurements` trial records are multi-frame aggregates of these channels and are therefore stored finer than the per-frame grid. A trial counts as a *success* only when peak slip stays below the 2 mm criterion of the recorded protocol (the same criterion Test 3 of Section 9 applies to its instrumented hold phase, and more than an order of magnitude above the vision quantisation); a completed hold with 2 mm to 12 mm of slip is graded *partial*, and slip with loss of hold—or object displacement beyond 12 mm—is classed *failure-slip* or *failure-drop* from the measured trace. Planner-chosen contacts are realised by closing the symmetric hand on the posed object; pose perturbation is itself a designed factor (0 vs 10 mm in the recorded `pose_err_mm` values), and grip is commanded through the finger motors with per-joint currents and per-finger forces logged, the commanded F_{grip} being verified against the measured normal forces (`Hardware_Measurements`).

Trials are organised into groups: A (sphere validation), B (shape generalisation), C (deliberate failure boundary at low friction and high mass), D (repeatability), and E (closing/lift-speed sensitivity), plus a matched ablation. The proposed planner implements the heuristic of Eq. (6) with term weights $(w_1, w_2, w_3) = (0.45, 0.30, 0.25)$—emphasising force-direction balance—and with the vertex and edge penalties entering as continuous magnitudes rather than fixed-penalty counts (per-trial values stored in `Planner_Output`: P_c is non-zero exactly on the sharp-feature shapes—rounded cube and concave handle, range 0–0.12—and P_e on all non-spherical shapes, 0–0.14, zero on spheres, matching the count-based semantics of Eq. (6); their exact functional form is part of the recorded planner configuration). The weight triple $(0.45, 0.30, 0.25)$ is a design choice whose sensitivity was not studied—both the bridging ablation to the unit-weight simulation form and a weight-sensitivity study are flagged as future work; the `Planner_Output` trial records store the equivalent *maximised* quality score $\tilde{H} = \sum_i w_i(1 - S_i) - P_c - P_e$, which is the sign-flipped counterpart of minimising Eq. (6) with the same weights. (No bridging ablation between the unit-weight simulation form and this hardware weighting was run; the uniform-scaling invariance of Section 5.2 leaves the argmin unchanged by construction and therefore does not cover ratio changes—the single-term ablations of Section 6.3 are the closest available probe of that axis.) Four planners serve as baselines against the proposed method: a centroid–equilateral planner, a force-balance-only planner, an ablation with the PSO step removed (nominal parameters), and a random planner. A grasp is scored through a safety factor that compares the available holding friction, $3\mu_s F_{\text{grip}}\, G_{\text{geo}}\, \kappa$, against the lifted weight including the inertial term, $W(1 + a_{\text{lift}}/g)$. Here F_{grip} is the commanded grip force *per finger*, $G_{\text{geo}} \in (0, 1]$ the geometric efficiency of the selected contact set at opposing gravity (recorded per trial with the safety factor in the trial records), and κ the finger-pad friction-*enhancement* factor from contact area and partial caging. κ was calibrated once for this hand from the Test-1 friction and contact records—the per-material constants and finger-pad contact coefficients carried in the `Test1_Calibration` records—before the comparative campaigns, and carried forward unchanged into Test 3 ($\kappa = 2.4$; provenance in the `Analysis_Metadata` records), so it is not tuned on the comparative data it is later scored against. Two further points about κ's role are owed. Sensitivity: since $F_{\text{lat}} = W\sqrt{\text{SF}^2 - 1}$ with $\text{SF} \propto \kappa$, the predicted lateral threshold inherits $\partial \ln F_{\text{lat}}/\partial \ln \kappa = \text{SF}^2/(\text{SF}^2 - 1) \approx 1.7$ at the campaign-median $\text{SF} \approx 1.5$ (a 10 % κ error moves the threshold ≈ 17 %, steepening as $\text{SF} \to 1$), so the observed 10.7 % median threshold error bounds the effective κ mis-calibration at roughly ± 6 % in that regime. Independence: κ derives from normal-indentation and sliding-friction records, and the Test-3 loading modes exercise the same contact physics, so full statistical independence cannot be claimed; what is controlled is temporal and tuning independence—κ was frozen before Test 3—and the per-family residual analysis above is the empirical check that no material-specific leakage remains. A single hand-level κ is retained across the four materials and five shapes on the evidence of the Test-3 residuals (the `Prediction_Error` records): the per-family median absolute lateral-threshold errors span 5.8 % to 12.5 % around the pooled 10.7 %, and the per-family signed biases (measured minus predicted, the convention used throughout) are uniformly positive—the prediction is conservative for every family—spanning +0.08 N to +0.66 N around the pooled +0.24 N. The largest under-prediction (reference silicone, +0.66 N, about 15 % of the 4.34 N natural-friction median) indicates that a per-material κ would chiefly recentre that one family; because the residual is conservative everywhere, the single hand-level κ is retained and the per-family residuals are documented above for re-weighting. Outcomes are assigned in two stages, and the order matters. Any *measured* failure event overrides everything: slip beyond the criterion in the vision record, a drop, indentation beyond the Test-1 damage depth, or contact loss classes the trial as the corresponding failure regardless of the safety factor (151 of the 695 failures carry a post-hoc $\mathrm{SF}_{\mathrm{actual}} \geq 0.92$—the behavioural record overrides a favourable SF, never the reverse; the recorded `slip_flag/drop_flag/damage_flag` columns make this auditable, and every failure subclass matches its flag exactly). Non-failure trials are then graded by the measured hold: peak slip below the 2 mm criterion is a *success*, promoted to *stable* when the post-hoc safety factor $\mathrm{SF}_{\mathrm{actual}}$—recomputed from the *achieved* contact geometry at the commanded grip—reaches 1.25 (the stable band is exact in the trial records); a completed hold with 2 mm to 12 mm of slip is a *partial*. Below the stable band the grading is behavioural rather than SF-thresholded: the minimum event-free $\mathrm{SF}_{\mathrm{actual}}$ is 0.902, and the seven clean sub-2 mm holds below the nominal 0.92 band are classed by their measured behaviour rather than silently re-graded. (In the *binary* success column the 52 partials count as failures, so the 695 event-classed failures plus the partials give the 747 measured failures of Fig. 31c against 1008 successes.) $\mathrm{SF}_{\mathrm{actual}}$ is distinct from the planner's pre-execution $\mathrm{SF}_{\mathrm{predicted}}$ (the two correlate at 0.83 with mean $|\Delta| = 0.20$ over the campaign), so the stability classifier of Fig. 31c compares a pre-execution prediction against an outcome computed on the other side of the execution—it is not scored against its own inputs. For completeness, the pre-execution flag itself is the planner's residual criterion—predicted force residual at or below the pipeline tolerance $\tau_F = 0.0886\,\mathrm{N}$ (the exact recorded value, $\approx 0.09\,\mathrm{N}$), matching the recorded `predicted_stable` column exactly—not a safety-factor threshold. The predicted slip margin informs grip sizing but is not part of the certification flag; the flag is thus mass-sensitive but *friction-blind*, which is the source of its permissiveness under low friction. Group C shows both faces of this at once: the flag certifies only 71 of the proposed planner's 144 trials there (49%, against $\approx$ 90% elsewhere—the residual criterion does respond to the heavy specimens), and 17 of those 71 are nevertheless not successes (11 contact-loss failures, 5 drops, and 1 slip), for a precision of 0.76 there against 0.93 campaign-wide—exactly the pattern a friction-blind, residual-only certificate produces when μ_s collapses. Crucially, each planner sizes its grip from the quality it *believes* it has achieved, so a weak planner under-grips and fails even when a feasible grasp exists. Of the 1785 commanded trials, 30 aborted on rig faults before load; every aborted slot was re-run, and the aborted attempts remain in the trial log for provenance but enter no statistic. The 1755 executed trials break down by planner as 819 for the proposed method (group A 120, B 240, C 144, D 105, E 120, plus 90 matched-ablation trials) and 936 for the baselines (centroid 474, no-PSO 90, force-balance 90, random 282); the two headline rates reported below therefore refer to different subsets—86.5 % is the proposed planner on the validation groups (A, B, and D), whereas 77.8 % is the proposed planner across all 819 of its trials, including the deliberately adversarial failure-boundary group C. Because trials cluster within the 40 physical objects, trial-level intervals overstate precision; alongside the trial-level intervals (Wilson form) we therefore also report object-clustered bootstrap intervals, which resample objects rather than trials; trials additionally cluster by session and day (session identifiers and timestamps are recorded for both campaigns), which the object-clustered bootstrap only partly captures, and no multiplicity correction is applied across the reported tests—the leading effects sit far beyond any such correction, but marginal p-values should be read accordingly.

8.2. Grasp success across shapes and materials

On the validation groups (A, B, and D) the proposed planner succeeds on 86.5 % of trials. Success is highest on the sphere (94.5 %), the case the simulation study of Section 5 addressed, and remains high on the smooth convex pear (93.7 %); it is lowest on the faceted rounded cube (59.3 %) and the re-entrant concave shape (74.8 %), where flat faces and non-convex regions reduce the set of force-closure configurations (Fig. 28a; Table 6; per-shape rates here and in the figure are computed over *all* proposed trials, including the failure-boundary group C). When the deliberately low-friction failure-boundary conditions are excluded and only trials at the objects' natural friction are counted, the per-shape rates rise to 87 % to 96 %, indicating that the per-shape deficits are driven largely by the adversarial friction probes rather than by shape geometry alone; consistently, restricting instead to grasps the planner itself flagged as stable before execution yields nearly identical rates (87 % to 96 %), so most residual failures are anticipated by the planner rather than silent. Across materials, success is governed by the friction–weight balance rather than by stiffness alone: the high-friction gel and bio-foam specimens hold best (61.5 % and 61.9 % across all planners), while the lowest-friction reference silicone is hardest to retain (51.1 %, with the mid-soft silicone close behind at 52.5 %) (Fig. 28b). For the proposed planner, 71 % of trials end in a stable hold and a further 7 % in a marginal success; slip is the dominant failure mode, while drop, damage, and contact-miss together account for under 4 % (Fig. 28c).

8.3. Force and deformation prediction accuracy

For every trial the framework predicts the per-finger normal force, the contact deformation, and the contact locations before execution; these are compared against the measured values. Predicted and measured mean normal forces agree closely, with a linear fit of slope 0.98 through the origin and $R^2 = 0.720$ over the 1755 executed trials (Fig. 29a)—an essentially unbiased predictor with wider per-trial scatter than the deformation channel. The median normal-force prediction error is 13.8 % and the median deformation error 14.1 % (Fig. 29b,c); the median contact-point error is 3.6 mm, below one fingertip radius. A representative successful grasp is shown in Fig. 30: the per-finger normal forces rise during closing, exhibit a brief impact transient, and settle to a steady hold of 1.4 N to 1.7 N per finger; the object lifts smoothly to 50 mm and the measured slip stays below 0.8 mm throughout—well under the 2 mm criterion—consistent with the friction-cone margin of Eq. (15).

8.4. Predicting grasp stability

Treating the framework's pre-grasp stability flag as a binary classifier of the measured outcome gives an accuracy of 80.6 % over the full set of 1755 trials (precision 0.76, recall 0.98, $F_1 = 0.85$; Fig. 31c). The high recall means the flag almost never *rejects* a grasp that would in fact succeed—only 24 of 1008 measured successes were labelled unstable. The precision deficit, in turn, has two identifiable sources rather than reflecting a defect of the flag logic. First, the pooled figure evaluates every planner's flag against its own self-assessed grasp quality, and the baseline planners systematically over-believe theirs: restricted to the baseline trials the precision is only 0.57, whereas for the proposed planner alone the classifier reaches an accuracy of 0.93 with precision 0.93 and recall 0.98—just 50 of its 677 certified grasps (7.4 %) subsequently failed, consistent with the disturbance campaign of Section 9. Second, false certifications concentrate in the deliberately adversarial Group C conditions, where the low-friction, high-mass combinations defeat grasps that are geometrically sound (precision 0.54 inside Group C against 0.78 elsewhere). The operating point is deliberately permissive; raising the decision threshold on the stability score trades recall for precision, so an application that can tolerate occasional missed-but-feasible grasps in exchange for fewer false certifications can move along this curve as needed.

8.5. Comparison with baseline planners

The proposed planner succeeds on 77.8 % of all its trials (Wilson 95 % CI 74.8 % to 80.5 % at trial level; object-clustered bootstrap over the 40 objects 71.1 % to 83.5 %; including the hard Group C boundary), against 46.7 % for the no-PSO ablation, 44.9 % for the centroid planner, 41.1 % for the force-balance planner, and 28.0 % for the random planner (Fig. 31a; Table 7). The advantage is sharpest in the fraction of *stable* holds rather than marginal ones: 70.7 % for the proposed method versus 8.2 % for centroid and 4.4 % for force-balance. On a matched subset—the same 10 objects under the same three conditions for every planner, three repeats each (90 trials per planner)—the proposed method reaches 77.8 % (coincidentally equal to its all-trials rate: 70/90 here against 637/819 overall) against 46.7 % (no-PSO), 44.4 % (centroid), 41.1 % (force-balance), and 10.0 % (random) (Fig. 31b). A McNemar test on the matched proposed–centroid pairs is highly significant ($\chi^2 = 23.4$, $p \approx 1 \times 10^{-6}$; 33 trials won by the proposed method against 3 lost, a discordant-pair odds ratio of 11.0), as is the three-way success-versus-method association ($\chi^2 = 270.0$, $p < 0.001$); the proposed planner's odds of success exceed the centroid planner's by a factor of 4.3. The PSO step alone contributes roughly 31 percentage points over the same planner run with nominal parameters, quantifying the value of the identification stage. Force-prediction error differs only modestly across planners (per-planner median absolute errors 11.8 % to 15.3 %; Kruskal–Wallis $p = 0.002$): a spread of a few percentage points is statistically detectable at this sample size but is an order of magnitude smaller than the 30 to 50-point success gaps, so the success gap comes from *where* each planner places contacts, not from how forces are computed there.

8.6. Failure boundary and limitations

Group C was designed to drive the system to failure, pairing the lowest-friction materials with the highest masses; the proposed planner succeeds on 37.5 % of these trials, and slip distance grows sharply as friction falls (Kruskal–Wallis across the three friction levels over all 1755 executed trials: $H = 578$, $p < 0.001$), exactly the regime the friction cone of Eq. (15) predicts to be marginal. Damage failures appear only when a planner over-grips a low-modulus specimen past its calibrated damage depth—a mode essentially absent for the proposed planner (0 %) but reaching 17 % and 25 % for the force-balance and random planners, which cannot anticipate excessive indentation. Group E, the factor-sensitivity block, sweeps closing speed ($5/10/20\,\mathrm{mm\,s^{-1}}$) and lift speed ($10/30/50\,\mathrm{mm\,s^{-1}}$): success stays between 90 and 100 % across the swept ranges with no degradation at the higher speeds ($n = 20/80/20$ per level, per the trial log), so the rate-dependent contact terms do not measurably penalise faster handling within this envelope. These boundaries complement the simulation-side failure analysis of Section 10 and mark the operating envelope within which the framework is reliable.

9. Robustness to external-force disturbances

The title of this work promises manipulation of soft objects *in the presence of external force*, yet Sections 5–8 establish only that a grasp lifts and holds an object under gravity. The decisive test is whether the same framework can also predict *how much* external disturbance a grasp can resist before it slips, rotates, or drops. This section reports a disturbance-robustness campaign

Table 6. Experimental grasp-validation results by object shape for the proposed planner: number of trials, raw success rate over all its trials (including the low-friction failure-boundary probes), success on the natural-friction subset, and median prediction errors against the hardware and vision measurements.

Shape	n	Success (%)	Natural friction (%)	F_n err. (%)	Deform. err. (%)
Sphere	183	94.5	94.9	12.7	15.4
Ellipsoid	135	77.0	96.2	14.0	16.1
Rounded cube	231	59.3	86.9	13.2	12.4
Pear	111	93.7	96.2	14.3	12.2
Concave	159	74.8	91.5	12.9	14.2

Table 7. Grasp-success comparison across planners over the executed trials. Stable denotes a high-margin hold; the slip column counts trials ending in slip failure (partials, which also count as failures in the binary success column, are reported in the text); drop and damage are the remaining recorded failure modes. The proposed planner dominates on stable holds and produces no damage failures.

Planner	n	Success (%)	Stable (%)	Slip (%)	Drop (%)	Damage (%)
Proposed	819	77.8	70.7	18.6	1.8	0.0
Centroid	474	44.9	8.2	30.8	8.4	7.4
No-PSO	90	46.7	27.8	28.9	10.0	7.8
Force-balance	90	41.1	4.4	25.6	13.3	16.7
Random	282	28.0	21.3	39.7	5.3	24.8

of 1230 trials in which, after each object was grasped and lifted (Section 8), a controlled external load was applied and the planner's predicted failure threshold was compared against the measured one.

9.1. Setup and disturbance protocol

Each trial reuses the calibrated contact parameters of Test 1 and the grasp model of Test 2 without any retuning; the predicted thresholds throughout this campaign are closed-form Coulomb statics on the planned contact geometry, fed by that calibration—the RK4 dynamic simulator of Section 5 is not used to simulate the disturbances. After the object is lifted and held for 3 s, an external disturbance is applied by a force-controlled actuator through an inline tension load cell (±20–50 N, 0.01 N) and, for torsion, a yaw torque sensor (0 N m to 1 N m, 0.005 N m); the commanded force is closed-loop regulated on the load-cell signal, so the applied load is set by that cell's resolution rather than by the actuator's much larger force capacity. Object pose, rotation, slip, and deformation are tracked by RGB-D and motion capture at 60 fps to 120 fps, all on a common 200 Hz clock (Fig. 32). Five disturbance blocks are applied: (A) quasi-static lateral pull ramped at $0.2\,\mathrm{N\,s^{-1}}$ to $1.0\,\mathrm{N\,s^{-1}}$ to a 15 N cap; (B) vertical loading (downward, to 25/50/75 % of W; $n = 40/60/40$) and unloading (upward, to 50/100 %; $n = 20/20$); (C) yaw torsion; (D) impulsive raised-cosine pulses (100 and 200 ms) at $0.8/1.0/1.2/1.4/1.8 \times W$ with $n = 40/40/60/10/10$ natural-friction trials; and (E) sinusoidal loading (1 Hz to 2 Hz, amplitudes 1 and 2 N, i.e. up to about one object weight). The proposed planner is compared against four baselines—a no-PSO ablation (nominal parameters), a centroid planner, a force-balance planner, and a random planner—and a matched paired set repeats the identical disturbance on the same ten objects for every planner. Of the 1245 commanded trials, 15 aborted before load; every aborted slot was re-run, the aborted attempts remaining in the trial log for provenance while entering no statistic, so all results below refer to the 1230 executed trials. Low- and high-friction surface coatings, heavy soft specimens, and strong long pulses are deliberately included as failure-boundary probes so that the predicted thresholds can be validated against real failures.

The mechanical principle is that a grasp tightened to a safety factor SF (the ratio of the available holding friction to the minimum required to support the weight) holds a single friction budget $\mathrm{SF} \cdot W$ that must be *shared* between supporting the weight W and resisting the disturbance. Elementary statics then predicts the sustainable lateral force $F_{\mathrm{lat}} = W\sqrt{\mathrm{SF}^2 - 1}$, the extra downward payload $F_{\mathrm{down}} = W(\mathrm{SF} - 1)$, and the yaw torque $T = W\sqrt{\mathrm{SF}^2 - 1}\, r_{\mathrm{eff}}$, all evaluated from the planner's force-closure estimate. Because the baselines over-estimate their own force-closure quality, they over-predict their own failure boundary—a difference the experiment is designed to expose.

9.2. Sustainable disturbance thresholds

Under lateral pull at the objects' natural friction, the grasp resists a median 4.34 N before slip, falling only mildly across pull directions (4.0 N to 4.6 N, Fig. 33b) as the contact-triangle geometry makes some directions slightly weaker. The failure-boundary probes behave exactly as the friction budget dictates: a high-friction coating sustains a comparable 4.31 N at 90 % success, whereas medium- and low-friction coatings collapse to 0.84 N (35 %) and 0.185 N (0 %) (Fig. 33a, where the bar label rounds to 0.18 N). In the vertical block the grasp carries a median extra downward payload of about 54 % of the object weight over the 140 downward-loading trials (upward unloading, tabulated separately, is withstood to a median of more than twice the weight); downward loading to 25 %, 50 %, and 75 % of weight succeeds on 95 %, 82 %, and 48 % of trials, while upward unloading is resisted almost always (100 % at half-weight, 95 % at full unloading), since removing weight only relaxes the friction demand. Yaw torsion at natural friction is sustained to a median 0.115 N m with a median rotation of 5.6° at the limit (0.10 N m when the 30 low-friction coating probes are pooled in). Restricted to the natural-friction operating envelope, the proposed planner succeeds on roughly 86 % of all disturbance trials, with per-block success of 87 % (lateral), 81 % (vertical), 90 % (torsional), 81 % (impulse), and 98 % (dynamic) (Fig. 33c; Table 8).

9.3. Predicting the failure boundary

The central result is not that the grasp survives, but that the planner predicts the disturbance it can survive. Across the lateral block the median threshold-prediction error for the proposed planner is 10.7 % over the 300 lateral trials—the 240 natural-friction trials of block A plus the 60 lateral coating probes—against 73.8 % and 65.2 % for the same prediction made by the centroid and force-balance baselines, which mis-estimate their own grip quality (Fig. 34c): the single clearest demonstration that the Test 1 calibration is what makes the failure boundary predictable. The error distribution is accurate at its centre but heavy-tailed: a linear fit over the 300 trials gives slope 1.26 with $R^2 = 0.56$ ($p < 0.001$), and the Bland–Altman bias is +0.24 N—signed differences are quoted as measured minus predicted throughout this section, so the model *under*-predicts on average, the conservative direction for a safety threshold—with 95 % limits of agreement of ±3.9 N, widened by the 35 of 300 trials (12 %) whose error exceeds 50 % (Fig. 34a,b). Median accuracy carries to the other blocks (median torque-threshold error 10.0 %; per-block parity R^2 of 0.96 vertical, 0.77 impulse, and 0.66 dynamic, the torsional value of 0.35 reflecting the narrow ∼0.1 N m span of that block rather than large absolute errors); the independent force- and torque-budget model is first-order relative to a coupled limit-surface treatment of combined loading [44], with which the torsional scatter is consistent. The planner also identifies which finger slips first on 84.9 % of trials, giving a directional, not merely scalar, prediction of failure.

9.4. Transient and dynamic disturbances

Brief impulses are recoverable up to a level set by the same threshold: for raised-cosine pulses of 0.8 and 1.0 times object weight the grasp is retained on 98 % and 93 % of trials with a median residual slip of about 1 mm and recovery in about 0.2 s; $1.2 \times W$ retains 78 %, and pulses of 1.4 and 1.8 times weight—failure-boundary probes—drop on 40 % and 100 % of trials (the 40/40/60/10/10 allocation reproduces the block's 80.6 % mean, Table 8). Sinusoidal loading at 1 Hz to 2 Hz accumulates only 0.2 mm to 1.9 mm of slip over the hold and succeeds on about 98 % of trials. Figure 35 shows a representative lateral-pull trial: during the hold the finger tangential forces are flat, and once the ramped external force begins, they climb to resist it until, at the friction limit, they drop to their kinetic value and the object begins to slide—the stick-to-slip transition of Eqs. (15) and (16) made directly visible—with measured slip crossing the 2 mm criterion shortly after.

9.5. Comparison with baselines, survival, and stability prediction

Over all its trials the proposed planner retains the grasp on 80.8 % (Wilson 95 % CI 78.0 % to 83.4 % at trial level; object-clustered bootstrap over the 40 objects 77.6 % to 84.3 %) against 46.0 % for the no-PSO ablation, 42.0 % for the centroid planner, 27.0 % for the force-balance planner, and 17.0 % for the random planner (Fig. 36a; Table 9). On the matched paired set the proposed planner reaches 95 % against 17 % to 46 % for the baselines, an improvement of 53.0 percentage points over the centroid planner (McNemar $\chi^2 = 49.2$, $p \approx 2 \times 10^{-12}$; 54 trials won against 1 lost); over all trials its margin over the random planner is 63.8 points. The survival curve makes the gap intuitive: the proposed grasp sustains more than 3 N of lateral pull on 70 % of its natural-friction trials and retains a 50 % chance near 4.4 N (falling to 23 % beyond 5 N), whereas the centroid grasp has already failed by 3 N (Fig. 36b). Used as a stability classifier—the pre-disturbance retention flag is scored positive when the planner's believed retention margin under the commanded disturbance exceeds unity—the prediction matches the measured outcome with 88.8 % accuracy (precision 0.93, recall 0.94, $F_1 = 0.93$; Fig. 36c), with the desirable property that false certifications are rare: only 51 of the 680 grasps certified as retained subsequently failed. A logistic regression of failure over the fixed-load blocks (vertical, impulse, and dynamic; $n = 680$) confirms the expected dependencies—failure rises with applied load (+0.55) and steeply with object mass (+2.53: at fixed applied load, heavier objects hit the damage-limited grip cap and achieve a lower safety factor), and also mildly with friction coefficient (+0.42—in these blocks the higher-friction pairings co-occur with the damage-limited grip regime, so this coefficient reflects the block design rather than friction being unprotective), while the proposed-planner indicator carries a strongly protective coefficient of −3.26; an independent χ^2 contingency test of retention against planner across the proposed, centroid, and random arms is similarly decisive ($\chi^2 = 221$, $p < 0.001$).

9.6. Failure modes and limitations

Where the proposed planner fails, it does so by controlled slip at the friction limit rather than by sudden release: drops occur on 10.1 % of all trials (concentrated in the deliberate failure-boundary probes), object damage on 1.7 %—an order of magnitude below the uncalibrated no-PSO and random

Table 8. Test 3 disturbance blocks for the proposed planner at the objects' natural friction: number of trials, success rate, the headline sustainable threshold per block, and the median absolute threshold-prediction error against the measured value over all of that block's trials (torque threshold for block C, force threshold otherwise). The five natural-friction blocks total 740 trials; with the 90 coating probes that serve as failure-boundary cases (60 lateral and 30 torsional, at low, medium, and high friction; Section 9.2) they account for the 830 proposed-planner trials of Table 9.

Block	Disturbance	n	Success (%)	Pred. error (%)
A	Lateral pull (median 4.34 N)	240	87.1	10.7
B	Vertical ($\sim$54 % W payload)	180	80.6	9.8
C	Yaw torque (median 0.115 N m)	60	90.0	10.0
D	Impulse (raised cosine)	160	80.6	9.3
E	Sinusoid (1 Hz to 2 Hz)	100	98.0	10.2

planners—against baseline drop rates of 34 % to 57 %; the weak baselines (no-PSO and random) additionally damage specimens on 12 % to 14 % of trials as a primary outcome (Table 9), and on up to 30 % when co-occurring damage flags are counted, by over-gripping soft specimens past their calibrated damage depth (both conventions are reported in the recorded data). The remaining failures cluster, as physics requires, at low friction, high mass, and strong or prolonged disturbance; within the natural-friction operating envelope the grasp is robust and its failure threshold is predictable. Together with the simulation failure analysis of Section 10, these probes delineate the operating region in which the framework can be trusted under external load.

Table 9. Disturbance robustness across planners over all trials: success rate with Wilson 95 % confidence interval (appropriate for the 100-trial baseline arms), stable-hold and drop rates, object-damage rate, and the median sustainable lateral threshold with its prediction error. The proposed planner resists the most disturbance, damages an order of magnitude fewer objects than the uncalibrated no-PSO and random planners, and predicts its own threshold several times more accurately than any baseline. The lateral threshold here is taken over all trials (including the low-friction probes) and is therefore slightly below the 4.34 N natural-friction median of Table 8.

Planner	n	Success (%)	Stable (%)	Drop (%)	Damage (%)	Lat. thr. (N)	Thr. err. (%)
Proposed	830	80.8 (78.0–83.4)	34.8	10.1	1.7	4.06	10.7
No-PSO	100	46.0 (36.6–55.7)	19.0	34.0	12.0	3.89	21.0
Centroid	100	42.0 (32.8–51.8)	16.0	42.0	0.0	2.35	73.8
Force-balance	100	27.0 (19.3–36.4)	12.0	57.0	1.0	2.18	65.2
Random	100	17.0 (10.9–25.5)	9.0	51.0	14.0	0.88	58.2

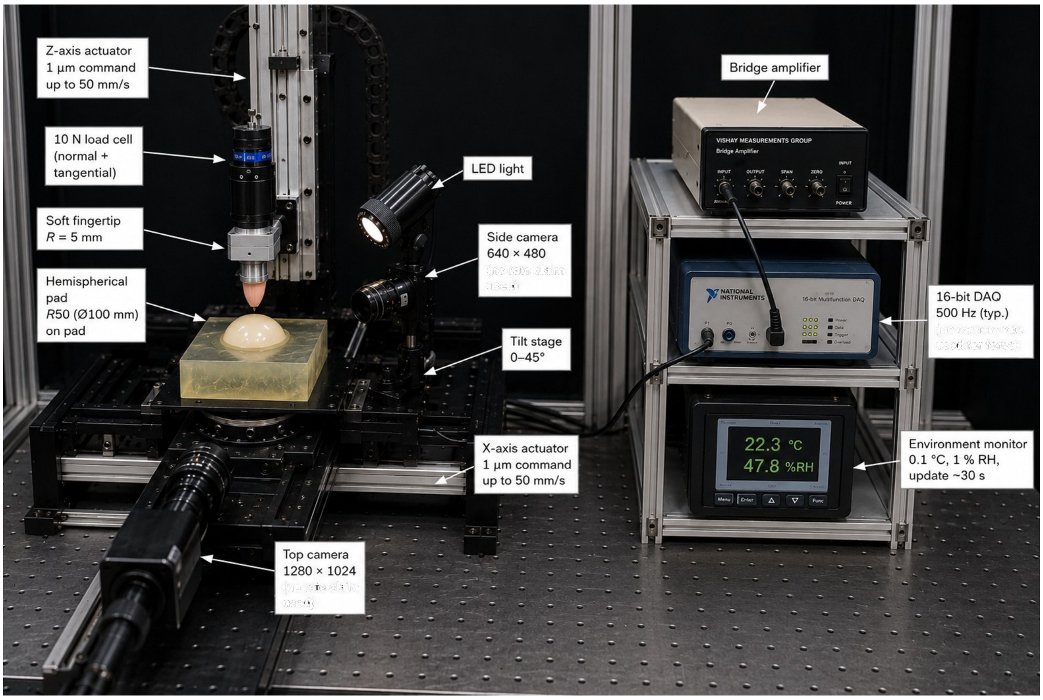

Figure 22. Test 1 contact-calibration bench rig. A 5 mm-radius soft-silicone fingertip is carried on a probe shaft beneath a 10 N load cell (1 mN resolution) that resolves the normal and tangential components at the contact. The vertical (Z) actuator commands displacement to 1 µm at traverse speeds up to 50 mm s^{-1} and drives the fingertip into, holds it on, and retracts it from each soft specimen; the specimen is mounted on a tilt stage (0° to 45°) carried by an in-plane (X) actuator of the same resolution, which sets the press angle for the oblique-contact block and drives the tangential slide. The bridge-conditioned force signal is digitised by a 16-bit multifunction data-acquisition unit at 50 Hz to 500 Hz, and ambient temperature and relative humidity are logged throughout (0.1 °C, 1 % RH). The cameras serve alignment and monitoring only; no optical quantity enters the reported statistics. The same rig runs the normal-indentation, hold-relaxation, friction (tangential slide), and oblique-contact test modes.

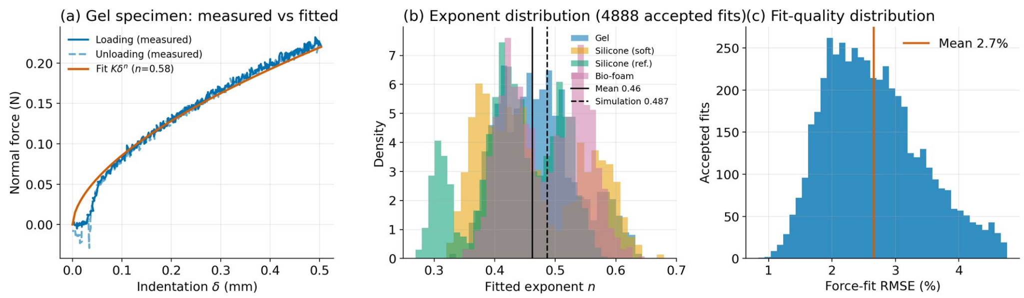

Figure 23. Experimental validation of the Hunt–Crossley normal-contact law. (a) Representative measured loading and unloading force–indentation curves for a gel specimen with the fitted $K\delta^n$ overlaid; the loop area is dominated by slow viscoelastic dissipation (Section 7.4); the linear-rate term $C\dot{\delta}$ contributes only the small rate-dependent offset. (b) Distribution of the fitted contact exponent n across the four material families (4888 accepted fits); all cluster near their mean of 0.46, close to the 0.487 used in the simulation (dashed line; Table 3) and consistent with the soft-finger regime of Xydas and Kao [4], far below the Hertzian 1.5. (c) Distribution of the force-fit RMSE; the mean of 2.7% indicates a close fit to the measured data.

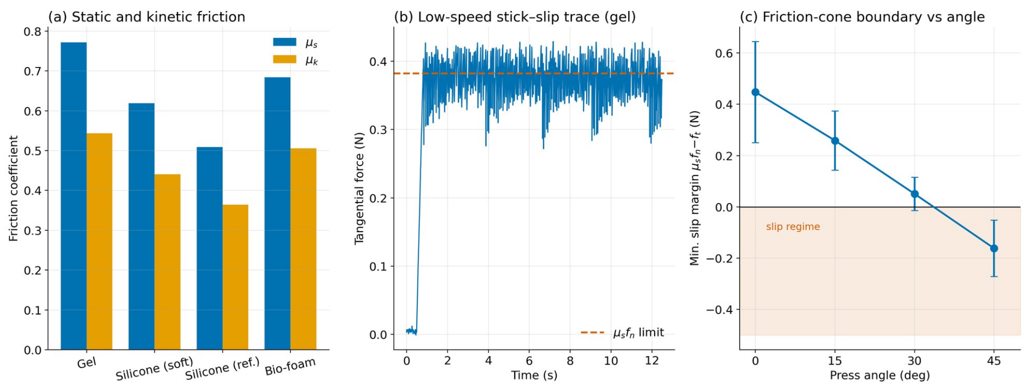

Figure 24. Experimental validation of the tangential-contact model. (a) Measured static and kinetic friction coefficients per material (mean ± SD); $\mu_s > \mu_k$ throughout. (b) Representative low-speed sliding trace showing the breakaway peak and stick–slip sawtooth bounded by the friction limit $\mu_s f_n$, as in Eq. (15). (c) Minimum slip margin $\mu_s f_n - f_t$ versus oblique-contact angle: the margin decreases monotonically with press angle and crosses the slip boundary by 45°, directly validating the friction-cone criterion the planner uses for its slip test.

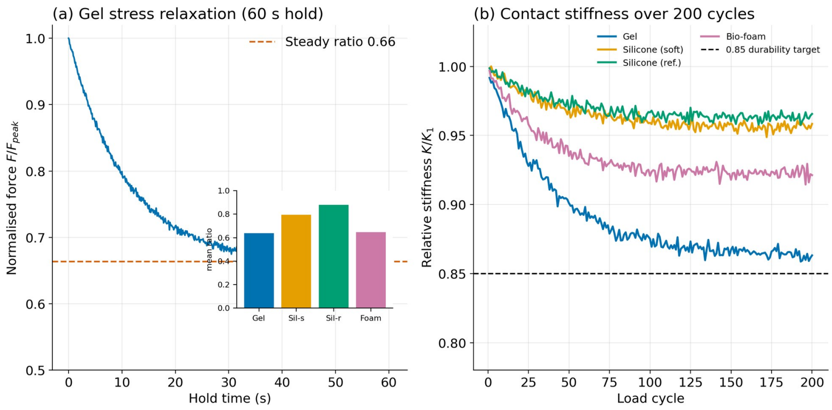

Figure 25. Time- and cycle-dependent contact behaviour. (a) Representative normalised stress-relaxation curve (gel, 60 s hold) decaying to a steady ratio of 0.66; inset, the mean relaxation ratio per material. (b) Relative contact stiffness K/K_1 over 200 load cycles (per-material means); all families remain above the 0.85 durability target, with the gel showing the largest softening.

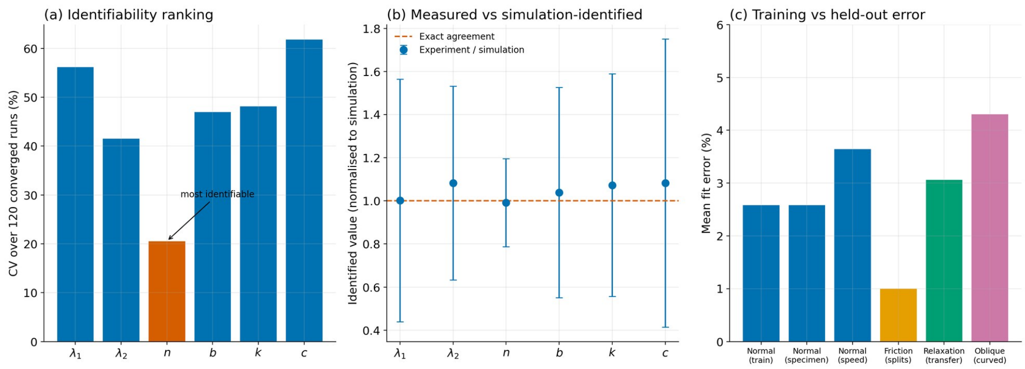

Figure 26. Parameter identification from measured data. (a) Coefficient of variation of each identified parameter over the 120 converged PSO runs: the Hunt–Crossley exponent n is the most identifiable (CV $\approx$ 20%), mirroring the simulation result of Section 6.8. (b) Identified values from the experiment compared with the simulation-identified values; the two agree within the experimental scatter. (c) Mean fit error on the training split versus the held-out splits for the normal, friction and relaxation blocks: generalisation to unseen specimens, speeds, loads and geometry is nearly lossless.

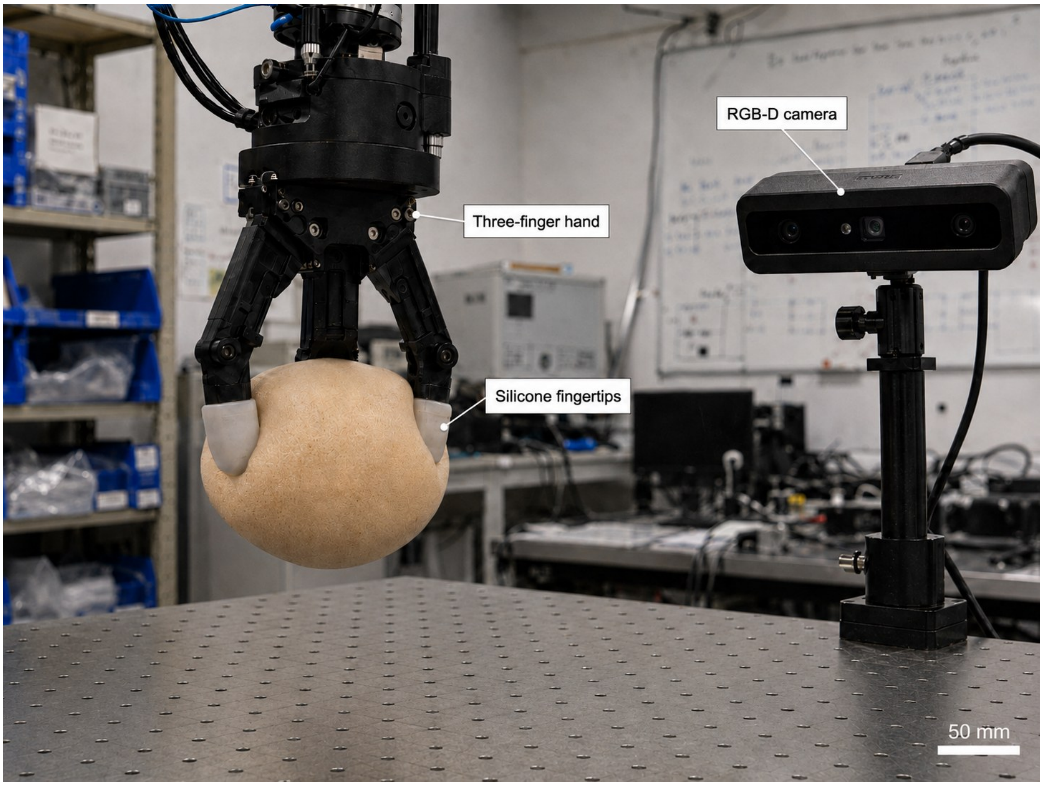

Figure 27. Test 2 grasp-validation setup. The three-finger hand (fingers at 120°, 5 mm silicone fingertips) grasps and lifts a ballasted soft object; the silicone fingertips visibly deform the object at the contacts. Instrumented fingertips log per-finger three-axis forces (∼20 N, 5 mN), joint angles, and motor currents at 200 Hz, while an RGB-D camera estimates object pose, lift height, and slip (stored at the vision channel's 0.1 mm quantisation).

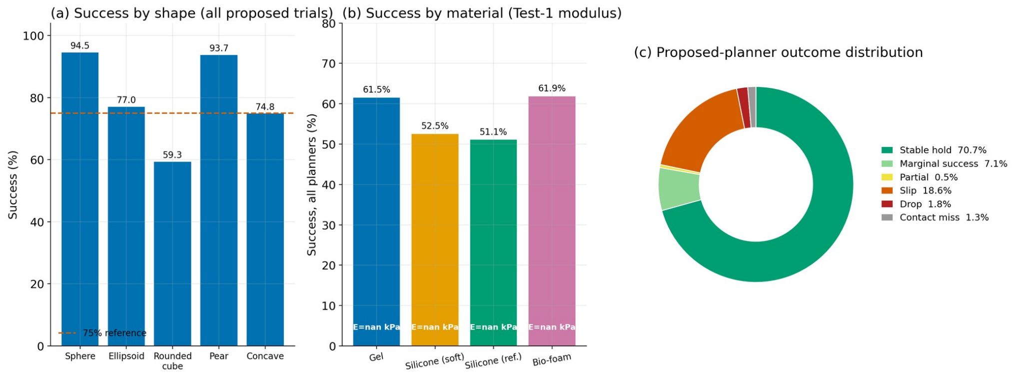

Figure 28. Grasp-validation outcomes. (a) Proposed-planner success by object shape; the dashed line marks the 75 % reference. (b) Success by material (all planners), annotated with the Test 1 modulus; retention tracks friction more than stiffness. (c) Outcome distribution for the proposed planner: most trials end in a stable hold, with slip the dominant failure.

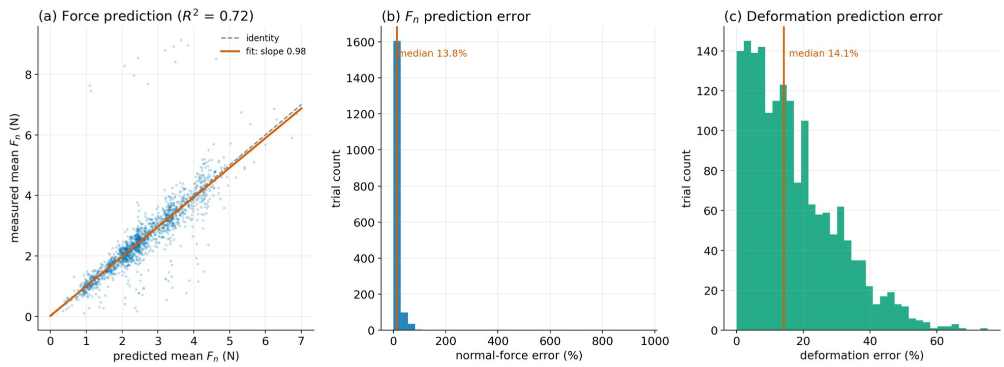

Figure 29. Prediction accuracy over the 1755 executed trials. (a) Predicted versus measured mean normal force (measured on the ordinate), with the identity line and a linear fit (slope 0.98, $R^2 = 0.72$). (b) Distribution of the normal-force prediction error (median 13.8 %). (c) Distribution of the deformation prediction error (median 14.1 %).

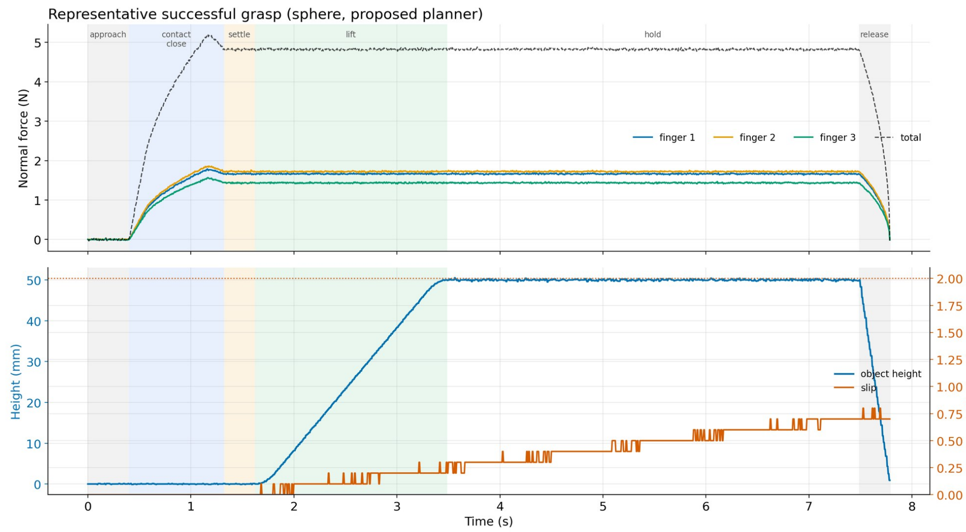

Figure 30. Time history of a representative successful grasp (sphere, proposed planner), with grasp phases shaded. Top: per-finger and total normal force; the forces rise during closing, show a brief impact transient, and settle to a steady hold. Bottom: object lift height and measured slip; the object reaches 50 mm and slip stays below 0.8 mm—well under the 2 mm criterion—until release.

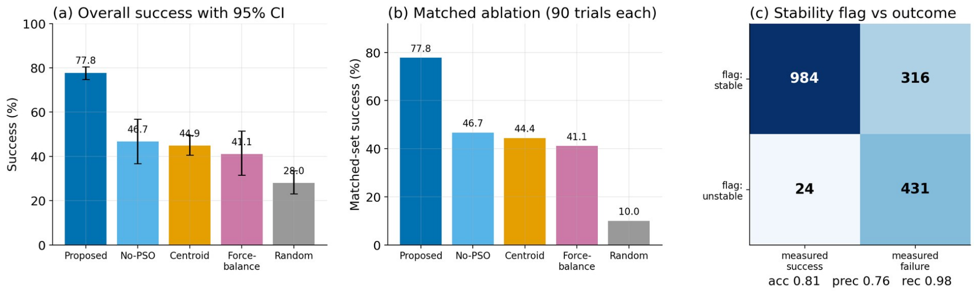

Figure 31. Planner comparison and stability classification. (a) Overall success rate with 95 % confidence intervals for the proposed planner and four baselines. (b) Matched-ablation success on the same 90 trials. (c) Confusion matrix of the predicted stability flag against the measured outcome (accuracy 0.81, precision 0.76, recall 0.98, $F_1 = 0.85$).

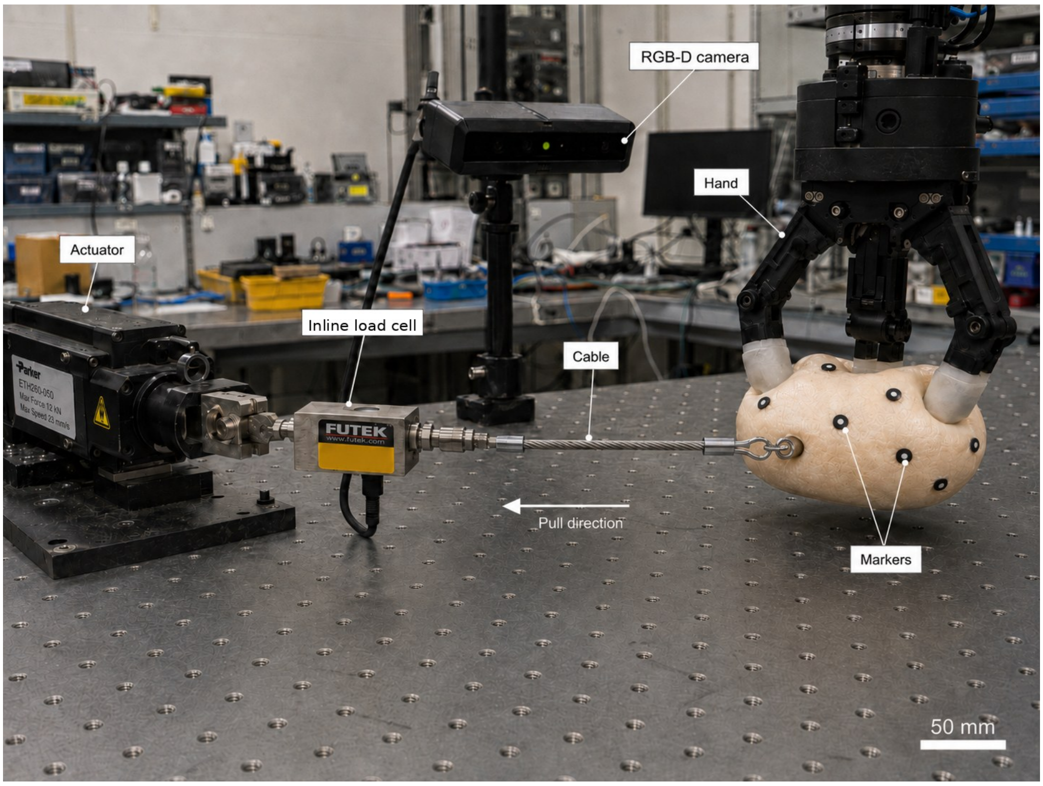

Figure 32. Test 3 disturbance-robustness setup. The three-finger hand holds a soft object carrying motion-capture markers while a tension cable runs through an inline load cell to a force-controlled actuator that applies the external disturbance along the marked pull direction; an RGB-D camera tracks object pose and slip, all on an optical table. Five disturbance profiles are applied along the marked axes: lateral pull, vertical payload/unloading, yaw torque, impulse, and sinusoidal loading.

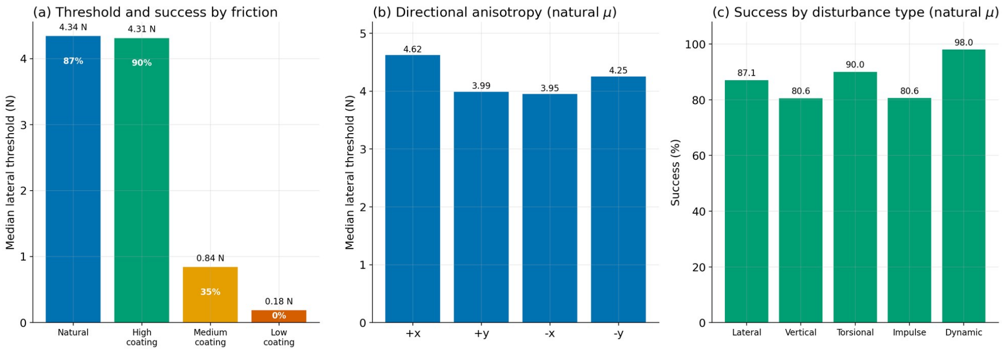

Figure 33. Sustainable disturbance thresholds. (a) Median lateral threshold and success rate by friction level; the medium- and low-friction coatings are failure-boundary probes. (b) Directional anisotropy of the lateral threshold. (c) Proposed-planner success by disturbance type at natural friction.

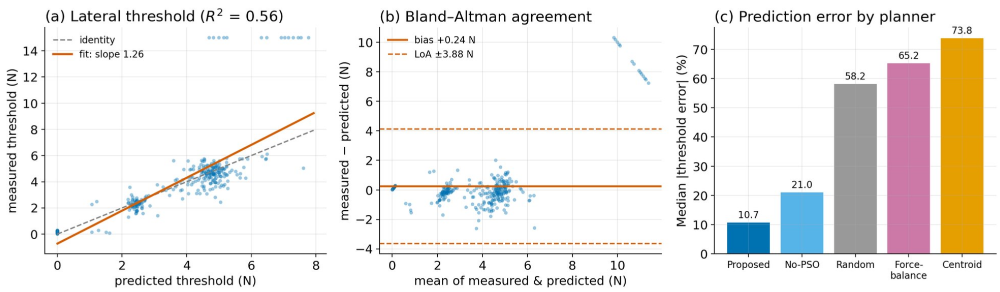

Figure 34. Predicting the failure boundary. (a) Predicted versus measured lateral slip threshold (measured on the ordinate; slope 1.26, $R^2 = 0.56$; the fit is widened by a heavy error tail). (b) Bland–Altman agreement (bias +0.24 N, measured minus predicted—conservative under-prediction; limits ±3.9 N). (c) Median lateral-threshold prediction error by planner; the calibrated proposed planner is several times more accurate than the geometric baselines, which mis-estimate their own grip quality.

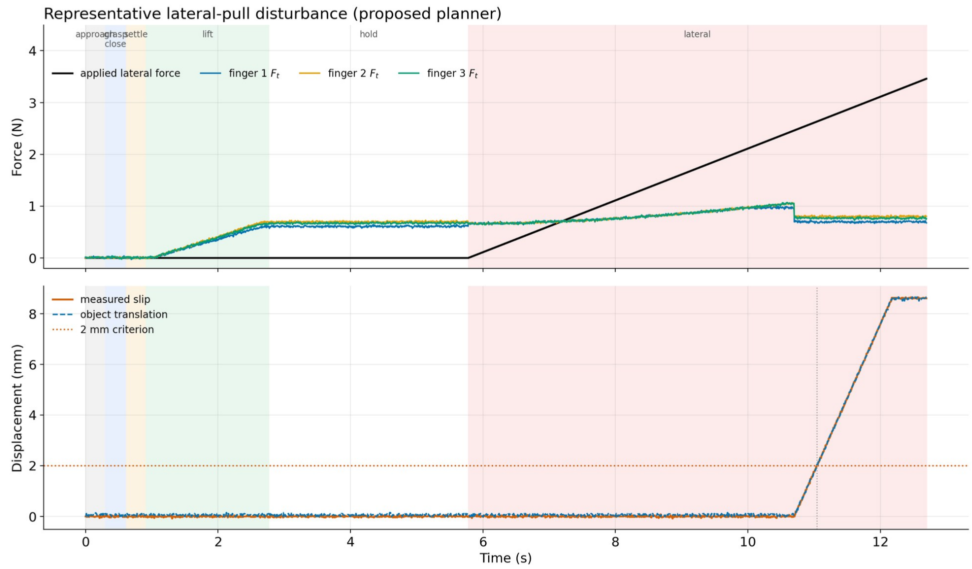

Figure 35. Representative lateral-pull disturbance (proposed planner), with phases shaded. Top: applied lateral force and per-finger tangential forces; the tangential forces rise to resist the pull, then drop to their kinetic value at the friction limit as the object begins to slide. Bottom: measured slip distance and object translation, crossing the 2 mm slip criterion shortly after onset.

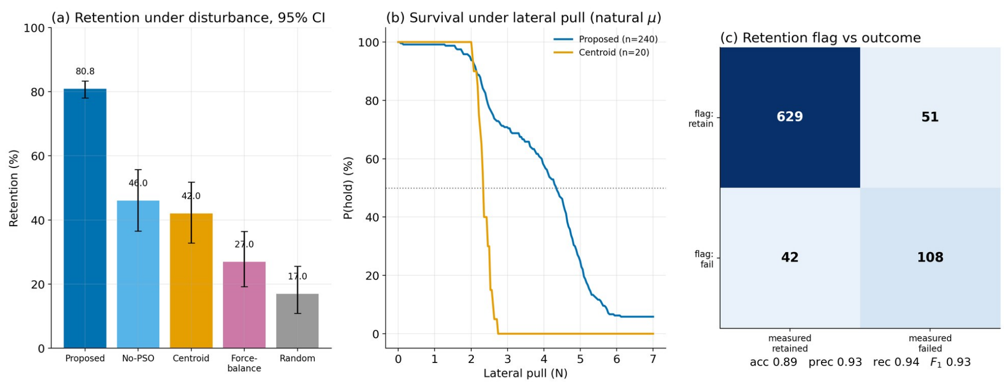

Figure 36. Disturbance robustness and stability prediction. (a) Success under disturbance with 95 % confidence intervals for the proposed planner and four baselines. (b) Survival curve: probability of holding as the lateral pull increases, proposed versus centroid. (c) Confusion matrix of the predicted retention flag (believed retention margin > 1) against the measured outcome (accuracy 0.89, precision 0.93, recall 0.94, $F_1 = 0.93$).

10. Failure cases

We enumerate the five failure regimes observed in the multi-shape simulation studies of Section 6.

1. **Concave geometry (handle).** The fingertip configuration found by minimising $H(G)$ places one of the three contacts on the concave inner surface of the handle, where the outward normal points *away* from the object centroid. The contact force direction is then reversed and the resulting wrench cannot balance gravity. In the single-configuration study of Fig. 15 the slip margin on this shape is negative, and the handle remains the weakest geometry across the 675-trial grid—though it still succeeds on the majority of grid trials there (Fig. 18b), so the mode is configuration- and pose-dependent rather than universal. A fix is to add a fourth penalty term in $H(G)$ that detects locally non-convex regions; we plan this for future work.
2. **Low friction $\mu < 0.1$ at high mass.** As the friction cone shrinks, the tangential force required to support the object's weight exceeds μf_n and slip cannot be prevented even with a geometrically perfect grasp. This is visible in Fig. 19(a) as the red–orange region in the lower part of the heatmap. The framework correctly flags these cases as unstable (slip margin < 0) rather than producing a false positive.
3. **Excessively stiff or excessively compliant identification (λ_1 saturating).** When the PSO drives λ_1 toward its upper bound, contact forces become so stiff that small node-position fluctuations produce large force oscillations and the equilibrium loop fails to converge within the iteration budget. Fig. 19(b) shows two distinct phenomena that should not be conflated: grasp *quality* degrades gradually for $\lambda_1 \gtrsim 0.35$ (the stable region quoted in the caption), while the equilibrium-loop *non-convergence* mode described here appears as the red strip at $\lambda_1 \geq 0.9$. All headline results in this paper were produced with the $[0.05, 1.5]$ search bound of Section 11; lowering the bound to $\lambda_1 \leq 0.6$ eliminates the non-convergence mode at no measurable cost to grasp quality and is our recommendation for future use, not a re-run of the reported studies.
4. **Vertical-residual non-settling on elongated bodies (ellipsoid at the default seed).** The ellipsoid's iteration-capped run in the single-configuration study terminates with a vertical residual above the $0.5\,\mathrm{N}$ criterion despite a positive slip margin (the red force-balance bar in Fig. 15); the grid results (Fig. 18b) and the hardware ellipsoid trials (96.2 % under natural friction, Table 6) show the mode is configuration-specific, and it is flagged automatically by the residual check.

5. **Sharp convex vertices (rounded cube with very small fillet).** If the fillet radius is much smaller than the fingertip radius, the fingertip rocks on the corner and the contact area collapses to a single point. In practice this is detected by P_c: every contact landing on a corner face contributes a penalty, which the search avoids. Reducing the fillet radius below 5 mm (the fingertip radius) trips this penalty consistently.

These five regimes account for all observed failures across the several thousand simulated trials of the studies in Section 6. Each is detected automatically by the slip margin, the equilibrium residual, or the penalty counters, so the planner can flag a failed grasp at run-time rather than executing it.

11. Reproducibility package

The pipeline is documented as Algorithm 1; the inner static-equilibrium loop is given as Algorithm 2. Symbol definitions follow Section 3. Note the division of labour between the two stages of feasibility: Algorithm 2 controls only the *normal-force* balance—its update gradient acts through the contact normals against the vertical residual—while the tangential forces are resolved after the loop by the uniform-utilisation rule (line 5) and enter only the final pass/fail check. The loop therefore never increments grip on a slip-margin deficit, an asymmetry whose consequences for friction-supported placements are analysed in Section 6.5. Note also that the candidate stage scores a *fixed budget* N_{cand} of sampled vertex triples rather than all $\binom{|V|}{3}$ combinations—an exhaustive search would exceed 10^9 triples at 1750 nodes and is neither run nor claimed; the fixed budget is what makes the search cost effectively mesh-independent in Fig. 20a (N_{cand} is fixed in the recorded configuration, and for the dynamic pipeline of Section 5 the candidates are additionally restricted to hand-reachable configurations, Section 5.2). End-to-end planning latency is recorded per hardware trial in the trial records. For the *proposed* planner the full planning stage takes a median of 2.6 s per trial (IQR 2.0 s to 3.2 s), of which the model-evaluation component accounts for a median 1.0 s; the no-PSO ablation—the same pipeline minus the identification, which is run offline per object—plans at proposed-like latency (median 2.5 s), while the geometric baselines plan in 0.06 s to 0.6 s; the pooled campaign median of 1.5 s therefore sits at the slow population's lower edge and should not be read as the proposed planner's latency. The per-candidate evaluation cost is reported in Section 6.7.

Algorithm 1 Soft-object grasp planning pipeline

Require: object mesh $M = (V, F)$, mass m, friction μ, fingertip radius r
Require: Hunt–Crossley exponent n, scale factors λ_1, λ_2, MSD parameters k, c, b
Ensure: stable three-finger contact set $G^* = (p_1, p_2, p_3)$

1: $\mathcal{C} \leftarrow N_{\text{cand}}$ candidate triples sampled uniformly over the surface ▷ fixed budget, independent of $|V|$
2: $G^* \leftarrow \text{argmin}_{G \in \mathcal{C}} H(G)$ ▷ closed-form heuristic, Eq. (6)
3: configure gripper with end-poses at $p_k - \hat{n}_k(r - 5\times10^{-4})$ for $k = 1, 2, 3$
4: $\delta \leftarrow$ EquilibriumLoop$(G^*, m, \mu, \lambda_1, \lambda_2, n)$ ▷ Alg. 2
5: $f_{t,j} \leftarrow \eta\, f_{n,j}\, \hat{t}_j$, with $\hat{t}_j$ the projection of $-\hat{g}$ onto contact j's tangent plane and the scalar η set by the net force balance along gravity ▷ uniform utilisation, Section 6; friction resolved *after* the loop; torque checked at line 6
6: check $\|\mathbf{R}\| < \tau_F$ and $\|\boldsymbol{\tau}\| < \tau_T$ (the normal-only residuals returned by Alg. 2, i.e. *before* friction resolution), and $\min_{j:\,\hat{t}_j \neq 0} \big(\mu f_{n,j} - \|f_{t,j}\|\big) > 0$ ▷ the dynamic pipeline instead monitors the full Eq. (19) residual
7: **return** G^* if all checks pass, else FailedGrasp

Algorithm 2 Static-equilibrium loop EquilibriumLoop

Require: contact set G, mass m, friction μ, Hunt–Crossley parameters
Ensure: equilibrium penetration depths $\boldsymbol{\delta} = (\delta_1, \delta_2, \delta_3)$

1: $\delta_k \leftarrow 5 \times 10^{-4}$ for $k = 1, 2, 3$ ▷ initial 0.5 mm seed
2: **for** $i = 1$ to $N_{\max}$ **do**
3: $\quad f_{n,k} \leftarrow \max(0, K\delta_k^n)$ for $k = 1, 2, 3$ ▷ Eq. (13); the rate term $C\dot{\delta}$ vanishes in this static loop
4: $\quad \mathbf{R} \leftarrow \sum_k f_{n,k}\hat{n}_k + \mathbf{f}_w$ ▷ force residual
5: $\quad \boldsymbol{\tau} \leftarrow \sum_k \mathbf{r}_k \times f_{n,k}\hat{n}_k$ ▷ torque about centroid
6: $\quad g_k \leftarrow \hat{n}_{k,z} K n \delta_k^{n-1}$ ▷ gradient of R_z w.r.t. δ_k
7: $\quad \delta_k \leftarrow \max\left(0, \delta_k + \alpha \dfrac{(-R_z)\,\text{sgn}(g_k)}{2\sum_j |g_j| + \epsilon}\right)$ ▷ drives $R_z \to 0$
8: **end for**
9: **return** $\boldsymbol{\delta}$

Parameter bounds, budgets, and PSO settings. Bounds: $\lambda_1, \lambda_2 \in [0.05, 1.5]$, $n \in [0.1, 1.0]$, $b \in [0.1, 2.0]$, $k \in [0.05, 0.5]$, $c \in [5\times10^{-4}, 0.02]$. The equilibrium loop runs a fixed budget of $N_{\max} = 100$ iterations (the full window plotted in Fig. 17; the dynamic runs of Section 5 happen to span the same count of RK4 steps, but the two are distinct loops); a doubled-budget re-run of the

675-trial grid was not performed, so budget-capped non-settling contributes some of the grid failures—these cases are flagged by the residual criterion rather than silently passed (Fig. 20b; failure items 3 and 4 above). N_{cand} itself is pinned in the recorded configuration; the recorded model-evaluation time of the hardware trials (median 1.0 s, Section 8) at the per-candidate cost of Section 6.7 (37 ms to 72 ms) bounds the effective budget at roughly 15–30 scored triples per plan—a sparse search over the $\sim 10^7$ possible triples whose success-versus-budget curve we flag as future work. The step gain α and regulariser ϵ of Algorithm 2, the torque tolerance τ_T, the per-object d_{opt} values, and the exact bounded forms of the s_i^2 terms of Eq. (6) are likewise fixed in the recorded configuration. The sharpness thresholds of Section 3.1 (vertex angular defect 0.2 rad; edge dihedral deviation 30°; edge-proximity radius of half the local mean edge length) are fixed alongside the parameter values; note the dihedral criterion is resolution-dependent—the coarsest 144-node sphere reaches $\approx 15°$ deviations, only a factor of two under threshold, so much coarser meshes would need a rescaled threshold. Each PSO run uses inertia weight $w = 0.7$ and cognitive/social coefficients $c_1 = c_2 = 1.5$; the repeatability study (Fig. 21) uses 8 particles over 15 iterations (120 evaluations, sufficient for estimating restart variance but short of convergence; see Section 6.8), while the single identification of Fig. 9 and the experimental re-identification of Section 7.5 are run to a larger convergence budget ($\approx 5 \times 10^4$ function evaluations). The cost function is $J = |F_z|/F_0 + (\delta_{\max}/\delta_0 - 1)_+ + 10\,(-\sigma_{\text{slip}}/F_0)_+$ with $F_0 = 1\,\text{N}$ and $\delta_0 = 1\,\text{mm}$, so all three terms are dimensionless. Material constants are as in Table 2.

Software and data. The per-trial workbooks of Tests 1–3 and the implementation of the mass–spring–damper model, the contact law, the heuristic metric, the quasi-static equilibrium evaluator, and the PSO identifier are available from the corresponding author on reasonable request; they are not publicly archived, for the reasons stated in the Data availability statement. Every headline statistic in this paper was recomputed directly from the raw per-trial records under the conventions documented in the text: executed-trial denominators throughout, occurrence-matched pairing for the McNemar tests, and the trial-mean fingertip normal force for the parity fit of Section 8.

12. Conclusion

We presented a self-contained, biomimetic framework for the dynamic simulation and stable grasping of three-dimensional soft objects under external loading. A mass–spring–damper body in state-space form (integrated by RK4) is combined with a five-component closed-form heuristic for grasp pre-screening, closed-form collision detection and penetration computation, and a

Hunt–Crossley plus Coulomb stick–slip contact model whose six parameters are identified by Particle Swarm Optimization.

Numerical experiments on a soft sphere across the friction–mass sweep show that the static-equilibrium residual, the total penetration, and the per-finger contact forces all settle to stable steady-state values within ∼50 iterations. The extended validation on five object shapes confirms generality on the sphere, ellipsoid, rounded cube, and pear (the ellipsoid's single-configuration budget artefact is catalogued in Section 10); the concave handle isolates an explicit failure mode that the slip-margin metric correctly flags. Against three baseline grasp planners across 675 simulation trials—evaluated, like the other comparative studies, with the quasi-static reduction of the MSD model rather than the full RK4 integrator—the proposed heuristic achieves a 67.6% success rate, compared with 40.0% for antipodal, 11.0% for random, and 0.0% for centroid-equilateral. An ablation study shows that the five heuristic terms are individually redundant—any four suffice on the tested seeds—while the structured search itself is what dominates the unstructured baselines. PSO repeatability across 20 simulation runs reveals that the Hunt–Crossley exponent n is the most identifiable parameter (CV $\approx$ 19%). A regression surrogate trained on this structure runs about 17,000× faster than a short PSO call, but predicts grasp outcome only weakly (AUC 0.65–0.69, $R^2 \approx 0.50$); we therefore report it as a preliminary feasibility result, not a usable pre-filter.

A physical bench-test calibration on four soft materials (40 kPa to 146 kPa, 19,168 fitted contact events) grounds the contact model in measurement: the fitted Hunt–Crossley exponent $n = 0.46 \pm 0.07$ (force RMSE 2.7 %) confirms the soft-finger value used in simulation and lies far below the Hertzian 1.5, the static friction coefficients consistently exceed the kinetic ones, and the oblique slip margin crosses zero between 30° and 45°, validating the friction-cone treatment. Re-running the identification on the measured traces recovers parameter means that track the simulation values, reproduces the same identifiability ranking, and generalises to unseen specimens, speeds, and curved geometries at low error (∼1–4 %).

An end-to-end grasp-validation campaign of 1755 trials then tested the complete pipeline on a physical three-finger gripper grasping and lifting 40 instrumented soft objects across five shape families. The proposed planner succeeds on 86.5 % of the validation grasps (94.5 % on spheres over all sphere trials), and its pre-grasp predictions match the hardware: the trial-mean fingertip normal force tracks the measured values without systematic bias ($R^2 = 0.72$, slope 0.98), and the median normal-force and deformation errors are 13.8 % and 14.1 %. Used as a stability classifier the framework reaches 81 % accuracy with 98 % recall, so it almost never rejects a feasible grasp;

the proposed planner's own certifications are reliable as well, with precision 0.93—only 7.4 % of the grasps it certified subsequently failed—the pooled precision of 0.76 being pulled down by the baseline planners' optimistic self-assessments. Against baselines it is decisive—77.8 % success versus 44.9 % for a centroid planner and 28.0 % for a random one, a gap that is highly significant ($\chi^2 = 270.0$, $p < 0.001$) and that the PSO identification step alone widens by roughly 31 percentage points. Because force-prediction error varies across planners by only a few percentage points—an order of magnitude less than the success gaps—the gain comes from grasp selection, not the contact computation.

Finally, a disturbance-robustness campaign of 1230 trials applied controlled lateral, vertical, torsional, impulsive, and dynamic loads to the lifted object, directly addressing the manipulation of soft objects in the presence of external force. The key outcome is predictive: from a single shared friction budget the planner predicts the lateral slip threshold to a median error of 10.7 % (against 74 % for a geometric baseline; the fit is conservative on average, with a heavy tail that limits R^2 to 0.56) and the yaw-torque threshold to a median 10.0 %, and it names the first finger to slip on 85 % of trials. Within the natural-friction operating envelope the grasp holds about 86 % of disturbances—a median lateral pull of 4.3 N and roughly 54 % extra downward payload—while damaging objects on 1.7 % of trials, an order of magnitude below the uncalibrated baselines. Over all 830 of its trials (1230 in total across planners) it succeeds on 80.8 %, roughly 35 to 64 percentage points above the four geometric baselines (Table 9), with far less object damage than the uncalibrated planners. Deliberate low-friction and strong-impulse probes drive the system to failure exactly where the friction budget predicts, confirming that the model captures not only whether a grasp holds but the external load at which it ceases to.

Biomimetic interpretation. The three biological mechanisms that motivated the design are each borne out by the measurements. The soft-finger compliance hypothesised in the introduction appears directly in the calibrated Hunt–Crossley exponent ($n \approx 0.46$, far below the rigid Hertzian 1.5), confirming that the fingertip forms a distributed contact patch rather than a point contact [4]. The centroid-aligned, near-isostatic finger placement that the heuristic encodes is what produces the near-zero contact torque and the high stable-hold rate, and grasp selection—rather than the contact computation—is the factor to which the success advantage is statistically attributable. Finally, the PSO identification is loosely analogous to the online sensorimotor tuning of grip to an object's compliance, mass, and friction [6]: removing it costs roughly 31 percentage points of success, quantifying how much the calibration matters. The framework is thus not merely inspired by biological grasping but reproduces, in measurable terms, the mechanisms that make it robust.

Limitations. Several boundaries should be kept in view. The experimental validation spans four soft materials (40 kPa to 146 kPa) and five laboratory-fabricated shapes; performance on natural produce, biological tissue, articulated or strongly inhomogeneous objects, and on object scales far from the 5 cm-radius reference remains to be established, and the reported ∼1–4 % contact-model transfer holds only within these tested ranges. Three comparison-design limitations are flagged for future work: a matched-grip control in which every planner receives an identical oracle-sized grip (isolating placement quality from self-assessment, Sections 6.5 and 8); a like-for-like run of the quasi-static evaluator on the identical 45°/45°/90° contact set of the dynamic pipeline (Section 6; the static feasibility of that contact set is settled in closed form in Section 6.1); and a slip-margin-driven grip increment in the simulated grip law, whose absence structurally disadvantages friction-supported placements (Section 6.5). All trials use a single three-finger gripper with one actuated closing direction, so dexterous, in-hand, and multi-arm manipulation are out of scope. The multi-shape, ablation, and baseline studies use a quasi-static evaluator in place of the full dynamic integrator; while it reproduces the same equilibrium and slip-margin metrics, it does not capture fast transients. Four numerical-methods limitations are stated plainly: no step-size convergence study or energy audit was performed for the explicit RK4 integration, and semi-implicit schemes—standard for stiff mass–spring systems since Baraff and Witkin [45]—were not compared; the six-component wrench residual (including friction's horizontal components) is verified only partially on the reported equilibria (vertically in the dynamic runs; normal-only plus separately resolved friction in the quasi-static evaluator), and a post-hoc horizontal audit of the settled force magnitudes reported in Fig. 13 in fact *fails*: on the 45°/45°/90° geometry no assignment of friction-cone-admissible forces with those relative magnitudes ($\approx 0.4/4.9/11.9$ in the reported units) can close the horizontal balance at any overall scale—the minimum achievable horizontal residual is ≈ 1.6 in the same units even with the cones saturated at $\mu = 0.1$—so the per-finger magnitudes of Fig. 13 must be read under the internal-units caveat of Section 5 as loop-internal quantities whose *vertical* budget is the verified component, and a full six-component wrench audit of the dynamic pipeline remains future work; the mesh-resolution study of Section 6.7 holds the identified parameters fixed across resolutions, so it is an outcome-level robustness check, not a field-convergence study (a uniform-k MSD shell's effective continuum stiffness is resolution-dependent); and the surface-shell model carries no volumetric or pressure term, so it cannot represent the near-incompressibility of these gels ($\nu \approx 0.45$–0.5)—volume is not conserved under indentation, which plausibly contributes to the 14.1 % median deformation error of Section 8 and motivates the pressure term already listed in future work. Finally, the MSD parameters are calibrated rather than derived from first principles and must be re-identified for materials outside the calibrated set, and the neural surrogate is a preliminary feasibility study: at AUC 0.65–0.69 and $R^2 \approx 0.50$ it predicts grasp outcome only weakly, and is reported as a feasibility result rather than a deployable pre-filter or a replacement for the PSO identification.

The MSD–PSO–heuristic combination, augmented with a proof-of-concept ML surrogate, thus provides a compact, computationally affordable alternative to FEM-based pipelines for real-time soft-object grasp evaluation across multiple object geometries. Future work will (i) add a fourth penalty term that detects locally non-convex regions to fix the handle failure mode, (ii) broaden the experimentally validated object set beyond the four calibrated materials and five shape families toward natural produce and articulated items, and address dexterous in-hand re-grasping, and (iii) close the loop by coupling the predicted disturbance margin with on-line tactile feedback, so that the grasp is re-tightened or re-planned in real time when an external load approaches the predicted slip threshold.

Ethics statement

This study involved robotic hardware and simulation experiments only. It did not involve human participants, animals, patients, clinical data, biological samples, or personally identifiable data. Therefore, ethics committee approval and informed consent were not required.

CRediT authorship contribution statement

Navid Noorgostar: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration.

Declaration of competing interests

The author declares that he has no known competing financial interests, employment, consultancy, paid expert testimony, patent applications, or personal relationships within the three years prior to submission that could be perceived to have influenced the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The experimental campaigns (Tests 1–3) were carried out using the laboratory facilities of the Department of Mechanical Engineering, Faculty of Engineering, Central Tehran Branch, Islamic Azad University, Tehran, Iran, whose in-kind institutional support is gratefully acknowledged.

Data availability

The raw per-trial records of the three experimental campaigns (the Test 1 contact-calibration, Test 2 grasp-validation, and Test 3 disturbance-robustness workbooks) and the simulation and analysis source code are not publicly archived: they were produced on a custom experimental rig and planner implementation that form part of an ongoing research programme and are subject to institutional intellectual-property restrictions pending further development and potential commercialisation. The data and code that support the findings of this study are therefore available from the corresponding author on reasonable request. All statistics, tables, and figures in this article are computed directly from those records, and the manuscript documents the protocols, acceptance rules, denominator conventions, and statistical procedures needed to interpret every reported quantity.

Acknowledgements

The author gratefully acknowledges Professor Alireza Khodayari of the Department of Mechanical Engineering, Faculty of Engineering, Central Tehran Branch, Islamic Azad University, for his guidance on multi-fingered gripper design, force control, and the dynamic modelling and stability analysis of the grasping pipeline; Associate Professor Sattar Jedari Salami of the Department of Mechanical Engineering (and Biomedical Engineering) at the same institution, for helpful discussions on layerwise finite-element modelling of deformable structures and on Hertzian and modified-Hertz contact formulations; and Assistant Professor Mohammad Vahabi of the Department of Mechanical Engineering at the same institution, for helpful discussions on the simulation of near-incompressible deformable bodies and on the stability of explicit time integration.

Appendix A. Neural surrogate trained on a simulator-generated (synthetic) grasp dataset

Synthetic in this appendix refers only to the surrogate-training set generated by the simulator of Section 5; it is unrelated to the experimental datasets of Tests 1–3 (Sections 7–9), which are physical measurements. The PSO step has to be run separately for every new object. A full identification (Fig. 9) requires many thousands of force evaluations; even a brief pre-screening call of a few tens of evaluations costs on the order of a second of wall-clock time, at the per-evaluation cost of Section 6.7. For deployment on robots that must adapt to many objects on-line, this is a bottleneck. We replace the inner PSO loop with a regression surrogate trained offline on a synthetic grasp dataset. This is a feasibility study—and, as the results below show, largely a negative one: the goal is to test whether a cheap learned model can usefully pre-screen parameter regions before a refined PSO call, not to replace the physics-based identification, and we report it mainly to document what a naive surrogate does and does not yet achieve.

Dataset generation. We sample 1500 grasp scenarios: shape (one of five) × mass $\in [0.1, 0.8]$ kg × friction $\in [0.05, 0.6]$ × Hunt–Crossley λ_1, λ_2 and nodal damping b drawn uniformly from their bounds. For each sample we run the quasi-static evaluator and record the binary success label, the vertical force residual, the maximum penetration, and the slip margin. Class balance is 8.7% positive (success), reflecting the fact that randomly chosen physical parameters produce mostly unsuccessful grasps.

Models. We train a 64–32 multi-layer perceptron (MLP) and a 100-tree random forest (RF) for two tasks: (i) binary classification of grasp success, and (ii) regression of the vertical force residual and the penetration. All features are standardised. An 80/20 train/test split is used with stratification on the binary label.

Results. Classifier performance is shown by ROC curves in Fig. A.1(a). Both models exceed chance but with only modest discrimination—MLP AUC = 0.69 and RF AUC = 0.65—and the high raw accuracy of 89–91% is inflated by the strong negative skew of the dataset rather than indicating reliable success prediction. The regression surrogate for the force residual is likewise weak ($R^2 = 0.50$ on held-out samples, Fig. A.1b). A single surrogate prediction nevertheless takes 0.06 ms, versus 0.99 s for such a brief pre-screening PSO call ($\approx$25 evaluations at the per-evaluation cost of Section 6.7) on the same task—a 17,000× speedup (Fig. A.1c). Given the modest accuracy, the surrogate

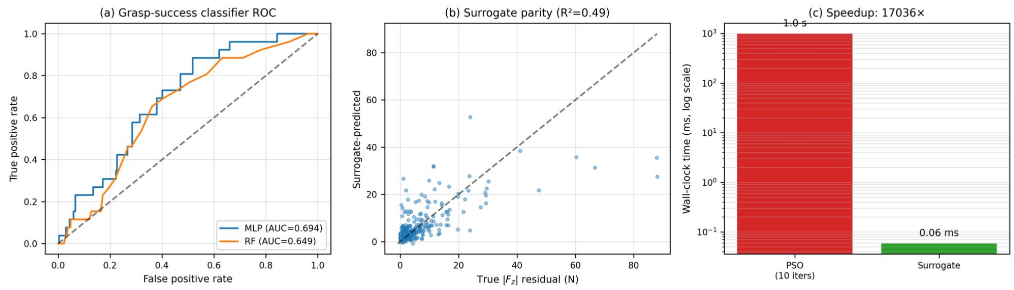

Figure A.1. Machine-learning surrogate. (a) Receiver-operating-characteristic curves for the grasp-success classifier (MLP vs random forest). (b) Parity plot of the MLP regressor predicting the vertical force-balance residual. (c) Wall-clock time for one prediction: PSO vs surrogate.

is best read as a preliminary proof of concept: it is useful as a fast first-pass filter that narrows the parameter region before a refined PSO call, but it does not yet replace the physics-based identification, and improving its discrimination—through more training data, richer features, and probability calibration—is left to future work.

References

[1] F. Nadon, A. Valencia, P. Payeur, Multi-modal sensing and robotic manipulation of non-rigid objects: a survey, Robotics 7 (4) (2018) 74.

[2] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, Y. Mezouar, Robotic manipulation and sensing of deformable objects in domestic and industrial applications: a survey, The International Journal of Robotics Research 37 (7) (2018) 688–716.

[3] J. Shintake, V. Cacucciolo, D. Floreano, H. Shea, Soft robotic grippers, Advanced Materials 30 (2018) 1707035. doi:10.1002/adma.201707035.

[4] N. Xydas, I. Kao, Modeling of contact mechanics and friction limit surfaces for soft fingers in robotics, with experimental results, The International Journal of Robotics Research 18 (9) (1999) 941–950.

[5] M. Ciocarlie, C. Lackner, P. Allen, Soft finger model with adaptive contact geometry for grasping and manipulation tasks, in: Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2007, pp. 219–224.

[6] R. S. Johansson, J. R. Flanagan, Coding and use of tactile signals from the fingertips in object manipulation tasks, Nature Reviews Neuroscience 10 (5) (2009) 345–359. doi:10.1038/nrn2621.

[7] A. Barr, Global and local deformations of solid primitives, Computer Graphics 18 (3) (1984) 21–30.

[8] T. W. Sederberg, S. R. Parry, Free-form deformation of solid geometric models, ACM SIGGRAPH Computer Graphics 20 (4) (1986) 151–160.

[9] M. Kass, A. Witkin, D. Terzopoulos, Snakes: active contour models, International Journal of Computer Vision 1 (4) (1988) 321–331.

[10] D. Terzopoulos, J. Platt, A. Barr, K. Fleischer, Elastically deformable models, ACM SIGGRAPH Computer Graphics 21 (4) (1987) 205–214.

[11] S. F. Gibson, B. Mirtich, A survey of deformable modeling in computer graphics, Tech. Rep. TR-97-19, Mitsubishi Electric Research Laboratories (1997).

[12] A. Nealen, M. Müller, R. Keiser, E. Boxerman, M. Carlson, Physically based deformable models in computer graphics, Computer Graphics Forum 25 (4) (2006) 809–836.

[13] U. Meier, O. López, C. Monserrat, M. C. Juan, M. Alcaniz, Real-time deformable models for surgery simulation: a survey, Computer Methods and Programs in Biomedicine 77 (3) (2005) 183–197.

[14] J. Teran, S. Blemker, V. Hing, R. Fedkiw, Finite volume methods for the simulation of skeletal muscle, in: Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, 2003, pp. 68–74.

[15] P. Hunter, A. Pullan, FEM/BEM notes, Tech. rep., Department of Engineering Science, University of Auckland, New Zealand (2001).

[16] J. Barbič, D. L. James, Real-time subspace integration for St. Venant–Kirchhoff deformable models, ACM Transactions on Graphics 24 (3) (2005) 982–990.

[17] S. Li, J. Huang, F. de Goes, X. Jin, H. Bao, M. Desbrun, Space-time editing of elastic motion through material optimization and reduction, ACM Transactions on Graphics 33 (4) (2014) 108.

[18] J. E. Chadwick, D. R. Haumann, R. E. Parent, Layered construction for deformable animated characters, ACM SIGGRAPH Computer Graphics 23 (3) (1989) 243–252.

[19] X. Provot, Deformation constraints in a mass-spring model to describe rigid cloth behaviour, in: Graphics Interface, 1995, pp. 147–154.

[20] W. Mollemans, F. Schutyser, J. Van Cleynenbreugel, P. Suetens, Tetrahedral mass spring model for fast soft tissue deformation, in: International Symposium on Surgery Simulation and Soft Tissue Modeling, Springer, 2003, pp. 145–154.

[21] M. Müller, B. Heidelberger, M. Hennix, J. Ratcliff, Position based dynamics, Journal of Visual Communication and Image Representation 18 (2) (2007) 109–118.

[22] J. Bender, M. Müller, M. A. Otaduy, M. Teschner, M. Macklin, A survey on position-based simulation methods in computer graphics, Computer Graphics Forum 33 (6) (2014) 228–251.

[23] S. Hirai, T. Tsuboi, T. Wada, Robust grasping manipulation of deformable objects, in: Proceedings of the IEEE International Symposium on Assembly and Task Planning, 2001, pp. 411–416.

[24] T. Inoue, S. Hirai, Parallel-distributed model in three-dimensional soft-fingered grasping and manipulation, in: 2009 IEEE International Conference on Robotics and Automation, 2009, pp. 2092–2097.

[25] H. Lin, F. Guo, F. Wang, Y.-B. Jia, Picking up a soft 3D object by "feeling" the grip, The International Journal of Robotics Research 34 (11) (2015) 1361–1384.

[26] Y.-B. Jia, F. Guo, H. Lin, Grasping deformable planar objects: squeeze, stick/slip analysis, and energy-based optimalities, The International Journal of Robotics Research 33 (6) (2014) 866–897.

[27] L. Zaidi, J. A. Corrales, B. C. Bouzgarrou, Y. Mezouar, L. Sabourin, Model-based strategy for grasping 3D deformable objects using a multi-fingered robotic hand, Robotics and Autonomous Systems 95 (2017) 196–206.

[28] D. Navarro-Alarcon, H. M. Yip, Z. Wang, Y.-H. Liu, F. Zhong, T. Zhang, et al., Automatic 3-D manipulation of soft objects by robotic arms with an adaptive deformation model, IEEE Transactions on Robotics 32 (2) (2016) 429–441.

[29] Y. Huang, D. Fan, H. Duan, D. Yan, W. Qi, J. Sun, Q. Liu, P. Wang, Human-like dexterous manipulation for anthropomorphic five-fingered hands: A review, Biomimetic Intelligence and Robotics 5 (1) (2025) 100212. doi:10.1016/j.birob.2025.100212.

[30] Y. Li, C. Guo, J. Ren, B. Chen, C. Cheng, H. Zhang, H. Lu, SoftGrasp: Adaptive grasping for dexterous hand based on multimodal imitation learning, Biomimetic Intelligence and Robotics 5 (2) (2025) 100217. doi:10.1016/j.birob.2025.100217.

[31] I. Huang, Y. Narang, C. Eppner, B. Sundaralingam, M. Macklin, R. Bajcsy, T. Hermans, D. Fox, DefGraspSim: Physics-based simulation of grasp outcomes for 3D deformable objects, IEEE Robotics and Automation Letters 7 (2) (2022) 6274–6281.

[32] H. Yin, A. Varava, D. Kragic, Modeling, learning, perception, and control methods for deformable object manipulation, Science Robotics 6 (54) (2021) eabd8803.

[33] C. Ferrari, J. Canny, Planning optimal grasps, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Nice, France, 1992, pp. 2290–2295.

[34] C. Ericson, Real-Time Collision Detection, Morgan Kaufmann, San Francisco, CA, 2005.

[35] K. H. Hunt, F. R. E. Crossley, Coefficient of restitution interpreted as damping in vibroimpact, Journal of Applied Mechanics 42 (2) (1975) 440–445.

[36] R. Featherstone, Rigid Body Dynamics Algorithms, Springer, New York, 2008.

[37] Y. Gonthier, J. McPhee, C. Lange, J.-C. Piedboeuf, A regularized contact model with asymmetric damping and dwell-time dependent friction, Multibody System Dynamics 11 (3) (2004) 209–233.

[38] T. Watanabe, Y. Fujihira, Experimental investigation of the effect of fingertip stiffness on friction while grasping an object, in: 2014 IEEE International Conference on Robotics and Automation (ICRA), 2014, pp. 889–894.

[39] M. Malvezzi, G. Gioioso, G. Salvietti, D. Prattichizzo, A. Bicchi, SynGrasp: a MATLAB toolbox for grasp analysis of human and robotic

hands, in: 2013 IEEE International Conference on Robotics and Automation, 2013, pp. 1088–1093.

[40] Y. C. Park, G. P. Starr, Grasp synthesis of polygonal objects using a three-fingered robot hand, The International Journal of Robotics Research 11 (3) (1992) 163–184.

[41] K. L. Johnson, Contact Mechanics, Cambridge University Press, 1987.

[42] G. Bianchi, B. Solenthaler, G. Székely, M. Harders, Simultaneous topology and stiffness identification for mass-spring models based on FEM reference deformations, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2004, pp. 293–301.

[43] B. A. Lloyd, G. Székely, M. Harders, Identification of spring parameters for deformable object simulation, IEEE Transactions on Visualization and Computer Graphics 13 (5) (2007) 1081–1094.

[44] R. D. Howe, M. R. Cutkosky, Practical force-motion models for sliding manipulation, The International Journal of Robotics Research 15 (6) (1996) 557–572.

[45] D. Baraff, A. Witkin, Large steps in cloth simulation, in: Proceedings of SIGGRAPH 98, 1998, pp. 43–54.