An improved version of this pre-print has been published in Nature: Machine Intelligence.
doi: 10.1038/s42256-021-00320-3

Real-world Embodied AI: A Morphologically Adaptive Quadruped Robot in the Wild

Tønnes F. Nygaard[1,2*], Kyrre Glette[1,3], Charles P. Martin[4], Jim Torresen[1,3], David Howard[5]

[1]Department of Informatics, University of Oslo, Norway
[2]The Norwegian Defence Research Establishment (FFI), Kjeller, Norway
[3]RITMO, University of Oslo, Norway
[4]Research School of Computer Science, Australian National University, ACT, Australia
[5]Cyber-Physical Systems Program, CSIRO, QLD, Australia

[*]To whom correspondence should be addressed; E-mail: tonnesfn@ifi.uio.no

Robots are traditionally bound by a fixed morphology during their operational lifetime, limited to adapting only their control strategies. Here, we present the first quadrupedal robot that can morphologically adapt to different environmental conditions in outdoor, unstructured environments. Our solution is rooted in embodied AI, and comprises two components; (i) a robot that permits in-situ morphological adaptation, and (ii) an adaptation algorithm that transitions between the most energy-efficient morphological configurations based on sensed terrain. First, we build a model that describes how the robot morphology affects performance on selected terrains. We then test continuous adaptation on previously unseen terrains while allowing the robot to constantly update its model. We show that the robot exploits its training to effectively transition between different morphological configurations, showing significant performance improvements over a non-adaptive approach. The demonstrated benefits of real-world morphological adaptation show the potential for a new embodied way of incorporating adaptation into future robotic designs.

Robots inspecting the damaged Fukushima reactor were presented with a daunting task: to pass through a narrow duct to enter the area, traverse gaps between platforms, move over and through various types of debris, and even swim through murky water. Designing a robot to work across such diverse and unstructured environments is challenging as task and environmental conditions may change, sometimes drastically, during operation. As such, technological limitations meant that the eventual solution required numerous highly specialized traditional robots, with correspondingly high numbers of deployments and extended mission times [1]. The challenges presented by Fukushima, chiefly multimodality and unpredictabilty, are characteristic of the type of unstructured environment that robotic systems as a whole continue to struggle with.

We postulate that the key to developing such flexible, adaptable robots may lie in a specific subfield of Machine Intelligence called Embodied Cognition [2]. The theory of Embodied Cognition states that the brain (software) is not the sole source of cognition, but rather that orchestration of interactions between brain (software), body (hardware), and environment are key to producing intelligent action [3]. Viewed through the lens of Embodied Cognition, the physical manifestation of a robot is a crucial adaption tool, which could be vital in achieving resilient robots that can operate across challenging real-world environments [4]. Indeed, in some cases, changing the robot's morphology might be the only viable option to elicit suitable in-environment behaviors [5].

Let's return to Fukushima, as it is a great example of the type of challenging environment we want robots to be able to operate in. An arguably more efficient and attractive (not to mention capable) solution to working in and around the reactor is a single 'Swiss army knife' robot. Capable of online morphological adaptation, this robot would be able to harness its variable morphology and variable controller to more strongly tie its behaviours to its immediate environment, which, considering Embodied Cognition, may increase its ability to perform tasks. Such a robot would be able to match its capabilities to its immediate needs: having at one time a large span to traverse gaps, yet at another time being able to shrink and squeeze through narrow openings in debris fields. Pure controller adaptation, as per most robots nowadays, cannot provide this ability. Shape-shifting (or morphologically adaptive) robots have long been a mainstay in our collective consciousness[1]. And with good reason; they represent an appealing future where robots have become masters over their environment, able to adopt a variety of configurations to meet their immediate and long-term needs and improve mission outcomes.

The underlying principle is that variable morphology provides additional degrees of freedom to adapt to a given environment compared to a static morphology, increasing the likelihood that the robot can adapt and survive in the face of unpredictable environmental conditions. In principle then, morphologically adaptive robots are a promising enabling technology to unlock operation in a broad swathe of unpredictable environments and tasks on the fly, without having to be redesigned and rebuilt each time they face something

[1]Often found in science fiction popular culture, e.g., The Transformers, or T-1000 from the Terminator series)

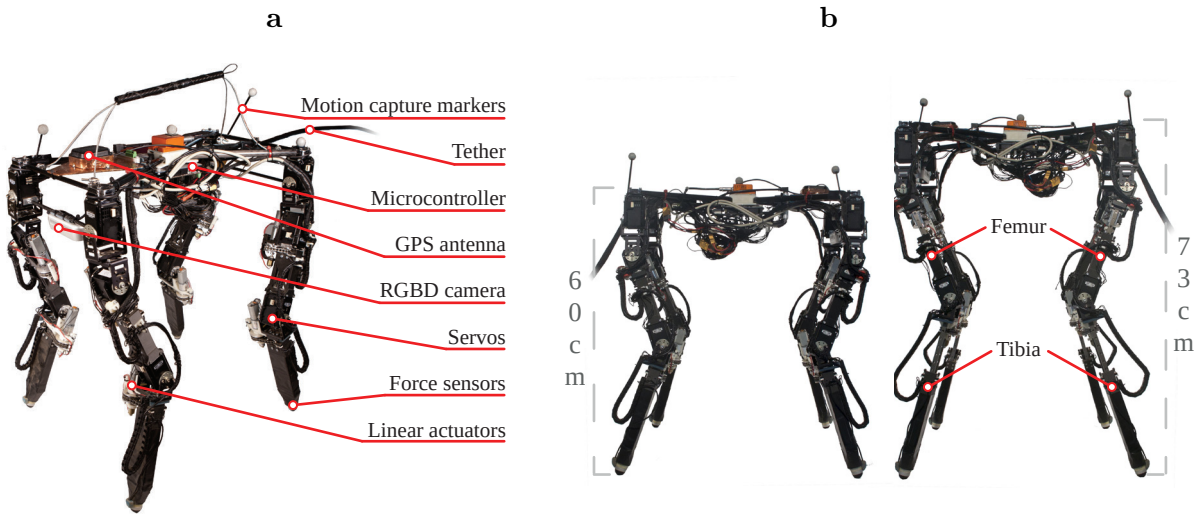

Fig. 1: **The morphologically adaptive robot used in this study.** (**a**) An overview of the main components of the robot. (**b**) The robot with the shortest (left) and longest (right) leg configuration.

unexpected. Due to this promise, morphological adaptation is an area of increasing scientific focus that encompasses a range of research from variable stiffness robot limbs [6] to elegant origami-inspired morphing structures[7].

In this paper, we present our own morphologically adaptive robot for unstructured environments, seen in Fig. 1a. The quadruped robot, DyRET[2][8, 9], provides a powerful proof of concept harnessing a variable morphology to adapt to realistic real-world conditions in outdoor settings. Morphological adaptation is provided through variable-length legs, whereby the length of both femur and tibia can be adjusted to enable different walking behaviours, whilst also tilting the central body, shown in Fig. 1b. A novel terrain-adaptation algorithm controls this morphing. Bootstrapped with knowledge from controlled experimentation in terrain boxes, it can continually alter the morphological configuration of the robot to optimize energy efficiency when traversing unstructured terrains based on sensed terrain characteristics.

This work is inspired by multiple fields, including legged robotics, embodied cognition/AI, and evolutionary robotics. We can broadly segregate the literature into three topics: (i) controller adaptation with static morphology, (ii) morphological adaptation offline, and (iii) morphological adaptation online.

Biologically-inspired legged robots are a promising solution for unstructured environments. Adaptation can be realized purely through software, primarily adapting gait patterns and foot-tip arcs. Techniques that allow locomotion on challenging terrain include evolutionary approaches [9, 10, 11], reinforcement learning [12, 13], and Bayesian opti-

[2]Dynamic Robot for Embodied Testing

mization [14, 15], as well as perception-less [16] and hybrid approaches [17, 18, 19]. Online adaptation to terrains of different compliance under aggressive maneuvers and external disturbances has been studied [20], as well as walking posture adaptation for navigation in confined spaces [21]. However, these approaches are implemented on a static morphology which limits the level of attainable environmental adaptation.

Evolutionary robotics and Artificial Life have deep links to embodied AI, and are concerned with investigating and understanding biological processes, including adaptive bodies [22, 23]. Recent work has studied co-optimization of control and morphology using gradient approaches in differentiable physics simulators [24], which has strong potential to efficiently couple control, morphology, and environment. These efforts are, so far, only simulated and run in simple environments. Indeed, most works in adaptive robotic morphology are carried out in physics simulation, and not on physical robots [25]. Examples include soft robots [26], modular robots [27], and legged robots [28].

The next step up from pure simulation of adaptive morphology is selecting a few virtual robots for real-world manufacture and testing [29, 30, 31, 32, 33]. However, the performance of these robots is often limited due to the inaccuracies in the simulation or models used, referred to as the reality gap [34, 35]. This discrepancy means that robots with morphologies optimized in simulation are not fully adapted to the intricate physical environments they will eventually operate in, but to a simplified version of it.

Our approach to morphology adaptation is performed exclusively in hardware, which is guaranteed to work in reality. Other examples where the body of a robot is optimized or changed in the real world directly are relatively rare, including manual assembly [36, 37] or an external mechanism for reconfiguration [38, 39]. Such approaches require excessive time, external apparatus, or human intervention and are not suitable for continuous adaptation during independent operation. There are some examples of robots with a built-in ability to change their own body during operation. Many of these robots are relatively small with no payload capacity, and are limited in their ability to function in real-world unstructured settings, e.g., [40]. More complicated robots possess a higher potential to solve real-world problems, including morphing drones [41, 42], and multi-modal legged-wheeled [43] and wheeled-flying robots [44]. These more advanced robots typically discretely change between a couple of pre-defined morphologies, whereas in this paper we sample morphologies from a continuous range.

Compared to the identified literature, our approach is the first to continuously optimize the morphology of a real legged robot with the capability to hold a reasonable payload and, in principle, carry out various missions, outdoors in the real world. It also makes DyRET the first 'fully featured' robot of its size, with software, sensing, and actuation, to close the embodiment brain-body-environment loop in a challenging real-world setting. To show this, we fill large boxes with real terrain material and train a simple regression model relating the sensed terrain to the performance of the different morphological configurations of the robot. We validate this model in a simple scenario indoors. Finally, we run the robot on previously unseen heterogeneous terrains outside, where we

test continuous adaptation of morphology while simultaneously updating the regression model 'in the wild'. Continuous adaptation outperforms a challenging baseline of the best static configuration discovered during the bootstrapping phase.

Results

To enable efficient adaptation in real-world environments, we start by gathering a baseline data set used to bootstrap the subsequent adaptive process. Our system is then evaluated in two different scenarios: (i) Adapting to previously seen homogeneous terrains, and (ii) Continually adapting to previously unseen heterogeneous terrains. In both, we compare the performance of the adaptation algorithm to the best performing static morphologies from the baseline data set.

Gathering the data set

A baseline data set was collected to pre-learn a model of how the robot's morphology affects its performance on different terrains, facilitating efficient adaptation in real-world environments by avoiding potentially poor-performing morphologies. It enables bootstrapping of the subsequent learning process and, in our final experiment, is updated continuously as the robot operates in new environments.

Wooden boxes were filled with terrain materials purchased from a landscaping supplier (Fig. 4a). Our robot senses both hardness and roughness of its terrain, so we selected three materials with a spread in these two terrain characteristics (details can be found in Supplementary Table 3). Sand is soft with low roughness, gravel is hard with high roughness, and a fiber-reinforced concrete sheet provides a hard surface with low roughness. Each box consists of two halves filled with different terrain materials. This allows the robot to walk on the separate terrains, as well as across various terrain transitions. The boxes were placed in a motion capture facility for high accuracy indoor positioning.

A minimum change in leg length is needed before seeing a notable effect on robot behavior, so each leg segment was limited to five uniformly sampled discrete lengths, giving 25 different morphological combinations in total. The robot walks with a forward velocity of about 2m/min for 15 seconds per morphology, covering all 25 combinations. Each morphology is repeated for five different starting locations per terrain type to cancel out any local variation in the surface. The robot does not traverse any transitions at this stage. The data set contains approximately 90 minutes of walking data.

The measured COT for each morphology on the three surfaces can be seen in Fig. 2. When walking on the concrete (Fig. 2a), the robot achieves the best energy efficiency with a long femur and short tibia, as well as a medium femur and medium tibia. On sand (Fig. 2b), the robot achieves a high efficiency for short to medium length tibias, with femur length having less of an effect. Much less consistency is seen in the gravel (Fig. 2c), but the best COT is seen for the shortest possible legs.

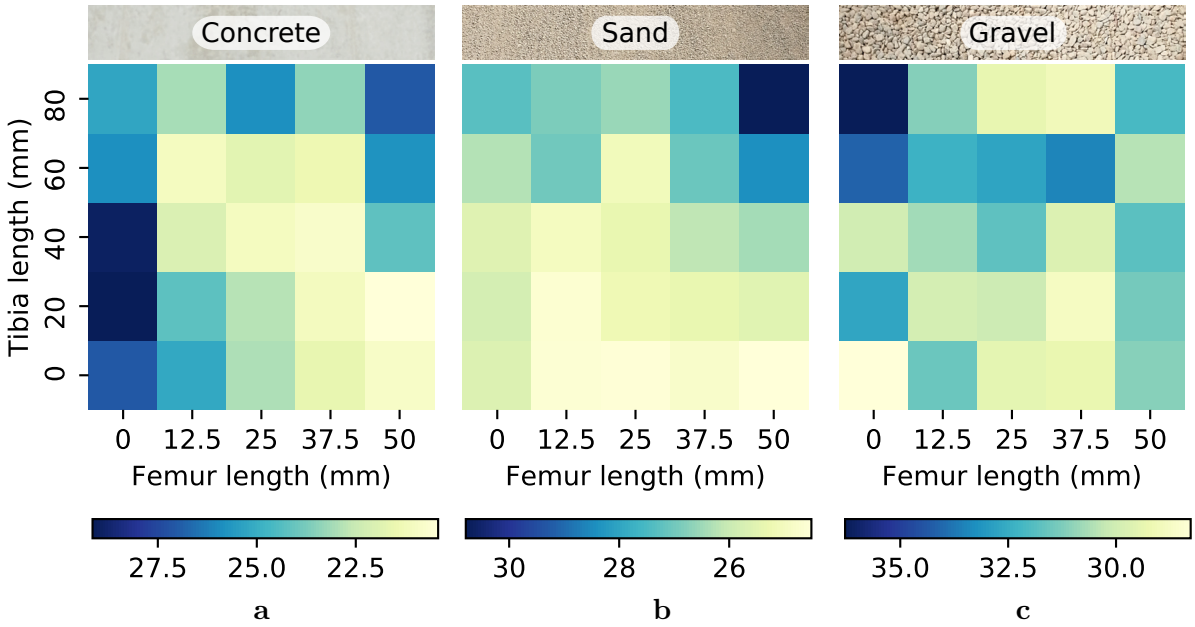

Fig. 2: **The cost of transport (COT) of different leg lengths on the three terrains.** The lower the COT is (yellow), the more efficiently the robot is walking. Please note that the ranges – as seen in the bottom of the figure – are not the same for each surface, to better highlight local differences within each terrain.

Adapting to previously seen homogeneous terrains

In this preliminary experiment, we demonstrate a simplified case where morphology is adjusted based on sensed terrain characteristics, but terrains are present in the training data and discretely-separated within terrain boxes (Fig. 3a). In this case there is also no need to continuously change the leg lengths, as the terrains are considered homogeneous. As the terrains are known, the adaptation algorithm takes the form of a classifier – see the Methods section. The robot is brought to a standstill before the morphology is changed, which is triggered by the onboard sensors detecting a step onto a new terrain type (details in section the Methods section). This serves as a simplified validation of our final, continuous, adaptation method.

We used the same terrain boxes used for collecting the baseline data set, seen in Fig. 4a. The first half was covered in the concrete sheet, with the rest comprising of gravel. The lowest-COT morphology for each surface is chosen from our baseline data set (femur 50mm, tibia 20mm for concrete; femur 0mm, tibia 0mm for gravel), and serves as a comparison for the adaptive morphology. Each morphology begins on concrete and walks onto gravel, triggering a change in morphology in the adaptive case. More details on the experiment design can be found in the Methods section.

Fig. 4b shows that morphologies specialized for one terrain do not transfer well to

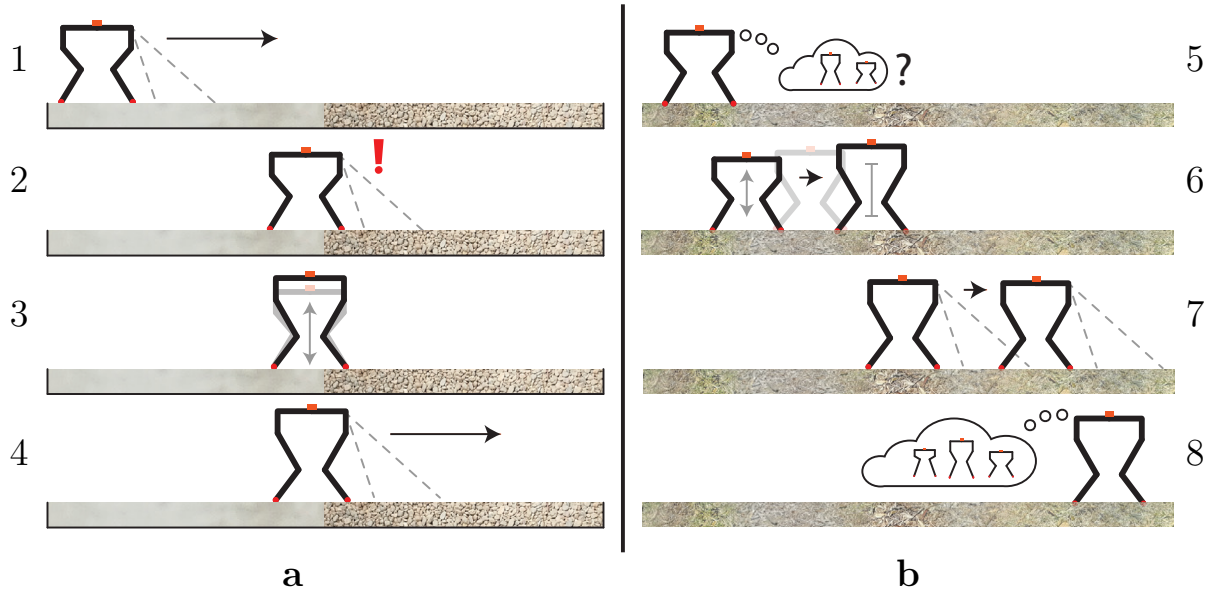

Fig. 3: **Diagram showing the two adaptation methods used.** (**a**) Adaptation principle for previously seen homogeneous terrains. 1: The robot walks forward while sensing its environment. 2: Once a terrain change has been detected, it stops walking. 3: The robot changes the length of its legs to the optimal morphology for the new terrain. 4: It starts walking with the new morphology, sensing for the next terrain transition. (**b**) Adaptation principle for previously-unseen heterogeneous terrains. 5: The robot predicts the best performing morphology based on sensor readings and its internal model. 6: It changes the length of its legs to this new morphology while walking. 7: When the legs have reached their goal length, the robot measures its performance and the terrain characteristics. 8: It adds the new measurements to its internal model, before repeating the process from step 5.

others, and that no single morphology is best across both terrains. This is expected given our terrain selection method tried to use terrains with different characteristics. The concrete-specialized morphology achieves a mean COT of 23 while walking on concrete, which rises to 37 after the transition, resulting in a reduction in energy efficiency of $\approx 60\%$. The gravel-specialized morphology starts with a mean COT of 36, but achieves 26 on gravel, showing an improvement of $\approx 70\%$ after stepping onto the optimal terrain for the morphology. The adaptive morphology is shown to perform consistently well across these known terrains, and the change detection algorithm causes a switch in morphology at the appropriate time.

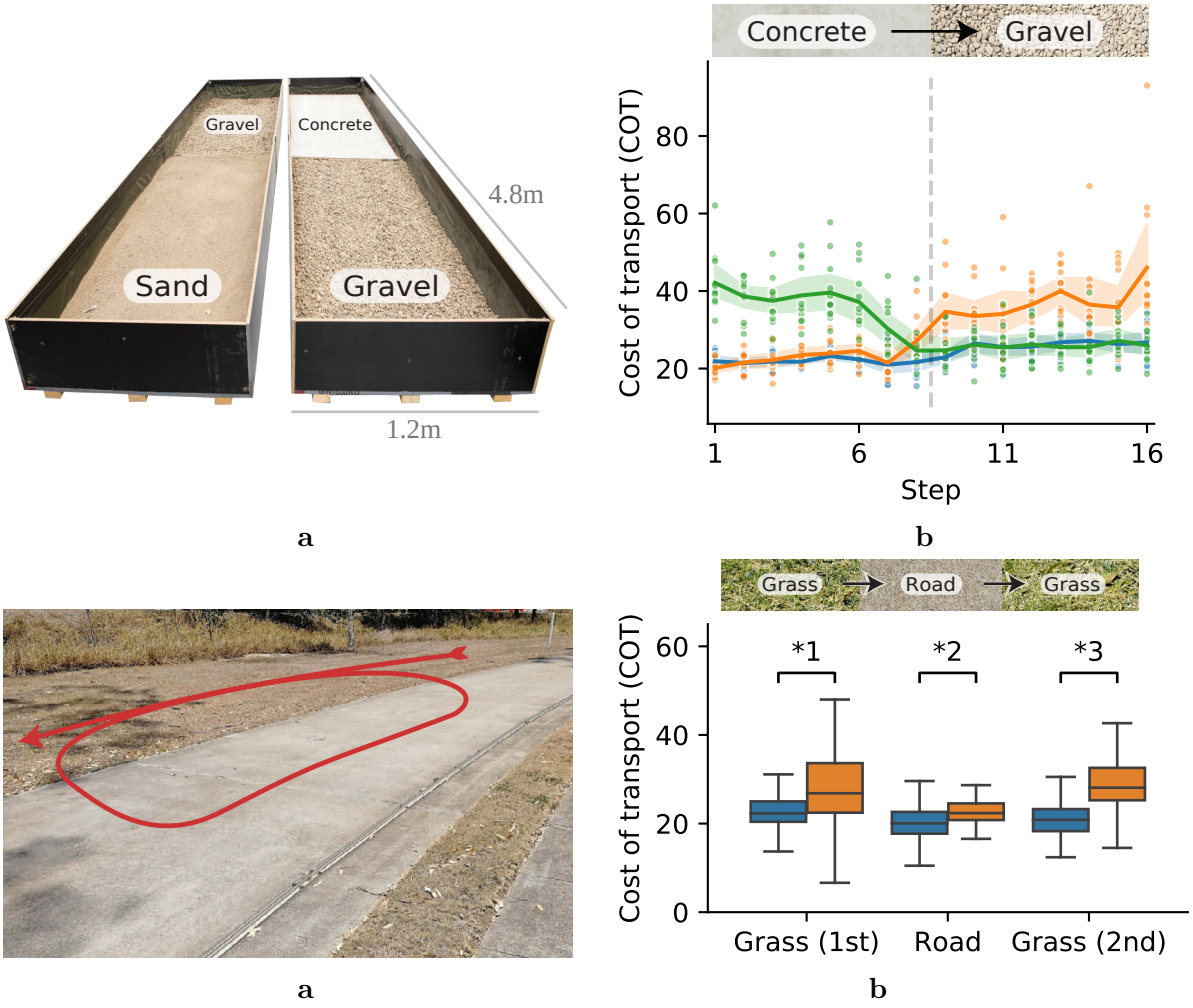

Fig. 4: **Previously seen homogeneous terrain experiments**. (**a**) The terrain boxes used for the experiment on previously seen homogeneous terrains. They contain sand, gravel and concrete. (**b**) The COT for the two best static morphologies (gravel-specialized in green and concrete-specialized in orange) and the adaptation (blue) when walking on concrete, then gravel in the boxes. The solid lines show the mean of 5 repeats, and the shaded areas are the 95% confidence interval. (**c**) The outdoor area used in the experiment on previously unseen heterogeneous terrains, with the red line showing a typical walking path for the robot. (**d**) Energy efficiency (COT) of the adapting (blue) and best all-rounder static morphology (orange) on the outside test track. **Statistically significant differences from two-sided Mann-Whitney U test on each parameter with Holm-Bonferroni p-value correction ($p < 0.01$, n1 = 160, n2 = 64, **1: U = 2540, **2: U = 747, **3: U = 2001).

Adapting to previously unseen heterogeneous terrains

Fig. 3b describes an extended method that takes into consideration the additional challenges of unseen heterogeneous terrains. As these terrains are less homogeneous and structured, we cannot assume that all terrains that the robot will encounter are present in the baseline data set. We therefore replace terrain classification (as in our previous approach) with characterization (details in the Methods section). The terrain is considered to be heterogeneous and might change substantially for every single step the robot takes. With this extended method, the robot does not stop to change morphology at any point, but operates continuously while the morphology slowly adapts and new experiences are added to the (now adaptive) model.

An outside test track was selected, seen in Fig. 4c. The route starts with a section of grass, before the robot steps on to a concrete road, then back on the grass. Returning to the same surface again shows to what degree the algorithm is able adapt its model based on previous experience of walking on grass.

The robot uses the model detailed in the Methods section to predict the best performing morphology on its current terrain. Since changing the length of the legs takes a considerable amount of time, only neighboring morphologies are considered (within 12.5mm for the femur, and 20mm for the tibia). The terrain and performance is not evaluated while the legs are changing length, but after the morphology has been achieved. The robot takes three steps per leg to get a representative measure, and we refer to this as an *evaluation*. It only reconfigures if any of the neighbors are predicted to outperform the efficiency it just achieved with it's current morphology. If not, it simply evaluates the same morphology again. Evaluations are therefore not based on discrete terrain changes or time passing, but are done continuously. Additional evaluations will improve the model, even for repeat measurements for the optimal morphology. The algorithm was allowed 32 evaluations on each terrain section, before being led onto the next. The best all-round static morphology (lowest COT across all 3 terrain types from the baseline data set: femur 37.5mm, tibia 20mm) served a comparison.

Fig. 4d shows the energy efficiency (COT) of every morphology evaluated while adapting. We see that for the first grass section, adaptation gives a median COT of 22, while the static morphology has 27, a reduction in efficiency of $\approx 17\%$. We see similar reductions in the median for the road and second grass surface of $\approx 10\%$ and 26%, respectively. The adaptation significantly outperforms the all-round best performing morphology on all three terrain sections.

Fig. 5a shows the difference between the predicted energy efficiency (COT) for the selected morphology and the actual efficiency measured after walking. The error in COT starts very high at ≈ 25, falls below 6 after trying 16 different morphologies, before ending at ≈ 2.5 at the end of the first grass section. The error spikes to ≈ 8 as the robot steps onto the road and ends up at ≈ 4 at the end of the section. When stepping back onto the grass, the error spikes up to 11, but that is still much less than initially encountered on

the first grass section. It very quickly converges and reaches below 6 in 8 more evaluations before ending at a COT of ≈ 5. These statistics are based on a locally-weighted regression on five repeats of the adaptive run, and contain uncertainty reflected in the confidence intervals in the figure.

Fig. 5b shows which morphologies are utilized on each terrain type. Mostly short femur with long tibia combinations are evaluated while the robot walks on the initial grass section, as seen to the left in the figure. When it steps onto the road, shown in the middle, there is a shift to long femur, long tibia combinations. The final grass section does not differ much from the road. The adaptation algorithm exploits almost the entire morphological range. It also delineates the benefit of having an adaptive morphology — the best generalist static morphology is consistently outperformed by this adaptation.

Fig. 6 shows how the model changes its understanding of the two outdoor terrains during the adaptation runs. The mean terrain characteristics from the outside test track (roughness 51.0 and hardness 143.4 for grass; roughness 16.5 and hardness 187.4 for the road) was used to visualize the model output at four different stages of the adaptation process. The initial maps generated solely on the baseline data set contain many extreme COT values, both at 0 and above 40. The optimal morphologies in the baseline data set achieved COT values in the range of approximately 18 to 25, while the worst morphologies were above 35, giving us a reference for realistic COT values. After walking on the first grass section, multiple prediction updates are seen in red in the second column of the figure. COT values for the updated grass model are in the range of 21-26, which can be considered realistic. The road prediction is more varied, with COT values as low as 12, which is considered unrealistic. After walking on the road, similarly large updates are seen to the road model in the third column, where we now have COT values between 19 and 25. The grass model is also slightly updated. After this section the model has experienced both terrain types, so when transitioned to the final grass section we see that only seven squares are updated for grass and five for the road. Only two of 25 possible leg-length combinations (femur 12.5mm, with tibia lengths 60mm and 80mm) were updated both in the first grass section and the last. This shows that the adaptation algorithm has successfully integrated previous experience of unseen terrains to rapidly generate low-error predictions.

Fig. 6 also serves as a demonstration of how the algorithm explores the space of available morphologies. We see from the grass map after the first grass section that the best predicted COT is in the top right area (femur 50mm, tibia 60mm). Because of the time taken to transition between morphologies, the adaptation algorithm is limited to only selecting the next morphology from neighboring cells. While this gives the benefit of being able to test a range of different morphologies in a short amount of time in a stable, controllable manner, it also means that areas like this are left unexplored since it is surrounded by low-COT cells. We see, however, that the area is visited when the robot returns to the grass section for a second time, and that it does in fact outperform the morphologies tested initially.

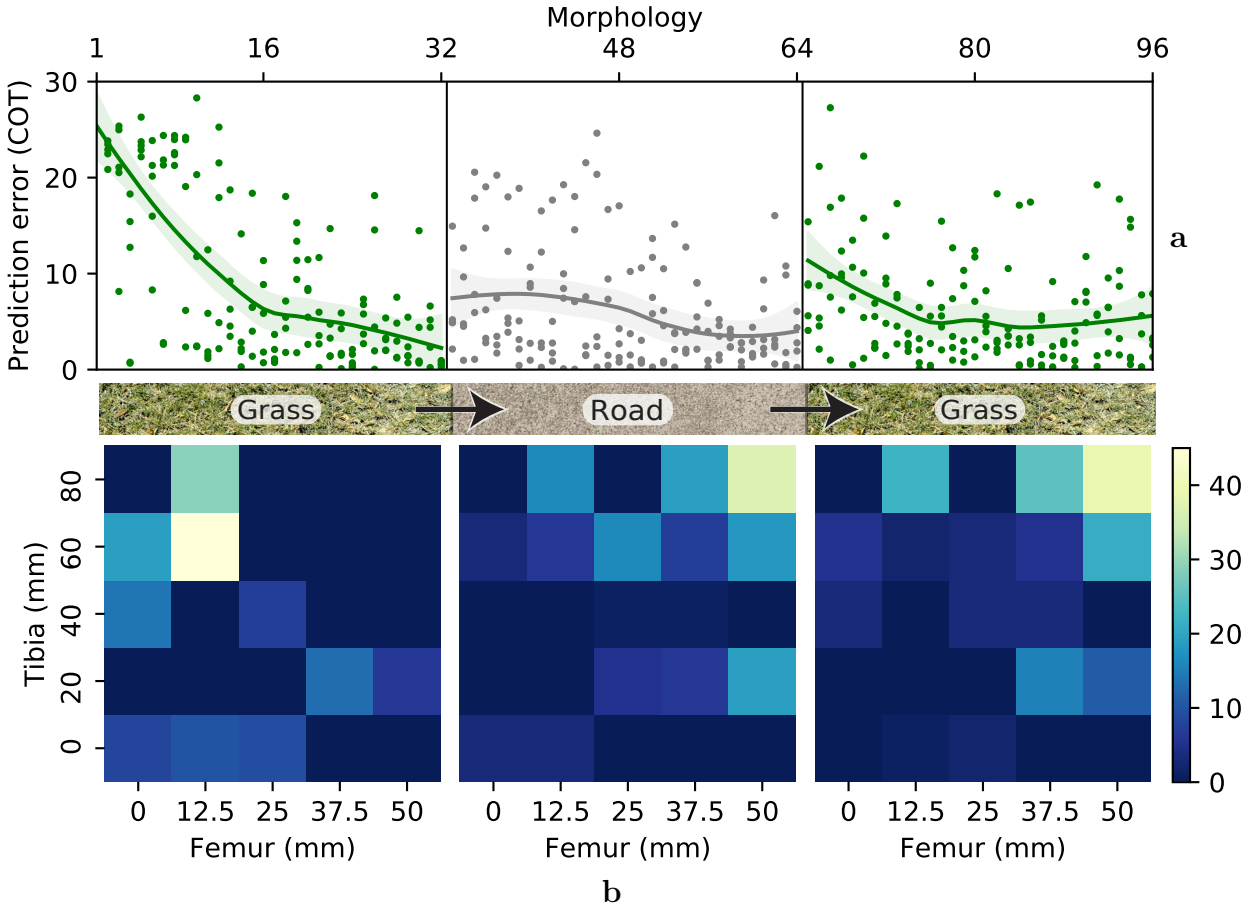

Fig. 5: **Analysis of the adaptation algorithm**. (**a**) The prediction error for each evaluated morphology as the robot walks on grass, road, then back on grass. The plot shows a locally-weighted regression with unweighted fit, while the shaded areas shows the 95% confidence interval of the regression. (**b**) The number of times each morphology was evaluated on the three terrain sections, summed over the whole experiment. The initial grass section is to the left, road in the middle, and the final grass section to the right.

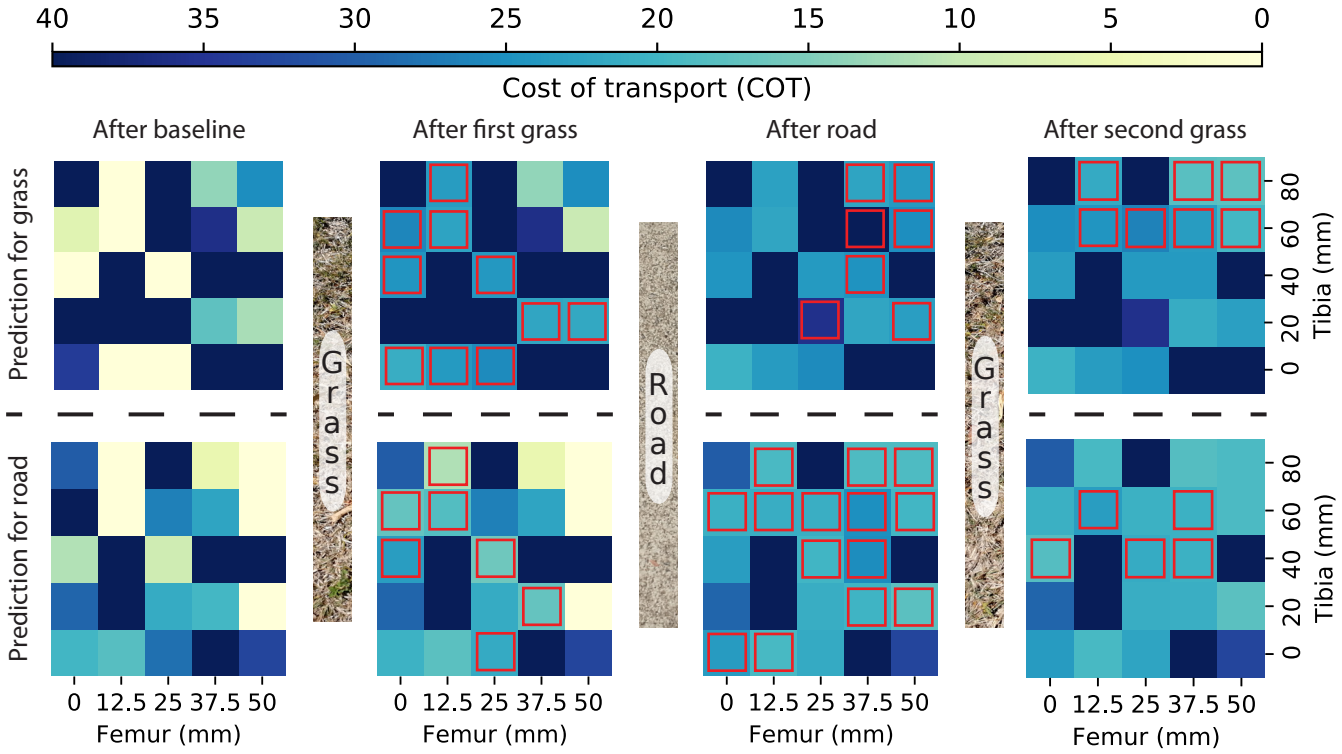

Fig. 6: **Model transformation during adaptation.** The first row shows the predicted energy efficiency (COT) map of all leg lengths for the grass surface, while the second shows for the road surface. The first column shows the initial models based solely on the baseline data set. The second shows the model after the walking on the first grass section, third after the road section, and the fourth after the final grass section. Data from all five iterations of the algorithm is included. All squares with a change in COT higher than 1 is marked with a red square.

Discussion

This work serves as an important step in the direction of morphologically adaptive robots and the overall goal of robotic operation in unstructured environments. We have demonstrated the validity of our approach, which harnesses morphological adaptation in an Embodied AI context to provide significantly improved performance compared to the best single static morphology.

Key experimental takeaways are (i) we can build a model of how cost of transport is affected by terrain and morphology, (ii) we can use this model to adapt to known homogenous terrains (a proof of concept), and most importantly (iii) we can combine the model with an adaptation algorithm, allowing the robot to continually vary its morphology in response to previously-unseen environments in live outdoor experiments over natural,

unstructured terrains, using the previously-learned model as a reference point. In our testing, the system quickly learned high-performance morphologies on grass, even though it had only previously experienced sand, gravel, and concrete. Importantly, this dynamic morphology strategy is shown to achieve better energy efficiency than any single static morphology during testing, and highlights adaptive morphology as advantageous trait for robots operating in unstructured terrains.

There are a few main limitations to this approach. Primarily, we use a 1-to-1 mapping of controller to morphology, rather than explicitly searching for effective body-brain combinations — a trade-off of faster adaptation speed for less behavioural diversity. Future efforts may focus on learning more sophisticated models on which to rapidly prototype control schemes before real-world rollout, as explicit controller adaptation may facilitate a more diverse behavioural repertoire for a broader range of terrains. We can also consider more advanced morphological adaptation mechanisms [6, 45, 46]. To balance reconfigurability with mechanical simplicity and stability, our adaptation actuator mechanism has a speed of $\approx$1mm/s. Faster adaptation would be advantageous in highly dynamic environments, where the robot has to constantly play 'catch up' between its instantaneous morphological configuration, and the best configuration as predicted by the model. In practice this was never an issue, as the nearest-neighbour adaptive approach was specifically designed to work on the hardware. Improvements to terrain modelling would also bring benefits. We used second-order polynomial regression models to facilitate analysis and gain an understanding the underlying mechanisms and effects of the adaptation process. We also chose to look at each morphology separately, instead of making one whole model that incorporates all data points.

Implications of these results are potentially far-reaching. We hope to inspire the design and adoption of similar mechanisms, for example, in commercially available platforms, to further increase their range, the tasks they can complete, and their possible operational environments. Our key takeaway is that morphological adaptation to real-world environments is a powerful and promising technique to conquering unstructured terrains, with significant benefits over the static morphologies that are ubiquitous within current robotics literature. We hope that our research helps to pave the way towards flexible hardware platforms that are capable of performing a variety of useful missions in outdoor, unstructured terrains.

Methods

Robot platform design

Our robotic platform, the Dynamic Robot for Embodied Testing (DyRET), can be seen in Fig. 1. It is a quadrupedal mammal-inspired robot with the ability to change the length

of its legs during operation and a fully certified open source hardware project[3]. The body weights approximately 5kg, measures 50cm by 30cm, and stands between 60cm and 73cm tall, depending on the pose and leg length. It has previously been used in laboratory settings, e.g., [47].

Fig. 1a shows the main components of the hardware design. The central body consists mainly of carbon fiber tubing, milled aluminum, and 3d printed plastic parts, as well as commercial-off-the-shelf available parts where possible. An RGBD camera is mounted at the front of the robot, pointing vertically down and measuring the roughness of the terrain surface under the front legs. Force sensors are mounted at the tip of each leg and report the measured surface hardness.

For indoors experiments, position is measured using a 26-camera motion capture system from Qualisys with four reflective markers placed on the robot. Outdoors, we use a Ublox c94-m8p differential GPS, mounted on the chassis, with the RTK base station placed consistently within 300m of the robot. Both these systems achieve a sub-cm precision in position, which is adequate for high accuracy localisation.

Attached to the chassis are four legs, each with three rotational joints. The proximal joint consists of a Dynamixel MX-64 servo from Robotis, while the two distal joints use MX-106 servos. Two prismatic joints vary the lengths of the femur and tibia, using a geared DC motor and custom linear actuator as shown in Fig. 1b. Each femur can lengthen by 50mm, and each tibia by 100mm. The longest transition, from minimum to maximum length of the tibia, takes approximately 90s at a speed of 1mm/s.

Our adaptive morphology mechanism alters the available workspace, as seen in Supplementary Figure 2b. The longest available leg length increases the workspace volume by $\approx 75\%$, and lifts the body ≈ 13cm away from the ground, consequently affecting the robot's balance. Only 11% of the workspace for the shortest legs and 6% for the longest legs overlap. This shared area is too small for an effective gait, making it impossible for a robot without adaptive leg-lengths to replicate the behavior of our platform.

Cost of Transport (COT) provides a straightforward means of assessing energy efficiency when walking, and is a standard indicator of performance applied to various robots as well as to biological life [48]. The formula for COT is given in Equation 1, where E is the energy, m is the mass of the robot, g is standard gravity, and d is the distance traveled. Energy is, in our case, solely based on the energy expended for locomotion by the servos, measured by an onboard current sensor in each servo. Power used for control and sensing is assumed to be independent of the morphological configuration, and therefore not included.

$$\mathrm{COT} = \frac{E}{mgd} \tag{1}$$

[3]Design, code, and documentation available on https://github.com/dyret-robot

Gathering the baseline data set

A small data set of walking performance for the different morphological configurations on selected terrains needed to be generated. To limit the extent of the data collection, we tested 25 leg length combinations of femur (0, 12.5, 25, 37.5 and 50mm) and tibia (0, 20, 40, 60, and 80mm). A full list of morphologies are available in Supplementary Table 2. The terrain boxes shown in Fig. 4a were half-filled along their length with gravel, sand, or a flat sheet of fiber-reinforced concrete to an approximate depth of 15cm. These materials were selected to give a wide spread in the hardness and roughness measures we use, as seen in Supplementary Table 3.

Each evaluation consists of the robot walking forwards on a single surface at a velocity of about 2m/min for 15 seconds. All of this time was spent on the same surface. This was done five times for each morphology, starting on different parts of each surface, to account for local variations in the terrain material. Twenty-five morphologies walking on three different surface types for five iterations of 15 seconds of walking each yielded approximately 90 minutes of walking data in total, which was collected over two consecutive days.

Indoor adaptation experiment

Our goal was to test the simple adaptation method on previously seen homogeneous terrains. This serves as a precursor to continuous adaptation to unstructured terrains outside by evaluating the feasibility of our methods in a controlled environment inside. One of the boxes in Fig. 4a was used with the first half covered with the concrete sheet and the second half with gravel. We compare the adaptation method, detailed below, to walking across the whole box with each of the two optimal static morphologies from the baseline data set (concrete-specialized, with femur 50mm and tibia 20mm; and gravel-specialized, with both femur and tibia 0mm).

The robot is initially positioned to walk eight steps on the concrete, before stepping onto the gravel for the last eight steps. When using the two optimal static morphologies, the robot walks the full 16 steps without stopping. In the adaptive case, the robot uses terrain classification to detect the transition between the two terrain types. When a change has been detected, the robot stops walking and changes the length of its legs. The new morphology is taken from the best performers in the baseline data set. Once the desired leg length has been reached, it recommences walking the rest of the 16 steps. The tests were repeated 12 times to get an accurate representation of the actual performance. 12 iterations of 16 steps gives a total of 192 steps for each morphology. Walking across the box with the adaptive and two static morphologies gives a total of 576 steps for the indoor experiment.

Pseudo-code for the experiment is available in Supplementary Methods 1.

Outdoor adaptation experiment

Our goal was to test the extended adaptation method in previously unseen heterogeneous terrains, which is the key experiment for this paper. This experiment was done on the outside test track shown in Fig. 4c, which consists of mixed terrain, dominated by grass and a concrete road. The grass is dry, with varying coverage, while the road includes cracks and small obstacles like rocks and sticks. No attempt was made to clean up or prepare the outdoor environments in any way. The adaptation method is compared to the best all-performing static morphology from the baseline data set (femur 37.5mm, tibia 20mm).

When adapting, the robot is initially positioned on the grass section with its most conservative morphology (femur 0mm, tibia 0mm). With a target forward velocity of about 2m/min, the robot is manually steered onto the grass, onto the road, then back on the grass. The adaptation algorithm is given the chance to change its morphology 32 times on each terrain section, referred to as *evaluations*. The static morphology evaluates the same morphology 64 times on each section without any reconfiguration, resulting in approximately the same time spent walking on each terrain section for the two approaches. This ensures similar battery conditions to not skew the results. The adaptation algorithm was tested five times since results can vary based on local variations in the terrain.

When choosing the next morphology to evaluate, the robot only considered neighboring morphologies (morphologies that only require changing the leg segment lengths by a single increment: 12.5mm for femur and 20mm for tibia). It decides on the morphology based on the model initialized with the baseline dataset, as well as the features of the current terrain. One evaluation comprises 3 full steps per leg, during which performance and terrain features are measured. This ensures accurate measurements in the noisy outdoor environment.

Pseudo-code for the experiment is available in Supplementary Methods 1.

Baseline modeling

Evaluating how each morphology performs in the real world can take a long time, so testing all possible leg lengths each time the terrain changes is impossible. A baseline model allows the robot to efficiently adapt its body during operation by providing some predictive knowledge of which morphologies might perform well. This model takes the terrain characteristics as input and predicts the performance for all possible morphologies as output. In our case, we limit the number of morphologies to 25 and treat them all independently when learning the model.

The whole model is a collection of 25 sub-models, one for each of the leg-length pairs in the data set. A diagram can be seen in Supplementary Fig. 1. Each sub-model is made using second-order polynomial regression to approximate the relationship between the two terrain characteristics and the energy efficiency of each leg-length pair. Each time the robot tests a new morphology, the result is added to the model to incorporate new

knowledge continuously. When a new terrain is encountered, the corresponding point from each of the 25 sub-models is used to generate a full predicted map for the given terrain. Examples of generated maps can be seen in Fig. 6, where they were made from data at different points throughout the adaptation algorithm.

Selection of the next morphology can be done globally in the generated map, but since changing the length of the legs can require a lot of time, we always select from neighboring morphologies in the model, where the new morphology is the neighbor with the best predicted efficiency (lowest predicted COT), given the measured terrain features. To prioritise more current information, the actual current performance is used when comparing to new morphologies, and not the theoretical prediction for the current morphology from the map.

Terrain sensing

Our system use two different methods to sense its terrain – classification, and characterization. In classification, the goal is to find out which class the perceived terrain belongs to, out of a few number of example terrains. In our indoor experiments, we only have three classes: concrete, sand, and gravel. In characterization, the goal is instead to measure some features of the terrain that are useful for the adaptation process. The perceived terrain is not classified as being a specific type but is given a set of quantitative measurements that describe it. Terrains may be characterised in a multitude of ways, here we use harness and roughness as they strongly inform the morphology and are easily sensed [49].

Hardness is calculated from the force sensors in the front feet (Optoforce OMD-20-SH-80N) at 100Hz. Back foot sensors are ignored to reduce ambiguities arising from walking between two terrain types (front feet on terrain A, back feet on terrain B). Raw sensor data is run through a median filter of size 5 for noise reduction, along with the removal of obvious erroneous force measurements of several times the weight of the complete robot (>100N). The final hardness value reported is the maximum value seen from both sensors in a six-second sliding window from the start of the measurement. This means that increases in hardness are immediately represented in the hardness feature, while reductions will take some time to propagate through the system.

Roughness is calculated using the point-cloud from an Intel Realsense D435 RGBD-camera, mounted at the front of the robot and pointing down, providing a point-cloud at 6hz. We fit a ground plane to the point-cloud and discard all points that have a distance to the plane of more than 35mm to filter out the legs and other robot parts. The measurement of roughness is the mean of the square distances from each point to the plane, where a roughness of zero implies a perfectly flat plane.

These methods for extracting terrain features are both relatively simple, but have been considered adequate for our needs and are well-studied in the literature. They were evaluated on a number of terrains inside and outside during development, and corresponded well to perceived terrain features by the researchers. Characterization, used in

the 'previously unseen, heterogeneous terrains' experiment, is via these two features directly. Classification, used in the 'previously seen homogeneous terrain' experiment, is done by calculating the euclidean distance to the mean value for each terrain group in the data set, and selecting the closest.

Robot control

The robot uses a high-level spline-based gait controller detailed in [50]. It describes a continuous, regular crawl gait where the body moves at a static speed, and only one leg is lifted at a time. Leg trajectories are identical and represented by a looping cubic Hermite spline. A balancing counter-movement is added to each step where the robot leans to the opposite side of the currently lifted leg. This allows statically stable gaits and is needed since each leg weighs approximately the same as the central body. Details on all parameters for the gait controller can be found in Supplementary Table 1.

Step height is a fundamental parameter, set to 100mm based on the terrains the robot will be operating in and a safety margin. Step length can then be maximized within the workspace of the legs (seen in Supplementary Figure 2b), and step frequency is calculated to keep the rotational velocity of the servos within a safe range. We also have to calculate a scaling factor given our variable leg lengths, as increasing the length of the legs allows longer steps with the same rotational speed for each joint. The step height was kept constant, but step length and width was scaled to keep the RPM of the servos consistent throughout the ground movement. The center of the trajectory spline was kept constant not to affect the balance of the robot. The scaling can be found in Supplementary Table 2.

Acknowledgements

This work was partially supported by The Research Council of Norway under grant agreement 240862 and its Centres of Excellence scheme, project number 262762.

Data Availability

All JSON-formatted measurements from robot evaluations in this study have been deposited in Figshare[4]. All other relevant data are available from the corresponding author on request.

[4]Available at https://figshare.com/s/d02297eec829271d26f7

Code Availability

The code for running the robot is available on the DyRET Github page[5], and is under a GPL v3 license. The code to setup the experiments and log the results is available is a separate Github repository[6].

Author Contributions

Author contributions classified in the CRediT taxonomy [51]:

T.F.N: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data Curation, Writing - Original Draft, Writing - Review and Editing, Visualization.
K.G: Conceptualization, Methodology,Writing - Review and Editing, Supervision.
C.P.M: Conceptualization, Methodology, Formal analysis, Writing – Review and editing, Supervision.
J.T: Writing – Review and Editing, Supervision, Funding acquisition.
D.H: Conceptualization, Methodology, Formal analysis, Resources, Writing – Review and Editing, Visualization, Supervision, Project administration, Funding Acquisition.

Competing Interests statement

The authors declare no competing interests.

References

[1] Kawatsuma, S., Fukushima, M. & Okada, T. Emergency response by robots to Fukushima-Daiichi accident: summary and lessons learned. *Industrial Robot: An International Journal* (2012).

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

[3] Wilson, A. D. & Golonka, S. Embodied cognition is not what you think it is. *Frontiers in psychology* **4**, 58 (2013).

[4] Zhang, T., Zhang, W. & Gupta, M. M. Resilient robots: Concept, review, and future directions. *Robotics* **6**, 22 (2017).

[5]Available at https://github.com/dyret-robot/dyret_documentation
[6]Available at https://github.com/tonnesfn/tonnesfn_experiments

[5] Picardi, G., Hauser, H., Laschi, C. & Calisti, M. Morphologically induced stability on an underwater legged robot with a deformable body. *The International Journal of Robotics Research* (2019).

[6] Baines, R., Freeman, S., Fish, F. & Kramer, R. Variable stiffness morphing limb for amphibious legged robots inspired by chelonian environmental adaptations. *Bioinspiration & Biomimetics* (2020).

[7] Paik, J. K., Byoungkwon, A., Rus, D. & Wood, R. J. Robotic origamis: Self-morphing modular robot. In *Proceedings of the 2nd International Conference on Morphological Computation* (2012).

[8] Nygaard, T. F., Martin, C. P., Samuelsen, E., Torresen, J. & Glette, K. Real-world evolution adapts robot morphology and control to hardware limitations. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '18, 125–132 (ACM, New York, NY, USA, 2018). URL `http://doi.acm.org/10.1145/3205455.3205567`.

[9] Nygaard, T. F., Martin, C. P., Howard, D., Torresen, J. & Glette, K. Environmental Adaptation of Robot Morphology and Control through Real-world Evolution. *arXiv e-prints* arXiv:2003.13254 (2020). `2003.13254`.

[10] Heijnen, H., Howard, D. & Kottege, N. A testbed that evolves hexapod controllers in hardware. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 1065–1071 (IEEE, 2017).

[11] Gong, D., Yan, J. & Zuo, G. A review of gait optimization based on evolutionary computation. *Applied Computational Intelligence and Soft Computing* **2010** (2010).

[12] Ha, S., Xu, P., Tan, Z., Levine, S. & Tan, J. Learning to walk in the real world with minimal human effort. *arXiv preprint arXiv:2002.08550* (2020).

[13] Kober, J., Bagnell, J. A. & Peters, J. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* **32**, 1238–1274 (2013).

[14] Calandra, R., Seyfarth, A., Peters, J. & Deisenroth, M. P. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence* **76**, 5–23 (2016).

[15] Rodriguez, D., Brandenburger, A. & Behnke, S. Combining simulations and real-robot experiments for bayesian optimization of bipedal gait stabilization. In *Robot World Cup*, 70–82 (Springer, 2018).

[16] Bellicoso, C. D., Gehring, C., Hwangbo, J., Fankhauser, P. & Hutter, M. Perception-less terrain adaptation through whole body control and hierarchical optimization. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 558–564 (IEEE, 2016).

[17] Hwangbo, J. *et al.* Learning agile and dynamic motor skills for legged robots. *Science Robotics* **4**, eaau5872 (2019).

[18] Kalakrishnan, M., Buchli, J., Pastor, P., Mistry, M. & Schaal, S. Learning, planning, and control for quadruped locomotion over challenging terrain. *The International Journal of Robotics Research* **30**, 236–258 (2011).

[19] Kaushik, R., Anne, T. & Mouret, J.-B. Fast online adaptation in robotics through meta-learning embeddings of simulated priors. *arXiv preprint arXiv:2003.04663* (2020).

[20] Fahmi, S. *et al.* Stance: Locomotion adaptation over soft terrain. *IEEE Transactions on Robotics* **36**, 443–457 (2020).

[21] Buchanan, R. *et al.* Walking posture adaptation for legged robot navigation in confined spaces. *IEEE Robotics and Automation Letters* **4**, 2148–2155 (2019).

[22] Sims, K. Evolving 3d morphology and behavior by competition. *Artificial life* **1**, 353–372 (1994).

[23] Long, J. *Darwin's devices: What evolving robots can teach us about the history of life and the future of technology* (Basic Books, 2012).

[24] Spielberg, A. *et al.* Learning-in-the-loop optimization: End-to-end control and co-design of soft robots through learned deep latent representations. In *Advances in Neural Information Processing Systems*, 8282–8292 (2019).

[25] Mouret, J.-B. & Chatzilygeroudis, K. 20 years of reality gap: a few thoughts about simulators in evolutionary robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1121–1124 (2017).

[26] Kriegman, S. *et al.* Automated shapeshifting for function recovery in damaged robots. *arXiv preprint arXiv:1905.09264* (2019).

[27] Marbach, D. & Ijspeert, A. J. Online optimization of modular robot locomotion. In *IEEE International Conference Mechatronics and Automation, 2005*, vol. 1, 248–253 (IEEE, 2005).

[28] Passault, G., Rouxel, Q., Fabre, R., N'Guyen, S. & Ly, O. Optimizing morphology and locomotion on a corpus of parametric legged robots. In *Conference on Biomimetic and Biohybrid Systems*, 227–238 (Springer, 2016).

[29] Lipson, H. & Pollack, J. B. Automatic design and manufacture of robotic lifeforms. *Nature* **406**, 974–978 (2000).

[30] Hornby, G. S., Lipson, H. & Pollack, J. B. Generative representations for the automated design of modular physical robots. *IEEE transactions on Robotics and Automation* **19**, 703–719 (2003).

[31] Ha, S., Coros, S., Alspach, A., Kim, J. & Yamane, K. Joint optimization of robot design and motion parameters using the implicit function theorem. In *Robotics: Science and systems* (2017).

[32] Kriegman, S. *et al.* Scalable sim-to-real transfer of soft robot designs. *arXiv preprint arXiv:1911.10290* (2019).

[33] Collins, J., Geles, W., Howard, D. & Maire, F. Towards the targeted environment-specific evolution of robot components. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 61–68 (2018).

[34] Jakobi, N., Husbands, P. & Harvey, I. Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, 704–720 (Springer, 1995).

[35] Erez, T., Tassa, Y. & Todorov, E. Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ode and physx. In *2015 IEEE international conference on robotics and automation (ICRA)*, 4397–4404 (IEEE, 2015).

[36] Sun, Y., Chen, X., Yan, T. & Jia, W. Modules design of a reconfigurable multi-legged walking robot. In *2006 IEEE International Conference on Robotics and Biomimetics*, 1444–1449 (IEEE, 2006).

[37] Guan, Y., Jiang, L., Zhangy, X., Zhang, H. & Zhou, X. Development of novel robots with modular methodology. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2385–2390 (IEEE, 2009).

[38] Vujovic, V., Rosendo, A., Brodbeck, L. & Iida, F. Evolutionary developmental robotics: Improving morphology and control of physical robots. *Artificial Life* (2017).

[39] Moreno, R. *et al.* Automated reconfiguration of modular robots using robot manipulators. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, 884–891 (IEEE, 2018).

[40] Zhakypov, Z. & Paik, J. Design methodology for constructing multimaterial origami robots and machines. *IEEE Transactions on Robotics* **34**, 151–165 (2018).

[41] Riviere, V., Manecy, A. & Viollet, S. Agile robotic fliers: A morphing-based approach. *soft robotics* **5**, 541–553 (2018).

[42] Bucki, N. & Mueller, M. W. Design and control of a passively morphing quadcopter. In *2019 International Conference on Robotics and Automation (ICRA)*, 9116–9122 (IEEE, 2019).

[43] Geilinger, M., Poranne, R., Desai, R., Thomaszewski, B. & Coros, S. Skaterbots: Optimization-based design and motion synthesis for robotic creatures with legs and wheels. *ACM Transactions on Graphics (TOG)* **37**, 1–12 (2018).

[44] Meiri, N. & Zarrouk, D. Flying star, a hybrid crawling and flying sprawl tuned robot. In *2019 International Conference on Robotics and Automation (ICRA)*, 5302–5308 (IEEE, 2019).

[45] Kriegman, S., Blackiston, D., Levin, M. & Bongard, J. A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences* (2020).

[46] Ritter, A. *Shape-Changing Smart Materials*, 46–71 (Birkhäuser Basel, Basel, 2007).

[47] Nygaard, T. F., Martin, C. P., Torresen, J. & Glette, K. Self-modifying morphology experiments with dyret: Dynamic robot for embodied testing. In *2019 IEEE International Conference on Robotics and Automation (ICRA)* (2019).

[48] Seok, S. *et al.* Design principles for energy-efficient legged locomotion and implementation on the mit cheetah robot. *Ieee/asme transactions on mechatronics* **20**, 1117–1129 (2014).

[49] Howard, A. & Seraji, H. Vision-based terrain characterization and traversability assessment. *Journal of Robotic Systems* **18**, 577–587 (2001).

[50] Nygaard, T. F., Martin, C. P., Torresen, J. & Glette, K. Evolving robots on easy mode: Towards a variable complexity controller for quadrupeds. In *Applications of Evolutionary Computation* (Springer International Publishing, 2019).

[51] Allen, L., O'Connell, A. & Kiermer, V. How can we ensure visibility and diversity in research contributions? how the contributor role taxonomy (credit) is helping the shift from authorship to contributorship. *Learned Publishing* **32**, 71–74 (2019). URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/leap.1210`. `https://onlinelibrary.wiley.com/doi/pdf/10.1002/leap.1210`.

Supplementary Materials

Supplementary Figure 1. Prediction model diagram.
Supplementary Figure 2. Reconfigurable mechanism and workspace difference.
Supplementary Methods 1. Pseudo-code for adaptation algorithms.
Supplementary Table 1. Parameters for the gait controller.
Supplementary Table 2. Morphologies and spline scaling.
Supplementary Table 3. Terrain characteristics for the terrain boxes.
Supplementary Video 1. Summary of the results and the method.

Supplementary Figures

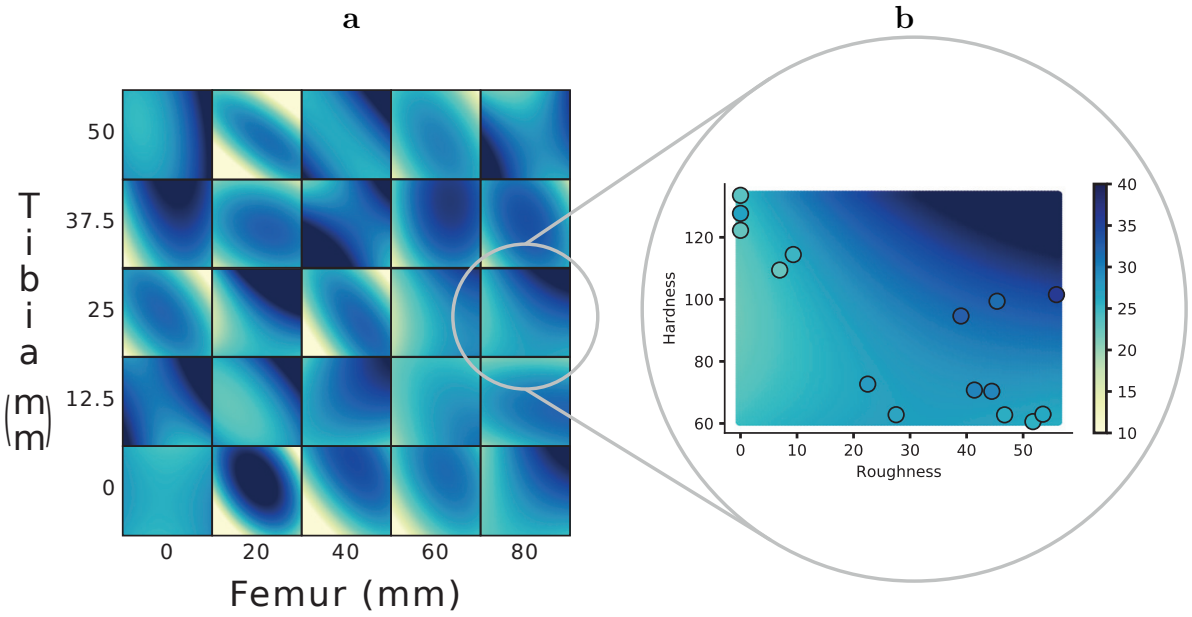

Supplementary Figure 1: **A diagram of the prediction model used, instantiated with the baseline data set measurements.** (**a**) The full model with the 25 sub-models for each leg length combination.) (**b**) One of the sub-models, which shows the predicted energy efficiency (COT) of that leg-length combination for different terrains. The dots are actual measurements from the baseline experiment.

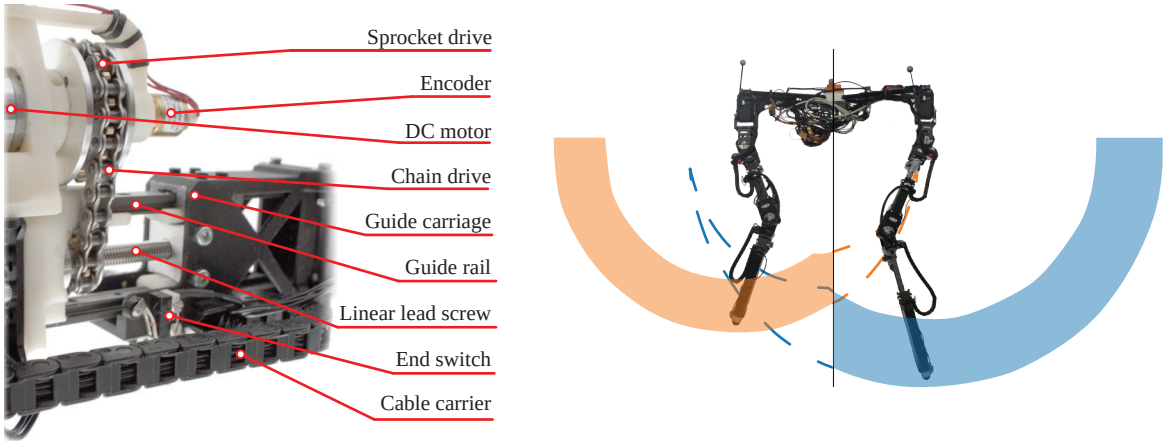

Supplementary Figure 2: **Reconfigurable mechanism and workspace difference.** (**a**) The components of the mechanical adaptation mechanism. (**b**) An indication of the difference in workspace for the two extreme leg lengths.

Supplementary Methods

Experiment 1: Adapting to previously seen homogeneous terrain in indoor conditions - concrete to gravel

```
Start with concrete-specialized morphology (femur 50mm, tibia 20mm)
repeat
    Take one step/leg forward, measuring terrain characteristics
    if Non-optimal terrain for current morphology detected then
        Stop walking
        Reconfigure to optimal morphology for detected terrain
    end
until 16 steps/leg has been walked in total;
```

The robot is initially positioned so that it will take 8 steps on the concrete, before stepping onto the gravel for the last 8 steps.

Experiment 2: Adapting to previously unseen heterogeneous terrain in outdoor conditions - grass to road to grass

```
Start with initial morphology (femur 0, tibia 0)
Walk for three steps/leg, measuring terrain characteristics
repeat
    Generate predicted map for current terrain from model
    if best predicted neighbor COT > current COT then
        Start changing morphology to best performing neighbor
    end
    repeat
        Take one step/leg
    until new leg lengths are achieved;
    Walk for three steps/leg, measuring terrain and energy efficiency
    Add measured terrain characteristics and COT to data set
    Regenerate model with newly experienced data point
until 96 morphologies tested;
```

The robot is initially positioned on the grass, before walking onto road, then back on grass. It is manually led onto the next terrain type after 32 morphologies have been tested on each terrain section.

Supplementary Tables

Supplementary Table 1: Parameters for the gait controller. * These parameters are linearly scaled as morphology changes, see Supplementary Table 2 for details.

frequency	0.2
lift duration	0.15
p0_x	0.0
p0_y	50.0
p1_x	0.0
p1_y	-80.0
p2_x	0.0
p2_y*	50.0
p2_z*	50.0
p3_x	0.0
p3_y*	-15.0
p3_z*	100.0
p4_x	0.0
p4_y*	-80.0
p4_z*	50.0
wagPhase	0.05
wagAmplitude_x	25.0
wagAmplitude_y	75.0

Supplementary Table 2: Morphologies and spline scaling for all leg length combinations used in our experiments. Optimal and best trade-off morphologies were found while generating our baseline data set.

ID	Femur (mm)	Tibia (mm)	Total	Scaling	Comment
0	0	0	0	100%	Optimal on gravel
1	0	20	20	103%	
2	0	40	40	106%	
3	0	60	60	109%	
4	0	80	80	112%	
5	12.5	0	12.5	103%	
6	12.5	20	32.5	106%	
7	12.5	40	52.5	109%	
8	12.5	60	72.5	113%	
9	12.5	80	92.5	116%	
10	25	0	25	106%	
11	25	20	45	109%	
12	25	40	65	113%	
13	25	60	85	116%	
14	25	80	105	119%	
15	37.5	0	37.5	109%	
16	37.5	20	57.5	113%	Best trade-off for all surfaces
17	37.5	40	77.5	116%	
18	37.5	60	97.5	119%	
19	37.5	80	117.	122%	
20	50	0	50	113%	Optimal on sand
21	50	20	70	116%	Optimal on concrete
22	50	40	90	119%	
23	50	60	110	122%	
24	50	80	130	125%	

Supplementary Table 3: Terrain characteristics of the three materials present in the indoor terrain boxes, shown with median and interquartile range.

	Roughness		Hardness	
	Median	IQR	Median	IQR
Concrete	5.1	7.6	135.5	13.3
Sand	25.2	18.5	61.3	8.2
Gravel	35.7	12.8	85.7	24.9