RFID-Guided Robots for Pervasive Automation

Travis Deyle, Hai Nguyen, Matthew S. Reynolds, and Charles C. Kemp

IEEE computer society

RFID-Guided Robots for Pervasive Automation

Using tags' unique IDs, a semantic database, and RF perception via actuated antennas, a robot can reliably interact with people and manipulate labeled objects.

Travis Deyle and Hai Nguyen
Georgia Institute of Technology

Matthew S. Reynolds
Duke University

Charles C. Kemp
Georgia Institute of Technology

During a trip to a department store, you purchase a robot and a roll of standard labels: "dish," "dish washer," "clothing," "washing machine," "toy," and "storage bin." You return home, apply the labels as directed, unbox the robot, and turn it on. Instantly, the robot can operate in the labeled world—loading the dishwasher with labeled dishes, putting away labeled toys, and washing labeled clothing. Additional functionality, such as delivering medicine, is just a few labels away.

This is a compelling vision, especially for the motor-impaired patients we work with in collaboration with the Emory ALS (Amyotrophic Lateral Sclerosis) Center. To help realize this vision, we developed *EL-E* (pronounced "Ellie"), a prototype assistive robot that uses novel RFID-based methods to perform tasks with ultrahigh-frequency (UHF) RFID-labeled objects (see Figure 1).

Exploiting UHF RFID Tags

Passive UHF RFID tags are well matched to robots' needs (see the "Related Work in RFID-Guided Robotics" sidebar). Unlike low-frequency (LF) and high-frequency (HF) RFID tags, passive UHF RFID tags are readable from across a room, enabling a mobile robot to ef-ficiently discover and locate them. Because they don't have onboard batteries to wear out, their lifetime is virtually unlimited. And unlike bar codes and other visual tags, RFID tags are readable when they're visually occluded. For less than $0.25 per tag, users can apply self-adhesive UHF RFID tags throughout their home. Because tags are thin and compact, they could also be embedded in objects.

EL-E's two body-mounted long-range antennas and four finger-mounted short-range ceramic microstrip antennas exploit two valuable UHF RFID tag properties:

- *Unique identifiers*. For most perceptual modalities, the object's identity is the culmination of extensive low-level sensory processing and comes with high uncertainty. RFID, on the other hand, provides an object a unique ID with extremely small uncertainty, whose location is inferred by lower-level sensory processing. Generation 2 UHF RFID tags use a challenge-response protocol to provide a globally unique ID with a virtually zero false-positive rate. Specifically, they provide a unique 96-bit ID at ranges of more than 5 meters indoors. A continuously operating reader at Duke University has received zero false-positive IDs in more than 60 days, for a false-positive rate of less than 1×10^{-6}. EL-E can read tags rapidly (500 per second) and in large groups (more than 200 at a time)[1] and, given a tag's ID, can access

Downward-facing laser range finder

In-hand camera

Short-range RFID antennas

Katana arm

Long-range RFID antennas

High-resolution camera

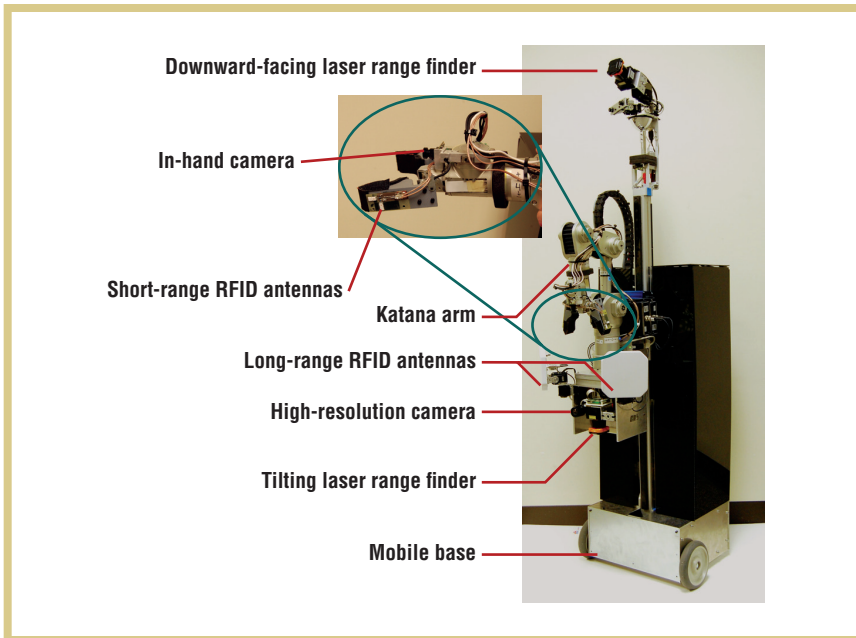Tilting laser range finder

Mobile base

Figure 1. EL-E, an autonomous mobile robot. EL-E uses its two body-mounted long-range ultrahigh-frequency (UHF) RFID antennas and four finger-mounted short-range ceramic microstrip UHF RFID antennas to detect tags.

ID: TV remote

Figure 2. EL-E's user interface (UI). On the basis of RFID-tagged objects it senses nearby, EL-E generates a context-aware remote UI on-the-fly. From this UI, users select an object and action for the robot to perform.

arbitrary associated data. For example, EL-E can use an ID to access a database with information about the tagged object such as its name and a photo, actions that EL-E can perform with it, and the actions' icons.

- *RF perception*. By perceiving the presence and strength of RF signals emanating from a tag, EL-E can estimate its location. If EL-E detects a tag with its long-range antennas, the tag is likely in the room. If EL-E detects the tag with its short-range antennas, the tag is likely close to its hand. EL-E's antennas are also directionally sensitive. If an antenna receives a strong signal from a tag when pointing in one direction but a weak signal from that tag when pointing in another direction, the tag will more likely be in the stronger signal's direction.

Interacting with People

EL-E's first goal is to discover which actions it can perform in a room. EL-E first scans the room for tags by panning its long-range antennas. It then uses each tag's unique ID to query a database and automatically populate a remote user interface (UI; see Figure 2). From this UI, users can select

an object and action for the robot to perform. Future systems might offer more complex interface options that exploit groups of tags. For example, a cooking UI could present dinner options with various types of cuisine, followed by particular dishes that the robot can produce with the available tagged ingredients.

Reliably Reading Tags

For these interfaces to work, EL-E must efficiently and reliably read tags from a variety of positions in the environment. So, we tested EL-E's ability to read 37 tags on a shelf from 36 different locations evenly spaced throughout a 3.7 × 7.3 m room (see Figure 3a). At each location, a complete scan took approximately 13 seconds and consisted of panning each long-range antenna back and forth, one at a time.

On average, EL-E read 23.75 tags at each location with a standard deviation of 4.03 tags. As Figure 3b shows, read reliability decreased with distance. For example, the two closest locations had the most reads (31 tags), whereas the farthest locations had the fewest (15 tags).

As Figure 3c indicates, EL-E read some tags more reliably than others. This variation relates to the tagged object, how we applied the tag, and the tag's pose relative to the antennas. For example, the two unread tags were a tag on a metallic toothpaste tube (the tags we used typically don't work on metal) and a misapplied tag on a medication bottle (the tag was wrapped around itself instead of in a nonoverlapping spiral). EL-E's reads of tags on objects containing electronics, such as the cell phone, cordless phone, and

remote control, were also less reliable. We can mitigate these issues by using multiple tags at different orientations and newly developed tags designed for use on metal.

Delivering and Receiving Objects

As Figure 4 shows, EL-E can deliver an object to someone wearing a tagged wristband and receive an object from that person. When EL-E receives a tagged object, its finger antennas identify it, which could facilitate future context-aware behaviors.

To evaluate EL-E's ability to deliver tagged objects, we performed 10 trials with a TV remote and medication bottle as test objects. Starting approximately 2 meters from the tagged person, EL-E tried to deliver one of the objects. EL-E released the object if its fingers detected forces and torques above a threshold. The person then handed back either the delivered object or another object.

We considered the trial successful if

- the person received the object while remaining seated and
- EL-E correctly identified the object that the person handed it.

EL-E succeeded in all 10 trials.

Approaching a Tagged Object

To deliver and receive objects and perform relevant tasks, EL-E must be able to approach the tagged objects. UHF RFID helps EL-E do this in several ways. With a long-range antenna, RF perception orients EL-E toward a tagged object and estimates its position. In addition, a database indexed by a tag's unique ID could provide information about the object's usual location, use history, and expected appearance.

Once EL-E decides to approach a tag, it specializes its queries to this tag alone. This process, called *singulation* (defined in the Generation 2 specification[1]) lets EL-E quickly perform RFID reads for a specific tag ID, even

Related Work in RFID-Guided Robotics

Researchers frequently discuss robots and ultrahigh-frequency (UHF) RFID tags as components of pervasive infrastructures.[1,2] Yet few researchers have used RFID sensing as an integral part of their robots. Because of factors such as reader cost and availability, research has often focused on the more mature low-frequency (LF) RFID at 125 kHz and high-frequency (HF) RFID at 13.56 MHz. Using these technologies, researchers have created robotic systems for waypoint navigation and object or person detection.[3] These technologies have also proven useful in nonrobotic systems for activity recognition.[4]

More relevant to the research reported in the main article, researchers have demonstrated that distributed HF RFID readers in ubiquitous sensing environments can inform mobile robots that manipulate objects.[5] Because LF and HF RFID have short read ranges (below 20 cm), they require both distributed readers and tags to emulate the capabilities we describe in the main article—an often impractical proposition.

An alternative to short-range LF and HF RFID is long-range UHF RFID. Researchers have demonstrated read ranges exceeding 50 m using active (battery-powered) tags,[6] but fully passive UHF RFID tags operating at 915 MHz cost less and are simpler to implement. To date, most UHF RFID tag research has focused on simultaneous localization and mapping (SLAM) techniques that map static tags' locations, often to subsequently localize the robot.[7–9] In contrast, we've taken a more object-centric approach to UHF RFID, in which long-range antennas help the robot find a tagged object from afar. Once the robot is closer to the tagged object, specialized short-range antennas provide capabilities analogous to LF and HF RFID, using the same UHF tag.

REFERENCES

1. B. Nath, F. Reynolds, and R. Want, "RFID Technology and Applications," *IEEE Pervasive Computing*, vol. 5, no. 1, 2006, pp. 22–24.

2. D. Estrin et al., "Connecting the Physical World with Pervasive Networks," *IEEE Pervasive Computing*, vol. 1, no. 1, 2002, pp. 59–69.

3. M. Shiomi et al., "Interactive Humanoid Robots for a Science Museum," *Proc. 2006 ACM SIGCHI/SIGART Conf. Human-Robot Interaction*, ACM Press, 2006, pp. 305–312.

4. J. Smith et al., "RFID-Based Techniques for Human-Activity Detection," *Comm. ACM*, vol. 48, no. 9, 2005, pp. 39–44.

5. R. Rusu, B. Gerkey, and M. Beetz, "Robots in the Kitchen: Exploiting Ubiquitous Sensing and Actuation," *Robotics and Autonomous Systems*, vol. 56, no. 10, 2008, pp. 844–856.

6. M. Kim, H.W. Kim, and N.Y. Chong, "Automated Robot Docking Using Direction Sensing RFID," *Proc. 2007 IEEE Int'l Conf. Robotics and Automation* (ICRA 07), IEEE Press, 2007, pp. 4588–4593.

7. D. Hahnel et al., "Mapping and Localization with RFID Technology," *Proc. 2004 IEEE Int'l Conf. Robotics and Automation* (ICRA 04), vol. 1, IEEE Press, 2004, pp. 1015–1020.

8. D. Joho, C. Plagemann, and W. Burgard, "Modeling RFID Signal Strength and Tag Detection for Localization and Mapping," *Proc. 2009 IEEE Int'l Conf. Robotics and Automation* (ICRA 09), IEEE Press, 2009, pp. 3160–3165.

9. P. Vorst et al., "Self-Localization with RFID Snapshots in Densely Tagged Environments," *Proc. 2008 IEEE/RSJ Int'l Conf. Intelligent Robots and Systems* (IROS 08), IEEE Press, 2008, pp. 1353–1358.
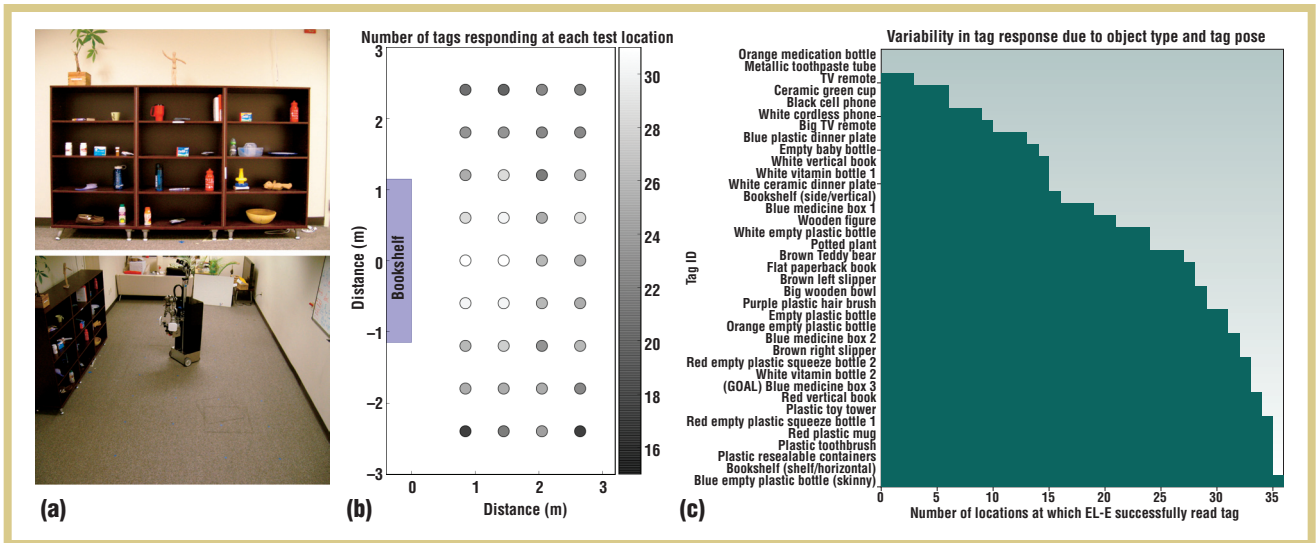
Figure 3. EL-E's read reliability. (a) The experimental setup. The top photo shows 37 uniquely tagged objects on the bookshelf; the bottom photo shows the bookshelf's position in the room. (b) The number of tags read at each test location in the room. (c) The number of times EL-E read each tag, demonstrating the variability in tag response due to object type, tag position and orientation, and material composition.

Figure 4. EL-E delivers a tagged medication bottle to a person wearing a tagged wristband. Tagging people and objects facilitates human-robot interaction and provides identities with high confidence, which is crucial for applications like delivering medicine.

in densely tagged environments. Most sensing techniques we describe here use this process.

To approach a tag, EL-E uses *received signal strength indicator* (RSSI), a scalar quantity describing a tag's response strength as seen by an RFID reader's antenna. First, EL-E pans its two long-range antennas to estimate the tag's bearing, smoothing the resulting RSSI values to filter out noise. It then selects the orientation with the maximum value and rotates toward this bearing.

Second, EL-E keeps its antennas at fixed orientations and "servos" toward the tag's position until its downward-facing laser range finder detects an object in its path. EL-E moves forward at a constant velocity (0.2 m/sec.) and rotates at an angular velocity proportional to the difference in the RSSI values received by its left and right antennas. Because of the long-range antennas' directional sensitivity (approximately 100-degree beamwidths), obtaining a higher RSSI is likely when the antennas are pointing in the tag's direction. So, if the right antenna receives a stronger signal, EL-E rotates right; if the left antenna receives a stronger signal, EL-E rotates left.

Finally, EL-E again estimates the tag's bearing and orients itself accordingly.

This method is efficient and effective. We evaluated EL-E's ability to approach a tagged object on a bookshelf from a grid of 36 distinct starting locations (see Figure 5). At each location, we initially oriented EL-E so that it faced the bookshelf. We considered a trial successful if EL-E stopped less than 1 meter from the tagged object and that object was fully visible from

its camera. This definition of success is well matched to the close-range methods we developed for EL-E.

This approaching behavior is a valuable foundation but has limitations. The current servoing method requires open space because EL-E stops when its range finder detects a potential collision and because large metal objects can significantly influence its path. Also, EL-E doesn't maintain an explicit representation of the tag's location and doesn't exploit other sensory modalities, such as vision. To address these issues, we developed two methods that we plan to incorporate in EL-E—a particle filter with an integrated multipath model and RSSI tag perception.

Estimating Tag Position with a Particle Filter
We've demonstrated probabilistic methods to estimate a tag's position on the basis of many sensor readings over time. We tested a particle filter approach that probabilistically estimates a tag's location in the sensor array's environment using readings from a circular array of antennas.[2] This approach incorporates motion to estimate the tagged object's likeliest location as the array moves through the environment.
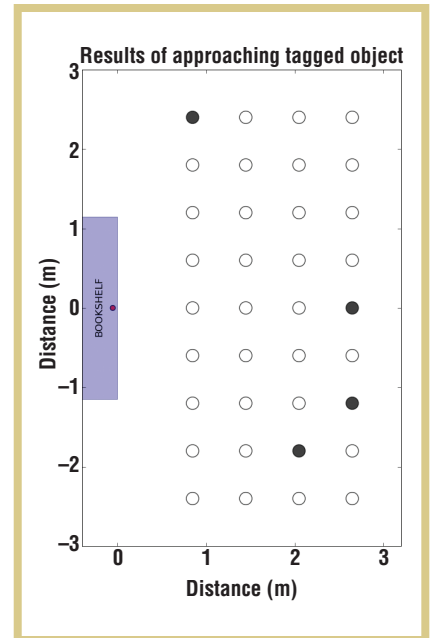
Figure 6 shows how the tag's location estimate improves as the array approaches a stationary tag. The coarse estimates' mean error was 0.4 m ($\sigma = 0.2$ m) in range and 5.1 degrees ($\sigma = 3.6$ degrees) in bearing. During this test, the reader was 1 to 4 meters from the tag in an office environment. We could potentially run this type of estimation in parallel with servoing to inform higher-level navigation methods.

Perceiving the Tag with RSSI Images
We've also developed a mode of perception that produces RSSI spatial-distribution images for each tagged object by mechanically panning and tilting the long-range antennas. Each pixel's intensity in the RSSI image is the interpolated and smoothed RF signal strength for a singulated tag in the corresponding direction. Figure 7 shows EL-E using RSSI images to track an object moving across the scene in corresponding camera images.[3]

RSSI images have three main benefits. First, they provide an intuitive visualization of a tag and antenna's RF properties in a given environment—in essence, showing what the RF signal looks like—which is helpful for development and debugging. Second, you can use them to estimate a tag's bearing in both azimuth and elevation. Third, you can fuse RSSI images with other sensor modalities, including camera

Results of approaching tagged object

Figure 6. A particle filter can be used to estimate a tag's position. (a) The particle filter is initialized when reading a tag. It outputs a position estimate (the pink circle) of the true tag position (the green circle) as an antenna array (the red circle) moves. (b) The estimate improves as the antenna array approaches the stationary tag. (c) The estimate approaches the tag's true position.
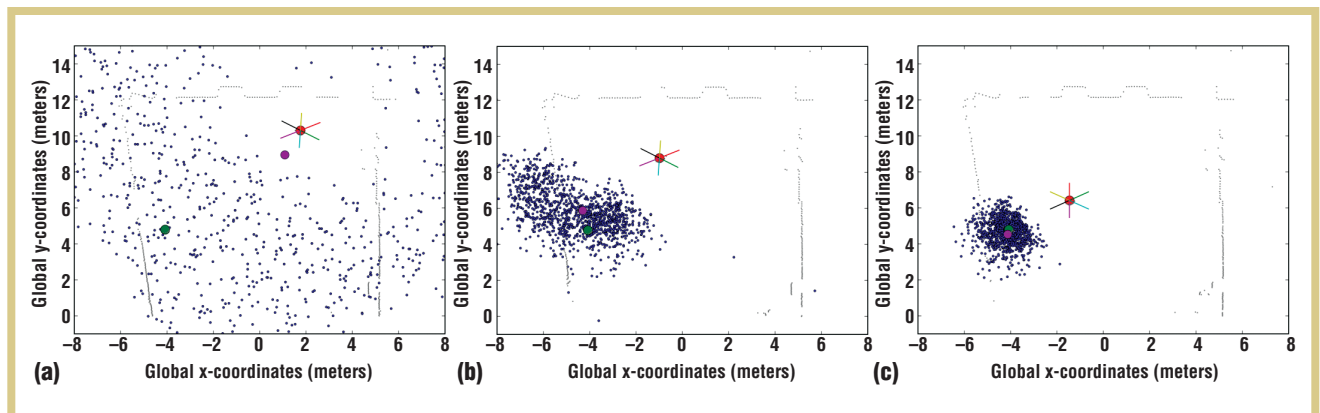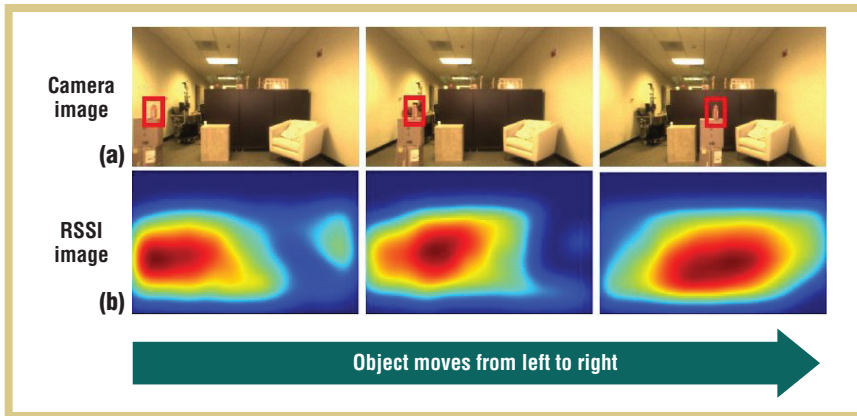
Figure 7. EL-E can use received signal strength indicator (RSSI) to track moving objects. (a) Camera images and (b) RSSI images show an RFID-tagged bottle (in the red square) moving from left to right, as captured by an early antenna rig. This technique can help EL-E locate tags in the environment.

images and lidar (light detection and ranging).

Our previous research showed an 11 percent improvement in object localization using an RSSI image in conjunction with a camera image and lidar (17 out of 18 trials) rather than a camera image and lidar alone (15 out of 18 tri-

als). We also performed tests in which EL-E fetched tagged objects by fusing these three sensing methods. Figure 8 shows a color histogram associated with the tag's ID that helped EL-E visually detect the object. EL-E successfully approached and grasped a selected object in each of three trials. These results suggest that RSSI images provide sensing that's complementary to vision and lidar.[3]

Antenna design is a big challenge for RSSI imaging. If the line-of-sight signal strength dominates signal strength from alternate paths (multipath interference), the bearing with the largest value will directly correspond to the tag's bearing. Highly directive antennas with a narrow angle of sensitivity can reject multipath interference and produce easy-to-use RSSI images, but such antennas can be quite large. EL-E's current antennas, which are

Figure 8. Sensor fusion using UHF RFID. (a) Raw data from three sensors. The desired tagged object is in the red square in the camera image. (b) EL-E loads sensor-specific probabilistic feature models from a semantic and perceptual database indexed by the tag's ID. (c) EL-E calculates and multiplicatively combines probabilities for each feature. (d) The sensor fusion results in the 3D location of the object.
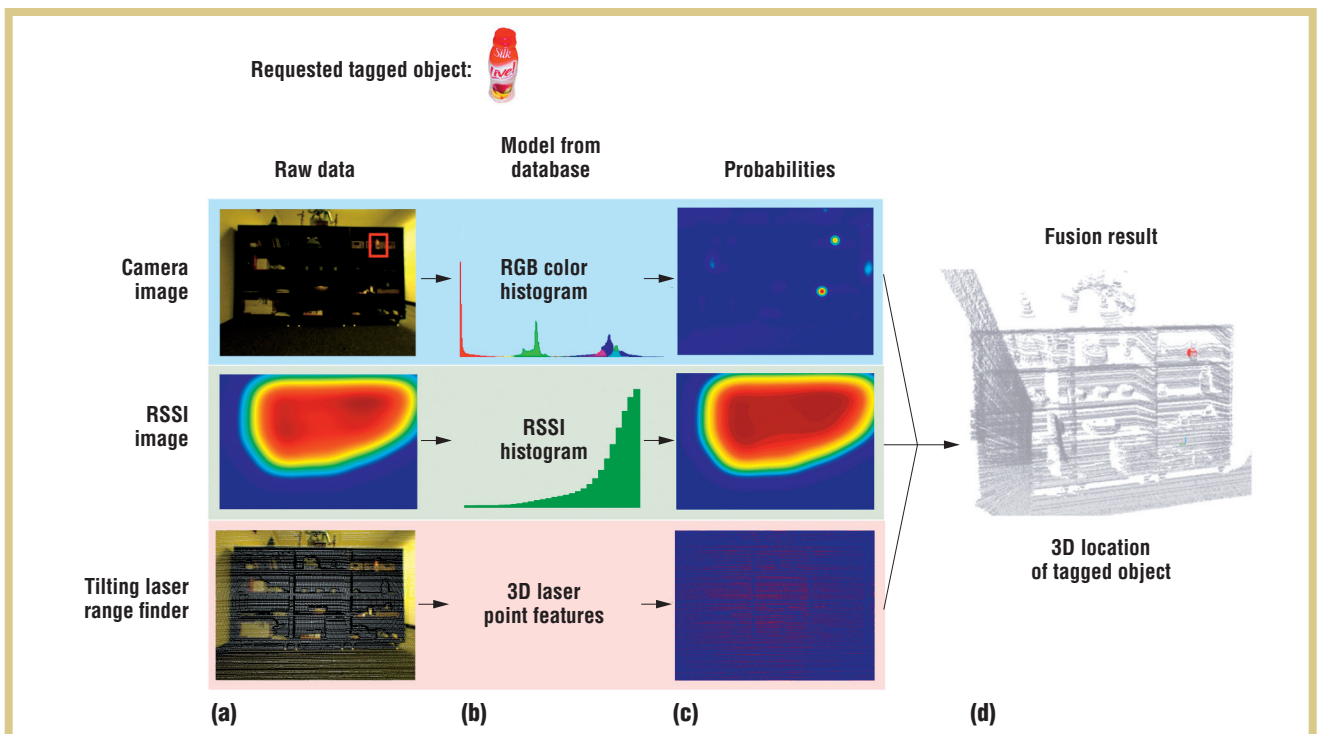
relatively compact (13 × 13 cm) and have 100-degree half-power beamwidths, don't perform as well as EL-E's previous antennas, which were large (26 × 26 cm) and had 65-degree half-power beamwidths.

## Manipulating a Tagged Object

When EL-E is within 1 meter of an object, RFID plays a different role in EL-E's mobile manipulation. High-precision localization becomes important, as does semantic information related to object manipulation.

Short-Range RF Perception

Because EL-E's long-range antennas aren't discriminative at short ranges, we developed finger-mounted antennas, which read the same UHF tags at a range of approximately 20 cm. These antennas excite UHF RFID tags in the magnetostatic near-field regime and allow EL-E to verify that the correct object is being manipulated. EL-E reads each antenna twice, checking the tag ID associated with the read possessing the largest RSSI value. EL-E uses the same method to identify an object someone places in its hand. Before performing the eight reads (two for each antenna), EL-E moves its hand up, away from potentially distracting tags.

We plan to integrate several other uses for these antennas by adapting our long-range-perception techniques to contactless short-range perception. For example, we servoed EL-E's arm on the basis of differential RSSI readings such that the end effector centered on a selected RFID-tagged object. EL-E could also distinguish and localize a particular tagged object among visually identical objects by monitoring RSSI values while its arm executed a trajectory that passed a short-range wrist-mounted antenna in front of objects. By correlating the RSSI values with lidar-provided 3D points, EL-E could locate and grasp a desired medication bottle next to visually identical bottles.[4]

Our results indicate that future finger-mounted antennas could help
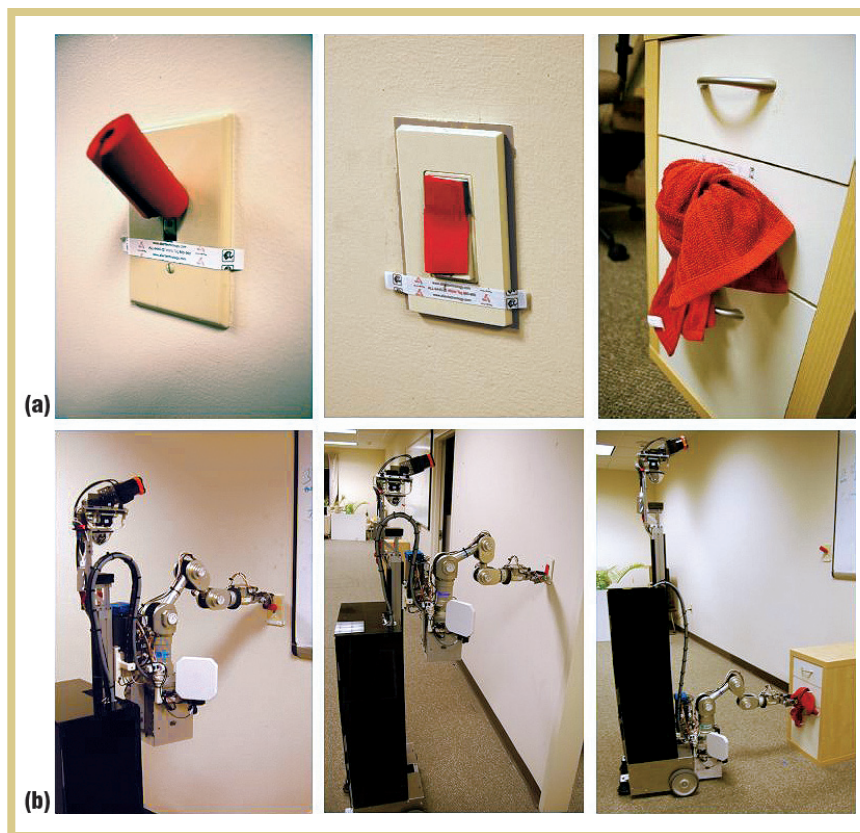
Figure 9. Physical, perceptual, and semantic (PPS) tags combine compliant and colorful materials with a UHF RFID tag to assist EL-E with physical manipulation, perception, and semantic understanding. (a) We affixed a PPS tag to a flip-type light switch, rocker-type light switch, and cabinet drawer. (b) EL-E manipulates the corresponding PPS tags.

finely position EL-E's hand with respect to tagged objects. In addition, close-range RFID sensing could complement our current lidar- and camera-based methods.

Semantic Databases for Physical Interaction

One particularly interesting use of the database at this range is to provide physically grounded semantics related to manipulation. We examined tags employing a more general class of environmental augmentation for task-relevant locations.[5] These tags helped EL-E physically interact with a location, perceive it, and understand its semantics. We call them PPS (physical, perceptual, and semantic) tags. Figure 9 shows three PPS tags, each of which

combines easy-to-manipulate compliant material with easy-to-perceive color and a UHF RFID tag.

The RFID tag provides a unique ID that indexes into a database storing information such as which actions EL-E can perform at the location, how EL-E can perform these actions, and which state changes EL-E should observe on task success. Figure 10 shows a sample top-level database entry for a rocker-type light switch.

Each database entry contains three main components. *Properties* store information about the tagged object not specific to any particular action. *Actions* map user-friendly names to the associated *behaviors*. An entry for each behavior relevant to the tagged object stores parameters, such as

```
{
        'properties':      {'type': 'ada light switch',
                            'name': 'A D A light switch 1',
                            'pps_tag': 'dycem',
                            'change': 'overall brightness',
                            'switch_travel': 0.02,
                            'height': 1.22,
                            'on_plane': True,
                            'direction': 'up',
                            'ele': {'color_segmentation':
                                    [[34, 255], [157, 255],
                                    [0, 11]]},
                            },

        'actions':         {'off': 'push_bottom',
                            'on': 'push_top'},

        'push_bottom':     {'force_threshold': 3.0,
                            'height_offset': -0.02
                            'ele': { 'gripper': 5}
                            },

        'push_top':        {'force_threshold': 3.0,
                            'height_offset': 0.02,
                            'ele': {'gripper': 5}
                            }

}
```

Figure 10. A semantic database entry for operating a rocker-type light switch. The current database is written in Python. The semantic database, indexed by a tag ID, facilitates robotic manipulation of tagged objects and locations.

EL-E's hand configuration when performing the action and the forces to expect. For example, in Figure 10, the properties, actions, and behaviors are as follows.

Properties. type stores the object class, such as ada light switch (ada stands for Americans with Disabilities Act). name stores a name specific to this particular object. pps_tag defines the physically helpful material used, such as dycem (a high-friction rubber sheet). change describes the state change when EL-E uses the object successfully, such as using the camera to detect overall brightness changes in lighting. direction tells EL-E where to look to observe this state change, such as up for the light switches. ele contains a table with information specific to EL-E. In this case, it holds the color boundaries that segment the PPS tag's red color with EL-E's camera.

Actions and behaviors. For ada light switch, the two associated actions are turning the light on and off, which map to the push_top and push_bottom behaviors, respectively. Each behavior also has an entry that stores parameters important to performing the behavior. For example, push_bottom holds information critical to pushing the bottom of the rocker switch to turn the light off. It has two entries. force_threshold, with a value of 3 Newtons, describes the force to apply when pushing. height_offset, with a value of −0.02 m, describes how far below

the PPS tag's center to push. The ele entry for push_bottom specifies the opening angle that EL-E's gripper should use when performing this action. Five degrees places the gripper in a pinching configuration useful for pushing the switch.

Evaluating Manipulation

We evaluated EL-E using the three PPS tags in Figure 9. For each tagged object, we conducted 10 trials with five initial locations evenly spaced by 35 cm along a line 1.7 m from the tagged object, running parallel to the wall.

EL-E initially faced the wall. In each trial, it tried to generate the UI, approach the user-selected tag, and manipulate the tag. Manipulation involved turning a flip light switch on or off, turning a rocker light switch on or off, or opening or closing a drawer.

EL-E performed successfully in nine of the trials for each tag, so its overall success rate was 27 out of 30—90 percent. In all 30 trials, EL-E correctly verified the tag ID before manipulation, using its finger-mounted antennas. It also correctly determined task success or failure (such as observing whether the lighting changed), using information from the database. So, EL-E could have tried again after its three recognized failures.

UHF RFID is a promising way to create useful robots in the near future. However, numerous challenges remain. Foremost is the need to test UHF RFID-guided robots in diverse real-world environments, such as homes and healthcare facilities. So far, we've tested EL-E only in a lab environment with limited clutter. The ability to read a tag and estimate its location depends on the object, the tag's pose relative to the antenna, and the environment's RF properties. How these factors affect real-world tasks remains an open question; serious considerations, such as antenna design, re-

quire further investigation. Likewise, we must study many usability issues, such as the ease of tagging objects, long-term tag reliability as tags are handled by people and by EL-E, and tagged objects' usability.

By circumventing long-standing object recognition problems, labeling the environment with UHF RFID tags opens up new avenues for robotics research. Robots could learn from their interactions with tagged objects over a long time period. The investigation of knowledge representations for manipulation changes from a somewhat esoteric subject to an area of immediate practical concern. A common sense knowledge repository with standardized labels for robots could be on the horizon, with robots of varying capabilities sharing their knowledge. We can imagine a smarter sensor-rich robot traveling through the environment, tagging locations, and recording relevant information for use by less sophisticated robots.

Although we ultimately hope to develop robots that don't require environment modification, we believe labeling the world with UHF RFID tags offers significant advantages at this time. This modest form of infrastructure could accelerate robot deployment in real-world applications, which could directly benefit society, make robots more affordable, and provide valuable research opportunities. ◘

REFERENCES

1. *Class 1 Generation 2 UHF RFID Protocol for Operation at 860 MHz–960 MHz, Ver. 1.0.9*, EPCglobal US, 2005; www.epcglobalus.org/Standards/EPCglobalStandards/Class1Generation2UHFAirInterface Protocol/tabid/322/Default.aspx.

2. T. Deyle, C.C. Kemp, and M.S. Reynolds, "Probabilistic UHF RFID Tag Pose Estimation with Multiple Antennas and a Multipath RF Propagation Model," *Proc. 2008 IEEE/RSJ Int'l Conf. Intelligent Robots and Systems* (IROS 08), IEEE Press, 2008, pp. 1379–1384.

3. T. Deyle et al., "RF Vision: RFID Receive Signal Strength Indicator (RSSI) Images for Sensor Fusion and Mobile Manipulation," *Proc. 2009 IEEE/RSJ Int'l Conf. Intelligent Robots and Systems* (IROS 09), IEEE Press, pp. 5553–5560.

4. T. Deyle et al., "A Foveated Passive UHF RFID System for Mobile Manipulation," *Proc. 2008 IEEE/RSJ Int'l Conf. Intelligent Robots and Systems* (IROS 08), IEEE Press, 2008, pp. 3711–3716.

5. H. Nguyen et al., "PPS-Tags: Physical Perceptual and Semantic Tags for Autonomous Mobile Manipulation," *Proc. IROS 2009 Workshop Semantic Perception for Mobile Manipulation*, 2009.

the **AUTHORS**

Travis Deyle is a PhD student at the Georgia Institute of Technology's School of Electrical and Computer Engineering (ECE). He's also a member of Georgia Tech's Healthcare Robotics Lab and Health Systems Institute. His research interests include novel sensor technologies, robotics, and ubiquitous computing. Deyle has an MS in ECE from Georgia Tech. Contact him at tdeyle@gatech.edu.

Hai Nguyen is a PhD student at the Georgia Institute of Technology's Robotics and Intelligent Machines Center. He's also a member of Georgia Tech's Healthcare Robotics Lab and Health Systems Institute. His research interests include machine learning for robotics, perception with novel sensors, and mobile manipulation. Nguyen has a BS in computer science from Georgia Tech. Contact him at haidai@gmail.com.

Matthew S. Reynolds is an assistant professor at Duke University's Department of Electrical and Computer Engineering. His research interests include the physics of sensors and actuators, RFID, and signal processing. Reynolds has a PhD from the Media Laboratory at the Massachusetts Institute of Technology. He's a member of the Signal Processing and Communications, and Computer Engineering groups at Duke as well as the IEEE Microwave Theory and Techniques Society. Contact him at matt.reynolds@duke.edu.

Charles C. Kemp is an assistant professor at the Georgia Institute of Technology's Wallace H. Coulter Department of Biomedical Engineering and an adjunct assistant professor in the School of Interactive Computing. He's also a member of Georgia Tech's Center for Robotics and Intelligent Machines and Health Systems Institute. His research interests include autonomous mobile manipulation, human-robot interaction, assistive robotics, healthcare robotics, bio-inspired approaches to robotics, and AI. Kemp has a PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. Contact him at charlie.kemp@bme.gatech.edu.

cn Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.