

Closing the Loop in the Lab:

Feasibility of a Physical Genetic Algorithm Approach for Optimising Thermoresponsive Hydrogel Morphologies

By

KIEWEK, LUCAS

MSc Biorobotics Dissertation



ETH zürich

School of Engineering Mathematics and Technology
UNIVERSITY OF BRISTOL
&
Department of Material Sciences
ETH ZÜRICH

A MSc dissertation submitted to the University of Bristol in
accordance with the requirements of the degree of MASTER OF
SCIENCE IN BIOROBOTICS in the Faculty of Engineering.

August 29, 2025

Declaration of own work

I declare that the work in this MSc dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

Lucas Kiewek 29/08/2025

Ethics statement

This project did not require ethical review as determined by my supervisor Dr. Helmut Hauser

Lucas Kiewek 29/08/2025

Closing the Loop in the Lab: Feasibility of a Physical Genetic Algorithm Approach for Optimising Thermoresponsive Hydrogel Morphologies

LUCAS KIEWEK

Abstract—The design space for hydrogel micro-actuators is expanding rapidly, increasing the need for systematic, physical optimisation beyond simulation. This work implemented a proof-of-concept genetic-algorithm (GA) pipeline that used real-world feedback to evolve poly(N-isopropylacrylamide-co-acrylic acid) (PNIPAm-co-AAc) micro-actuators. Candidate morphologies were encoded, fabricated, and actuated; fitness was then assigned by quantifying non-reciprocal deformation from tracked trajectories. The pipeline produced 16 generations (25 candidates each; 400 actuators) over 11 working days. Population medians did not improve and per-generation maxima declined, while the top actuator in each generation showed increasing path amplitude. Analysis indicated a measurement-selection misalignment where motion-based seed selection emphasised large displacement, whereas the fitness rewarded geometric non-reciprocity. The concurrent rise in voxel clustering is consistent with morphological changes accompanying the observed trends. This work provides a practical basis for GA-driven micro-scale morphological optimisation once the alignment between seed selection and fitness is corrected.

I. INTRODUCTION

Genetic algorithms (GAs) are search and optimisation methods inspired by Darwinian evolution, in which populations of encoded solutions undergo fitness-based selection, recombination, and mutation so that fitter individuals are more likely to propagate. Originating in Holland’s mathematical formalism for adaptive systems, GAs provided a theoretical basis that catalysed decades of work on evolving solutions rather than hand-designing them [1]. This evolutionary approach has been progressively adapted in robotics because it effectively navigates large, multimodal, or non-differentiable design spaces where gradients are unavailable or deceptive [2]. This work is focused on the implementation of an evolutionary pipeline for evolving morphologies at the micro-scale, where perfecting parameters can be particularly perplexing.

A. Towards Genetic Algorithms for Microrobotics

Early research in evolutionary robotics used GAs to generate virtual, macro-scale morphologies which could be cheaply evaluated through simulation at the expense of approximated real-world physics [3]. Advancements in fabrication techniques such as additive manufacturing enabled simulated morphologies to be realised physically [4]. However, these physical morphologies are often less fit than their virtual counterparts due to “the reality gap” [5]. Current work seeks to close this gap by incorporating real-world evaluation and iteration [6]. At the micro-scale, GA-driven morphology

searches largely remain simulation-based because iterative microfabrication and measurement are technically demanding [7], [8].




Fig. 1. Closed-loop pipeline for in situ micro-actuator evolution: voxel genome, 2PP-DLW fabrication, global actuation, video measurement, and selection.

B. Motivation

Poly(N-isopropylacrylamide) (PNIPAm) is a thermoresponsive hydrogel whose reversible swelling–collapse is routinely harnessed to realise microactuators [9]. Recent work from ETH Zürich’s Soft Materials and Interfaces (SMI) group broadens the palette for micro-scale actuation, demonstrating tunable thermoresponsive hydrogels, multi-material constructs, and chemically programmable hysteresis [10], [11]. These modalities expand the design space and create a need for systematic methods to navigate it. Furthermore, two-photon polymerisation direct laser writing (2PP-DLW), the fabrication method used here, can exhibit process-dependent distortions, such as polymerisation blur and layer warping, that decouple fabricated actuators from their virtual blueprints, a micro-scale analogue of the reality gap [12]. Together, these factors motivate an evolutionary approach that uses real-world feedback to discover functional hydrogel microactuator designs.

C. Objective and Approach

The goal of this project was to explore the feasibility of a GA-based, in situ optimisation pipeline for evolving PNIPAm microactuators by implementing one practically in a laboratory setting. This involved solving challenges in iteration-friendly microfabrication, mapping physical measurements to genome performance, and mitigating unintended selection pressures. The GA design was heavily inspired by “Evolving and generalising morphologies for locomoting micro-scale robotic agents” [7]. The laboratory methods and materials followed Shen et al. and van Kesteren et al., simplified and adapted to support reliable, iterative fabrication [11], [10]. Insights from early explorations were progressively incorporated until the loop connecting virtual and physical processes was closed.




Fig. 2. Schematic of the isotropic, reciprocal thermal response of a typical PNIPAm hydrogel structure.

D. The Pipeline

The pipeline links encoding, fabrication, actuation, measurement, and selection into a single closed loop. Generations of candidate morphologies are encoded as binary genomes and fabricated as anchored PNIPAm-based microactuators by 2PP-DLW in array format to maximise throughput and simplify identification. A global thermal stimulus elicits repeatable swelling–collapse cycles that are recorded under a microscope. The recordings are processed to extract deformation trajectories from optically tracked points (seeds) on each actuator. Each actuator’s fitness is computed as the average of scalars that quantify non-reciprocity of seed trajectories by calculating the geometric separation between its swelling and collapse trajectories. This criterion, rather than net displacement, was chosen to test whether the pipeline can optimise a non-trivial behaviour with practical uses at low-Reynolds-number. While locomotion is not demonstrated here, success on this objective could inform the design of hydrogel-based, Purcell-style swimmers because non-reciprocal motion is essential for overcoming the “scallop theorem” [13], [11]. Selection, crossover, and newcomer injection then produce the next population, enabling an in situ evolutionary search. The sections that follow detail each module of the pipeline and evaluate feasibility, throughput, and the alignment between the measured fitness and its intended objective.

II. RESEARCH METHODOLOGY

A. Study design

An in situ genetic-algorithm pipeline for PNIPAm microactuator optimisation, as outlined in subsection I-D, was implemented as a proof of concept. Each generation proceeded through genome encoding, 2PP-DLW of arrays of anchored actuators, global thermal actuation, video acquisition of swelling–collapse cycles, and selection to generate subsequent genomes. Experimental conditions for fabrication, actuation, imaging, and analysis were standardised so that only the genome varied between generations. Before multi-generation evolution, non-reciprocal deformation of an anchored actuator was demonstrated under these conditions to verify that the target behaviour was attainable and measurable (§II-B). Fitness quantified path non-reciprocity (“higher is better”; §II-H) from tracked deformation trajectories and was used to drive selection and crossover. Detailed materials, parameter values, and algorithms are provided in §§II-B–I; mid-study adjustments are listed in §II-L.




Fig. 3. Exploratory PNIPAm morphologies with positive/negative features showing pronounced shape change under actuation.

B. Groundwork

Printing readiness, material formulation, and morphology heuristics were established before multi-generation evolution. 2PP-DLW parameters were selected to produce repeatable anchored prints suitable for actuation. Poly(N-isopropylacrylamide-co-acrylic acid) (PNIPAm-co-AAc) with 10 mg BIS crosslinker and 5 wt % AAc was adopted to realise a comparatively looser network and larger swelling–collapse amplitude than the tighter mesh of PNIPAm with higher cross-linkage. Prior exploratory experiments, depicted in Fig. 3, indicated that asymmetric distributions of positive (printed) and negative (void) space elicited more pronounced deformation and cycle hysteresis than uniformly filled geometries; this informed the genome-to-voxel mapping used in §II-E.

C. Demonstration of non-reciprocal deformation

Non-reciprocal deformation was demonstrated on anchored single actuators fabricated from PNIPAm-co-AAc (BIS 10 mg). Three exploratory morphology families were

printed; one exemplar exhibited a clear, approximately elliptical trajectory in the image plane over a swelling-collapse cycle when tracked by simple point-based analysis (Fig. 4). The result confirmed that the target behaviour was attainable and measurable with the chosen setup; subsequent sections describe the evolutionary pipeline that operated under these fixed conditions.




Fig. 4. Demonstration of non-reciprocal deformation: (A) tracked trajectories forming a loop; (B) microscope image with tracked points; (C) three morphology families (bumps; bumps+holes; holes), with bumps yielding the exemplar.

D. Materials and apparatus

Prepolymer formulations. Two hydrogel formulations were used. For parameter screens and early experiments, PNIPAm was prepared by dissolving N-isopropylacrylamide (NIPAM, 200 mg), N,N'-methylenebisacrylamide (BIS, 15 mg), and diphenyl(2,4,6-trimethylbenzoyl)phosphine oxide (TPO, 50 mg) in 225 μ L 1,2-propanediol; 0.5 mg fluorescein was included as a tracer. For the evolutionary runs, PNIPAm-co-AAc comprised NIPAM (200 mg), BIS (10 mg), TPO (50 mg), 1,2-propanediol (225 μ L), and acrylic acid (2.9 μ L); no dye was added. The reduced crosslinker content and dye removal produced larger, faster swelling near $\text{pH} \approx 6$ and were therefore adopted for evolution.

Substrates and surface preparation. Borosilicate glass slides were vinyl-functionalised to promote adhesion of anchored actuators. The functionalisation used vinyltrimethoxysilane (VTMS, 98%), ethanol, acetone, and nitrogen gas for drying; slides were oxygen-plasma activated before silanisation. The preparation protocol is reported in §II-F.

Fabrication system. Microstructures were fabricated by 2PP-DLW on a Nanoscribe Photonic Professional GT2 (model-specific optics and print parameters in §II-F).

Actuation and imaging platform. Actuation and imaging were performed on a Nikon Eclipse Ti2 inverted microscope. A simple capillary chamber was assembled as described in §II-G; video acquisition settings are given in §II-H.

Software. Geometry generation was performed in Blender; print slicing and job preparation were done in DeScribe; video analysis used Python with OpenCV.

E. Genome encoding and morphology generation

Each candidate morphology was encoded as a fixed-length 200-bit genome that addressed a $5 \times 5 \times 8$ lattice of 6 μ m cubic voxels (overall envelope $30 \times 30 \times 48 \mu$ m). A bit value of 1 specified a printed voxel and 0 a void. Bits were read in a fixed order with z varying fastest, then y, then x, so that reconstruction of the virtual 3D occupancy from the genome was deterministic.




Fig. 5. Genome-to-geometry workflow: codes.csv read in Blender, voxelised and tiled as a 5×5 array, then exported as STL for DeScribe.

A $54 \times 54 \times 2.05 \mu$ m base was appended beneath the lattice with a slight overlap into the body to promote adhesion during development and handling. To ensure unambiguous spatial indexing on circular slides, a small generation label was printed just outside the 5×5 array. The label fixed the array orientation so that the [0,0] actuator corresponded to the upper-left position, thereby preserving the mapping from array position to actuator to genome across fabrication, imaging, and analysis, while remaining outside the imaging field.

For each generation, a codes.csv file listed 25 individuals (GlobalID, Code). A bespoke Blender script read this file, reconstructed each geometry as specified above, and tiled the actuators into a 5×5 array with 80 μ m centre-to-centre spacing (Fig. 5). The scene was exported as STL for slicing in DeScribe prior to 2PP-DLW fabrication (§II-F). Designs were not pre-filtered for manufacturability and were recorded if they failed.

Design rationale. A voxel representation was adopted because voxel-based encodings are widely used in evolutionary morphology and have been demonstrated at the microscale [7]. It also allowed the direct control of positive versus negative space, which aligned with preparatory observations about cycle hysteresis. The representation also allowed single-point crossover to act on spatially coherent substructures.

F. Fabrication by 2PP-DLW

Substrate functionalisation (vinyl-silane). Glass slides were rinsed with ethanol and dried under nitrogen, oxygen-

plasma activated, and placed in a sealed chamber containing 0.4 mL vinyltrimethoxysilane (VTMS, 98%) for overnight chemical vapour deposition at < 0.05 kPa. After silanisation, substrates were rinsed with acetone, dried under nitrogen, and mounted with the treated face oriented toward the resin during printing.

Print file preparation. Genome-derived geometries (§II-E) were exported as STL and imported into DeScribe for slicing and job preparation. The profile “IP Dip 63× Fused Silica (3D SF)” was used with 0.2 μm slicing, 0.2 μm hatching, contours disabled, hatching angle 90° with 30° offset, block size $80 \times 80 \times 50$ μm with block offset +30, +30, 0 μm , and zig-zag write order.

Writing. Structures were fabricated on a Nanoscribe Photonic Professional GT2 with a 63× oil-immersion objective. The generated GWL script was edited to set 90% laser power and 10 000 $\mu\text{m s}^{-1}$ scan speed. Silanised slides were mounted on the holder’s oil side and secured; a PDMS spacer window on the top face was filled with the prepolymer (§II-D). Immersion oil was applied to the opposite face. After locating the resin-glass interface, a debris-free write field was selected. Anchors and bases were written first, followed by the voxelised bodies.

Development and hydration. Following exposure, samples were developed in PGMEA (10 min), rinsed in IPA (3 min), and dried gently under nitrogen gas. Prints were then equilibrated in water before chamber assembly (§II-G).

Acceptance. Arrays advanced to actuation when at most five actuators failed to polymerise or adhere to the glass slide and no gross defects or occluding debris were observed. Generations failing these criteria were reprinted.

G. Chamber assembly and thermal actuation

Chamber assembly. A 0.12 mm double-sided adhesive spacer was centred over the printed array. A gold-coated slide was placed gold side down to close the chamber. 10 μL of pH 6 Milli-Q water was introduced by micropipette as the working fluid; capillarity filled the cavity. Two sides were left open for air egress, and complete immersion was verified visually.

Thermal stimulus. Thermal actuation was provided by the microscope’s Lumencor fluorescence channels at 100% intensity in 3 s on / 3 s off cycles (5 repeats). Samples were illuminated for approximately 1 minute prior to beginning the cycles to pre-bias them near the PNIPAm transition thus eliciting an immediate response. Illumination served as a global stimulus.

H. Imaging and dataset

Imaging configuration. Imaging was performed on a Nikon Eclipse Ti2 inverted microscope with a 20× objective and a 1.5× intermediate magnifier. Brightfield illumination was DIA LED at 100% (Shutter: FL Lo). The filter train was set to Pinkel Quad on the turret (DAPI/FITC/TRITC/Cy5) with Cy3 emission on the EM wheel. The field of view was positioned to frame the 5×5 array with the orientation label kept out of view.

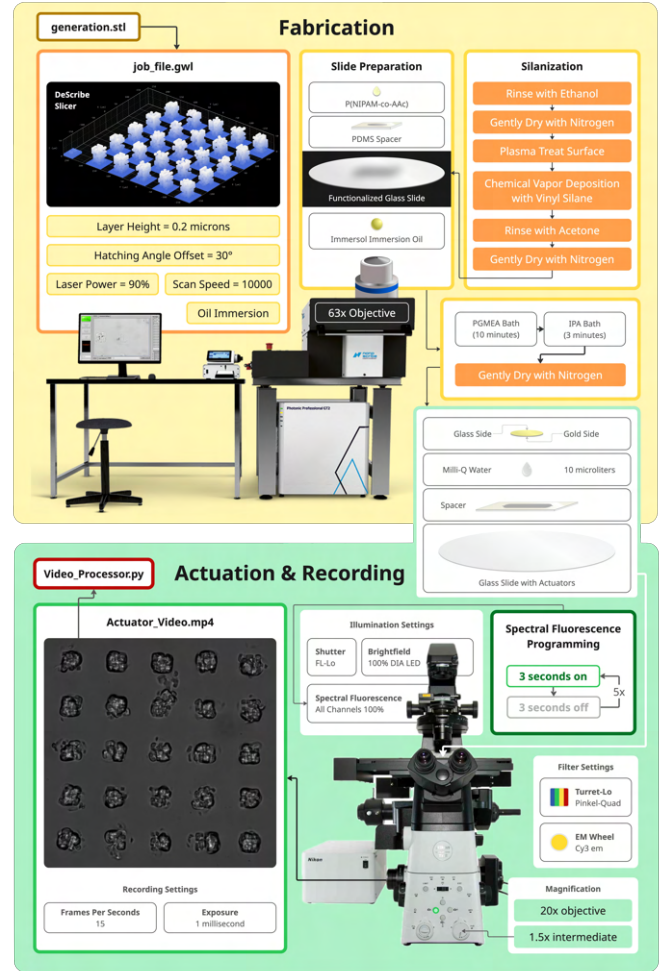


Fig. 6. Physical protocol overview: slide functionalisation, 2PP-DLW writing, development, hydration, and capillary-chamber assembly.

Acquisition settings. Videos were recorded at 15 fps with 1 ms exposure. For each generation, a single video of the entire array was acquired while running the thermal actuation programme (§II-G).

I. Trajectory extraction and cycle segmentation

Preprocessing and phase detection. Videos were converted to grayscale and thresholded with Otsu’s method to separate the actuators from the background [14]. A light morphological open–close removed speckle while preserving boundaries. For each frame, the area of the largest external contour was used as a simple proxy for degree of actuator swelling. The resulting area–time trace was smoothed (Savitzky–Golay) and analysed for alternating growth and shrinkage phases that matched the fluorescence-driven thermal actuation schedule (§II-G) [15]. Simple timing rules were applied to suppress spurious flips from noise or debris, namely a minimum time between peaks and troughs, minimum peak prominence, and a short “debounce” so a phase had to persist before switching. Minor robustness adjustments made after early runs are listed in §II-L; settings were then fixed for all subsequent generations.

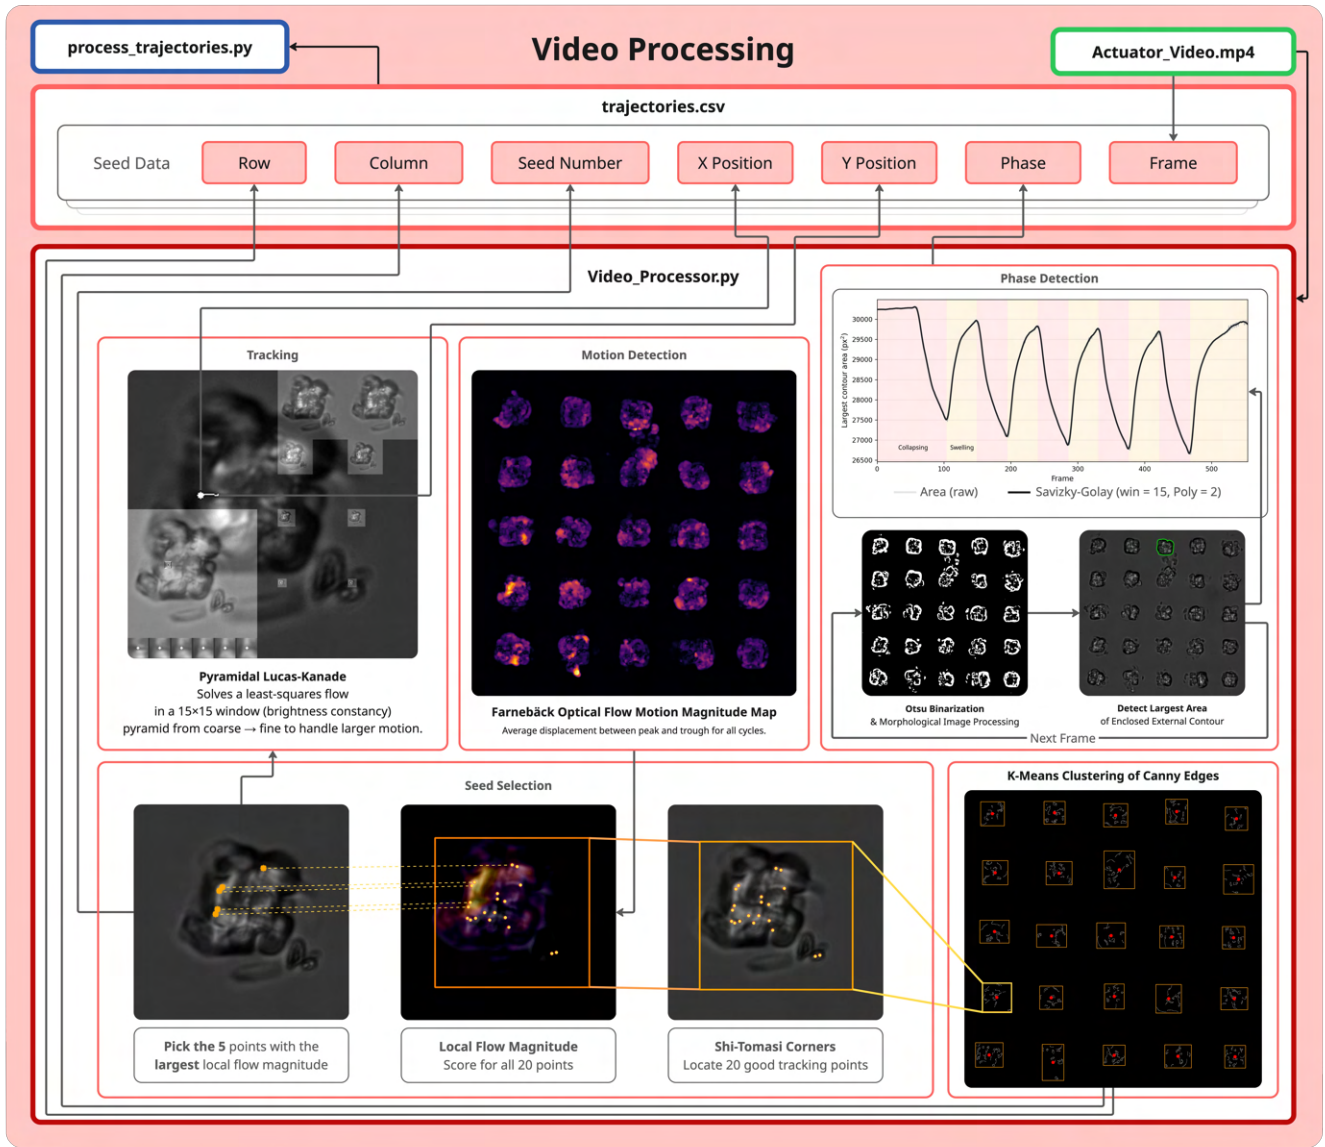


Fig. 7. Video-analysis pipeline: preprocessing and phase detection, ROI inference, feature selection and tracking, and seed-trajectory extraction.

Actuator localisation. The field of view was partitioned into 25 regions of interest (ROIs), one per actuator in the 5×5 array. On the frame with maximal swelling, Canny edge detection provided high-contrast boundaries, which were clustered by k-means ($k = 25$); the bounding box of each cluster defined an ROI which was ordered in row-major order to match the indexing of the physical actuators [16]. Missing or damaged actuators were masked via a manual exclusion list. When edge contrast was insufficient, a uniform 5×5 tiling was used as a fallback.

Feature selection and tracking. Within each ROI, up to 20 candidate points were detected using the Shi–Tomasi corner measure, which favours well-defined, trackable image features [17]. To select for points with greater movement over a half-cycle, optical flow was computed which estimates the apparent motion of image patterns between consecutive

frames. A dense Farneback method provided a per-pixel motion field; its magnitude was averaged over the swell and over the collapse half-cycles to form a motion-saliency map [18]. The five highest-saliency points were retained as seeds. Seed trajectories were then obtained with pyramidal Lucas–Kanade optical flow, which estimates point motion by local image matching and handles larger displacements via image pyramids [19]. Tracked (x, y) positions with phase labels were exported for fitness computation (§II-J). A processed video with overlays was generated for quality assurance.

Algorithm choice justification. The pipeline used broadly adopted, well-characterised computer-vision methods that are fast, reproducible, and require minimal manual tuning.




Fig. 8. Fitness computation: target swell/collapse half-cycles selected, paths aligned by DTW, mean-squared separation computed, and the log updated.

J. Fitness definition and validation

Phase selection. For each actuator, only the final collapse and the immediately preceding swell half-cycles were used for scoring, as trajectories were most stable by this point (§II-G).

Definition. For each tracked seed (§II-I), the swell trajectory was compared with the time-reversed collapse trajectory. The two paths were aligned by dynamic time warping (DTW) to allow small timing mismatches. The per-seed score was the mean squared Euclidean distance between DTW-aligned point pairs in the image plane (pixels). Actuator fitness was the arithmetic mean of the five seed scores. Actuators with no valid seed trajectories were assigned a fitness of 0.0. Fitness was maximised; larger values indicate greater cycle non-reciprocity.

Rationale. At low Reynolds number, motion is effectively kinematic: the outcome depends on the sequence of shapes, not the rate at which they are executed (Scallop Theorem) [13]. A suitable metric for non-reciprocal motion should therefore emphasise path geometry and be insensitive to time parameterisation. DTW provides this by aligning swell and (time-reversed) collapse such that the resulting squared-

distance captures geometric separation of the two loops while discounting small timing offsets from imaging or actuation.

K. Evolutionary loop

Population and encoding. Each generation contained 25 individuals with genomes comprised of 200 bits (§II-E).

Initialisation. The seed generation was created by independent sampling at each locus (Bernoulli $p = 0.5$ per bit).

Fitness evaluation. Fitness was computed from tracked trajectories as described in §II-J.

Selection and variation. An elitist generational scheme was used. The top two individuals (highest fitness) were carried unchanged (elites). Parents were chosen by tournament selection ($k = 5$). Single-point crossover produced 13 offspring. No per-bit mutation was applied.

Exploration. To maintain diversity, 10 randomly generated 200-bit genomes (Bernoulli $p = 0.5$ per bit) were added each generation.

Replacement. The next generation comprised 2 elites, 13 crossover offspring, and 10 newcomers ($n = 25$).

Termination. The experiment proceeded sequentially until the available time budget elapsed.




Fig. 9. Genetic-algorithm flow: population initialisation, evaluation, elitism and tournament selection with single-point crossover, newcomer injection, and next-generation creation.

Algorithm choice justification. Operator choices were adapted from Uppington et al. for a fully physical loop with limited, noisy evaluations [7]. Elitism carried forward the top designs despite noisy scores; tournament selection applied steady selection pressure via simple relative comparisons; and single-point crossover exchanged contiguous gene blocks, maintaining spatial coherence in the voxel encoding. Exploration was guaranteed by injecting ten random newcomers each generation in place of per-bit mutation, avoiding extra hyperparameters.

L. Protocol adjustments

Working fluid volume. Immersion volume was standardised to 10 μL of pH ≈ 6 Milli-Q water for all subsequent runs. Earlier variability in volume may have negatively impacted fitness; volume was fixed thereafter (§II-G).

Phase detection robustness. The phase-labelling pipeline was made more robust by (i) applying speckle suppression and smoothing the area trace before peak finding, (ii) enforcing cycle-aware timing rules matched to the actuation programme, (iii) adding a debounce so that a new phase had to persist before a switch was accepted, and (iv) terminating in swelling to match the actuation schedule (§II-I).

ROI fallback and exclusions. When edge contrast was insufficient for Canny and k-means ROI estimation, a uniform 5×5 tiling of the field of view was used as a fallback. Missing or damaged actuators were masked via a manual exclusion list; ordering remained row-major to preserve array indexing (§II-I).

III. RESULTS




Fig. 10. Distribution of actuator fitness across generations ($n=25$ per generation). Population trend and top-actuator series shown (top); within-generation right-skew visible (bottom).

A. Data yield & operating context

The evolutionary pipeline was operated over eleven working days, during which 16 generations, comprising 400 micro-actuators, were fabricated. Because the project ran on

a compressed timeline, issues were identified and corrected during operation. Early production was limited by fault catching, failed prints, and operator error, yielding fewer than one successful generation per day; throughput later stabilised at approximately three generations per day.

B. Population distribution

Figure 10 visualised the distribution of actuator fitness across generations. The lower panel indicated right skew of the fitness distributions, consistent with a coefficient of variation of approximately 1 across generations. For this reason, the median was used to summarise population-level behaviour, while the maximum was used to characterise selection behaviour, as shown in the upper panel.




Fig. 11. Median fitness by origin type (elite, crossover, newcomer) for each generation; no improvement observed.

C. Preliminary Trends

When the per-generation median fitness was plotted together with medians stratified by origin type (elite, crossover, random) in Figure 11, no positive trend was observed. A slight negative association was indicated by Kendall's $\tau = -0.467$ with $p = 0.0115$. This lack of improvement in the medians suggested that overall fitness did not increase across generations.




Fig. 12. Best-of-generation fitness across 16 generations (linear fit shown), indicating a declining trend.

D. Selection Behaviour

Top-per-generation fitness was then regressed against generation in Fig. 12 to test whether the algorithm selected for fitness itself. No increase was observed (slope = -0.097 , $R^2 = 0.46$, $p = 0.004$), indicating that higher fitness, and thus non-reciprocal motion, was not being optimised.




Fig. 13. Start-normalised seed trajectories for top actuators across generations; dots mark path amplitude (pixel distance from start).




Fig. 14. Increase in path amplitude of top seeds across generations with linear fit.

E. Trajectory Analysis

To investigate what was favoured, trajectories of the best seeds within the top actuator of each generation were examined. Fig. 13 shows the flower-like data created when plotting these paths normalised to the same starting point. The darker petals created by the seeds of earlier generations appear to

have a more squat, rounded shape compared to the longer, thinner petals from later generations.

The paths lengthened over time. Quantifying this as the distance from the starting point to the farthest point on the path (amplitude) and plotting amplitude against generation in Fig. 14 yielded a positive linear trend (slope = 0.381, $R^2 = 0.48$, $p = 0.00292$).




Fig. 15. Relationships for top actuators: voxel spatial autocorrelation (Moran's I) versus generation, amplitude, and fitness.

F. Morphological Trends

To relate this behaviour to morphology, Moran's I was computed for the most fit actuator of each generation and compared with generation, fitness, and amplitude (Figure 15). Moran's I quantifies spatial autocorrelation in each actuator's $5 \times 5 \times 8$ binary morphology to measure the extent to which occupied voxels formed clusters after accounting for overall occupancy. A 3D king's-move weight matrix (face/edge/corner adjacency, open boundaries) was applied to reflect local physical coupling in the lattice. Moran's I increased with generation (Spearman $\rho = 0.594$, $p = 0.015$), correlated positively with amplitude (Spearman $\rho = 0.381$, $p = 0.145$), and correlated negatively with fitness (Spearman $\rho = -0.430$, $p = 0.096$). These relationships suggested that morphologies became more spatially structured in a manner associated with larger path excursions rather than with higher fitness. Overall, the evidence indicated that the pipeline exerted selection pressure on amplitude rather than on non-reciprocal actuation.

IV. DISCUSSION AND CONCLUSIONS

This work advanced evolutionary morphology by demonstrating an iterative microfabrication process that linked physical performance to morphological selection. Methodologically, it contributed a practical package for physical evolution: a voxel genome compatible with 2PP-DLW, array-level printing and printed orientation labels to preserve identity, and automated video analysis for trajectory extraction and fitness computation.

A. Key Findings

The pipeline was executed repeatedly at a useful throughput (~ 3 generations per day) for 16 generations (400 actuators) with stable, repeatable operation, confirming practical throughput under laboratory constraints. At the population level, fitness (non-reciprocity) did not improve, and per-generation maxima declined, while the highest-scoring actuators showed a systematic increase in path amplitude; voxel morphologies also became more spatially clustered over time.

A plausible explanation was that the seed-selection heuristic emphasised net displacement rather than trajectory geometry. Seeds were chosen at locations with large optical-flow magnitude between the peak swollen and most-collapsed frame; this quantity approximated a single end-to-end displacement vector, not the intervening path. Fitness, however, was computed from path non-reciprocity. This constituted a measurement–selection misalignment: the criterion used to decide where to sample promoted a quantity that is inversely related to the quantity used to judge success. As a result, selection pressure was steered toward large net displacements (amplitude), which narrowed loop geometry and depressed the non-reciprocity score.

The observed rise in amplitude and voxel spatial clustering, alongside declining per-generation maxima, was consistent with this mechanism; environmental variability may have modulated magnitudes but did not explain the coherent direction of these trends.

B. Limitations

This work was a proof-of-concept exploration rather than a systematic optimisation study. Throughput was prioritised, which reduced time for slide cleaning. Increased debris was sometimes classified as foreground by Otsu thresholding; when a swelling actuator contacted debris and contours merged, the contour-area-based proxy for phase estimate spiked. A manual option to select a clean reference actuator for phase estimation was added, but residual sensitivity to debris remained.

Environmental control was incomplete early on. Immersion volume varied prior to Generation 05, altering heating dynamics and limiting actuation amplitude; Generation 05 was refabricated and $10 \mu\text{L}$ was fixed thereafter. Evaporation and pre-bias timing were not instrumented, so unmeasured drift may have persisted despite procedural standardisation.

The video pipeline changed mid-study to stabilise phase labelling: introducing speckle filtering and smoothed area traces before peak finding, cycle-aware thresholds, and a ~ 0.5 -cycle debounce was applied. These updates improved consistency but created non-stationarity between early and late generations.

Actuator fitness was defined as the mean of five seed scores to guard against rare tracking failures. Given the robustness observed in tracking, this averaging likely diluted genuinely strong non-reciprocal behaviour when expressed locally by a single seed. A field-level metric derived directly from optical flow would avoid this aggregation effect.

Finally, the evaluation budget was limited relative to the search space. Sixteen generations of 25 individuals were too few for converging a 200-bit morphology code. Results were specific to PNIPAm-co-AAc, the chosen 2PP-DLW recipe, and global thermal actuation; generalisation to other materials, stimuli, or fabrication settings was not assessed.

C. Implications for future work

Non-reciprocal deformation of anchored PNIPAm actuators was demonstrated and then used as the optimisation target. While the non-reciprocity did not improve, it was achieved through geometry-exclusive tuning with one-step 2PP-DLW. The use of positive and negative space distribution for non-reciprocal actuation, even if not evolved, has implications for hydrogel-based Purcell swimmers.

The future iterations of this pipeline should make deliberate efforts to ensure that the morphology selection process and fitness function are aligned. The evaluation of fitness directly from an optical flow motion field should be explored as a more holistic metric. Seed selection could also be explored as an actionable lever for selection. Environmental parameters such as pre-bias timing should be fixed or logged, and a simple temperature readout should be added to contextualise actuation magnitude. If again faced with a limited evaluation budget, the search space should be reduced to fewer voxels to have a greater chance of converging. With these adjustments, this proof-of-concept pipeline can be adapted to outcome-oriented optimisation over longer runs. Furthermore, this evolutionary pipeline is potentially transferable to other micro-fabrication systems that accept geometry as input and provide time-series imaging as output, reducing reliance on designer intuition and enabling less biased exploration of design spaces.

REFERENCES

- [1] J. H. Holland, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. The MIT Press, Apr. 1992. [Online]. Available: <https://direct.mit.edu/books/book/2574/Adaptation-in-Natural-and-Artificial-Systems>
- [2] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. The MIT Press, Nov. 2000. [Online]. Available: <https://direct.mit.edu/books/book/4241/Evolutionary-RoboticsThe-Biology-Intelligence-and>
- [3] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson, “Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding,” 2013.
- [4] H. Lipson and J. B. Pollack, “Automatic design and manufacture of robotic lifeforms,” 2000.
- [5] S. Koos, J.-B. Mouret, and S. Doncieux, “The Transferability Approach: Crossing the Reality Gap in Evolutionary Robotics,” *IEEE Transactions on Evolutionary Computation*, vol. 17, no. 1, pp. 122–145, Feb. 2013. [Online]. Available: <http://ieeexplore.ieee.org/document/6151107/>
- [6] Y. Wang, X. Han, J. Liu, M. Wang, J. Wang, Y. Zhou, and J. Ren, “Closed-loop optimization of tetris-inspired robotic morphology evolution,” *Materials & Design*, vol. 257, p. 114533, Sept. 2025. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0264127525009530>
- [7] M. Uppington, P. Gobbo, S. Hauert, and H. Hauser, “Evolving and generalising morphologies for locomoting micro-scale robotic agents,” *Journal of Micro and Bio Robotics*, vol. 18, no. 1-2, pp. 37–47, Dec. 2022. [Online]. Available: <https://link.springer.com/10.1007/s12213-023-00155-8>
- [8] S. Kriegman, D. Blackiston, M. Levin, and J. Bongard, “A scalable pipeline for designing reconfigurable organisms,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 4, pp. 1853–1859, Jan. 2020. [Online]. Available: <https://pnas.org/doi/full/10.1073/pnas.1910837117>
- [9] M. Hippler, E. Blasco, J. Qu, M. Tanaka, C. Barner-Kowollik, M. Wegener, and M. Bastmeyer, “Controlling the shape of 3D microstructures by temperature and light,” *Nature Communications*, vol. 10, no. 1, p. 232, Jan. 2019. [Online]. Available: <https://www.nature.com/articles/s41467-018-08175-w>
- [10] S. Van Kesteren, X. Shen, M. Aldeghi, and L. Isa, “Printing on Particles: Combining Two-Photon Nanolithography and Capillary Assembly to Fabricate Multimaterial Microstructures,” *Advanced Materials*, vol. 35, no. 11, p. 2207101, Mar. 2023. [Online]. Available: <https://onlinelibrary.wiley.com/doi/10.1002/adma.202207101>
- [11] X. Shen, S.-H. Jung, N. Stauffer, S. Van Kesteren, and L. Isa, “Programmable Microscale Actuation of Hydrogels with Tunable Response via Printing on Particles,” July 2025. [Online]. Available: <https://chemrxiv.org/engage/chemrxiv/article-details/68792041728bf9025eff7c2a>
- [12] X. Zhou, Y. Hou, and J. Lin, “A review on the processing accuracy of two-photon polymerization,” *AIP Advances*, vol. 5, no. 3, p. 030701, Mar. 2015. [Online]. Available: <https://pubs.aip.org/adv/article/5/3/030701/279889/A-review-on-the-processing-accuracy-of-two-photon>
- [13] E. M. Purcell, “Life at low Reynolds number,” in *AIP Conference Proceedings*, vol. 28. AIP, 1976, pp. 49–64, iSSN: 0094243X. [Online]. Available: <https://pubs.aip.org/aip/acp/article/28/1/49-64/718477>
- [14] N. Otsu, “A Threshold Selection Method from Gray-Level Histograms,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, Jan. 1979. [Online]. Available: <http://ieeexplore.ieee.org/document/4310076/>
- [15] A. Savitzky and M. J. E. Golay, “Smoothing and Differentiation of Data by Simplified Least Squares Procedures,” *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, July 1964. [Online]. Available: <https://pubs.acs.org/doi/abs/10.1021/ac60214a047>
- [16] J. Canny, “A Computational Approach to Edge Detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986. [Online]. Available: <https://ieeexplore.ieee.org/document/4767851>
- [17] Jianbo Shi and Tomasi, “Good features to track,” in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition CVPR-94*. Seattle, WA, USA: IEEE Comput. Soc. Press, 1994, pp. 593–600. [Online]. Available: <http://ieeexplore.ieee.org/document/323794/>
- [18] G. Farneback, “Two-Frame Motion Estimation Based on Polynomial Expansion,” in *Image Analysis*, G. Goos, J. Hartmanis, J. Van Leeuwen, J. Bigun, and T. Gustavsson, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, vol. 2749, pp. 363–370, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/3-540-45103-X_50
- [19] B. Lucas, “An Iterative Image Registration Technique with an Application to Stereo Vision,” 1981.