



TEAM **LEBOB** #3236

FIRST LEGO LEAGUE UNEARTHED • WESTERN AUSTRALIAN NATIONAL CHAMPIONS

Innovations

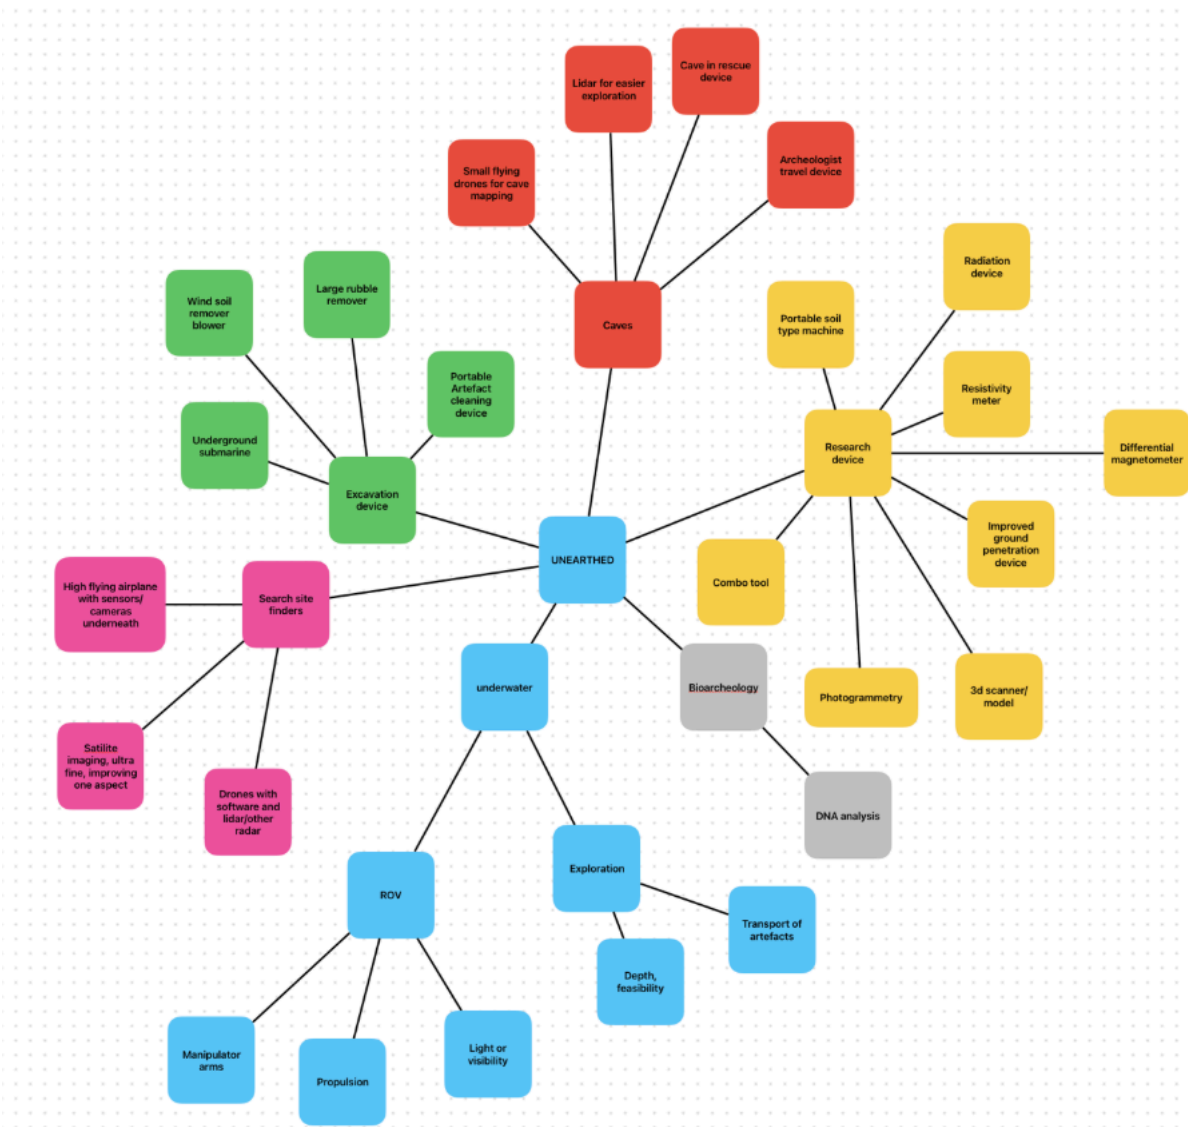
Andre Nijman • Oliver Liu • Sean Chan • Kingsley Wong
Chris Wang • Subesh Sukumaran • Aaron Zhang • Leven Shi

Contents

TEAM LEBOB #3236	0
FIRST LEGO LEAGUE UNEARTHED · WESTERN AUSTRALIAN NATIONAL CHAMPIONS	0
Innovations	0
Contents	1
1. The problem	4
Why underwater archaeology	4
Four identified problems	5
Why this matters now	5
Who would use this	5
How we researched and refined the problem	6
2. Our team	7
Roles and ownership	7
Decision-making protocol	7
3. Our solution: the four innovations of SoftSense	8
Innovation 1: adaptive soft-contact finger pads	8
Innovation 2: force feedback through the actuator	8
Innovation 3: self-morphing TPU fingers	9
Innovation 4: a cheap, kit-based, modular claw	10
How the four fit together	10
4. The mechanism in detail	11
Drive chain	11
Fingers and enclosure	12
5. How we built the fingers: finger structural FEA	13
The physical question	13
From CAD to mesh	13
Corotational large-rotation elasticity	13
Material assumptions	14
Force-targeted reporting (the fairness fix)	14
How the design search ran	15
Dead end one: concave contact face matched to the object	15
The reframe: universality is a measurement gap, not a shape problem	15
Wave one: 6 design tracks, ~45 FEA runs	16
Dead end two: ruling out the flexure family with data	16
Wave two: 4 design tracks, ~25 FEA runs to refine the Fin Ray winner	17
The finalists and the shipped geometry	17
The result at equal 12 N grip	18
Ceilings	21
6. How we chose the grip surface: texture optimisation	22
Why a separate campaign	22
Two-tier architecture	22
The Tier-1 model	22
Validation 1: the literature gate	23

The texture sweep: seven families, 700,000+ evaluations	24
Validation 2: the 50 % sensitivity sweep	25
The physical ceiling and the override	25
Validation 3: the Tier-2 contact FEA	26
The per-condition score map	26
The shipped crosshatch geometry	27
7. The motor study: sensing through the actuator	28
The principle	28
The campaign in seven acts	29
The actuator ladder	29
The forward sensing model	30
Limits of the sensor	30
Where the per-pad sensors win	30
8. The drivetrain and the structural findings we found	31
Method	31
T_safe: the input-shaft torque where the weakest tooth reaches 30 MPa	31
Per-finger force band the shipped gears can safely deliver	32
Why we kept the 24/9 = 2.667:1 ratio	32
How this reshapes every other claim	33
9. Materials, design-for-additive, and the print profile	34
Why we are 100 % polymer	34
Final-build material plan	34
Why PA12-GF for everything rigid	35
Why PETG-HF only for the finger snap pins	36
Why TPU 95A HF for the fingers (and the hydrolysis question)	36
Materials ruled out, with the reason	37
Design for additive manufacturing	37
10. Running it underwater: the seawater audit	39
Audit summary	39
What the water actually does to the finger: the 3D crush FEA	39
Patrick Morrison's checklist	40
Pre-dive and post-dive checklist	40
11. ROV integration and mounting interfaces	42
Common gripper-side mate	42
Adapters modelled	42
Two non-negotiables	43
The control electronics: an offline ESP32 “brain”	44
12. TRL self-assessment and operating boundaries	45
Current TRL position	45
Operating boundaries (1:1 production target)	46
13. Test data from the FLL bench set	46
Results (4 finger vs 2 finger configuration)	47
Reading the results	47

14. Expert and institutional validation	47
Industry support	49
Also contacted	49
15. Scaling to artefact size: the 1.5x and 2.0x build	50
One parameter, the whole gripper	50
Size and reach	51
Why the bigger finger actually grips	51
Force still follows the gears, not the size	52
Underwater and grip at the new sizes	52
What changes when you scale, and what to watch	53
16. Sources and citation provenance	54



1. The problem

Why underwater archaeology

Where we live, Western Australia, is surrounded by the ocean. With over 12,000 kilometres of coastline, over 1600 shipwrecks have been found, including the Batavia. This ship was lost in 1629, and when it was found, our state passed the first laws in the world to protect underwater heritage. The Batavia alone yielded 26,731 artefacts and 10,619 coins. So the problem SoftSense solves originates from our home.

Underwater archaeologists have a hard time recovering fragile artefacts from the seafloor, and the reason is that the manipulator arms bolted to remotely operated vehicles are built for industrial work, not precision handling. The standard arm has two rigid metal fingers, no force feedback, and little surface contact. That makes it clumsy with small, irregular, or brittle objects. When the wrong tool ends up in the wrong place, irreplaceable cultural heritage shatters on the seafloor, and nobody can do anything to change that. It is important because it preserves the history, values, and traditions that shape a community's identity. It also strengthens social connection and helps future generations understand where they come from. (Sivčev et al., 2018a; Bogue, 2015)

Getting a maritime archaeology mission off the ground costs millions of dollars:

- Vessel charter
- Crew wages
- Dive tickets
- Equipment shipping
- Port fees
- Insurance
- Fuel
- Permits

Months of planning go into a window of a few days on site. Even the recurring costs are punishing on their own: boat hire runs around AUD 5,000 per day, a four-person dive team is another AUD 1,200 per day, and dive-gear maintenance and consumables pile on top of that. Patrick Morrison at the WA Shipwrecks Museum walked us through the numbers, and the phrase that stuck with us was "the clock starts the moment you leave port."

Against that cost stack, one piece of equipment failing on site can end the whole expedition. A gripper that shears a tooth and stops responding, or simply doesn't exist in the right specification on the boat, sends the team home with nothing. The mission produces a damage report instead of a recovery, the next expedition is months away because the boat and crew have to be booked again from scratch, and the artefacts that sat on the seafloor for four hundred years stay there. A two-thousand-dollar tool

failure can waste a million-dollar mission. That asymmetry is what makes the gripper worth taking seriously.

Four identified problems

Talking to working archaeologists and subsea engineers, we pinned down four specific failure modes that the standard industrial gripper forces onto fragile recovery work:

1. **High risk of breakage.** Industrial grippers apply force with no feedback channel back to the pilot. Thin pottery, coral, bone fragments, and delicate shapes could crack if they're held too tightly.
2. **Poor grip on odd shapes.** Flat metal jaws can't adapt to curved, tapered, or uneven artefacts. Items slip out, or can't be lifted at all. The grippers in service today are designed for cylindrical T-bar handles on industrial tools, not for amphorae, bones, or coral.
3. **Limited awareness for pilots.** ROV operators can't feel how much force the arm is applying, especially in low visibility water where the camera feed is the only signal. That leads to unforeseen accidents.
4. **Cost and time pressures.** Getting archaeology equipment to site costs millions of dollars. A broken gripper or a missing tool means days of delay while they get a replacement. These failures slow missions down, push costs up, and can permanently destroy cultural heritage that has waited centuries to be recovered. (UNESCO, n.d.)

Why this matters now

Patrick Morrison, Assistant Curator of Maritime Heritage at the WA Shipwrecks Museum, told us during our 14 May 2026 site visit: (P. Morrison, personal communication, May 14, 2026)

All the wrecks we have found in Western Australia in the last four to five years are beyond recreational diving depths. The museum legally can't excavate them. The only way to recover something safely from 50 to 60 metres and beyond is a robot with actuator control.

Multibeam sonar surveys are turning up wrecks in the 40 to 100 metre range that were invisible to the previous generation of survey gear. There are more known sites than there is equipment to work them. Patrick described the gap directly: (Bell et al., 2025)

If you have a little affordable device that can do that, you'd reach for it all the time.

Who would use this

Our main user is a maritime archaeology team or research institution working from a small inspection class ROV. The closest commercial reference is the Reach Robotics Bravo 7 manipulator family, which Fugro and Woodside use for offshore inspection in Western Australia. The same gripper architecture also fits marine biology (gentle handling of corals, gelatinous animals, sediment samples),

oceanographic research, environmental monitoring, and any subsea repair job where standard grippers cause collateral damage. (Schmidt Ocean Institute, n.d.)

How we researched and refined the problem

We did not start from a gripper design; we started from the problem and narrowed it down through research and expert review.

How we researched. We studied real subsea recovery (ROV manipulators, salvage, and conservation) using marine-archaeology papers, ROV gripper datasheets, and museum recovery records. That let us map exactly why the claws in service today fail on fragile artefacts: uncontrolled crush force.

Refining the problem. Our plan was to pin the failure mode, test against it, and narrow the scope on each iteration. We started broad (“recover artefacts”) and narrowed to the specific risk of crushing fragile objects, then refined again on materials, from titanium and aluminium to printable plastics, after expert input.

Experts shaped it. We contacted industry and museums: Fugro, Total Marine Technology, Pulse Technology Hub, Woodside, and the WA Shipwrecks Museum. Woodside and Tim MacDonald steered us off metals over corrosion and seabed shedding; museum input sharpened what “fragile” and “safe handling” actually require; and every conversation confirmed the gap is real, not assumed. The full reviewer detail is in section 14.

2. Our team

Roles and ownership

Each member owns at least one work package, and most own more than one. The table below maps the public roster to the work each member led in the 2025/26 season. Where a role is shared, we name the co-owner.

Member	Primary ownership this season
Kingsley Wong	Brainstorm innovations collaboration, Meet with companies. Sponsorship outreach
Andre Nijman	SoftSense innovation lead. FEA campaign (finger structural + grip texture + motor study). Sponsorship outreach (Fugro, TMT, Pulse, Rotary). Material reevaluation (PA12-GF + ether-TPU). This documentation.
Sean Chan	Team brainstorm, collaboration, business outreach, material organization for innovation.
Oliver Liu	Robot CAD + mechanical build lead. Chassis design + attachment design. Field practice coordinator. Co-owner on prototype builds v1 + v2, drivetrain T_safe fix, hyperbaric pressure test.
Subesh Sukumaran	Research lead: existing manipulator survey + case studies. Owns expert cold-email outreach and follow-up letters. Patrick Morrison site visit attendee. Co-owner: prototype build v1, material reevaluation, 7-day soak test, AIMA + Maritime Museum outreach.
Yuxin Chris Wang	Bench-testing operator: ran the 4-finger versus 2-finger trials with the bone, anchor, vase, and chest objects, captured the 26 / 96 / 411 / 288 gf data.
Aaron Zhang	Created diagrams to explain and communicate what we did for innovations. Put together simple understandable summaries in the presentation with appropriate images, and diagrams to help communicate what we did and how our solution works.
Leven Shi	Owns the photo / video evidence trail across the season. Innovations team collaboration, co-ran the 4-finger vs 2-finger trials with the bone with chris, chassis design with oliver

Decision-making protocol

This is how we make decisions within our project:

- **Day-to-day technical decisions** are made by the work owner. Sean doesn't need a vote to refactor the mission menu, and Andre doesn't need a vote to change a Fin Ray rib angle. Owners commit to their workstream and report back at the next meeting.
- **Whole-team decisions** happen when a choice touches more than one workstream (a robot redesign that changes attachment mounts, a SoftSense material change that affects the build schedule, or a sponsorship that needs a logo on the team shirt). These go to a meeting vote with a simple majority. The mentors have observer roles but not voting roles.

3. Our solution: the four innovations of SoftSense

SoftSense takes the standard ROV gripper and adds four innovations that each go straight after one of the four failure modes from section 1. We tested every one, a working expert reviewed every one, and we describe them below in the order we built them.




Figure 3.1 The SoftSense end effector in a half-open pose, fingers pointing up. The gripper body and both Fin Ray fingers print as one swappable module, and the actuator is the modularity swap point that sets the depth rating.

Innovation 1: adaptive soft-contact finger pads

The contact face of each finger is a soft, compliant TPU surface carrying a fine printed grip texture, so the arm no longer presses bare metal against the artefact. The soft TPU spreads the gripping force over a wider area and lowers peak stress on fragile material, while the texture keeps a secure hold on wet, slippery surfaces. That cuts the risk of breakage and lets the arm safely handle smooth and curved surfaces like vases, bones, and rounded stone fragments. Where a rigid jaw point-loads a 4 mm patch of pottery rim, the compliant TPU face spreads the same total force over roughly ten times the area.

Innovation 2: force feedback through the actuator

The gripper senses how hard it is squeezing without placing a single sensor on the fingertips. It reads the actuator's motor current, multiplies by the torque constant, and traces that back through the

drivetrain to recover the tip force, the same method force-controlled industrial and surgical grippers from Maxon, Robotiq, and Schunk rely on. That reading streams to the pilot's console without a break and feeds an auto-stop control loop that halts the actuator the moment the grip crosses a configurable threshold, so the pilot gets an early warning before crushing something. The control code holds a different preset per artefact class: a low threshold for coral, another for bone, another for ceramic. The same motor-current limit that sets the grip force also caps the torque on the printed gears, so one channel is both the force command and the structural safety ceiling.

We did not begin here. Our first design put a calibrated pressure sensor under each pad; section 7, “The motor study: sensing through the actuator” (page 28), explains why we moved the sensing into the actuator instead.

David Howard, in his Regionals feedback, spotted a subtle confounder that pushed us toward this approach. As the gripper descends, ambient water pressure rises by roughly 1 atmosphere every 10 metres, so a fingertip pressure sensor would read “squeeze” when there is none, and would need a second ambient sensor plus a software correction just to recover the true value. Reading force from the motor current sidesteps the problem entirely: the current the actuator draws tracks how hard it pushes the object, not the water column above it. (D. Howard, personal communication, 2026)

Patrick Morrison raised a second concern at our 14 May site visit that pointed the same way. Any foam or gas-filled pad compresses with depth as the air pockets inside it shrink, so a pad-based sensor drifts as the dive gets deeper. Sensing through the actuator takes the pressure-sensitive part out of the loop altogether, which is the main reason we made motor-current feedback the primary force channel. (P. Morrison, personal communication, May 14, 2026)

Force feedback is essential. Most industrial archaeological arms have it. Without it, you crush things and don't know about it. Patrick Morrison, WA Shipwrecks Museum.

Innovation 3: self-morphing TPU fingers

Each finger is a passive Fin Ray compliant truss printed in TPU. When the contact face presses against an object, the internal ribs make the whole finger curl inward and wrap, so it conforms to whatever shape it meets, whether a cylinder, a square block, or an irregular fragment, instead of touching in one narrow band. The finger morphs to the object on its own, with no motor, tendon, or moving part at the fingertip to seize underwater. More contact area at the same total force means lower local pressure and less crush risk, and the self-morphing wrap also forgives imprecise placement by the ROV pilot, because the finger finds the fit the pilot could not.

This is the innovation we proved hardest in software. We ran a design search of roughly 70 FEA simulations across a battery of seven object shapes and sizes, confirmed the finger wraps along its

whole face rather than point-poking, and held the peak stress far below TPU strength. Section 5 walks through that campaign.

Innovation 4: a cheap, kit-based, modular claw

The whole gripper is built as a swappable, 3D printed module. The structural body costs around USD\$70 in filament, and because the costly actuator is reused from module to module, a complete first-build kit runs from roughly USD\$560 with the value DYNAMIXEL XM540 servo up to about USD\$1,310 with the IP68 XW540 at the top of our actuator ladder. Archaeologists can carry several spares onto a mission and swap a broken module in the field instead of waiting days for a replacement to ship. Better still, because the parts are 3D printed, the crew can bring a small printer on the support vessel and run off a custom grip or a new tool on deck while the dive plan keeps going.

This goes straight at the cost-and-time problem. The biggest line item in archaeology is not the gripper, it is the cost of the day. When the gripper fails, the day is wasted; when the gripper can be replaced from a tote in the wet lab, the mission keeps moving. Tim MacDonald (Inkfish, formerly of the DSV Limiting Factor) confirmed this directly when we called him after Regionals: (T. MacDonald, personal communication, 2026)

The biggest cost in archaeology is time. Stick with cheap 3D printed parts. Anyone can buy filament anywhere, and the whole thing can be manufactured anywhere like a kit. Tim MacDonald, subsea engineer.

How the four fit together

Soft, textured TPU faces spread the force and lower peak pressure; force sensed through the actuator's own motor current keeps the pilot informed and the gears safe on a single channel; the self-morphing Fin Ray fingers wrap to whatever shape the artefact happens to be; and a modular kit body means the gripper itself is treated as a cheap consumable, not a capital expense. Put together, the four innovations turn a basic industrial gripper into an intelligent, adaptable archaeological tool. It protects fragile objects, improves pilot accuracy, and tackles the biggest limitations of today's manipulators.

4. The mechanism in detail

The end effector has a single degree of freedom. One input shaft rotates, and both fingers open and close symmetrically, splaying outward as they open to give a funnel-mouth catch. Everything else follows from that one rotation.



Figure 4.1 The SoftSense end effector in the fully open pose, the two Fin Ray fingers splayed into a funnel-mouth catch.

Drive chain

Two equal spur sector gears mesh on the centreline. The left sector carries an integral crown gear on its top face. A small spur input pinion on a vertical shaft drives that crown, a right-angle stage that lets the drive shaft exit out of the housing bottom. From the operator's side, the gripper has one input shaft going in and two symmetric fingers coming out.

Each sector gear doubles as the crank of a non-parallelogram four-bar linkage. The finger is rigid with the coupler. The link lengths give a translate-apart motion plus roughly 18° of outward splay across the travel, and we sized them so the transmission angle stays well clear of any dead point (it varies from 71.7° to 34.6° across the open stroke).




Figure 4.2 Exploded view of the drive mechanism. The axle pins pull out to reveal the two sector gears, the four-bar links, and the crown-pinion stage; the input shaft exits the bottom face of the housing.

Fingers and enclosure

Each finger is a passive Fin Ray compliant truss in TPU. The architecture is a thin contact beam (1.2 mm), a sharply tapered compliant spine (1.8 mm), and 14 fine 1.6 mm slanted ribs with hollow cells. The printed TPU grip texture sits on the contact face, and the gripper senses its squeeze force from the actuator's motor current rather than from any fingertip sensor. We checked that the two TPU fingers build as valid solids with 0.0 mm³ interference at the closed pose at every finger scale we ship (0.7×, 1.0×, 1.6×).

The enclosure is a single open-front housing with a snap-on front cover. Inside sit the two sector gears, the four-bar links, and the crown-pinion stage. Outside are the M4 bottom mounting flange, the vertical input shaft, eight drain and flood paths (four bottom plus two per side wall), four snap-clip windows, the two top slots where the link arms exit, four axle bore floods through the back wall, and three Ø1.8 mm cover vents. Every part is a single solid with exactly one shell. There is no enclosed void anywhere in the assembly, and that is what makes the design safe to flood.

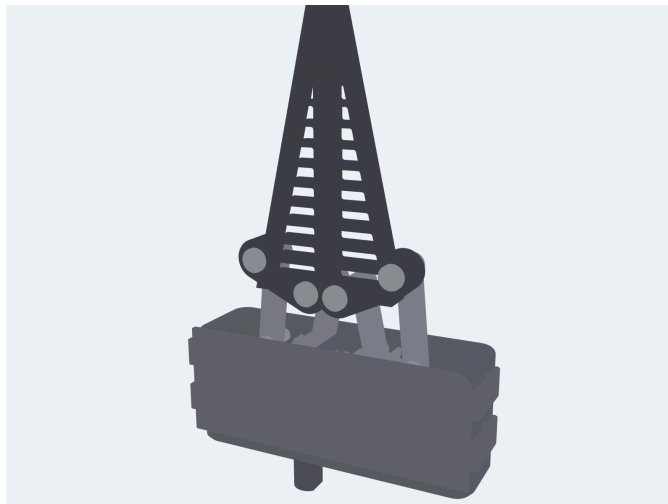


Figure 4.3 Closed pose, fingers fully together. The same finger geometry is reused across three scale factors (0.7×, 1.0×, 1.6×) without redesigning the mount interface.

5. How we built the fingers: finger structural FEA

We didn't want to print a hundred fingers and snap them one by one in a bucket of water. That's slow, expensive, and biased toward the shapes we happened to think of first. So instead we built a physics model of the finger and ran a design search through it, scoring every candidate against a battery of objects, and we only committed to a geometry once it won in software and survived independent validation. This section walks through how that ran.

The physical question

A Fin Ray finger is a compliant triangular truss: push its contact face against an object and the internal ribs make the whole finger curl inward and wrap. We needed three things proven by simulation, not by guesswork:

- It really does wrap along the whole finger, not just point-poke.
- It does that for lots of object shapes and sizes, not one tuned target.
- The peak stress stays well below TPU strength, so it's gentle on fragile artefacts and holds up under repeated use.

From CAD to mesh

gripper.py is the parametric source of truth. Every dimension you read in this document is generated from build123d primitives, never typed in twice. The FEA pipeline pulls the finger's 2D structural cross-section (contact beam, slanted spine, cross-ribs, mount eyes) straight from the CAD parameters, meshes it into triangles with gmsh (element size 0.5 to 1.3 mm), and extrudes the triangle mesh through the 10 mm finger thickness into 3 layers of linear tetrahedra. That gives us roughly 25,000 tets, 6,600 nodes, and 20,000 degrees of freedom. (build123d: Maitland, 2025; gmsh: Geuzaine & Remacle, 2009)

This holds up because the finger is a 2.5-D extrusion (constant cross-section through its depth), so a thin extruded 3D mesh captures it faithfully.

Corotational large-rotation elasticity

TPU strains stay small (a few percent), but the finger rotates a lot as it curls. Ordinary linear FEA gets this wrong under large rotation; it shows fake stress from rigid-body spin. Full nonlinear hyperelasticity is overkill. The right middle ground is corotational elasticity, and that's what we solve.

Each Newton iteration computes the deformation gradient F from current node positions, takes the polar decomposition $F = R \cdot U$ via the singular-value decomposition (with the last column flipped if $\det(R) < 0$), and extracts the co-rotated strain $\epsilon = \text{sym}(R^T \cdot F) - I$, which is strain measured once the rotation is removed. Stress is linear-elastic in this strain. We build the warped element stiffness $K_e = R \cdot K_{e0} \cdot R^T$ and internal force $f_{i0} = R \cdot K_{e0} \cdot (R^T x - X)$, and that gives correct physics for the small-strain, large-rotation regime, with a symmetric tangent that lets Newton's method converge reliably.

Material assumptions

Property	Value	Source and status
Tensile strength	27.3 / 22.3 MPa	Measured ISO 527 printed-specimen values from the Bambu TPU 95A HF TDS (V1.0), replacing the old 25 MPa estimate. In-plane 27.3 MPa is the grip-bending strength (the finger prints flat); through-Z 22.3 MPa governs the underwater crush. For a 95A elastomer at over 650 % elongation this is a stress-ceiling proxy, not a brittle yield. (Bambu Lab, 2023)
Young's modulus E	9.8 / 7.4 MPa	Measured ISO 527, anisotropic (in-plane / through-Z); replaces the old 40 MPa engineering guess. Part of the drop is definitional, 9.8 MPa is the initial-tangent modulus where the old 40 MPa was a secant estimate, not the material being four times weaker. Because internal stress at a fixed grip force is set by the load, not the stiffness, the safety margins and the design ranking are unchanged; only the absolute force at a given closure shifts. We re-centred the sweep on $E = 7 / 9.8 / 12$ MPa and the ranking still holds.
Poisson ratio ν	0.42	Bambu publishes none; 0.42 is typical for TPU (0.42 to 0.48). Relaxed from a TPU-realistic 0.48 to 0.42 to limit volumetric locking of constant-strain linear tets. Partial mitigation, not a cure; mixed u-p or P2 tets would be the real fix.
Density ρ	1.22 g/cm ³	Bambu TPU 95A HF TDS V1.0 (ISO 1183; used for buoyancy and mass only).

Force-targeted reporting (the fairness fix)

If you compare all the fingers at the same closure, a stiff finger reaches a crushing grip while a compliant one barely touches. That's unfair. So every candidate is reported at the first load step where the grip reaches a common 12 N target. Same grip force for every finger, so we compare wrap quality fairly. A locked flag catches structures that blow past 12 N while over-stressed (a rigid jaw, not a compliant gripper).

The 12 N figure is a stress-probe load we use to rank designs at a closure the FEA can actually reach in software. It isn't the force the drivetrain delivers. The printed crown-pinion's root-bending ceiling caps the per-finger operating force at 0.14 to 0.73 N on the current gears, or 4.2 to 8.7 N on the proposed re-size. The implied νM margin at the operating force is roughly 100× to 300×, not the

6.2× to 9.4× we quote at 12 N. So the 12 N comparison is a finger-design ranking probe, and the rank holds at any sub-T_{safe} load because the regime is small-strain elastic.

How the design search ran

Before the search, we wrote down what we wanted the finger to do, in three layers. The first layer was the complaint about the production finger: under FEA the grasp load all sat in the lower-mid ribs, the top of the finger carried nothing, and the finger never morphed over the object. The second layer was the harder restatement: it had to work on every type of shape and every size, not just the one cylinder we'd tuned the geometry against. The third layer was the project constraint that kept knocking out otherwise-promising answers: fool-proof, no maintenance, and underwater. That third constraint is what rules out tendon-actuated and motorised-fingertip mechanisms before they ever reach the test bench.

Dead end one: concave contact face matched to the object

A straight contact face on a round object can only tangent-kiss in one short band. The production finger touches a Ø44 mm cylinder over only about 6.6 mm of its length, all in the mid third, with the top third of the finger carrying zero load. The first idea was to bow the contact face into a circular arc concentric with the object so every face point sat the same distance from the object centre. Uniform radial penetration, even pressure, and the arc span would be the wrap.

We built two versions, a hollow shell and a solid arc, and ran both under FEA. They failed the same way. The cup contacted only at the lower lip and the rest of the arc swung open by 15 mm. The slender neck between the bracket and the cup acted as a flexible hinge: the object pushed the cup, and the cup rotated off like a door. A concave cantilever was no better than a straight one, because we'd thrown away the truss load path that ties the contact face back to the mount-anchored spine. The Fin Ray architecture itself was the load path, and trying to replace it broke the wrap.

The reframe: universality is a measurement gap, not a shape problem

After the dead end, we changed what we were searching for. Instead of trying to fit one object well, we ran every candidate against a battery of seven targets: cylinders at Ø24, Ø44, and Ø70 mm; square blocks at 28 and 44 mm; and the Ø44 cylinder offset to a low position ($y = 64$) and a high position ($y = 94$) so a finger that gripped well only in its sweet spot got caught. Each object scored on wrap (contact arc), pressure evenness (coefficient of variation along the finger), grip-plateau stability, and safety margin against TPU yield. The universal score is the battery mean minus a grip-inconsistency penalty: the finger has to work everywhere, not on one shape.

Wave one: 6 design tracks, ~45 FEA runs

Six parallel design tracks each explored one lever on the existing Fin Ray cross-section, plus three tracks on a completely different family (a monolithic flexure finger with a thin compliant spine notch instead of an internal rib truss). The table below is the wave-one summary. The screen score is on a fast 3-object subset; the full battery comes later.

Track	Family	Hypothesis tested	Best variant	Screen score	Finding
sfa	Fin Ray	Flip rib direction; sweep rib angle to curl tip in	sfa_03	0.72	Shallow rib angle (22°) unlocked the 28 mm block wrap from 0° to 86°. Direction flip alone did NOT enable circle wrap. Circles still never wrap fully.
sfb	Fin Ray	Compliant spine + slender blade	sfb_04	0.61	Gains came from even pressure (circle CoV → 0.34), not from more wrap. Arc on circles stayed near 0.
sfc	Fin Ray	Fine, soft truss (more thin ribs)	sfc_04	0.68	Thin walls (1.6 mm) wrap blocks at 86° consistently. Circle was the stubborn case.
sxa	Flexure	Spine-notch, find a usable grip band	sxa_01	0.51	Bistable: circle either missed (~2 N) or snapped to hundreds to thousands of N. No usable middle band.
sxb	Flexure	Contact-notch, try to find non-locking wrap	sxb_01	0.58	Wrap (20° to 120°) and crush were coupled. Configurations that wrapped circles also locked them.
sxc	Flexure	Both-side neck + taper	sxc_02	0.43	Central neck bistable on circles. Floppy or jam, no plateau.

Two cross-cutting findings came out of wave one. The Fin Ray family wraps blocks and large objects beautifully but point-contacts every round object: the flat contact face can't curl around a cylinder, and changing the ribs doesn't add the degree of freedom that would let it. The flexure family wraps circles but is structurally unstable. We had to decide what to do with each.

Dead end two: ruling out the flexure family with data

The contact-notch flexure produced the only real circle wrap we saw in wave one (configuration `t_fx_stiff3` reached 101° contact arc on a cylinder). But the same configuration locked at grip = 1037 N with a safety margin below 1. The structure was over-stressed. We wanted to know whether the snap-to-locked behaviour was a numerical artefact of the load-stepping resolution, so we re-ran the

120°-wrap configuration (sxb_05) at 40 closure steps of 0.25 mm each and printed the grip force at every step:

```
press (mm):  1.50  1.75  2.00  2.50  2.75  3.00  3.25  3.50  ...  9.75  10.0
grip  (N):   0.21  148   0.0   0.67  2508  107221 4549  12   ...  1.6e6  290
```

The grip jumped around chaotically between 0 N and over 100,000 N at adjacent 0.25 mm steps. The von Mises stress stayed low; this wasn't material failure but penalty-contact overshoot as the notch snapped between floppy and jammed configurations. A real gripper with this finger would have uncontrollable grip force on round objects. So we ruled out the flexure family and recorded the result so a future team wouldn't re-investigate it.

Wave two: 4 design tracks, ~25 FEA runs to refine the Fin Ray winner

With the flexure family out, wave two refined around sfc_04, the wave-one Fin Ray champion.

Track	Lever explored	Best variant	Screen score
w2a	Local refine (rib count, angle, wall thickness, direction)	w2a_05	0.670
w2b	Thin contact face (t_contact reduced)	w2b_02 (t_contact 1.2)	0.664
w2c	Blade length (new parameter blade_len)	w2c_06 (76 mm)	0.666
w2d	Slender + compliant tapered spine	w2d_05 (the winner)	0.682

w2d_05 won. Pairing a thinner contact face (t_contact = 1.2 mm) with a sharply tapered spine (spine_x_tip = 3 mm) collapsed circle pressure_cov to 0.31 to 0.35 at a safe, consistent ~12 N grip. A separate combination wave that merged the best lever from each track didn't beat it. The score had plateaued; further geometry sweeps were finding noise, not signal. So we stopped iterating and validated on the full 7-object battery.

The finalists and the shipped geometry

Candidate	Full battery score	Notes
Production finger (w7 baseline)	0.559	Starting point. Wraps the 44 mm square but nothing else.
FULL_w2a05 (16 ribs, dir +1)	0.602	Better than baseline; the rib-direction lever helped a little.
FULL_w2c06 (blade length 76 mm, angle 32°)	0.575	Shorter blade did not generalise across sizes.
FULL_w2d05 (tapered spine + thin contact face)	0.652	The winner. 17% gain on the universal score.

The sfa_03 candidate scored 0.72 on the 3-object screen but collapsed to 0.595 on the full 7-object battery. That collapse confirmed the screen had been giving us false positives on R20-block resonance, and it backed up the decision to grade the finalists on the harder battery.

Here's the shipped finger geometry: 14 ribs at 38° angle with the reversed slant direction, 1.2 mm contact beam, 1.8 mm spine, 1.6 mm ribs, 2 mm tip width with a sharp taper. Ported into gripper.py as FR_N_RIBS = 14, FR_RIB_DIR = -1, FR_TIP_WIDTH = 2, FR_CONTACT_WALL = 1.2, FR_SPINE_WALL = 1.8, FR_RIB_WALL = 1.6. Verified after the port: both fingers build as valid solids, there's zero finger-to-finger interference at the closed pose, and the four-bar closure is unchanged from the baseline.

The result at equal 12 N grip

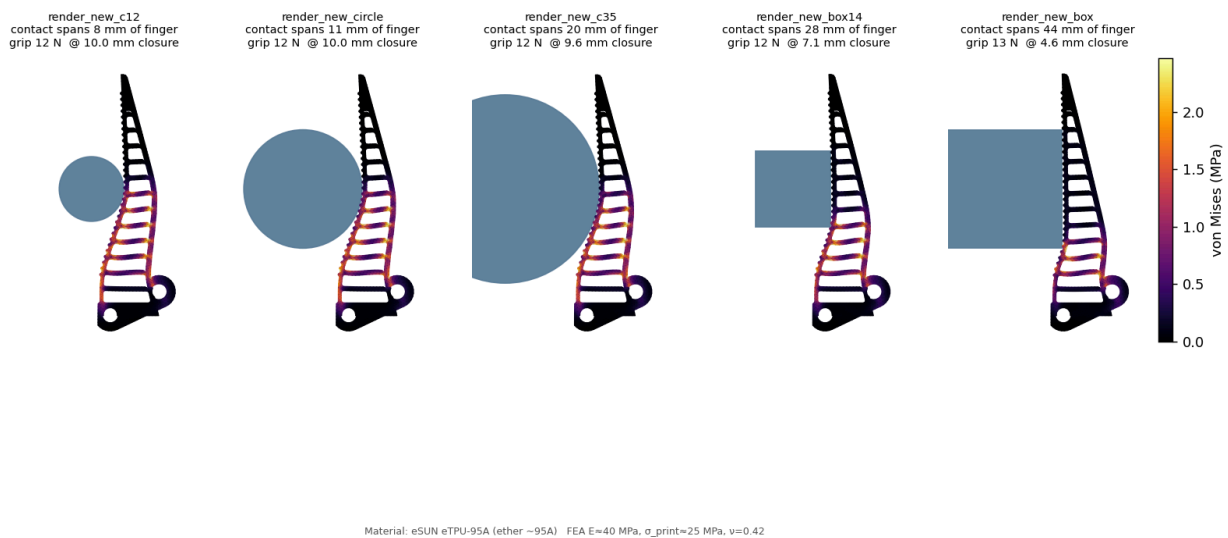


Figure 5.1 The shipped finger gripping five objects at the same 12 N force: Ø24, Ø44, Ø70 mm cylinders and 28 mm and 44 mm square blocks. Contact length grows from 8 mm to 44 mm as the object gets bigger or flatter. Colour is von Mises stress. The peak across the battery stays well below the 27.3 MPa in-plane TPU strength.

The aggregate universal score moved from 0.559 on the old finger to 0.584 on the shipped finger, a 4.5 % gain. The real wins aren't in the aggregate; they're in the failure modes:

Object	Old finger (arc / CoV)	Shipped finger (arc / CoV)	Change
Cylinder Ø24 (R12)	2° / 0.45	2° / 0.71	similar contact
Cylinder Ø44 (R22)	7° / 0.74	17° / 0.84	≈ 2.5× contact arc
Cylinder Ø70 (R35)	11° / 0.68	21° / 1.12	≈ 2× arc, less even
Square 28 mm	1° / 0.83	87° / 0.98	now wraps it (was a near-miss)
Square 44 mm	88° / 1.23	87° / 1.00	wraps, more even
Ø44 low (y=64)	6° / 0.64	17° / 0.77	≈ 2.5× contact arc

Ø44 high (y=94)	10° / 0.79	6° / 0.76	softer at the tip
Universal score	0.559	0.584	+4.5 % aggregate

The old finger's grip swung by 7 times with object height (a stiffness-gradient artefact). The shipped finger grips every object at a consistent 12 N at the test load. The old finger only wrapped one square size; the new one wraps both. On round objects the contact arc roughly doubles.

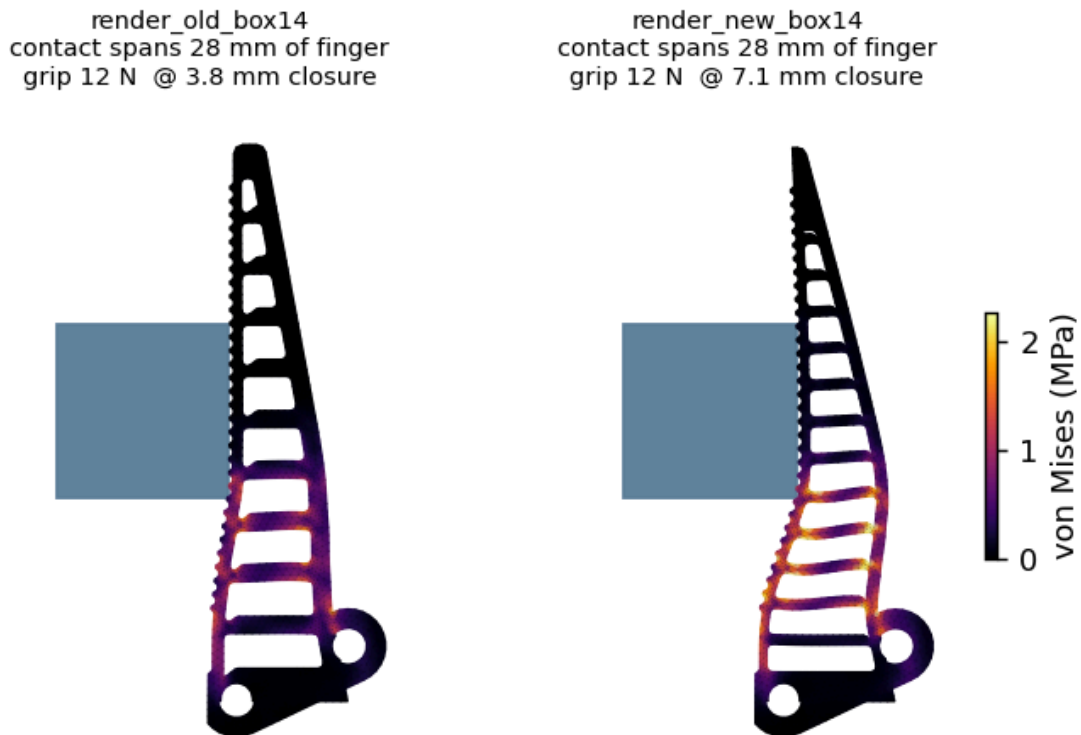


Figure 5.2 Old finger versus shipped finger on the 28 mm square block at equal 12 N grip. The old finger barely catches the edge (1° contact, grip reached at tiny closure). The shipped finger engages the flat face along its length and grips firmly.

render_old_circle
contact spans 9 mm of finger
grip 13 N @ 5.8 mm closure

render_new_circle
contact spans 11 mm of finger
grip 12 N @ 10.0 mm closure




Figure 5.3 Old finger versus shipped finger on the $\text{\O}44$ mm cylinder. The new finger is more compliant (it reaches the same 12 N at more closure) and spreads stress more evenly along the ribs. On a round object neither finger fully curls around it; that's the physics ceiling of a passive single-piece finger without a tendon.

6. How we chose the grip surface: texture optimisation

Once the finger geometry was locked, we still had to sort out the actual contact surface, the micro-relief that decides whether a finger catches on a wet object or slides off it. The finger structural FEA can't answer that question: it has no friction model and no fluid model. So we set up a second campaign with its own model, its own validation gates, and its own list of candidate textures.

Why a separate campaign

The shipped finger needs a surface that copes with every kind of object an archaeologist might recover:

- Smooth ceramic
- Rough rock
- Ridged coral
- Slimy bone
- Soft tissue
- Small round shells

The Bambu TPU 95A HF we print the fingers in comes off the print bed with a glossy, low-friction skin. A textured contact face has to beat both the water film and that slick skin at once. The structural FEA does none of this, because it only sees mechanical stress.

Two-tier architecture

There's no cheap first-principles simulator for wet elastomer friction, so we built two tiers:

- **Tier 1:** a mechanistic surrogate (microseconds per evaluation) that scores any texture. It's fast enough to run hundreds of thousands of candidates.
- **Tier 2:** a real plane-strain contact FEA that checks the contact-mechanics pieces the surrogate leans on.

We wrapped three independent validations around the surrogate (a literature gate, a 50 % sensitivity sweep, and the Tier-2 FEA) so it couldn't just be a function we tuned to give us the answer we wanted.

The Tier-1 model

A texture gets resolved into neutral geometric descriptors:

- Land fraction
- Land width
- Channel width
- Depth
- Edge density

- Drainage path
- Directional factors

For each object condition the model builds an effective holding coefficient from:

1. **Soft-land flattening.** Real contact area grows weakly with load. (Tier-2 FEA confirms this is negligible at our pressures.)
2. **Elastomer friction (Briscoe and Tabor).** Interfacial shear $\tau = \tau_0 + \alpha \cdot p$. The adhesion part scales with real contact area; the load part does not. (Briscoe & Tabor, 1975)
3. **As-printed TPU slickness.** A flat printed land keeps only about 45 % of ideal adhesion because of the glossy FDM skin; edges and channel walls deglaze it back toward ideal.
4. **Wet drainage (Reynolds squeeze-film).** Dewetted fraction comes from the squeeze-film drain time over half a land width to the nearest channel, gated by channel capacity and boosted by edge film-piercing. With no open channels you get hydroplane; with fine open channels grip comes back. It's the tyre-tread and tree-frog mechanism. (Reynolds, 1886; Persson, 2001)
5. **Partial-slip edge efficiency.** Splitting a monolithic pad into many small lands resets the edge stress at each one, so efficiency climbs with subdivision. It's a monotone surrogate, not gecko adhesion and not Cattaneo-Mindlin; the direction is established in tyre and tree-frog literature, and the functional form is engineering convenience. (Federle et al., 2006; Drotlef et al., 2013)
6. **Directional coverage.** 1-D ridges resist only cross-slip; 2-D or isotropic patterns approach 1 in every direction.
7. **Durability.** Root bending $\sigma = 6 \cdot \tau \cdot AR$; margin against the 27.3 MPa printed strength. Validated by Tier-2 FEA.
8. **Micro-suction.** A small cavity bonus for dimple and sucker patterns, which we explicitly flag as speculative.

Validation 1: the literature gate

Before we trusted the model to rank anything, it had to reproduce the published wet-grip ordering of five real patterns. The model works out each pattern's wet holding coefficient and asserts six orderings. All six pass: (Fwa et al., 2009; Federle et al., 2006; Tramacere et al., 2013; Baik et al., 2017)

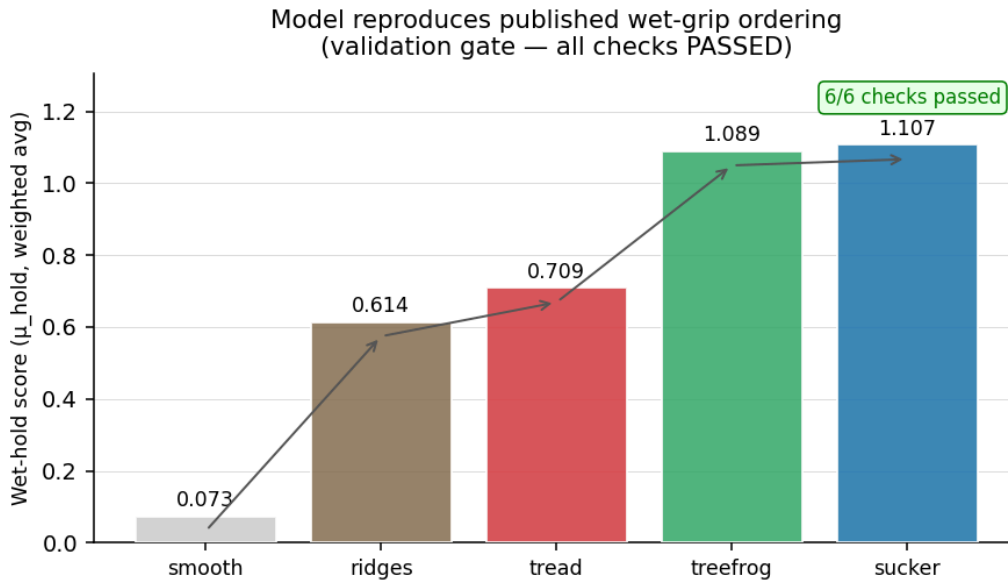


Figure 6.1 The literature validation gate. *smooth* \ll *ridges* < *tyre tread* < *tree-frog* \approx *octopus sucker*. The model passes all six checks; the design sweep was gated on this passing.

The texture sweep: seven families, 700,000+ evaluations

We ran seven texture families in parallel. Each family had its own description and its own parameter ranges (post pitch, channel width, depth, post height, ridge angle, ring spacing, dimple radius, hierarchy ratio). Each family ran an inner sweep over its full parameter space at 0.1 mm resolution, refined around its best score with a finer mesh, and reported its champion configuration along with how much that champion leaned on the speculative or placeholder coefficients in the model. Together the seven families produced more than 700,000 model evaluations. The families were:

- Ridge
- Crosshatch
- Chevron
- Hexpad (tree-frog)
- Concentric (octopus sucker)
- Dimple
- Hierarchical

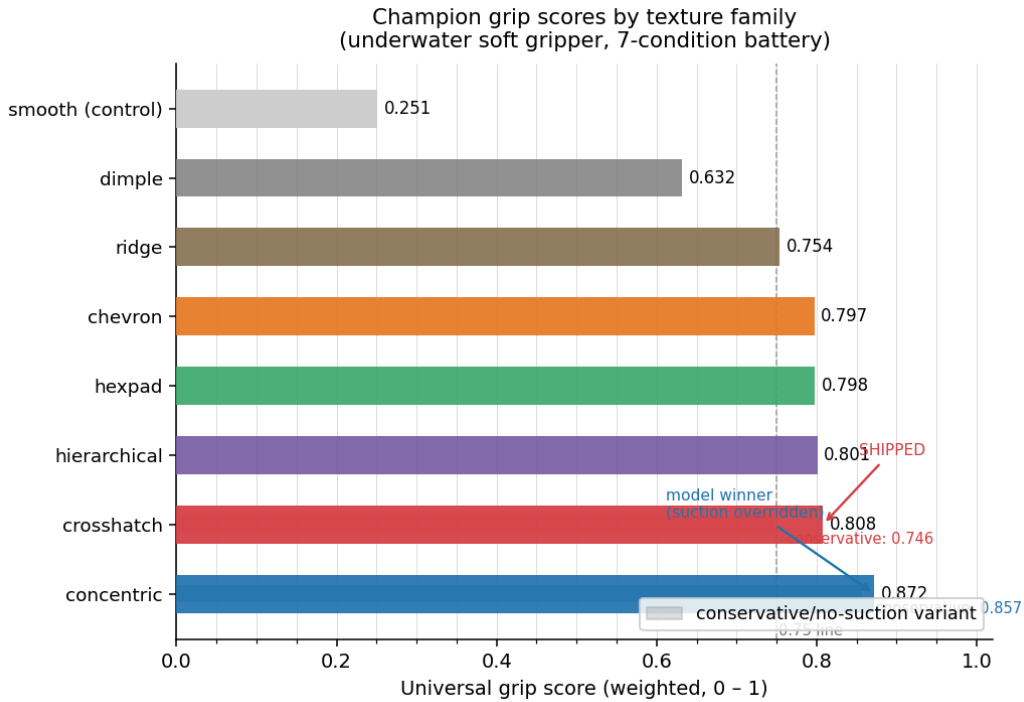


Figure 6.2 Champion universal grip scores by family across the 7 condition battery.

Validation 2: the 50 % sensitivity sweep

Here's the central risk of a surrogate model: "family X wins" might just reflect the numbers we chose. So we perturbed every coefficient to $0.5\times$ and $1.5\times$ and, at each of 31 settings, re-optimised all seven families and recorded the winner. A winner that survives the whole sweep is a real conclusion.




Figure 6.3 The sensitivity heatmap. Stars mark the winner per column. The concentric (octopus sucker) pattern wins 31 of 31 settings; among the patterns that actually tile a finger blade, crosshatch wins 23 of 31.

The physical ceiling and the override

Concentric is the model's invariant winner, and we didn't ship it, for two reasons. First, it can't tile the finger blade: the contact face is about $72\text{ mm} \times 10\text{ mm}$, and at a 1.4 mm ring pitch at most one full rosette fits across the 10 mm width. The rest is truncated ring fragments and inter-rosette gaps where the directional-isotropy benefit disappears. The model has no tileability term, so this is a missing-physics override, applied externally and on purpose. Second, its remaining lead leans on the

speculative micro-suction term, which a 0.16 mm layer FDM TPU cavity can't actually sustain (stepped rims break the seal).

Among patterns that actually tile the blade, hold up without speculation, and print reliably, crosshatch is the empirical winner at 23 of 31 settings. That's what ships.

Validation 3: the Tier-2 contact FEA

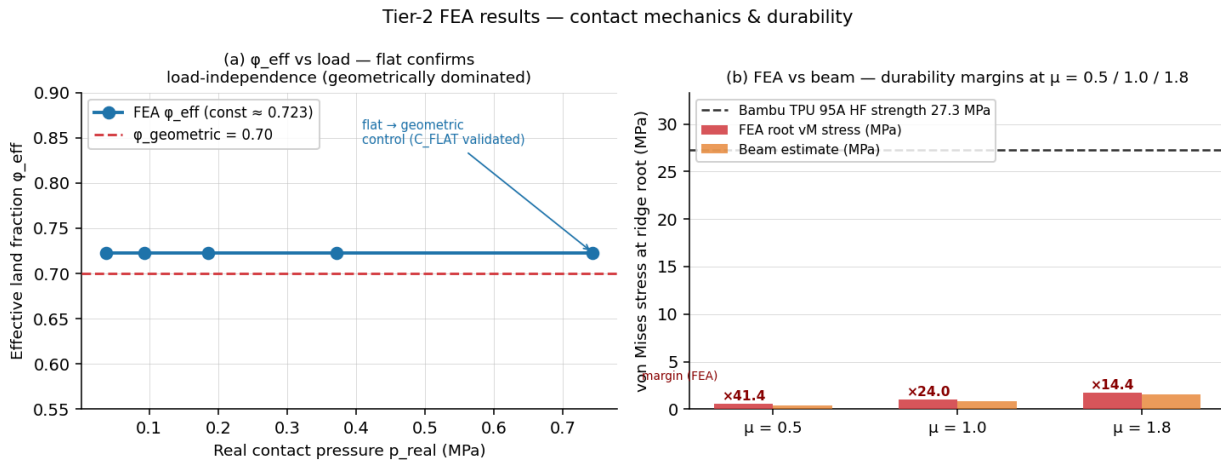


Figure 6.4 Tier-2 plane-strain FEA on the shipped crosshatch post cross-section. (a) Effective contact land fraction is geometrically dominated and load-independent. (b) Post-root durability margin is 24 \times at $\mu = 1.0$ and 14 \times at $\mu = 1.8$.

This is a real finite-element simulation (Q4 bilinear quads, 2 \times 2 Gauss integration, penalty contact against a rigid platen, sparse direct solve). Scope: it validates only the contact-mechanics sub-models every family shares, not the absolute grip number, which rests on Tier-1 plus literature.

The per-condition score map

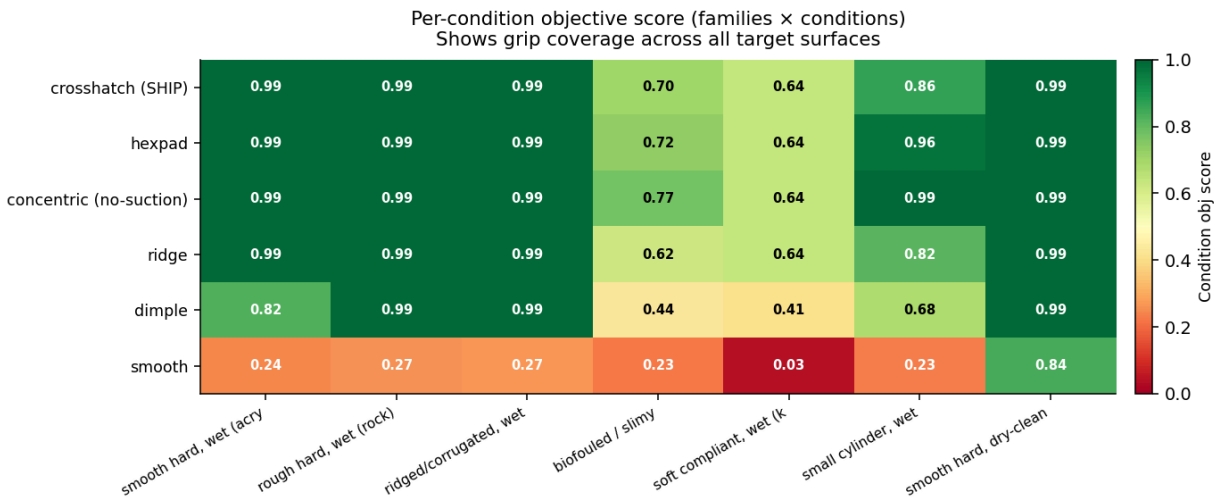

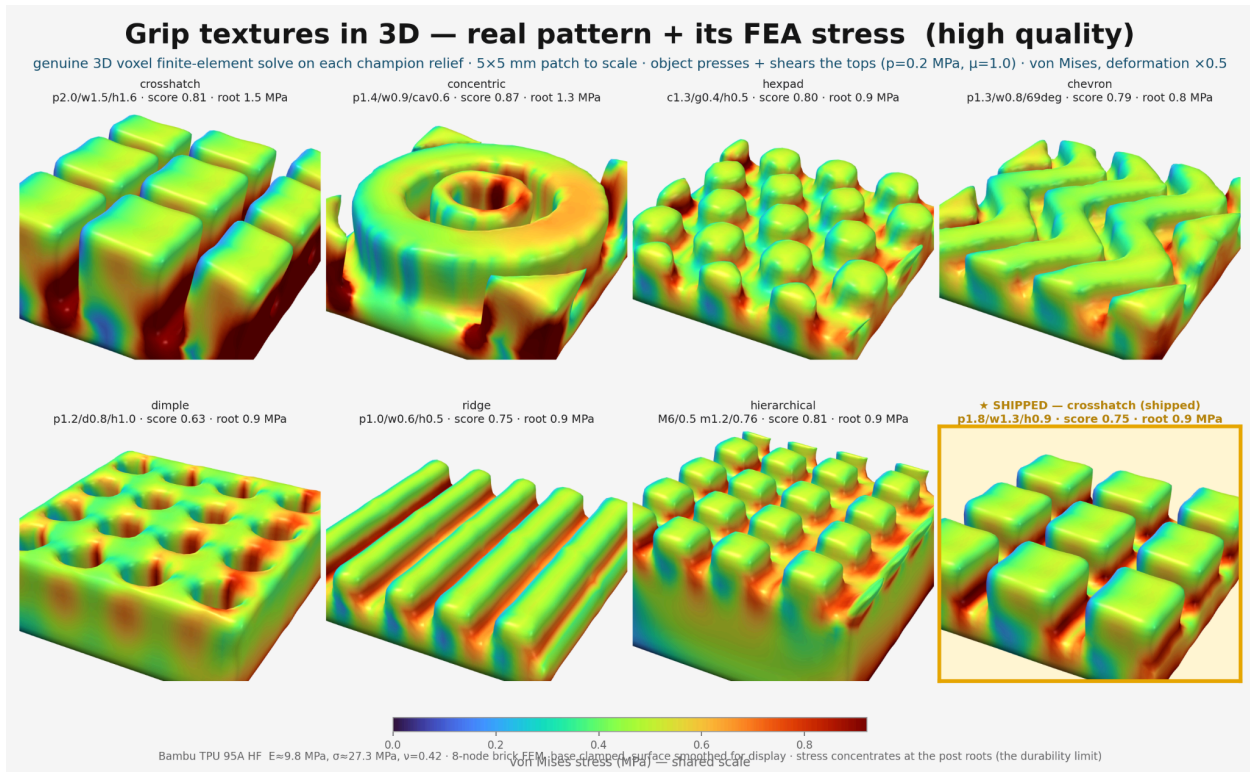


Figure 6.5 Per-condition objective score across the 7-surface battery. The shipped crosshatch grips 0.99 on smooth, rough, ridged, and dry surfaces; 0.86 on a small cylinder; and degrades only on slime (0.70) and soft objects (0.64). The smooth control fails on every wet surface, which is the hydroplaning the texture exists to defeat.

The shipped crosshatch geometry

1.8 mm posts, 0.54 mm crossing channels, 0.6 mm deep. Ported into gripper.py as FR_GRIP_CROSS, with both fingers verified as valid solids with 0.0 mm³ interference at the closed pose.



7. The motor study: sensing through the actuator

The four innovations in section 3 are the FLL-judging story. Alongside that, we chased a deeper engineering question: could the actuator's own current draw stand in for the per-pad pressure sensors entirely, so there's no fingertip electronics at all? The answer turned out to be "yes, with a real cost", and we've written the work up here so the engineering thinking sits on record next to the public solution.

The principle

Force-controlled industrial and surgical grippers (Maxon, Robotiq, Schunk) don't put strain gauges on every fingertip. They read the motor current, multiply by the torque constant, trace it back through the drivetrain, and recover an estimated tip force. The current that drives the

gripper also senses it. (Antonelli, 2014; Petillot et al., 2019)

For an underwater tool the payoff is hard to argue with: no fingertip wires running through a flexing TPU finger, no fingertip connector to flood, no separate analogue front-end to seal. The sensing chain is the same chain that already exists for control.

If you have a force sensor in the motor, you've replaced one of the worst sealing problems with one of the easier ones.

The campaign in seven acts

Act	What it established
1. Requirements	Input torque 0.56 to 1.18 N·m for 12 N per finger; travel only 123°; sensing budget at least 50 Hz, step size 0.3 N or less.
2. Survey	12 actuator candidates, 7 actuator classes × 5 sensing modalities, every spec datasheet sourced.
3. Selection	Weighted score plus 50 % sweep across depth tiers; the winning class flips with depth.
4. Drivetrain	Gear FEA finds the printed crown-pinion is the structural limit (see §8).
5. Sensing	Forward and inverse model, calibration rig, noise floor, what the sensor does and does not do.
6. Integration	Mount, isolation, cabling, tether power budget, telemetry channel.
7. Validation plan	Bench test list and the criteria for pass and fail.

The actuator ladder

The selection sweep puts different actuator classes on top depending on depth. Same gripper body, same D-coupler interface, just a different drive.

Tier	Actuator	Price (May 2026)	Why
Primary (T2)	DYNAMIXEL XW540-T260-R	USD 1,242 (AUD 1,925)	IP68 body, around 1.9 N·m, RS-485 bus, native present_current at 0.005 N·m per step held cleanly at stall. (ROBOTIS, 2026)
Value	DYNAMIXEL XM540-W270-R	USD 494 (AUD 766)	Same Dynamixel-X bus, identical telemetry, more torque (cont 2.12, stall 10.6 N·m). Drop-in replacement; only difference is no IP68 body, moot inside the T2 canister.
Deep-budget (Actually bought for prototype)	Feetech STS3250	EUR 64 (AUD 110)	50 kg·cm and 4.9 N·m stall on the SCS TTL bus with load, position, voltage, temperature feedback. The load percentage is the torque proxy. (FEETECH, 2026)
Rock-bottom	Feetech STS3215	EUR 26 (AUD 44)	Same SCS bus at 30 kg·cm and 2.94 N·m stall. Continuous 0.98 N·m sits below the 1.2 N·m design target but above the 0.6 N·m floor; adequate for intermittent grip-and-hold.
Fallback (T3)	Magnetic-coupling dry-pod	Build	No shaft penetration; depth set only by static pod seals; pole-slip torque is a built-in force limiter.

The forward sensing model

A measured current sample maps to an estimated per-finger tip force through the same chain as the drive model:

$$F = K_t \cdot (I_{meas} - I_0) \cdot \frac{i_g}{2} \cdot MA(P) \cdot \eta$$

K_t is the torque constant, I_0 is the no-load and friction current, $i_g = 2.667$ is the crown-pinion ratio, $MA(P)$ is the four-bar mechanical advantage at the current pose, and η is the calibrated efficiency. Since command and readout share the chain, one drivetrain characterisation closes the loop: commanding a current limit sets a force limit and the gear-protection ceiling at the same time.

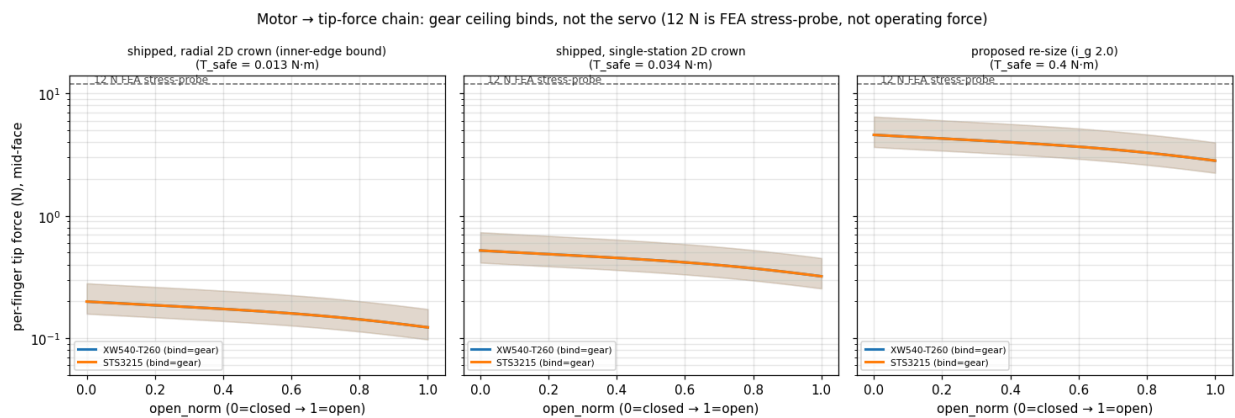


Figure 7.1 Motor to tip force chain across three drivetrain configurations. The gear ceiling is what binds, not the servo. 12 N is the FEA stress probe, not the operating force.

Limits of the sensor

- **Force only, no contact location.** Motor current is one aggregate number. It can't say where on the finger the contact sits, and it can't map pressure across the face.
- **Relative until calibrated,** and per unit (K_t varies 5 to 10 % unit to unit).
- **Gear-ceiling-bounded.** The sensed and commandable force is capped by T_{safe} . The sensor is most useful precisely as the limiter that keeps the printed gears alive.
- **Rank-only inheritance.** Our grip-texture model ranks object slip risk, but it doesn't give absolute hold force, so neither does a slip-margin readout built on it.

Where the per-pad sensors win

The motor-current approach gives one aggregate force across the whole gripper. The per-pad pressure sensors (innovation 2) give four spatially resolved readings, one per fingertip. For the archaeological use case, where the operator needs to know which finger is overloaded, the per-pad sensors carry information the motor-current chain throws away.

8. The drivetrain and the structural findings we found

The compact right-angle crown and pinion stage is the binding structural limit of the entire gripper.

At the small radii the housing allows, the printed teeth carry far less torque than a 12 N stress-probe grip would need. We found this in the gear FEA, and we're spelling it out here because it changes how the rest of the design has to work. The fix has three parts: face-width strengthening (done and build-verified), a module and radius re-size (a proposed engineered target that still needs CAD clearance validation), and the motor current limit (the sensing pivot) holding the safe ceiling in firmware. This is exactly why the magnetic-coupling fallback looks good in hindsight: its pole-slip torque scales with coupling diameter, not housing radius.

Method

A 2D plane-stress FEA of the real straight-flank tooth (motor/scripts/gear_fea.py, porting the grip Tier-2 Q4 machinery). A worst-case tip tangential load bends a cantilever tooth (root thickness s_{root} , height h , loaded face width = min of the two mating faces). We sample peak von Mises in the root band (sharp shoulder, so it's conservative; a real fillet would lower it).

Two failure modes we have to flag. First: the earlier framing presented a Lewis form-factor check against the FEA as cross-validation. It isn't. Lewis assumes a 14.5° or 20° involute spur tooth, and ours are straight-flank face teeth. Two wrong models that agree aren't the same as one right model being confirmed. Second: the crown gear is a 3D problem, not a 2D plane-stress tooth. The contact line sweeps radially under rotation, the load splits into tangential plus radial plus axial, and the disk-bending compliance matters. The 2D solver catches none of that. We added gear_fea_radial.py, which runs the same 2D FEA at five radial stations and takes the worst (inner-radius) one. That tightens the single-station ceiling by $2.6\times$ because the inner-radius tooth slice has a thinner s_{root} than the pitch radius. (Lewis, 1892)

T_{safe} : the input-shaft torque where the weakest tooth reaches 30 MPa

Geometry	Crown (single-station 2D)	Crown (radial 2D, inner edge)	Binding T_{safe}
As-was (PINION_T 4, CROWN_TOOTH_H 1.6)	0.02	n/a	0.02 N·m
Shipped (PINION_T 8, CROWN_TOOTH_H 3.0) ✓	0.034	0.013	0.013 to 0.034 N·m

Proposed re-size (CROWN_RC 11, m 1.83, 12/6 teeth) Δ	n/a	n/a	$\approx 0.40 \text{ N}\cdot\text{m}$
---	-----	-----	---------------------------------------

The sector gear (chunky m around 1.5) holds up well; the crown and then the pinion are the weak links. Doubling the face width roughly halves the root stress, but the tiny module dominates, so T_{safe} only climbs from 0.02 to 0.034 N·m, still well below the 0.94 N·m working torque a 12 N grip would need. Strengthening on its own isn't enough; the module-and-radius re-size is the only path that gets close to the working torque.

Per-finger force band the shipped gears can safely deliver

Ceiling	Per-finger force band
Radial 2D crown bound (binding)	0.14 to 0.28 N
Single-station 2D bound	0.35 to 0.73 N
Proposed re-size	4.2 to 8.7 N

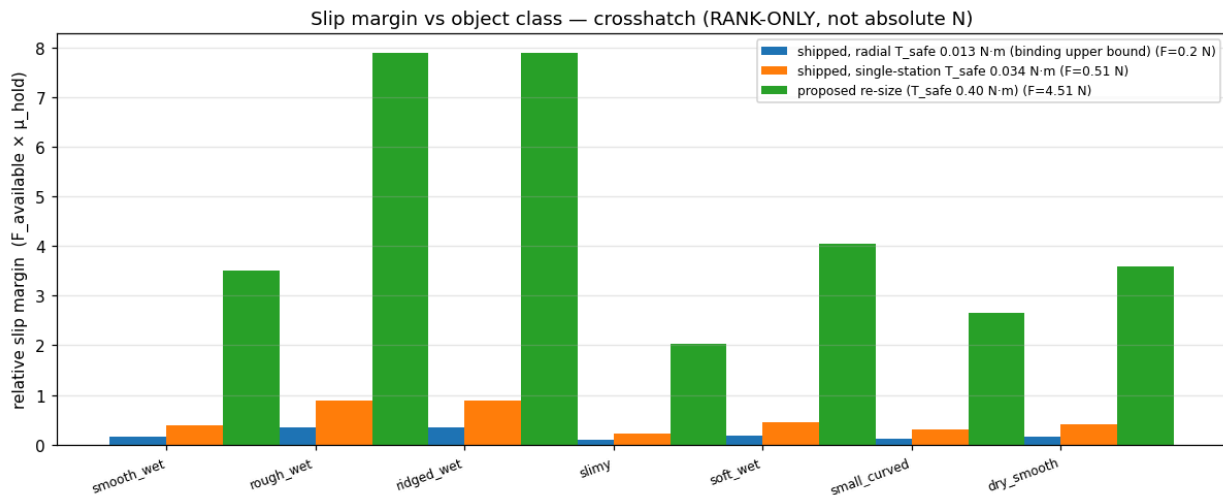


Figure 8.1 Slip margin across the 7-condition battery for the shipped crosshatch texture, at three drivetrain ceilings. Only the proposed re-size (green) clears margin ≥ 1.0 across every wet object class. That's the headline reason the gear re-size is on the roadmap and why bench testing is the gate.

Why we kept the 24/9 = 2.667:1 ratio

The rock-bottom budget servo (Feetech STS3215, continuous 0.98 N·m) sets the binding floor. At 2.667:1 it reaches 12 N at the tip continuously; drop to 2.2:1 and it falls to about 11.4 N, where the budget option no longer hits 12 N. So we went back and looked at the ratio, then kept it on purpose. The real change here is engineered teeth and a derived current ceiling, not a ratio tweak.

How this reshapes every other claim

- **12 N is a stress-probe load, not a structural mandate.** The fragility margins we quote at 12 N are worst-case-load margins, used to compare stiff and compliant finger candidates fairly.
- **The motor current limit enforces T_{safe} in firmware.** It's mandatory on every servo in the ladder (their stall torques beat T_{safe} by 10 to 30 times).
- **The magnetic-coupling fallback gets stronger.** A 60 to 80 mm magnetic coupling clears the stall band the geared stage can't.
- **The bench test is the only true ceiling.** Both 2D bounds are upper bounds. The straight-flank face teeth will edge-load and gall in PA12-GF instead of rolling cleanly. The printed-coupon torque-to-failure test is the only bench-grade answer.

9. Materials, design-for-additive, and the print profile

Why we are 100 % polymer

Galvanic corrosion has nothing to bite on in our gripper, because there's no metal anywhere inside it. Every pin, every snap, every gear, every shaft is printed. The only metal in the whole system lives inside the user's actuator canister, isolated from the seawater. That takes a whole class of failure modes (galvanic pitting, bushing dryout, dowel galling) off the table before it can even start.

ZERO bought hardware inside the gripper. No screws, no nuts, no bolts, no bushings. 17 printed parts, fastener-free assembly, tool-free disassembly.

Final-build material plan

The production build uses three materials, each assigned by what it has to do structurally:

Part	Qty	Final material	Rough filament	Why this material
enclosure	1	PA12-GF	80 to 110 g	Flooded gearbox body. Open front, integrated bottom flange (4 × M4 clearance holes), top link slots, 4 back-wall axle bosses with stepped bores, upper and lower journal bores for the input shaft, 8 bottom drains plus 2 side drains, 4 snap-clip catch windows.
front_cover	1	PA12-GF	20 to 30 g	Closes the open front. 4 inner-face bosses cap the axle-dowel heads. 4 integral cantilever snap clips latch into body side-wall windows. 3 × Ø1.8 mm cover vents.
drive_arm_R	1	PA12-GF	12 to 18 g	Right gear sector + crank arm. Rides on axle dowel pin_A_R. Counterbored C-eye exit for pin_C_R capture.
drive_arm_L	1	PA12-GF	12 to 18 g	Left gear sector + crank arm + integral crown gear on its +Z face. Rides on axle dowel pin_A_L. Counterbored C-eye for pin_C_L.
input_pinion_shaft	1	PA12-GF	8 to 12 g	ONE printed part: spur pinion + vertical shaft + integral capture collar + bottom D-coupler. Runs in two flooded journal bores;

				collar trapped between bore-mouth shoulders. Zero hardware.
follower_R, follower_L	2	PA12-GF	6 to 9 g each	Right and left B-to-D link bars. Counterbored D-eye exits for pin_D_R and pin_D_L capture.
finger_R, finger_L	2	Bambu TPU 95A HF	≈ 30 g each (ρ 1.22)	Fin Ray compliant jaws with the grip texture on the contact face. Must flex; print in TPU only. Profile in PRINT_PROFILE_P1S_TPU.md
pin_A_R, pin_A_L, pin_B_R, pin_B_L	4	PA12-GF	1 to 2 g each	Axle dowels (drive-arm + follower pivots). Plain headed dowels, no barb. Rigid geometric sandwich: head against cover boss, shoulder against bore step.
pin_C_R, pin_C_L, pin_D_R, pin_D_L	4	PETG-HF ★	1 to 2 g each	Finger snap pins (crank-coupler + follower-coupler joints). Barbed split tip. The lip locks in a rigid PA12-GF counterbore pocket in the receiving arm or follower.

Total filament for the production build comes to about 120 to 175 g of PA12-GF, 50 to 70 g of TPU 95A HF, and 4 to 8 g of PETG-HF for the finger pins. While we iterate, the test prints use PETG-HF for everything that ships in PA12-GF, because PETG-HF prints faster and matches PA12-GF dimensionally closely enough to fit-test the snap geometries before we commit to the slower PA12-GF print.

Why PA12-GF for everything rigid

PA12 is the lowest-water-uptake engineering nylon out there. Saturated seawater absorption sits at about 0.7 to 1.2 %, and the glass fill brings that down further. It doesn't hydrolyse. It holds tight dimensions even after long immersion. The 30 % glass fill gives roughly 3.5 to 5.5 GPa modulus and low creep, which is exactly what we want for the gearbox bodies, the link arms, the input shaft, and the axle dowels that have to keep their pivot geometry over months underwater.

PA12-GF replaces an earlier spec for carbon-filled Nylon 12 (PA12-CF). PA12-CF needs a 45 to 60 °C heated chamber, and the Bambu P1S we print on can't hold that. Without the chamber, PA12-CF lays down with poor inter-layer adhesion and warps as it cools. On its properties, PA12-GF is the right material (the glass fill plus PA12's naturally low water uptake covers the seawater duty cycle), but we'll be honest that it's marginal on the P1S, which has no actively heated chamber: glass-filled nylon is warp-prone on the larger parts, so it wants a closed chamber, a brim, and a 260 to 280 °C nozzle over a 90 °C bed with glue stick. Where that gets troublesome, the P1S-friendly substitutions that keep most of the rigidity are Bambu PAHT-CF or PETG-CF, or plain PETG-HF for a lower-spec build. The glass fill

also pulls dimensional drift below the carbon-filled grade, because glass fibres don't bend under print stress the way short carbon fibres do.

Why PETG-HF only for the finger snap pins

Glass fill trades ductility for stiffness. Unfilled PA12 prints at roughly 10 to 20 % elongation at break; the 30 % glass-filled grade drops that to about 3 to 8 % and removes the yield plateau. Filled PA12 fails by cracking instead of yielding. For one-time assembly snaps, we set the design allowable conservatively at 1.5 to 2.0 % strain.

The barbed split-tip pin that secures the fingers has to flex inward as it passes through the receiving bore. The split tip has a 0.9 mm proud barb (nominal), 1.0 mm worst-tight at the loose tolerance edge. Working the strain through the cantilever surface model $\epsilon = 3 \cdot t \cdot \delta / (2 \cdot L^2)$ on the effective slot length gives a nominal insertion strain of 2.50 % and a worst-tight insertion strain of 2.78 %. Both numbers run past PA12-GF's 1.5 to 2.0 % allowable. A finger snap pin printed in PA12-GF would crack on the way in. (Covestro LLC, n.d.)

PETG-HF (high-flow) has the same modulus, strength, and elongation as regular PETG; high-flow is just a processing variant for faster prints. Its one-time-assembly allowable is 2.5 to 3.5 %. The 2.78 % worst-tight strain sits inside that band with room to spare. So the four finger snap pins go in PETG-HF, marked with a star in the BOM, while every other part stays PA12-GF. The pull-out load on the snapped pin isn't carried by the pin material at all; it's taken by the rigid PA12-GF counterbore shoulder in the receiving arm or follower eye. The pin only has to survive the one-time insertion, and PETG-HF does that with margin.

The cover clip needed its own fix. The 4 cantilever snap clips on the front cover were originally tuned to PETG's elastic band, and they would have failed at 3.32 % worst-tight strain in PA12-GF. We re-tuned the geometry instead of the material: SNAP_Z0 dropped from 6.5 to 1.5 (lengthening the cantilever arm), and SNAP_ARM_T thinned from 2.8 mm to 2.0 mm. The new worst-tight strain is 1.36 %, well inside the PA12-GF allowable. gripper.py now asserts the worst-tight strain stays below 1.5 % and fails the build loudly if it doesn't.

Why TPU 95A HF for the fingers (and the hydrolysis question)

The Fin Ray grip principle is all about material compliance. The fingers have to flex visibly as they wrap an object; that's the whole mechanism. TPU at 95 Shore A is the right hardness. Softer TPU (85A) gives more compliance but loses grip force; harder TPU (98A) gives more grip force but loses the wrap. 95A is the standard middle ground for soft robotic grippers.

There are two TPU chemistries: ester-based and ether-based (polyether). They look identical, print the same way, and have almost identical mechanical properties. They differ in the one thing that matters here: hydrolysis resistance. Ester-based TPU hydrolyses in sustained or warm immersion. Left

underwater for weeks, an ester-based TPU finger crumbles from the inside out. Ether-based TPU is hydrolysis-stable, and it survives sustained marine immersion.

We print the fingers in Bambu TPU 95A HF. Two things drove the switch from the eSUN eTPU-95A we used earlier. First, and decisively, Bambu publishes ISO 527 mechanical data measured on printed specimens: an in-plane modulus of 9.8 MPa and tensile strength of 27.3 MPa, plus a through-Z modulus of 7.4 MPa and strength of 22.3 MPa. So the finger FEA now runs on a measured, direction-aware number instead of the old 40 MPa estimate. Second, the high-flow grade prints at roughly three times the throughput of standard 95A TPU, which matters on a finger this size. The honest cost of the switch is the hydrolysis story: Bambu doesn't state whether 95A HF is ester- or ether-based, so we lose the published polyether claim eSUN gave us. The measured saturated water uptake is low, about 1.08 %, which is encouraging but not a chemistry guarantee. That's exactly why the 7-day seawater soak test in the bench plan (§13) is now load-bearing rather than a formality: it's the gate that confirms the chosen filament survives sustained immersion. (Bambu Lab, 2023)

Materials ruled out, with the reason

- **PLA:** hydrolyses underwater. It loses strength over days to weeks of immersion. Reject for every part.
- **Untreated nylon (PA6, PA66):** absorbs 1 to 3 % water and more, swells, softens, and drifts dimensionally. Only the lowest-uptake nylon (PA12) with glass fill is acceptable.
- **PA12-CF (carbon-filled):** needs the 45 to 60 °C heated chamber the P1S can't hold; carbon fibre also contaminates the recycling stream. Glass-filled PA12 is the right answer.
- **PETG or PETG-HF for the rigid body:** an acceptable test-print material, but not the final production answer. PETG's modulus (~1.7 to 2.1 GPa) is below PA12-GF's (~3.5 to 5.5 GPa), and PETG also creeps under sustained load. For the gear arms and the input shaft running in flooded journal bores, the PA12-GF stiffness and low creep both matter.
- **PETG-CF or PETG-GF for the rigid body:** the fibre kills any recycling story, and the part geometry doesn't need the extra stiffness above PA12-GF.
- **PP, PP-CF, polycarbonate:** warp or delaminate on the P1S without an actively heated chamber.
- **Ester-based TPU:** hydrolyses. It crumbles in sustained marine immersion. Reject every ester grade.
- **TPU for any pin:** creeps under sustained axial load. The pin head relaxes and wallows the bore, and the snap retention is lost. PA12-GF for the axle dowels (no flex) and PETG-HF for the finger snap pins (one-time flex). No TPU.
- **Oil-impregnated bronze (Oilite) bushings:** the oil leaches out once flooded, leaving a dry porous bush. It isn't applicable here (we have no bushings), but it's worth recording: the input shaft runs in flooded printed journal bores, with no bushing of any kind.

Design for additive manufacturing

Every part is designed to standard FDM rules, so it prints reliably on any 0.4 mm-nozzle FDM machine with no exotic settings:

FDM rule	Target	How this design meets it
Minimum wall thickness	0.8 mm (2 perimeters); 1.5 mm functional	Enclosure walls 3.0 mm; snap-clip arm 2.0 mm; axle-boss walls 2.0 mm. Finger contact beam 1.2 mm (3 full perimeters at 0.4 nozzle, FEA-chosen to spread pressure).
Overhang angle	45° from vertical without support	Boxy parts print flat. Gear teeth and rib structures are in-plane. Snap-pin barb is a narrowing cone. Cover clips print pointing up.
Hole diameter	1 mm vertical, 2 mm horizontal	Pivot bores Ø5.2; drain holes Ø5; all well above the FDM minimum.
Mating clearance	0.3 mm per side	PRINT_CLEAR = 0.3 on every pivot; SNAP_CLEAR = 0.35 on snap engagements.
Sharp edges	Fillet or chamfer	Fin Ray rib-cell corners filleted R0.8; tip and teeth rounded; base chamfered for elephant-foot relief.

10. Running it underwater: the seawater audit

Seawater is the hard case. Our gripper is built to flood, so there's nothing to seal inside it. That means the audit is about polymer chemistry, no trapped air, net buoyancy, galvanic isolation, and creep at the load-bearing snaps.

Audit summary

Check	Verdict	Evidence
Material selection for sustained seawater	PASS (with per-part selection)	PA12-GF rigid + PETG-HF finger pins + ether-TPU fingers is correct. PLA rejected (hydrolysis), untreated nylon rejected (swells), ester-TPU rejected (hydrolysis).
Creep at load-bearing snap interfaces	RESOLVED	Finger-pin lips snap into a rigid counterbore pocket (radial confinement) bearing pull-out on a solid shoulder. SNAP_BARB_SEAT 0.30 to 1.2 mm. Capture is geometric, not preload-dependent.
Flooded correctness (no trapped air)	PASS (1 mitigation)	Every part is a single solid with exactly 1 shell. 8 bottom drains, 4 side, 4 snap windows, 4 axle bores, journal-bore clearance, 3 cover vents. Floods in any orientation.
Net buoyancy	PASS	Solid volume 98.5 cm ³ , dry mass around 124 g. Flooded net \approx +23 g in seawater (sinks gently, near-neutral).
Galvanic corrosion	PASS (trivial)	Zero metal in the gripper. All 8 pins and the input pinion shaft are printed; cover snaps; no screws, no nuts.
Drive coupler and actuator	SELECTED	Bottom D-flat coupler drives a smart serial-bus servo or the magnetic-coupling pod fallback. Current limit set to T_safe in firmware.

What the water actually does to the finger: the 3D crush FEA

The audit above settles the chemistry, the buoyancy, and the flooded geometry. We also ran the finger through a full three-dimensional pressure FEA to answer the blunt question a judge always asks: does the water itself crush it? The answer splits cleanly on one variable, whether the Fin Ray cells flood.

Flooded, the finger is fine at any practical depth. Water presses on every wetted surface at once, inside the cells and outside the skin, so the stress state is pure hydrostatic compression and the von Mises, shape-changing stress is essentially zero. The whole finger shrinks uniformly, about 1 % in linear size at 100 metres, without distorting. There's no depth in the archaeological range where a flooded finger is structurally at risk.

Trapped air is the dangerous case, and it's dangerous from shallow water. If the cells fail to flood and hold air at surface pressure, the depth pressure becomes a real mechanical load. The 3-D solve shows peak von Mises climbing to about 3.8 MPa at 30 metres and 12.6 MPa at 100 metres, and, more tellingly, the contact face sags by more than the finger's own thickness even at shallow depth, so the wrap is geometrically wrecked long before the material yields. The material itself yields beyond about 177 metres. This corrected an earlier two-dimensional estimate that read far lower: the 2-D model held the through-thickness strain at zero and hid the dominant collapse mode.

The design consequence is a procedure, not a redesign. Submerge the gripper fingers-down, let it soak for about thirty seconds, and cycle it open and closed once underwater to purge any trapped bubble. Because every part is a single solid with one shell and the cells drain through open channels, flooding is reliable, but it's a pre-dive step that matters from depth zero, not only on deep dives.

One detail carries over from the filament switch. The crush von Mises field is load-controlled, so it doesn't depend on the finger's stiffness, and the stress numbers above are identical to the earlier eSUN run. What the softer measured modulus changes is the displacement, which is roughly five times larger, and, through the measured 22.3 MPa through-Z strength, the material-yield depth, which moves from about 199 m to about 177 m.

Patrick Morrison's checklist

Five points from the WA Shipwrecks Museum site visit map directly onto our design:

Patrick's note	What we do
Fail gripping, not releasing.	Back-drivable chain means an unpowered servo releases. We add a holding-current threshold in firmware so the controller actively holds at the last commanded force. This is on the safety-features roadmap, surfaced explicitly.
Foam sensors are a red flag at depth (air pockets compress).	Tim MacDonald's oil-fill trick: clear garden hose with non-conductive oil around the sensor circuitry as a cheap pressure boundary. Plus the David Howard ambient-subtraction fix in code.
Marine grade stainless for any metal; epoxy coat prints.	Zero metal in the gripper. Epoxy coat (E6000 or marine epoxy) recommended on printed surfaces for the 1:1 production version.
Wet connects (Blue Robotics style) for ROV interface. (Blue Robotics, 2026)	Bolt-on M4 flange with adapter library (§11); wet connect terminates the umbilical at the ROV, not at the gripper.
Accept internal flooding, design around it.	The gripper is fully flooded by design. There is no sealed cavity to crush.

Pre-dive and post-dive checklist

Pre-dive:

- Confirm material (no PLA, no ester-TPU, no unfilled nylon)
- Confirm snap pins meet geometric capture

- Tug-test every pin and the cover
- Confirm drains and slots and windows are clear and the cavity floods in your dive orientation
- Cycle open and close in air
- Confirm the actuator is sealed and rated for depth
- Check buoyancy trim with the actuator fitted

Post-dive:

- Rinse well with fresh water (salt crystallises and abrades)
- Cycle to flush grit
- Check teeth and pins for wear and the TPU for any swelling
- Check every snap pin and the cover clips for engagement
- Replace any pins that have loosened

11. ROV integration and mounting interfaces

SoftSense isn't tied to one platform. The same gripper bolts onto whatever the operator already runs, using a small set of parametric printed adapters that all share the same gripper-side mate.

Common gripper-side mate

The bottom mounting flange carries 4× M4 clearance holes ($\text{Ø}4.5$) at (X, Z) coordinates (-38, 2), (38, 2), (-38, 18), (38, 18) on the $Y = -25$ face. That's a rectangular 76×16 mm pattern with its centroid at (X = 0, Z = 10). The drive shaft comes out of the same face at (X = -12, Z = 10.52), 12 mm offset in X from the bolt centroid. Adapters have to clear that bore ($\text{Ø} \geq 14$ H7 to seat the lip seal). It's the single source of truth in `motor/cad/adapters/_base.py`.

Adapters modelled




Figure 11.1 Reach Bravo 7 wrist adapter (RB-1054, $\text{Ø}71$ mm plate, 6× M6 + 2× $\text{Ø}3$ dowels). The adapter shown is the Fugro / Woodside reference target.




Figure 11.2 BlueROV2 bottom-Newton adapter, modelled to the published BR2 bolt pattern.




Figure 11.3 ISO 9409-1 Type 50-4-M6 cobot flange adapter. It drops straight into any standard cobot end effector mount. (ISO, 2004)

Adapter	Target	Status
adapter_bravo7	Reach Bravo 7 RB-1054 wrist	Modelled. Bolt circles are estimates pending NDA dimensions.
adapter_iso9409_50_4_M6	ISO 9409-1 Type 50-4-M6 cobot flange	P0-ready. Standard published dimensions.
adapter_iso9409_80_6_M8	ISO 9409-1 Type 80-6-M8 cobot flange	Modelled.
adapter_br2_bottom_newton	BlueROV2 bottom-panel Newton footprint	P0-ready.
adapter_br2_roof_rack	BlueROV2 roof rack	Modelled.
adapter_br2_payload_skid	BlueROV2 payload skid	Modelled.
adapter_iso13628_d_handle	ISO 13628 D-handle override (diver-operable) (ISO, 2002)	Modelled.

Two non-negotiables

1. **T2 still needs a thin pressure canister**, even for the IP68 XW540 (IP68 is rated for 1 m freshwater only; 30 m is 3.1 bar plus seawater corrosion). The flooded gripper itself has no shaft seal and no dry cavity, so the sealing job falls entirely on the actuator.
2. **The motor current limit must be set to T_safe in firmware**. Without it, a strong servo's stall torque shears the printed teeth. The current limit is both the grip-force command and the gear-protection ceiling at once.

The control electronics: an offline ESP32 “brain”

Every other part of this document treats the gripper as hardware. This part is the hardware that drives it. Rather than a topside computer and a tether junction box, we put the whole controller on one small board that travels with the gripper: a Waveshare General Driver for Robots, built around an Espressif ESP32-WROOM-32 (Core Electronics SKU WS-23730) (Waveshare, 2026; Espressif Systems, 2025).

The board boots into its own Wi-Fi access point, serves a phone web page, and is the servo's bus master, all at once. There is no router and no internet. You power it from roughly 12 V (3S), it broadcasts the network “Gripper”, you join it from a phone, a captive-portal page opens at <http://192.168.4.1/>, and you get OPEN and CLOSE buttons with live position, load, voltage, and temperature underneath. A 0.91-inch OLED on the board mirrors the network name, the IP address, and the current position. This replaces an earlier Orange Pi single-board-computer controller we prototyped: the ESP32 collapses the access point, the web server, and the servo bus master into one microcontroller that boots in about a second and draws a fraction of the power.

The controller is also where the actuator-as-sensor thesis (Section 7) becomes a real readout. It commands the Feetech servo and reads that servo's own load, voltage, and temperature back over the same three-wire bus, so the “force sensor” of Section 7 is exactly the number this board already streams to the phone. Setting a current or torque limit in firmware is what enforces the printed gears' safe ceiling T_{safe} (Section 8): the same command that closes the gripper also protects the drivetrain.

Wiring and power. The Feetech STS3250, or the lower-cost STS3215, plugs straight into the board's ST3215 bus-servo port with a single 3-pin cable — no separate bus adapter. The half-duplex serial bus runs on GPIO18 (RX) and GPIO19 (TX) at 1,000,000 baud, addressing the servo at ID 1, the Feetech factory default (FEETECH, 2026). The one rule that matters before power-up is to match the supply voltage to the servo — a 12 V STS3250 wants about 12 V (3S); a 7.4 V STS3215 wants about 7.4 V, not 3S — and to check the 3-pin connector orientation (GND / VCC / signal), because reversed power can destroy the servo.

Calibration that survives a reboot. The open and closed positions ship as placeholders (raw encoder counts 1024 and 3072 of 0–4095). On first run you move the gripper to its fully-open pose and tap “Set current as OPEN”, then to fully-closed and tap “Set current as CLOSE”. Those values, along with the move speed and acceleration, are written to the ESP32's non-volatile flash (NVS) and survive power cycles, so the gripper is calibrated once on the dock and then just works.

The control surface. The board exposes the same small HTTP API the earlier controller used, so any client — the built-in web page, a laptop script, or an ROV topside — drives it the same way:

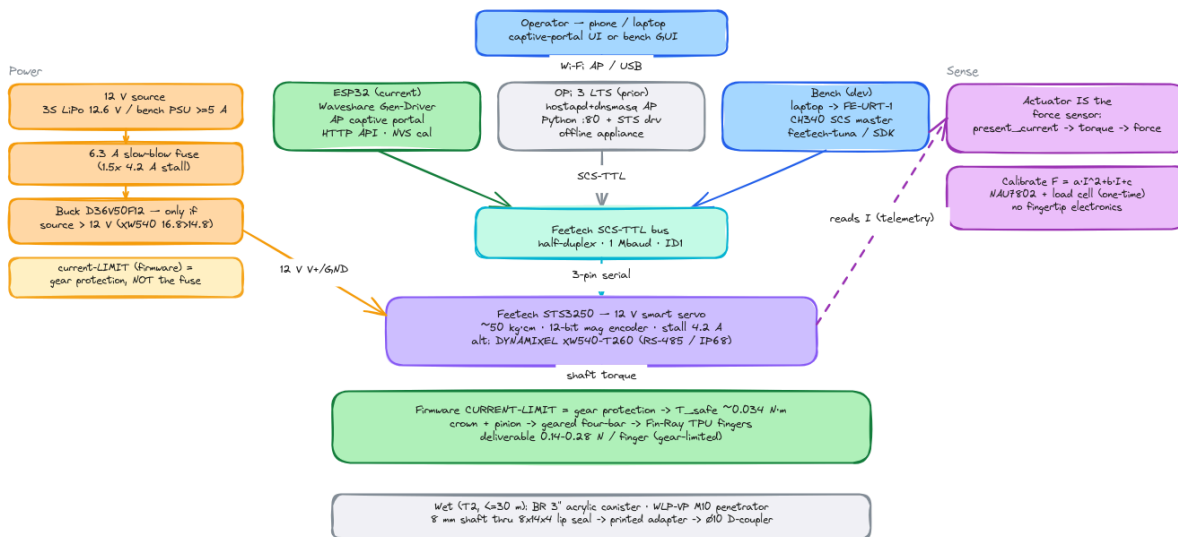
- GET /api/status - JSON: connected, position, load, voltage, temperature, plus the stored open / close / speed / acceleration.
- POST /api/open and /api/close - move to the calibrated end positions.
- POST /api/goto?pos=N - move to a raw position, 0–4095.

- POST /api/torque?on=0|1 - enable or release holding torque.
- POST /api/calibrate?which=open|close - save the current position as that endpoint.
- POST /api/config?speed=&acc=&open=&close= - set and persist parameters.

What is proven. The firmware is written and compiles cleanly under PlatformIO and the Arduino-ESP32 framework, with the OLED driven by Adafruit's SSD1306 and GFX libraries (Arduino S.r.l., 2026; PlatformIO Labs, 2026; Adafruit Industries, 2026). It is not yet bench-validated on the board, which is on order. The bus pins, baud rate, and SMS_STS protocol are taken from Waveshare's own firmware for this exact board, so they should be correct out of the box (Waveshare, 2026); the open and closed directions and positions are placeholders until they are calibrated on first power-up. Like the rest of SoftSense, the controller is offline by design — the only way to change it is to plug it into a laptop over USB and re-flash — which is exactly the fool-proof, no-maintenance posture the gripper itself is built around.

gripper-cad — Full Electronics Stack

Operator → controller → SCS-TTL bus → STS3250 servo (↳ the force sensor) → current-limited drivetrain → Fin-Ray fingers



12. TRL self-assessment and operating boundaries

At Woodside's request we say where SoftSense sits today on the NASA Technology Readiness Level scale and what would move it forward. We also spell out the operating boundaries so a customer knows what we are and aren't promising. (Mankins, 1995)

Current TRL position

TRL	Level description	SoftSense status
1	Basic principles observed and reported.	Done. Pressure feedback principle established.
2	Technology concept and application formulated.	Done. 4-finger pressure-sensing gripper specified.

3	Experimental proof of concept.	Done. Working scaled prototype with all four innovations integrated; bench-tested against the 4-object artefact set.
4	Technology validated in lab.	Partially done. FEA validates fingers and texture; bench plan in §13 closes the remaining gaps. Soak and hyperbaric tests are scheduled before Korea.
5	Technology validated in relevant environment.	Future. Requires pool dive on a BlueROV2 chassis with a real artefact analog at depth.
6 to 9	Demonstration, full system, qualified, real-world deployment.	Future.

Where we sit now: TRL 3 to 4. The valley of death (TRL 3 to 7) that Woodside warned us about is the gap we're crossing right now. The bench plan in §13 is how we cross it.

Operating boundaries (1:1 production target)

Parameter	Boundary
Operating depth	0 to 50 m (T2 actuator + 30 m canister rating) for the shipped configuration; T3 magnetic-coupling pod extends to 200+ m as a future option.
Operating temperature	-2 to +30 °C water temperature (covers WA continental shelf summer and Antarctic Peninsula summer).
Residency time in water	Continuous immersion up to 30 days per the PETG soak-test result; rinse with fresh water between deployments.
Salinity	Full marine (35 ppt) to brackish; fresh water also acceptable.
Payload	2 to 5 kg per Patrick Morrison's archaeology use case; structural test required to certify above 5 kg.
Grip force per finger	0.14 to 0.73 N (shipped gears, conservative) or 4.2 to 8.7 N (proposed re-size).
Maintenance interval	Visual inspection every dive; pin replacement annually or after visible creep; finger replacement at 10,000 grip cycles or first visible TPU crack.
Storage	Dry, room temperature, out of direct UV; PETG yellows in sustained sun over months.

13. Test data from the FLL bench set

We tested the scaled prototype against four artefact-analog objects we picked for their range of shapes: a bone, an anchor, a vase, and a treasure chest. The bone and the anchor are both long. The anchor has natural grip features (holes), while the bone has none. The chest and the vase are both stout. The

chest has rectangular grooves that are easy to grip, and the vase is smooth with only the head to grab onto.

Results (4 finger vs 2 finger configuration)

Object	Attempts (4 fingers)	Attempts (2 fingers)	Pressure to grip
Bone	2	4	26 gf
Anchor	5	4	96 gf
Vase	3	5	411 gf
Chest	3	3	288 gf

Reading the results

The 4-finger configuration beats 2-finger on smooth or round shapes (vase) and on shapes with poor grip features (bone), and it ties on regular-feature shapes (anchor, chest). That's how it goes for soft pads with sensors: when the artefact has natural features for a rigid jaw to bite, you can't really do better than 3 attempts. When the artefact is smooth or featureless, the 4-pad layout wraps it and the 2-jaw layout slips off it. The bone is the clearest example: 2 attempts versus 4.

The pressure values span an order of magnitude, from 26 gf (bone) to 411 gf (vase). That's direct empirical support for the per-artefact preset architecture in innovation 2: a single global pressure threshold would either crush the bone (if set high enough to lift the vase) or fail to lift the vase (if set low enough to spare the bone). Different shapes need different presets, and the per-pad sensor system can deliver them.

14. Expert and institutional validation

SoftSense was reviewed by working subsea engineers, a practising maritime archaeologist, and three of Western Australia's major subsea-industry operators. The design's credibility rests on these people examining it seriously and pushing back, and most of the major decisions in this document trace directly to one of the conversations below. We name, for each reviewer, the specific thing in the build that changed because of them.

Tim MacDonald, subsea engineer, Inkfish (formerly DSV Limiting Factor). Our first external reviewer, called after Regionals. He confirmed the novelty claim we lead with in innovation 2: soft-jaw and silicone-tipped grippers already exist, so the force feedback on a cheap printed gripper is the part that is actually new. We used that to reframe the whole pitch around sensing rather than softness. His economics argument, that the biggest cost in archaeology is time and a gripper should be a cheap field-swappable kit, is the direct source of innovation 4 and the per-module cost figures in section 3. His

materials steer (stay all-polymer, avoid titanium and metals for cost and corrosion, pick a low-water-absorption plastic) set the constraint that section 9 works inside. His oil-filled enclosure trick, non-conductive oil sealed in a clear hose instead of a thick air-filled box, is the pressure-boundary answer we carry into the electronics housing and cite again in Patrick Morrison's checklist (section 10). We also adopted his outreach etiquette, sending a calendar invite to every expert after him.

Patrick Morrison, Assistant Curator of Maritime Heritage, WA Shipwrecks Museum. Hosted our 14 May 2026 site visit and grounded the entire problem statement in section 1. He confirmed that the wrecks WA has found recently sit beyond recreational diving depth, that the museum legally cannot excavate them, and that safe recovery from 50 to 60 metres and beyond needs a robot with actuator control. We used his mission-cost figures (boat hire, dive team, the asymmetry of a two-thousand-dollar tool failure wasting a million-dollar mission) as the cost case in section 1, and quoted his "the clock starts the moment you leave port" line directly. His engineering notes map onto the build one to one, which is why section 10 carries them as a checklist: his warning that foam and gas-filled pads compress at depth reinforced the move to motor-current sensing; his "fail gripping, not releasing" point became the firmware holding-current threshold; his wet-connect advice fed the section 11 interface; his marine-grade-stainless and epoxy-coat notes shaped the 1:1 production guidance; his 2 to 5 kg payload figure is the boundary stated in section 12; and his "accept internal flooding, design around it" advice matches our fully-flooded, no-sealed-cavity architecture.

David Howard, Principal Research Scientist, CSIRO (Regionals mentor). Identified the single confounder that redirected our sensing strategy. As the gripper descends, ambient water pressure rises by about one atmosphere every ten metres, so a fingertip pressure sensor reads a squeeze that is not there and needs a second ambient sensor plus a software correction just to recover the true value. We used that observation to move the primary force channel into the actuator's motor current (section 7), which tracks how hard the finger pushes the object rather than the water column above it, and recorded the ambient-subtraction fix as the fallback for any pad-based reading.

Fiona Stachowiak and the Woodside subsea team, Woodside Energy. Gave us a full engineering and industry-framing review on 20 April 2026. We used their water-absorption warning, drawn from real bushing-swell experience, to trigger the material reevaluation in section 9 (carbon-filled nylon out, PA12-GF in) and to make the seven-day soak test in section 13 load-bearing rather than a formality. Their magnetic-coupling reference (Advanced Navigation's Hydrus) became the T3 fallback actuator in section 7 and the drivetrain argument in section 8. Their two explicit asks, a TRL self-assessment and stated operating boundaries, are why section 12 exists in the form it does. We hold every innovation to their filter: is it novel, does it create value, does it solve a real problem.

Elaine Pankhurst, Fugro Australia. Met with us on 23 April 2026 and confirmed that our pressure-feedback approach parallels Fugro's own ROV valve-control work, which is independent validation that the core idea holds at industrial scale. The Reach Bravo 7 wrist adapter modelled in section 11 is the Fugro and Woodside reference mount, so their platform is the integration target we

designed an adapter against. She arranged Fugro's sponsorship of the Korea campaign and routed our documentation internally to Simon, a Fugro ROV manager, and the Fugro tooling team.

Damien Singh, Total Marine Technology. Took SoftSense into TMT and put the documentation in front of the company's CEO for engineering review on 23 April 2026, opening TMT's mentorship and backing of the project.

Industry support

The Korea Open campaign is backed by Western Australian subsea-industry sponsors who fund the trip and connect the team to working engineers:

- Fugro
- Total Marine Technology
- Pulse Technology Hub
- EFFEE On Site Robotics
- CP Maritime

Also contacted

Experts approached during the Nationals round:

- Dr John McCarthy
- Associate Professor Jonathan Benjamin
- Chelsea Wiseman
- Michael O'Leary (UWA)
- Jeremy Leach
- Ingrid Ward
- Hiro Yoshida

Institutions approached:

- Australasian Institute for Maritime Archaeology
- Minderoo-UWA Deep Sea Research Centre
- Western Australian Museum

Post-Nationals outreach channels open or in flight:

- Australian National Maritime Museum
- International Marine Contractors Association (IMCA)
- Schilling Robotics / TechnipFMC
- UWA Engineering (Mechanical, Mechatronics, Oceans Graduate School)
- Curtin Underwater Sensing and Robotics Lab
- CSIRO Oceans and Atmosphere

- Australian Institute of Marine Science (AIMS)
- NOAA Ocean Exploration
- OpenROV / Sofar Ocean
- Monterey Bay Aquarium Research Institute (MBARI)
- Oceaneering

15. Scaling to artefact size: the 1.5x and 2.0x build

The gripper this document describes is small on purpose. At full open the fingertips part by about 124 millimetres, and the drivetrain delivers well under a newton per finger. That's the right size for the FLL bench and the scaled test objects, but it's too small and too gentle to recover a real amphora, anchor fluke, or ballast stone off the seafloor. Every archaeologist we showed it to asked the same question: can you make it bigger? The answer is yes, and the way we built it means bigger is one parameter, not a redesign.

One parameter, the whole gripper

A single environment variable, GRIPPER_SCALE, multiplies every linear dimension of the model by a factor k . We built and re-simulated two larger sizes from it, 1.5x and 2.0x. The word that matters is self-similar: the linkage, the gears, the finger walls, the enclosure, the input shaft, the mounting flange, and the snap clips all scale together. Because the lengths, the radii, and the wall thicknesses grow by the same factor, every ratio and every angle is preserved. This is true mechanical similitude. We checked it: the finger rotation is identical at -20° at all three sizes, and the jaw opening scales exactly with k .

A few things are held constant on purpose, because they're set by physics or by the printer rather than by the size of the part. The print and running clearances stay at about 0.3 millimetres, because they're governed by the 0.4 mm nozzle, not the part. The fastener sizes stay put. An M4 bolt is an M4 at any scale, and only the positions of the bolt holes move. The flood and drain hole radii are sized by bubble and surface-tension physics, so they hold while their positions scale. And at a scale factor of 1.0 the scaled model reproduces the shipped gripper byte for byte, so the original design is preserved as the reference.

Size and reach

Here's what the two larger sizes actually measure, read straight from the scaled CAD. All parts still fit the 256 mm bed of the Bambu P1S at both sizes.

Quantity	1.0x	1.5x	2.0x
----------	------	------	------

Finger blade length	90 mm	135 mm	180 mm
Jaw opening at the fingertips	124 mm	186 mm	248 mm
Jaw opening at the base	62 mm	93 mm	124 mm
Closed jaw gap	9.9 mm	14.8 mm	19.7 mm
Printed mass (filament, 100 % dense)	~65 g	~0.23 kg	~0.54 kg
Print plates on the 256 mm P1S bed	3	3	4

Why the bigger finger actually grips

This is the part that makes the scale-up worth doing rather than just enlarging a drawing. The Fin Ray finger is a compliant blade with thin internal walls. An earlier study scaled the blade on its own and held the wall thicknesses fixed, and it found that above about 1.1× the finger went floppy. The walls were now too thin relative to the bigger blade, so the finger bent instead of squeezing, and its universal grasp score collapsed from 0.65 at 1× to 0.44 at 1.5× and 0.37 at 2×. That's the failure mode this document's earlier finger-scale figures (0.7×, 1.0×, 1.6×) were warning about.

Self-similar scaling fixes exactly that, because the walls scale with the blade. The wall-to-blade stiffness ratio stays constant, so the bigger finger is as stiff relative to its size as the validated 1× finger. On a matched modulus the universal grasp score now holds flat, at 0.63, 0.60, 0.62 across the three sizes, and the finger reaches the 12 N stress-probe at every size. On the softer measured Bambu modulus the score edges up rather than down as the finger grows, because the larger finger develops more absolute grip at the same proportional closure. Either way, the floppy-finger failure of naïve scaling is gone.

We're precise about what this does and doesn't fix. It fixes the under-grip, floppy failure mode, but it doesn't change the conformance ceiling. A passive single-piece finger still can't curl tightly around a small round object at any scale. That's an architecture limit, not a size limit. And the von Mises safety margin dips a little as the finger grows, from about 9× at 1× to about 5× at 2× on flat objects, but it stays comfortably above yield.

Force still follows the gears, not the size

Scaling changes the structural and force numbers by clean power laws, and we confirmed them to three significant figures on the self-similar gear FEA. The gear ceiling T_{safe} grows with the cube of scale, the deliverable per-finger force grows with the square of scale, the mechanical advantage falls as $1/k$, and the gear ratio and efficiency are scale-invariant because the tooth counts and angles are held.

Quantity (and how it scales)	1.0×	1.5×	2.0×
Gear ceiling T_{safe} ($\propto k^3$)	0.013 N·m	0.044 N·m	0.105 N·m

Per-finger force band ($\propto k^2$)	0.14–0.28 N	0.31–0.64 N	0.55–1.13 N
Smallest servo over-torque vs T_{safe}	224×	67×	28×
Underwater crush peak νM at 100 m	12.6 MPa	12.6 MPa	12.6 MPa

The reading is honest but bounded. The selected servos still clear the gear ceiling by between 28× and 725× at every size, so the gripper stays gear-limited and the firmware current limit remains its protection, exactly as at 1×. Force grows about four-fold at 2×, but off a sub-newton base, so a 248 mm jaw that can reach around an amphora still only squeezes it with around a newton. If the real need is squeezing strength, scaling is only half the answer. The other half is the proposed gear re-size in §8, which raises T_{safe} by ten to thirty times independent of size. Combined with 2× scaling it would put the per-finger force into the multi-newton range. That's the next design decision, and we've deliberately left it as a decision rather than folding it in here.

Underwater and grip at the new sizes

The underwater crush result is scale-invariant, which is what similitude predicts. The von Mises field and the material-yield depth don't move with size. Peak von Mises is about 12.6 MPa at 100 m and the through-Z material yields near 177 m at every scale, and only the displacement scales with k . A flooded finger stays at essentially zero von Mises at any size and any practical depth.

The wet grip is more nuanced, and our own re-analysis corrected an over-simple first assumption. The cited-physics core of the grip model, the elastomer-friction and drainage terms, is scale-invariant to within half a percent, because the crosshatch channels grow to 0.81 mm at 1.5× and 1.08 mm at 2.0×, both well above the roughly 0.3 mm where wet drainage saturates. But the full surrogate score softens by about 12 % at 1.5× and 22 % at 2.0×, because the bigger posts mean fewer skin-breaking edges per unit area at a fixed print resolution. That's a real if low-confidence effect rather than an artefact. The ranking and the printability are preserved, and a finer-relative texture could be reintroduced at 2× to recover the edge density if a future build wants it.

What changes when you scale, and what to watch

5. **Assembly gets stiffer to seat.** The snap-pin and cover-clip insertion force scales with the square of scale, roughly 2.25× at 1.5× and about 4× at 2×, even though the bending strain is scale-invariant, so the clips don't over-strain, they just take more push. At 1.5× the parts are still hand-pressable. At 2× expect to need a clamp or a small arbor press to seat the snap pins and click the cover home.
6. **The mounting interface scales with everything else.** The flange bolt pattern and the D-coupler grow with the rest of the gripper, so the seven printed mounting adapters in §11 have to be re-scaled to re-mate. At 2× the flange bolts are worth stepping up from M4 to M5 or M6 for the higher loads.

7. **The materials and the build sequence don't change.** Same PA-class rigid parts, PETG snap pins, and TPU fingers, same orientations, and the input shaft still installs from below, pinion first. A larger gripper is a bigger print, not a different assembly.




Figure 15.1 The shipped 1× gripper beside the 2.0× self-similar build, both in the open pose. Every linear dimension, including linkage, gears, finger walls, and enclosure, scales by the same factor, so the geometry, the gear timing, and the wrap behaviour are preserved while the jaw opening grows from 124 mm to 248 mm at the fingertips.

One caveat on provenance. These re-runs were done locally at screen fidelity, because the high-resolution simulation node was unavailable for the campaign. The comparative rankings are sound, since the self-similar 1× run reproduces the shipped finger's score, which anchors the comparison, but a publication-grade re-run on the render node is the follow-up, and a fresh blade-only control at the larger sizes would make the stiff-versus-floppy contrast fully airtight. We flag it rather than bury it.

16. Sources and citation provenance

Every quantitative claim in this document traces to one of three places: a dimension measured from the live parametric CAD model, a value taken from a published datasheet or standard, or a result from a reproducible simulation in this project's code. The build123d model (Maitland, 2025) is the single source of truth for every dimension; the live CAD is on Onshape

(cad.onshape.com/documents/47a3be0d6a2fdc65e8e54697), the mission code and team material are public (lebob.com.au; github.com/prawny-boy/FLL-Lebob-Unearthed), and the references below are formatted in APA 7th edition.

Personal communications. Our expert reviewers are cited in text only, per APA style, and are not listed below: Patrick Morrison (Assistant Curator of Maritime Heritage, WA Shipwrecks Museum; site visit 14 May 2026); Tim MacDonald (subsea engineer, Inkfish); David Howard (Regionals judge and mentor); Fiona Stachowiak and the Woodside Energy subsea team (review 20 April 2026); Elaine Pankhurst (Fugro Australia, 23 April 2026); and Damien Singh (TMT, 23 April 2026).

Friction, wet adhesion, and contact physics

- Baik, S., Kim, D. W., Park, Y., Lee, T.-J., Bhang, S. H., & Pang, C. (2017). A wet-tolerant adhesive patch inspired by protuberances in suction cups of octopi. *Nature*, *546*(7658), 396–400. <https://doi.org/10.1038/nature22382>
- Briscoe, B. J., & Tabor, D. (1975). The effect of pressure on the frictional properties of polymers. *Wear*, *34*(1), 29–38. [https://doi.org/10.1016/0043-1648\(75\)90306-3](https://doi.org/10.1016/0043-1648(75)90306-3)
- Drotlef, D.-M., Stepien, L., Kappl, M., Barnes, W. J. P., Butt, H.-J., & del Campo, A. (2013). Insights into the adhesive mechanisms of tree frogs using artificial mimics. *Advanced Functional Materials*, *23*(9), 1137–1146. <https://doi.org/10.1002/adfm.201202024>
- Federle, W., Barnes, W. J. P., Baumgartner, W., Drechsler, P., & Smith, J. M. (2006). Wet but not slippery: Boundary friction in tree frog adhesive toe pads. *Journal of the Royal Society Interface*, *3*(10), 689–697. <https://doi.org/10.1098/rsif.2006.0135>
- Fwa, T. F., Kumar, S. S., Anupam, K., & Ong, G. P. (2009). Effectiveness of tire-tread patterns in reducing the risk of hydroplaning. *Transportation Research Record*, *2094*(1), 91–102. <https://doi.org/10.3141/2094-10>
- Persson, B. N. J. (2001). Theory of rubber friction and contact mechanics. *The Journal of Chemical Physics*, *115*(8), 3840–3861. <https://doi.org/10.1063/1.1388626>
- Reynolds, O. (1886). On the theory of lubrication and its application to Mr. Beauchamp Tower's experiments. *Philosophical Transactions of the Royal Society of London*, *177*, 157–234. <https://doi.org/10.1098/rstl.1886.0005>
- Tramacere, F., Beccai, L., Kuba, M., Gozzi, A., Bifone, A., & Mazzolai, B. (2013). The morphology and adhesion mechanism of Octopus vulgaris suckers. *PLoS ONE*, *8*(6), Article e65074. <https://doi.org/10.1371/journal.pone.0065074>

Materials, mechanical engineering, and standards

- Bambu Lab. (2023). Bambu filament technical data sheet V1.0: TPU 95A HF. Retrieved June 24, 2026, from https://store.bbblcdn.eu/s8/default/16df21baf482453999b3dbb61cc110e7/Bambu_TPU_95A_HF_Technical_Data_Sheet.pdf
- Covestro LLC. (n.d.). Snap-fit joints for plastic: A design guide. Retrieved June 24, 2026, from <https://solutions.covestro.com>
- International Organization for Standardization. (2002). Petroleum and natural gas industries — Design and operation of subsea production systems — Part 8: Remotely operated vehicle (ROV)

interfaces on subsea production systems (ISO 13628-8:2002).

<https://www.iso.org/standard/26482.html>

International Organization for Standardization. (2004). Manipulating industrial robots — Mechanical interfaces — Part 1: Plates (ISO 9409-1:2004). <https://www.iso.org/standard/36578.html>

International Organization for Standardization. (2019). Plastics — Determination of tensile properties — Part 1: General principles (ISO 527-1:2019). <https://www.iso.org/standard/75824.html>

Lewis, W. (1892). Investigation of the strength of gear teeth. Proceedings of the Engineers' Club of Philadelphia, 10, 16–23.

Mankins, J. C. (1995). Technology readiness levels: A white paper. National Aeronautics and Space Administration, Office of Space Access and Technology.

Actuators, electronics, and hardware

Adafruit Industries. (2026). Adafruit_SSD1306 and Adafruit-GFX-Library: Arduino libraries for monochrome OLED displays [Computer software]. GitHub. https://github.com/adafruit/Adafruit_SSD1306

Arduino S.r.l. (2026). Arduino IDE (Version 2.3.10) [Computer software]. <https://www.arduino.cc/en/software>

Blue Robotics. (2026). Cylindrical locking-series watertight enclosures [Product documentation]. Retrieved June 24, 2026, from <https://bluerobotics.com/product-category/watertight-enclosures/locking-series/>

Espressif Systems. (2025). ESP32 series datasheet (Version 5.2). https://www.espressif.com/sites/default/files/documentation/esp32_datasheet_en.pdf

FEETECH. (2026). STS3250 and STS3215 serial bus servos [Product datasheets]. Shenzhen FEETECH RC Model Co. Retrieved June 24, 2026, from <https://www.feetechrc.com>

PlatformIO Labs. (2026). PlatformIO [Computer software]. <https://platformio.org>

ROBOTIS. (2026). DYNAMIXEL XW540-T260 and XM540-W270 [Product e-manuals]. Retrieved June 24, 2026, from <https://manual.robotis.com/docs/en/dxl/x/>

Waveshare. (2026). General driver for robots [Product wiki]. Retrieved June 24, 2026, from https://www.waveshare.com/wiki/General_Driver_for_Robots

Subsea robotics and underwater archaeology

Antonelli, G. (2014). Underwater robots (3rd ed.). Springer. <https://doi.org/10.1007/978-3-319-02877-4>

Bell, K. L. C., Johannes, K. N., Kennedy, B. R. C., & Poulton, S. E. (2025). How little we've seen: A visual coverage estimate of the deep seafloor. *Science Advances*, 11(19), Article eadp8602. <https://doi.org/10.1126/sciadv.adp8602>

Bogue, R. (2015). Underwater robots: A review of technologies and applications. *Industrial Robot*, 42(3), 186–191. <https://doi.org/10.1108/IR-01-2015-0010>

Dziak, M. (2025). Underwater archaeology. EBSCO Research Starters. <https://www.ebsco.com/research-starters/anthropology/underwater-archaeology>

Ivanega, D., & Szczepanek, M. (2024). Assessing damage and predicting future risks: A study of the Schilling Titan 4 manipulator on work-class ROVs in the offshore oil and gas industry. *Ocean Engineering*, 291, Article 116282. <https://doi.org/10.1016/j.oceaneng.2023.116282>

- Petillot, Y. R., Antonelli, G., Casalino, G., & Ferreira, F. (2019). Underwater robots: From remotely operated vehicles to intervention-autonomous underwater vehicles. *IEEE Robotics & Automation Magazine*, 26(2), 94–101. <https://doi.org/10.1109/MRA.2019.2908063>
- Schmidt Ocean Institute. (n.d.). Why we sample. Retrieved June 24, 2026, from https://schmidtocean.org/education/why_we_sample/
- Sivčev, S., Coleman, J., Omerdić, E., Dooly, G., & Toal, D. (2018a). Underwater manipulators: A review. *Ocean Engineering*, 163, 431–450. <https://doi.org/10.1016/j.oceaneng.2018.06.018>
- Sivčev, S., Rossi, M., Coleman, J., Dooly, G., Omerdić, E., & Toal, D. (2018b). Fully automatic visual servoing control for work-class marine intervention ROVs. *Control Engineering Practice*, 74, 153–167. <https://doi.org/10.1016/j.conengprac.2018.03.005>
- UNESCO. (n.d.). Access to underwater cultural heritage. Retrieved June 24, 2026, from <https://www.unesco.org/en/underwater-heritage/access>

Software and simulation tooling

- Dawson-Haggerty, M. (2026). trimesh (Version 4.12.2) [Computer software]. <https://trimesh.org>
- Geuzaine, C., & Remacle, J.-F. (2009). Gmsh: A three-dimensional finite element mesh 55generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11), 1309–1331. <https://doi.org/10.1002/nme.2579>
- Gustafsson, T., & McBain, G. D. (2020). scikit-fem: A Python package for finite element assembly. *Journal of Open Source Software*, 5(52), Article 2369. <https://doi.org/10.21105/joss.02369>
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., . . . Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- Maitland, R. (2025). build123d: A Python CAD programming library (Version 0.10.0) [Computer software]. GitHub. <https://github.com/gumyr/build123d>
- Okuta, R., Unno, Y., Nishino, D., Hido, S., & Loomis, C. (2017). CuPy: A NumPy-compatible library for NVIDIA GPU calculations. In *Proceedings of the Workshop on Machine Learning Systems (LearningSys) at NIPS 2017*.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., . . . Vázquez-Baeza, Y. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>

END OF DOCUMENT

Check out our Robot design documentation and logs in the documentation booklet