npj Robotics

https://doi.org/10.1038/s44182-026-00092-1

Article in Press

High-throughput soft robot design via an adaptive experimental platform

Received: 28 November 2025

Accepted: 5 May 2026

Cite this article as: Oliveira, R., Pinskier, J., Wang, X. *et al.* High-throughput soft robot design via an adaptive experimental platform. *npj Robot* (2026). https://doi.org/10.1038/s44182-026-00092-1

Rafael Oliveira, Josh Pinskier, Xing Wang, Lois Liow, Sarah Baldwin, James Brett, Vinoth Viswanathan, Richard Scalzo & David Howard

High-throughput soft robot design via an adaptive experimental platform

Rafael Oliveira[1], Josh Pinskier[1], Xing Wang[1,2], Lois Liow[1], Sarah Baldwin[1,3], James Brett[1], Vinoth Viswanathan[1], Richard Scalzo[1], and David Howard[1]*

[1]The Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia
[2]University of Canberra, Australia
[3]Advanced Robotic Manufacturing (ARM) HUB, Australia
*Corresponding Author: david.howard@csiro.au.

Abstract

Soft devices critically require design approaches that can fully realise their embodied potential. Soft robots are an especially problematic soft system as behaviour is complex and data is difficult to obtain, leading to narrow exploration of potential embodiments and inaccurate behavioural assessment. We create the first high-throughput design approach for soft robotics to address these challenges. Scaleable and automated, our approach adaptively combines simulated and experimental assessment to efficiently explore a design space of soft grippers in an automated closed loop. In a prototype study using this high-throughput regime, we demonstrate discovery of higher-performing grippers than comparative methods, and automatically identify and close the simulation-to-reality gap, as well as recording an order of magnitude more experimental grasps than comparative approaches in the literature. Our experimental regime is significantly differentiated from the current literature, offering a realistic route to turn soft robotics into an increasingly data-rich domain and opening up previously unattainable opportunities for the design of soft systems. Data for this paper is publicly available through CSIRO's Data Access Portal: https://data.csiro.au/collection/csiro:65672.

Introduction

Soft matter systems, including polymers, gels, and granular materials, display a dizzying variety of behaviours, which can be leveraged to create functional systems with a range of diverse, high-impact applications [1]. One specific success story of soft systems is in soft robotics [2], which

continues to push the boundaries of robotic ability and represents the state of the art in a range of high value domains where resilience, adaptation, human safety, and environmental robustness are key [3].

The critical question for soft robotics is 'how do I design a soft robot for my application?' – in other words, how do I effectively harness the positive embodied potential [4] of soft robotics to solve a given problem? Despite a growing record of success, the prevailing notion is that the impact of soft robotics remains limited by a lack of tools that discover and design suitable embodiments. Due to the huge range of material and geometric possibilities enabled by soft materials, coupled with a wide variety of embodied behavioural modalities resulting from self- and environmental- interactions, searching a given design space for high-performing solutions remains *the* key unsolved challenge for the field [5]. Methods to search this space and discover devices with 'positive embodied characteristics' would more strongly position soft systems to deliver impact across a range of valuable domains.

Human-centric engineering approaches [6] do not permit the necessary scaling to meaningfully sample these expansive design spaces. Generative approaches [7–12] offer more expansive design space exploration and are therefore more scaleable in principle, however they typically rely heavily on simulation and modelling to achieve that scalability. Due to the field's reliance on hard-to-model environmental interactions, tight body-brain control loops, and material deformation [13, 14], generative design are approaches particularly susceptible to the reality gap. Experimentation offers physical veracity, but even contemporary approaches lack the sheer throughput required to scale. The field critically requires design approaches that efficiently combine the speed of simulation with the veracity of experimentation, allowing search to be both expansive and representative of real world performance. We present a solution to this challenge that removes the critical data bottleneck in design space search and simultaneously allows us to quantify and close the reality gap.

In seeking inspiration to answer this challenge, it makes sense to look to comparable transformations in other fields. In particular, fields such as chemistry, materials, bioengineering, and drug design, have, until recently, been held back by similar issues - lack of data, costly manual experimentation, and inaccurate modelling. For these fields, a type of experimentation known as *high throughput*, meaning that data is produced in such quantity that algorithmic interventions are required to analyse it effectively, catalysed the transformation of these research labs into automated factories of discovery [15], whose experiments are orchestrated by robot scientists [16]. Rapid assessment and large-scale data capture of samples occurs through robotics, e.g., an arm or Cartesian manipulator. The resulting data seeds a predictive model of the design space, where the software component (a Design of Experiments algorithm [17]) decides on subsequent experiments for the robot to carry out based on the hypothesis, past history of experiments, and previously generated data.

High-throughput is a scale-based solution uses AI-based experimental control to turn data-poor domains into data-rich domains. Soft robotics is currently a data poor domain crying out for more fundamentally scalable design methodologies. Hence we investigate a proof of concept high throughput system for soft robotics (see Figure 1) that is in principle applicable across a range of soft matter systems. We call this systems Fast Accurate Soft Testing (FAST). Our approach offers a simultaneous solution to several longstanding issues in the field:

1. FAST finds more performant grippers than the two main design approaches in the field (pure simulation, pure experimentation), using a combined simulation-experimentation optimisation,

2. FAST provides easy access to accurate and repeatable large-scale data for soft robotics applications,

3. FAST automatically characterises and closes the reality gap between simulation and lab-

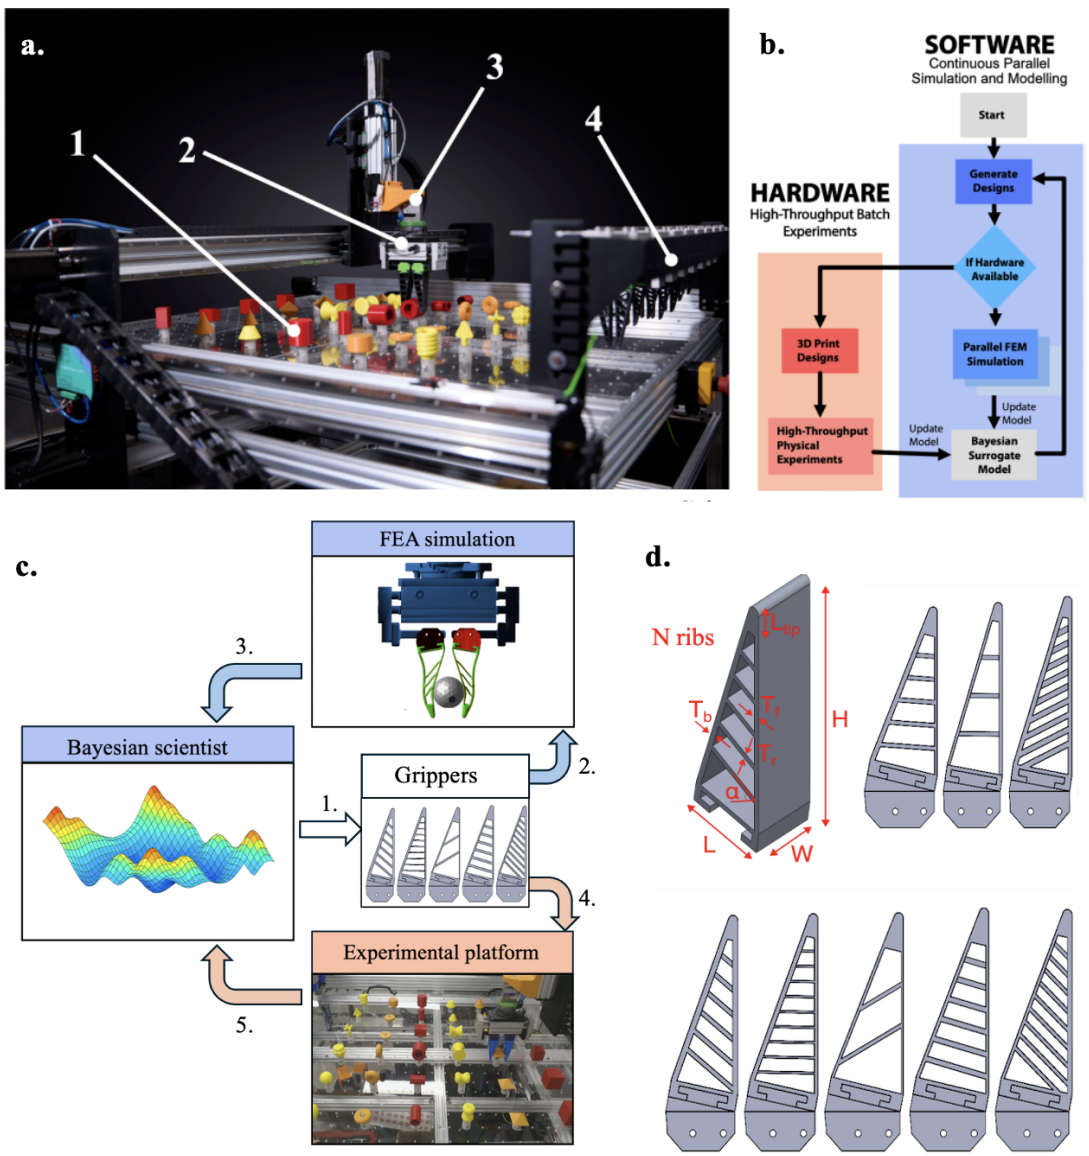

Figure 1: **The FAST platform.** **a** Platform diagram including (1) object tray containing 36 objects, (2) tool changing head with attached gripper, (3) load cell, (4) rack 10 tool changing heads.**b** Delineation of hardware and software components, **c** process diagram showing the interacting parts of the system, where (1) the Bayesian scientist generates batches of grippers based on its surrogate model. Grippers are either assessed in simulation (2), providing one fidelity of feedback (3), or 3D printed tested (4), providing higher fidelity data. The Bayesian scientist updates its model before proposing a subsequent batch of grippers.**d** Sample fin rays.

oratory experimentation,

FAST represents a new type of experimental regime for the field; a scaleable hybrid (*in silico* and *in materia*) approach that is significantly differentiated from the current state of the art. The hardware component provides automated, rapid, repeated robotic assessment of gripper performance [18], providing significant volumes of experimental data over 1 order of magnitude higher than the current state of the art. Our secondary evaluation method is a finite-element model of the gripper whose reality gap is automatically identified and closed and provides fast and increasingly accurate evaluation. Finally, we use a novel multi-fidelity Bayesian optimiser to integrate these two evaluation methods, creating a data-driven surrogate that decides which grippers to test next. We follow a design of experiments methodology [19], where the autonomous system balances and combines experimental evaluations with simulation-based evaluations to rapidly model a complex and high-dimensional design space of soft grippers. A key differentiator compared to the state of the art is that both simulation and experimental data are actively used in the optimisation loop, which has not been done before in soft robotics.

FAST is the first instantiation of the 'high throughput' regime applied to soft robotic design. Automated experimentation is seen across chemical and biological sciences [16], and has applied to the automated assembly and evaluation of (rigid modular) robots with limited design freedom [20–22]. Our approach to automated testing is much higher throughput, and optimises a design domain of soft embodied devices that is more complex, freeform, and challenging. Bayesian approaches [23, 24] are sporadically seen in computational robotic design, however used only as an optimiser and are typically simulation only. Design of Experiments (DoE) (e.g., [25]) is seen across a range of application areas. Experimental data is combined with domain knowledge and prior information via probability distributions which regularize learning and design optimization [19, 26–28]. Combining adaptive experimental design algorithms with robotic platforms then enables the development of automated experimentation systems, which

have been shown in different areas of science and engineering [29]. Application of these techniques to design robots has previously been highlighted as a promising avenue of research [18], but never realised. Perhaps most similar to FAST is BEAR [30], a state of the art platform for Bayesian adaptive experimental design of soft mechanical structures, which was further extended to incorporate simulation data [31] which was used to train a model offline. FAST has several advantages over this approach: we learn a multi-fidelity Gaussian process (GP) model jointly observing simulations and real data as different fidelities [32,33] online during the experimental loop. Moreover, at each iteration, our algorithm proposes a batch of ten experimental designs, instead of single designs per trial as BEAR, which is key to enabling high-throughput experimentation. Our approach is also shown to handle a design space over twice as large as in [31].

We run a series of experiments to investigate the prototype high throughput system, totalling 50608 simulated grasps and 8532 experimental grasps across all experiments. Our approach improves grasp strength of the discovered grippers compared to conventional approaches, which holds even with when using a very coarse uncalibrated simulation. Importantly, the size of the datasets we are able to collect lets us identify important grasp and object properties from only a single data stream (force). Combined, these results highlight the potential of high-throughput as an experimental paradigm for soft robotic design, and the design of soft systems more generally. The potential to dial up embodied devices with complex behavioural properties 'to order' can potentially cataylse the development of soft robotics as a field. Techniques to support the transition from a 'small data' to a 'big data' field opens further research in the application of previously-unsuitable machine learning techniques to further accelerate discovery and enhance the maturity and interconnectedness of the field, in a similar way to the change evidenced when materials and biomedical sciences adopted high-throughput techniques themselves.

Results

System Description

FAST combines 3D printing, multi-object batch evaluation of robotic grippers, Finite Element Modelling (FEM) simulation with gap closure, Bayesian experimental control and adaptation, and a software infrastructure which allows for closed-loop integration of experimental data and simulations (Figure 1.a-b., also see Supplementary Video).

FAST is controlled by a Bayesian 'scientist', that maintains a multi-fidelity Gaussian surrogate model of the design space and proposes new candidate designs for assessment. Assessment is via either a (faster, cheaper) simulator or experimentally using the Cartesian robot. A multi-fidelity Bayesian framework [19] integrates the two data sources, [10] by controlling the execution of simulated and experimental evaluations [34]. The hybrid adaptive experimental design method is summarised in Figure 1.c.

The physical component of FAST comprises a Cartesian gantry robot and testing bed, which has 36 test objects mounted (Figure 1.a). Test objects are fixed during an experiment, but can be easily interchanged depending on experimental requirements – in this case a range of different geometries embody a variety of diverse gripping challenges; see Figure 2.c-d. More information on object selection can be found in Materials and Methods. One end of the testing bed has a rack containing ten tool changing heads where different soft grippers can be mounted. The Cartesian robot can select one head at a time, which mechanically locks to the robot and provides electrical and pneumatic pass-through to the mounted gripper. The geometry of each gripper is encoded in 9 parameters (Figure 1.d) and decided by the Bayesian scientist. The robotic head is equipped with a load cell which allows us to characterise grasp strength (N). Data collection is automatic, rapid and repeatable: each grip is evaluated multiple times (in this case 3), allowing us to create large, reliable data sets. Data is recorded in real time to an

attached server; the Bayesian scientist is automatically informed of the new set of data upon experimental run completion and uses it to decide on the next set of grippers to investigate. The Supplementary Video shows an example physical test conducted on the experimental platform.

Simulation is via FEM in SOFA [35], which allows us to simulate the response of the gripper and efficiently find approximate solutions for deformation that occurs upon contact. SOFA is especially suited to the large-scale deformations inherent to the modelling of soft robots [36]. Full implementation of the FEM can be found in Materials and Methods, and in Supplementary Figure 5.

During an experiment, FAST uses a hybrid simulation and robotic experimentation, as well as updating a data-driven surrogate model after each design iteration, which the Bayesian scientist uses for experimental planning. The algorithm initially generates a batch of designs for experimental evaluation. Designs are 3D printed and mounted on the tool changers, with experimentation then proceeding automatically. Since the physical experimentation process can take relatively long periods of time (practically, each batch takes $\approx$ 1 day - although most of this is 3D printing and post-processing), the algorithm continues optimising designs in simulation and updating its internal surrogate model with lower-fidelity simulated data while the physical experiment is running. Simulation is asynchronous and handled by a scheduler. Augmenting the surrogate with simulated data allows for a more comprehensive exploration of the design space than solely relying on real experiment data, which is crucial in our setting due to the scarcity of real data and the high dimensionality of the design domain. Once data from the real experiment is available, the higher-fidelity data points are added to the model, and a new batch of designs for printing and experimentation is generated with simulation continuing to update the model in between experimental runs. The multi-fidelity optimiser automatically balances the relative impact of data from the two sources, i.e., placing greater weight on experimental data. The process then repeats, starting a new loop of simulations while the next batch of real

data is collected. A full description of the Bayesian system can be found in *Methods*.

Problem definition

We optimise gripping performance of a popular bio-inspired class of soft robotic grippers known as Fin Rays. Compared to previous approaches, we open up a larger design space, allowing for a range of designs as shown in Figure 1.**d**. The 9 gripper parameters are: length L (19-35 mm), width W (28-35 mm), height H (90-120 mm), tip length L_{tip}(20-30 mm), front surface thickness T_f, back surface thickness T_b, rib thickness T_r (2-3 mm), plus rib number N (1-10) and rib angle α (-40 to 40 degrees). The Bayesian scientist guarantees that new grippers respect these ranges. The parameters are selected to give a broad search space for the Bayesian scientist whilst remaining physically meaningful, that is manufacturable and durable. Hence, the minimum thickness of all materials was set to 2 mm to prevent buckling, tearing and excessive stress in the thin features. The minimum height of 90 mm ensured that there was at least 20 mm clearance between the object and the palm during grasping to ensure compliant and durable grasps, whilst the upper limit of 120 mm prevented needlessly long designs. A relatively large minimum width was used to minimise undesirable torsion and out of plane bending on the finger, whilst a relatively low upper limit was selected to restrict excessive material usage. The length range was similarly selected to ensure that the designs were guaranteed to be compatible with FAST's mounting, whilst restricting material usage. Finally, the number and angle of ribs was set to a wide range to permit broad search over the relatively unknown and unrestricted domains.

Experimentally, for each object grasp, pull-off force is measured by closing the grippers on the object horizontally and then raising the gripper while recording force and position data. Peak pull-off force is calculated for each object and averaged across the 36 objects to provide a measure of performance.

In simulation, we evaluate the maximum contact force exerted on an object as a simply-measured proxy for the peak grasping force. FAST allows us to use such an indirect proxy by automatically accounting for discrepancies between simulations and real data via a multi-fidelity Gaussian process model – See *Methods* for more detail.

Experimental regime

FAST's experimental regime is unique in the field of soft robotic design as it incorporates both experimental and simulation-based evaluation with a data-driven surrogate that drives experimental planning within the optimisation loop. The state of the art in the field [11, 37, 38] predominantly uses pure simulation, with experimentation only for verification post optimisation, or pre-optimisation to perform e.g., material model fitting. For pure hardware approaches, we see 100 (10 grippers 10 objects 1 repeat) and 300 (75 grippers 4 objects 5 repeats) experimental evaluations [39, 40] respectively, with 500 evaluations for a modular approach with prefabricated blocks (and hence limited morphological diversity) [41]. In contrast, we run 5400 grasps in the optimisation loop on printed grippers, which is an order of magnitude higher. FAST allows us to perform a rapid, accurate, and highly diverse (grippers $\times$ objects) exploration of soft material design spaces in hardware, which is critical as soft systems are often unintuitive and not amenable to simulation.

As the experimentation runs inside the optimisation loop, it allows us to progressively ground our optimisation process through the iterative construction of the surrogate, even without explicit simulator tuning. The high throughput regime provides a mixture of real world data and scaleable simulated experimentation, which results indicate provides an enhanced optimisation compared to pure simulation or pure experimentation regimes seen in the literature.

Design Optimisation

The primary goal of our hybrid experimental platform is to identify high-performing soft gripper designs within complex optimisation landscapes.

FAST finds better candidates compared to pure simulation and pure experimental approaches, which are the two most common approaches in soft robotics (Figure 2.**a**). FAST outperforms these challenging baselines with an initially-uncalibrated simulator. Despite producing demonstrably inaccurate results, FAST correlates the low-fidelity simulation data with high-fidelity experimental data, and learns to compensate for simulation errors and make inferences about true parameter values, therefore crossing the reality gap. FAST uses a novel bi-fidelity Max-value Entropy Search [42, 43] to choose design points which maximise information gain about the maximum of the high-fidelity objective function (real-data average pull-off force) and incorporates the low-fidelity (but high-volume) data to effectively explore the high-dimensional optimisation domain. FAST therefore provides an efficient multi-fidelity Bayesian approach that is robust to potential discrepancies between simulations and real observation, effectively harnessing experimental and simulation data.

Using a random initial sample of 30 gripper designs (totalling 3132 grasps), 75% were sub-optimal and gave an average grasping force below $20\,\mathrm{N}$, with a mean of $16.30\,\mathrm{N}$. Using only simulations, 150 designs were evaluated. The 10 highest-performing simulated designs were then 3D printed and experimentally validated. We see an improvement on the mean ($21.53\,\mathrm{N}$), but high variance ($\sigma = 4.18\,\mathrm{N}$) due to the simulation inaccuracies and model misspecification. In simulation, these same designs were much more tightly grouped (although biased from their true value), with a mean peak contact force of $15.71\,\mathrm{N}$ and standard deviation of $1.14\,\mathrm{N}$.

Using only real data yields both a substantially higher mean performance and lower variance ($22.99 \pm 2.12\,\mathrm{N}$) than simulation, but also a lower peak performance. This results from its higher-quality, but lower volume data. Exploiting the best of simulation and real data, FAST

Parameter	
Height	$96.96\,\mathrm{mm}$
Width	$33.01\,\mathrm{mm}$
Length	$30.83\,\mathrm{mm}$
Tip length	$24.78\,\mathrm{mm}$
Back thickness	$2.89\,\mathrm{mm}$
Front thickness	$2.61\,\mathrm{mm}$
Ribs thickness	$2.81\,\mathrm{mm}$
Number of ribs	5
Ribs angle	3.63^o

Table 1: Optimised Parameters of best performing candidate

outperforms all baselines, which represent the two common methods of optimising grippers, with the highest mean gripper performance ($26.61 \pm 2.03\,\mathrm{N}$) and finding the best overall design. The best overall design was B07T8 (batch 7, tool 8), which had an average pull-off force of $29.68\,\mathrm{N}$. The optimised design parameters can be found in Table 1, and further details about the experimental procedure can be found in Materials and Methods.

Figure 2.**c**) shows the stability of the optimised gripper designs compared against their randomly generated initial counterparts across the different test objects. The plots clearly show that grasp stability has significantly improved with optimisation. For objects with a flat surface where we observed generally stable grasps, such as cubes and prisms (objects 30 to 35), the plots show large improvements. For the cones (objects 24 and 25), we still observe significant increases in overall grasp stability, yet to a relatively lower magnitude. Some of the most noticeable improvements are on objects presenting a large contact surface during grasping, such as the tube (object 14).

Reality Gap Closure

Although the uncalibrated model provided sufficient information for our FAST to outperform benchmark methods, further gains can be achieved through simulator calibration. Luckily,

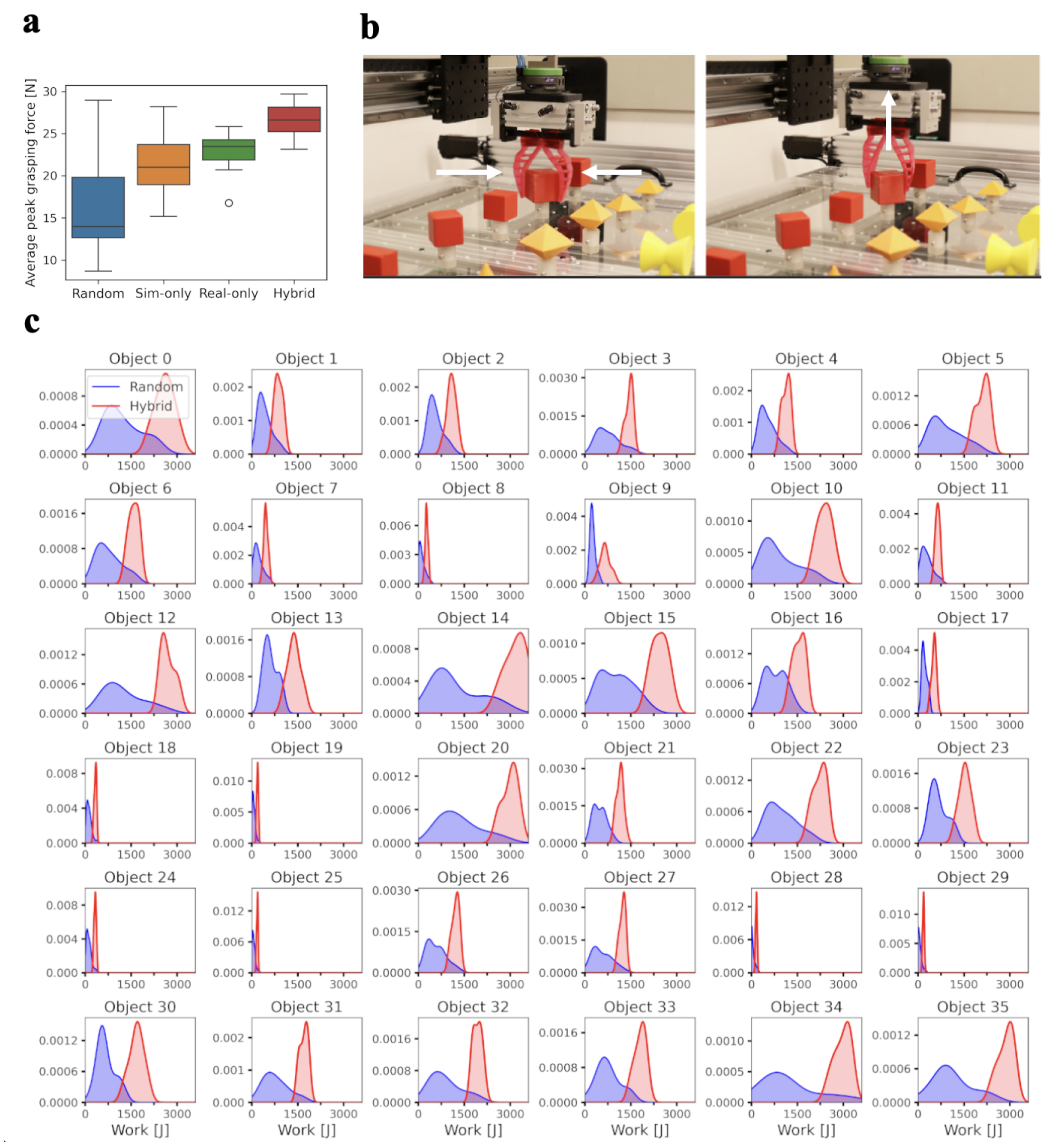

Figure 2: **Experimental analysis:** **a** Demonstrating the ability of FAST to create better-optimised grippers than comparative approaches including full-simulation and full-experimental, which represent the two major approaches in the field. **b** Representative closure and pull-up during a grasp, and **c** grasp stability across all objects: the distribution of the area under the force-by-displacement curves across gripper designs in a batch, i.e., the amount of work (in Joules) done to break a grasp for the initial random designs (blue) and the designs in the last batch (red). Higher values indicate more stable grasps.

FAST generates large amounts of data. Learning physical properties from large datasets is a key feature of FAST, which allows us to improve simulation quality and reduce the reality gap from simulation to experimental data.

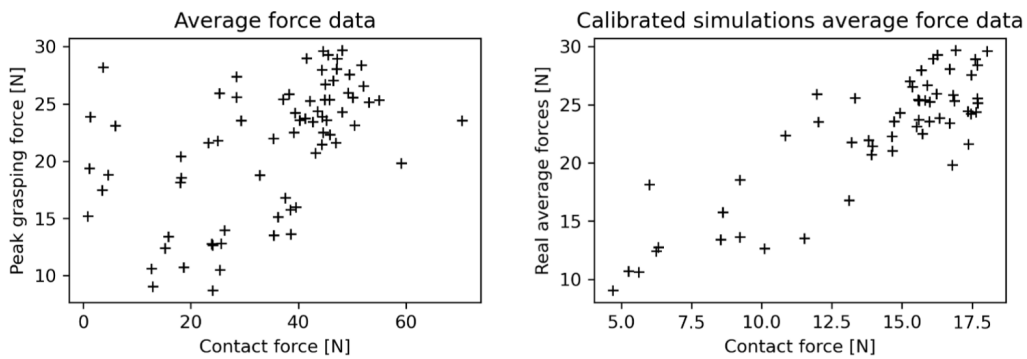

(a) Before calibration (b) After calibration

Figure 3: **Discrepancies between simulations and real data for the peak pull-off force averaged across objects before (a) and after (b) calibration.** Each data point corresponds to a gripper design evaluated in physical experiments. On the x-axis, we have the simulated pull-off force while on the y-axis we have the real peak grasping force. Ideally, the data points should follow a line due to the expected correlation between the simulated force and the real force. The figure contrasts the variability in this relationship before and after the calibration process.

Figure 3.**a** shows discrepancies between real and simulated grasping forces. The simulated contact forces show weak correlation with the real measured grasping forces, suggesting significant discrepancies between the two – this is an instantiation of the reality gap, a critical issue in the field soft robotics.

We selected 3 physics parameters that directly drive simulation behaviour: Young's modulus (E), coefficient of friction (CF) and Poisson ratio ν. We fit them to the experimental data with Bayesian calibration. Fitting the parameters to the data allows us identify the simulation parameters which minimise error in the grasp simulation. Often these fitted parameters do not perfectly match their physical values, as they can be adjusted to compensate for model misspecification. The large datasets generated by FAST enable both the calibration of physical

	Mean	Std. dev.	5%	Median	95%	Mode
Young's modulus (MPa)	9.86	5.04	3.70	8.81	19.02	6.76
Friction coefficient	0.57	0.13	0.37	0.56	0.79	0.50
Poisson ratio	0.33	0.05	0.25	0.33	0.42	0.35

Table 2: **Posterior descriptive statistics**, showing the posterior mean, standard deviation, 5 % quantile, median, 95% quantile and mode estimates, respectively, for each parameter.

parameters and the identification of model misspecification, as specific errors can be isolated from the data.

Using Latin hypercube sampling, we uniformly sample across the 9 soft gripper design parameters and 3 calibration parameters ($E \in [2, 200\,\mathrm{MPa}]$, $CF \in [0.02, 2]$, $\nu \in [0.2, 0.5]$) to generate 1469 simulation designs. A Gaussian process with Gaussian likelihood model correlated the peak simulated grasping force for each object with experimental data [44], with prior probability distributions uniquely chosen for each parameter according to their expected behaviour (see Methods for a detailed discussion).

The resulting joint posterior distribution over the calibration parameters is plotted in Figure 4. It shows a clear anti-correlation between the friction coefficient and the Young's modulus, while the Poisson ratio estimates remain independent. This is an expected result of the simulation configuration; increasing the material's stiffness increases the normal force on the object and hence increase grasp force. However, this can be balanced by reducing the coefficient of friction which relates normal force to grasp force. The resultant mean estimates for Young's modulus, friction and Poisson are $9.86\,\mathrm{MPa}$, 0.57, and 0.33, with full descriptive statistics presented in Table 2. In comparison, the experimentally derived values were $11.5\,\mathrm{MPa}$ for Young's modulus and 0.42 for coefficient of friction. Both parameter estimates fall within the support of the posterior distribution in Figure 4, validating the ability of our system to learn physical parameters from large datasets. In our case, inference is over simulation parameters and undertaken through exact Bayesian inference, however, given the volume of data provided, the

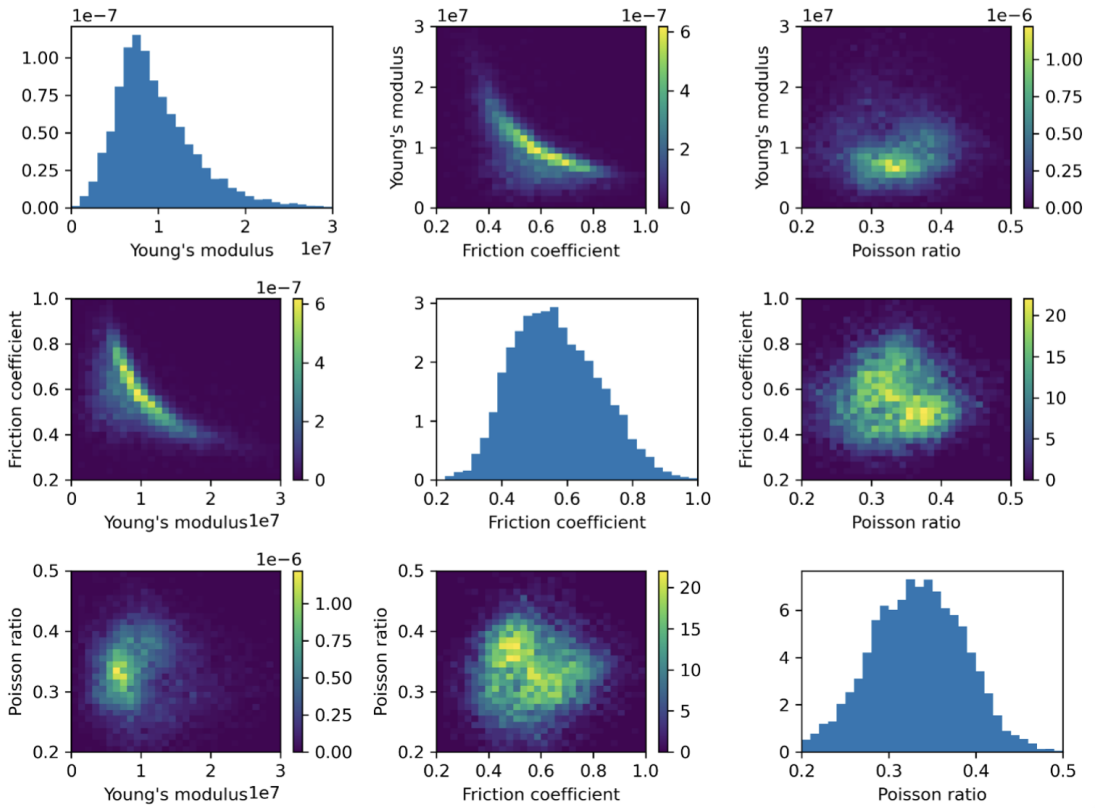

Figure 4: **Joint posterior probability distribution over the calibration parameters** and their marginals given the real and the simulated data after the Bayesian calibration procedure. The pair plots reveal how parameter estimates correlate. While the friction coefficient and the Poisson ratio estimates are mostly uncorrelated, a clear non-linear anti-correlation structure appears between the Young's modulus and the friction coefficient.

approach could be extended to other machine learning methods including flow and diffusion methods.

We additionally inferred the variance of the observation noise and other hyper-parameters of the Gaussian process model used within our Bayesian calibration framework, which provide insights into the underlying physics of grasping with soft grippers. Plots for these additional calibration results are presented in Supplementary Note 1. The results show low measurement noise levels with the most probable value for the standard deviation estimated at $0.06\,\mathrm{N}$. Regarding the GP hyper-parameters, posterior estimates for the covariance function lengthscales

revealed the design variables most affecting the grasping force. Simulated forces were most sensitive to the grippers length and the settings for the calibration parameters, while the discrepancy between real and simulated forces was more influenced by the gripper thickness. The latter may be due to rib buckling (i.e., lateral instability under compression), which could lead to numerical instabilities, especially for a low number of ribs. Being able to identify specific causes of the reality gap makes us more able to solve those gap issues.

Grasp and Object Identification

FAST's other key feature is the ability to generate large quantities of data, which allow us to directly identify and learn salient features of the experiment, objects, and grasps. The Bayesian optimisation framework estimates the sensitivity of designs and design parameters to each object, enabling features to be extracted directly from the experimental data obtained during the optimisation and calibration process, without the need for any targeted experiments.

In the context of soft grasping, three features are of interest: Tool Quality, Grasp Quality, and Object identification. Tool quality being the aggregate grasp performance across objects and the primary optimisation objective, where grasp quality is the stability of a particular tool-object-pose grasp combination. Object identification is the ability to discern object features directly from the grasp data, and creates insights about grasp pose or tool design.

The design optimisation process begins with a set of random designs; as a broad set of tool designs with diverse behaviour, which converge over time (Figure 5.**a-b**).

FAST's ability to systematically generate and analyse data creates the potential for targetted improvements across different dimensions. Grasp quality can be approximated by the integral of force with respect to time during the grasp experiment (i.e., the area under the grasp force curve). Whilst each object-tool pairing creates a distinct grasp force profile (Figure 5.**b**, Batch 1), trends emerge in aggregate which enable stable shapes and poses to be identified. We illus-

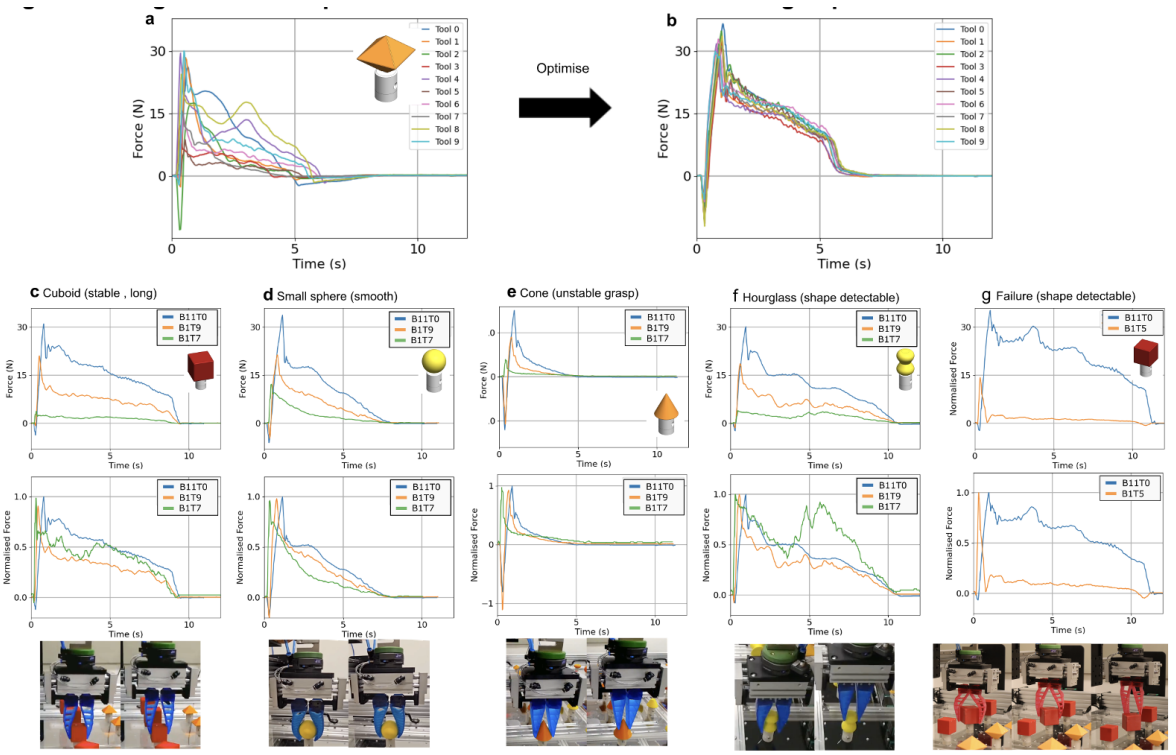

Figure 5: **Force plots and grasp feature extraction: a** The initial (random) and **b** final (optimised) batches of tools, showing object-based specialisation of the designs. **c-g** Grasp quality and object analysis: The volume of data collected by FAST enables reliable identification of tool, object and grasp, using only a single continuous data stream (force). Shown by comparing the softest and stiffest tools in the initial batch (Batch 1, Tool 7 and Batch 1, Tool 9, respectively), with a representative optimised tool (Batch 11, Tool 0), shown as different coloured lines. Objects with a flat **c** or smoothly curved **d** surface result in a stable grasp across tool designs. **e** Inwards tapering objects present unstable grasps. **f** Objects with pronounced geometries have those geometries reconstructed in the force data when using soft tools. **g** Grasp anomalies present as instabilities within otherwise stably grasped objects, indicating poor designs needing further investigation, with an illustrated example of torsional slip.

trate these by comparing the softest tool in the dataset (Batch 1, Tool 7), the stiffest tool of the first, random batch (Batch 1, Tool 9) and a representative tool from the optimised final batch (Batch 11, Tool 0). Flat-sided shapes result in stable grasps, as the objects have a large, smooth surface area to hold (Figure 5.**c**). Convex objects also give stable grasps (e.g sphere, diamond), but have a stronger preference for grasps which fully enclose the object and a steady decline

in grasp force outside that regime (Figure 5.**d**). Where flat objects see a relatively firm grasp held until contact with the object is lost, with holding force declining with contact area, curved objects exhibit a steady decline in holding force as the tool's initial conformance to the object is gradually lost.

In contrast, objects that taper inwards along the vertical axis create very unstable grasps with a narrow force peak and zero or negative forces outside that point (Figure 5.**e**). The inherently challenging geometry presents as a narrow peak in grasp force. Tools initially generate grasp force through compression and adhesion, which is rapidly lost when the tool's movement induces slip.

Whilst the optimization process successfully generated stronger grasps (Figure 5.**c-e**, top row), the grasp force curve are relatively consistent across the sampled tools (Figure 5.**c-e**, bottom row).

Distinctions emerge between tools and objects which can be seen as a deviation from the standard patterns. Notably, softer tools exhibit a greater ability to conform to an object's surface and recover shapes. In complex geometries, this presents as a force curve which tracks the object's contour, as seen in the twin-peaked curve of the hourglass object when using the softest tool (Batch 1, Tool 7) in Figure 5.**f**. In the Hourglass object, concavities show as a twin-peaked force curve in soft tools as force declines when the tools lose contact with the lower surface then recover as they wrap around the upper surface.

More generally, this allows them to grasp objects without an initial negative force (pushing the object into the table). The vast majority of evaluated tools apply a negative force to hold the object before it is held with positive force. However, soft tools can apply positive forces immediately, reducing the chance of object ejection on pickup. Finally, anomalous grasps stand out in the data, for example the slippage caused in Figure 5.**g**. Here the low torsional stiffness of Batch 1-5 causes the grasp to slip off the edge of the rotated cube, which presents as a narrow

peak in the force curve where others have a stable, long duration grasp.

Resource usage

FASTs theoretical maximum throughput is 240 grasps per hour in the current configuration, discounting a negligible amount of head transit and tool changing time. One grasp takes ≈15 seconds, starting with the head above the object, performing a grasp, and finishing with the head above the object. We perform 3 repeats per gripper/object combination, which takes ≈45 seconds. FAST performed over 8000 experimental grasps in the study, with 5400 in the high-throughput regime and the remainder gathering data for the baselines. FAST provides significantly better designs than solely relying on the same amount of real data while taking the exact same amount of time to run all the experiments, as simulations are run in parallel to the physical experiments — this highlights a key feature of our platform that can use uncalibrated simulation data usefully to drive the design process as discovery is hybridised with periodic experimental data.

Discussion

The development of novel design workflows is essential for the continued growth and uptake of soft robotics. We describe a first-of-its-kind prototype platform for conducting high-throughput experimentation for the design of soft-robotic grippers, and offers a route towards turning soft robotics from a small-data to a large-data research field. In contrast to existing approaches, our adaptive hybrid experimentation framework combines real and simulated data, and a data-driven model to inform ongoing experimental planning, which is carried out by a robotic platform in a largely automated manner. FAST provides efficient mapping of the design domain, offering a scalable route to creating large data sets which can be used throughout the optimisation process. We show the that bi-fidelity optimisation via FAST provides a clear performance advantage

when compared to contemporary optimisation approaches. The generated data can be used post-experiment for other purposes, such as the calibration of the simulation environment, as we demonstrated, or creating emulators and surrogates, or deploying traditionally data-hungry machine learning techniques.

In this work, we optimise soft grippers for reliably and safely picking complex objects. Using our adaptive hybrid experimentation framework, we found a gripper configuration which increased grasp performance by almost 2x compared to random baselines. Compared to commercial designs, it uses thicker material sections (ribs and faces) and a larger base area (length and width) to increase the overall stiffness of the design, allowing larger forces to be exerted on the object. However it uses fewer ribs than most standard designs, with the ribs rotated from the horizontal. The resulting space between ribs facilitates localised conformance of the gripper to object and enhances overall grasp strength.

The reality gap is a longstanding issue in soft robotics, and again FAST provides a differentiated capability: Bayesian calibration estimates reveal insights about the correlations across the physics parameters and the effect of different design parameters on the simulation error. Our multi-fidelity Bayesian approach is robust to discrepancies between simulations and real observations (Figure 3), allowing it to benefit from even uncalibrated simulation data as the Gaussian process model captures bias and scaling effects in the simulations relative to the real data [44].

However, we note a number of limitations with the approach which must be addressed before this potential can be fully realised. Firstly, the platform requires human intervention, predominantly in the post-processing required given the polyjet printing technique which allowed easy access to multi-material designs. Fused Deposition Modelling (FDM) printing is attractive from an automation standpoint (see, e.g., BEAR [30]), however imposes limits on the design domain in terms of both geometry and materials. Here we evaluate single-material planar fingers, but in the future we envisage a broader search over free-form 3D designs with multiple

materials and functional grading (as in [45]), which require polyjet printing. Although FDM printers would have been able to produce the designs tested here, polyjet was preferred for its throughput. An entire batch of 10 fingers could be printed overnight and processed together. This is both substantially faster than FDM printing, and minimises the effort in setting up and processing prints, as multiple FDM prints would be needed per batch. Secondly, the platform itself is limited by its design as to the range and type of sensors that can be supported. We decided to focus on the scalable data collection over provision of multi-modal data as the key novelty, but it should be noted that this is trade-off requires consideration. This is particularly important as multi-modal sensing is known to be beneficial in machine-learning scenarios and thus may be preferred depending on the type of downstream tasks targeted. Cost-reducing re-designs focused are also required to make platforms such as FAST more readily available to the research community.

Generalising this approach to other soft robotics problems is feasible, but not trivial and will require the community to develop reusable and reconfigurable platforms which can adapt to the requirements of different problems including sensor and actuator design and characterisation, mechanism and metamaterial design, grasping and manipulation, and entire mobile soft robots. In most cases, Cartesian gantry platforms or systems based on robot arms offer sufficient workspace and speed for evaluation. However, substantial engineering effort is required to integrate the required sensing modalities, design suitable experiments, and tie these into the Bayesian scientist. FAST offers a first-time demonstration of such an integrated system, which was the result of over 2 years of development in our lab. As more research moves towards data-centric and high-throughput development, we expect to see many more platforms emerging and ultimately a convergence towards widely available standardised designs.

Overall, findings show the potential benefit of adopting high-throughput technologies in the soft robotics domain. Platforms like FAST can support the transition of soft robotics into

a more data-available domain, which is critical if we wish to unlock and explore advanced, complex, and multi-faceted design domains. Currently we can produce with more flexibility than we can design, and we envision new types of design algorithm that are tuned to exploit the heightened amount of available data, as well as more comprehensive studies into related areas including definition and quantification of concepts such as embodied intelligence and morphological computing [46].

Finally, we note the power of platforms like FAST in both (i) more fully harnessing ongoing advances in machine learning through access to sufficiently large data sets, and (ii) addressing significant issues in the field associated with benchmarking and reproducibility, which becomes increasingly relevant as the field of soft robotics transitions to a more mature stage [47]. Soft matter systems generally could benefit from high throughput systems like FAST, due to the increased ratio of experimental to simulated evaluations and ability to harness imperfect simulations whilst still efficiently performing discovery in these complex design spaces.

Methods

The objective of this study is to characterise the operation of our high-throughput platform for soft robot design, FAST. We run through a representative experiment, highlighting key differences with respect to pure simulation and pure experimental approaches, which are standard in the field. We focus on performance of the discovered grippers, efficiency of the optimisation process, and other phenomenological differences brought on by the use of a hybrid evaluation regime. The key components of the study are the Cartesian testing robot, FEM-based simulation, data-driven surrogate, gripper 3D printing, and batch Bayesian experimental design. Post experiment, Bayesian calibration is performed to identify and close the reality gap.

Additive manufacturing of soft-robotic grippers

The fin-ray grippers are fabricated using a Stratasys J850 Prime 3D printer. Each multi-material fin-ray finger consists of two parts: (1) the compliant finger and its (2) rigid mounting base, which are 3D printed in one go without need for assembly, as shown in Figure 1.**d**. The fin-ray fingers are printed as composite of PolyJet materials Agilus30 and VeroVivid to achieve a composite of Shore A 85 hardness. The finger's base is printed in VeroVivid which is Shore D.

Each iteration, a batch of 10 grippers is 3D printed together on the same tray using High Mix Mode with SUP706 soluble support material. A glossy finish is selected for all the grippers. After printing, supports are removed using a waterjet machine and are left to dry completely before being mounted onto the gantry tool changers. This is the only manual part of the experimental procedure.

Experimental platform

The platform is a modified 3-axis Cartesian machine with a work space of 1500x1500x180 mm. The gantry bed is a sheet of perspex with 50mm hole spacing to accommodate the affixing of test objects at 150 mm spacing. The machine features a tool changer with ten Wingman interchangeable heads, each with their own double acting pneumatic gripper and pneumatic pass through. The tool changer head is mounted to a a 20 kg Zemic S-type load cell on the z Axis to measure the grippers' pull-off force. A Click Programmable Logic Controller supplies expandable GPIO and pneumatic valve control, and a Laumas weight transmitter is used to interface with the load cell. Gantry control is via an Arduino Mega board running GRBL, and a custom ROS package communicates with the control board, Click PLC, and Laumas weight transmitter.

Test objects

Many robotic grasping performance protocols and benchmarks have been previously proposed, such as the YCB [48] and ARCV [49] benchmarks, which comprise various household items. Typically these objects can be large, making mounting on the gantry platform infeasible. Hence for the gantry grasping experiments, 36 primitive objects are selected to evaluate the pull-off force performance of the generated fin-ray designs. The total number was selected to maximise the number of objects tested per finger by filling the FAST table. The object set comprises cubes, cones, diamonds, pyramids, diabolos, hourglasses, tubes, cylinders, toruses, and spheres, of varying sizes and orientations. Hence, spanning a set of key geometric features including single and double curvature, smooth and sharp objects. Irregular shapes such as a spiral, triangle, jack, and ellipsoid are also included. These 36 primitive objects are selected to (i) provide a range of grasping challenges, and (ii) induce different contact regions between the gripper and objects, creating different grasp types such as power grasps, pinch grasps, and fingertip grasps [50, 51]. An enveloping grasp is often seen when grasping objects with large surface area of contact, such as cubes and spheres. Pinch grasping is often performed with inclined objects such as the cone and pyramid. Fingertip grasping is performed with grasping flat objects, such as the disks and torus (when placed in a horizontal position). Furthermore, the use of primitive objects allows ease of information decoding, allowing more direct inferences to be drawn between an object's features (size, shape, and orientation) and the collected force data. As the objects are 3D printed from models, the model can be easily incorporated into the simulator.

The objects are 3D printed in rigid Acrylonitrile Styrene Acrylate (ASA) material on a Bambu Fused Deposition Modelling (FDM) printer, and mounted 33 mm above the gantry's perspex sheet. Each of the 36 objects fits within a 50x50x50 mm^3 volume. Objects are mounted in a 6x6 array in the centre of the gantry platform, as shown in Figure 1.**a**. A total of 8 batches of pull-off grasping experiments is performed. For each batch, the grasping performance of 10

grippers is evaluated by measuring the pull-off forces for all 36 primitive objects. Each grasp is repeated 3 times per object. The pneumatic grippers are actuated with 30 psi positive air pressure to enable grasp closure, followed by a vertical displacement at 5 mm/s until the gripper separates from the object, during which force is measured. When complete, the pneumatic valve allows atmospheric pressure into the cylinder, returning the gripper to its fully-open state.

Design Optimisation and Baselines

The hybrid optimisation algorithm is initialised with data from 30 randomly generated designs, which are evaluated on the real hardware, totalling 30 grippers and 300 average peak grasping force measurements (900 individual measurements), and simulations of 50 randomly generated designs. After that, the algorithm successively generated batches of 10 real designs, with 100 simulated designs between each physical experiment. This gives enough data to characterise a representative operating regime of the high-throughput system.

The number of experimental rounds is comparable to [52], which ran 3 experimental rounds for each of the two main strategies they followed. Their real experiments, though, were solely guided by the real data, with simulation-based experiments being run separately beforehand. In contrast, our simulations are run online and interleaved with experimental gripper evaluations. The number of simulation batches was set to 10, each with 10 tool designs, which were run in-between the real-data experiments. The batch size of 10 was set to mimic the experimental batches (FAST has 10 tool changers), and the number of 10 rounds of batch simulation accounted for the availability of our high-performance computing facilities in the time available in-between real experiments.

We compare our approach to two single-fidelity baselines representing the two common approaches to soft robot design [19]: purely simulated (with final experimental evaluation), and purely experimental optimisation. The first approach exclusively runs BO in simulation, guided

by the expected improvement criterion [53]. We then evaluate top-10 best performing simulated gripper designs on the real platform. The second approach solely relies on real data and runs batch BO with the single-fidelity max-value entropy search [42] on batches of 10 designs per iteration. This approach was initialized with the same 30 randomly generated designs that our hybrid algorithm used, generating batches of 10 grippers which were evaluated on the gantry. Lastly, we also evaluated randomly generated batches as a trivial baseline. Note that these are both challenging baselines as Bayesian simulation-based optimisation is a current state of the art approach [11, 23].

This mix of simulation and reality, and inclusion of simulation and experimentation directly in the optimisation loop, is representative of the experimental regime that we wish to investigate and allows us to assess the feasibility and impact of high-throughput as a design approach for soft robotics.

Soft gripping simulation

The mechanical behaviour of soft robots is governed by the principles of continuum mechanics, for which analytical solutions are typically intractable. Consequently, numerical techniques such as nonlinear FEM are a well-suited tool for obtaining converging approximate solutions to their deformation. The utility of this approach is demonstrated by FEMs capacity to model large-scale deformations in both elastic and hyperelastic materials, to permit the rapid validation of design iterations, and to simulate the complex, non-linear contact interactions between the soft robot and its operational environment.

We used FEM provided by the SOFA framework to compute two reaction forces that characterise a soft grasping process. This process involves an initial grasping phase, where the two-fingered gripper is actuated to pinch grasp and hold an object at equilibrium, followed by a vertical pull-off test that concludes when the object slips from the grasp. The two forces com-

puted from these respective phases were the horizontal contact force (global x-axis) and the vertical pull-off force. Both forces were formulated as potential fitness values for subsequent optimisation algorithms.

For the initial design optimisation, the contact force was chosen over the pull-off force as the fitness metric to ensure robust convergence. The pull-off simulation was deemed less suitable for initial experiments because it involves significant changes in boundary conditions and contact points, leading to high nonlinearity that compromises solution stability. Further details regarding the simulation setup can be found in Supplementary Note 2.

A linear elastic material model is used for the soft grippers, with the Young's modulus found using the ASTM D412 tensile testing standard on an Instron 34SC-5 universal testing machine on 3D printed samples. Taking the average of 5 results and fitting a linear stress-strain curve, the Stratasys Polyjet Shore 85A material was characterised by Young's modulus of $11.5\,\mathrm{MPa}$ ($\sigma = 1.75\,\mathrm{MPa}$). As it is assumed to be nearly incompressible, the Poisson ratio is estimated as 0.45. The coefficient of friction was similarly ascertained, with the static and dynamic coefficients of friction between the Agilus 85A and ASA 3D printed materials calculated as 0.41 and 0.99, respectively. Data was collected using the Instron Friction Testing station. Further information can be found in the Supplementary Material, specifically Supplementary Table 2.

Batch Bayesian optimisation

Our experimental design goal is to maximise an objective function $f : \mathcal{X} \to \mathbb{R}$ mapping gripper designs $\mathbf{x}$ in a given design space $\mathcal{X}$ to an experiment outcome $f(\mathbf{x}) \in \mathbb{R}$, which in our case corresponds to the average peak grasping force across the set of test objects. Due to limited sensing precision in physical systems, we assume we do not observe $f(\mathbf{x})$ directly, but instead observe a noisy outcome, i.e., $y = f(\mathbf{x}) + \nu$, where $\nu \sim \mathcal{N}(0, \sigma_\nu^2)$ is assumed to be i.i.d. Gaussian noise.

We run up to T iterations of a batch optimisation algorithm, with R real experiments and $S > R$ simulations. Each iteration of the real experiment evaluates a batch of B different designs using a novel batch Bayesian optimisation algorithm [26] in a multi-fidelity setting [32]. The algorithm is designed to minimise the need for physical experiments and make efficient use of time by continuing the optimisation process via simulation in between real experiment trials. This process is illustrated in Figure 1.**b**. Simulated and experimental data are combined via a Gaussian Process (GP) regression model [54], which is designed to properly account for the discrepancy between real and simulated data.

Bayesian optimisation relies on a Bayesian probabilistic model to infer the value of the objective function at a candidate design point. To perform Bayesian inference about the value of real experiment outcomes with the aid of simulation data in our hybrid setting, we apply a bi-fidelity Gaussian process model inspired by Bayesian calibration [44]. Letting $h : \mathcal{X} \to \mathbb{R}$ denote a simulator mapping design points $\mathbf{x} \in \mathcal{X}$ to a simulated outcome $\hat{y} \in \mathbb{R}$, we assume the real objective function $f : \mathcal{X} \to \mathbb{R}$ can be modelled as:

$$f(\mathbf{x}) = \rho h(\mathbf{x}) + \varepsilon(\mathbf{x}) \,, \quad \mathbf{x} \in \mathcal{X} \,, \tag{1}$$

where $\varepsilon : \mathcal{X} \to \mathbb{R}$ represents the error (or discrepancy) between simulations and reality, and $\rho \in \mathbb{R}$ denotes a scaling factor.

We jointly model real and simulated outcomes by introducing an auxiliary fidelity variable $s \in \{0, 1\}$. The highest fidelity $s = 1$ corresponds to the real experiment data, which is more costly to obtain. The lowest fidelity $s = 0$ corresponds to the simulated data, which is more easily obtainable, but with limited accuracy. This formulation is captured by the combined objective:

$$\hat{f}(\mathbf{x}, s) = \begin{cases} \rho h(\mathbf{x}) + \varepsilon(\mathbf{x}), & s = 1 \\ h(\mathbf{x}), & s = 0 \,. \end{cases} \tag{2}$$

In this formulation, if h is modelled as a zero-mean GP with covariance function $k_h : \mathcal{X} \times \mathcal{X} \to$

$\mathbb{R}$, i.e., $h \sim \mathcal{GP}(0, k_h)$, and the error model also follows a GP with $\varepsilon \sim \mathcal{GP}(0, k_\varepsilon)$, it turns out that $\hat{f}$ is also a Gaussian process with covariance function given by:

$$k(\mathbf{x}, s, \mathbf{x}', s') = k_\rho(s, s')k_h(\mathbf{x}, \mathbf{x}') + ss'k_\varepsilon(\mathbf{x}, \mathbf{x}')\,, \tag{3}$$

where $k_\rho(s, s') := (1 + s(\rho - 1))(1 + s'(\rho - 1))$.

Now, given N observations $\mathbf{y}_N \in \mathbb{R}^N$ and the corresponding set of inputs $\{\mathbf{x}_i, s_i\}_{i=1}^N$ containing a mix of real and simulated data, the GP model predicts the value of $\hat{f}$ at any given $\mathbf{x} \in \mathcal{X}$ as a normal distribution:

$$\hat{f}(\mathbf{x}, s) \mid \mathbf{y}_N \sim \mathcal{N}(\mu_N(\mathbf{x}, s), \sigma_N^2(\mathbf{x}, s))\,. \tag{4}$$

The predictive mean and variance are given by:

$$\mu_N(\mathbf{x}, s) = \mathbf{k}_N(\mathbf{x}, s)^\mathsf{T}(\mathbf{K}_N + \mathbf{\Sigma}_\nu)^{-1}\mathbf{y}_N \tag{5}$$

$$\sigma_N^2(\mathbf{x}, s) = k(\mathbf{x}, s, \mathbf{x}, s) - \mathbf{k}_N(\mathbf{x}, s)^\mathsf{T}(\mathbf{K}_N + \mathbf{\Sigma}_\nu)^{-1}\mathbf{k}_N(\mathbf{x}, s)\,, \tag{6}$$

where $\mathbf{k}_N(\mathbf{x}, s) = [k(\mathbf{x}, s, \mathbf{x}_1, s_1), \ldots, k(\mathbf{x}, s, \mathbf{x}_N, s_N)]^\mathsf{T} \in \mathbb{R}^N$, the GP covariance matrix is denoted by $\mathbf{K}_N = [k(\mathbf{x}_i, s_i, \mathbf{x}_j, s_j)]_{i,j=1}^N \in \mathbb{R}^{N\times N}$, and $\mathbf{\Sigma}_\nu$ is the noise covariance matrix. In our case, $\mathbf{\Sigma}_\nu$ is set as a diagonal matrix with entries given by σ_ν^2 for real data observations ($s = 1$) and a small nugget value close to 0 for simulated outcomes ($s = 0$), which are deterministic, except for limitations due to finite numerical precision.

We use Matérn kernels with smoothness parameter set to 2.5 for both the simulation k_h and error k_ε kernels, each applying independent lengthscales to every design variable for automatic relevance determination [54]. Hyper-parameters are learnt and adapted online during the BO process via maximum-a-posteriori estimation with weakly informative priors.

Typical BO algorithms sequentially propose candidates after each data point by maximising an acquisition function a_t, which quantifies the expected utility of running an experiment at a

point $\mathbf{x} \in \mathcal{X}$ given the data available at time t. In multi-fidelity settings, it is common to combine the acquisition function with a cost function that determines the value of evaluations with each fidelity [32], making the algorithm avoid proposing evaluations with the highest fidelity if they are not necessary. However, in our setting, such cost functions are not needed, since the real hardware for the physical experiments cannot be requested arbitrarily due to limited availability, and we run a real experiment whenever the hardware is available. Fidelity evaluations, therefore, follow a fixed cycle in our setting, and the acquisition function solely evaluates designs for the fidelity available at that point in time.

To propose informative optimal designs for our experiments, we make use of the multi-fidelity max-value entropy search acquisition function [43]. This acquisition function estimates the expected information gain about the maximum value of the objective function:

$$f^* := \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \tag{7}$$

if we run an experiment at a given $\mathbf{x} \in \mathcal{X}$ at fidelity $s \in \{0, 1\}$, i.e.:

$$a_t(\mathbf{x}, s) = \mathbb{I}(y; f^* | \mathbf{x}, s, \mathcal{D}_t)\,, \tag{8}$$

where $\mathbb{I}(\dots)$ denotes the mutual information between random variables [55], and $\mathcal{D}_t$ represents the data available at time t. Using this criterion, the acquisition function guides BO to designs which are both informative (i.e., reduce uncertainty in the GP model) and near-optimal. To generate batches, the mutual information in Eq. 8 is extended to cover multiple samples, by conditioning on a batch of points $\{\mathbf{x}_i\}_{i=1}^{B}$. Technical details can be found in the original paper for multi-fidelity max-value entropy search [43].

Our particular implementation was adapted from the BoTorch package [56] and the Ax adaptive experimentation platform (`https://ax.dev/`). We used an asynchronous scheduler-based candidate generation scheme for simulations where each candidate design was generated based on the available data as soon as a new observation becomes available without the need

to wait for all the simulations in a batch to finish. Hence, we can continuously run FEM simulations in parallel while we wait for real data. This scheme, first proposed by Kandasamy et al. [57], allows for more efficient use of simulation resources, as some simulations might take longer to finish due to the complexity of the proposed designs.

Bayesian calibration

Bayesian calibration is used to close the reality gap in the FEA simulator. To calibrate simulations against an available set of real data $\mathcal{D}_R := \{\mathbf{x}_i, y_i\}_{i=1}^{R}$, we follow the approach proposed by Kennedy and O'Hagan [44] by varying the simulation physics parameters $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^d$ to estimate the settings which best approximate reality. A randomised set of simulations $\mathcal{D}_S := \{\hat{\mathbf{x}}_i, \hat{\boldsymbol{\theta}}_i, \hat{y}\}_{i=1}^{S}$ is sampled to cover a representative subspace of the design space $\mathcal{X}$ and the calibration parameters space Θ. As a probabilistic model, we extend our bi-fidelity GP model to account for the different simulation parameter settings $\boldsymbol{\theta} \in \Theta$ as:

$$\hat{f}(\mathbf{x}, \boldsymbol{\theta}, s) := \begin{cases} \rho h(\mathbf{x}, \boldsymbol{\theta}) + \varepsilon(\mathbf{x}), & s = 1 \\ h(\mathbf{x}, \boldsymbol{\theta}), & s = 0\,, \end{cases} \tag{9}$$

where $\rho \in \mathbb{R}$ is an optional parameter accounting for differences in scale between simulations and real data, as described above. The corresponding covariance function is given below:

$$k((\mathbf{x}, \boldsymbol{\theta}, s), (\mathbf{x}', \boldsymbol{\theta}', s')) := k_\rho(s, s') k_h((\mathbf{x}, \boldsymbol{\theta}), (\mathbf{x}', \boldsymbol{\theta}')) + s s' k_\varepsilon(\mathbf{x}, \mathbf{x}') \tag{10}$$

which extends the model used for optimisation by including simulation calibration parameters as inputs to k_h.

Assuming that there is a particular unknown set of parameters $\boldsymbol{\theta}^* \in \Theta$ with which the simulations best approximate the real data, we place a prior distribution $p(\boldsymbol{\theta}^*)$ over such parameters and perform Bayesian inference to estimate $\boldsymbol{\theta}^*$ given the available data $\mathcal{D} := \mathcal{D}_R \cup \mathcal{D}_S$. The goal of Bayesian inference is to estimate a probability distribution, called the Bayesian poste-

rior, which provides parameter estimates and allows us to quantify uncertainty about them:

$$p(\boldsymbol{\theta}^*|\mathcal{D}_N) = \frac{p(\mathcal{D}_N|\boldsymbol{\theta}^*)p(\boldsymbol{\theta}^*)}{p(\mathcal{D}_N)}\,, \tag{11}$$

where $p(\mathcal{D}_N|\boldsymbol{\theta}^*)$ is the likelihood, and $p(\mathcal{D}_N) = \int_{\boldsymbol{\theta}\in\Theta} p(\mathcal{D}_N|\boldsymbol{\theta}^*)p(\boldsymbol{\theta}^*)\,\mathrm{d}\boldsymbol{\theta}^*$ is the evidence, which is usually intractable. The likelihood model is given by the GP equations as:

$$p(\mathcal{D}_N|\boldsymbol{\theta}^*) = \mathcal{N}(\mathbf{y}_N; \mathbf{0}, \mathbf{K}_N(\boldsymbol{\theta}^*) + \boldsymbol{\Sigma}_\nu)\,, \tag{12}$$

where $\mathbf{K}_N(\boldsymbol{\theta}^*)$ represents the GP covariance matrix using $\boldsymbol{\theta}^*$ as the calibration parameters for the real data given by $\boldsymbol{\theta}^*$, i.e., the covariance matrix is computed over inputs $\{\mathbf{x}_i, \boldsymbol{\theta}_i, s_i\}_{i=1}^N$, and $\boldsymbol{\theta}_i = \boldsymbol{\theta}^*$ whenever $s = 1$. Additional information is provided in the Supplementary Material: Supplementary Figure 1 estimates the amount of measurement noise in the real data, and Supplementary Figure 2 shows the posterior distribution of the sim-to-real error amplitude. Supplementary Figure 3 presents the sensitivity of the error kernel with respect to different design parameters. Lastly, Supplementary Figure 4 shows the sensitivity of the simulation kernel with respect to both design and calibration parameters.

The Young's modulus is modelled as a log-normal prior distribution with mode at $10.2\,\mathrm{MPa}$ and a variance which concentrated the probability mass within ± 1 order of magnitude of the mode (i.e., between $1\,\mathrm{MPa}$ and $100\,\mathrm{MPa}$). The mode was estimated by converting the material's shore hardness to Young's modulus [58]. The CF was estimated from literature at 0.56, which was then used as the mode of a gamma prior for the coefficient that concentrated most of the mass between 0 and 1. For Poisson ratio, we allowed our priors to be very broad and uninformative, since this was one of the parameters we were most uncertain about. Therefore, we placed a uniform prior on the interval between 0.2 and 0.5. The exact parameters for these priors are presented in Table 3, with sampling ranges in Supplementary Table 1.

Output observations are taken as the average peak grasping force across the test objects for both the real and simulated data. When there is missing data for simulations of a given test

Calibration parameter	Prior
Young's modulus	Log-Normal $(\mu = 16.138, \sigma^2 = 1)$
Friction coefficient	Gamma $(\alpha = 6.6, \beta = 10)$
Poisson ratio	Uniform $(0.2, 0.5)$

Table 3: Prior probability distributions for each calibration parameter.

object due to failures in the simulation process (e.g., when there are errors meshing a particular design or simulations time out), observations are imputed with the average observation value for that object's simulation output force. We also adopt a heteroscedastic noise model for the GP, which allows for the variance of the observation noise to be informed. In the case of missing data, the noise variance of the grasping force on the missing object is estimated as the sample variance for that object and then propagated to the average force observation.

To estimate a posterior distribution over the calibration parameters given the probabilistic model described above, we run Markov chain Monte Carlo (MCMC) with the No-U-Turn sampler [59]. For each posterior, we performed 8000 posterior draws, after an initial burn-in phase of 8000 steps, across 8 parallel independent chains.

Acknowledgments

This work was supported by CSIRO's Future Digital Manufacturing Fund. High Performance Compute resources were provided and supported by CSIRO's IMT Scientific Computing Team. This research is supported by the Science and Industry Endowment Fund SIEF MEP3-02, who funded the purchase of the 3D printer used in this work.

Data and materials availability: Data for this paper is publicly available through CSIRO's Data Access Portal: https://data.csiro.au/collection/csiro:65672. Code is specific to the platform but is available on request for adaptation.

Author Contributions: Conceptualization: R.O, J.P, R.S and D.H; Data acquisition: R.O,

J.P, X.W, L.L, S.B, J.B, V.V; Data Analysis and Interpretation: R.O, X.W, J.P and R.S; Manuscript Drafting and Editing: R.O, J.P, R.S and D.H All authors reviewed the manuscript.

Competing interests: The authors declare no competing financial or non-financial interests.

References

[1] S. R. Nagel, "Experimental soft-matter science," *Reviews of modern physics*, vol. 89, no. 2, p. 025002, 2017.

[2] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.

[3] M. S. Xavier, C. D. Tawk, A. Zolfagharian, J. Pinskier, D. Howard, T. Young, J. Lai, S. M. Harrison, Y. K. Yong, M. Bodaghi *et al.*, "Soft pneumatic actuators: A review of design, fabrication, modeling, sensing, control and applications," *IEEE Access*, vol. 10, pp. 59 442–59 485, 2022.

[4] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.

[5] J. Pinskier, X. Wang, L. Liow, Y. Xie, P. Kumar, M. Langelaar, and D. Howard, "Diversity-Based Topology Optimization of Soft Robotic Grippers," *Advanced Intelligent Systems*, vol. 6, no. 4, 2024.

[6] S. Kim, C. Laschi, and B. Trimmer, "Soft robotics: a bioinspired evolution in robotics," *Trends in biotechnology*, vol. 31, no. 5, pp. 287–294, 2013.

[7] F. Chen and M. Y. Wang, "Design Optimization of Soft Robots: A Review of the State of the Art," *IEEE Robotics and Automation Magazine*, no. December, pp. 27–43, 2020.

[8] L. Smith and R. MacCurdy, "SoRoForge: End-to-End Soft Actuator Design," *IEEE Transactions on Automation Science and Engineering*, vol. 20, no. 3, pp. 1475–1486, 2023.

[9] S. E. Navarro, T. Navez, O. Goury, L. Molina, and C. Duriez, "An Open Source Design Optimization Toolbox Evaluated on a Soft Finger," *IEEE Robotics and Automation Letters*, vol. 8, no. 9, pp. 6044–6051, 2023. [Online]. Available: http://arxiv.org/abs/2304.07260

[10] R. Garnett, *Bayesian Optimization*. Cambridge University Press, 2023. [Online]. Available: https://bayesoptbook.com/

[11] Y. Yao, Y. Chen, L. He, and P. Maiolino, "Design Optimization for Bellow Soft Pneumatic Actuators in Shape-Matching," *2023 IEEE International Conference on Soft Robotics, RoboSoft 2023*, pp. 1–7, 2023.

[12] Y. Xie, J. Pinskier, X. Wang, and D. Howard, "Evolutionary seeding of diverse structural design solutions via topology optimization," *ACM Transactions on Evolutionary Learning*, vol. 4, no. 4, pp. 1–23, 2024.

[13] J. Collins, S. Chand, A. Vanderkop, and D. Howard, "A review of physics simulators for robotic applications," *IEEE Access*, vol. 9, pp. 51 416–51 431, 2021.

[14] J. C. Case, E. L. White, and R. K. Kramer, "Soft material characterization for robotic applications," *Soft Robotics*, vol. 2, no. 2, pp. 80–87, 2015.

[15] A. Angelopoulos, J. F. Cahoon, and R. Alterovitz, "Transforming science labs into automated factories of discovery," *Science Robotics*, vol. 9, no. 95, p. eadm6991, 2024.

[16] R. D. King, K. E. Whelan, F. M. Jones, P. G. Reiser, C. H. Bryant, S. H. Muggleton, D. B. Kell, and S. G. Oliver, "Functional genomic hypothesis generation and experimentation by a robot scientist," *Nature*, vol. 427, no. 6971, pp. 247–252, 2004.

[17] B. Durakovic, "Design of experiments application, concepts, examples: State of the art," *Periodicals of Engineering and Natural Sciences*, vol. 5, no. 3, 2017.

[18] T. Howison, S. Hauser, J. Hughes, and F. Iida, "Reality-assisted evolution of soft robots through large-scale physical experimentation: a review," *Artificial Life*, vol. 26, no. 4, pp. 484–506, 2020.

[19] T. Rainforth, A. Foster, D. R. Ivanova, and F. B. Smith, "Modern Bayesian Experimental Design," *Statistical Science*, vol. 39, no. 1, pp. 100 – 114, 2024. [Online]. Available: https://doi.org/10.1214/23-STS915

[20] M. Hale, E. Buchanan Berumen, A. Winfield, J. Timmis, E. Hart, G. Eiben, W. Li, and A. Tyrrell, "The are robot fabricator: How to (re) produce robots that can evolve in the real world," in *International Society for Artificial Life: ALIFE2019*, 2019, pp. 95–102.

[21] L. Brodbeck, S. Hauser, and F. Iida, "Morphological evolution of physical robots through model-free phenotype development," *PloS one*, vol. 10, no. 6, p. e0128444, 2015.

[22] V. Vujovic, A. Rosendo, L. Brodbeck, and F. Iida, "Evolutionary Developmental Robotics: Improving Morphology and Control of Physical Robots," *Artificial Life*, vol. 23, no. 2, pp. 169–185, 4 2017. [Online]. Available: https://doi.org/10.1162/ARTL_a_00228

[23] X. Wang, B. Wang, J. Pinskier, Y. Xie, J. Brett, R. Scalzo, and D. Howard, "Fin-bayes: A multi-objective bayesian optimization framework for soft robotic fingers," *Soft Robotics*, vol. 11, no. 5, pp. 791–801, 2024.

[24] T. S. Yook Min, L. Y. Lee, and S. G. Nurzaman, "Bayesian optimization of pneumatic soft grippers via reconfigurable modular molds," in *2023 IEEE International Conference on Soft Robotics (RoboSoft 2023)*. IEEE, 2023.

[25] K. Korovina, S. Xu, K. Kandasamy, W. Neiswanger, B. Poczos, J. Schneider, and E. Xing, "Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, S. Chiappa and R. Calandra, Eds., vol. 108. PMLR, 2020, pp. 3393–3403.

[26] S. Greenhill, S. Rana, S. Gupta, P. Vellanki, and S. Venkatesh, "Bayesian Optimization for Adaptive Experimental Design: A Review," *IEEE Access*, vol. 8, pp. 13 937–13 948, 2020.

[27] A. Giovagnoli, "The Bayesian design of adaptive clinical trials," *International Journal of Environmental Research and Public Health*, vol. 18, no. 2, pp. 1–15, 2021.

[28] R. Oliveira, D. Sejdinovic, D. Howard, and E. V. Bonilla, "Bayesian adaptive calibration and optimal design," *Advances in Neural Information Processing Systems*, vol. 37, pp. 56 526–56 551, 2024.

[29] L. M. Roch, F. Häse, C. Kreisbeck, T. Tamayo-Mendoza, L. P. E. Yunker, J. E. Hein, and A. Aspuru-Guzik, "Chemos: Orchestrating autonomous experimentation," *Science Robotics*, vol. 3, no. 19, p. eaat5559, 2018. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.aat5559

[30] A. E. Gongora, B. Xu, W. Perry, C. Okoye, P. Riley, K. G. Reyes, E. F. Morgan, and K. A. Brown, "Bayesian experimental autonomous researcher for mechanical design," *Science Advances*, vol. 6, no. 15, p. eaaz1708, 2020. [Online]. Available: https://www.science.org

[31] A. E. Gongora, K. L. Snapp, E. Whiting, P. Riley, K. G. Reyes, E. F. Morgan, and K. A. Brown, "Using simulation to accelerate autonomous experimentation: A case study using mechanics," *iScience*, vol. 24, no. 4, 2021. [Online]. Available: https://doi.org/10.1016/j.isci.2021.102262

[32] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, and S. Trimpe, "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with Bayesian optimization," in *IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, 2017, pp. 1557–1563.

[33] L. Brevault, M. Balesdent, and A. Hebbal, "Overview of gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems," *Aerospace Science and Technology*, vol. 107, p. 106339, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S127096382031021X

[34] S. Sarkar, S. Mondal, M. Joly, M. E. Lynch, S. D. Bopardikar, R. Acharya, and P. Perdikaris, "Multifidelity and multiscale Bayesian framework for high-dimensional engineering design and calibration," *Journal of Mechanical Design*, vol. 141, no. 12, pp. 1–11, 2019.

[35] F. Faure, C. Duriez, H. Delingette, J. Allard, B. Gilles, S. Marchesseau, H. Talbot, H. Courtecuisse, G. Bousquet, I. Peterlik *et al.*, "Sofa: A multi-model framework for interactive physical simulation," *Soft tissue biomechanical modeling for computer assisted surgery*, pp. 283–321, 2012.

[36] M. S. Xavier, A. J. Fleming, and Y. K. Yong, "Finite element modeling of soft fluidic actuators: Overview and recent developments," *Advanced Intelligent Systems*, vol. 3, no. 2, p. 2000187, 2021.

[37] S. Ghosh, R. Maurya, and S. Roy, "Pagdesa: An automated generative design of multi-modal soft pneumatic actuators," *IEEE Robotics and Automation Letters*, 2025.

[38] J. Pinskier, X. Wang, L. Liow, Y. Xie, P. Kumar, M. Langelaar, and D. Howard, "Diversity-based topology optimization of soft robotic grippers," *Advanced Intelligent Systems*, vol. 6, no. 4, p. 2300505, 2024.

[39] T. S. Y. Min, L. Y. Lee, and S. G. Nurzaman, "Bayesian optimization of pneumatic soft grippers via reconfigurable modular molds," in *2023 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2023, pp. 1–6.

[40] D. Howard, J. O'Connor, J. Letchford, J. Brett, T. Joseph, S. Lin, D. Furby, and G. W. Delaney, "Getting a grip: In materio evolution of membrane morphology for soft robotic jamming grippers," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2022, pp. 531–538.

[41] L. Brodbeck, S. Hauser, and F. Iida, "Morphological evolution of physical robots through model-free phenotype development," *PloS one*, vol. 10, no. 6, p. e0128444, 2015.

[42] Z. Wang and S. Jegelka, "Max-value entropy search for efficient Bayesian optimization," in *34th International Conference on Machine Learning, ICML 2017*, vol. 7, Sydney, Australia, 2017, pp. 5530–5543.

[43] S. Takeno, H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama, "Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization," in *37th International Conference on Machine Learning (ICML 2020)*, ser. Proceedings of Machine Learning Research, vol. 119. Online: PMLR, 2020, pp. 9276–9287.

[44] M. C. Kennedy and A. O'Hagan, "Bayesian calibration of computer models," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 63, no. 3, pp. 425–464, 2001.

[45] J. Pinskier, J. Brett, and D. Howard, "Towards Bespoke Soft Grippers through Voxel-Scale Metamaterial Topology Optimisation," in *Robosoft 2024*. San Diego: IEEE, 2024, pp. 1–8.

[46] G. Mengaldo, F. Renda, S. L. Brunton, M. Bächer, M. Calisti, C. Duriez, G. S. Chirikjian, and C. Laschi, "A concise guide to modelling the physics of embodied intelligence in soft robotics," *Nature Reviews Physics*, pp. 1–16, 2022.

[47] N. Obayashi, D. Howard, K. L. Walker, J. Jørgensen, M. Gepner, D. Sameoto, A. Stokes, F. Iida, and J. Hughes, "A democratized bimodal model of research for soft robotics: Integrating slow and fast science," *Science Robotics*, vol. 10, no. 99, p. eadr2708, 2025.

[48] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.

[49] J. Leitner, A. W. Tow, N. Sünderhauf, J. E. Dean, J. W. Durham, M. Cooper, M. Eich, C. Lehnert, R. Mangels, C. McCool *et al.*, "The acrv picking benchmark: A robotic shelf picking benchmark to foster reproducible research," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4705–4712.

[50] C. B. Teeple, T. N. Koutros, M. A. Graule, and R. J. Wood, "Multi-segment soft robotic fingers enable robust precision grasping," *The International Journal of Robotics Research*, vol. 39, no. 14, pp. 1647–1667, 2020.

[51] M. R. Cutkosky *et al.*, "On grasp choice, grasp models, and the design of hands for manufacturing tasks." *IEEE Transactions on robotics and automation*, vol. 5, no. 3, pp. 269–279, 1989.

[52] A. E. Gongora, B. Xu, W. Perry, C. Okoye, P. Riley, K. G. Reyes, E. F. Morgan, and K. A. Brown, "A bayesian experimental autonomous researcher for mechanical design," *Science advances*, vol. 6, no. 15, p. eaaz1708, 2020.

[53] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2016.

[54] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: The MIT Press, 2006.

[55] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, 2005.

[56] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, "BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization," in *Advances in Neural Information Processing Systems 33*, 2020. [Online]. Available: http://arxiv.org/abs/1910.06403

[57] K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Poczos, "Parallelised Bayesian Optimisation via Thompson Sampling," in *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, Lanzarote, Spain, 2018. [Online]. Available: http://arxiv.org/abs/1705.09236

[58] I. M. Meththananda, S. Parker, M. P. Patel, and M. Braden, "The relationship between Shore hardness of elastomeric dental materials and Young's modulus." *Dental materials :*

official publication of the Academy of Dental Materials, vol. 25, no. 8, pp. 956–959, aug 2009.

[59] M. D. Hoffman and A. Gelman, "The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo," *Journal of Machine Learning Research*, vol. 15, pp. 1593–1623, 2014.

Figure captions:

Figure 1.**The FAST platform. a** Platform diagram including (1) object tray containing 36 objects, (2) tool changing head with attached gripper, (3) load cell, (4) rack 10 tool changing heads.**b** Delineation of hardware and software components, **c** process diagram showing the interacting parts of the system, where (1) the Bayesian scientist generates batches of grippers based on its surrogate model. Grippers are either assessed in simulation (2), providing one fidelity of feedback (3), or 3D printed tested (4), providing higher fidelity data. The Bayesian scientist updates its model before proposing a subsequent batch of grippers.**d** Sample fin rays.

Figure 2. **Experimental analysis: a** Demonstrating the ability of FAST to create better-optimised grippers than comparative approaches including full-simulation and full-experimental, which represent the two major approaches in the field. **b** Representative closure and pull-up during a grasp, and **c** grasp stability across all objects: the distribution of the area under the force-by-displacement curves across gripper designs in a batch, i.e., the amount of work (in Joules) done to break a grasp for the initial random designs (blue) and the designs in the last batch (red). Higher values indicate more stable grasps.

Figure 3. **Discrepancies between simulations and real data for the peak pull-off force averaged across objects before (a) and after (b) calibration.** Each data point corresponds to a

gripper design evaluated in physical experiments. On the x-axis, we have the simulated pull-off force while on the y-axis we have the real peak grasping force. Ideally, the data points should follow a line due to the expected correlation between the simulated force and the real force. The figure contrasts the variability in this relationship before and after the calibration process.

Figure 4.**Joint posterior probability distribution over the calibration parameters** and their marginals given the real and the simulated data after the Bayesian calibration procedure. The pair plots reveal how parameter estimates correlate. While the friction coefficient and the Poisson ratio estimates are mostly uncorrelated, a clear non-linear anti-correlation structure appears between the Young's modulus and the friction coefficient.

Figure 5.**Force plots and grasp feature extraction: a** The initial (random) and **b** final (optimised) batches of tools, showing object-based specialisation of the designs. **c-g** Grasp quality and object analysis: The volume of data collected by FAST enables reliable identification of tool, object and grasp, using only a single continuous data stream (force). Shown by comparing the softest and stiffest tools in the initial batch (Batch 1, Tool 7 and Batch 1, Tool 9, respectively), with a representative optimised tool (Batch 11, Tool 0), shown as different coloured lines. Objects with a flat **c** or smoothly curved **d** surface result in a stable grasp across tool designs. **e** Inwards tapering objects present unstable grasps. **f** Objects with pronounced geometries have those geometries reconstructed in the force data when using soft tools. **g** Grasp anomalies present as instabilities within otherwise stably grasped objects, indicating poor designs needing further investigation, with an illustrated example of torsional slip.