

A Vision for Spatial-Reasoning Commodity Robots

Zeke Koziol, Sabreen Lakhani, Anatole Paine, and Zachary Dodds, Harvey Mudd College, Claremont, CA
{zkoiol, slakhani, apaine, dodds}@cs.hmc.edu

Abstract—This work seeks to augment commodity robots to enable service tasks that require global spatial reasoning, e.g., delivery, surveillance, and map maintenance. In addition, we seek to do this as accessibly and inexpensively as possible. To that end we present a design and three prototype platforms that illustrate educational, service-level, and research applications atop the iRobot Create. The prototypes significantly improve the robotic capabilities broadly available at the \$500 price range, particularly for pedagogical applications. This work informs a facet of the design space for next-generation commodity robots. Though the prototypes' first audience is likely to be educators, we hope they pique the interest of those who design low-cost consumer and service robots, as well.

1. MOTIVATION

The millions of service and entertainment robots in households demonstrate the remarkable success and solid foothold of commodity robotics today. The next generation of ~\$250 commodity platforms, we believe, will add the capability of *spatial deliberation* to the robustness that characterizes currently available systems. Maintaining environmental maps will open up delivery, surveillance, and other service capabilities to a broad audience.

To investigate the feasibility of low-cost robots that can reason about more than their local environment, we present in this work prototype extensions of the iRobot Create [1]. The key design choice is to manage all sensor processing off-board via wireless router. An ethernet camera and short-range IR sensors provide excellent data, complementing the Create's, for map maintenance. Currently the platform uses provided maps in order to complete point-to-point navigation tasks, e.g., for pickup and delivery. In addition, it serves as an accessible basis for research and experiments in vision-based navigation.

2. CAPABILITIES AND COMPONENTS

On top of the proprioceptive and actuation capabilities provided by the Create [1] a suite of six Sharp GP2D12 IR sensors and an Axis207 ethernet camera send a stream of data via an off-the-shelf wireless router. Any router will do;

we have used Linksys WRT320N and 546L models. Router-to-robot control occurs through an Arduino microcontroller and its accompanying ethernet interface. Voltage regulators and their heat sinks are the only additional wiring required. In total, the cost beyond the Create is less than \$400: \$200 for the camera and \$200 for the other components. These reflect prototype costs; commercial production would yield substantial savings. Figure 1 shows our integrated platforms and their components.

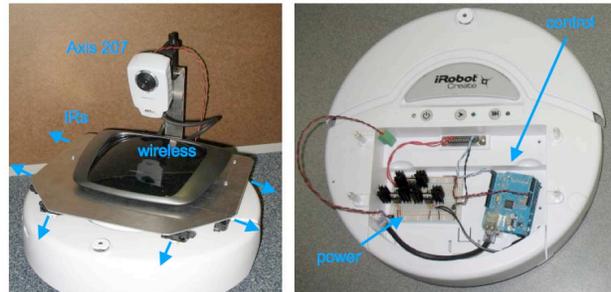


Figure 1 - (left) At left is the assembled platform and its external sensors, six infrared range sensors and an Ethernet camera. Sensor processing and control is off-board; lag is not an issue; the limiting factors are the Create's serial bandwidth and the image-processing time, not data- or image-transfer costs. **(right)** With the shelf removed, the simplicity and robustness of the design is apparent: a power-distribution circuit of three voltage regulators and a control interface using an Arduino support the sensors and base.

A second design possibility we have pursued replaces the router and camera with a netbook computer, following the lead of CMU's Tekkotsu-on-the-Create [2] and designs for telepresence by researchers at Southern Illinois University, Edwardsville. The OS-on-board offers an easy interface to additional devices, but camera quality and software development suffer atop a netbook: those conditions will likely improve in the future.

3. INTERFACE SOFTWARE

To highlight the capabilities of these platforms, we have created several interfaces:

- (1) a remote-control interface using a game controller
- (2) autonomous exploration routines for indoor use
- (3) spatial-reasoning software supporting navigation
- (4) an OpenCV wrapper for vision-based capabilities

At the low level, the robot presents itself as a network device to which ASCII strings are passed back and forth. Although any language can be used, to date we have built a

Python interface that exposes all of the motor and sensory capabilities except vision. That web-accessible image stream, 640x480 at 20fps, is available to whatever tools a developer may wish to use. We use an OpenCV-based interface written in C++.

Inspired by work at Brown University [4], our first software interface leverages pygame (www.pygame.org) to provide direct control of the robot through a Wiimote game controller. The result is engaging and challenging interactions appropriate for elementary-school audiences – and enjoyable for any age!



Figure 2. The resulting platform is seven-year-old-proof. Inexpensive enough to use in outreach activities, the software and Wiimote-control allows elementary school students to navigate the robots using only sensor-provided data. Racing is popular, too (left)!

As highlighted in Figure 2, we have developed pedagogical activities that allow young students to experience hands-on how it is different to "be" a robot. Navigating to a destination with only the video feed or, more difficult, following the walls with only the IRs deepens appreciation of the challenges inherent in bridging sensing and control. The addition of a Nerf launcher has been a wonderful ice-breaker, and is equally effective among researchers attending IJCAI and elementary school children.



Figure 3. Having an OS onboard makes the addition of peripherals straightforward. Indeed, the actuation itself *is* a peripheral. The robots, at home in hallways (left) were also exhibited at IJCAI '09 (center); in addition, a Nerf missile launcher can act as a universal ice-breaker (right).

Fast, inexpensive autonomy via IR sensing

The second interface - autonomous exploration - mimics the Create's ability to quickly cover a large environment by wall-following. Using the IRs increases efficiency, however, leading to an average indoor wall-following speed of 39.1

centimeters per second, including concave and convex corners. The Create's built-in wall-following behavior runs at 24.3 cm/sec. As Figures 3 and 4 show, the sensor suite is well-suited to navigating within indoor environments.

Atop this autonomous exploration, we have added a spatial-reasoning module. Similar in spirit to the pioneering platforms of [5] and [6], but at a cost orders of magnitude lower, the software allows a user to click on a goal location within the robot's known map. The system then autonomously plans a path to that goal and executes it. Figure 4 contrasts the odometric estimate (in purple) and map-corrected locations (in red along straightaways, green when turning, and blue upon task completion) of the robot throughout a short run. This "delivery" capability will be tested in November 2009's robotics innovation competition and conference (RICC).

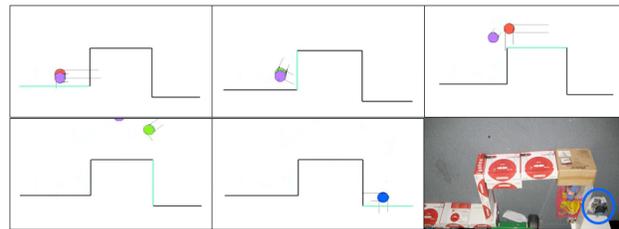


Figure 4. Snapshots from the robot's odometric position (purple) and actual position through a multi-segment wall-following task. The walls themselves enable correction of the robot's location, culminating in a correct estimate of its pose even though the odometric estimate has long left the field of view. Color represents task state: red for wall-following, green for turning corners, and blue at the desired destination. IR values are also visible. The wall that the platform believes it is currently following is highlighted in cyan. The final snapshot shows the actual location of the robot at the end of this run, very near its estimated pose.

An accessible platform for real-time vision experiments

The fourth interface makes the platform an accessible basis for open-ended research. An OpenCV interface to the video stream supports interaction with arbitrary libraries and external code. We use a combination of Matlab, C++ neural network libraries, and Python to investigate algorithms for learning range from monocular vision, following work such as [7] and [8]. We describe the approach and the results of this research in the following section.

Obtaining the software

All of this software is available from our public subversion repository [3]. Please contact the last author with concerns or follow-up inquiries. Note that other freely available libraries, such as [9] would allow a natural interface through popular packages such as Matlab.

4. RESEARCH APPLICATIONS

Perhaps the greatest potential for this platform design, as it exists today rather than as a prototype for future applications, is as a testbed for research and education. Accessible both in terms of cost and the shallow learning curve of its software, the Create does not impose any unwanted layers between researchers and the data collected onboard.

Background: Range from Texture

Thus, images such as those to the left of Figure 5 stream at 20-30 frames per second back to a host computer, from which they are accessible, for example, via OpenCV. Seeing these, we realized that monocular vision offers at least the *promise* of an advantageous alternative to laser range-finding across several axes: cameras are less power-hungry, less heavy, less bulky, less range-limited, and, perhaps most importantly, less expensive.

Less is more, however, when it comes to computation. Extracting range from pixel intensities requires far more algorithmic and computational lifting than extracting range from time-of-flight. Usual range-from-vision approaches use temporal feature correspondence across a monocular image stream to deduce distance from pixels [12]. This corpus of work is mature, but it is worth noting that these techniques are most successful when significant spatial context is used across the image stream. Large patches of pixels facilitate accurate and precise correspondence.

Recent approaches have boldly asked, "What can we deduce from *only* those patches, and not the correspondence at all!?" For instance, Hoiem et al.'s Photo pop-out [10] software and Saxena et al.'s Make3d system [8] yield range at each of a single image's pixels. Spatial grouping, e.g., into the edges between ground and vertical planes enable the compelling visualizations those groups have produced.

Such work has seen robot applications: in [10] Hoiem et al. show confidence levels in terrain navigability, in [8] Saxena et al. drive an RC car safely and quickly through rough, natural terrain. Those projects emphasized machine-learning contributions over the resulting range accuracy. More recently, Plagemann, et al., [7] quantified the range accuracy of the learned mapping from *columns* of pixels from omnidirectional images to range. That work used Gaussian Processes to achieve ~1m precision sufficient to support off-the-shelf SLAM algorithms under assumptions common to indoor environments. Such results underscore the power of range-from-texture approaches, pioneered in Horswill's Polly [11], and used in many systems since.

Here, we learn indoor range-to-obstacle directly from visual texture. We contrasted two machine learning approaches, multilayer perceptrons (MLP) using backpropagation and restricted Boltzmann machines (RBM) [14].

Figure 5 illustrates how we collected data in order to train these two learning algorithms.

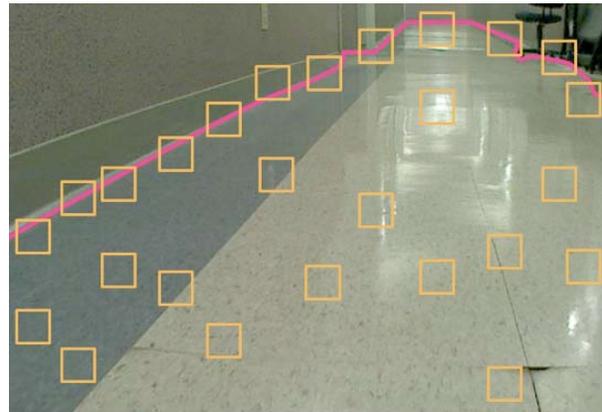


Figure 5. In contrast to traditional approaches' feature correspondence and camera calibration, our range-from-texture approach learns from labeled images to distinguish the source of 8x8 pixel patches: floor vs. wall. The pink line is human-labeled ground truth and the yellow squares indicate patches used to train. Although the learning itself is off-line, the resulting classifier finds groundplane within the live image stream, as documented in Figure 6.

Architecture

The MLP and RBM both have 64 visible inputs: the raw pixel values from an 8x8 patch of one of the images. The MLP had a single output that was reinforced toward 1.0 for obstacle patches and 0.0 for floor patches. The RBM had two distinct labels: one for floor and one for obstacles. RBMs do not output a traditional binary classification signal; rather, they reconstruct the input using both the learned models, floor and obstacle. The fidelity of each reconstruction allows our system to decide which to classify each patch. In our implementation, we weighted each component equally with opposite signs. Both networks were 5 layers with 3 20-neuron hidden layers.

Results

The fundamental differences between multilayer perceptrons and restricted Boltzmann machines are apparent in Figure 6, below. The MLP carefully hews to the environment in which the classifier was trained: as a result it is more accurate in segmenting ground plane from obstacles in those environments. For example, the MLP segments the alternating bicycle wheels and wall segments as well or better than a human observer can.

Because the RBM seeks parsimony in its constructed representation of the floor vs. wall patches, it does not classify known environments as well as the MLP. However, it generalizes far better when patches very different from training data appear. In Figure 6, an observer looks into the camera of the robot: such texture had never appeared in the training set, and the RBM segments it better than the MLP.



Figure 6. The results of our ground-plane segmentation based on both multilayer perceptron and restricted Boltzmann machine networks. The images to the left in each case have been segmented by an MLP. Those on the right are segmented by an RBM. The greater precision of the MLP in familiar environments is evident in the top two images, whereas the greater ability of RBMs to generalize is apparent in the bottom two images: images containing a person did *not* appear in the training data.

5. PERSPECTIVE

We believe that commodity robots are the foundation for a coming generation of capable autonomous agents. Certainly start-ups such as Heartland, Willow Garage, and The Droid Works - among many others - will "change the game" with novel platforms. Yet this work shows that even modest resources can augment available hardware to create accessible and capable systems capable of autonomous spatial reasoning – and the tasks that build upon it.

ACKNOWLEDGMENTS

The authors recognize and thank the generous support of funding from NSF REU #0451293, NSF CCLI #0536173, The Rose Hills Foundation, and Harvey Mudd College.

REFERENCES

[1] The iRobot Create Open Interface, from www.irobot.com/filelibrary/pdfs/hrd/create/create%20Open%20Interface_v2.pdf.

[2] Nickens, Glenn V., Tira-Thompson, Ethan J., Humphries, Thorna, and Touretzky, David (2009) An inexpensive hand-eye system for undergraduate robotics instruction *SIGCSE Bulletin* 41(1) ACM Press, pp. 423-427. Online guide available at chiara-robot.org/Create/.

[3] Software repository with public access to the interface: <https://svn.cs.hmc.edu/svn/robotics/Summer09/roombaLib>

[4] M. Lapping-Carr, O. Jenkins, D. Grollman, J. Schwertfeger, and T. Hinkle. (2008) Wiimote interfaces for lifelong robot learning. In *AAAI Symposium on Using AI to Motivate Greater Participation in Computer Science*, Palo Alto, CA, USA, Mar 2008, AAAI Press, pp. 61-66.

[5] Simmons, Reid G., Goodwin, Richard, Haigh, Karen Zita, Koenig, Sven, O'Sullivan, Joseph, and Veloso, Manuela M. (1997) Xavier: experience with a layered robot architecture *SIGART Bulletin* 8(1-4) ACM Press, pp. 22-33.

[6] Thrun, Sebastian, Beetz, Michael, Bennewitz, Maren, Burgard, Wolfram, Creemers, A.B., Dellaert, Frank, Fox, Dieter, Hahnel, Dirk, Roy, Nicholas, Schulte, Jamieson, and Schulz, Dirk (2000) Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva *International Journal of Robotics Research*, 19(11) Sage Science Press, pp. 972-999.

[7] Plagemann, C., Enres, F., Hess, J., Stachniss, C., and Burgard, W. (2008) Monocular Range Sensing: A non-parametric learning approach. *Proceedings, ICRA 2008* May 19-23, Pasadena, CA IEEE Press, pp. 929-934.

[8] Saxena, Ashutosh, Chung, Sung H., and Ng, Andrew (2008) 3-D Depth Reconstruction from a Single Still Image *International Journal of Computer Vision* 76(1) Kluwer Academic Publishers, pp. 53-69.

[9] Esposito, Joel M. and Barton, O. Matlab toolbox for the iRobot Create, at URL www.usna.edu/Users/weapsys/esposito/roomba.matlab/ 2008.

[10] Hoiem, D., Efros, A. A., and Hebert, M. Recovering Surface Layout from an Image. *IJCV*, Vol. 75(1) Oct. 2007.

[11] Horswill, Ian. Analysis of Adaptation and Environment. *Artificial Intelligence*, Vol. 73, pp. 1-30. 1995

[12] Kanade, T., Kanade, B. Y., and Morris, Daniel D. Factorization methods for structure from motion. *Phil. Trans. of the Royal Society of London, Series A.* 356, pp.1153-1173, 2001.

[13] Rosenblatt, F. The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain, *Cornell Aeronautical Laboratory, Psychological Review*, v65, No. 6, pp. 386-408, 1958.

[14] Hinton, G. E. and Brown, A. Spiking Boltzmann Machines. *Advances in Neural Information Processing Systems* 12, MIT Press, Cambridge, MA, 2000.