VisLab and the Evolution of Vision-Based UGVs

Massimo Bertozzi, Alberto Broggi, and Alessandra Fascioli
VisLab, Università degli Studi di Parma

Unmanned ground vehicles will shape our future by providing enhanced safety and improved mobility. Four decades after researchers first developed the concept, they're still building prototypes, but are now moving toward demonstrations, indicating that the early ideas have real applicability.

Unmanned ground vehicle technology has evolved in fits and starts. Researchers first pondered the idea in the 1960s, but the technology wasn't mature then, and it wasn't until the mid-1980s that the military developed a UGV prototype intended to help automate its ground fleet. Researchers' interest in the civilian sector picked up in the late 1980s after governments worldwide launched their first projects. And in the late 1990s, the automotive industry jumped aboard after the development and road testing of the first autonomous vehicles.

We're now seeing a flurry of intelligent-vehicle research initiatives. However, technological issues and legal considerations associated with fully automated vehicles have prompted the automotive industry to focus more on supervised systems and advanced driver assistance systems (ADAS).

At the same time, research on UGVs has slowed, since industry and governments no longer view UGVs as a primary strategic area of investment. Transportation departments worldwide are concerned with social, economic, or environmental objectives aimed at enhancing fuel and road network efficiency and quality of life.

Recently, the automotive industry's success with ADAS has induced the military to reconsider its ground-fleet-automation goal. DARPA took a considerable step forward with its 2004 and 2005 Grand Challenges (www.grandchallenge.org), which drew entrants from top-level research institutes competing for seven-figure prizes.

Researchers are considering unmanned-vehicle technology for many other applications. Ever-increasing personnel costs are prompting consideration of the technology for agricultural, demining, rescue, and other dangerous applications. There would be a decrease in the number of individuals put at risk and an increase in operational efficiency if a vehicle could move autonomously, plant seeds, enter minefields, or perform dangerous missions. However, most common and attracting the most industry interest is the automation of road vehicles.

PRECOMPETITIVE RESEARCH

On the heels of military organizations' first experiments with ground and aerial applications, academic and private research centers started a precompetitive research stage for intelligent-transportation systems (ITS). National and international organizations launched programs to reduce traffic congestion and increase safety. Among the initiatives were

- the US's Mobility 2000 and Automated Highway System,
- Japan's Road/Automobile Communication System (RACS) and Advanced Mobile Traffic Information and Control System (Amtics), and
- Europe's Drive and Europe-Wide Network for Market-Oriented Research and Development (Eureka) programs.

0018-9162/06/$20.00

Figure 1. Vamp prototype. Universität der Bundeswehr developed Vamp in 1995 with an automatic trip from Munich, Germany, to Odense, Denmark.

The real-time requirements of systems installed onboard vehicles forced most research groups to develop custom architectures by assembling off-the-shelf components or starting from silicon. Devices like radar or acoustic sensors capable of direct measurement were a common choice since they acquire small amounts of data and don't require high computational capabilities. Researchers also considered vision sensors, which had the advantage of providing a rich description of the environment without needing specific road infrastructure. However, since vision sensors acquire a large amount of data, they needed a complex processing phase, so this was indeed a challenging choice.

Lessons from early projects

The Università degli Studi di Parma has been working on UGV research for 15 years through its Vision and Intelligent Systems Lab (VisLab; www.vislab.it), closely following worldwide developments in the field and setting milestones in the history of intelligent vehicles.[1]

VisLab developed the Mobile Laboratory (Mob-Lab) vehicle, within the framework of the Program for European Traffic with Highest Efficiency and Unprecedented Safety (Prometheus), part of the Eureka program that Europe's automotive and automotive-supplier industry carried out between 1989 and 1994.

The Mob-Lab prototype, a Fiat Ducato 18 Maxi, was suited to study, develop, and test real-time intelligent systems that featured computer vision. Black-and-white and color analog cameras installed in the vehicle's front and back analyzed the external environment.

We developed a custom architecture based on the single instruction, multiple data (SIMD) paradigm. This system, named Parallel Processor for Image Checking and Analysis (Paprica), interfaced directly to a stereo-vision system that could directly display processing results on an external monitor. Moreover, Paprica featured hardware support for pyramidal-based image processing and was therefore suitable for multiresolution image processing.

We used Mob-Lab to test early solutions and algorithms conceived for two basic functionalities, lane and obstacle detection. Lane detection used a lane model at different resolutions to exploit the Paprica architecture's pyramidal capabilities. Similarly, we used a rectangular-shaped template for obstacle detection.

These initial results, while embryonic, demonstrated that using vision for intelligent vehicles was a viable path. Dedicated hardware enabled the two detection systems to run in real time. In both cases, we processed gray-level images and tested color images, but concluded the setup was unaffordable because it required processing a larger amount of data.

Initially, we tested algorithms on the massively parallel Connection Machine CM-2, and then ported them onto the Paprica architecture. In addition, the use of a dedicated architecture required us to develop an operating environment consisting of code assembler, debugger, and graphic libraries.

The Paprica system demonstrated that researchers could use custom architectures to solve the inadequacies of computers available at that time. The SIMD paradigm implemented on Paprica was a successful choice, with Intel later adopting a similar solution: the MMX extensions, a SIMD extension for image-and-sound processing.

PIONEER PROJECTS

Given the promising perspectives that the precompetitive phase generated, a few research centers embraced pioneer projects aimed at developing concept vehicles that researchers could later transform into intelligent-vehicle prototypes. Although the scientific community was becoming aware of ITS's potential, the automotive industry had not developed a similar interest. Single groups or institutions willing to invest in high-risk research were advancing the work. Two examples of such pioneer projects are the Vamp prototype,[2,3] shown in Figure 1, that Germany's Universität der Bundeswehr demonstrated and Carnegie Mellon University's NavLab vehicle,[4] shown in Figure 2.

Many of these projects used machine vision as the main sensing device. The technological solutions available for the processing system differed: Some research groups chose general-purpose hardware or customized processing engines, while others adopted a mixed solution. The main challenge these projects faced was implementing and field-testing prototype vehicles equipped with vision-based driving-assistance systems.

VisLab carried out two pioneer projects in the late 1990s: ARGO and the Surface Antarctic Robot (RAS).

The ARGO project

ARGO[5] was one of the first prototypes to demonstrate autonomous driving capabilities on public roads. It integrated the main results of the preceding few years' research conducted on the algorithms and the architectures for vision-based automatic road-vehicle guidance.

We had been developing, testing, and tuning several solutions for autonomous navigation, particularly for the basic obstacle- and lane-detection functionalities.

We integrated the most promising approaches for both functionalities into the Generic Obstacle and Lane Detection (GOLD) system,[6] which acted as ARGO's automatic driver. Stereo vision detected and localized obstacles on the road, while the processing of a single monocular image allowed extraction of road geometry in front of the vehicle. The two functionalities shared the same underlying approach: inverse-perspective mapping (IPM) to remove the perspective effect from the acquired image. Lane detection was performed through extraction of specific image features, followed by the use of lane-marking models. Applying IPM to stereo images, in conjunction with a priori knowledge of the road shape, let researchers compute the free space and detect generic obstacles through a simple match with a model representing the environment without obstacles.

Figure 2. NavLab prototype. Developed by Carnegie Mellon University, NavLab traveled in automatic mode across the US in 1995.

Figure 3. ARGO prototype. VisLab tested ARGO during the thousand-mile MilleMiglia in Automatico Tour in 1998.

Figure 3 shows ARGO, a passenger car equipped with a vision system composed of two low-cost, black-and-white cameras mounted on the interior top corners of the windshield. The system fed the processing result to the driver through a set of output devices providing acoustical warnings, visual feedback, and steering wheel actuation.

The ARGO prototype was used to test ITS concepts such as:

- *Manual driving*. The system monitored driver activity, using acoustic and optic signals to warn the driver of dangerous situations.
- *Supervised driving*. The system took control of the vehicle in dangerous conditions to keep it in a safe condition.
- *Automatic driving*. The system drove automatically, following lanes, localizing obstacles, and performing lane changes.

Since the main challenge was fielding the vehicle in a real environment, researchers assessed the vehicle's autonomous capabilities during the 1998 thousand-mile MilleMiglia in Automatico Tour. In this test, ARGO drove itself autonomously along the Italian highway network, passing through flat and hilly regions, including viaducts and tunnels, and quickly varying road scenarios with changing weather conditions and intense traffic. The system demonstrated high robustness with respect to horizontal road signs, guardrails, forks, junctions, highway exits, and heavy traffic conditions. Moreover, high temperatures, different light conditions, and high speeds didn't influence the system's stability and robustness, for either hardware or software.

The RAS project

RAS, one of the first efforts toward automating driving functionalities in extreme environments, tested a snowcat's automatic maneuvering performance during Italian scientific missions to the South Pole.

The project aimed to apply artificial vision to the autonomous driving of a platoon of snowcats used to transport people and goods in Antarctica. The final goal was to develop a vehicle that could automatically follow a leader vehicle. Since ice cracks can make it dangerous to deviate even slightly from the driving path, researchers selected vision as the sensing capability that could deliver the most precise performance in localizing tracks the previous vehicle left.

Antarctica presents extreme environmental conditions that make this application challenging and far different from driving unmanned vehicles on highways or urban roads. We considered several approaches and developed specific filters to cope with problems typical of this environment: very low temperatures, possible snow-

storms, high brightness, low contrast, sharp shadows, direct sunlight and reflections, absence of ground reference, or unknown terrain slope.

The automotive research group's experience suggested the approach we used to solve the artificial-vision problem. We preferred a simple method that we could easily use with the snowcat's simple processing engine.

Following the approach that drove development of the pioneer prototypes, we considered alternative strategies to evaluate different algorithmic solutions. The first solution we investigated was based on traditional feature-extraction techniques and exploited available libraries and previously developed experience. We also tested an innovative approach based on an evolutionary technique. The second solution confirmed the applicability to vision for vehicles of this emerging technique, which still constitutes a hot research topic.

> **Full automation will require governments to address legal and liability issues.**

We demonstrated the vehicle on a test site in the Italian Alps and in real environmental conditions in Antarctica. Track detection succeeded in different situations even in the presence of noisy or critical conditions such as shadows, sun reflections, unknown terrain slope, and dark objects.

Real-world lessons

At the time VisLab developed these projects, intelligent prototypes commanded the interest of only a few far-seeing groups. For example, VisLab alone conceived, designed, implemented, and financed ARGO. On the other hand, other research centers invested a large amount of money in innovative and risky research on unexplored topics.

Environmental noninvasiveness dictated the decision to use only passive sensors to perceive the surroundings in both the road and Antarctic environments.

Because we envisioned deploying ADAS in common road vehicles, sensor cost was a main concern in the ARGO project.

We chose low-cost standard processing hardware for these projects, completing the entire MilleMiglia in Automatico Tour with a 200-MHz Intel Pentium-based system. We believe that these and other projects of a similar age taught the scientific community that it could begin considering commercial hardware to perform vision-based automotive tasks.

We developed stand-alone software applications aimed at realizing specific driving functionalities for these projects, basing the software on three main modules: data acquisition, image processing for feature detection (lane markings, obstacles and vehicles, tracks), and output to devices for warning the driver or actuators. These applications represent the core from which a complex software framework subsequently evolved.

The ARGO and RAS vehicles provided two of the first prototypes with intelligent capabilities to be demonstrated in the field. Confronting the real world, the vehicles faced problems such as reflections and road construction, temperature variations, interference from real vehicles, and human-machine interface issues. Becoming aware of such problems and solving some of them were the test's main accomplishments.

SUPERVISED SYSTEMS

Following the outcome of the precompetitive phase and the achievements of the pioneer projects, and as sensors and processors evolved, the automotive industry developed an interest in intelligent vehicles. Because full automation is highly complex, researchers continue to pursue initiatives aimed at developing supervised, rather than fully automated, systems. Car manufacturers are pushing in the ADAS direction in an effort to reach this new market early.

Full automation will require governments to address legal and liability issues. In addition, complete automation of civil transportation requires a new infrastructure and a thorough redesign of mobility. But governmental agencies have a high interest now in road safety, environmental protection, and sustainable mobility. For instance, the eSafety initiative that the European Commission launched in 2002 aims to reduce the number of road fatalities in half and have driving-assistance systems in one-fifth of cars by 2010. Governments are also trying to regulate and standardize such systems.

Some supervised systems are already commercialized or close to market. These systems address the simplest aspects of driving. They're based on the most elementary or settled sensing technologies, or a combination of the two, and feature basic output. Examples are maneuver assistants, blind-spot detectors, lane-departure warnings, and enhanced night vision. Conversely, researchers must further assess the most complex systems, such as protecting vulnerable road users or stop-and-go driving.

Research now focuses on three main streams:

- applying emerging technologies that are available at an affordable cost, such as infrared cameras for night and day vision or laser scanners for robust all-weather obstacle detection;
- using multiple sensors with complementary characteristics and capabilities and fusing their data to obtain a more reliable system; and
- engineering systems that have been previously tested and demonstrated as effective.

Developing robust and reliable prototypes that can be transformed into products remains a challenge for

researchers. For example, products need quick and automatic sensor recalibration during system operation.

Current VisLab research

VisLab developed several research prototypes based on different technologies in collaboration with automotive partners and research centers.

Vulnerable road users' protection. One research area that already appeared promising at the beginning of this decade is protecting vulnerable road users. In this field, VisLab has been collaborating with the Volkswagen research center to develop a prototype pedestrian-detection system for precrash or driver-assistance applications.

We used thermal infrared as the main sensing device, since it seems promising for pedestrian applications. However, we've also investigated near-infrared, along with specific illuminators, because they're less expensive. We've considered fusing the system with radar to add strength.

The US Department of Defense plans to automate one-third of ground military forces by 2015.

Preventive safety. Europe's APALACI-PReVENT project aims to develop and demonstrate preventive applications and technologies to improve road safety. In this framework, VisLab is working with Volvo to develop a start-inhibit system for large trucks. The system uses stereo vision to detect the presence of pedestrians or obstacles in the forward blind spot, warning the driver and preventing the vehicle from taking off.

Within the same European project, we're also collaborating with the Fiat research center to develop a road-obstacle-classification system based on the fusion of radar and monocular vision. The system is aimed at classifying vehicles, guardrails, and pedestrians.

Enhanced vision. A further collaboration involves Hella, a large German automotive supplier, dealing with using near-infrared headlights and cameras for obstacle localization.

The automotive industry is also demonstrating interest in driver-assistance applications based on color vision. Although color entails the analysis of a larger quantity of raw data, current digital signal processing technology makes it viable. As an example, we're working for a large Italian automotive supplier, Magneti Marelli, on a project using color vision for road-sign recognition.

Military applications. Leading-edge technologies are feasible for military applications, in which expense is less of an issue. For example, VisLab collaborated with the vetronics group of the US Army's Tank Automotive Research, Development and Engineering Center (TARDEC) to deploy a four-camera system, named tetravision, to improve safety in robotic vehicles. Initially developed as a daylight stereo-vision system to localize human shapes, tetravision now features simultaneous use of far-infrared and visible-camera stereo pairs. The main idea is to exploit the advantages of both far-infrared and visible cameras.

We're also designing another night vision application for military operations for Oshkosh Truck Corp.

According to the general trend, we applied the fusion of multiple sensors in all these investigations to obtain robust systems. All these projects generally share the need for a real-time software system that can acquire data from different vehicle sensors and perform playback in the laboratory for algorithm development. Each project needs specific data acquisition procedures.

We've developed a complex software framework featured by several acquisition modules tailored to different devices: the vehicle CAN bus, a variety of cameras, other sensors such as radars and laser scanners, and a network file system and disk acquisition for laboratory postprocessing. We integrated the driving-functionality applications into this framework as plug-ins, allowing the system to be used as a laboratory-development environment for the algorithms and as the ADAS software engine on the prototype vehicle.

We exploited multithreaded processing to boost performance overlapping data acquisition, processing, and output. We used graphical software libraries and hardware acceleration to enhance the visual output. This software framework provides an easy tool for use in prototyping many applications in the automotive domain and has provided the base for all the current and future projects in our laboratory.

A NEW PUSH TOWARD UGVS

As automakers deploy the first supervised systems, the military is demonstrating a renewed interest in autonomous vehicles. For example, the US Department of Defense plans to automate one-third of ground military forces by 2015. In addition, DARPA organized the 2004 and 2005 Grand Challenges, two long races for autonomous ground vehicles, and, in 2005, the German Federal Armed Forces held the Land-Robot Trial to demonstrate autonomous or semiautonomous vehicles.

DARPA offered a $1 million prize to the winner of the 2004 Grand Challenge and $2 million to the 2005 winner to boost the development and fielding of autonomous robotic vehicles. The price was a winning idea, since much of the development was auto-financed or externally sponsored. DARPA's investment paid off in the teams' results. The 2004 race drew 100 teams, and the 2005 competition attracted 195. Moreover, DARPA let teams select their own sensing devices and technology, allowing the exploration of various solutions. DARPA's evolutionary approach let entrants test

different solutions so that the most effective approach won.

Taking place in the inhospitable Mojave Desert and requiring computing engines that could operate continuously without human intervention, the races forced teams to address reliability and performance in extreme environments. In particular, the races required systems that could perform in high temperatures and dusty conditions. Teams selected reliable sensors, placing them in positions where accidents or rocks were less likely to damage them. The systems also required cleaning devices to keep sensors free of dust, water, mud, and other debris.

VisLab at the Grand Challenge

Claudio Caraffi and Stefano Cattani
VisLab, Università degli Studi di Parma

VisLab, together with Oshkosh Truck Corp. and Rockwell Collins, developed and successfully fielded TerraMax, an autonomous vehicle that reached the finish line of the 2005 DARPA Grand Challenge. TerraMax uses artificial vision, laser scanners, GPS, inertial sensors, and map databases to sense and understand its environment.[1] VisLab developed its vision system for obstacle detection and drivable path sensing.

TerraMax's three-camera system allows precision and efficient computation at a wide range of viewing distances. The cameras sit on a rigid bar over the vehicle hood. By selecting two cameras at a time, the system can get stereo pairs with different baselines—that is, intercamera distances.

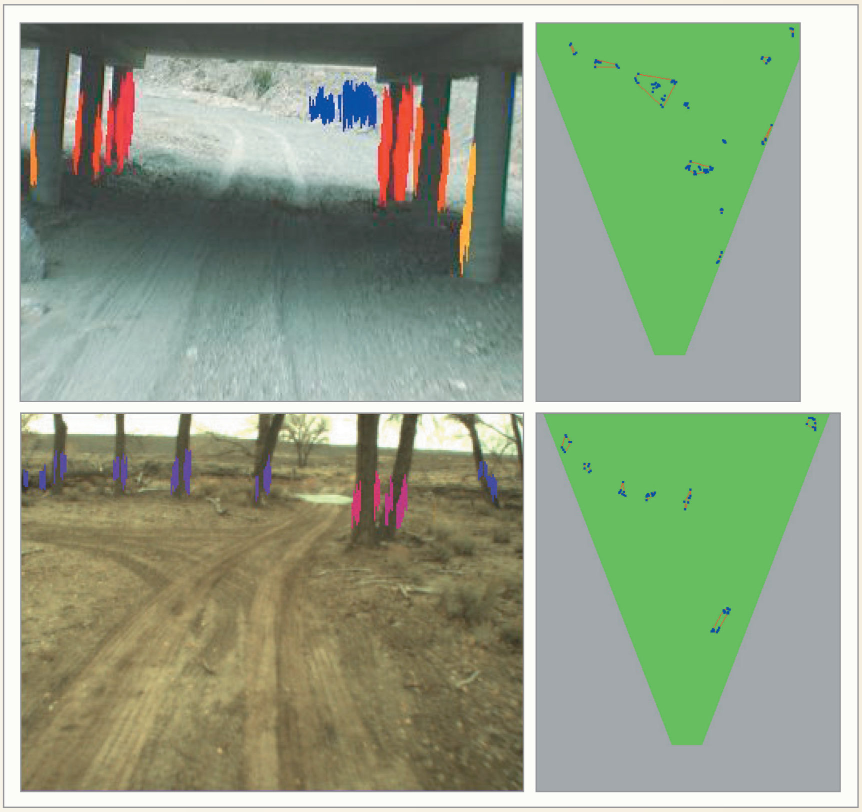

Figure A. Obstacle detection for different scenarios. The color of the obstacles varies with distance. The images on the right map the obstacles in a bird's eye view. The camera's field of view is plotted in green.

During the DARPA Grand Challenge, TerraMax selected the baseline based on vehicle speed. Higher speeds required greater sensing distances and thus wider baselines. We developed image stabilization to overcome vehicle oscillations from off-road environment terrain bumps.

The system relies on a two-step approach for full 3D stereo reconstruction. The first step exploits V-disparity image properties to stabilize images. The V-disparity images are 3D graphical representations of the similarity measures between left and right image rows depending on the shifts used to compare them.[2] V-disparity images contain basic information about the ground's position and shape, taking slope changes into consideration. In this way, the system can estimate the vehicle's pitch at the time of acquisition and stabilize images.

As a second step, using the information about the ground, the system addresses the correspondence problem by computing a disparity search image (DSI). The correspondence problem is the process of finding which pixel in the left image matches which pixel in the right image of a stereo image pair. We compute the DSI using small confrontation windows to allow detection of thin obstacles that other sensors missed. We apply a series of filters to the DSI, highlighting disparity concentrations that are detected as obstacles. We compute the obstacle 3D world coordinates via stereo triangulation. Figure A shows an example of this process.

Code optimizations exploited the processors' MMX and Streaming SIMD Extensions (SSE) instruction set. During the DARPA Grand Challenge, obstacle detection ran with a guaranteed 15-Hz throughput, although it could perform the entire computation within 30 milliseconds on a Pentium IV using a 2.8-GHz processor system, allowing other image-processing routines to run on the same machine.

Along with obstacles, the vision system also provides drivable path

Five prototypes completed the 2005 race. The winning team, Stanford University's "Stanley," mainly used laser-scanner and GPS sensors. H1ghlander and Sandstorm employed a similar approach and technology; in addition to the information coming from onboard sensors, they also exploited a precise premission-planning strategy. The Kat-5 vehicle used lasers as primary sensors as well.

As the "VisLab at the Grand Challenge" sidebar describes, the TerraMax prototype—a Medium Tactical Vehicle Replacement (MTVR) truck that Oshkosh Truck Corp., Rockwell Collins, and VisLab developed—

information. Most of the algorithms used for path detection look for a single homogeneous road surface in front of the vehicle. Since the hypothesis of homogeneity becomes a huge limitation because it bounds the set of detectable roads to the case of medium/well-structured environments, we tried to generalize the problem, considering that roads can also be made of patches of heterogeneous surfaces. To find these potentially heterogeneous surfaces, the algorithm looks for a variable number of small homogeneous terrain portions. They can represent any kind of natural or artificial environment elements, such as gravel or paved roads, grass, water puddles, oil stains, drivable rocks, lane markers, and shadows. Consequently, it's possible to summarize the path-detection algorithm as a two-step process:

1. Divide the image in homogeneous regions made of connected pixels.
2. Decide which combination of the obtained regions could represent the road surface with the highest probability.

Researchers have successfully studied the first step, clustering, using both evolutionary and traditional approaches. However, the Grand Challenge's real-time constraints led to the adoption of a simple—but fast and easily tunable—clustering algorithm as a good tradeoff between performance and computational requirements.

The second step falls into the class of decision problems. The decision process we developed tries to minimize the risk of wrong classifications, taking into account the current vehicle state. The underlying idea is that each cluster belongs to the road with a given probability depending only on its own intrinsic properties: homogeneity, size, shape factors, and covered free-space area. Figure B shows the clusters classified as road and off-road. However, a high probability of being road isn't sufficient, and sometimes not even necessary, to be finally classified as road. In fact, a sensible road classifier's principal goal is to minimize the risk associated with an incorrect classification on the basis of the current vehicle state. The following rule applies to classification decisions: Those with higher risks (requiring sudden changes in vehicle behavior) need higher probabilities of being correct before they can be assigned. Decision networks[3] extend Bayesian networks and provide a general methodology for rational decision making that fits the problem of deciding about the set of clusters that belongs to the road surface.

Path detection provided results in accordance with the output of the laser scanner and extended the perception range beyond the limits of the laser scanner, which is affected by vehicle pitch.

References

1. D. Braid, A. Broggi, and G. Schmiedel, "The TerraMax Autonomous Vehicle," *J. Field Robotics*, vol. 23, no. 9, 2006, pp. 693-708.
2. A. Broggi et al., "The Single-Frame Stereo Vision System for Reliable Obstacle Detection Used During the 2005 DARPA Grand Challenge on TerraMax," *Proc. Intelligent Transportation System Conf.* (ITSC 06), IEEE Press, 2006, pp. 745-752.
3. A. Broggi et al., "A Decision Network-Based Framework for Visual Off-Road Path-Detection Problem," *Proc. Intelligent Transportation System Conf.* (ITSC 06), IEEE Press, 2006, pp. 951-956.

Claudio Caraffi *is a PhD student in the Department of Information Technology at Università degli Studi di Parma. Contact him at caraffi@ce.unipr.it.*

Stefano Cattani *is a PhD student in the Department of Information Technology at Università degli Studi di Parma. Contact him at cattani@ce.unipr.it.*

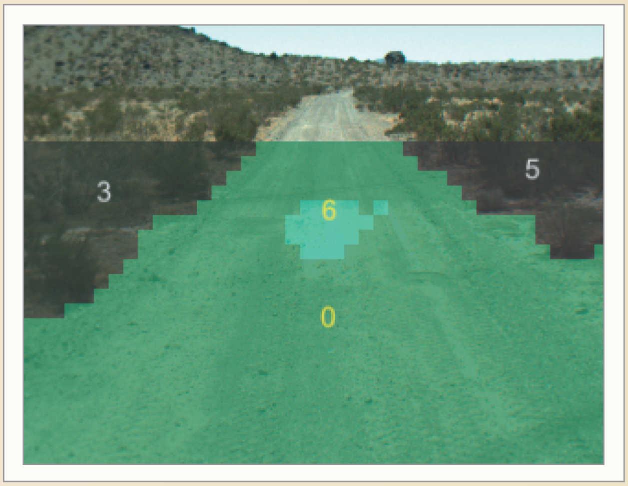

Figure B. White numbers denote off-road clusters, and yellow numbers denote road clusters. The interest area is limited to the region within 50 meters from the vehicle. This information is obtained from the preliminary processing of the obstacle-detection module.

Figure 4. TerraMax concludes the 2005 DARPA Grand Challenge.

reached the finish line relying primarily on vision. TerraMax, shown in Figure 4, featured an innovative trinocular vision system with different intercamera distances that behaved better than approaches that relied on moving camera heads.

During the race, TerraMax reached a maximum speed of 68 km/h—an impressive performance not only because of the truck's size and weight, but also because it avoided obstacles, recognized paths, and planned trajectories at this speed. In compliance with race rules, the vehicle engine and sensors suite remained operational while TerraMax paused during the night. Due to DARPA having to remove preceding vehicles from the course, TerraMax ran the race in two days and was the only vehicle that spent the night in the desert. The following morning, the vehicle proceeded, reaching the finish line after 28 hours of continuous activity. The prolonged operational time demonstrated the entire system's robustness and reliability.

The DARPA Grand Challenge 2005's success led to the agency launching an even more challenging competition: the November 2007 Urban Challenge. In this competition, fully autonomous vehicles must complete a 60-mile-long race in an urban environment in less than six hours, obeying road rules and negotiating traffic. VisLab will again contribute to the worldwide research efforts by fielding a vision system for a new vehicle.

Meeting the goals of the 2007 Urban Challenge will further open the road to a technological transfer from the military field to the car market. We still need legislation that takes into account liabilities and rules for automatically driven vehicles. Standardization is also important since UGVs must communicate with one another and will be part of a network with intelligent infrastructures.

The Urban Challenge will provide some feel for how long it will be before we sit in our own automatic cars. For certain, positive results will stimulate the automotive industry to invest more in automatic driving products. ■

References

1. M. Bertozzi et al., "Artificial Vision in Road Vehicles," *Proc. IEEE*, special issue on technology and tools for visual perception, July 2002, pp. 1258-1271.
2. V. Graefe and K.-D. Kuhnert, "Vision-Based Autonomous Road Vehicles," *Vision-Based Vehicle Guidance*, I. Masaki, ed., Springer-Verlag, 1991, pp. 1-29.
3. E.D. Dickmanns, "The Development of Machine Vision for Road Vehicles in the Last Decade," *Proc. IEEE Intelligent Vehicles Symp.*, vol. 1, 2002, pp. 268-281.
4. C.E. Thorpe, ed., *Vision and Navigation: The Carnegie Mellon Navlab*, Springer, 1990.
5. A. Broggi et al., *Automatic Vehicle Guidance: The Experience of the ARGO Vehicle*, World Scientific, 1999.
6. M. Bertozzi and A. Broggi, "Vision-Based Vehicle Guidance," *Computer*, July 1997, pp. 49-55.

***Massimo Bertozzi** is an assistant professor in the Department of Information Technology at Università degli Studi di Parma and a researcher on the ARGO project. His research interests are applying image processing to real-time systems and vehicle guidance, optimizing machine code at assembly level, and parallel and distributed computing. Bertozzi received a PhD in information technology from Università di Parma. Contact him at bertozzi@ce.unipr.it.*

***Alberto Broggi** is a professor in the Department of Information Technology at Università degli Studi di Parma and the founder of VisLab. His research interest is artificial vision techniques in automotive applications. Broggi received a PhD in information technology from Università di Parma. He is a member of the IEEE. Contact him at broggi@ce.unipr.it.*

***Alessandra Fascioli** is a researcher in the Department of Information Technology at Università degli Studi di Parma. Her research interests are real-time computer vision and computer architectures for automotive applications. Fascioli received a PhD in information technology from Università di Parma. Contact her at fascal@ce.unipr.it.*