Journal of Biomechanics 87 (2019) 28–36

Contents lists available at ScienceDirect

Journal of Biomechanics

journal homepage: www.elsevier.com/locate/jbiomech
www.JBiomech.com

Adapting stiffness and attack angle through trial and error to increase self-stability in locomotion

Kathryn Walker *, Helmut Hauser

University of Bristol, Bristol BS8 1TH, United Kingdom
Bristol Robotics Laboratory, Stoke Gifford, Bristol BS16 1QY, United Kingdom

ARTICLE INFO

Article history:
Accepted 11 February 2019

Keywords:
Learning
Legged locomotion
Trial and error
SLIP model
Control
Morphology

ABSTRACT

Biological systems are outperforming machines in legged locomoting under almost any conditions. This is partly due to their capability of learning from failure and adapting their control approach and morphological features. This paper proposes an approach that extends the spring-loaded inverted pendulum (SLIP) model with the capability to adapt its attack angle (control) and stiffness (morphology) based on previous locomotion attempts. A set of different update rules, i.e., how this experience is used to adapt, are systematically investigated. The results suggest that modifying either attack angle, or stiffness, or both is beneficial with respect to achieve stable locomotion. Particularly, if the current system configuration (control and morphology) outperforms the previous one, the results suggest that increasing the angle and decreasing the stiffness of the system leads to more stable solutions. Consequently, the basic SLIP model extended by the proposed learning capabilities is able to reach stable locomotion over a much wider range of parameter combinations simply through trial and error.

1. Introduction

Despite the huge success of robotics in general, there are almost no machines that are capable to stably locomote in rough, unknown terrain. As a consequence, most of today's robots work only in well-defined environments like assembly lines and on the factory floor. Even state-of-the-art robots are still not able to deal with the uncertainty, complexity and variety typical of natural environments and our living and working places. Simply put, if they can't model it, they'll fail. On the other hand, animals, including us humans, are particularly good at locomotion and we outperform robots in almost any category, including, energy efficiency, stability, robustness, agility, and many others (Alexander, 2003; Wieber et al., 2016). This is partly due to the fact that a lot of the locomotion process happens without the need of an exact model of the environment. For example, part of self-stabilisation can be carried out by the morphological structure of the legs. One could say control is outsourced to the morphology. In robotics, this is often referred to as morphological computation (Hauser et al., 2011, 2012, 2014). Another important difference is that animals are remarkable adaptive. They can locomote over a wide range of different terrains (Vu et al., 2013; Hurst and Rizzi, 2005; Dickinson et al., 2000; Cham et al., 2004). Key properties are their capability of learning from previous experience and their capacity to adapt, i.e., change their control strategy and/or change their morphological properties, e.g., like adapting stiffness. While numerous mechanisms have been proposed to adapt stiffness, typically referred to as variable compliant mechanisms (Vanderborght et al., 2013; Wolf et al., 2016), only few have been used to actively increase locomotion performance. For example, Quy et al. used adaption of stiffness to improve energy efficiency in single-legged hopper (Vu et al., 2013).

In this work we propose an approach that implements these learning and adaptation capabilities on the well-established spring-loaded inverted pendulum model (SLIP) model, which is a prevalent model for analyzing running and hopping (Blickhan, 1989). This simple model is surprisingly general as it works over a wider range of species and locomotion types. In its basic form it describes the action of the leg by representing it as a lossless, linear spring of constant stiffness k and rest length l_o as shown in Fig. 1. The leg has no mass and the body is represented by a point mass. The SLIP model is self-stabilizing, i.e., the system can tolerate small perturbations without loosing its periodic locomotion pattern (Ernst et al., 2010). Seyfarth et al. showed the range of stiffness and angle combinations that resulted in self-stability for a simple SLIP model (Seyfarth et al., 2002; Geyer et al., 2002).

* Corresponding author at: University of Bristol, Bristol BS8 1TH, United Kingdom.
E-mail address: kw16876@bristol.ac.uk (K. Walker).

https://doi.org/10.1016/j.jbiomech.2019.02.009
0021-9290/

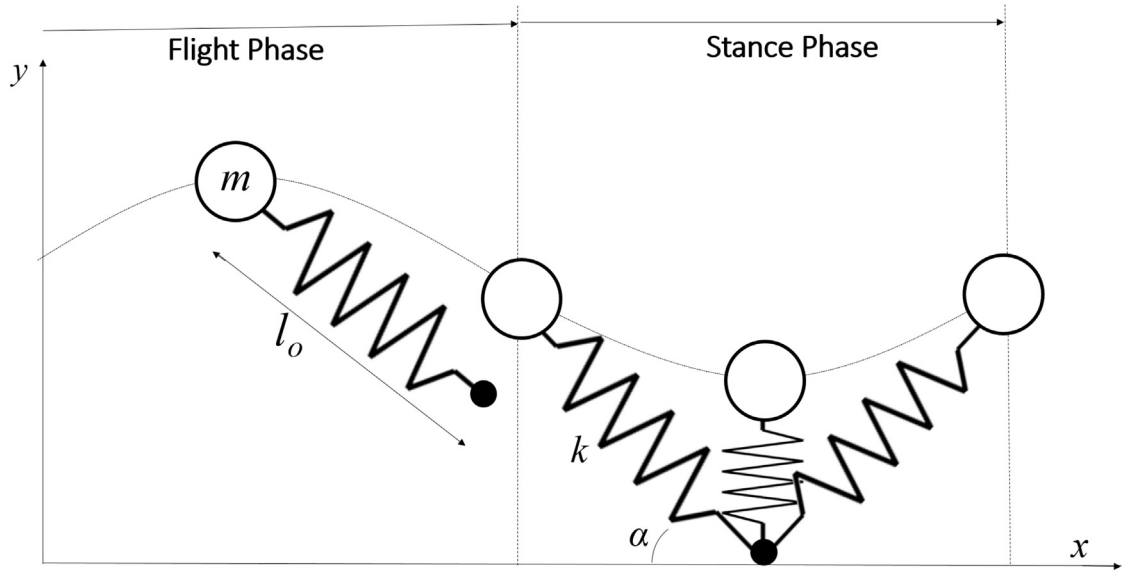

Fig. 1. Figure shows the two stages in the SLIP model – flight and stance. During the flight phase the length of the massless spring remains constant and the pendulum follows a ballistic trajectory until the foot comes in contact with ground level. During stance phase the foot position of the leg remains constant. The mass rotates around the foot position driven by the kinetic energy of the mass. The model switches to flight phase again, when the spring has reached it's resting length represented by l_o. Other parameters the mass m, the stiffness k, and the attack angel α.

Naturally, researchers have tried to find ways to widening this self-stabilizing region. Some have adapted the control of the SLIP model, such as Seyfarth et al. (2003, 2002) who use the stiffness and horizontal velocity values to influence the retraction of the leg during the flight phase and thus change the attack angle at touch down. Similarly, Schmitt (2006) changes the attack angle based on the angle at lift off and desired angle.

Another adopted approach is to change the morphology (stiffness) of the model. Owaki et al. investigated different nonlinear stiffness function and identified the ones that are beneficial

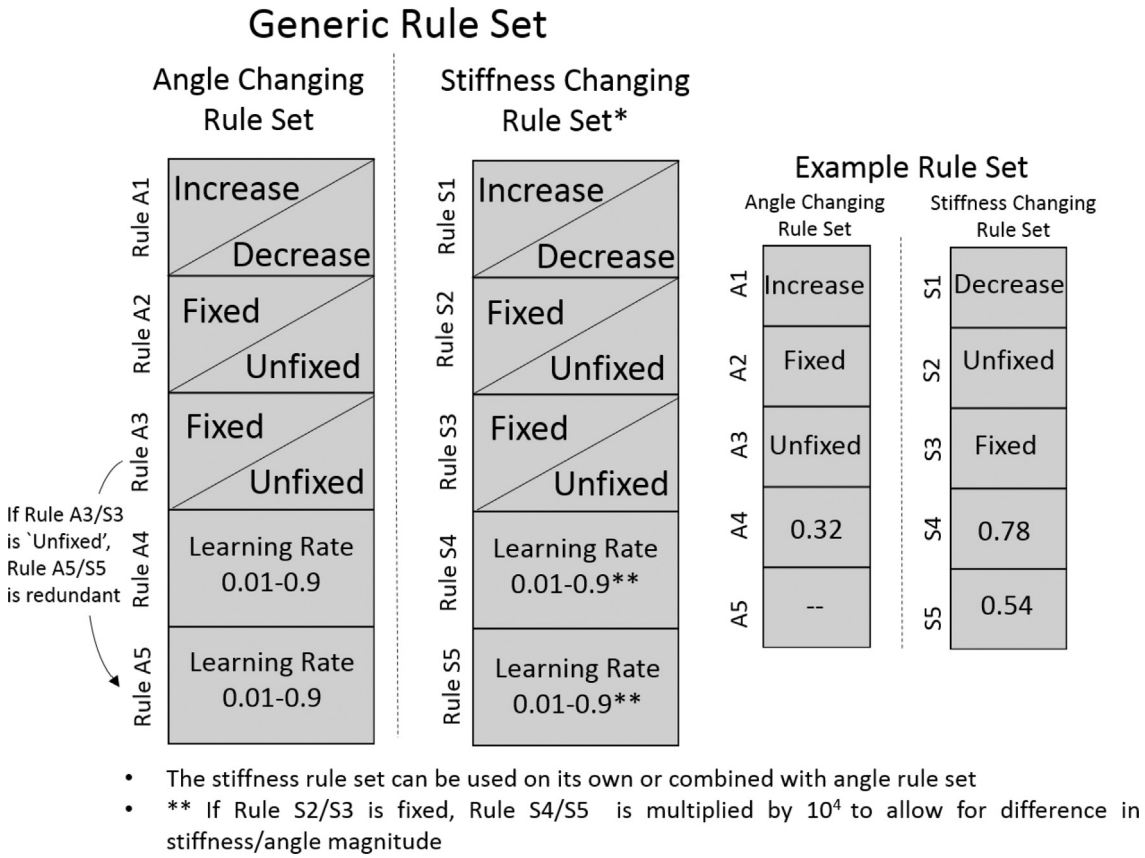

Fig. 2. Generic rule set used throughout this paper. The left side show the set of rules to adapt the attack angle α (rules A1-A5) and the right to adapt stiffness k (rules S1-S5). The figure also shows the range for the different rules. Note that Increase/Decrease and Fixed/Unfixed are binary. Also included in this figure is an example rule set for added clarity.

(Owaki and Ishiguro, 2006). Similarly, Karssen and Wisse (2011) optimized offline the stiffness function with respect to disturbance rejection. Also using increasing the number of leg segments, like in Rummel et al. (2008) and Rummel and Seyfarth (2008), lead to nonlinear stiffness functions and a bigger region of self-stabilization.

Blum et al. (2010) combined both previously mentioned methods by adding control strategies to change both angle and stiffness in the flight phase, ready for touch down. Iida and Tedrake (2007) use reinforcement learning to tune the motor frequency of a one legged robot based on the SLIP model.

While all these approaches are able to widen the region of stable pairs of attack angle and stiffness, they don't allow to learn from unsuccessful starting parameter (stiffness/attack angle) combinations. If these starting parameters are outside the area of self-stability there is no possibility of becoming stable.

We propose a method that allows the SLIP model to adapt its attack angle (control parameter) and stiffness (morphological parameter) based on previous experience. It uses the information from failed runs to guide the change, i.e., how the angle and/or stiffness of the model has to be adapted to eventually get to the stable region. We investigate systematically various rules and evaluate their performance.

In the next section, we briefly explain the SLIP model and described the used methodology. In Section 3 we present results and discuss their implications in the Section 4. Finally, we present a future outlook.

2. Method

2.1. The SLIP model

This SLIP model has been introduced in Blickhan (1989) (and discussed further in McMahon and Cheng (1990)) as generic model of legged locomotion. Historically, the movement of the leg is formed of two phases; the flight phase and the stance phase, as shown in Fig. 1. We follow the simulation layout from Seyfarth et al. (2003).

The leg was simulated in Matlab using the Forward Euler method for integration with a time step of $\Delta t = 0.001$ s. The SLIP model was simulated until either the leg fell over (the vertical height of the mass was less than ground level, i.e. $y < 0$) or the maximum time of $t = 10$ s ($= 10{,}000$ simulation time steps) was achieved. If the leg had not fallen over after 10 s it was considered stable. For this paper, the time from the point the leg is starting to move to when it falls over or it reaches 10 s is called as an *episode*. Therefore, the maximum length of one episode is 10 s. Since, our approaches allows to learn from previous experience, e.g. either from a failed episode of $t < 10$ s or from a successful episode, we

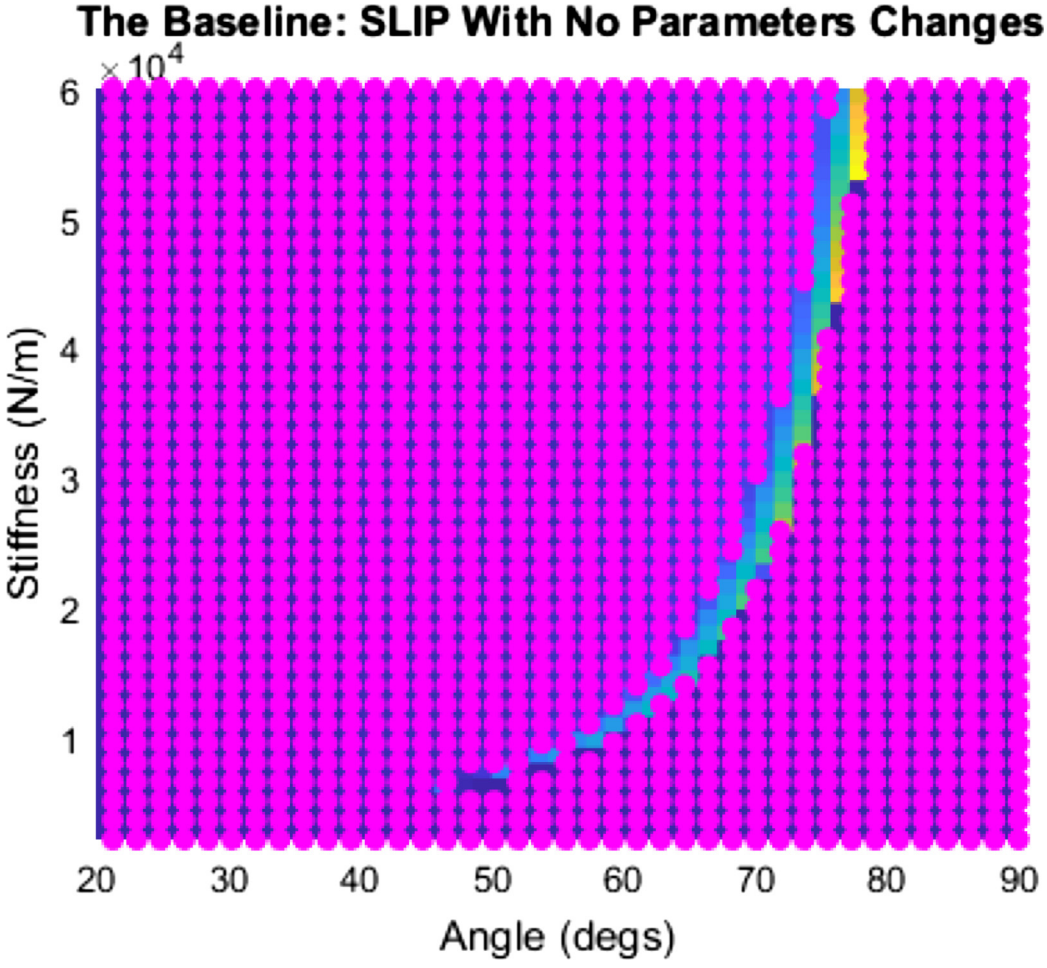

Fig. 3. Figure shows the amount of stable parameter combinations for all possible attack angle and stiffness for basic SLIP model that does not adapt its parameters. A pink dot indicates that no stable solution could be found with this starting position. It can be seen that such a fixed SLIP model has very few stable starting points. The color coding in the background reflects the average traveled distance for this given angle/stiffness combination. The lighter the color the further it was able to locomote. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

assume the model can "get up" again and start with slightly changed parameters (either different attack angle α or stiffness k, or both). We allow the simulation to run a total of 100 episodes to learn from experience. This 100 episodes are referred to as one life time. The resting length of the leg in the simulation was 1 m, the point mass was 80 kg. This is consistent with the values used in literature, see (Seyfarth et al., 2003).

2.2. Rule set design

This section introduces the basic set of rules that was used to learn from one episode to another. Depending on the rules the SLIP model will either change the attack angle α, the stiffness k, or both. We systematically investigate all possible combinations of this basis rule set to obtain the best solutions. The goal was to have a

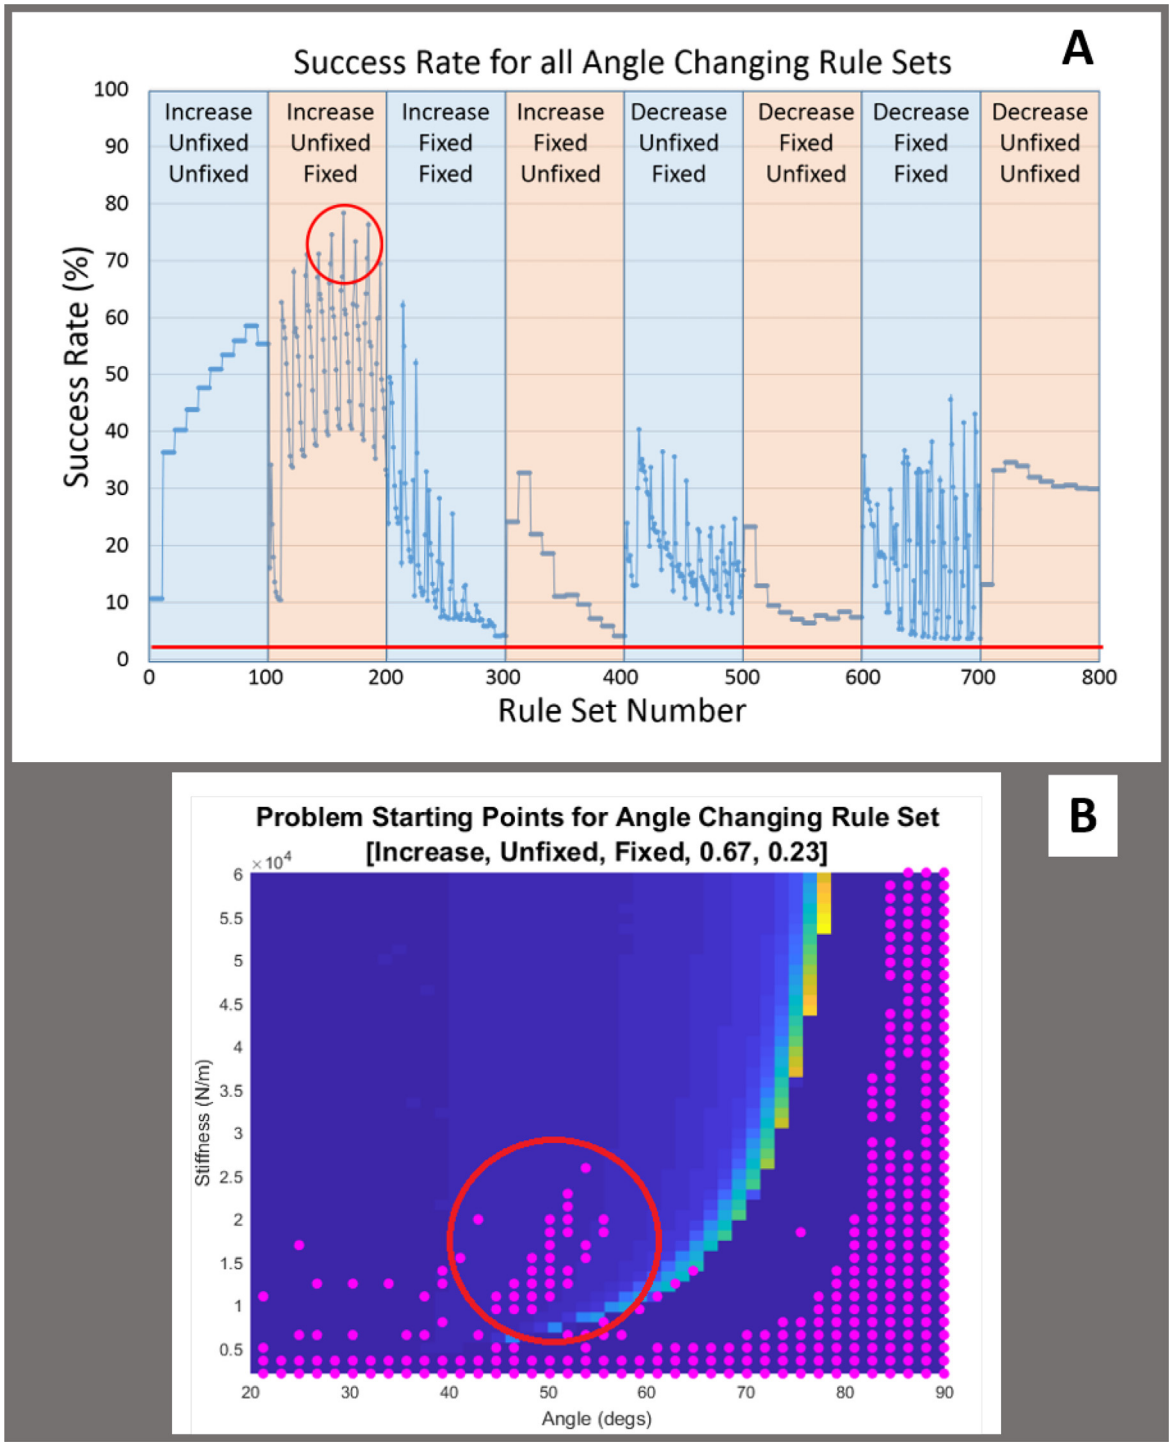

Fig. 4. (A) shows the success rate (as defined in the text) for all 800 tested angle changing rule sets. Each column represents 100 rule sets that share the same first three binary rules. The red circle shows the location of the best sets of rules that are also detailed in Table 4. The red line on the bottom shows the base line success rate (i.e. the success rate for a non-learning model, see Fig. 3). Areas of similar success rates seen in 1st, 4th, 6th and 8th columns are due to the redundancy of the Rule A5 in these cases. (B) Shows the performance of the best angle changing rule set *[Increase, Unfixed, Fixed, 0.67, 0.23]*. A pink dot indicates that no stable solution could be found with this starting position. The red circle highlights an area of unexpected problem starting points discussed in the text. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

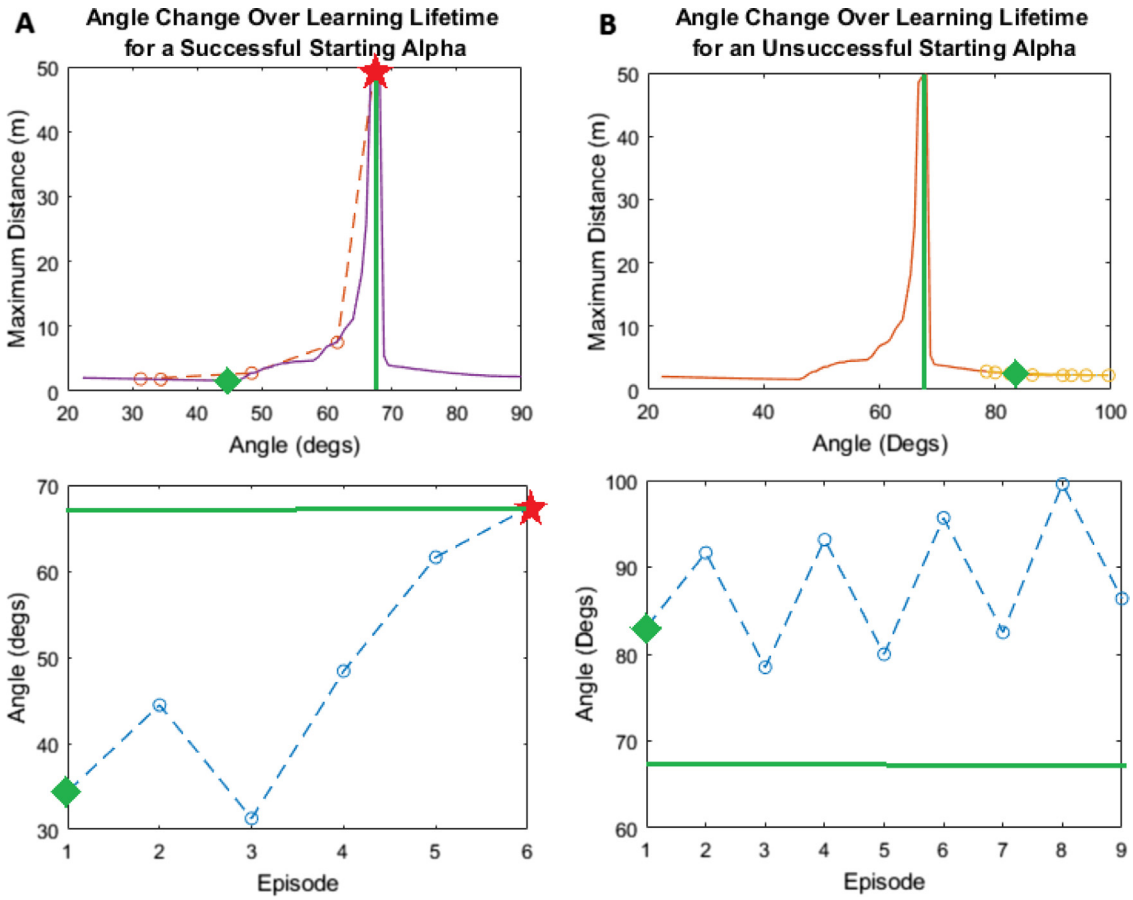

Fig. 5. The graphs show how the top angle changing rule set attempts learn a stable solution. The top figures shows how the rule adapts its angles over multiples steps with respect to the underlying cost function landscape. The bottom graphs show how the angle is changed over time. (A) Shows results for a successful starting angle. (B) Shows results for an unsuccessful starting position (from a pink dot). In both cases the green diamond indicates the start position, a red star the final position. The stiffness in both cases is kept constant at 20,000 N/m. The green lines on the bottom graphs show the angle where the model would be stable. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

SLIP model that can, through learning from failure and corresponding adaptation, reach stable solutions. The generic rule set as well as an example rule set are shown in Fig. 2.

The sets of rules for exploring the adaptation of attack angle and stiffness are similar. Each set consists of five individual rules. At the end of each episode the adaptation is based on the achieved distance D_i and the distance achieved in the previous episode D_{i-1}. Note that for a successful run the achieved distance is the distance the SLIP model has traveled after 10 s. For an unsuccessful run it is the distance traveled before model has fallen over.

The key parameter for the adaptation process is the difference between these two distances, i.e., $\Delta D = D_i - D_{i-1}$.

The first rule dictates whether, upon a positive ΔD, the parameter (either α or k) should increase or decrease. Symmetry is employed, so if ΔD is negative the opposite occurs. Rule 2 dictates whether, when the parameter is to *increase*, this increase should be dependent on ΔD or a fixed value. Similarly, Rule 3 dictates whether, when the parameter is to *decrease*, this amount should be fixed or unfixed.

If the parameter change is *unfixed* (regardless of whether it is increasing/decreasing or) the algorithm uses following Equation to calculate the parameter change.

$$\text{parameter change} = \frac{\mu}{2 + \Delta D}, \tag{1}$$

where $\Delta D = D_i - D_{i-1}$ and μ is given by A4/S4. Note that if A3/S3 is unfixed, A5/S5 is redundant. The parameter change is then added or removed from the current parameter depending on Rules 2/3.

If a fixed value is to be used Rules 4 and 5 dictate this value, 4 if the new parameter configuration has performed better, 5 if it has performed worse.[1]

If the parameter is changed according to the rules to a value that is outside of the range ($< 20°, > 90°$ or a negative stiffness) the learning process is terminated (the lifetime is finished) and the rule set is determined unsuitable.

Based on these basic rules, we investigated all possible combinations. Considering that Rules A1, A2, and A3 are binary (i.e., either fixed/unfixed or increase/decrease) and we explored Rule A4 and Rule A5 for values from 0.01 to 0.9 in 10 discretized steps, we investigated at total of $2^3 * 10 * 10 = 800$ combinations (see for

Table 1
Top 5 successful angle changing rule sets and their corresponding success rates. These top rules can also be seen in Fig. 4 within the red circle.

Rule set	Success rate
Increase, Unfixed, Fixed, 0.67, 0.23	78.93%
Increase, Unfixed, Fixed, 0.78, 0.23	73.91%
Increase, Unfixed, Fixed, 0.45, 0.12	71.12%
Increase, Unfixed, Fixed, 0.34, 0.12	71.12%
Increase, Unfixed, Fixed, 0.56, 0.12	69.31%

[1] Note that varying stiffness rules S4 and S5 are multiplied by a factor of 10^4 to allow for the difference in magnitude between stiffness and attack angle.

example Fig. 4A). The same number of different rule set were explored for the stiffness adaptation using Rules S1-S5.

For each of the 800 combinations we systematically tested a wide range of possible pairs of attack angles α and stiffness values k. The range encompassed for α, values from 20° to 90° (in 40 discrete steps) and for k from 2000 to 60,000 N/m (also in 40 discrete steps). This results in 1600 different angle/stiffness pairs. This range was selected as it fully encompasses the stable region (determined by Seyfarth et al. (2002)), but still includes a large enough area where stability is not available with the basic SLIP model. For a given set of rules we simulated the model with all 1600 different starting pairs and observed if the system was able to achieve stable locomotion through applying the rule over a lifetime (i.e., over 100 episodes).

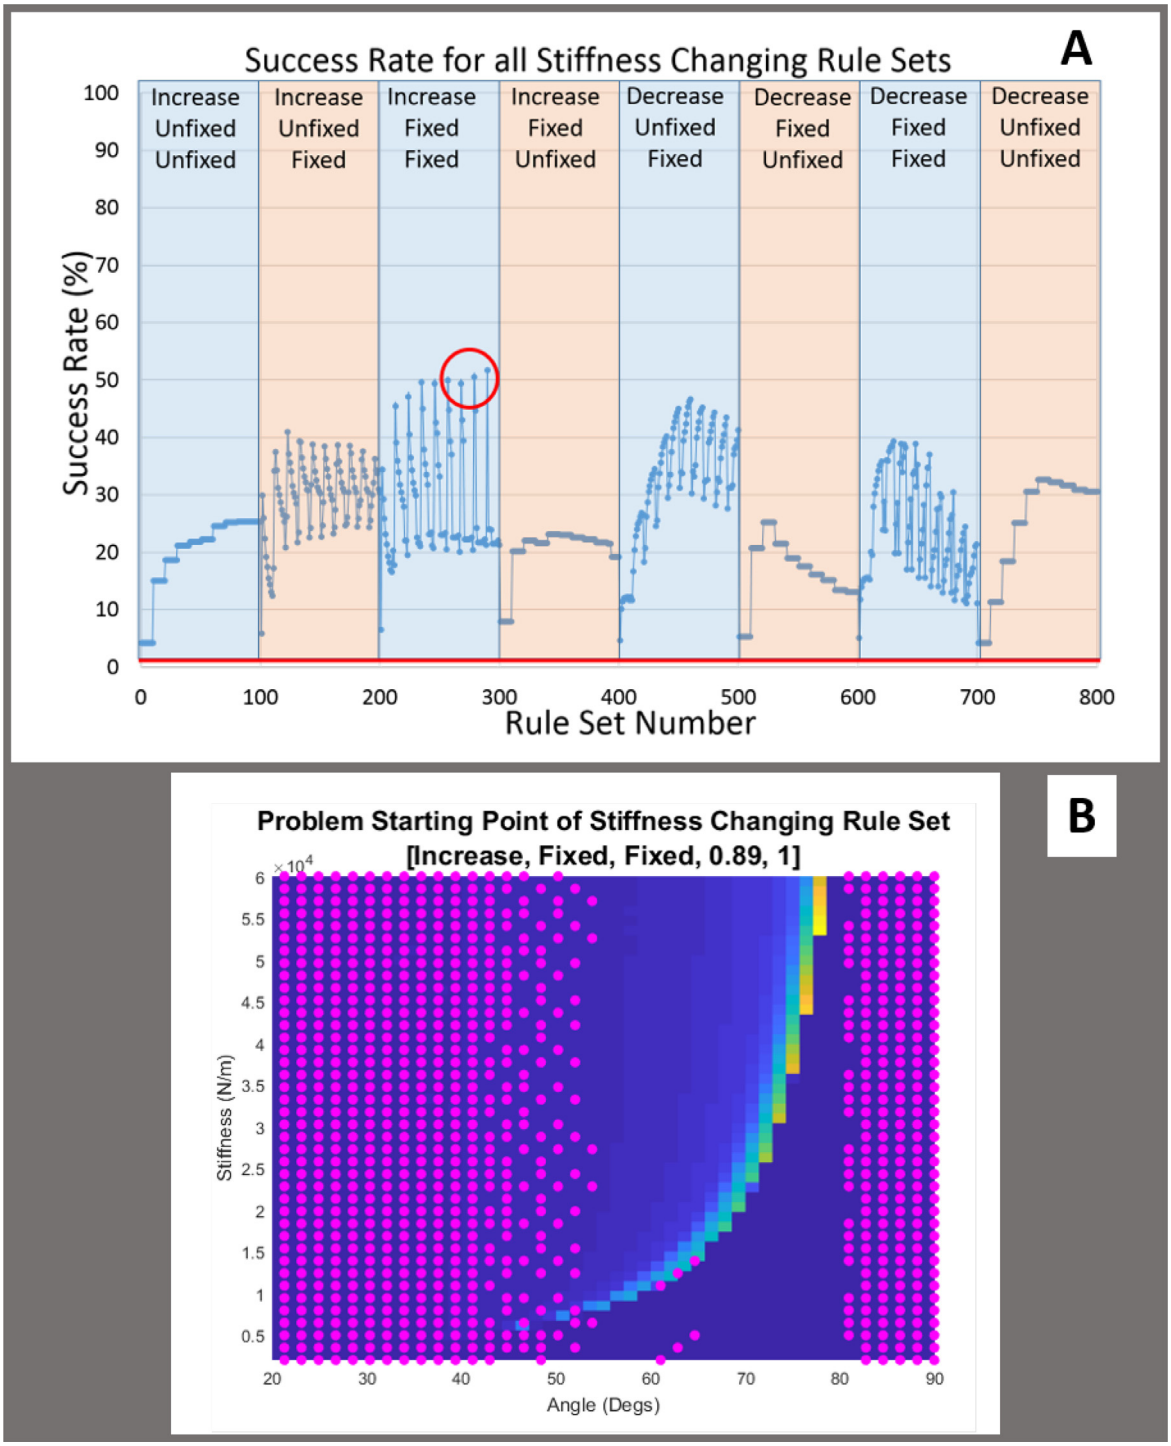

Fig. 6. (A) shows the success rate (as defined in the text) for all 800 tested stiffness changing rule sets. Each column represents 100 rule sets that share the same first three binary rules. The red circle shows the location of the best sets of rules that are also detailed in Table 4. The red line on the bottom shows the base line success rate (i.e. the success rate for a non-learning model, see Fig. 3). Areas of similar success rates seen in 1st, 4th, 6th and 8th columns are due to the redundancy of the Rule S5 in these cases. (B) Performance of the best angle changing rule set *[Increase, Fixed, Fixed, 0.89, 1]*. A pink dot indicates that no stable solution could be found with this starting position. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

To assess whether a rule set is good or not, a *success rate* was defined. The *success rate* is number of points out of all 1600 possible starting points that lead to a stable solution (in percentage). For example, a rule set always able to find a stable solution would have a success rate of 100%. Since we use the same number of simulations for all rules, we can quantitatively compare them within the explored parameter range.

3. Results

3.1. The baseline: SLIP with no parameter changes

First, we simulated the standard SLIP model (i.e., no online adaptation of parameters) to establish a baseline. We tested the model over the previously described 1600 angle and stiffness combinations. Since no learning was included the stiffness and angle parameters were unable to change.

Fig. 3 shows the performance. A pink dot indicates a starting condition where a stable solution could not be obtained throughout the model's life time. The figure clearly shows that a standard SLIP model only a small set of stiffness/angle combinations are leading to stable locomotion. This area of stability is in literature often referred as J-Figure (e.g., Seyfarth et al., 2002.) The success rate for the non-learning SLIP model was 3.75% (60 out 1600 parameter pairs were stable).

3.2. Adapting the attack angle

This section describes the result from testing rules that change the attack angle α. Every combination (a total of 800) was tested and the results in form of success rates are presented in Fig. 4. Table 1 summarizes the top 5 rule sets (they are also highlighted in red in Fig. 4A). The best rule obtained a success rate of 78.93%, which is much higher compared to the standard SLIP model with 3.75% (the corresponding base line is the red line in Fig. 4A).

The top scoring rule sets from Table 1 all follow a similar learning mechanism. If the current distance D_i is larger than the previously achieved distance D_{i-1} the angle of attack is increased. The amount is inversely proportional to δD. This allows the agent to climb the gradient of the SLIP model without overshooting the global optimum (which is in the J-Figure). Accordingly, if the leg performance was worse than the previous run, i.e., $\Delta D < 0$, the angle was reduced by a fixed amount.

Fig. 4A also shows that any combination of our proposed rules for changing the attack angle is performing better than the default SLIP model, i.e., all success rates are higher than the base line in red. Similar to the Fig. 3 we tested all starting combinations of α and k for the best rule set, see Fig. 4B. It can been seen clearly that the region that leads to stable solutions is much bigger (success rate 78.93%) compared to the standard SLIP (success rate 3.75%, see Fig. 3). But there are also some limitations visible in Fig. 4B. It seems for starting points left to the J-figure (i.e., smaller angles) most of the points reach a stable solution (i.e., not a pink dot), while on the right side (higher angle values) the performance is limited. To further investigate that we looked at starting parameters specifically on the right side and on the left side and observed the change of α over time. If the starting angle is on the right hand side of the J-Figure, the learning mechanism will zig-zag away from the optimal region.

By looking at the underlying gradient we can see that either side of the peak stability region the gradient is negative. Consequently, when on the right side of the solution the update rules will lead away from the stable region, see Fig. 5B.

Given this gradient arrangement, a rule set that would be able to distinguish between "left" and "right" would perform even better. However, our proposed setup did not allow for this, because we assume the system does not have the information in which direction the stable region is located. Another problem starting point area is shown by the red circle in Fig. 4B. By looking at the development of the attack angle over time we observed oscillation between an angle slightly below the stable region and one slightly above it. Interestingly, this behaviour is unique to the top rule set and is not observed in the other top 4 rules listed in Table 1. However, the other rules have lower success rates due more unstable points on the right-hand side (i.e., at higher attack angles).

3.3. Adapting the stiffness

In addition to the angle changing rule set testing, we also tested 800 stiffness changing rule. Fig. 6A summarizes the obtained success rates. Table 2 shows the top 5 stiffness changing rule sets.

For the best stiffness adaptation rule we plotted the region of success starting points, see Fig. 6B. The corresponding success rate was 51%. One can see that the parameter regions that lead to stable solutions is much bigger than for the SLIP without learning (compare Fig. 3, success rate 3.75%), but it's smaller compared to the best angle adaptation rule (see Fig. 4, success rate 78.93%). Again, any of the tested adaptation rules performs better than the standard SLIP model with fixed parameters (compare red baseline in Fig. 4A).

As with the angle changing rule sets, the top five stiffness changing rule sets all follow the same three starting rules. The stiffness value is increased if the leg performs better than in the previous episode $\Delta > 0$, and decreased if $\Delta D < 0$. Both the increase and decrease of the stiffness is fixed regardless of the magnitude of ΔD.

In contrast to the attack angle adaptation rule sets, where one rule seems to stand out with respect to performance, for stiffness adaptations there seems be various rule sets with the different first three rules (for example *[Decrease, Fixed, Fixed]*, and even *[Decrease, Unfixed, Fixed]*, see Fig. 6A) that are also very successful.

3.4. Adapting and attack angle and stiffness simultaneously

Finally, the adaptation of both, attack angle and stiffness, was investigated. Due to the high number of possible combinations (640,000) not every single possible rule set was tested. In fact, it would take just under 1 year of simulation time to complete all this testing. Instead a simpler approach was taken.

First, a stiffness and angle changing rule set was created by combining the best angle changing rule set (Table 1) with the best stiffness changing rule set (Table 2), i.e. *[Increase, Unfixed, Unfixed, 0.67, 0.23]* for adapting the attack angle, and *[Increase, Fixed, Fixed, 0.89, 1]* for adapting the stiffness. Fig. 7 shows the corresponding performance. The success rate was 75.56%, which is a little bit lower than the best angle adaptation rule (78.93%), but higher than the best stiffness adaptation rule (51.0%).

Although the overall success rate of the combined rule set is lower than the top angle changing rule set, an advantage of this combination is that the problem area shown in Fig. 7 does not show here. A disadvantage of this combination is the region high-

Table 2
Top 5 successful stiffness changing rule set and their corresponding success rates. These top rules can also be seen in Fig. 6 within the red circle.

Rule set	Success rate
Increase, Fixed, Fixed, 0.89, 1.00	51.0%
Increase, Fixed, Fixed, 0.78, 0.89	50.5%
Increase, Fixed, Fixed, 0.56, 0.67	49.9%
Increase, Fixed, Fixed, 0.45, 0.56	49.3%
Increase, Fixed, Fixed, 0.67, 0.78	49.3%

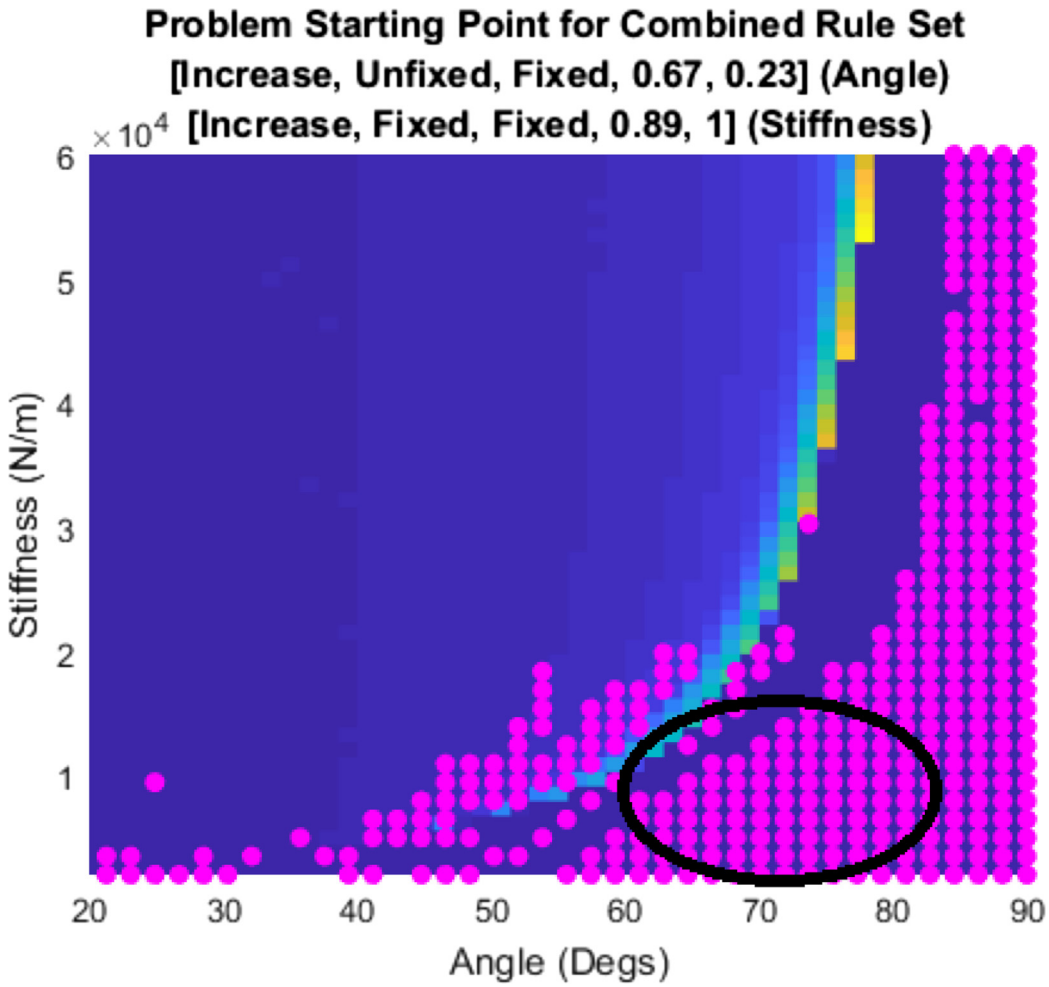

Fig. 7. Performance when the best stiffness and angle changing rule sets are combined, i.e., *[Increase, Unfixed, Fixed, 0.67, 0.23 (angle) Increase, Fixed, Fixed, 0.89, 1 (stiffness)]*. The black circle highlights an interesting area of unexpected problem starting points. See text for discussed.

Table 3
Table showing top 5 of the 25 different rule set combinations tested.

Angle rule rank	Stiffness rule rank	Success rate
1	4	76.31%
1	3	76.31%
1	2	76.25%
1	5	75.81%
1	1	75.56%

lighted by the black circle in Fig. 7, which seems to emergence because of the combination. The approach of simply combing the best angle attack and stiffness adaptation rule sets is rather naive. Potentially, there are other rule sets that perform better. Due to the huge search space, however, we are constrained. However, we tested all combinations of our top 5 rules (i.e. 25 combinations). The summary of the top five combinations can be found in Table 3. Interestingly, they all use the best angle rule set, but use inferior stiffness rule sets to achieve higher success rates.

4. Discussion and further work

The paper shows that by allowing the SLIP model to learn from previous experience, it is able to recover from a much wider range of possible starting combinations of control parameter (attach angle) and morphological parameter (spring stiffness). This add a new level of adaptivity. While in this work the environment was fixed, it is clear that having this capability to learn, can enable SLIP-model-based machines to deal with a wide range of changing environments. In addition, it would allow to cope with morphological changes, e.g., wearing off stiffness in the locomotion system due to age. In this paper we only explored a small number of possible rule sets with a brute force approach. Our results in combining adaptation of angles and stiffness suggest however that even in this small set of rules, potentially, there is still numerous combinations that might lead to even better performing systems. Nonlinear optimization approaches, like Genetic Algorithms, could be used to explore this big parameter space. Clearly, extending the rules with additional features will potentially improve performance as well. Both will be part of future work.

Conflict of interests

We, the authors, declare that we have no financial or personal relationships with other people or organizations that could inappropriately influence (bias) our work.

Acknowledgements

This work was supported by the EPSRC Centre for Doctoral Training in Future Robotics and Autonomous Systems (FARSCOPE)

EP/L015293/1. The research was also partly supported by the Leverhulme Trust Research Grant RPG-2016-345.

References

Alexander, R.M., 2003. Principles of Animal Locomotion. Princeton University Press.

Blickhan, R., 1989. The spring-mass model for running and hopping. J. Biomech. 22 (11-12), 1217–1227.

Blum, Y., Lipfert, S.W., Rummel, J., Seyfarth, A., 2010. Swing leg control in human running. Bioinspiration Biomimetics 5 (2), 026006.

Cham, J.G., Karpick, J.K., Cutkosky, M.R., 2004. Stride period adaptation of a biomimetic running hexapod. Int. J. Robot. Res. 23 (2), 141–153.

Dickinson, M.H., Farley, C.T., Full, R.J., Koehl, M., Kram, R., Lehman, S., 2000. How animals move: an integrative view. Science 288 (5463), 100–106.

Ernst, M., Geyer, H., Blickhan, R., 2010. Spring-legged locomotion on uneven ground: a control approach to keep the running speed constant. In: Mobile Robotics: Solutions and Challenges. World Scientific, pp. 639–644.

Geyer, H., Blickhan, R., Seyfarth, A., 2002. Natural dynamics of spring-like running: emergence of self-stability. In: 5th International Conference on Climbing and Walking Robots. Professional Engineering Publishing Ltd, Suffolk, England, pp. 87–91.

Hauser, H., Ijspeert, A.J., Füchslin, R.M., Pfeifer, R., Maass, W., 2011. Towards a theoretical foundation for morphological computation with compliant bodies. Biol. Cybernet. 105 (5–6), 355–370.

Hauser, H., Ijspeert, A.J., Füchslin, R.M., Pfeifer, R., Maass, W., 2012. The role of feedback in morphological computation with compliant bodies. Biol. Cybernet. 106 (10), 595–613.

Hauser, H., Füchslin, R.M., Pfeifer, R., 2014. Opinions and outlooks on morphological computation.

Hurst, J.W., Rizzi, A.A., 2005. Physically variable compliance in running. In: Climbing and Walking Robots. Springer, pp. 123–133.

Iida, F., Tedrake, R., 2007. Motor control optimization of compliant one-legged locomotion in rough terrain. In: Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on. IEEE, pp. 2230–2235.

Karssen, J.D., Wisse, M., 2011. Running with improved disturbance rejection by using non-linear leg springs. Int. J. Robot. Res. 30 (13), 1585–1595.

McMahon, T.A., Cheng, G.C., 1990. The mechanics of running: how does stiffness couple with speed? J. Biomech. 23, 65–78.

Owaki, D., Ishiguro, A., 2006. Enhancing self-stability of a passive dynamic runner by exploiting nonlinearity in the leg elasticity. In: SICE-ICASE, 2006. International Joint Conference. IEEE, pp. 4532–4537.

Rummel, J., Seyfarth, A., 2008. Stable running with segmented legs. Int. J. Robot. Res. 27 (8), 919–934.

Rummel, J., Iida, F., Smith, J.A., Seyfarth, A., 2008. Enlarging regions of stable running with segmented legs. In: 2008 IEEE International Conference on Robotics and Automation, pp. 367–372. https://doi.org/10.1109/ROBOT.2008.4543235.

Schmitt, J., 2006. A simple stabilizing control for sagittal plane locomotion. J. Comput. Nonlinear Dyn. 1 (4), 348–357.

Seyfarth, A., Geyer, H., 2002. Natural control of spring-like running–optimized self-stabilization. In: Proceedings of the Fifth International Conference on Climbing and Walking Robots. Professional Engineering Publishing Limited, pp. 81–85.

Seyfarth, A., Geyer, H., Günther, M., Blickhan, R., 2002. A movement criterion for running. J. Biomech. 35 (5), 649–655.

Seyfarth, A., Geyer, H., Herr, H., 2003. Swing-leg retraction: a simple control model for stable running. J. Exp. Biol. 206 (15), 2547–2555.

Vanderborght, B., Albu-Schäffer, A., Bicchi, A., Burdet, E., Caldwell, D.G., Carloni, R., Catalano, M., Eiberger, O., Friedl, W., Ganesh, G., et al., 2013. Variable impedance actuators: a review. Robot. Auton. Syst. 61 (12), 1601–1614.

Vu, H.Q., Hauser, H., Leach, D., Pfeifer, R., 2013. A variable stiffness mechanism for improving energy efficiency of a planar single-legged hopping robot. In: 2013 16th International Conference on Advanced Robotics (ICAR). IEEE, pp. 1–7.

Wieber, P.-B., Tedrake, R., Kuindersma, S., 2016. Modeling and control of legged robots. In: Springer Handbook of Robotics. Springer, pp. 1203–1234.

Wolf, S., Grioli, G., Eiberger, O., Friedl, W., Grebenstein, M., Höppner, H., Burdet, E., Caldwell, D.G., Carloni, R., Catalano, M.G., et al., 2016. Variable stiffness actuators: review on design and components. IEEE/ASME Trans. Mech. 21 (5), 2418–2430.