

A Survey of Research on Cloud Robotics and Automation

Ben Kehoe, *Student Member, IEEE*, Sachin Patil, *Member, IEEE*, Pieter Abbeel, *Senior Member, IEEE*, and Ken Goldberg, *Fellow, IEEE*

Abstract—The Cloud infrastructure and its extensive set of Internet-accessible resources has potential to provide significant benefits to robots and automation systems. We consider robots and automation systems that rely on data or code from a network to support their operation, i.e., where not all sensing, computation, and memory is integrated into a standalone system. This survey is organized around four potential benefits of the Cloud: 1) **Big Data**: access to libraries of images, maps, trajectories, and descriptive data; 2) **Cloud Computing**: access to parallel grid computing on demand for statistical analysis, learning, and motion planning; 3) **Collective Robot Learning**: robots sharing trajectories, control policies, and outcomes; and 4) **Human Computation**: use of crowdsourcing to tap human skills for analyzing images and video, classification, learning, and error recovery. The Cloud can also improve robots and automation systems by providing access to: a) datasets, publications, models, benchmarks, and simulation tools; b) open competitions for designs and systems; and c) open-source software. This survey includes over 150 references on results and open challenges. A website with new developments and updates is available at: <http://goldberg.berkeley.edu/cloud-robotics/>

Note to Practitioners—Most robots and automation systems still operate independently using onboard computation, memory, and programming. Emerging advances and the increasing availability of networking in the “Cloud” suggests new approaches where processing is performed remotely with access to dynamic global datasets to support a range of functions. This paper surveys research to date.

Index Terms—Big data, cloud automation, cloud computing, cloud robotics, crowdsourcing, open source.

I. INTRODUCTION

As illustrated in Fig. 1, the Cloud has potential to enhance a broad range of robots and automation systems. The National Institute of Standards and Technology (NIST) defines the Cloud as “a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable resources (e.g., servers, storage, networks, applications, and

services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” [115]. An example is the online word processing capabilities offered by Google Docs. One can send Microsoft Word documents over the Internet, but Google Docs differs in that the document and software does not reside locally. The data and code is stored in the Cloud using remote server farms with shared processors and memory. This is helpful because one does not have to worry about maintenance, outages, and software or hardware updates. The Cloud also provides economies of scale and facilitates sharing of data across applications and users [122].

Cloud Robot and Automation systems can be broadly defined as follows: *Any robot or automation system that relies on either data or code from a network to support its operation, i.e., where not all sensing, computation, and memory is integrated into a single standalone system.* This definition is intended to include future systems and many existing systems that involve networked teleoperation or networked groups of mobile robots such as UAVs [100], [116] or warehouse robots [44], [96] as well as advanced assembly lines, processing plants, and home automation systems, and systems with computation performed by humans [132], [155]. Due to network latency, variable quality of service, and downtime, Cloud Robot and Automation systems often include some capacity for local processing for low-latency responses and during periods where network access is unavailable or unreliable. This is not a binary definition; there are degrees to which any system will fit under this definition.

The Google self-driving car exemplifies the idea. It indexes maps and images collected and updated by satellite, Streetview, and crowdsourcing from the Cloud to facilitate accurate localization. Another example is the Kiva Systems pallet robot for warehouse logistics. These robots communicate wirelessly with a local central server to coordinate routing and share updates on detected changes in the environment.

In 2010, James Kuffner coined the term “Cloud Robotics” and described a number of potential benefits [98]. An article in IEEE Spectrum quickly followed [70] and Steve Cousins summarized the concept as “No robot is an island.” The next section considers the history of this important idea.

This survey is organized around four potential benefits from the Cloud: 1) **Big Data**: access to remote libraries of images, maps, trajectories, and object data; 2) **Cloud Computing**: access to parallel grid computing on demand for statistical analysis, learning, and motion planning; 3) **Collective Robot Learning**: robots sharing trajectories, control policies, and outcomes; and

Manuscript received March 15, 2014; revised September 15, 2014; accepted November 06, 2014. Date of publication January 12, 2015; date of current version April 03, 2015. This paper was recommended for publication by Associate Editor S. Sarma and Editor S. Sarma upon evaluation of the reviewers' comments. This work was supported in part by Google and Cisco and by the U.S. National Science Foundation under Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems.

The authors are with the University of California, Berkeley, CA 94720 USA (e-mail: benk@berkeley.edu; sachinpatil@berkeley.edu; pabbeel@berkeley.edu; goldberg@berkeley.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TASE.2014.2376492



Fig. 1. The Cloud has potential to enable a new generation of robots and automation systems to use wireless networking, Big Data, Cloud Computing, statistical machine learning, open-source, and other shared resources to improve performance in a wide variety of applications including assembly, inspection, driving, warehouse logistics, caregiving, package delivery, housekeeping, and surgery.

4) Human computation: using crowdsourcing access to remote human expertise for analyzing images, classification, learning, and error recovery. This survey also cites examples where the Cloud can enhance robotics and automation systems by facilitating access to: a) datasets, publications, models, benchmarks, and simulation tools; b) open competitions for designs and systems; and c) open-source software.

II. A BRIEF HISTORY

The value of networking to connect machines in manufacturing automation systems was recognized over 30 years ago. In the 1980s, General Motors developed the Manufacturing Automation Protocol (MAP) [80]. A diverse set of incompatible proprietary protocols were offered by vendors until a shift began in the early 1990s when the World Wide Web popularized the HTTP over IP protocols [121].

In 1994, the first industrial robot was connected to the Web with an intuitive graphical user interface that allowed visitors to teleoperate the robot via any Internet browser [63]. In the mid and late 1990s, researchers developed a series of web interfaces to robots and devices to explore issues such as user interfaces and robustness [64], [65] that initiated the subfield of “Networked Robotics” [66], [112].

In 1997, work by Inaba *et al.* on “remote brained robots” described the advantages of remote computing for robot control [79].

In May 2001, the IEEE Robotics and Automation Society established the Technical Committee on Networked Robots [10] which organized a number of workshops. Chapters of the Springer Handbook on Robotics were focused on Networked Tele-robots (where robots are operated remotely by humans using global networks) and Networked Robots (where robots communicate with each other using local networks), respectively [99], [143].

In 2009, the RoboEarth project was announced. It envisioned “a World Wide Web for robots: a giant network and database repository where robots can share information and learn from each other about their behavior and environment” [22], [156]

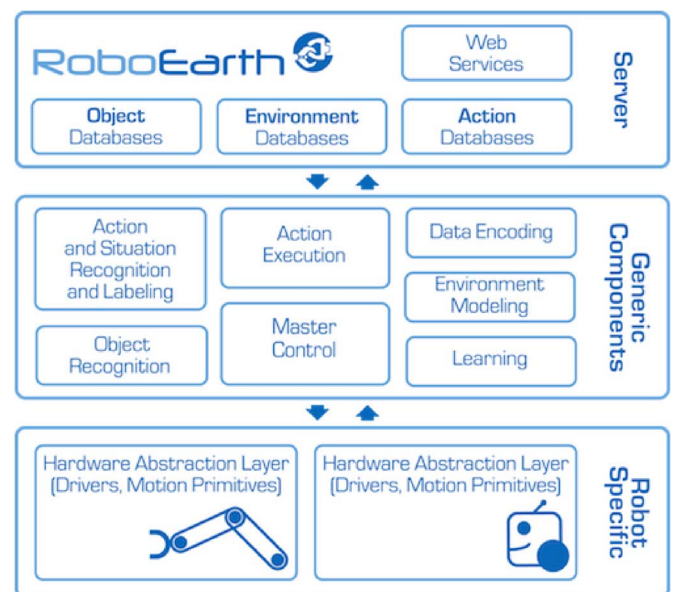


Fig. 2. The RoboEarth systems architecture designed to allow robots to share data and learn from each other [22], [156]. (Image reproduced with permission from authors.)

as illustrated in Fig. 2. Under a major European Union grant, the RoboEarth research team developed a series of system architectures for service robotics [30], [53], developing Cloud networking [74], [87], and computing resources [78] to generate 3D models of environments, speech recognition, and face recognition [148].

As noted in the previous section, James Kuffner introduced the term “Cloud Robotics” in 2010. This broader term supplanted earlier terminology and has been adopted by many researchers including the organizers of this Special Issue of the IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING.

Cloud Robotics and Automation is related to several other new initiatives. The “Internet of Things” [32], a term also introduced in 2010, describes how RFID and inexpensive processors




Fig. 3. Data can be collected from many sources as shown in this schematic architecture for the Mobile Millennium, a Cloud-based transportation system that combines streaming data from taxis, maps, and road-based sensors [77]. Mobile Millennium uses the Big Data and Collective Robot Learning aspects of Cloud Robotics and Automation. (Image reproduced with permission from authors.)

could be incorporated into a vast array of robots and physical objects from inventory items to household appliances [110] to allow them to communicate and share information.

The term “Industry 4.0,” introduced in Germany in 2011, predicts a fourth industrial revolution that will use networking to follow the first (mechanization of production using water and steam power), the second (mass production with electric power), and the third (use of electronics to automate production) industrial revolutions [11].

In 2012, General Electric introduced the term “Industrial Internet,” to describe new efforts where industrial equipment such as wind turbines, jet engines, and MRI machines connect over networks to share data and processing for industries including energy, transportation, and healthcare [55], [93]. For example, GE is using sensor readings from aircraft engines to optimize fuel consumption under a myriad of conditions [59]. The power of Big Data and Cloud Computing are extensively being used by many industries [25], [113], and are being harnessed to optimize water usage for irrigation [52].

Many related projects are emerging. In August 2014, Ashutosh Saxena announced his goal to build a “RoboBrain: a large-scale computational system that learns from publicly available Internet resources, computer simulations, and real-life robot trials.”

III. BIG DATA

The Cloud can provide robots and automation systems with access to vast resources of data that are not possible to maintain in onboard memory. “Big Data” describes “data that exceeds the processing capacity of conventional database systems” [54] including images, video, maps, real-time network and financial transactions [102], and vast networks of sensors [159].

A recent U.S. National Academy of Engineering Report summarizes many research opportunities and challenges created by Big Data [41] and other challenges are summarized in [29] and [165]. For example, sampling algorithms can provide reasonable approximations to queries on large datasets to keep run-

ning times manageable [38], but these approximations can be seriously affected by “dirty data” [58].

Hunter *et al.* [77] presents algorithms for a Cloud-based transportation system called Mobile Millennium, which uses the GPS in cellular phones to gather traffic information, process it, and distribute it and also to collect and share data about noise levels and air quality (see Fig. 3).

Large datasets can facilitate machine learning, as has been demonstrated in the context of computer vision. Large-scale image datasets such as ImageNet [50], PASCAL visual object classes dataset [56], and others [142], [151] have been used for object and scene recognition. By leveraging Trimble’s SketchUp 3D warehouse, Lai *et al.* reduced the need for manually labeled training data [101]. Using community photo collections, Gammeter *et al.* created an augmented reality application with processing in the Cloud [57]. Combining internet images with querying a local human operator, Hidago-Pena *et al.* provided a more robust object learning technique [72]. Deep learning is a technique using many-layered neural networks that can take advantage of Big Data [49], and has been used for computer vision [97], [140] and grasping [104].

Grasping is a persistent challenge in robotics: determining the optimal way to grasp a newly encountered object. Cloud resources can facilitate incremental learning of grasp strategies [40], [120] by matching sensor data against 3D CAD models in an online database. Examples of sensor data include 2D image features [75], 3D features [68], and 3D point clouds [39], as well as demonstrations [94].

Google Goggles [9], a free image recognition service for mobile devices (see Fig. 4), has been incorporated into a Cloud-based system for robot grasping [89], as illustrated in Fig. 5.

The RoboEarth project stores data related to objects and maps for applications ranging from object recognition to mobile navigation to grasping and manipulation (see Fig. 2) [156]. The Columbia Grasp dataset [67], the MIT KIT object dataset [88], and the Willow Garage Household Objects Database [40] are available online and have been used to evaluate different aspects of




Fig. 4. Google's object recognition system combines an enormous dataset of images and textual labels with machine learning to facilitate object recognition in the Cloud [9], [98]. (Image reproduced with permission.)

grasping algorithms, including grasp stability [45], [46], robust grasping [162], and scene understanding [128]. Dalibard *et al.* attach “manuals” of manipulation tasks to objects [43].

One research challenge is defining cross-platform formats for representing data. While sensor data such as images and point clouds have a small number of widely-used formats, even relatively simple data such as trajectories have no common standards yet but research is ongoing [129], [147], [149]. Another challenge is working with sparse representations for efficient transmission of data, e.g., algorithms for sparse motion planning for robotic and automation systems [51], [105].

Large datasets collected from distributed sources are often “dirty” with erroneous, duplicated, or corrupted data [6], [158], such as 3D position data collected during robot calibration [111]. New approaches are required that are robust to dirty data.

IV. CLOUD COMPUTING

Massively-parallel computation on demand is now widely available [29] from commercial sources such as Amazon's Elastic Compute Cloud [1], [2], Google's Compute Engine [8], and Microsoft's Azure [12]. These systems provide access to tens of thousands of remote processors for short-term computing tasks [105], [106]. These services were originally used primarily by web application developers but have increasingly been used in scientific and technical high-performance computing (HPC) applications [20], [86], [114], [152].

Uncertainty in sensing, models, and control is a central issue in robotics and automation [62]. Such uncertainty can be modeled as perturbations in position, orientation, shape, and control. Cloud Computing is ideal for sample-based Monte-Carlo analysis. For example, parallel Cloud Computing can be used to compute the outcomes of the cross-product of many possible perturbations in object and environment pose, shape, and robot response to sensors and commands [154]. This idea is being explored in medicine [157] and particle physics [141].

Cloud-based sampling can be used to compute robust grasps in the presence of shape uncertainty [90]–[92] (see Fig. 6). This grasp planning algorithm accepts as input a nominal polygonal outline with Gaussian uncertainty around each vertex and the center of mass and uses parallel-sampling to compute a grasp quality metric based on a lower bound on the probability of achieving force closure.

Cloud Computing has potential to speed up many computationally-intensive robotics and automation systems applications such as robot navigation by performing SLAM in the Cloud [134], [135] as illustrated in Fig. 7 and next-view planning for object recognition [125]. Cloud-based formation control of ground robots has also been demonstrated [153].

For optimal sampling-based motion planning methods such as RRT*, Cloud Computing is useful to generate the graphs; it is also important to recognize that these graphs can grow rapidly so algorithms for graph reduction are needed to facilitate data transfer, as illustrated in Fig. 8.

The Cloud also facilitates video and image analysis [123], [137], and mapping [119], [136] (see Fig. 7. Image processing in the Cloud has been used for assistive technology for the visually impaired [36] and for senior citizens [58].

Bekris *et al.* [33] propose an architecture for efficiently planning the motion of new robot manipulators designed for flexible manufacturing floors in which the computation is split between the robot and the Cloud.

It is important to acknowledge that the Cloud is prone to varying network latency and quality of service. Some applications are not time sensitive, such as decluttering a room or pre-computing grasp strategies or offline optimization of machine scheduling, but many applications have real-time demands [83] and this is an active area of research [26], [27], [95], [109].

V. COLLECTIVE ROBOT LEARNING

The Cloud facilitates sharing of data for robot learning by collecting data from many instances of physical trials and environments. For example, robots and automation systems can share initial and desired conditions, associated control policies and trajectories, and importantly: data on the resulting performance and outcomes.

The “Lightning” framework (see Fig. 10), proposes a framework for Collective Robot Learning by indexing trajectories from many robots over many tasks and using Cloud Computing for parallel planning and trajectory adjustment [35].

Such systems can also be expanded to global networks to facilitate shared path planning, including traffic routing, as shown in Fig. 9.

For grasping [37], grasp stability of finger contacts can be learned from previous grasps on an object [45]. Sharing data through Collective Robot Learning can also improve the capabilities of robots with limited computational resources [69].

The MyRobots project [13] from RobotShop proposes a “social network” for robots: “In the same way humans benefit from socializing, collaborating and sharing, robots can benefit from those interactions too by sharing their sensor information giving insight on their perspective of their current state” [21].

The RoboEarth and RoboBrain databases in Section III are designed to be updated with new information from connected robots. The RoboBrain project “learns from publicly available Internet resources, computer simulations, and real-life robot trials.” [16]

KIVA Systems [44], [96] uses hundreds of mobile platforms to move pallets in warehouses using a local network to coordinate motion and update tracking data.

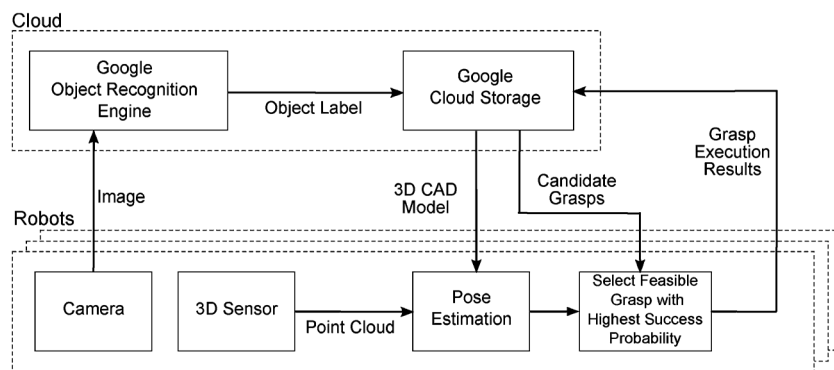


Fig. 5. System Architecture for Cloud-based object recognition for grasping. The robot captures an image of an object and sends via the network to the Google object recognition server. The server processes the image and returns data for a set of candidate objects, each with precomputed grasping options. The robot compares the returned CAD models with the detected point cloud to refine identification and to perform pose estimation, and selects an appropriate grasp. After the grasp is executed, data on the outcome is used to update models in the Cloud for future reference [89]. This project uses the Big Data, Cloud Computing, and Collective Robot Learning aspects of Cloud Robotics and Automation. (Image reproduced with permission.)

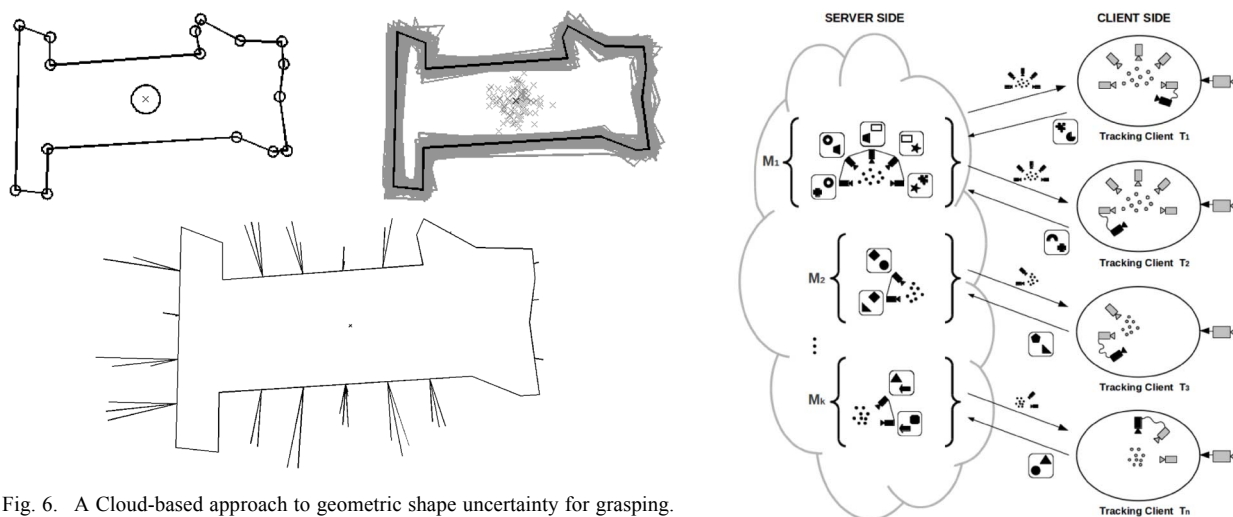


Fig. 6. A Cloud-based approach to geometric shape uncertainty for grasping. (Top) Uncertainty in object pose and shape. (Bottom) Computed push grasps. Kehoe *et al.* use sampling over uncertainty distributions to find a lower bound on the probability of success for grasps [90]–[92].

VI. HUMAN COMPUTATION: CROWDSOURCING AND CALL CENTERS

Human skill, experience, and intuition is being tapped to solve a number of problems such as image labeling for computer vision [40], [87], [98], [155], learning associations between object labels and locations [138], and gathering data [48]. Amazon's Mechanical Turk is pioneering on-demand “crowdsourcing” with a marketplace where tasks that exceed the capabilities of computers can be performed by human workers. In contrast to automated telephone reservation systems, consider a future scenario where errors and exceptions are detected by robots and automation systems which then contact humans at remote call centers for guidance.

Research projects are exploring how this can be used for path planning [73], [82], to determine depth layers, image normals, and symmetry from images [61], and to refine image segmentation [84]. Researchers are working to understand pricing models [145] and apply crowdsourcing to grasping [144] (see Fig. 12). Knowledge-based solutions are being explored for industrial automation as well [146].




Fig. 7. A Cloud framework for robot navigation using cooperative tracking and mapping (C²TAM). Riazuelo *et al.* demonstrate computer intensive bundle adjustment for navigation using simultaneous localization and mapping (SLAM) performed in the Cloud [134]–[136]. (Image reproduced with permission.)

Networked robotics has a long history of allowing robots to be remotely teleoperated by humans over the web [63], and the expanded resources of the Cloud enables new research into remote human operation [103], [144], [164] (see Fig. 11).

VII. OPEN-SOURCE AND OPEN-ACCESS

The Cloud supports the evolution of Cloud Robotics and Automation by facilitating human access to: a) datasets, publica-




Fig. 8. Distributed sampling-based motion planning. A roadmap of trees for motion planning in high-dimensional spaces. Plaku *et al.* show that their planner can “easily solve high-dimensional problems that exhaust resources available to single machines” [126]. (Image reproduced with permission.)

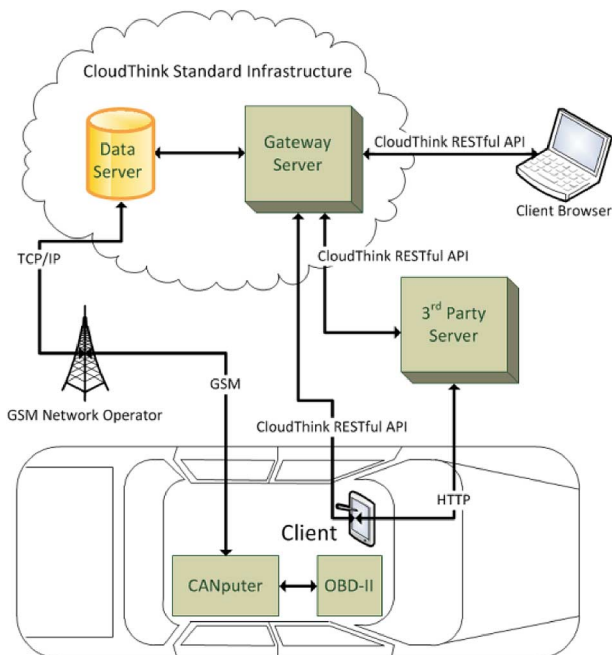


Fig. 9. Schematic architecture of CloudThink. Wilhem *et al.* developed an open-standard for self-reporting sensing devices such as sensors mounted in automobiles. Cloud-enabled storage of sensor network data can enable collaborative sharing of data for traffic routing and other applications [163]. CloudThink uses the Collective Robot Learning aspect of Cloud Robotics and Automation. (Image reproduced with permission from authors.)

tions, models, benchmarks, and simulation tools; b) open competitions for designs and systems; and c) open-source software.

The success of open source software [42], [124] is now widely accepted in the robotics and automation community. A primary example is ROS, the Robot Operating System, which provides libraries and tools to help software developers create robot applications [14], [18], [131]. ROS has also been ported to Android devices [19]. ROS has become a standard akin to Linux and is now used by almost all robot developers in research and many in industry, with the ROS Industrial project created to support these users [17].

Additionally, many simulation libraries for robotics are now open source, which allows students and researchers to rapidly set up and adapt new systems and share the resulting software. There are many open source simulation libraries, including Bullet [4], a physics simulator originally used for video games,

OpenRAVE [15] and Gazebo [7], simulation environments geared specifically towards robotics, OOPSMP, a motion-planning library [127], and GraspIt!, a grasping simulator [117]. The open source nature of these libraries allows them to be modified to suit applications and they were not originally designed for.

Another exciting trend is in open source hardware, where CAD models and the technical details of construction of devices are made freely available [47]. The Arduino project [3] is a widely-used open source microcontroller platform with many different sensors and actuators available, and has been used in many robotics projects. The Raven [71] is an open-architecture laparoscopic surgery robot developed as a research platform an order of magnitude less expensive than commercial surgical robots [23]. Recent advances in 3D printing (also known as additive manufacturing) are poised to have a major impact on many fields, including development of open source hardware designs [60], [81], [108].

The Cloud facilitates open challenges and design competitions that can draw on a diverse and geographically distributed population of innovators.

The DARPA Robotics Challenge (DRC) is “a competition of robot systems and software teams vying to develop robots capable of assisting humans in responding to natural and man-made disasters,” supported by NIST and the Southwest Robotics Institute (SwRI) [24]. The DRC simulator is provided to all contestants through CloudSim, an open-source Cloud-based simulation platform for testing the performance of the Atlas humanoid robot (shown in Fig. 13) on a variety of disaster response tasks [5], [7]. The Cloud permits running interactive, real-time simulation tasks in parallel for purposes such as predicting and evaluating performance, validating design decisions, optimizing designs, and training users [28].

Another example of an open competition is the “Ultra-Affordable Educational Robot Challenge” organized by the African Robotics Network with support from the IEEE Robotics and Automation Society in the Summer of 2012. It attracted 28 designs from around the world including the Grand Prize winning design shown in Fig. 14, where a modified surplus Sony game controller uses the vibration motors to drive wheels and lollipops as inertial counterweights for contact sensing by the thumb switches. This robot can be built from surplus parts for US \$8.96 [150].

VIII. CHALLENGES AND FUTURE DIRECTIONS

Using the Cloud for robotics and automation systems introduces many new challenges. The connectivity inherent in the Cloud raises a range of privacy and security concerns [133], [139]. These concerns include data generated by Cloud-connected robots and sensors, especially as they may include images or video or data from private homes or corporate trade secrets [130], [161]. Cloud Robotics and Automation also introduces the potential of robots and systems to be attacked remotely: a hacker could take over a robot and use it to disrupt functionality or cause damage. For instance, researchers at University of Texas at Austin demonstrated that it is possible to hack into and remotely control UAV drones via inexpensive

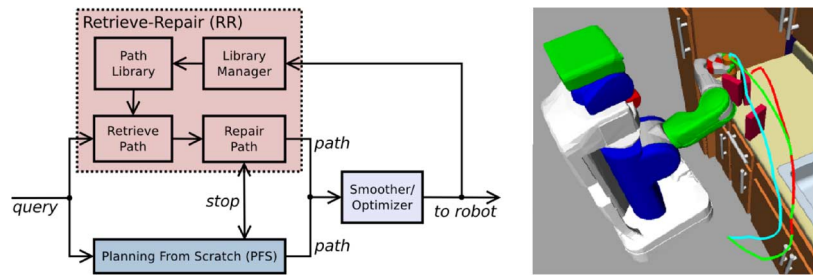


Fig. 10. (Left) Schematic architecture of the Lightning path planning framework. Berenson *et al.* show a system that is able to learn from experience from pre-computed motion plans, which could be stored in the Cloud. The planner attempts to find a brand-new plan as well as find an existing plan for a problem similar to the current one. Whichever finishes first is chosen [35]. Lightning uses the Big Data, Cloud Computing, and Collective Robot Learning aspects of Cloud Robotics and Automation. (Image reproduced with permission from authors.)

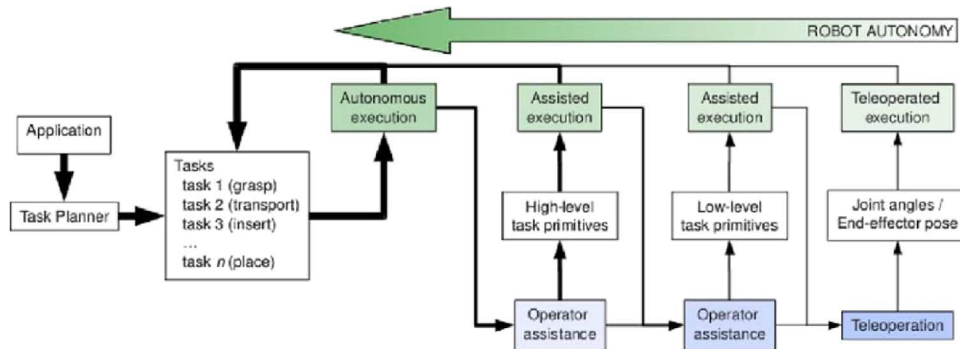


Fig. 11. Tiered human assistance using Cloud-based resources for teleoperation. Leeper *et al.* developed an interface for operators to control grasp execution using a set of different strategies. The results indicate humans are able to select better and more robust grasp strategies [103], [164]. (Image reproduced with permission.)

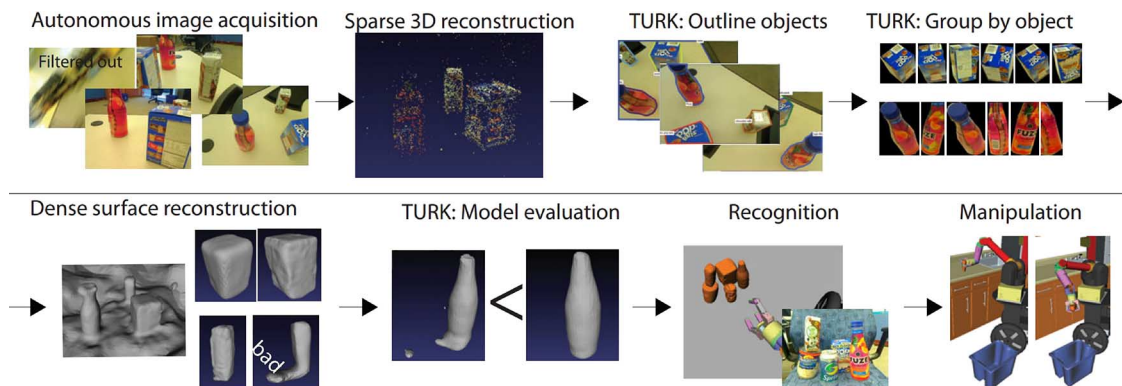


Fig. 12. Crowdsourcing object identification to facilitate robot grasping. Sorokin *et al.* developed a Cloud robot system that incorporates Amazon's Mechanical Turk to obtain semantic information about the world and subjective judgments [144]. This work uses the Human Computation aspect of Cloud Robotics and Automation. (Image reproduced with permission from authors.)

GPS spoofing systems in an evaluation study for the Department of Homeland Security (DHS) and the Federal Aviation Administration (FAA) [76]. These concerns raise new regulatory, accountability and legal issues related to safety, control, and transparency [107], [130]. The “We Robot” conference is an annual forum for ethics and policy research [160].

On the technical front, new algorithms and methods are needed to cope with time-varying network latency and Quality-of-Service. Faster data connections, both wired Internet connections and wireless standards such as LTE [31], are reducing latency, but algorithms must be designed to degrade gracefully when the Cloud resources are very slow, noisy, or unavailable [34]. For example, “anytime” load balancing algorithms for speech recognition on smart phones send the speech signal to the Cloud for analysis and simultaneously

process it internally and then use the best results available after a reasonable delay. Similar algorithms will be needed for robotics and automation systems [35].

New algorithms are also needed that scale to the size of Big Data, which often contain dirty data that requires new approaches to clean or sample effectively [6], [158]. When the Cloud is used for parallel-processing, it is vital that algorithms oversample to take into account that some remote processors may fail or experience long delays in returning results. When human computation is used, algorithms are needed to filter unreliable input and balance the costs of human intervention with the cost of robot failure.

Moving robotics and automation algorithms into the Cloud requires frameworks that facilitate this transition. The Cloud provides three possible levels at which a framework could be



Fig. 13. The DARPA Robotics Challenge (DRC) used CloudSim, an open-source Cloud-based simulation platform for testing the performance of the Atlas humanoid robot (shown) on a variety of disaster response tasks [5], [7]. The Cloud permits running interactive, real-time simulation tasks in parallel for purposes such as predicting and evaluating performance, validating design decisions, optimizing designs, and training users. This competition also resulted in enabling sharing of robotics research efforts. (Image reproduced with permission.)




Fig. 14. Lollibot, designed by Tom Tilley of Thailand, won the Grand Prize in the \$10 Educational Robot Design Challenge organized by the African Robotics Network. This design can be built from surplus parts for US \$8.96. [150]. (Image reproduced with permission.)

implemented [115]. The lowest level is Infrastructure as a Service (IaaS), where bare operating systems are provided on (possibly virtualized) machines in the Cloud. The second level, Platform as a Service (PaaS), provides more structure, including application frameworks and database access, while restricting the choice of programming languages, system architectures, and database models that can be used. Software as a Service (SaaS), the highest level of structure, is exemplified by the difference between Google Docs, a Cloud-based word processor, and Microsoft Word, which must be downloaded and installed locally.

The RoboEarth project includes a Cloud Computing platform called Rapyuta [118], which is a Platform as a Service (PaaS) framework for moving computation off of robots and into the Cloud. It also connects to the RoboEarth knowledge repository, integrating the Big Data aspect. We believe that this PaaS approach can be extended to use the Software as a Service (SaaS) paradigm, which offers many potential advantages for robots

and automation systems [85]. With SaaS, an interface allows data to be sent to a server that processes it and returns outputs, which relieves users of the burden of maintaining data and software and hardware and allows companies to control proprietary software.

We call this approach *Robotics and Automation as a Service* (RAaaS). To illustrate the concept, consider two scenarios for a graduate student setting up a robot workcell. The workcell contains a 7-DoF Fanuc industrial arm with parallel-jaw gripper and a Microsoft Kinect RGBD sensor. The purpose of the workcell is to pick up and inspect parts as they come down an assembly line, requiring object recognition and localization, grasp planning, and motion planning.

In Scenario 1 (today with ROS), the software runs locally. ROS (Robot Operating System), the well-known open-source library of robotics software [131], provides access to over 2000 open-source ROS packages. Currently however, ROS is only supported on the Ubuntu Linux operating system. While Ubuntu is popular, the computers available to the graduate student run OSX. Many stable ROS packages are provided as packages, which simplifies installation, but some software is only available as a source distribution, which requires the download and installation of dependencies. The graduate student must set up a new machine with Ubuntu and resolve all library dependencies, including those that conflict with other packages.

In contrast, Scenario 2 (in the future with RAaaS), the analysis and planning software runs in the Cloud. The graduate student visits a website to input the robot, sensor, and gripper models. She then selects her desired object recognition and localization, motion planning, and grasping algorithms, and uses a graphical interface to connect these algorithms into a pipeline. Her robot begins sending up data in the form of point clouds from the Kinect. The robot receives and executes motion plans and grasps, reporting back outcomes to the Cloud-based pipeline, which are combined with feedback from other robots to improve the Cloud-based software parameters over time. We are excited about the potential of such a system and actively working with others on developing its components.

This survey is based on research available in August 2014. A repository for new developments and updates is available at: <http://goldberg.berkeley.edu/cloud-robotics/>

ACKNOWLEDGMENT

The authors thank C. Agero, A. Aydemir, A. Bayen, George Bekey, K. Bekris, D. Berenson, G. Bradski, M. Ciocarlie, J. Civera, S. Cousins, D. Culler, R. D'Andrea, S. Gentner, B. Gerkey, E. Guizzo, T. Hunter, M. Ani Hsieh, V. Isler, R. Katz, J. Kuffner, V. Kumar, J. Malec, M. Mason, D. Portugal, B. Recht, L. Riazuelo, J. Salmern-Garca, S. Sarma, A. Saxena, D. Song, C. Sutter, R. Voyles, A. Waibel, L. Wang, and J. Wiegley for ongoing insights and advice on this topic.

REFERENCES

- [1] "Amazon Elastic Cloud (EC2)." [Online]. Available: <http://aws.amazon.com/ec2/>
- [2] "Amazon Web Services." [Online]. Available: <http://aws.amazon.com>
- [3] "Arduino." [Online]. Available: <http://www.arduino.cc>
- [4] "Bullet Physics Library." [Online]. Available: <http://bulletphysics.org>

- [5] "CloudSim." [Online]. Available: <http://gazebosim.org/wiki/CloudSim/>
- [6] "For Big-Data Scientists, Janitor Work Is Key Hurdle to Insights." [Online]. Available: <http://www.nytimes.com/2014/08/18/technology/for-big-data-scientists-hurdle-to-insights-is-janitor-work.html>
- [7] "Gazebo." [Online]. Available: <http://gazebosim.org/>
- [8] "Google Compute Engine." [Online]. Available: <https://cloud.google.com/products/compute-engine>
- [9] "Google Goggles." [Online]. Available: <http://www.google.com/mobile/goggles/>
- [10] "IEEE Networked Robots Technical Committee." [Online]. Available: <http://www-users.cs.umn.edu/~isler/tc/>
- [11] "Industry 4.0." [Online]. Available: <http://www.bmbf.de/en/19955.php>
- [12] "Microsoft Azure." [Online]. Available: <http://www.windowsazure.com>
- [13] "MyRobots.com." [Online]. Available: <http://myrobots.com>
- [14] "Open Source Robotics Foundation." [Online]. Available: <http://www.osrfoundation.org>
- [15] "OpenRAVE." [Online]. Available: <http://openrave.org/>
- [16] "RoboBrain." [Online]. Available: <http://robobrain.me>
- [17] "ROS-Industrial." [Online]. Available: <http://rosindustrial.org>
- [18] "ROS (Robot Operating System)." [Online]. Available: <http://ros.org>
- [19] "Rosjava, an Implementation of ROS in Pure Java With Android Support." [Online]. Available: <http://cloudrobotics.com>
- [20] "TOP500." [Online]. Available: <http://www.top500.org/list/2012/06/100>
- [21] "What is MyRobots?" [Online]. Available: <http://myrobots.com/wiki/About>
- [22] "What is RoboEarth?" [Online]. Available: <http://www.roboearth.org/what-is-roboearth>
- [23] "An Open-Source Robo-Surgeon," *The Economist*, 2012. [Online]. Available: <http://www.economist.com/node/21548489>
- [24] "DARPA Selects Southwest Research Institute to Support DARPA Robotics Challenge," *Market Watch*, 2013.
- [25] Accenture Inc., "A New Era for Energy Companies: Cloud Computing Changes the Game." 2012. [Online]. Available: <http://www.accenture.com/SiteCollectionDocuments/PDF/Accenture-New-Era-Energy-Companies-Cloud-Computing-changes-Game.pdf>
- [26] B. Addad, S. Amari, and J.-J. Lesage, "Analytic calculus of response time in networked automation systems," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 7, no. 4, pp. 858–869, Oct. 2010.
- [27] B. Addad, S. Amari, and J.-J. Lesage, "Client-server networked automation systems reactivity: Deterministic and probabilistic analysis," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 8, no. 3, pp. 540–548, Jul. 2011.
- [28] C. Aguero, N. Koenig, I. Chen, H. Boyer, S. Peters, J. Hsu, B. Gerkey, S. Paepcke, J. Rivero, J. Manzo, E. Krotkov, and G. Pratt, "Inside the virtual robotics challenge: Simulating real-time robotic disaster response," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 494–506, Apr. 2015.
- [29] M. Armbrust, I. Stoica, M. Zaharia, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, and A. Rabkin, "A view of cloud computing," *Commun. ACM*, vol. 53, no. 4, p. 50, Apr. 2010.
- [30] R. Arumugam, V. Enti, L. Bingbing, W. Xiaojun, K. Baskaran, F. Kong, A. Kumar, K. Meng, and G. Kit, "DAvinCi: A cloud computing framework for service robots," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, 2010, pp. 3084–3089.
- [31] D. Astély, E. Dahlman, A. Furuskär, Y. Jading, M. Lindström, and S. Parkvall, "LTE: The evolution of mobile broadband," *Commun. Mag.*, vol. 47, no. 4, pp. 44–51, 2009.
- [32] L. Atzori, A. Iera, and G. Morabito, "The Internet of things: A survey," *Comput. Netw.*, vol. 54, no. 15, pp. 2787–2805, Oct. 2010.
- [33] K. Bekris, R. Shome, A. Krontiris, and A. Dobson, "Cloud automation: Precomputing roadmaps for flexible manipulation," *IEEE Robot. Autom. Mag.: Special Issue on Emerging Advances and Applications in Automation*, 2014, under review.
- [34] W. Beksi and N. Papanikolopoulos, "Point cloud culling for robot vision tasks under communication constraints," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Chicago, IL, USA, 2014.
- [35] D. Berenson, P. Abbeel, and K. Goldberg, "A robot path planning framework that learns from experience," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2012, pp. 3671–3678.
- [36] B. Bhargava, P. Angin, and L. Duan, "A mobile-cloud pedestrian crossing guide for the blind," in *Proc. Int. Conf. Adv. Comput. Commun.*, 2011.
- [37] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—A survey," *IEEE Trans. Robotics (T-RO)*, no. 99, pp. 1–21, 2013.
- [38] V. Chandrasekaran and M. I. Jordan, "Computational and statistical tradeoffs via convex relaxation," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 13, pp. E1181–E1190, 2013.
- [39] M. Ciocarlie, K. Hsiao, E. G. Jones, S. Chitta, R. Rusu, and I. Sucan, "Towards reliable grasping and manipulation in household environments," in *Proc. Exp. Robot. Int. Symp. (ISER)*, Delhi, India, 2010.
- [40] M. Ciocarlie, C. Pantofaru, K. Hsiao, G. Bradski, P. Brook, and E. Dreyfuss, "A side of data with my robot," *IEEE Robot. Autom. Mag.*, vol. 18, no. 2, pp. 44–57, Jun. 2011.
- [41] Committee on the Analysis of Massive Data; Committee on Applied and Theoretical Statistics; Board on Mathematical Sciences and Their Applications; Division on Engineering and Physical Sciences; National Research Council, "Frontiers in massive data analysis," The National Academies Press. Washington, DC, USA, 2013.
- [42] L. Dabbish, C. Stuart, J. Tsay, and J. Herbsleb, "Social coding in github: Transparency and collaboration in an open software repository," in *Proc. ACM Conf. Comput. Supported Cooperative Work*, 2012, p. 1277.
- [43] S. Dalibard, A. Nakhaei, F. Lamiroux, and J.-P. Laumond, "Manipulation of documented objects by a walking humanoid robot," in *Proc. IEEE-RAS Int. Conf. Humanoid Robot.*, Dec. 2010, pp. 518–523.
- [44] R. D'Andrea, "Guest editorial: A revolution in the warehouse: A Retrospective on KIVA systems and the grand challenges ahead," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 9, no. 4, pp. 638–639, Oct. 2012.
- [45] H. Dang and P. K. Allen, "Learning grasp stability," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2012, pp. 2392–2397.
- [46] H. Dang, J. Weisz, and P. K. Allen, "Blind grasping: Stable robotic grasping using tactile feedback and hand kinematics," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2011, pp. 5917–5922.
- [47] S. Davidson, "Open-source Hardware," *IEEE Design Test Comput.*, vol. 21, no. 5, pp. 456–456, Sep. 2004.
- [48] G. de Croon, P. Gerke, and I. Sprinkhuizen-Kuyper, "Crowdsourcing as a methodology to obtain large and varied robotic data sets," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Chicago, IL, USA, 2014.
- [49] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, Q. V. Le, and A. Y. Ng, "Large scale distributed deep networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Eds., 2012, pp. 1223–1231.
- [50] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2009, pp. 248–255.
- [51] A. Dobson, A. Krontiris, and K. E. Bekris, "Sparse roadmap spanners," *Algorithmic Foundations of Robotics X*, pp. 279–296, 2013.
- [52] Droplet, "Cloud Powered Water Sprinkler System," 2014. [Online]. Available: <http://smartdroplet.com/>
- [53] Z. Du, W. Yang, Y. Chen, X. Sun, X. Wang, and C. Xu, "Design of a robot cloud center," in *Proc. Int. Symp. Auton. Decentralized Syst.*, Mar. 2011, pp. 269–275.
- [54] E. Dumbill, "Defining Big Data," [Online]. Available: [Forbes Data Driven Blog](http://ForbesDataDrivenBlog.com), 2014
- [55] P. C. Evans and M. Annunziata, "Industrial Internet: Pushing the boundaries of minds and machines," General Electric, Tech. Rep., 2012.
- [56] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The PASCAL visual object classes (VOC) challenge," *Int. J. Comput. Vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [57] S. Gammeter, A. Gassmann, L. Bossard, T. Quack, and L. Van Gool, "Server-side object recognition and client-side object tracking for mobile augmented reality," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, Jun. 2010, pp. 1–8.
- [58] J. García, "Using cloud computing as a HPC platform for embedded systems," 2011, Unpublished work. [Online]. Available: <http://www.atc.us.es/descargas/fmHPCcloud.pdf>
- [59] General Electric, "The case for an industrial big data platform: Laying the groundwork for the new industrial age," 2014. [Online]. Available: <https://gesoftware.com/IndustrialBigDataPlatform.pdf>
- [60] N. Gershenfeld, *Fab: The Coming Revolution on Your Desktop—from Personal Computers to Personal Fabrication*. New York, NY, USA: Basic Books, 2007.
- [61] Y. Gingold, A. Shamir, and D. Cohen-Or, "Micro perceptual human computation for visual tasks," *ACM Trans. Graphics*, vol. 31, no. 5, pp. 1–12, Aug. 2012.
- [62] J. Glover, D. Rus, and N. Roy, "Probabilistic models of object geometry for grasp planning," *Robotics: Science and Systems (RSS)*, vol. IV, pp. 278–285, 2008.
- [63] K. Goldberg, "Beyond the Web: Excavating the real world via mosaic," in *Proc. 2nd Int. World Wide Web Conf.*, 1994, pp. 1–12.

- [64] K. Goldberg and B. Chen, "Collaborative control of robot motion: Robustness to error," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2001, vol. 2, pp. 655–660.
- [65] K. Goldberg, M. Mascha, S. Gentner, N. Rothenberg, C. Sutter, and J. Wiegley, "Desktop teleoperation via the world wide web," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, Nagoya, Japan, 1995, pp. 654–659.
- [66] *Beyond Webcams: An Introduction to Online Robots*, K. Goldberg and R. Siegwart, Eds. Cambridge, MA, USA: MIT Press, 2002.
- [67] C. Goldfeder, M. Ciocarlie, and P. Allen, "The Columbia grasp database," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2009, pp. 1710–1716.
- [68] C. Goldfeder and P. K. Allen, "Data-driven grasping," *Auton. Robot.*, vol. 31, no. 1, pp. 1–20, Apr. 2011.
- [69] B. Gouveia, D. Portugal, D. Silva, and L. Marques, "Computation sharing in distributed robotic systems: A case study on SLAM," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 410–422, Apr. 2015.
- [70] E. Guizzo, "Cloud robotics: Connected to the cloud, robots get smarter," *IEEE Spectrum Automaton Blog*, 2011.
- [71] B. Hannaford, J. Rosen, D. Friedman, H. King, P. Roan, L. Cheng, D. Glozman, J. Ma, S. Kosari, and L. White, "Raven II: Open platform for surgical robotics research," *IEEE Trans. Biomed. Eng.*, pp. 954–959, Apr. 2013.
- [72] E. Hidalgo-Pena, L. F. Marin-Urias, F. Montes-Gonzalez, A. Marin-Hernandez, and H. V. Rios-Figueroa, "Learning from the Web: Recognition method based on object appearance from Internet images," in *Proc. ACM/IEEE Int. Conf. Human-Robot Interaction*, 2013, pp. 139–140.
- [73] J. Higuera, A. Xu, F. Shkurti, and G. Dudek, "Socially-driven collective path planning for robot missions," in *Proc. Conf. Comput. Robot Vision*, May 2012, pp. 417–424.
- [74] G. Hu, W. Tay, and Y. Wen, "Cloud robotics: Architecture, challenges and applications," *IEEE Network*, vol. 26, no. 3, pp. 21–28, 2012.
- [75] K. Huebner, K. Welke, M. Przybylski, N. Vahrenkamp, T. Asfour, and D. Kragic, "Grasping known objects with humanoid robots: A box-based approach," in *Proc. Int. Conf. Adv. Robot.*, Munich, Germany, 2009.
- [76] Humphreys, Todd, "Cockrell School researchers demonstrate first successful 'spoofing' of UAVs," 2012. [Online]. Available: <http://www.engr.utexas.edu/features/humphreysspoofing>
- [77] T. Hunter, T. Moldovan, M. Zaharia, S. Metzgui, J. Ma, M. J. Franklin, P. Abbeel, and A. M. Bayen, "Scaling the mobile millennium system in the cloud," in *Proc. 2nd ACM Symp. Cloud Compu. (SOCC)*, Cascais, Portugal, 2011, pp. 28:1–28:8.
- [78] D. Hunziker, M. Gajamohan, M. Waibel, and R. D'Andrea, "Rapyuta: The roboearth cloud engine," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Karlsruhe, Germany, 2013, pp. 438–444.
- [79] M. Inaba, "Remote-brained robots," in *Proc. Int. Joint Conf. Artif. Intell.*, 1997, pp. 1593–1606.
- [80] J. D. Irwin, *The Industrial Electronics Handbook*. Boca Raton, FL, USA: CRC Press, 1997.
- [81] P. F. Jacobs and D. T. Reid, *Rapid Prototyping & Manufacturing: Fundamentals of Stereolithography*. Dearborn, MI, USA: Soc. Manuf. Eng., 1992.
- [82] A. Jain, D. Das, and A. Saxena, "PlanIt: A crowdsourcing approach for learning to plan paths from large scale preference feedback," *Tech. Rep.*, 2014.
- [83] N. K. Jangid, "Real time cloud computing," in *Data Manage. Security*, 2011.
- [84] M. Johnson-Roberson, J. Bohg, G. Skantze, J. Gustafson, R. Carlson, B. Rasolzadeh, and D. Kragic, "Enhanced visual scene understanding through human-robot dialog," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Sep. 2011, pp. 3342–3348.
- [85] S. Jordan, T. Haidegger, L. Kovacs, I. Felde, and I. Rudas, "The rising prospects of cloud robotic applications," in *Proc. Int. Conf. Comput. Cybern. (ICCC)*, Jul. 2013, pp. 327–332.
- [86] G. Juve, E. Deelman, G. B. Berriman, B. P. Berman, and P. Maechling, "An evaluation of the cost and performance of scientific workflows on amazon EC2," *J. Grid Comput.*, vol. 10, no. 1, pp. 5–21, Mar. 2012.
- [87] K. Kamei, S. Nishio, N. Hagita, and M. Sato, "Cloud networked robotics," *IEEE Network*, vol. 26, no. 3, pp. 28–34, 2012.
- [88] A. Kasper, Z. Xue, and R. Dillmann, "The KIT object models database: An object model database for object recognition, localization and manipulation in service robotics," *Int. J. Robot. Res. (IJRR)*, vol. 31, no. 8, pp. 927–934, May 2012.
- [89] B. Kehoe, A. Matsukawa, S. Candido, J. Kuffner, and K. Goldberg, "Cloud-based robot grasping with the google object recognition engine," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, Karlsruhe, Germany, 2013, pp. 4263–4270.
- [90] B. Kehoe, D. Berenson, and K. Goldberg, "Estimating part tolerance bounds based on adaptive cloud-based grasp planning with slip," in *Proc. IEEE Int. Conf. Autom. Sci. Eng. (CASE)*, Seoul, Korea, 2012, pp. 1106–1113.
- [91] B. Kehoe, D. Berenson, and K. Goldberg, "Toward cloud-based grasping with uncertainty in shape: Estimating lower bounds on achieving force closure with zero-slip push grasps," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, St. Paul, MN, USA, May 2012, pp. 576–583.
- [92] B. Kehoe, D. Warrier, S. Patil, and K. Goldberg, "Cloud-based grasp planning for toleranced parts using parallelized Monte Carlo sampling," *IEEE Trans. Autom. Sci. Eng. (TASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 455–470, Apr. 2015.
- [93] J. Kelly, "The industrial internet and big data analytics: Opportunities and challenges," 2013. [Online]. Available: http://wikibon.org/wiki/v/The_Industrial_Internet_and_Big_Data_Analytics%3A_Opportunities_and_Challenges
- [94] D. Kent and S. Chernova, "Construction of an object manipulation database from grasp demonstrations," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Chicago, IL, USA, 2014.
- [95] W.-J. Kim, K. Ji, and A. Ambike, "Real-time operating environment for networked control systems," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 3, no. 3, pp. 287–296, 2006.
- [96] KIVA, "KIVA Systems," 2014. [Online]. Available: www.kivasystems.com
- [97] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Eds., 2012, pp. 1097–1105.
- [98] J. Kuffner, "Cloud-enabled robots," in *Proc. IEEE-RAS Int. Conf. Humanoid Robot.*, Nashville, TN, USA, 2010.
- [99] V. Kumar, R. Daniela, and G. S. Sukhatme, "Networked robotics," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Berlin, Germany: Springer-Verlag, 2004.
- [100] V. Kumar and N. Michael, "Opportunities and challenges with autonomous micro aerial vehicles," *Int. J. Robot. Res. (IJRR)*, vol. 31, no. 11, pp. 1279–1291, 2012.
- [101] K. Lai and D. Fox, "Object recognition in 3D point clouds using web data and domain adaptation," *Int. J. Robot. Res. (IJRR)*, vol. 29, no. 8, pp. 1019–1037, May 2010.
- [102] Q. Le, M. Ranzato, R. Monga, M. Devin, K. Chen, G. Corrado, J. Dean, and A. Ng, "Building high-level features using large scale unsupervised learning," in *Proc. Int. Conf. Mach. Learn.*, Vancouver, BC, Canada, 2012, pp. 8595–8598.
- [103] A. E. Leeper, K. Hsiao, M. Ciocarlie, L. Takayama, and D. Gossow, "Strategies for human-in-the-loop robotic grasping," in *Proc. ACM/IEEE Int. Conf. Human-Robot Interaction*, 2012, pp. 1–8.
- [104] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *Int. J. Robot. Res. (IJRR)*, to be published.
- [105] Z. Li, L. O'Brien, H. Zhang, and R. Cai, "On the conceptualization of performance evaluation of IaaS services," *IEEE Trans. Services Comput.*, to be published.
- [106] Z. Li, H. Zhang, L. O'Brien, R. Cai, and S. Flint, "On evaluating commercial Cloud services: A systematic review," *J. Syst. Softw.*, vol. 86, no. 9, pp. 2371–2393, 2013.
- [107] P. Lin, K. Abney, and G. A. Bekey, *Robot Ethics: The Ethical and Social Implications of Robotics*. Cambridge, MA, USA: MIT Press, 2011.
- [108] H. Lipson and M. Kurman, *Fabricated: The New World of 3D Printing*. New York, NY, USA: Wiley, 2013.
- [109] B. Liu, Y. Chen, E. Blasch, and K. Pham, "A holistic cloud-enabled robotics system for real-time video tracking application," in *Future Inform. Technol.*, 2014, pp. 455–468.
- [110] C.-H. Lu and L.-C. Fu, "Robust location-aware activity recognition using wireless sensor network in an attentive home," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 6, no. 4, pp. 598–609, 2009.
- [111] J. Mahler, S. Krishnan, M. Laskey, S. Sen, A. Murali, B. Kehoe, S. Patil, J. Wang, M. Franklin, P. Abbeel, and K. Goldberg, "Learning accurate kinematic control of cable-driven surgical robots using data cleaning and Gaussian process regression," in *Proc. IEEE Int. Conf. Autom. Sci. Eng. (CASE)*, Taipei, Taiwan, 2014.
- [112] G. McKee, "What is networked robotics?," *Informatics in Control Automation and Robotics*, vol. 15, pp. 35–45, 2008.
- [113] McKinsey Inc., "Big data and the opportunities it creates for semiconductor players," 2012. [Online]. Available: http://www.mckinsey.com/~/media/McKinsey/dotcom/client_service/Semiconductors/Issue/Autumn2012/PDFs/Big_data_and_the_opportunities_it_creates_for_semiconductor_players.aspx

- [114] P. Mehrotra, J. Djomehri, S. Heistand, R. Hood, H. Jin, A. Lazanoff, S. Saini, and R. Biswas, "Performance evaluation of Amazon EC2 for NASA HPC applications," in *Proc. 3rd Workshop on Scientific Cloud Computing Date—ScienceCloud'12*, 2012, pp. 41–50.
- [115] P. Mell and T. Grance, "The NIST definition of cloud computing," *National Institute of Standards and Technology*, 2009.
- [116] N. Michael, D. Mellinger, Q. Lindsey, and V. Kumar, "The GRASP multiple micro UAV testbed," *IEEE Robot. Autom. Mag.*, vol. 17, no. 3, pp. 56–65, 2010.
- [117] A. Miller and P. Allen, "GraspIt! A versatile simulator for robotic grasping," *IEEE Robot. Autom. Mag.*, vol. 11, no. 4, pp. 110–122, Dec. 2004.
- [118] G. Mohanarajah, D. Hunziker, M. Waibel, and R. D'Andrea, "Rapyuta: A cloud robotics platform," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 12, no. 2, pp. 481–493, Apr. 2015.
- [119] G. Mohanarajah, V. Usenko, M. Singh, M. Waibel, and R. D'Andrea, "Cloud-based collaborative 3D mapping in real-time with low-cost robots," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 423–431, Apr. 2015.
- [120] M. Moussa and M. Kamel, "An experimental approach to robotic grasping using a connectionist architecture and generic grasping functions," *IEEE Trans. Syst., Man, Cybern., Part C*, vol. 28, no. 2, pp. 239–253, May 1998.
- [121] M. Narita, S. Okabe, Y. Kato, Y. Murakawa, K. Okabayashi, and S. Kanda, "Reliable cloud-based robot services," in *Proc. Conf. IEEE Ind. Electron. Soc.*, 2013, pp. 8317–8322.
- [122] National Science Foundation, "Enabling a new future for cloud computing," 2014. [Online]. Available: http://nsf.gov/news/news_summ.jsp?cntn_id=132377
- [123] D. Nister and H. Stewenius, "Scalable recognition with a vocabulary tree," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, 2006, vol. 2, pp. 2161–2168.
- [124] D. Nurmi, R. Wolksi, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov, "The eucalyptus open-source cloud-computing system," in *Proc. IEEE/ACM Int. Symp. Cluster Comput. Grid*, 2009, pp. 124–131.
- [125] G. Oliveira and V. Isler, "View planning for cloud-based active object recognition," Dept. Comput. Sci., Univ. Minnesota, Minneapolis, MN, USA, Tech. Rep. 13-025, 2013.
- [126] E. Plaku, K. E. Bekris, B. Y. Chen, A. M. Ladd, and L. E. Kavraki, "Sampling-based roadmap of trees for parallel motion planning," *IEEE Trans. Robotics (T-RO)*, vol. 21, no. 4, pp. 597–608, 2005.
- [127] E. Plaku, K. E. Bekris, and L. E. Kavraki, "OOPS for motion planning: An online, open-source, programming system," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, Apr. 2007, pp. 3711–3716.
- [128] M. Popovic, G. Koostra, J. A. Jorgensen, D. Kragic, and N. Kruger, "Grasping unknown objects using an early cognitive vision system for general scene understanding," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Sep. 2011, pp. 987–994.
- [129] E. Prestes, J. L. Carbonera, S. Rama Fiorini, V. A. M. Jorge, M. Abel, R. Madhavan, A. Locoro, P. Goncalves, M. E. Barreto, M. Habib, A. Chibani, S. Gérard, Y. Amirat, and C. Schlenoff, "Towards a core ontology for robotics and automation," *Robot. Autom. Syst.*, vol. 61, no. 11, pp. 1193–1204, 2013.
- [130] A. A. Proia, D. Simshaw, and K. Hauser, "Consumer cloud robotics and the fair information practice principles: Recognizing the challenges and opportunities ahead," *Minnesota J. Law, Sci. Technol.*, to be published.
- [131] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: An open-source robot operating system," in *Proc. ICRA Workshop on Open Source Software*, Kobe, Japan, 2009.
- [132] A. J. Quinn and B. B. Bederson, "Human computation: A survey and taxonomy of a growing field," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2011, ser. CHI '11, pp. 1403–1412.
- [133] K. Ren, C. Wang, and Q. Wang *et al.*, "Security challenges for the public cloud," *IEEE Internet Comput.*, vol. 16, no. 1, pp. 69–73, 2012.
- [134] L. Riazuelo, J. Civera, and J. Montiel, "C2TAM: A cloud framework for cooperative tracking and mapping," *Robot. Autom. Syst.*, vol. 62, no. 4, pp. 401–413, 2013.
- [135] L. Riazuelo, M. Tenorth, D. D. Marco, M. Salas, L. Mosenlechner, L. Kunze, M. Beetz, J. D. Tardos, L. Montano, and J. Montiel, "RoboEarth web-enabled and knowledge-based active perception," in *Proc. IROS Workshop on AI-Based Robotics*, Tokyo, Japan, 2013.
- [136] L. Riazuelo, M. Tenorth, D. Marco, M. Salas, D. Galvez-Lopez, L. Mosenlechner, L. Kunze, M. Beetz, J. Tardos, L. Montano, and J. Montiel, "RoboEarth semantic mapping: A cloud enabled knowledge-based approach," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 432–443, Apr. 2015.
- [137] J. Salmeron-Garcia, F. Diaz-del Rio, P. Inigo-Blasco, and D. Cagigas, "A tradeoff analysis of a cloud-based robot navigation assistant using stereo image processing," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 444–454, Apr. 2015.
- [138] M. Samadi, T. Kollar, and M. Veloso, "Using the web to interactively learn to find objects," in *Proc. AAAI Conf. Artif. Intell.*, 2012, pp. 2074–2080.
- [139] C. Schmitt, "Security and privacy in the era of big data," 2014. [Online]. Available: <http://www.renci.org/wp-content/uploads/2014/02/0313WhitePaper-iRODS.pdf>
- [140] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun, "OverFeat: Integrated recognition, localization and detection using convolutional networks," *CoRR*, 2013.
- [141] M. Sevier, T. Fifield, and N. Katayama, "Belle Monte-Carlo production on the amazon EC2 cloud," *J. Phys.: Conf. Series*, vol. 219, no. 1, Apr. 2010.
- [142] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: Exploring photo collections in 3D," *ACM Trans. Graphics (TOG)*, vol. 25, no. 3, pp. 835–846, 2006.
- [143] D. Song, K. Goldberg, and N. Y. Chong, "Networked telerobots," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Berlin, Germany: Springer-Verlag, 2004.
- [144] A. Sorokin, D. Berenson, S. S. Srinivasa, and M. Hebert, "People helping robots helping people: Crowdsourcing for grasping novel objects," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Oct. 2010, pp. 2117–2122.
- [145] A. Sorokin and D. Forsyth, "Utility data annotation with amazon mechanical turk," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR) Workshops*, Jun. 2008, pp. 1–8.
- [146] M. Stenmark, J. Malec, K. Nilsson, and A. Robertsson, "On distributed knowledge bases for small-batch assembly," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 519–528, Apr. 2015.
- [147] M. Tenorth and M. Beetz, "KnowRob: A knowledge processing infrastructure for cognition-enabled robots," *Int. J. Robot. Res. (IJRR)*, vol. 32, no. 5, pp. 566–590, 2013.
- [148] M. Tenorth, A. Clifford Perzylo, R. Lafrenz, and M. Beetz, "The RoboEarth language: Representing and exchanging knowledge about actions, objects, environments," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2012, pp. 1284–1289.
- [149] M. Tenorth, A. C. Perzylo, R. Lafrenz, and M. Beetz, "Representation and exchange of knowledge about actions, objects, environments in the roboearth framework," *IEEE Trans. Autom. Sci. Eng. (T-ASE)*, vol. 10, no. 3, pp. 643–651, Jul. 2013.
- [150] The African Robotics Network (AFRON), "'Ten dollar robot' design challenge winners," 2012. [Online]. Available: http://robotics-africa.org/design_challenge.html
- [151] A. Torralba, R. Fergus, and W. T. Freeman, "80 million tiny images: A large data set for nonparametric object and scene recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 11, pp. 1958–1970, 2008.
- [152] R. Tudoran, A. Costan, G. Antoniu, and L. Bougé, "A performance evaluation of azure and nimbus clouds for scientific applications," in *Proc. Int. Workshop on Cloud Computing Platforms—CloudCP'12*, 2012, pp. 1–6.
- [153] L. Turnbull and B. Samanta, "Cloud robotics: Formation control of a multi robot system utilizing cloud infrastructure," in *Proc. IEEE Southeastcon*, 2013, pp. 1–4.
- [154] J. van den Berg, P. Abbeel, and K. Goldberg, "LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information," *Int. J. Robot. Res. (IJRR)*, vol. 30, no. 7, pp. 895–913, Jun. 2011.
- [155] L. von Ahn, "Human computation," Ph.D. dissertation, Carnegie Mellon Univ., Pittsburgh, PA, USA, 2009.
- [156] M. Waibel, M. Beetz, J. Civera, R. D'Andrea, J. Elfring, D. Gálvez-López, K. Häussermann, R. Janssen, J. Montiel, A. Perzylo, B. Schieble, M. Tenorth, O. Zweigle, and R. De Molengraft, "RoboEarth," *IEEE Robot. Autom. Mag.*, vol. 18, no. 2, pp. 69–82, Jun. 2011.
- [157] H. Wang, Y. Ma, G. Pratz, and L. Xing, "Toward real-time Monte Carlo simulation using a commercial cloud computing infrastructure," *Phys. Med. Bio.*, vol. 56, no. 17, pp. N175–N181, Sep. 2011.
- [158] J. Wang, S. Krishnan, M. J. Franklin, K. Goldberg, T. Kraska, and T. Milo, "A Sample-and-clean framework for fast and accurate query processing on dirty data," in *Proc. ACM SIGMOD Int. Conf. Manage. Data, Snowbird, UT, USA*, 2014.

- [159] L. Wang, M. Liu, and M. Q.-H. Meng, "Real-time multi-sensor data retrieval for cloud robotic systems," *IEEE Trans. Autom. Sci. Eng. (T-ASE): Special Issue Cloud Robot. Autom.*, vol. 12, no. 2, pp. 507–518, Apr. 2015.
- [160] We Robot, "We robot conference," 2012. [Online]. Available: <http://robots.law.miami.edu/>
- [161] R. H. Weber, "Internet of things-new security and privacy challenges," *Comput. Law Security Rev.*, vol. 26, no. 1, pp. 23–30, 2010.
- [162] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2012, pp. 557–562.
- [163] E. Wilhelm, J. Siegel, S. Mayer, J. Paefgen, M. Tiefenbeck, M. Bicker, S. Ho, R. Dantu, and S. Sarma, "CloudThink: An open standard for projecting objects into the cloud," 2013. [Online]. Available: <http://cloud-think.com/>
- [164] Willow Garage, "Personal service robotics with tiered human-in-the-loop assistance," NSF SBIR Award 1256643 (PI: Matei Ciocarlie), 2013.
- [165] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: State-of-the-art and research challenges," *J. Internet Services App.*, vol. 1, no. 1, pp. 7–18, 2010.



Ben Kehoe received the B.A. degree in physics and mathematics from Hamline University, Saint Paul, MN, USA, in 2006. He is now working towards the Ph.D. degree at the Department of Mechanical Engineering, University of California, Berkeley, CA, USA.

His research interests include cloud robotics, medical robotics, controls, and grasping.



Sachin Patil received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Bombay, India, in 2006, and the Ph.D. degree in computer science from the University of North Carolina at Chapel Hill, NC, USA, in 2012.

He is now a Postdoctoral Researcher with the Department Electrical Engineering and Computer Science, University of California, Berkeley, USA. His research interests include motion planning, cloud robotics, and medical robotics.



Pieter Abbeel received the B.S. and M.S. degree in electrical engineering from KU Leuven, Leuven, Belgium, and the Ph.D. degree in computer science from Stanford University, Stanford, CA, USA, in 2008.

He joined the faculty at the University of California, Berkeley, CA, USA, in the Fall of 2008, with an appointment in the Department of Electrical Engineering and Computer Sciences. His current research focuses on robotics and machine learning with a particular focus on challenges in personal robotics, surgical robotics and connectomics. He

Dr. Abbeel has won numerous awards including Best Paper Awards at ICML and ICRA, the Sloan Fellowship, the Air Force Office of Scientific Research Young Investigator Program (AFOSR-YIP) Award, the Office of Naval Research Young Investigator Program (ONR-YIP) Award, the DARPA Young Faculty Award (DARPA-YFA), the Okawa Foundation Award, the TR35, the IEEE Robotics and Automation Society (RAS) Early Career Award, and the Dick Volz Best U.S. Ph.D. Thesis in Robotics and Automation Award.



Ken Goldberg (S'84–M'90–SM'98–F'05) received the Ph.D. degree in computer science from Carnegie Mellon University, Pittsburgh, PA, USA, in 1990.

He is a Professor of Industrial Engineering and Operations Research with the University of California, Berkeley, CA, USA, with appointments in Electrical Engineering, Computer Science, Art Practice, and the School of Information. He was appointed Editor-in-Chief of the IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING (T-ASE) in 2011 and served two terms (2006–2009)

as Vice-President of Technical Activities for the IEEE Robotics and Automation Society. He has published over 200 refereed papers and awarded eight U.S. patents.

Prof. Goldberg is Founding Co-Chair of the IEEE Technical Committee on Networked Robots and Founding Chair of the (T-ASE) Advisory Board. He is the recipient of the NSF Presidential Faculty Fellowship (1995), the Joseph Engelberger Award (2000), and the IEEE Major Educational Innovation Award (2001).