TTI VANGUARD
ADVANCED TECHNOLOGY RESEARCH MEMBERSHIP

ROBOTS, AI, AND THE FUTURE OF WORK HIGHLIGHTS

June 13–14, 2019 • Pittsburgh, Pennsylvania

prepared by Nancy Kleinrock nk33@cornell.edu

Why Drive Autonomously when You Can Fly Autonomously?—Dr. Sanjiv Singh, Carnegie Mellon University and Near Earth Autonomy

- Sitting in a traffic jam, who hasn't fantasized about one's vehicle elevating above the congestion and hightailing it airborne—Jetsons-like—to the intended destination?
  - But if you could, so too could others. What would that 3-D transportation environment look like?
  - Singh imagines urban aviation to encompass personal conveyance, likely from building top to building top (via "vertiports"), along with delivery drones, inspection drones, and conventional commercial flight.
  - This might seem like science fiction, but drive–fly concept vehicles are nothing new; the U.S. Army began experimenting with its hybrid AirGeep in the late 1950s.
  - Multiple firms are developing electrically powered, lightweight, quiet, and safe intracity rotor or fixed-wing aerial vehicles.
  - Singh lays out the safety and vehicular requirements that will enable a future of mixed-mode autonomous transportation.
- Anticipating the appearance of intracity aircraft in Los Angeles, Dallas, and Melbourne as soon as 2020, NASA and the Federal Aviation Administration are developing a framework to manage this new transportation mode.
  - Initially, personal intracity aircraft would be human-flown, but the transition to autonomously flown vehicles would be swift (Singh predicts 2028, with NASA/FAA planning for as soon as 2023).
    - The shortage of trained pilots demands that autonomy dominate this transportation domain.
  - The framework requires an autonomous aircraft to "fly safely, land safely, even when things go wrong, and do so without GPS," says Singh.
    - To fly and land safely will require forward- and down-looking onboard sensors to develop a real-time understanding of the terrain and to build a representation for on-the-fly decisions.
      - Vision analysis systems would have to distinguish terrain types and recognize the presence of obstacles.
    - GPS is too unreliable for it to underpin an autonomous aircraft's positioning system.
      - GPS outages routinely occur roughly once in every 1000 hours of flight time.
      - Moreover, "Jamming GPS is highly illegal but very easy to do," says Singh. "And GPS spoofing—making the aircraft think it is somewhere it isn't—is very easy to do."
    - Similarly, autonomous aircraft should not rely on ground control or other external communication for obstacle avoidance or control.
      - Communication to the aircraft is not inherently bad, but it is essential to limit the potential for hacking.
  - In 2010, Singh was involved in mock casualty simulations involving the first autonomous helicopter: It would fly independently to the location of an injured war fighter; there, a medic would load the injured individual onto the aircraft. It would then fly back to a safe zone and autonomously choose where to land.
    - Autonomous navigation, including flight, has benefited mightily over the past decade from the plummeting cost of sensors coupled with advancement in machine-learning-based vision.
  - In 2015, a Singh graduate student demonstrated the potential of non-GPS navigation with real-time point map generation by running through the Carnegie Mellon campus and up and over a parked vehicle, while wearing a helmet with a laser-based sensor on it.
  - With similar hardware, a quadcopter has been shown to independently reconstruct the details of an agrarian or suburban environment and safely fly, maneuver, and land among obstacles (barn, overhead wires, etc.) or mixed terrain (pond, home, woods, gazebo, and lawn), and safely land on the open grass.

        - As is the case with autonomous land-based vehicles, the vision system uses a combination of cameras and laser scanners to recognize that an object is moving (e.g., ambulating human, animal, car) and therefore requires real-time tracking.
- When considering risk tolerance, to date, there is appreciable societal resistance to cars that drive themselves without active human intervention or oversight, largely because people have grown accustomed to the significant risks associated with human-driven cars, but are—arguably irrationally—afraid of machines taking the wheel and behaving in more predictable but perhaps different ways than human drivers.
    - Automakers and auto insurers have long-term data on which to base risk.
        - When considering risks associated with new a technology, it is important to pose questions carefully.
            - In the self-driving car domain, the moral dilemma should not be whether to program the car to “drive into six nuns, or swerve and drive the passenger into the river,” but rather to instead drive more slowly from the outset and never face the horrific choice.
    - Greater societal assurance will be needed that self-flying vehicles will not fall out of the sky.
        - The real Achilles’ heel of flight—or, more aptly, its Icarus’ wing—is the ability to respond in the case of failure.
            - On the road, a car can swerve into a field if its brakes fail; in the air, it’s a long way down to a hard landing if a system component fails.
        - This is already true for commercial aviation: The FAA permits no more than one Class A mishap over 1M hours of flight, and airlines raise the bar for themselves by requiring no more than one such mishap in 1B hours of flight.
    - To raise comfort levels with autonomous aircraft, Singh anticipates rolling out applications in order of increasing consequence in the case of a mishap: first transport cargo in rural areas, then transport people in rural areas, then transport cargo in urban areas, and finally transport people in urban areas.
    - Still, the litigious nature of America could delay the rollout of cars with higher levels of autonomy as well as of autonomous aircraft.
- The firms that are betting big on autonomous cars are doubling down on the sky.
    - “There is an ecosystem of moving people and things from one place to another,” says Singh about Uber, which seeks Amazon-like dominance over “everything that moves.”
        - Flight sharing of electric autonomous aircraft should offer efficient transportation through high vehicle utilization, without the need for a horde of trained pilots.
        - “Aerial ride sharing is going to be a thing if this is going to be real,” says Singh.
- But the autonomous-driving problem is imperfectly solved: “We have solved the 3 am problem, but not the 3 pm problem,” says Singh, referring to the complexity inherent in driving to, say, Pittsburgh International Airport from downtown when the roadways are nearly devoid of obstacles or other vehicles versus when uncertainty is rife—uncertainty in the time to destination and uncertainty in the nature of objects (A person or a topiary? A bicycle or a person in a wheelchair? A child or a short adult? Members of each pair move differently.).
    - Cities with little-to-no non-car traffic (e.g., Mesa, Arizona, suggests Singh) should be able to progress to full automotive autonomy sooner than those with lively 3 pm-like activity.
    - Still, for full automation, cars must gain semantic understanding, given the appreciable probability of encountering signs that are easy for a human to decipher, but not for today’s AI (with the possible exception of Cyc).
        - Example problematic road sign: *bridge closed .0 miles ahead; local traffic only*, appended with a *detour* arrow.
            - Did a digit fall off before the decimal point? If not, will local traffic, if it proceeds, immediately end up in the drink?
    - Alternative environments, with clear boundaries and highly predictable behavior, have already been successfully automated.
        - Example: As Jeff Legault shared with TTI/Vanguard in July 2016, mines in Australia, where the cost of labor is high, the work is dangerous, vehicular activities are regimented, and the environment is free of obstacles and pedestrians.
- In some respects, the problem of autonomous flight is easier than that of autonomous driving (notably, prediction is easy in the flight domain, but hard when faced with the 3 pm problem), but motion planning is harder for flight.

- Just as self-driving cars struggle somewhat with inclement weather, so too do autonomous aircraft; this need for relatively arid environments informed the location of initial pilot programs.
- The pilot cities (LA, Dallas, Melbourne) are all investing heavily to design and build out aerial infrastructure at a range of scales, from ultralocal vertiports to neighborhood-scale vertiports, helipads, or airports, as well as regional and full-scale international airports to support the size and variety of aircraft and density of air traffic.

Key Challenges and Directions in Robotics—Dr. Martial Hebert, Carnegie Mellon University

- As the world's foremost academic robotics program, Carnegie Mellon University's Robotics Institute encompasses more than 800 faculty, staff, students, and research professionals spread across the university's main campus, the National Robotics Engineering Center, and the Field Robotics Center.
  - As its director, Hebert has a multitude of examples to choose from when sharing stories of the institute's contributions to the field—and, more generally, to society.
  - His choices span the gamut of how machines act, see, learn, autonomously decide, and interact within either specific or generalized environments, depending on the nature and goals of the robot under consideration.
  - Notably, across all facets of robotics, the key to effective design pertains less to the robot itself and more to how well the encoded model predicts the behaviors of the people with whom the robot interacts.
- Act:
  - The ultimate goal is to devise robots capable of acting on their environment in complex and useful ways.
    - CHIMP, CMU's third-place-winning entrant in DARPA's Robotics Challenge, successfully and rapidly completed all eight disaster response tasks, including turning a valve, cutting a hole in a wall, and climbing a ladder.
  - Still, significant progress remains in the areas of dexterity, agility, and modularity/reusability.
    - Dexterity:
      - "The sad truth is that our robots are still extremely limited in terms of the tasks that they can carry out, and part of the reason is because they lack the level of dexterity that we have," says Hebert."
      - While mechanical design figures prominently in the construction of a dexterous robotic hand, machine learning is even more important.
        - The key to progress is to leverage a large number of observations of human hands in action and extract from them guidance of how a robotic hand should move.
      - An example from "foodbotics": A robotic tomato slicer that grips and positions the delicate fruit before expertly sectioning it.
        - "This is a hard operation, because the system has to understand the squishiness of that tomato," says Hebert. "The physics behind it is very hard. If you try to do this explicitly, you can't program this. This is done by learning from observing people."
    - Agility:
      - Preconceived notions of a robot's physical form can overlook structures that are ultimately optimal for a particular task.
      - Example: A snake-inspired robot to climb poles or wriggle into tight spaces.
    - Modularity:
      - Just as the modularity inherent in software libraries has made reusability a hallmark of software development, the use of modular robotic parts eliminates the need to reinvent the wheel—or joint or other component—with each new robot.
        - "We are democratizing robotics and making it accessible to people who do not have the training to design a robotic system or the training to program it," says Hebert. "Remember, robotics is still a luxury activity."
      - Example: Modular components of Choset's Robot Operating System (ROS)-based snakebot, which connect automatically.
      - Example: Medrobotics has received FDA approval to use a scaled-down version of the snakebot in its Flex System, enabling minimally invasive heart surgery.
      - Example: Connect modules to form a fully operational robotic arm within half an hour.
- See:

- In collaboration with Pittsburgh-based self-driving car company Argo AI, the Robotics Institute is making "tremendous progress" in machine vision and scene understanding.
  - Still, machine-learning techniques suffer from the need for massive amounts of training data and system deficiencies when there is a mismatch between the characteristics of the current environment and that represented by the training data.
    - "This is simply unsustainable," says Hebert. "It is fine if you are Google or Facebook or Apple, and you have systems that you can train offline, but it is not okay if you are dealing in robotics or other systems that need to adapt rapidly."
    - As such, much of CMU's vision work centers around the use of minimal training data.
  - Part of this effort pertains to real-time construction of 3-D maps, including CMU's ongoing efforts in the DARPA Subterranean (SubT) Challenge in which researchers are tasked with mapping a mine in Kentucky with the ultimate intention of having robots navigate autonomously within the underground environment.
    - "There is no GPS or other way to do localization," says Hebert.
- Machines finding their way through novel environments is one thing, but to be able to map, anticipate, and evaluate human movements is another facet of machine vision altogether.
  - "To interact with people, you need to see them and understand them," says Hebert, which is why CMU has efforts underway to visually identify human expressions, internal state, intent, level of alertness, and so forth.
    - CMU's IntraFace project has automated the real-time tracking of facial features, estimation of head pose, recognition of facial attributes, and analysis of facial expressions.
      - These capabilities have been commercialized to assess driver alertness and to distinguish between engagement with the driving task and distraction.
    - OpenPose looks not specifically at faces, but at whole-body positioning and movement, performing real-time estimates of multiperson positioning and pose in two dimensions, with granularity extending to the motion of individual fingers.
    - "We can look at this as having a complete understanding of what people are doing, both their gross motion and their fine-grained activities," says Hebert. "These can be looked at as signals about people that can be used to do interaction, prediction, etc."
    - This work can be used in the development of robots that operate in close proximity with people, but also to track the bodily and facial responses of children to assess autism spectrum disorder, by picking up on a child's subtle movement/behaviors that could be indicative of an issue worth addressing.
      - "Imagine if you could do this from an app on your smartphone at home, without having to rely on highly trained professionals," marvels Hebert.
- Certainly, all vision systems rely on a camera or other sensor (e.g., laser range-finder) for input.
  - The consumer electronics sector has stimulated great advances in sensor technology over the past years, but exciting work is still coming out of university research labs.
    - For instance, as part of a DARPA program, work at CMU involves a camera that makes use of reflected light rays that permit visual sensing of what appears around the corner and otherwise out of range.
      - The computational load for image reconstruction is heavy, but with single-photon perception and capable algorithms, it is becoming possible to solve the "massive optimization problem" to determine the path of each photon, says Hebert.
      - Applications include greater situational awareness for warfighters, as well as the ability for self-driving cars to know what is around a curve before having an open field of view.
    - An operational application of higher order reflection of photons, using only off-the-shelf componentry, is to visually assess in the field the ripeness of fruit (e.g., grapes); another—still experimental—potential application is to peer within the body to noninvasively monitor real-time blood flow down to the level of the capillary.

- Learn:
  - In the same way as a child learns through repeated trial and error, CMU's robot-learning initiative entails perhaps 50K tries over 700 hours; arduous, yes, but with this training under its mechanical belt, the robot can successfully pick up a wide range of objects that it has never before encountered.
    - Similarly, a drone that persists through many hours of flying, crashing, and flying again learns to navigate within the test environment, never having been explicitly programmed to map input to navigational control.

    - Hebert emphasizes that real-world applications would provide the robot with a clue about what it will encounter and how to manage itself within the environment, but the ability to learn online gives it a leg up when placed in conditions only slightly different from what it encountered during extensive training sessions.
- Decide:
    - Decision-making by autonomous systems span a variety of dimensions:
        - Speed—For a drone to fly through the woods at high speed, it is necessary to imbue it with computationally lightweight onboard algorithmic techniques so that the system might keep up with the drone's translational motion.
        - Challenging environments—An environment can be challenging in various ways, such as the presence of multiple obstacles, uneven terrain, low light, tight clearance, fog, rain, or dust.
            - The robot must have a strategy to manage whatever challenges it encounters.
        - Teamwork—In some environments, the primary challenge is coordinate among a large number of robots; example: an Amazon fulfillment center, which has tens of thousands of robots.
            - "Decisions, now, are not just for individual robots, but to coordinate those robots to execute target tasks," says Hebert.
        - Complexity—The broader the deployment of robots, the more complex will be the environment they encounter.
            - For instance, self-driving cars will have to share the road with not only others of their ilk but also human drivers, pedestrians, bicycles, and other roadway users—each with its own quirks and ways of being (un)predictable.
            - "It is not enough for a self-driving car to drive safely—following the rules of the road and avoiding collisions—it has to drive in a way that is understandable by the other drivers and that is natural in some way," says Hebert. "It will have to have driving behavior that is consistent with expected behavior, including nonverbal communication. This is still a completely open area of research."
- Interact (with people):
    - For a robot to interact appropriately with humans, it must first be able to accurately predict how the people around it will behave.
        - People have mental models of how others will move, given a relevant context; machines need to acquire similar models.
            - Example: A system that predicts the path to be taken by an urban pedestrian can then interact safely/appropriately.
            - An "extreme example" of shared autonomy: A paralyzed individual with a direct mind-to-robot brain implant can learn to train a manipulator arm/hand to feed her.
                - "She can train herself to move the manipulator, but she cannot control it to the level of executing useful tasks," says Hebert.
                - To close the loop, the manipulator has a vision system that recognizes objects (e.g., table, food on the table, the woman's head and mouth), as well as an interim-prediction system that contextually learns the intentions associated the thoughts she produces.
                    - "When the system understands that the person's intent is to eat a particular food, it can now help the person by controlling the manipulator," says Hebert.

The Need for Situated Autonomy: Lessons from the 737 MAX and Extreme Environments—Dr. David Mindell, Massachusetts Institute of Technology and Humatics Corporation

- "Aviation has been a world of collaborative human–machine interaction, including quite a lot of automation, for a very long time," says Mindell.
    - Commercial flight has matured into a very safe system with few surprises, in no small part because of the system of controls that supports it.
        - Although he doesn't delve into any of the particulars, Mindell alludes to the fact that even a small change in software can lead to deadly outcomes, as was the case with the Boeing 737 MAX.
    - One question is how to avoid future disasters in the aviation realm; another is to get in front of similar conflicts between autonomous systems and the humans who are ultimately responsible for their operation.
    - Now, in the initial stages when autonomous cars are finding themselves on public roadways, it is time to build situational awareness into the built environment to minimize unforeseen glitches.

        - Toward this end, Humatics presents its intercommunicating microlocation sensors for ultrawideband radiofrequency ranging, which establish the 3-D location of each associated tracking device with centimeter-or-better precision.
            - In an urban environment, the location of vehicles with tracking devices can be known with pinpoint accuracy with no reliance on GPS; similarly for indoor positioning, provided sensors are distributed appropriately around the facility.
    - Mindell discusses the need for improved positioning and describes how the Humatics technology enables it.
- To set the scene for Humatics' work, Mindell first shares a story of the importance of human–computer interface design.
    - In the 1960s, NASA requested that the MIT Instrumentation Lab design the autonomous guidance system for the lunar lander; following delivery, NASA came back to MIT inquiring after the interface through which astronauts would interact with it.
        - The joke—reflecting that all the smarts were in the system—was a pair of buttons, one labeled *go to moon* and, the other, *take me home*.
            - "Today, we would call this nearly full autonomy," says Mindell.
        - Ultimately, the interface was a capsule-wide array of gauges, dials, meters, and lights that provided detailed information about the instantaneous status of the system and—importantly—to enable relevant humans to control the degree of autonomy they were giving over to the system at any given time.
            - People might trust the systems on which they rely, but also desire a window into their inner workings and the ability to override them.
        - In the case of Apollo, the astronauts in the lunar lander were, essentially, in an autonomous robot, while also in communication with a near-range option for intervention (via the orbiter) and far-range intervention (via Houston's ground control center).
            - "What you ended up with on Apollo was a 250K-mile-long, semiautonomous, teleoperated autonomy system composed of people, machines, and infrastructure," says Mindell. "That's 50 years ago."
        - Of course, in the interim, much progress has been made on machines, infrastructure, and the interfaces of people to them.
            - Still relevant is the pervasiveness of so-called autonomous systems that are embedded in very human systems.
            - Example: The Autonomous Aerial Cargo Utility System (AACUS)—an autonomous helicopter designed for the Marine Corps to deliver goods—pilot-free—to military outposts.
                - "The whole point of it is to bring stuff to people who you care about, so to make it useful it has to have this collaborative merging with this human environment," says Mindell, who adds that Marines on the ground paired with this largely autonomous aircraft via a to guide the system during landing according to the particulars of the real-time environment.
        - Autonomous systems not only interact with humans, but they also interact with one another.
            - "We live in a world in which our autonomous systems are much more collaborative and networked," says Mindell.
- "There are many, many cases of how autonomy is most successful when it is best integrated with the human environments that are around it," says Mindell—environments simultaneously composed of people, infrastructure, and robots. "I call this *situated autonomy*."
    - Humatics strives to establish situated autonomy from the get-go for land-based robotics as an important component of autonomous systems that are embedded in the built environment.
        - "I founded Humatics to bring new technologies to bear specifically for the purpose of bringing autonomous systems into precise relationships with what is around them," says Mindell. "We think of ourselves as a robotics company that doesn't make robots."
        - Instead, Humatics focuses on enabling better local decision-making when moving components are in close proximity, whether governed by humans or autonomous systems.
    - At the heart of situated autonomy is the need for robots to navigate within the built environment; to do so requires non-GPS-based microlocation.
        - "GPS is a very vulnerable infrastructure," says Mindell, "and, by the way, it doesn't work in cities, it doesn't work indoors, and it doesn't work underground. Guess where most robots work these days?"

    - Using time-of-flight ranging, Humatics' cellphone-sized, mesh-networkable RF beacons achieve 2-cm precision over 500 m; they use "tens-to-hundreds of milliwatts when they are running, but they run at fairly low duty cycles."
        - The firm's software integrates separate ranging measurements to establish 3-D positioning with up to millimeter precision.
    - In a New York City Department of Transportation trial, the same beaconing technology, affixed to light posts at 60-m intervals along one mile of Sixth Avenue, created a geopositioning corridor for any vehicle outfitted with a paired tracking sensor.
        - Although GPS considered his vehicle to be a full block away from his actual position, the Humatics network tracked Mindell's path with great accuracy.
            - "We even detected potholes just from the actual motion of the sensor on top of the vehicle."
    - In a more contained trial, this technology can enable precise navigation and tracking in interior environments, such as a busy e-commerce fulfillment warehouse, where carts, forklifts, and workers might all be tagged with sensors for mutual situational awareness and autonomy.
    - Other deployments:
        - Rail/transport: Canadian Freightways has reduced its fatal accident rate from as many as ten per year to zero by alerting conductors to the presence of maintenance activities down the tracks.
            - "We are now doing extremely exciting pilot [programs] in operating subway systems," says Mindell, with safety benefits for on-track workers wearing vests with credit card-sized trackers and to provide greatly improved positioning data of the trains in the tubes, which permits trains to run safely yet more frequently during rush hour.
                - Workers being tracked recognize the safety benefits of this technology and consider it more socially acceptable than camera-based tracking.
        - Ports/cranes: The addition of positioning creates opportunities to exploit autonomous features in such environments.
        - Landing zones for the types of intracity autonomous aircraft introduced by Sanjiv Singh ("vertiports"): The smaller the footprint of vertiports, the more useful autonomous aircraft will be in the urban environment; beaconizing them with Humatics' technology will create a necessary backup positioning system to GPS and accelerate this soon-to-appear transportation mode.
    - Humatics designs, makes, and sells its beacons and trackers, as well as the associated software to generate the ultrawideband-based navigation solution.
        - To deploy a microlocation system, the cost of infrastructure is modest; for example, a 10K-sq-ft facility with just ten beacons provided a 3-D grid with centimeter-scale precision.
- "If you tell any robot where it is, you have done half the work for it," says Mindell.
    - Looking to the future, Humatics anticipates commercializing its millimeterwave microlocation system, which provides millimeter-scale 3-D precision throughout a ballroom-sized space with a single beacon the size of a ceiling tile, with trackers no bigger than a fingertip; the initial application is for precise positioning of robotic manipulators or other facets of industrial automation.

AI: That's Entertainment—Mr. Jon Snoddy, Walt Disney Company

- Humans engage one another through stories, and what interests us is neither the mundane nor the perfect, but rather the oddity, the imperfection, the unlikely event or relationship.
    - "It is the flaws that power the story," says Snoddy.
    - Although doing so is within its technological and artistic capability, Disney does not strive for photorealism, but instead for creative exaggeration designed to engage the audience, whether it reaches that audience through a screen or at one of its theme parks.
        - "If entertainment is your business, perfection is not the goal," says Snoddy.
    - Snoddy presents a suite of robotic and interactive Disney characters created and instantiated by the talented artists and engineers of Walt Disney Imagineering.
        - The most advanced of these—the stuntronic acrobatic robot—began with the simple idea of controlling the motion of a foldable stick tossed into the air; amazing things can arise from humble beginnings combined with Newtonian mechanics, a bit of sensing, computation, and a story to tell.
- The robots and their stories:
    - Lucky the Dinosaur (2003)—Disney's first free-roaming audio-animatronics figure.
        - Eight-foot-tall and green, the segnosaurus named Lucky is guided by his handler, Chandler, and pulls a cart that houses not only his batteries and computers, but also his puppeteer.

    - As a hybrid character that locomotes unaided but under the direct control of the puppeteer, Lucky actively interacts with people who approach him, whether previously at a theme park or museum, or now on a tour of Imagineering.
- Muppet Mobile Lab (2007)—Perched atop a Segway platform, this rocket ship-shaped, free-roaming audio-animatronics figure encompasses a lab and two Muppets: Dr. Bunsen Honeydew and his assistant Beaker.
    - In contrast to Lucky the Dinosaur's onboard puppeteer and nearby handler, the Muppet Mobile Lab is remotely controlled by puppeteers who can be miles away, but gain situational awareness through the many cameras, microphones, and speakers distributed across the platform.
    - The people—adults and children, alike—who encounter the Muppet Mobile Lab in a theme park become captured by the magic of the illusion: They know the "scientists" aren't alive, but visitors engage with them enthusiastically and enjoy their antics as though they were.
- Destini the Fortune Teller (2010)—This stationary (as opposed to free-roaming) character engages in conversational banter with a selected audience member; it lacks speech recognition capability but can recognize when the person starts/stops talking.
    - The tag line of the fortune-telling exhibit—*Destini can sense your presence*—is a purposeful misnomer, since the audio-animatronics figure cannot understand its conversational partner.
    - Nevertheless, Destini's visitors gain the impression of a true give and take, and they remain engaged and caught up in the illusion of reality.
- Pascal (2017)—Imagineering has embodied the cute little chameleon from the *Tangled* franchise as an animatronics figure the fits in the palm of its handler's hand and is controlled through electronics in the handler's glove.
    - As in the films, the robotic Pascal does not engage in speech, but instead produces gecko-like vocalizations to engage with the people its handler encounters.
        - It takes both the robot and the handler to make the magic happen.
- Tiny Life: Vyloo (2018)—This glass-encapsulated trio of bird-like creatures, native to the planet Berhart in Marvel's *Guardians of the Galaxy Vol. 2*, are outgoing and responsive to theme park visitors, except when they behave shyly.
    - With a camera between their eyes and programmed to be attracted to faces, each Vyloo makes eye contact with visitors and interacts autonomously, aside from overall demeanor, which their operator dictates in real time.
    - "People love interacting with them," says Snoddy.
- Woodland animals at the Shanghai Disneyland Castle—As a component of the castle walkthrough, visitors pass by a garden with a variety of robotic woodland animals, which they can coax to come toward them.
    - The original conception was for guests to point to an animal and direct its path through the wooded area with continued pointing, but engagement with the animatronic characters was higher when the goal became for visitors to lure the animals in their direction.
- Conversational Dory—Not a physically instantiated robot, this character is an animated chatbot—powered by Doug Lenat's Cyc—that engages children on cruise ships by conversing with them about family.
    - Young children engage with it as though it were real, while older kids consider it a character in a videogame and attempt to suss out the rules.
- Jake Droid (2019)—This free-roaming example of Disney's "automatronics" is a *Star Wars*-like character, although not one specific to any movie in the franchise; it wanders unattended among guests through Disneyland's new Galaxy's Edge attraction and therefore must be unerringly safe.
    - Tween girls tend to drape an arm over Jake's shoulder while walking along it, while tween boys attempt to upend Jake with a shove; it must survive both experiences without rolling over the girls' feet or rewarding boys with attention for their shenanigans.
        - Jake ignores badly behaving youngsters, which curtails their activities.
        - "They have to have all the lidar and other sensors in your Tesla, but nobody hugs your Tesla," says Snoddy.
    - Early software used generalized facial recognition to promote engagement with visitors, but once theme park guests interrupted the view of their face with a camera, Jake would turn away, no longer recognizing the face-plus-camera as a human.
        - To overcome this deficit, Snoddy's team adjusted the facial recognition to include cameras.
- Stuntronics (not yet deployed)—Beginning as a project dubbed Stickman, after these robot acrobats launch from a trapeze they use laser range-finders, an inertial measurement unit, and onboard

computation to determine—on the fly—when to tuck and when to extend to perform amazing feats of aerial acrobatics, sticking the landing every time.
- "We are all going to be in this business of using robots for storytelling," says Snoddy.
  - Key to the successful engagement of characters with their audience over the long term is for each character to have a distinctive and persistent personality; *The Simpsons*—and each individual Simpson—has perfected this.
  - Snoddy rises above internal controversy at Disney about the use of technology to entertain by always putting storytelling first.
    - In every age, artists have used technology to advance their craft—whether pen and ink, motion pictures, computer-generated imagery, AI, or robotics—but the story always remains central.
  - The Disney ethos is entertainment, but other uses of robotics/AI are utilitarian.
    - Snoddy expects the best personal assistants will strike a balance between being endearing/engaging—e.g., a truly personalized Alexa or Google Assistant—and perfunctory.
      - Sometimes people will want a digital companion, but sometimes they will just want to know when the next train will arrive, so the AI will have to be able to discern human intent.
- Imagineering is a welcoming environment for experimentation; to pitch a project, imagineers need only "state what they believe" to receive funding.
  - Snoddy ensures that projects have time to try, fail, and eventually progress.
    - A company goal is to encourage internal sharing and collaboration, rather than accelerating the pace of work.
      - "Tools that allow you to make decisions more quickly gets you to bad decisions very quickly"; instead, good decisions often require time to incubate.

Robots That Both Walk and Drive—Dr. John Suh, Hyundai Motor Company

- Hyundai Motor Company develops and markets 17 distinct models of passenger vehicle that it sells in 193 countries.
  - Most of these are entirely conventional sets of wheels, but the firm's venture and innovation operation, CRADLE, has an expansive vision of what a vehicle might be.
  - The Hyundai Elevate project began with a simple—if odd—question: "What if a car had legs and could walk?"
    - "Legs and wheels simultaneously together is a relatively unexplored robotic research area," says Suh.
  - He relates the story of how his team took this question and ran with it—well, walked, rolled, waddled, and climbed with it.
    - The still-active journey has entailed an active imagination coupled with the ever-present need to make down-to-earth tradeoffs and decisions.
  - The result is a concept vehicle like no other: the Elevate—which debuted at CES 2019, albeit in a combination of videos, computer graphics, and a small-scale partial prototype—is an ingenious combination of a wheeled electric vehicle with extendible legs and a chassis that could potentially serve variously as a base for a two-passenger senior mobility microcar, a family-sized hyper-SUV (i.e., a Jeep Wrangler-sized subcompact), or an intentionally go-anywhere search-and-rescue or emergency response vehicle.
    - In all cases, the goal has been to provide alternative modes of locomotion, enabling the Elevate to extricate itself from a snowy ditch, deliver a mobility-impaired passenger to the top of a brownstone's stoop, pick its way across a flooded section of roadway, navigate through off-road debris or over post-disaster-scene rubble, or simply lay down some highway miles.
      - SUVs have been touted as never-get-stuck vehicles, but nature has a way of foiling that vision.
      - The Elevate brings additional motional components to bear, leading Suh to call it the "ultimate mobility vehicle (UMV)."
      - "If you could get up and walk to safety on robotic legs to improve the mobility of cars, that would be of great utility," says Suh.
- At the outset of the project, the team took a no-holds-barred brainstorming approach to what a UMV might be like.
  - Cars that could roll over one another, pneumatic appendages to coax the vehicle and like an "active jack" when stuck, six wheels for extra stability, legs that could unfold a foot during walking mode—each of these was considered, but ultimately rejected.

- The team made the choice for the Elevate to retain its inherent car-ness, but to incorporate additional features to enhance its utility and mobility.
  - It would remain a wheeled vehicle, but with tucked-away legs that would come out only when functionally necessary to go over an obstacle or extricate itself (while nevertheless conferring a bit of extra shock absorption when folded).
    - That is not to say that legs would be hidden from sight: “The legs themselves start to take part in the overall aesthetic of that vehicle,” says Suh. “There is something on the exterior that indicates to the user and others that there is something different about this vehicle.”
- Nature served as inspiration for the vehicle’s walking modes:
  - Grasshopper legs fold away efficiently, yet provide great power and range.
  - Horses employ several distinct gaits, some for nimble mobility and others for speed.
  - Reptiles spread their legs wide for very stable ambulation—“Imagine being asked to walk while doing a pushup,” says Suh. “The motion is that you swing [your limbs] out to the side.”—including the use of large feet.
- Putting these together, the multimodal Elevate can, as mentioned, roll like a normal car with legs tucked up against the chassis, or it can extend the legs and walk over rough terrain with a mammalian gait, or maneuver closer to the ground using an extra-stable reptilian-like gait with not only its legs splayed, but also its wheels.

- Getting down to the nitty-gritty of automotive leg design required early consideration of the number of degrees of translational freedom for the Elevate.
  - A conventional vehicle has two degrees of translational freedom: wheel rotation and steering.
  - The initial Elevate model added three more: incremental forward motion (inchworm-like creeping), an additional rotational axis to yield minimal but jerky walking capability, and, with a fifth degree of freedom—reptilian-style motion through the “shoulder joint, so it can go up and forward”—a comfortable walking gait becomes achievable.
- Actuation has been the Elevate project’s thorniest problem; the team considered three alternatives.
  - Hydraulic—This mode was implemented in the first prototype, due to the benefits of high force, high torque, and proven ability to be robust in a variety of harsh environments; moreover, an important safety consideration is that hydraulic components hold their position when deenergized.
    - On the downside, they require considerable maintenance, tend toward leakage, and carry the overhead associated with auxiliary components, such as pumps and reservoir tanks.
    - Working with University of Michigan’s RAMLab, Suh’s team simulated a vehicle with hydraulically powered legs to surmount obstacles (ramps, walls, downed trees) and span gaps to evaluate the potential for a vehicle of reasonable size and mass to function as desired using currently available technology.
    - For a microcar configuration—two passengers, one seated in front of the other—RAMLab simulations that considered pressures, flow rates, torques, and so forth, provided positive confirmation for hydraulic legs that would be capable of converting from rolling to walking/stepping over a one-meter-high obstacle (at the pace of a casual stroll), and then rolling once more.
    - With the senior citizen market in mind, such a vehicle could enable elders to remain in their homes, with no need to install wheelchair ramps, since the Elevate could load passengers directly.
      - “I call this the house of the future,” says Suh. “It is the house you live in today.”
      - Similarly for the “vacation of the future,” during which adventurers need not worry about getting stuck in a snowbank, knowing the vehicle could switch to reptilian mode to extricate itself if necessary.
  - Electric—Quieter than pneumatics, and more capable than hydraulics for providing precise/repeatable motion using modular components, electric actuators offer intriguing advantages.
    - But electric actuation is disadvantaged, relative to hydraulics, regarding force-to-volume output.
    - Still, electrics would enable a valuable sixth degree of freedom—the possibility of combined flexion/rotation—to the five previously described.
      - The hydraulic model offered longitudinal shoulder motion; the electric model added a perpendicular axis, as well: “The first degree of freedom is rotation before adding this flexion,” says Suh.
        - Once implemented, little modification is necessary to enable the vehicle chassis to lower to ground level to accept the entry of rolling units, such as wheelchairs or carts.

        - Additionally, this extra degree of freedom permits the use of car tires (with a square profile), rather than needing to resort to motorcycle tires to ensure a strong platform with adequate traction when using the cantilevered reptilian mode.
    - Pneumatic—Suh rejected this mode due to lack of the technology's maturity to actuate and control a vehicle of this scale.
- Taking actuation considerations into account, the Elevate concept that Suh's team ultimately pursued is a larger, intentionally off-road-capable vehicle ("a Jeep on steroids") with electric actuators, six degrees of translational freedom in the wheel–leg assembly and intended use cases of search-and-rescue, disaster relief, or as a mobile living space (the car (like a garage) can become an extension of the home)—each mode would be supported by a distinct body type that would rest, interchangeably, on the same base structure.
    - "We couldn't quite nail down the exact application," says Suh, "so we explored whether we could make a modular vehicle architecture that could do different things."
        - "We have a body on chassis on wheels," he says, conceding that more work (e.g., structural and material analyses) has yet to be done to confirm the ability to develop a lightweight, rigid chassis that would complete the concept's modularity and also—being lined with batteries—serve as the vehicle's power source.
    - Strength: The body, of whichever sort, would have external framing to confer strength as part of its design aesthetic, as well as four-sided access to maximize flexible use.
    - Suspension: When in conventional driving mode and not activating its specialty ambulatory capabilities, Elevate makes use of coil-over-damper suspension for absorption of vertical movement in the zero-to-three-inch range.
        - Over more severe terrain, the legs extend to perform in "active-knee" mode for variable clearance within the 3–50" range, keeping the chassis/body horizontal.
        - To retain a car-like exterior aesthetic, each leg–wheel assembly is attached to the chassis using a cantilevered offset, rather than a motorcycle-like fork.
- With all of this in mind, for demonstration purposes, the team constructed a one-fifth-scale model of a search-and-rescue Elevate that has traveled to various car and consumer electronics shows.
    - The introductory video describes the six degrees of freedom somewhat differently than above: vertical hip, horizontal hip, knee, ankle, steering, and wheel.
    - Battery capacity in the chassis specs out at 66 kWh.
    - Interchangeable body configurations depict a 1+1-passenger sport model, six outward-facing passengers in the SUV model, and an end-loaded gurney in the search-and-rescue model.
    - The concept video of terrain management is perhaps overly ambitious, with an Elevate expertly navigating a rough, rocky, steep mountainside, yet a video of a one-eighth-scale model of the base-plus-chassis—created using off-the-shelf parts specifically for testing purposes—successfully rolled up to and climbed over a coffee table with none of its wheels slipping during the maneuver.
    - Autonomous control of locomotion is, as yet, "an unexplored area," says Suh, but he anticipates that the vehicle will make its way with the sole use of software and sensors.
        - "I expect it will be somewhat like riding a horse," he says. "You don't tell your horse where to place its feet; you just tell it where to go."
    - This is not to say that Suh anticipates Elevate to be fully autonomous, at least not out of the gate; instead, concept drawings include a human operator, perhaps standing in a central position within the vehicle, akin to the operator of a San Francisco trolley.
- At present, Suh has no estimate of ultimate pricing, if the vehicle ever makes it to production.
    - "To be honest, this is a science experiment," he says. "It is probably overbuilt, but we will learn a lot about how to activate legs and that sort of thing."

Making Robots Less Brittle—Dr. Doug Lenat, Cycorp and TTI/Vanguard Advisory Board

- Anyone who had witnessed any of Lenat's prior presentations to TTI/Vanguard would know that he is addressing robots' cognitive brittleness, rather than defects in their material composition.
    - The current popular approach to imbuing robots with smarts is to leverage statistics via induction or machine learning.
    - While recognizing the value of this right-brain style of thought, which focuses on pattern recognition, Lenat has instead devoted his career to the left-brain strategy of "capturing pieces of knowledge, encoding them as rules of thumb, and letting the program do a kind of logical inference—logical deduction."

- Just as humans benefit from both left- and right-brain thinking, so too can robots; Lenat cautions not to become caught up in the deep-learning frenzy, but rather to also give logic its due.
- Lenat's talk has been 35 years in the making; this is how long Cycorp has been priming the pump of the Cyc knowledgebase with hand-coded nuggets of common sense.
- It has matured into what can rightly be called a thinking machine that largely relies on logic, but resorts to statistics when expedient, as it successfully answers complex questions and provides background reasoning—in plain English.
- Lenat offers examples of how statistics-only-based AIs have lulled us into doing the work for them and shares the potential of an AI that puts logical reasoning front and center.
  - "These two approaches have enormous strengths and enormous weaknesses," says Lenat. "The good news is that they are different strengths and weaknesses, so one could theoretically put them together. And, not just theoretically, they are put together now in some applications, and that is where I think the future of AI lies."
  - Lenat does not, however, credit Google's Knowledge Graph or IBM's Watson with having achieved the blended processing he touts.
    - "They are taking a baby step in the right direction, but in fact something significantly more has to be done, can be done, and is being done already," he says.

- Simple examples expose the brittleness of each brain hemisphere:
  - Right-brain reasoning, as performed by the Google search engine, does well in one-step reasoning tasks, such as presenting the correct date when asked when Theresa May was born (Oct. 1, 1956).
    - Not only does Google leave the exercise to the questioner to determine "who was prime minister of the UK when Theresa May was born," but neither can it straightforwardly answer the more targeted question of who the prime minister was on Oct. 1, 1956.
      - Google's top hit to the latter question is the Wikipedia page listing UK's leaders, beginning with Sir Robert Walpole (1721–1742), including Sir Anthony Eden (1955–1957), and concluding with Theresa May, who persists in that position through this writing.
    - That is, Google (and many other convolutional-net-based AIs) have trained its users to do the heavy lifting of coordinating facts into knowledge.
    - "Out of the million hits," says Lenat, "none of them answers the question, because Google has trained us not to expect even one step of logical or arithmetic deduction to take place."
    - More generally, machine learning suffers when there is insufficient relevant data, when two matters are considered simultaneously (as in the prime minister example, or when a student makes two distinct errors when solving a math problem), for scenario generation, when the goal is to weigh the pros and cons of various lines of reasoning, or when the questioner poorly poses a question.
      - These are all areas in which Cyc performs admirably.
  - This is not to say that left-brain reasoning is unassailable, as Lenat exposes with an example from a skin disease diagnosis system.
    - *System*: "Are there spots on the body?" *Answer*: "Yes."
    - *System*: "What color are the spots?" *Answer*: "Reddish brown."
    - *System*: "Are there more on the trunk than elsewhere?" *Answer*: "No."
    - This exchange produced a diagnosis of measles with nearly 93% confidence; however, the "patient" was an old car with a rusted-out door.
    - "The system didn't have the piece of common sense that mechanical devices don't catch human diseases—whether or not they have been vaccinated," quips Lenat.
    - More generally, Cyc underperforms if the necessary expertise to answer the question does not exist (i.e., human inability to introspect and articulate), rules are misapplied due to different terminology/context/assumptions, or data is contextual or entails exceptions or contradictions.
    - Symbolic-based AI is also costly to bootstrap (as per Cyc's 35-year lead time) and can be slow, due to the huge search space for proofs.
  - Based on experience with many examples akin to those above, Lenat concludes that AIs that rely on either left- or right-brain approaches have only the veneer of intelligence or of understanding, but not truly the real thing.
    - Joe Weizenbaum's 1960s-era ELIZA program, which posed as a psychotherapist engaged in the practice of reflection, parroting back the core of each patient's revelations to help the patient discover more about him/herself.
      - The system worked reasonably well until it came upon language it didn't understand and produced a non sequitur.

            - How well do today's natural language systems—e.g., Alexa, Siri, Cortana, Google Assistant—do in comparison? You be the judge:
                - *User*: Call me an ambulance; *Siri*: From now on, I'll call you "An Ambulance". OK?
                - *User*: I think I have alcohol poisoning what do I do; *Siri*: I found 7 liquor stores fairly close to you.
    - Nevertheless, within appropriate domains, statistics-based reasoning does a good job by leveraging the wisdom of the crowd or large and relevant datasets.
        - Moreover, open-source libraries (e.g., Google's TensorFlow) and public datasets (e.g., ImageNet) make it easy and cheap to tap into the power of statistics.
    - And a logic-based reasoning system provides step-by-step explanations as a consequence of its operation, which enables the human to evaluate the assumptions at each step and/or offers insight on how to experimentally test a theory.
        - Demo: Lenat poses the following question to Cyc: "Can a can cancan?" Answer: No.
            - Among the justifications are that a cancan is a dance entailing leg kicks, cans lack legs, cans lack brains, and cans are not animals.
            - Cyc's reasoning is fast, with most of the time between question and answer devoted to converting the logic into English language sentences.
- Lenat concedes that the story he has told thus far should be met with a big yawn, saying, "Systems based on neural nets, on the one hand, and expert-system-type rules, on the other, have been well understood for about 40–50 years."
    - Statistics-based AI has been waiting for Moore's Law to catch up and offer computers that are fast enough and data that is ubiquitous enough; it has taken those advances plus the technological breakthrough of convolution to bring neural net-based AI into its own.
    - Logic-based AI has been awaiting heuristics to overcome the challenges of context, inconsistency, and dozens of other matters, as well as the buildup of factual and commonsense knowledgebases over which to reason.
        - Actually, logic-based AI hasn't been waiting; Lenat's team at Cycorp—heavily weighted with PhD philosophers introspecting on the unstated assumptions of text—has been steadily addressing both heuristics and common sense, such that Cyc is now an able reasoner.
        - Focusing on expressiveness at the expense of the efficiency of limited logics (such as Bayesian networks, knowledge graphs, triple-stores, partitioned semantic networks, description logics, and so forth), Cycorp has stayed the course and built out its system to include over 24M axioms and an equally ambitious inference engine that includes more than 1K special-case reasoners (to increase the system's speed and efficiency) and can tap into a combination of general knowledge and specific data, facts, terms, and observations to draw conclusions and explain them.
        - The upshot is an ability to reason deeply.
            - Lenat uses the example of answering deep questions—and providing explicit reasoning—related to Shakespeare's *Romeo and Juliet*.
                - Examples: *Does Juliet, at the time she drinks the feign-death potion, believe that Romeo will believe that she is alive from that time through to when she awakens?* and *Which of Juliet's lines are satirical? Metaphorical? Hyperbolic?*
            - Answering such questions is only possible because Cyc represents meta-level information, tactical and strategic information, contexts, and so forth—and all in the language in which questions are posed and answers and explanations are provided.
            - Of course, the system can just as successfully be applied to contemporary news stories, matters of geopolitical importance, medicine, whether cans can dance, and so much more.
- "Both machine learning and symbolic reasoning, like Cyc does, are available and are being harnessed today," says Lenat; tying them together is the crux of "exploiting their latent synergy."
    - The process is as follows:
        - The machine-learning and knowledge-based system mutually agree on an API through which to interface.
        - The knowledge-based system generates a vast amount of training data for to train the machine-learning system.
        - Thus trained, the machine-learning system generates hypotheses.
        - The knowledge-based system individually evaluates those hypotheses by generating plausible causal scenarios.
    - Cycorp has applied this to a medical application—a gene-wide association study related to osteoporosis.

        - Part of the multistep scenario included the expectation of elevated vitamin D levels—objectively evaluable by bloodwork—as Lenat detailed in his address to TTI/Vanguard in Austin, February 2016.
        - “Along the way, you generated one or two or three or four independently testable predictions,” says Lenat. “You are using statistics to confirm or disconfirm the predictions that your causal engine made. It is that kind of synergy that we’re talking about.”
- Looking toward the future, Cycorp will be involved with DARPA’s Machine Common Sense program, which will commence late June 2019.

Humans, Robots, and Cognitive Economics—Mr. Leigh Caldwell, The Irrational Agency

- Today’s dominant form of artificial intelligence involves the use of deep neural networks to perform machine learning over big data.
    - With enough exposures to applicable data, an AI learns to categorize according to patterns within the data, but individual nodes in the neural net do not correspond to specific characteristics.
        - Example: A neural net learns to recognize cats, but there are not nodes that respectively evaluate for whiskers, pointy ears, and other components of cat-ness.
- Contrast this model-free strategy with the model-based approach relevant to systems with known rules—like games, such as chess or Go—in which an AI can iteratively play against itself and become increasingly adept at the narrow task.
- Systems like IBM’s Watson take a page from each of these books to gain additional sophistication, yet most problems that would benefit from intelligent analysis remain out of reach of even this combined approach.
- The field of cognitive economics takes into account a distinct data type: that which resides in people’s minds—expectations, beliefs, hypothetical choices, and attitudes—to introduce the notion of human attention to the AI computational toolkit.
    - Caldwell emphasizes the importance of attention to decision-making and how people are simultaneously adept and irrational in how they apportion it.
        - A deep-learning AI that incorporates mental models into its training phase should outperform competitors when predicting human preferences and decision-making.
    - Caldwell demonstrates the power of cognitive economics with a shopping example.
- “AI can’t solve what you are likely to do tomorrow, because your decision is not based on big data or rules,” says Caldwell.
    - Instead, human choice is based on what percolates within that particular person’s head; beliefs, biases, and opinions all play into the disposition of attention.
    - Cognitive goods therefore matter greatly to digital-heavy firms that compete for users’ attention.
        - Things that focus mental attention—what excites you, who you care about, your self-conception—do not fit the economic concept of goods that can be bought or sold, but they consume us mentally and therefore confer value.
            - The boundary for goods covered by cognitive economics is the boundary of the body.
                - For instance, conventional economics encompasses buying and eating food, but *thoughts* about tasty food lie within the cognitive domain.
        - Rewards come from not only external stimuli (eating something delicious), but also from internal stimuli (thinking about the food, or even about the food’s packaging or the store where it is bought, due to previously established mental associations).
            - “You get a lot of pleasure from what goes on inside your head,” says Caldwell.
            - People give their attention to what provides them with rewards/pleasure.
                - Implicit in cognitive economics is recognition of human creativity.
                    - Caldwell believes creativity is a third system of thought, distinct from system 1 (automatic/fast/unconscious) or system 2 (effortful/slow/controlled).
                - “The unconscious stories in your imagination can teach machines to dream,” says Caldwell.
- Caldwell incorporates cognitive economics into the machine-learning process by using a directed graph trained on a model of human emotion as the hidden layer between decision-making inputs and outputs, as he demonstrates with a Primark fast-fashion shopping example.
    - That is, Caldwell plugged the human mental model into the machine to generate a cyborg output.

- Example component of the cognitive-economic hidden middle: *variety*, which is influenced by *choice*, *good value*, and *trend-led fashion*; and which influences *makes you feel good* and *everything under one roof*.
- The result for a given customer will be individualized suggestions of items to purchase.
- Doug Lenat expresses frustration that this version of AI diminishes the basis for decision-making into increasingly primitive logic: only nodes and directions, without even labeling graph edges.
  - Caldwell counters that it is an improvement on many of today's commercial AI implementations.

- Example applications of cognitive economics, all of which are intangibles: attention, imagination, planning for the future, empathy, consumption of fiction, symbolic objects (beliefs, identity, sports teams, political parties).
  - Future work will strive to find meaning in specific layers of the cognitively enhanced neural net with a goal of making such systems reusable for related problems without wholesale retraining.
  - "If we can give the android an electric mind, we can give it its dreams," says Caldwell.
- Just as the basis of cognitive psychology is how thoughts influence behaviors (rather than considering behaviors per se), cognitive economists recognize that people behave irrationally relative to their self-interests, yet reasonably when considering how beliefs influence behaviors (i.e., according to an articulable model).

Humans in the Loop: Lessons from ARMOR1—Mr. Gabriel Goldman, Carnegie Mellon University

- Automation is most usefully applied when humans performing the same work would be in danger, or when automation is economically advantageous.
  - The purpose of Goldman's work with the U.S. Army Corps of Engineers to automate the process of laying concrete mats to shore up the banks and the riverbed of the ever-meandering Mississippi River is to make working conditions safer.
    - Like many mighty waterways, the Mississippi has spontaneously changed its course throughout history, which wreaks havoc on enterprising farmers, industrialists, and city-dwellers along its banks, who might be advantageously positioned for river-aided irrigation, transportation, or power generation one year, and then be either high-and-dry or underwater the next.
    - To mitigate such risks, the Army Corps has been working tirelessly—and largely successfully—throughout the low-water season, year in and year out since 1930, to keep the river within consistent banks, but at considerable risk to members of the Corps and at considerable expense to U.S. taxpayers.
    - The practice of sinking mats has proven effective for erosion control over hundreds of years, beginning with the use of woven mats by the Japanese.
  - Goldman explains the need, the process, and the execution of this mammoth automation project, explicitly promoting the benefit of devising the robotic system to work in tight coordination with human workers who provide supervision, backup, troubleshooting, and dexterity.
    - "Automated systems don't always have to be 100% automatic," says Goldman. "Automation with humans physically in the loop can make for a more effective overall process."
  - Although only partially built out to one-sixth scale, the physical scope of ARMOR1 already exceeds anything previously built at Carnegie Mellon University's National Robotics Engineering Center.
- To motivate the need for automation, Goldman first lays out the details of the current, largely manual, 1950s-era process of manufacturing concrete mats, transferring them to a barge on the river, and sinking them in the Mississippi.
  - Each 3600-lb section of matting is a 4' x 25' x 3" flexible slab of concrete that is precast offsite (with fabric wires adhering adjacent subsections) and transported to a supply barge in batches of 40 stacks, each 13 slabs high.
  - In continual rotation, resupply barges approach the mat barge, while workers in close proximity to huge cranes align the cranes with strips of mat to transfer them stack-wise, onto the mat barge; meanwhile, other workers manually use pry bars and tie guns (previously pneumatic and more recently electric) to connect the mats to one another in 35-mat-wide strips that together form a huge mattress of flexibly connected concrete, which progresses—conveyor-belt-like—onto the riverbank and into the drink as a revetment.
    - The "tie" consists of fabric wire, launch cable, and tie rod; the tie gun wraps the tie rod around the cable and wire within the open space—dubbed the scarf box—that lies at the intersection of each pair of properly aligned adjacent mats.
      - When scarf boxes are poorly aligned, insecure ties are likely.

            - Tying proceeds even as already-tied segments of matting progressively roll into the river beneath the feet of mat barge workers.
            - To lay down each 35-mat-wide section of concrete mat, the mat barge retreats a few hundred feet toward the center of the river.
                - Once in position, workers cut 36 wires, under tension, to release the final strip of mats, completing the sink of that section—a highly dangerous procedure.
        - In total, this undertaking requires more than 250 workers and support crew—barge motion operators, on-deck workers, under-deck workers, supervisors, maintenance personnel, and support staff for food, living quarters, and transportation of people and equipment.
            - The annual cost for maintenance alone is $5M–6M.
- NREC's automation does not strive to greatly alter the entrenched, successful mat-laying processes, but rather to roboticize the most dangerous components, retaining the human workforce for supervision, maintenance, and stepping in as real-time backup when the machinery fails to effectuate a tie or other procedure; the intention is for the new system to be operational for the next four decades.
    - That is, the intent is to leverage the fact that people are flexible and resilient, and adept at using tools in multiple ways and making corrections when things go awry.
        - The automated system is therefore designed to perform all the heavy, manual, dangerous work that Army Corps members have historically undertaken, but with an anticipated failure rate; workers would intervene as needed.
            - To do so effectively requires a vision/sensor system that accommodates imperfect alignment between vessels and imperfect (e.g., droopy) lifts of mat strips from the resupply barge to the mat barge, and provides alerts to the supervisory team for visual occlusion of a scarf box or a poorly executed tie.
    - NREC's assessment allots 120 seconds to pick, place, and tie each mat section—assuming 7.5 seconds per side tie)—plus 60 seconds to launch, amounting to three minutes per launch.
        - Only the robotic system would occupy the mat storage and mat-loading zones, while humans would also be present in the mat-tying and completed-mat zones for manual backup in the event of mat-tying irregularities; below-deck operations would remain as human-operated zones, as with the current, unautomated system.
        - Jobs to be required under the new automation regime: loading operators, tying operators, barge operators, on-deck tying operators, on-deck supervisor, below-deck operators, and maintenance crew, plus auxiliary staff for food, lodging, transportation, etc.
    - As designed, the six-armed system will total hundreds of degrees of freedom, including six placement arms (6 dof each), each with a 37-dof automated lifting frame, an automated deck conveyance (36 dof), a 1-dof automated tie gantry with 9 13-dof end tie modules and 36 6-dof side tie modules.
        - The arms move linearly, picking up a line of mats from the resupply barge, transferring it to the tie gantry, and progressing the collection of tied mats to the mat-sinking conveyance.
            - Each lifting frame arm performs passive mechanical alignment using force equalization on either side of the stack of mats on the resupply barge.
            - Once aligned, a mechanism shoots down through the scarf box indentations to compress and lift a double-high set of concrete sheets and move them linearly to the mat barge.
            - Visual feedback (green, yellow, red) informs supervisors of readiness for each lift.
        - Once the full complement of mats have been arranged 35 sheets across, a mechanism shoots down through the aligned scarf boxes, row by row, to perform the ties.
            - A leading machine vision sensor verifies scarf box alignment; a trailing sensor confirms tie quality and informs workers of the need for a manual tie, for instance if the scarf box is visually occluded or notably misaligned.
            - The system components are intentionally adaptable to failure; if a part goes down, people will take up the slack until the mechanism is again functional.
    - NREC has assembled a full-sized, one-arm prototype at its Lawrenceville location; final construction will take place close to the Mississippi River.
- As previously noted, the intention is neither to replace workers nor notably accelerate operations, but rather to improve working conditions and safety.
    - Worker feedback has been integral to the development of the robotic system.
        - It is expected that the new system will reduce on-deck personnel by 25%.
    - "We are removing the people from the most unsafe portions of the work environment," says Goldman.

Robots, Immigration, and Work—Dr. Richard Freeman, Harvard University

- Each wave of technological advancement has carried with it societal concerns about worker displacement.
  - Yet, transitions from the agricultural to industrial age, and from the industrial to computational age, have left plenty for workers to do, with the economy continuing to grow and plenty of opportunity for workers to advance.
  - Now, with AI and robotics claiming not only the tasks of blue-collar workers but also those of highly educated professionals, are things different?
  - Freeman believes that the current technological transition is indeed different than those that have come before, as evidenced by wage stagnation, the declining share of national income enjoyed by the labor force, and the escalation of wage inequality.
    - "This is the first time [a scare associated with technological change] is being driven by wage inequality and media reports about automation, not by low employment," says Freeman.
      - At the macro level, with high deficits and low interest rates, the U.S. Federal Reserve has few tools to manage the next recession when it descends.
    - In the age of automation, the key to thriving will be to own the source of production—that is, to own the robots.
      - In contrast, education alone will no longer provide a leg up in society.
    - Freeman offers two strategies for societal wellbeing rooted in spreading the wealth:
      - redistribution of income as the result of taxing the rich;
      - altering the ownership model to give workers a share of the companies that employ them.
- To put the current technological transformation into perspective, Freeman fast-forwards 20 years to take the vantage point of the baby being born today.
  - No longer will this newly minted adult be able to count on the benefits of higher education to secure a well-paying job: machines will be gaining a comparative advantage, not only as truck drivers and proof readers and bookkeepers, but also in the skilled and creative work of programmers and novelists (one class of AI would spew out novels, while another would evaluate them for virtuosity).
  - At the very least, economic logic dictates that the "robolution" will depress wages across many sectors of society.
- This is not to say that today's workers believe Freeman's scenario will directly affect them.
  - Instead, most—benignly or even naïvely—consider that their jobs will remain intact, while only those less skilled than they will suffer from the effects of AI/roboticization.
  - Moreover, macroeconomists are not sounding the alarm, but are instead adhering to the conventional wisdom that, on the one hand, product innovation will lower prices, encouraging more robust sales; and, on the other hand, process innovation will reduce the cost of production, reducing prices, again spurring sales.
    - As such, joblessness induced by technological change will be absorbed by the demand for workers in new jobs, although perhaps at lower wages.
- What today's babies will encounter when they enter the workforce will be jobs that humans can perform at a lower cost than machines—robots or software—can.
  - Two decades hence, people will still likely surpass machines in their ability to walk, run, and kick a ball with precision, but machines could well outdo humans when it comes to neural activity.
    - Consider the accelerating pace with which machines can self-train using generative adversarial networks and other forms of reinforcement learning.
    - As more aspects of work become digital, more jobs will be at risk for automation.
    - "The biggest software investment is in the brain-intensive industries," cautions Freeman. "It can assist us—or replace us."
  - Freeman believes that there will continue to be jobs available for people, but questions whether these will be good jobs or bad jobs in the wake of improving robots and AI.
    - That is, to a significant extent, as machines gain a competitive advantage over people, robots will employ human workers, rather than the other way 'round.
    - "The issue is what people are doing and what money they are making," says Freeman.
- In light of these considerations, Freeman offers his "three laws of robo-economics":
  - "Robots and humans become better substitutes."
    - Technological advancement will lead to higher elasticity of substitution between people and robots/AI.

                - Example of human-like robots: use of biomimicry to improve robotic hands, including artificial skin with tactile sensors and stretchable electronics.
                - Example of robot-like human: brain–computer interfaces.
            - Substitutions will be driven by a combination of gains to firms and gains to society; i.e., society incentivizing firms based on relative merits/demerits of implementing substitutional technological advancement.
                - Substituting robots in dangerous jobs is a societal boon.
        - “Technological change reduces costs of robot substitutes for humans over time, bounding wages.”
            - Human wages will not exceed the production/operational cost of a robot/AI substitute.
            - The cost of robots can drop in two ways: decreasing costs for a fixed skill set, and fixed or modestly increasing price for an improved skill set.
            - As robots/AI and human labor become functionally interchangeable, firms will “employ” whichever costs less, driving down wages for jobs held by human workers.
        - “The effect on incomes depends on who owns the robots.”
            - If a worker owns the robot/AI that performs that person’s work, the benefit will go to the worker, but when an external entity—nominally, the employer—owns the robot/AI, the worker suffers.
            - “Who owns the robots rules the world!” says Freeman.
- Freeman also explores the relative impact on U.S. jobs of robots and immigrants:
    - The near-term shock of industrial robots on employment is considerably lower than that of new immigrants.
        - Data:
            - The year 2017 saw the purchase of 30K new industrial robots compared to the influx of 900K new immigrant employees.
            - Cumulatively, there are 28M immigrant workers, amounting to 17% of the U.S. workforce, and 350K robots deployed.
                - The robot stock is currently concentrated in the automotive sector.
        - Assumptions:
            - One robot replaces two-to-three workers (but one-to-one replacement by immigrants).
            - Supplantation by robots largely affects workers without a college degree, as is the case with immigrants.
    - This analysis considers only the conventional definition of a robot—i.e., a machine capable of performing autonomous physical tasks—and does not include digital-only automation.
        - The impact on jobs would more severe were AI to be considered, as well.
- A future society in which work increasingly shifts to machines will be characterized by more severe income inequality and will therefore be fundamentally unstable—unless the model of capital ownership changes.
    - Freeman’s solutions:
        - Tax relatively heavily and redistribute income, perhaps with universal basic income for all.
        - Change the model of capital ownership by shifting ownership in the direction of employees, whether through the broadscale expansion of employee stock ownership plans (ESOPs) or by implementing sovereign fund ownership, akin to a mutual fund with the assets being the nation’s stock of robots/AI.
    - Society will regain stability and be broadly affluent if everyone has an ownership stake in the nation’s resources.
        - Education alone will not lift workers to greater levels of affluence and the fulfillment of the American dream when AIs gain the upper hand at task performance.

Bio-Inspired Robots—Dr. David Hu, Georgia Institute of Technology

- Life has had eons to devise clever ways to attend to its basic needs, providing intriguing inspiration for roboticists.
    - As a professor of mechanical engineering and biology, Hu is well-positioned to explore the particularities of animal behavior that have provided them with evolutionary advantage, and his curiosity has led him to explore widely.
        - In just the past couple of years he has published on moles, black soldier flies, jellyfish, frogs, anteaters, honeybees, snakes, leaf beetles, elephants, cats, and fire ants (not to mention the hydrodynamics of defecation and urination).

- Although his research has extended to the development of bio-inspired robots, for TTI/Vanguard he limits his discussion to the intricacies of elephants' use of their trunks, cats' use of their tongues, and fire ants' use of one another to form floating rafts and towers.
  - In each case, the detailed biomechanics could be used to inform the construction of robots that perform the tasks specific to the animal model.
- Elephants:
  - It is not news that African elephants are large; to remain viable in the wild, they have to eat a lot—more than 200 kg of vegetable matter daily, which means constant consumption throughout their waking hours.
    - An elephant needs its dexterous and versatile trunk to be an efficient enough eater to survive.
  - At 100–150 kg, an elephant trunk is nature's largest land-based muscular hydrostat—a boneless organ composed of a 3-D arrangement of interdigitated muscle fibers.
    - Given its implicit heft, the trunk must be strong, and to perform its ingestion roles, it must be able to variously grab, wrap, blow, or suck, depending on the type of material the animal is collecting and the trunk's proximity to it; an elephant's dexterity is greatly amplified when able to see the object(s) it seeks to gather.
      - Grab—Less force is necessary to grab large objects (e.g., 32-mm vegetable cubes) than small objects (e.g., 10-mm cubes or 2-mm grain granules).
        - The trunk remains fully extended to grab a small number of large cubes, with no need to leverage the force of gravity; it instead relies on end-of-trunk dexterity.
        - The smaller the unit of food, the greater the end-of-trunk angle, as the elephant creates a knee-like bend toward the extremity of the trunk to effectively grab a large number of small particles; moreover, suction plays a role when collecting small objects.
          - "The elephant first sweeps it, then scoops it, pushes it down, and finally grabs it," says Hu, with the elephant "building a taller pillar for smaller particles."
      - Wrap—As an alternative to grabbing, the trunk can curl around the cubes to gather them together, then squeeze the pile collectively within a loop to convey the food to the mouth.
        - Dramatically, an elephant can readily be taught to demonstrate the raw strength of its trunk by lifting a barbell by wrapping its trunk around the bar.
          - Experiments using different weights made clear that the heavier the weight (up to 60 kg), the more vertical the trunk and the greater the grip area when wrapping around the bar to lift it.
      - Blow—Before grabbing/sucking a collection of small particles, the animal can use a blowing action to gather the food particles into a pile, as Darwin documented in 1874.
      - Suck—Elephants approach large-area, lightweight food items, such as a fried tortilla, differently: Since applying as little as 1 kg of pressure would crush the chip, only suction is applied.
        - Mammalian noses can apply 10 kPa of pressure, on average; small-nosed mammals (like humans) must be very close to the object of suction to be effective, but the large surface area of an elephant trunk means that can lift a 15-g tortilla from a considerable distance with just an intake of breath.
        - Of course, elephants use their trunks to suck more than tortilla chips; they also suck water and release it with great force.
          - An elephant can intake 6 L of water in 1.5 seconds—aided by 12–16% expansion of the nostrils—and expel it with half the pressure of a fire hose.
        - A trunk siphons air at 70 m/s (160 mph); compare to the 110-mph exhalation of a human sneeze.
- Cats:
  - "Cats are championship groomers," says Hu—and not only do housecats spend one-third of their waking hours licking themselves, but lions, tigers, bobcats, snow leopards, and cougars are also fastidious in their self-care.
  - With the housecat as model species, Hu studied the anatomy and physiology of feline grooming.
  - Each 0.25-second-long lick commences with the extension of the tongue and proceeds with expansion, sweep, and retraction steps.
  - Crucial to a cat's ability to groom is the presence on the tongue of rear-facing spines—papillae—that have a hollow cavity at the tip that uses capillary action to wick saliva from the mouth, then efficiently deposit it on the hairs with each lick.

        - The saliva serves two purposes: cleaning fluid and cooling agent, since a cat only has sweat glands on its paws and instead uses the combined action of hair separation and fluid application to regulate its temperature.
    - The structural design of combs and brushes has been largely static for thousands of years, but this new understanding of cat tongues could prove a game-changer, were the biologically proven model applied to manufactured grooming tools.
        - Not only are papillae effective for unidirectional grooming, but—collectively—they shed hair in with a simple swipe.
- Fire ants:
    - Fire ants collectively form amazing structures, ranging from floating rafts to towers.
        - Rafts—Although individual ants cannot survive underwater, when a swarm of hundreds, thousands, or hundreds of thousands of individuals self-assemble into an intertwined mass, the raft can float for as long as months.
            - The ants grip one another using mandibles and tarsi, as well as adhering with the help of natural glue that exudes from the bottom of each foot and can stick to up to 20 connection points on large ants and to roughly eight on small individuals.
                - Each such connection can exert a force of 150 times the individual's body weight.
            - An ant's exoskeleton is hydrophobic, enabling the insect to surround itself with a small bubble of air; even ants on the downward-facing side of the raft have sufficient buoyancy to float and sufficient air to breathe.
            - Collectively, a fire ant raft is water-repellant due to the enhanced effective contact angle between the raft and the water surface (compared to that of individuals), as per the Cassie-Baxter law of wetting.
            - Ants, as a collective, can be both springy and dissipative, able to survive both hard and soft impacts.
                - On the one hand, they can store elastic energy; on the other hand, they can exhibit flow through viscous dissipation.
            - A "snowball" of ants, if dropped into a pool of water, will upon impact spread out from a mass roughly 30 ants high into a two-ant-high raft, with each ant having eight nearest neighbors, on average; this process typically takes less than two minutes.
            - The arrangement of individuals within a raft is hardly static, with positions generally rotating from bottom to top, albeit with a random component.
            - An ant raft that must float for long periods lacks a source of food; the colony's solution is to eat the babies first.
            - Any ant that escapes from the raft rarely makes it back.
        - Towers—As noted above, fire ant rafts are bilayers, as an ant cannot withstand more than twice their weight for any length of time; however, when a raft anchors to vegetation, they collectively build dynamic—yet stable—tower-like structures that are bell-shaped, through the even distribution of load.
        - A tower can grow to exceed 30 ants in height over the course of about 25 minutes.
        - The purpose of a tower is to protect the colony from rain and water currents after a flood until it can build a new network of underground tunnels.
        - When forming a tower, each ant is individually exploring the environment via random walk.
            - In a confined space, this leads to ants climbing over one another.
            - As hundreds-to-thousands of individuals engage in this activity, a "tower of constant strength" naturally forms, says Hu.
            - In lab experiments, each new layer of ants (i.e., a ring around the rod at the center of the tower) becomes established in four minutes, on average.
            - Following roughly ten minutes of tower growth, interior ants exit from the bottom of the tower and rejoin others moving on the tower's exterior surface.
        - Towers are not conical, but rather the slope of the tower is steeper toward the apex.
        - Impinging water droplets flow off the structure, rather than penetrating it.
        - The colony behaves as both a liquid and a solid, simultaneously, as individual ants are constantly in motion.
    - Hu hopes his research on fire ant colonies might inform the effective use of swarming robots.

What Shouldn't We Automate?—Dr. John Leslie King, University of Michigan; Robert Charette, ITABHI Corporation; and Tae Wan Kim, Carnegie Mellon University

- The question of whether to automate a process separates into two components: Can we? and Should we?
  - Even when the technological capability exists, proceeding with automation entails both an economic calculation and an ethical evaluation.
    - Even when an economic case exists for some type of automation, is it the right thing to do?
    - This is the root of the issues that Steven Cherry raises with members of the panel.
- Unique concerns that AI/machine-learning raise:
  - *King*: Concerns arise now that AI-based automation is penetrating cognitive domains that have been the sole purview of people.
  - *Kim*: "AI is an attempt to imitate human intelligence," says Kim, leading his MBA students to focus on the question of what it means to be human.
    - Biased training data jeopardizes the quality of a machine's intelligence (e.g., Microsoft's chatbot Tay evidenced misogynistic and racist behaviors; facial recognition systems trained largely on Caucasian faces perform poorly when presented with black faces).
- Unintended social consequences of AI systems:
  - *Charette*: Referencing Jon Snoddy's entertainment robots, Charette raises the issue of multiple objective functions.
    - The primary role of a cute audio-animatronics figure might be to engage with and entertain theme park visitors, but also be tasked with keeping lines moving or alerting security if observing bad behavior among visitors.
    - "We have to consider the unintended consequences of the technology as it is implemented," says Charette, a matter that can be compounded when additional technology is brought in to overcome deficiencies in existing technologies.
- Envisioning the eventual creation of general artificial intelligence, the matter arises of when robots/AI might get rights—and be assigned responsibilities:
  - *King*: Although there is little agreement about what defines humanity, rights are being granted to nonhumans, notably animals (e.g., experiments that are forbidden to perform on chimpanzees, but are fine for cockroaches).
  - *Kim*: Remaining sanguine about potential AI overreach, he says, "We can always unplug the AI."
    - Kim, reporting on the ability of an AI to learn human emotions, believes human-like AI might be closer than many think.
      - An eight-year-old CMU experiment tasked a deep-learning system with distinguishing fMRI traces associated with different emotional states, with data provided by emoting actors.
        - The system achieved 75% accuracy when evaluating the emotional state of off-the-street test subjects based on real-time fMRI scans.
        - Given that an AI can recognize (and therefor imitate) human emotion, can intelligence be far behind?
    - Yet Kim waxes philosophical, noting that Socrates touted the intellectual achievement associated with the knowledge of not knowing—something that AIs have yet to demonstrate.
      - A TTI/Vanguard participant offers a reminder of Aristotle's modes of persuasion—ethos (persuasion based on ethical arguments), pathos (emotional persuasion), and logos (logic-based persuasion)—noting that AIs struggle with logos and have little to offer in the way of ethos or pathos.
        - Any consideration of AI rights must also consider associated responsibilities.
  - *Charette*: From the perspective of risk, the question becomes what rights and responsibilities to confer on a "sentient" AI/robot.
    - Corporations are legal persons, but are governed by a set of rights/responsibilities distinct from those that apply to human beings.
    - Rights/responsibilities should be assigned on the basis of a risk analysis, but the relevant parameters have yet to be defined.
    - Moreover, early consideration should be given to which aspects of technology require an ethics-oriented risk assessment.
    - For every new technological advance, risks should be overtly considered.
      - For instance, who should be held liable when an autonomous vehicle causes an accident? The manufacturer? The software designer? This is hardly an academic question.

        - *Kim*: Liability resides with the causal actor associated with the bad act.
        - *King*: Historically, assignment of responsibility comes down to risk assessment.
        - *Charette*: When bad things occur, society as a whole—or a responsible industry—covers costs, but proactive consideration is essential to not be caught without a preestablished contingency plan/fund.
            - In the case of self-driving cars, Charette reports that one in every three lines of code are for proactive error correction.
- Teaching machines ethics:
    - *Kim*: On the one hand, he has an initiative underway to translate ethical principles into a mathematical/machine-readable language; on the other hand, he is teaching ethics to people with the goal of building it into automation.
    - *King*: Organizations should treat ethical considerations as first-class entities that guide the progress of technological innovation, rather than be relegated to the back burner.
- Competitive advantage of risk taking:
    - *Charette*: Firms/countries that are willing to take bigger technological/ethical risks will advance more quickly, but will also be subject to the downsides of risk, both economically and societally.
    - *Kim*: For firms, ethics is not the number-one consideration; rather, progress/profit is.
        - However, Millennials take ethics seriously and are pushing firms to invest in ethics.
- The ethical effect of automation on jobs:
    - *Cherry*: Pessimistically, he worries that the U.S. has long put off a national conversation about the distribution of profits that a handful of autonomous-truck manufacturers are expected to enjoy to the detriment of hundreds of thousands of soon-to-be-unemployed truckers.
        - Would the ethical solution be for profits to be shared with today's truckers?
    - *Kim*: This problem will not be unique to the trucking industry: Economists predict that half of the people desiring to work in 30–40 years will not find available jobs.
        - Yes, new jobs will be created, but people will compete with robots for them.
- Speaking on behalf of the robots, Doug Lenat relates a Marvin Minsky story of a multistep punch card-based program:
    - The first operation was to run for 20 seconds.
    - The second operation was to request 30 additional seconds of run time from the operator.
    - The third operation was to request an additional minute, and so forth.
    - Eventually, the program had racked up a $15K computer time bill, but Minsky objected, declaring that, yes, the program asked for more time, but it was the human operator who kept granting it.

General-Purpose versus Bespoke Robots—Mr. Tim Enwall, Misty Robotics

- The bulk of robots currently deployed are dedicated to particular tasks, be it vacuuming the floor, welding car bodies, scanning grocery shelves for inventory purposes, aligning/fastening/sinking concrete mats onto the river bottom, or otherwise.
    - Want to do anything distinct from what is on the robot's preprogrammed menu? Good luck!
        - "You need a roboticist to solve a problem with a robot," says Enwall.
    - Task-specific robots have their place, but Enwell and Mistsy Robotics founder Ian Bernstein (of Sphero) have a different vision: democratizing the distribution and use of this class of technology.
        - Toward this end, Misty Robotics is selling a general-purpose robot that can be readily programmed by any competent software developer to perform desired tasks.
    - While describing Misty's hardware specs and extensibility, Enwell also explains the design trade-offs required to make an affordable, multipurpose, humanoid-like robot.
- As a sophisticated robotic platform for developers, Misty might be just 14" tall and six pounds, but "she" is replete with advanced components for her $2400 price tag (discountable by up to $400).
    - Front-facing head-mounted components: Occipital Structure Core depth sensor for mapping, 4K camera for face and object recognition, 4.3" LCD display (default display: eyes), and lights (flashlight and indicator LEDs).
    - Side- or rear-facing head-mounted components: array of three far-field microphones, six capacitive touch panels, and removable magnetized headpiece (for extensibility (e.g., affix a FLIR camera)).
    - Neck with three degrees of freedom.
    - Front- or side-facing body components: two hi-fi speakers, light (RGB LED), four front bump sensors, three front and four edge/downward time-of-flight sensors (for obstacle avoidance), and two detachable arms with sockets (for extensibility; CAD specs provided), and trailer hitch (for towing).

    - Rear-facing body components: USB and serial ports secreted behind a magnetically attached backpack (for extensibility), two rear bump sensors, and one rear time-of-flight sensor (for obstacle avoidance).
    - Tread-based drive system (for locomotion over smooth surfaces, as well as rugs, electrical cords).
    - Computation: Snapdragon 820 mobile processor (Windows IoT Core OS, chosen both for security and self-management of software updates) and Snapdragon 410 processor (Android 8 OS for navigation and computer vision); programming interfaces for Python, JavaScript, REST, with C# coming soon.
    - Connectivity: Wi-Fi and Bluetooth.
    - Power: battery recharges via inductive charging when paired with the charging pad.
- All told, Misty has sight (with programmed-in human face detection), speech, autonomy, and hearing, but no native capability for climbing, manipulation, or lifting.
    - Using the provided API, Misty can be trained to recognize specific faces or other objects of interest.
    - The provided arms can pivot, but are otherwise nonfunctional; remove them with a single screw and replace with bespoke 3-D printed arms to enable, say, a pincer grip, laser pointer, or other functionality, when augmented with power and control through the back-of-body ports.
    - Misty's potential task space is limited, but sufficiently interesting, believes Enwell, to make her useful and to serve as a good jumping-off point for follow-on versions of the robot.
        - He recognizes that Misty does not meet the mental model of *Star Wars'* C3PO and is therefore not the general-purpose robot of people's dreams.
- As with other general-purpose platforms (PCs, smartphones), the platform creators cannot do it all, but require an ecosystem of software and hardware add-ons for a thriving, evolving outcome.
    - As such, Misty Robotics faces cost-conscious trade-offs: what to include out of the box, what to support as add-ons as the ecosystem builds out, and what to put off for future hardware generations.
        - Autonomy—Choices were made regarding vision, sound, bump, and time-of-flight sensors to keep costs moderate while also enabling high-quality situational awareness.
            - Nevertheless, Misty's capabilities might prove insufficient for robust face- or language-recognition in chaotic environments like homes or offices.
        - Charging—Eschewing the choice of wired charging, it was deemed important for Misty to independently proceed to/from her charging station with no need for human intervention.
        - OS—The propriety choice of Windows IoT Core was determined to be approachable by more developers, compared to open-source Robot Operating System (ROS), which requires developers to have robotics-specific training.
        - Appearance—A small, cute robot is unthreatening; "If multipurpose robots are going to be successful, humanity is going to have to accept them. They can't be intimidating," says Enwell.
            - Misty's large-range-of-motion neck (patent-pending) enables her to make "eye contact" with people, whether standing, seated, or at Misty's eye level.
                - Bernstein collaborated with the Disney design team on Misty's welcoming appearance.
            - Still, Enwell worries that Misty's size will be judged too diminutive for constructive tasks.
        - Extensibility—In addition to third-party programmability of Misty's innate hardware components, her detachable arms, removable earpiece, removable backpack, and tow hitch make this platform rather flexible.
        - Price—Misty's price point is purposefully similar to that of a developer's laptop; compare with SoftBank Robotics' Pepper, which sets a purchaser back more than $10K.
        - Battery life—Depending on operating mode, Misty runs for 4–8 hours on a single charge; compare with (the much larger) Pepper, which can run for 12–18 hours.
        - Physical robustness—Misty is purposefully durable; she can survive being kicked.
- Anticipated applications with out-of-the-box capabilities:
    - Autonomous patrol—examples:
        - continuous visual scan data center resources, and send an alert when a component is in an anomalous state;
        - snap photos of retail store endcaps to establish compliance with vendor contracts;
        - provide in-home reminders to take medications as prescribed.
    - Security guard: Just a few hours of JavaScript development yielded a skill that tied together a select number of face–name combinations; when an unknown person comes within Misty's visual frame, she gives it the opportunity to identify itself and, if unsuccessful, generates an English-language warning to the interloper while also snapping a photo and sending an intrusion alert.

- As developers expand the task-specific and reusable code base, Enwell anticipates the growth of a sharing economy or app store to extend capabilities across deployed Mistys.
  - "I am optimistic about the inventiveness of humanity," says Enwell.
  - "This isn't going to be able to do the 50 different things you might want out of the box," he says; it will take the developer community to make Misty interesting and viable.
  - Doug Lenat thinks sizing Misty akin to a baby might set appropriately low expectations.
- The Misty Robotics business model is to sell the robots and to encourage a sharing economy for developers, but not to take a cut of their software proceeds.

Microbots—Dr. Marc Miskin, University of Pennsylvania

- Pick up a vial of water and it looks clear and pristine, but it is teeming with microscopic life.
  - The intricacies and diversity of microorganisms have always fascinated Miskin, but instead of thrusting him toward a career in biology, he pursued physics and engineering with the goal of designing, making, and deploying robots on the same size scale as cellular structures.
  - His timing is right, with MEMS and Moore's Law enabling the fabrication of CMOS-based 3-D structures that not only compute or store data, but that can also serve as sensors, actuators, and output modes—with nanoscale components.
  - Miskin has successfully overcome three thorny challenges in his journey to build his microbots:
    - choosing appropriate actuators and 3-D structures (compare with previous microscale mechanisms, which have been quite primitive and do not rightly earn the moniker of *machine*, such as a corkscrew-shaped device that self-propels in the presence of a magnetic field, but is otherwise devoid of control or autonomy);
    - integrating such parts with the host of resources related to information technologies (the corkscrew swimmer is a standalone object with no interaction or communication capability).
    - releasing the devices from their substrate to set them free in the world.
  - Miskin describes his microrobots and the applications they might enable.
    - Not surprisingly, many uses are in the medical domain, for what better way to interface with cells than with a cell-sized robot?
      - Miskin notes that, although with little additional engineering effort, he could make his robots smaller—in fact, as much as two-orders-of-magnitude smaller—he has chosen to settle in at cellular scale.
        - At single-micron scale, objects succumb to the physical force of water molecule's impact: "You start diffusing randomly," says Miskin, "which for a robot is very annoying."
          - He offers the conjecture that evolution established cells at their observed size to be as small as possible, but to not be buffeted by their ambient environment.
  - "This robot that I built is about 100 μm on a side, it walks for locomotion, and it integrates a piece of silicon microelectronics—in this case, a sort of solar cell that it uses to power its legs," he says. "And, yes, my robots are about the same size as the ship in *Fantastic Voyage*."
- Miskin is not the first to downsize computation to cubic-millimeter proportions.
  - As Dennis Sylvester shared with TTI/Vanguard in San Francisco, December 2015, the Michigan micro mote has dimensions of just 300 μm, made possible in large part because of optimization for low power.
    - Still, while capable of sensing, communication, computation, and memory, the micromote lacks functional 3-D structures or components for actuation and locomotion.
  - "What we're really great at doing is making stuff for processing information at small scales," says Miskin, "but at the machines bucket, we're not so good"; he is changing this equation.
- Four requirements of a microrobot's actuators:
  - the moving parts of a microbot should bend but not break, despite their small scale;
  - the device should operate at low voltage and low power to smoothly integrate with conventional microelectronics;
  - the moving parts should be subjectable to high forces, to be able to reliably initiate movement;
  - the components individually, and the robot as a unit, must be mechanically and chemically robust to survive intact under its operational load within the environment in which it functions.
  - Miskin's actuator that simultaneously satisfies all of these needs is a 7-nm-thick strip of platinum (70 atoms, +/–1 atom) that, in its native form, curls increasingly tightly, corkscrew-like, in the presence of an electric field (200 mV, 10 nA), relaxing again in the absence of the field.

- "When you make something very thin, it is easy to bend it," says Miskin. And when you make it out of platinum, place it in water, and apply a voltage, electrochemistry generates forces on the surface of the strip, causing its deformation.
  - "When you are very, very small and just 7 nm thick, your surface area compared to your volume is astronomical, so these forces are really big," he says.
- Moreover, since the voltage relevant to platinum's electrochemistry matches that of semiconductor microelectronics, integration between the actuator and conventional chip fabrication comes for free.
- "This is a phenomenally good technology," says Miskin. "Oftentimes when you work at the nanoscale with nanomaterials it is easy to make one object, but hard to make millions of objects. But this just works: We have something like 96% yield" (a number he anticipates rising to in excess of 99% with professional fabrication, rather than his lab-based efforts).

- Integration with silicon microelectronics:
  - Using a 600-mV, 1-μA silicon photovoltaic laser as a power source, voltage and current align well with the actuator's need, with plenty of resources left over to power the little robot's computational components.
    - "You can now check all the boxes for an actuation technology that has all the mechanical components you need that is plug-and-play with microelectronic elements.
- 3-D structure and motion:
  - Differential bending is insufficient for a physically structured robot capable of performing specific tasks (like walking).
  - To define the microrobot's form factor, Miskin adopted a page from the work of Daniela Rus, who spoke to TTI/Vanguard in Boston, April 2014; that is, the use of rigid panels that dictate—origami-like—*where* the device's segments will/won't bend.
    - "What you don't want are bends; you want folds—specific regions that are going to bend where you tell them to," says Miskin.
    - Applying a strategic-overlay approach to the platinum strips integrated into the planar fabrication scheme on the CMOS wafer, Miskin can define how each part of a microbot can fold and therefore how the overall mechanical critter can move.
  - This is not to say that the details of combining the silicon microelectronics front-end with the actuator attachment back-end proceeded flawlessly from the get-go, but after two years of toil, Miskin established a reliable multistep fabrication procedure.
    - In the university cleanroom, his efforts eventually yielded a shiny four-inch wafer covered in 1M micro-scale robots in a trio of designs that vary by the number of actuator legs and the number of photovoltaics (in both cases, either two or four).
- "Release the robot army!"
  - To release the robots from the substrate requires etching away the surface, which Miskin shows with a real-time video of the critters breaking free from the 2-D plane and folding up into their intended structures.
    - "You can probe them with your laser, and off they go," he says, speaking of imagery of them crawling the length of microscope slides.
- Miskin is quick to point out that the scope of designs for cell-sized robotics is limited only by the imagination.
  - In addition to incorporating photovoltaics, his lab has also produced microbots with some combination of voltmeters, thermometers, and LEDs (as an output device).
    - If fabrication were upgraded to a high-end commercial foundry, it would be possible to incorporate nearly 1M transistors in the "area occupied by a paramecium," says Miskin, who has been working instead with 55-nm technology—in other words, phenomenal potential for sensing, I/O, computation, and memory, not to mention bespoke nanoactuators, depending on the application.
    - Doug Lenat suggests that building up the microelectronics of the robot in three dimensions by using multilayered crystalline-silicon, instead of CMOS, as the basis for fabrication; the small size of the robots would likely overcome the problem of heat dissipation that would otherwise make such an approach untenable.
  - To consider what those applications might be, Miskin first offers the reminder that the cost–benefit trade-offs for diminutive robots differs markedly for that of their macroscopic cousins in aspects beyond heat management.

- Power is plentiful (“Basically, walking is free,” he says), cost per robot is negligible (0.1¢/robot), assembly is streamlined and fully automated, and they stand up to extreme environments (e.g., strong acid/base, or 400°C).
  - “You have a monetary incentive for miniaturization: Your cost per machine goes down,” says Miskin. “Everything gets easier, provided you can build the thing in the first place.”
- Not only is the scope of designs effectively unlimited, so too is the sheer volume of robots that could be manufactured, since production would use the existing infrastructure for microelectronics fabrication.
  - “We are not maxed out in any capacity for these robotic systems,” says Miskin.
    - Thinner components, smaller folds, more complex structures, denser electronics, the addition of memory, heavier payloads (up to 30x the robot’s body weight; up to 100 μm thick), bio-integration—all of these and more are possible.
- To date, each capability of Miskin’s microrobots is designed in at the point of fabrication, but the size scale also permits programmability: “The rough estimate is that, in the area of one of these robots, you could build a four-bit microcontroller that has about 1000 instructions on it,” he says. “The question that we have is whether or not you have to do that. For the applications that we have considered, there is a good argument for hard-coding, but intellectually, we should definitely make a programmable one.”

- So, what might those applications look like? As is the case with robots of any size, applications fall into two broad categories:
  - Telemetry, in which the robot travels to an otherwise inaccessible location and sends back information (e.g., Mars rovers):
    - Potential medical examples for microrobots include cellular-scale chemical sensing of disease biomarkers, or physical sensing of cell stiffness to evaluate for cancer.
  - Autonomy, in which the robot performs a useful task without human intervention (e.g., Roomba):
    - Potential medical examples for microrobots include highly localized drug delivery, or crawling into tissue to make measurements.
- Neural interfacing spans both of these.
  - Today’s state-of-the-art approaches to brain sensing and stimulation share the characteristic of being pointy and physically invasive—“They rely on being shoved through the skull,” says Miskin.
  - His lab is in the early stages of tailoring microbots to interface with individual neurons to measure neural activity.
  - Components include a voltmeter for sensing, photovoltaic LEDs for power and communication, and a robotic exoskeleton “to hunt down nearby neurons and get as close as possible to maximize the signal transmission, and beam that data back to you in light.”
    - Future work would presumably include real-time stimulation of relevant neurons.
- As noted, Miskin has, to date, been supplying power to his robots with light, but this only works with line-of-sight access; translated into the medical domain, only millimeter-scale penetration into the body would be possible.
  - However, commonplace medical technologies suggest other modes of power delivery: ultrasound can reach 1 mm–10 cm, and magnetic fields penetrate even deeper—deep enough to reach any bodily tissue.
    - “As we develop this, we think realistically that we can supply enough power to these robots so that they could operate virtually anywhere,” says Miskin.
- Circling back to the first sentence of this report, a vial of water teeming with Miskin’s robots appears no different than any other vial of water.
  - The robots are only visually apparent under a microscope.
  - This leads Miskin to offer an analogy: “Since they are aqueous and in a solution, from the perspective of the human, they are identical to chemicals. You can mix them, pour them, spray them—there is no difference whatsoever.”
  - The consequences of this mode of thinking is to consider microbots as active reagents, capable of sensing a chemical reaction and then driving a reaction in response, while also being able to actively navigate to targeted regions where the robot’s presence and actions would be therapeutic or otherwise desirable.
  - Alternatively, it would be possible to post-process manufactured materials of all kinds with a coating or embedding of microrobots to make those materials smart.
  - “Unfortunately, when I pipette robots, they all land on their backs. They don’t walk after injection,” says Miskin. “I have yet to solve that”—but presumably he will.

    - Another perhaps-temporary disappointment for Miskin is that, as much as he looked forward to witnessing bacteria and his robots go head to head, he has yet to see a microorganism or white blood cell attempt to attack one of his critters.
  - A nonmedical application that taps into this mode of thinking is a microrobot that would pare back the dendrites that inevitably grow on the anodes of lithium–ion batteries.
    - Miskin is actively pursuing this research with UPenn colleagues James Pikul and Mark Yim: “We want to build a robot that lives on the lithium anode. This robot would be powered by the battery itself,” says Miskin. “It would have circuitry onboard to hunt down the dendrites by looking for field gradients, and would use chemical actuators that would allow it to etch the lithium back into solution. You wouldn’t have to rethink how you make batteries; this would simply be an additive that you stick in at the last step of manufacturing.”
- Miskin’s bold conclusion: “We are entering an age where the micro world is no longer the privileged domain of biology.”
  - With his robots, human-designed control, autonomy, and precision become possible at the fundamental scale of biology.
  - What he does not seek to do is enable in his microrobots is self-reproduction.
    - When meddling with biological equivalents, the ability to confine them is essential, so he is adamant about retaining the power to turn off the light, turn off the ultrasound, or turn off magnetic field and know that the robots will not rise again.