In Proceedings of SPIE AeroSense 2001.

Robotic Technologies for Outdoor Industrial Vehicles

Anthony Stentz
National Robotics Engineering Consortium
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15201
tony@cmu.edu

Abstract

The commercial industries of agriculture, mining, construction, and material handling employ a wide variety of mobile machines, including tractors, combines, Load-Haul-Dump vehicles, trucks, paving machines, fork trucks, and many more. Automation of these vehicles promises to improve productivity, reduce operational costs, and increase safety. Since the vehicles typically operate in difficult environments, under all weather conditions, and in the presence of people and other obstacles, reliable automation faces severe technical challenges. Furthermore, the viable technology solutions are constrained by cost considerations. Fortunately, due to the limited application domain, repetitive nature, and the utility of partial automation for most tasks, robotics technologies can have a profound impact on industrial vehicles. In this paper, we describe a technical approach developed at Carnegie Mellon University for automating mobile machines in several applications, including mass excavation, mining, and agriculture. The approach is introduced via case studies, and the results are presented.

1. Introduction

The opportunities for the automation of mobile equipment used in outdoor industrial applications are immense. Mobile machines are used in agriculture, surface mining, underground mining, quarrying, construction, material handling, and other applications. The machines navigate about, performing a wide range of operations, including digging, hauling, crushing, boring, planting, tilling, harvesting, baling, spraying, mowing, painting, stripping, paving, and grading. Automation promises to increase productivity, reduce operational costs, and improve safety.

Unfortunately, the automation of outdoor machines for industrial purposes is very difficult and poses great challenges. The work places are typically unstructured with the full richness of a natural setting. In many cases, the ground is uneven or otherwise treacherous. The terrain includes hazards such as potholes, ruts, thin branches, rocks occluded by tall grass, and steep slopes, which are difficult to detect but necessary for ensuring safe passage. The place may be uncontrolled in the sense that people, animals, and other machines can enter the environment. Rigging the environment with infrastructure to assist vehicle automation may be difficult or expensive; for example, an underground coal mine would require new infrastructure continually, as material is progressively removed. Weather and ground conditions can be highly variable and unpredictable; many applications require the machines to operate in rain, snow, fog, dust, day, night, etc. The sensors used for automation must be functional under these conditions and robust to the water, dirt, mud, vibration, and shock they encounter as the machine engages the environment. The range of tasks performed by one of these machines may be quite large, making full automation a difficult task at best. Finally, it is not enough for the automated machines to function—they must *perform*. The competition is the human-driven equivalent, so unless the automated system is faster, less expensive, or safer, it won't be cost effective. This competition is stiff, since even so-called unskilled labor is quite adept at the motor skills robotics seeks to mimic.

Fortunately, a number of factors work to our advantage. First, mobile machines employed in industrial tasks operate in a specific context. The machines navigate over terrain, encounter obstacles, and handle materials of known types or from a limited set of classes. This factor simplifies the perception, planning, and control problems. Second, the machines are likely to engage in repetitious tasks that are limited in scope and complexity. This factor reduces the need for sophisticated planning and offers the opportunity to learn from mistakes and fine-tune performance. Third, the machines need not be fully autonomous to add value. For many tasks, the proverbial 80/20 rule applies. Automation can handle the 80% of the task that is easy, and human intervention can cover the remaining 20% that is difficult. This arrangement can be cost effective if the overall productivity is increased or if the total time of human intervention per machine is reduced.

At Carnegie Mellon University's National Robotics Engineering Consortium, we have developed an approach that leverages these factors to introduce automation into outdoor vehicles. This paper discusses the technical elements of this approach via several case studies: 1) an autonomous hydraulic excavator for mass excavation; 2) a continuous mining machine for underground coal mining; and 3) an autonomous tractor for spraying, mowing, and other agricultural operations. We present these three case studies and draw conclusions.

2. Hydraulic Excavator for Mass Excavation

The digging machines used in surface mines, quarries, and construction projects are prime candidates for robotic automation. The rapid removal of material, called *mass excavation*, is a key operation for these applications. Digging machines such as wheel loaders and hydraulic excavators are used to load dirt, shot rock, and other materials into trucks for transportation to a processing plant or landfill. In one configuration (see Figure 1), an excavator sits atop a bench, digs soil from the face in front of it, and loads trucks that queue to the side of it. During the course of a shift, an operator typically loads hundreds of trucks. Because loading is the bottleneck of the operation, there is much to gain by increasing the machine's productivity.

Figure 1: Hydraulic excavator loading trucks at a construction site.

Mass excavation is performed under all weather conditions, so the sensors, computers, and other equipment must be robust to rain, snow, and ice. Additionally, loading and dumping creates considerable dust, which can obscure the sensor's field of view, and high forces, which can shock and vibrate the devices. To our advantage, the domain is limited to piles of rock, soil, machines, and people, so sensing geometric shape is key to automation. Additionally, the materials loaded are compliant and diffuse, so force sensing at the bucket is critical for guiding the digging process. The excavators move infrequently over groomed and well-understood terrain, so relatively simple obstacle detection suffices for safeguarding. The operation is largely repetitive. Depending on the relative sizes of the digging and hauling machines, each truck requires between three and six buckets for a full load, where each bucket requires between 15 and 20 seconds to dig and dump. In addition to digging and loading, the excavator operator is required to spot the truck, clean the floor between trucks, reposition on the bench, and deal with "hard spots" on the face. Taken together, these other operations consume a small percentage of the operator's time.

Automated excavation is an active field of research. Many researchers have automated the digging part of the operation, where a human operator selects the starting point for the dig and the control system takes over to complete the dig [6] [8] [14] [15] [17]. Others use active sensing to measure the terrain and automatically select the dig points [19] [24]. Still others sequence many digging operations [3] [9] [20]. We developed a fully autonomous, 25-ton hydraulic excavator for loading trucks with soft materials such as dirt [23]. The machine is the only one capable of end-to-end automation, starting with recognizing the truck and ending with the truck completely loaded. The machine uses a pair of scanning laser rangefinders to image the dig face, recognize the truck, detect obstacles, and monitor the digging and loading process (see Figure 2). The laser rangefinders are single-axis, scanning 360 degrees in a vertically oriented circle, acquiring 12kHz range samples.

Additionally, each scanner is mounted on a pan table, which allows the scan pattern to be swept left and right, thereby covering all space of interest. One scanner is used to image the dig face and the other the truck. Each laser rangefinder is a time-of-flight, last-pulse device, meaning that the sensor measures range by sending a pulse of laser light and recording the time of the last return. This configuration provides the sensor with a dust penetrating capability, since the early returns from multiple returns are likely to be reflections from the dust cloud. In our tests, the sensors were able to see through dust clouds about as well as a human. The excavator arm is equipped with force sensors to measure the load on the bucket during digging, specifically to determine when to curl the wrist and "break out" of the soil. The system is capable of loading a truck at nearly the speed of an expert human operator (for soft materials), and can detect obstacles, clean the floor, and reposition itself on the bench. To date, the machine has loaded hundreds of trucks fully autonomously.

Figure 2: CMU's computer-controlled excavator uses a pair of scanning laser rangefinders (black boxes) to recognize the truck, detect obstacles, monitor the dig face, and monitor soil in the truck bed.

The excavator is a successful prototype for a variety of reasons. The selected sensors take advantage of the limited domain. From load to load, the excavator has strong expectations about what is in its work space and where. The soil face, trucks, and rocks are easily detected and measured by the range sensors. Any objects not matching the expected profile are likely to be obstacles (such as people) and cause the machine to pause or shut down. Although heavy dust and other airborne obscurants can be problematic for the sensors, in general the devices can see at least as well as humans, thus making them a competitive alternative. The bucket loading cycle is a repetitive operation that varies slightly from pass to pass. Since the soil face and truck bed occupy relatively small spaces, the possible dig and dump points in the excavator's workspace comprise a fairly small set. Over time, the machine automatically learns the best sequence of control actions to move the bucket quickly between dig and dump points, thus minimizing load time and maximizing productivity. For our excavator, we were able to reduce the cycle time by 25% or more [13].

Even though the prototype is fully autonomous, it is unlikely to appear commercially as such, at least not initially. The excavator is autonomous for specific operations, not the full range of possible uses. It is able to load soft materials under some adverse weather conditions, but not all. It is able to detect most obstacle conditions, handle most trucks, and perform many ancillary operations such as floor cleaning, but the machine could not be left alone to work a complete shift without incident. In spite of these shortcomings, the machine could be a productive asset under semi-autonomous control. Since the dig/load cycle dominates the shift and is repetitious enough to be optimized, it is a prime candidate for automation. The operator could remain on the machine to monitor for safety, and to manually handle less frequent operations, provided the performance gain in the dig/load cycle more than offsets the cost of automation. Alternatively, a remote operator could oversee the operation of several machines via a wireless link. The operator could check the truck recognition results, monitor digging and loading for even results, verify plans for repositioning, and check the quality of floor cleaning, by viewing images of the scene. If the operation does not proceeding as desired, the remote operator could "steer it" by designating intermediate data such as the truck bed location and digging and dumping points.

3. Continuous Mining Machine for Underground Coal Mining

As with surface mining, mobile machines are used for the underground mining of hard materials, such as iron and copper ore, and soft materials, such as coal. For hard materials, the ore is shot with explosives, excavated by Load-Haul-Dump (LHD) machines, and transported out of the mine by haul trucks. For soft materials, shearing machines rip the material from the seam, and it is hauled via shuttle car or conveyer out of the mine. Room-and-pillar mining is a common way to extract soft materials from the seam, using an ensemble of machines to cut a lattice network. In this method, a continuous mining machine (CM) uses a rotating cutter head for shearing coal from the seam (see Figure 3). The sheared coal is deposited into a shuttle car or mobile conveyer belt. A roof bolting machine supports the roof by drilling holes and inserting bolts. The same ensemble is used to develop the support entries for another common method of underground mining: longwall. In this case, the entry development is the bottleneck, so the continuous mining machine is a good candidate for automation for both types of mining.

Since coal is flammable and mine ventilation is limited, internal combustion engines are infeasible and the machines are electrically driven. For recently manufactured machines, the operator controls the machine off-board using a “button box”. For safety and space reasons, the operator stands behind the machine. Due to darkness, dust in the air, and occlusion by the machine itself, the operator has difficulty ensuring that the right amount of coal is cut to match the capacity of the shuttle car and that the entry is cut straight. If entries are not straight, they must be re-cut or additional roof bolts must be inserted—both are a costly recourse. If the wrong amount of coal is cut, productivity suffers.

Figure 3: A continuous mining machine extracts coal from a seam. The cutter head spins to shear the coal from the face. The machine collects the coal and deposits it in a shuttle car or mobile conveyor that follows (courtesy of Joy Mining Machinery).

For the continuous miner, the rotating drum shears the coal in front of the machine as it travels forward, so the desired mine geometry is achieved by appropriately navigating the CM. Thus, knowing the global position of the machine is crucial. Global Positioning System (GPS) sensors cannot be used due to occlusion from satellites. Inertial sensors drift over time and accrue substantial error over the course of a mining shift. It is possible to install infrastructure, such as beacons, reflectors, or transmitters to assist in position estimation, but this method is inconvenient at best and impractical at worst. The problem is that the mine changes as material is removed, requiring new infrastructure. It is possible to add sensors to the CM itself to measure position, but the devices must be rugged, since the CM experiences high shock, constant vibration, and exposure to dirt, dust, and water.

Most research in underground mining automation has centered on hard rock mining using some form of infrastructure (e.g., beacons, light tubes) or inertial sensors [1] [7] [10] [16]. Early work with a laser rangefinder measuring the natural mine infrastructure showed that position estimation with this type of sensor was promising [18], but to date the devices have not been rugged enough to survive in the environment long term. We developed some onboard sensor components suitable for automating a CM [22]. The components are sump depth measurement (SDM) and global heading measurement (GHM). SDM measures the forward motion of the vehicle, primarily to determine how far the cutter head has driven into the coal face. SDM consists of one or more pairs of stereo cameras that look to the side at the coal ribs (walls), or upwards at the roof (see Figure 4). The stereo cameras measure the range to the nearest coal surface and convert the side-to-side image flow into machine translation and rotation. SDM provides the operator with a forward motion estimate that is accurate to 2% of the distance traveled. The second component, GHM, uses a roof-mounted laser that defines the desired heading for a new entry. The laser emits a vertically oriented plane of light that intersects a pair of planar targets on the CM (see Figure 4). A camera images the lines of laser light striking the planar targets and converts the data into a measurement of heading and lateral offset. The operator uses this information to cut a straight entry, correcting both for heading and lateral offset errors. GHM is accurate to a third of degree in heading and two centimeters in lateral offset. The two components were integrated and tested on a CM above the ground. The measurements were fed into the CM's control system and were used to sump to the proper depth and to steer straight, both autonomously.

Figure 4: Left photo: SDM stereo cameras mounted on the side of CMU's continuous miner. Right photo: GHM planar targets with imaging camera (small box with cable) mounted on top of the continuous miner.

The sensor components show great promise, and we believe they will see commercial usage. The components take advantage of the limited domain of the application. The high shock and vibration experienced by a CM rule out non-solid state sensors, such as mechanically scanned laser rangefinders. The low ambient light in an underground mine favors the use of solid state cameras, which above ground are subject to saturation from sunlight and intensity contrast from strong shadows. The components make zero use of *new* infrastructure—GHM uses a roof-mounted laser that surveyors currently emplace to help guide the machine. The components also take advantage of the limited scope of the core operation, the sump and shear cycle. Only three degrees of freedom for the CM need to be measured: relative motion forward, absolute heading and lateral offset. SDM and GHM use methods engineered specifically to measure just these parameters to high accuracy, rather than resorting to a more costly, and potentially less accurate, six degree-of-freedom measurement system.

Finally, the sensor components can add value to a CM short of full automation. SDM and GHM provide additional information to the operator to remove some of the guesswork, thus maximizing productivity and minimizing mistakes. By feeding the SDM/GHM data to the machine controller, we expect additional gains by reducing the delays that often occur under human control when the machine is sequenced from one operation to the next. The human operator can monitor the machine while it is cutting in autonomous mode, and then manually assume control for less frequent operations, such as turning a crosscut or back driving the machine to a new section of the mine.

4. Autonomous Tractor for Farm Operations

The tractor is the workhorse for the farm. When properly outfitted with implements, it can till, plant, weed, fertilize, spray, haul, mow, and harvest. The tractor needs to navigate properly, generally in a coverage pattern, to conduct the operation. For example, tractors spray pesticides on orange trees by pulling an airblast sprayer and driving between the rows of trees (see Figure 5). Tractors till a field by driving parallel swaths in the field to disc the soil. Two reasons for automating tractors and other agricultural machines are productivity and safety. Driving coverage patterns for an entire shift is fatiguing, and the result is that the operator tends to slow down. Some operations, such as the spraying of pesticides, are hazardous, both to the machine operator and to other workers in the field. Automated spraying can be conducted at night when there are no other workers in the field, eliminating the hazard to all parties.

High production farms are typically located in flat, open areas with full access to sunlight. With low probability of satellite occlusion, GPS is the most effective position estimation sensor. But since farms are less controlled and more accessible to the public than either surface or underground mines, vehicle safeguarding is a greater challenge. People, rocks, ditches, potholes, animals, other machines, water, and mud are of particular concern. The crop itself complicates the problem; it may occlude rocks and other hazards. The machine may need to drive over harvested plants but avoid standing plants. In general, differentiating between soft objects, such as vegetation, and hard objects, such as rocks, is a difficult task.

Figure 5: a manually-driven tractor pulls an airblast sprayer to apply pesticides to citrus trees.

Research into the automation of agricultural vehicles is widely pursued. Some researchers have looked at GPS- or radio-based techniques for guidance [4] [11], others have pursued computer vision approaches [2] [5] [21], and others have investigated a combination of vision with other techniques [12] [25]. Few agricultural researchers are investigating what we believe to be the hardest problem: safeguarding, the avoidance of human injury and damage to the machines and environment. We are developing an autonomous tractor that is addressing these issues. The tractor is a 90-horsepower, Deere Model 6410, equipped with a GPS receiver, heading gyroscope, Doppler radar, and four-wheel odometry for positioning and a pair of cameras for obstacle detection. The software provides a teach/playback capability, whereby the system records GPS waypoints as a human operator drives the tractor along a path. During playback, the tractor automatically re-drives the taught path. As the tractor drives autonomously, it examines the terrain in front for obstacles. The stereo cameras provide both range and appearance data which are fed into a neural network for classification. Due to the complexity of the obstacle detection problem, we believe that many types of sensor data are needed to discriminate obstacles from normal terrain conditions. Shape and appearance (from visible light) are two primary types, but others could be employed as well, including infrared (IR). When the tractor detects something that *may* be an obstacle, it stops and sends an image over a wireless link to a human supervisor. If the “obstacle” is a false alarm, the human authorizes the machine to continue; otherwise, he/she must attend to the problem. The system was tested for a pesticide spraying application. The tractor was transported to Florida, outfitted with an airblast sprayer, and tested in an orange grove (see Figure 6). The system drove about 7 km autonomously at speeds ranging from 5 to 8 kph. The obstacle detection system was evaluated in a variety of ways, including its ability to detect humans and classify grove vegetation as tree or grass. The results are promising, but much more work is needed in this area.

The prototype is promising for commercialization for several reasons. Good global positioning information for reliably and accurately executing coverage patterns is the key to successful agricultural operations. Due to the open nature of the terrain, GPS provides position estimates of sufficient accuracy, especially when combined with dead reckoning sensors (e.g., odometry and heading) to smooth the estimates and enable navigation to continue during data dropout. Agricultural operations, such as spraying, mowing, and harvesting, are repetitious in the sense that the entire field must be processed several times a year. This characteristic enables the scenario where a human operator can teach the system the entire task by performing it once and then re-playing the task multiple times. The task itself is narrow in focus: follow a sequence of GPS waypoints as accurately and reliably as possible.

The difficult part is safeguarding people, the machines, and the environment. This task is where human supervision can be most fruitful. When designing an obstacle detection system, there are two types of errors that can occur: the detection of obstacles that are not present (false positives) and the failure to detect obstacles that are present (false negatives). The former cause the machine to stop unnecessarily and therefore adversely affect productivity. The latter have worse consequences: they result in injury to people or damage to the machine. Attempts to minimize one type of error increase the chances that the other will occur. By keeping a human in the control loop for verification purposes, we can design the system to be very cautious, that is, to err toward false positives, provided a human can quickly intervene to confirm the false detection and resume driving. The wireless link enables this rapid and efficient intervention. Furthermore, since the time required to confirm or refute an obstacle is minimal, a single operator can oversee several machines.

5. Conclusions and Future Work

Although the automation of outdoor industrial vehicles is challenging, robotics is making steady inroads by capitalizing on the restricted domains, repetitious tasks, and opportunities for partially automated systems. Surface mining, underground mining, and agriculture are examples of large industries that will benefit immensely.

Figure 6: CMU's autonomous tractor sprays water in a grove in Florida. The tractor is guided by GPS with cameras for obstacle detection.

To continue making progress toward this objective, we need technical advances in three key areas. First, sensors and image processing algorithms continue to be crucial. Given the broad range of demands on such sensors, there is no single device that suffices for all purposes. We need to engineer sensors for particular applications, rely on multi-sensor systems with complementary failure modes, and fuse data to maximize accuracy and reliability. Second, given the variability in the machines' tasks and environments, it is unlikely that we can engineer them to be fully autonomous when they leave the factory. Instead, we must produce machines that can be trained in the field (by the customer) to meet the specific requirements of a set of tasks. Finally, human-machine interfaces are becoming very important. Human operators can play a role by timesharing control on multiple machines, assuming control for difficult tasks, and supervising operations at a variety of different levels. Suitable interfaces are needed for humans and computers to share control in a symbiotic and productive way.

Acknowledgements

The author would like to thank the research teams for the autonomous loading system project, the underground mining project, and the autonomous spraying project. The work described here was supported by Caterpillar Inc., Joy Mining Machinery, Deere & Company, and NASA under grant number NCC5-223.

References

1. Baiden, G., “Multiple LHD Teleoperation and Guidance at INCO Limited,” Proc. of the International Mining Conference, 1993.
2. Billingsley, J., Schoenfisch, M., “Vision Guidance of Agricultural Vehicles,” Autonomous Robots, Vol 2, 1995.
3. Bullock, D.M., Apte, S., Oppenheim, I.J., “Force and Geometry Constraints in Robot Excavation,” Proc. of Space 90: Engineering, Construction and Operations in Space, 1990.
4. Erbach, D.C., Choi, C.H., Noh, K., “Automated Guidance for Agricultural Tractors,” Automated Agriculture for the 21st Center, ASAE, 1991.
5. Gerrish, J.B., Fehr, B.W., Van Ee, G.R., Welch, D.P., “Self-Steering Tractor Guided by Computer Vision,” Applied Engineering in Agriculture, Vol. 13, No. 5, 1997.
6. Huang, X.D., Bernold, L.E., “Control Model for Robotic Backhoe Excavation and Obstacle Handling,” Proc. of Robotics for Challenging Environments, 1994.
7. Hurteau, R., St-Amant, M., Laperriere, Y., Chevrette, G., “Optical Guidance System for Underground Mine Vehicles,” Proc. of the International Conference on Robotics and Automation, 1992.
8. Lever, P., Wang, F., Chen, D., “Intelligent Excavator Control for a Lunar Mining System,” Proc. of ASCE Conference on Robotics for Challenging Environments, 1994.
9. Lever, P. J., Ntaimo, L., "Autonomous Digging Control for a Wheel Loader: Implementation and Results," Proc. of SPIE AeroSense, April 2000.
10. Makela, H., Lehtinen, H., Rintanen, K., Koskinen, K., “Navigation System for LHD Machines,” Intelligent Autonomous Vehicles, 1995.
11. O’Conner, M., Bell, T., Elkaim, G., Parkinson, B., “Automatic Steering of Farm Vehicles Using GPS,” Proc. of the 3rd International Cnference on Precision Agriculture, 1996.
12. Pilarski, T., Happold, M., Pangels, H., Ollis, M., Fitzpatrick, K., Stentz, A., “The DEMETER System for Autonomous Harvesting,” Proc. of the 8th International Topical Meeting on Robotics and Remote Systems, April 1999.
13. Rowe, P., “Adaptive Motion Planning for Autonomous Mass Excavation,” Carnegie Mellon Robotics Institute doctoral dissertation, Technical Report CMU-RI-TR-99-09, January 1999.
14. Salcudean, S.E., Tafazoli, S., Lawrence, P.D., Chau, I., “Impedance Control of a Teleoperated Mini Excavator,” Proc. of the 8th IEEE International Conference on Advanced Robotics (ICAR), July 1997.
15. Sameshima, M. Tozawa, S., “Development of Auto Digging Controller for Construction Machine by Fuzzy Logic Control,” Proc. of Conference Japanese Society of Mechanical Engineers, 1992.
16. Scheding, S., Nebot, E.M., Stevens, M., Durrant-Whyte, H., Roberts, J., Corke, P., Cunningham, J., Cook, B., “Experiments in Autonomous Underground Guidance,” Proc. of the International Conference on Robotics and Automation, 1997.
17. Seward, D., Bradley, D., Mann, J. Goodwin, M., “Controlling an Intelligent Excavator for Autonomous Digging in Difficult Ground,” Proc. of 9th International Symposium on Automation and Construction, June 1992.
18. Shaffer, G., Stentz, A., “A Robotic System for Underground Coal Mining,” Proc. of the International Conference on Robotics and Automation, 1992.
19. Singh, S., “Synthesis of Tactical Plans for Robotics Excavation,” Ph.D. Thesis, Robotics Institute, Carnegie Mellon University, January 1995.
20. Singh, S., Multi-Resolution Planning for Earthmoving, Proc. of the International Conference on Robotics and Automation, May 1998.
21. Southall, B., Hague, T., Marchant, J.A., Buxton, B.F., “Vision-aided Outdoor Navigation of an Autonomous Horticultural Vehicle,” Proc. of the First International Conference on Vision Systems (ICVS), 1999.
22. Stentz, A., Ollis, M., Scheding, S., Herman, H., Fromme, C., Pedersen, J., Hegadorn, T., McCall, B., Bares, J., Moore, R., “Position Measurement for Automated Mining Machinery,” Proc. of the 1999 International Conference on Field and Service Robotics,” August 1999.
23. Stentz, A., Bares, J., Singh, S., Rowe, R., “A Robotic Excavator for Autonomous Truck Loading,” Autonomous Robots, Vol. 7, No. 2, September 1999.
24. Takahashi, H., Kamata, H., Masuyama, T., Sarata, S., “Autonomous Shoveling of Rocks by Using Image Vision System on LHD,” Proc. of the International Symposium on Mine Mechanization and Automation, June 1995.
25. Zhang, Q., Reid, J.F., Noguchi, N., “Agricultural Vehicle Navigation Using Multiple Guidance Sensors,” Proc. of the International Conference on Field and Service Robotics, August 1999.