Research Article

Evolving Controllers for a Transformable Wheel Mobile Robot

Anthony J. Clark ¹, **Keith A. Cissell**,¹ and **Jared M. Moore** ²

¹Missouri State University, Springfield, Missouri, USA

²Grand Valley State University, Allendale, Michigan, USA

Correspondence should be addressed to Anthony J. Clark; anthonyclark@missouristate.edu

Received 17 August 2018; Revised 12 November 2018; Accepted 21 November 2018; Published 16 December 2018

Guest Editor: Xiaofeng Liu

Copyright © 2018 Anthony J. Clark et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Unmanned ground vehicles (UGVs) are well suited to tasks that are either too dangerous or too monotonous for people. For example, UGVs can traverse arduous terrain in search of disaster victims. However, it is difficult to design these systems so that they perform well in a variety of different environments. In this study, we evolve controllers and physical characteristics of a UGV with transformable wheels to improve its mobility in a simulated environment. The UGV's mission is to visit a sequence of coordinates while automatically handling obstacles of varying sizes by extending wheel struts radially outward from the center of each wheel. Evolved finite state machines (FSMs) and artificial neural networks (ANNs) are compared, and a set of controller design principles are gathered from analyzing these experiments. Results show similar performance between FSM and ANN controllers but differing strategies. Finally, we show that a UGV's controller and physical characteristics can be effectively chosen by examining results from evolutionary optimization.

1. Introduction

Autonomous unmanned ground vehicles (UGV) provide an excellent solution to tasks that require searching or monitoring in environments deemed too remote or dangerous for humans. Consider *search and rescue*: after a natural disaster a UGV can be used by first responders to help locate victims in unstable and hazardous locations. UGVs have long operating durations, can carry heavy payloads (e.g., sensors), and can search in narrow and covered places such as forests and caves.

Ensuring that a UGV can handle many different types of terrain is an ongoing challenge. Researchers have invented several different methods for addressing the issue of mobility in varied terrain. Specifically, robots have been designed with treaded wheels, tracks, legs [1], legged-wheels (wheels are rimless, wheel spokes make contact with the ground) [2–5], wheeled-legs (wheels are on the end of legs and suspensions can be actuated) [6–8], and transformable wheels [9–12]. Although these systems provide an advantage over traditional wheeled robots, optimization is not performed in the vast majority of these studies. Moreover, as identified by Mintchev and Floreano [13], most researchers in the area of transformable wheels currently focus on the mechanical

design and leave control and decision making to future work. For example, most robots with transformable wheels are controlled remotely [11, 14], and Kim et al. [9] designed a passive triggering mechanism that does not require any controller input.

The device in this study, the *Adabot* (see Figure 1), includes transformable wheels that can smoothly be converted from a round wheel, to a wheel with tire *studs*, to a legged-wheel. Wheel transformations are performed by extending wheel struts radially outward from the center of the wheel (see Figure 2). *Adabot* has been optimized using an evolutionary algorithm such that its physical characteristics and its controller are better able to handle terrain that includes obstacles of varying sizes. In previous work [15], a similar system was optimized to maximize forward velocity over uneven terrain. The present study differs in two main ways: (1) here we evolve controllers for a more difficult task: *way-point following*, and (2) we analyze results from evolving two types of feedback controllers (rather than feed-forward).

In this study, we evolve the robot's chassis dimensions, wheel radius, the number of wheel struts, along with either a finite state machine (FSM) controller or an artificial neural network (ANN) controller. The best evolved FSMs and ANNs

FIGURE 1: Adabot, a UGV with transformable wheels.

FIGURE 2: Illustration of Adabot's wheel extension mechanism, where the struts are fully retracted (left), partially extended like tire studs (middle), and fully extended like legged-wheels (right).

are analyzed and compared. For this initial work, to ensure that we are able to effectively analyze the ANN, the network only has three input nodes, zero hidden nodes, and three output nodes. The inputs are fully connected to the outputs. The network is only slightly more complex than a Type 2 Braitenberg vehicle [16]. Conclusions drawn from our analysis are used to create a set of design principles for a new controller that takes advantage of both techniques. In particular, it is attractive to design a controller that is not a black-box like an ANN but less rigidly defined than an FSM. Source code has been made available on GitHub (<https://github.com/anthonyjclark/adabot02-ann>).

2. Related Work

In the field of evolutionary robotics (ER), an evolutionary algorithm (EA) optimizes free variables of a given system [17]. ER methods have been successfully applied to many different types of robotic systems (aerial, aquatic, walking, etc.). For example, we have previously used differential evolution to evolve adaptive neural networks and morphologies for a robotic fish [18, 19], and Moore et al. [20] evolved hierarchical controllers for segmented worm-like animats. Although evolution has been regularly utilized at an abstract level to optimize wheeled-robot navigation processes (for example, see Gomes et al. [21] and Lehman and Stanley [22]), it has not often been used to directly evolve UGV morphologies, and to the best of our knowledge this is the first study in which the characteristics of a transformable wheel are evolved.

A large number of ER studies utilize ANNs to control mobile robots, including *Evolving Virtual Creatures* [23],

which is considered one of the first ER works. ANNs provide several benefits when using an evolutionary method. First, since ANNs are so-called universal approximators [24], evolution often produces *novel* and sometimes unintuitive results that may not have been found when creating a controller by hand [25]. And second, ANNs require a minimal amount of user design. Specifically, an evolutionary algorithm can automatically decide the importance of each input (sensor values) in the calculation of each output (actuation mechanisms) [26]. The primary disadvantage of using an ANN is that it is considered a black-box system. That is, *how* an ANN achieves its results is not often clear or analyzed. Recently, however, some researchers have attempted to extract state machines from evolved neural networks. For example, Yaqoob and Wróbel [27] automatically generated a state machine with the same properties of an evolved spiking neural network.

3. Adabot

Hardware. The Adabot, pictured in Figure 1, is a prototype device that includes a Raspberry Pi 3 Model B (RPI) as its main control board. The RPI was chosen for its ability to run the Robot Operating System (ROS) [28, 29], which Adabot uses to deploy its software systems. The size of an RPI constrains the minimum dimensions of the Adabot's chassis. Specifically, the chassis must be at minimum 8 cm by 8 cm. Table 1 lists all configurable parameters for Adabot's physical characteristics, where *WheelBase* and *TrackWidth* denote the distance between the front and rear axles and the lateral distance between wheels, respectively, and *StrutCount* parameter indicates the number of struts per wheel.

FIGURE 3: An example environment including randomly generated obstacles. The current way-point is shown as a dark gray sphere, and the robot starts at the origin facing in the positive x-axis (away from the way-point, along the red axis line).

TABLE 1: Adabot physical parameters.

Name	Range
<i>WheelBase</i>	8 to 16 cm
<i>TrackWidth</i>	8 to 16 cm
<i>WheelRadius</i>	2 to 3 cm
<i>StrutCount</i>	0 to 7

Each wheel is driven by its own DC gear-motor with magnetic encoders. Likewise, each wheel includes a set of struts that can be extended and retracted by a linear servo. For sensing, Adabot includes three forward facing distance sensors and an IMU (3-axis gyroscope, 3-axis accelerometer, and 3-axis magnetometer). Finally, it uses a 2.4 GHz wireless communication module and is powered by a 2200 mAh battery pack, which provides roughly two hours of operating time.

Strut Extension. Figure 2 depicts the strut extension process. This mechanism enables the wheel to exhibit a range of characteristics. With the struts fully retracted, the wheels operate conventionally; when extended a small amount, the struts act as tire studs; and with the struts fully extended, each wheel resembles a legged-wheel. Due to limitations of the design, the maximum extension of the struts is equal to the wheel's radius minus 1 cm ($MAXEXT = WheelRadius - 1$ cm). For a more detailed discussion of Adabot's software and wheel extension mechanism, and an example of evolving Adabot with ROS and Gazebo (a simulation environment tightly coupled with ROS), see our preliminary study [15].

Simulation. An image of the simulation environment is shown in Figure 3. The environment is populated by generating 40 boxes with random dimensions, positions, and densities. These boxes act as obstacles that the simulated robot must traverse. If a newly generated box collides with an existing box it is removed from the simulation. We see on average 31 boxes placed in the environment. Box heights range from 2 to 5 cm, which is high enough (compared to *WheelRadius* values) to drastically reduce mobility for a

wheeled robot [30]. Moreover, rather than each box being in a fixed position, it is possible for the Adabot to push a box (depending on its size and density).

For this study, we are using the Dynamic Animation and Robotics Toolkit (DART) (<https://dartsim.github.io/index.html>). DART is specifically designed for robotics applications, and is comparable in speed (if not faster) than common alternatives [31].

Way-Point Navigation Control. Adabot is a skid-steer style robot—it turns by rotating its left and right wheels at different rates. Although each wheel and wheel strut set can be controlled independently, in this study we only have three control outputs: (1) an angular rate for the left wheels, (2) an angular rate for the right wheels, and (3) a single extension amount for all four sets of struts. Although it may be beneficial to control each wheel independently, for this study we have chosen to synchronize both left wheels and both right wheels. This reduces the number of evolved control parameters and enables us to use a differential drive model for predicting the robots dynamics. In the future, we will explore the effects of controlling each wheel independently.

For Adabot to aid during a search and rescue operation, it must be able to successfully *cover* (completely search) its designated area. A simplified version of this task, called *way-point navigation*, is considered during evolutionary optimization. For this task, a UGV must visit a set of way-points in sequence.

FSM Control. The hand-designed FSMs for this task are depicted in Figure 4. This FSM includes two independent actions: (a) directing the robot towards the next way-point by controlling the left and right wheels, and (b) extending the struts when the robot is experiencing reduced mobility due to an obstacle. Essentially, the robot remains in the *Forward* state as long as the angle between the heading of the UGV and the direction to the target (α_{target}) is within some threshold. Once the threshold is surpassed, the FSM transitions to either the *Left* or *Right* state. In the *Left* and *Right* states, the robot will rotate in place until α_{target} is greater than $L.ToForwardThresh$ or less than $R.ToForwardThresh$,

FIGURE 4: (a) The FSM controlling the direction of the robot, (b) the single state controlling wheel struts, and (c) a diagram depicting the angles used to trigger transitions between states.

respectively, after which point the FSM transitions back to *Forward*. Threshold angles are shown in Figure 4(c).

To determine when, and by how much, wheel struts should be extended, we use a simple differential drive model and compare expected speeds with measured speeds. Specifically, we calculate an expected linear (v) and angular (ω) velocity (based on the wheel rates) using the following model:

$$v = \frac{V_r + V_l}{2}, \quad (1)$$

$$\omega = \frac{V_r - V_l}{TrackWidth} \quad (2)$$

where V_r and V_l are the left and right wheel linear velocities, respectively, and $TrackWidth$ represents the distance between wheels on the same axle line (front or rear axles). These calculated values (expected based on the differential drive model) are then subtracted from the actual (**measured**) linear and angular velocities values. The actual speed of the simulated robot is provided by the simulator, and in a real-world environment it can be measured using an overhead camera system. The difference values (the error between expected and actual velocities) are then scaled between 0 and 1 to produce v_e and ω_e , which are the scaled linear and angular velocity errors, respectively. These two error values are then filtered using exponential smoothing. Finally, they are used in the following to calculate the extension amount of all struts:

$$ext_v = m_{ext} \cdot v_e + b_{ext}, \quad (3)$$

$$ext_\omega = m_{ext} \cdot \omega_e + b_{ext}, \quad (4)$$

$$ext_{\%} = \max[ext_v, ext_\omega], \quad (5)$$

$$ext = MAXEXT \cdot ext_{\%} \quad (6)$$

where ext_v and ext_ω denote the extension amount calculated due to the linear and angular speed values, respectively. These two values are calculated using a linear equation with a configurable slope (m_{ext}) and intercept (b_{ext}). The final extension amount (ext) is based on the maximum of these two values, and is calculated as a percentage of the maximum possible extension ($MAXEXT$). In essence, the struts will be extended by an amount that is linearly proportional to the current error in speed (maximum between linear and angular error). Thus, when Adabot encounters an obstacle that reduces its mobility (compared to that predicted by the differential drive model), it will extend the struts in an effort to climb over the obstacles.

Table 2 shows all configurable parameters for the FSM (hand-chosen values are shown in parentheses). Aside from the first m_{ext} and b_{ext} , each name in the table takes the following form: a capital letter representing a state in Figure 4(a) (*Forward*, *Left*, or *Right*), followed by a period, followed by either an angular wheel rate or a angle threshold value also described in Figure 4(a). Finally, to reduce vibration and potential damage to the wheel struts, the maximum angular

TABLE 2: Adabot FSM parameters (Hand-Chosen).

Name	Range
m_{ext}	(0.5) 0 to 1
b_{ext}	(0.5) 0 to 1
$F.Speed$	(20) 0 to 20 rad s ⁻¹
$F.ToLeftThresh$	(10) 0 to 90°
$F.ToRightThresh$	(-10) -90 to 0°
$L.LeftSpeed$	(-20) -20 to 20 rad s ⁻¹
$L.RightSpeed$	(20) -20 to 20 rad s ⁻¹
$L.ToForwardThresh$	(5) 0 to 90°
$R.LeftSpeed$	(20) -20 to 20 rad s ⁻¹
$R.RightSpeed$	(-20) -20 to 20 rad s ⁻¹
$R.ToForwardThresh$	(-5) -90 to 0°

rate of the wheels is linearly scaled down from 20 rad s⁻¹ to 4 rad s⁻¹ when the struts are fully extended.

ANN Control. As an alternative to the FSM controller, we evolve an ANN for the same task. The neural network receives three inputs (each scaled between 0 and 1): (1) α_{target} , (2) v_e , and (3) ω_e . Essentially, the ANN is given the same information as the FSM, and produces the same three output values (left and right wheel rates and an extension amount). In our preliminary work, we found hidden nodes were unnecessary for this task (the same strategies and fitness values were attained with and without hidden nodes). The genome for our ANN includes 13 values: one integer value representing the activation function (logistic, hyperbolic tangent, or the rectified linear unit) and 12 values for the neural network weights (three inputs plus one bias for each of the three outputs).

Evolution. For this study, we employ the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [32]. In particular, we use pycma (developed by Hansen [33]), which works well on real-valued problems and has support for handling integer values such as *ExtCount*.

4. Discussion and Results

In this section we provide our results from evolving the Adabot. Specifically, we evolved the Adabot in two environments (with and without obstacles) and with two different controllers. Each of these four experiments is repeated 20 times. Finally, we discuss principles that can be learned from these experiments.

4.1. Fitness Calculations. Here, Adabot's goal is to visit a set of coordinates (way-points) in sequence. During a single simulation, the device has 30 (t_{max}) seconds to visit **four** predefined way-points, but the simulation will terminate as soon as the fourth way-point is reached. Fitness is calculated as follows (pycma is used to maximize this function):

$$f = 2w + \left(1 - \frac{\min[d, d_{max}]}{d_{max}}\right) + \frac{(t_{max} - t)}{t_{max}} \quad (7)$$

where w represents the number of way-points reached, d and d_{max} denote the distance to the next way-point and a scaling factor for distances, respectively, and t denotes the time transpired. This function is meant to provide a smooth gradient for generating controllers that quickly navigate to all way-points in order. The first part of the equation ensures that the CMA-ES algorithm heavily favors any controller that reaches even a single way-point; values for this component range from 0 to 8. Next, a distance component is added to reward solutions that drive near the next way-point in sequence, but do not reach all four. This is particularly useful at the beginning when solutions are at an early stage of evolution. The distance component results in a value scaled between 0 and 1. Since the simulation ends once all four way-points have been reached, the time component will be a value between 0 (zero time remaining) and 1 (all four way-points are reached in an instant). The time component is meant to favor any controllers that solve the task quickly. Thus, the maximum possible fitness is 10.

4.2. Evolution without Obstacles. In our first experiment, *FSM-0-1*, we evolve the fifteen parameters found in Tables 1 (physical) and 2 (control) in an environment without obstacles. The naming scheme for our experiments indicates the controller type (*FSM* or *ANN*), the maximum number of potential obstacles (0 or 40), and the number of trials per fitness evaluation (1 or 2). Plots of fitness vs iteration are shown in Figure 5 (this figure shows the fitness values for both experiments not containing obstacles). In this first experiment, there are zero obstacles and therefore the environment will always be the same. In later experiments, each fitness evaluation includes two trials with randomly generated obstacles. As shown in the figure, in all replicate experiments the Adabot reaches all four way-points in approximately 10 seconds, which corresponds to a fitness value of 9.7. The population quickly converges on a final value, likely because this experiment was seeded with a hand-designed set of parameters known to achieve good results (see Table 2). The evolved results, however, quickly outperform the hand-chosen values. This experiment serves as a convenient baseline with which the others can be compared.

The second experiment, denoted *ANN-0-1*, also reaches a fitness value of 9.7, which shows that the an ANN can effectively perform the task of navigating the robot to a sequence of points. For this experiment, 17 total parameters were evolved: the four physical characteristics listed in Table 1 and the 13 ANN parameters discussed in the previous section. Although both of these experiments reach the same final fitness value, an examination of Figure 5 shows that the ANN result takes longer to evolve—roughly 120 iterations compared with less than 10 iterations for the FSM. This can be explained by the lack of a seed controller and the fact that, unlike an FSM, an ANN must learn the entire solution from scratch.

Figure 6 depicts the trajectories taken by the best performing controllers from these two experiments. Although these trajectories look similar, there is one key difference: the ANN actively controls only one wheel. FSMs, on the other

FIGURE 5: Plots for the maximum fitness found in the two experiments without obstacles. Shaded regions indicate confidence intervals of one standard deviation from the mean for the 20 replicates of each experiment. The maximum possible fitness is 10, and fitness values above 2 indicate that Adabot was able to reach at least one way-point.

FIGURE 6: Trajectories of the best evolved individual for the first two experiments: FSM-0-1 and ANN-0-1. No obstacles were present for these trajectories.

hand, can rotate in-place both clockwise and counterclockwise, which is why there are sharper turns in the left plot.

Figure 7 shows the wheel speeds for the best FSM and ANN controllers. The evolved ANN perpetually sets the right wheel to its maximum speed. The ANN moves forward by setting its left wheel to the same value, and turns by making the left wheel rotate in the opposite direction. Effectively, the ANN can only turn left, however, this is not a problem for the relatively simple task at hand.

Figure 8 provides a comparison of the fitness values and evolved physical characteristics for these two experiments. This figure only shows results for the combined final populations of all replicate experiments. From this figure, we can establish what will be good physical characteristics for the Adabot when it does not face any obstacles. Specifically, *WheelBase* and *TrackWidth* should be 8.5 cm and 11.5 cm, respectively, *WheelRadius* should be 3 cm, and *StrutCount* does not matter since the struts are not extended.

FIGURE 7: Left and right wheel speeds and strut extension amounts for the best evolved solutions for *FSM-0-1* (a) and *ANN-0-1* (b). The left vertical axis shows values for wheel speeds (the solid red line and the dash-dot blue line), and the right vertical axis shows the scale for wheel strut extension amounts (the orange dashed line).

FIGURE 8: Distributions for the evolved fitness values and physical characteristics for the combined final populations of the *FSM-0-1* (left side) and *ANN-0-1* (right side) experiments. The y-axis limits are the parameter limits allowed during evolution.

Finally, Figure 9 plots the distributions for all evolved FSM parameters. It is worth noting that in the absence of obstacles, neither the evolved FSMs or the ANNs extend the struts by a significant amount. This result is not unexpected as any extension would result in a reduced speed due to scaling of linear velocity mentioned previously, and the struts are not needed when obstacles are not present. Also of interest is the evolved symmetry of the FSM. Specifically, the threshold values and speeds evolved for the *Left* and *Right* states are nearly perfect mirror images of each other.

4.3. Evolution with Obstacles. The final two evolutionary experiments are referred to as *FSM-40-2* and *ANN-40-2*.

These experiments differ from the previous two in two respects. First, each fitness value is calculated as the average of two trials (where each trial lasts at most 30 seconds), and second, each fitness trial occurs in an environment with around 31 randomly generated obstacles. Utilizing multiple trials during fitness evaluations improves the robustness of the evolved results [34]. The fitness plots for these experiments appear in Figure 10. Of note is that the ANNs evolved with obstacles have a greatly reduced maximum fitness. A few individuals achieve a fitness above 9, however, we found that this was only when the randomly generated environment did not pose much difficulty. Videos (and interactive animations) for high fitness individuals can

FIGURE 9: Distributions for all evolved FSM parameters for the *FSM-0-1* experiment. These parameters are described in Figure 4(a) and Table 2.

be found here: *FSM-40-2*: <https://youtu.be/VXnrwwpE598> (<https://goo.gl/NtoVYe>) and *ANN-40-2*: <https://youtu.be/q8PFqQps5e4> (<https://goo.gl/2xjh6X>).

Similar to Figure 8, Figure 11 shows the distributions for the evolved physical characteristics. These distributions have a larger spread due to the randomly generated environments. The values found in these distributions indicate that the presence of obstacles does not have a drastic effect on the evolution of physical characteristics. At first this was unexpected, however, analyzing these values (and visualizing their resulting behaviors) reveals a few basic principles: (1) for a skid steer robot it is important for the *WheelBase* to be less than the *TrackWidth* (this will reduce the amount of *skidding* and improve controllability), (2) to maximize velocity *WheelRadius* should be maximized (since we are evolving wheel angular rate a larger wheel will result in a higher velocity), and (3) as long as the number of struts is greater than 4 the system will be able to navigate the generated environments. The first and second principles match results

that we have seen on the physical prototype, and we intend to investigate the third principle in the near future.

While the physical characteristics are similar between the two sets of experiments, control strategies have been adjusted to handle the obstacles. Figure 12 shows the control patterns for two solutions randomly selected from the best performing individuals of the *FSM-40-2* and *ANN-40-2* experiments. Note that since environments are randomly generated, even though the evolved ANN does not reach all four way-points for this test, it does not mean that it did not do so during fitness evaluation. The two most striking features of the plots in Figure 12 are that the evolved controllers ANNs are operating at reduced speeds and that with the addition of obstacles to the simulation the wheels struts are being extended for both controllers. For the evolved FSM, the wheel struts are extended when the first obstacle is reached, and they remain roughly halfway extended for the duration of the evaluation. The ANN controller uses a slightly different strategy. The wheel struts are fully extended at the beginning

FIGURE 10: Plots for the maximum fitness found in the two experiments including obstacles. Shaded regions indicate confidence intervals of one standard deviation from the mean for the 20 replicates of each experiment. The maximum possible fitness is 10, and fitness values above 2 indicate that Adabot was able to reach at least one way-point.

of the simulation and remain so throughout. This means that the top speed of the UGV must be reduced for safety (see Section 3).

Examining the evolved FSM values, we see that nearly identical values are discovered for all parameters except b_{ext} (a set of distributions similar to Figure 9 has been omitted to save space). In the experiment with no obstacles, b_{ext} converged to zero; however, for this experiment b_{ext} converged to 0.45. A higher value for b_{ext} results in the struts always being extended (even when no obstacle has been encountered). Thus, these behaviors are slower because the struts are required to climb obstacles.

Directly examining the evolved weights of a neural network provides only a limited view of the resulting behavior. Likewise, comparing each input's effect on each output in isolation obscures the resulting behaviors. For example, some output values are only active when some combinations of multiple input values are provided. Thus, in Figure 13 we provide all pairwise input relationships on the output for the speed of the left wheels in the form of heat-maps. These heat-maps were generated using a parameter sweep over all possible input combinations. Each square represents the output value given the two input values on the x- and y-axes averaged over all possible values for the remaining input. As was the case for the ANN-0-1 experiment, all navigation is handled by driving the left wheel at different speeds, and so we have not provided heat-maps for the wheel strut and right speed outputs. Examining the figure shows that the left wheel's speed has a positive linear relationship with both ω_e and α_{target} and that α_{target} has the greatest effect on control (since it is used to turn the robot towards the target).

In both experiments including obstacles, the evolved controllers extended the struts and never fully retracted them. However, there is a clear advantage to retracting the struts: the robot has a higher maximum allowed speed. Thus, it is likely

an issue with using the differential drive model to calculate the error. We have identified two sources of error with the simple model: (1) it does not take into account that when the struts are extended the wheel has a larger effective radius, and (2) the model does not take into account the noisy nature of skid steering and extended struts.

For our final comparison between these two control models, we took five best performing individuals from each replicate experiment and evaluated them on three new environments. The new environments required the mobile robot to drive four times further and handle twice as many obstacles. The simulation time was also increased from 30 s to 90 s. Results from these evaluations are shown in Figure 14. As shown in the figure, the FSM controllers were still on average able to reach two way-points, while the ANN controllers frequently failed to reach even one.

In summary, regarding the optimization of the Adabot system we found that

- (1) Similar physical characteristics are optimal with and without obstacles in the environment.
- (2) The speed of the left and right wheels should have a linear relationship with α_{target} (rather than a discrete relationship as is the case with the current FSM). This will enable the robot to *veer* towards the target.
- (3) The task can be solved by controlling only a single side of wheels, though, this is likely not a desirable trait. In future work, we plan to add an evolutionary pressure so that the evolved ANNs turn in both directions, for example, by creating environments and way-points that require both left and right turns.
- (4) Controlling the strut will require a more complex model of the robots dynamics. Once the struts are extended, it is difficult to discern when they should be retracted. In our future work, we will investigate vision-based methods and parameter identification for measuring and detecting poor mobility.

Taking these observations into account, we developed a hybrid two-state controller. The controller is in *Left* when α_{target} is greater than zero and in *Right* otherwise. Equations for these states are as follows:

$$\alpha_{scale} = 2 \cdot \left(1 - \frac{\alpha_{target}}{\pi}\right) - 1 \quad (8)$$

$$Left_{left} = -MAXRAD \cdot \alpha_{scale} \quad (9)$$

$$Left_{right} = MAXRAD \quad (10)$$

$$Right_{left} = MAXRAD \quad (11)$$

$$Right_{right} = MAXRAD \cdot \alpha_{scale} \quad (12)$$

where α_{scale} is α_{target} scaled between -1 and 1. This simple hybrid controller is able to visit all way-points in 9.9 seconds, which is one tenth of a second faster than the evolved controllers reported above. The controller also works well in the presence of obstacles when the struts are extended 10%. Overall, this hybrid controller provides a smoother

FIGURE 11: Distributions for the evolved fitness values and physical characteristics for the combined final populations of the *FSM-40-2* (left side) and *ANN-40-2* (right side) experiments. The y-axis limits are the parameter limits allowed during evolution.

FIGURE 12: Left and right wheel speeds and strut extensions for the best evolved FSM (a) and ANN (b) in a randomly generated environment that includes obstacles. Similar to Figure 7, the left and right vertical axes show scales for wheel speeds and strut extension amounts, respectively.

motion and good performance. For future work, we intend to evolve this hybrid controller along with a more sophisticated approach to handling strut extension as mentioned in point (4) above.

5. Conclusion

UGVs are becoming more prevalent. Likewise, their envisioned environments are becoming more dynamic and varied. We have evolved a UGV so that it is better able to handle obstacles of varying sizes. Specifically, we compared and analyzed FSM and ANN controllers with and without

obstacles in the environment while simultaneously evolving the physical characteristics of our UGV. In comparing these two techniques we were able to find design principles that incorporate the advantages of both. Specifically, we found that a mixture of the two strategies seems able to maintain the strengths of both approaches. For example, an advantage of the FSM designed for this study is that it turns in both directions, but there was insufficient evolutionary pressure for this behavior to evolve in the ANNs. On the other hand, ANNs evolved a more continuous nature to their turning. Instead of turning in place, they tend to *veer* towards the target. Our final, hand-designed controller incorporates both

FIGURE 13: Heat-maps showing the relationship between input and output for the left wheels of the best evolved ANN. A light shade indicates that the wheel is at its maximum forward rate, and a dark color indicates that it is at its maximal reverse rate. All input and output values are scaled between 0 and 1. Outputs area only shown for the left wheel has it was the only ANN output that exhibited a variety of different values.

FIGURE 14: The best evolved controllers from each replicate were evaluated in new environments. FSM controllers were still able to reach some way-points, but most ANN controllers failed to reach even one.

of these strategies, but it may not have been obvious to design such a controller without first evolving the FSMs and ANNs.

Although a direction controller is straightforward to optimize, the complex dynamics associated with climbing over obstacles makes it more difficult to design a controller for extending the Adabot's struts. Specifically, the differential drive model used to predict the robot linear and angular speed does not take into account obstacles, wheel slipping, or the extension of wheel struts. Our future work will focus both on optimizing the hybrid controller and investigating different strategies for extending and retracting the struts so that the robot is able to more effectively gain the benefits of both wheeled and legged-wheel locomotion.

One possibility for improving control is to use a recurrent neural network (RNN) for control. Doing so may provide a means by which the robot can sense that it has transitions from one type of terrain to another. Evolving an RNN, however, will require a more careful selection of evolutionary pressures, and it may require a more gradual increase in task

difficulty. A technique such as Lexicase selection [35] could be used to evolve RNNs that work well in many types of terrain.

Data Availability

All code used to produce our results and all data generated by the evolutionary algorithm used to support the findings of this study has been deposited in the following repository: <https://github.com/anthonyjclark/adabot02-ann>.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

This work was supported by NSF Grant no. PHY-9723972.

References

- [1] D. W. Haldane, K. C. Peterson, F. L. Garcia Bermudez, and R. S. Fearing, "Animal-inspired design and aerodynamic stabilization of a hexapedal millirobot," in *Proceedings of the 2013 IEEE International Conference on Robotics and Automation, ICRA 2013*, pp. 3279–3286, Germany, May 2013.
- [2] U. Saranli, M. Buehler, and D. E. Koditschek, "RHex: a simple and highly mobile hexapod robot," *International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, 2001.
- [3] R. D. Quinn, J. T. Offi, D. A. Kingsley, and R. E. Ritzmann, "Improved mobility through abstracted biological principles," in *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2652–2657, Switzerland, October 2002.
- [4] M. Eich, F. Grimminger, F. Kirchner, and D. Bremen, "A versatile stair-climbing robot for search and rescue applications," in *Proceedings of the 2008 IEEE International Workshop on Safety, Security and Rescue Robotics, SSRR 2008*, pp. 35–40, Japan, November 2008.
- [5] G. Keneally, A. De, and D. E. Koditschek, "Design Principles for a Family of Direct-Drive Legged Robots," *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 900–907, 2016.

- [6] C. Grand, F. Benamar, F. Plumet, and P. Bidaud, "Stability and traction optimization of a reconfigurable wheel-legged robot," *International Journal of Robotics Research*, vol. 23, no. 10-11, pp. 1041-1058, 2004.
- [7] C. Zheng, J. Liu, T. E. Grift et al., "Design and analysis of a wheel-legged hybrid locomotion mechanism," *Advances in Mechanical Engineering*, vol. 7, no. 11, 2015.
- [8] L. M. Smith, R. D. Quinn, K. A. Johnson, and W. R. Tuck, "The Tri-Wheel: A novel wheel-leg mobility concept," in *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4146-4152, Hamburg, Germany, September 2015.
- [9] Y.-S. Kim, G.-P. Jung, H. Kim, K.-J. Cho, and C.-N. Chu, "Wheel transformer: a wheel-leg hybrid robot with passive transformable wheels," *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1487-1498, 2014.
- [10] S.-C. Chen, K.-J. Huang, W.-H. Chen, S.-Y. Shen, C.-H. Li, and P.-C. Lin, "Quattroped: a leg-wheel transformable robot," *IEEE/ASME Transactions on Mechatronics*, vol. 19, no. 2, pp. 730-742, 2014.
- [11] W.-H. Chen, H.-S. Lin, Y.-M. Lin, and P.-C. Lin, "TurboQuad: A Novel Leg-Wheel Transformable Robot with Smooth and Fast Behavioral Transitions," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1025-1040, 2017.
- [12] Z. Wei, G. Song, G. Qiao, Y. Zhang, and H. Sun, "Design and implementation of a leg-wheel robot: Transleg," *Journal of Mechanisms and Robotics*, vol. 9, no. 5, 2017.
- [13] S. Mintchev and D. Floreano, "Adaptive morphology: A design principle for multimodal and multifunctional robots," *IEEE Robotics and Automation Magazine*, vol. 23, no. 3, pp. 42-54, 2016.
- [14] . Yu She, C. J. Hurd, and H. Su, "A transformable wheel robot with a passive leg," in *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4165-4170, Hamburg, Germany, September 2015.
- [15] A. J. Clark, "Evolving adabot: A mobile robot with adjustable wheel extensions," in *Proceedings of the 2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1-8, Honolulu, HI, November 2017.
- [16] D. C. Dennett and V. Braitenberg, "Vehicles: Experiments in Synthetic Psychology," *Philosophical Review*, vol. 95, no. 1, p. 137, 1986.
- [17] F. Silva, M. Duarte, L. Correia, S. M. Oliveira, and A. L. Christensen, "Open issues in evolutionary robotics," *Evolutionary Computation*, vol. 24, no. 2, pp. 205-236, 2016.
- [18] A. J. Clark, P. K. McKinley, and X. Tan, "Enhancing a model-free adaptive controller through evolutionary computation," in *Proceedings of the 16th Genetic and Evolutionary Computation Conference, GECCO 2015*, pp. 137-144, Spain, July 2015.
- [19] A. J. Clark, X. Tan, and P. K. McKinley, "Evolutionary multiobjective design of a flexible caudal fin for robotic fish," *Bioinspiration & Biomimetics*, vol. 10, no. 6, 2015.
- [20] J. M. Moore and A. Stanton, "Tiebreaks and Diversity: Isolating Effects in Lexicase Selection," in *Proceedings of the The 2018 Conference on Artificial Life*, pp. 590-597, Tokyo, Japan, July 2018.
- [21] J. Gomes, P. Mariano, and A. L. Christensen, "Devising effective novelty search algorithms: A comprehensive empirical study," in *Proceedings of the 16th Genetic and Evolutionary Computation Conference, GECCO 2015*, pp. 943-950, Spain, July 2015.
- [22] J. Lehman and K. O. Stanley, "Abandoning objectives: evolution through the search for novelty alone," *Evolutionary Computation*, vol. 19, no. 2, pp. 189-222, 2011.
- [23] K. Sims, "Evolving 3D Morphology and Behavior by Competition," *Artificial Life*, vol. 1, no. 4, pp. 353-372, 1994.
- [24] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks*, vol. 2, no. 5, pp. 359-366, 1989.
- [25] J. C. Bongard, "Evolutionary Robotics: Taking a biologically inspired approach to the design of autonomous, adaptive machines," *Communications of the ACM*, vol. 56, no. 8, pp. 74-83, 2013.
- [26] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary Computation*, vol. 10, no. 2, pp. 99-127, 2002.
- [27] M. Yaqoob and B. Wrobel, "Very small spiking neural networks evolved to recognize a pattern in a continuous input stream," in *Proceedings of the 2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1-8, Honolulu, HI, November 2017.
- [28] S. Cousins, B. Gerkey, K. Conley, and W. Garage, "Sharing Software with ROS [ROS Topics]," *IEEE Robotics & Automation Magazine*, vol. 17, no. 2, pp. 12-14, 2010.
- [29] S. Cousins, "Welcome to ROS topics," *IEEE Robotics and Automation Magazine*, vol. 17, no. 1, pp. 13-14, 2010.
- [30] R. D. Quinn, G. M. Nelson, R. J. Bachmann et al., "Parallel complementary strategies for implementing biological principles into mobile robots," *International Journal of Robotics Research*, vol. 22, no. 3-4, pp. 169-186, 2003.
- [31] J.-B. Mouret and K. Chatzilygeroudis, "20 Years of reality gap: A few thoughts about simulators in evolutionary robotics," in *Proceedings of the 2017 Genetic and Evolutionary Computation Conference Companion, GECCO 2017*, pp. 1121-1124, Germany, July 2017.
- [32] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)," *Evolutionary Computation*, vol. 11, no. 1, pp. 1-18, 2003.
- [33] N. Hansen, "The CMA evolution strategy: a tutorial," <http://arxiv.org/abs/1604.00772>.
- [34] E.-L. Ruud, E. Samuelson, and K. Glette, "Memetic robot control evolution and adaption to reality," in *Proceedings of the 2016 IEEE Symposium Series on Computational Intelligence, SSCI 2016*, Greece, December 2016.
- [35] J. M. Moore, A. J. Clark, and P. K. McKinley, "Effect of animat complexity on the evolution of hierarchical control," in *Proceedings of the 2017 Genetic and Evolutionary Computation Conference, GECCO 2017*, pp. 147-154, Germany, July 2017.