

# Real-world Embodied AI Through a Morphologically Adaptive Quadruped Robot

Tønnes F. Nygaard<sup>1,2\*</sup>, Charles P. Martin<sup>3</sup>,  
Jim Torresen<sup>1,4</sup>, Kyrre Glette<sup>1,4,6</sup>, David Howard<sup>5,6</sup>

<sup>1</sup>Department of Informatics, University of Oslo, Norway

<sup>2</sup>The Norwegian Defence Research Establishment (FFI), Kjeller, Norway

<sup>3</sup>Research School of Computer Science, Australian National University, ACT, Australia

<sup>4</sup>RITMO, University of Oslo, Norway

<sup>5</sup>Cyber-Physical Systems Program, CSIRO, QLD, Australia

<sup>6</sup>These authors share senior authorship

\*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 morphologies based on the currently-sensed terrain. First, we build a model that describes how the robot morphology affects performance on selected terrains. We then test continuous adaptation in realistic outdoor terrain 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. These challenges, chiefly multimodality and unpredictability, are characteristic of the type of unstructured environment that robotic systems as a whole continue to struggle with. In Fukushima, technological limitations meant that the eventual solution required numerous highly specialized traditional robots, with correspondingly high numbers of deployments and extended mission times [1]. Shape-shifting, or morphological adaptation, presents a more attractive, albeit more technically challenging, solution that is capable of achieving more complex mission outcomes. Able to vary both body and controller, a morphologically adaptive 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. The underlying principle is that variable morphology provides additional degrees of freedom to more strongly tie a robot’s behaviour to its immediate environment for improved mission performance, 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 to adapt to a broad swathe of unpredictable environments and tasks on the fly, without having to be redesigned and rebuilt each time they face something 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 [2] to elegant origami-inspired morphing structures [3].

We postulate that the key to developing such flexible, adaptable robots may lie in a specific subfield of Machine Intelligence called Embodied Artificial Intelligence [4]. Embodied AI is a subfield of Embodied Cognition, which 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 [5]. Viewed through the lens of Embodied Cognition, the physical manifestation of a robot is a crucial adaptation tool, which could be vital in achieving resilient robots that can operate across challenging real-world environments [6]. Indeed, in some cases, changing the robot’s morphology might be the only viable option to elicit suitable in-environment behaviors [7].

In this paper, we present our own morphologically adaptive robot for unstructured environments, seen in Fig. 1a. The quadruped robot, DyRET (the Dynamic Robot for Embodied Testing), provides a powerful proof of concept harnessing a variable morphology to adapt to realistic real-world conditions in outdoor settings [8, 9]. Morphological adaptation is provided through variable-length legs, whereby the length of both femur and tibia can be adjusted on the fly to enable different walking behaviours whilst also tilting the central body (Fig. 1b). A novel terrain-adaptation algorithm tailors morphol-

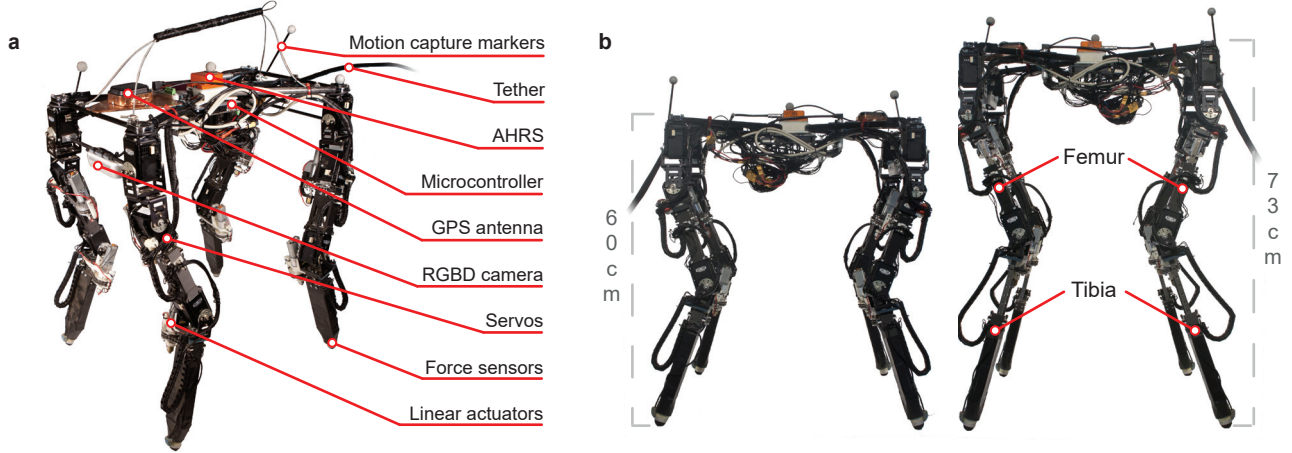


Fig. 1: **The morphologically adaptive robot used in this study** (a) An overview of the main components of the robot. (b) The robot with the shortest (left) and longest (right) leg configuration.

ogy to the current terrain. Bootstrapped with knowledge from controlled experimentation in indoor terrain boxes, it alters the morphological configuration of the robot online 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 and 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 optimization [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]. While a robotic adaptation process would ideally be embodied [23], most works in adaptive robotic morphology are carried out in physics simulation, and not on physical robots [24]. Examples include soft robots [25], modular robots [26], and legged

robots [27]. Recent work has studied co-optimization of control and morphology using gradient approaches in differentiable physics simulators [28], which has strong potential to efficiently couple control, morphology, and environment. These efforts are, so far, only simulated and run in simple environments.

The next step up from pure simulation of adaptive morphology is selecting a few virtual robots for real-world manufacture and testing. The manufacturing process can involve 3D printing [29, 30, 31], assembly from building blocks [32, 33], or even silicone mold techniques for soft robots [34]. 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 [35, 36]. 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 morphological 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 [37, 38, 39] or an external mechanism for reconfiguration [40, 41, 42]. Such approaches require extensive time, external apparatus, or human intervention and are not suitable for continuous adaptation during independent operation. Few examples exist of robots with a built-in ability to morphologically adapt, mainly due to the challenges associated with designing, building and maintaining a robot with a complex dynamic morphology [43]. Many of these robots are therefore relatively small with no payload capacity, and are limited in their ability to function in real-world unstructured settings, e.g., [44]. More complicated robots possess a higher potential to solve real-world problems, including morphing drones [45, 46], multi-modal legged-wheeled [47] and wheeled-flying robots [48]. 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 demonstrate this ability, 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 in realistic 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

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 in a controlled indoor environment, and (ii) Adapting in a realistic outdoor environment. 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. The baseline model bootstraps 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, as seen in Fig. 4a. Our robot senses both the 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. Please refer to the Methods section for our definitions of roughness and hardness. Each box consists of two halves filled with different terrain materials. This allows the robot to walk on the separate terrains, as well allowing us to test 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 forward for 15 seconds per morphology, covering all 25 combinations. The velocity varies greatly, but the theoretical speed is 1.56m/min for the shortest legged robot, and 1.95m/min for the longest legged robot. The walking surface typically reduces the speed, but it can in some cases also increase it due to the complex effects of the dynamics of the mechanical system. Each morphology is tested at 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 pure walking data.

The measured cost of transport (COT) for each morphology on the three surfaces can be seen in Fig. 2. When walking on the concrete, 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, 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, but the best COT is seen for the shortest possible legs.

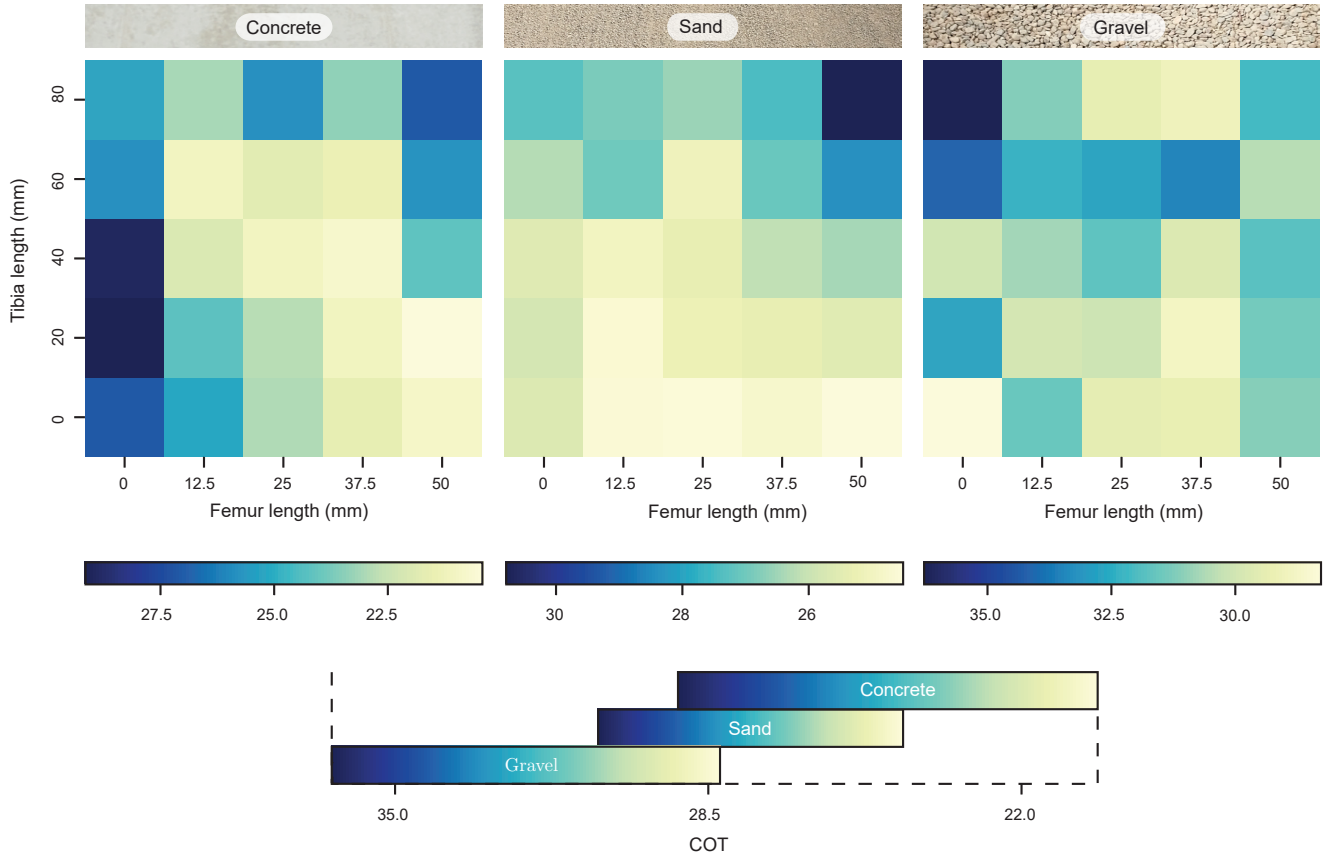


Fig. 2: **The COT of different leg lengths on the three terrains** The lower the COT is (yellow), the more efficiently the robot is walking. Please note that the ranges – as seen in the bottom of the figure – are not the same for each surface, to better highlight local differences within each terrain.

### Adapting in a controlled indoor environment

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 highly uniform throughout. As the terrains are known, the adaptation algorithm takes the form of a classifier (see the Methods section for details). 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. This serves as a simplified validation of our final, continuous, adaptation method done in a realistic outdoor environment.

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




Fig. 3: **Diagram showing the two adaptation methods used** (a) Adaptation principle for the controlled indoor environment. 1<sub>a</sub>: The robot walks forward while sensing its environment. 2<sub>a</sub>: Once a terrain change has been detected, it stops walking. 3<sub>a</sub>: The robot changes the length of its legs to the optimal morphology for the new terrain. 4<sub>a</sub>: It starts walking with the new morphology, repeating from 1<sub>a</sub>. (b) Adaptation principle for realistic outdoor terrains. 1<sub>b</sub>: The robot predicts the best performing morphology based on sensor readings and its internal model. 2<sub>b</sub>: It changes the length of its legs to this new morphology while walking. 3<sub>b</sub>: When the legs have reached their goal length, the robot measures its performance and the terrain characteristics. 4<sub>b</sub>: It adds the new measurements to its internal model, before repeating the process from step 1<sub>b</sub>.

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 the other, 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 triggers a switch in morphology at the appropriate time.

### Adapting in a realistic outdoor environment

Fig. 3b describes an extended method that takes into consideration the additional challenges of a realistic outdoor environment. As these terrains are highly dynamic and unstructured, we cannot assume that all terrains that the robot will encounter are present in the baseline data set. We therefore replace terrain classification (used in our previous approach) with characterization (details in the Methods section). Realistic outdoor terrains can also change substantially over even small areas, potentially 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 to 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 its 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 as 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

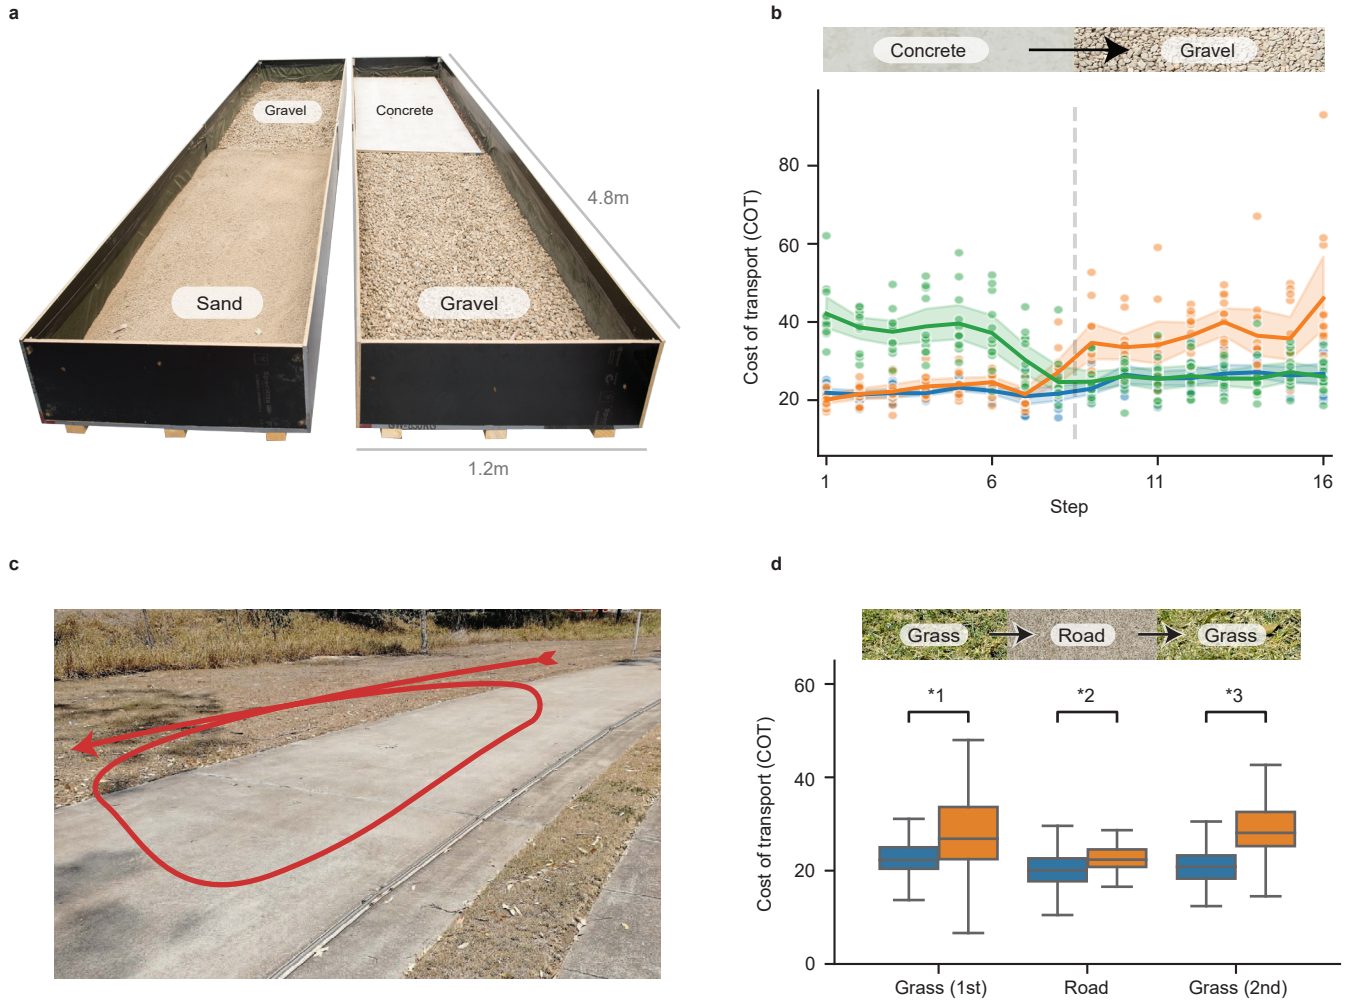


Fig. 4: **Environments and results for indoor and outdoor experiments** (a) The terrain boxes used for the experiment in the controlled indoor environment. They contain sand, gravel and concrete. (b) The COT for the two best static morphologies (gravel-specialized in green and concrete-specialized in orange) and the adaptation (blue) when walking on concrete, then gravel in the boxes. The solid lines show the mean of 12 repeats, and the shaded areas show the 95% confidence interval. The vertical gray dashed line denotes the approximate point where the middle of the robot crosses between the two terrain types. The front legs step onto the new terrain a few steps before. (c) The outdoor area used in the experiment in a realistic outdoor environment, with the red line showing a typical walking path for the robot. It starts on grass, walks onto the road, then back onto the grass. (d) Standard boxplot of energy efficiency (COT) of the adapting (blue) and best all-rounder static morphology (orange) on the outside test track. The boxes extend from the lower to upper quartile, with the median shown within with a black line. The whiskers on both sides of the boxes show the extreme values. \*Statistically significant differences from two-sided Mann-Whitney U test on each parameter with Holm-Bonferroni p-value correction ( $p < 0.01$ ,  $n_1 = 160$  (adaptive),  $n_2 = 64$  (static all-rounder), \*1:  $U = 2540$ , \*2:  $U = 747$ , \*3:  $U = 2001$ ).

starts very high at  $\approx 25$ , falls below 6 after trying 16 different morphologies, before ending at  $\approx 2.5$  at the end of the first grass section. The error spikes to  $\approx 8$  as the robot steps onto the road and ends up at  $\approx 4$  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 eight more evaluations before ending at an error of  $\approx 5$ . 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 used 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 morphologies tested on the final grass section are similar to those used on the road. The adaptation algorithm exploits almost the entire morphological range. The algorithm 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’s understanding of the two outdoor terrains changes during the adaptation runs. The mean terrain characteristics from the outside test track (roughness  $51.0\text{mm}^2$  and hardness  $143.4\text{N}$  for grass; roughness  $16.5\text{mm}^2$  and hardness  $187.4\text{N}$  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 transitioning 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  $12.5\text{mm}$ , with tibia lengths  $60\text{mm}$  and  $80\text{mm}$ ) were updated both in the first grass section and the last. The fact that not more morphologies were updated shows that the adaptation algorithm has successfully integrated the experience from new terrains with the baseline data set 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  $50\text{mm}$ , tibia  $60\text{mm}$ ). Because of the time taken to transition between morphologies, the adaptation algorithm is limited to

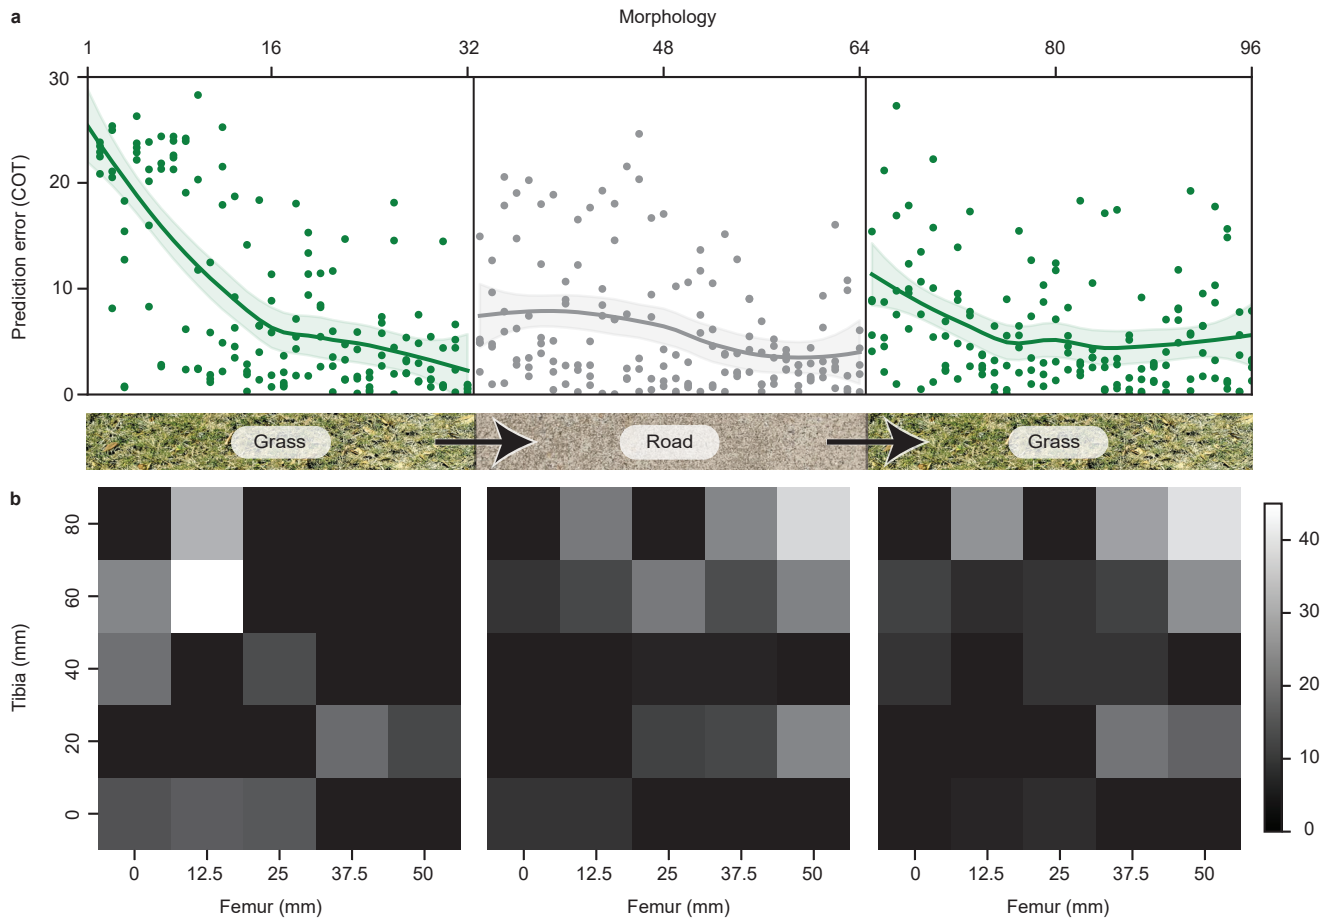


Fig. 5: **Analysis of the adaptation algorithm** (a) The prediction error for each evaluated morphology as the robot walks on grass, road, then back on grass. The plot shows a locally-weighted regression with unweighted fit, while the shaded areas shows the 95% confidence interval of the regression. (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.

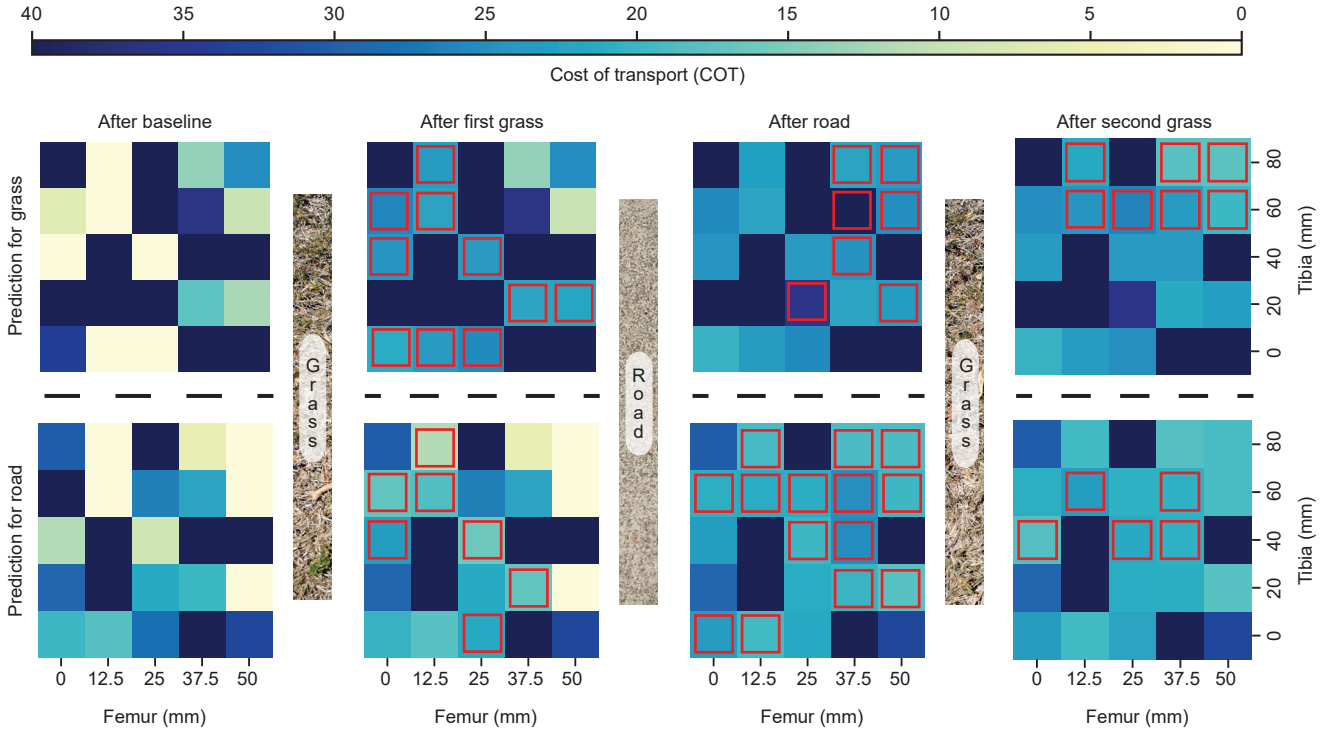


Fig. 6: **Model transformation during adaptation** The first row shows the predicted energy efficiency (COT) map of all leg lengths for the grass surface, while the second shows for the road surface. The first column shows the initial models based solely on the baseline data set. The second shows the 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. To indicate where the model was significantly updated, all visited cells resulting in a change in COT higher than 1 are marked with a red square.

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.



## 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 in a controlled indoor environment (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 an advantageous trait for robots operating in unstructured terrains.

In contrast to the relatively controllable indoor experiments, the realistic outdoor environment exposed our system to many sources of noise and inaccuracy. Outdoor experimentation also precluded the use of motion capture equipment for position tracking. To compensate for these difficulties, we used separate regression models for each morphology to keep noisy measurements contained. Each of these sub-modules are continually updated with new data, increasing the accuracy of the model with each new evaluation. Overly optimistic measurements would likely be selected for a new evaluation and subsequently be adjusted to a more realistic level. Overly pessimistic measurements can quickly be overlooked in favour of neighbouring morphologies, leading to premature convergence to local optima. This is addressed by only using the previously recorded COT when comparing to the predicted performance of the neighbours, enabling a single measurement to force exploring an otherwise ignored neighbour. Noisy terrain measurements also aid the algorithm, as walking outside exposes the robot to a range of different terrain examples very quickly, and the model will therefore gradually be filled out until it has the coverage needed for realistic prediction for its entire operating environment. This ability to exploit noise and inaccurate measurements is a strength of our adaptation mechanism allowing it to cope with complex and changing outdoor environments.

There are a few main limitations to our 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 [2, 49, 50]. To balance reconfigurability with mechanical simplicity and stability, our adaptation actuator mechanism has a speed of  $\approx 1\text{mm/s}$ . 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 constructing one model that incorporated 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 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 [51]. The body weights approximately 5kg, measures 50cm by 30cm, and stands between 60cm and 73cm tall, depending on the pose and leg length. The robot is powered by a three cell 5Ah LiPo battery supplying the robot with an unregulated voltage of 12.6V when fully charged. It has previously been used in laboratory settings, e.g., [9, 52].

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 perceived surface hardness of the robot.

For indoors experiments, position is measured using a 26-camera motion capture system from Qualisys with four reflective markers placed on the robot. This achieves a sub-cm precision. 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. The outdoor system typically achieves a precision of less than two cm, which is considered adequate for accurate speed estimation.

Four legs are attached to the chassis, 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 Extended Data Fig. 2a. 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 about 1mm/s. Increasing the length of the legs acts as a mechanical gearing of the servos. Longer legs mean higher movement speed at the end of the leg given the same rotational velocity at the joint. It does, however, come at the cost of less force at the end of the leg, given the same torque at the joint. Being able to change the length of its legs during operation allows the robot to select this trade-off between speed and stability as needed.

Our adaptive morphology mechanism alters the available workspace, as seen in Extended Data Fig. 2b. The longest available leg length increases the workspace volume by  $\approx 75\%$ , and lifts the body  $\approx 13\text{cm}$  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 static robot without adaptive leg-lengths to replicate the behavior of our platform.

Cost of Transport (COT) provides a straightforward and informative means of assessing energy efficiency when walking, taking into account key metrics including energy expended as well as mass and distance travelled. A dimensionless parameter, COT allows for ready comparison to other robots as well as to biological life [53], and is specifically very popular in the legged robotics community, e.g., [54]. 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.

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

## Terrain sensing

Our system uses 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 measure-

ments that describe it. Terrains may be characterised in a multitude of ways, here we use hardness and roughness as they strongly inform the morphology.

Hardness and roughness are often sensed indirectly, e.g. using vision [55]. Precise definitions also vary slightly across the literature. Here, we evaluate perceived hardness by measuring the force of the impact as the front feet hit the ground. Roughness is evaluated by measuring the deviation from a perfectly flat ground plane for a number of points in front of the robot. Our methods for terrain measurement are informed by our available sensing hardware.

Hardness is inferred from the impact force measured by the sensors in the front feet (Optoforce OMD-20-SH-80N) at 100Hz. The back foot sensors are ignored to reduce ambiguities arising from crossing 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 ( $>100\text{N}$ ). The final hardness value reported is the summed maximum value measured on each of the force axes on both sensors in a six-second sliding window from the start of the measurement. Details are available in the Supplementary Materials. Only looking at the maximum from a sliding window means that increases in hardness are immediately represented in the hardness estimate, while reductions will take some time to propagate through the system. There is a slight dependency on environmental factors like the friction of the surface, but these are considered minor contributors compared to the high level of noise in the measurements and natural variance in the terrains the robot operates in.

Roughness is inferred using the point-cloud from an Intel Realsense D435 RGBD-camera mounted at the front of the robot and pointing down, providing a 3D representation of the ground at 6hz. A ground plane is fit to the measured points and all points that have a distance to the plane of more than 35mm are discarded to filter out the legs and other parts of the robot from the scene. The mean of the square distances from each point to the plane is used as a roughness estimate, where zero would imply walking on a perfectly flat surface.

These methods for extracting terrain features are both relatively simple, but have been considered adequate for our needs. Details on the sensors and measurement ranges are available in Supplementary Table 4, and a plot showing the distribution of terrain measurements for the indoor boxes are shown in Extended Data Fig. 3. Sensor measurements from walking for 16 seconds on the different boxed terrains are also available in Extended Data Fig. 5. The measurement methods were evaluated on a number of terrains inside and outside during development, and corresponded well to perceived terrain features by the researchers. Varying the morphology does have some impact on the measurements of the walking surface, but quite large changes to the morphology is needed before significant differences are found. For hardness, the increased speed of the legs, the reduced force from the servos, and more flexibility in the mechanism as the distance between the body and the end of the leg increase are important factors. For roughness, the distance

to the ground has an effect, as well as increase motion blur from higher speeds. Extended Data Fig. 4 shows terrain measurements for different speeds, which is directly correlated with overall leg length.

Terrain characterization, used for adaptation in realistic outdoor environments, is done via these two features directly. Classification, used for adaptation in the controlled indoor environment, is done by calculating the Euclidean distance to the mean value for each terrain group in the data set, and selecting the closest match. Pseudo-code for both terrain measurement methods are included in the Supplementary Methods.

### 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 contains a mix of gravel, sand, and concrete surfaces. The gravel and sand was filled to an approximate depth of 15cm, but half a box with concrete was not possible due to the high weight. Instead, a flat sheet of fiber-reinforced concrete was placed on top of highly compacted dirt and mulch. These materials were selected to give a wide spread in the hardness and roughness features 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. 25 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.

### 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. While walking, the sensed terrain characteristics are used to generate a map of predicted performance for all possible morphologies. 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 fully independent sub-models, one for each of the leg-length pairs in the data set. A diagram can be seen in Extended Data 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. The output is clamped to provide a COT prediction between 0 and 40, as

values beyond this is considered unrealistic. The parameters for all the sub-models after indoor data collection can be seen in Supplementary Table 5. Each time the robot tests a new morphology, the terrain measurements and performance are 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.

### **Indoor adaptation experiment**

Our goal was to test the simple adaptation method in a controlled indoor environment. This serves as a precursor to continuous adaptation in unstructured terrains outside by evaluating the feasibility of our methods on a simpler problem. 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 the Supplementary Methods.

### **Outdoor adaptation experiment**

Our goal was to test the extended adaptation method in a realistic outdoor environment, 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 grassy areas and a concrete road. As experiments were conducted in the middle of summer in Australia, the earth was very hard and dry, and the light covering of grass did not contribute much in terms of perceived hardness. 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 for all tests. 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 data set, 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.

## Robot control

The robot uses a high-level spline-based gait controller detailed in [56]. 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 Extended Data Fig. 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 (the sideways movement) was scaled to keep

the RPM of the servos consistent throughout the ground movement. The center of the trajectory spline was kept constant so as not to affect the balance of the robot. Since the frequency of the gait is kept the same, the increased step length for taller robots leads to slightly higher walking speeds (up to 25%). The scaling can be found in Supplementary Table 2.

## Correspondence and requests for materials

All correspondence and request for materials should be addressed to author **T.F.N.**

## 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 [57]. All other relevant data are available from the corresponding author on request.

## Code Availability

The code for running the robot is available on the DyRET Github page [51], and is under a GPL v3 license. The code to setup the experiments and log the results is available in a separate Github repository [58].

## Author Contributions

Author contributions classified in the CRediT taxonomy [59]:

**T.F.N:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data Curation, Writing - Original Draft, Writing - Review and Editing, Visualization.

**C.P.M:** Conceptualization, Methodology, Formal analysis, Writing – Review and editing, Supervision.

**J.T:** Writing - Review and Editing, Supervision, Funding acquisition.

**K.G:** Conceptualization, Methodology, Writing - Review and Editing, Supervision.



**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] Baines, R., Freeman, S., Fish, F. & Kramer, R. Variable stiffness morphing limb for amphibious legged robots inspired by chelonian environmental adaptations. *Bioinspiration & Biomimetics* (2020).
- [3] 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).
- [4] Pfeifer, R. & Bongard, J. *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT press, 2006).
- [5] Wilson, A. D. & Golonka, S. Embodied cognition is not what you think it is. *Frontiers in psychology* **4**, 58 (2013).
- [6] Zhang, T., Zhang, W. & Gupta, M. M. Resilient Robots: Concept, Review, and Future Directions. *Robotics* **6**, 22 (2017).
- [7] 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).
- [8] Nygaard, T. F., Martin, C. P., Samuelson, 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* (ACM, 2018).
- [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 *RoboCup 2018: Robot World Cup XXII*, 70–82 (Springer, 2018).
- [16] Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V. & Hutter, M. Learning Quadrupedal Locomotion over Challenging Terrain. *Science Robotics* **5** (2020).
- [17] Hwangbo, J. *et al.* Learning agile and dynamic motor skills for legged robots. *Science Robotics* **4** (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] Long, J. *Darwin’s Devices: What Evolving Robots can Teach Us about the History of Life and the Future of Technology* (Basic Books, 2012).

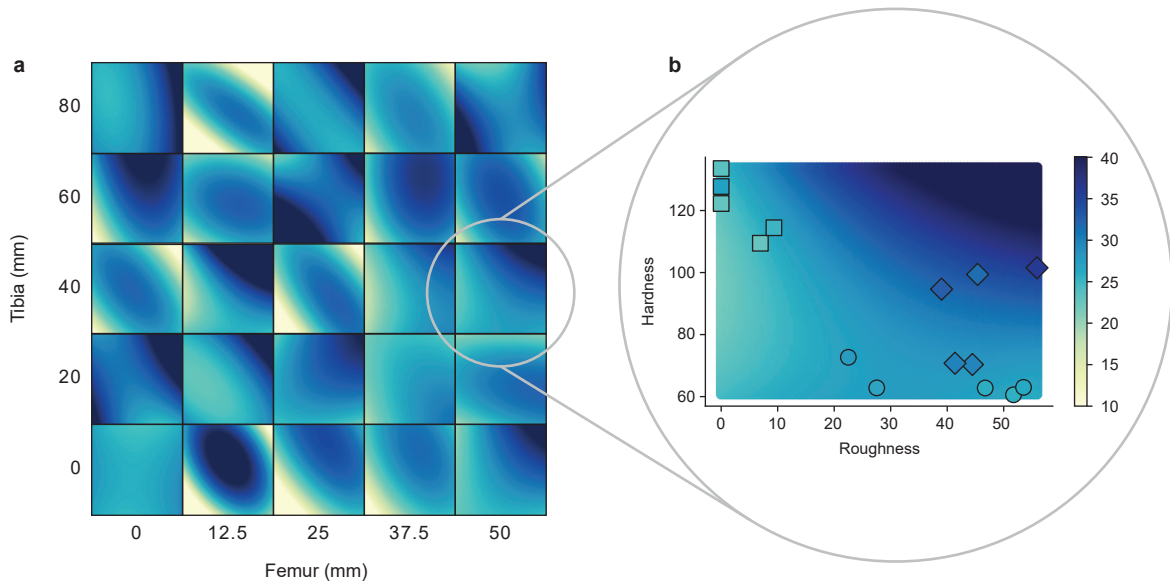
- [23] Eiben, A. E., Kernbach, S. & Haasdijk, E. Embodied Artificial Evolution: Artificial Evolutionary Systems in the 21st Century. *Evolutionary Intelligence* **5**, 261–272 (2012).
- [24] Mouret, J.-B. & Chatzilygeroudis, K. 20 years of reality gap: a few thoughts about simulators in evolutionary robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1121–1124 (2017).
- [25] Cheney, N., MacCurdy, R., Clune, J. & Lipson, H. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, 167–174 (2013).
- [26] 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).
- [27] 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).
- [28] 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).
- [29] Lipson, H. & Pollack, J. B. Automatic design and manufacture of robotic lifeforms. *Nature* **406**, 974–978 (2000).
- [30] 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).
- [31] 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).
- [32] 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).
- [33] Auerbach, J. *et al.* Robogen: Robot generation through artificial evolution. *Artificial Life Conference Proceedings* 136–137 (2014).
- [34] Kriegman, S. *et al.* Scalable sim-to-real transfer of soft robot designs. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, 359–366 (IEEE, 2020).

- [35] 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).
- [36] 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).
- [37] 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).
- [38] 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).
- [39] Jelisavcic, M. *et al.* Real-World Evolution of Robot Morphologies: A Proof of Concept. *Artificial Life* (2017).
- [40] Brodbeck, L., Hauser, S. & Iida, F. Morphological Evolution of Physical Robots Through Model-free Phenotype Development. *PLoS ONE* **10**, 1–17 (2015).
- [41] Vujovic, V., Rosendo, A., Brodbeck, L. & Iida, F. Evolutionary Developmental Robotics: Improving Morphology and Control of Physical Robots. *Artificial Life* (2017).
- [42] 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).
- [43] Nygaard, T. F. *et al.* Experiences from Real-World Evolution with DyRET: Dynamic Robot for Embodied Testing. In *Symposium of the Norwegian AI Society*, 58–68 (Springer, 2019).
- [44] Zhakypov, Z. & Paik, J. Design Methodology for Constructing Multimaterial Origami Robots and Machines. *IEEE Transactions on Robotics* **34**, 151–165 (2018).
- [45] Riviere, V., Manecy, A. & Viollet, S. Agile Robotic Fliers: A Morphing-Based Approach. *Soft Robotics* **5**, 541–553 (2018).
- [46] 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).

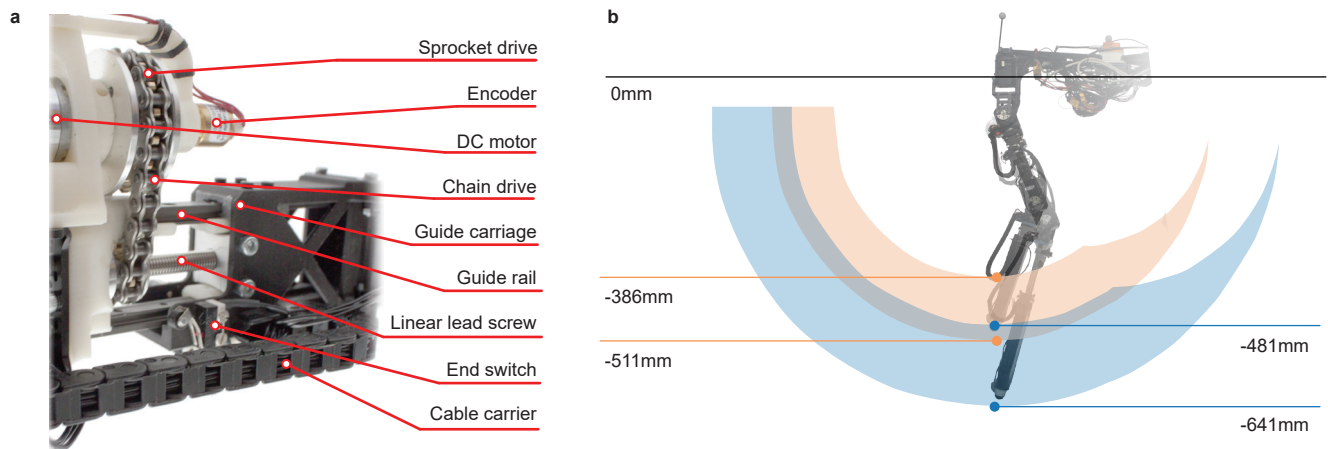
- [47] Geilinger, M., Poranne, R., Desai, R., Thomaszewski, B. & Coros, S. Skaterbots: Optimization-based design and motion synthesis for robotic creatures with legs and wheels. *ACM Transactions on Graphics (TOG)* **37**, 1–12 (2018).
- [48] Meiri, N. & Zarrouk, D. Flying STAR, a Hybrid Crawling and Flying Sprawl Tuned Robot. In *2019 International Conference on Robotics and Automation (ICRA)*, 5302–5308 (IEEE, 2019).
- [49] Kriegman, S., Blackiston, D., Levin, M. & Bongard, J. A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences* (2020).
- [50] Ritter, A. *”Shape-Changing Smart Materials”*, 46–71 (Birkhäuser Basel, Basel, 2007).
- [51] Nygaard, T. F. & Nordmoen, J. Dyret Documentation GitHub repository. [https://github.com/dyret-robot/dyret\\_documentation](https://github.com/dyret-robot/dyret_documentation).
- [52] 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).
- [53] 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).
- [54] Xi, W., Yesilevskiy, Y. & Remy, C. D. Selecting gaits for economical locomotion of legged robots. *The International Journal of Robotics Research* **35**, 1140–1154 (2016).
- [55] Howard, A. & Seraji, H. Vision-based terrain characterization and traversability assessment. *Journal of Robotic Systems* **18**, 577–587 (2001).
- [56] 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).
- [57] Nygaard, T. F. Dataset hosted on figshare. [doi.org/10.6084/m9.figshare.12661619](https://doi.org/10.6084/m9.figshare.12661619).
- [58] Nygaard, T. F. Tonnesfn Experiments GitHub repository. [https://github.com/tonnesfn/tonnesfn\\_experiments](https://github.com/tonnesfn/tonnesfn_experiments).
- [59] 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).

- [60] Nygaard, T. F. Terrain Characterizer GitHub repository.  
[https://github.com/dyret-robot/terrain\\_characterizer](https://github.com/dyret-robot/terrain_characterizer).

## Extended Data Figures




Extended Data Figure 1: **A diagram of the prediction model used, instantiated with the baseline data set measurements from the indoor boxes** (a) The full model with the 25 sub-models for each leg length combination. (b) The sub-model for femur 50mm tibia 40mm, which shows the predicted energy efficiency (COT) of that leg-length combination for different terrains. The 15 points are the actual measurements from the indoor data collection (square from concrete, circle from sand, and diamond from gravel). When the robot encounters a new terrain, each sub-model is queried at the given roughness/hardness values to form a predicted 5x5 COT map similar to the ones shown in Figure. 2 in the main paper.




Extended Data 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.

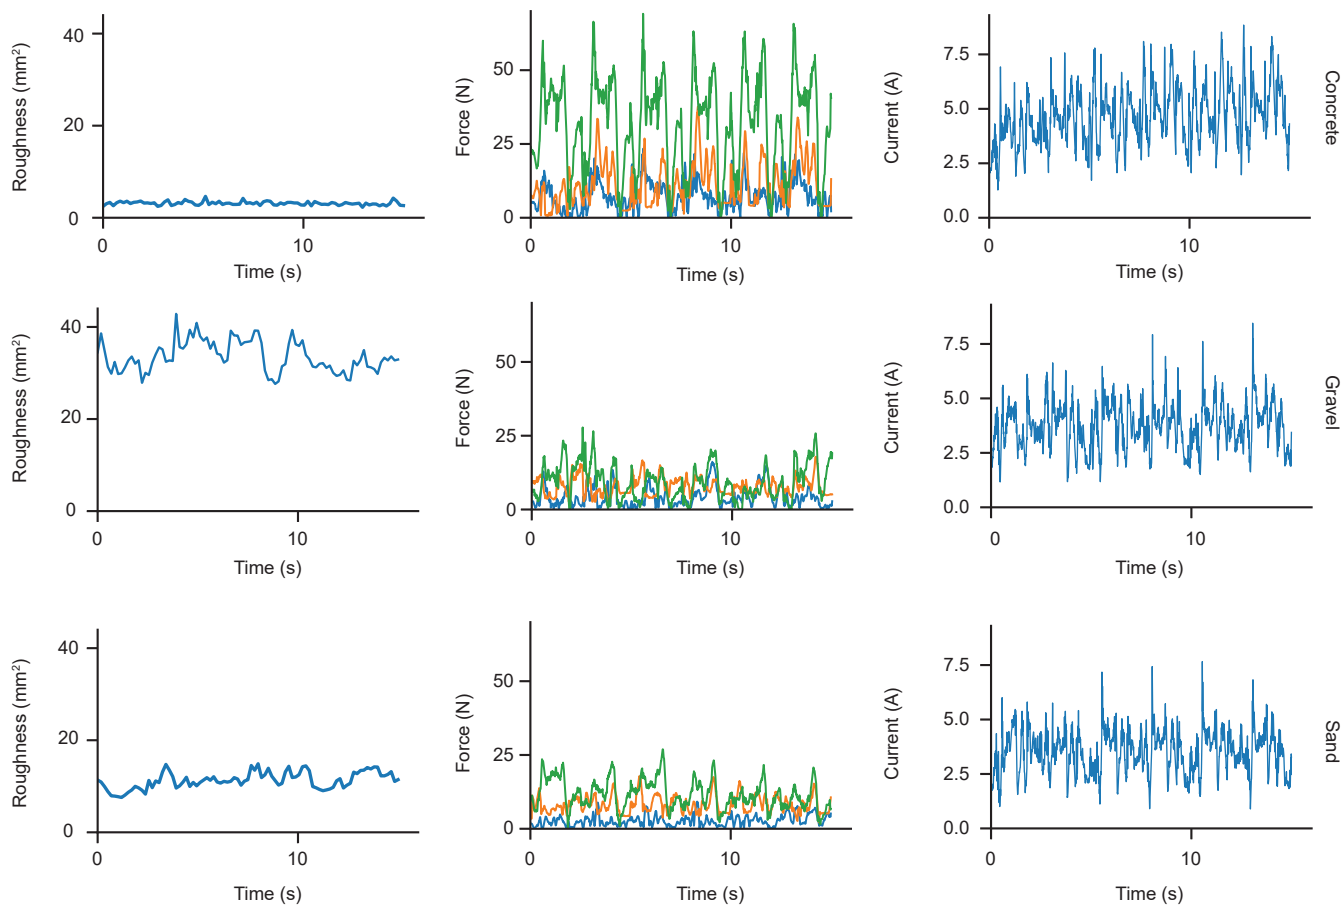




Extended Data Figure 3: **Terrain features for indoor boxes** Kernel density estimate of the terrain features for the three different surfaces in the indoor terrain boxes, measured during the data set collection. This includes all the different morphologies. Default `seaborn.jointplot` parameters are used for the estimate.



Extended Data Figure 4: **Effect of speed on terrain measurements** Terrain characteristics of the three surfaces in the indoor terrain boxes, shown for different walking speeds. The solid lines show the mean, and the shaded areas show the standard deviation. (a) Perceived roughness from the depth camera in the front of the robot. (b) Perceived hardness from the force sensors in the feet.



Extended Data Figure 5: **One walking sequence on the indoor terrains** Generated from single 16 second walking sessions with shortest possible leg length. Roughness to the left is calculated from the depth camera. The middle plot shows the absolute, filtered three axis force measurements used to infer hardness, summed for the two front leg sensors (x-axis, sideways, in blue; y-axis, lengthwise, in orange; and z-axis, up, in green). The right plot shows the total current reported by the servos.

# Supplementary Materials

## Supplementary Methods

---

**Experiment 1:** Adapting in controlled indoor environments. Stop when sensing a change in terrain, reconfigure the morphology, then start walking again.

---

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 in realistic outdoor environments - Adapt morphology and measure terrain and performance without stopping.

---

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.

---

**Algorithm 1:** Sensing terrain roughness with the Intel Realsense RGBD camera.

---

Find 3d plane best approximating the point cloud

Inliers = all points within a 35mm distance of the plane

Square error = 0

**for** *all inliers* **do**

| Square error += squared distance between current inlier and plane

**end**

Mean square error = Square error / number of inliers

**if** *More than 30000 inliers:* **then**

| Roughness = Mean square error

**else**

| Roughness = 0

**end**

---

Standard methods from the PCL library are used for plane extraction and segmentation. The RGBD camera used had problems with flat surfaces without discerning visual features, so a roughness of 0 was assumed when a very low number of points were returned by the sensor. The functionality for roughness sensing is mainly implemented in the `pointCloudPlaneFitter.cpp` file in the `terrain_characterizer` repository [60].

---

**Algorithm 2:** Sensing terrain hardness with the Optoforce leg sensors.

---

```
for both front legs do
  for all three axes do
    if current force measurement > 100N then
      | discard measurement
    end
  end

  filtered x-axis force = median(last 5 x-axis measurements)
  filtered y-axis force = median(last 5 y-axis measurements)

  if current z-axis force measurement < 0 then
    | current z-axis force measurement = 0
  end

  filtered z-axis force = median(last 5 z-axis measurements)
end

hardness = 0
for both front legs do
  for all three axes do
    | hardness += max filtered absolute force over last 6 seconds
  end
end
```

---

The force sensors exhibit some high frequency noise, and a median filter of size 5 has been used to reduce it. Erroneous values over 100N are also removed due to being above the nominal capacity of the sensors. Due to the construction of the sensors, it sometimes reports negative forces in the z-direction. These are also filtered away, as they do not reflect actual forces. The full datasheet for the sensors has been uploaded to the dyret-documentation repository [51].

## 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%	Best trade-off for all surfaces
16	37.5	20	57.5	113%	
17	37.5	40	77.5	116%	
18	37.5	60	97.5	119%	
19	37.5	80	117.	122%	Optimal on sand Optimal on concrete
20	50	0	50	113%	
21	50	20	70	116%	
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 (mm <sup>2</sup> )			Hardness (N)		
	Median	IQR	Max	Median	IQR	Max
Concrete	5.1	7.6	14.0	135.5	13.3	155.2
Sand	25.2	18.5	63.9	61.3	8.2	82.8
Gravel	35.7	12.8	58.1	85.7	24.9	119.7

Supplementary Table 4: Values and ranges for all sensors used on the robot. \* This is the nominal range, and the sensor can return values outside of this, depending on calibration.

Sensor	Measurement	Rate (Hz)	Unit	Value
Servo (x 12)	Current	100	A	[-9.2, 9.2]
Optoforce leg sensors (x 4)	Force x-direction	100	N	*[-40, 40]
	Force y-direction	100	N	*[-40, 40]
	Force z-direction	100	N	*[0, 80]
Realsense depth camera	3D point cloud	6	mm	
Motion Capture rig	Position, x-direction	100	m	[0, 8]
	Position, y-direction	100	m	[0, 8]
	Position, z-direction	100	m	[0, 6]
RTK GPS	Position, x-direction	8	m	
	Position, y-direction	8	m	

Supplementary Table 5: Regression parameters for models trained on the indoor data set. Linear regression from the Scikit Learn library was used to generate and deploy the models. Femur and tibia are measured in mm, roughness in mm<sup>2</sup>, and hardness in N.

Femur	Tibia	independent term	roughness	hardness	roughness <sup>2</sup>	roughness * hardness	hardness <sup>2</sup>
0	0	23.111	-0.292	0.124	0.011	-0.001	0
0	20	101.816	-5.604	-0.392	0.085	0.023	0
0	40	-34.391	2.236	0.903	-0.028	-0.012	-0.03
0	60	-34.656	2.999	0.513	-0.049	-0.005	-0.001
0	80	42.523	-0.835	-0.218	0.022	0.003	0.001
12.5	0	-93.561	4.471	1.636	-0.065	-0.017	-0.006
12.5	20	51.254	-1.030	-0.495	0.015	0.008	0.002
12.5	40	64.163	-0.607	-1.038	-0.009	0.018	0.005
12.5	60	-47.784	1.681	1.251	-0.019	-0.007	-0.006
12.5	80	-160.639	3.238	2.799	-0.017	-0.021	-0.011
25	0	-62.823	2.991	1.040	-0.034	-0.012	-0.003
25	20	11.706	-0.204	0.296	-0.009	0.009	-0.002
25	40	-98.220	4.566	1.426	-0.045	-0.021	-0.004
25	60	126.911	-2.581	-1.117	0.008	0.021	0.002
25	80	77.427	-1.355	-0.713	0.008	0.010	0.002
37.5	0	-27.108	1.801	0.626	-0.023	-0.006	-0.002
37.5	20	1.300	1.180	0.223	-0.011	-0.004	-0.001
37.5	40	29.676	0.633	-0.405	-0.009	0.002	0.002
37.5	60	-18.316	0.874	0.625	-0.009	-0.001	-0.002
37.5	80	-32.512	1.157	0.914	-0.008	-0.006	-0.004
50	0	23.989	-0.111	-0.065	-0.008	0.009	0
50	20	-34.658	0.788	1.175	-0.005	-0.004	-0.006
50	40	43.179	-0.147	-0.449	-0.003	0.005	0.002
50	60	-55.32	1.509	1.382	-0.014	-0.007	-0.006
50	80	60.060	-1.578	0.022	0.009	0.011	-0.003