2020 IEEE International Conference on Robotics and Automation (ICRA)
31 May - 31 August, 2020. Paris, France

Sim2real gap is non-monotonic with robot complexity for morphology-in-the-loop flapping wing design

Kent Rosser[1+], Jia Kok[1], Javaan Chahl[2] and Josh Bongard[3]

***Abstract*—**

Morphology of a robot design is important to its ability to achieve a stated goal and therefore applying machine learning approaches that incorporate morphology in the design space can provide scope for significant advantage. Our study is set in a domain known to be reliant on morphology: flapping wing flight. We developed a parameterised morphology design space that draws features from biological exemplars and apply automated design to produce a set of high performance robot morphologies in simulation. By performing sim2real transfer on a selection, for the first time we measured the shape of the reality gap for variations in design complexity. We found for the flapping wing that the reality gap changes non-monotonically with complexity, suggesting that certain morphology details narrow the gap more than others, and that such details could be identified and further optimised in a future end-to-end automated morphology design process.

Keywords - morphology ; simulation to reality ; evolution ; bio-inspired

*This work was supported by the DST Group Trusted Autonomous Systems Strategic Research Initiative and the Vermont Advanced Computing Core

[1]Defence Science and Technology Group, Australia
[2]University of South Australia, Australia
[3]University of Vermont, VT USA
[+]kent.rosser@dst.defence.gov.au

I. INTRODUCTION

In many applications, the design of a robot is intuitive, fast and achievable by human hand. A fixed morphology robot platform design is then produced and characterised as a simulation model, allowing control policies to be developed that achieve a desired task in simulation and then subsequently transferred to reality [14], [27], [18], [25], [12], a process known as sim2real transfer. For some applications however, the intuitive design approach to morphology may not be available as construction materials, actuation and the environmental forces being experienced by the robot may not be within the realm of traditional design experience [28]. Without a guiding intuition, manual design can degrade into a systematic form of trial and error, and applying machine learning approaches to morphology design becomes appropriate.

Evolutionary algorithms are one such machine learning approach capable of producing a population of designs from which to select [32], [31]. Inspired by the ability of natural evolution to produce diverse and complex structures, evolutionary robotics has developed tools that facilitate morphological search to produce robots with novel morphologies

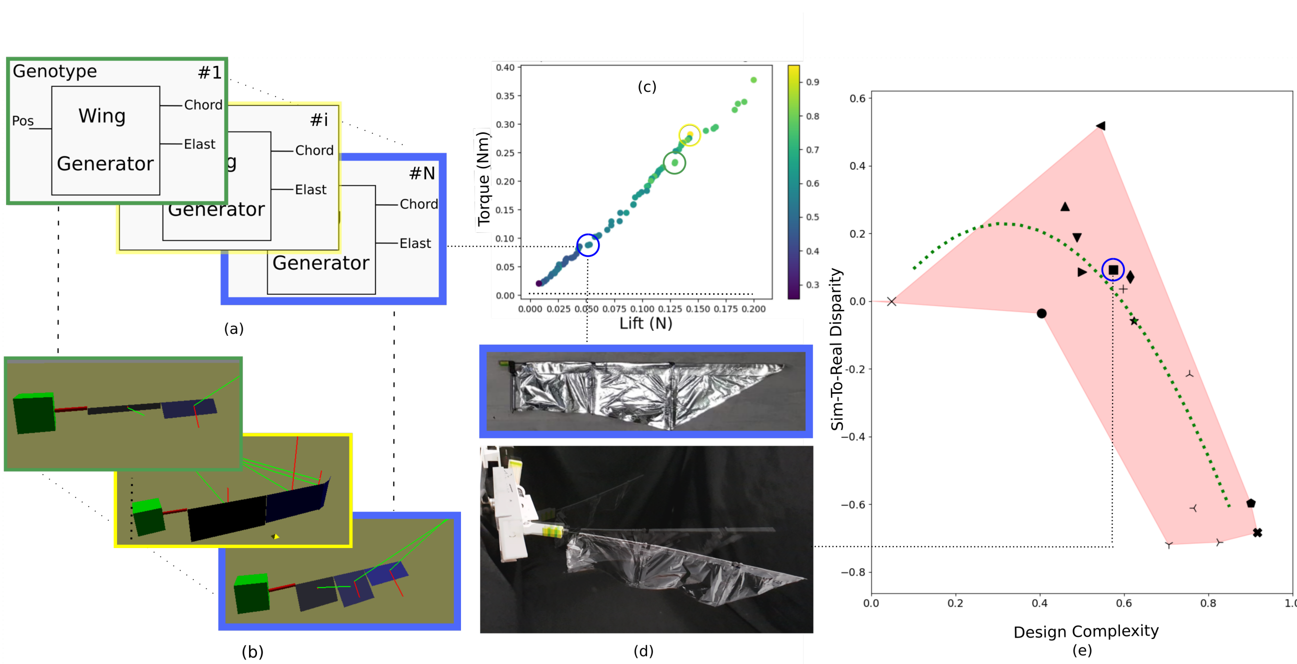

Fig. 1. Automatic design of flapping wing morphologies using simulation produces a set of high performance designs that are transferred to reality to measure the shape of the simulation to reality gap. (a) An evolved population of designs (genotypes); (b) designs are evaluated in simulation; (c) evolutionary algorithm produces a Pareto optimal set of designs; (d) selected individual is transferred to reality and evaluated; (e) plot of the simulation-to-reality gap defined by transferred designs.

[6] that a human designer may not consider [20].

Automated design of these robots in simulation is challenged by the "reality gap" problem [14]. The reality gap is the observed difference between the performance of a robot with high fitness in simulation and its real world equivalent, and remains a critical issue for robotics [11], [14], [22], [18]. The ability to produce simulations that match reality for arbitrary behaviour and morphology exceeds currently available physics simulation packages, and when we consider the effect of material properties, non-homogeneous materials, stiffness, aerodynamics and friction (to name a few) it is easy to imagine a range of ways that robots designed in simulation can produce behaviours that will differ from reality.

There is a growing literature reporting machine learning design methods that automatically achieve successful sim2real transfer without using human intuition to improve the simulation or undertake the design. Jakobi [13] produced early work in sim2real transfer for fixed morphology robots by applying noise to the simulation of the components that would likely transfer poorly. More recently Hwangbo *et al* [12] showed that for a fixed morphology robot the reality gap can be mapped into a deep neural network facilitating sim2real transfer. Lipson *et al* [22] reported the first example of sim2real design with both morphology and control. Their work demonstrated the existence of the reality gap for morphological design, as only a small fraction of their realised robots performed to expectation. Bongard *et al* [3] successfully crossed the reality gap by developing a method for continuous morphology self-modelling while operating on the physical robot and co-evolving a set of controllers to suit in simulation. Koos *et al* [18] produced the transferability method to automatically achieve sim2real transfer for controllers. That method undertakes sim2real transfer of selected designs, quantifies the difference between simulation and real performance and for any poor transfer, it probabilistically excludes the local region of the design space from further search. While all of these examples have made contributions to the automatic design of robots, none of them have provided any machine learning methods that address reliable sim2real transfer when morphology is included in the design loop, nor provided any explanation as to why certain regions of the design space do not transfer well that can be exploited by an automated design process.

To begin to address these gaps when morphology is not fixed (which we call morphology-in-the-loop design), we seek an understanding of the 'shape' of the reality gap. In variable morphology simulation, it is typical to model a design as a set of relevant finite elements such as rigid body components, voxels, splines etc [22], [5], [7] and accumulate the localised effects experienced by each element to produce an overall robot behaviour. We know from FEA (Finite Element Analysis) tools used extensively in human directed computer aided design to deconstruct larger problems into smaller elements, that simulation accuracy is dependant on the number of elements and the overall size of a design [23], both of which vary in a morphology-in-the-loop design process. Using a finite element approach to morphology design adds morphological structure by adding one or more additional finite elements, so it is consistent to expect that simulation accuracy will relate to the total number of finite elements as well as the spacing of those elements across the size of the design. In this study, we use the term morphological simulation complexity to represent that combination of elements and size of a robot design within a specific simulation approach. Qualitatively, when a robot morphology with known simulation error is extended by adding another finite element representing additional structure, the initial simulation error is expected to be retained and the new element will incorporate its own model imperfections as an additional quantity of error, intuitively suggesting that simulation error would grow monotonically with morphological simulation complexity.

We explore this reality gap shape within the bio-inspired field of flapping wing flight. Understanding flapping wing aerodynamics is challenging problem [30] and customisation of morphology can be intuitively expected based on the specific objective to be optimised. While only a few man-made fliers exist in this domain [26], natural fliers have solved the problem of staying airborne through evolution, and their morphology is known to be important and tailored to their environmental niche. Consider some examples of natural fliers. Dragonflies have four long, high aspect ratio wings that provide agility; butterflies have short low aspect ratio wings flapped at a lower rate allowing long flight times; and birds have wing geometries dependant on species that are tailored to speed, manoeuvrability or endurance. Beyond just wing shape, the distribution of mass, elasticity and actuation are also important. For example, Yin *et al* [33] showed that mass of the wing is critical for flapping wing systems, and Li *et al* [21] showed the substantial amount of angular twist (60^o) across the span of a dragonfly wing during maneuver, which can reasonably be expected to be enabling rather than detrimental to its performance. Examples such as these suggest morphological design is crucial for gaining a niche in an environment, and equally that no control scheme can overcome limitations caused by a poor morphology.

Human designers often seek inspiration from biological examples that natural evolution has honed over aeons to achieve spectacular refinement and capability. Rather than direct mimicry of the morphology of biological wings, we approached the challenge of morphological design by identifying salient characteristics of biological examples to use as building blocks in a machine learning approach than can automatically specify new morphological designs.

This paper presents our sim2real morphology-in-the-loop design within our parameterised morphospace (the region spanned by all specifiable morphology design solutions) of a flapping wing aerodynamics problem. We have created an evolvable design specification (genotype) that can be expressed to create a single structured flapping wing robot (phenotype) for analysis in a simulation engine or to form the specification for the manufacture and assessment in reality. Our manufacturing process allows each wing design to be produced using readily available materials in a short time frame and at low cost, allowing this study to include substantially more morphology transfers than previous sim2real studies cited above.

Here, for the first time, we measure the shape of the morphology design reality gap as a function of morphological complexity. For our flapping wing system the shape is non-monotonic, which has implications for machine learning approaches to simulator development and morphology-in-the-loop design, and sets future research questions relating to the generality of the shape and the factors that drive it.

II. Materials and Methods

This section describes the bio-inspired flapping wing morphospace used in this study, From the introduction, morphological shape, size, stiffness and inertia are known to be important for natural fliers so we made these features available for selection within our morphospace including the simulator, real world fabrication and evolutionary design approaches.

A. Simulator

Within flapping wing research, many studies use computational fluid dynamics (CFD) supported by wind tunnel testing (WTT) [16], [26] for the analysis of a single wing design, however those tools are poorly suited to evaluation of large numbers of different designs due to computation time, finite element meshing variation and fluid structural interaction [15], and the difficulty will increase if the design space extends to include multiple wings to match natural flyers. We are driven towards approaches relevant to automated design, and for that reason we hypothesise that CFD may not be required. Instead we propose a lower fidelity but faster simulation process that allows a much wider exploration of the search space. Provided a simulation provides a fitness landscape with a gradient towards high performance solutions in reality, it is suitable for automated design.

For our flapping wing simulation, we modelled a flapping wing design as a set of spanwise aligned connected flat ”blade” elements. Each individual blade includes local geometry (span and chord) values that collectively determine the shape, span and inertia distribution of the entire wing. The blades are sequentially connected by joints that allow elastic deformation in both chordwise (twisting) and spanwise (bending) directions when the wing is under load. Within the simulation, all size and stiffness characteristics were quantised to match those available in the real wing fabrication. Figure 2(a) shows a full wing constructed in the simulator.

Each blade embodied the Sane *et al* quasi-static model for flapping wing aerodynamic forces and torques [29]. That model solves time domain forces as the sum of translational and rotational effects. It excludes wake capture, an interaction of the wing with the disturbance it created in a previous stroke, which has only limited characterisation in the literature.

The simulation was encapsulated in the PYROSIM robotics simulator [19] which was extended to compute

(a) (b) (c) (d)

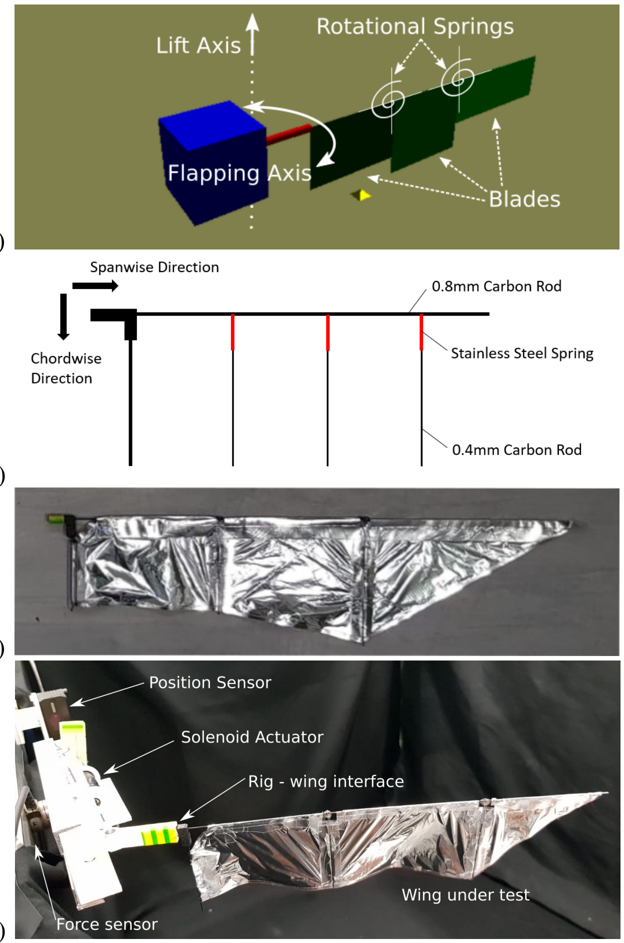

Fig. 2. Robot wing development pipeline (a) simulation phenotype defined by genotype for evaluation (b) construction of elastic wing ribs (c) realised wing after manufacture (d) wing mounted to flapping wing test apparatus.

quasi-static forces for each blade within a simulation, while the underlying physics simulator resolved the inertial and elastic response of the wing. The wing is connected by a wing root (red cylinder) to a rigid base (blue cube). Rotational actuation of our flapping wing simulation was applied at the root of the wing around the z-axis, with angle controlled in accordance with a provided control waveform. Figure 1(b) shows simulation of three differing individuals and visualises the instantaneous quasi-static normal and axial forces (green and red lines) for each blade.

Actuation causes interaction between the environment and the wing and therefore morphology dependent amounts of lift and drag as the wing translates and deforms. The base block in the simulator measures the force and torque experienced during actuation cycles to determine overall lift (vertical force) experienced and drive forces applied.

B. Manufacture and Test

To support sim2real transfer, we developed a cost-effective manufacturing method that produces a wing that allows controlled variation of span, shape, elasticity and inertia.

From the wing root, a carbon fibre rod formed the main spar of the wing and its length set the span of the physical wing. One of the morphological parameters in the morphospace is the elasticity along the wing, which impacted on the amount of aero-elastic deformation that the wing experiences when under load. It is difficult and costly to accurately produce a smooth elasticity profile along the wing, so we discretize the elasticity by attaching ribs to the spar of the wing oriented in the chordwise direction, connected with spring elements as illustrated in Figure 2(c). The spring elements used were stainless steel music wire of varying diameters (0.35-0.45mm) creating selectable levels of elasticity. Spring elements were isolated to a 15mm section of the rib attached to the spar, while the remainder of the rib element was formed by a rigid carbon rod stiffener of dimension such that the rib length is the same as the wing shape geometry at that position on the wing. Using this, we obtain a definable shape, elasticity and inertia profile along the span without adding significant manufacturing complexity. The aerodynamic skin was aluminised Mylar film, which was selected for its robustness. A completed wing approximately weighed 1 gram, cost US$1 in consumables and required 2 hours of assembly time (excluding adhesive cure time).

To assess real world performance, a flapping wing test rig was developed and is depicted with a mounted wing in Figure 2(d). Our test rig drives the wing repeatedly through a single flapping axis using a linear solenoid actuator with magnetic plungers similar to the design demonstrated by Kok *et al*. [17]. Position feedback was provided by a laser displacement sensor (Keyence IL100) that measured the position of the magnetic plunger which defined angular position of the wing. An Arduino Due [1] measured the position feedback and commanded the actuator such that real time angular position of wing stroke could be accurately controlled. The actuator and wing assembly was mounted on an ATI Nano17 Force/Torque transducer such that forces generated by the actuated wing were recorded by a National Instruments data acquisition system [2] at a sampling frequency of 100kHz. From this, for each wing we measured mean and variation of lift over several cycles. The flapping test rig was capable of producing controlled oscillation beyond 10Hz dependant on the specific wing attached over a angular range of $\pm 40^o$. This study set a fixed sinusoidal pattern at 5Hz to facilitate arbitrary morphology-in-the-loop design without risking damage to either the wing or apparatus.

C. Evolutionary design

To conduct automated design using evolution, we defined a genotype structure that could describe designs that span of the morphospace. Figure 1(a) shows an image of a population of individuals that are each able to specify a unique wing morphology by evaluating their own wing generator at a set of evolved positions along the wing span. The result is a wing design that includes variations in shape, elasticity and inertia along its span which matches the bio-inspired design space made available in simulation and manufacture.

We placed the development of a population of these genotypes under evolutionary control in simulation. The algorithm used for this study was the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [10] evolutionary optimisation method. Multi-objective optimisation is commonly used to maintain a diverse set of high performance designs within an evolving population. In this case, we searched for designs that maximised Lift produced while minimising drive power and torque. In optimising across more than one dimension, rather than determining a single best morphology, the search tool produces a set of individuals known as a non-dominated front (NDF). Within a NDF, each individual represents a single design that is superior to all the others in the population on at least one objective. Figure 1(c) shows one example NDF produced using this method.

D. Experiment

For our experiment we were interested in the variation between simulation and reality of high performance solutions as we vary the structure of the design specification in our parameterised design space. To focus on morphological effects, we constrained the control behaviour in both simulation and reality to be a fixed 5Hz sinusoid, while we varied the wing design to produce efficient lift in simulation. We transferred high performing designs to reality and compared the lift as measured from the flapping wing test rig to determine transferability.

In their study, Koos *et al* [18] defined the Simulation To Reality (STR) disparity measure to assess transferability of controllers as the difference between the phenotypic performance measures in simulation and reality. Their definition of STR included an built in bias which matched their observation that only reduced performances were observed following simulation to reality transfer - i.e. the controller always performed worse in reality than simulation. We propose a refined definition to allow for the possibility that reality may outperform simulation in morphological design.

For our flapping wing experiment, the phenotypic performance metric we use is lift. For any robot design m, we calculates STR as per Equation 1, where L_S and L_R are the mean inter-cycle lift per flapping cycle averaged over 5 flapping cycles for simulation and reality respectively. L^{max} is the maximum value of L_S of all wing designs considered.

$$STR(m) = \frac{L_R(m) - L_S(m)}{L^{max}} \tag{1}$$

We are assessing the variation in STR with respect to morphological simulation complexity, C_{MS}. From the introduction, we identified that the number of finite elements and their location across the size of the design were important to simulation accuracy, so we define this measure for a design m in our flapping wing morphospace as per Equation 2.

$$C_{MS}(m) = \frac{1}{2}\Big(\frac{B(m)}{B^{max}} + \frac{S(m)}{S^{max}}\Big) \tag{2}$$

where $B(m)$ is the number of blade finite elements in the simulation and $S(m)$ is the span of the wing. B^{max} and S^{max} are the maximum value of those measures across all designs.

Wing Label	Design	C_{MS}	$B(m)$	$S(m)$ (mm)	L_S (g)	L_R (g)	STR
MIN		0.13	1	50	0.0	0.0 $\pm$ 0.1	0.0
CD-A		0.4	2	250	0.0	-0.5 $\pm$ 0.2	-0.04
CD-B		0.49	3	250	3.8	6.4 $\pm$ 0.1	0.19
EV-A	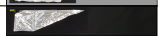	0.46	2	308	3.4	7.3 $\pm$ 0.3	0.28
EV-B		0.5	2	352	3.8	5.0 $\pm$ 0.4	0.09
EV-C		0.54	3	308	2.9	10.1 $\pm$ 0.7	0.52
EV-D	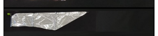	0.57	3	341	4.6	6.0 $\pm$ 1.0	0.1
EV-E		0.6	3	365	7.0	6.3 $\pm$ 0.1	-0.05
EV-F	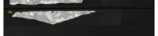	0.61	3	383	5.2	6.2 $\pm$ 4.8	0.07
EV-G		0.62	3	393	8.0	7.2 $\pm$ 0.2	-0.06
EV-H	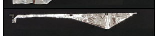	0.71	3	479	12.3	2.3 $\pm$ 0.4	-0.72
EV-I		0.75	4	443	5.4	2.4 $\pm$ 0.5	-0.22
EV-J		0.76	4	454	13.0	4.5 $\pm$ 0.2	-0.61
EV-K		0.82	4	517	13.1	3.2 $\pm$ 0.1	-0.71
EV-L		0.9	5	509	10.1	1.8 $\pm$ 0.1	-0.6
EV-M		0.92	5	526	13.9	4.4 $\pm$ 0.1	-0.68

TABLE I

TRANSFERRED WING DESIGNS AND METRICS

We performed sim2real transfer on a set of wing designs selected from the evolved NDF that covered the range of morphological simulation complexity and predicted lift. Additionally, as an experimental control against our automated design process, we custom designed a limited number of wing morphologies based on studies that have explored flapping wing aerodynamics. The first design was a square ribbed plate with rigid ribs inspired by Sane *et al* [29]. The second design is representative of the DelFly clapping wing construction [9] and included increased elasticity in the wing structure to allow deflection under aerodynamic loads, which is intuitively needed to create lift. Finally we created a minimal wing design representing the simplest design possible in the morphospace. These wings were mapped into our paramaterised wing descriptor format for fair comparison in sim2real transfer.

III. RESULTS

The design process was undertaken and produced three non-dominated fronts, one of which is shown in figure 1(c). The color of each individual represents $C_{MS}(m)$ for that design and shows that larger values of C_{MS} allow larger predicted lift. The non-dominated front shown included more than 100 morphology designs, produced efficient lift between 10mN and 200mN and represent a spread across a large portion of the complexity space.

Figure 3 shows our key result from this study which is the scatter plot of $STR(m)$ against $C_{MS}(m)$ of the transferred wings. That plot includes a bounded region (shaded in pink) set by connecting the upper and lower bound of STR between adjacent transferred wings in the evolved search space. Overlaid on the plot is the regression selected polynomial line of best fit (second order) for the transferred points (green dashed line). The shape of that region and line of best fit represents the empirically measured shape of the reality gap for the flapping wing design.

In our discussion of the reality gap shape from the introduction, we qualitatively predicted that STR would degrade monotonically with C_{MS}, and we plot a best fit linear decay relationship extending from the minimal solution (red dotted line) in Figure 3. It is clear from inspection that the empirically measured shaded region and polynomial best fit is substantially different to that prediction. The shape of the reality gap that we have measured identifies a non-linear relationship, where STR is small (less than ± 0.2) up to a threshold C_{MS} value of approximately 0.6, at which point STR falls off at a much faster rate. We interpret the rapid degradation in STR for designs above the C_{MS} threshold as a compounding of the smaller errors within the finite elements of our morphology simulation.

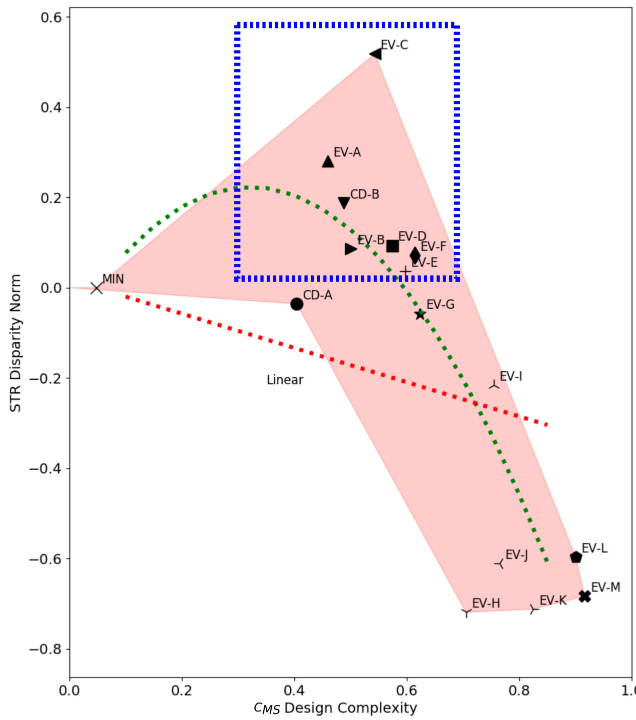

Fig. 3. Scatter plot of STR verses C_{MS} for sim2real transfers. Green dotted line - Polynomial best fit shape. Shaded region - bound of upper and lower STR values of neighbouring sim2real transfers. Red dotted line - monotonic decay prediction for STR drawn for comparison. Blue bounded region - region where simulation under represents performance in reality.

Further, in this example, we identify a region of the plot where substantially improved performance was obtained in reality compared to simulation. This region is highlighted in blue, and represents an area where the model is insufficient to capture a crucial positive operating effect, which we interpret to be the unmodelled wake capture. Investment in simulation improvements in this zone are likely to result in improved search results, and development of methods to automatically improve the modelling environment in this zone would extend the automatic design capability.

Regardless of STR, Table I shows that the wing with the best real lift is "EV-C" followed by "EV-A" and "EV-G" all of which outperformed the best hand designed wing, "CD-B". These wings all have lower aspect ratio than many of the other wings suggesting that there may be a contribution of that morphological similarity that directly affects lift. That same morphology group also have positive STR and are exploiting one or more effects that do not appear in our simulator. Conversely, wings with higher aspect ratio tended to have negative STR values and lower lift which suggests that this type of morphology leads to over-fitting to inaccuracies of the simulation that are not available to realised systems. In the future, these types of insights may be inferred automatically from transfer results and used to guide automatic finite element model improvement.

IV. Discussion

For the first time in this study we have measured the shape of the reality gap that emerges during morphology-in-the-loop robotics design. In this instance, we found the novel result that the shape is non-monotonic. Up to a threshold complexity level, real behaviour of the robot matches or even exceeds simulation performance. Above that threshold, STR quickly reduces resulting in the real performance of the robot being much worse than predicted. This study has produced one of only a few examples of morphological evolutionary search across the sim2real gap [22], [4], [24], [8].

Morphology is important for many robotics applications and there are many domains where human intuition may not be available for design. Rather than using bio-mimicry, we have shown that a machine learning design process can create a diverse set of morphologies tailored for a task from an underlying set of bio-inspired characteristics. Our approach led us to define a genotype that described a wing using features of shape, span, elasticity and inertia for both simulation and reality implementations of the flapping wing aerodynamics problem. We defined morphological simulation complexity based on measures relevant to the expected accuracy of finite element simulation. Using an evolutionary design process, we searched for high performance designs and showed that predicted lift in simulation increased with C_{MS}, but that real lift peaked at a threshold value of ≈ 0.6 and degraded STR and lift was found in the more complex designs. We interpret this rapid degradation as the compounding rather than monotonic additive effect of smaller errors in the finite element morphology simulation structure.

We hypothesised that CFD would not be suitable for sim2real design due to computational overheads limiting search in simulation, and instead developed a fast, comparatively low fidelity simulation to facilitate automated search. We found this revised approach was able to design high performance wings for successful sim2real transfer up to a threshold level of C_{MS}, and that those designs out-performed hand designed wings. This suggests that our simulation approach retains a gradient that evolution can follow towards high performing designs in reality. Whether a similar reality gap shape would be found for a different simulation approach like CFD is a topic for future research.

The results are encouraging for the utility of machine learning based sim2real morphology design using finite element simulation. Had the shape of the gap been found to degrade proportionally (or worse) with design complexity, then an automated design target that produced maximum lift with minimum complexity would be the obvious search objective, but could only produce near trivial designs. For our problem, we have shown that a richer search space of wing morphologies that narrow the sim2real gap is available for exploration automatically for specific tasks.

The reality gap shape for morphology-in-the-loop design observed in this study is the first such measurement reported and it remains to be seen if this shape holds for other morphology domains. Our finite element simulation approach was applied linearly along the span of a flat wing, and an interesting question would be to determine how dimensionality of the finite element structure would affect the relationship. Applying a similar analysis to a 3 dimensional design problem, such as those in [22], would assess how compounding of inaccuracies affects the sim2real relationship.

We focused this study on morphology effect of sim2real transfer and used a fixed control policy throughout. It remains to be seen how the sim2real gap would be affected when controllers are also included into the search space. For instance, are the morphologies we found in this study better able to cross the reality gap and produce lift than wings designed using different methods, or at random, when combined with tailored controllers?

Robust automatic robot design through the inclusion of sim2real feedback with our morphology exploration process is a topic for continued research and we identify two options for development. First, we will investigate extension of the transferability approach [18] to morphology design by incorporating the reality gap shape and its relationship to C_{MS} as means to tune the size of design space to suppress in the region surrounding poor transfers. Secondly, we identified in this application a region where performance in reality was better than in simulation. While we believe that this is due to the non-modeled wake capture effect and that directly incorporating this complex feature within the simulation would apply correction, our desire for an automatic design process suggests that investment in automated sim2real feedback to produce modelling updates of the finite element morphology simulation might reduce the need for human led model enhancement. Using this, we hope to demonstrate the ability of biological evolution to build upon fortuitous morphological design features that improve performance without the requirement of human intuitive understanding.

Our intent is to take automated wing design from this one degree of freedom example to include multi-wing, multi-degree of freedom and variable control policies similar to the properties observed in natural fliers. The scalability of our initial approach is suited to this. Our test rig design in this study uses small sensors, low cost actuators and quickly constructed wings, and will allow us move from evaluating forces on a single tethered wing to multi-wing morphologies and to use varying control policies. Using this revision, we aim to undertake morphology and control selection to achieve lift as well as the force and torque vectoring demonstrated in stabilised hover and flight of biological fliers.

References

[1] http://www.arduino.cc.

[2] National instrument. ni usb-6225,6003. http://www.ni.com/datasheet/pdf/en/ds-10. 2016.

[3] Josh Bongard, Victor Zykov, and Hod Lipson. Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121, 2006.

[4] Daniel Cellucci, Robert MacCurdy, Hod Lipson, and Sebastian Risi. 1d printing of recyclable robots. *IEEE Robotics and Automation Letters*, 2(4):1964–1971, 2017.

[5] Nicholas Cheney, Jeff Clune, and Hod Lipson. Evolved electrophysiological soft robots. In *ALIFE*, volume 14, pages 222–229, 2014.

[6] Nick Cheney, Josh Bongard, Vytas SunSpiral, and Hod Lipson. Scalable co-optimization of morphology and control in embodied machines. *Journal of The Royal Society Interface*, 15(143):20170937, 2018.

[7] Jack Collins, Wade Geles, David Howard, and Frederic Maire. Towards the targeted environment-specific evolution of robot components. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 61–68. ACM, 2018.

[8] Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W Sumner, Wojciech Matusik, and Bernd Bickel. Computational design of mechanical characters. *ACM Transactions on Graphics (TOG)*, 32(4):83, 2013.

[9] GCHE De Croon, KME De Clercq, Remes Ruijsink, Bart Remes, and Christophe De Wagter. Design, aerodynamics, and vision-based control of the delfly. *International Journal of Micro Air Vehicles*, 1(2):71–97, 2009.

[10] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.

[11] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E Gusz Eiben. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2:4, 2015.

[12] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019.

[13] Nick Jakobi. Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive behavior*, 6(2):325–368, 1997.

[14] Nick Jakobi, Phil Husbands, and Inman Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, pages 704–720. Springer, 1995.

[15] Martin Alexander Jones. *CFD analysis and design optimization of flapping wing flows*. PhD thesis, North Carolina Agricultural and Technical State University, 2013.

[16] Manas Khurana and Javaan Chahl. Bioinspired corrugated airfoils for micro air vehicles. In *Bioinspiration, Biomimetics, and Bioreplication 2013*, volume 8686, page 86860X. International Society for Optics and Photonics, 2013.

[17] Jia Ming Kok and Javaan Chahl. Design and manufacture of a self-learning flapping wing-actuation system for a dragonfly-inspired mav. In *54th aiaa aerospace sciences meeting*, page 1744, 2016.

[18] Sylvain Koos, Jean-Baptiste Mouret, and Stéphane Doncieux. The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation*, 17(1):122–145, 2013.

[19] Sam Kriegman, Collin Cappelle, Francesco Corucci, Anton Bernatskiy, Nick Cheney, and Josh C Bongard. Simulating the evolution of soft and rigid-body robots. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1117–1120. ACM, 2017.

[20] Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Belson, David M Bryson, Nick Cheney, et al. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *arXiv preprint arXiv:1803.03453*, 2018.

[21] Chengyu Li and Haibo Dong. Wing kinematics measurement and aerodynamics of a dragonfly in turning flight. *Bioinspiration & biomimetics*, 12(2):026001, 2017.

[22] Hod Lipson and Jordan B Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974, 2000.

[23] Yucheng Liu and Gary Glass. Effects of mesh density on finite element analysis. Technical report, SAE Technical Paper, 2013.

[24] Ankur M Mehta, Joseph DelPreto, Benjamin Shaya, and Daniela Rus. Cogeneration of mechanical, electrical, and software designs for printable robots from structural specifications. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 2892–2897. IEEE, 2014.

[25] Igor Mordatch, Kendall Lowrey, and Emanuel Todorov. Ensemble-cio: Full-body dynamic motion planning that transfers to physical humanoids. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 5307–5314. IEEE, 2015.

[26] T Nakata, H Liu, Y Tanaka, N Nishihashi, X Wang, and A Sato. Aerodynamics of a bio-inspired flexible flapping-wing micro air vehicle. *Bioinspiration & biomimetics*, 6(4):045002, 2011.

[27] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018.

[28] Daniela Rus and Michael T Tolley. Design, fabrication and control of soft robots. *Nature*, 521(7553):467, 2015.

[29] Sanjay P Sane and Michael H Dickinson. The aerodynamic effects of wing rotation and a revised quasi-steady model of flapping flight. *Journal of experimental biology*, 205(8):1087–1096, 2002.

[30] Wei Shyy, Hikaru Aono, Satish Kumar Chimakurthi, Pat Trizila, C-K Kang, Carlos ES Cesnik, and Hao Liu. Recent progress in flapping wing aerodynamics and aeroelasticity. *Progress in Aerospace Sciences*, 46(7):284–327, 2010.

[31] Karl Sims. Interactive evolution of dynamical systems. In *Toward a practice of autonomous systems: Proceedings of the first European conference on artificial life*, pages 171–178, 1992.

[32] Karl Sims. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22. ACM, 1994.

[33] Bo Yin and Haoxiang Luo. Effect of wing inertia on hovering performance of flexible flapping wings. *Physics of Fluids*, 22(11):111902, 2010.