

DOI:10.1145/2428556.2428574

The MAGIC 2010 robot competition showed how well multi-robot teams can work with human teams in urban search.

**BY EDWIN OLSON, JOHANNES STROM, ROB GOEDEL,
RYAN MORTON, PRADEEP RANGANATHAN,
AND ANDREW RICHARDSON**

Exploration and Mapping with Autonomous Robot Teams

URBAN RECONNAISSANCE AND search-and-rescue missions are ideal candidates for multi-robot teams due to the potential hazard the missions pose to humans and the inherent parallelism that can be exploited by teams of cooperating robots. However, these domains also involve challenging problems due to having to work in complex, stochastic, and partially observable environments. In particular, non-uniform and cluttered terrain in unknown environments is a challenge for both state-estimation and control, resulting in complicated planning and perception problems. Limited and unreliable communications



further complicate coordination among individual agents and their human operators.

To help address the problems, the Multi-Autonomous Ground robot International Challenge (MAGIC), held November 2010 in Adelaide, Australia, brought together five teams, including nearly 40 robots, to pursue more than \$1 million in prize money in a competition organized and funded by the Australian government's Defence Science and Technology Organisation (DSTO, <http://www.dsto.defence.gov.au/MAGIC2010/>) and the U.S. Army's Research, Development and Engineering Com-



Figure 1. Team Michigan robots. Michigan deployed 14 custom-made robots that cooperatively mapped a 500m × 500m area; each included a color camera and laser range finder capable of producing 3D point clouds.

mand (RDECOM, www.rdecom.army.mil/). The teams were instructed to explore and map a large indoor-outdoor area while recognizing and neutralizing threats (such as simulated bombs and enemy combatants). Although the contest showcased the ability of teams to coordinate autonomous agents in a challenging environment it also reflected the limitations of the state of the art in state estimation and perception (such as map-building and object recognition) (see Figure 1).

MAGIC is the most recent of the robotics Grand Challenges, following in the tradition of well-known com-

petitions sponsored by the Defense Advanced Research Projects Agency (DARPA) tracing back to a 2001 U.S. congressional mandate requiring one-third of all ground combat vehicles to be unmanned by 2015. Over the course of the DARPA challenges, teams developed technologies for fully autonomous cars, including the ability to drive in urban settings, navigating moving obstacles and obeying traffic laws.^{18,19} Moreover, they fostered devel-

» key insights

- Human operators can help a robot team be more efficient and recover from errors.
- A good state estimate, in the form of a map, is the most critical piece of information for a team of robots—and the most difficult to obtain.
- Grand Challenge competitions like MAGIC highlight challenging open problems and provide a venue for evaluating new approaches.

opment of new methods for planning, control, state estimation, and perhaps most important, robot perception and sensor fusion.

Unfortunately, however, these advances were not mirrored in smaller robots (such as those used by soldiers searching for and neutralizing improvised explosive devices, or IEDs) or for robots intended to help first responders in search-and-rescue missions. Instead, tele-operation, or remote-joystick control by a human, remains the dominant mode of interaction (see Figure 2). These real-world systems pose several challenges not in the DARPA challenges:

Limited/unreliable GPS. The Global Positioning System (GPS) is often unreliable or inaccurate in dense urban environments or indoors. GPS can also be jammed or spoofed by an adversary; the winning DARPA vehicles relied extensively on GPS;

Multi-robot cooperation. Robots are individually generally less capable than humans, with their potential

arising from multi-robot deployments that explicitly coordinate with one another; and

Humans in the loop. By allowing humans to interact with robot teams in real time, the system becomes more effective and adaptable to changes in mission objectives and priorities; this ability entails developing visualization methods and user-interface abstractions that allow humans to understand and manipulate the state of the team.

MAGIC focused on increasing the effectiveness of multi-robot systems by increasing the number of robots a single human operator can manage effectively. This is in contrast to more-traditional robot systems that typically require one or more operators per robot.² Participants were required to deploy a team of cooperating robots to explore and map a hostile area, recognize and catalog the location of interesting objects (such as people, doorways, IEDs, and cars), and perform simulated neutralization of IEDs using a laser pointer. Two human operators were allowed to interact with

each robot team, but interaction time was measured and used to calculate a penalty to each team's final score.

The contest attracted 23 teams from around the world, a number reduced through a series of competitive down selects to five finalists invited to the final competition at the Adelaide Showgrounds, a 500m × 500m area including indoor and outdoor spaces. Aerial imagery provided by contest organizers was the only prior knowledge. While previous DARPA challenges provided detailed GPS waypoints describing the location and topology of safe roads, MAGIC robots would have to figure out such information on their own. Whereas other search-and-rescue robotics contests typically focus on smaller environments with significant mobility and manipulation challenges (such as the RoboCup Rescue League, <http://www.robocuprescue.org/>), MAGIC was at a much larger scale, with greater focus on autonomous multi-robot cooperation.¹⁶

To succeed, a team had to combine robot perception, mapping, planning, and human interfaces. Here, we highlight some of the key decisions and algorithmic choices that led to Michigan's first-place finish.¹⁴ Additionally, we highlight how Michigan's mapping and state-estimation system differed from the other competitors, one of the key differences setting Michigan apart.

System Design

The Michigan system was largely centralized: A ground-control station near the center of the competition area collected data from individual robots, fused it to create an estimate of the current state of the system (such as position of robots and location of important objects), then used it to assign new tasks to the robots. Most robots focused on exploring the large competition area, a task well suited to parallelization. However, other robots were able to perform additional tasks (such as neutralizing IEDs). The discovery of such a device would cause a "neutralize" task to be assigned to a nearby robot. Each team's human operators were positioned at the ground-control station where they could view current task assignments, a map of the operating area, and (perhaps most important) guide the system by vetting sensor data

Figure 2. Finalist robots.




(a)



(b)

Each team used a unique robot platform (in ranked order, left to right): Michigan had 14 custom-built robots; Penn had seven custom-built robots; the Reconnaissance and Autonomy for Small Robots team, principally organized by Robotics Research LLC and QinetiQ, based its seven robots on the Talon commercial platform; MAGICian, a coalition of Australian schools, adapted a commercial base for its five robots; and Cappadocia, a coalition based mainly in Turkey, had six custom-built robots.



(c)



(d)

or overriding task assignments.

Michigan's robots received their assignments via radio and were responsible for executing their tasks without additional assistance from the ground-control station; for example, robots used their 3D laser range finder to identify safe terrain and avoid obstacles on their own. They were also responsible for autonomously detecting IEDs and other objects. The information they gathered (including object-detection data and a map of the area immediately around the robot) was heavily compressed and transmitted back to the ground-control station; in practice, these messages were often relayed by other robots to overcome the limited range of Michigan's radios. With the newly collected information, the ground-control station updated its map and user interfaces and computed new (improved) tasks for each robot. This process continued until the mission was complete.

Such a system poses many challenges: How does the ground-control station compute efficient tasks for the robots in a way that maximizes the efficiency of the team? How can humans be kept informed about the state of the system? How can humans contribute to the performance of the system? How do robots reliably recognize safe and unsafe terrain? How do robots detect dangerous objects? How can the information collected by robots be compressed sufficiently so it can be transmitted over a limited, unreliable communications network? And how does the ground-control station combine information from the robots into a single globally consistent view?

Recognizing that many of these tasks rely on an accurate, detailed map of the world, Michigan focused on fusing robot data into a globally consistent view. The accuracy of the map was a primary evaluation criterion in the competition, as well as a critical component in effective multi-agent planning and the human-robot interface; for example, where should robots go next if one does not know where they are now or where they have already been.

A notable difference between Michigan and the other teams was the accuracy of the maps it produced. Map quality pays repeated dividends throughout the Michigan system, with correspond-



While MAGIC posed many technical challenges, mapping and state estimation were arguably most critical.



ing improvement in human-robot interfaces and planning. The variability in map quality from team to team is a testament to the difficulty and unsolved nature of multi-robot mapping. Michigan began with a state-of-the-art system, but it was inadequate in terms of both scaling to large numbers of robots and dealing with the errors that inevitably occur. New methods, both automatic and humans in the loop, were needed to achieve adequate performance; the following section explores a few of them.

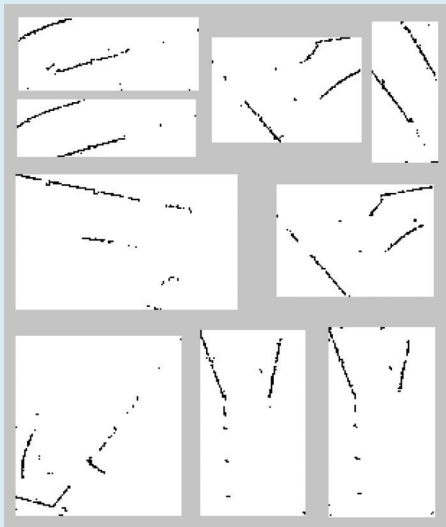
Technical Contributions

While MAGIC posed many technical challenges, mapping and state estimation were arguably most critical. Using GPS may seem like an obvious starting point, but even under best-case conditions, it cannot provide a navigation solution for the significant fraction of the time robots are indoors. Outdoors, GPS data (particularly from consumer-grade equipment) is often fairly good, within a few meters, perhaps. GPS can also be wildly inaccurate due to effects like multi-path. In a combat situation, GPS is easily jammed or even spoofed. Consequently, despite having GPS receivers on each robot, Michigan ultimately opted not to use GPS data, relying instead on its robots' sensors to recognize landmarks. This strategy was not universally adopted, however, with most teams using GPS in some way.


Mapping and state estimation. Conceptually, map building can be viewed as an alignment problem, with robots periodically generating maplets of their immediate surroundings using a laser scanner. The challenge is to determine how to arrange the maplets so they form a large coherent map, much like the process of assembling a panoramic photo from a number of overlapping photos (see Figure 3). Not only can the system assemble a map this way but also the position of each of the robots, since each is at the center of its own maplet.

Michigan's state-estimation system was based on a standard probabilistic formulation of mapping in which the desired alignment can be computed through inference on a factor graph; see Bailey and Durrant-Whyte¹ and Durrant-Whyte and Bailey⁷ for a survey of other approaches. The Michigan


Figure 3. Mapping overview. Individual maplets (a) are matched in a pairwise fashion.



(a)



(b)



(c)

The resulting network of constraints can be represented through a factor graph similar to (c) in which circles represent robot positions and squares probabilistic constraints. The final map (b) is computed by re-projecting all sensor observations according to maximum-likelihood robot positions.

factor graph included nodes for unknown variables—the location of each maplet—and edges connecting nodes when something is known about the relative geometric position of the two nodes. Loosely speaking, an edge encodes a geometrical relationship between two maplets; that is “maplet A is six meters east and rotated 30 degrees from maplet B”; none of these relationships is known with certainty, so edges are annotated with a covariance matrix. A map commonly contains many of these edges, with many of them subtly disagreeing with one another.

More formally, let the position of all maplets be represented by the state vector x , which can be quite large, as it contains two translation and one rotation component for each maplet, and there can be thousands of maplets. Edges convey a conditional probability distribution $p(z_i|x)$, where z_i is a sensor measurement. This quantity is the measurement model; given a

particular configuration of the world, it predicts the distribution of the sensor; for example, a range sensor might return the distance between two variable nodes plus some Gaussian noise the variance of which can be characterized empirically.

Our goal was to compute $p(x|z)$, or the posterior distribution of the maplet positions, given all sensor observations. Using Bayes’s rule, and assuming we have no a priori knowledge of what the map should look like (or $p(x)$ is uninformative), we obtain:

$$p(x|z) \propto \Pi p(z_i|x) \quad (1)$$

The goal was to find the maplet positions x that has maximum probability $p(x|z)$. Assuming all the edges are simple Gaussian distributions of the form $e^{-(z_i - \mu)^T \Sigma^{-1} (z_i - \mu)}$, this computation becomes a nonlinear least-squares problem. Specifically, we can take the logarithm of both sides, which con-

verts the right-hand side into a sum of quadratic losses. We maximize the log probability by differentiating with respect to x , resulting in a first-order linear system. The key idea is that maximum likelihood inference on a Gaussian factor graph is equivalent to solving a large linear system; see Thrun et al.¹⁷ for a more detailed explanation. The solution to this linear system yields the position of each maplet.

The resulting linear system is extremely sparse, as each edge typically depends on only two maplet positions. In the Michigan system, each maplet was generally connected to from two to five other maplets. Sparse linear algebra methods can exploit this sparsity, greatly reducing the time needed to solve the linear system for x . The Michigan method was based on sparse Cholesky factorization;⁶ we could compute solutions for a graph with 4,200 nodes and 6,300 edges in about 250ms on a standard laptop CPU. New data is

always arriving, so this level of performance allowed the map to be updated several times per second.

An important advantage of using the factor graph formulation is that it is possible to retroactively edit the graph to correct errors; for example, if a sensing subsystem erroneously adds an edge to the graph (incorrectly asserting that, say, two robot poses are a meter apart), we could “undo” the error by deleting the edge and computing a new maximum likelihood estimate. Such editing is not possible with methods based on, say, Kalman filters. In this case, we relied on human operators to correct these relatively rare errors, discussed later.

Scan matching and loop validation.

The Michigan mapping approach depended on identifying high-quality edges; more edges generally result in a better map since the linear system becomes over-constrained, reducing the effect of noise from individual edges.

The system used several different methods to generate edges, including dead reckoning (based on wheel-encoder odometry and a low-cost inertial measurement unit, or the set of sensors that measures acceleration and rotation of the robot) and visual detection of other robots using their 2D “bar codes,” as in Figure 1.¹⁰ But the most important source of edges in the Michigan system (by far) was its scan-matching system, attempting to align two maplets by correlating them against each other, looking for the translation and rotation that maximize their overlap. One such matching operation (see Figure 4) includes the probability associated with each translation and rotation computed in a brute-force fashion.

This alignment process is computationally expensive, and in the worst case, each maplet had to be matched with every other maplet. In practice, the robot’s dead-reckoning data can help rule out many false matches. But with 14 robots operating simultaneously, and with each one producing a new maplet every 1.4 seconds, hundreds or thousands of alignment attempts per second are needed.

The Michigan approach to mapping was based on an accelerated version of a brute-force scan-matching system.¹¹ The key idea is a multi-resolution matching system, generating low-res-

olution versions of the maplets and a first attempt to align them. Because they are smaller, the alignment is much faster. Good candidate alignments are then attempted at higher resolution.

While simple in concept, a major challenge was ensuring the low-resolution alignments did not underestimate the quality of an alignment that could occur with higher-resolution maplets. The Michigan solution relied on constructing the low-resolution maplets in a special way; rather than apply a typical low-pass-filter/decimate process (which would tend to obliterate structural details), we used a max-decimate kernel to ensure matches between low-resolution maplets never underestimate the overlap that could result from aligning full-resolution maplets. When aligning low-resolution maplets, we never underestimated the overlap that could result from aligning the full-resolution maplets.

Michigan’s earlier scan-matching work¹¹ considered two different resolutions, allowing approximately 50 matches per second. This would be adequate for a small number of robots, but for a larger team, it becomes a bottleneck. For MAGIC, we modified the approach to consider matches over a full pyramid of reduced-resolution images that resulted in matching rates of approximately 500 matches per second. The quality of the resulting map ultimately depended on the number of matches found, and a higher processing rate increases the likelihood


of finding these matches. Our fast-matching system was pivotal in keeping up with our large robot team. Other teams used similar maplet-matching strategies though were not as fast; for example, the Australian MAGICian team reported its GPU-accelerated system was capable of seven to 10 matches per second.

The improvement in our matching speed allowed us to consider a large number of possible matches in real time to support our global map. However, our state-of-the-art method had a non-zero false-positive rate, aligning maplets based on similar-looking structures, even if the maplets are not actually near each other.

There is a fundamental trade-off between number of true positives and increased risk of false positives. Increasing the threshold for what constitutes a “good-enough” match also increases the likelihood that similar looking, but physically distinct, locations are matched incorrectly. Such false-positive matches can cause the inference method to distort the map to explain the error.

To reduce the false-positive rate to a usable level, we performed a loop-validation step on candidate matches before the system could add them to the factor graph. The basic idea of loop validation is to require that multiple matches “agree” with each other.^{4,12,13} Consider a topological “loop” of matches: A match between node A and B, another match between B and

Figure 4. Brute-force search for best maplet alignment. The search space is 3D (two translation, one rotation) represented as a series of 2D cross-sections.



C, and a third match between C and A. If the matches are correct, then the composition of their rigid-body transformations should approximately be the identity matrix. When this occurs, the system adds the matches to the factor graph.

Human-robot interfaces. In simple environments (such as an indoor warehouse), the combination of loop-validation and automatic scan-matching presented here were sufficient for supporting completely autonomous operation of Michigan's robot team (see Figure 5). However, in a less-structured environment (such as many of the outdoor portions of MAGIC 2010), mapping errors would still occur; for example, the MAGIC venue included numerous cable conduits that caused robots to unknowingly get stuck, causing severe dead-reckoning estimation error. At the time of MAGIC 2010, Michigan's system was not able to handle such problems autonomously.

However, map errors are relatively obvious to human operators. Michigan thus developed a user interface that allowed human operators to look for

errors and intervene when necessary. With new (validated) loop closures being added to the graph at a rate of two to three per second, a human operator could easily be overwhelmed by asking for explicit verification of each match.

Instead, human operators would monitor the entire map. When an error occurred (typically visible as a distortion in the map), the operator could "roll back" automatically added matches until the problem was no longer present. The operator could then ask the mapping system to perform an alignment between two selected maplets near where the problem had been detected. This human-assisted match served as additional a priori information for future autonomous matching operations, making it less likely the system would repeat the same mistake.

Michigan found this approach, which required only a few limited interactions to remove false positives, a highly effective use of humans to support the continued autonomy of its robots' planning system. Michigan was the only team to build a user interface that allowed direct supervision of the real-time state estimate; other teams

handled failures in automatic state estimation by requiring humans to track the global state manually, then intervene at the task-allocation level. Early versions of the system lacked a global-mapping system, with human operators providing separate map displays for each robot. Michigan's experience with this approach indicated that operators could not effectively handle more than five or six robots this way. Maintaining a global map is critical to scaling to larger robot teams, and the Michigan user interface was a key part of maintaining map consistency.


Evaluation

The main evaluation metrics for an autonomous reconnaissance system are quality of the final map produced and amount of human assistance required to produce it and were also the primary metrics the MAGIC organizers used to determine the winner and subsequent ranking of the finalists, as in Figure 2. While the specific performance data used in the competition was not made public, we present selected results we obtained by processing our logs; we also compare with other teams' published results where possible.

Lacking detailed ground truth for the MAGIC venue, the best evaluation of map quality is necessarily subjective; Figure 6 compares post-processed maps for the Michigan team against the mapping software of MAGICian (fourth place) applied to the data collected by the Penn team (second place); additionally, the map produced by the Michigan system (inset), includes distortions resulting from erroneous matches that, in the interest of time, the human operators chose not to correct. This result shows that high-quality maps can be produced in this domain; Michigan's competition-day results showed our state estimation was sufficiently good to be useful for supporting online planning. The system allowed us to completely explore the first two phases of the competition while simultaneously performing mission objectives relating to dynamic and static dangers (such as IEDs and simulated mobile enemy combatants).

We would also would like to measure the frequency of human inter-

Figure 5. Indoor-storage-warehouse map. In uncluttered environments posing few mobility challenges, Michigan's team of 14 robots could explore and map with little human intervention.



action required to support our state estimation system during MAGIC. However, Michigan did not collect the data necessary to evaluate this metric during our run; we thus replicated the run by playing back the raw data from the competition log and having our operator reenact his performance during the competition. These conditions are obviously less stressful than competition but are still representative of human performance. The result (see Figure 7) was the addition of 175 loop closures, on average two interactions per minute, which generally occurred in bursts. However, at one point, the operator did not interact with the system for 5.17 minutes.


Our evaluation shows we were able to support cooperative global state estimation for a team of autonomously navigating robots using only a single part-time operator. Yet there remain significant open problems, including how to reduce human assistance even further by improving the ability of the system to handle errors autonomously. Additional evaluation of the system, as well as technical descriptions of the other finalists, can be found elsewhere, including in Boeing et al.,³ Butzke et al.,⁵ Erdener,⁸ Lacaze et al.,⁹ and Olson et al.¹⁴

Discussion


MAGIC's focus was on increasing the robot-to-human ratio and efficiently coordinating the actions of multiple robots. Key to reducing the cognitive load on operators is how to increase the autonomy of robots; for a given amount of cognitive loading, more robots can be handled if they simply require less interaction. We identified global state estimation as a key technology for enabling autonomy and believe the mapping system we deployed for MAGIC outperforms the systems of our competitors. While this was one of the key factors differentiating Michigan from the other finalists, it was not the only important point of comparison. In fact, many of the other choices we made when developing the system also had a positive effect on our performance.

In particular, we made a strategic decision early on that we would emphasize a large team of robots. This is reflected in the fact that we brought

Figure 6. Minimally post-processed maps from Michigan's robots (a) and MAGICian's mapping algorithm using Penn's data (b) from Reid and Brauni.¹⁵ The map produced online by the Michigan robots is inset top-left.



(a)




(b)

twice as many robots to the competition as the next largest team. This strategy ultimately affected the design of all our core systems, including mapping, object identification, and communication. Given that we had a finite budget, it also forced us to deploy economical robot platforms with only the bare necessities in sensing to complete the challenge. The result was our robots were also the cheapest of any of the finalists (by a significant margin), costing only \$11,500 each.

One approach to detecting dangerous objects is to, say, transmit video feeds back to human operators and

rely on them to recognize the hazard. Given a design goal of maximizing the number of robots, such a strategy is unworkable; there is neither the bandwidth to transmit that many images nor could humans be expected to vigilantly monitor 14 video streams. The system simply had to be able to detect dangerous objects autonomously, whereas other teams with fewer robots could be successful with less automation. At the same time, handling more tasks autonomously also meant our human operators had more time to assist with mapping tasks.

Figure 7. Map interaction experiment. Michigan’s mapping operator reenacted the supporting role for the phase 2 dataset to measure the frequency of interaction required to maintain a near-perfect state estimate; see Figure 6 for resulting map. The human workload was modest, averaging only two interactions per minute.



An interesting question relates to the optimal size of the robot team; larger teams have greater parallelism-exploiting potential but also place greater demands on human operators. These demands are strongly dependent on the reliability and autonomous capabilities of the robots; less-capable robots would place greater demands on the operators. An area of future work involves trade-offs between the size of a robot team and the cognitive load on human operators and how the autonomous capabilities of the robots affect this trade-off.

Conclusion

MAGIC resulted in significant progress toward urban search using teams of robots aided by human operators. Michigan’s approach, emphasizing accurate mapping, helped maximize the autonomous capabilities of its robots and maintain the operator’s situational awareness, allowing two humans to effectively control a larger team of robots. However, MAGIC also highlighted the shortcomings of state-of-the-art methods. It remains difficult to maintain a consistent map for large numbers of robots; Michigan’s competition-day maps still show distortions due to errors in the system’s matching capability. The system coped with these errors at the cost of greater operator workload we continue to target in our ongoing work.

Competitions like MAGIC highlight open technological challenges in areas often viewed as “solved.” MAGIC thus

brought the prospect of cooperative teams of robots and humans closer than ever, but also highlighted the challenging research problems that remain.

Acknowledgments

Team Michigan was a collaboration between the University of Michigan’s APRIL Robotics Laboratory (<http://april.eecs.umich.edu/>) and Soar Technology (<http://www.soartech.com/>). In addition to the authors here, core team members included Mihai Bulic, Jacob Crossman, and Bob Mariner; we were also supported by more than two dozen undergraduate researchers. We thank the MAGIC contest organizers who ran the competition and prepared the contest venue. And special thanks go to our liaison, Captain Chris Latham of the 9th Combat Service Support Battalion of the Australian Army. Our participation would not have been possible without the help of our sponsors at Intel and Texas Instruments.

References

1. Bailey, T. and Durrant-Whyte, H. Simultaneous localization and mapping (SLAM): Part II state of the art. *Robotics and Autonomous Systems* 13, 3 (Sept. 2006), 108–117.
2. Barnes, M. and Jentsch, F. *Human-Robot Interactions in Future Military Operations*. Ashgate Publishing Company, Brookfield, VT, 2010.
3. Boeing, A., Boulton, M., Brunl, T., Frisch, B., Lopes, S., Morgan, A., Ophelders, F., Pangen, S., Reid, R., and Vinsen, K. WAMbot: Team MAGICian’s entry in the Multi-Autonomous Ground-Robotic International Challenge 2010. *Journal of Field Robotics* 5 (Sept./Oct. 2012), 707–728.
4. Bosse, M.C. *ATLAS: A Framework for Large-Scale Automated Mapping and Localization*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, Feb. 2004.
5. Butzke, J., Danilidis, K., Kushleyev, A., Lee, D.D., Likhachev, M., Phillips, C., and Phillips, M. The University of Pennsylvania MAGIC 2010 Multi-Robot

- Team. *Journal of Field Robotics* 5 (Sept./Oct. 2012), 745–761.
6. Dellaert, F. and Kaess, M. Square root SAM: Simultaneous localization and mapping via square root information smoothing. *International Journal of Robotics Research* 25, 12 (Dec. 2006), 1181–1203.
7. Durrant-Whyte, H. and Bailey, T. Simultaneous localization and mapping (SLAM): Part I, the essential algorithms. *Robotics and Autonomous Systems* 13, 2 (June 2006), 99–110.
8. Erdener, A., Ari, E.O., Ataseven, Y., Deniz, B., Ince, K.G., Kazancioglu, U., Kopanoglu, T.A., Koray, T., Kosaner, K.M., Ozgur, A., Ozkok, C.C., Soncul, T., Sirin, H.O., Yakin, I., Biddlestone, S., Fu, L., Kurt, A., Ozguner, U., Redmill, K., Aytikin, O., and Ulusoy, I. Team Cappadocia design for MAGIC 2010. In *Proceedings of the Land Warfare Conference* (Brisbane, Australia, Nov. 15–19, 2010).
9. Lacaze, A., Murphy, K., Giorno, M.D., and Corley, K. The Reconnaissance and Autonomy for Small Robots (RASR) MAGIC 2010 Challenge. In *Proceedings of the Land Warfare Conference* (Brisbane, Australia, Nov. 15–19, 2010).
10. Olson, E. AprilTag: A robust and flexible visual fiducial system. In *Proceedings of the IEEE International Conference on Robotics and Automation* (May). IEEE, 2011, 3400–3407.
11. Olson, E. Real-time correlative scan matching. In *Proceedings of the IEEE International Conference on Robotics and Automation* (Kobe, Japan, June). IEEE, 2009, 4387–4393.
12. Olson, E. Recognizing places using spectrally clustered local matches. *Robotics and Autonomous Systems* 57, 12 (Dec. 2009), 1157–1172.
13. Olson, E. *Robust and Efficient Robotic Mapping*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2008; <http://april.eecs.umich.edu/papers/details.php?name=olson2008phd>
14. Olson, E., Strom, J., Morton, R., Richardson, A., Ranganathan, P., Goeddel, R., and Bulic, M. Progress towards multi-robot reconnaissance and the MAGIC 2010 competition. *Journal of Field Robotics* 29, 5 (Sept. 2012), 762–792.
15. Reid, R. and Braund, T. Large-scale multi-robot mapping in MAGIC 2010. In *Proceedings of the IEEE Conference on Robotics, Automation, and Mechatronics* (Sept. 17–19). IEEE, 2011, 239–244.
16. Saenbunsiri, K., Chaimuengchuen, P., Changlor, N., Skolapak, P., Danwiang, N., Ppoosuwan, V., Tienkum, R., Anan Rakrajathulthum, P., Nitisuchakul, T., Bumrungjitt, K., Tunsiri, S., Khairid, P., Santi, N., and Yan Primee, S. *RobotCupRescue 2011: Robot League Team iRAP JUDY (Thailand)*. Technical Report, 2011.
17. Thrun, S., Burgard, W., and Fox, D. *Probabilistic Robotics*. MIT Press, Cambridge, MA, 2005.
18. Thrun, S. et al. Stanley: The robot that won the DARPA Grand Challenge. In *The 2005 DARPA Grand Challenge, Vol. 36 of Springer Tracts in Advanced Robotics*. Springer Berlin/Heidelberg, 2007, 1–43.
19. Urmson, C. et al. Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics* 25, 8 (Aug. 2008), 425–466.

Edwin Olson (ebolson@umich.edu) is an assistant professor in the Department of Computer Science and Engineering at the University of Michigan, Ann Arbor.

Johannes Strom (jhsstrom@umich.edu) is a Ph.D. student focusing on communications and mapping in the Department of Computer Science and Engineering of the University of Michigan, Ann Arbor.

Rob Goeddel (rgoeddel@umich.edu) is a Ph.D. candidate focusing on exploration and planning in the Department of Computer Science and Engineering of the University of Michigan, Ann Arbor.

Ryan Morton (rmorton@umich.edu) is a Ph.D. candidate focusing on terrain classification and human interaction in the Department of Computer Science and Engineering of the University of Michigan, Ann Arbor.

Pradeep Ranganathan (rpradeep@umich.edu) is a Ph.D. student focusing on learning and inference in the Department of Computer Science and Engineering of the University of Michigan, Ann Arbor.

Andrew Richardson (chardson@umich.edu) is a Ph.D. candidate focusing on computer vision in the Department of Computer Science and Engineering of the University of Michigan, Ann Arbor.