Walker, K., & Hauser, H. (2019). Adapting stiffness and attack angle through trial and error to increase self-stability in locomotion. *Journal of Biomechanics*, 87, 28-36.
<https://doi.org/10.1016/j.jbiomech.2019.02.009>

Peer reviewed version

License (if available):
CC BY-NC-ND

Link to published version (if available):
[10.1016/j.jbiomech.2019.02.009](https://doi.org/10.1016/j.jbiomech.2019.02.009)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Elsevier at <https://www.sciencedirect.com/science/article/pii/S0021929019301393?via%3Dihub>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Adapting Stiffness and Attack Angle through Learning Rule Sets to Increase Self-Stability in Locomotion

Kathryn Walker^{1,2,*}, Helmut Hauser^{1,2}

¹*University of Bristol, Bristol, BS8 1TH*

²*Bristol Robotics Laboratory, Stoke Gifford, Bristol BS16 1QY*

kw16876@bristol.ac.uk

** Corresponding Author*

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.

Keywords: Learning; Legged locomotion; Trial and Error; SLIP Model;

Control; Morphology

Word Count: 3229

12 1. Introduction

13 Despite the huge success of robotics in general, there are almost no ma-
14 chines that are capable to stably locomote in rough, unknown terrain. As a
15 consequence, most of today’s robots work only in well-defined environments like
16 assembly lines and on the factory floor. Even state-of-the-art robots are still
17 not able to deal with the uncertainty, complexity and variety typical of natu-
18 ral environments and our living and working places. Simply put, if they can’t
19 model it, they’ll fail. On the other hand, animals, including us humans, are
20 particularly good at locomotion and we outperform robots in almost any cate-
21 gory, including, energy efficiency, stability, robustness, agility, and many others
22 [1][24]. This is partly due to the fact that a lot of the locomotion process hap-
23 pens without the need of an exact model of the environment. For example, part
24 of self-stabilisation can be carried out by the morphological structure of the
25 legs. One could say control is outsourced to the morphology. In robotics, this
26 is often referred to as morphological computation [9, 10, 8]. Another important
27 difference is that animals are remarkable adaptive. They can locomote over a
28 wide range of different terrains [23, 11, 5, 4]. Key properties are their capability
29 of learning from previous experience and their capacity to adapt, i.e., change
30 their control strategy and/or change their morphological properties, e.g., like
31 adapting stiffness). While numerous mechanisms have been proposed to adapt
32 stiffness, typically referred to as variable compliant mechanisms [22, 25], only
33 few have been used to actively increase locomotion performance. For example,
34 Quy et al. used adaption of stiffness to improve energy efficiency in single-legged
35 hopper [23].

36 In this work we propose an approach that implements these learning and adapta-
37 tion capabilities on the well-established spring-loaded inverted pendulum model
38 (SLIP) model, which is a prevalent model for analyzing running and hopping
39 [2]. This simple model is surprisingly general as it works over a wider range of
40 species and locomotion types. In its basic form it describes the action of the
41 leg by representing it as a lossless, linear spring of constant stiffness k and rest

length l_o as shown in Figure 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 losing its periodic locomotion pattern[6]. Seyfarth et al. showed the range of stiffness and angle combinations that resulted in self-stability for a simple SLIP model [20, 7].

Many 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. [21, 19] 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 [18] 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 [15]. Similiary, Karssen et al. [13] optimized offline the stiffness function with respect to disturbance rejection. Also using increasing the number of leg segments, like in [16] and [17], lead to nonlinear stiffness functions and a bigger region of self-stabilization.

Blum et al. [3] combined both previously mentioned methods by adding control strategies to change both angle and stiffness in the flight phase, ready for touch down. Iida et al. [12] 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 change in kinetic energy between strides through use of learning rules. This approach has a great potential to be used in variety of applications where learning/adaption over time is necessary. This ranges from locomotion rehabilitation devices to helping stroke patients, to supportive walking devices for elderly, to adaptive,

73 autonomous, legged robots. In all these cases, there is a need for learning to
 74 adapt to changes in either the human user (e.g. due to improvement through
 75 training or deterioration by aging) or the robot and environment (due to wear
 76 or changes in the ground dynamics) .
 77 In the next section, we briefly explain the SLIP model and described the used
 78 methodology. In Section 3 we present results and discuss their implications in
 79 the Section 4. Finally, we present a future outlook.

80 2. Method

81 2.1. The SLIP model

82 **Insert Figure 1 here**

83 This SLIP model has been introduced in [2] (and discussed further in [14]) as
 84 generic model of legged locomotion. Historically, the movement of the leg is
 85 formed of two phases; the flight phase and the stance phase, as shown in Figure
 86 1. We follow the simulation layout from [21].

87 The leg was simulated in Matlab using the Forward Euler method for inte-
 88 gration with a time step of $\Delta t = 0.001$ seconds. The SLIP model was simulated
 89 until either the leg fell over (the vertical height of the mass was less than ground
 90 level, i.e. $y < 0$) or the maximum time of $t = 10$ seconds ($= 10'000$ simulation
 91 time steps) was achieved. If the leg had not fallen over after 10 seconds it was
 92 considered stable. For this paper, the time from the point the leg is starting
 93 to move to when it falls over or it reaches 10 seconds is called as an *episode*.
 94 Therefore, the maximum length of one episode is 10 seconds. Note that po-
 95 tentially there are parameter combinations that could lead to falling over after
 96 10 seconds. However, we found that 10 seconds for the duration of an episode
 97 was a reasonable trade-off between simulation time and stability. Since, our
 98 approaches allows to learn from previous experience, e.g. either from a failed
 99 episode of $t < 10$ seconds or from a successful episode, we assume the model
 100 can "get up" again and start with slightly changed parameters (either different
 101 attack angle α or stiffness k , or both). We allow the simulation to run a total of

102 100 episodes to learn from experience. This 100 episodes are referred to as one
103 life time. Similiarly to our choice for the duration of one episode, a life time of
104 100 episodes was found, after some testing, to be a reasonable trade-off between
105 simulation time and stability. The resting length of the leg in the simulation
106 was 1 m, the point mass was 80 kg. This is consistent with the values used in
107 literature, see [21].

108 2.2. Rule Set Design

109 This section introduces the basic set of rules that was used to learn from one
110 episode to another. Depending on the rules the SLIP model will either change
111 the attack angle α , the stiffness k , or both. We systematically investigate all
112 possible combinations of this basis rule set to obtain the best solutions. The
113 goal was to have a SLIP model that can, through learning from failure and
114 corresponding adaptation, reach stable solutions. The generic rule set as well
115 as an example rule set are shown in Figure 2.

116 **Insert Figure 2 here**

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

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

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

131 If the parameter change is *unfixed* (regardless of whether it is increasing/decreasing

132 or) the algorithm uses following Equation to calculate the parameter change.

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

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

136 If a fixed value is to be used Rules 4 and 5 dictate this value, 4 if the new
 137 parameter configuration has performed better, 5 if it has performed worse¹.

138 If the parameter is changed according to the rules to a value that is outside of
 139 the range ($< 0^\circ$, $> 90^\circ$ or a negative stiffness) the learning process is terminated
 140 (the lifetime is finished) and the rule set is determined unsuitable.

141 Based on these basic rules, we investigated all possible combinations. Con-
 142 sidering that Rules A1, A2, and A3 are binary (i.e., either fixed/unfixed or
 143 increase/decrease) and we explored Rule A4 and Rule A5 for values from 0.01
 144 to 0.9 in 10 discretized steps, we investigated at total of $2^3 * 10 * 10 = 800$
 145 combinations (see for example Figure 4A). The same number of different rule
 146 set were explored for the stiffness adaptation using Rules S1-S5.

147 For each of the 800 combinations we systematically tested a wide range of
 148 possible pairs of attack angles α and stiffness values k . The range encompassed
 149 for α , values from 20° to 90° (in 40 discrete steps) and for k from 2,000 to 60,000
 150 N/m (also in 40 discrete steps). This results in 1,600 different angle/stiffness
 151 pairs. Note that this range (i.e., lower and upper bounds) has been chosen
 152 pragmatically; first, to fully encompassed the stability region as determined
 153 by [20] and, second, to explore enough interesting parameter combination to
 154 demonstrate the validity as well the limitations of our approach. For a given
 155 set of rules we simulated the model with all 1,600 different starting pairs and
 156 observed if the system was able to achieve stable locomotion through applying
 157 the rule over a lifetime (i.e., over 100 episodes).

¹Note 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.

158 To assess whether a rule set is good or not, a *success rate* was defined. The
 159 *success rate* the percentage out of all 1,600 possible starting points that lead
 160 to a stable solution (in percentage) applying the learning rules. For example, if
 161 a rule set would be able to find a stable solution for all 1,600 combination, it
 162 would have a success rate of 100%. Since we use the same number of simulations
 163 for all rules, we can quantitatively compare them within the explored parameter
 164 range. Note that the *success rate* gives us a comparative measurement on the
 165 usefulness of any given set of learning rules. The higher the *success rate*, the
 166 bigger is the region of parameter combinations that our set of learning rules can
 167 recover from, i.e. it can lead successfully the system through learning to the
 168 stable region in the parameter space. This behaviour is particularly useful if
 169 we have to deal with real systems that change their dynamics, e.g. changes in
 170 stiffness due to age or failure.

171 3. Results

172 3.1. The Baseline: SLIP with No Parameter Changes

173 First, we simulated the standard SLIP model (i.e., no online adaptation of
 174 parameters) to establish a baseline. We tested the model over the previously
 175 described 1,600 angle and stiffness combinations. Since no learning was included
 176 the stiffness and angle parameters were unable to change.

177 Figure 3 shows the performance. A pink dot indicates a starting condition
 178 where a stable solution could not be obtained throughout the model's life time.
 179 The figure clearly shows that a standard SLIP model only a small set of stiff-
 180 ness/angle combinations are leading to stable locomotion. This area of stability
 181 is in literature often referred as J-Figure (e.g., [20].) The success rate for the
 182 non-learning SLIP model was 3.75% (60 out 1600 parameter pairs were stable).

183 **Insert Figure 3 here**

184

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 Figure 4. Table 1 summarizes the top 5 rule sets (they are also highlighted in red in Figure 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 Figure 4A).

Insert Table 1 here

Insert Figure 4 here

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 the previous run, i.e., $\Delta D < 0$, the angle was reduced by a fixed amount.

Figure 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 Figure 3 we tested all starting combinations of α and k for the best rule set, see Figure 4B. It can be 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 Figure 3). But there are also some limitations visible in Figure 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

216 stability region the gradient is negative. Consequently, when on the right side
217 of the solution the update rules will lead away from the stable region, see Figure
218 5B.

219 **Insert Figure 5 here**

220 Given this gradient arrangement, a rule set that would be able to distinguish
221 between "left" and "right" would perform even better. However, our proposed
222 setup did not allow for this, because we assume the system does not have the
223 information in which direction the stable region is located. Another problem
224 starting point area is shown by the red circle in Figure 4B. By looking at the
225 development of the attack angle over time we observed oscillation between an
226 angle slightly below the stable region and one slightly above it. Interestingly,
227 this behaviour is unique to the top rule set and is not observed in the other top
228 4 rules listed in Table 1. However, the other rules have lower success rates due
229 more unstable points on the right-hand side (i.e., at higher attack angles).

230

231 *3.3. Adapting the Stiffness*

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

235 **Insert Table 2 here**

236 **Insert Figure 6 here**

237 For the best stiffness adaptation rule we plotted the region of success starting
238 points, see Figure 6B. The corresponding success rate was 51%. One can see
239 that the parameter regions that lead to stable solutions is much bigger than
240 for the SLIP without learning (compare Figure 3, success rate 3.75%), but it's
241 smaller compared to the best angle adaptation rule (see Figure 4, success rate
242 78.93%). Again, any of the tested adaptation rules performs better than the
243 standard SLIP model with fixed parameters (compare red baseline in Figure
244 4A).

245 As with the angle changing rule sets, the top five stiffness changing rule sets all

246 follow the same three starting rules. The stiffness value is increased if the leg
 247 performs better than in the previous episode $\Delta > 0$, and decreased if $\Delta D <$
 248 0. Both the increase and decrease of the stiffness is fixed regardless of the
 249 magnitude of ΔD .

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

255 3.4. Adapting and Attack Angle and Stiffness Simultaneously

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

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

267 **Insert Figure 7 here**

268 Although the overall success rate of the combined rule set is lower than the
 269 top angle changing rule set, an advantage of this combination is that the the
 270 problem area shown in Figure 9 does not show here. A disadvantage of this com-
 271 bination is the region highlighted by the black circle in Figure 7, which seems
 272 to emergence because of the combination. In this case the stiffness and angle
 273 change work against each other. The changes they introduce have opposite ef-
 274 fects resulting in a zig-zag pattern in a direction perpendicularly away from the
 275 area of stability. This suggests that the approach of simply combing the best

angle attack and stiffness adaptation rule sets is rather naive. Nevertheless, potentially, there are other combinations that perform better. Due to the huge search space we only tested combinations of all 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.

Insert table 3 here

4. Discussion and Further Work

This paper demonstrates 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 adds a new level of adaptivity with the great potential to use it in a wide range of applications. For example, the approach could be employed in rehabilitation devices to help stroke patients, where stable walking is often a challenge. Another application could be supportive walking devices (e.g. exoskeleton) for elderly who are particularly prone to falling over. Having a SLIP-model-based machine capable of adapting itself over time (i.e during a training phase supervised by a medical professional) would significantly improve the performance. Furthermore, it would be able to adapt itself over time in case of degenerative diseases, where locomotion capabilities of the user naturally decreases over time. In addition, the same approach would be beneficial in the context of autonomous locomotion for robotic systems, e.g., in changing environments or changes in the dynamics due to extensive use.

Our results suggest that updating the angle is more successful in achieving robust locomotion than changing the stiffness. This is due to the fact that the angle rules establish a much steeper gradient close to the region of stability, compared to the one based on stiffness. Therefore, changing the angle has a bigger, positive effect (i.e., getting faster to the region of stability). This be-

305 came particularly obvious when we combined the top stiffness and the top angle
 306 changing rule sets. Furthermore, even adding small angle change to any given
 307 stiffness rule improved the success rate significantly.
 308 In this work we have only explored a small number of possible rule sets and
 309 combination with a *brute force* approach. While the presented results are en-
 310 couraging, various nonlinear optimization approaches, like Genetic Algorithms,
 311 clearly could be used to explore a much bigger space of combination and possible
 312 rules. This will be a part of future work.

313 5. Figures

Figure 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 .

Figure 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.

Figure 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.

Figure 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 Figure 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. It should be noted that the difference between the stepped and spikey patterns is due to the particular arrangement of the rule sets. If the rules were sorted differently a different arrangements of patterns would be seen. 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.

Figure 5: The graphs show how the top angle changing rule set attempts to learn a stable solution. The top figures show how the rule adapts its angles over multiple 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.

Figure 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 Figure 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. It should be noted that the difference between the stepped and spikey patterns is due to the particular arrangement of the rule sets. If the rules were sorted differently a different arrangements of patterns would be seen. 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.

Figure 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 discussion.

314 6. Tables

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%

Table 1: Top 5 successful angle changing rule sets and their corresponding success rates. These top rules can also be seen in Figure 4 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%

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

315 7. Acknowledgements

316 This work was supported by the EPSRC Centre for Doctoral Training in
317 Future Robotics and Autonomous Systems (FARSCOPE) EP/L015293/1. The
318 research was also partly supported by the Leverhulme Trust Research Grant
319 RPG-2016-345.

320 8. Conflict of Interests

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

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%

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

our work.

References

- [1] Alexander, R. M. (2003). *Principles of animal locomotion*. Princeton University Press.
- [2] Blickhan, R. (1989). The spring-mass model for running and hopping. *Journal of Biomechanics*, 22(11-12):1217–1227.
- [3] Blum, Y., Lipfert, S. W., Rummel, J., and Seyfarth, A. (2010). Swing leg control in human running. *Bioinspiration & Biomimetics*, 5(2):026006.
- [4] Cham, J. G., Karpick, J. K., and Cutkosky, M. R. (2004). Stride period adaptation of a biomimetic running hexapod. *The International Journal of Robotics Research*, 23(2):141–153.
- [5] Dickinson, M. H., Farley, C. T., Full, R. J., Koehl, M., Kram, R., and Lehman, S. (2000). How animals move: an integrative view. *science*, 288(5463):100–106.
- [6] Ernst, M., Geyer, H., and Blickhan, R. (2010). Spring-legged locomotion on uneven ground: a control approach to keep the running speed constant. In *Mobile Robotics: Solutions and Challenges*, pages 639–644. World Scientific.

- 340 [7] Geyer, H., Blickhan, R., and Seyfarth, A. (2002). Natural dynamics of
341 spring-like running: Emergence of selfstability. In *5th International Confer-*
342 *ence on Climbing and Walking Robots*, pages 87–91. Suffolk, England: Pro-
343 fessional Engineering Publishing Ltd.
- 344 [8] Hauser, H., Fuchslin, R. M., and Pfeifer, R. (2014). Opinions and outlooks
345 on morphological computation.
- 346 [9] Hauser, H., Ijspeert, A. J., Fuchslin, R. M., Pfeifer, R., and Maass, W.
347 (2011). Towards a theoretical foundation for morphological computation with
348 compliant bodies. *Biological cybernetics*, 105(5-6):355–370.
- 349 [10] Hauser, H., Ijspeert, A. J., Fuchslin, R. M., Pfeifer, R., and Maass, W.
350 (2012). The role of feedback in morphological computation with compliant
351 bodies. *Biological cybernetics*, 106(10):595–613.
- 352 [11] Hurst, J. W. and Rizzi, A. A. (2005). Physically variable compliance in
353 running. In *Climbing and Walking Robots*, pages 123–133. Springer.
- 354 [12] Iida, F. and Tedrake, R. (2007). Motor control optimization of compliant
355 one-legged locomotion in rough terrain. In *Intelligent Robots and Systems,*
356 *2007. IROS 2007. IEEE/RSJ International Conference on*, pages 2230–2235.
357 IEEE.
- 358 [13] Karssen, J. D. and Wisse, M. (2011). Running with improved distur-
359 bance rejection by using non-linear leg springs. *The International Journal*
360 *of Robotics Research*, 30(13):1585–1595.
- 361 [14] McMahon, T. A. and Cheng, G. C. (1990). The mechanics of running: how
362 does stiffness couple with speed? *Journal of biomechanics*, 23:65–78.
- 363 [15] Owaki, D. and Ishiguro, A. (2006). Enhancing self-stability of a passive
364 dynamic runner by exploiting nonlinearity in the leg elasticity. In *SICE-*
365 *ICASE, 2006. International Joint Conference*, pages 4532–4537. IEEE.

- [16] Rummel, J., Iida, F., Smith, J. A., and Seyfarth, A. (2008). Enlarging regions of stable running with segmented legs. In *2008 IEEE International Conference on Robotics and Automation*, pages 367–372.
- [17] Rummel, J. and Seyfarth, A. (2008). Stable running with segmented legs. *The International Journal of Robotics Research*, 27(8):919–934.
- [18] Schmitt, J. (2006). A simple stabilizing control for sagittal plane locomotion. *Journal of Computational and Nonlinear Dynamics*, 1(4):348–357.
- [19] Seyfarth, A. and 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*, pages 81–85.
- [20] Seyfarth, A., Geyer, H., Günther, M., and Blickhan, R. (2002). A movement criterion for running. *Journal of Biomechanics*, 35(5):649–655.
- [21] Seyfarth, A., Geyer, H., and Herr, H. (2003). Swing-leg retraction: a simple control model for stable running. *Journal of Experimental Biology*, 206(15):2547–2555.
- [22] 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. *Robotics and autonomous systems*, 61(12):1601–1614.
- [23] Vu, H. Q., Hauser, H., Leach, D., and Pfeifer, R. (2013). A variable stiffness mechanism for improving energy efficiency of a planar single-legged hopping robot. In *Advanced Robotics (ICAR), 2013 16th International Conference on*, pages 1–7. IEEE.
- [24] Wieber, P.-B., Tedrake, R., and Kuindersma, S. (2016). Modeling and control of legged robots. In *Springer handbook of robotics*, pages 1203–1234. Springer.

393 [25] Wolf, S., Grioli, G., Eiberger, O., Friedl, W., Grebenstein, M., Höppner,
394 H., Burdet, E., Caldwell, D. G., Carloni, R., Catalano, M. G., et al. (2016).
395 Variable stiffness actuators: Review on design and components. *IEEE/ASME*
396 *transactions on mechatronics*, 21(5):2418–2430.