

A Roadmap for US Robotics – From Internet to Robotics 2020 Edition

Suggested Citation: H. I. Christensen, N. Amato, H. Yanco, M. Mataric, H. Choset, A. Drobni, K. Goldberg, J. Grizzle, G. Hager, J. Hollerbach, S. Hutchinson, V. Krovi, D. Lee, W. Smart and J. Trinkle (2021), “A Roadmap for US Robotics – From Internet to Robotics 2020 Edition”, Foundations and Trends® in Robotics: Vol. 8, No. 4, pp 307–424. DOI: 10.1561/23000000066.

H. I. Christensen

UC San Diego

N. Amato

Univ. of Illinois Urbana Champaign

H. Yanco

Univ. of Mass. - Lowell

M. Mataric

Univ. of South. California

H. Choset

Carnegie Mellon Univ.

A. Drobni

Computing Research Assoc.

K. Goldberg

UC Berkeley

J. Grizzle

Univ. of Michigan

G. Hager

Johns Hopkins Univ.

J. Hollerbach

Univ. of Utah

S. Hutchinson

Georgia Inst. of Tech.

V. Krovi

Clemson Univ.

D. Lee

Cornell Univ.

W. Smart

Oregon State Univ.

J. Trinkle

Lehigh Univ.

This article may be used only for the purpose of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval.

now

the essence of knowledge

Boston — Delft

Contents

1	Executive Summary	308
1.1	Introduction	308
1.2	COVID-19	309
1.3	Main Findings	312
1.4	The Roadmap Document	314
2	Societal Drivers	315
2.1	Manufacturing	315
2.2	Logistics and E-Commerce	319
2.3	Transportation	323
2.4	Quality of Life	326
2.5	Clinical Healthcare	330
2.6	Feeding the planet	333
2.7	Security and Rescue Robotics	335
3	Mapping Societal Drivers to Research Challenges	340
3.1	Identifying challenges to growth/progress	340
3.2	Mapping challenges to research needs	347
4	Research Challenges	353
4.1	Architectures and Design Realizations	353
4.2	Locomotion	358
4.3	Grasping and Manipulation	363
4.4	Perception	367

4.5	Planning and Control	370
4.6	Learning and Adaptive Systems	375
4.7	Multi-Robot Systems	378
4.8	Human-Robot Interaction	382
5	Technology Context	390
5.1	Additive Manufacturing	391
5.2	Model Based Programming	391
5.3	Configuration Lifetime Management	392
5.4	Collaborative Systems	394
5.5	Internet of Things / Industry 4.0	394
5.6	Big Data and Analytics	396
6	Workforce Development	399
6.1	Introduction	399
6.2	Strategic Findings	400
6.3	Near Term Opportunities and Factors Affecting Immediate Deployments	402
7	Shared Infrastructure in Robotics	404
7.1	Flexible Research Platforms	405
7.2	Community Consensus Validation Benchmark Frameworks	405
7.3	Reference Open-Access Testbeds	406
8	Legal, Ethical, and Economic Context	408
8.1	Safety	409
8.2	Liability	410
8.3	Impact on Labor	411
8.4	Social Interaction	412
8.5	Privacy and Security	413
8.6	Recommendations	414
9	Contributors	415
	Acknowledgements	418
	References	419

A Roadmap for US Robotics – From Internet to Robotics 2020 Edition

H. I. Christensen, N. Amato, H. Yanco, M. Mataric, H. Choset, A. Drobnis, K. Goldberg, J. Grizzle, G. Hager, J. Hollerbach, S. Hutchinson, V. Krovi, D. Lee, W. Smart and J. Trinkle

ABSTRACT

Recently, the robotics industry celebrated its 60-year anniversary. We have used robots for more than six decades to empower people to do things that are typically dirty, dull and/or dangerous. The industry has progressed significantly over the period from basic mechanical assist systems to fully autonomous cars, environmental monitoring and exploration of outer space. We have seen tremendous adoption of IT technology in our daily lives for a diverse set of support tasks. Through use of robots we are starting to see a new revolution, as we not only will have IT support from tablets, phones, computers but also systems that can physically interact with the world and assist with daily tasks, work, and leisure activities. The present document is a summary of the main societal opportunities identified, the associated challenges to deliver desired solutions and a presentation of efforts to be undertaken to ensure that US will continue to be a leader in robotics both in terms of research innovation, adoption of the latest technology, and adoption of appropriate policy frameworks that ensure that the technology is utilized in a responsible fashion.

H. I. Christensen, N. Amato, H. Yanco, M. Mataric, H. Choset, A. Drobnis, K. Goldberg, J. Grizzle, G. Hager, J. Hollerbach, S. Hutchinson, V. Krovi, D. Lee, W. Smart and J. Trinkle (2021), “A Roadmap for US Robotics – From Internet to Robotics 2020 Edition”, Foundations and Trends® in Robotics: Vol. 8, No. 4, pp 307–424. DOI: 10.1561/23000000066.

1

Executive Summary

1.1 Introduction

Recently the robotics industry celebrated its 60-year anniversary. We have used robots for more than six decades to empower people to do things that are typically dirty, dull and/or dangerous. The industry has progressed significantly over the period from basic mechanical assist systems to fully autonomous cars, environmental monitoring and exploration of outer space. We have seen tremendous adoption of IT technology in our daily lives for a diverse set of support tasks. Through use of robots we are starting to see a new revolution, as we not only will have IT support from tablets, phones, computers but also systems that can physically interact with the world and assist with daily tasks, work, and leisure activities.

The “old” robot systems were largely mechanical support systems. Through the gradual availability of inexpensive computing, user interfaces, and sensors it is possible to build robot systems that were difficult to imagine before. The confluence of technologies is enabling a revolution in use and adoption of robot technologies for all aspects of daily life.

Thirteen years ago, the process to formulate a roadmap was initiated at the Robotics Science and Systems (RSS) conference in Atlanta. Through support from the Computing Community Consortium (CCC) a roadmap was produced by a group of 120 people from industry and academia. The roadmap was presented to the congressional caucus and government agencies by May 2009. This in turn resulted in the creation of the National Robotics Initiative (NRI), which has been an interagency effort led by the National Science Foundation. The NRI was launched 2011 and had its ten-year anniversary. The roadmap has been updated 2013 and 2016 prior to this update.

Over the last few years we have seen tremendous progress on robot technology across manufacturing, healthcare applications, autonomous cars and unmanned aerial vehicles, but also major progress on core technologies such as sensors, communication systems, displays and basic computing. All this combined motivates an update of the roadmap. With the support of the Computing Community Consortium three workshops took place 11-12 September 2019 in Chicago, IL, 17-18 October 2019 in Los Angeles, CA and 15-16 November 2019 in Lowell, MA. The input from the workshops was coordinated and synthesized at a workshop in San Diego, CA February 2020. In total the workshops involved 79 people from industry, academia, and research institutes. The 2016 roadmap was reviewed, and progress was assessed as a basis for formulation of updates to the roadmap.

The present document is a summary of the main societal opportunities identified, the associated challenges to deliver desired solutions and a presentation of efforts to be undertaken to ensure that US will continue to be a leader in robotics both in terms of research innovation, adoption of the latest technology, and adoption of appropriate policy frameworks that ensure that the technology is utilized in a responsible fashion.

1.2 COVID-19

Over the last few months we have seen some major changes to society. The COVID-19 (Center for Disease Control and Prevention (CDC), [2020](#)), or the more accurate name for the infection Sars-CoV-2, has

changed many things. It has already infected more than 27 million people with 6+ million of them in the US alone (by September 2020) (Worldometer, 2020).

The outbreak of the pandemic has had a number of effects. First of all, the healthcare system has been challenged. People have also been quarantined at home for extended periods of time. A large number of people have been laid off in USA (and globally). In addition, people have almost stopped traveling. An obvious question is how robotics and automation can assist in such a scenario.

In the healthcare sector there are quite a few obvious use-cases. i) there is a need to increase the frequency of testing people to get a nuanced view of the degree of infection and the speed of infections (R_0). Laboratory robots allow for faster processing of samples and return of answers to people. Laboratory robots can automate the testing and allow for extensive testing. Many healthcare professionals have been exposed to COVID due to their front-line jobs. There is a real need to use automation to acquire samples from patients, but also to enable a doctor at a distance to examine a patient and acquire basic information such as temperature, blood pressure, pulse, etc. Using tele-presence robots, it is possible to increase the social distancing between patients and medical personnel for routine tasks and through this reduce the risk of exposure for professionals. There are numerous use-cases for medical robots beyond the well-known examples in surgery.

Manufacturing has declined significantly during COVID-19, which is partly due to changes in market needs, but also due to the economic recession gaining momentum after the start of the pandemic. Total industrial production is seeing a downturn. We have seen automotive sales go down by as much as 50% (Federal Research Economic Data, 2020). When isolated at home the traffic patterns change dramatically. Retail sales was down by 20+% during September 2020 and food/drink sales were down by 50% in September 2020. At the same time e-commerce continued to have significant growth. Sales of goods in the traditional retail sector was shifting from brick-and-mortar shops to the web.

E-Commerce has seen tremendous growth over the last year. The growth is both in US with major companies such as Amazon and Walmart, but also internationally by companies such as Alibaba, JD

and Tmall. Already today Alibaba with Taobao is 50% larger than Amazon and is expected to continue to grow. Amazon has deployed more than 200,000 mobile platforms in their warehouses (the number is more like 300,000 by now) (Tech Crunch, 2020). In addition, we are also seeing major progress on automated object pick-up / handling with companies such as Covariant.AI, Righthand Robotics and Berkshire Grey. As people desire a minimum of contact for items entering their house, we will see higher automation at distribution centers. There is significant interest in the last-mile problem of delivering from the truck to the front door in a domestic setting. The last mile could be solved using a traditional mobile platform as seen by Amazon's Scout (Amazon, 2020) another solution is clearly humanoid robots such as digit by Agility Robotics (Agility Robotics, 2020) or traditional services such as May Mobility (Maymobility Corp, 2020). Leaving the ground for a minute the drone market is considered for last mile deliveries as seen by Amazon Prime Air (Wikipedia, 2020) or the experiments by UPS (Drone Life, 2020).

Cleaning is another important topic. This includes cleaning and disinfection beyond the hospital and the home. iRobot has seen a major uptick in sales of vacuum cleaners and floor scrubbers during the pandemic and shares are up 65% year to date. Additional cleaning is important to many households. A flood of UV-C disinfection robots has also been announced. Using UV-C lighting it is possible to achieve a high degree of disinfection with more than 99.9% of the virus eliminated when more than 10 micro watt / cm² is radiated onto a surface. In many cases, a high-power source is used to allow even indirect illumination to kill the virus. There are already more than 100 companies worldwide pursuing this market. Keenon has developed a robot that uses both UV-C lighting and a vaporizer to disinfect an area. The vapor will get to areas that may not be directly exposed by the UV-C light and provide redundant security. These two robots are merely examples of the vast number of new robots entering this market. The first place to see deployment of these UV-C robots were hospitals and care facilities. High-traffic use-cases such as airports have also seen deployments. One would expect other use-cases to include hotels, malls, cruise ships, and eventually they may enter your house as supercharged home cleaning

robots. This is a new robotics segment that was unrealistic just a few months ago.

COVID has exposed a number of opportunities for robotics from cleaning/disinfection over e-commerce to manufacturing and transportation. Robots are primarily designed to empower people to do things better, in some cases in terms of accuracy in other cases as power or sensory extensions, and access. In the aftermath of the 2009 recession adoption of robotics grew significantly. In a post-COVID world we will see new behavior patterns for social interaction, cleaning, collaboration and delivery. There are thus many new opportunities for utilization of robot technology to enhance many of everyday life.

1.3 Main Findings

Over the last decade a tremendous growth in utilization of robots has been experienced. Manufacturing has in particular been impacted by the growth in collaborative robots. There is no longer a need for physical barriers between robots and humans on the factory floor. This reduces the cost of deploying robots. In the US the industrial robotics market has grown 10+% every year and the market has so far seen less than 10% penetration. We are thus far away for full automation of our factories. US is today using more robots than it has even done before.

A major growth area over the last decade has been in use of sensor technology to control robots. More digital cameras have been sold the last decade than ever before. When combined with advanced computing and machine learning methods it becomes possible to provide robust and more flexible control of robot systems.

A major limitation in the adoption of robot manipulation systems is lack of access to flexible gripping mechanisms that allow not only pick up but also dexterous manipulation of everyday objects. There is a need for new research on materials, integrated sensors and planning / control methods to allow us to get closer to the dexterity of a young child.

Not only manufacturing but also logistics is seeing major growth. E-commerce is seeing annual growth rates in excess of 40% with new methods such as Amazon Express, Uber Food, . . . these new commerce

models all drive adoption of technology. Most recently we have seen UPS experiment with use of Unmanned Vehicles for last mile package delivery. For handling of the millions of different everyday objects there is a need of have robust manipulation and grasping technologies but also flexible delivery mechanisms using mobility platforms that may drive as fast as 30 mph inside warehouses. For these applications there is a need for new R&D in multi-robot coordination, robust computer vision for recognition and modeling and system level optimization.

Other professional services such as cleaning in offices and shops is slowly picking up, this is in particular true given the recent COVID-19 pandemic. The layout of stores is still very complex and difficult to handle for robots. Basic navigation methods are in place, but it is a major challenge to build systems that have robust long-term autonomy with no or minimal human intervention. Most of these professional systems still have poor interfaces for use by non-expert operators.

For the home market the big sales item has been vacuum and floor cleaners. Only now are we starting to see the introduction of home companion robots. This includes basic tasks such as delivery services for people with reduced mobility to educational support for children. A major wave of companion robots is about to enter the market. Almost all these systems have a rather limited set of tasks they can perform. If we are to provide adequate support for children to get true education support or for elderly people to live independently in their home there is a need for a leap in performance in terms of situational awareness, robustness and types of services offered.

A new generation of autonomous systems are also emerging for driving, flying, underwater and space usage. For autonomous driving it is important to recognize that human drivers have a performance of 100 million miles driven between fatal accidents. It is far from trivial to design autonomous systems that have a similar performance. For aerial systems the integration into civilian airspace is far from trivial but does offer a large number of opportunities to optimize airfreight, environmental monitoring, etc. For space exploration it is within reach to land on asteroids as they pass by earth or for sample retrieval from far away planets. For many of these tasks the core challenge is the flexible integration with human operators and collaborators.

The emergence of new industrial standards as for example seen with Industry 4.0 and the Industrial Internet facilitates access to cheap and pervasive communication mechanisms that allow for new architectures for distributed computing and intelligent systems. The Internet of Things movement will facilitate the introduction of increased intelligence and sensing into most robot systems and we will see a significant improvement in user experience. The design of these complex systems to be robust, scalable, and interoperable is far from trivial and there is a new for new methods for systems design and implementation from macroscopic to basic behavior.

As we see new systems introduced into our daily lives for domestic and professional use it is essential that we also consider the training of the workforce to ensure efficient utilization of these new technologies. The workforce training has to happen at all levels from K-12 over trade schools to our colleges. Such training cannot only be education at the college level. The training is not only for young people but must include the broader society. It is fundamental that these new technologies must be available to everyone.

Finally, there is a need to consider how we ensure that adequate policy frameworks are in place to allow US to be at the forefront of the design and deployment of these new technologies but it never be at the risk of safety for people in their homes and as part of their daily lives.

1.4 The Roadmap Document

The roadmap document contains sections specific to societal drivers, mapping these drivers to main challenges to progress and the research needed to address these. Sections are also devoted to workforce development and legal, ethical and economic context of utilization of these technologies. Finally, a section discusses the value of access to major shared infrastructure to facilitate empirical research in robotics.

2

Societal Drivers

2.1 Manufacturing

Manufacturing, from handicraft to high tech, is the staged transformation of raw materials into finished goods on a large scale using human-labor, machines, tools, chemical or biological processing on a large scale.

Manufacturing output accounts for some \$2 trillion in the United States. It represents about 12% of the GDP. Every dollar worth of manufacturing goods generates \$1.4 in output in other sectors of the economy. The U.S. is second only to China in Manufacturing (Kotkin, 2018).

Today, U.S. manufacturing companies face the twin challenges of an aging population and a shortage of skilled workers. As a result, our manufacturing competitiveness is at risk.

Robots keep U.S. manufacturing competitive by allowing them to improve product quality, increase productivity, get products to market faster and lowering the overall costs. As a result, manufacturing jobs are growing as more robots are adopted in the U.S. Since 2010, some 180,000 robots have been shipped to U.S. companies during the same period 1.2 million new manufacturing jobs have been created. At the same time,

robots are making the workplace safer by performing dangerous tasks that people should not be doing.

This allows people to do higher value, higher-paying tasks. Because of robotics, U.S. companies are now bringing some manufacturing jobs back to the U.S. According to the Reshoring Institute, about 78,000 jobs have been returned since 2010. Figure 2.1 shows the relation between robot sales and employment, with a strong positive correlation between sales and employment.




Figure 2.1: Relation between robot sales and employment in US (data from FRED and IFR World Robotics 2019)

Perhaps more importantly, manufacturing jobs that might have been outsourced to take advantage of low-cost labor are now being performed in the U.S. In addition to growing manufacturing jobs and output, these efforts also help revitalize communities that were hard hit by job losses when U.S. factories were closed.

U.S. competitors recognize that adopting robots is critical to manufacturing success. China, the E.U., Japan and Korea are “all in” on robotics with well-established government-funded programs to ensure they remain leaders in the use and development of robotics technology.

The production-line, a key innovation of the industrial revolution, set the stage for the development of the modern deployments but is in great need of an overhaul to accommodate dramatic paradigm shifts/megatrends in manufacturing: including mass production while

permitting customization in lot-sizes-of-one, digitalization (digitally-enabled insight into a traditionally opaque analog world), cloud-manufacturing systems, and need for scaling up production of highly integrated smart intelligent consumer products.

Worthy of note, manufacturing operations are increasingly becoming lean with just-in-time supply-chain and logistics operations in order to keep them economically feasible. At the macro-level autonomous transportation (transportation) promises revolutionary improvements in speed, efficiency, safety and reliability along with concomitant benefits for society and economy.

Inasmuch, it is useful to view a manufacturing shop floor from the lens of a “microcosm of a smart city”. Success and productivity depend upon synchronized orchestration of humans and automation which can occur at various spatio-temporal scales. There is a significant need for movement of people and materials between multiple physical locations – in the past, this was accomplished by high-cost and inflexible fixed-automation (conveyor-belts etc.) with implicit lock-in once selected.

Over the past decades, fixed infrastructure deployments (robots in cages) have made way for emerging classes of robots (e.g., mobile manipulators) and human-robot collaboration in shared spaces. In as much, the modern production-floor now offers an interesting sandbox to examine: alternate methods of realizing production (flexible automation) coupled with alternate provisioning of ancillary support between fixed (production-line), flexible (mobile robotic agents) and built infrastructure (WIFI, localization beacons).

Industrial robotics grew in deployments building upon a *general-purpose manipulator* capable of being *reprogrammed flexibly for multiple tasks*. While the former aspect is well-exercised, current deployments do not fully exploit the re-programmability (due to a variety of reasons including complexity). Nevertheless, sales for traditional industrial robots has grown at a CAGR of 19% from 2013 to 2019 even just in well-understood manufacturing use-case settings. In 2018, global robot installations increased by 6% to 422,271 units worth USD 16.5 billion bringing the operational stock of robots to about 2.44M units (+15%). With software, peripherals, and systems engineering included, the value is approximately 50 billion USD. For the eighth year in a row, robot

installations in the United States reached a new peak level (40,373 units; +22%) but still remains in 3rd place after China and Japan. The sales across regions during the last decade is also shown in Figure 2.2.




Figure 2.2: Number of robots shipped/sold per region (America, Asia and EU) over the period 2007 - 2018. (Source: IFR World Robotics 2019).

Newer paradigms such as collaborative robots - also called cobots - (designed to work together with humans) accounted for less than 14,000 out of more than 422,000 industrial robots installed in 2018. Despite strong media attention of cobots, the number of units installed is still low with a share of 3.24% of annual installation. Their growth rate was slightly higher (23%) as compared to traditional robots for reasons including the lack-of-awareness, change-management and lack of effective technology use-case performance or business ROI evaluations.

Nevertheless, there is both considerable excitement and trepidation about the latent potential of next-generation robotics (enable shorter production runs, smaller factories, and higher productivity) to transform production-systems and its ability to power growth around

the world (Atkinson, 2019). AI-enhanced robotics (e.g., with better machine vision) with other technological advances (better sensors/compute/actuation), promises to see significantly improved pricing and performance over the next decade.

The “Advanced Manufacturing Partnership” (AMP) recognized robotics as a key-transformative technology that can revolutionize manufacturing and embodied/deployed via the Manufacturing USA institutes. These Manufacturing USA institutes including DMDII and ARM have sought to build out an ecosystem of industry-SMEs-academia-government constituents (~200-400 members) to develop 3–4 year horizon technology roadmaps, updated annually and deploy them in technology investment strategies. However, due to the focus on higher TRL-levels, the opportunities for cross-pollination and translation of latent unrealized potential of approaches developed in other application spaces are not being fully realized in the manufacturing setting.

Greater national-level coordination is needed to capture the productivity and competitiveness benefits of robotics driven by: (i) Shortage of labor in key high-tech manufacturing sectors; (ii) need to compensate for the deficit in manpower by improving workforce productivity; (iii) gain a technological-multiplier to maintain leadership in a more competitive export market; while (iv) offsetting effects of national-level technology-investment efforts across the globe. Targeted national-scale investments in translating early-stage R&D efforts in robotics and automation into key manufacturing sectors of national interest – aerospace, apparel, electronics, machining, and automotive – would create significant opportunities for productivity gains.

2.2 Logistics and E-Commerce

According to 2019 figures released by the US Department of Commerce, total US retail sales are \$1.36 trillion, of which \$150B (10.7%) is e-commerce. E-commerce offers an unprecedented inventory of products at competitive prices to customers throughout the U.S. Customers expect their orders to be delivered promptly and reliably. E-commerce has doubled since 2015 and is growing by 12–20% annually; it is expected to accelerate in the next 5 years with increasing adoption for food and pharmaceuticals (Hadad, 2017).

New warehouses are being built, but it is extremely difficult to hire and retain human workers who can keep up with the pace of packing an increasing volume of diverse orders. Combined with similar demands in the upstream logistics of wholesale shipping, handling, sorting, storage and retrieval throughout the supply chain drives a pressing US need for robots and automation. “The broader market for warehouse and logistics automation topped \$53 billion last year and is forecast to exceed \$80 billion in 2023”, said Jeremie Capron, head of research at ROBO Global LLC, a research and investment advisory firm (Smith, 2020)

There is a clear need to have a comprehensive view of the supply chain for logistics and e-commerce, The Autonomous Supply Chain includes smart manufacturing, distribution and fulfillment centers, vehicles, and people whose primary task is to deliver products to consumers.

There is clear evidence that products and systems are being developed that support a 10-year vision of a Highly Autonomous Supply Chain. A 20-year vision may support a Fully Autonomous Supply Chain. This is supported by the investment and early emergence of: Autonomous (A) Air Cargo, A Supply Cargo ships, A tractor trucks, A delivery trucks, vans, cars, A flying drones and A ground-based drones. In addition, Industry 4.0, the SMART City initiative, Autonomous Manufacturing, Additive Micro Factory technology and Cyber Agriculture are changing the placement and speed of good and supplies to consumers. The integrated eco-system is exemplified in Figure 2.3.

Reshoring, population densification, sustainability and shortages in general labor and skilled labor are driving a need to transform how we think about supply chains. Industry is already facing shortages in general labor, labor retention, asset utilization, human labor safety injuries and high energy costs.

The perceived benefits to adding autonomy to supply chain activities include:

- Speed – products are delivered faster
- Safety
- Cost savings
- Demand satisfaction

- Address labor shortage concerns
- Sustainability
- New job opportunities
- Fuel and energy savings – autonomous vehicles have been shown to demonstrate a 15% fuel gain




Figure 2.3: The Eco-System for Logistics and Material Handling (drawing from UML workshop)

Robots can help if they can achieve Universal Picking: the ability to grasp any object in a huge and diverse range of shapes and sizes. Despite over 40 years of research, this problem remains unsolved. The difficulty stems from the inherent uncertainty in physics, perception, and control. Sensor noise and occlusions obscure the exact shape and position of objects in the environment, and object properties such as the center of mass and friction cannot be observed directly. Many groups are experimenting with new gripper and suction hardware, but still missing are algorithms and software that can rapidly compute robot grasp positions robust to uncertainty without requiring painstaking engineering expertise.

Existing software approaches to Universal Picking can be categorized as either analytic or empirical (Bohg *et al.*, 2014). Analytic methods rely on precise data about object properties and can work for uniform

picking of the identical objects from a bin, these methods cannot scale to many diverse objects. Empirical methods, on the other hand, learn to compute grasps based on data from physical experiments. This data is used to train function approximators such as deep neural networks, which generalize well to new objects. However, physical experiments require months to perform and are specific to one robot, gripper, camera, and set of objects.

A number of approaches to robot picking have emerged such as Right Hand Robotics, Berkshire Grey, SoftRobotics, Kindred, Osaro, Covariant and Kinema Systems (Hodson, 2018). Many study uniform picking, where all objects are identical in shape, which is distinct from universal picking, the challenge of grasping unknown and varying objects that is required for warehouse order fulfillment. Approaches can be considered in four categories – learning from demonstration, reinforcement learning, 3D registration, or hardware-centric methods. Learning from demonstration and reinforcement learning both require substantial data collection and data cleaning for each new environment or task. 3D registration methods require a-priori 3D object models and generally only work when the object is known explicitly. Hardware-centric methods explore novel gripper models such as pneumatics that work well with certain object classes.

One challenge is to reduce the time required for each sensing, computing, and robot motion cycle to match or exceed human performance of 500 pick attempts per hour (7.2 seconds per cycle). To increase both reliability and range, Continuous Learning might help distributing dataset generation and learning across virtual instances in the Cloud.

Commerce is also driving demand for robotics for last-mile delivery. A McKinsey & Co. report estimated that last-mile delivery costs amount to more than \$86 billion per year. There is also an increasing need for robots for retail front-of-store operations like monitoring and re-stocking shelves. Amazon Go stores opened in several markets using advances in robot vision to eliminate checkout lines by automatically sensing as each customer removes products in real-time.

2.3 Transportation

Efficient and safe transportation systems are a critical need of our society, enabling frictionless movements of both people and goods across town and urban centers as well as along long-distance interstate routes. New forms of transportation in the US have continuously evolved, from the horse and buggy, to locomotive, and then to automobiles. As transportation systems have changed over the years to address critical societal and economic needs, they have become plagued with congestion and still suffer from accidents, leading to loss of both time and lives.¹

Can we imagine a future where people and goods are able to move about with revolutionary improvements in speed, efficiency, safety and reliability? The use of novel robotic technologies holds the promise of providing for future advances in transportation systems. These positive changes will lead to numerous benefits across a number of sectors, including public safety, land use, supply chains, logistics, manufacturing, and quality of life.

2.3.1 Current Trends

Transportation systems connect us to our homes, our work, and our families and friends. As our population grows and changes, our transportation needs will also change. It is anticipated that over the next 25 years, the population of the US will grow to 390 million people, an increase of 70 million representing more people than the current populations of Texas, New York and Florida combined. Unfortunately, our transportation system has not kept up with this increase in demand. The capacity of our roads and airports is limited, leading to record levels of traffic congestion and frequent delays in aviation.

Similarly, congestion and inefficiency in our freight system lead to enormous economic costs. A robust multimodal freight transportation system is needed to meet the expectations of consumers and industry and for the nation to compete in the global market. The recent rise of online shopping services has drastically driven up demand for small package

¹For example, the average speed of transportation dropped in London from horse and buggy in 1916 (17 mph) to automobiles in 2016 (11.8 mph).

home delivery of food, clothing, electronics and other consumer goods. At current rates, online purchases requiring fast package deliveries to homes are predicted to represent over 20% of all retail purchases by 2025. This will result in a critical need to handle the transportation and delivery of an enormous number of packages in a rapid, efficient, and sustainable manner. Figure 2.4 exemplifies how both road-users and infrastructure are getting connected to provide dynamic situation awareness.



Figure 2.4: Example of how interconnected transportation systems will soon become omni-present

2.3.2 Personal Transportation

Typical households in the US today exhibit mobility patterns based upon the usage of a small number of owned vehicles. Examples include transporting children to school and commuting to work on a daily basis. Future autonomous vehicle transport could drastically improve mobility for children and for elderly and handicapped persons who are currently dependent upon human assistance for their access to transportation services. Robotic technologies that will drive the future development of near-autonomous and autonomous vehicles include better sensing and perception, especially under bad weather and hazardous conditions. The problem of handover from autonomous control to human operation

is a problem of human-robot interaction, sharing situational awareness of the road between robot and human operators.

The use of connected communication systems including vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), and with other entities (V2X), provides an opportunity to enable better transportation services by integrating information across larger scales than is currently possible. By augmenting the local sensors onboard individual vehicles with knowledge about road and traffic conditions beyond the line of sight, future connected vehicles will be endowed with the ability to navigate dense and congested areas safely and more efficiently. Critical robotic technologies for connected vehicles include better mapping capabilities and multi-agent planning and coordination techniques.

Robotic technology has the potential to impact and transform public transit systems. Traditional modes of public transportation (buses, trains, subways, light rail, etc.) will be enhanced with autonomous technology and also augmented with shared autonomous vehicles for shorter trips. Technology to better measure traffic patterns and predict demand will be used to optimize dynamic routes and schedules. There is also a critical need to provide safety guarantees from real-time sensors and controllers. Development of large-scale remote presence systems are also needed for efficient monitoring and to provide for rapid response in emergency situations.

2.3.3 Emerging personal mobility systems

Innovative forms of personal mobility platforms have recently emerged in recent years, such as scooters, e-bikes, and personal aircraft. These new platforms will require robotic technologies to facilitate widespread adoption, including autonomous stabilization controllers and accurate localization and mapping for user-friendly and safe navigation. Another important technology will be improved batteries along with intelligent power management systems that deftly plan and manage the power systems using onboard sensor information. When these new forms of transportation are deployed in dense population centers, efficiently coordinating and safely planning the movements of multiple vehicles in congested byways will also be a critical need.

2.3.4 Freight

Robotic technology will facilitate the future movement of freight across the US using a variety of transportation modes. It is certainly possible that air freight systems will be remotely operated incorporating more automated systems in the near term future. Benefits of these air freight systems include not having to pressurize pilot compartments, setting trajectories that do not have to account for human comfort parameters, and enabling more flexible routing schedules.

Currently, freight being transported by trucks involves high fuel costs in addition to releasing emissions that are harmful for the environment. Robotic technology will expedite tightly coupled truck conveys across long distances to significantly reduce air resistance for freight delivery and will significantly mitigate these costs.

2.3.5 Home delivery systems

The ever increasing demand for home delivery from the rapid growth of online retail in the US implies that future delivery systems will have to become more automated and utilize novel forms of locomotion. Future possibilities include wheeled robots for curbside delivery, legged robots to carry packages to the door, and aerial robots for deliveries through the sky. These robots will also need to be tightly coupled to highly efficient logistical systems in order to reduce package delivery times [See Section 2.2].

2.4 Quality of Life

Robots can be used to improve the quality of life for Americans. There is an increasing need for robot systems to assist people in their homes with activities of daily living (ADLs), with education, and with their healthcare and wellness needs.

The population of the United States continues to age. In 2030, 21% of the population will be older Americans (65+), as opposed to 15% in 2018 (US Census Bureau, 2020, see also Figure 2.5). By 2060, over 25% of the U.S. population will be over 65, with a tripling of people who are 85 and older. Robotics technologies can help our aging population to

age in place (at home) while maintaining their quality of life, including enhancing mobility and strength, providing transportation, home-based healthcare, physically and socially assistive technologies, and robot-based rehabilitation.




Figure 2.5: The age pyramid for United States (US Census Bureau, 2018)

Additionally, there are increasing caregiving and education needs as more children are diagnosed with developmental disabilities. Approximately one in 6 (17%) children in the United States have some form of developmental disabilities; the rate has risen steadily for over two decades Zablotsky *et al.*, 2019. In 2017, one in 59 children aged 8 were diagnosed with autism spectrum disorder (ASD) . Robots systems can improve educational outcomes for children with disabilities and learning differences (e.g., by teaching social skills to children with ASD) and provide support for ADLs.

Quality of life advancements in home and service domains can be organized around three dimensions: (1) addressing the hierarchy of needs, (2) supporting the complete human lifespan, and (3) enhancing human capabilities (whether for independence or to develop beyond the current limits of human capability), as shown in Figure 2.6.

At the base of the hierarchy of needs Maslow, 1943 are fundamental physiological needs like food, water, and sleep; the tasks that address these basic needs are described as Activities of Daily Living (ADLs).

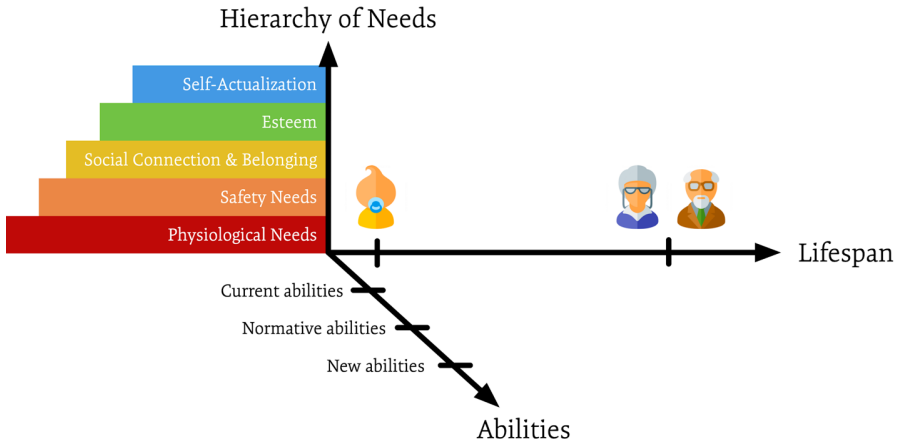


Figure 2.6: The multiple dimensions of users and needs (developed by road-mapping team)

The canonical ADL list includes eating, dressing, hygiene, transferring or walking, bathing, and continence. Robots have the potential to improve people’s ability to independently complete *all* ADLs, reducing caregiver burnout, enhancing and making care more accessible, thereby improving quality of life for both users and caregivers.

Robotics research has already made some advances in these areas. For example, researchers have investigated how robots can make eating a more seamless process for people with severe upper motor impairments. Researchers have investigated the use of robots for dressing and bed transfer.

Assisting with these basic physiological needs requires close contact with people. Therefore, issues of safety come strongly into play. Robots must be strong and capable enough to manipulate objects, but safe enough to operate around people with limited mobility. Privacy challenges also come into play here, because robots will be privy to extremely personal situations. There are also challenges around adoption at the policy level (e.g., FDA approval) and at the personal level (e.g., will people want robots to help them with this task).

Robots can both enable people with disabilities to have improved access and increased independence, and to bring new abilities and new

Robots also have the potential to enable people to do work they could not (alone) otherwise. Examples include home improvements, carrying heavy objects, etc. In many circumstances, the ability to perform this work can provide a sense of independence and satisfaction with being able to perform tasks without asking for help from others.

Robots may eventually be designed with capabilities that can help achieve an individual's self-actualization or the full realization of talents and potential of human beings. Robots like these might take the form of a Yoda-like figure, a life coach, a sage or a spiritual advisor.

Ultimately, robots can increase access for people, as they provide a mechanism for people to engage and interact with the world in ways they may not be able to do on their own. To be useful in practice, robots need to be usable by people of all abilities and communities. This ranges from basic ensuring interactions with robots are accessible to inclusion and cultural appropriateness.

2.5 Clinical Healthcare

Robots can assist hospital personnel in many of the stages of caring for patients, including surgical assistants, nurse assistants, therapist assistants, and direct patient assistants.

Robot manipulators on mobile bases will need to efficiently plan and consistently perform fine manipulation and grasping tasks in unstructured and constrained environments. Precision in motion and sensing allows more accurate orthopedic procedures. Minimally invasive surgery tasks can be partially automated for teleoperated da Vinci surgical robots, such as suturing. The new minimally invasive surgical method of steerable needles and concentric tube robots can be more precisely controlled by robots. Magnetic microrobots can perform surgical tasks such as inspection, drug delivery, and cutting. Camera pills with added magnets can be precisely controlled through the GI tract, including backing out when stuck.

Many routine tasks for a patient's hospital care can be performed or assisted with by robots, thereby offloading demands on aides and nurses throughout day and night.

- Patient assistant. A robot can fetch items for a patient or pick up dropped items. A robot can assist a patient to reposition, get into and out of bed, go to the bathroom, or get up and walk around safely.
- Nurse assistant. A robot can perform surveillance on a patient, to monitor activity throughout the day or note any difficulties the patient may be having.

After surgery or other interventions, a physical therapist can be assisted by a robot to perform routine repetitive motions and measure progress quantitatively through motion tracking. The robot can assist with fall prevention and lifting the patient to protect the therapist from heavy loads.

With robots able to perform these tasks, a natural next step is telemedicine and remote treatment.

2.5.1 Motivational statistics

Healthcare in the United States comprises nearly 18% of the GDP, making it larger than any other sector of the economy. In-patient care accounts for about $\frac{1}{3}$ of this total – more than 1 trillion dollars annually. Healthcare is also highly labor intensive with 12% of the US workforce directly employed in healthcare, which does not account for related industries such as laboratory services or medical device companies. Healthcare touches everyone -- over 20% of the world's population has a motor, cognitive or sensory impairment, the average American will have more than 9 surgical or interventional procedures, and In 2012, there were 36.5 million hospital stays in the United States, with an average length of stay of 4.5 days and an average cost of \$10,400 per stay.

Over the past decade robotics has begun to make inroads in in-patient healthcare. More than 5,000 robots have been deployed in hospitals for logistics tasks such as delivery of meals, beddings, and medicine (International Federation of Robotics, 2019). Systems designed to provide support or treatment for the disabled, those undergoing rehabilitation, and the aged are also seeing transition into practice.

Robotics is poised to have even larger transformational impacts in the healthcare industry due to their ability to extend, augment and quantify healthcare activities. Since its inception roughly 30 years ago, robotic surgery has grown to the point that it is now the standard of care in multiple common procedures. Growth continues to be rapid with an increase of 18% from 2018 to 2019 with well over 1 million robotic minimally invasive surgical procedures performed. Looking ahead, robotics provides both a platform for new forms of surgical augmentation and quality improvement. Robotics effectively creates physical and computational aids that project the eyes, brain, and hands of an interventionist into the human body. However, in the process of doing so, a robot becomes a mechanism for measuring and quantifying surgical performance itself.

The ramification of this new form of “surgical data science” is manifold. From the patient perspective, it is a mechanism for precisely quantifying what happened to that patient in that surgery -- effectively the “dose” of surgery provided to that patient. From the clinician perspective, this provides a platform for both learning and augmentation using data aggregated across thousands of surgeons and millions of procedures. This will improve the training and learning curve of surgeons, as well as providing immediate decision support and means for retraining and quality improvement in surgery. At the system level, this provides a means to study the effectiveness of surgery and to relate performance to patient outcome. This in turn will have immense implications for both the science used to define models for care, and for studies that determine the cost-benefit tradeoff for reimbursement of interventional procedures.

Robotic systems such as the MIT-Manus (commercially, InMotion), Lokomat (Hocoma) and Proficio (Barrett Medical) are also successfully delivering physical and occupational therapy. Rehabilitation robots enable a greater intensity of treatment that is continuously adaptable to a patient’s needs. They hold the potential to amplify the impact of physical therapists through a greater number of hours spent in therapy, and in some scenarios have already proven more effective than conventional approaches, especially in assisting recovery after stroke, the leading cause of permanent disability in the US.

Such systems can play a therapeutic role not only for movement disorders (such as those resulting from stroke, traumatic brain injury, and other trauma) but also as intervention and therapeutic tools for social and behavioral disorders including autism spectrum disorder, ADHD, and other pervasive and growing disorders among children today. We also have seen emergence on the commercial market of human-operated wheelchair-mounted robotic arms with FDA-approval (e.g. the JACO from Kinova Robotics).

A large part of the cost of hospitalization is in patient care. After the patient returns home, it becomes difficult to continue the quality of care received in the hospital, due to the absence of trained nursing staff and adequate facilities. Both the cost of hospitalization and the disruption in care at home can be alleviated by transporting a nursing assistant robot home with the patient, after the patient has familiarized with the robot care during the hospital stay. A nursing assistant robot might handle tasks including but not limited to fetching objects, feeding the patient, taking the patient to the lavatory, dressing the patient, washing the patient, and general assistance with the activities of daily life. Robot assistants should be designed in a way that upholds the dignity of the patient in the same, if not more prudent, way that a nursing assistant would. Daily tasks that the patient feels sensitive handled by a judging human eye can be better handled by a robot, provided that the interaction is designed to be positive, constructional, yet impersonal to an extent.

2.6 Feeding the planet

Food represents about 12.9% of the average household expenditure in the US and accounts for close to 6% of the GDP or \$1.053 trillion (2019). Food and agriculture also represent about 11% of the employment in US (according to USDA (USDA - Economic Research Service, [2020](#)). More than 9 billion chickens, 241 million turkeys, 131 million hogs, 33 million cattle and calves were processed during 2018. That is 317 animals processed per second. Meat, wine and dairy are the three major sectors of the domain in terms of employment. In comparison the fishing industry is the smallest sector. In particular the fruit and vegetable industry has a strong reliance on a migrant workforce.

Over the last couple of decades there has been a growing interest in use of robots as part of food processing. The spectrum covers all aspects of food processing for planting seed in the ground over weed removal to picking mature fruit/vegetables. More than 10 years ago John Deere presented the concept of a driverless harvester / tractor (John Deere, 2020). The idea was to enable farmers to task a vehicle to maintain crops without requiring a driver to be inside the vehicle mainly supervising a large autonomous operation. So far these vehicles have seen little real deployment. In general robots have been applied to precision agriculture, weed control, nursery automation and harvesting. Precision agriculture is used to monitor crops, collect data and apply fertilizer (as shown in Figure 2.8). Weed control is either mechanical weeding or delivering small amounts of herbicides to weeds. Nursery automation includes managing weed, transportation and monitoring. Drones are starting to be used to map out orchards to monitor the state of growth. The cost of a drone is typically too high for a single farmer but as “drone as a service” evolves it is increasingly an economically viable option for farmers. Today the estimate is that farmers can afford \$5 per acre for such services (International Federation of Robotics, 2019). As such a modest amount of automation has so far been introduced for use in the field. There is a significant potential for automation in the field, especially in a time when a migrant workforce is harder to get by.

It may be surprising to some but one of the largest areas in field robotics is milking of cows. This is a billion-dollar industry. More than six thousand milking robots are sold annually. The milking robots typically reduce the cost of milking a cow by 10% and at the same time the machines allow for milking 24/7 with minimum supervision. This is an industry that is seeing 10% annual growth.

Processing of meat is another area that has seen some growth. The main challenge has been the high rates of processing and to compete with skilled labor. As progress on vision and force-torque sensing evolves there is a tremendous opportunity to increase automation.

The World Bank states that the world will have to produce 50% more food by 2050 if the global population continues to rise at its current pace. Over the next 15 years, global demand for meat is expected to increase by 40% triggered by a growing number of people adopting



Figure 2.8: Using UAVs for maintenance of a field (Source: Shutterstock)

protein-rich diets. Crop yields will have to rise by at least as much as crop demand to avoid further encroachment of cropland into natural habitats (Blomqvist and Douglas, 2016).

Access to affordable automation that will allow increase in productivity will be essential to enable wider use of robotics and automation in the food sector.

2.7 Security and Rescue Robotics

2.7.1 National Security

A primary role of the US Government is to protect the wellbeing and standard of living of its citizens and permanent residents. Some of the large and important tasks that must be done to carry out this role are to: prevent illegal immigration and drug smuggling, maintain robust national infrastructure, and find safe efficient travel zones during large-scale disasters (such as wildfires, hurricanes, and floods). A prerequisite for doing each of these tasks effectively and efficiently is real-time surveillance and inspection of millions of acres of land and waterways and public infrastructure built upon them along our borders, in fire and flood zones, and the spaces in which public infrastructure are built.

While these areas can be seen by satellites, swarms of sensorized robots coupled with 5G communications offer the possibility of obtaining data that satellites cannot gather, and gathering it at high resolution and in real-time. For example, satellites cannot “see” under trees or bridges, measure concentrations of hazardous chemicals, or track the front of a wildfire in real-time. Future semi-autonomous robots will be able to gather such data and use it to aid public servants who keep our borders secure and our infrastructure operational; and do so more quickly and efficiently, thus saving lives, reducing economic loss, and reducing the cost of the operation, as shown in Figure 2.9.



Figure 2.9: Example application of a security robot (Source: Shutterstock)

2.7.2 Infrastructure Inspection and Maintenance

The American Society of Civil Engineers estimates that to rehabilitate the Nation’s infrastructure to a level necessary to maintain global economic competitiveness and public safety by 2025 the Nation must spend \$200 billion per year more than it is currently spending (American Society of Civil Engineers (ASCE), 2020) This calculation covers all core infrastructure, including transportation networks, waste disposal, and

fresh water and power distribution. Robots can help with rehabilitation by inspecting existing many facilities with greater coverage and detail than human workers. They can inspect structural members of bridges and powerline towers, find defects in road surfaces and above-ground pipelines, spot dangerous debris floating in waterways, and many other things. Swarms of future autonomous robots with advanced perception algorithms and natural human interfaces have the potential to allow workers to inspect all critical infrastructure. The data gathered could be sent to a command center for processing to identify high-priority repairs or upgrades and which require repair by humans. Preliminary processes could be done on location to identify simple repairs that robots could do autonomously.

2.7.3 Evacuation in Large-Scale Disasters

The number of acres burned worldwide is increasing with global warming. Firefighting costs and damages are rising steadily. Cost of fighting U.S. wildfires topped \$2 billion in 2017. Over the last 20 years, the fraction of the budget of the National Forest Service's spent on firefighting has increased from 15% to 50%. One consequence of this spending shift has been the reduction of fire prevention work, such as controlled burns, which increases the chance of large wildfires (Zuckerman, 2017). In California alone, in 2018, more than 58,083,000 wildfires (Insurance Information Institute, 2020) burned 8.8 million acres causing an estimated total economic loss of \$400 billion to the state of California, making it the most expensive natural disaster in the history of the United States (Accuweather, 2019). The Camp Fire in Northern California wiped out the town of Paradise, killing 85 people in the process, making it the deadliest in California history (Cal Fire, 2019). It was the world's costliest natural disaster in 2018 causing \$16.5 billion in damage (Amadeo, 2020)

Swarms of future autonomous robots with advanced perception algorithms have the potential to allow real-time tracking of fire fronts, prediction of firefront movements, and road and terrain conditions to compute the safest evacuation routes. Evacuation routes that could be planned with the data are not restricted to roads, but could be planned

through rough terrain that could help firefighters reach safe zones, as shown in Figure 2.10. This could eliminate a leading cause of firefighter death – entrapment (National Wildfire Coordination Group, 2017).



Figure 2.10: Scenario for rescue workers collaborating with robots (Source: Ye Zhao)

2.7.4 Border Surveillance

Billions of dollars are spent annually to secure the Nation’s borders. The US Department of Homeland Security (DHS) spends about \$4 billion annually on its border patrol operations and the Coast Guard’s annual budget is about \$11 billion per year (American Immigration Council, 2021). The border of the US is thousands of miles long - too long for thorough surveillance by active-duty enlisted members of the Coast Guard and border patrol agents. Therefore, despite these expenditures, large amounts of illegal drugs enter the US every year, decreasing the Nation’s productivity and increasing crimes and their associated costs to the public. The economic cost of drug abuse was estimated in 2007 (the last available estimate) at \$193 billion; \$121 billion is due to lost productivity (time in incarceration, drug abuse treatment centers, etc.), \$11 billion due to associated healthcare costs, and \$61 billion due to criminal justice, incarceration, victim costs, etc. (Office of National Drug Control Policy, 2016).

Swarms of future autonomous robots with advanced perception algorithms and natural human interfaces have the potential to allow workers to monitor the entire border, from the air and underwater. Border security personnel will be able to command missions by fleets of autonomous robots and ensure complete accurate coverage of the border. Using advanced AI techniques, the robots' perception systems will learn to continuously improve their abilities to identify people and boats attempting to cross the border illegally.

3

Mapping Societal Drivers to Research Challenges

In this section the main obstacles to societal progress are identified and discussed based on the discussion in Section 2. Initially the various application domains are analyzed and based on this the obstacles to progress are mapped onto current/future research challenges.

3.1 Identifying challenges to growth/progress

3.1.1 Manufacturing

A key factor in the introduction of automation into manufacturing is always cost. The business-case has to make sense and the margins/rates in a manufacturing plant are often a challenge. There is thus a need to consider how the cost of installation, operation, and maintenance can be optimized.

In recent years there has been a push towards a higher degree of customization. As mentioned in Section 2.1 cars are now available in millions of different configurations. As such manufacturing is very much becoming a high-mix/low volume environment. There is a large number of different variations and every item manufactured is different from the previous unit.

In manufacturing safety is always a major objective. The introduction of collaborative robots around 2005 changed the setup of factories. Prior to that there was typically a physical barrier between robots and humans on the factory floor. The new collaborative robot systems allow for a more flexible cooperation between humans and robots. They can exist in the same space and it is possible to dynamically interact through careful design. Today collaborative robots are almost exclusively used for smaller payload tasks. Safety will remain a major focus as the possible set of applications is expanded. The drivers are summarized in Figure 3.1.




Figure 3.1: The main driver for automation in manufacturing

Today there are only about 1 robot for every 50 workers in manufacturing. A challenge to adoption in particular for small and medium sized companies is the time for setup for a new task. There is a need to make the system easy to set up for a new task. The setup time should ideally be shorter than the task time.

Finally, there is a need to make it simple to use automation/robotics for all the users. Today it is often required that operators spend significant time in training to be able to operate and/or do basic teach-in. Ideally the process of use of technology should be effortless.

3.1.2 Logistics

Logistics has seen tremendous growth over the last decade. One of the big challenges in logistics is the tremendous variability across the items handled. As an example, a normal US grocery store has more than 45,000 items in inventory all of them are delivered through a standard supply chain. FedEx handles 10 million packages per day (FedEx, 2020). Amazon at it peak ship more than 25 million units per day. Consequently, there is a tremendous need to handle high variability and every item is likely to be different from the next.

Obviously as the logistics is expanded it is essential that the system be safe at all levels from handling items in distribution center to last mile deliveries on residential streets. The COVID-19 pandemic has clearly shown a significantly increase in home deliveries for basic groceries and meals to all things e-commerce. The main business drivers are outlined in Figure 3.2.




Figure 3.2: The main drivers for increased automation in logistics

The logistics sector has traditionally seen a significant turn-over in the workforce. Some companies see a new person in every sorting

position 3.5 times year. Consequently, there is limited if any time available for training of the workforce and it is essential that use of equipment is effortless.

Introduction of automation has the potential to provide significant savings in terms of throughput, 24/7 operation, and streamlining operations. However, it is essential that the systems have a robust performance. Minutes of downtime can be extremely costly and given how streamlined processing is the down-stream disturbances can be significant.

With the expansion of logistics to handle packages, meals, groceries, etc the diversification drives down prices. The cost of any new technology will have to be carefully considered to ensure capitalization over a reasonable period of time.

3.1.3 Transportation

Over the last 10 years a small revolution in transportation has taken place. Cars with some degree of autonomy have entered the market such as the Tesla. TuSimple are testing level 4 autonomy for daily for logistics transport between Phoenix, AZ and Houston, TX. In addition, UAVs have started deliveries of critical items such as blood samples. UAVs are also tested for early response as part of the 911 system (early testing in Chula Vista, CA). The transportation sector is predicted to see major changes over the next 5–10-year period. Waymo is promising level 5 autonomy. Amazon is testing Scout last mile delivery vehicles in multiple metropolitan areas.

An essential part of delivery of autonomous vehicles at level 3–5 is the need to carefully consider safety. Human drivers on average drive 10 million miles between major accidents. Nonetheless close to 40,000 people are killed on the road of USA every year. Safety regulations vary across states and across highways so there is a need to harmonize safety regulations, but also to provide unified safety regulations for ground vehicles. For UAVs the technology has been in place for 5+ years but there is a lack of a regulatory framework that allows for more widespread utilization of unmanned aerial vehicles.

For operation on public roads and in public airspace there is a need to ensure that the vehicles that have a robustness and reliability that is at least as good as human operated vehicles. On the ground as mentioned above the requirement is at least 10 million miles between major accidents and in the air the normal requirement cited is 7 years between fatal accidents given present traffic volume. Consequently there is a need to consider how we can build systems that have long-term robustness 24/7 under all kinds of environmental conditions.

In transportation it will also be important to consider the usability of new technology. As some of the technology will be on our residential roads and some vehicles will transport young and elderly people it will be essential that the technology is effortless to use.

3.1.4 Quality of Life

Quality of life has many meanings and as such is a categorical label that covers most use-cases to assist people in their daily lives. We have already seen 10 million plus vacuum cleaners in people's homes, but the use-cases cover social access, mobility assistance, rehabilitation, cleaning, entertainment, etc. Given the diversity in the possible use-cases it is essential that the technology is able to accommodate a wide variety of residential setups from studios/apartments to large houses, that may include multi-level setups. As such the technology must be able to enter a large variety of setups with a minimum of adaptation.

Obviously for assistance to people in their homes and essential aspect is safety. The units will in some cases be in directly physical contact with people as must have both physical and psychological safety mechanisms not to mention the need for strong cyber security and ways to explicitly address privacy issues.

In particular for Quality of Life technologies the units will be deployed to assist people that may never have used a robot before and in most cases, they will not be technology experts. As such it is essential that interaction / embedding of the technology is effortless. It is unrealistic to assume that much if any training can be required.

For deployment of quality of life technologies, it will also be essential that the technology can be made accessible to as many people as possible.

Cost is consequently going to be a major factor in the design and delivery of next generation robot systems.

3.1.5 Medical/Clinical Support

Medical robots have already seen major successes. The use of robots for minimally invasive surgery is a great success story. The use cases for cancer treatment and for radiation therapy are well documented. As more flexible instruments are invented the number of new applications will grow significantly. During COVID-19 it was also demonstrated how robots can be used for medical professionals to consult a patient using tele-presence. A doctor could be in another building/city or just 25 ft away, but the robots allow appropriate social distancing even in a clinical setting. There is no doubt this trend will gain momentum over the next 5–10 year period. For this is it of course essential that the systems are robust. As services are provided it is not acceptable to have systems that are no robust as the consequences could be significant.

Safety is also essential for the use of this technology as in many cases the robots will be in directly physical contact with patients and personnel. Already today FDA has a strong regulatory framework in place. Considering design of systems that are inherently safe is a hallmark for next generation systems.

Every client is different from the next one. As such there is a need to design systems where the setup time is minimal. Early systems required extra time for setup and calibration, which resulted in a reduced throughput. In many cases that it not acceptable and as such flexibility in the face of accuracy will be a key driver in the design of new technology.

3.1.6 Feeding the planet

Agriculture is going through a major resolution in use of technology as mentioned in Section 2. At the same time fewer and fewer people are engaged in food production. Traditionally some areas have relied on a migrant workforce for maintenance of fields and harvesting. In the present climate there is a significant shortage of migrant workers. In addition, food safety is playing a bigger role. There is a need for

traceability if there is a contamination of food. The planet is still far away for being able to feed its population.

A challenge in design of technology for feeding the planet has been natural variation across crops, vegetables, and fruits, but also the variation in the size of animals, fields of crop, etc. There is a need to design technology that can handle the significant variation in the tasks and products. In many cases the volume of a particular product is too small to warrant a particular solution. I.e. the volume of strawberry is too small to warrant a specific solution. It is necessary to design solutions that can cover multiple products or product groups.

The margins for production of food are very limited and as such the cost of new solutions must be competitive with manual labor or the productivity increase must finance the new technology. As such it is essential to consider innovation that has a broad enough set of applications and require a small enough investment to make it viable for small producers.

Vertical farming is also starting to get traction which typically is run by very small producers for private or public use, but it has to be provided at a minimal cost.

Given the natural variations there is a need to build technologies that are robust enough to be used 24/7 across the continent including weather and climate changes. In addition the systems have to be simple/effortless to use. Most of the employees have little to no interest in becoming technology experts.

3.1.7 Security & Safety

Gradually robot technology is becoming a core technology to assist first responders. Robots have a long history of assisting with dismantling of explosive devices and inspection of improvised devices. In addition, UAVs are used for situation reporting from partially collapsed structures and early reporting from accident sites. More recently robots are also being used for inspection of critical infrastructure such as bridges, damms, electrical power lines etc. In almost all cases the scenario is unique. No two bridges are alike, and a potential threat device was designed to be unique. Consequently, there is a need for robustness and

handling of major variations across scenarios. In many cases the time to setup base and address the particular challenge is limited and as such there is a need to find new solutions that are drop and execute without a lot of prior knowledge. The solutions must be robust as failure can be very costly and, in some cases, result in loss of lives.

Another major consideration for these systems is safety. As lives are at stake it is essential that the systems provide the best possible protection of operators and by-standers but also provide the best long-term safety to ensure that as few as possible lives are at risk and that environments are preserved to the extent possible. Failure is not an option in almost all cases.

3.2 Mapping challenges to research needs

3.2.1 Cost

For most use-cases the cost of providing a solution is a major driver. There are professional applications where cost is secondary. For the US to remain competitive in an international context the solutions provided must be cost-effective. Typically, the cost must be considered in a systems context where design, installation, operation and maintenance are all considered jointly.

New materials are opening up opportunities to build systems that are most agile and have a richer set of sensors for operation. In some cases, such new mechanisms can be 3D printed, which in turn provided a new degree of agility.

Cost of installation can sometimes be reduced through flexible use of planning methods that allow on-the-fly organization of production systems for manufacturing of one-off products or designing medical treatments that are patient specific.

The introduction of planning, adaptive methods/learning, and common sense has the potential to significantly reduce cost. In addition, by using learning by demonstration it is possible to reduce the setup time and allow domain experts to teach a system with little to no effort.

A limiting factor is often the need to train operators in the use of a new system. Some systems may take days or weeks to master, which

poses a major cost during introduction of new technology. There is a need to design technology with user interfaces that are easy to use such that training is kept at a minimum without loss of performance or safety.

3.2.2 High Mix / Low Volume

The world is changing to an environment where each task is a one-off. The concept “lot size one” has been used to describe such a setup. To enable implementation of such an approach there is a need to endow the robots with rich perception capabilities to allow the robot to detect and adjust to product / task variations on the fly. Progress in sensor technology is allow for integration of multiple modalities such a sight, tactile, force, temperature, etc. into a coherent model for estimating the state of the robot and the external environment for more effective interaction with objects and users. The key challenges to progress for flexible manufacturing are shown in Figure 3.3.




Figure 3.3: Main challenges for flexible manufacturing

It is becoming possible to design mechanisms that have a much higher flexibility in handling of objects. It is possible to design end-

effectors that allow the system to interact with a large variety of possible objects, but also to provide maximum safety. So far, the use of new materials and design concepts is still in its infancy but it offers major new opportunities for added agility for robots.

Methods in planning and adaptation will also play a key role in the design of systems that have a high degree of adaptive to different objects and tasks. Planning under a high degree of uncertainty will be required for some of these tasks while providing a robust performance.

Autonomy has been studied for some time and it is a key aspect to building systems that are robust in the presence of significant variations in the environment. For some use cases the high-mix is deterministic, but in other cases such as agriculture or security it may not be possible to describe all the possible variations up-front, which is where autonomy will play a major role.

3.2.3 Safety

The importance of safety cannot be over-emphasized. One accident is impacting the entire industry. Consequently, safety has always been a major factor in the design of next generation systems. Safety is ingrained in all aspects of a robot system from the mechanism over perception and planning/adaptation to the user interface.

Collaborative robots have demonstrated that the integration of materials, control and mechanism allows for design of systems that are inherently safe for human interaction. The design of collaborative systems is still in its infancy. There is a need to design systems for large payload applications. There is also a need to consider how safety can be provided for autonomous vehicles. How can one design mobile systems that have a safety record that is better than systems operated by humans?

In some cases, it is advantageous to use multiple vehicles to address a situation as it provides more flexibility, but it may also improve safety. One such example in radiation therapy where is it possible to use one robot to move the patient bed and another to move the radiation source. The joint system reduced the amount of radiation dosage for healthy tissue. Consequently, it is essential as part of safety to consider the best possible design of a system.

The design of interactions with humans is also an essential part of robot studies. For physical interaction it must be inherently safe to use the mechanism. For non-contact applications it must still be secure to operate a system. Mobile platform will have to be used for example for logistics applications driving on sidewalks. It must be safe to operate such vehicles with little to no-instructions to humans.

3.2.4 Effortless Usage

In some cases, such as safety and security it is reasonable to assume that the operator has received significant training. The same is true for operators of large unmanned aerial vehicles. Medical professionals are also likely to have received training. However, for a majority of use-cases it is necessary to move towards a future where the interaction with a system is effortless. How can we design systems where it takes no training to start to use a new robot? I.e., a vacuum cleaner with only one button – “clean”. Future robots will have a rich set of perceptual modalities for gesture, body-motion, speech, facial expression, etc. interaction that will make it evident to the robot the intent of the user and other people in the vicinity of the robot. It must be possible to have a notion of drop and deliver where the system is unpacked and immediately ready to perform its expected function. An example of such a codeless deployment from Ready Robotics is shown in Figure 3.4.

The effortless functionality will require that the robot has a rich set of perceptual capabilities to understand the state and intent of objects and people in its vicinity. It must be able to reason about tasks and the ability to execute the task. It must be able to use common sense/preexisting plans to generate solutions without needing the user to train it for hour or days.

The need for effortless operation is particularly important for multi-robot systems. If 10 UAVs are needed to fight a fire or deliver packages in the local neighborhood it is unrealistic for an operator to directly control all of them. It must be possible to organize the team in such a manner that intent and task objectives are communicated, and the rest is executed with a minimum of human/operator interaction.



Figure 3.4: Example of a system optimized for simple programming (source: Ready Robotics)

3.2.5 Setup Time

It has already been noted that sometimes it can be tedious to setup a robot for a particular task. As high-mix takes over and an increased flexibility is introduced it will be essential to find mechanisms to reduce the time to setup a system.

An important part is the design of flexible user interfaces that allow people to organize the task at hand with minimum interaction with a system. When we enter our car in the morning to go to work it does not require a lot of setup time. A similar paradigm is needed for robot systems.

Planning, adaptation, and machine learning is offering an opportunity to reduce the setup time. It is possible to build libraries of plans that can be adopted to particular use-cases. It is also possible to leverage learning by demonstration to allow for adoption to a particular task after a few demonstrations. This will require that the robot has the perceptual capabilities to understand its environments and changes to the setup.

Another option is to endow the systems with a higher degree of autonomy. Through use of methods from autonomous systems it is

possible for the system to handle richer variation in task specification and still generate robust solutions. The plug-n-produce model for setup of robot systems has a lot of potential. Through design of effective models for user interaction it is possible to provide such solutions.

3.2.6 Robustness

Robustness is essential to the design of next generation robot systems. It is possible to make robustness a key design criterion for new systems. Progress on new materials and on mechanism design can ensure an increased degree of robustness. As an example, the new generation of legged robots allow for last mile deliveries in residential neighborhoods with a need to navigate stairs at the entrance. For many other use-cases it is possible to design systems that are robust and potentially much simpler to control. Such approaches are important to progress.

Robustness permeates an entire system. There is a need to not only have flexible mechanisms but also robust perception. Having a model of the environment of a robot allows it to plan how to proceed to accomplish its task and how to control the system in a manner that optimizes mission success. Robustness is also about integration of planning under uncertainty to ensure that knowledge about uncertainty in estimating the state and the uncertainty in task execution both are an integral part of the reasoning about the execution of a task.

Robustness can also take learning and adaptation into account. If a task is executed multiple times, then the performance can be optimized over time through use of learning and adaptation.

In some cases, it is also possible to leverage multiple robots to accomplish a task such as monitoring a building that is about to collapse or mapping out a forest fire. In such cases the utilization of multiple robots not only improves the speed of execution, but it also provides a method to become failsafe in the presence of failure by one or more robots. It is thus of interest to study how multi-robot systems can be deployed for overall robustness in task execution.

4

Research Challenges

Successful deployment of robot systems requires careful consideration of the business drivers, the main obstacles to progress and the required research efforts to be undertaken to deliver as illustrated in Figure 4.1.

4.1 Architectures and Design Realizations

For over 60 years now, the robotics “Sense-Think-Act” paradigm has enabled extending the reach of humans for manipulating, interacting with and transforming the world. During this time, the methodologies, methods, and materials used to build robots have been gradually changing -- away from the traditional low DOF rigid-link architectures with discretely sense/actuated joints and centralized controllers to variable-topology reconfigurable high internal DOF systems with distributed/integrated multimodal sensing/actuation.

A next-generation Distributed Networked Robotic System paradigm has also been slowly emerging from (i) decomposing the traditional monolithic robotic-system via the Digital Redesign paradigm; and/or (ii) composing loosely interconnected heterogeneous components in a system-of-systems approach (Christensen *et al.*, 2016). Emerging middleware paradigms (such as ROS) have supercharged the creation of

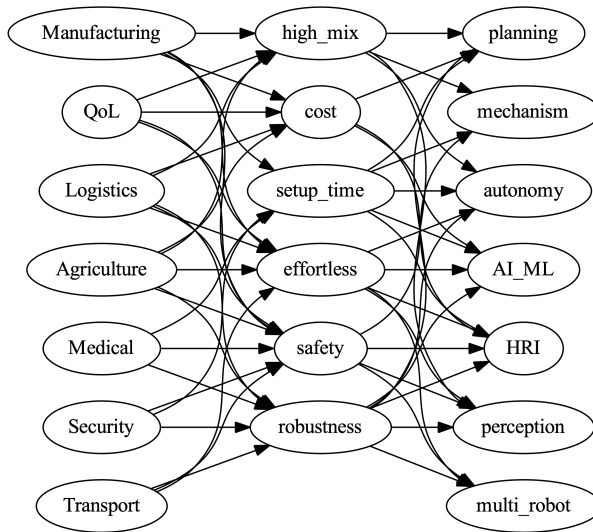


Figure 4.1: Relation between business verticals, obstacles to progress and research challenges

such modularly composed networked systems -- helping to ameliorate the challenge of building every system from the ground-up. In such a networked world robotic systems-of-systems can not only access resources from each other but also include various infrastructure elements (e.g. cloud compute) in ways not previously possible.

At every stage, advances have capitalized on the ambient technology advances, in terms of convergence of computation, communication and miniaturization, for embedding intelligence; new materials and construction paradigms; new manufacturing techniques to generate increasingly compact, capable and energy-efficient subsystem- and system-level integrated realizations. The characteristic inherent feature is *complexity* – arising both from the diversity of disciplines engendered, the integrated technologies and the increased scale/numbers due to the distributed paradigm.

While deployments of end-to-end operationalized systems in application spaces are emerging, systematic engineering of high-performance/high-confidence operational capacities, in the presence of uncertainties, has proven challenging. Verification and validation and re-engineering performance in such loosely interconnected distributed networked dynamical systems-of-systems remains a significant challenge. Significant emphasis needs to be placed on lifecycle treatment (design, analysis, refinement, prototyping, and validation) of such Distributed Networked Robotic Systems with the goal of realizing tangible enhancements in functionality, performance, and cost-effectiveness.

The core requirements especially for successful fielding of robots in real-world settings include: semi-autonomous operations, continuous adaptation to its environment, data-driven learning and control together with energy-efficiency and zero down-time.

Reenvisioning cyber-physical system-architectures: Currently, every area in robotics builds on a foundational element -- the availability of one or more reliable, robust networked cyber-physical platform(s) with adequate real-time computational intelligence from an appropriate suite of sensing, computation, networking and actuation. As new application spaces and use-cases (e.g., operating inside a body, emergency response scenarios) emerge, there has been a movement to enhance the internal architectures and degrees of freedom of the underlying articulated mechanical systems to provision greater dexterity and mobility. As human-robot interactions are on the rise with semi-autonomous/autonomous agency (bilateral power-exchange with environment), it is critical to assure security (physical and cyber) and safety from the ground-up for the core foundational cyber-physical systems. There is an opportunity for synergy from the basic embodied/materialized realizations, extended through individual component, integrated subsystem and networked system-architecture selection process to create novel CPS platform capable of exploiting new capabilities through advanced algorithmic controls. In addition to the gaining of traction of seamless and safe integration of humans and robots, there is also a significant emphasis on novel materials, electromechanical design architectures and modular decision-making and control software.

Engineering the data-interface to/from the analog/digital domains in such distributed asynchronous system-of-systems faces numerous challenges: from sensor/actuator phenomenology to calibration-drift at the component-, subsystem- and system-levels. As such this forms the tip-of-the spear -- utilizing such unsynchronized spatio-temporal data-streams as inputs for information-extraction and inferencing can create significant challenges to the relevance, robustness of developed situational-awareness. Distributed actions within such a loosely interconnected framework based on uncertain inferencing can lead to degraded performance (as compared even to monolithic benchmark counterparts where they exist). Yet teamwork in robots -- with the ability to distribute operations in a decentralized setting-- remains the gateway to scaling operations.

Real-time Robotic Digital Information Architectures: is the gateway to analysis (via AI) and performance (in operational tasks). While computational simulation offers an early surrogate data-source, our ability to capture the complexity of the real-world remains limited. Sensing (as packaged into modular networked sensor subsystems) coupled with action (active-sensing paradigm) still remains the best lens into the traditional opaque world. However, ensuring the provenance and quality of the raw spatio-temporal data streams from the multiple spatially distributed and asynchronously temporally sampled sensors is critical. Core to the robot-supported active-introspection (into the traditionally opaque analog world) are mounted sensor-suites on individual robots or across the system can produce a significant amount of spatio-temporal information about the world. Coupled with information-enhanced real-time/interactive mobility and manipulation this empowers a range of advanced algorithms -- loosely called Collaboration for X. All the challenges of Big Data (5Vs: Velocity, Veracity, Variety, Volume and ultimately Value) manifest as these robotic systems-of-systems act as sensitive instrumented probes to gather data to inform decision-making in application-verticals (from agriculture to infrastructure inspection).

Multifunctional Modular Integration of sensing/actuation, mechanism and control: There is an intricate interplay between the underlying electromechanical architecture (sensing/actuation) and the algorithmic complexity of controlling them. Intelligent electrome-

chanical design with carefully-configured passive-dynamics can greatly simplify the control challenge of the next generation of systems-of-systems (Design for Control paradigm). New technologies for actuator, manufacturing, and construction paradigms will synergistically enable progress, as the line between control algorithm, hardware, and actuation blur.

New materials paradigms: 3D printed parts and softer polymers formed in 3D-printed molds, sometimes formed with other materials in a composite structure, have the potential to create a new paradigm of robot design that is more similar to soft biological machines and less similar to hard metal machines. While this field is in its early stages, it is clear that soft materials are far more effective than hard materials for gripping, manipulation, traction, and many physical interaction tasks. The strength and the challenge of soft materials are the complex dynamics of the materials; while compliance in a robot finger may be useful for gripping, it is also challenging to model, sense, and actuate. Continued development will yield new sensor paradigms, new actuators and transmissions (such as hydraulic bladders), and greater integration of the dynamics afforded by soft materials with the control methods for robot motion.

New Manufacturing Techniques: Additive manufacturing (3D printing) techniques have unshackled the traditional constraints on robot geometry and form (complex shapes and structures) but also deployment of new materials (multifunction materials) and embedding of sensor/actuator integration within robotic structural elements. Hybrid additive/subtractive manufacturing methods coupled with can be used not only to produce useful components but also as a part of the manufacturing process to generate molds for other materials or forms for composite structures. 2D planar manufacturing processes, such as laser-cutting, are being used to create complex 3D geometries using origami-inspired methods. MEMS-based fabrication techniques make it possible to fabricate truly microscale robotic elements.

The following are intended to serve as potential anchoring examples to highlight the vision for 5,10, and 15 years.

Merging the design of actuator, mechanism and control

- 5 Years** Encapsulated Design: Miniaturized Distributed Sensors
- 10 Years** Tradespace for hardware/software realization e.g. Cheetah vs Series Elastic Systems
- 15 Years** Adjustable n-rotors evolved into various other forms?

New materials and construction paradigms

- 5 Years** Subsystem level: piezoelectric sensing/power Nitinol Needles
- 10 Years** Distributed Compliant Arms/Exos with energy recycling during gait cycle
- 15 Years** Origami based design & Soft robots (Vijay Kumar, Rob Wood and UCSD work)

New manufacturing techniques

- 5 Years** Low cost Dextrous sensor-enhanced hands – integrated discrete components/sensing/actuation into hands to get manipulation data-sets
- 10 Years** Multifunction printers enable Sprawlita-like- design for no-assembly, conductive ink-traces
- 15 Years** Distributed Compliant Macro- and Micro-System designs with distributed sensing/actuation e.g. continuum robots

4.2 Locomotion

4.2.1 Legged Robots

Motivation: Legs are the most effective mobility solution for many environments, indoors and out. Evidence for this abounds in nature, where legged animals have populated everything from trees to the most extreme boulder-strewn mountain peaks of the planet, places where

locomotion by wheeled and tracked robots is infeasible. And while tracked and wheeled robots could be fashioned to move effectively on stairs in homes, most homeowners would not find that an acceptable solution.

Why bipedal robots and not quadrupedal robots, since the latter are inherently more stable? Section 2.7 of this roadmap outlines key steps toward developing the technology for a robot to navigate and traverse a disaster site. Humans are upright, narrow walkers and we design our homes, factories, and affordances to accommodate us: viz, tight turning areas, manholes, ladders, levers and valves that must be reached to be activated. An important contribution of first-responder robots would be to make preliminary maps of a site so that human first responders could be sent into the most promising areas to provide aid or prevent further damage. It is imperative that the robots can move effectively in human spaces.

The ability of robots to move freely in human environments will also enable them to work with humans, for humans, and around humans, in human environments. Combined with intelligence, perception, and manipulation, legged robots have the potential to become as ubiquitous in our world and in our spaces as cars are on the road. In other words, understanding and implementing legged locomotion as a general discipline in robotics will be one of the enabling technologies for robots to really impact human quality of life in a positive way.

State of the art: Companies are beginning to provide legged locomotion platforms. The miniSpot robot by Boston Dynamics is perhaps the most well-known quadruped at the present date, but the platforms ANYmal by Anybotics and the Vision 60 and Minitaur robots by Ghost Robotics are commercially available, and still others, such as DogBot by UK-based React Robotics, are coming. All of these robots work out of the box and come equipped with adequate control laws for traversing terrain typical of a golf course: paved paths, short grass, roots around trees, and sand traps (if not too steep).

In the area of bipedal robots, prior to the startup Agility Robotics, the available robots exhibited quasi-static gaits and poor energy efficiency. In 2018, Agility Robotics introduced its Cassie series of robots that were torque controlled, walked and balanced dynamically, and

came with built-in control laws that would allow the robot to walk indoors in uncluttered environments and outdoors on sidewalks, grass and gravel paths. Weighing 32 kg, its 4kg LiPo battery provides 3 hours of walking autonomy (shown in Figure 4.2).



Figure 4.2: The Digit-1 robot developed by Agility Robotics for last mile delivery (Source: Agility Robotics)

Spinoffs into Rehabilitation Robotics and Exoskeletons: Mechanical and algorithmic technologies being developed for legged robots are impacting the area of medical exoskeletons and lower-limb prostheses. At present approximately 4.7 million people in the United States would benefit from an active lower-limb exoskeleton due to the effects of stroke, polio, multiple sclerosis, spinal cord injury, and cerebral palsy (Dollar and Herr, 2008). Moreover, by 2050, an estimated 1.5 million people in the United States will be living with a major lower-limb amputation (Ziegler-Graham *et al.*, 2008). Such individuals expend up to twice the metabolic effort to walk at half the speed of able-bodied persons, experience higher-risk of falls, and have secondary pathological conditions such as osteoarthritis, back pain, and depression (Waters *et al.*, 1976, Pell *et al.*, 1993, Miller *et al.*, 2001).

Companies such as Ekso Bionics and Wandercraft are offering exoskeletons that allow patients with paraplegia to walk again without the

use of crutches for lateral balance. State-of-the-art lower-limb prostheses in research labs now have powered knees and ankles, and control software that can effectively coordinate their motion in specific situations. Lower-limb exoskeletons serve as assistive devices by providing support and balance to wheelchair users and enabling them to perform normal ambulatory functions such as standing, walking and climbing stairs. Lower-limb exoskeletons have also been utilized for gait training and rehabilitation purposes. The control laws for these devices are based on control work in bipedal robots.

Relation to Engineering Education: For students, learning how to control and utilize legged machines leads to learning most of the fundamental components of robotics which apply to any actuated physical mechanism. To be clear, the dynamic behaviors of legged robots are quite different from industrial robot arms, which have more in common with rigid CNC machines. Legged robots exhibit forces and compliant behaviors similar to those in animals. Importantly, controlling such systems is a huge challenge, and requires new approaches to control. The realities of controlling such a dynamic system is pushing researchers to explore machine learning approaches for dynamic systems, develop ultra-fast simulation tools, and new optimization approaches.

Key Challenges: Due to their small size, the companies building legged platforms are forced to focus on the design, mechanical hardware, motors, and power electronics required to realize a walking machine. Mostly, the legged robots have no inherent autonomy: they are operated over a standard hobbyist RC transmitter/receiver and hence must be within line of sight of an operator. In the rare cases where the robot is equipped with perception, it is not effectively protected, meaning it is easily damaged in falls. The mechanical and electrical reliability of legged robots needs to be enhanced. Anecdotal evidence suggests that their mean-time-to-failure is on the order of four to six hours. The present companies are not integrated with those that produce robot arms and hands. Hence, in terms of mobile manipulation platforms, legged robots are still in their infancy.

Legged robots need to be equipped with perception and computing packages that are much lighter, smaller, and less energy-intensive than those used in autonomous vehicles (AV). While an AV can carry a 100

kg load of GPUs, SSDs and networking switches with a multiple KW power budget, the corresponding weight, and power budgets for legged machines in the 30 to 100 Kg range are much much smaller. Meeting these needs requires innovation on both the hardware and algorithmic sides of the sensing, reasoning, and acting pipeline of intelligent machines. The big AV companies, such as Waymo and Cruise, are registering their vehicle into pre-built maps. They are not building multi-layer maps, in real-time, on the basis of onboard sensors. For robots to go into unstructured areas such as disaster zones, they must build the maps in real-time.

- 5-year goals: A legged robot can fall off a three-foot platform, remain operational, and right itself. It can walk at 0.8 m/s on level, uncluttered ground for at least five hours. It has enough real-time perception, mapping, and reasoning capability to autonomously navigate a typical university campus while remaining on sidewalks.
- 10-year goals: A legged robot can tumble down a flight of stairs, right itself, and continue its mission. Given an approximate map of an environment, legged robots are able to autonomously and robustly synthesize and execute a given mobility or manipulation task. This may potentially require navigating through and exploring multiple floors of a building or modifying the environment in some way to complete a task, such as lifting debris to search for survivors. A pair of legged robots should be able to coordinate their actions and build a semantic-metric map of a neighborhood that had been severely damaged stricken by gale-force winds. The map would allow first responders to do X, Y, and Z.
- 15-year goals: Humanoid robots are able to operate autonomously and robustly in completely unstructured, dynamic environments. This will require perceiving and understanding their environment to continuously plan in order to achieve some global task objective. In particular, they must be able to account for unexpected environmental changes while walking and be able to replan in the face of initial failure.

4.3 Grasping and Manipulation

Grasping: Robots have speed and strength far superior to human hands, but they cannot reliably grasp unfamiliar objects -- robots remain remarkably clumsy. Almost all applications, from manufacturing to service to security, would benefit if robots were to achieve the ability to grasp any object among a diverse range of shapes and sizes from rigid to deformable and under a variety of frictional conditions. Despite over 40 years of research, this problem remains unsolved (Rimon and Burdick, 2019). The difficulty stems from the inherent uncertainty in physics, perception, and control. Sensor noise and occlusions obscure the exact shape and position of the object in the environment, and the object's physical properties, such as center of mass and friction, cannot be observed directly. Many research groups are experimenting with new gripper and suction hardware, but still missing are algorithms and software that can rapidly compute and execute robot grasps that are robust to uncertainty. Related challenges include placing fingers around an object to contain it ("caging"), using inertia or gravity to hold parts, finding, grasping, and extracting a specific target object that may be fully or partially occluded among heterogeneous objects (sometimes called "mechanical search") (Danielczuk *et al.*, 2019).

Manipulation: Once an object has been securely grasped, it can be transported and dropped into a bin or box. However, for many applications such as inspection, assembly, and machining, robots must also manipulate the object, bringing it into new positions and orientations which may require re-grasping, applying forces, or inserting it into an assembly. More complex manipulation involves grasping tools such as keys and screwdrivers and inserting and applying appropriate forces and torques. "Dextrous manipulation" is the deliberate changing of the position or orientation of the grasped object within the hand, as is typically done when picking up a key and inserting it into a lock or when solving a Rubik's cube with one hand. Manipulation can also involve more than one object or hand. The World Robotic Summit runs an assembly competition¹ to spur the development of innovative

¹<https://worldrobotssummit.org/en/wrc2018/industrial/assembly.html>

planning and control algorithms to solve basic assembly tasks including simultaneous multi-part insertions and mounting of belts on pulleys.

Robot grasping and manipulation are relevant to almost all aspects of manufacturing, where object shapes are known and motions are often repetitive. Grasping can facilitate e-commerce as described above and sorting recycling and search and rescue operations, all where object shapes are not known in advance. For mobile robot manipulators in home and service applications, additional uncertainty in perception and imperfection in control increases the difficulty of grasping, for example, to declutter floors in homes, machine shops, and retail stores. For agricultural robots to harvest produce like strawberries and for surgical-assist robots, grasping is extremely challenging due to the delicate, deformable, and viscous nature of produce and human tissues.

4.3.1 Key Questions

Grasping: Several research groups and companies are exploring novel approaches for robot grasping and progress has been made in competitions such as the Amazon Picking Challenge (Morrison *et al.*, 2018; Hodson, 2018). Grasping from bins where object shape is known and all objects are identical has to some degree been solved. Approaches to grasping unfamiliar objects fall into four categories -- 3D registration, hardware-centric, learning from demonstration, and reinforcement learning methods. 3D registration methods require a priori 3D object models and generally only work when the object is rigid and its pose is known explicitly. Hardware-centric methods explore novel gripper designs, which for example, could include soft, multi-chambered, inflatable fingers that extend and curl when some of the chambers are inflated. Such grippers can work well with certain object classes.

New approaches to robot grasping are empirical (Bohg *et al.*, 2014), learning to grasp based on data from physical (Pinto and Gupta, 2016; Levine *et al.*, 2016) or simulated (Mahler *et al.*, 2017; Mahler *et al.*, 2019) grasp experiments. This data is used to train control policies or hyper-parametric function approximators such as deep neural networks, which show promise to generalize to unfamiliar object shapes. Grasping performance is measured by the “3 R’s”: Rate, Reliability, and Range

(the diversity of objects and materials) (Mahler *et al.*, 2018). For rate, research is needed to reduce the time required for each sensing, computing, and robot motion cycle to match or exceed human performance (e.g., 500 grasp attempts per hour (7.2 seconds per cycle) by a human box-packer in a typical fulfillment center). To increase both reliability and range, deep continuous learning has the potential to improve grasping reliability using data from thousands of robots connected via the Cloud. Learning from demonstration and reinforcement learning both require substantial data collection and data cleaning for each new environment or task. The collection and postprocessing of sufficient real-world data is currently too costly to be practical. Public databases to support the development and fair comparison of grasping algorithms do not exist.

Roadmap Goals for reliable grasping:

- 5 years: Inexpensive robot grippers (parallel-jaw grippers with two contacts and suction cups with one contact) could be capable of grasping novel rigid objects with planar faces from cluttered bins with reliability approaching that of humans and will be used in e-commerce and manufacturing.
- 10 years: Inexpensive grippers could be capable of grasping a broad array of rigid and deformable objects from cluttered bins with reliability exceeding that of humans. Robots could also be capable of reliably locating and extracting specified target objects from bins.
- 15 years: A broad range of grippers could be capable of reliably grasping and extracting any rigid or deformable object from a cluttered bin except those that are extremely adversarially-shaped (eg, Apple Airpod Pro earbuds).

4.3.2 Manipulation

In October 2019, OpenAI announced that it had used deep learning with massive amounts of simulation data to train a single five-fingered robot hand to solve a Rubik's cube without putting it down. However, in physical experiments, it was successful only 20% of the time.

The human hand is generally capable. A robotic equivalent, or superior grasping ability, would avoid the added complexity of robot interfaces on objects, and provide a sensate tool change-out capability for specialized tasks. Dexterity can be measured by the range of grasp types, scale, strength, and reliability. Challenges include fundamental 1st principles of physics in the development of actuation and sensing. Other challenges include 2 point discrimination, contact localization, extrinsic and intrinsic actuation, back-drivability vs. compliance, speed/strength/power, hand/glove coverings that do not attenuate sensors/motion but are rugged when handling rough and sharp objects.

4.3.3 Full immersion, telepresence with haptic and multi-modal sensor feedback

Telepresence is the condition of a human feeling they are physically at a remote site where a robot is working. Technologies that can contribute to this condition include fully immersive displays, sound, touch and even smell. Challenges include 1st principles of physics in the development of systems that can apply forces to human fingers, displays that can be endured for long periods of telepresence immersion, and systems that can be used by people while walking or working with equipment concurrently with the telepresence tasks.

- 5 years: Robots with two simple hands covered by tactile sensor arrays will be able to perform grasp adjustment and re-grasping of simply-shaped objects without putting the object down.
- 10 years: Robots with medium complexity hands covered by tactile sensor arrays will be able to perform dynamic grasp adjustment without relinquishing its grasp (a.k.a. re-grasping).
- 15 years: High complexity hands with tactile sensing arrays with densities and sensitivities approaching that of humans will be capable of high-speed whole-hand grasp acquisition of novel objects and fine control enabling execution of dexterous manipulation tasks with skill and reliability approaching that of humans.

Key Challenges:

- Develop articulated hands covered by tactile sensors that are robust enough to be used for thousands of hours without repair.
- Develop techniques for identifying minimal unbiased sets of data for learning grasping and manipulation tasks for a given class of objects and manipulation tasks.
- Efficient retraining (to cross the reality gap) of neural-network-based controllers with real data to allow them to work robustly in the real world.

4.4 Perception

Perception is what grounds robotics in the physical world. It would be impossible to perform most daily tasks without vision, haptics, tactile perception, and hearing. Robotic perception also includes sensors that are not biological analogs, but instead are designed for specific situations and tasks. For the purposes of this section, perception is taken to include the acquisition and interpretation of data from sensors which includes sensors that produce images (RGB, IR, Depth, plus medical modalities), domain specific modalities such as OCT or radar, haptic and tactile perception, sound, and potentially other unstructured information channels.

4.4.1 State of the art

Computer vision is often viewed as the primary sensory modality in robotics. It has manifold uses -- it can be used to compute geometry (where), identify structures in the environment (what), analyze movement, or support control. Some of these problems overlap with the objectives of the broader computer vision community, and others are unique to robotics.

No matter what the problem, computer vision (and image interpretation broadly) has been transformed by machine learning (in particular deep learning) over the past decade. For example, the error rate on

image classification performance on ImageNet, a standard benchmark, has gone from more than 25% a decade ago to less than 2.5% today -- a factor of two better than human performance. Similar trends have been observed on many related problems -- video activity recognition, object detection, image captioning, cancer detection, semantic segmentation, etc. At the same time, the ability to embed these capabilities into a low-power platform has been accelerated by the adoption of these technologies in mobile phones and automated driving systems. Finally, access to these capabilities has been transformed by the development of open-source tools such as Pytorch. It is now quite possible for someone with rudimentary skills in Python to read online tutorials, download code, and field a state-of-the-art vision system. Two examples of use of machine learning for training of models and subsequent estimation of object locations or human face pose are shown in Figure 4.3.



Figure 4.3: Visual recognition and post estimation has made significant advances (Source: Shutterstock)

While vision has advanced rapidly, many of these advances have been tied to the availability of large, curated and labeled data sets. While in principle these same advances can be applied to other tasks and other sensing modalities, progress has been more limited due to structural limitations (e.g regulatory limitations in the medical area, more limited deployment of IR or depth cameras, or simply lack of a cost-effective mechanism to obtain labels). This has motivated substantial interest in other means for obtaining data and/or transferring learned models from domains that are data-rich to domains that are data-poor.

Haptics has seen more incremental progress over the past decade. This is due to a number of factors including the lack of wide-spread

tactile sensors and displays as well as the complexities of embedding haptics and force-feedback in control. Much of the recent interest in this area has been driven by the desire to use haptics in commercial devices such as vibrotactile displays in mobile phones or in cars. Research is equally focused on understanding human tactile and haptic perception as on technology development. However, commercial tactile sensors and haptic displays are largely unchanged over the last decade.

4.4.2 Key questions

Computer vision will continue to advance in terms of breadth of problems and available tools due to the diverse set of application areas. However, robotics poses unique challenges for computer vision in terms of reliability and speed. Many of the latest computer vision systems have performance that is remarkable compared to the past but nonetheless operates at an error rate that is well below what is necessary to support reliable long-term robot operation. For example, an automated driving system that fails to detect a person in front of the car 1% of the time or which operates with a one-second delay, or a home robot that mis-grasps an object 5% of the time, or a medical robot that mistakes the liver for the spleen 1% of the time is unacceptable.

There are several key areas of progress in perception needed to advance robotics. These include the following:

- **Active Task Performance from Video:** Advances in activity recognition have led to impressive capabilities for recognizing large ranges of human activities. However, moving from observing an activity to performing similar activities requires far more fine-grained representations that support active control of action
- **Active Perception:** At the same time, computer vision explores data in a passive manner -- it does not take advantage of the ability to actively sense and/or capture redundant information. Part of creating systems that act in the environment will be creating systems that are able to actively observe to improve their performance.

- **Complex, high dimensional inference:** The broader computer vision community doesn't always concern itself with classes of problems relevant to robotics. For example, predicting grasps on objects from images is a high dimensional continuous problem. High-performance approaches and architectures (and data sets) for such problems will likely differ from those popular for recognition or detection tasks.
- **Open-world Performance:** Most computer vision systems adopt a closed-world assumption -- because they are learned from data, the data set represents the totality of examples the system is trained for. Robotics will often be faced with stimuli that have never been experienced, or task variations that are entirely new. Being able to generalize to new contexts and tasks is an open problem.
- **Integratable with Systems:** To integrate vision with other systems, it they need to be able to provide an assessment of their internal performance. This includes both methods for verifying or validating a vision component or vision-based system, and methods for systems to return something related to their reliability and uncertainty.
- **Systems Structure:** It is possible to perform reinforcement learning of a task from images in an end-to-end fashion. However, such an approach is not amenable to transfer to similar tasks or similar contexts. Conversely a more traditional approach would be to separately train computer vision modules from action modules. However, adapting modern learned computer vision modules to action or planning suffers from the reliability limitations outlined above.

4.5 Planning and Control

Robots of the future will need more advanced control and planning algorithms capable of dealing with single and multi-agent systems with greater uncertainty, and larger numbers of degrees of freedom

than current systems can handle. They will need to work safely and robustly in all settings -- ranging from fully autonomous operation in extreme environments to collaboration with humans at home or at work. Robot manipulators on mobile bases will need to efficiently plan and consistently perform fine manipulation and grasping tasks in unstructured and constrained environments. These robots might have a dozen or more degrees of freedom. Anthropomorphic humanoid robots, on the other hand, will have many more degrees of freedom to control and coordinate. At the far extreme, multi-agent and swarm robotics, while physically decoupled, require the coordination of a few to thousands of agents.

While in the past, control and planning were considered separate problems, modern control and motion planning are increasingly addressed in unison. Efficient planning methods that consider low-level controllers of their agents (whether arms, rovers, drones, etc.) and their tasks (manipulation, locomotion, flight, etc.) will use new computational techniques (including sampling based planning methods and optimization-based approaches) to effectively search the relevant high-dimensional spaces that define their environments and interactions.

4.5.1 Task and Motion Planning Under Uncertainty

Robots use sensors to observe the environment and situate themselves in it, and then plan actions to attain a goal configuration. Due to the lack of precise sensors, algorithms must be designed so that robots operate safely and robustly in the presence of uncertainty. While progress has been made in recent years, current methods can only handle simple tasks in fairly structured environments. More research is needed to develop algorithms for planning that can handle realistic problems in unstructured environments. These methods must be capable of real time operation in close proximity to and cooperation with humans. They need to provide safety and robustness guarantees while accommodating incomplete, inaccurate, and intermittent sensor data. Finally, although they have traditionally been studied separately, a principled integration of task and motion planning incorporating uncertainty is required to reach the level of autonomy needed for robots to become useful partners in unstructured settings such as the home.

4.5.2 Manipulation

Manipulation and grasping are fundamental capabilities for operating in the physical world – they are needed to open and close doors and drawers, to pick up, move or push objects, to use tools, to manipulate a steering wheel, or to otherwise reconfigure or interact with the environment. Current algorithms can only handle relatively simple scenarios, such as low degree of freedom problems with small, regular geometries and quasi-static motion. Research is needed to develop grasp planning and metrics for complex and unique geometries. This is increasingly important as end-effector technologies improve, presenting a larger range of grippers, including multi-fingered hands. Improved techniques are also needed for contact tasks, for manipulating deformable objects, for non-prehensile actions and tool use, and for dynamic motion. Strategies for robustness and failure detection and recovery are required for safe and secure operation.

Mobile manipulation robots (mobile robot platforms equipped with a multi-link robot arm and end effector) are becoming increasingly common. These devices inherit all of the difficulties associated with manipulation and bring extra difficulty due to their high degree of redundancy (e.g., moving the hand can be accomplished by moving the mobile base, reconfiguring the arm, or by a combination of these). To be efficient and effective, we will need methods to optimize motions with respect to task requirements (e.g., to move a heavy object, a qualitatively different motion might be used than for manipulating a fragile object), including safety, since these systems will increasingly be deployed in environments shared by humans.

4.5.3 Complex and Dynamic Environments

Dynamic environments encompass manipulation tasks in sensitive environments, with humans or other robots, and moving obstacles for which the robot does not have explicit knowledge of their underlying motions. Currently, strides have been made in modeling dynamic environments on a small scale with one or a few objects and agents in the environment generally following known trajectories or having repetitive behaviors that can be modeled using simple process models. Robots

in these low-dimensional environments can effectively plan over long time horizons. Yet, a challenge that needs to be addressed is scalability (numerous, heterogeneous dynamic objects and agents) and uncertainty (complex or unpredictable dynamics) that may require re-planning and adaptation by robot systems in real-time.

Extend planning and control methods to consider more complex environments, including highly deformable and uncertain environments. Robots based on a greater variety of mechanisms (e.g., operating inside a body, emergency response scenarios) can autonomously plan and control their motions in a variety of challenging applications (e.g., operating inside a body, emergency response scenarios).

Constraints to robot control planning present themselves in many different forms, whether they are physical constraints on a robot's reach, obstacles that constrain their workspaces, force constraints when interacting with sensitive materials, power/resource constraints of the robot, or dynamical constraints limiting robot actuation. Currently, efforts in constrained optimization approaches have been used to demonstrate effective optimization of these tasks for short tasks and small motions, in static environments with reasonable certainty. However, situations such as in surgery, service, and manufacturing involve long durations, fluid sequences of tasks, and dynamic environments. The next work in constrained optimization for robotics will be to roll the constrained tasks effectively into planning algorithms that provide continuous and connected motions, that can anticipate and react to dynamic constraints, and do so over long periods of time where its performance remains stable.

4.5.4 Control of high-dimensional, highly dynamic, hybrid systems

This area is currently being driven by the legged locomotion area where models with multiple dynamics phases arise from alternating leg contact with the ground and high-dimensionality arises from the large number of degrees of freedom (40 or more) in a humanoid. Model-simplification methods employed widely today, in order to “overcome these obstacles,” come at a cost: they limit the flexibility and capability of control solutions to yield dynamically stable, terrain robust, energy efficient

gaits. Model simplification techniques also require much hand-tuning of the feedback loops on the actual robot before stable operation is achieved.

Much faster trajectory optimization is required to allow robot designers and control engineers to explore the dynamic capability of models with multiple dynamic phases and without predetermining the order of traversal of the phases. The methods should readily handle non-trivial state spaces, such as $SE(3) \times R^3$. For control algorithm design, stability assessment techniques need to grow beyond Lyapunov and Poincare methods so as to provide disturbance-to-state stability with provable guarantees.

In terms of safety, reachable set analysis can only treat systems with polynomial dynamics and even then, dimension six is a stretch. Current efforts to overcome these limitations must be greatly extended, such as Kupman linearization, lifting the dynamics to an infinite-dimensional linear system where semidefinite programming can be applied. There are budding efforts to extend formal methods from Computer Science to Control Systems, where differential equations rule. Linear temporal logic is being extended to Signal Temporal Logic which recognizes the existence of real numbers. Control barrier functions have been introduced as an extension of barrier certificates in order to provide probable safety; in order to maintain as much performance as possible, QP-CBF-CLFs are being developed. Much remains to be done in regards to building control synthesis solutions around these promising directions.

4.5.5 Planning and Control for Sliding Autonomy

Today's robots are often limited in the scenarios and environments in which they can successfully operate autonomously or semi-autonomously. We need control, planning, and decision-making methods that can apply to a variety of robot types (e.g., soft/flexible robots, high-DOF or untethered medical robots, emergency response robots) and challenging environments (e.g., operating inside a body, emergency response scenarios). Depending on the needs of the domain, these methods should support sliding levels of autonomy, from reduced autonomy with substantial human input to fully autonomous systems. These methods

should be able to consider highly deformable and uncertain environments and consider application-specific objectives. These methods can also enable the design of novel robots, and we should also design robots to effectively operate with the controllers/planners/decision methods that we are available.

4.6 Learning and Adaptive Systems

In the future, robots will no longer be used for executing only a specific single task, but rather will be faced with thousands of different tasks that will rarely be repeated in ever changing environments. It is infeasible to preprogram all possible tasks and scenarios in future robots, and robots will need to learn and adapt by themselves or with the help of humans. They will need to automatically adjust to stochastic, dynamic and non-stationary environments and compensate for hardware degradation due to wear and tear. Machine learning and related AI techniques hold the promise to achieve this high degree of autonomy and reliability.

4.6.1 Promise and Challenges of Machine Learning

In recent years, machine learning using neural networks has demonstrated impressive achievements in computer vision, speech recognition and in playing games such as Go and Pong. For a number of these tasks, the use of machine learning techniques results in remarkable performance, in some cases exceeding human performance. We are beginning to observe the potential of applying machine learning to areas of robotics; however, the exuberance for applying machine learning to build better robotic systems should be tempered by the realization that a number of research problems still need to be solved to fully realize its promise.

In image processing and speech recognition tasks, predictions can be made on datasets that have a clearly defined notion of ground truth that machine learning algorithms are able to continuously optimize towards to achieve improved performance. The predictions do not impact the input data distribution and the target goal remains stationary. On a robot, however, predictions become complex physical actions that

quickly change its surroundings, influencing the consequences of the next action. In some cases, actions can irreversibly affect the state of the environment and significantly alter the overall task, such as knocking over and breaking a needed object. Developing techniques to have a robot learn rapidly and safely under these conditions is a critical research area still in its nascent stages.

- High-dimensional continuous state and action spaces for robots
- Incomplete information and partial observability of state
- Difficulty of obtaining training data
- Current ML algorithms have high sample complexity
- Current ML lacks transparency and ability to connect low-level of latent representations to explicit models/high level representations

4.6.2 Robot Learning Approaches

The most immediate application of machine learning for robotics is to leverage better perception algorithms using neural networks for object recognition and pose estimation in robot applications (see Section 4.4). However, maximizing the benefits of learning to robotics requires learning full policies that integrate planning and control in addition to perception (see Section 4.5). One example of robot learning leverages imitation, giving robots the ability to recognize and reproduce human actions. Imitation learning is a technique that reduces the complexity of high dimensional state and action spaces for learning through direct observation of both good and bad trajectory examples. Future advances in imitation learning have the potential for novice human users to be able to customize robot behavior in a very natural manner.

Models are essential tools used to design and program robots. A model describes essential predictive information about the surrounding environment and the influence of robot agents on the environment and are critical to provide performance and safety guarantees for robots. While classical robotics relies on manually generated models based upon expert human insights, future robotic systems will be able to use machine

learning to build accurate models in a data-driven manner. Using learned models in robotics will be critical for the future development of robots that incorporate soft and compliant materials exhibiting highly nonlinear and complex dynamics. Another potential application of robot learning will be to learn better predictive models of environmental interactions with high uncertainty, such as when robots interact with humans (see Section 4.8).

Another approach to robot learning is to have robots automatically discover optimal behaviors through trial and error interactions with their environment. Reinforcement learning does not require an explicit teaching signal but rather robot behaviors are tuned by optimizing a scalar objective function that measures the overall performance of the robots. However, current reinforcement learning techniques require an enormous number of trials to learn, predicating the need to train on simulators rather than on physical robots. Unfortunately, the gap between present-day simulators and on fielded robots is vast. Research advances to mitigate this gap and the development of learning techniques that require far fewer examples to learn will be necessary for the future practical development of robots that use reinforcement learning.

4.6.3 Datasets and Benchmarks

In contrast to other machine learning domains, robot learning suffers from a variety of complex real-world data problems. It is difficult to gather large amounts of training data of trajectories on physical robots; real-world training time is limited and relatively few executions of a task can be generated during individual experiments. These episodes are also typically quite noisy and cannot completely cover all possible scenarios and every reaction to external stimuli.

It is thus imperative that more comprehensive and openly available datasets be made available to the broader research community to spur future progress in robot learning. Concomitant with these datasets, a new set of benchmarks, evaluations, and associated infrastructure will also be needed. These datasets and benchmarks should support different robot systems and a variety of tasks, with realistic and changing environments. Open-source tools and platforms need to be provided that

allow for easy data and code sharing, enabling researchers to exchange and quantitatively compare their approaches to more rapidly drive future innovation.

4.6.4 Vision for 5 Years

- Applications of robot learning to existing conventional robot platforms
- Continuous performance improvements on defined tasks in fairly constrained environments

4.6.5 Vision for 10 Years

- Increasing diversity of robot platforms using novel materials and architectures, micro/large scale and increasing numbers of robots, requiring robot learning to handle increasing complexity
- Development of large scale and realistic datasets, benchmarks, etc.

4.6.6 Vision for 15 Years

Seamless operation of robots in collaborative environments, e.g. Level 5 autonomous driving deployed and commercially available to wide segment of population

4.7 Multi-Robot Systems

4.7.1 State of the Art: Innovations and Key Questions

In many tasks there is strength in numbers from the perspective of both robustness and effectiveness. Being able to deploy a team of robots, as opposed to a single robot, has distinct advantages in applications such as planetary-scale exploration, city-scale security, and large-scale warehouses. If a robot breaks down, the mission will still continue, by spreading robots over a large spatial domain, more effective coverage is achieved, or by distributing capabilities across multiple platforms, more flexible and adaptive robot systems are produced.

Multi-robot systems have been successfully deployed in manufacturing and warehouse management, disaster monitoring, construction, and agriculture. Much of the work in multi-robot coordination has been inspired by nature. Evolutionary algorithms and decentralized intelligence have produced complex behaviors, e.g., for making teams of robots assemble geometric shapes, covering areas, tracking boundaries, or finding and tracking intruders. These are challenged when required to converge in short periods of time and have generally been applied to homogeneous agents, whereas realistically, deploying heterogeneous agents is more practical and adds flexibility to the planning process. Centralized intelligence-based methods are limited in their ability to work in real-time, with individual robots needing local controllers to react quickly to unexpected events.

Despite significant recent advances, a number of research issues remain to be solved for large teams of robots to be reliably and robustly deployed in real environments over sustained periods of time. We expect research to focus on real-time, scalable coordination methods that take advantage of the full spectrum from centralized intelligence to local behavior, formal methods that prove convergence to some optimal behavior, and planning for heterogeneous teams for executing complex sequences of tasks.

4.7.2 Key Challenges

Distributed Control and Decision Making: The design and deployment of multi-robot systems is fundamentally challenging because the individual robots typically have access to limited information – usually what they can measure themselves and what neighboring robots may share. As such, distributed decision making algorithms must be developed that take this limited information into account in such a way that the desired global behaviors emerge from a collection of local rules. A number of examples of such mechanisms are present in the literature, but we still lack an effective framework that can take as input high-level specifications of what the team should be doing, and produce low-level, distributed control algorithms. The development of such a framework is the central challenge for multi-robot systems to become truly flexible and useful across a number of application domains.

Hybrid Decentralized/Centralized Mechanisms: Traditionally, small teams of robots have been controlled by a centralized decision maker while larger teams are programmed to act in a truly decentralized manner. However, in practice one may desire a mix of these two extremes, where a centralized (possibly, cloud-based) node may gather information over time and intermittently inject centralized information into the system. There are currently no effective abstractions or systematic algorithms for describing or taking advantage of hybrid information exchange mechanisms where the information flows may be operating at different time scales. Subsequently, no systematic understanding exists for what information needs to be shared in a centralized manner and what information can be kept local. A related unresolved question is how a human operator should be interacting with teams of robots (e.g., a farmer interacting with a team of autonomous tractors or a pilot controlling a large team of unmanned aerial vehicles). In these two scenarios, it is fundamentally not understood what constitutes effective interaction modalities, both from a cognitive workload and from a bandwidth management perspective. As multi-robot systems are deployed, these questions must be addressed for people to be able to effectively use these systems effectively.

Heterogeneous Robot Teams: One of the benefits of deploying robot teams is that one can distribute capabilities across the different robots. For example in a disaster response setting, aerial robot drones provide a high-level overview of the environment, while robots on the ground may be able to navigate under rubble. The resulting system is a heterogeneous team. Heterogeneity includes dynamical configurations, sensing capabilities, spatial footprints, or behavioral strategies. But, beyond isolated examples (mainly air-ground coordination), we still do not fully understand how to take advantage of heterogeneity in a fundamentally sound manner. What types of robots are needed, given a particular task? How heterogeneous is a collection of robots? How heterogeneous should a team be to provide maximal flexibility in terms of the tasks it can perform? We lack both a framework to pose and answer these questions systematically.

Communication and Sensing in Multi-Robot Systems: In order to realize the full potential of a multi-robot system, we need a foundational

understanding of the interplay between sensing, communication and acting in these systems. For instance, communication between robots or between robots and humans or between robots and the cloud are degraded due to several factors such as path loss, shadowing or multipath fading. Furthermore, path planning at every robot not only affects its own sensing but also impacts its ability to communicate. This is a challenge for designing robust decision-making strategies for multi-robot systems and is an open problem in networked robots, when considering realistic communication links.

4.7.3 A Roadmap for Progress in 5, 10, and 15 years

With the appropriate research investments, we expect that the next 5 years will see significant progress in distributed control and decision making leading to progress in the automatic generation of distributed control algorithms from high-level, global specifications. We will see the beginnings of a formalism for handling multiple time scales leading to the first hybrid decentralized/centralized multi-robot systems at scale. Richer examples of effective interaction modalities will be prevalent that allow small numbers of humans to interact and control large-scale robot teams. We will also see a formal characterization of heterogeneity across a number of dimensions, including functional, spatial, and temporal aspects. Finally, we expect that more effective models for trading off mobility, sensing, and communications will emerge.

Over a 10-year period we expect to see robustly deployed large-scale teams in real environments and ‘teams of teams’ that learn and collaborate with each other via cloud-based architectures. We expect to see commercially available, human-centric swarm systems available (e.g., like one can buy a single drone today, we expect you will be able to buy a team of drones that will work out of the box). We expect that heterogeneous solutions to complex tasks via dynamic team composition will emerge. These systems will optimize power consumption across sensing, mobility, and communications.

Over the 15-year period we will see commercial penetration and large scale deployment across a number of industries, including agriculture, manufacturing, warehousing, and environmental monitoring. We expect

a true ‘Internet-of-Robotic-Things’ ecosystem to be in place along with robustly deployed, mix-and-match collaborative robotic teams at low cost.

4.8 Human-Robot Interaction

Across many application domains, robots are expected to work in human environments, side by side with people. Interactions between robots and their users will take many forms, from a trained operator supervising several industrial robots, to an older adult receiving care from a rehabilitation robot, to a child safely practicing social, cognitive, or emotional skills with a readily available socially assistive robot. The users will vary substantially in background, training, physical and cognitive abilities, and readiness to adopt technology. Robotic products are expected to not only be intuitive, easy to use, and responsive to the needs and states of their users, but they must also be designed with these differences in mind, making human-robot interaction (HRI) a key area of research.

4.8.1 State of the art

Over the past decade, the state of the art in HRI has advanced significantly. Social HRI systems have been used to improve the quality of life for children with autism and for older adults with dementia. Learning from demonstration has enabled robots to learn new skills by being shown how to do it by a person. Teleoperation interfaces have been developed to monitor and control larger numbers of robots, while adding augmented and virtual reality control to more traditional screen-based interfaces. Assistive robot systems, from exoskeletons to rehabilitation robots, have been deployed in the real world. Brain-control interfaces (BCI) have allowed people with quadriplegia to feed themselves with a robot arm by thinking about moving the robot. Despite this progress, there is a need for much more research and development. Most current HRI systems are only used for short periods of time, with an inability to adapt to the person or people using the system. A diverse set of possible use-cases for human-robot collaboration is shown in Figure 4.4.




Figure 4.4: Examples of the many different aspects of interacting with robot systems (Source: Shutterstock)

Current Issues *How do we develop robots that understand humans?* People have trouble understanding each other’s views, beliefs, intentions, and actions. Machines are much worse at it than people. For robots, much of what drives human behavior is hidden (e.g., histories, hopes, dreams) and the observed behavior may be confusing (e.g., mixed signals, sarcasm). Human behavior is also highly varied and diverse across numerous dimensions including context, culture, familiarity, fatigue, etc. resulting in unpredictability. The quest to understand people better in order to interact more effectively spans robotics, machine vision, speech and other signal processing, and machine learning. Robotics offers a unique enabler by providing embodied social partners that can be physically present around people in order to collect data from natural interactions in order to further research into understanding people better. However, such large multi-modal datasets require extensive time and resources to annotate and analyze; there is a great need for effective open-source tools for automated multi-modal human behavior data annotation. Further advances in sensory modalities (e.g., unencumbering wearables) and signal processing, as well as many more data sets, will be critical for progress in this area.

How do we develop robot systems that can interact at a variety of timescales? Some areas of robot control deal with very fast interactions with the environment, but HRI is unique in requiring

a broad spectrum of temporal dynamics: interactions that happen very quickly (a wink or a twitch of the mouth), interactions that happen very slowly (gradually getting used to a pattern of behavior), and interactions that change unexpectedly (due to context or intent inaccessible to the robot). To be effective work and social partners, robots must be able to perceive, understand, react, and adapt to such interaction at the right timescales, from providing an interaction repair immediately based on a missed social cue or almost dropping a fragile object being carried together. Humans can make fast impressions and snap judgements that are difficult to shift once formed, can be open to adapting to valued collaborators and partners, and can be tolerant, empathetic and compassionate to vulnerable social partners. Robots will have to fit into appropriate roles in HRI, project realistic expectations, motivate interactions, and manage the inevitable occasional failures. Finally, robots meant for long-term interaction, such as in-home service and companion robots, must utilize new types of machine learning to adapt to each individual user and the dynamics with that user over time.

How do we foster appropriate levels of trust in robot systems, so that systems are used correctly -- and not over- or under- trusted? In order to trust robots, users must intuitively understand the robot's capabilities and intent and feel they have appropriate authority and autonomy relative to the robot. To relate to people, robots must appropriately assess what users to trust and how much to trust them, then must communicate their level of trust to users. While some research has been conducted to determine what factors influence robot trust, there is a need to develop more robust models to enable the development of trust.

How do we design HRI to facilitate user acceptance? Acceptance and adoption of robots in human society is part of the age-old broader challenge of societal acceptance of technology. Unlike automation, HRI keeps the human in the process. The acceptance and adoption of HRI therefore hinges on its ability to meet the users' expectations and needs. This is a complex challenge, since sometimes repeatable and reliable robot behavior, as specified, may become boring while unexpected robot behavior, even if inaccurate, may be entertaining and preferred in some contexts. The entertainment industry has historically

driven human expectations of intelligent systems to be unreasonable. A careful design that generates the right expectations is essential.

4.8.2 Key challenges to progress?

Interactions with humans are varied and often difficult to evaluate objectively. There has been some success in prior research with specifying design guidelines or boundaries, but this is insufficient to give us the end goal at this time. People do not know what the value of a robot will be for people. There is not much consistency in what kind of HRI is required across different domains -- in some cases, people are fully teleoperating a robot; in other cases, people are occasionally checking in but letting a robot be generally autonomous. The same diversity applies for social interactions in HRI. Beyond these issues, Hollywood and other sources can create unrealistic expectations about interaction and human awareness.

Lack of robot platforms for HRI research: A major barrier to progress in HRI research is a lack of robot platforms that are: 1) designed for interactions with people and 2) affordable. Various areas of HRI are difficult or even impossible to pursue due to the lack of appropriate robot platforms. Most research platforms are still designed mostly for mobility (in 2D on the ground, in 3D in the air); interaction with people requires careful design of the robot's affordances for interaction: both the form (what the robot looks like) and function (what it can perceive and do). For researchers interested in humanoid robots for social HRI, there are no expressive and affordable humanoid robots available for HRI research. (The best alternatives are the Softbank Pepper robot, which is very expensive and has no facial features.) Similarly, there is a lack of affordable tabletop robot platforms for social and socially assistive HRI that can be used to develop, deploy and test sufficiently large user studies. Various potentially promising platforms emerge in startups (e.g., Jibo, Kuri) and then rapidly disappear. An investment in properly designed shared platforms would be a major enabler for HRI research.

Lack of available datasets from interaction scenarios: In an age when big data and machine learning are enabling major leaps across

many areas of research and industry, HRI data with real-world users are still very difficult to access. To advance work on HRI for the elderly, we must have data involving interactions with the elderly, analogously with children, with stroke patients, with individuals on the autism spectrum, and so on. Such data are very difficult to collect, and also difficult to share due to privacy constraints. As a result, few researchers work with real datasets and scalable studies. Robotics has the potential to be a platform for collecting unique and invaluable interaction datasets, but to make that possible, there is a need for facilitating access to data with real-world populations and then making it possible to share such datasets.

Lack of access to real-world evaluation domains: Setting up collaborations with real-world contexts, from nursing homes to schools to hospitals to inform HRI design and evaluation is extremely difficult and time-consuming, resulting in very few groups across the nation doing research with real-world populations. Most HRI research is being done with university students as study participants, resulting in biased results. An investment in test sites with real world users that would allow for data collections and comparisons of methods and results would significantly boost HRI progress. Ecological validity and transfer out of the lab - coming up with natural experiments that are constrained enough for robots to operate but realistic enough to allow us to make conclusions that transfer to the real world.

5 years:

- Robots are able to reliably determine basic context-appropriate behavioral and emotional expressions (e.g., handing an object, smiling, pointing, thanking...)
- Robots are able to maintain interactions with, learn from, and adapt to their users in the timeframe of months in semi-controlled settings.
- Augmented and interactive reality systems will allow subject matter experts (e.g., civil engineers, first responders) to exploit and manipulate incoming data from robots in real-time.

- Robots are able to provide explanations about their behaviors that help people to trust them.
- Robots are judged as largely acceptable in larger-scale research studies.
- Community testbeds, metrics and evaluation metrics for HRI systems allow for direct comparisons of different approaches.
- Robots are increasingly research-tested in structured real-world environments (e.g., classrooms, rehabilitation centers, retirement homes).

10 years:

- Robots can learn and update user and task models on the fly and handle perception, modeling, and adaptation in semi-structured tasks and environments, allowing them to adapt their actions to reasonably changing environments and users.
- Robots can handle contingencies in dialogue, adapt to user states, and seamlessly integrate social behaviors in semi-controlled settings (e.g., the lobby of a corporate building).
- Robots function in semi-structured settings for a year and accommodate the needs of multiple users.
- Robots are able to model people's expectations and prior understanding in order to provide only the most relevant information to foster trust.
- Shared autonomy systems can determine user goals and continually adapt and engage the user in supporting those goals and choosing the appropriate level of autonomy.
- User expectations (beyond safety) about how robots should function in human environments are developed, applied to deployed robot systems, and tested for effectiveness.

15 years:

- Robots use information from past interactions and publicly available data to adapt their behaviors to individuals.
- Trust is developed between people and robots in the same way as between two people.
- Robots are increasingly found in daily life in workplaces, public areas, and homes performing specific roles and tasks safely.
- Authoring and programming tools for robot interfaces support standards, scalability objectives, and address safety, privacy, and security.
- Robots can perceive, model and adapt to complex user behaviors, actions, and intent in semi-structured tasks and environments, and transfer learned models across domains and environments.
- Robots can adapt their interactions to diverse groups of users using dialogue and social interaction strategies in uncontrolled settings (e.g., public spaces, field and disaster settings).
- Robots can not only recognize but also predict contingencies, user error, and changing capabilities of human collaborators and take action toward preventing or minimizing their effects.
- Shared autonomy systems can integrate and fuse various forms of implicit and explicit user input, model user goals and error on the fly, and vary levels of autonomy as necessary while communicating with their users.
- Robots will maintain adaptive functionality in uncontrolled environments in the timeframe of several years with an arbitrary number of users with varying degrees of capability
- Norms and standards will be refined based on an understanding of long-term use and interaction with deployed systems toward seamlessly integrating robots into society and enabling safe, effective,

and acceptable robot functioning in uncontrolled environments with diverse groups of users.

5

Technology Context

Technology is evolving rapidly. Manufacturing is seeing a number of significant changes that impact production. Some of these big changes are:

- Additive Manufacturing
- Model Based Programming
- Configuration Lifetime Management
- Collaborative Systems
- Internet of Things / Industry 4.0
- Big Data and Analytics

These technologies are in fast evolution and in most cases there are significant programs in place that impact manufacturing. An important consideration for a roadmap is how these efforts may be leveraged to speed up deployment and reduce costs. Each of the efforts are briefly discussed below.

5.1 Additive Manufacturing

The area of additive manufacturing, also sometimes referred to as 3D printing, though only covering a subset of the process, has seen tremendous progress over the last decade. The US administration has launched the National Network of Manufacturing Institutes (NNMI) and the first institute is focused on additive manufacturing. The additive manufacturing genome project is trying to capture the idea of a fully integrated methodology from design to delivery where all the key data are captured in a “genome”. The latest annual report for NNMI is from 2018.¹

The key aspects considered for design are related to energy efficient manufacturing and bio-inspired systems from cellular structures to products that are energy optimized from creation to retirement. For materials the key aspects considered are “non-ad-hoc” processes and characterization of materials properties. For process advancement the key ideas promoted are multi-material delivery and deposition, next generation machines and improved real-time temperature control. For value chain improvement key ideas considered include advanced sensing, machine control, rapid inspection and standards. The additive manufacturing institute has been sponsored by DOD/NIST.

No doubt additive manufacturing is going to advance rapidly. General Electric has recently started to build components for aerospace using 3D printing and Boeing has a number of internal projects. One would expect to see many critical components and support structures to be manufactured using an additive process as speed and consistency ramps up.

5.2 Model Based Programming

Progress on software systems for automated planning, verification and code-generation has been significant over the last decade. Initial progress was driven by academic research but with limited complexity systems. Over the last few years major progress has been achieved through a

¹<https://doi.org/10.6028/NIST.AMS.600-5>

number of major projects. The most well-known is probably the Adaptive Vehicle Make (AVM) program sponsored by DARPA (Adaptive Vehicle Make, 2020), where the objective is to manufacture a military vehicle directly from the engineering design files. The project has since then become part of the Digital Manufacturing NNMI institute, which has significant support from Boeing. Several projects across the world, but very much dominated by the automotive sector, are driving automatic generation of software for manufacturing processes. As the project variation, while potentially large, is deterministic it is possible to design a process that is relatively deterministic. For Boeing where each airplane appears to have much large variability it is less clear how/if the same tools can be directly applied. The NNMI institute on Digital Manufacturing has yet to release a technology roadmap for general industries.

In Europe there are a number of major efforts underway as part of the Horizon 2020 program. Again most of the programs are driven forward by the automotive industry. The vision for Europe has been proposed by the HYCON network (Highly-complex and networked control systems, 2020) and the follow-up CPSoS² support action. More recently the big driver has been the Horizon 2020 - Factory of the Future program, which has its emphasis across design, manufacturing, deployment and maintenance. The program is funded at \$1.2B over 2014–2020. The roadmap is available online (European Union, n.d.). So far limited emphasis was been devoted to software generation for low-rate manufacturing processes.

5.3 Configuration Lifetime Management

There is a strong trend to deliver fully customized products across many different sectors. As an example Audi has 10³¹ different car configurations. Several industries such as energy, automotive, and electronics are moving towards business processes that are driven by configurations data. Every station in a manufacturing system defines its work based on the particular configuration of the final product. Audi has 128 different

²<http://www.cpsos.eu/>

configurations for the A3 steering wheel, which challenges assembly, supply chain and maintenance.³ As an example, BMW uses an RFID unit attached to the chassis as it moves through the assembly process to define the SOW for each work cell. An active RFID module can capture configuration data, calibration information etc., to allow for a highly adaptive process. There is a strong push towards 1-off manufacturing. The current market trend is towards full customization, which in turn requires 1-off manufacturing. The use of configuration specific data to drive supply chain and assembly in each cell allows for a lean process and higher flexibility in process definition.

Configuration Lifecycle Management (CLM)⁴ is the management of configuration definitions and individual configurations across all involved business processes. A CLM system is integrated into a number of central business processes such as:

- Product development
- Marketing communications
- Sourcing of materials and parts
- Manufacturing processes
- Maintenance and service processes

A CLM solution:

- Streamlines and aligns all processes in relation to configurable products
- Supports both the product definition's lifecycle and the ordered product's lifecycle
- Provides cross-area collaboration between otherwise separate business functions
- Users range from back-office engineers and financial controllers to service technicians and customers

³<http://www.clmsummit.com>

⁴<http://configit.com>

- Is used for planning, analysis and optimization as well as daily operations
- Aligns processes typically supported by other systems: PLM, CRM, ERP, various bespoke (typically outdated) systems, spreadsheets and text documents

5.4 Collaborative Systems

The manufacturing process is becoming more and more human centered. Humans play a key role in the management of ever-increasing complexity, for processes that require significant cognitive reasoning and rapid evolution in product definition or mix.

Future workers will use multi-modal interfaces, intuitive and user-experience driven workflows, to safely plan, program, operate, and maintain manufacturing systems. Mobile and ubiquitous technology will allow workers to remotely control and supervise manufacturing operations. New safety systems will allow full adaptation of worker-robot collaboration that will enhance competitiveness and compensate for age- or inexperience-related worker limitations. Dynamic reallocation of tasks and changes in automation levels will enable human-automation symbiosis and full deployment of the skills of the workforce. Enhancement and support of the workers' cognitive skills will become increasingly important to create human-centered workplaces.

Human-machine interaction has evolved significantly through new and emerging safety standards such as ISO 10218.6 and R15.06. The clear definition of models and methods for interaction allows design of systems at a much lower cost and with improved performance as seen for collaborative robotics. A major challenge is the need for application specific safety certification.

5.5 Internet of Things / Industry 4.0

According to Gartner Internet of Things is “the network of physical objects that contain embedded technology to communicate and sense or interact with their internal states or the external environment” (Gartner

Group, 2020). The emergence of highly interconnected manufacturing lines opens a number of opportunities. Primarily in terms of real-time access to data at all levels from line buffers over processing speed to quality control information. The interconnected factory opens new opportunities for cell, line, facility and enterprise access to performance data. It provides a new level of access to data about the process that can be used for optimization of the lean enterprise and downtime/diagnostics. Almost every item on a manufacturing line can be tagged and potentially have a network address. The risk of drowning in data must be considered and the need for strong analytics will be essential to ensure appropriate data is available. In addition, cyber security must be considered as an integral of the design, operation and maintenance of such connected manufacturing lines.




Figure 5.1: The many diverse areas covered by Internet of Things

There is an emergence of two different approaches to architectures for industrial IoT systems. One is through Siemens where the approach is termed Industry 4.0 (Siemens, 2019) while GE has a similar approach termed the Industrial Internet. The GE approach has been expanded

into an international consortium – the industrial internet consortium (Industrial Internet Consortium, 2020)

IoT/Industrial Internet/Industry 4.0 facilitate a major revolution in access to manufacturing data that can be used to optimize the production line. The German National Academy of Science estimate that the reduction in cost will be in excess of 35% and the reduction in the errors will be in excess of 20%. The number of possible use-cases for Internet of Things is vast and a small selection of them is shown in Figure 5.1.

5.6 Big Data and Analytics

We have seen a tremendous growth in the availability of sensors for monitoring of processes over the last decade. In addition, we have seen exponential growth in the availability of computer power for data processing. Figure 5.2 illustrates how Graphical Processing Units (GPU) have emerged as desktop mini-super-computers for advanced tasks such as signal/image processing.⁵




Figure 5.2: Evolution in computing power for CPUs and GPUs over the last decade

The progress on teraflops is only part of the story. The chip technology has moved away from only studying 7- or 10-nm production

⁵<https://github.com/mgalloy/cpu-vs-gpu>

technology to consider the evolution in the number of cores use to build processors as shown in Figure 5.3.⁶




Figure 5.3: Another perspective on trends in microprocessor evolution

The amount of data available per person has double every 40 month since 1980. Year 2012 the amount of data generated every day was 2.12 exabyte (2.5×10^{18}). It is anticipated that the big winner in terms of utilization of data will be in manufacturing due to improved process monitoring and optimization of the supply chain.⁷

See (Lee, Bagheri, & Kao, 2015) for a discussion of recent progress on big data architectures for manufacturing.

Big Data processing and the use of Graphical Processing Units (GPUs) has already revolutionized image processing. The area of machine learning termed deep learning⁸ has facilitated a new level of performance in image based diagnostics and recognition, which has

⁶<https://github.com/karlrupp/microprocessor-trend-data>

⁷<http://www.tcs.com/SiteCollectionDocuments/White%20Papers/Big-Data-Analytics-Manufacturing-0914-1.pdf>

⁸<http://deeplearning.net/>

motivated companies such as Facebook, Google and Microsoft to make major investments in these technologies. It is important to recognize that there is an abundance of data and processing power but this far limited progress has been achieved on turning data into actionable information. The biggest challenge remains model-based data processing for monitoring and controlling tasks in real-time.

Several technologies of direct relevance are mapped out in the Gartner 2019 Hype Cycle shown in Figure 5.4.




Figure 5.4: Gartner 2019 Hype Cycle for Emerging Technologies (Source: Gartner Group)

It is interesting to see nano 3D printing is included, but it is interesting to see that autonomous vehicles at levels 4–5 are considered 10+ years out. Light weight cargo drones are 5–10 years away and 3D sensing cameras are 2–5 years away and considered a mature R&D area by now. Nonetheless it does give an outside perspective on the maturity of different technologies. It is also worth noting the Gartner 2020 hype cycle for emerging technologies does not mention robots or robot technology.

6

Workforce Development

6.1 Introduction

Popular news reporting has helped create the impression that robotics technologies are proceeding at a rapid rate and in many industries will increasingly replace human workers. As with the entire history of industrial manufacturing and automation, there will continue to be shifts from unskilled to skilled labor. However, in reality the expanded use of robotics and automation will likely create many off setting skilled labor and engineering types of jobs. The workforce challenges that face all sizes of industry, but especially small to medium size enterprises, are complex and have limited the successful and cost-effective utilization of robotics and automation. Applications demand extensive use of engineers and specialized technicians to implement, start up, and sustain such operations. Unique combinations of embedded computing, software, and electronics skills, to name a few, are required and are generally expensive and in short supply. New ideas and programs are needed to address this the future workforce in this sector.

6.2 Strategic Findings

From their early beginnings as tele-operators and manipulators, robotic systems have served to *extend the reach of humans* in interacting, manipulating and transforming the world around us. Since then, the enormous growth in numbers, diversity and complexity has been driven by their usefulness in enhancing human manipulation capabilities over various spatial and temporal scales (nano to mega) and for automation of the 4D (dull, dirty, dangerous and dumb) tasks. At the same time, the applications-oriented field of robotics and intelligent machines has offered a means of tangible embodiment of ideas and algorithms for a host of scientific disciplines – system design, control engineering, computational science, and artificial intelligence – among others. The diversity of application arenas stands as testament to its interdisciplinary nature and ultimately its immense potential – however setting the scope of robotics activities and developing forward-looking roadmaps for the field is a useful exercise, especially given the blurring of the boundaries between robotics and its constituent scientific disciplines.

These drivers are already revolutionizing the robotics landscape with attention being focused on developing all facets of the research, development, educational and ultimately logistics and commercial infrastructure to support this enterprise. And the commercial industry-led driving impetus behind these ventures bodes well for the long-term success of these efforts.

However, in the midst of these developments, another revolution has been silently underway that is fundamentally transforming the landscape of robotics. The archetypical robotic system is being transformed into a system-of-systems created as heterogeneous collections of physical and information resources coupled together by intricate connections and interactions. The enablers of this revolution are the same science and technology drivers that have made robotic devices smaller, smarter, easier to use and more connected with each other, with people, and with the environment. None of these changes are remarkable in of themselves but the net effect has been to enable new capabilities and remove barriers in ways that previously inconceivable. The ramifications of this

transformation are so immense that to (mis)quote Thomas Friedman:¹

“Overnight, while the world slept the robotics field has been flattened.”

Over the past 5 years, Bill Gates’ vision of a “robot in every home” paving the way emergence of a new robotics industries with even greater potential to revolutionize the way we live. There are the striking parallels between the personal-computer and the personal-robot industries in their early years – in terms of the fragmented state-of-existence (diversity of platforms/software), the inflexible operational paradigms (monolithic solutions) and the newer hardware and software trends (modularity, open-source) that paved the way for the revolution. And, as in the personal-computer industry, the evolutionary pressures from *evolving application focus*, *rapid technological/scientific progress* and *technological crosspollination* are driving and constantly reshaping the landscape.

- First, the archetypical PUMA manipulators, central to the manufacturing-floor automation of the heavy industries, could be viewed as the equivalent of the mainframes of the past era. Today, the growth in robotic systems is focused more the non-manufacturing application arenas and principally in the service robotics sector. Even here, the high-cost specialized devices from computer-assisted surgeries (da Vinci comes to mind), space explorations (NASA Mars Rover), and military robots in hostile combat environments (disposal of roadside bombs in Iraq), and robots to assist the search of trapped miners, form only a small fraction principally due to low-volumes. The most significant growth comes from the low-cost, high-market volume domestic and personal robotics market.
- Second, technological advances in sensing/actuation/computing on one hand and improved fundamental scientific understanding and algorithmic implementation on the other have contributed significantly to the growth of robotic systems of various shapes, sizes and functionality. Modularity and standardization in hardware, software and tools and the coupling of commercial interest

¹https://en.wikipedia.org/wiki/The_World_Is_Flat

with open-source movement is beginning to reshape the robotics arena much in the same way that personal-computing has been transformed.

- Last, the technological cross-pollination that occurs with each new round of innovation, improves not only existing robotic systems but opens up other avenues where intelligent mobile robots can be employed, effectively creating new markets. For example, the student developing advances for a robotic unmanned ground vehicle could go on to develop your neighbor's robotic lawn mower. Or for that matter, the years of research on safe and stable teleoperation can serve to enhance driving feel in tomorrow's drive-by-wire automobile!!

6.3 Near Term Opportunities and Factors Affecting Immediate Deployments

The rapid increase in the number of formal undergraduate and graduate degree programs in robotics in recent years motivates the need to develop a model robotics curriculum. Such a curriculum consisting of unified courses and projects utilizing standard robotics software and hardware will accelerate the creation of robotics programs to support the ever-increasing demand from the industry for engineers with multidisciplinary skills.

In the past decade, major attention has been given to STEM initiatives seeking to attract K-12 students into all science and engineering areas, especially minorities and women. Typical initiatives seek to make K-12 students aware of the wide range of STEM opportunities, and often use "robotics" as the attention getting focus. Such efforts are critically important, but they will not necessarily produce the results needed to assure the skilled labor force that will enable the expanded use of increasingly complex robotics and automation solutions. Some are beginning to talk about "K-to-Gray", bringing attention to need to consider the entire "life cycle" of workforce issues. New ideas and programs are needed to attract younger workers to the opportunities in robotics and automation, and to seek ways to better transfer and couple

the knowledge base of experienced workers into the emerging workforce – this need applies to all aspects of the robotics and automation future workforce from technicians to engineers and computer scientists. Specifically, new ideas and programs are needed to integrate the engineering and skilled labor training domains so that future technician training keeps track with the rapidly advancing use of intelligent systems and robotics. Better communications and collaborations are needed between professional societies, universities and community colleges are needed to assure that skilled technicians are being trained in critical areas such as mechatronics and embedded computer controls. Such interactions exist in many instances, but in general the depth, comprehensiveness, and real-time integration must go to another level.

7

Shared Infrastructure in Robotics

As robotics continues to expand to more application domains, the development and maintenance of suitable experimental facilities are becoming bottlenecks in the innovation process. In fact, there is a significant gap between the theoretical foundations that are being broadly pursued, and a focused, application-driven transition from small-scale experiments to robust and high impact deployments. This gap is both scientific and practical. By having researchers from different institutions, disciplines, and backgrounds come together around a common testbed, there is potential to accelerate innovation and to build on past findings in a more effective manner than what is currently done. The development and maintenance of meaningful, large-scale robotic testbeds is a resource-intensive undertaking, which is why it is particularly well-suited to a shared and even remote-access format.

Some specialized robot testbeds exist (e.g., UMass Lowell's NERVE Center, the Southwest Research Institute's Small Robotic Vehicle Evaluation and Applications Group, courses for response robots and manufacturing robots at NIST in Gaithersburg, MD, and Texas A&M's Disaster City), but these testbeds lack a variety of robot systems; researchers must bring their own robot systems, limiting the ability to test the gen-

eralizability of algorithms and preventing people without robot systems from being able to test their theoretical results in the real world. Georgia Tech's Robotarium, currently in development, will result in a robot testbed for swarm robotics, with both testing environments and robot hardware. In order to accelerate the development and effective testing of robot systems, shared community resources of testbeds for a variety of application domains with a variety of robot systems must be developed throughout the country, each with a particular application focus (e.g., agriculture, marine, manufacturing, medical). To maximize the use of available resources, existing facilities could be expanded to create a comprehensive shared infrastructure (i.e., one with both testbeds and shared robot systems) while developing testbeds for application domains where no such facilities yet exist.

7.1 Flexible Research Platforms

In order to be a truly useful remote-access research testbed, it is vitally important that the testbed is structured in such a manner that it allows for a number of different research questions and experiments to be pursued. Moreover, the testbed itself must evolve over time to remain relevant to the changing research trends and directions. It must be possible to automatically specify experimental setups and scenarios, which calls for research to be done on modular interoperable hardware (plug-n-play etc.) as well as software (robotic experiment description languages etc.) to facilitate their inclusion into larger eco-systems and downstream commercialization. Given that this infrastructure is a community resource, each facility should have a user committee to allocate site usage and suggest facility updates. While each facility will have a specific application focus, there should be collaboration between these facilities in order to prevent duplicating efforts and to share best practices.

7.2 Community Consensus Validation Benchmark Frameworks

Various research groups have developed in-house methods for quantitative performance assessments of both robotic systems and the human-robot interaction. Individual groups have begun the process to collect

and share their data sets and best-practices. However, efforts remain fragmented and disconnected due to the lack of realistic and relevant test environments (physical and virtual benchmarks). A multipronged validation regimen (e.g. supporting both virtual and physical testing; staged evaluation of components, subsystems and systems; device vs. user testing) is crucial. The development of such frameworks for open-access creation, collection and curation of the appropriate reference environments and data-corporuses against which quantitative performance can be assessed would significantly speed the process of technology development as well as transfer. Past efforts in the robotics community have strengthened the argument that potentially posing these as a competition or grand challenge could help focus the energy of both the academic and industrial communities, while also opening doors for subsequent standardization efforts. Some robotics domains have already been developing standard test methods and metrics through ASTM, IEEE and ASME; these efforts should continue while efforts in new application domains begin. The Robotics-VO could provide oversight for community discussions to develop these shared resources.

7.3 Reference Open-Access Testbeds

The enormous growth in the field has created an explosion in the number and variants of the solutions presented. For example, the range of manipulator arms, mobile bases and grippers commercially available can be mind-boggling. It is becoming increasingly difficult for a researcher working on grasping algorithms to obtain access to a variety of manipulators with different types of grippers in order to evaluate the effectiveness of their algorithms. However, the lack of access to truly industry-grade test-beds with interoperable hardware and software modules is beginning to impede innovation. Efforts at creating open-source platforms are underway and represent a good starting point. A broad and inclusive program needs to be supported through roadmapping workshops and study-groups to facilitate development of open-source community-vetted standards. A good example in the robotics software

arena is the ROS Framework.¹ The accessibility to such plug and play frameworks will let research groups focus on their subtopics while still contributing to a broader coherent community effort. Additionally, system interoperability and synergistic technical tools (e.g. programming, hardware, communication) are critical and will benefit academia and industry alike for hastening robotic system research and development.

Flexible Research Platform

- 5 Years** Coordination framework among academic researchers to create modular shareable hardware (CAD repositories) and software (APIs) for the next-gen robotic systems (e.g. wearables, soft-suits etc.)

Community Consensus Validation Benchmark Frameworks

- 5 Years** Coordination with Professional Societies (IEEE, ASME) to host competitions and roadmapping workshops
- 10 Years** Involvement of Standards organizations (IEEE, ISO, ANSI, ASME)

Reference Open-Access Testbeds

- 5 Years** Robot testbeds spread across the nation. Standup a coordination framework (perhaps like Robotics VO or CPS VO)
- 10 Years** Transition of viable hardware/software to pre-competitive TRL hardening via NNMIIs

¹<http://www.ros.org>

8

Legal, Ethical, and Economic Context

The Roadmap to Robotics is primarily a technical document. Its central purpose is to describe the present and anticipated state of the art in robotics in the United States and to help the American government set levels and priorities for support.

It is clear to the authors, however, that the development of robotics in the United States and elsewhere takes place against a backdrop of law, policy, ethics, and economics—among other social, cultural, and political forces. The purpose of the following chapter is to acknowledge this broader context. The chapter raises some of the more pressing non-technical challenges for robotics and directs the reader’s attention to ongoing efforts and resources to address these issues, where such efforts exist.

This chapter is not meant to be comprehensive, nor does it purport to articulate a consensus in the legal, policy, ethical or other communities as to what official policy toward robotics should be. Rather, we aim to raise certain key challenges that have repeatedly surfaced in the literature, in workshops, and in public discourse.¹ In addition, we

¹These include the National Science Foundation and Department of Homeland Security *Policy for Automation* workshop, the *Future of AI: Opportunities and Challenges*, the White House series *Preparing for the Future of Artificial Intelligence*,

articulate our commitment as a community to participate in and support this dialogue, which is by necessity deeply interdisciplinary, as well as to recommend that government and academia work to actively remove barriers interrogating robotics' broader societal context.

The remainder of this chapter consists of short discussion of key issues followed by our recommendations.

8.1 Safety

Robots have to be safe. But how safe is safe enough? There are many possible configurations, but a key role for government is to help set the safety thresholds or standards for a variety of robotic systems with the capacity to do physical harm to people or property out in the world. Thus, the Federal Aviation Administration will have to set safety thresholds for delivery of goods using unmanned aerial systems and the National Highway Transportation Safety Administration will have to set expectations around autonomous vehicles. Having set these thresholds or standards, techniques are then needed to test and validate that they are being met.

Special considerations may arise where robots are performing task usually performed by people with specialized training. Each profession that today certifies its own professionals will need to confront whether and how their standards can be translated into technical systems performing comparable tasks. While it may not make sense to give autonomous vehicles driving tests, clearly the medical profession will need to sign off on robots that eventually perform surgery autonomously.

For robotics to remain as safe and accountable as possible, there should also be a role for independent researchers. Academics and others may be in a position to help determine if systems are behaving or will behave as intended. To secure their participation, however, inde-

the Stanford University *AI 100* inaugural report, and the annual robotics law and policy conference *We Robot*, with more efforts in progress and on the horizon. Several of the authors of this report have participated in these and other efforts to identify and address the legal, policy, economic, and ethical issues robotics and AI may present. These efforts are focused on the United States; there are, if anything, more and longer standing efforts abroad, especially Europe (Italy, the UK, Germany), Japan, and South Korea.

pendent researchers need to know exactly what sorts of activities the law permits. The concern is that existing law—such as the Computer Fraud and Abuse Act, which disallows unauthorized access to many technical systems, and the anti-circumvention provisions of the Digital Millennium Copyright Act—may be read to prohibit research activities that ultimately serve the goals of public safety (Stone *et al.*, 2016). Lawmakers and enforcement agencies should clarify that, for instance, reverse engineering or otherwise examining software or hardware for the purpose of assessing its safety is permitted under all relevant laws.

8.2 Liability

Wise investment in robotics is likely to mean continued gains in public safety. Robots can perform inherently dangerous tasks, for instance, and perform risky tasks with greater precision. But robots, like humans, may find themselves in situations where harm is unavoidable. Courts and perhaps lawmakers will need to establish liability rules by which to compensate victims of robot-related hazards while preserving incentives for innovation.

Consider, for example, a home robot built by one company that injures a person while running software the robot's owner purchased from another company through a robot app store. From the victim's perspective, a robot built by a company with deep pockets caused an injury. But is it wise or fair to hold the manufacturer of a robot that—like a computer, tablet, or smartphone—is open to third party innovation by design? (Calo, 2011) Or consider a robot that, alone or in interaction with other systems, causes a kind of harm no one could reasonably anticipate. It should be clear, for example, that the manufacturer of a fully autonomous vehicle will be liable where that vehicle causes a traffic accident by turning without signaling. But now imagine an autonomous vehicle designed to find ways to maximize fuel efficiency through experimentation. The system might perform functions—such as running its engine in an enclosed garage to recharge its battery—that no one intended or anticipated, but which end up causing serious harm. Such events could pose a challenge to tort law, which is premised on the notion that courts should only compensate

injuries that are intended or foreseeable (Calo, 2015). Closely related to the task of determining liability is understanding the role of insurance. Market forces seems already to be responding to the acceleration of robotics; companies whose business models rely upon the use of robots are better able to secure insurance than they were a decade ago. But there also may be a role for government. The widespread availability of autonomous vehicles, for example, may present the need to revisit the utility of no-fault insurance, which has been declining in popularity among state lawmakers (Engstrom, 2012)

8.3 Impact on Labor

Many commentators have articulated a concern that the risks associated with the use of artificial intelligence to make decisions about consumers and citizens will fall disproportionately upon the vulnerable, i.e., those in society with the least capacity to mitigate technology's effects (see Crawford and Calo, 2016). These concerns also apply to robotics. For example, we might worry that greater reliance on robotics by police would disproportionately affect, and come to further alienate, low income or minorities communities (Joh, 2016).

The prospect that robots may take low or high skill jobs is of particular concern to the public. The concern is sometimes overstated. Even a simple analysis reveals that robots will both displace and create jobs at an individual level. There will be diminished need for captains, pilots, or truck drivers if companies automate long-haul transportation. At the same time, the burgeoning unmanned aerial vehicle or drone industry is already hiring. There is also evidence that automation in manufacturing has, to date, correlated with job *creation* in the United States, as mentioned in Section 2.

Finally, many analyses of robotics' impact on labor all but ignore the extensive and growing area of human augmentation. In contrast to automation, augmentation aims to enhance human abilities and create collaborations with machines, so that people are empowered, not replaced. Examples include rehabilitation robotics, socially assistive robotics, and collaborative robotics, to name a few.

Nevertheless, greater reliance on robotics is likely to have impacts in the short, medium, and long term and that must be managed responsibly (Brynjolfsson and McAfee, 2014). One solution that has been advanced in response to the prospect of widespread automation of jobs is the idea of a universal income, i.e., a basic income for every American subsidized by the gains in productivity and efficiency from automation. Another variant recommends imposing an obligation on employer to pay for retraining of workers displaced by robots.

There are several challenges around universal income, including that it may not be politically palatable, that income guarantees do not resolve other issues—such as idleness or inequity—that follow from unemployment, and that robotics might never so thoroughly transform our economy to permit redistribution of wealth on this scale.² A requirement that firms provide or subsidize retraining also brings challenges, such as the potential disincentive to adopt robots where doing so would constitute a net gain for productivity, safety, or both.

We are ultimately hopeful as a community that, if handled well at the level of policy, advances in robotics are likely to improve the overall health, resilience, and well-being of American society.

8.4 Social Interaction

An extensive literature evidences the ways in which people tend to react to anthropomorphic technology such as robots as though the robot were a social entity e.g. (Reeves and Nass, 1996). Designers of personal and service robots are well aware of this tendency and many have made considerable efforts to ensure a positive and respectful interaction between people and robots. Commentators worry, however, that the propensity people have to form social bonds with robots will prove problematic. Sherry Turkle and others have expressed concern that robotic interaction will substitute for far richer interpersonal relationships, as when an elderly relative is left in the care of a home robot (Turkle, 2012). Others worry that anthropomorphic robots will be capable of exploiting

²For an early argument that robots are unlikely to cause massive economic change, see Simon, 1965

our social reactions to nudge us toward corporate or other goals at odds with our own (Hartzog, 2015; Nourbakhsh, 2013).

There was widespread agreement among authors—many of whom work in the field of Human-Robot Interaction (HRI)—that both the positive and negative effects of robots on people need to be carefully researched and considered. Indeed, a limitation of current funding models is that too few resources are directed specifically toward studying social impact as itself a technical challenge of robotics. Participants further suggested the establishment of one or more testing facilities designed to emulate the real world. These would consist of instrumented environments where researchers can study human-robot interaction and compare share and compare results. Models for such robot spaces already exist in the United States and abroad.³

8.5 Privacy and Security

Closely related to the social interaction concern is the set of privacy and security challenges robots inevitably raise. Ryan Calo has argued that robots raise at least three categories of privacy issues: (1) robots make it easier to engage in surveillance, as when police use drones to monitor a protest; (2) robots create access to spaces historically reserved for solitude, as when government or black hat hackers compromise a home robot; and (3) as alluded to above, anthropomorphic technology such as robots occasion in people the perception of being observed. The potential for security vulnerabilities is rendered more acute by the prospect that a compromised robot could cause physical harm. Today a number of groups within civil service, industry, academia, and government are working to address some of these privacy and security issues.⁴ Government is in a position to better support this important work going forward, including by removing barriers to research.

³For example, the University of Michigan has a test range (M-City) where researchers test driverless cars and the Federal Aviation Administration has designed 10+ areas of the country for unmanned aerial vehicle testing.

⁴One example is the National Institute on Standards and Technology working group around privacy for unmanned aerial systems.

8.6 Recommendations

To reiterate: this document is primarily a technical roadmap. Its central purpose is to update decision makers on the state of the art in robotics and to help policymakers determine where to channel resources in order to realize robotics' great promise as a technology. Robotics develops against the background of a legal, policy, ethical, economics, and social context. This chapter has identified some of the challenges that recur in ongoing discussion of that context.

With this in mind, we conclude by tentatively offering a handful of recommendations aimed at preserving, fostering, and expanding the discussion of how robotics interact with society:

- *Greater expertise in government.* In order to foster innovation in robotics, maximize its potential for social good, and minimize its potential for harm, government at all levels should continue to accrue expertise in cyber-physical systems.
- *Support of interdisciplinary research in government and academia.* Few issues in robotics, or any other context, are amendable to resolution by reference to any one discipline. Government and academia should actively work to support and incentivize interdisciplinary research and breakdown siloes between expertise.
- *Removal of research barriers.* As alluded to above, independent researchers should be assured that efforts to understand and validate systems for the purpose of accountability and safety do not carry legal risk under existing law or doctrine.

9

Contributors

Name		Institution
Henny	Admoni	Carnegie Mellon Univ.
Ron	Alterovitz	Univ. of North Carolina
Nancy	Amato	Univ. of Illinois Urbana-Champaign
Jon	Battles	Amazon
Gregory	Brown	UPS
Jeff	Burnstein	A3 / Robot Industry Assoc.
Stefano	Carpin	Univ. of California, Merced
Howie	Choset	Carnegie Mellon Univ.
Henrik	Christensen	Univ. of California, San Diego
Jacob	Crandell	Brigham Young Univ.
Karthik	Dantu	State Univ. of New York - Buffalo
Joseph	Davidson	Oregon State Univ.
Pat	Davison	Material Handling Industry
Emel	Demircan	Cal State Long Beach
Jory	Denny	Univ. of Richmond
Aaron	Dollar	Yale Univ.
Katherine Rose	Driggs-Campbell	Univ. of Illinois Urbana-Champaign
Ann	Drobnis	Computing Community Consortium
Aleksandra	Faust	Google
Naomi	Fitter	Oregon State Univ.
Maria	Gini	Univ. of Minnesota

Name		Institution
Ken	Goldberg	Univ. of California, Berkeley
Mike	Goodrich	Brigham Young Univ.
Jessy	Grizzle	Univ. of Michigan
Yu	Gu	Western Virginia Univ.
Yan	Gu	Univ. of Massachusetts, Lowell
SK	Gupta	Univ. of Southern California
Gregory	Hager	Johns Hopkins Univ.
Ross	Hatton	Oregon State Univ.
John	Hollerbach	Univ. of Utah
Robert	Howe	Harvard Univ.
Jonathan	How	Massachusetts Institute of Technology
Seth	Hutchinson	Georgia Institute of Technology
Kshitij	Jerath	Univ. of Massachusetts, Lowell
Ewart	J. de Visser	US Air Force Academy
Jeff	Johnson	Maeve Automation
Chris	Jones	iRobot
Lydia	Kavraki	Rice Univ.
Zsolt	Kira	Georgia Institute of Technology
Jana	Kosecka	George Mason Univ.
Venkat	Krovi	Clemson Univ.
Amy	LaViers	Univ. of Illinois Urbana-Champaign
Johnny	Lee	Google
Daniel	Lee	Cornell Univ. / Cornell Tech
Nico	Lingg	Abcellera
Kevin	Lynch	North Western Univ.
Maja	Mataric	Univ. of Southern California
Roland	Menassa	Amazon
Elena	Messina	Nat. Inst. of Standards and Tech.
Robin	Murphy	Texas A&M
David	Naffin	John Deere
Nikos	Papanikolopoulos	Univ. of Minnesota
Jancie	Phillippus	Univ. of Illinois Urbana-Champaign
Elizabeth	Phillips	US Air Force Academy
Aaron	Prather	FedEx
Paul	Robinette	Univ. of Massachusetts, Lowell
Daniela	Rus	Massachusetts Institute of Technology
Junead	Satter	Univ. of Minnesota
Brian	Scasselatti	Yale Univ.
Matthias	Scheutz	Tufts Univ.
Elaine	Short	Tufts Univ.
Reid	Simmons	Carnegie Mellon Univ.

Name		Institution
Bill	Smart	Oregon State Univ.
Aaron	Steinfeld	Carnegie Mellon Univ.
Gaurav	Sukhatme	Univ. of Southern California
Daniel	Sula	Automation Focus
Carmillo Jose	Taylor	Univ. of Pennsylvania
Masayoshi	Tomizuka	Univ. of California, Berkeley
Jeff	Trinkle	Lehigh Univ.
Marynel	Vázquez	Yale Univ.
Yue	Wang	Clemson Univ
Melonee	Wise	Fetch Robotics
Zack	Woodruff	North Western Univ.
Srihari	Yamanoor	Independent
Holly	Yanco	Univ. of Massachusetts, Lowell
Michael	Yip	Univ. of California, San Diego
Milios	Zefran	Univ. of Illinois Chicago
Mabel	Zhang	Open Robotics Software Found.
Ye	Zhao	Georgia Inst. of Techn.

Acknowledgements

The authors gratefully acknowledge the support received from the Computing Community Consortium (CCC), Amazon Robotics, UC San Diego, University of Massachusetts – Lowell, University of Illinois Urbana Champaign, and University of Southern California. The support of all the contributors is also gratefully acknowledged, without their engagement the roadmapping process would not be possible.

References

- Accuweather. (2019). “AccuWeather predicts 2018 wildfires will cost California total economic losses of \$400 billion”. URL: <https://www.accuweather.com/en/weather-news/accuweather-predicts-2018-wildfires-will-cost-california-total-economic-losses-of-400-billion/432732>.
- Adaptive Vehicle Make. (2020). web. URL: <http://cps-vo.org/group/avm>.
- Agility Robotics. (2020). “Agility Robotics - Go where humans go”. web. URL: <https://www.agilityrobotics.com>.
- Amadeo, K. (2020). “How wildfires impact the economy”. The Balance Report - Web. URL: <https://www.thebalance.com/wildfires-economic-impact-4160764>.
- Amazon. (2020). “What is next for the Amazon Scout”. web. URL: <https://www.aboutamazon.com/news/transportation/whats-next-for-amazon-scout>.
- American Immigration Council. (2021). “The cost of immigration enforcement and border security”. web. URL: <https://www.americanimmigrationcouncil.org/research/the-cost-of-immigration-enforcement-and-border-security>.
- American Society of Civil Engineers (ASCE). (2020). “Report Card for American Infrastructure”. web. URL: <https://www.infrastructurereportcard.org>.

- Atkinson, R. D. (2019). “Robotics and the Future Production and Work”. ITIF. URL: <https://itif.org/publications/2019/10/15/robotics-and-future-production-and-work>.
- Blomqvist, L. and D. Douglas. (2016). “Is precision agriculture the way to peak cropland?” *Tech. rep.* The Breakthrough Institute.
- Bohg, J., A. Morales, T. Asfour, and D. Kragic. (2014). “Data-driven grasp synthesis—a survey”. *IEEE Transactions on Robotics*. 30(2): 289–309.
- Brynjolfsson, E. and A. McAfee. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- Cal Fire. (2019). “Pleasant Fire”. URL: <https://www.fire.ca.gov/incidents/2018/2/18/pleasant-fire>.
- Calo, R. (2011). “Open Robotics”. *Maryland Law Review*. 70(3).
- Calo, R. (2015). “Robotics and the Lessons of Cyberlaw”. *California Law Review*. 103: 513.
- Center for Disease Control and Prevention (CDC). (2020). “COVID-19”. web. URL: <https://www.cdc.gov/coronavirus/2019-ncov/index.html>.
- Christensen, H., A. Okamura, M. Mataric, V. Kumar, G. Hager, and H. Choset. (2016). “Next-Generation Robotics”. *Tech. rep.* CCC.
- Crawford, K. and R. Calo. (2016). “There is a blind spot in AI research”. *Nature News*. 538(7625): 311.
- Danielczuk, M., A. Kurenkov, A. Balakrishna, M. Matl, D. Wang, R. Martín-Martín, A. Garg, A. Savarese, and K. Goldberg. (2019). “Mechanical Search: Multi-Step Retrieval of a Target Object Occluded by Clutter”. In: *International Conference on Robotics and Automation (ICRA)*. Montreal, Canada.
- Dollar, A. and H. Herr. (2008). “Lower extremity exoskeletons and active orthoses: challenges and state-of-the-art”. *IEEE Trans. on Robotics*. 24(1): 144–158.
- Drone Life. (2020). “UPS Drone Delivery: DroneUp Flies to Prove the Case for Coronavirus Response”. web. URL: <https://dronelife.com/2020/04/21/ups-drone-delivery-droneup-partners-fly-to-prove-the-case-for-coronavirus-response>.
- Engstrom, N. F. (2012). “An Alternative Explanation for No-Fault’s Demise”. *DePaul Law Review*. 61(303).

- European Union. “Future of Manufacturing - Roadmap”. web. URL: <http://www.effra.eu/attachments/article/129/Factories%5C%20of%5C%20the%5C%20Future%5C%202020%5C%20Roadmap.pdf>.
- Federal Research Economic Data. (2020). “Industrial Production - Manufacturing - Automotive Vehicles”. web. URL: <https://fred.stlouisfed.org/series/IPG3361T3S>.
- FedEx. (2020). “FedEx Corporation Q4 Fiscal 2020 Statistics”. web. URL: https://s1.q4cdn.com/714383399/files/doc_financials/quarterly/2020/q4/FedEx-Q4-FY20-Stat-Book.pdf.
- Gartner Group. (2020). “Internet of Things”. web report. URL: <http://www.gartner.com/it-glossary/internet-of-things/>.
- Hadad, J. (2017). “E-Commerce & Online Auctions in the US”. IBIS World Industry Report 45411a:
- Hartzog, W. (2015). “Unfair and Deceptive Robots”. *Maryland Law Review*. 74(4): 785.
- Highly-complex and networked control systems. (2020). “EU FP7 Project on complex networked control”. web. URL: <http://www.hycon2.eu/>.
- Hodson, R. (2018). “How robots are grasping the art of gripping”. *Nature*. 557(7704).
- Industrial Internet Consortium. (2020). “White Papers”. web page. URL: <http://www.industrialinternetconsortium.org/white-papers.htm>.
- Insurance Information Institute. (2020). “Facts + Statistics: Wildfires”. web. URL: <https://www.iii.org/fact-statistic/facts-statistics-wildfires>.
- International Federation of Robotics. (2019). *World Robotics*. Vol. 2 - Service Robotics. VDMA.
- Joh, E. (2016). “Policing police robots”. *UCLA L. Rev. Discourse*. 64: 516.
- John Deere. (2020). “The Future of Farming”. web. URL: <https://www.deere.co.uk/en/agriculture/future-of-farming/>.
- Kotkin, J. (2018). “Where U.S. Manufacturing Is Thriving In 2018”. Forbes. URL: <https://www.forbes.com/sites/joelkotkin/2018/05/23/where-u-s-manufacturing-is-thriving-in-2018/?sh=7d174f2f53b3>.

- Levine, S., P. Pastor, A. Krizhevsky, and D. Quillen. (2016). “Learning hand-eye coordination for robotic grasping with large-scale data collection”. In: *International Symposium on Experimental Robotics*. Springer Verlag. 173–184.
- Mahler, J., J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg. (2017). “Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics”. arXiv:1703.09312.
- Mahler, J., M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg. (2019). “Learning Ambidextrous Robot Grasping Policies”. *Science Robotics*. 4(26).
- Mahler, J., R. Platt, A. Rodriguez, M. Ciocarlie, A. Dollar, R. Detry, M. A. Roa, H. Yanco, A. Norton, J. Falco, K. van Wyk, E. Messina, J. Leitner, D. Morrison, M. Mason, O. Brock, L. Odhner, A. Kurenkov, M. Matl, and K. Goldberg. (2018). “Robot Grasping Benchmarks, Protocols, and Metrics (Guest Editorial)”. *IEEE Transactions on Automation Science and Engineering (T-ASE)*. 15(4).
- Maslow, A. H. (1943). “A theory of human motivation”. *Psychological Review*. 50(4): 370–96.
- Maymobility Corp. (2020). “Maymobility”. web. URL: <https://maymobility.com/>.
- Miller, W. C., A. B. Deathe, M. Speechley, and J. Koval. (2001). “The influence of falling, fear of falling, and balance confidence on prosthetic mobility and social activity among individuals with a lower extremity amputation”. *Archives of physical medicine and rehabilitation*. 82(9): 1238–1244.
- Morrison, D., A. W. Tow, M. McTaggart, R. Smith, N. Kelly-Boxall, S. Wade-McCue, J. Erskine, R. Grinover, A. Gurman, and T. Hunn. (2018). “Cartman: The low-cost cartesian manipulator that won the amazon robotics challenge”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 7757–7764.
- National Wildfire Coordination Group. (2017). “Report on Wildland Firefighter Fatalities in the United States: 2007-2016”. *Tech. rep.* No. PMS 841. NWCG. URL: <https://www.nwcg.gov/sites/default/files/publications/pms841.pdf>.

- Nourbakhsh, I. R. (2013). *Robot futures*. Boston, MA: MIT Press.
- Office of National Drug Control Policy. (2016). “How Illicit Drug Use Affects Business and the Economy”. ONDCP Web. URL: <https://obamawhitehouse.archives.gov/ondcp/ondcp-fact-sheets/how-illicit-drug-use-affects-business-and-the-economy>.
- Pell, J. P., P. T. Donnan, F. G. R. Fowkes, and C. V. Ruckley. (1993). “Quality of life following lower limb amputation for peripheral arterial disease”. *European journal of vascular surgery*. 7(4): 448–451.
- Pinto, L. and A. Gupta. (2016). “Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours”. In: *Intl. Conf. of Robotics and Automation (ICRA)*. 3406–3413.
- Reeves, B. and C. Nass. (1996). *The media equation: How people treat computers, television, and new media like real people*. Cambridge university press Cambridge, UK.
- Rimon, E. and J. Burdick. (2019). *J. Mechanics of Robot Grasping*. Cambridge University Press.
- Siemens. (2019). “Industry 4.0”. web. URL: <https://press.siemens.com/global/en/pressrelease/leading-industry-40-players-collaborate-help-manufacturers-accelerate-digital>.
- Simon, H. (1965). *The shape of automation for men and management*. Vol. 13. Harper & Row New York.
- Smith, J. (2020). “Warehouse Robotics Startups Drawing Bigger Investor Backing”. Wall Street Journal. URL: <https://www.wsj.com/articles/warehouse-robotics-startups-drawing-bigger-investor-backing-11578394802>.
- Stone, P., R. Brooks, E. Brynjolfsson, R. Calo, O. Etzioni, G. Hager, J. Hirschberg, S. Kalyanakrishnan, E. Kamar, S. Kraus, K. Leyton-Brown, D. Parkes, W. Press, A. Saxenian, J. Shah, M. Tambe, and A. Teller. (2016). “Artificial Intelligence and Life in 2030”. *Tech. rep.* One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA. URL: <http://ai100.stanford.edu/2016-report>.
- Tech Crunch. (2020). “Amazon says it has deployed more than 200,000 robotic drives globally”. web. URL: <https://techcrunch.com/2019/06/05/amazon-says-it-has-deployed-more-than-200000-robotic-drives-globally/>.

- Turkle, S. (2012). “In Constant Digital Contact, We Feel’Alone Together””. *Alone Together*.
- US Census Bureau. (2018). “An Aging Nation: Projected Number of Children and Older Adults”. web. URL: <https://www.census.gov/library/visualizations/2018/comm/historic-first.html>.
- US Census Bureau. (2020). “The Greying of America: More Older Adults than Kids by 2035”. web. URL: <https://www.census.gov/library/stories/2018/03/graying-america.html>.
- USDA - Economic Research Service. (2020). “Food and Beverage Manufacturing”. web. URL: <https://www.ers.usda.gov/topics/food-markets-prices/processing-marketing/manufacturing/>.
- Waters, R. L., J. Perry, D. Antonelli, and H. Hislop. (1976). “Energy cost of walking of amputees: the influence of level of amputation”. *Jour of Bone Joint Surgery*. 58(1): 42–46.
- Wikipedia. (2020). “Amazon Prime Air”. web. URL: https://en.wikipedia.org/wiki/Amazon%5C_Prime%5C_Air.
- Worldometer. (2020). “CoronaVirus”. web. URL: <https://www.worldometers.info>.
- Zablotsky, B., L. I. Black, M. J. Maenner, L. A. Schieve, M. L. Danielson, R. H. Bitsko, S. J. Blumberg, M. D. Kogan, and C. A. Boyle. (2019). “Prevalence and Trends of Developmental Disabilities among Children in the United States: 2009–2017”. *Pediatrics*. 144(4).
- Ziegler-Graham, Kathryn, E. J. MacKenzie, P. L. Ephraim, T. G. Trivison, and R. Brookmeyer. (2008). “Estimating the prevalence of limb loss in the United States: 2005 to 2050.” *Archives of physical medicine and rehabilitation*. 89(3): 422–429.
- Zuckerman, L. (2017). “Cost of fighting U.S. wildfires topped \$2 billion in 2017”. Reuters. URL: <https://www.reuters.com/article/us-usa-wildfires/cost-of-fighting-u-s-wildfires-topped-2-billion-in-2017-idUSKCN1BQ01F>.