This Is Not Your Ordinary Ivy League

How Do We Design and Make Robots That Can Mimic Every Aspect of Nature?

HOD LIPSON, MECHANICAL AND AEROSPACE ENGINEERING/COMPUTING AND INFORMATION SCIENCE

Design and Make

I look for ways to use evolution to design things. Robotics is the visible part of what I do, but underlying the robotics are two basic questions: Can we design machines that can design other machines? Can we make a machine that can make other machines? These are fundamental engineering questions. Engineering is about designing and making things. The ultimate thing we can design and make is the thing that can design and make other things. If we can do this, then we will have understood the underlying engineering processes well enough to automate them. It is the ultimate engineering challenge.

As time moves forward, the kinds of things that the engineering community needs to design and make become increasingly complex. The only way to sustain this growth and complexity of products, the next generation of iPods for example, is by automating some of the processes. We must elevate the generic process from designing and making the products themselves to designing the processes that can in turn make new products.

The Biology Connection

Nature is my inspiration. Nature is always designing and making things through the evolutionary process—through development, growth, and learning. It is amazing. Its accomplishments dwarf those produced by the best teams of human engineers. If it is done in nature, however, I know it is possible, and there must be a way to automate some of it. Engineers can only do a small portion of this now, so we have a long way to go.

While these questions permeate many engineering fields, we focus on robotics, as these represent a notoriously difficult, visible, and potentially high-impact challenge.

Growing Robots

You have seen plenty of robots in factories and in Hollywood movies. They are made out of rigid components, and while you might be able to reprogram them to do different things, the actual physical body is fixed. It cannot be changed. If something breaks, it cannot be fixed—game over. Most of the products today are like that: their morphology cannot be changed. This is not the case in nature. Over a lifetime, bodies grow and change. If something breaks, it heals. If one muscle is used more, it expands. Over a developmental timescale, a baby develops and matures. Over evolutionary timescales, the body plans of a species change to match new environments and adapt to new tasks.

How can we make robots that can change their physical structure? Grow? Adapt? We are trying to make machines that literally do that. We look at mechanisms that will allow materials to develop in an adaptive way—to heal, respond, and adapt to the situation. I am very excited about this project.

The Body of a Robot

We take two basic approaches to the problem: a materials approach and an approach based on small modules, or small artificial cells that join together to compose a structure. First, how do we find materials that are adaptive to their environment? We use 3-D printing to work with combinations of materials that together provide functionality. We build a robot like an onion. We gradually layer materials that provide structure, conduct electricity, and do sensing and actuation—adding more and more layers with different materials and in different locations to create a working device. We spray on these materials, and the object grows into a robot.

We also build robots and machines made of very small interchangeable components, like cells in the human body. Humans live for a long time, but the cells are constantly replaced. New cells develop, and old cells die. Although our cells are constantly replaced, an individual—a collection of cells—remains the same person with every cell in the body having changed multiple times and none of the original cells left. We want to make a machine that does the same.

We have various projects exploring how to make large-scale machines like robots from particles on the microscopic scale. We are trying to make components that are the size of a grain of sand. With a million or billion of these, we can make a machine. We can remove some of them or add more of them to get variability—the adaptation in shape, just as in nature. Even machines that can metabolize other machines and reuse their components—the ultimate form of recycling—are possible.

Mimicking Nature’s Process

How do we know which materials will mimic nature’s process? We start with a concept, and we try out different ideas. We fall off some dead ends, and these do not get published. We go back to the concept, refine it, and try again. We are close to making a complete robot from raw materials using the layering process. We have been able to make robots out of large modules, using small modules, and now we can make them smaller and smaller and make more of them. Eventually, we want to make a robot out of a million small cells that can adapt and change.

> We build a robot like an onion. We gradually layer materials that provide structure, conduct electricity, and do sensing and actuation—adding more and more layers with different materials and in different locations to create a working device.

The Mind of a Robot

The first project I described was about the robot’s body—how the body changes. This second set of problems is about the mind of a robot—how the mind changes. How can we make processes that analyze and discover—processes that allow machines to change their behavior or adapt their behavior over time to new situations? This question has many different approaches. It is basically an artificial intelligence (AI) question, but it is interesting to try to address it in the context of a physical body, not just an abstract algorithm playing chess. Most machines of automatic systems currently adapt to things that are foreseen by their designers. If a robot’s designer-engineer foresees situations that the robot might need to tackle, then these are programmed into the machine. Can we make machines that can change their behavior in response to drastic new situations? Can they adapt to new, unforeseen situations?

We are looking at ways to make machines adapt like nature—like animals that adapt their behavior to drastically new situations and opportunities, such as collaboration and competition. Envision this: A robot loses a leg, and it begins to limp, or it discovers a new opportunity that it is curious about. Can we make machines like this? These are common themes in science fiction. (You can probably name a famous robot with these attributes, right?)[1]

I believe that the key challenge is self-reflection—getting a robot to create a self-model, an internal image of itself, so it can anticipate what is going to happen to it. When a robot has a representation of itself, it can use that representation to anticipate and react in advance, like humans who anticipate a consequence of actions, from moving a hand to more complex long-term actions. We are able to make robots that can do this rudimentarily. We have a robot that can move and internally form a self-image—an image of what it looks like and of what it can and cannot do—and use these attributes to make decisions.

The next step is to see if a robot can create an image of other robots. Humans collaborate and compete with each other because of the ability to understand other people’s intentions, goals, and actions. Robots are not capable of this yet. The ability to model oneself and others is fundamental in cognitive science, and I want robots to be able to do the same thing, which is key to embodying intelligence. This is our goal.

How Far Along Are We?

We have a robot that can learn how to walk by randomly wiggling its feet and creating a self-image. For example, is it a spider or a snake? In the beginning, it does not know. But it creates an internal image and then figures out how to walk. Then we remove a leg from the robot, but we do not tell it. The robot figures out its self-image, adapts, loses a leg internally, and begins to limp. We did not preprogram it. It is all self-image, and it is spontaneous because it creates a model of itself. Right now, this is only for locomotion, which is simple.

We are working on robots that can look at other robots’ behavior and by just observing

Lipson in the lab with (l .to r.) Michael D. Schmidt, Jonathan D. Hiller, Evan Malone, Floris van Breugel, Zhi Ern Teoh, and Viktor Zykov

them, anticipate what they are going to do and how to work with them. Our earlier self-assembling robots, the ones that could change their shapes, were made of only four large cells. Now, we are working towards hundreds of microscale cells. Also with self-reflection, we are now seeing robots that model their own thinking, a sort of mental self-reflection.

> We have a robot that can learn how to walk by randomly wiggling its feet and creating a self-image. For example, is it a spider or a snake? In the beginning, it does not know.

Biology, a Treasure Box for Engineers

[A copy of Thomas Eisner's book, *For Love of Insects*, lies on a table in the lab.] *For Love of Insects*, by Thomas Eisner, Neurobiology and Behavior, is inspirational. Biology is a treasure box for engineers. Just take any animal, and see the fantastic things nature has figured out. We are so far away from doing anything near it. Every page of the book is a potential research project. For example, we have been working on insects' flapping and hovering in the same place. We ask, "Can we make something like that?" The answer resounds, "We cannot."

We reply, "Why?" Immediately our questioning leads to interesting speculations. What is it about how insects work that we cannot figure out? If we could, what kind of insight would we gain? One example is the undergraduate research project of Floris Van Breugel (now a PhD student at California Institute of Technology). Van Breugel made a flapping, hovering machine like a mosquito that weighed 28 grams.

There are many fantastic phenomena—insects, animals, cells, and primates—across all scales that can give us ideas about how to do engineering differently. I think evolution is the mother of designers. We use very crude methods to make things in engineering—cut and assemble parts made out of single materials—whereas nature grows complicated structures with multiple materials simultaneously. We analyze that, and bring new ideas to engineering from nature.

Machine Creativity

I have been interested in machine creativity for many years. A lot of work has been done to make computers smart. They can play chess or drive a vehicle across the dessert. These are hallmarks of intelligence, but not the ultimate artificial intelligence we are looking for. What uniquely defines human intelligence? I think the ultimate challenge is creativity and curiosity. Trying to understand what it means to be creative or curious in a way that we can imitate has been a long fascination. Can computers augment creativity or curiosity? Can computers ask intelligent questions? Generate new ideas? This is the epitome of AI. How can we make computers with these characteristics?

Undergraduate Researchers, Awesome in the Lab

Cornell's unusually high quality of students, particularly undergraduates, has been remarkably rewarding. If I compare Cornell to other top schools, certainly we have excellent graduate students, but what is unique here is that we also have undergraduate students who are encouraged to do high-quality research. It is part of the curriculum. I have students who come into the lab during their freshman year and do research all four years, almost like a PhD student. Undergraduate researchers are a tremendous help, especially for a new faculty, because recruiting students and finding funds for them is one of the hardest things to do, and we spend so much of our time on it, along with building and funding a research program. Having an abundance of high-caliber undergraduate students capable of doing research and with the time and motivation to spend four years on research is exceptional. Undergraduate researchers have done some of my best work. Their work has even led to a paper in *Nature*. These days in order to get research grants, we have to do a significant portion of the work up front—before we get the funds.

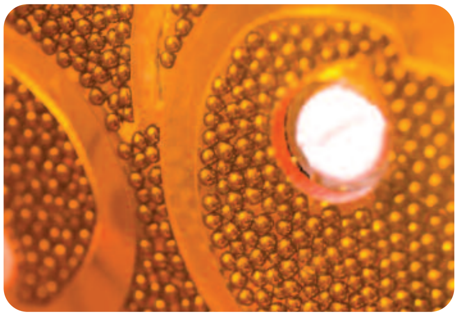

Fascinating!

- Nature is always designing and making things through the evolutionary process—through development, growth, and learning. My research lab designs and makes robots to test ideas from nature.
- How can we make robots that can change their physical structure? Grow? Adapt? We are trying to make machines that literally do that.
- We are trying to make components that are the size of a grain of sand. With a million or billion of these, we can make a machine. Machines could even metabolize each other.
- This second set of problems is about the mind of a robot—how the mind changes. How can we make processes that analyze and discover—processes that allow machines to change their behavior or adapt their behavior over time to new situations?
- Is it a spider or a snake? In the beginning, the robot does not know. But it gradually creates an internal image and then figures out how to walk.

University Photography

Having undergraduates to help produce results before starting the ultimate research is an extraordinary resource for faculty, and a tremendously beneficial learning experience for the students. Many of my undergraduates have gone on to do PhDs, accepted into excellent schools because of their research and published papers.

and security—flexibility, creativity, and at the same time job security. This is usually a difficult balance to achieve. People who are creative, artists and authors for example, often do not have much job security. They take risks in order to do the work they love. A faculty position, if you are tenured, offers both.

> The suspension bridge over Fall Creek Gorge near the Johnson Museum of Art is an enchantment. I like to stop there and take a moment. It is so beautiful.

A Better Way

Before I did a PhD, I spent five years in industry—in the Israeli navy and then in the software industry. I saw engineers design and make things in the early 1990s. Although computers were ubiquitous in the workplace, they were very passive, like word processors for engineers. The computer-aided design tools could record engineers' thoughts and ideas, but they could not generate ideas or ask questions. I looked at these machines and thought there must be a better way. If we could crack the question of how to make machines creative and enable them to ask questions, we would gain a lot of leverage. As designers in society, we could focus on higher-level questions and relegate most mundane day-to-day things to computers and do more as engineers. This was the seed of my interest.

We have recently applied this concept to scientists as well. We created a robotic scientist of sorts that can look at experimental data and deduce the underlying laws governing the system. It has been discovering things we cannot yet explain. But this form of AI will certainly be essential in our ability to make progress in science in the future, as the problems we tackle become increasingly complex.

Why Academia?

My father is in academia. I remember, as a teenager, thinking that I would be anything except a professor. And here I am. I chose this way of life after having done other things. It is the best combination of freedom

Taking the Research to the Public

Connecting with the General Public. Interaction with the public has taken me by surprise. Much of our research in the lab is done internally, but we put one of our projects, 3-D Printing of Multiple Materials, on the web as an open source project. We call it the Fab@Home project. We placed the blueprints for one of our basic pieces of research equipment online, showed how to make it, gave the list of parts needed and where to purchase them, and said, "Assemble it, and in a weekend you can have your own machine, print these things, and do research at home." It exploded. Ten million downloads! For an academic site, this is a lot. Many people have built these machines at home, and it has been fascinating to watch. People are exploring, and they send us e-mails and pictures of their machines. This is a huge amplification of the research. I think there's a huge potential for at-home research. Just as universities are making teaching broadly accessible, research will go that way too.

We were surprised by so much interest in our 3-D printers for printing—something I never would have guessed. It is not robots. What do you think people like to make with these printers?[2] When we put this technology in the hands of the public, we get new ideas. Although not every project is appropriate, I look forward to involving the public more in the future with our research projects. To allow the public to participate in these ideas touches millions of people, in addition to the students in the lab. By opening our research to the

About Lipson

Years as Cornell faculty
7

Came to Cornell from
MIT/Brandeis, Boston

Favorite spot on campus
Suspension bridge over Fall Creek Gorge

Cornell's research distinction
Freedom to pursue unconventional research

Cornell's trademark
The union of Ivy League and state college

I am also
An improvisational jazz pianist

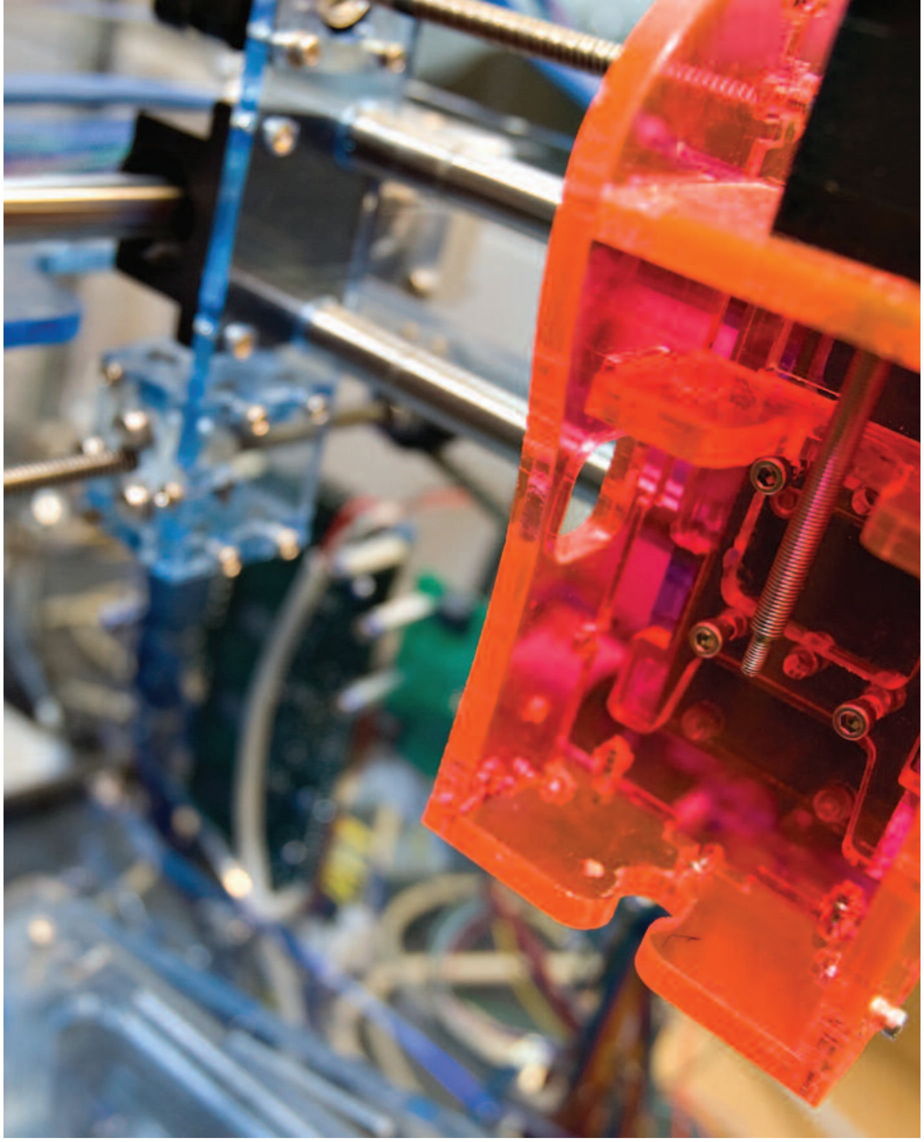

public, we create incentives for the new generation of engineers.

Kids, Teachers, Parents, and Video Games. Marvin Minsky (MIT) said if we compare schools 100 years ago and today, we find they have not changed much. We still have the teacher standing in front of the students at the blackboard dictating information and students sitting at desks. If we follow the students home, we see them spend three hours gaming, using computers, and expending a lot of energy thinking while playing games. But we are not tapping this use of time and energy for anything useful. One challenge is to figure out how to harness this immense attraction and investment that the kids are willing to put into games for teaching and education.

As we try different approaches to get kids involved in science, we are beginning to develop educational games that are really fun to play. The current software is mostly "feel good" software for parents. Imagine software that kids can't wait to play, and they do not notice that they are learning. To design this kind of software requires a lot of thought. It is hugely untapped and has lots of potential as a resource for teaching. I look forward to doing some of this work in the future.

Favorite Spot on Campus

The suspension bridge over Fall Creek Gorge near the Johnson Museum of Art is an enchantment. I like to stop there and take a moment. It is so beautiful. People drive a long way to get to a spot a like this. I cross it every day. The campus is highlighted with many spots like this.

Living in Ithaca

Choosing Countryside over City. I enjoy walking and love experiencing the four seasons. Coming from Israel, I am still excited about snow. Although I enjoy the city and the countryside, if I had to choose one, it would be the countryside. I enjoy living in Ithaca. It has character.

A Deeper Experience. If I had to wish for something, then 40 minutes to one of the big cities would be my ideal. There is a saying, "It's a hole in the woods, surrounded by reality," which captures the essence of the area. Ithaca is isolated and insulated, but it is isolated enough so that we can rise above the immediate concerns of the day and look at more significant things. On the one hand, this insulation is good. On the other hand, because it is in upstate New York, it is connected to reality in ways metropolitan areas are not. It is authentic and grounded in a deeper experience.

If I'm Not in the Lab ... I am with the kids. I have two children, ages 12 and 5 years, and sometimes the kids are with me here in the lab. The Legos under the table are for all of us—the kids, the people in the lab, and me. A wonderful advantage is that my children can understand what I do in the lab. I build robots. I can easily explain my work to them. This is my hobby—the kids and my research lab. I also play jazz piano and love to improvise.

The Last Word

On Cornell

I have worked at many places, but here I really have fun coming to work. Now, what is it? Is it the atmosphere of the department? Is it the students? Cornell has become a personal entity. It is more than a job or a place. I feel I belong to it. Cornell has the ambience of a state school and an Ivy League school, of a down-to-earth sensibility and an ivory tower. It is an indescribable aura, and it is fun to be here.

For more information:
E-mail: hl274@cornell.edu
Website: www.mae.cornell.edu/lipson/
Video online: http://ccsl.mae.cornell.edu/research/golem/golem480x240.wmv

---

[1] R2D2 in Star Wars

[2] Food. Yes, they like to print food. They will put chocolate frosting, peanut butter, cheese, you name it, into these printers and make food constructs. Some companies are interested in printing wedding cakes.