Advancing the State of Urban Search and Rescue Robotics through the RoboCupRescue Robot League Competition

Raymond Sheh, Adam Jacoff, Ann-Marie Virts, Tetsuya Kimura, Johannes Pellenz, Sören Schwertfeger, and Jackrit Suthakorn

Abstract The RoboCupRescue Robot League is an international competition that has grown to be an effective driver for the dissemination of solutions to the challenges posed by Urban Search and Rescue Robotics and accelerated the development of the performance standards that are crucial to widespread effective deployment of robotic systems for these applications. In this paper, we will discuss how this competition has come to be more than simply a venue where teams compete to find a champion and is now "A League of Teams with one goal: to Develop and Demonstrate Advanced Robotic Capabilities for Emergency Responders."

1 Introduction

Competitions have often been used as a way of encouraging collaboration and directing research energy towards particular goals. In this paper, we discuss how the RoboCupRescue Robot League (RoboCup RRL), through competitions and other events, is an effective driver for its three main objectives: increasing awareness of the challenges inherent in Urban Search and Rescue (USAR) Robotics, providing objective performance evaluations for these robots, and promoting collaboration, research, and dissemination of best-in-class solutions to the challenges posed by this

Raymond Sheh, Adam Jacoff and Ann-Marie Virts
National Institute of Standards and Technology, Maryland, USA
e-mail: raymond.sheh@robolit.com, [adam.jacoff|ann.virts]@nist.gov

Tetsuya Kimura
Nagaoka University of Technology, Niigata, Japan
e-mail: kimura@mech.nagaokaut.ac.jp

Johannes Pellenz
Bundeswehr Technical Center for Engineer and General Field Equipment, Koblenz, Germany
e-mail: pellenz@uni-koblenz.de

Sören Schwertfeger
Jacobs University Bremen, Bremen, Germany
e-mail: s.schwertfeger@jacobs-university.de

Jackrit Suthakorn
Mahidol University, Nakorn Pathom, Thailand e-mail: egjst@mahidol.ac.th

Presented at the 8th International Conference
on Field and Service Robotics
(FSR 2012), Matsushima, Japan, 16-19 July 2012

domain. In particular, the League has become a competition and community that is geared not to competition against rivals, but rather to moving forward, as a world-wide league of teams, towards solutions to very real challenges. There are many other application domains in the area of field and service robotics where a similar approach may be used to focus development, collaboration, and dissemination among the research community.

The RoboCup RRL was started in response to the 1995 Kobe earthquake and began in association with the broader RoboCup Competitions in 2001. Each year over one hundred teams from academic institutions all over the world compete in regional open competitions that culminate in the International RoboCup RRL Competition. These events are built around the demonstration of advanced robotic capabilities for emergency response applications. These include advanced sensing for situational awareness, mapping and surveying, articulated arms for manipulation and directed perception, mobility platforms for a wide variety of terrains, autonomous behaviours for operator assistance and survivability, algorithms for data interpretation and map building, and human-robot interfaces for highly complex real-time control of robots.

During the course of a week-long competition, 15 to 25 teams, qualified through a competitive regional open and paper submission process, field robots that undertake a series of autonomous and remotely teleoperated search missions to locate simulated victims in a specially constructed arena, an example of which is shown in Figure 1. The robots must generate a map with useful information on the state of victims and hazards for emergency responders. This arena consists of a maze of terrains and challenges based on current and emerging standard test method apparatuses for response robots, developed by the US National Institute of Standards and Technology (NIST) through the ASTM International Committee on Homeland Security Operations; Operational Equipment; Robots (E54.08.01). The competition acts as a stepping stone between the laboratory and wider deployment. The apparatuses represent the real world operational requirements gathered from emergency responders, distilled into a form that allows for the repeatable evaluation of robot performance. They reflect the many different challenges that a response robot may encounter. Teams vying for the championship must demonstrate reliable, cutting edge capabilities across the entire arena. Teams may also win Best-in-Class Champion awards for demonstrating Best-in-Class performance in specific problems in USAR robotics.

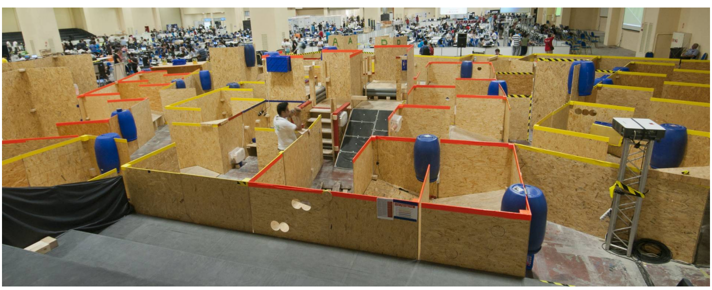

Fig. 1 The RoboCup RRL Arena from the 2011 competition. The tops of the walls are color coded corresponding to the arena areas (the Red and Black arenas are hidden behind the raised wall sections). A map of this arena appears in Figure 2.

The RoboCup RRL forms part of the development and dissemination process for standard test methods for response robots, led by NIST [3]. The standards process is driven from the operational requirements of emergency responders who benefit from the development of these tools, such as police, explosive ordinance disposal (EOD) technicians, fire fighters, and rescue and military personnel. At the competition, these test methods are further refined in the presence of the wide variety of robotic solutions to the challenges presented. The competition also helps to disseminate the test methods and apparatuses as teams recreate the arenas in their own labs. Valuable input from researchers, working on the next generation of capabilities, are also gathered and fed back into the standards process. Likewise, exposure to the operationally relevant apparatuses gives teams valuable exposure to real world requirements and highlights the gaps in current capabilities. The competition also gives NIST the opportunity to expose the emergency responders to demonstrations of next generation capabilities. NIST invites teams with best-in-class implementations to response robot evaluation exercises, held at search and rescue training facilities, in order to demonstrate their capabilities to emergency responders, vendors, and officials, with the aim of generating the demand that will shorten the time taken for developments in the lab to reach deployment. As we have focused this paper on how the RoboCup RRL is an effective driver of research and collaboration, we have restricted our reporting of such topics as rules, history, specific test methods, arena construction, and results. For an overview, we invite the reader to refer to the current League Overview [8]. The reader is also invited to refer to the current Rules Outline, Arena Construction Guide, and Community Wiki [4, 9, 6].

2 The RoboCupRescue Robot League

Several unique features of the RoboCup RRL Competition make it particularly effective in promoting research and collaboration. First and foremost, the RoboCup RRL is presented as a competition that does not consider teams as rivals. Rather, the rivalry is between the teams and the application, represented by the arena and the mission specifications. The competition is carefully structured to give teams every chance possible to demonstrate and disseminate their advanced capabilities to their full potential. The RoboCup RRL also specifically encourages teams with well integrated, near-deployable robotic systems, as well as teams with experimental systems that push the boundaries in specific areas and demonstrate Best-in-Class capabilities. It is usually the case that major advances in specialized capabilities come from groups that have specialized and devoted the majority of their energies into that area. Making their inclusion in the competition a priority allows those advances to be demonstrated alongside the more mature, but perhaps less advanced implementations. This makes the competition a unique venue whereby researchers may disseminate their research to others and accelerate the progress of these technologies towards eventual deployment.

The competition starts with a preliminary round where all teams have the opportunity to fully and equally compete in. The aim of this round is to provide all teams with the opportunity to demonstrate their capabilities and obtain statistically significant performance data on their systems. The preliminary round takes up one-half of the week and is held in two half-sized arenas to allow all teams to complete six 15-minute missions. Team performances in the Preliminary missions decide their qualification for the Championship and Best-in-Class competitions. Teams may compete in all of these competitions or they may decide to focus on one or a few. The Championship competition consists of a semi-final and final round and decides the overall 1st, 2nd, and 3rd place winners. The qualification process, discussed in Section 2.3, ensures that teams demonstrating competent, statistically significant performance across the whole arena qualify for the Championship. However, there are also opportunities for more specialized teams to demonstrate their capabilities in the Championship. The Best-in-Class awards, of which there are currently three, are decided on the basis of both the Preliminary missions and a dedicated set of Best-in-Class missions, described in Section 2.4. In the rest of this section, we will present in further detail how the different components of the competition work together to foster collaboration among teams with very different capabilities, allow teams to demonstrate their work, and advance the work of the community as a whole. A more general discussion of the competition logistics and the current outline of the rules appears in the current Rules Outline [4].

2.1 Arena

The arena in which the competition is held represents a building, such as a small house or apartment, in various stages of collapse. The arena that was used in the 2011 international competition is shown in Figures 1 and 2 and consists of current and emerging standard test method apparatuses embedded in a maze. The RoboCup RRL is geared towards showcasing the abilities of robots that are still in the research lab and that are rarely hardened for deployment. Therefore, many of the less practical or unsafe challenges such as dust, water, mud, fire, chemical hazards, truly unstable terrain, and radio disruptions have been omitted or replaced with more practical substitutes. The simulated victims that teams must locate in order to score points are distributed throughout the arena in such a way as to require them to demonstrate mastery of these apparatuses. In this section we will describe how the design and layout of the arenas and victims were influenced by the objectives of the League. Further details of the makeup of the arena elements and victims may be found in the Arena Construction Guide [9]. As this evolves from year to year, the reader is also invited to refer to the Community Wiki [6]. The arena is divided into three main arenas: Yellow, Orange, and Red, shown in Figure 3. These form a continuum of increasing difficulty and focus on different current and future challenges faced by Urban Search and Rescue robotics. In addition to the three major arena areas, three specialty arenas exist to encourage developments in specific challenges relevant to USAR Robotics. They are the Blue and Black arenas, shown in Figure 3, and the Radio Dropout Zone. By dividing the continuous arena into discrete areas, specialist teams are able to focus their energies on the arenas that lie within their area of expertise, where they are able to compete without disadvantage against the Championship teams. It also provides the League with a context in which to introduce new test method apparatuses that reflect the refined operationally relevant requirements gleaned from emergency responders through the ongoing standards development process.

The Yellow arena, in which all missions begin, represents a building to navigate that is largely intact and is designed to showcase teams that are able to navigate and identify victims autonomously. It is arranged as a maze with continuous but non-flat flooring in order to encourage teams to implement robust, 3D-aware mapping and localisation. Although it is the easiest of the three arena areas, robots that find and score victims in this arena must do so without any human control. For teleoperated robots, it represents a navigation speed test. The faster they can reach the other arenas, the more time they will have to perform the rest of their mission. The Orange arena represents an area that has suffered from additional collapse and showcases the ability of robots to cover a wide variety of structured terrains and accurately place their sensors in hard-to-reach locations. It is laid out as a maze with multiple levels and discontinuous flooring. Embedded in the Orange arena are structured mobility test method apparatuses such as steep ramps, stairs, steps, and mismatched ramps. All robots may score victims in the Orange arena, however all of these victims are challenging. For instance, they may be placed in difficult locations such as very close to the ground or against steep ramps and only accessible through small holes, or high on walls next to sloped flooring. The Red arena represents complex terrain such as a rubble pile and showcases the abilities of advanced mobility robots and the systems that allow operators to effectively drive these robots in highly unstructured environments. It consists of a stepfield [2], designed to replicate unstructured terrain in a way that is reproducible. This arena tests advanced mobility where the primary challenge is getting to the victim and staying in location while the victim is surveyed – victims are usually easy to view once the robot is co-located. Barriers and victims are placed around the arena so robots must demonstrate well controlled movement. The Black arena, an extension of the Red arena, represents a less standardized part of the arena where robots can demonstrate their abilities in more realistic scenarios. It usually takes the form of victims hidden inside a full sized vehicle.

The Blue arena tests the ability of robots to manipulate objects and carry them as they traverse the terrain in the arena. It consists of shelves containing objects, such as water bottles, radios, and boxes, and grids of holes in which the objects may be placed. In the main competition, teams earn additional points by taking objects from

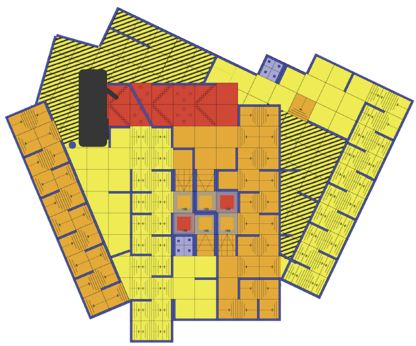

Fig. 2 Map of an example RoboCup RRL Arena showing the Red, Orange, and Yellow arenas, the Black arena (car with victims), Yellow/Black Radio Dropout Zone, and Blue manipulation area. This arena is approximately 15 m×19 m.

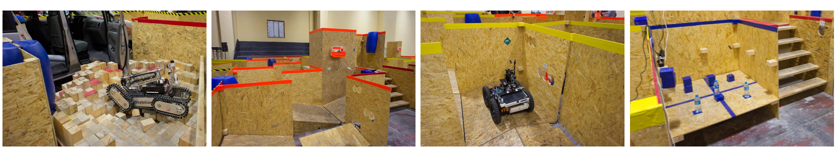

Fig. 3 The Red and Black arenas, Orange arena, Yellow arena and Blue arena as configured for the Preliminary and Championship missions.

Fig. 4 Examples of simulated victims that appear throughout the arena.

the shelves and placing them with any victims that they find. As mobile manipulation capabilities are still in their infancy among League teams, teams are allowed to start the mission with a single object in the robot's possession. The Radio Dropout Zone, also called the Black/Yellow arena, represents an area of a largely intact building that is experiencing radio interference and encourages teams to incorporate both mobility and autonomy on a single robot. Although it resembles the Yellow arena, it is placed beyond the Orange arena and thus robots require some degree of mobility to reach it. While teams may teleoperate their robots up to the Radio Dropout Zone, once inside the robot must navigate autonomously. On reaching the end of the Zone, teams may resume teleoperation to score victims. As a further incentive, successful autonomous navigation back to the start of the Zone doubles the team's score for finding the zone's sole victim.

Simulated victims are the objectives that all of the robots seek in the arena. They are evenly distributed throughout the three main arenas; an additional victim is assigned to the Radio Dropout Zone. In order to score a victim, the team's robot must be co-located with the victim (usually defined as the centre of the robot's base within 1.2m of the victim) and have line-of-sight to the victim with its sensors. To score well, robots often must deploy their sensor package right up against the victim, requiring slender, dexterous arms and compact sensor packages. Examples of simulated victims are shown in Figure 4. They consist of several co-located signs of life: a human form, heat, motion, sound, and carbon dioxide. Several of these signs of life taken together are used to identify the state of the simulated victim: unconscious, semi-conscious, or aware. These signs of life have been chosen to be detectable with widely available equipment and yet require sensing and processing that reflects that required to sense real objects of interest. Alongside the simulated victims are test method artifacts that evaluate the ability of the robots to survey the environment, such as hazardous materials labels and eye charts, placed in hard-to-see areas.

Victims are placed in boxes that are either open or accessed through holes of varying sizes, changing the level of difficulty. Victims located in areas that pose greater challenges, such as advanced mobility or full autonomy, are easier to access while victims in areas that have easier terrain are placed in smaller holes that require the robots to be positioned in difficult locations and require novel ways of rapidly and reliably directing perception. Reaching a victim and reporting their condition is only part of the challenge. Teams must also locate the victim on a human-readable georeferenced Geospatial Tag Image File Format (GeoTIFF) map of the arena that their robot autonomously generates as it traverses the arena. Points are scored for the accuracy of the map and its usefulness to an emergency responder who should be able to use it to locate the victim [7].

2.2 Missions

In this section we will first discuss the overall structure of the missions, followed by a brief discussion of the particular features of the missions, which make up the Preliminary and Championship rounds, that help to achieve the goals of the League, as described earlier in this paper. A more complete discussion of the mission and rules appears in the Rules Outline [4]. Each time-limited mission represents a search-and-rescue deployment. Teams deploy their robots to the start point in the arena and their operator interface to the operator station, which is out of sight of the arena. Once set up, the single robot operator uses their single teleoperated robot and multiple autonomous robots to navigate and map the arena, locate the victims, determine their signs of life, and deliver objects. The mission ends when their allotted time has expired, after which they have a short period of time to produce an automatically generated map of the arena with victims identified and localized. The team's score depends on the quality of this map and the extent and accuracy of the information they are able to gather about the victims.

The League uses the scoring process to encourage teams to implement solutions to gaps in current capabilities, as identified through the ongoing standards process. For example, points are awarded for the ability to obtain thermal information about the victim. While teams are able to score some points using non-contact thermometers, more points are available to teams that are able to return a thermal image, resulting in some innovative approaches to performing thermal image surveys of the arena using low cost sensors. Likewise, automatic mapping, 3D perception, and high resolution imaging are encouraged using the scoring metric. As in the rest of the competition, the missions have been designed to facilitate experimentation, allow teams to disseminate their developments to other teams and the general public, collaborate, and assist in the refinement of the standard test methods. The League's objectives are best satisfied if teams are confident that they will be able to showcase their capabilities to their full potential. However, the implementations seen in the competition are rarely deployment-hardened and are highly experimental. The competition has been carefully structured to reduce the risk that luck or random factors result in a team not reaching their full potential. Teams qualify for the Championship rounds based on overall points scored during the preliminaries. To encourage teams to push their robot and to reward teams for bringing innovative, but possibly imperfect, implementations, the worst preliminary mission score for each team is discarded. Furthermore, the qualification cutoff is decided based on the distribution of scores after the preliminary round. There should be a clear performance gap between the best eliminated team and the worst qualified team.

Prior to each mission, the team's operator, or their representative, is asked to walk the arena, identify all of the simulated victims, and check their signs of life. This ensures that bad luck, in the form of the operator randomly choosing a poor path through the maze or a broken simulated victim, plays no part in the ability of the team to showcase their capabilities. While it also means that the operator knows the location and signs of life of each victim, they must still navigate the arena and present convincing evidence to the referee, by way of their user interface, in order to score points. Once the mission has begun, teams are encouraged to have a "spotter" in the arena, who is responsible for ensuring the safety of the robots. Although the mission ends if the spotter touches the robot or interferes with its sensors, it allows teams to experiment and push their robots without fear of damaging them. Teams are also allowed to reset their missions after an intervention, if the robot becomes immobilized, or at any other point of their choosing. This entails moving the robots back to the starting location, resetting their maps and other internal state, repairing, modifying, or replacing the robots as necessary and beginning again. However, while their mission score is set to zero, the clock is kept running so teams effectively have as many missions as they like, within the 15 or 20 minutes they have been allotted. At the end of the time, the mission in which they scored the most points is recorded. In this way, teams may experiment with new settings or approaches, knowing that they will be able to salvage the mission and recover to still earn points. It also reduces the cost of failures early in a mission, before they have had a chance to demonstrate their abilities. During the course of the mission, other teams and members of the public are encouraged to watch the operator and the robots, both of which are projected onto a large screen. This is further facilitated by swapping the order in which teams run between days so that all teams have the opportunity to observe all other teams while not preparing for their own mission. As a further incentive to disseminate their work, a small but significant number of points are given to teams that produce an informative slide and assign a representative to provide a live commentary to the audience during their mission. Following the mission, the maps that each team produces to document their mission is scored by the referees. These are publicly posted, along with their scores and the salient features that either earned points or showed scope for improvement. This allows teams to see how each others' mapping sensors and algorithms have affected their maps and scores and identifies the Best-in-Class implementations to aspire towards.

2.3 Championship

At the start of the semi-final rounds, all team scores are reset to zero. This allows high scoring teams that are confident of qualifying to experiment during the preliminary rounds, where the competition is less fierce, without jeopardising the higher stakes semi-final and final rounds. Depending on the number of qualified teams, the semi-finals are held in either a half-sized or full-sized arena and consume another day of competition. Following the semi-finals, the best performing three or five teams move on to a finals round in the full-sized arena that usually consists of two short missions with scores carrying over from the semi-finals and the final 3 places decided on points. Teams that present integrated systems that competently perform in all aspects of the arena always qualify for the Championship rounds. Very inexperienced and highly specialized teams are often eliminated. These teams are encouraged to use the remaining one or two days to practice for the Best-in-Class competitions on the final day. Indeed, many specialized teams come to the competition only to compete in the preliminaries and the Best-in-Class competitions with no intention of qualifying for Championship. Yet these teams are not out of the main competition. To further encourage the entry of specialist teams and to foster a spirit of collaboration, each qualified team is strongly encouraged to pair with a team that has been eliminated and progress through to the Championship as a combined team. Each member of the combined team is awarded should the combined team receive the 1st, 2nd, or 3rd place awards. In recent years, it has been common for teams that focus on mechatronic engineering, and who often perform well across the Orange and Red arenas, to qualify and merge with teams that focus on artificial intelligence, who dominate the Yellow autonomous arena and have excellent mapping but lack the resources to tackle the more physically challenging arenas.

2.4 Best-in-Class

As previously discussed, the field of USAR robotics is so large that many of the breakthroughs in the advancement of capabilities come from specialized research groups. To recognize and reward these groups, who may lack the resources to demonstrate broad competence in enough areas to win the Championship, the League also provides three Best-in-Class awards. Equal in status to the Championship, these awards focus on specific challenges that exist in this field, as identified through consultation with emergency responders as part of the broader standards development process. The current Best-in-Class competitions focus on Mobility, Autonomy, and Mobile Manipulation. Qualification is based on the preliminary missions and performance in these rounds contributes one-half of the Best-in-Class score. The second half comes from a dedicated Best-in-Class round. There is also a separate award, decided by the Organizing Committee, for the team that demonstrates the most remarkable, operationally relevant innovation in user interface development.

The Best-in-Class Mobility competition challenges robots with advanced mobility to demonstrate their ability to rapidly traverse unstructured terrain. Half of a robot's Best-in-Class Mobility score consists of the number of victims that were found by the robot in the Red arena during the preliminary round; teams must find at least one victim in this arena to qualify. The second half of the Best-in-Class Mobility score is based on the standard test method for advanced mobility: stepfields [2]. It takes the form of a race against the clock around the Red arena stepfields. In a fixed time period (usually 10 minutes) teams must drive their robot in a figure-of-8 pattern around the stepfield as many times as possible. While being one of the oldest challenges in the competition, many real world terrains continue to thwart deployed solutions. Through this challenge, teams are encouraged to not only develop highly mobile, robust mobility platforms, but also to develop effective ways of controlling the robots through careful camera placement for situational awareness, reliable low latency communications, and intuitive, low cognitive load user interfaces.

The Best-in-Class Autonomy competition was developed to encourage teams to pursue autonomous victim identification and navigation in continuous but non-flat flooring and in the presence of difficult terrain that must be avoided. Half of the robot's score consists of the number of victims that were found autonomously during the Preliminary missions. The second half of the Best-in-Class Autonomy score is earned during a single Best-in-Class Autonomy mission. The arena is reconfigured so that the Yellow arena becomes one large loop around the whole arena and the mission is for the robot to map as much of the arena as possible within a set time limit. The maps are then evaluated based on coverage and accuracy [7]. There is only a small, albeit growing, subset of teams that are competitive in the Best-in-Class Autonomy competition. However, the technologies required to perform well in this challenge have greater applicability in augmenting existing capabilities. For example, autonomous victim identification drives the development of sensor fusion and object class recognition. This may be equally applied to recognizing not only actual victims but also other objects of interest, such as hazardous material placards or shipping labels across a wide area, something that a robot operator, concentrating on the task at hand, can easily miss. Autonomous navigation can assist the operator by helping to perform long downrange traverses and enhance robot survivability by allowing it to act appropriately should radio communication be impaired. Automatic mapping is also a valuable technology that is only just starting to reach deployment;

this competition serves as a valuable proving ground for the test methods that are being developed to evaluate these new capabilities.

The Best-in-Class Mobile Manipulation competition is the newest of the three and encourages the development of dexterous, intuitive mobile manipulation capabilities. Half of the score for this competition comes from the number of objects that were placed with victims during the Preliminary missions. The other half comes from the dedicated Best-in-Class Mobile Manipulation competition, based on the emerging standard test methods for mobile manipulation. Teams use their robots to take as many objects as possible from the shelf and place them in a grid of holes in the Blue Arena within 10 minutes. There is a very real need for dexterous mobile manipulators with intuitive control in the field of USAR robotics. Apart from a few salient examples, robots used in these applications have rudimentary dexterous manipulation capabilities when compared to those in other domains such as telemedicine. Clearly this is a domain where the performance of the bulk of deployable solutions, and to a large part the expectations of the end users, trail the state-of-the-art in the academic community by a significant margin. This competition plays an important part in not only evaluating the performance of the next generation of implementations, but also in helping disseminate their progress to end users and vendors and assisting their rapid integration into fieldable implementations.

One of the main areas of open research in USAR robotics is in the quality of the user experience. This includes all aspects of the robotic system that influence what the user observes through the interface, from the interface itself right through any assistive technologies to the sensor selection, placement, and direction capabilities. It also includes the whole process by which the user controls the robot. To encourage innovations that push the state-of-the-art in this area, the League presents an Innovative User Interfaces award, selected at the end of the competition by the Organizing Committee based on the implementations demonstrated during the week of competition. This award is given to teams that demonstrate implementations that solve an operationally relevant user interface and experience problem in the field of USAR robotics. The League community is made up of groups with a wide variety of approaches to the user interface experience. In the past, this award has been presented to teams for developing wearable interfaces, interfaces based on metaphors such as car driving, and interfaces that use autonomous behaviours to assist the operator.

2.5 Rules and Administration

The RoboCup RRL is primarily administered by NIST as part of its work on standard test methods for response robots. NIST guides the design of the arena and the embedded test method apparatuses, provides the general outline of the competition rules and scoring metrics, defines the Best-in-Class challenges, and manages various competition logistics [4]. Embedded in the arena design, standard test method apparatuses and rule outlines are the operationally relevant requirements that NIST has distilled from its close collaboration with emergency responders and industry. This provides NIST with a powerful way of using its position at the crossroads of researchers, vendors, and emergency responders, to guide the competition towards the challenges that exist in the real world, in a way that is feasible for research-grade implementations to address.

The focus of the RoboCup RRL is not on finding a winner but rather to encourage research and collaboration. To this end, the rules themselves are not set in advance.

Rather, based on the outline, the rules are presented, discussed, amended, and agreed upon during meetings at the start of the competition and again at the start of every day with leaders from every competing team. The rules are allowed to evolve from one day to the next to respond to new developments, close loopholes, and adjust scoring to best reward promising implementations. This approach means that teams have little incentive to try to "game" the system – any exploitation of the rules will simply result in that loophole being closed. Instead, teams are encouraged to bring general solutions to the types of challenges that exist in the field of USAR robotics and as emphasised by the published arena design guides and rule outlines. During the competition, much like in a real deployment, teams that do well are the ones that are able to improvize, adapt, and perform rapid development as the conditions, apparatuses, and rules evolve.

3 Dissemination of Solutions

The RoboCup RRL has expanded to host regional open competitions through the year, often serving as qualifying competitions that feed into the international competition. Particularly significant regional opens include those held in Germany, Japan, Iran, and Thailand. The Iranian community, inspired by the 2003 earthquake in Bam, has grown to be one of the major forces in the international competition, fed by a vibrant domestic competition of several hundred teams. The Thai competition is particularly notable, in only a few short years it has rapidly grown to encompass over 100 domestic teams with a championship trophy that is granted and presented by the Thai royal family. As a result of resources and support provided to the domestic Thai RoboCup RRL community, Thai teams dominate the world competition, winning the Championship several years in a row. In the process they have demonstrated many practical advancements that have been disseminated throughout the league, particularly in mobility and directed perception.

It is rare for teams to finish the competition without having made significant improvements to their robot's capabilities in response to its performance in the arenas and the dissemination of best-in-class implementations from other teams. However, it is a high pressure environment, where teams are ultimately concerned with demonstrating their capabilities to their full potential. In order to further encourage collaboration and the dissemination of best-in-class solutions, especially among groups that are considering entering the League, the RoboCup RRL community hosts annual international Rescue Robotics Camps and Summer Schools. Representatives from teams that have exhibited Best-in-Class capabilities in previous competitions are invited to present their work and lead tutorials to help other teams and researchers build on their work. Topics that these events have been particularly effective at disseminating in the past include 2D and 3D mapping, user interface design, and autonomy. These events, attended by around 30 participants and led by around 10 speakers and practical leaders each year, have been instrumental in the development of several new teams who now rank as Best-in-Class in their own right. The League also brings in distinguished speakers in relevant fields to introduce the League community to new developments and encourage research to proceed in the direction of gaps in current capabilities.

Through competition experience, some teams have been able to refine their implementations and move beyond the confines of the relatively safe, structured arenas. Teams that perform well in the competition are invited to demonstrate their implementations at Response Robot Evaluation Exercises, held regularly at fire and rescue training facilities to evaluate the performance of deployable systems. These events represent the next stepping stone between the lab and deployment, with emergency responders, procurement officials, and vendors in attendance. While they are still safe, structured events, the environments are more realistic with real concrete and wood rubble, dirt, dust, and water. Sensors are further challenged with a wider variety of surfaces, longer sight lines, and sunlight, while communications are hampered by steel reinforcing and other structures.

Many teams in the League have implementations that demonstrate capabilities several generations ahead of currently deployed capabilities across the spectrum of user interfaces, manipulation, directed perception, autonomy, mapping, sensing, sensor fusion, and mobility. By demonstrating these capabilities in scenarios that emergency responders are familiar with, the League can spark demand for these technologies and, with vendors also in attendance, hasten their integration into deployable systems. At the same time, researchers are able to see more clearly how their work fits with the overall operational scenarios and identify where and how further improvements may be made. Several teams have already taken advantage of these opportunities, which have proven to be valuable to all parties involved. During subsequent meetings with responders it has been clear that exposure to these cutting-edge implementations has changed their way of thinking about how their problems might be solved. Likewise, researchers exposed to these events have found the experience critical in shaping their research directions to be more application-focused. For procurement officers it has also helped them gain a clearer insight into the capabilities that exist outside of a corporate marketing officer's brief.

Many teams have also formed collaborations with local emergency responder communities and fielded their technologies at Urban Search and Rescue facilities in Australia, Germany, Japan, the United States, and beyond. For example, teams from Japan have long been involved in the domestic Japanese emergency responder community. The design of their robots bear the hallmarks of development in collaboration with emergency responders, with unique capabilities in manipulation and mobility, and unusual and specialized mechanisms such as snake-like robots. German RoboCup RRL teams have also participated in events such as the European Land Robot Trials involving military and civilian applications. Thailand, with its vibrant RoboCup RRL community, has seen teams working with emergency responders to deliver robots for operational purposes. These collaborations with the emergency responder community help to validate both the approaches that the teams have taken, as well as the overall direction of the RoboCup RRL. In an ideal world, the advances made by the teams in the RoboCup RRL would not be required. However, disasters do happen and several teams have already seen their robots deployed in full operational scenarios. The 2011 competition saw the sudden withdrawal of several teams from Japan, when their robots and expertise were required in the aftermath of the 2011 East Japan Earthquake. Robots that were developed for the competition played an instrumental role in the disaster response. This included the Quince robot (CIT, Tohoku University, IRS and NEDO), which was initially designed for traversing the stepfields in the competition. Its proven ability to overcome these terrains resulted in it being chosen for inspecting the buildings housing the reactors at the tsunami-damaged Fukushima Dai-1 nuclear power plant [5].

4 Future Directions

The RoboCup RRL is constantly evolving as the state of the art in implementations and test methods change. In particular, the focus in past years on improving mobility, 2D mapping, and sensing have resulted in most teams achieving a very high level of performance in these areas. The next challenge will be in 3D mapping, manipulation, autonomy, and user interfaces – technologies where a small number of teams have demonstrated mastery. The League organization is also refining the links between these focus areas and the real world challenges through collaboration with emergency responders as part of the standards process and other projects and competitions. As these are reflected in the apparatuses, rules, scores, and Best-in-Class awards, it is anticipated that these focus areas will continue to improve.

New test methods on the horizon are also slated for possible inclusion into the League. In particular, standards for other, related, classes of robots, such as those designed to fight fires, are being considered for adaptation into the League. By carefully matching these new challenges to the skill set already present in the community, the League organization is able to leverage the existing capabilities. This helps to drive innovation in a new, much needed area of research, increases awareness among the research community of a new set of research challenges, and helps to advance a new set of standard test methods by which the performance of future implementations may be evaluated. For the competition as a whole, it also allows new teams a niche in which they can grow and excel and compete with the entrenched players; in the process they benefit the League community by bringing and sharing their unique expertise into the community while themselves learning from the other teams in the League. The Explosive Ordinance Disposal (EOD) community has also been interested in leveraging the RoboCup RRL to further research in their domain. This interest fits well with the increasing focus on manipulation tasks in the competition. The RoboCup RRL is working with the EOD community to ensure the relevance of new competition challenges.

While the RoboCup RRL is focused on robots for disaster response, another league that has its origins in the RoboCup RRL community, RoboCup@Home, focuses on service robots in the home. The two leagues share a common approach to promoting research and collaboration, as well as many common research frontiers. The communities are quite distinct, with the RoboCup RRL community dominated by teams with mechanical and mechatronic backgrounds and the RoboCup@Home community, with its focus on human-robot interactions, heavily focused on artificial intelligence. Going forward, the two leagues plan on working together to encourage the RoboCup RRL community to pursue more complex autonomy in wider, more semantically complex environments, while encouraging the RoboCup@Home community to improve their mechanical platforms to cope with more varied and unstructured terrain. In the process, it is hoped that collaboration and dissemination of technologies between the two communities, and a better understanding of each others' challenges, can be further strengthened.

The RoboCupRescue Virtual Robot League was also formed out of the RoboCup RRL community in order to lower the barrier of entry into this field, especially for teams that are heavily focused on artificial intelligence. This league shares the same challenges as the RoboCup RRL, deployed in a validated physics simulator. Another competition that builds on the RoboCup RRL experience is the DARPA Robotics Challenge [1]. This multi-year competition encourages the development of semi-autonomous robots for disaster response that can operate in a designed-for-human world. It incorporates many of the concepts that make the RoboCup RRL a success and leverages many of the same test methods and processes.

Finally, additional outreach events, such as new summer schools, are planned that aim to bring together not only RoboCup RRL based teams, but also other students and early career researchers to work intensively on focused problems that exist in the league. These include software architectures that improve the ability of teams to share software implementations across different robot platforms and the software modules that will allow teams with little software engineering expertise to begin to compete in this space. Other focus areas for these events include 3D mapping, user interfaces, manipulation, and autonomy. The RoboCup RRL has focused the community on the challenges that matter to the future deployment of Urban Search and Rescue Robotics. Through the combination of a solid grounding in the requirements of emergency responders, integration with the wider standards development process that promotes fair and rigorous evaluation, and an active decision to encourage experimentation, research, collaboration, and dissemination as opposed to rivalry purely for the sake of competition, the RoboCup RRL has succeeded in becoming a "League of Teams with one goal: to Develop and Demonstrate Advanced Robotic Capabilities for Emergency Responders."

Acknowledgements The authors would like to thank past and present members of the RoboCup RRL Organizing Committee, administrators, and teams. The Department of Homeland Security sponsored the production of this material under an Interagency Agreement with the National Institute of Standards and Technology. Certain commercial equipment, instruments, or materials are identified in this document. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products identified are necessarily the best available for the purpose.

References

1. DARPA: DARPA seeks robot enthusiasts (and you) to face off for $2M prize! `http://www.darpa.mil/NewsEvents/Releases/2012/04/10.aspx` (2012)
2. Jacoff, A., Downs, A., Virts, A., Messina, E.: Stepfield Pallets: Repeatable Terrain for Evaluating Robot Mobility. In: Proceedings of the 2008 Performance Metrics for Intelligent Systems Workshop (PerMIS'08), NIST Special Publication 1090, August (2008)
3. Jacoff, A., Huang, H.M., Messina, E., Virts, A., Downs, T.: Comprehensive Standard Test Suites for the Performance Evaluation of Mobile Robots. In: Proceedings of the 2010 Performance Metrics for Intelligent Systems Workshop (PerMIS'10), NIST Special Publication 1113, September (2010)
4. Jacoff, A., Tadokoro, S., Mihankhah, E., Kimra, T., Pellenz, J., Birk, A., Suthakorn, J., Hofbauer, M., Gutierrez, A.L.: RoboCupRescue Robot League: Rules 2011.2. `http://www.nist.gov/el/isd/upload/Robocup_Rules_2011.pdf` (2011)
5. Nagatani, K., Kiribayashi, S., Okada, Y., Tadokoro, S., Nishimura, T., Yoshida, T., Koyanagi, E., Hada, Y.: Redesign of rescue mobile robot Quince: Toward emergency response to the nuclear accident at Fukushima Daiichi Nuclear Power Station on March 2011. In: International Workshop on Safety, Security and Rescue Robotics (2011)
6. RoboCup Rescue Robot League Organizing Committee: RoboCup Rescue Wiki. `http://wiki.robocup.org/wiki/Robot_League` (2011)
7. Schwertfeger, S., Jacoff, A., Scrapper, C., Pellenz, J., Kleiner, A.: Evaluation of Maps using Fixed Shapes: The Fiducial Map Metric. In: Proceedings of the 2010 Performance Metrics for Intelligent Systems Workshop (PerMIS'10), NIST Special Publication 1113, September (2010)
8. Sheh, R., Kimura, T., Mihankhah, E., Pellenz, J., Schwertfeger, S., Suthakorn, J.: The RoboCupRescue Robot League: Guiding Robots Towards Fieldable Capabilities. In: International Workshop on Advanced Robotics and Social Impacts (2011)
9. Virts, A., Jacoff, A., Downs, A.: RoboCupRescue Arena Assembly Guide 2011. `http://www.nist.gov/el/isd/upload/2011_Assembly_Guide.pdf` (2011)