

Profile of a Winner: Kansas State University

The Kansas State University (KSU) robotics programming team won the Find-the-Remote event. The team's software was able to find, recognize, and retrieve all six items used in the preliminary round. Because there was no other competitor for the final round, it was turned into a demonstration with four items found and retrieved.

The team developed its winning software program on WILLIE, a NOMAD200 robot from Nomadic Technologies, Inc. WILLIE is a black cylinder approximately 2 feet in diameter and 3-feet tall, which weighs about 200 pounds. It is equipped with three wheels. On board the NOMAD is a PENTIUM computer with a hard drive. The NOMAD used in the competition is equipped with 2 rings of 16 sonar sensors. For the contest, an arm and a color camera were purchased. The LINUX operating system is used on board the robot. The competition software was written in C++.

The team's success was based on its software-engineering approach. In the requirements phase, the team identified some critical issues. These issues included learning to use the arm and the color camera, which were both new to the students. Algorithms for line, edge, and ellipse detection and camera calibration were investigated. The issues of mapping the environment, path planning, and robotic motion in the environment were familiar from class exercises as well as previous competitions.

The team needed a robust architecture for the robot. A layered architecture based on abstractions of the tasks was chosen. There were three levels in the object model: First, the bottom layer interfaced with the arm, motors, and sensors. Second, the middle layer contained item recognition, path planning, and item manipulation. Third, the top layer controlled the overall strategy. Each command reported success or failure to the calling method in the higher level. The calling method would retry the method or call an alternative method as appropriate. This approach allowed recovery from many errors, such as misalignment with the table or failure to pick up an already identified item.

The robot maintained a metric map of the environment. A number of locations were designated as viewing stations. Because each item was constrained to be in a few locations, the possible viewing stations for each item were associated



Courtesy, Karen Gustafson

WILLIE Positioned to Pick Up the Small Green Cup.

with the item. When the robot was trying to retrieve a specified item, it looked up the item and determined the nearest viewing station associated with the item. It then moved to the viewing station. Before looking for the item, it checked its position relative to the nearest table and compared it with the map to readjust its global position. If it found the item at the site, it aligned with the item, picked up the item, and returned to the starting location. If the item was not found, it moved to the next-closest viewing station associated with the item. If it exhausted all the viewing stations associated with an item without finding the item, it would go on to the next item on the list.

For use of the camera in recognition of the items, many approaches were investigated. An identification approach was determined for each item. The approach used color or simple size and shape, such as diameter of the ellipse at the top of the cup to distinguish the small cup from the large cup. Items were recognized in two distinct ways: First, many items were recognized by color. The recognition of these

items was done with HSV (hue, saturation, luminosity) thresholding and numerous filters to remove noise. Second, these thresholded images were then evaluated for a silhouette of the proper color and dimensions of the desired item. A cylinder, such as a cup, was recognized by finding the ellipse on the top of the cylinder. Once the ellipse was detected, the three-dimensional (3D) space mapping was used to find the exact diameter of the cylinder and, thus, a good indication of the identity of the item.

Because the camera and the arm are on opposite sides of the robot, the issue of positioning the robot precisely so that the robot can turn and pick up an item was important. An edge-detection algorithm was written using Lowe's algorithm (Rosin and West 1995) to detect the major edge of the table. The camera was calibrated (Pratt 1991), and trigonometry was used to create a mapping from any (x, y) point in a camera image whose height is known to its corresponding (x, y, z) position in 3D space. The exact position of the camera was critical for the algorithm, and it took about two hours to recalibrate the camera every time it was moved.

The height of each table was then used to determine the position of the table edge in 3D space relative to the robot. The 3D space mapping was also used to determine the exact point in 3D space of the item so that the arm could locate and retrieve the item without having to take additional images when approaching the item. Then the robot could be positioned to pick up the item (see the figure) with an accuracy of about a half inch.

At the contest site, the actual map for the contest environment was made; adjustments were made for lighting and, in particular, for glare on the tables from the overhead lights; and the vision-recognition routines were specialized for the actual items.

The KSU team for the 1997 AAAI Mobile Robot Competition and Exhibition consisted of Mike Novak, Todd Prater, Brian Rectanus, and Steve Gustafson. David Gustafson was the adviser.

– David Gustafson