

A Survey of Research on Cloud Robotics and Automation

Ben Kehoe *Student Member, IEEE*, Sachin Patil *Member, IEEE*,
Pieter Abbeel *Senior Member, IEEE*, 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—Cloud Automation, Cloud Robotics, Big Data, Cloud Computing, Open Source, Crowdsourcing

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



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, house-keeping, and surgery.

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 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 1980’s, 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 1990’s 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 1990’s, 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] 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

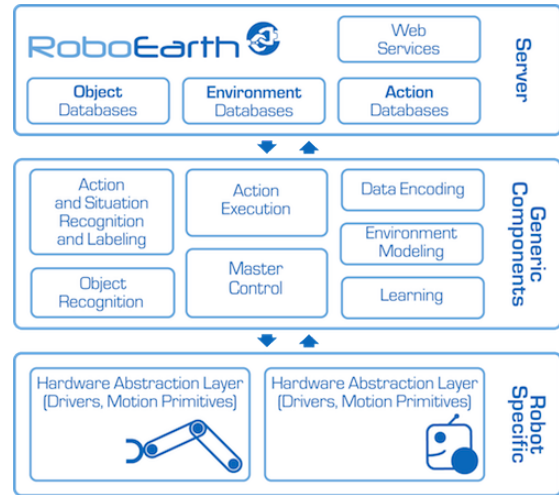


Fig. 2. The RoboEarth systems architecture designed to allow robots to share data and learn from each other [22], [156]. (Image reproduced with permission from authors).

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 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].




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).

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], [165]. For example, sampling algorithms can provide reasonable approximations to queries on large datasets to keep running times manageable [38], but these approximations can be seriously affected by “dirty data” [158].

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. 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).

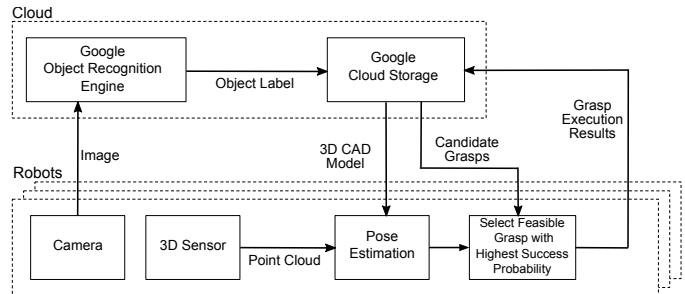


Fig. 5. System Architecture for Cloud-based object recognition for grasping. The robot captures an image of an object and sends via the network to the Google object recognition server. The server processes the image and returns data for a set of candidate objects, each with pre-computed grasping options. 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. 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 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.




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].

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],





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).




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).

[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].




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).

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.




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).

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.

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].

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, publications, 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




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).

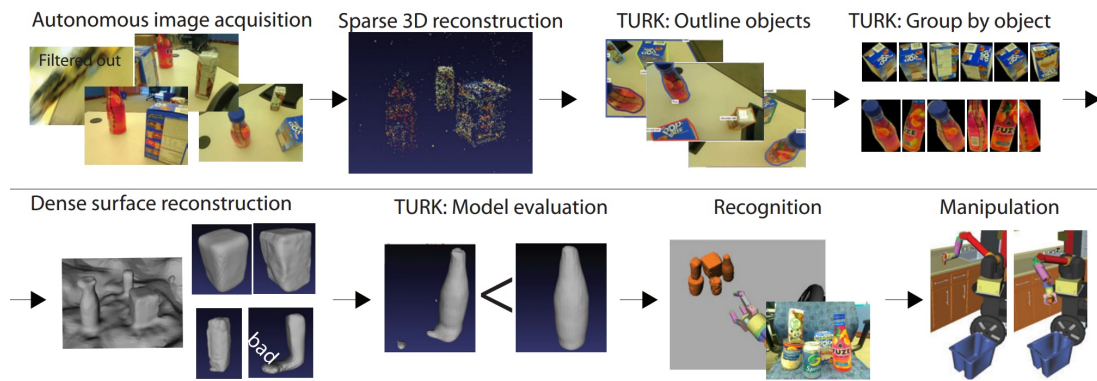


Fig. 12. Crowdsourcing object identification to facilitate robot grasping. Sorokin et al. developed a Cloud robot system that incorporates Amazon’s Mechanical Turk to obtain semantic information about the world and subjective judgments [144]. This work uses the Human Computation aspect of Cloud Robotics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



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).

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




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).

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 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 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/>

ACKNOWLEDGMENTS

This research was supported in part by funding from Google and Cisco and by the U.S. National Science Foundation under Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems.

The authors thank Carlos Agero, Alper Aydemir, Alexandre Bayen, George Bekey, Kostas Bekris, Dmitry Berenson, Gary Bradski, Matei Ciocarlie, Javier Civera, Steve Cousins, David Culler, Raffaello D'Andrea, Steven Gentner, Brian Gerkey, Erico Guizzo, Timothy Hunter, M. Ani Hsieh, Volkan Isler, Randy Katz, James Kuffner, Vijay Kumar, Jacek Malec, Matt Mason, David Portugal, Ben Recht, Luis Riazuelo, Javier Salmern-Garca, Sanjay Sarma, Ashutosh Saxena, Dezhen Song, Carl Sutter, Richard Voyles, Alex Waibel, Lujia Wang, and Jeff 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.robearth.org/what-is-robearth>
- [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." [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 Transactions on Automation Science and Engineering (T-ASE)*, vol. 7, no. 4, pp. 858–869, 2010.
- [27] —, "Client-Server Networked Automation Systems Reactivity: Deterministic and Probabilistic Analysis," *IEEE Transactions on Automation Science and Engineering (T-ASE)*, vol. 8, no. 3, pp. 540–548, 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 Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, 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," *Communications of the 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 *International Conference on Robotics and Automation (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," *Comm. Mag.*, vol. 47, no. 4, pp. 44–51, 2009.
- [32] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A Survey," *Computer Networks*, 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 Robotics & Automation Magazine: Special Issue on Emerging Advances and Applications in Automation*, p. Under Review, 2014.
- [34] W. Beksi and N. Papanikolopoulos, "Point cloud culling for robot vision tasks under communication constraints," in *International Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [35] D. Berenson, P. Abbeel, and K. Goldberg, "A Robot Path Planning Framework that Learns from Experience," in *International Conference on Robotics and Automation (ICRA)*, May 2012, pp. 3671–3678.
- [36] B. Bhargava, P. Angin, and L. Duan, "A Mobile-Cloud Pedestrian Crossing Guide for the Blind," in *International Conference on Advances in Computing & Communication*, 2011.
- [37] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-Driven Grasp Synthesis—A Survey," *IEEE Transactions on Robotics (T-RO)*, no. 99, pp. 1–21, 2013.
- [38] V. Chandrasekaran and M. I. Jordan, "Computational and statistical tradeoffs via convex relaxation," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 110, no. 13, pp. E1181–90, 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 *International Symposium on Experimental Robotics (ISER)*, 2010.
- [40] M. Ciocarlie, C. Pantofaru, K. Hsiao, G. Bradski, P. Brook, and E. Dreyfuss, "A Side of Data With My Robot," *IEEE Robotics & Automation Magazine*, vol. 18, no. 2, pp. 44–57, June 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, 2013.
- [42] L. Dabbish, C. Stuart, J. Tsay, and J. Herbsleb, "Social Coding in GitHub: Transparency and Collaboration in an Open Software Repository," in *ACM 2012 Conference on Computer Supported Cooperative Work*, 2012, p. 1277.
- [43] S. Dalibard, A. Nakhaei, F. Lamiraux, and J.-P. Laumond, "Manipulation of Documented Objects by a Walking Humanoid Robot," in *IEEE-RAS International Conference on Humanoid Robots*, 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 Transactions on Automation Science and Engineering (T-ASE)*, vol. 9, no. 4, pp. 638–639, 2012.
- [45] H. Dang and P. K. Allen, "Learning Grasp Stability," in *International Conference on Robotics and Automation (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 *International Conference on Robotics and Automation (ICRA)*, May 2011, pp. 5917–5922.
- [47] S. Davidson, "Open-source Hardware," *IEEE Design and Test of Computers*, vol. 21, no. 5, pp. 456–456, Sept. 2004.
- [48] G. de Croon, P. Gerke, and I. Sprinkhuizen-Kuyper, "Crowdsourcing as a methodology to obtain large and varied robotic data sets," in *International Conference on Intelligent Robots and Systems (IROS)*, 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. Curran Associates, Inc., 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [51] A. Dobson, A. Kroutiris, and K. E. Bekris, "Sparse Roadmap Spanners," in *Algorithmic Foundations of Robotics X*, 2013, pp. 279–296.
- [52] Droplet, "Cloud Powered Water Sprinkler System." [Online]. Available: <http://smardroplet.com/>
- [53] Z. Du, W. Yang, Y. Chen, X. Sun, X. Wang, and C. Xu, "Design of a Robot Cloud Center," in *International Symp. on Autonomous Decentralized Systems*, Mar. 2011, pp. 269–275.
- [54] E. Dumbill, "Defining Big Data," Forbes Data Driven Blog, 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," *International Journal of Computer 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 *IEEE Conference on Computer Vision and Pattern Recognition*, June 2010, pp. 1–8.
- [58] J. García, "Using Cloud Computing as a HPC Platform for Embedded Systems," 2011.
- [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/Industrial_Big_Data_Platform.pdf
- [60] N. Gershenfeld, *Fab: The Coming Revolution on Your Desktop—from Personal Computers to Personal Fabrication*. Basic Books, 2007.
- [61] Y. Gingold, A. Shamir, and D. Cohen-Or, "Micro Perceptual Human Computation for Visual Tasks," *ACM Trans. on 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," in *Robotics: Science and Systems (RSS)*, 2008.
- [63] K. Goldberg, "Beyond the Web: Excavating the Real World via Mosaic," in *Second International World Wide Web Conference*, 1994, pp. 1–12.
- [64] K. Goldberg and B. Chen, "Collaborative control of robot motion: Robustness to error," in *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, vol. 2. IEEE, 2001, 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 *International Conference on Robotics and Automation (ICRA)*, 1995.
- [66] K. Goldberg and R. Siegwart, Eds., *Beyond Webcams: An Introduction to Online Robots*. MIT Press, 2002.
- [67] C. Goldfeder, M. Cioarlie, and P. Allen, "The Columbia Grasp Database," in *International Conference on Robotics and Automation (ICRA)*, May 2009, pp. 1710–1716.
- [68] C. Goldfeder and P. K. Allen, "Data-Driven Grasping," *Autonomous Robots*, 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 Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, 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. Biomedical Engineering*, pp. 954–959, 2012.
- [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 *ACM/IEEE International Conference on Human-Robot Interaction*. IEEE, 2013, pp. 139–140.
- [73] J. Higuera, A. Xu, F. Shkurti, and G. Dudek, "Socially-Driven Collective Path Planning for Robot Missions," *Conference on Computer and Robot Vision*, pp. 417–424, May 2012.
- [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 *International Conference on Advanced Robotics*, 2009.
- [76] Humphreys, Todd, "Cockrell School Researchers Demonstrate First Successful "Spoofing" of UAVs." [Online]. Available: <http://www.engr.utexas.edu/features/humphreysspoofing>
- [77] T. Hunter, T. Moldovan, M. Zaharia, S. Merzgui, J. Ma, M. J. Franklin, P. Abbeel, and A. M. Bayen, "Scaling the Mobile Millennium System in the Cloud," in *2nd ACM Symposium on Cloud Computing (SOCC)*, 2011, p. 28.
- [78] D. Hunziker, M. Gajamohan, M. Waibel, and R. D'Andrea, "Rapyuta: The roboearth cloud engine," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2013.
- [79] M. Inaba, "Remote-brained robots," in *International Joint Conference on Artificial Intelligence*, 1997, pp. 1593–1606.
- [80] J. D. Irwin, *The Industrial Electronics Handbook*. CRC Press, 1997.
- [81] P. F. Jacobs and D. T. Reid, *Rapid Prototyping & Manufacturing: Fundamentals of Stereolithography*. Society of Manufacturing Engineers, 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 Management & 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 *International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2011, pp. 3342–3348.
- [85] S. Jordan, T. Haidegger, L. Kovacs, I. Felde, and I. Rudas, "The rising prospects of cloud robotic applications," in *International Conference on Computational Cybernetics (ICCC)*, July 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," *Journal of Grid Computing*, 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," *International Journal of Robotics Research (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 *International Conference on Robotics and Automation (ICRA)*, 2013.
- [90] B. Kehoe, D. Berenson, and K. Goldberg, "Estimating Part Tolerance Bounds Based on Adaptive Cloud-Based Grasp Planning with Slip," in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2012.
- [91] —, "Toward Cloud-based Grasping with Uncertainty in Shape: Estimating Lower Bounds on Achieving Force Closure with Zero-slip Push Grasps," in *International Conference on Robotics and Automation (ICRA)*, May 2012.
- [92] B. Kehoe, D. Warrier, S. Patil, and K. Goldberg, "Cloud-Based Grasp Planning for Toleranced Parts Using Parallelized Monte Carlo Sampling," *IEEE Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, 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 *International Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [95] W.-J. Kim, K. Ji, and A. Ambike, "Real-time operating environment for networked control systems," *IEEE Transactions on Automation Science and Engineering (T-ASE)*, vol. 3, no. 3, pp. 287–296, 2006.
- [96] KIVA, "KIVA Systems." [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. Curran Associates, Inc., 2012, pp. 1097–1105.

- [98] J. Kuffner, "Cloud-Enabled Robots," in *IEEE-RAS International Conference on Humanoid Robots*, 2010.
- [99] V. Kumar, R. Daniela, and G. S. Sukhatme, "Networked Robotics," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer Berlin Heidelberg, 2004.
- [100] V. Kumar and N. Michael, "Opportunities and Challenges with Autonomous Micro Aerial Vehicles," *International Journal of Robotics Research (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," *International Journal of Robotics Research (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 *International Conference in Machine Learning*, 2012.
- [103] A. E. Leeper, K. Hsiao, M. Ciocarlie, L. Takayama, and D. Gossow, "Strategies for Human-in-the-loop Robotic Grasping," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2012, pp. 1–8.
- [104] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *International Journal of Robotics Research (IJRR)*, p. To appear, 2014.
- [105] Z. Li, L. O'Brien, H. Zhang, and R. Cai, "On the Conceptualization of Performance Evaluation of IaaS Services," *IEEE Transactions on Services Computing*, vol. X, no. X, pp. 1–1, 2014.
- [106] Z. Li, H. Zhang, L. O'Brien, R. Cai, and S. Flint, "On evaluating commercial Cloud services: A systematic review," *Journal of Systems and Software*, 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*. The MIT Press, 2011.
- [108] H. Lipson and M. Kurman, *Fabricated: The New World of 3D Printing*. 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 Information Technology*, 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 Transactions on Automation Science and Engineering (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 *IEEE International Conference on Automation Science and Engineering (CASE)*, 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." [Online]. Available: http://www.mckinsey.com/~media/McKinsey/dotcom/client_service/Semiconductors/Issue\Autumn\2012/PDFs/Big_data_and_the_opportunities_it_creates_for_semiconductor_players.ashx
- [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 *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 Robotics & Automation Magazine*, vol. 17, no. 3, pp. 56–65, 2010.
- [117] A. Miller and P. Allen, "GraspIt! A Versatile Simulator for Robotic Grasping," *IEEE Robotics & Automation Magazine*, 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 Transactions on Automation Science and Engineering (T-ASE)*, pp. 1–13, July 2014.
- [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 Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, Apr. 2015.
- [120] M. Moussa and M. Kamel, "An Experimental Approach to Robotic Grasping using a Connectionist Architecture and Generic Grasping Functions," *IEEE Trans. on Systems, Man and Cybernetics, Part C*, vol. 28, no. 2, pp. 239–253, May 1998.
- [121] M. Narita, S. Okabe, Y. Kato, Y. Murakwa, K. Okabayashi, and S. Kanda, "Reliable cloud-based robot services," in *Conference of the IEEE Industrial Electronics Society*. IEEE, 2013, pp. 8317–8322.
- [122] National Science Foundation, "Enabling a new future for cloud computing." [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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 2006, pp. 2161–2168.
- [124] D. Nurmi, R. Wolski, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov, "The Eucalyptus Open-Source Cloud-Computing System," in *IEEE/ACM International Symp. on Cluster Computing and the Grid*, 2009, pp. 124–131.
- [125] G. Oliveira and V. Isler, "View Planning For Cloud-Based Active Object Recognition," Department of Computer Science, University of Minnesota, Tech. Rep., 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 Transactions on 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 *International Conference on Robotics and Automation (ICRA)*, Apr. 2007, pp. 3711–3716.
- [128] M. Popovic, G. Kootstra, J. A. Jorgensen, D. Kragic, and N. Kruger, "Grasping Unknown Objects using an Early Cognitive Vision System for General Scene Understanding," in *International Conference on Intelligent Robots and Systems (IROS)*, Sept. 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," *Robotics and Autonomous Systems*, 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 Journal of Law, Science & Technology*, p. to appear, 2014.
- [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 *ICRA Workshop on Open Source Software*, 2009.
- [132] A. J. Quinn and B. B. Bederson, "Human Computation: A Survey and Taxonomy of a Growing Field," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '11, 2011, pp. 1403–1412.
- [133] K. Ren, C. Wang, Q. Wang, et al., "Security Challenges for the Public Cloud," *IEEE Internet Computing*, 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," *Robotics and Autonomous Systems*, 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 *IROS Workshop on AI-based Robotics*, 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 Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, Apr. 2015.
- [137] J. Salmeron-Garcia, F. Diaz-del Rio, P. Inigo-Blasco, and D. Cagigas, "A Trade-off Analysis of a Cloud-based Robot Navigation Assistant using Stereo Image Processing," *IEEE Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, Apr. 2015.
- [138] M. Samadi, T. Kollar, and M. Veloso, "Using the Web to Interactively Learn to Find Objects," in *AAAI Conference on Artificial Intelligence*, 2012, pp. 2074–2080.
- [139] Schmitt, Charles, "Security and Privacy in the Era of Big Data." [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. Seivior, T. Fifield, and N. Katayama, "Belle Monte-Carlo Production on the Amazon EC2 Cloud," *Journal of Physics: Conference Series*, vol. 219, no. 1, Apr. 2010.
- [142] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: exploring photo collections in 3D," *ACM transactions on 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. Springer Berlin Heidelberg, 2004.
- [144] A. Sorokin, D. Berenson, S. S. Srinivasa, and M. Hebert, "People Helping Robots Helping People: Crowdsourcing for Grasping Novel Objects," *International Conference on Intelligent Robots and Systems (IROS)*, pp. 2117–2122, Oct. 2010.
- [145] A. Sorokin and D. Forsyth, "Utility Data Annotation with Amazon Mechanical Turk," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2008, pp. 1–8.
- [146] M. Stenmark, J. Malec, K. Nilsson, and A. Robertsson, "On Distributed Knowledge Bases for Small-Batch Assembly," *IEEE Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, Apr. 2015.
- [147] M. Tenorth and M. Beetz, "KnowRob: A Knowledge Processing Infrastructure for Cognition-Enabled Robots," *International Journal of Robotics Research (IJRR)*, vol. 32, no. 5, pp. 566–590, 2013.
- [148] M. Tenorth, A. Clifford Perzlyo, R. Lafrenz, and M. Beetz, "The RoboEarth language: Representing and exchanging knowledge about actions, objects, and environments," in *International Conference on Robotics and Automation (ICRA)*, May 2012, pp. 1284–1289.
- [149] M. Tenorth, A. C. Perzlyo, R. Lafrenz, and M. Beetz, "Representation and Exchange of Knowledge about Actions, Objects, and Environments in the Roboearth Framework," *IEEE Transactions on Automation Science and Engineering (T-ASE)*, vol. 10, no. 3, pp. 643–651, 2013.
- [150] The African Robotics Network (AFRON), "'Ten Dollar Robot' Design Challenge Winners." [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 Transactions on Pattern Analysis and Machine Intelligence*, 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 *International Workshop on Cloud Computing Platforms - CloudCP '12*. ACM Press, 2012, pp. 1–6.
- [153] L. Turnbull and B. Samanta, "Cloud robotics: Formation control of a multi robot system utilizing cloud infrastructure," in *Proceedings of IEEE Southeastcon*. IEEE, 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," *International Journal of Robotics Research (IJRR)*, vol. 30, no. 7, pp. 895–913, June 2011.
- [155] L. von Ahn, "Human Computation," Ph.D. dissertation, Carnegie Mellon University, 2009.
- [156] M. Waibel, M. Beetz, J. Civera, R. D'Andrea, J. Elfring, D. Gálvez-López, K. Häussermann, R. Janssen, J. Montiel, A. Perzlyo, B. Schieß le, M. Tenorth, O. Zweigle, and R. De Molengraft, "RoboEarth," *IEEE Robotics & Automation Magazine*, vol. 18, no. 2, pp. 69–82, June 2011.
- [157] H. Wang, Y. Ma, G. Pratz, and L. Xing, "Toward Real-time Monte Carlo Simulation using a Commercial Cloud Computing Infrastructure," *Physics in Medicine and Biology*, vol. 56, no. 17, pp. N175–81, Sept. 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 *ACM SIGMOD International Conference on Management of Data*, 2014.
- [159] L. Wang, M. Liu, and M. Q.-H. Meng, "Real-time Multi-sensor Data Retrieval for Cloud Robotic Systems," *IEEE Transactions on Automation Science and Engineering (T-ASE): Special Issue on Cloud Robotics and Automation*, vol. 12, no. 2, p. To appear, Apr. 2015.
- [160] We Robot, "We Robot Conference." [Online]. Available: <http://robots.law.miami.edu/>
- [161] R. H. Weber, "Internet of Things—New Security and Privacy Challenges," *Computer Law & Security Review*, 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 *International Conference on Robotics and Automation (ICRA)*. IEEE, 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," *Journal of Internet Services and Applications*, vol. 1, no. 1, pp. 7–18, 2010.



Ben Kehoe received his B.A. in Physics and Mathematics from Hamline University in 2006. He is now a Ph.D. student in the ME department at University of California, Berkeley. His research interests include cloud robotics, medical robotics, controls, and grasping.



Sachin Patil received his B.Tech degree in Computer Science and Engineering from the Indian Institute of Technology, Bombay, in 2006 and his Ph.D. degree in Computer Science from the University of North Carolina at Chapel Hill, NC in 2012. He is now a postdoctoral researcher in the EECs department at University of California, Berkeley. His research interests include motion planning, cloud robotics, and medical robotics.



Pieter Abbeel received a BS/MS in Electrical Engineering from KU Leuven (Belgium) and received his Ph.D. degree in Computer Science from Stanford University in 2008. He joined the faculty at UC Berkeley in Fall 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 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 is Professor of Industrial Engineering and Operations Research at UC Berkeley, 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. Goldberg earned his PhD in Computer Science from Carnegie Mellon University in 1990. Goldberg is Founding Co-Chair

of the IEEE Technical Committee on Networked Robots and Founding Chair of the (T-ASE) Advisory Board. Goldberg has published over 200 refereed papers and awarded eight US patents, the NSF Presidential Faculty Fellowship (1995), the Joseph Engelberger Award (2000), the IEEE Major Educational Innovation Award (2001) and in 2005 was named IEEE Fellow.