Improve Quadrupedal Locomotion with Actuated or Passive Joints?

Jared M. Moore[1] and Anthony J. Clark[2]

[1]School of Computing and Information Systems,
Grand Valley State University, Allendale, MI, USA
[2]Computer Science Department,
Missouri State University, Springfield, MO, USA
moorejar@gvsu.edu

Abstract

Animals interact with their environment softly through interaction of muscles, tendons, and rigid skeleton. By incorporating flexibility, they reduce ground impact forces and improve locomotive efficiency. Flexibility is also beneficial for robotic systems, although it remains challenging to implement. In this paper, we explore the addition of passive flexibility to a quadrupedal animat; we measure the impact of flexibility on both locomotive performance and energy efficiency of movement. Results show that spine and lower limb flexibility can significantly increase distance traveled when compared to an animat with no flexibility. However, replacing passively flexibile joints with actively controlled joints evolves more effective individuals with similar efficiency. Given these results, the number of joints and joint configuration appear to drive performance increases rather than just the addition of passive flexibility.

Introduction

Animals exhibit a remarkable ability to adapt locomotion to varying conditions. Gaits are driven by responses from the central nervous system and the morphology of the organism itself. Often, characteristics of the musculoskeletal system, such as elasticity of the tendons, contribute to their movements. For example, Alexander and Vernon (1975) found that large hind limb tendons in kangaroos and wallabies allow them to efficiently conserve energy during locomotion. Muscle-tendon systems in bipeds and quadrupeds act as energy storage contributing to running gaits in vertebrates (Alexander, 1984) while the spine has been shown to conserve energy during galloping (Alexander et al., 1985).

Robotic systems typically comprise rigid-body components, connected with single degree-of-freedom (DOF) actuators such as servo motors and linear actuators. These systems are often bioinspired, drawing upon the morphology and behaviors of biological organisms. Mechanical components, however, lack the flexibility of their natural counterparts, so compliant components are often added to these systems. Ackerman and Seipel (2013) added elasticity through springs, reducing the energetic cost of legged locomotion in a hexapedal robot. Increasing flexibility in the hexapod damped vertical movement of the torso as compared to a fully rigid-body robot. The addition of a flexible spine increased locomotive performance in a quadruped animat (Moore et al., 2015). It remains an open question whether the performance gains were due to flexibility or the increase in the DOF in the animat. Would performance increase if we replaced passive components with actively actuated joints?

In our preliminary study (Moore and Clark, 2018), we found that additional degrees of freedom improve the walking speed of our animats. In this paper, we further investigate the differences between passive and active joints, and we explore the impacts of these configurations on efficiency as well as speed. We conduct seven treatments with different animat configurations. We first examine performance, based on distance traveled, of a quadruped with legs actuated by hinge joints and no passive flexibility. Next, we increase the flexibility of the animat by adding sliding joints to the lower limbs (acting as shock absorbers) and a flexible spine. Finally, we replace the passive sliding joints with actively controlled hinge joints in the lower limbs, maintaining the DOF but reducing flexibility. We investigate both a passively flexible and active spine for this new animat.

We find that the addition of passive flexibility, whether it is in the spine or legs, significantly increases the distance traveled over a fully rigid-body quadruped. The highest performing platform with passive flexibility includes a flexible spine and lower limb sliders. Still, replacing flexibility in the lower sliders with actively controlled hinge joints produces the furthest distances traveled. The most effective individuals across all treatments include both an actively controlled spine and actively controlled lower hinge joints. Efficiency does not significantly change between passively flexible and actively controlled joints. This suggests that while flexibility can increase the performance of a robotic system, the real factor for performance increases is likely a combination of increasing DOF in the animat and joint configurations.

Background and Related Work

In evolutionary robotics (ER) (Nolfi and Floreano, 2000; Doncieux et al., 2015) both control and morphology of robotic systems are optimized using concepts derived from biological evolution. Evolutionary approaches are well

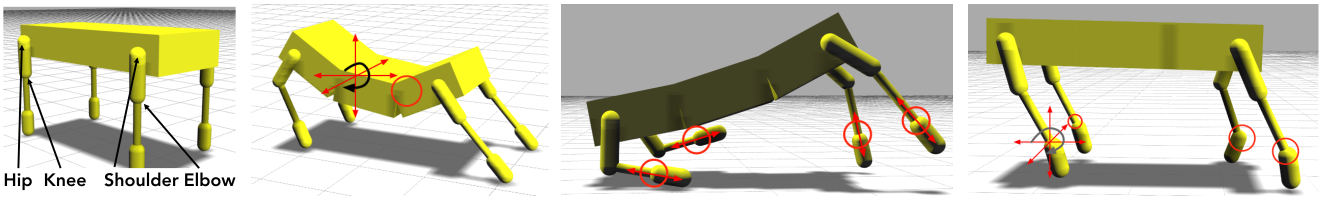

(a) Base Animat (b) Actuated Spine (c) Flexible Sliders (d) Lower Hinges

Figure 1: Different joint configurations explored in this study. (a) Base quadruped configuration with no sliders on the lower limbs. (b) Spines can be active or passive with three torso segments connected by two hinge joints. Axis of rotation indicated on rear spine, circle indicates second spine joint. (c) Flexible sliders on the lower limbs allow for a dampening of interaction with the ground. Note the compressed slider in the right rear leg compared to the extended slider in the front right leg. Arrows indicate axis of slider translation. (d) Actively controlled hinge joints on the lowest joint in each limb. Rear right leg indicates axis of rotation for all lower hinge joints.

suited to problems where an algorithm for deriving the optimal solution is not known *a priori* (Li and Miikkulainen, 2014). They have proven effective for optimizing controllers in wheeled (Fischer and Hickinbotham, 2011) and legged (Cully and Mouret, 2013; Stanton and Channon, 2013) systems. Evolutionary algorithms are especially useful when exploring morphology (Funes and Pollack, 1998). The body plays an important role in movement, performing an innate control prior to engagement of higher-level control from the brain (Valero-Cuevas et al., 2007). Auerbach and Bongard (2010) demonstrated that optimizing brain and body together produces effective systems exploiting integration between aspects of morphology and control.

Passive flexibility plays an important role in biological organisms, helping to reduce the energetic cost of locomotion by storing energy in spring-like tendons (Baudinette et al., 1992; Ruina et al., 2005). In robotics, Rieffel et al. (2010) demonstrated that even in the absence of a higher level controller the spring systems comprising tensegrity robots can be harnessed to realize locomotion. In traditional rigid-body robots, compliant joints enable robotic systems to mimic the passive flexibility of animals (Vanderborght et al., 2013). Passive compliance in robotic systems improves climbing ability (Seo and Sitti, 2013) and swimming (Clark et al., 2014). While augmenting robotic systems with passive flexibility can improve performance, it remains difficult to determine whether it is due to the elasticity itself, or if it is perhaps the additional DOF added to the system. Evolutionary methods coupled with simulation enable exploring many configurations that would not be practical with physical systems.

Methods

Simulation Environment The Open Dynamics Engine (ODE) (Smith, 2013) is used to conduct simulations. ODE is a 3D rigid-body physics simulation engine that models forces such as gravity, friction, and collisions between objects. Actively controlled actuators include single DOF hinges and linear motors, among others. In ODE we model flexibility by connecting rigid-bodies with spring-like joints that can be active or passive. The environment is a flat, high-friction surface. Animats are evaluated for 10 seconds of simulation time with a timestep of 0.005 seconds.

Quadruped Animat The base quadruped animat is shown in Figure 1a. The torso is composed of three segments connected by fixed joints. Each leg is three segments with hinge joints at the hip and knee. In the base treatment, the joint connecting the lowest component to the mid-leg is fixed, effectively creating a short upper segment and a longer lower leg segment.

Other animats are derived from the base treatment by adding a passive or active spine combined with a passive leg slider or active leg hinge. Figure 1b shows a quadruped animat with spine joints that are passive or actuated depending on treatment. Here, we replace the rigid joints in the torso with hinge joints that actuate along the lateral planes of the animat. Figure 1c shows the addition of flexible slider joints between the two lower limb segments. They compress during locomotion acting as shock absorbers. In the figure, the slider on the right rear leg is at maximum compression. Figure 1d shows a quadruped animat with actively controlled hinge joints on the lowest joint of each leg.

Controller The controller in this experiment is a conventional sinusoid, and each joint has its own evolved sinusoid parameters. Joint control signals are determined by the time of the simulation, evolved control signal modifiers per joint, and the maximum force output potential for each type of joint. Equation 1 generates the movement command for a single joint at each timestep in the simulation.

$$sin(-2\pi ft + (2\pi(\phi_{leg} + \phi_{joint}))) \quad (1)$$

f is the oscillation frequency common across the joints of the animat determining how quickly the sine wave oscillates represented by a real-value ranging from 0 to 2, t is the current simulation time, ϕ_{leg} and ϕ_{joint} are the phase offsets. Phase offsets are one of 16 set values ranging from 0 to 1.875, which corresponds to shifting the phase of the oscillating signal in $\frac{1}{8}$ increments. Each leg has its own phase offset relative to the common signal (ϕ_{leg}). Each joint type (e.g. shoulders, elbows, hips and knees) has an offset as well (ϕ_{joint}). Together, the two offsets produce common control signals for the the rear legs and the front legs. Each leg pair (front and rear) can have a common behavior (specified by ϕ_{joint}) that is then shifted temporally by ϕ_{leg}. Thus a common behavior for the front or rear legs can be out of phase, similar to walking in animals as symmetry and coordination can evolve between limbs.

We evolve forces acting on the shoulder, elbow, hip, and knee, allowing for joints to be entirely passive or force limited. Under this configuration, the oscillating signal sent to a joint and its maximum force output potential determine the response of the joint to a command. For example, if a joint evolves a low force output, it will passively flex under the force of gravity and not actively assist in locomotion. However, should the joint evolve a high force output, it will not deviate from its specified motion even when large external forces are applied, such as when making contact with the ground where such forces would typically hinder normal actuation of a joint.

Passively flexible joints are governed by spring and damper constraints parameterized in ODE as ERP and CFM. ERP values evolve in the range of 0.4 to 1.0. CFM values evolve from 0.0001 to 0.15. Together, the two parameters specify the stiffness and damping of a joint. In general, high CFM values and low ERP values result in flexible joints whereas the opposite lead to stiff joints.

Treatments Seven treatments are conducted in this study.

1. *No Sliders* (NS) - 8 DOF
   Base quadruped animat with no passive flexibility.
2. *Flexible Spine, No Sliders* (FSpNS) - 10 DOF
   Flexible spine, no sliders on lowest joint.
3. *Rigid Spine, Flexible Sliders* (FS) - 12 DOF
   Flexible sliders on lowest joint with a rigid spine.
4. *Flexible Spine, Flexible Sliders* (FSpFS) - 14 DOF
   Flexible spine and sliders.
5. *Rigid Spine, Active Lower Hinge* (HL) - 12 DOF
   Actively controlled hinges on lowest joint with a rigid spine.
6. *Flexible Spine, Active Lower Hinge* (HLFSp) - 14 DOF
   Actively controlled hinges on lowest joint and flexible spine.
7. *Active Spine, Active Lower Hinge* (HLASp) - 14 DOF
   Actively controlled hinges on lowest joint and spine.

The first four treatments (*NS*, *FSpNS*, *FS*, *FSpFS*) evaluate varying degrees of passive flexibility in the animat. *NS* is the base quadruped with 8 DOF and no flexibility. *FSpNS* adds spine flexibility, increasing the DOF to 10. *FS* adds flexible lower sliders to each leg while maintaining a rigid spine. There are 12 DOF in this animat. *FSpFS* combines both flexible spine and flexible sliders with 14 DOF; 6 DOF more than the base animat configuration.

The final three treatments (*HL*, *HLFSp*, *HLASp*) replace the lower passive slider joints with actively controlled hinge joints. *HL* has a rigid spine and 12 DOF similar to the *FS* treatment. *HLFSp* has a passively flexible spine 14 DOF and the spine configuration of the *FSpFS* treatment. *HLASp* has an actively actuated spine and 14 DOF.

Evolutionary Algorithm For each treatment, we evolve 120 individuals over 4,000 generations using the DEAP framework (De Rainville et al., 2012) with a conventional genetic algorithm. DEAP is an open-source framework implementing many common evolutionary algorithms. 20 replicate runs, each seeded with a unique random number, are conducted per treatment. Fitness is the horizontal Euclidean distance from the starting point to the center of the torso after 10 seconds of simulation time. Selection of a parent is performed through a tournament of 4 randomly chosen individuals. Crossover is performed (two parents selected) with a rate of 50% and mutation of 4% per gene. An individual replicate took approximately 5 hours on a Blade system parallelized across 24 cores and a clock speed of 2GHz.

Genome Table 1 lists the genes in this study. The composition of each genome varies depending on the treatment and is indicated in the right columns of the table. All treatments have 10 genes for generating the control signals and 5 genes for joint forces (back, shoulder, elbow, hip, knee) described previously. Depending on the combination of genes included, treatments range in value from 15 (*NS*) to 23 (*FSpFS*, *HLFSp* and *HLASp*) genes.

Table 1: Genes defining the quadruped animat. "l/r sym" denotes that left/right symmetry is enforced.

Description		Treatments						
	# Genes	NS	FSpNS	FS	FSpFS	HL	HLFSp	HLASp
Oscillation Frequency	1	●	●	●	●	●	●	●
Max Joint Velocity	1	●	●	●	●	●	●	●
Phase Offset Per Leg	4	●	●	●	●	●	●	●
Joint Phase Offset (l/r sym)	4	●	●	●	●	●	●	●
Joint Max Force (l/r sym)	5	●	●	●	●	●	●	●
Slider Flexibility (l/r sym)	4			●	●			
Spine Joint Flexibility	4		●		●		●	
Hinge Phase Offset (l/r sym)	2					●	●	●
Hinge Max Force (l/r sym)	2					●	●	●
Active Spine Phase Offset	2							●

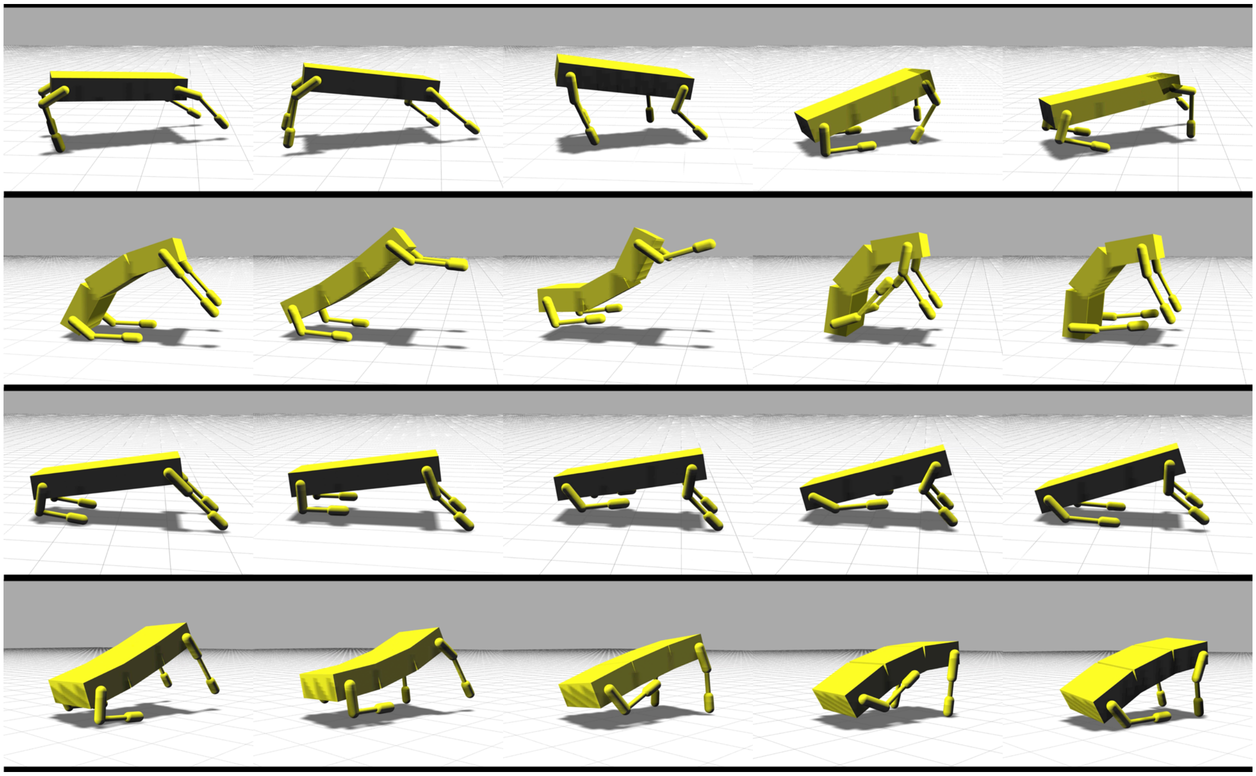

Figure 2: Sample evolved gaits. (Top) Galloping gait evolved in the *NS* treatment. (Top-Mid) Hopping gait from the *FSpNS* treatment. (Bottom-Mid) Bounding gait from the *FS* treatment. (Bottom) Bounding gait from the *FSpFS* treatment.

Results

Bounding, galloping, canters, and trots evolve across treatments, and samples of evolved gaits from all treatments can be seen at `https://youtu.be/UCNxJ3pmmkc`. Our analysis focuses on two main questions. First, does a quadruped animat with passive flexibility significantly outperform a fully rigid-body animat? If so, what combination of lower leg slider, spine, or both, leads to the most effective individual? Second, how does replacing passive joints with actively actuated hinges alter performance?

Passive Flexibility Figure 2 highlights a few of the gaits that evolve across the initial treatments. Figure 3 plots the maximum distance traveled at each generation averaged across twenty replicates per treatment over evolutionary time. Shaded regions represent the 95% confidence intervals for each treatment. For reference, the animat's body length is 3 units. A fully rigid body animat (*NS*) yields the lowest distances traveled. The other treatments show a significant improvement over the *NS* treatment with a flexible spine (*FSpNS*) slightly better. Flexible sliders (*FS*) on the lower limbs increase distance traveled, with the combination of spine and slider flexibility (*FSpFS*) yielding the farthest traveling individuals.

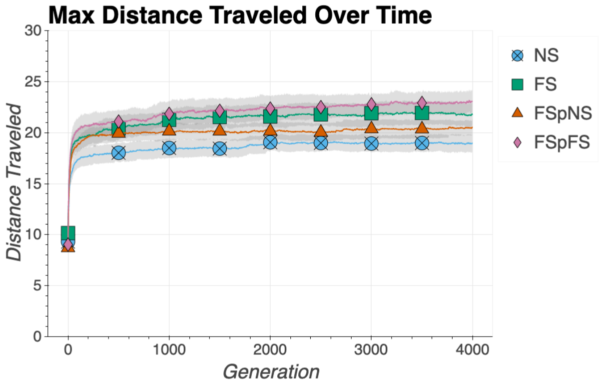

Figure 3: Maximum distance traveled across twenty replicates per treatment over evolutionary time for the passive flexibility treatments. Shaded areas represent the 95% confidence intervals.

Figure 4 plots the distribution of the farthest traveling individual per replicate across the treatments. We conduct the following statistical tests using a Wilcoxon Rank-Sum Test performed on the distribution of farthest traveling individual per replicate for each treatment. Pairwise results across all treatments can be seen in Table 2. For the two treatments without sliders, *NS* and *FSpNS*, there is no significant difference in performance ($p = 0.07$). Whereas both treatments with sliders have significantly higher performance. *FS* versus *NS* ($p < 0.01$), *FS* versus *FSpNS* ($p < 0.03$), and ($p < 0.01$) for *FSpFS* versus *NS* and *FSpNS*. There is no significant difference between *FSpFS* and *FS* ($p = 0.08$).

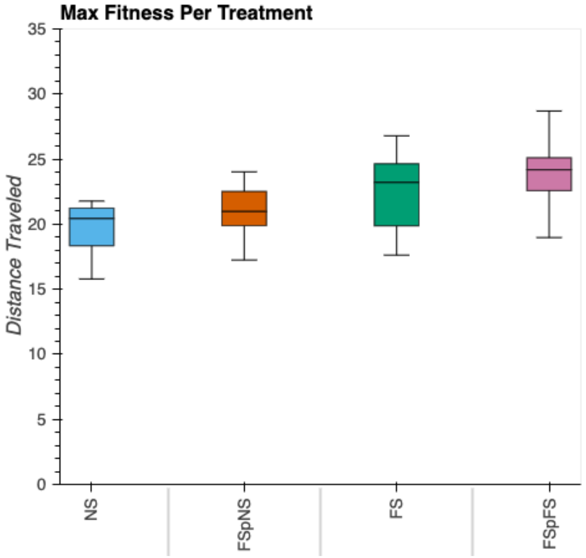

Figure 4: Distribution of the farthest traveling individual per replicate across the passive flexibility treatments.

It appears that the addition of spine flexibility alone is not enough to significantly improve locomotive performance even though two additional DOF have been added to the animat. Although the distance traveled of the best individuals for the *FS* treatment is significantly higher, performance also varies considerably. Animats with both flexible spine and flexible sliders exhibit the highest average performance across these treatments. However, the *FSpFS* treatment has 6 more DOF than the *NS* treatment. Differences in performance might not solely be due to flexibility and instead could be attributed to the possible additional behaviors the increased mobility allows.

Hinge Joints By replacing the flexible sliders with actively controlled hinge joints, we address the question of whether the increase in DOF or flexibility drives performance increases. We chose a passive slider and active hinge as these are the most effective actuator for their respective control type. Figure 5 shows three variations of bounding gaits that evolve across the three active lower hinge treatments. Figure 6 shows the maximum distance traveled per generation over evolutionary time for all seven treatments conducted in this study. *HL* has similar evolutionary performance as compared to the best of the passive flexibility treatments, *FSpFS*. The addition of the lower hinge joints, with spine mobility, produces the farthest traveling individuals observed. A combination of actively controlled lower hinges and a passive spine flexibility (*HLFSp*) outperforms just replacing the sliders with hinges while an active spine leads to the highest average maximum distance traveled across all twenty replicates.

Figure 7 plots the distribution of the farthest traveling individual across replicates for all seven treatments conducted in this study. Changing the lower joint from passive sliders (*FS*) to actively controlled hinges (*HL*) does not significantly increase distance traveled ($p = 0.12$). Furthermore, there is no significant difference between the *FSpFS* and *HL* treatments. Alone, actively controlled hinges don't generate significant improvements in performance compared to the flexible slider treatments. Further modifying the active lower hinge animat by adding either a flexible spine (*HLFSp*) or an active spine (*HLASp*) does significantly improve performance over the flexible slider treatments.

Table 2: Wilcoxon Rank-Sum Test comparing fitnesses of the farthest traveling individuals per replicate across treatments.

	NS	*FSpNS*	*FS*	*FSpFS*	*HL*	*HLFSp*	*HLASp*
NS	-	$= 0.07$	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
FSpNS		-	< 0.03	< 0.01	< 0.01	< 0.01	< 0.01
FS			-	$= 0.08$	$= 0.12$	< 0.01	< 0.01
FSpFS				-	$= 0.88$	$= 0.01$	< 0.01
HL					-	< 0.01	< 0.02
HLFSp						-	$= 0.10$
HLASp							-

Performance increases appear to be due to a combination of the number of DOF and the use of actively controlled hinge joints versus flexible sliders. The top two treatments, *HLFSp* and *HLASp*, have 14 DOF and are significantly better than any other treatment. Replacing flexible sliders with a hinge joint does not significantly increase performance, but it maintains similar performance to the *FSpFS* treatment, which has 14 DOF compared to *HL*s 12 DOF. The two lowest performing treatments, *NS* and *FSpNS*, have 8 and 10 DOF, respectively. Increases in distance traveled for the animat configurations in this study appear to be influenced more by the DOF, and addition of actively controlled joints, than the addition of passive flexibility.

Efficiency Flexibility in natural organisms can lower the energetic cost of locomotion (Alexander, 1984). In this study, our sole objective is to maximize distance traveled, but efficiency might differ between animats because of including passive flexibility. We measure efficiency as the distance traveled divided by the total power expended through actively controlled joints in an animat. Total power is the summation of force exerted by each actively controlled joint at each time step as reported by the physics engine. Passive joints replicate spring systems and therefore are not included in the total power calculation. Figure 8 plots the distribution of efficiency for the farthest traveling individual in each replicate across treatments. Here, there appears to be little difference in efficiency. Table 3 lists the pairwise Wilcoxon-Rank sum tests comparing efficiency between treatments. While some pairs are significantly different, there is predominantly no significant difference in efficiency across treat-

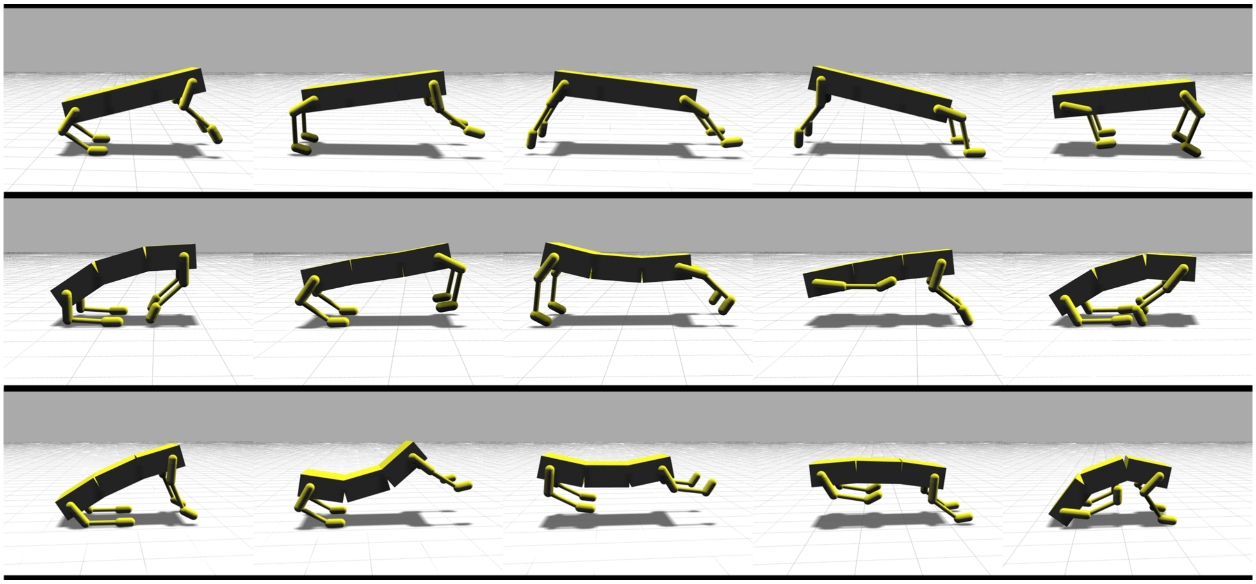

Figure 5: Sample evolved bounding gaits from the active hinge joint on lowest limb segment treatments. (Top) *HL* treatment. (Mid) *HLFSp* treatment. (Bottom) *HLASp* treatment.

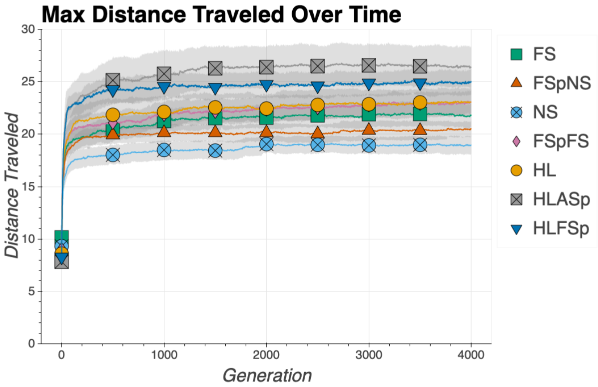

Figure 6: Maximum distance traveled across twenty replicates per treatment over evolutionary time for all treatments. Shaded areas represent the 95% confidence intervals.

ments. It appears that a flexible spine, compared to similar animats that have a rigid, or actively controlled spine, typically have lower efficiency, but this is only a significant difference in 1 out of 3 cases (*NS* vs. *FSpNS*, $p < 0.04$). When evaluating replicates on both efficiency and distance, we found that the single lowest performing individual is in the *HLASp* treatment. However, many of the highest performing individuals are also in the *HLASp* treatment. No clear advantage is apparent when passive flexibility is included in an animat in terms of efficiency.

Conclusions

In this paper, we investigated the effect of adding additional DOF to a rigid-body quadrupedal animat in terms of performance and efficiency. Adding additional DOF in the form of passive flexibility or actively controlled hinge joints significantly increases performance in terms of distance traveled. However, efficiency remains unaffected when it is not directly included as a selective pressure during evolution.

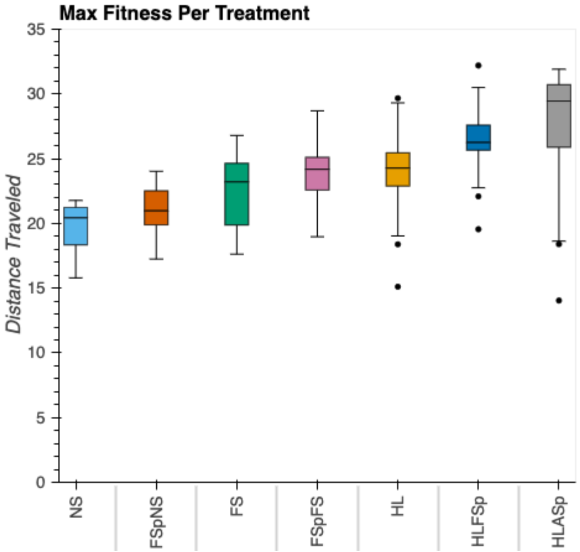

Figure 7: Distribution of the farthest traveling individual per replicate across the treatments. Adapted from Moore and Clark (2018).

Adding flexibility to the animat in the spine and lower sliders significantly increases distance traveled versus the base animat configuration. Animats with both spine and lower slider flexibility are the farthest traveling individuals among those with passive flexibility. This result supports those of earlier works (Seo and Sitti, 2013; Lessin et al., 2014; Clark et al., 2012) where flexibility aids in the per-

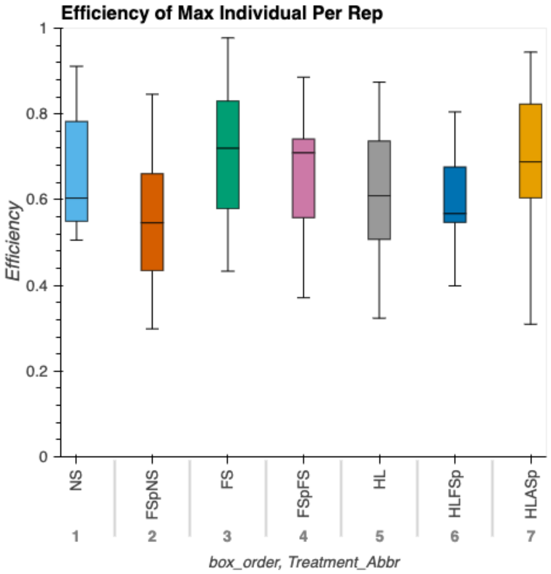

Figure 8: Efficiency of the best individual per replicate across treatments.

Table 3: Wilcoxon Rank-Sum Test comparing efficiency of the farthest traveling individual per replicate across treatments.

	NS	*FSpNS*	*FS*	*FSpFS*	*HL*	*HLFSp*	*HLASp*
NS	-	< 0.04	$= 0.26$	$= 0.55$	$= 0.31$	$= 0.33$	$= 0.43$
FSpNS		-	< 0.01	$= 0.11$	$= 0.55$	$= 0.18$	< 0.03
FS			-	$= 0.19$	< 0.02	< 0.05	$= 0.60$
FSpFS				-	$= 0.37$	$= 0.28$	$= 0.09$
HL					-	$= 0.88$	$= 0.25$
HLFSp						-	$= 0.09$
HLASp							-

formance of robotic systems. Thus, when building robotic systems, examining the incorporation of flexible components, such as springs, may be worthwhile to increase performance of a rigid-body robot. It also further clarifies and expands (Moore et al., 2015), in that flexibility is likely not the sole driver of performance increases, rather, both flexibility and an increase in the DOF positively impact performance.

Actively controlled joints lead to even higher performing individuals across all treatments. Replacing the lower sliding joints with actively controlled hinge joints results in the three highest performing treatments, out of the seven conducted. Adding spinal mobility in the form of a passively flexible spine, or actively controlled one, further increases distance traveled. This suggests that while passive flexibility improves performance over the base animat in this study, the increase is likely not due directly to including passive flexibility. Rather, the increase in DOF drives improvements in distance traveled.

In terms of robotics systems, incorporating flexibility could still be beneficial depending on the platform and problem constraints. Here, we find that including flexibility increases performance over a fully rigid-body robot with no flexibility. Although active control produces the highest performing individuals, it may be that a designer does not want further increases in control complexity nor the additional hardware (servos, wiring, batteries, etc) required to coordinate additional actively controlled actuators. Instead, a controller can use the dynamics of passive joints to improve performance as demonstrated here, and in other work mentioned previously. Furthermore, passive flexibility may reduce wear on other mechanical components by dampening locomotive forces.

Future extensions to this study will investigate how flexibility and active control affect performance in other animat platforms such as hexapods. We plan to introduce more complex high-level controllers such as artificial networks (ANN) to see how these features are integrated in control logic. Furthermore, we plan to expand the scope of evolvability in terms of morphological components to evolve, along with exploring multi-objective algorithms.

Source Code

The source code for running these experiments is provided at `https://github.com/jaredmoore/Evo_Flex_Quadruped_Code`.

Acknowledgements

The authors would like to thank our reviewers for their thoughtful commentary to improve this work as well as Philip K. McKinley, Craig P. McGowan, and Byron DeVries for their support in this work. Additional thanks to Grand Valley State University and Missouri State University for supporting this research through computing resources and financial support.

References

Ackerman, J. and Seipel, J. (2013). Energy efficiency of legged robot locomotion with elastically suspended loads. *IEEE Transactions on Robotics*, 29(2):321–330.

Alexander, R. and Vernon, A. (1975). The mechanics of hopping by kangaroos (macropodidae). *Journal of Zoology*, 177(2):265–303.

Alexander, R. M. (1984). Elastic energy stores in running vertebrates. *American Zoologist*, 24(1):85–94.

Alexander, R. M., Dimery, N. J., and Ker, R. F. (1985). Elastic structures in the back and their rôle in galloping in some mammals. *Journal of Zoology*, 207(4):467–482.

Auerbach, J. E. and Bongard, J. C. (2010). Dynamic resolution in the co-evolution of morphology and control.

In *Proceedings of the Twelfth International Conference on Artificial Life*, pages 451–458, Odense, Denmark.

Baudinette, R. V., Snyder, G. K., and Frappell, P. B. (1992). Energetic cost of locomotion in the tammar wallaby. *American Journal of Physiology - Regulatory, Integrative and Comparative Physiology*, 262(5):R771–R778.

Clark, A., Wang, J., Tan, X., and McKinley, P. (2014). Balancing performance and efficiency in a robotic fish with evolutionary multiobjective optimization. In *2014 IEEE International Conference on Evolvable Systems*, pages 227–234, Orlando, Florida, USA.

Clark, A. J., Moore, J. M., Wang, J., Tan, X., and McKinley, P. K. (2012). Evolutionary design and experimental validation of a flexible caudal fin for robotic fish. In *Proceedings of the 13th International Conference on the Simulation and Synthesis of Living Systems*, pages 325–332, East Lansing, Michigan, USA.

Cully, A. and Mouret, J.-B. (2013). Behavioral repertoire learning in robotics. In *Proceedings of the 2013 ACM Genetic and Evolutionary Computing Conference*, pages 175–182, Amsterdam, Netherlands. ACM.

De Rainville, F.-M., Fortin, F.-A., Gardner, M.-A., Parizeau, M., and Gagné, C. (2012). Deap: A python framework for evolutionary algorithms. In *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation*, pages 85–92, Philadelphia, Pennsylvania, USA. ACM.

Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. G. (2015). Evolutionary robotics: What, why, and where to. *Frontiers in Robotics and AI*, 2(4).

Fischer, V. and Hickinbotham, S. (2011). Evolving a metabolic subsumption architecture for cooperative control of the e-puck. *Memetic Computing*, 3(4):231–244.

Funes, P. and Pollack, J. (1998). Evolutionary body building: Adaptive physical designs for robots. *Artificial Life*, 4(4):337–357.

Lessin, D., Fussell, D., and Miikkulainen, R. (2014). Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, pages 705–712, Vancouver, BC, Canada. ACM.

Li, X. and Miikkulainen, R. (2014). Evolving multimodal behavior through subtask and switch neural networks. In *Proceedings of The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, New York, NY.

Moore, J. M. and Clark, A. J. (2018). Bend and flex: Passive flexibility or active control in a quadruped animat. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 113–114, Kyoto, Japan. ACM.

Moore, J. M., McGowan, C. P., and McKinley, P. K. (2015). Evaluating the effect of a flexible spine on the evolution of quadrupedal gaits. In *Proceedings of the 13th European Conference on Artificial Life*, pages 166—173, York, UK.

Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*. The MIT Press.

Rieffel, J. A., Valero-Cuevas, F. J., and Lipson, H. (2010). Morphological communication: Exploiting coupled dynamics in a complex mechanical structure to achieve locomotion. *Journal of The Royal Society Interface*, 7(45):613–621.

Ruina, A., Bertram, J. E. A., and Srinivasan, M. (2005). A collisional model of the energetic cost of support work qualitatively explains leg sequencing in walking and galloping, pseudo-elastic leg behavior in running and the walk-to-run transition. *Journal of Theoretical Biology*, 237(2):170–192.

Seo, T. and Sitti, M. (2013). Tank-like module-based climbing robot using passive compliant joints. *IEEE/ASME Transactions on Mechatronics*, 18(1):397–408.

Smith, R. (2013). Open Dynamics Engine, http://www.ode.org/.

Stanton, A. and Channon, A. (2013). Heterogeneous complexification strategies robustly outperform homogeneous strategies for incremental evolution. In *Proceedings of the 12th European Conference on Artificial Life*, pages 973–980, Taormina, Italy.

Valero-Cuevas, F., Yi, J.-W., Brown, D., McNamara, R., Paul, C., and Lipson, H. (2007). The tendon network of the fingers performs anatomical computation at a macroscopic scale. *IEEE Transactions on Biomedical Engineering*, 54(6):1161–1166.

Vanderborght, B., Albu-Schäffer, A., Bicchi, A., Burdet, E., Caldwell, D., Carloni, R., Catalano, M., Eiberger, O., Friedl, W., Ganesh, G., Garabini, M., Grebenstein, M., Grioli, G., Haddadin, S., Hoppner, H., Jafari, A., Laffranchi, M., Lefeber, D., Petit, F., Stramigioli, S., Tsagarakis, N., van Damme, M., ten Ham, R., Visser, L., and Wolf, S. (2013). Variable impedance actuators: a review. *Robotics and autonomous systems*, 61(12):1601–1614.