Cloud Robotics and Automation: A Survey of Related Work

Ken Goldberg
Ben Kehoe

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-5
http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-5.html

January 27, 2013

Acknowledgement

The authors thank James Kuffner, Steve Cousins, Rafaello D'Andrea, Alex Waibel, Gary Bradski, Vijay Kumar, Erico Guizzo, Dmitry Berenson, and Pieter Abbeel for ongoing insights and advice on this topic. This work was supported in part by NSF Award IIS-1227406.

Contact
Ken Goldberg, craigslist Distinguished Professor of New Media
Professor, IEOR and EECS, College of Engineering, UC Berkeley
Dept of Art Practice and School of Information, UC Berkeley
Professor, Department of Radiation Oncology, UC San Francisco

Contact: 425 Sutardja Dai Hall, Berkeley, CA 94720-1758
twitter: @Ken Goldberg | g+: http://j.mp/Ken_Goldberg
(510) 643-9565 | goldberg@berkeley.edu | http://goldberg.berkeley.edu

1 Cloud Robotics and Automation

What if robots and automation systems were not limited by onboard computation, memory, or programming? This is now practical with wireless networking and rapidly expanding Internet resources. In 2010, James Kuffner at Google introduced the term "Cloud Robotics" [54] to describe a new approach to robotics that takes advantage of the Internet as a resource for massively parallel computation and real-time sharing of vast data resources. The Google autonomous driving project exemplifies this approach: the system indexes maps and images that are collected and updated by satellite, Streetview, and crowdsourcing from the network to facilitate accurate localization. Another example is Kiva Systems new approach to warehouse automation and logistics using large numbers of mobile platforms to move pallets using a local network to coordinate planforms and update tracking data. These are just two new projects that build on resources from the Cloud. Steve Cousins of Willow Garage aptly summarized the idea: "No robot is an island." Cloud Robotics recognizes the wide availability of networking, incorporates elements of open-source, open-access, and crowdsourcing to greatly extend earlier concepts of "Online Robots" [36] and "Networked Robots" [35, 56].

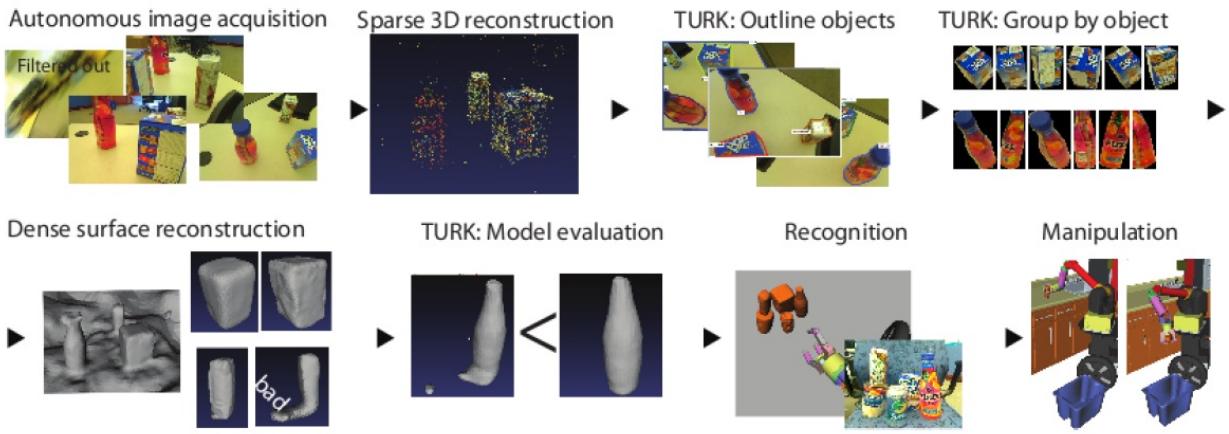

Figure 1: A cloud robot system that incorporates Amazon's Mechanical Turk to "crowdsource" object identification to facilitate robot grasping [68]. (Image reproduced with permission from authors).

The Cloud has been used as a metaphor for the Internet since the inception of the World Wide Web in the early 1990's. As of 2012, researchers are pursuing a number of cloud robotics and automation projects [39] [70] . New resources range from software architectures [19] [30] [42] [48] to computing resources [44]. The RoboEarth project [74] aims to develop "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 their environment" [15]. Cloud Robotics and Automation is related to concepts of the "Internet of Things" [20] and the "Industrial Internet," which envision how RFID and inexpensive processors can be incorporated into a vast array of objects from inventory items to household appliances to allow them to communicate and share information.

This report reviews five ways that Cloud Robotics and Automation has potential to improve performance: 1) providing access to global libraries of images, maps, and object data, eventually annotated with geometry and mechanical properties, 2) massively-parallel computation on demand for demanding tasks like optimal motion planning and sample-based statistical modeling, 3) robot sharing of outcomes, trajectories, and dynamic control policies, 4) human sharing of "open-source" code, data, and designs for programming, experimentation, and hardware construction, and 5) on-demand human guidance ("call centers") for exception handling and error recovery.

Updated information and links are available at: `http://goldberg.berkeley.edu/cloud-robotics/`

1.1 Big Data

The term “Big Data” describes data sets that are beyond the capabilities of standard relational database systems, which describes the growing library of images, maps, and many other forms of data relevant to robotics and automation on the Internet. One example is grasping, where online datasets can be consulted to determine appropriate grasps. The Columbia Grasp dataset [37] and the MIT KIT object dataset [49] are available online and have been widely used to evaluate grasping algorithms [28] [27] [76] [64].

Related work explores how computer vision can be used with Cloud resources to incrementally learn grasp strategies [24] [59] by matching sensor data against 3D CAD models in an online database. Examples of sensor data include 2D image features [43], 3D features [38], and 3D point clouds [23]. Google Goggles [7], a free network-based image recognition service for mobile devices, has been incorporated into a system for robot grasping [50] as illustrated in Figure 2.

Dalibard et al. attach “manuals” of manipulation tasks to objects [26]. The RoboEarch project stores data related to objects maps, and tasks, for applications ranging from object recognition to mobile navigation to grasping and manipulation (see Figure 5) [74].

As noted below, online datasets are effectively used to facilitate learning in computer vision. By leveraging Google’s 3D warehouse, [55] reduced the need for manually labeled training data. Using community photo collections, [31] created an augmented reality application with processing in the cloud.

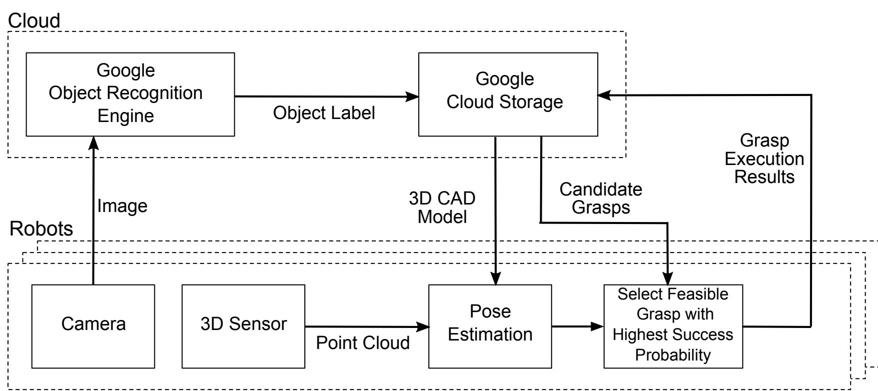

Figure 2: System Architecture for cloud-based object recognition for grasping. The robot captures an image of an object and sends via the network to the Google object recognition server. The server processes the image and returns data for a set of candidate objects, each with pre-computed grasping options. 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 [50]. (Image reproduced with permission from authors).

1.2 Cloud Computing

As of 2012, Cloud Computing services like Amazon’s EC2 elastic computing engine provide massively-parallel computation on demand [18]. Examples include Amazon Web Services [2] Elastic Compute Cloud, known as EC2 [1], Google Compute Engine [6], Microsoft Azure [8]. These rovide a large pool of computing resources that can be rented by the public for short-term computing tasks. These services were originally used primarily by web application developers, but have increasingly been used in scientific and technical high performance computing (HPC) applications [47] [57] [71] [13].

Cloud computing is challenging when there are real-time constraints [45]; this is an active area of research. However there are many robotics applications that are not time sensitive such as decluttering a room or pre-computing grasp strategies.

There are many sources of uncertainty in robotics and automation [34]. Cloud computing allows massive sampling over error distributions and Monte Carlo sampling is “embarrassingly parallel”; recent research in fields as varied as medicine [75] and particle physics [67] have taken advantage of the cloud. Real-time video and image analysis can be performed in the Cloud [55] [60] [62]. Image processing in the cloud has been used for assistive technology for the visually impaired [22] and for senior citizens [32]. Cloud computing is ideal for sample-based statistical motion planning under uncertainty, where it can be used to explore many possible perturbations in object and environment pose, shape, and robot response to sensors and commands [72]. Cloud-based sampling is also being investigated for grasping objects with shape uncertainty [51] [52] (see Figure 3). A grasp planning algorithm accepts as input a nominal polygonal outline with Gaussian uncertainty around each vertex and the center of mass to compute a grasp quality metric based on a lower bound on the probability of achieving force closure.

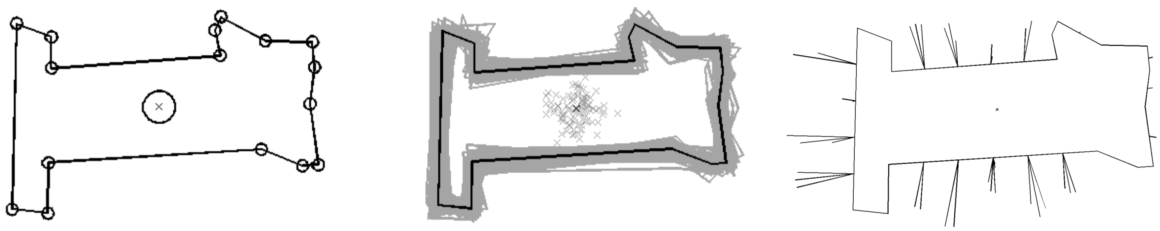

Figure 3: A cloud-based approach to geometric shape uncertainty for grasping [51] [52]. (Image reproduced with permission from authors).

1.3 Collective Robot Learning

The Cloud allows robots and automation systems to "share" data from physical trials in a variety of environments, for example initial and desired conditions, associated control policies and trajectories, and importantly: data on performance and outcomes. Such data is a rich source for robot learning.

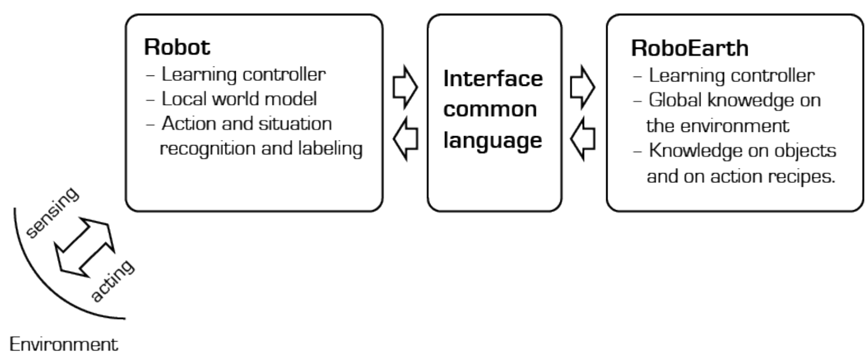

Figure 4: RoboEarth architecture [74]. (Image reproduced with permission from authors).

One example is for path planning, where previously-generated paths are adapted to similar environments [21] and grasp stability of finger contacts can be learned from previous grasps on an object [27].

The MyRobots project [9] 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" [14].

1.4 Open-Source and Open-Access

The Cloud facilitates sharing by humans of designs for hardware, data, and code. The success of open-source software [25] [40] [61] 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 [11] [65]. ROS has also been ported to Android devices [12]. ROS has become a standard akin to Linux and is now used by almost all robot developers in research and many in industry.

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. Open-source simulation libraries include Bullet [4], a physics simulator originally used for video games, OpenRAVE [10] and Gazebo [5], simulation environments geared specifically towards robotics, OOPSMP, a motion-planning library [63], and GraspIt!, a grasping simulator [58].

Another exciting trend is in open-source hardware, where CAD models and the technical details of construction of devices are made freely available [29] [66]. The Arduino project [3] is a widely-used open-source microcontroller platform, and has been used in many robotics projects. The Raven [53] is an open-source laparoscopic surgery robot developed as a research platform an order of magnitude less expensive than commercial surgical robots [16].

The Cloud can also be used to facilitate open challenges and design competitions. For example, the African Robotics Network with support from IEEE Robotics and Automation Society hosted the "$10 Robot" Design Challenge in the summer of 2012. This open competition attracted 28 designs from around the world including a winning entry from Thailand that modified a surplus Sony game controller, adapting its embedded vibration motors to drive wheels and adding lollipops to the thumb switches as inertial counterweights for contact sensing, which can be built from surplus parts for US $8.96 [17].

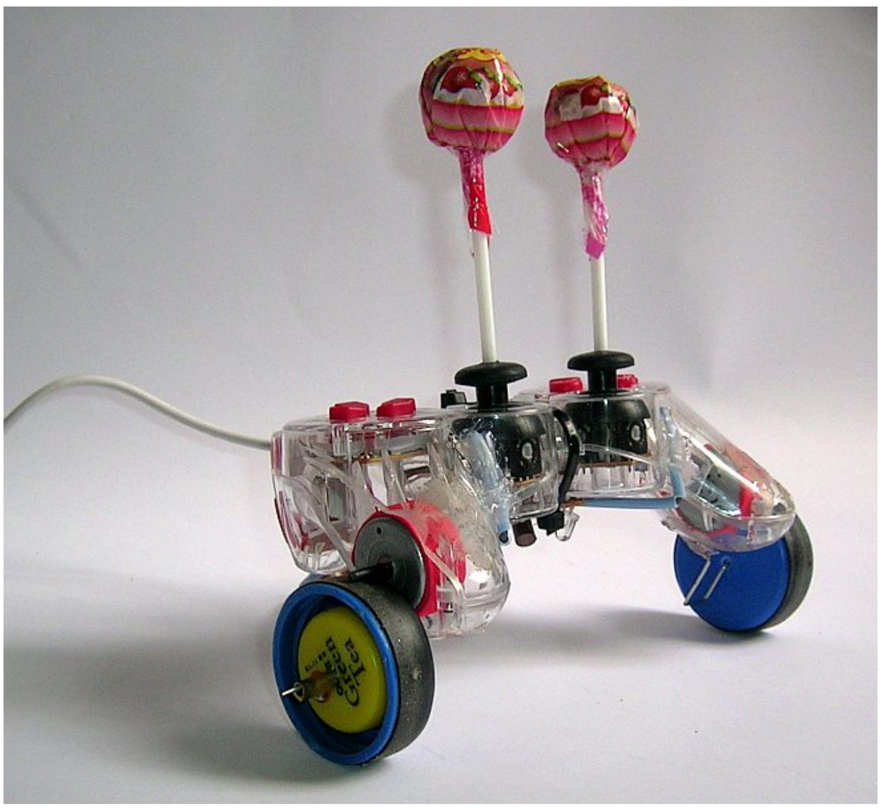

Figure 5: Suckerbot, designed by Tom Tilley of Thailand, a winner of the $10 Robot Design Challenge [17]. (Image reproduced with permission from authors).

1.5 Crowdsourcing and Call Centers

In contrast to automated telephone reservation and technical support systems, consider a future scenario where errors and exceptions are detected by robots and automation systems, which then access human guidance on-demand at remote call centers. Human skill, experience, and intution is being tapped to solve a number of problems such as image labeling for computer vision [73] [24][48] [54]. Amazon's Mechanical Turk is pioneering on-demand "crowdsourcing" that can draw on "human computation" or "social computing systems". Research projects are exploring how this can be used for path planning [41], to determine depth layers, image normals, and symmetry from images [33], and to refine image segmentation [46]. Researchers are working to understand pricing models [69] and apply crowdsourcing to grasping [68] (see Figure 1).

1.6 Acknowledgements

The authors thank James Kuffner, Steve Cousins, Raffaello D'Andrea, Alex Waibel, Gary Bradski, Vijay Kumar, Erico Guizzo, Dmitry Berenson, and Ben Kehoe for ongoing insights and advice on this topic.

1.7 Contact

Ken Goldberg, craigslist Distinguished Professor of New Media
Professor, IEOR and EECS, College of Engineering, UC Berkeley
Dept of Art Practice and School of Information, UC Berkeley
Professor, Department of Radiation Oncology, UC San Francisco
Contact: 425 Sutardja Dai Hall, Berkeley, CA 94720-1758
twitter: @Ken_Goldberg — g+: http://j.mp/Ken_Goldberg
(510) 643-9565 — goldberg@berkeley.edu — http://goldberg.berkeley.edu

2 References

References

[1] Amazon Elastic Cloud (EC2). http://aws.amazon.com/ec2/.

[2] Amazon Web Services. http://aws.amazon.com.

[3] Arduino. http://www.arduino.cc.

[4] Bullet Physics Library. http://bulletphysics.org.

[5] Gazebo. http://gazebosim.org.

[6] Google Compute Engine. https://cloud.google.com/products/compute-engine.

[7] Google Goggles. http://www.google.com/mobile/goggles/.

[8] Microsoft Azure. http://www.windowsazure.com.

[9] MyRobots.com. http://myrobots.com.

[10] OpenRAVE. http://openrave.org/.

[11] ROS (Robot Operating System). http://ros.org.

[12] rosjava, an implementation of ROS in pure Java with Android support. http://cloudrobotics.com.

[13] TOP500. http://www.top500.org/list/2012/06/100.

[14] What is MyRobots? http://myrobots.com/wiki/About.

[15] What is RoboEarth? http://www.roboearth.org/what-is-roboearth.

[16] An open-source robo-surgeon. *The Economist*, 2012.

[17] The African Robotics Network (AFRON). "Ten Dollar Robot" Design Challenge Winners. http://robotics-africa.org/design_challenge.html.

[18] Michael Armbrust, Ion Stoica, Matei Zaharia, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, and Ariel Rabkin. A View of Cloud Computing. *Communications of the ACM*, 53(4):50, April 2010.

[19] Rajesh Arumugam, V.R. Enti, Liu Bingbing, Wu Xiaojun, Krishnamoorthy Baskaran, F.F. Kong, A.S. Kumar, K.D. Meng, and G.W. Kit. DAvinCi: A Cloud Computing Framework for Service Robots. In *IEEE International Conference on Robotics and Automation*, pages 3084–3089. IEEE, 2010.

[20] Luigi Atzori, Antonio Iera, and Giacomo Morabito. The Internet of Things: A Survey. *Computer Networks*, 54(15):2787–2805, October 2010.

[21] Dmitry Berenson, Pieter Abbeel, and Ken Goldberg. A Robot Path Planning Framework that Learns from Experience. *IEEE International Conference on Robotics and Automation*, pages 3671–3678, May 2012.

[22] Bharat Bhargava, Pelin Angin, and Lian Duan. A Mobile-Cloud Pedestrian Crossing Guide for the Blind. In *International Conference on Advances in Computing & Communication*, 2011.

[23] Matei Ciocarlie, Kaijen Hsiao, E. G. Jones, Sachin Chitta, R.B. Rusu, and I.A. Sucan. Towards Reliable Grasping and Manipulation in Household Environments. In *Intl. Symposium on Experimental Robotics*, pages 1–12, New Delhi, India, 2010.

[24] Matei Ciocarlie, Caroline Pantofaru, Kaijen Hsiao, Gary Bradski, Peter Brook, and Ethan Dreyfuss. A Side of Data With My Robot. *IEEE Robotics & Automation Magazine*, 18(2):44–57, June 2011.

[25] Laura Dabbish, Colleen Stuart, Jason Tsay, and Jim Herbsleb. Social coding in GitHub: transparency and collaboration in an open software repository. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, page 1277, New York, New York, USA, 2012. ACM Press.

[26] Sebastien Dalibard, Alireza Nakhaei, Florent Lamiraux, and Jean-Paul Laumond. Manipulation of documented objects by a walking humanoid robot. In *IEEE-RAS International Conference on Humanoid Robots*, pages 518–523. Ieee, December 2010.

[27] Hao Dang and Peter K. Allen. Learning grasp stability. In *IEEE International Conference on Robotics and Automation*, pages 2392–2397. IEEE, May 2012.

[28] Hao Dang, Jonathan Weisz, and Peter K. Allen. Blind grasping: Stable robotic grasping using tactile feedback and hand kinematics. In *IEEE International Conference on Robotics and Automation*, pages 5917–5922. Ieee, May 2011.

[29] S. Davidson. Open-source hardware. *IEEE Design and Test of Computers*, 21(5):456–456, September 2004.

[30] Zhihui Du, Weiqiang Yang, Yinong Chen, Xin Sun, Xiaoying Wang, and Chen Xu. Design of a Robot Cloud Center. In *International Symposium on Autonomous Decentralized Systems*, pages 269–275. IEEE, March 2011.

[31] Stephan Gammeter, Alexander Gassmann, Lukas Bossard, Till Quack, and Luc Van Gool. Server-side Object Recognition and Client-side Object Tracking for Mobile Augmented Reality. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, number C, pages 1–8. Ieee, June 2010.

[32] JJS García. Using Cloud Computing as a HPC Platform for Embedded Systems. 2011.

[33] Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. Micro perceptual human computation for visual tasks. *ACM Transactions on Graphics*, 31(5):1–12, August 2012.

[34] Jared Glover, Daniela Rus, and Nicholas Roy. Probabilistic Models of Object Geometry for Grasp Planning. In *Robotics: Science and Systems*, Zurich, Switzerland, 2008.

[35] Ken Goldberg, Michael Mascha, Steve Gentner, Nick Rothenberg, Carl Sutter, and Jeff Wiegley. Desktop teleoperation via the World Wide Web. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 654–659. IEEE, 1995.

[36] Ken Goldberg and Roland Siegwart, editors. *Beyond Webcams: An Introduction to Online Robots*. MIT Press, 2002.

[37] C. Goldfeder, M. Ciocarlie, and P.K. Allen. The Columbia Grasp Database. In *IEEE International Conference on Robotics and Automation*, pages 1710–1716. IEEE, May 2009.

[38] Corey Goldfeder and Peter K. Allen. Data-Driven Grasping. *Autonomous Robots*, 31(1):1–20, April 2011.

[39] Eric Guizzo. Cloud Robotics: Connected to the Cloud, Robots Get Smarter, 2011.

[40] Alexander Hars. Working for free? Motivations of participating in open source projects. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*, volume 00, page 9. IEEE Comput. Soc, 2001.

[41] Juan Camilo Gamboa Higuera, Anqi Xu, Florian Shkurti, and Gregory Dudek. Socially-Driven Collective Path Planning for Robot Missions. *2012 Ninth Conference on Computer and Robot Vision*, pages 417–424, May 2012.

[42] Guoqiang Hu, WP Tay, and Yonggang Wen. Cloud Robotics: Architecture, Challenges and Applications. *IEEE Network*, 26(3):21–28, 2012.

[43] Kai Huebner, Kai Welke, Markus Przybylski, Nikolaus Vahrenkamp, Tamim Asfour, and Danica Kragic. Grasping Known Objects with Humanoid Robots: A Box-Based Approach. In *International Conference on Advanced Robotics*, 2009.

[44] Dominique Hunziker, Mohanarajah Gajamohan, Markus Waibel, and Raffaello D Andrea. Rapyuta: The RoboEarth Cloud Engine. 2013.

[45] Nitesh Kumar Jangid. Real Time Cloud Computing. In *Data Management & Security*, 2011.

[46] Matthew Johnson-Roberson, Jeannette Bohg, Gabriel Skantze, Joakim Gustafson, Rolf Carlson, Babak Rasolzadeh, and Danica Kragic. Enhanced visual scene understanding through human-robot dialog. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3342–3348. IEEE, September 2011.

[47] Gideon Juve, Ewa Deelman, G. Bruce Berriman, Benjamin P. Berman, and Philip Maechling. An Evaluation of the Cost and Performance of Scientific Workflows on Amazon EC2. *Journal of Grid Computing*, 10(1):5–21, March 2012.

[48] Koji Kamei, Shuichi Nishio, Norihiro Hagita, and Miki Sato. Cloud Networked Robotics. *IEEE Network*, 26(3):28–34, 2012.

[49] 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. *The International Journal of Robotics Research*, 31(8):927–934, May 2012.

[50] B. Kehoe, A. Matsukawa, S. Candido, J. Kuffner, and K. Goldberg. Cloud-Based Robot Grasping with the Google Object Recognition Engine. 2013. Submitted to IEEE International Conference on Robotics and Automation, 2013.

[51] Ben 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*. IEEE, 2012.

[52] Ben Kehoe, Dmitry Berenson, and Ken Goldberg. Toward Cloud-based Grasping with Uncertainty in Shape: Estimating Lower Bounds on Achieving Force Closure with Zero-slip Push Grasps. In *IEEE International Conference on Robotics and Automation*, pages 576–583. IEEE, May 2012.

[53] H Hawkeye King, Lei Cheng, Philip Roan, Diana Friedman, Sina Nia, Ji Ma, Daniel Glozman, Jacob Rosen, and Blake Hannaford. Raven II: Open Platform for Surgical Robotics Research. In *The Hamlyn Symposium on Medical Robotics*, 2012.

[54] James J. Kuffner. Cloud-Enabled Robots. In *IEEE-RAS International Conference on Humanoid Robots*, Nashville, TN, 2010.

[55] K. Lai and D. Fox. Object Recognition in 3D Point Clouds Using Web Data and Domain Adaptation. *The International Journal of Robotics Research*, 29(8):1019–1037, May 2010.

[56] G. McKee. What is Networked Robotics? *Informatics in Control Automation and Robotics*, 15:35–45, 2008.

[57] Piyush Mehrotra, Jahed Djomehri, Steve Heistand, Robert Hood, Haoqiang Jin, Arthur Lazanoff, Subhash Saini, and Rupak Biswas. Performance evaluation of Amazon EC2 for NASA HPC applications. In *Proceedings of the 3rd workshop on Scientific Cloud Computing Date - ScienceCloud '12*, page 41, New York, New York, USA, 2012. ACM Press.

[58] A.T. Miller and P.K. Allen. GraspIt! A Versatile Simulator for Robotic Grasping. *IEEE Robotics & Automation Magazine*, 11(4):110–122, December 2004.

[59] M.A. Moussa and M.S. Kamel. An experimental approach to robotic grasping using a connectionist architecture and generic grasping functions. *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, 28(2):239–253, May 1998.

[60] D Nister and H Stewenius. Scalable Recognition with a Vocabulary Tree. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 2161–2168. IEEE, 2006.

[61] Daniel Nurmi, Rich Wolski, Chris Grzegorczyk, Graziano Obertelli, Sunil Soman, Lamia Youseff, and Dmitrii Zagorodnov. The Eucalyptus Open-Source Cloud-Computing System. In *IEEE/ACM International Symposium on Cluster Computing and the Grid*, pages 124–131. IEEE, 2009.

[62] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object Retrieval with Large Vocabularies and Fast Spatial Matching. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. Ieee, June 2007.

[63] Erion Plaku, Kostas E. Bekris, and Lydia E Kavraki. OOPS for Motion Planning: An Online, Open-source, Programming System. In *IEEE International Conference on Robotics and Automation*, number April, pages 3711–3716. IEEE, April 2007.

[64] Mila Popovic, Gert Kootstra, Jimmy Alison Jorgensen, Danica Kragic, and Norbert Kruger. Grasping unknown objects using an Early Cognitive Vision system for general scene understanding. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 987–994. IEEE, September 2011.

[65] Morgan Quigley and Brian Gerkey. ROS: an open-source Robot Operating System. In *ICRA Workshop on Open Source Software*, number Figure 1, 2009.

[66] Erik Rubow. Open Source Hardware. Technical report, 2008.

[67] Martin Sevior, Tom Fifield, and Nobuhiko Katayama. Belle monte-carlo production on the Amazon EC2 cloud. *Journal of Physics: Conference Series*, 219(1):012003, April 2010.

[68] A Sorokin, D Berenson, S S Srinivasa, and M Hebert. People helping robots helping people: Crowd-sourcing for grasping novel objects. *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2117–2122, October 2010.

[69] Alexander Sorokin and David Forsyth. Utility Data Annotation with Amazon Mechanical Turk. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, number c, pages 1–8. IEEE, June 2008.

[70] Moritz Tenorth, Alexander Clifford Perzylo, Reinhard Lafrenz, and Michael Beetz. The RoboEarth language: Representing and exchanging knowledge about actions, objects, and environments. In *IEEE International Conference on Robotics and Automation*, number 3, pages 1284–1289. Ieee, May 2012.

[71] Radu Tudoran, Alexandru Costan, Gabriel Antoniu, and Luc Bougé. A performance evaluation of Azure and Nimbus clouds for scientific applications. In *Proceedings of the 2nd International Workshop on Cloud Computing Platforms - CloudCP '12*, pages 1–6, New York, New York, USA, 2012. ACM Press.

[72] Jur van den Berg, Pieter Abbeel, and Ken Goldberg. LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. *The International Journal of Robotics Research*, 30(7):895–913, June 2011.

[73] Luis von Ahn. Human Computation. In *Design Automation Conference*, page 418, 2009.

[74] Markus Waibel, Michael Beetz, Javier Civera, Raffaello D'Andrea, Jos Elfring, Dorian Gálvez-López, Kai Häussermann, Rob Janssen, J.M.M. Montiel, Alexander Perzylo, Björn Schieß le, Moritz Tenorth, Oliver Zweigle, and René De Molengraft. RoboEarth. *IEEE Robotics & Automation Magazine*, 18(2):69–82, June 2011.

[75] Henry Wang, Yunzhi Ma, Guillem Pratx, and Lei Xing. Toward real-time Monte Carlo simulation using a commercial cloud computing infrastructure. *Physics in medicine and biology*, 56(17):N175–81, September 2011.

[76] Jonathan Weisz and Peter K. Allen. Pose error robust grasping from contact wrench space metrics. In *IEEE International Conference on Robotics and Automation*, pages 557–562. IEEE, May 2012.