

Exploiting Dynamical Complexity in a Physical Tensegrity Robot to Achieve Locomotion

Mark Khazanov¹, Ben Humphreys¹, Willam Keat² and John Rieffel¹

¹Computer Science Department, Union College, Schenectady NY 12308

²Mechanical Engineering Department, Union College, Schenectady NY 12308
rieffelj@union.edu

Abstract

The emerging field of *morphological computation* seeks to understand how mechanical complexity in living systems can be advantageous, for instance by reducing the cost of control. In this paper we explore the phenomenon of morphological computation in tensegrities – unique structures with a high strength to weight ratio, resilience, and an ability to change shape. These features have great value as a robotics platform, but also make tensegrities difficult to control via conventional techniques. We describe a novel approach to the control of tensegrity robots which, rather than suppressing complex dynamics, *exploits* them in order to achieve locomotion. Our robots are physically embodied (rather than simulated), evolvable, and locomote at higher speeds (relative to body size) and with fewer actuators than those controlled by more conventional approaches.

Introduction

Traditional engineering approaches to design of structures and the control of robots attempt to avoid, or at least actively suppress, complex system dynamics such as vibration and dynamic coupling among components. By reverting them to generally rigid and holonomic systems, their kinematics can be modeled using classical mechanics (Murray et al., 1994; Sciavicco and Siciliano, 2000), and these models can in turn be used to generate desired movements.

By contrast, the bodies of many natural organisms (and robots which attempt to mimic them) are by their very nature high dimensional dynamic systems with an essentially infinite number of degrees of freedom. Properties of living systems such as elasticity and deformability come at the cost of resonances and tight dynamic coupling between components (Trimmer, 2007) – properties which are often assiduously avoided in conventional engineering approaches to robotic design. This precludes the use of most of the traditional kinematic and inverse-dynamics approaches described above (Craig, 1989). While some methods exist for the control of non-holonomic (under-controlled) mechanical systems, many are incredibly computationally expensive, and difficult to transfer from simulation to reality (Hannan and Walker, 2003; Fung, 1993; Vogel, 2003).

How then, are dynamically complex biological systems so controllably robust and agile? The emerging field of *morphological computation* (Paul, 2006; Pfeifer et al., 2007; Pfeifer and Bongard, 2006) conjectures that “outsourcing” the computation into the mechanics of the structure allows related neural pathways to devote their resources to higher level tasks (Valero-Cuevas et al., 2007) – a type of “intelligence by mechanics” (Blickhan et al., 2007). These phenomena have been shown in the physiology of animals such as wallabies (Biewener et al., 2004) and guinea fowl (Daley and Biewener, 2006) and cockroaches (Ahn and Full, 2002).

Biological morphological computation has served as inspiration for robotic control in several recent works. Iida and Pfeifer (Iida and Pfeifer, 2006) explored how the body dynamics of a quadruped robot can be exploited for sensing. Watanabe (Watanabe et al., 2008) demonstrated how inducing long distance mechanical coupling in a snake robot improves its ability to learning a crawling motion. All of these systems were largely composed of rigid elements. Our interest is expanding these principles into the realm of soft materials and structures, where the complexity, and therefore the potential for beneficial exploitation, is significantly higher.

This paper demonstrates morphological computation in tensegrity robots. Tensegrities, pre-stress stable structures composed of rigid struts and tensile springs, possess many appealing traits, but exhibit high degrees of mechanical coupling, and are therefore difficult to control through conventional means. We implement an alternative and novel approach to tensegrity locomotion, one which seeks to exploit, rather than suppress, vibration and dynamical coupling between components. The resulting robot is quite simple, with open-loop actuation by low-voltage vibrating pager motors, and yet is capable of robust and controllable motion. We believe this to be both the fastest and the smallest physically embodied tensegrity robot yet to be developed. We begin the paper by describing the design of the robot. We then demonstrate vibrationally-actuated gaits which produce linear and rotational behavior. These gaits can be sequenced together in order to generate controllable trajectories. We conclude by discussing more sophisticated control algorithms and op-

timization techniques.

Tensegrity Robots

A tensegrity structure (Figure 1) is a self-supporting structure consisting of a set of disjoint rigid elements (struts) whose endpoints are connected by a set of continuous tensile elements (springs). Despite the fact that none of the rigid elements touch, tensegrities are able to maintain their structure due to a synergistic interplay of compressive and tensile forces (Wang, 1998). Because of this *pre-stress stability*, they are able to quickly return to form when perturbed by an outside force. (Connelly and Back, 1998).

Examples of tensegrity can be found in structures ranging from camping tents to sports stadiums. These same principles are also found in the biological realm, at all scales from the structure of proteins (Ingber, 1998) and cellular cytoskeleton (Wang et al., 2001) up to the tendinous network of the human hand (Valero-Cuevas et al., 2007).

What makes tensegrities particularly appealing as a robotic platform is their high strength-to-weight ratio and resilience, as well as their ability to change shape by altering the resting length of the tensile elements. As a result, tensegrity structures are increasingly being used for applications such as smart structures and soft robots (Tibert and Pellegrino, 2003; Tibert, 2002; Motro, 2003; Sultan, 1999; Matsuda and Murata, 2006).

Unfortunately, the pre-stress stability of tensegrities imposes complex nonlinear dynamics, even for relatively small structures (Skelton et al., 2001). Conventional approaches to tensegrity robots therefore attempt to dampen the vibrational modes of the robots before controlling them. Skelton *et al.* have been able to demonstrate both active vibration damping (2004) and open-loop control of simple structures (2004).

In most cases, once the vibration and dynamical coupling of a tensegrity robot has been reduced either actively or passively, deformation and control are achieved by changing the rest lengths of the tensile elements, for instance by attaching strings to a reeled servo motor (Paul et al., 2006, 2005). Even so, the majority of tensegrity robotics has occurred in simulation (Aldrich et al., 2003; Paul et al., 2006; Graells Rovira and Mirats Tur, 2009; Iscen et al., 2013) rather than reality (Shibata et al., 2009). One notable recent contribution is Caluwaerts *et al.*'s work (2013) on physical reservoir computing in a simulated tensegrity robot, in which they demonstrate learnable gaits produced by relatively simple central pattern generators (CPGs).

There are a few published examples of physically embodied tensegrity robots moving. Paul *et al.* (2006) built a three-bar tensegrity robot with 0.4m struts, with three of its tensile elements actuated by servomotors. Using a gait derived from simulation, the physical robot was able to achieve speeds of around 60cm/min. Shibata *et al.* (2009) built a 6-bar tensegrity with 15cm struts, with its tensile elements actuated by shape-memory alloy wires. (The speed of the




Figure 1: Tensegrities consist of a set of rigid elements (rods) joined together by tensile elements (springs). They maintain their shape through a synergistic interplay of forces. (Photo by Steven Stangle)

resulting robot was not published). Later, Koizumi *et al.* (2012) built a six-bar robot with 0.6m struts with tensile elements actuated by 24 pneumatic McKibben actuators, which moved by rolling. While the speed of the robot is not published, an online video shows 15 rolls in the space of 45 seconds (1 roll per 3 seconds). We estimate from the video that the robot is moving at approximately 25cm per roll, or 5m/min.

A new way to move tensegrities

The approaches describe above have essentially treated tensegrity-based robots as quasi-static and non-oscillating structures. And yet, tensegrities are by their very nature highly dynamic – anecdotally, the tensegrities we have built in our lab will readily oscillate as the table is bumped, or even as someone types on a keyboard.

The motivation for our work, therefore, lies in striving to exploit, rather than suppress, this inherent dynamical complexity as an advantage – making tensegrities move by vibrating, rather than suppressing their vibrations. Since the dynamics of real-world tensegrity structures are incredibly difficult to model in simulation, we chose to avoid simulation entirely and perform all experiments in the physical world. Quoting Rodney Brooks, “the world is its own best model” (Brooks, 1990).

Design

Our ambition was to design a small (< 15 cm) tensegrity that was powered by vibration alone. It would also have to be robust enough to endure the long hours of testing and, given the practical challenges of constructing a mechanical device on so small a scale, it would be advantageous if it were easy to manufacture and repair.

The resulting design, based upon a canonical six-bar tensegrity shape, is shown in Figure 1. The geometry is




Figure 4: Max distance (over 10 evaluations, each lasting 7 seconds) traveled by the robot during a sparse sweep of voltages for motors 1 and 2 between 2.1V and 2.9V, in increments of 0.3V (motor 3 was fixed at 0v). As can be seen, the resulting space is complex, and distance does not correlate directly with motor speed.




Figure 5: Variation across trials for a sweep of motor 1 values while keeping voltages for motors 2 and 3 fixed (the bottom row of Figure 4))



Figure 6: An example of how distance traveled was automatically measured. Position was calculated by subtracting an “empty” arena image from one containing the robot. Distance could then be measured by comparing beginning (left) and final (center) robot positions (right).



Figure 2: The robot testing arena is on the floor in order to maximize the navigable area. (Photo by Steven Stangle)



Figure 3: The body of the tensegrity robot consists of rigid plastic rods and metal springs. It is actuated by simple DC vibrating motors attached at the midpoint of three of its six struts. A closeup of the robot illustrated the springs and the DC motors. (Photo by Steven Stangle)

defined by six equal length composite struts which are connected to each other via 24 identical helical springs, with four springs emanating from each strut end.

This 6-bar tensegrity has three orthogonal planes of symmetry. As a result of this natural symmetry, each spring wants to stretch an equal amount in the fully assembled equilibrium configuration, greatly simplifying design and analysis. There is however a slight loss of symmetry when the tensegrity is placed in contact with the ground because none of the stable contact positions are aligned with the planes of symmetry. Sagging of the springs under the weight of the spars and motors further disrupts the symmetry. Three small vibration motors were mounted on composite struts and oriented such that each shaft axis is perpendicular to one of the planes of symmetry.

Few actual machining operations are required to produce the tensegrity. The 9.4 cm long composite struts are cut from 6.35 mm square graphite composite tubes. The motors were mounted to the flat outer surface of the struts using epoxy. The struts, while square on the outside, are hollow circular on the inside. Therefore, both ends of each strut could be tapped to allow for insertion of 10-24 nylon screws. These screws provide a smooth contact surface for the tensegrity to rest on and are used to fasten nylon washers to the ends of the struts. The hooked ends of the helical springs are attached directly to the nylon washers via 4 equally-spaced drilled holes.

Selection of the springs is critical to overall performance. The basic strategy was to try to produce the smallest possible natural frequencies under the presumption that small natural frequencies would lead to large displacement amplitudes, thus enhancing the chances that the tensegrity might roll. We strove to achieve this goal by minimizing spring stiffness and spring preload subject to the constraint of keeping static deflections within an acceptable limit (so that the tensegrity would not lose its basic shape). A single vertical strut was modeled as supported by eight linear springs oriented at 45° in order to limit the maximum static deflection

to 5% of total strut length. This calculation led to selection of a helical spring with a spring constant of 0.209 N/cm.

Attempts were also made to optimize the selection of a vibration motor. While the need for a small motor with a relatively large offset mass was clear, choosing the best range of operating speeds involved compromise. On the one hand, the motor speeds have to be large enough to produce sufficient centrifugal force, but also small enough to excite the lower (high energy) natural frequencies of the tensegrity. These considerations led to selection of the Precision Microdrives vibration motor Model 312-107 (Figure 3) which operates between 100 and 260 Hz.

To check the suitability of this operating range, the tensegrity was modeled using a matrix structural analysis code. One end of one strut was assumed fixed in space and the associated natural frequencies were determined. The fundamental (lowest) frequency was found to be 7.8 Hz, well below the operating range of the motor. However, physical testing of the tensegrity with these motors indicated the existence of natural frequencies that were within the operating range.

Evaluation

Having designed the robot in order to maximize resonance, we then evaluated its ability to locomote via vibration. Early trials indicated a wide diversity of motions was possible, and that small changes in motor frequencies could lead to large changes to the resulting gait. Below we seek to quantify this.

(Videos of the robot moving can be seen at the corresponding author's web page: www.cs.union.edu/~rieffelj/videos/)

Setup

The robot was placed into an 85cm x 80cm arena (Figure 2) on the floor of our lab (the robot had a habit of falling off of tables) and tethered to the power supply and motor controllers. In order to reduce the effects of the tether, which might undesirably constrain the motion, the tether was built from narrow gauge magnet wire. Two cameras, a USB camera and a small handheld video camera, were placed 130cm above the arena. The USB camera was used for distance measures and the video camera was used to document results.

The voltage (speed) of each of the three motors was controlled by USB motor controllers connected to a host computer.

As illustrated by Figure 6, the process of distance measure was automated using the overhead USB camera connected to the control computer. The location of the tensegrity in a frame was determined by subtracting the image of an "empty" arena from an image containing the robot and then finding the centroid of the remaining pixels. The tether is visible in some frames, but has a negligible impact upon positional measurements. Distance could then be calculated

by comparing the pre- and post-evaluation locations. The arena was large enough that multiple evaluations could often be performed before manually returning the robot to the center of the arena.

Two-Motor Gaits

In order to demonstrate the diversity of gaits produced by the tensegrity robot, we ran a sparse sweep of motor voltages for two motors between 2.1V and 2.9V, keeping the third motor fixed at 0V (voltages below 2.0V do not produce motion in the motors). Each frequency set was evaluated over ten trials, each lasting 7 seconds. Figure 4 illustrates the maximum distance traveled by the robot at each measured frequency pair.

Figure 5 shows the variation in distances between trials when sweeping through frequencies for motor 1 while keeping the other two motor frequencies fixed at 2.1V and 0V respectively (corresponding to the the bottom row of the heat map).

Combined, these results hint at the complexity of the underlying space of distances achievable by the full range of three motor voltages: even when only using two of the three motors, there is significant variation in distances traveled, and there is a non-linear relationship between motor frequencies and distance traveled. Lacking any analytical approach to mapping motor frequencies to corresponding gaits, this suggests that automated trial and error via hill climbing or genetic algorithm might be the best way to discover effective gaits.

Three-Motor Gaits

Having demonstrated the diversity and complexity of two-motor gaits, we then manually explored three-motor gaits, searching for frequency sets which led to interesting and effective behaviors. We were able to discover sets which produced consistent clockwise, anti-clockwise, and linear locomotion.

Figure 7 illustrates the motion of the tensegrity over the course of 6 seconds of motion, at two-second intervals.

Maximum linear locomotion speed was on the order of 4cm/sec, or 2.4m/min. On an absolute basis this is four times faster than Paul's robot and half the speed of Kuizumi's. When normalized to body size, however, our tensegrity is considerably faster, while using simpler means of locomotion.

Videos of these gaits are available at the corresponding author's web page.

Controllable Motion

Most compellingly, these frequency sets which result in diverse gaits can be sequenced in order to *steer* the tensegrity in a controllable fashion. Since one gait can be used to propel the robot forward, and a second to rotate it, these gaits can be sequenced in order to produce controllable motion.




Figure 7: Video frames, taken at two-second intervals, illustrating motion of the vibrating tensegrity. Given a single set of motor frequencies, the robot is able to exploit its vibrational tendencies in order to move quite quickly – at speeds over 4cm/sec. Videos available on the author’s web site.

To demonstrate this we created an alternating sequence of forward-propagating and rotational gaits which caused the tensegrity robot to traverse a path within the arena. A video of this trajectory is provided on our web page.

This method of steering a tensegrity robot simply by changing its mode of vibration is unique, and a valuable example of how adding morphological complexity can sometimes simplify the task of control, allowing aspects of control to be “outsourced” into the complex mechanics of a structure. This is, therefore, a valuable example of *morphological computation* in a tensegrity robot.

As we continue our studies, and scale to increasingly large and complex tensegrity robots, we hope to uncover an even more diverse range of gait behaviors, leading to more interesting and effective means of controlled locomotion.

Discussion

There are several improvements we would like to make to move the system forward. Foremost among them is automating the discovery of effective motor frequencies with a physically embodied evolutionary algorithm, following in the footsteps of Harvey *et al.* at Sussex (Harvey *et al.*, 1997), Watson *et al.* (Watson *et al.*, 1999), and more recently Zykov (2004) and Yosinski (2011).

Like all physically embodied evolutionary robotics, however, we must deal with the issue of noisy evaluation, reliability, and consistency between trials. Some solutions lie at the algorithmic level (for instance, via multiple trials (Fitzpatrick and Grefenstette, 1988)), and some lie at the hardware level, for instance by using a more consistently smooth surface for evaluation. However, we want to avoid “sterile” surfaces – perfectly flat, perfectly smooth – since our ambition is to evolve robots capable of robust performance in rough and uncertain environments.

There is also the matter of the evolvability of the system itself. While the dynamics of tensegrities are complex enough to justify the use of automated optimization techniques (as opposed to analytical modeling) to generate gaits, our current control scheme is not very amenable to evolution. The use of a single control parameter (frequency) means that genotypes contain only three floating




Figure 8: *top*: A model of a larger 15-bar tensegrity, from (Rieffel *et al.*, 2010). *bottom*: a rig we have built to enable construction of the physical 15-bar tensegrity.

point loci, making the system somewhat too simple to explore with GAs. In the near term we are interested in more complex open-loop gaits, for instance, allowing each motor’s control voltage to change during a gait, and being able to specify its phase, amplitude, and frequency. In the longer term, we are interested in closed-loop control via Artificial Neural Networks (ANNs), with feedback provided by off-the-shelf micro-scale accelerometers (such as those used in smart phones).

We are also interested in using high speed video to analyze in more detail the varying resonant modes exhibited during different gaits.

Ultimately our goal is to build tensegrity robots with considerably more structural elements (while maintaining our relatively short strut length), such as the one shown in Figure 8. As tensegrities become more complex and irregular, they become increasingly difficult to analyze or model, fur-

ther necessitating our embodied approach.

Conclusion

Tensegrities are an appealing platform for modern robotics. They are robust, agile, and can quickly change shape, lending themselves to promising applications ranging from urban search-and-rescue to biomedical devices. However, these properties also make them exceedingly difficult to control through conventional means, particularly as the complexity of the robot increases. We have described a means of actuating and controlling tensegrity robots which treats their dynamical complexity as a feature to be exploited rather than as a liability to be suppressed. By designing the structure in order to maximize resonant possibilities, we can make the robot move simply by vibrating it at specific frequencies. This leads to a tensegrity robot which is much smaller and much simpler than existing designs, and yet outperforms in many regards.

More valuably, we have demonstrated how we can affect behavioral change merely by changing the frequencies at which our robot vibrates. Achieving behavioral diversity by exploiting mechanical complexity in this manner is a valuable example of *morphological computation*, in which increasing dynamical coupling can, paradoxically, reduce the cost of control. Given the pervasiveness of both tensegrity and dynamical coupling in biological systems, our hope is that this can lead to a deeper understanding of how mechanically complex living systems at all scales of life move and interact with the world.

Acknowledgments

The authors would like to thank RJ Hojnacki, Steve Stangle, and Julian Jocque, all of whom assisted in this research as part of undergraduate practica at Union College. Funding for undergraduate summer research by authors Mark Khazanov and Ben Humphreys was provided through the generosity of the Union College Summer Research Fellowship Program. Materials were paid for in part by several Union College Internal Education Fund (IEF) grants to the authors.

References

- Ahn, A. N. and Full, R. (2002). A motor and a brake: two leg extensor muscles acting at the same joint manage energy differently in a running insect. *Journal of Experimental Biology*, 205.
- Aldrich, J., Skelton, R., and Kreutz-Delgado, K. (2003). Control synthesis for a class of light and agile robotic tensegrity structures. In *American Control Conference, 2003. Proceedings of the 2003*, volume 6, pages 5245–5251. IEEE.
- Biewener, A., McGowan, C., Card, G., and Baudinette, R. (2004). Dynamics of leg muscle function in tammar wallabies (*m. eugenii*) during level versus incline hopping. *Journal of Experimental Biology*, 207.
- Blickhan, R., Seyfarth, A., Geyer, H., Grimmer, S., and Wagner, H. (2007). Intelligence by mechanics. *Phil Trans R. Soc. A.*, pages 199–220.
- Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6:3–15.
- Caluwaerts, K., D'Haene, M., Verstraeten, D., and Schrauwen, B. (2013). Locomotion without a brain: Physical reservoir computing in tensegrity structures. *Artificial life*, 19(1):35–66.
- Chan, W. L., Arbelaez, D., Bossens, F., and Skelton, R. E. (2004). Active vibration control of a three-stage tensegrity structure. In *Proceedings of SPIE 11th Annual International Symposium on Smart Structures and Materials*.
- Connelly, R. and Back, A. (1998). Mathematics and tensegrity. *American Scientist*, 86.
- Craig, J. (1989). *Introduction to Robotics*. Addison-Wesley.
- Daley, M. and Biewener, A. (2006). Running over rough terrain reveals limb control for intrinsic stability. *Proceedings of the National Academy of Sciences*, 103:15681–15686.
- Fitzpatrick, J. M. and Grefenstette, J. J. (1988). Genetic algorithms in noisy environments. *Machine learning*, 3(2-3):101–120.
- Fung, Y. (1993). *Biomechanics: mechanical properties of living tissues*. Springer, 2 edition.
- Graells Rovira, A. and Mirats Tur, J. M. (2009). Control and simulation of a tensegrity-based mobile robot. *Robotics and Autonomous Systems*, 57(5):526–535.
- Hannan, M. and Walker, I. (2003). Kinematics and the implementation of an elephant's trunk manipulator and other continuum style robots. *J Robot Syst*, 20:45–63.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the sussex approach. *Robotics and autonomous systems*, 20(2):205–224.
- Iida, F. and Pfeifer, R. (2006). Sensing through body dynamics. *Robotics and Autonomous Systems*.
- Ingber, D. E. (1998). The architecture of life. *Scientific American*.
- Iscen, A., Agogino, A., SunSpiral, V., and Tumer, K. (2013). Controlling tensegrity robots through evolution. In *Proceedings of the 2013 Conference on Genetic and Evolutionary Computation (GECCO '13)*.
- Koizumi, Y., Shibata, M., and Hirai, S. (2012). Rolling tensegrity driven by pneumatic soft actuators. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1988–1993. IEEE.
- Masic, M. and Skelton, R. E. (2004). Open-loop control of class-2 tensegrity towers. In *Proceedings of SPIE 11th Annual International Symposium on Smart Structures and Materials*.
- Matsuda, T. and Murata, S. (2006). Stiffness distribution control-locomotion of closed link robot with mechanical softness. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 1491–1498. IEEE.

- Motro, R. (2003). *Tensegrity: Structural Systems for the Future*. Kogan.
- Murray, R., Li, Z., and Sastry, S. (1994). *A Mathematical Introduction to Robotic Manipulation*. CRC Press.
- Paul, C. (2006). Morphological computation: A basis for the analysis of morphology and control requirements. *Robotics and Autonomous Systems*, 54.
- Paul, C., Valero-Cuevas, F. J., and Lipson, H. (2005). Redundancy in the control of robots with highly coupled mechanical structures. In *Int. Conf. on Intelligent Robots and Systems*, pages 802–808.
- Paul, C., Valero-Cuevas, F. J., and Lipson, H. (2006). Design and control of tensegrity robots for locomotion. *IEEE Transactions on Robotics*, 22(5).
- Pfeifer, R. and Bongard, J. (2006). *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press.
- Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, pages 1088–1093.
- 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.
- Sciavicco, L. and Siciliano, B. (2000). *Modelling and Control of Robot Manipulators*. Springer.
- Shibata, M., Saijyo, F., and Hirai, S. (2009). Crawling by body deformation of tensegrity structure robots. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, pages 4375–4380. IEEE.
- Skelton, R. E., Pinaud, J. P., and Mingori, D. (2001). Dynamics of the shell class of tensegrity structures. *Journal of the Franklin Institute*.
- Sultan, C. (1999). *Modeling, design, and control of tensegrity structures with applications*. PhD thesis, Purdue University.
- Tibert, A. (2002). *Deployable Tensegrity Structures for Space Applications*. PhD thesis, Department of Mechanics, Royal Institute of Technology (KTH).
- Tibert, A. and Pellegrino, S. (2003). Review of form-finding methods for tensegrity structures. *International Journal of Space Structures*, 18(4):209–223.
- Trimmer, B. (2007). New challenges in biorobotics: incorporating soft tissue into control systems. In *IEEE International Conference on Robotics and Automation*.
- Valero-Cuevas, F., Yi, J., Brown, D., McNamara, R., Paul, C., and Lipson, H. (2007). The tendon network of the fingers performs anatomical computation at a macroscopic scale. *IEEE Trans Biomed Eng.*, 54:1161–6.
- Vogel, S. (2003). *Comparative biomechanics : life's physical world*. Princeton University Press.
- Wang, B.-B. (1998). Cable-strut systems: Part i - tensegrity. *J. Construct. Steel Res.*, 45(3).
- Wang, N., Naruse, K., Stamenović, D., Fredberg, J. J., Mijailovich, S. M., Tolić-Nørrelykke, I. M., Polte, T., Mannix, R., and Ingber, D. E. (2001). Mechanical behavior in living cells consistent with the tensegrity model. *Proceedings of the National Academy of Sciences*, 98(14):7765–7770.
- Watanabe, W., Sata, T., and Ishiguro, A. (2008). An efficient decentralized learning by exploiting biarticular muscles - a case study with a 2d serpentine robot. In *Robotics and Automation, IEEE International Conference*, pages 3826–3831.
- Watson, R. A., Ficici, S. G., and Pollack, J. B. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In Angeline, P. J., Michalewicz, Z., Schoenauer, M., Yao, X., and Zalzal, A., editors, *Proceedings of the Congress on Evolutionary Computation*, volume 1, pages 335–342, Mayflower Hotel, Washington D.C., USA. IEEE Press.
- Yosinski, J., Clune, J., Hidalgo, D., Nguyen, S., Zagal, J., and Lipson, H. (2011). Evolving robot gaits in hardware: the hyperneat generative encoding vs. parameter optimization. In *Proceedings of the 20th European Conference on Artificial Life, Paris, France*, volume 8, page 12.
- Zykov, V., Bongard, J., and Lipson, H. (2004). Evolving dynamic gaits on a physical robot. In *Proceedings of Genetic and Evolutionary Computation Conference, Late Breaking Paper, GECCO'04*.