SOFT ROBOTICS
Volume 2, Number 2, 2015
© Mary Ann Liebert, Inc.
DOI: 10.1089/soro.2015.28999.btr

INTERVIEW

An Interview with NASA Principal Investigator Vytas SunSpiral: Expert Opinion on the Advantages and Limitations of Soft Robotics

Interview by Barry Trimmer

Vytas SunSpiral is an entrepreneurial researcher moving fluidly between leading startups and building research labs to explore cutting edge robotic technologies. During the last 20 years, he has been the founder and CTO of multiple startups and launched a number of robotics projects at NASA. He has served as an advisor and consultant to startups, and he is currently the Principle Investigator of the Dynamic Tensegrity Robotics Lab (DTRL) at NASA Ames Research Center and is a Fellow of the NASA Innovative Advanced Concepts (NIAC) program. Vytas graduated from Stanford University (1998) with a BA in Symbolic Systems and an MS in Computer Science, with a robotics focus in both.

Good morning, Vytas. Thank you for making the time to speak with us about your work in robotics. We are particularly interested in getting some of your thoughts and ideas about how soft robots might factor into space applications. I would like to start by having you introduce yourself. Please tell me about your current role at NASA.

Sure. Thanks, Barry. Really pleased to be sharing my thoughts with the *Soft Robotics* community today. I currently work in the Intelligent Robotics Group (IRG) at NASA Ames Research Center. I am a contractor with SGT, Inc., and I lead the Dynamic Tensegrity Robotics Lab, (DTRL), which is primarily funded by a NASA Innovative Advanced Concepts (NIAC) project to develop a tensegrity robot for planetary landing and exploration.

We are looking at the development of these robots from the full range of hardware and control perspectives, and how to properly simulate them so that we can explore the robots' use for a wide range of different applications in space and on Earth.

As such, I have ended up engaging very heavily with a lot of different university partners across the nation and have even had international students coming to spend time in our lab to pursue their interest in tensegrity robotics. So really my role ends up being sort of the classic PI/entrepreneur, setting the tone and vision, attracting the students, researchers, and funds, and trying to push the grand vision forward in all directions at once.

Right, absolutely. And I think the sort of robotics you do involves a lot of different disciplines. What exactly is your background training?

This is an extremely multidisciplinary effort. My original undergraduate degree from Stanford was, appropriately enough, a very multidisciplinary degree called Symbolic Systems. It is actually in the Humanities department and is a combination of computer science, psychology, computational linguistics, and philosophy. The aspect of the Symbolic Systems program that attracted me was its focus on artificial intelligence (AI). My underlying interest all along was understanding the human experience and understanding intelligence as part of that and how we think and feel and do all the things we do.

But the process of studying AI led me into robotics. The idea of just symbol manipulation, sort of pure information processing as a means to understanding intelligence, eventually felt like a dead-end because it was ungrounded. Every symbol could be defined by another symbol. A tree has branches, branches have leaves, and leaves have photosynthesis. But I do not really need to know about photosynthesis to know what a tree is if I can hug a tree, if I can sit in the shade of a tree, if I can be out of the rain from being under a tree. So if you can experience things then you know what they are; you do not need to go down the full intellectual path. You just need to have some experiential grounding.

That was really where I got into robotics. I got my Masters in Computer Science focusing on robotics, also at Stanford, and just carried on from there working in the robotics industry, initiating a couple of startups and eventually moving to NASA and continuing to pursue these ideas. And along the way, getting deeper and deeper into the mechanisms, I sort of combined this intellectual interest in AI and all this robotics experience with my own personal experience in human motion.

I have done a lot of martial arts, dance, yoga, and other sorts of human motion training throughout my life. I really saw from experience that the traditional way of building robots was very fragile compared to how we actually move. Though we make the robots out of strong metal, this very rigidity makes their control and balance fragile compared to what biology is capable of in its flexible, dynamic, tensile ways.

I started combining these ideas in recent years to really understand how the human body works as a tensile network and to apply that to the robotics field in order to unlock the advantages we see in this sort of tensile, soft materials approach. Of course it is a challenge for engineering, but here we are doing it. So even though I started with an interest in AI, I am now thinking heavily about structures and materials and actuators along with controls, neuroscience, machine learning, and even physics simulation engines.

It is very interesting that there are a lot of folks who have come into this area from a computer science, AI background. I am thinking particularly of Rolf Pfeifer and his concepts of embodiment. It sounds as if you have a parallel trajectory. You are clearly interested in modifying the way we think about the design and operation of these devices. I guess that is where your interest in complex structural pieces of robots comes from.

I know that you have worked on a variety of nontraditional robots, wheel-on-leg type machines and tensegrities. So what is the motivation behind those novel systems?

The challenge here, and I think the reason why you see folks with my kind of background ending up in this new mechanisms area, is that robotics is a system—you need to design your structures and your controls to work well together. You cannot actually separate them out into individual problems and assume that you have some control theory that works on any structure. As an extreme example, the best AI or control theory in the world is not going to make a brick fly. Mechanism and control are always very interrelated with each other, and yet the challenge is, of course, that these are all deep disciplines in their own right; it is very hard to fuse them all together.

And so a couple of things have really interested me as I have been studying mechanisms as well as the human body. One is noting that our robotics technology as a field has been largely built upon a foundation of position control. That was our early successes with robotic manipulators and factory automation. If you could have a very stiff and precise mechanism you could control its position exactly, and then you could accomplish prescribed tasks. Since then we have been trying to add things like force control and compliance on top of that foundational layer of technologies.

If you look at biology, it is the other way around. Everything starts with forces. Forces are what matter. Position is something that comes much, much later. Fine control position is a recent evolution, if you will. But it is really the forces that dictate what you are doing and what you are trying to do. If you look at the early evolution of motion, the first step in moving in the world is to apply forces to the environment. You later figure out how to apply forces to the environment such that you end up in a specific location or some specific part of your body, say the end of a fin or the end of an arm, ends up in a specific location. So those position controls are the more recent evolutions.

The limitations of our current approaches really became apparent to me when I was working with the ATHLETE robot (*editor comment:* All-Terrain Hex-Limbed Extra-Terrestrial Explorer), which is that wheel-on-leg robot you mentioned, and which was built by Brian Wilcox, a brilliant engineer down at the Jet Propulsion Laboratory (JPL). It is a very large robot with six legs, each with seven degrees of freedom (DOF), and the whole thing can stand 4 m tall. The legs on ATHLETE are standard serial chain arms with joints that are rotary motors.

Its sheer size and mass really made evident challenges that could be ignored on smaller robots. Specifically one of the things I really saw from it was how the whole thing acted like a giant lever arm.

I was developing walking algorithms for ATHLETE, and we would have it out in the JPL Mars Yard, which is a flat dirt yard with some rocks you can move around. I kept finding situations in which we would be on essentially flat ground, with some subtle undulation or a light bump or two here and there that we as humans would ignore. And yet one of the ankle joints, for instance, would be completely over-torqued, and we would have to stop the program and reconfigure the robot to reduce the torque or that joint so that we could move it again.

This is really challenging and also very surprising because it is not like it was doing anything challenging. The robot, under careful control could step up and over a very large obstacle and even climb over little cliffs and stuff like that. Yet here we were on flat ground getting saturated.

What became apparent was that because of the rigid connections we were essentially getting a 6-m-long lever arm from one side of the robot to the other. The robot is about 4 m tall, but it is wider than it is tall, so if you have just got the right configuration the tiniest little bump can be magnified through that rigid leverage and completely saturate a joint.

That was really an eye-opener, and got me to start thinking about the forces that were flowing through the structure and the leverage that was being induced and how hard it was to manage that. At the same time, I was, as I mentioned, very interested in human motion studies and as a result of my activities spent a lot of time getting physical therapy. I have had a couple of knee surgeries, and that taught me a lot.

As I spent time working with the physical therapists I heard from them their theories of how the body works. Their theories were a lot more about long chains of tensile connection that would go from end to end in the body, and it is via these tensile chains that the body transfers force around. That was really interesting to start thinking about, that through a tensile chain you can redistribute force rather than accumulate it in a joint. That distribution of force allows the whole body to participate in absorbing stresses and absorbing interactions with the terrain. This tensile chain model also allows multiple muscles to participate, and you end up with these long kinematic chains that work together to generate motion. And of course in the human body the dysfunction comes when somewhere along that chain, something gets too tight or compressed or whatever, and then the other parts of that kinematic chain are overloaded and you end up with a sore back or aching shoulders, etc.

But this is a very different model to think about how forces flow through a structure compared to a traditional robot. And so that is the origin of my interest in tensegrity structures, which are an extreme example of that tensile network in which none of the rigid rods directly touch each other. It really forces you to think deeply about this tensile network and how to manage it and model it and control it.

That is a really nice description and justification for heading in that direction, and I completely understand it. But to play the devil's advocate, if you use a soft, highly deformable material or you go sticking a nonlinear cable system into a robot to solve these issues, some people might argue that we are adding just as much uncertainty and complexity to the system. So we are replacing one complex control problem with another, less tractable mathematical problem. What do you think to that? Is there a path to easier, better robot control using soft structures?

I am not sure if I would say it is going to be easier control, because we know that the control of these very compliant systems is difficult and requires developing whole new approaches. But I do think it will result in better and more capable systems in the long run. While it is nice to think of traditional control approaches being precise and exact, it is worth considering that once you hit seven and more DOF, you have an infinite possible solution set for the inverse kinematics, and you stop having clean, analytical solutions anyway. So everything we do beyond simple six DOF mechanisms ends up with a bunch of heuristics in the solution of the controls.

This idea of having clean, highly precise, closed-form solutions actually goes out the window very fast for any type of approach. If you are going to do that anyway, let us open up the gates even further and just ask the question of what is the right way to control stuff. That is one observation.

The other observation is that there are aspects of how a structure responds that you want to take advantage of at the mechanical layer rather than the control layer and our soft robots. These tensegrity robots all tap into this idea that you want to have passive dynamics in your system that are going to respond faster than whatever control system you may come up with. The idea is if you build a rigid structure, yes, you can come up with examples of a controller that gets it to respond very quickly to interactions with the environment, but there is always going to be some sort of interaction that is even faster than your control is capable of doing.

Consider high frequency impacts as experienced when walking. Every time you take a step and you have a big impact that propagates through the structure, you either get really stiff and heavy to manage that or you have a very fast active controller. When you start running or jumping, you start doing really dynamic actions, and the frequencies get higher and your controllers must get even faster. It just becomes very difficult.

Since I mentioned the passive dynamics, it is worth considering that the most energy-efficient robots all take advantage of passive dynamics in their mechanical structure, much like biological systems do. You will never be able to achieve those levels of efficiency with a fully actively controlled system.

Beyond energy efficiency, if you are trying for a full active control of a very high DOF robot, then you need to get information about the entire structure from all the sensors so that you can make rational, centralized decisions about the overall kinematics of the entire structure. We know that doing inverse kinematics for high-DOF structures is computationally expensive and involves many heuristics. You are going to do that at some hundredths or thousandths of hertz so you can deal with high-frequency impacts, and that starts getting really challenging and still ultimately kind of fragile. This is because your structure has the wrong mechanical properties for what you are trying to do. That is really at the heart of the issue—what are the mechanical properties of the structure? If you want a system that moves, build something dynamic with inherent motion. If you want a system that stays still, like a house, build something rigid with static mechanical properties.

But taking that approach also requires that you have got the appropriate sensors collecting information quickly enough and handing it off for the closed-loop control. In animals, the rate of information transfer for any given sensor system is relatively slow. We are talking about a kilohertz in a fast neuron. A lot of parallelism is built into animal senses, but biology is pretty slow compared to what can be achieved in silicon electronics. So animals have to solve that problem mechanically. And we see animals using mechanical compliance to achieve control without any direct motor input to the system.

And that is the heart of it, right? Even though in the short term, going down this path the controls are harder because we just do not yet know how to do task-specific activities the way humans do. But in the long run what you see is that if you use the right mechanical properties, you simplify the control problem. The key we see in biology is that you distribute the control solution.

I think it is very important for us to recognize that the work on decerebrated animals, in which the spinal cord has been severed below the brainstem and the animal has been kept alive. There is absolutely no involvement of the brain, no involvement of the motor cortex, and yet the animal executes complex coordinated behaviors. They can walk. They can trot. They can go through different gait cycles depending on the speed of treadmill that they are on.

And this is happening in a distributed decentralize manner in the spine and below, with very slow-moving neurons. In the spine there is no centralized point where all the sensors from the whole body come together, yet rather it is an emergent property of the system that enables these behaviors. We want to develop this low-lying, body-based capability of basic locomotion, which does not have any sense of direction, does not have any purpose, does not accomplish any tasks, but it just gives this quality of being able to locomote. With that capability, your high-level controller suddenly becomes a lot easier, right? If you want to take this hundreds of DOF system and navigate over complex terrain, you have a much easier job of just saying, ''Hey, go this way, go that way,'' at a high level.

This underlying very robust body and embedded control system take care of a lot of the details. And that, I think, is really what we want to see in the future—this sort of decomposition into multiple dynamical control systems at different levels of control resolution working with each other. Roughly what we see are these the low-level mechanical controls with responsive motions and patterned behaviors in the spine, while in the brain, we find higher-level controls focused on task planning and the ability to modify those underlying patterns and steer them and shape them to accomplish goal-specific tasks. I think any one part of that puzzle becomes a lot easier once you have the right physical system in place.

So there is a hierarchical aspect to robot control. First, the basic mechanical structures and materials then a highly distributed local set of controls including both mechanical and neural components. And on top of that are built the central organizing principles for transitioning between behaviors and making decisions on path finding, navigation, and that sort of thing. You have got this hierarchy of control systems and they are all working together.

Perhaps an important point is that by thinking of robots in this way, we are going to design them differently. Although we may still encounter phenomenally difficult problems of control by starting with a better design based on these principles we are likely to get better outcomes.

And here is an interesting thing to add to that, which is that you can only be successful at these truly distributed, emergent control approaches if you start with a soft robot body. In my case, I particularly think the tensegrity body is a good one, but other forms of soft robots should be perfectly good, too. And here is why. If you look at the history of distributed controls for robotics, you see lots of people have been trying this over time. Inevitably most of the efforts end up dying under the weight of information sharing. People have essentially taken a centralized control model and tried to distribute it across a lot of controllers, and then they end up having to send information back and forth between all these distributed controllers to try to maintain some state, some sense of global state.

As the system gets more complex, that just bogs down. One of the reasons for this is because if you have a rigidly designed robot, the actions of any local controller can have disproportionate global impact. Go back to my example of the ATHLETE robot, which could step on a very small bump, causing a 6-m-long lever arm to induce massive forces into one of the joints. That is the problem with the rigid robot, right? A small action at a local joint or a local controller, due to the rigid leverage of the structure, can impart huge forces elsewhere in the system.

You take a soft robot, on the other hand, and just to use my tensegrity robots as an example, every actuator is integrated into this tensile network, but they are disintermediated, if you will. They are in this compliant network, and if any local actuator does something, its action is primarily local. It does have a global impact; it will change the global structure, but it changes it by modifying the balance of forces. The whole structure sort of adapts to the lowest energy state for the action of that local controller. Thus, you do not end magnifying the effect of the local control via leverage to cause a giant global impact.

It is this natural and passive rebalancing of the forces that is the truly valuable physical property of tensegrity structures. In some respect, it does not matter whether that force is an external force or the internal force of a muscle or actuator. No matter how the force is applied, the structure is constantly adapting to it in a natural way and diffusing it through the soft structure, through the tensile network.

That makes distributed controls realistic and possible. And possible in a way in which you do not need a lot of explicit information sharing. The only information that really matters to any one of the controllers is the forces it is currently experiencing. And you can sense that locally; you do not need to be passing around global state or global information. You just need to know what the local environment is like, and then you can base your actions on that. You can get coherent behaviors out of that.

Now, to get to the point of doing goal-directed behaviors, you will start needing to send down some high-level information, but it will not require a global state being passed through all the local controllers.

Right. So this is very interesting to me because there is a parallel language used in animal neuromechanics in which people have been thinking about synergies. The idea is that a nervous system does not necessarily have to control individual actuators but instead they are collectively operated at the task level. This frees up the central controller to do more high-order planning of complex motions. I think there is a lot of similarity.

So do you see yourself motivated by the way nervous systems control bodies? I know we have talked a little bit about coupled oscillators and central pattern generators. Is that something that you think can be used in the designs of controls for robots?

Absolutely. I am very, very interested and motivated by that. And again, you know, this is the existent proof that we have, that by using these approaches, biology is able to do really amazing stuff. And so let us understand the properties at play there.

We have been using Central Pattern Generators (CPGs) in a number of our research efforts. We are using the abstracted dynamical systems versions of it where we model the CPGs as oscillators that are then coupled with each other, rather than getting all the way down to the modeling individual neurons. Most of the time when you are modeling the neuron it is still an abstraction anyway; those things are really quite complex little animals.

There are many things that I love about CPGs. I think that the most important aspect is this idea of using rhythm and synchronization. There is this great book out there, I am sure many folks have read it, by Steven Strogatz called *Sync: How Order Emerges From Chaos*. He is a mathematician, and he talks about how rhythmic systems that interact can synchronize.

This is the heart of distributed computing and distributed controls, and I think the heart of how biology functions, which is that if you have rhythmic things, such as CPGs, which are rhythmic oscillating controllers, then they have a mathematical property that they can synchronize. So this is fundamentally a mathematical property that allows you to have emergent coordination from a distributed set of controllers.

And this is brilliant, right, because as you know, in biology we have no central control. There is no single central CPU in our bodies anywhere. Some people say, ''Oh, the brain, that is our central processing unit.'' Well, it is not because the brain is made of millions of individual neurons, and each neuron is its own living animal; they are working together collaboratively to come up with a coherent behavior.

And of course we already know that you can get very coherent behavior out of the spine with no brain involvement whatsoever. And the spine is well known to be decentralized. It is modular. And so I think that rhythm and oscillation must be at the heart of really complex and coordinated biological controls because they enable coherent behavior to emerge out of a decentralized system. It is powerful. It is fascinating.

Synchronization-based computing should also have the properties of being more robust to noise and continuously adaptable. A classic quality of CPGs and similar dynamical systems controls is that they exhibit smooth transitions between ''modes''—which is exactly what we see in biological behaviors. Every control system I've seen that uses discrete modes suffers for complexity in how the system transitions between modes. While there are some hacks that make it work, they rarely scale to large systems with many modes. The advantage of smooth transitions is that they allow for robustness to noise and edge cases.

The concept of synchronization-based computing replicates and ripples throughout everything you see in human psychology even and human behavior. Things that have traditionally challenged us in the field of AI, like classification and set creation and, generalization of concepts, those are capabilities that you can imagine easily falling out of a computational process built around synchronization. A computational system built around synchronization means that things that are similar enough click together, and that is your foundation of your computation, rather than binary logic being your foundation. Then all the stuff that we see in human behavior, from how we see things, how we classify things, how we just pick up each other's dialects, how we like to build relationships with each other, how we like to see connections between people and ideas and all that, all these types of very human qualities become the natural outputs of that type of control theory.

In fact, it is amazing that we do math at all, right? That is something binary logic is great at, and we humans figured out how to do that, but it does not seem to be at the foundation of how our computing system works. You know, our neurons did not come around to do math; our neurons came around to move. That is the reason neurons exist—to control motion. Math is a much later evolution, very recent really.

Philosophically, the interesting point is that we built technologies and developed computing in order to solve problems that are hard for us—like complex mathematics. Now, we are learning how to build the machinery and computing for tasks that are easy for us, like moving in the real dynamic world, and it turns out that is a harder problem.

And so, going back to that earlier critique, that by using compliant materials and tensile networks, we have made the modeling of the system more difficult, I would answer that by saying that we don't need to make explicit models of the system. Instead I am going to develop a system that is going to create the behaviors I want without ever explicitly modeling the system. This seems like something that can be achieved through these rhythmic, synchronizing, distributed systems like CPGs. Of course, we have a lot to learn still, and you end up with big challenges like machine learning and parameter tuning. But I currently have a 3-week-old baby, and I can tell you he is busily using some form of parameter exploration and is tuning his system.

This is interesting and you have explained it very well. It is interesting because in my own research we have come to some of the same conclusions. In fact, we just started to develop a model-free control based on state transitions for our little soft, crawling, caterpillar-like robot.

Oh, good! Just a last thought to finish up on the CPG ideas. People have studied CPGs in the past and most of them have been applied to rigidly connected robots because that is just where the technology has been. That has been fine; it has shown a lot of value in these types of controllers. But I think the real value, the real power of these model-free rhythmic control networks gets tapped when you attach them to a soft robot. Once again, it is important that the properties of the controller match the properties of the mechanism.

And something like the tensegrity robot is inherently oscillatory. It is inherently rhythmic and vibratory. I think a lot of our other soft robots have those properties. The beauty here is when you have a compliant, rhythmically oscillating physical structure, and you combine it with a control approach that is inherently built on rhythm and synchronization. It enables that synchronization quality to start bridging between the physical body and the control system.

Now you can entrain the controller and the body together and allow that property of emergent coordination, emergent synchronization to happen between your information domain and your physical domain. This goes beyond just the controller and the robot. If you do this right, it will enable the controller to entrain with the combined dynamics of the robot interacting with the environment, enabling natural adaptation to complex natural terrains. That is something you need to enable the use of soft compliant systems. Again, it is like all these components have been studied in isolation, but putting it all together in the right way I think is when the system really starts to be enabled.

I agree, and I think one of the challenges is to get folks who are extremely good at designing proficient engineered systems to let go a little bit and realize that this robot is going to have to develop a little bit of learning in order for it to work properly.

Turning to another area of your expertise. Is there a special aspect to extraterrestrial space applications that you think feed into this whole concept?

The key question is what are the advantages that are being enabled? I think soft robots play a role in space exploration for the new mission concepts they enable, and for the robustness and reliability they could enable. So far, the primary area of our tensegrity robotics research within NASA has been the SUPERball planetary lander and exploration robot. It is our spherical, icosahedron tensegrity ball that provides a new mission capability; it is lower mass, and it reduces the risk of planetary exploration. We are enabling a new mission concept by taking advantage of the structural ability to absorb impact forces of landing, and also providing the mobility system. This combines what is normally two different features into one system, and thus it saves mass for the overall mission.

So traditionally a mission uses an airbag (or other landing system) once during landing and discards that airbag and all the mass associated with it, and then there is a separate robot that goes driving around on the planetary surface. In our SUPERball project, we are saying, hey, what happens if you build a robot that is compliant and adaptable, in this case the tensegrity sort of sphere-like robot, such that when it lands it can land at 15 m a second, equivalent to how the MER rovers landed in their airbags on Mars. It can hit that hard and fast and survive. And then you do not need to discard anything—the robot itself is its landing system.

Once it has landed, it can move around and explore. Well, that changes two things. One, you have potentially saved some system mass because you did not have to carry along this extra airbag that you only used once. Second, it completely changes the risk profile for exploration! Right now if one of our robots on Mars fell down a 2-foot cliff, people would freak out and would not be sure if it would survive.

With our current rovers, we do not want to get too close to the edges of cliffs or other places where it might slip or fall. And, of course, there are limitations to the steepness of slopes that you might want to explore with the rover because you might get to the point where the soil starts slipping, or maybe the rocks that its wheels were on suddenly slip and tumble.

So if you have a robot that can fall from orbit and land safely, that changes everything. Now if you want to explore the edge of a cliff, you might be willing to do so. You might even be willing to roll off the edge of the cliff, you know, to get to some point at the bottom that you cannot otherwise reach. You may be willing to take all sorts of other risks that would be unacceptable with today's rovers.

So that is one avenue that we are pursuing that I think has a lot of impact on future space missions. And really at the heart of it is this idea that as we go deeper and deeper into space and we explore planetary surfaces and subterranean surfaces further and further away, we are not always going to be able to predict what is going to happen. If we move slowly and carefully on flat, stable ground, we can usually predict what is going to happen. But if we want to go into more complex terrains, into boulder fields, or traversing across crazy icy terrains on the surface of Europa where the ground has been constantly cracking and changing, we are going to have really new challenges ahead of us.

Having a system for which you can sort of guarantee the fundamental reliability of its locomotion capabilities I think would be awesome. I think it would be really enabling, kind of like that decerebrated cat—it is always going to be able to walk and it is going to keep its balance and land on its feet if it falls—that is what we want at the lowest level of a mechanism. It is just going to be safe.

Then we have to figure out how tell it what to do, you know, and that is the currently the tricky thing. A lot of people think about autonomous systems for space exploration. They think about robots that are just going to go exploring and finding science on their own and doing it and all that. I think that is unlikely. We are always going to have human controllers in the loop because you can have armies of scientists down here that are going to look at the data and say, ''Hey, that rock over there, that is interesting!'' or ''Hey, that looks like a fossil!'' Those types of insights are going to be hard to automate.

But what we can automate is that sort of physical reliability so that no matter what happens, the system is likely to be okay. To do that, we need both the controls and the physical platform to have that reliability and to have that adaptability so that it is not fragile to the unexpected. That is really where I think biology has won over every machine we have made so far. Biology does not just have energy efficiency or task completion as a goal, it also has fundamental physical survivability as a high-level goal.

So that is something that we need to start bringing into our machines if we want to get out of highly controlled and prescribed environments, be it the factory floor or the laboratory or very contrived settings. If we want to have robots that are able to work with us in our day-to-day lives, in the very complex and messy urban settings and natural environments that we move through, those systems need to be very physically robust—and not robust in the big, heavy manner but robust in the graceful ''if I get bumped I am going to be okay'' manner.

I think that level of robustness applies both here on Earth and applies to space. Robustness is really the key for planetary exploration—how do we make systems that are going to survive the unexpected?

I think these are all exciting prospects, and I really hope that research in this area is going to be increasingly funded. So what do you think we should be focusing on right now? If there were one or two things that we should be investing in as an economy, what would be the biggest ''bang for the buck'' in moving this field forward?

We need to keep pushing for what does not exist right now, which is a really good artificial muscle. We have lots of bits and pieces of it, and there are actuators that are almost compelling, but nothing really wins. Nothing really compares to real mammalian muscle. That is going to be our limitation. It is going to be very hard to replicate the physical dynamics of a biological system without having actuators that have similar dynamics. The physical properties really matter.

Some of the amazing aspects of the human muscle, or the mammalian muscle, are its ability to hold position without using energy, its elasticity, and its complex force generation behaviors across length and position and strain.

An important insight that is often missed is that we should not think about the human or animal muscle at the level that we normally talk about muscles: biceps, hamstrings, and so forth. A bicep or a hamstring is actually a very large conglomeration of the actual underlying actuators. The individual actuator units are the muscle fibers, and there are hundreds or thousands of fibers in any one muscle like a bicep.

It is the muscle fibers that matter. Those are the active components. And by combining a bunch of fibers into different patterns, both parallel and in-series, and by controlling the angles at which they relate to each other and the angles at which they are imbedded within the elastic material that they are imbedded in, that is what creates the very wide range of top-level muscles that we see in all their amazing different capabilities.

So we really need to be focused at that muscle fiber level, which means small. They do not necessarily have to have a high force production in any one fiber because hundreds or thousands of them will work together in parallel and in series to produce the larger capabilities that we see in the gross scale skeletal muscle.

The thing about the muscle fibers is that it is going to have to be cheap. It is going to have to be easy. It is going to have to be easy to control, in terms of sending probably just an electrical signal to them, rather than having lots of hydraulics or lots of valves, which all end up being heavy and expensive right now. Of course, some people are working on cheap flexible valves, so maybe they will solve it that way.

This highlights something that I think is an emerging theme in this field. Another huge advantage of muscle is that it directly turns chemical energy into mechanical energy. Muscle is not especially efficient, but because it sits in its own fuel source and synchronizes all those little molecular motors, it is a very effective system. I think once we figure out a way to engineer something like that, we will have actuators even better than muscle. I sincerely believe that once we know how to make something and we have the capability, we can engineer things that nature will never evolve because nature is stuck with the bits it already has, whereas we can use materials that would never find their way into biology.

Yes, exactly, and I think that one of clues along the way to understanding how the muscles work is to look at the large patterns of how they are deployed in the body, and there I feel that you see a lot of spiraling that is involved. Spirals both at the physiology level and down at the muscle fiber level as patterns similar to twisted cable actuation. Inside the sarcomeres, when we talk about myosin and actin sliding past each other, that is actually happening in a spiraling manner, which is then twisting up the protein Titin, which acts kind of like a passive twisted cable. So, going back to your question, I think good artificial muscles are going to be a combination of material, structure, and control, and once we understand it, we will be able to tweak it. Actually, if you look around nature, it already does a lot of that tweaking, and we might be able to combine the best innovations from different species together.

Since spirals are part of my name, I strongly recommend looking for them, and if you do, you see spirals all throughout the body, not just at the muscle fiber level, but even in the way our overall body is organized. If you look at the activation patterns of muscles coming down our limbs, they are big spirals, and they all come together right at our belly button, which is the spiral of our umbilical cord that we branched out of. So you see it as a very foundational pattern of the connective tissue, and the fascial and muscle alignments of our bodies. And if you understand that pattern, and you can control the spirals, then you really can control the motion effectively. As someone who has studied a lot of dance and martial arts, I can tell you that if you want to control what you hand is doing, start by controlling your belly-button—the source of all the spiral patterns in your body.

This has been a wonderful discussion. Do you have any comments about things we might not have touched on that you think would be important to this area?

Yeah, I would like to share an idea that I have been stewing on and have not yet fully articulated in writing. I would like to shed some light on an area of active debate about modeling the gross physiology of animals as tensegrity structures. As I've discussed, there are a lot of reasons to model physiology this way, starting with the fact that our bones are not rigidly hinged to each other, which allows them to move in complex ways relative to each other, constrained by the tensile network of our soft tissue. Yet there is always this ongoing debate about the details of the bones touching and passing compressive forces, or actually floating slightly apart, like a true tensegrity structure. The debate is really between the views of the bio-tensegrity purist and those who are confident that the bones bear compressive loads across the joints. And my claim here will be that it doesn't really matter. That really boils down to the question of ''Does the tensile network completely control forces of five or six axes across a joint?'' In many ways the distinction is minor, and we have a lot to learn before we even need to settle that finally.

So, to explain further, the thing that has struck me of late is that if you look at how our bones move relative to each other, you have these complex facets that are rolling, gliding, and sliding in relation to each other in very complex maneuvers as our joints do things. And there is a very low friction surface. And there is no hinge that holds them together.

So you see that there is actually this six DOF of how any two bones relate to each other. They can twist a little bit, and they can roll, and so forth. These DOF are elastically constrained by the tensile network, but they are never eliminated. And if you have ever gone to a chiropractor you know how a very small modification of the relationship of two bones to each other can completely change the rest of your tensile network and your experience of freedom to move and in terms of pain or dysfunction in your body.

So what you are really doing is using this tensile network to move these bones relative to each other, and very small subtle shifts and twists and rotation of these bones change how those two complex curved surfaces relate and move relative to each other. It is important to notice how very small changes in the joint can have very large impacts in how the distal ends of the bones move relative to each other.

And so I really think of this as we can push off to the future this question of whether the bones are touching or not. In some respects, it does not matter. Imagine that two bones are touching in a joint while you are holding some static pose. If you think of a frame of reference fixed to that point of contact and assign the z-axis to be along the line of compressive force transfer between the bones at that location. To hold that pose, the tensile network of your body must hold all five of the other DOF under control, with forces zeroed out, otherwise there would be motion between the bones given the very low friction of the joints. The tensegrity purist would say that the z-axis is also zeroed out, and there is no compressive load transfer. I think we need can first focus on figuring out how to control those first five DOF with a tensile network, and then we will be in a better position to ask if the body also controls that final z-axis of compressive load transfer or not.

That is really at the heart of it is that this tensile network is very precisely trying to control the forces at the point of contact, or almost contact, between the bones, and that relates to controlling the position of the bones relative to each other.

If I understand correctly, you are suggesting that careful engineering of the connections, the nodes, if you will, where the friction varies depending on the angle at which the bones are approaching one another.

It is not just the friction environment. I guess what I am trying to say is it matters that the bones are not hinged and that is actually a very useful quality that they move in six DOF relative to each other because it goes back to this idea of enabling the system to passively adapt to the loads it is experiencing.

Yeah, okay.

And so allowing the bones to shift relative to each other within their tensile network and to respond to forces gives you a very, very wide range of workspace and response space that you can deal with. And some of that shifting is very subtle. It is slight rotations in the bones relative to each other, which changes the trajectory across the articular surfaces that relative motion between those bones is controlled by.

And small changes have a big impact. If you imagine these two complex curved surfaces of the bone facets, if you transcribe a slightly different path across that, you can have a very large impact in terms of how those bones move. So very subtle control of how those surfaces relate to each other is very important and really enables you to do a much wider range of things than if you had a rigid hinge between those bones.

So if we understand it well enough we can perhaps engineer them to essentially reconfigure the robot, to make the robot perform in a completely different way, given a certain set of circumstances.

Right, and so not only is it the surfaces that matter—given those surfaces, it is the tensile network that manages motion across that joint.

And the force distribution.

And that is why we need to think about artificial muscles at the level of the muscle fiber because you need thousands, or maybe millions of actuators to control all the angles, all the degrees of motion that are possible.

If you look at any one muscle in the human body, the easy idea is to think, oh, I am either engaging my bicep or I am not. But that is not true. What you actually see is differing activation at the muscle fiber level within one muscle. You can have some muscle fibers in the hamstring contracting, while other muscle fibers are extending at the same time, depending on what the muscle is trying to do. It is the fibers that are the actual actuators, and we have thousands and thousands and probably millions of them. I have never counted.

And you know, if you look at any of the large fan-shaped muscles, like the trapezius, it is very clear that different parts of that fan shape are being activated depending on the angle of the joint and where the load is currently being carried. You do not necessarily activate the entire muscle at once; you activate some parts of it depending on the angle.

That is really why we need cheap, massively parallel soft actuators like the muscle fibers, in order to be able to have that range of subtle control over angles and forces.

What you are fundamentally saying is that the extraordinary mechanical features of the device can be very complicated and nonlinear, yet that is a property we can exploit to make the robots work better, rather than treating it as something to remove because we cannot control it.

Yeah, right. Absolutely.

I think it is a very important aspect.

I think there are good reasons to just, you know, bite the bullet and accept some of these areas of the complexity in order to enable physical capabilities and physical properties that are beyond what any of our current robots can do.

Well, I think that is a fantastic place for us to draw things to a close. We have touched on a lot of issues and in quite a lot of depth. I really thank you for your time.

Okay, thank you.

Address correspondence to:
Vytas SunSpiral, MS
Dynamic Tensegrity Robotics Lab
Intelligent Robotics Group
Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 94035

E-mail: vytas.sunspiral@nasa.gov