Agentic AI for Optimization of Magnetic Robots

Pengsong Zhang*

Abstract

Magnetic milli-spinner design spans a large coupled geometry space, yet manual CAD construction, mesh preparation and repeated fluid simulations make exploration labour-intensive. Here we develop an Agentic AI workflow connecting literature-derived constraints, parametric geometry, computational fluid dynamics and persistent research memory. Three Missions progressively reconstruct a published benchmark, optimize within its parameter family and explore new spinner–ring geometries. Matched screening calculations identify a within-family design with improved equilibrium speed at both 100 and 160 Hz. Expanded exploration discovers asymmetric and threefold-symmetric speed leaders. The speed-leading ring-nose design reaches 20.04 and 36.63 cm s^{-1}, exceeding the reconstructed published reference by 10.3% and 12.1% while also increasing its internal pressure contrast. Spatial flow diagnostics identify a distinct interaction-potential-leading design. These findings demonstrate physics-guided agentic structural discovery and identify improved computational designs with distinct propulsion and internal-flow advantages.

Computational device design spans a large coupled space of shapes and operating conditions. For fluid-driven devices, each proposed geometry must be converted into a valid CAD model, discretized into a usable mesh and evaluated through repeated flow calculations. Manual construction, mesh correction and simulation setup make this process time-consuming and labour-intensive. Computational inverse design and data-driven fluid mechanics offer powerful ways to explore such spaces [1, 2]. Yet selecting another parameter vector is only part of the work: a change in geometry can alter numerical feasibility, the physical interpretation of a comparison and the next useful experiment. Connecting these activities is therefore central to turning simulation capacity into useful design exploration.

Miniature medical robots illustrate the importance of this connection. Their potential roles span delivery, sensing and intervention [3], while locomotion depends on how geometry and actuation interact with the surrounding medium [4]. Magnetic surface microrollers have demonstrated navigation against blood flow [5]; magnetic manipulation also enables telerobotic navigation of guidewires through tortuous vasculature [6]. These distinct platforms emphasize the need to match a design to its operating environment rather than compare speeds across unlike devices. Advances in three-dimensional printing broaden the geometries that can be fabricated [7], increasing the value of computational methods that identify which designs warrant further evaluation.

Magnetically actuated milli-spinners provide a demanding design problem in which these issues converge. A hollow body, helical fins and a magnetic ring jointly influence propulsion and local flow. The platform has demonstrated preclinical thrombectomy and endovascular navigation capabilities [8, 9]. A subsequent study combined CFD and experiments to investigate fin number, through-hole radius, helical angle and slit width [10]. Its final configuration provides a concrete physical benchmark: three fins, a through-hole ratio of 1.25, a helical angle of 60° and

*OpenAGS, UofT

a normalized slit width of 0.75. These measurements and geometric constraints support both reconstruction and the search for improved designs.

The variables are coupled. Increasing through-hole radius shortens the fins under a fixed spinner radius; it also changes normalized slit width when absolute slit width is held constant. The ring contributes its own hydrodynamic load, so stronger spinner thrust need not imply faster motion of the complete assembly. Propulsion and internal suction are also different objectives. In the benchmark study, equilibrium propulsion speed measures the opposing flow that the device can balance, whereas the internal pressure contrast characterizes a local low-pressure response. An informative search must therefore connect shape changes to force balance and field structure rather than rank designs by an isolated geometric feature.

Derivative-free and Bayesian optimization offer established approaches to expensive simulation-based design [11, 12]. Bayesian reaction optimization illustrates how scientific evaluations can be selected from a much larger set of possible conditions [13], and deep active optimization combines learned surrogates with adaptive exploration [14]. These methods guide the selection of evaluations, but a complete geometry-to-CFD research cycle also requires source interpretation, model construction, numerical recovery and comparable evidence. Here the design problem includes deciding which representation to expand and which matched comparison will distinguish competing explanations, as well as selecting points in a fixed parameter space. The search procedure must connect numerical outcomes to a physical question that can change as the design evolves.

AI increasingly supports not only prediction but also the organization of scientific discovery [15]. Earlier robot scientists connected hypothesis generation to automated testing [16], and mobile robotic laboratories coupled experimental execution to iterative selection [17]. These examples establish a broader research direction: computational decisions can be linked to executable experiments and their feedback. For simulation-based design, the analogous opportunity is to connect the choice of a physical modification to the sequence of modelling operations needed to evaluate it.

Agentic AI provides a computational interface for this purpose. Language-model agents can interleave external actions with reasoning and retain feedback across attempts [18, 19]. Coscientist and ChemCrow connect language models to specialist tools and experimental execution [20, 21], while SciToolAgent addresses the integration of multiple scientific tools [22]. Scientific-agent systems have also connected literature, hypotheses and experimental analysis [23–25]. The relevant capability is not language generation alone, but the use of external computational results to guide a sequence of research actions.

Physics-based enzyme engineering provides a recent example of translating design intent into executable modelling [26]. In CFD, ChatCFD and Foam-Agent address case preparation, execution and recovery [27, 28]. These developments motivate extending executable simulation assistance into an iterative physical-design programme. Such a programme must carry geometric identity, operating conditions and performance definitions across successive calculations, so that each qualified result informs the next geometry or matched comparison. It must also distinguish an improvement within an established parameterization from a new mechanism made accessible by changing the geometry representation.

Here we develop an Agentic AI workflow that links literature-derived constraints, feature-level CAD, mesh preparation, accepted CFD and persistent research memory. We organize the design programme into three progressively broader Missions: reconstructing the published benchmark, optimizing jointly within its parameter family and exploring additional spinner–ring geometries. The retained figures label these projects Aims 1–3. A common screening model enables matched comparisons at 100 and 160 Hz. The within-family search identifies a

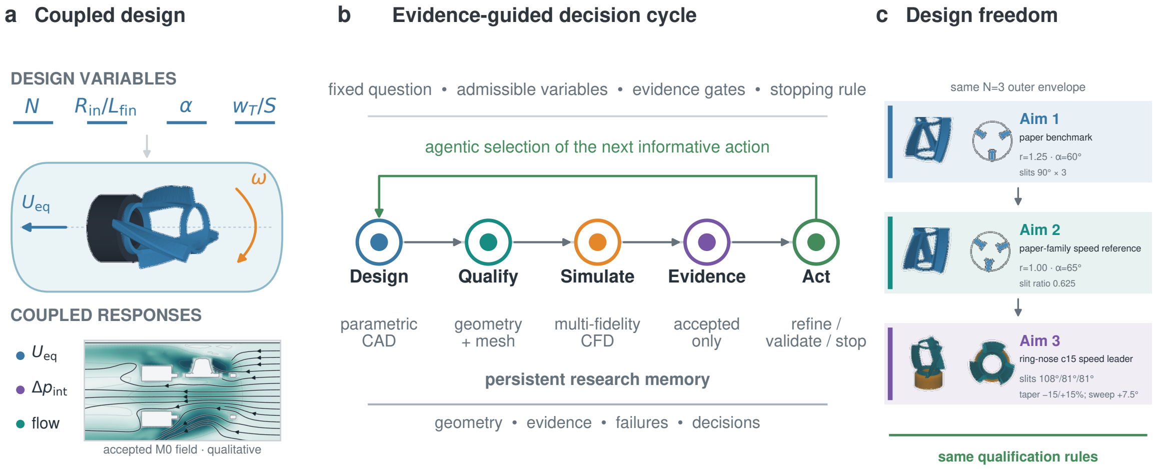

Figure 1: **Agentic AI workflow for milli-spinner design.** **a**, Device geometry, equilibrium propulsion and internal pressure response. The flow field is a representative CFD calculation. **b**, Literature-grounded constraints connect CAD, mesh qualification, CFD, accepted-result assessment and next-action selection. Research memory retains decisions and failures. **c**, Representative structures from the benchmark, within-family search and expanded geometry search. Labels Aim 1–3 identify the three design Missions.

dual-frequency speed-improving design and conditional geometric interactions. The expanded search discovers asymmetric and threefold-symmetric speed leaders, including a ring-nose design that improves both speed and internal pressure contrast relative to the reconstructed published reference. Spatial diagnostics further distinguish a speed-leading design from an interaction-potential-leading design. Together, these results demonstrate how Agentic AI can connect benchmark reconstruction to interpretable structural discovery within one computational design workflow.

Results

An agentic loop connects literature, geometry and physical computation

The workflow connects literature inspection, parametric CAD, mesh generation, CFD and result interpretation (Figs. 1 and 2). Each comparison specifies design variables, fixed physical conditions and a reference geometry. Numerical tools assess model validity and extract loads and fields; qualified results inform the next design comparison. Persistent research memory supports neighbourhood tests, matched controls, operating-condition comparisons and selected resolution checks. Numerical feasibility is evaluated separately from hydrodynamic performance.

The source reconstruction contains 227 structured configuration–result records, 1,748 normalized result records, 50 parametric models and 67 formal CAD views. Reported, digitized, derived and unreported quantities are distinguished. The authors' native CAD was unavailable; missing geometric details remain declared reconstruction choices. Two representative reconstructions achieved view-specific silhouette intersection-over-union values of 0.966/0.901 and 0.946/0.907, exceeding the 0.90 criterion in all four views (Supplementary Fig. S1). Solid, dimensional and topology checks complement the projection scores; silhouette agreement alone does not identify a unique three-dimensional geometry.

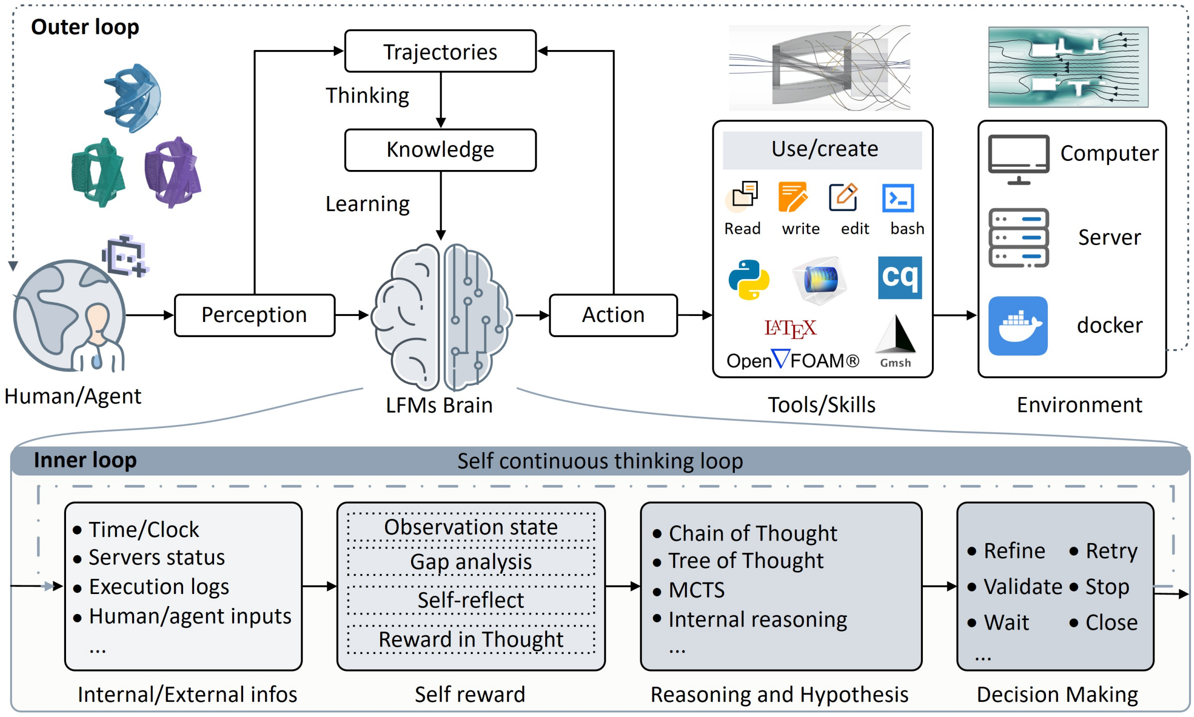

Figure 2: **Agent architecture and decision cycle.** The outer loop connects project state, research memory, a foundation-model agent, tools and computing resources. CAD models and flow fields illustrate tool outputs. The inner loop links state observation, evidence gaps, hypotheses and permitted actions. Numerical admission and evidence use are determined by executable checks.

Mission 1 reconstructs the published benchmark and its principal trends

All 59 benchmark anchors have at least one accepted M0 fixed-speed sample (Fig. 3a). The set comprises 33 fin-ratio, five helix-angle, five slit-width, nine frequency, two viscosity and five no-ring conditions. Eleven configurations required local mesh correction under unchanged quality limits. Fixed-speed coverage is distinguished from equilibrium-root coverage.

Across eight assessed trend families, four show full qualitative agreement, three show partial agreement with incomplete equilibrium coverage and one remains discrepant (Supplementary Table S7). Two-fin ratio, frequency and viscosity responses recover the reported trend directions. Four-fin, helix-angle and slit-width comparisons retain peak structure in the evaluated regions. The three-fin ratio comparison does not recover the published interior maximum near 1.25–1.5. At 100 Hz, nine accepted near-root points and two visibly distinguished accepted-endpoint secants span the 11 ratios; a shallow maximum near 0.75 precedes the overall decline. At 160 Hz, 10/11 ratios have accepted near-root values, decreasing over 0.75–4.0; ratio 0.5 has accepted evidence but no root.

Fixed-speed force profiles at additional frequencies show maxima that shift with frequency, but they are not equilibrium-speed curves (Fig. 3e). They provide complementary diagnostics of force generation. Two pressure profiles at the published speed and representative flow topology provide separate checks, not quantitative root-level field validation.

The discrepancy is sensitive to both assembly construction and discretization. Complete-annular connection tests changed force cancellation at ratios 1.0 and 1.25, but did not restore the interior peak. At ratio 1.25 and 160 Hz, increasing mesh fidelity at a common speed changed assembly force from 2.10 to 42.00 µN. Six matched-resolution pairs gave M0 equilibrium estimates 12.0–16.8% below their M1 counterparts. These observations establish sensitivity, not a correction

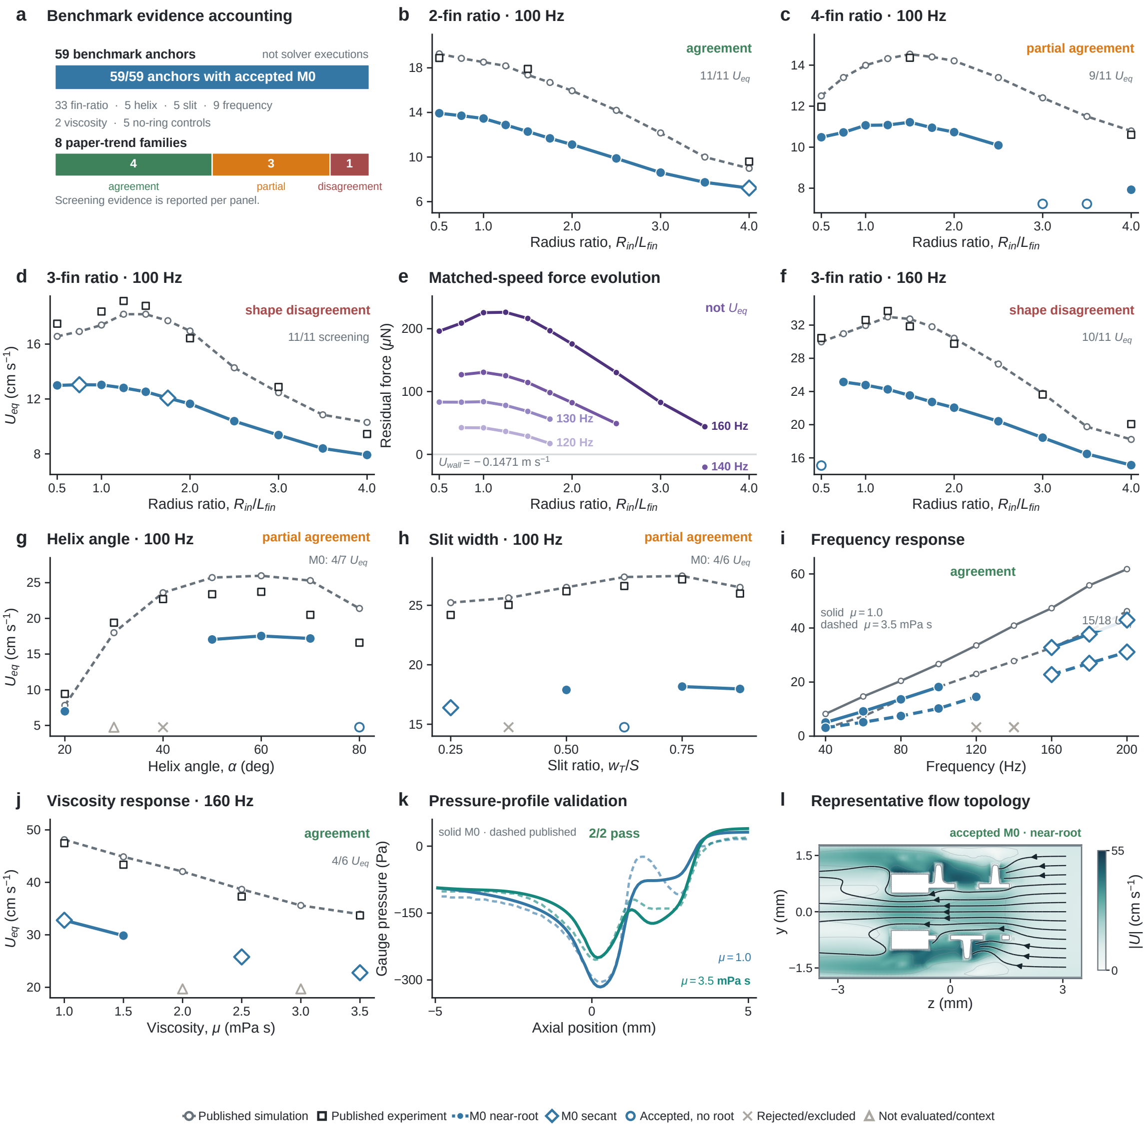

Figure 3: **Published benchmark comparison and screening limits.** **a**, Accepted fixed-speed samples at all 59 anchors. **b–d**, Fin-ratio comparisons at 100 Hz. Panel d contains nine accepted near-root values and two accepted-endpoint secants. **e**, Fixed-speed force diagnostics; the ratio-3 point at 140 Hz is omitted because it did not satisfy numerical convergence criteria. **f**, Three-fin 160-Hz comparison with 10/11 near-root values and an accepted no-root condition at ratio 0.5. **g–j**, Helix, slit, frequency and viscosity comparisons. **k,l**, Pressure profiles at the paper speed and representative near-root topology. Grey curves and open grey circles denote published CFD; black squares denote published experiments; blue filled circles denote near-root screening values; diamonds denote diagnostic secants; open blue circles denote accepted evidence without a root; crosses denote rejected or excluded records; triangles denote unevaluated or context conditions. Published, reconstructed and interpolated evidence remain distinct. The higher-resolution root is reported separately in Supplementary Section H.

factor or a resolved cause. A 9.20-million-cell no-layer calculation of the final published geometry at 100 Hz yielded an accepted root of 23.8194 cm s^{-1}, still 10.68% below the digitized published CFD target of 26.6667 cm s^{-1} (Supplementary Section H). This result is kept separate from the M0 rankings.

Mission 2 identifies a dual-frequency within-family speed winner

The initial 100-Hz joint screen combines five fin angles and five radius ratios. Mesh correction produced 25/25 mesh-admissible cells and 24/25 accepted M0 equilibrium estimates (Fig. 4a). The remaining cell at $\alpha = 50°$ and ratio 1.50 contains accepted fixed-speed samples without a root. The grid fixes $N = 3$ and absolute total slit width at 1.83 mm, so normalized slit width changes with radius. Its response surface therefore compares coupled geometries rather than isolating angle and radius effects.

Later normalized-slit studies evaluated matched cases at 100 and 160 Hz. At normalized slit width 0.75, raising angle from 65° to 67.5° reduced common-speed force at 100 Hz but increased it at 160 Hz. Complete angle-by-slit 2×2 comparisons at both frequencies, and a radius-by-slit comparison at 100 Hz, support conditional rather than additive effects. Common-speed force tests and independently accepted equilibrium estimates are kept separate.

Design R ($\alpha = 65°$, ratio 1.00, normalized slit width 0.625) reaches 18.66 and 33.40 cm s^{-1}, compared with 18.17 and 32.68 cm s^{-1} for the reconstructed final published design. The gains are approximately 2.7% and 2.2%, identifying a dual-frequency speed-improving within-family design in the common screening model. Its internal pressure contrasts are 58.5 and 141.7 Pa: larger than the published-reference reconstruction at 100 Hz and slightly smaller at 160 Hz. Speed improvement does not imply improvement in every pressure metric.

R leads the displayed normalized-slit cohort at 100 Hz, whereas Y10 reaches 33.43 cm s^{-1} at 160 Hz, only 0.085% above R. Together, these results identify a dual-frequency improvement relative to the published-reference reconstruction and a leading design region. A matched 62.5° comparison maintained speed while lowering internal pressure contrast by 8.05%; this change concerns the local pressure response rather than hydraulic loss. Candidate values, representative-root selection and matched interactions are detailed in Supplementary Section D.

Mission 3 discovers asymmetric and symmetric designs with higher speed

The out-of-family search releases slit allocation, taper, sweep and ring-shape variables while retaining declared comparison physics. The paired comparison evaluates six expanded designs and three reference structures at both frequencies.

The asymmetric ring-nose c15 design reaches 20.04 and 36.63 cm s^{-1}, the largest speeds in this paired comparison. Relative to the M0 reconstruction of the final published geometry, the gains are 10.3% and 12.1%. Its internal pressure contrast also rises from approximately 44.4/143.6 to 51.9/170.3 Pa. It therefore improves both speed and this localized pressure proxy relative to that reference at the two tested frequencies, although it does not dominate every candidate and reference on every quantity. Ring-nose c10 has nearly the same speed and slightly greater internal pressure contrast.

The ring-nose c15 spinner allocates slit sectors as 108°/81°/81° and combines a -15% fin taper, +15% bore taper and +7.5° forward sweep with its ring-nose modification. It is therefore not the same geometry as the negative-bore, counter-swept composite below. The explicit feature representation allows the expanded search to alter local passage shape and ring treatment instead of only fin angle or radius. Dimensions and feature conventions are retained in Supplementary Section A.

A second asymmetric composite combines nonuniform slit sectors, negative fin and bore taper, a $-3.75°$ counter-sweep and a volume-matched dual-chamfer ring. It reaches 19.95/36.22 cm s^{-1}, exceeding the reconstructed published design, the benchmark reference and R at both frequencies. Its lower centreline pressure contrast does not imply weaker spatial interaction potentials, as shown below. Because multiple features change together, the composite's gain is not assigned

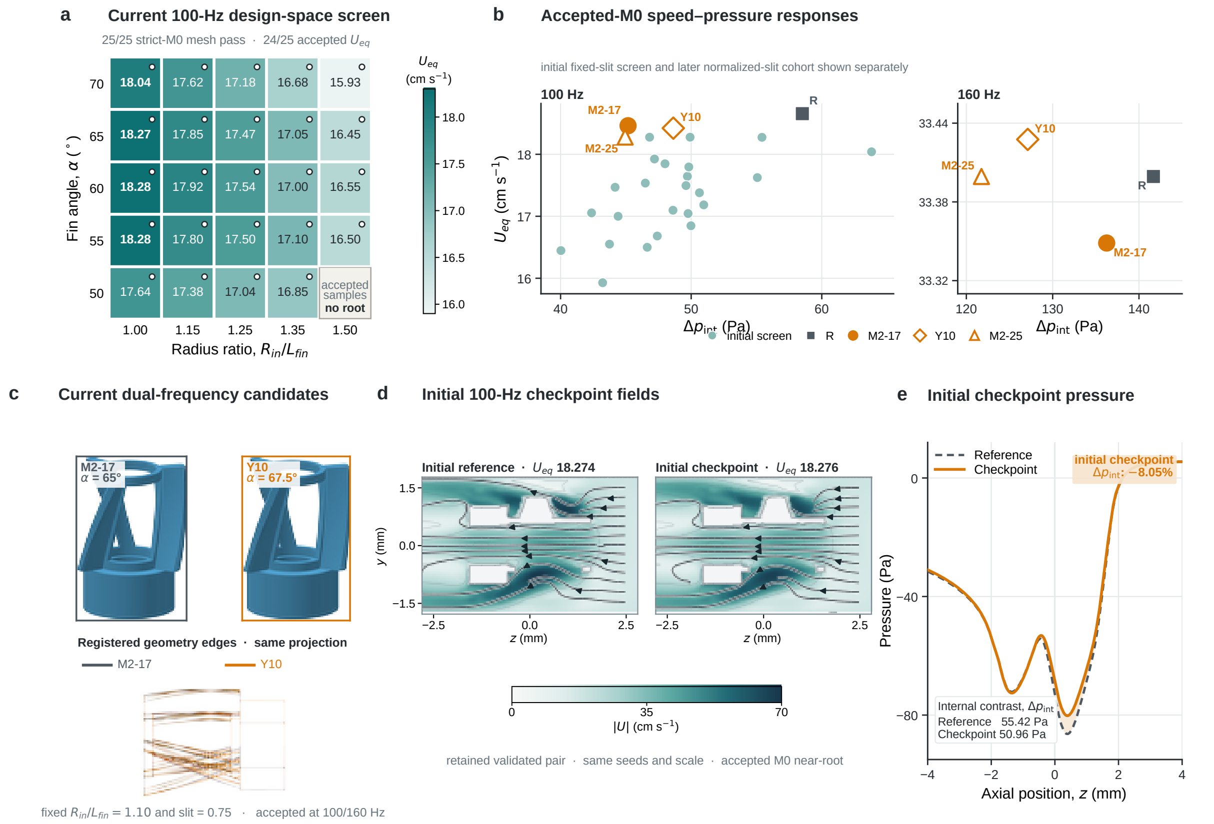

Figure 4: **Within-family search and paired-frequency responses.** **a**, Repaired joint screen with 25/25 admissible meshes and 24/25 accepted equilibrium estimates. **b**, Accepted M0 speed and internal pressure contrast at 100 and 160 Hz. R is the selected within-family speed-improving design; Y10 is marginally faster at 160 Hz. Fixed-total-slit and normalized-slit cohorts are distinct. **c**, Matched M2-17 and Y10 CAD at ratio 1.10 and normalized slit width 0.75. **d,e**, Velocity fields and centreline pressure profiles for the 100-Hz 65° reference and 62.5° comparison. A decrease in internal pressure contrast is not a demonstrated hydraulic-loss benefit. Rankings are limited to the declared same-method screening cohort.

causally to one feature.

A symmetry-constrained branch retains the M2-17 spinner and changes the coaxial ring while conserving ring volume and contact position. OD1.875 reaches 19.78/35.51 cm s^{-1}, the largest tested C3-symmetric speeds. The nearby OD1.88 design has a larger pressure contrast with a small speed penalty. The ring-only intervention provides a more localized design comparison than the asymmetric composite, with geometric symmetry imposed as an engineering constraint (Fig. 5).

Spatial diagnostics distinguish pressure extrema from interaction potentials

Postprocessing retained accepted fields produces three S0 estimates at a common 0.25-mm virtual gap: signed pressure-holding potential, excess inward flow and excess inward momentum. The calculation revolves a single diametral slice; it is not full-three-dimensional surface integration, particularly important for asymmetric designs. Subtracting the wall speed defines a background-relative flow descriptor rather than an isolated causal estimate of rotation.

The counter-sweep design leads all three estimates in the nine-structure paired comparison.

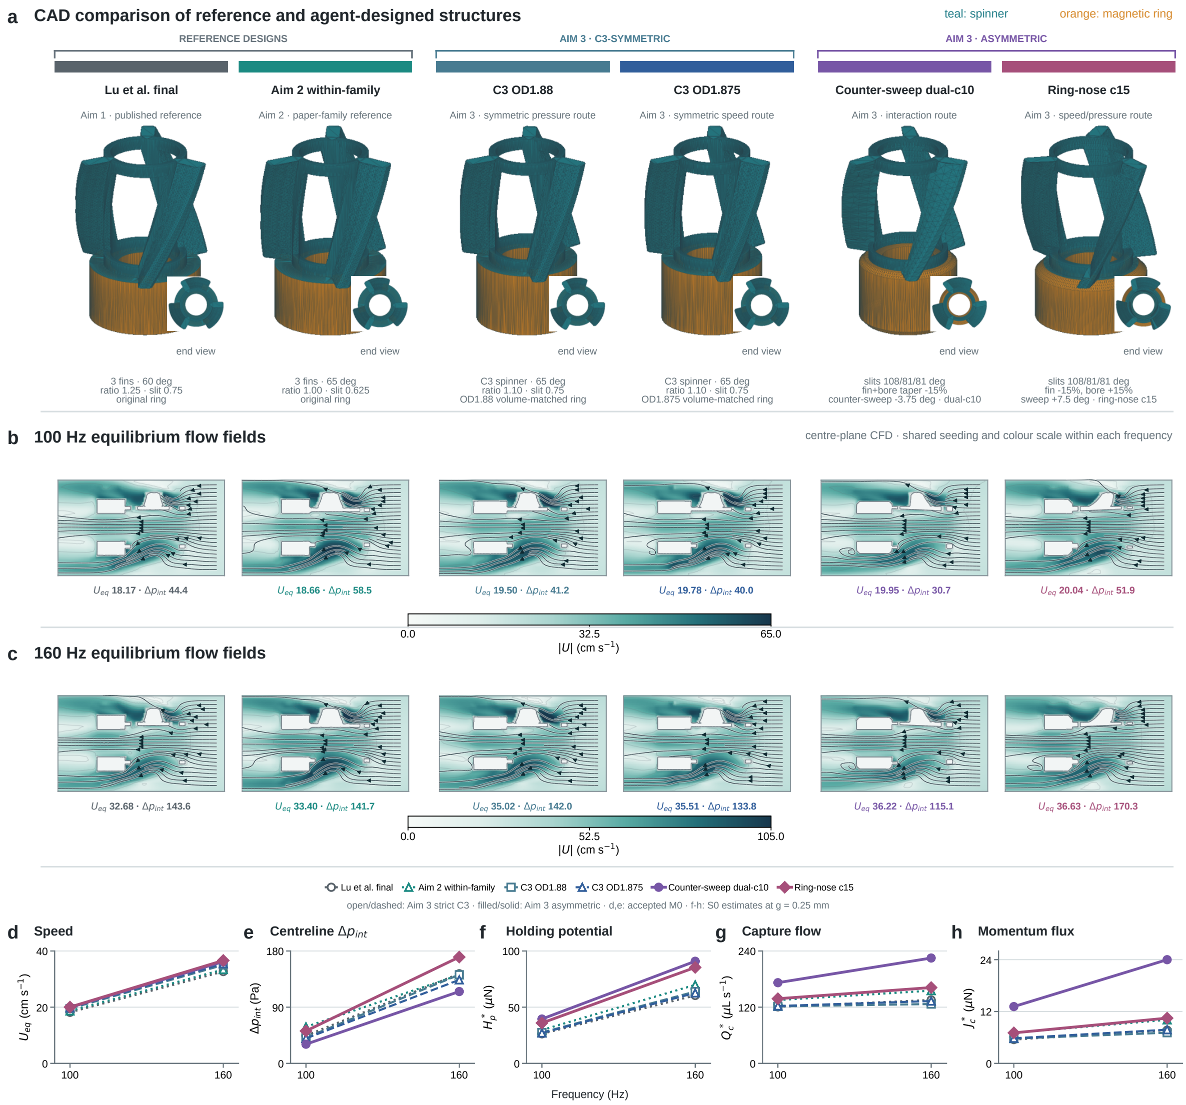

Figure 5: **Expanded designs and distinct hydrodynamic advantages. a**, CAD comparison of the reconstructed published geometry, R, two C3-symmetric candidates and two asymmetric candidates. **b,c**, Accepted M0 equilibrium flow fields at 100 and 160 Hz with shared display settings within each frequency. **d,e**, Equilibrium speed and centreline internal pressure contrast. Larger contrast indicates a stronger localized low-pressure proxy, not greater conventional hydraulic loss. **f–h**, S0 slice-derived estimates of pressure-holding potential, excess inward flow and momentum at a 0.25-mm virtual gap. These are open-flow diagnostics, not full-three-dimensional integrals or clot-present forces. Connecting segments identify paired observations and are not continuous-frequency fits.

Relative to the reconstructed final published design, pressure-holding potential increases by approximately 50% at both frequencies, excess inward flow by 43–66% and inward momentum by 139–212%. Conversely, the slit90–taper−40–sweep−7.5 candidate has the largest centreline internal pressure contrast but smaller spatial interaction estimates. A single pressure minimum therefore does not determine the ranking of these slice-derived descriptors. The symmetric ring redesign primarily improves speed while its interaction estimates remain close to the published-reference reconstruction.

These comparisons motivate separate reporting of propulsion, localized pressure contrast and spatial interaction potentials rather than an arbitrary weighted score. Ring-nose c15 favours speed, the counter-sweep design favours the S0 interaction estimates and the C3 branch preserves symmetry. Their performance differences support interpretable design alternatives within the screening model.

Discussion

This study connects benchmark reconstruction, within-family optimization and out-of-family structural discovery through an Agentic AI workflow. The three Missions build on one another: literature-grounded reconstruction defines a common physical comparison, joint parameter studies identify interactions and improved designs, and an expanded geometry representation exposes additional design opportunities. The resulting discoveries are concrete. Mission 2 identifies a design with higher accepted equilibrium speed at both 100 and 160 Hz than the reconstructed published reference. Mission 3 increases this advantage through asymmetric and C3-symmetric structures, demonstrating that useful improvements extend beyond the original parameter family.

The expanded designs also reveal different routes to improvement. Ring-nose c15 gives the highest paired speeds in the reported cohort and increases internal pressure contrast at both frequencies relative to the reconstructed published design. The counter-sweep composite exceeds the three declared speed references while leading the spatial interaction estimates. The C3 ring redesign provides a symmetry-preserving route that also exceeds the within-family speed reference. These findings support alternative design choices rather than requiring one geometry to maximize every observable. Multiobjective autonomous experimentation similarly distinguishes improvements in individual properties from the trade-offs among them [29]. A velocity winner relative to named references is a meaningful result even when another candidate is nearly tied or leads a different pressure descriptor.

The physical distinction between a pressure extremum and a spatially distributed response is particularly useful. The internal centreline contrast describes a localized pressure deficit, not conventional vessel loss. The counter-sweep composite has a smaller centreline contrast yet larger slice-derived pressure-holding, inward-flow and momentum estimates. Conversely, the largest internal contrast does not identify the largest spatial interaction estimates. Thus, the expanded search provides more than a speed ranking: it identifies how propulsion, local suction proxies and symmetry constraints can favour different geometries.

The computational contribution is an executable connection across the design cycle. Rather than treating CAD, CFD and interpretation as disconnected tasks, the agent uses retained results to move between parameter neighbourhoods, matched controls and new geometry features. Numerical feedback also distinguishes an unusable mesh from a poor-performing structure, allowing mesh correction without altering the physical question. This connects the broader direction of scientific agents [23, 26] to geometry-based discovery. Contemporary work on turning papers into interactive agents further emphasizes the value of making methods and data executable [30]. Here, the connection extends from reconstructing a published physical system to testing modifications beyond its original design family.

Several limitations define the next stage of this work. First, broad benchmark coverage does not yet provide a fully calibrated reproduction of the published CFD. The three-fin interior peak remains discrepant, and the 9.20-million-cell calculation gives a final-reference speed of 23.8194 cm s^{-1} at 100 Hz, still 10.68% below the digitized published target. The authors' native CAD and complete numerical settings were unavailable, so discretization sensitivity and reconstruction choices remain coupled. Mesh refinement reduces the speed gap but has not established mesh independence or isolated the cause of the peak mismatch. Convergence tolerances and force-tail variability do not replace discretization-error estimates [31].

The immediate numerical priority is a compact three-anchor validation: the final published geometry at 100 and 160 Hz, and a neighbouring published geometry at 100 Hz held out from settings selection. Only the first has a converged finer-mesh equilibrium result; the frequency-transfer and neighbouring-geometry tests remain incomplete, and the first result still falls short of paper alignment. After these tests establish a transferable method, the settings should be fixed and applied to the five shortlisted designs using identical comparison conditions. This sequence separates inexpensive screening from final confirmation, consistent with the cost–information distinction in multifidelity research [32, 33]. It does not require repeating the entire benchmark survey. Until this comparison is completed, candidate gains remain same-method M0 findings rather than higher-fidelity or experimental superiority. Near-tied within-family designs also preclude a unique global ordering without negating their improvement over the named reference.

Second, physical validation must connect constrained-axis hydrodynamics to device behaviour. The two tested frequencies address the published operating conditions, not a continuous frequency band. Fixed-axis MRF calculations cannot resolve lateral motion or precession and do not establish free-motion stability, even for a symmetric geometry. Matched ring volume does not prove identical magnetic torque or attainable rotation frequency. Fabrication tolerance, minimum printable features and material response require joint consideration with actuation [34]. Similarly, the slice-derived S0 quantities are open-flow interaction proxies, not measurements of clot adhesion, capture or debulking. Three-dimensional field integration and subsequent clot-present calculations or benchtop tests are needed to establish whether those proxies translate into functional advantages. These investigations extend the reported hydrodynamic comparisons to manufacturing, motion and clot interaction.

Third, efficiency and transferability need independent evaluation. Uncertainty-calibrated language-model optimizers illustrate a complementary route to candidate selection [35], but no matched-budget comparison with such methods, derivative-free search or surrogate-guided optimization has been performed here [11, 14]. Nor has the contribution of the foundation model or individual reasoning/search components been isolated. A controlled evaluation should include geometry construction, numerical recovery, unsuccessful evaluations and human effort, as well as completed CFD samples. The supplementary neural-operator/PINO pilot likewise needs independent, physical-geometry-held-out testing with separation of related simulation states before supporting candidate selection. It remains separate from the discovery provenance of the reported designs.

These outstanding tasks concern quantitative calibration, physical realization and measured computational efficiency. They do not change the demonstrated same-method improvements: a dual-frequency within-family design, asymmetric and symmetry-preserving speed leaders, and distinct internal-flow advantages. Overall, the three-Mission programme demonstrates physics-guided agentic design discovery and identifies concrete structures for targeted numerical and experimental validation.

Methods

Literature reconstruction and parametric geometry

The benchmark is Lu *et al.* and its Supplementary Information [10]. Reported values, raster-digitized values and quantities derived from published dimensions are distinguished. Native CAD and complete numerical settings were unavailable; unspecified reconstruction choices are described in Supplementary Section A. The spinner outer radius is 1.25 mm, body length 2.15 mm and intervening wall thickness 0.12 mm. With $r = R_{\mathrm{in}}/L_{\mathrm{fin}}$,

$$R_{\mathrm{in}} + 0.12 \text{ mm} + L_{\mathrm{fin}} = 1.25 \text{ mm}. \tag{1}$$

Normalized total slit width is $w_T/S = N\theta_N/(2\pi)$, where θ_N is one slit's angular width. The final published configuration is $N = 3$, $r = 1.25$, $\alpha = 60°$ and $w_T/S = 0.75$.

CAD construction uses CadQuery 2.8.0 and OCP 7.9.3.1.1. Checks assess solid validity, dimensions, surface closure, orientation and intersections with the fluid domain. For projection validation, published silhouettes are selected before scoring and registration excludes non-rigid deformation. Equivalent geometries are reconciled, and wall speeds within 10^{-7} m s^{-1} are treated as duplicates.

Agentic design cycle

The agent connects literature inspection, parametric CAD, mesh generation, CFD and analysis through executable tools. Its inputs comprise design variables, fixed physical conditions, comparison references and previously evaluated results. Available actions include parameter-neighbourhood comparisons, matched controls, additional operating conditions and changes to the geometry representation. CAD tools return geometric validity and mesh tools return discretization quality; CFD supplies loads, pressure and velocity fields. Qualified outcomes inform subsequent design choices, while unsuccessful calculations are distinguished from low-performing geometries.

Persistent research memory links design choices to physical conditions and numerical outcomes. Independent designs can be evaluated concurrently; each force-balance chain is updated sequentially using qualified samples. Executable numerical tests, rather than language-model predictions, determine which samples enter equilibrium searches and performance comparisons. Reasoning/search labels in Fig. 2 describe conceptual actions rather than a distinct Monte Carlo tree-search implementation.

Flow model and boundary conditions

OpenFOAM v2606 solves steady incompressible Newtonian laminar flow using a stationary multiple reference frame (MRF) and SIMPLE pressure–velocity coupling [36, 37]. The common saline-like condition uses density $\rho = 1000$ kg m^{-3}, dynamic viscosity $\mu = 0.001$ Pa s and kinematic viscosity $\nu = 10^{-6}$ m^2 s^{-1}. Other benchmark viscosities are specified for their respective comparison conditions.

The tube has inner diameter 3.5 mm and length 300 mm, with zero gauge pressure at both axial ends. The no-slip tube wall has a prescribed axial velocity; spinner and attached ring are rotating no-slip surfaces. Rotation is about the device axis at $\omega = 2\pi f$. The MRF zone has radius 1.40 mm and 0.25-mm axial clearance at both ends of the assembly; its exact extent is a reconstruction choice.

M0 uses bounded first-order convection, no prism layers and approximately 0.68 million cells, with geometry-dependent counts. M1 is a finer no-layer mesh family. The selected 9,201,337-cell calculation is analyzed separately. Mesh labels M0 and M1 denote resolution, not the three design Missions.

Numerical quality and convergence

Numerical verification and validation against physical observations address different questions [38]. The mesh and solution criteria below qualify computational samples; comparisons with published CFD and experiments assess agreement with the reconstructed benchmark.

Mesh criteria include positive cell volumes, required patch and rotating-zone coverage, maximum non-orthogonality $\leq 65°$, maximum skewness ≤ 4, concave-cell fraction ≤ 0.01, low-determinant-cell fraction $\leq 5 \times 10^{-5}$ and zero low-weight or nonstandard edge faces. Local corrections preserve CAD geometry and these limits.

A qualified solution requires finite output, velocity and pressure residuals $\leq 10^{-4}$, normalized mass imbalance $\leq 10^{-3}$, assembly-force closure and stable force histories. Force-tail and conservation definitions are given in Supplementary Sections B and C. The initial screening calculations use a 500-iteration startup and 250-iteration evaluation intervals, with 1500–2000 iterations; pressure-damped comparisons use pressure relaxation 0.1 and a 4000-iteration maximum. These numerical settings are distinguished by calculation cohort. Results failing the applicable quality or convergence criteria are excluded from quantitative comparisons.

Equilibrium speed and internal pressure contrast

Hydrodynamic traction is integrated separately over spinner and ring surfaces. Their sum defines the assembly force; tube and MRF-interface patches are excluded. The spinner and attached ring form one rigid body, so assembly axial force determines equilibrium. A bracketed search uses qualified opposite-sign samples, with secant prediction

$$U^*_{\text{wall}} = U_1 - F_1 \frac{U_2 - U_1}{F_2 - F_1}. \tag{2}$$

An M0 equilibrium estimate requires a simulation at the estimated speed that satisfies the numerical criteria and $|F_{z,\text{assembly}}| \leq 5$ µN. Then $U_{\text{eq}} = -U_{\text{wall}}$. The selected finer-mesh result uses a 1 µN tolerance. Interpolated zeros without a qualified simulation at that speed are shown separately and excluded from candidate rankings. When several qualified near-root results exist, the representative is selected at the smallest absolute assembly force, with force-tail variation as a tie-breaker.

Kinematic pressure is multiplied by density. Centreline internal pressure contrast is defined as

$$\Delta p_{\text{int}} = p(z_{\text{front}}) - \min_{z \in \text{device interior}} p(z). \tag{3}$$

It characterizes local low pressure rather than inlet–outlet energy loss. Directly measured and interpolated pressure values remain distinguishable.

Design comparisons and spatial diagnostics

Candidate and reference comparisons match frequency, fluid, boundaries, mesh resolution and observables. The initial within-family grid fixes absolute slit width; normalized-slit studies are analyzed separately. Force-tail variation is a convergence diagnostic, not a confidence interval for ranking near-tied designs.

S0 descriptors use a virtual disk of radius 1.25 mm at a 0.25-mm gap from the common front envelope and upstream reference pressure at $z = 3.80$ mm. They are estimated by diametral revolution of an $x = 0$ flow-field slice. With $v_{\text{in}} = \max[-(u_z - U_{\text{wall}}), 0]$,

$$H^*_p = \int_A (p_{\text{ref}} - p)\, \mathrm{d}A, \tag{4}$$

$$Q^*_c = \int_A v_{\text{in}}\, \mathrm{d}A, \qquad J^*_c = \rho \int_A v_{\text{in}}^2\, \mathrm{d}A. \tag{5}$$

The pressure integral is signed. These are background-relative open-flow estimates, not full-three-dimensional integrals or forces on a real clot. A 0.50-mm plane provides distance context; a 0.10-mm plane intersects an asymmetric geometry and is excluded.

Geometry definitions, physical conditions and source calculations remain associated with derived outputs to support traceable data reuse [39]. Software and computing environments are specified in Supplementary Section F; the exploratory neural-operator study is confined to Supplementary Section G.

Data availability

All data presented in this study are included in the article and its Supplementary Information.

Code availability

The code will be released on GitHub soon.

Competing interests

The author declares no competing interests.

AI assistance

For manuscript preparation, AI tools were used to polish the language and provide suggestions for figure design.

References

[1] Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12:659–670, 2018. doi: 10.1038/s41566-018-0246-9.

[2] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020. doi: 10.1146/annurev-fluid-010719-060214.

[3] Jinxing Li, Berta Esteban-Fernández de Ávila, Wei Gao, Liangfang Zhang, and Joseph Wang. Micro/nanorobots for biomedicine: Delivery, surgery, sensing, and detoxification. *Science Robotics*, 2(4):eaam6431, 2017. doi: 10.1126/scirobotics.aam6431.

[4] Stefano Palagi and Peer Fischer. Bioinspired microrobots. *Nature Reviews Materials*, 3: 113–124, 2018. doi: 10.1038/s41578-018-0016-9.

[5] Yunus Alapan, Ugur Bozuyuk, Pelin Erkoc, Alp Can Karacakol, and Metin Sitti. Multifunctional surface microrollers for targeted cargo delivery in physiological blood flow. *Science Robotics*, 5(42):eaba5726, 2020. doi: 10.1126/scirobotics.aba5726.

[6] Yoonho Kim, Emily Genevriere, Pablo Harker, et al. Telerobotic neurovascular interventions with magnetic manipulation. *Science Robotics*, 7(65):eabg9907, 2022. doi: 10.1126/scirobotics.abg9907.

[7] Sajjad Rahmani Dabbagh, Misagh Rezapour Sarabi, Mehmet Tugrul Birtek, Siamak Seyfi, Metin Sitti, and Savas Tasoglu. 3D-printed microrobots from design to translation. *Nature Communications*, 13:5875, 2022. doi: 10.1038/s41467-022-33409-3.

[8] Yilong Chang, Shuai Wu, Qi Li, Benjamin Pulli, Darren Salmi, Paul Yock, Jeremy J. Heit, and Ruike Renee Zhao. Milli-spinner thrombectomy. *Nature*, 642:336–342, 2025. doi: 10.1038/s41586-025-09049-0.

[9] Shuai Wu, Yilong Chang, Sophie Leanza, Jay Sim, Lu Lu, Qi Li, Diego Stone, and Ruike Renee Zhao. Magnetic milli-spinner for robotic endovascular surgery. *Advanced Materials*, 38 (4):e08180, 2026. doi: 10.1002/adma.202508180.

[10] Lu Lu, Luca Higgins, Jack Bernardo, and Ruike Renee Zhao. Optimization of magnetic milli-spinner for robotic endovascular intervention. *Advanced Robotics Research*, 2(4):e70121, 2026. doi: 10.1002/adrr.70121.

[11] Jeffrey Larson, Matt Menickelly, and Stefan M. Wild. Derivative-free optimization methods. *Acta Numerica*, 28:287–404, 2019. doi: 10.1017/S0962492919000060.

[12] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016. doi: 10.1109/JPROC.2015.2494218.

[13] Benjamin J. Shields, Jason Stevens, Jun Li, et al. Bayesian reaction optimization as a tool for chemical synthesis. *Nature*, 590:89–96, 2021. doi: 10.1038/s41586-021-03213-y.

[14] Ye Wei, Bo Peng, Ruiwen Xie, Yangtao Chen, Yu Qin, Peng Wen, Stefan Bauer, Po-Yen Tung, and Dierk Raabe. Deep active optimization for complex systems. *Nature Computational Science*, 5:801–812, 2025. doi: 10.1038/s43588-025-00858-x.

[15] Hanchen Wang, Tianfan Fu, Yuanqi Du, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620:47–60, 2023. doi: 10.1038/s41586-023-06221-2.

[16] Ross D. King, Jem Rowland, Stephen G. Oliver, et al. The automation of science. *Science*, 324(5923):85–89, 2009. doi: 10.1126/science.1165620.

[17] Benjamin Burger, Phillip M. Maffettone, Vladimir V. Gusev, et al. A mobile robotic chemist. *Nature*, 583:237–241, 2020. doi: 10.1038/s41586-020-2442-2.

[18] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2210.03629. ICLR camera-ready version.

[19] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-0377. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html.

[20] Daniil A. Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624:570–578, 2023. doi: 10.1038/s41586-023-06792-0.

[21] Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D. White, and Philippe Schwaller. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6:525–535, 2024. doi: 10.1038/s42256-024-00832-8.

[22] Keyan Ding, Jing Yu, Junjie Huang, Yuchen Yang, Qiang Zhang, and Huajun Chen. SciToolAgent: A knowledge-graph-driven scientific agent for multitool integration. *Nature Computational Science*, 5:962–972, 2025. doi: 10.1038/s43588-025-00849-y.

[23] Kyle Swanson, Wesley Wu, Nash L. Bulaong, John E. Pak, and James Zou. The virtual lab of AI agents designs new SARS-CoV-2 nanobodies. *Nature*, 646:716–723, 2025. doi: 10.1038/s41586-025-09442-9.

[24] Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, et al. Accelerating scientific discovery with Co-Scientist. *Nature*, 655:487–496, 2026. doi: 10.1038/s41586-026-10644-y.

[25] Ali E. Ghareeb, Benjamin Chang, Ludovico Mitchener, et al. A multi-agent system for automating scientific discovery. *Nature*, 655:497–505, 2026. doi: 10.1038/s41586-026-10652-y.

[26] Qianzhen Shao, Yinjie Zhong, Sebastian Stull, Xinchun Ran, Ning Ding, Kieran Nehil-Puleo, Ruizhe Yao, Han Xu, and Zhongyue J. Yang. MutexaGPT: An intuition-to-design translator for physics-based enzyme engineering. *Nature Computational Science*, 2026. doi: 10.1038/s43588-026-01049-y. Published online 10 September 2026.

[27] E. Fan, Kang Hu, Zhuowen Wu, Jiangyang Ge, Jiawei Miao, Yuzhi Zhang, He Sun, Weizong Wang, and Tianhan Zhang. ChatCFD: A large language model-driven agent for end-to-end computational fluid dynamics automation with structured knowledge and reasoning. *Advanced Intelligent Discovery*, 2(3):e202500174, 2026. doi: 10.1002/aidi.202500174.

[28] Ling Yue, Nithin Somasekharan, Tingwen Zhang, Yadi Cao, and Shaowu Pan. Foam-Agent 2.0: An end-to-end composable multi-agent framework for automating CFD simulation in OpenFOAM. arXiv:2509.18178, 2025. Preprint, version 2.

[29] Benjamin P. MacLeod, Fraser G. L. Parlane, Connor C. Rupnow, et al. A self-driving laboratory advances the Pareto front for material properties. *Nature Communications*, 13: 995, 2022. doi: 10.1038/s41467-022-28580-6.

[30] Jiacheng Miao, Joe R. Davis, Yaohui Zhang, Jonathan K. Pritchard, and James Zou. Reimagining research papers as interactive and reliable AI agents. *Nature*, 2026. doi: 10.1038/s41586-026-11044-y. Published online 16 September 2026.

[31] I. B. Celik, U. Ghia, P. J. Roache, C. J. Freitas, H. Coleman, and P. E. Raad. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *Journal of Fluids Engineering*, 130(7):078001, 2008. doi: 10.1115/1.2960953.

[32] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, 2018. doi: 10.1137/16M1082469.

[33] Víctor Sabanza-Gil, Riccardo Barbano, Daniel Pacheco Gutiérrez, et al. Best practices for multi-fidelity Bayesian optimization in materials and molecular research. *Nature Computational Science*, 5:572–581, 2025. doi: 10.1038/s43588-025-00822-9.

[34] Chuanrui Chen, Shichao Ding, and Joseph Wang. Materials consideration for the design, fabrication and operation of microscale robots. *Nature Reviews Materials*, 9:159–172, 2024. doi: 10.1038/s41578-023-00641-2.

[35] Bojana Ranković, Ryan-Rhys Griffiths, and Philippe Schwaller. Large language models as uncertainty-calibrated optimizers for experimental discovery. *Nature Machine Intelligence*, 8:1466–1477, 2026. doi: 10.1038/s42256-026-01283-z.

[36] Fadl Moukalled, Luca Mangani, and Marwan Darwish. *The Finite Volume Method in Computational Fluid Dynamics: An Advanced Introduction with OpenFOAM and Matlab*, volume 113 of *Fluid Mechanics and Its Applications*. Springer, 2016. doi: 10.1007/978-3-319-16874-6.

[37] OpenCFD Ltd. OpenFOAM documentation: simpleFoam. Official solver documentation, n.d. URL https://doc.openfoam.com/2306/tools/processing/solvers/rtm/incompressible/simpleFoam/. Accessed 19 September 2026; background documentation, not the project's v2606 case definition.

[38] William L. Oberkampf, Timothy G. Trucano, and Charles Hirsch. Verification, validation, and predictive capability in computational engineering and physics. *Applied Mechanics Reviews*, 57(5):345–384, 2004. doi: 10.1115/1.1767847.

[39] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, et al. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3:160018, 2016. doi: 10.1038/sdata.2016.18.

[40] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2010.08895.

[41] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021. doi: 10.1038/s42256-021-00302-5.

[42] Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 6:320–328, 2024. doi: 10.1038/s42254-024-00712-5.

[43] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

[44] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021. doi: 10.1038/s42254-021-00314-5.

[45] Dmitrii Kochkov, Jamie A. Smith, Ayya Alieva, Qing Wang, Michael P. Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021. doi: 10.1073/pnas.2101784118.

[46] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023. URL https://jmlr.org/papers/v24/21-1524.html.

[47] Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. arXiv:2111.03794, 2021. Preprint.

Supplementary Information

Agentic AI for Optimization of Magnetic Robots

This Supplementary Information provides geometric definitions, numerical methods, detailed benchmark and design comparisons, discretization sensitivity and an exploratory neural-operator study. Aims 1–3 in the figures denote the three design Missions, whereas M0 and M1 denote mesh resolutions. Directly evaluated equilibrium values and interpolated estimates are distinguished throughout. All design rankings compare calculations at the same frequency, fluid properties and mesh resolution.

A Geometry Definitions and Design Variables

The geometry used in this study is a parametric reconstruction of the published milli-spinner because the authors' native CAD was unavailable [10]. Reported dimensions and topology were retained, and details absent from the article and Supplementary Information were specified as reconstruction choices. All dimensions below are in millimetres unless stated otherwise.

Reference geometry and parameterization

The spinner has outer radius $R_{\mathrm{out}} = 1.25$ and body length $L_{\mathrm{b}} = 2.15$. A cylindrical through-hole of radius R_{in} is separated from the fins by a wall of thickness $t_{\mathrm{w}} = 0.12$. The reported radial constraint is

$$R_{\mathrm{in}} + t_{\mathrm{w}} + L_{\mathrm{fin}} = R_{\mathrm{out}}, \tag{6}$$

so the ratio $r = R_{\mathrm{in}}/L_{\mathrm{fin}}$ uniquely determines both R_{in} and L_{fin}. For example, $r = 1.25$ gives $R_{\mathrm{in}} = 0.627778$ and $L_{\mathrm{fin}} = 0.502222$.

The principal design variables are defined in table S1. The total slit width is evaluated on the wall centreline, $R_{\mathrm{c}} = R_{\mathrm{in}} + t_{\mathrm{w}}/2$, as

$$w_T = N R_{\mathrm{c}} \theta_N, \qquad \frac{w_T}{S} = \frac{N\theta_N}{2\pi}, \tag{7}$$

where θ_N is the angular width of one slit in radians and $S = 2\pi R_{\mathrm{c}}$. In the reconstruction, fins and slits share the same axial twist. The reported helical angle α is interpreted at the outer radius through

$$\tan\alpha = \frac{L_{\mathrm{b}}}{R_{\mathrm{ref}}\Phi}, \qquad R_{\mathrm{ref}} = R_{\mathrm{out}}, \tag{8}$$

with total twist Φ in radians. The choice of R_{ref} is a reconstruction convention because neither the native helix equation nor the pitch was reported.

The common reconstructed dimensions are body diameter 2.50, body length 2.15, wall thickness 0.12, fin thickness 0.25, slit axial length 1.75 and planar edge chamfer 0.05. The attached ring has inner diameter 1.0, outer diameter 2.0 and axial length 1.0. For the published optimized geometry $(N, r, \alpha, w_T/S) = (3, 1.25, 60^\circ, 0.75)$, each slit spans 90°.

The two representative Aim 3 geometries are specified in table S2. Ring volume and contact position are held fixed within each matched redesign.

Supplementary Table S1: Geometric and operating variables used in the three aims.

Symbol	Definition	Published-family values	Role in this study
N	Number of helical fins	2, 3 and 4	Reconstructed in Aim 1; jointly selectable in Aim 2; not restricted to 2–4 in Aim 3.
r	$R_{\text{in}}/L_{\text{fin}}$ under equation (6)	0.5–4.0	Reported sweep in Aim 1; jointly varied with α in the Aim 2 grid.
α	Fin helical angle	$20°$–$80°$	Reported sweep in Aim 1; $50°$–$70°$ in the Aim 2 grid; variable-pitch forms permitted in Aim 3.
w_T/S	Normalized total slit width	0.25–0.875	Reported sweep in Aim 1; retained or coupled to new slit patterns in Aims 2 and 3.
f	Rotation frequency	40 Hz to 200 Hz	Operating variable, not a geometric degree of freedom.
μ, ρ	Dynamic viscosity and density	Condition-specific	Fixed within each matched comparison.
m	Ring-magnet state	Present or absent in reported series	Held fixed within a comparison; registered assembly and ring dimensions may vary across Aim 3 designs.

Supplementary Table S2: Geometry of the two representative Aim 3 designs.

Design route	Spinner backbone	Released spinner variables	Ring redesign
Performance-first	$N = 3$, $\alpha = 60°$, $r = 1.25$; $R_{\text{in}} = 0.62778$, $L_{\text{fin}} = 0.50222$	Slit sectors $108°/81°/81°$; fin and bore taper -15%; counter-sweep $-3.75°$	ID 1.00, OD 1.85, length 1.26784; volume matched; two 0.10 chamfers
C3-symmetric	M2-17 spinner: $N = 3$, $\alpha = 65°$, $r = 1.10$, $w_T/S = 0.75$	Spinner unchanged; exact threefold symmetry retained	ID 1.00, OD 1.875, length 1.19255; volume matched; fixed contact position

Agentic CAD reconstruction and projection validation

Reported views, dimensions, topology and design variables were encoded as geometric constraints. Using parametric modelling tools, the agent autonomously generated the corresponding parametric three-dimensional CAD model and rendered the published views. Unreported features were documented as reconstruction choices. Registration allowed rigid pose, orthographic or bounded perspective projection, and at most $\pm 3\%$ isotropic scaling. Anisotropic deformation, reflection and non-rigid warping were excluded.

For each selected view, the paper silhouette was selected before CAD scoring as mask P and compared with the CAD projection C using

$$\mathrm{IoU}(P, C) = \frac{|P \cap C|}{|P \cup C|}. \tag{9}$$

Each reconstruction had to pass solid, dimensional and topology checks and achieve IoU > 0.90 in every selected view. This verifies consistency with the published silhouettes but cannot establish a unique three-dimensional reconstruction.

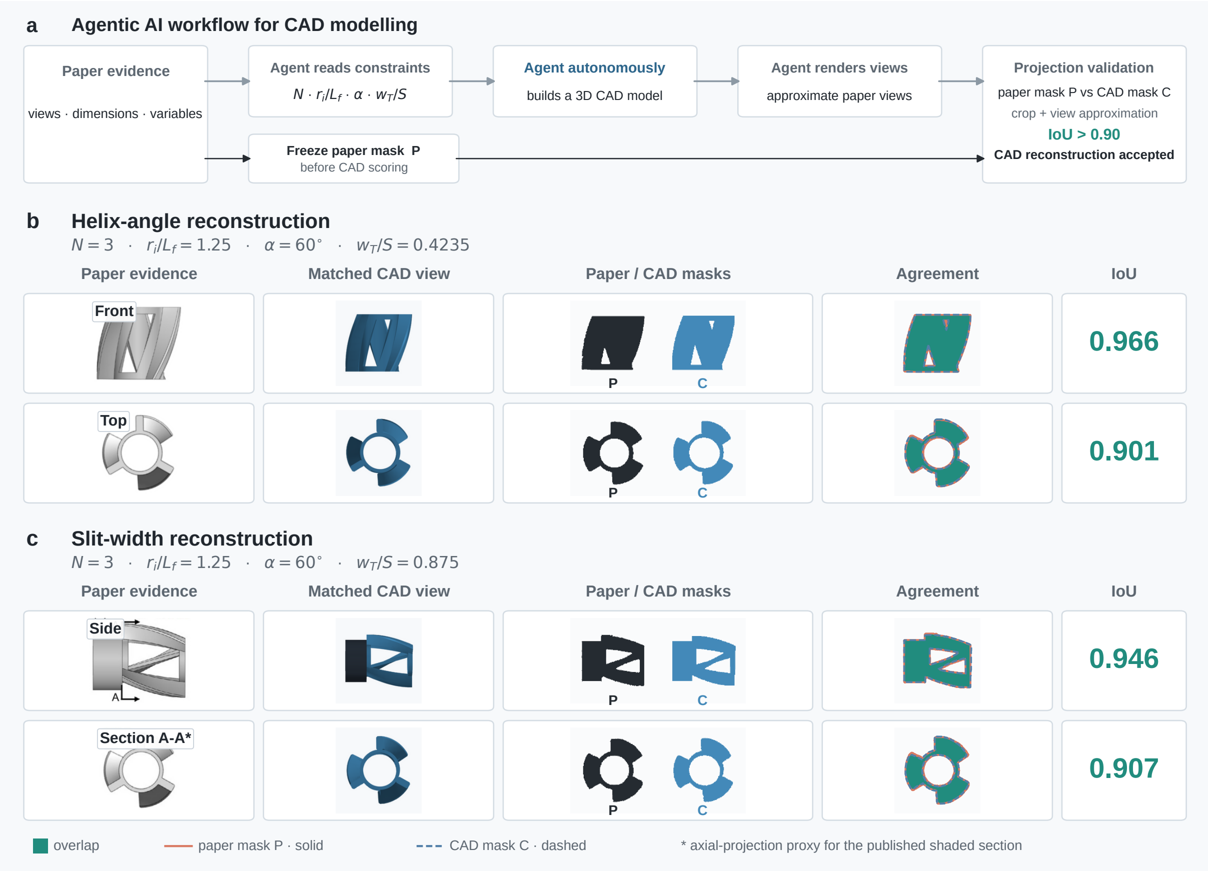

Supplementary Figure S1: **Agentic CAD modelling and multi-view projection validation. a**, The agent converts reported views, dimensions and variables into a three-dimensional parametric CAD model, renders the corresponding views and compares each CAD mask C with a paper mask P fixed before scoring. Each selected view must satisfy IoU > 0.90. **b**, Helix-angle reconstruction with $(N, r_i/L_f, \alpha, w_T/S) = (3, 1.25, 60°, 0.4235)$; front and top IoUs are 0.966 and 0.901. **c**, Slit-width reconstruction with $(N, r_i/L_f, \alpha, w_T/S) = (3, 1.25, 60°, 0.875)$; side and Section A–A IoUs are 0.946 and 0.907. Teal denotes mask overlap, the coral solid contour denotes P, and the blue dashed contour denotes C. The side view includes the ring; the Section A–A comparison follows the ring-free published shaded section and uses an axial-projection proxy in place of a physical plane cut. The top-view match uses a near-orthographic 1° perspective fit at the registered search boundary. Paper evidence is reproduced from the article and Supplementary Information [10]. All four views pass the criterion, showing consistency with the reported structures. Residual differences may reflect source-image cropping and approximate camera registration.

Coordinate and sign conventions

The device axis is $+z$. In the paper-registered frame, the rear face of the ring is $z = 0$, the ring occupies $0 \leq z \leq 1.0$, and the spinner body occupies $1.0 \leq z \leq 3.15$. The solver frame centres the complete assembly at the origin: the ring occupies $-1.575 \leq z \leq -0.575$ and the spinner body $-0.575 \leq z \leq 1.575$. Hence

$$z_{\text{paper}} = z_{\text{solver}} + 1.575 \text{ mm}. \tag{10}$$

The MRF axis is $+z$ through $(x, y) = (0, 0)$. Positive force denotes the $+z$ component of the traction on the rigid wet assembly. The imposed tube-wall velocity is negative during counter-flow calculations, and the reported equilibrium propulsion speed is $U_{\text{eq}} = -U_{\text{wall}} > 0$. The internal pressure contrast, Δp_{int}, is defined as the centreline pressure at the distal device front minus the minimum centreline pressure inside the device. It differs from the difference between inlet and outlet gauge pressures.

Design constraints and degrees of freedom

Supplementary Table S3: Design constraints and degrees of freedom for each aim.

Aim	Controlled variables	Released variables
1	Published vessel, spinner-body envelope, ring state, fluid and frequency for each anchor	One reported geometry or operating variable at a time, following the published sweep.
2	Published parameterization, vessel, ring, fluid and matched operating point	Coupled combinations of N, r, α and slit geometry within the published family. The preliminary grid samples $\alpha \times r$ at $N = 3$, fixed $w_T = 1.83\,\text{mm}$ and $100\,\text{Hz}$; w_T/S co-varies with r.
3	Vessel geometry, spinner-body envelope and actuation condition	Interpretable features beyond the paper family, including fin taper and sweep, variable pitch, nonuniform slits, hole taper, registered assembly and ring dimensions, chamfers, serration, concavity and staged gaps.

Each generated spinner must form a valid single solid with a watertight, consistently oriented surface. It must also satisfy the prescribed envelope and avoid intersections with the fluid domain. These checks assess geometric and numerical admissibility.

Reconstruction boundary

The article does not provide a feature tree, native CAD, helix pitch, exact adhesive geometry, the gap between the magnet and spinner, or the complete set of CFD boundary selections. The reconstruction places the slit on the angular bisector between adjacent fins, uses the same twist for fins and slits, treats the ring as coaxial and in zero-gap contact, and applies the 0.05 edge chamfer in the two-dimensional fin profile without a second three-dimensional end-face chamfer. These choices preserve the reported dimensions and topology and are recorded as reconstruction settings.

B CFD Model and Boundary Conditions

Physical model

The screening model assumes steady, incompressible Newtonian flow with constant properties and represents rotation using a stationary multiple reference frame (MRF). The paper reports constant properties, stationary laminar flow and a frozen-rotor treatment but does not specify a constitutive equation [10]. In the stationary frame, the governing equations are written schematically as

$$\nabla \cdot \boldsymbol{u} = 0, \tag{11}$$

$$\rho(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\nabla p + \mu \nabla^2 \boldsymbol{u} + \boldsymbol{f}_{\mathrm{MRF}}, \tag{12}$$

where $\boldsymbol{f}_{\mathrm{MRF}}$ is restricted to the cylindrical rotating zone. This expression identifies the model class; software-specific MRF interface terms follow the archived OpenFOAM case definition.

For the common 100 Hz saline condition, $\rho = 1000\,\mathrm{kg/m^3}$, $\mu = 0.001\,\mathrm{Pa\,s}$ and $\omega = 2\pi f = 628.319\,\mathrm{rad/s}$. Other frequency and viscosity anchors use their recorded condition values. Each comparison cohort uses one set of fluid properties.

Domain, boundaries and rotation

The reported computational vessel is a straight tube with inner diameter 3.5 mm and length 300 mm. Both axial ends are set to zero gauge pressure. The tube is no-slip with a prescribed axial wall velocity during the force-balance search; the spinner and attached ring are no-slip rotating surfaces. The rotation axis is the device longitudinal axis. The MRF cylinder used for the preliminary evidence has radius 1.40 mm and 0.25 mm axial clearance at each end of the assembly. Its exact extent is a protocol choice because the paper describes the rotating cylinder only as slightly larger than the device.

The paper does not report the original element type, mesh size, boundary-layer count, rotating-interface averaging, solver settings, discretization order, stabilization, convergence tolerances or mesh-independence study. These settings are therefore specified for the present reconstruction. It uses a bounded first-order convection scheme and SIMPLE coupling between pressure and velocity for screening.

Force and pressure observables

Axial force is integrated from the hydrodynamic traction,

$$F_z = \boldsymbol{e}_z \cdot \int_{S_{\mathrm{wet}}} \left[-p\boldsymbol{I} + \mu\left(\nabla\boldsymbol{u} + \nabla\boldsymbol{u}^{\mathsf{T}}\right)\right]\boldsymbol{n}\,\mathrm{d}S. \tag{13}$$

Forces on the spinner, ring and complete assembly are reported separately, and $F_{z,\mathrm{assembly}} = F_{z,\mathrm{spinner}} + F_{z,\mathrm{ring}}$ is verified numerically. The spinner and ring form one rigid body, so assembly force defines the root. Tube patches and patches at the MRF interface are excluded. The authors' exact COMSOL boundary IDs were unavailable, making this force selection a reconstruction choice.

OpenFOAM's kinematic pressure is multiplied by the recorded density before reporting pascals. With $p(z)$ sampled on the tube centreline in the registered paper frame, the reported internal pressure contrast is

$$\Delta p_{\mathrm{int}} = p(z_{\mathrm{front}}) - \min_{z \in \text{device interior}} p(z). \tag{14}$$

This quantity is a localized low-pressure proxy; it is not the conventional upstream-to-downstream hydraulic loss of the vessel. Measured pressure at an accepted fixed point and pressure linearly interpolated between accepted bracket endpoints are retained as different evidence types.

Supplementary Table S4: Mesh tiers used in the preliminary study. Cell counts vary slightly with geometry.

Tier	Surface spacing	Surface level	Typical size	Scientific role
M0	50 µm	1; no prism layers	Approximately 0.68 million cells	High-throughput trend screening, accepted common-speed evidence and provisional root estimation.
M1	25 µm	2; same no-layer family	Approximately 1.55–1.64 million cells	Selected resolution checks and tighter equilibrium calculations; not a universal correction for M0.

Fidelity roles

M0 and M1 use identical geometry, fluid properties, boundary conditions and observable definitions; they differ only in spatial discretization. Broad screening uses M0, while six M1 equilibrium calculations provide a matched resolution comparison.

Convergence assessment and equilibrium estimation

A screening solution must use a valid mesh, finite output and a continuous iteration sequence. After a 500-iteration startup, numerical criteria are evaluated every 250 iterations, beginning at iteration 1000. Initial M0 calculations use 1500 iterations, with selected calculations extended to 2000; M1 screening uses up to 2000 iterations. Pressure-damped comparisons use pressure relaxation 0.1 and up to 4000 iterations. The selected finer-mesh test is described separately in appendix H.

For each iteration, r_U is the largest final initial residual among U_x, U_y, U_z, and r_p is the initial residual from the final pressure solve after the non-orthogonal corrections. Both final values must be no larger than 10^{-4} and, for each field, $\mathrm{median}(r_{n-49:n})/\,\mathrm{median}(r_{n-99:n-50}) \leq 1.05$. Normalized mass imbalance between the inlet and outlet must be no larger than 10^{-3}. Assembly-force closure requires

$$|F_{z,\mathrm{assembly}} - F_{z,\mathrm{spinner}} - F_{z,\mathrm{ring}}| \leq \max\left(10^{-12}\,\mathrm{N}, 10^{-8}\max_j |F_{z,j}|\right). \tag{15}$$

The force-stability statistic is $\delta F = \max(P_{100}^{\mathrm{prev}}, P_{100}^{\mathrm{final}}, |\bar{F}_{100}^{\mathrm{final}} - \bar{F}_{100}^{\mathrm{prev}}|)$, where P denotes peak-to-peak assembly force. It must satisfy

$$\delta F \leq \max\left(2 \times 10^{-7}\,\mathrm{N}, 0.02\max(|\bar{F}_{100}^{\mathrm{final}}|, 10^{-6}\,\mathrm{N})\right). \tag{16}$$

Detailed mesh and claim gates are summarized in appendix C.

An accepted M0 screening equilibrium estimate requires a qualified fixed point satisfying $|F_{z,\mathrm{assembly}}| \leq 5 \times 10^{-6}\,\mathrm{N}$, for which $U_{\mathrm{eq}} = -U_{\mathrm{wall}}$ and Δp_{int} is measured at that point. Points satisfying $|F_{z,\mathrm{assembly}}| \leq 1 \times 10^{-6}\,\mathrm{N}$ are identified as a tighter subset, but the tighter threshold is not the general M0 acceptance rule. Accepted endpoints (U_1, F_1) and (U_2, F_2) with $F_1 F_2 < 0$

Supplementary Table S5: Mesh-quality criteria used for the screening calculations.

Metric	**Limit**	**Purpose**
Maximum non-orthogonality	$\leq 65°$	Limits face-normal misalignment.
Maximum skewness	≤ 4.0	Limits cell-face skew.
Concave-cell fraction	≤ 0.01	Limits widespread concavity.
Low-determinant-cell fraction	$\leq 5 \times 10^{-5}$	Limits strongly distorted cells.
Low-weight faces	0	Excludes interpolation-weight defects.
Nonstandard edge faces	0	Excludes invalid edge connectivity.
Topology and volumes	valid	Required patches, rotating-zone coverage and positive volumes.

may define the diagnostic in-bracket secant prediction

$$U^*_{\text{wall}} = U_1 - F_1 \frac{U_2 - U_1}{F_2 - F_1}, \qquad U_{\text{eq}} = -U^*_{\text{wall}}. \tag{17}$$

The corresponding Δp_{int} is linearly interpolated and labelled derived. The prediction becomes accepted M0 equilibrium evidence only after a simulation at the predicted speed terminates as an accepted fixed point within the 5×10^{-6} N gate. A secant is rejected if either endpoint is unaccepted, the force signs do not oppose, the root lies outside the bracket or the local response is evidently non-monotonic. Mesh-QC, numerical and solver-rejected records never enter equation (17) or a ranking.

C Numerical quality and comparison criteria

Calculation types

Fixed-speed simulations provide hydrodynamic loads and pressure fields at a prescribed opposing flow. An equilibrium-speed estimate additionally requires a converged calculation with assembly force within the specified tolerance. Interpolation between converged, opposite-sign force samples can locate a candidate zero, but is distinguished from a directly evaluated equilibrium point. Calculations with invalid meshes, non-finite outputs or inadequate convergence are excluded from quantitative comparisons.

Literature values are distinguished as reported, digitized or derived. Digitized curves retain the limitations of raster extraction and are not treated as original solver arrays.

Mesh quality

Geometry must form a valid solid with a watertight, consistently oriented surface. The mesh must contain the required boundary patches and rotating zone, including the cells adjacent to moving surfaces. All cells must have positive volume. The mesh-quality limits are listed in table S5.

Local mesh corrections preserve CAD geometry and these numerical limits. Mesh quality and hydrodynamic performance are evaluated separately.

Equilibrium and performance comparisons

An M0 equilibrium estimate requires convergence and $|F_{z,\text{assembly}}| \leq 5 \times 10^{-6}$ N; the 1×10^{-6} N subset uses a tighter force tolerance. Residual, mass-conservation and force-stability criteria are specified in appendix B. Only converged force samples enter the bracketed root search.

Supplementary Table S6: Evidence summary for the three aims.

Aim	Analysis set	Evidence partition	Performance-bearing subset
1	59 benchmark anchors	59/59 configurations with qualified fixed-speed simulations	Three-fin screening U_{eq}: 11/11 at 100 Hz (9 measured, 2 derived secants) and 10/11 accepted fixed points at 160 Hz.
2	25-cell screen plus controlled refinements	Strict-M0 mesh coverage 25/25; accepted M0 U_{eq} 24/25; matched two-frequency and factorial comparisons completed	Preliminary within-family speed winner, near-tied dual-frequency candidates and conditional geometry–pressure responses.
3	Expanded geometry families	Accepted M0 roots at 100 and 160 Hz for the reported asymmetric and C3-symmetric leaders	Functional leaders for speed, local interaction and C3 symmetry; no global, increased-fidelity or experimental optimum.

When several converged near-root calculations are available for the same geometry and operating condition, the representative speed is selected at the smallest absolute assembly force, with force-tail variation as a tie-breaker. An available speed range or bracket is considered where it could affect ranking. Force-tail variation describes numerical convergence, not a complete confidence interval for equilibrium speed.

Design comparisons match fluid properties, frequency, ring state, boundaries and mesh resolution. Common-speed force comparisons are used to interpret hydrodynamic changes, not substituted for equilibrium-speed rankings. Increased-resolution calculations are reported as sensitivity checks and are not combined with screening values in one ranking.

Equilibrium speed is the primary propulsion objective. Internal pressure contrast is a complementary local-suction proxy: a larger contrast denotes a stronger centreline pressure deficit, not greater conventional vessel loss or demonstrated clot capture. Reference-relative improvement at 100 and 160 Hz is distinguished from global optimality or continuous-band performance.

D Screening cohorts and detailed results

This section reports benchmark trend coverage, within-family comparisons and the paired-frequency performance of the expanded designs.

Simulation dataset and comparison sets

Figure S2 summarizes the numbers of geometries, CFD cases, solver executions and qualified equilibrium estimates in the initial screening dataset. The detailed comparisons below include subsequent targeted design studies.

An accepted common-speed calculation provides the force sign at that speed but does not establish an equilibrium root. Measured near-root values and secants derived from accepted

Supplementary Figure S2: **Initial screening dataset.** **a**, The dataset comprises 156 geometries, at least 1,939 unique CFD cases and at least 2,134 solver executions. Of these geometries, 136 reached the solver and 20 did not meet presolver quality criteria. At least 144 points provided equilibrium evidence, including at least 111 directly evaluated roots under the dataset definitions. Repeated or restarted calculations increase solver executions but not unique-case counts. **b**, CFD cases and solver executions by Mission. **c**, Qualified equilibrium evidence and directly evaluated roots. Counts are lower bounds for the initial screening dataset, rather than the combined total of all detailed design studies.

opposite-sign endpoints are displayed separately. Only numerically qualified simulations enter interpolation or performance rankings.

Aim 1 benchmark results

The accepted M0 screen showed full qualitative agreement for four of eight published trend families, partial agreement for three with incomplete U_{eq} coverage, and one disagreement (table S7). All eight published peak neighbourhoods contain accepted common-speed evidence. Exact common-speed endpoint coverage is incomplete for three families, and the five without-ring anchors remain discrete matched checks rather than a connected sweep.

Six matched-physics M1 equilibrium calculations can be paired with accepted M0 estimates. Across these cases, M0 estimates are 11.98–16.78% lower than M1 estimates, with a signed mean difference of -14.33%. The cases cover only two radius-ratio families, so the mean describes only the observed fidelity difference and is neither an uncertainty estimate nor a correction factor.

At 100 Hz, nine of the eleven three-fin values are accepted fixed points and two are in-bracket secants derived from accepted endpoints. At 160 Hz, ten ratios have accepted fixed-point equilibrium estimates and only ratio 0.5 remains unresolved. The 160-Hz curve decreases over the accepted range. Fixed-speed profiles recover an interior force maximum at higher frequencies, but do not replace the equilibrium comparison.

Supplementary Table S7: Aim 1 family comparison. Coverage counts M0 screening U_{eq} estimates on each plotted comparison grid; measured and derived values remain distinct.

Family	Coverage	Published shape	M0 shape	Assessment
2-fin radius ratio	11/11	monotonic decrease	monotonic decrease	agreement
4-fin radius ratio	9/11	unimodal peak	unimodal peak	partial
3-fin radius ratio, 100 Hz	11/11	interior peak	shallow edge maximum, then decrease	disagreement
Helical angle	4/7	unimodal peak	unimodal peak	partial
Slit width	4/6	unimodal peak	unimodal peak	partial
Frequency, $\mu = 0.001\,\mathrm{Pa\,s}$	7/9	monotonic increase	monotonic increase	agreement
Frequency, $\mu = 0.0035\,\mathrm{Pa\,s}$	8/9	monotonic increase	monotonic increase	agreement
Viscosity at 160 Hz	4/6	monotonic decrease	monotonic decrease	agreement

Supplementary Table S8: Selected accepted-M0 values from the normalized-slit Aim 2 two-frequency study.

Geometry	U_{100}	U_{160}	$\Delta p_{\text{int},100}$	$\Delta p_{\text{int},160}$	Interpretation
Reference R: 65°, 1.00, 0.625	0.186560	0.333994	58.512	141.672	same-method reference
M2-17: 65°, 1.10, 0.75	0.184612	0.333487	45.167	136.258	representative near-zero-force solution
Y10: 67.5°, 1.10, 0.75	0.184217	0.334276	48.631	127.113	frequency-dependent speed response
M2-25: 67.5°, 1.10, 0.80	0.182781	0.333994	44.941	121.746	lower Δp_{int} with 100-Hz speed loss

Aim 2 within-family screen

The joint grid combines helical angles 50°, 55°, 60°, 65° and 70° with radius ratios 1.00, 1.15, 1.25, 1.35 and 1.50. It fixes $N = 3$, total slit width $w_T = 1.83\,\mathrm{mm}$ and 100 Hz. With these settings, w_T/S decreases from 0.4660 to 0.3947 across the radius-ratio range. The grid therefore compares complete geometries and cannot isolate the effects of α and r. Versioned mesh repairs recovered all eight initially excluded cells under the original quality limits. The current grid has 25/25 admissible meshes and 24/25 accepted M0 equilibrium estimates. The remaining 50°, $r = 1.50$ cell has accepted fixed-speed evidence but no accepted equilibrium estimate.

The initial 62.5°, $r = 1.00$ refinement preserved 100-Hz speed within 0.0113% and lowered Δp_{int} by 8.05% relative to its matched 65°, $r = 1.00$ reference. Later normalized-slit studies introduced a separate reference and controlled angle, slit and radius comparisons. Selected two-frequency values are given in table S8.

Speeds are in m/s, internal pressure contrasts are in Pa and all four rows are accepted M0 fixed points at both frequencies. Raising angle from 65° to 67.5° at normalized slit width 0.75 reverses the common-speed force response between 100 and 160 Hz. A separate 100-Hz radius-by-slit branch produced a 0.374% speed gain together with a 1.679% increase in Δp_{int}. These conditional effects define separate speed and pressure-response rankings rather than one universal within-family optimum.

Aim 3 beyond-family screen

The Aim 3 search examined nonuniform slit patterns, fin and hole taper, sweep, chamfers and coupled spinner–ring variants. A consistent comparison of three reference structures and six dual-frequency-closed Aim 3 candidates is given in table S9.

Supplementary Table S9: Paired-frequency comparison of the reference and Aim 3 structures. Arrows indicate the favourable direction; bold values are the largest in each metric at both frequencies.

Structure	**Role; symmetry**	$\boldsymbol{U_{\mathrm{eq}}}$ ↑ (cm s^{-1})	$\boldsymbol{\Delta p_{\mathrm{int}}}$ ↑ (Pa)	$\boldsymbol{H_p^*}$ ↑ (μN)	$\boldsymbol{Q_c^*}$ ↑ (μL s^{-1})	$\boldsymbol{J_c^*}$ ↑ (μN)
				100 Hz / 160 Hz		
Reconstructed Lu *et al.* Fig. 8	published reference; C3	18.17 / 32.68	44.4 / 143.6	26.3 / 60.4	120.7 / 135.5	5.51 / 7.68
Aim 1 ratio-peak benchmark	benchmark reference; C3	12.81 / 24.26	36.7 / 131.9	17.9 / 36.8	97.6 / 83.6	3.75 / 3.65
Aim 2 preliminary within-family winner	paper-family reference; C3	18.66 / 33.40	58.5 / 141.7	29.7 / 69.9	135.1 / 155.3	7.07 / 10.06
C3 OD1.88	Aim 3; strict C3	19.50 / 35.02	41.2 / 142.0	27.4 / 63.9	121.8 / 126.7	5.74 / 7.10
C3 OD1.875	Aim 3; strict C3	19.78 / 35.51	40.0 / 133.8	27.0 / 62.7	122.6 / 133.0	5.81 / 7.79
Counter-sweep dual-c10	Aim 3; asymmetric	19.95 / 36.22	30.7 / 115.1	**39.5 / 90.9**	**172.3 / 225.0**	**13.14 / 23.99**
Ring-nose c15	Aim 3; asymmetric	**20.04 / 36.63**	51.9 / 170.3	35.9 / 85.4	138.1 / 161.9	7.06 / 10.47
Ring-nose c10	Aim 3; asymmetric	19.98 / 36.56	52.1 / 171.2	35.9 / 85.3	137.0 / 158.0	6.97 / 10.04
Slit90–taper−40–sweep−7.5	Aim 3; asymmetric	19.08 / 33.97	**57.1 / 176.3**	17.8 / 38.0	88.6 / 86.3	3.11 / 3.85

All U_{eq} and Δp_{int} entries are derived from accepted M0 equilibrium states. The remaining quantities are consistently postprocessed S0 estimates at a 0.25-mm virtual gap. Here Δp_{int} denotes the centreline pressure contrast between the spinner-front reference and the internal low-pressure region; a larger value is interpreted as a stronger localized-suction proxy, not as a favourable increase in conventional hydraulic loss. The asymmetric counter-sweep design combines slit sectors of 108°/81°/81°, negative fin and through-hole taper, a −3.75° counter-sweep and an OD1.85 volume-matched dual-chamfer ring. It exceeded all three declared reference categories in all six frequency-by-reference speed comparisons and produced the largest S0 holding, capture-flow and momentum-flux estimates in the nine-structure comparison (table S10).

Supplementary Table S10: Counter-sweep Aim 3 speed and internal-pressure-contrast changes relative to the three declared reference categories.

Reference	$\boldsymbol{\Delta U_{100}}$	$\boldsymbol{\Delta p_{\mathrm{int},100}}$	$\boldsymbol{\Delta U_{160}}$	$\boldsymbol{\Delta p_{\mathrm{int},160}}$
Reconstructed Lu *et al.* Fig. 8 design	+9.83%	−30.95%	+10.83%	−19.85%
Aim 1 peak-region reference	+55.79%	−16.33%	+49.30%	−12.76%
Aim 2 within-family speed winner	+6.96%	−47.55%	+8.45%	−18.77%

The negative Δp_{int} changes in this table indicate a weaker centreline low-pressure contrast, not reduced hydraulic loss. The same design nevertheless leads the independently postprocessed H_p^*, Q_c^* and J_c^* estimates, demonstrating that Δp_{int} alone does not determine the local-interaction ranking.

The C3-symmetric route retains the M2-17 spinner and changes only the ring from OD2.000 to OD1.875 while conserving ring volume and contact position. Relative to that exact parent, speed increased by 7.12% and 6.47% at 100 and 160 Hz, while Δp_{int} decreased by 10.75% and 1.83%. The nearby OD1.88 design also closed accepted roots at both frequencies and supports the ring-diameter neighbourhood, but is slower than OD1.875 at 160 Hz. Paired-taper and no-sweep controls provide additional context for the expanded geometry comparisons.

Ring-nose c15 is the accepted-M0 speed leader, while the counter-sweep design leads the three S0 interaction estimates. OD1.875 is the fastest tested C3-symmetric design, and OD1.88 trades a small speed reduction for a larger internal pressure contrast. No structure maximizes every reported quantity.

E Discretization sensitivity and the three-fin response

This analysis compares matched mesh resolutions and examines the unresolved three-fin trend. Local mesh corrections preserve the intended geometry and numerical quality limits; unsuitable meshes are not interpreted as low-performing designs.

Matched-resolution comparison

M0 is lower for all six conditions, although the magnitude varies and the sample is neither broad nor randomly selected. The 14.33% mean cannot define a bias, error bar or correction. It provides scale context for the near ties in Aim 2. The asymmetric Aim 3 speed gains relative to the reconstructed Fig. 8 design are of similar order to these cross-tier differences. The same-M0 comparison is internally consistent for screening, but matched increased-fidelity confirmation is required to confirm the quantitative margin. Pressure response is a separate objective and is not corrected by these speed pairs.

Three-fin trend discrepancy

The published 3-fin radius-ratio response is unimodal [10]. At 100 Hz, nine accepted M0 near-root points and two accepted-endpoint secants cover all eleven ratios, but they show only a shallow edge maximum near 0.75 followed by an overall decrease. At 160 Hz, ten accepted near-root points cover ratios 0.75–4.0 and decrease across that range; ratio 0.5 remains unresolved. The disagreement therefore persists after near-complete root coverage. Possible explanations include:

1. the native helix pitch, end treatment, slit placement or ring gap differs from the explicit CAD reconstruction;
2. the unreported rotating-domain extent, rotating-interface treatment or force-surface selection differs between the published COMSOL model and the OpenFOAM MRF model;
3. first-order, no-layer M0 discretization shifts the cancellation between spinner and ring pressure and viscous forces; and
4. digitization of a published raster and missing original solver arrays limit the quantitative benchmark.

Fixed-speed force profiles recover an interior maximum at higher frequencies, but do not resolve the equilibrium discrepancy. A same-speed ratio-1.25 test at 160 Hz changed assembly force from 2.10 to 42.00 µN and Δp_{int} from 131.90 to 111.23 Pa between M0 and mesh-M1. The present data still do not distinguish among the hypotheses above. At ratio 1.25 and 100 Hz, a complete-annular connection changed converged same-speed assembly force from +1.02 to −1.62 µN, an effect of −2.64 µN, while Δp_{int} changed by only +0.22 Pa. At ratio 1.0, the corresponding converged annular calculation had an assembly force of +1.93 µN, compared with +5.51 µN for the matched fixed-speed control. The annular effects were therefore −3.58 and −2.64 µN at ratios 1.0 and 1.25, respectively. Their ratio-1.25-minus-ratio-1.0 difference-in-differences was +0.94 µN. The connection changes force cancellation, but the stronger correction at ratio 1.0 does not support preferential recovery of the reported ratio-1.25 peak. The comparison does not identify which model or discretization choice accounts for the remaining peak discrepancy.

Supplementary Table S11: Exact-physics pairs between M1 equilibrium calculations and qualified M0 estimates.

Case	$U_{\mathrm{eq}}^{\mathrm{M1}}$ (m s^{-1})	$U_{\mathrm{eq}}^{\mathrm{M0}}$ (m s^{-1})	$(U_{\mathrm{M0}} - U_{\mathrm{M1}})/U_{\mathrm{M1}}$ (%)
2-fin, $r = 0.75$	0.158603	0.137060	−13.58
2-fin, $r = 1.75$	0.137550	0.116761	−15.11
2-fin, $r = 2.50$	0.117000	0.098786	−15.57
2-fin, $r = 3.00$	0.103411	0.086059	−16.78
3-fin, $r = 0.75$	0.148127	0.130381	−11.98
3-fin, $r = 1.25$	0.147113	0.128081	−12.94

F Software and computational environment

Parametric geometry was constructed using CadQuery 2.8.0 and OCP 7.9.3.1.1 and exported as STEP and STL. Flow simulations used OpenFOAM v2606 in double precision with 32-bit labels. Postprocessing and quantitative plots used Python 3.10, NumPy and Matplotlib. Each analyzed value is associated with its geometry parameters, fluid properties, frequency, mesh resolution and numerical settings.

The local computational node contained two Intel Xeon Platinum 8358P processors, 64 physical cores, 128 logical CPUs and 503 GiB RAM. Two additional containerized CPU environments each exposed 224 logical CPUs, with a project allocation of 56 CPUs and 120 GiB memory. Screening calculations generally used 16 pinned MPI ranks; selected earlier local calculations used 28. These rank counts describe the screening calculations, not all increased-resolution tests.

The exploratory neural-operator study used one NVIDIA RTX 4090 with 24 GB memory per run, PyTorch 2.10.0 and CUDA 12.8. The reported flow solver ran on CPUs.

Geometry models, mesh definitions, force histories, pressure data, selected flow fields and analysis scripts are retained together with the processed numerical tables. Duplicate physical configurations are reconciled before analysis. Original and derived quantities remain distinguishable so that plotted values can be related to their source calculations. Public data and code availability are addressed in the manuscript declarations.

G Exploratory neural-operator study

An independent exploratory study

Operator learning approximates mappings between functions rather than solving each new state through a complete iterative CFD calculation. Fourier neural operators and DeepONet provide established examples [40, 41]. Their potential use in simulation and design extends beyond interpolation within one fixed numerical problem [42], but generalization remains an empirical question for each dataset and application. Physics-informed neural networks introduce differential-equation constraints into learning [43], within the broader framework of physics-informed machine learning [44]. Such constraints are distinct from the choice of an operator architecture.

Learning can also be incorporated within a numerical solver, as demonstrated by machine-learning-accelerated CFD [45]. That approach is distinct from replacing a full calculation with a geometry-conditioned field predictor. These references motivate the supplementary pilot; their reported performance is not attributed to this project, and no surrogate prediction replaces an accepted flow calculation in the main design comparisons.

The neural-operator/PINO pilot reused archived CFD without generating new solver evidence [46, 47]. The broad archive inventory contains 561 flow-field states from 71 physical geometries; it is not the training denominator for every experiment. The illustrated study pretrained on 178

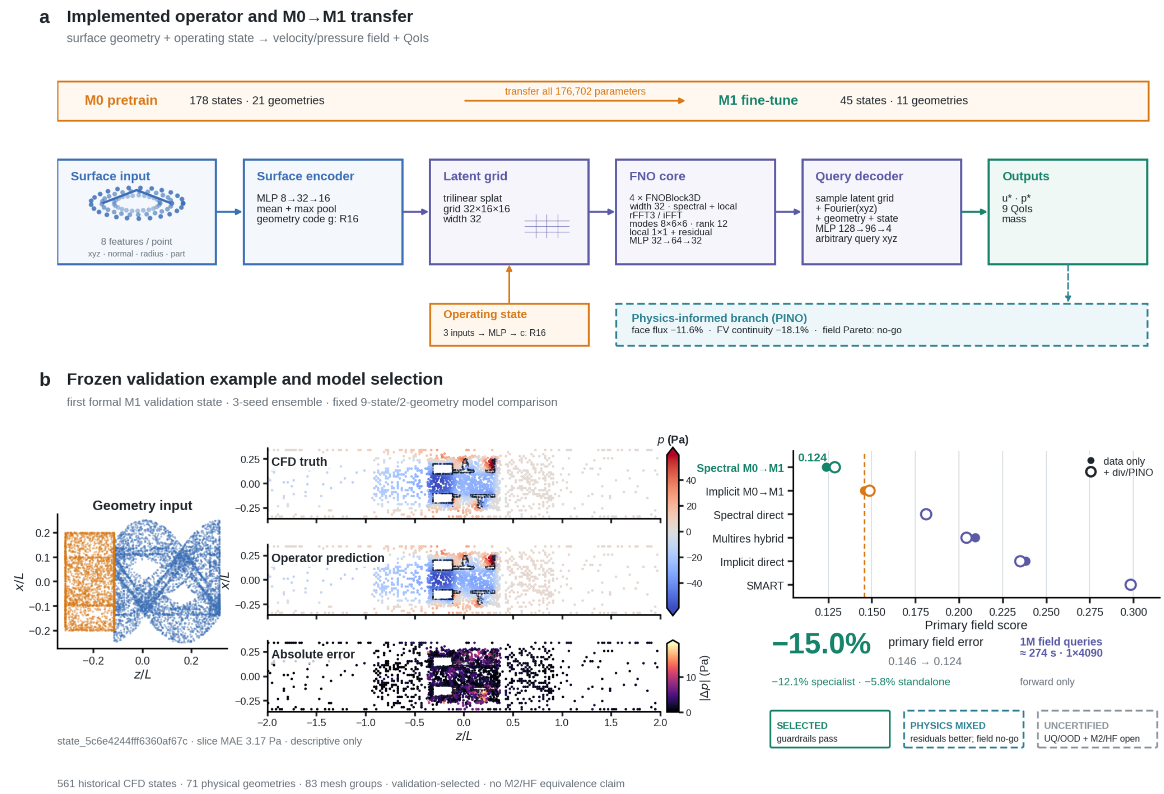

Supplementary Figure S3: **Historical-CFD geometry neural-operator/PINO pilot. a**, Geometry-conditioned prediction with M0 pretraining and M1 fine-tuning. The illustrated training subsets are 178 M0 states from 21 geometries and 45 M1 states from 11 geometries. Physics-informed regularization improved selected conservation residuals but failed the field-accuracy Pareto criterion. **b**, M1 validation example and validation-only comparison of six operator families. Selection used nine states from two physical geometries. One-million-query timing measures forward inference only, excluding preprocessing, ensemble and uncertainty costs. The pilot reports exploratory surrogate performance, not CFD equivalence or independently validated design discovery.

M0 states from 21 geometries and fine-tuned on 45 M1 states from 11 geometries. Model selection used nine M1 validation states from two physical geometries. This is a small development-stage comparison, not an independent held-out test demonstrating topology-wide generalization.

The prototype maps geometry and operating conditions to arbitrary-query velocity and pressure fields and auxiliary quantities. M0-to-M1 transfer reduced a validation-selected ensemble composite error by 31.3% relative to direct M1 training in its development cohort. This composite metric is distinct from the primary field score, for which spectral transfer improved 0.146 to 0.124 relative to an implicit-transfer incumbent. Baselines, cohorts and denominators must remain separate. The PINO branch reduced relative face-flux and finite-volume continuity residuals by 11.6% and 18.1%, but did not improve the joint field-accuracy criteria. Better conservation residuals did not establish uniformly better field prediction.

Forward inference for one million queries required minutes on one RTX 4090, excluding preprocessing and ensemble or uncertainty costs. This is an inference-only measurement; queries are not independent CFD simulations. The reported candidates were obtained from the CFD design programme, independently of this pilot.

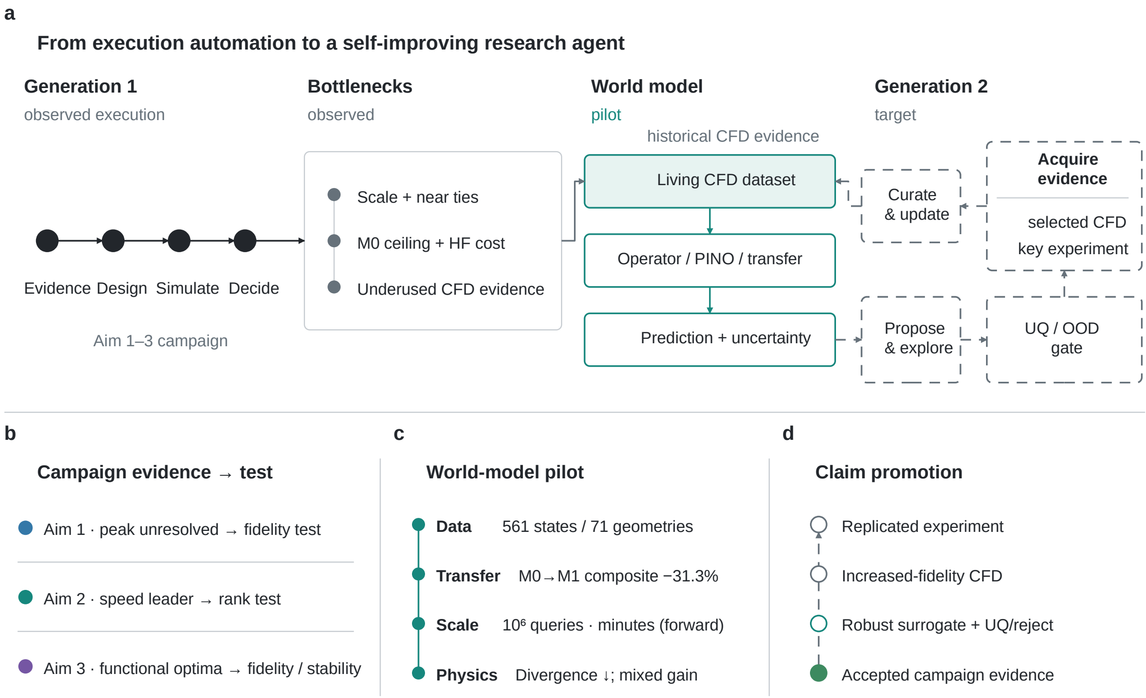

Supplementary Figure S4: **Exploratory data reuse and prospective workflow extensions.** **a**, Existing CAD/CFD execution and a target uncertainty-routed research loop. **b**, Evidence gaps and prospective discriminating tests. **c**, Archive-scale and validation-selected pilot signals, with cohort and metric qualifications in the text. **d**, Conditional progression toward selected fidelity checks and experiments. Future components are not completed results, and the neural-operator pilot is not evidence of CFD-equivalent accuracy or end-to-end search efficiency.

Architecture diagram legend

In Supplementary Fig. S4, dashed learned-screening and experiment-routing components denote proposed capabilities rather than evaluated modules; these are distinguished from the CFD workflow and pilot results.

H Resolution sensitivity of the final published geometry

Selected higher-resolution calculations examine the absolute equilibrium speed of the reconstructed final published geometry. They are separate from the common-M0 candidate comparisons.

At 100 Hz, the finer no-layer mesh contains 9,201,337 cells and satisfies the mesh-quality limits in table S5. The converged equilibrium calculation gives

$$U_{\mathrm{eq}} = 23.819432\ \mathrm{cm\,s^{-1}}, \tag{18}$$

$$F_{z,\mathrm{assembly}} = +0.096374\ \mu\mathrm{N}, \tag{19}$$

$$\Delta p_{\mathrm{int}} = 67.250492\ \mathrm{Pa}. \tag{20}$$

Spinner and ring forces are approximately +152.077864 and −151.981491 µN, respectively. Their near cancellation emphasizes the importance of resolving both components of the rigid assembly. The calculation uses a 1×10^{-6} N equilibrium-force tolerance.

The digitized published CFD value is 26.666667 cm s^{-1}, giving an absolute difference of 2.847235 cm s^{-1}, or 10.6771%. An intermediate M1-resolution calculation gives approximately

21.8003 cm s^{-1}, 18.25% below the same target. The finer-resolution value is closer to the published CFD value.

The equilibrium value above is the sole reported root in this higher-resolution dataset; the 160-Hz condition and neighbouring geometry have no reported converged higher-resolution roots. The finer-resolution reference is not substituted into comparisons with M0 candidates.