by Eagle Gamma

ALMOST NATURAL

ROBOTS AND HUMANS LIVE UNDER ONE ROOF AT THE CARIS LABORATORY

The PR2 Humanoid robot designed by Willow Garage, which does technical research to help humanity.

A robot arm gradually adjusts according to a cybernetic algorithm, and its hand smoothly transfers an object to a human hand. The transfer serves as an analogy for robots more generally, as they adapt their position in human society through gradual feedback. Right now, we stand on the brink of a transition in how robots and humans relate.

Already many automata serve in military adventures, not to mention the robots doing duty at international borders, in factories and nursing homes,

Two robot arms are shown that are used in research on precision movement.

and – of course – on lawns. With increasing use, robots more frequently interact with individual humans, and with society. In the transition, we will have to figure out how to get along, humans and robots side-by-side. It all comes down to that handover.

The artificial arm participating in the handover belongs to Charlie, a partially humanoid robot at the Collaborative Advanced Robotics and Intelligent Systems (CARIS) Laboratory, at the University of British Columbia (Canada). The spacious lab hosts robots of hugely varying dimensions, which intermingle with researchers and students. Developments here delve into individual behaviors, such as balance and manipulation and vision, to enable safe human-robot interactions. Neatly stacked bookshelves serve as dividers, and every direction and height sports a scattering of plants and bicycles and cables.

Here's looking at you! Detail of the PR2 shows nearly 10 imaging sensors.

HOW DO YOU DO WHAT YOU DO?

Elizabeth Croft, who directs the CARIS lab, and Dr. Mike Van der Loos, the associate director, show me the lab and describe current research and development. Croft believes that robotics must incorporate both robots and people. "The whole lab is about how people and robots work together, how we can use the technology from the robotics area and our understanding of human physiology and human psychology to bring the two sides together."

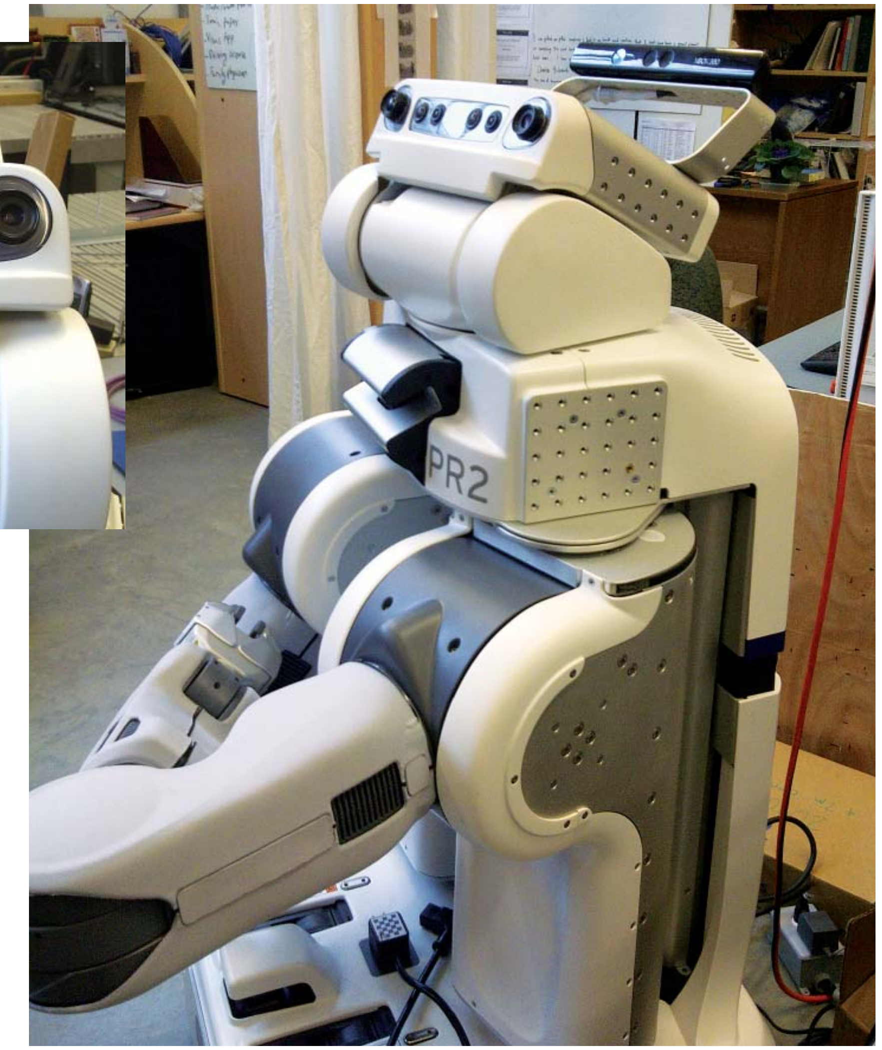

PR2 sideview.

How can predictable robots work effectively in an unpredictable world? Many of the lab's methods come down to observing robots and humans together under carefully monitored conditions, isolating the essence of a specific behavior, then developing usable implementations of those models.

"I think the interesting thing that we focus on, we focus on behavior. The things that we're doing are applicable not just to one handover one time, or one standing-up occasion one time, or one hesitation. It's an embedded behavior that can be used in all sorts of different circumstances. We're looking at the lower-level behavior that we need to give robots so that we will be able to interact with them in an efficient and human-like way."

The main goal of the lab is to get robots to behave flexibly enough to adapt to many human situations. When I ask whether that extends to being flexible enough to decide when to go shopping, or to moral questions, the roboticists quickly add that it's in progress. Indeed, some of the research already touches on these issues.

In addition to Charlie, the lab hosts a gigantic robot capable of simulating human balance, which I get to ride, and lightweight robot arms that help to investigate cooperation between humans and robots. The robot arms conveniently exercise fine motor control, says Croft.

"They're really nice experimental robots that you can use for different human-robot interaction experiments, because you can pick them up, they're not that heavy, you can backdrive them, and we use them both for industrial applications as well as rehabilitation applications."

A few boxes of marbles sit between us and a robot arm. Some compartments have mixes of marbles, while other compartments are neatly sorted already. Sorting mimics to some extent an industrial environment, for example, a robot working in an automobile manufacturing plant. In a plant, some tasks involve a robot picking parts out of a bin. If a robot could sort through and pick up the right parts, then it would no longer require separate boxes for separate parts, a costly setup. Can robots learn to sort well, in tandem with people? To do so, they must take turns.

ENTRAINMENT

When people do an activity, such as walking or talking, we naturally fall into sync with each other, in a process called entrainment. Croft illustrates the point with a favorite example. "I was doing my Ph.D. and holding my baby, and I was doing this," she says, gesturing as if to rock her baby back and forth. "I was talking to a professor and they started swaying. Another professor came up and they saw me, and they were talking to me and they started swaying. I had ten professors swaying in the room!"

Croft now wants to know "whether robots can use entrainment to control the cadence of people." Potential applications use the robot to adjust a person's speed to unfolding conditions, such as relatively busy or free periods. "We're just trying to explore this aspect."

A robot able to guide a human by gesturing could supervise or teach a worker. The machine would monitor the person for any errors, and provide corrective tips. Rather than flashing lights or a computer warning, the worker would see a human-like gesture. The gesture could take the form of a pointing finger, an outstretched hand, or simply an articulated arm. Croft points out a hefty benefit of a robot arm even in the absence of a hand: "Do people understand it if you don't have [the hand]? Because if I can take the hand off, I've just saved myself fifty-thousand dollars."

To study how humans and robots exchange information, the lab records a very brief interaction, at around the speed of human vision, the rate that tunes many of our behaviors. "Basically, the way we measure [the interaction] is we instrument people, and we instrument the robots, and we see how fast they react." Observations help to discover what makes a robot gesture useful, especially in situations like automotive assembly.

The author rides the Riser at Caris Lab. The device can be used to tilt subjects while varying the angle of the platform surface on which the subject stands, for use in medical research and rehab.

"The robot can be doing other stuff but keeping an eye on what the worker is doing, and saying, 'hey, you forgot the handle, get it on there, that one there right now.' The thing is, it has to be done quickly, because the turnaround on each operation is on the order of twenty to sixty seconds, so we really need this on-the-job inspection. People habituate to flashing lights, and sirens, and beeps, and stuff like that. But – people do not habituate to motion, because we are trained to be able to see the tiger out of the corner of our eye, and run away from it."

I mention that our reactions even depend on neural hardware, from the evolution of our bodies and minds in complex three-dimensional environments. Croft agrees: "We're wired for motion. Robots provide 3D motion that we're wired to see. We're really, really triggered on gesture, and that's why it's so important."

BALANCING ACT

A large robot right at the lab entrance serves to conduct studies on human balance. The Robot for Interactive Sensory Engagement and Rehabilitation (RISER) has a motion platform to stand on and uses force/torque sensors to measure the interaction between a person's feet and the platform. RISER then adjusts its position appropriately using highly tunable control software. The results inform not only how we ourselves stand – on Earth, on the Moon, under any conditions – but also how robots stand. Lessons from human balance and movement data can guide robotic implementations, to make the machines more relatable.

With RISER, Croft notes that "you could experience balancing on a surfboard, you could experience balancing on the moon, you could experience different kinds of dynamic challenges that we wouldn't necessarily be able to do if you were just standing here." Punch in new settings on the control terminal, and the robot behaves completely differently. Perform the right experiments, and discover how to keep a body vertical.

"We're working with people in human physiology, and we're trying to understand how the human balance system works." A better understanding of balance helps to develop robots. It also helps to rehabilitate people with damage to the vestibular system, our inner ear's source of spatial sense. "We're all inverted pendulums. We stand up and we don't fall over, that's really awesome!"

Humans stand up through an ongoing effort to counteract the physical forces that would otherwise knock us over. We sway back and forth, actively adjusting our balance. (When we sit, our body sends the signal to stop monitoring sway.) The CARIS lab aims to find out how exactly we detect and adapt to the frequent changes. "We're trying to understand what the sensing is, what the physiology is in terms of the actuation that goes on when we stand."

To study balance, subjects stand on RISER, and balance on top of the robot. However, because the robot runs a subtle control system, constantly monitoring conditions, the subject actually balances an artificial version of the physical setup, not simply the subject's own body. The robot tilts in response to how the person stands, whether pushing down on toes or heels, and the human rides the motions.

In one experiment, the researchers want to test the validity of their models of human balance. One way to verify a model is to disable a human's sense of balance, then let the robot take control. By electrically stimulating the vestibular system, or anaesthetizing the leg nerves and then stimulating the leg muscles, researchers can knock out portions of a person's balance system, and replace the functions with controls from the computer. Then, the robot balances, along with the human body, "so it feels like you're standing on the ground, but instead of you rocking back and forth, the whole robot rocks back and forth," Croft says.

"So you're standing up in here, your leg is anaesthetized, and instead of your brain controlling that big muscle there, we put functional electrical stimulation into the leg to control the contractions of the muscle. The computer now is saying, 'Oh! I see that you are falling forward. I need to provide a torque so that you will stand up again.' We had the computer balancing the person."

The experiment allows for comparisons among different mathematical models of balance. A control model inspired by humans makes a series of discrete shifts, whereas a robot servo motor can make smoothly continuous adjustments. A control model can operate in real-time, or it can incorporate predictive information about its trajectory. Each approach offers tradeoffs, in terms of performance, energy efficiency, and style.

Humans stand by falling forward a little bit, then pushing back with leg and foot muscles, and repeating. Lacking the ability to push ourselves forward in the same way, if we lean too far backward we simply fall all the way. The constant exertion of energy and calibration makes activities that require balance – like standing, walking, and riding a bicycle – difficult for children to learn.

We perceive our balance through the inner ear canal, which should ideally match our actual position. The CARIS lab currently tries to quantify how the information travels between perception and action. How we sense our movements through space may in fact differ from what our bodies do. Shocking the balance system subjects it to experiment.

"To feel different kinds of sway – we can zap behind their ears, and make people feel like they're twisting this way, or turning that way – that's electro-stimulation of the vestibular system. We can take that, then we can physically move you and see if we can cancel them out! If they cancel out, then basically we determine how the stimulation of the nerves is actually mapped to what you're physically experiencing. Isn't that funky? Every day I wake up and go, 'oh my goodness!'"

TO RIDE A ROBOT

I get the opportunity to experience RISER, to ride a robot. First I have to sign a consent form, which notifies me of several safety considerations, including two emergency stop buttons: one for me, and one for the experimenter. Climbing up a small ladder, I step onto the large robot. Above its dynamically tilting base, a platform supports torque plates which measure foot forces from the rider. I step onto the plate markings, and the experimenter straps me in.

We start by setting the control parameters – the highly configurable directions for the robot's behavior – to react realistically for my body dimensions. The experimenter asks me to balance normally, while the robot counterbalances. I lean back slightly on my heels, and the robot leans forward by the same amount. I lean forward slightly on my toes, and the robot leans back by the same amount. At first it confuses my senses a little bit, and I make mistakes like leaning in the wrong direction or overcompensating. After a few seconds, though, my body adapts and it feels almost natural to stand up straight.

Next, the experimenter gets the robot to rotate the pressure plates instead of remaining parallel with the floor, so that my feet now tilt with the rest of my body and the robot. When I lean forward, I now rotate my whole body, instead of bending at the ankles. Lacking balance feedback from my lower legs, I have a slightly harder time standing up straight, although I still manage through my redundant systems: vestibular and visual.

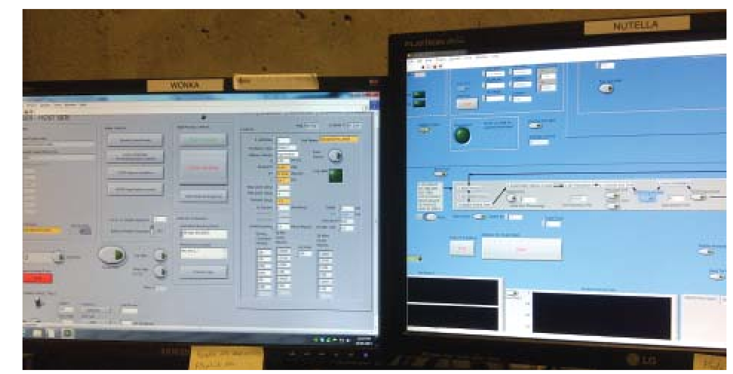

Wonka and Nutella control systems are displayed.

The experimenter then asks me to close my eyes. With two out of my three balance systems knocked out, I have to rely on my inner ear canals alone. All of a sudden the challenge grows real, like staying up on a unicycle. (I surprise the experimenter with my balance, presumably due to my bicycle riding.) We play around with the robot's settings. It amazes me to experience my focus shift as we switch the setup around. I palpably feel my brain reroute the responsibility for balance from my ankles to my eyes to my inner ears. Like a roving camera, my attention pans and zooms and fades to follow the action.

The robot, RISER, is firm yet supple. With one profile, the robot moves ploddingly against your pressures. A simple switch of the settings, and RISER follows you lightly and briskly. When we first interact, the robot responds gingerly to touch. A delicate finger to the force plate suffices to place the platform where one pleases, like a dentist's lamp. Yet, push any part of the robot without sensors, using all your might, and the behemoth won't budge.

CHARLIE DON'T SURF... YET

A little further into the lab, Croft points to Charlie. Next to RISER, an imposing, industrial-looking figure, Charlie looks downright residential: likeable, touchable, humanoid. Charlie has grayscale skin, a geometrically simple body plan, and plenty of signs of roboticity, yet he has a relatively endearing demeanor. Merely standing in front of Charlie, an early-model Willow Garage PR2, I struggle against the temptation to reach out and feel him – and the researchers definitely use a human pronoun.

The CARIS lab uses Charlie to test how people interact with robots in everyday scenarios. Do we treat robots as humans or as machines? How can we coordinate actions? How should robots behave?

In experiments, Charlie takes turns with human subjects on shared activities. One study proposes a robotic implementation of human-like hesitation gestures, to communicate uncertainty in order to prevent collisions. Another study teaches Charlie how to hand over objects. Up to now, handovers between robots and humans feel forced, and often fail – or as Croft says, "what you'll observe is a tug-of-war."

Again, the lab covers people and robots with measuring instruments, and carefully observes behavior. When humans hand over objects, the giver and receiver tacitly coordinate the action. Multiple repetitions let the lab examine the grip force of the handlers and the load force of the object. Learning how we perform a handover provides insight to teach robots the same maneuver. Let go too soon and the object drops. Let go too late and there is no handover.

"Charlie's been programmed to be a good giver," Croft says. It turns out that a good giver takes responsibility for the safety of a transfer, while a good receiver takes responsibility for the timing. In test handovers between a robot and a human, people expressed their level of confidence in the robot. When the giver holds on too tightly, it feels like the robot does not want to let go. At the other extreme, it seems like a risky transaction, or as Croft says, "I'm not going to hand the robot my Fabergé egg."

Humans naturally have a rapid proficiency at interactive tasks. Van der Loos points out that "we are very adept, and we can feel force changes very quickly. We also have these great shear sensors on our fingertips, so we are very good at negotiating all that stuff in half a second." The CARIS lab analyzes our behaviors, looking at the psychology and physiology. Croft concludes that by distilling the control of actions, people can develop cooperative robots.

VISUAL SERVOING

In an approach called "visual servoing," the lab aims to enable humans to communicate with robots by showing the machines what to do: monkey see, artificial monkey do. Traditionally, people control robots with "teach pendants," remote controls for robots. Instead, Croft says, "we're imagining that what you should do is take a robot, and show it what to do by grabbing the end." After a few lessons, the robot navigates even in the face of variations in the environment.

One emerging area of robotics where visual servoing can help, telepresence lets people interact with each other over arbitrary distances through robotic actors. Commercial robot suppliers already sell telepresence models, such as "Texai" from Willow Garage. A human in one location communicates via robot at the remote location. Likewise, the distant human could have a three-dimensional avatar at the near side.

A telepresence robot embodies all of the controlling person's actions, and perceives the scene. However, if an object comes between robotic camera and target, then it becomes a challenge for the robot to relocate the target. Moving targets further complicate tracking efforts. The CARIS lab strives to give robots enough sense to track the essence of visual cues – often using off-the-shelf hardware from video game devices.

The lab has also started a new project called FEATHERS, which has yet to accumulate a lot of hardware, to apply robotic control ideas to health care. For stroke victims, people with cerebral palsy, and others with hemiplegia (paralysis on one side of the body), a critical step to recovery involves intense physical rehabilitation. Van der Loos likens it to becoming an Olympic athlete.

The serious efforts – "thousands and thousands of repetitions over weeks or months, to get any kind of an effect" – allow people to retain their independence. Loss of independence means skyrocketing costs, personally and socially.

To assist people with rehabilitation, a small device runs robotics control software, and provides feedback for a video game system. Think of a joystick with a hint of artificial intelligence.

The lab hopes to produce an affordable device that responds to the individual movements of each user. The device would help to provide incentives to do the difficult physical regimens. Currently, the researchers prepare prototypes for use by dozens of patients in nearby medical clinics.

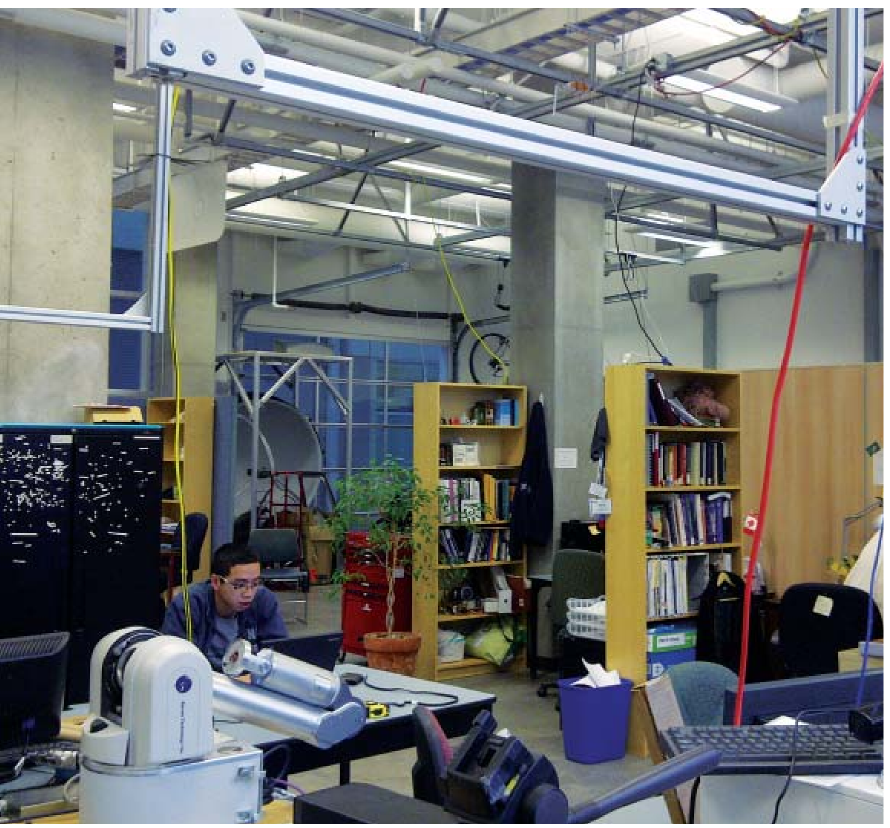

The Caris lab is partitioned for informal access to a variety of research desks and lab benches.

EXPLORING THE HUMAN-ROBOT FRONTIER

Experiments in the lab raise questions not just about what we can teach robots, but also about what robots can teach us. Ultimately, we learn through our interactions with robots about what makes us human. Robots increasingly play prominent roles in our society. From conception in ancient times through a rapid rise in the last century, and perhaps in the future toward mainstream acceptance, robots grow in importance. This growth already puts them in direct contact with humans. Now, at this juncture we can discover what people expect in robots, and how to realize those expectations.

Human morality stems largely from interactions. The more we bump into each other, the more we inspect our behaviors. As humans and robots increasingly live together, new moral questions arise. Is it acceptable to build robots that could kill people? Fatal accidents with industrial and military robots have already happened. What about military robots designed to kill? What determines a robot's rights and obligations? How should robots negotiate the human world, both physically and socially?

Fundamentally, many questions seek to know what a robot *is*. Is a robot a person, a pet, a machine, or in between? Modern research asks whether humans see robots as more closely resembling living beings or vending machines. Actually, we may treat robots as a whole new category. (Interestingly, a minority of test subjects don't see *humans* as living beings!)

Currently robots work in plants, but in clearly marked, separate territories. It is unsafe for a person to enter, even to do quick repairs, without shutting down the whole area. The CARIS lab hopes to integrate robot and human workers in the same environment.

Many of the lab's endeavors focus on transplanting human behaviors into robots. I ask about "native" robot behaviors: actions that ideally suit robots, rather than imitations of humanity. Croft cites certain industrial robots, which operate efficiently, if not safely in the presence of humans.

"If we talk about the paint shop and the welding shop, where exclusively robots – with big gates [enclosing them] – go zooming around, there's no people involved. It's all gated. Even when you go to repair, you have to shut the whole thing down to go in and do changes. Because of all the rules, and because those robots aren't safe. We're interested in safe robots."

Safe robots require a consideration of human expectations. "It's about the robots being predictable, and understandable, and polite, and following the rules of the road that we follow. That's why the human focus, understanding of how people do things, is a good metaphor for how we want the robots too."

CONCLUSION

For the most part, humans and robots live in different worlds. Humans don't go safely in robot zones, and robots don't go competently in human zones. The CARIS lab wants to change that. By programming robots to interact safely and cooperatively, we gain the best of both worlds. A safe and cooperative robot *gets* human behavior.

With enough effort on the robot side and on the human side, Croft believes that we will discover how to get along. "We've figured out how phones work, we've figured out how computers work, we've come up with new etiquettes of interacting with these devices – although some people aren't as good at those etiquettes as others! We're going to start to develop both ethics and etiquettes around interacting with robots."

In the lab I repeatedly hear the word "explore," and the CARIS researchers revel in exploring the future of human-robot relations. I get a palpable feeling of impending change, with an optimistic overtone. Already, robots probably hold far more sway than most people believe. Soon, perhaps, robots will leave the factory floor in droves, walking into society with a remarkably human-like gait.

Links

Collaborative Advanced Robotics and Intelligent Systems (CARIS) Laboratory, at the University of British Columbia (Canada), http://caris.mech.ubc.ca

www.popsci.com/technology/gallery/2010-07/gallery-rise-helpful-machines?image=6

For more information, please see our source guide on page __.