

ARGO and the MilleMiglia in Automatico Tour

Alberto Broggi, University of Pavia, Italy
Massimo Bertozzi and Alessandra Fascioli, University of Parma, Italy

ARGO IS THE EXPERIMENTAL autonomous vehicle developed at the Department of Information Engineering of the University of Parma, Italy. It is a passenger car with a vision-based system for extracting road and environmental information from the acquired images, using different output devices to test the automatic features. ARGO integrates the main results of our last few years' research on algorithms and architectures for vision-based automatic road vehicle guidance. (Its name is based on the name of the ship of the mythical Greek hero Jason and on Argus, the many-eyed Greek god.)

Thanks to ARGO's availability, we've developed, tested, and tuned several different approaches to autonomous navigation, particularly for obstacle and lane detection. We've integrated the most promising approaches into the GOLD (Generic Obstacle and Lane Detection) system,¹ which acts as ARGO's automatic driver. GOLD's primary task is road following. It uses stereo vision to detect and localize obstacles on the road, and processes a single monocular image to extract the road geometry in front of the vehicle. This article focuses on lane detection; we describe obstacle detection elsewhere.¹

To extensively test ARGO and GOLD under different traffic situations, road environments, and weather conditions, we car-

THE ARGO CAR INCORPORATES A SYSTEM THAT DETECTS OBSTACLES AND LANE MARKINGS IN REAL TIME AND AUTONOMOUSLY STEERS THE VEHICLE. IT DROVE AUTONOMOUSLY FOR APPROXIMATELY 2,000 KM ON PUBLIC ROADS IN REAL TRAFFIC CONDITIONS.

ried out the MilleMiglia in Automatico tour, a 2,000-km journey from 1 to 6 June 1998. The tour demonstrated GOLD's robustness in different terrain, weather, and traffic conditions, and helped to detect situations in which the system has problems, with the aim of enhancing its abilities.

ARGO's systems

Figure 1 shows ARGO and its equipment, which consists of the data-acquisition, processing, and output systems.

Data acquisition. ARGO uses only passive sensors (cameras and a speedometer) to sense the surrounding environment, because they offer noninvasive data acquisition—that is, without altering the environment. (For background on vision-based sensing and

vehicle guidance, see the related sidebar on page 61).

The vision system. One of the project's aims is to develop a system inexpensive enough to ease its integration into a large number of vehicles. So, the use of low-cost acquisition devices was a clear starting point. ARGO's stereoscopic vision system consists of two small (3.2×3.2 cm), low-cost (\$100 each) cameras that can acquire pairs of gray-scale images. They feature a 6.0-mm focal length and 360-line resolution, and can be synchronized by an external signal.

The cameras are inside the car at the top corners of the windshield, so that the longitudinal distance between them is maximal. The optical axes are parallel. To handle non-flat roads, the cameras capture part of the scene over the horizon. However, the framing of a portion of the sky might adversely affect

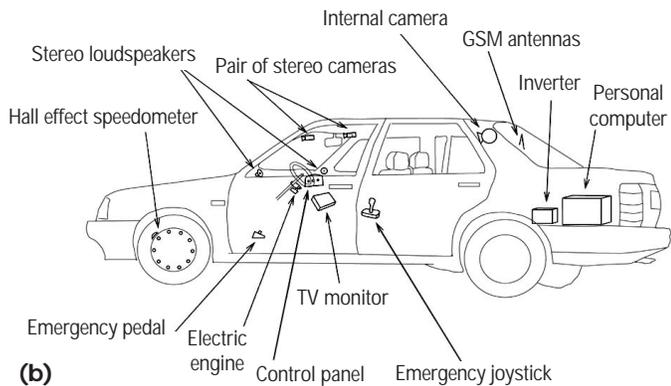
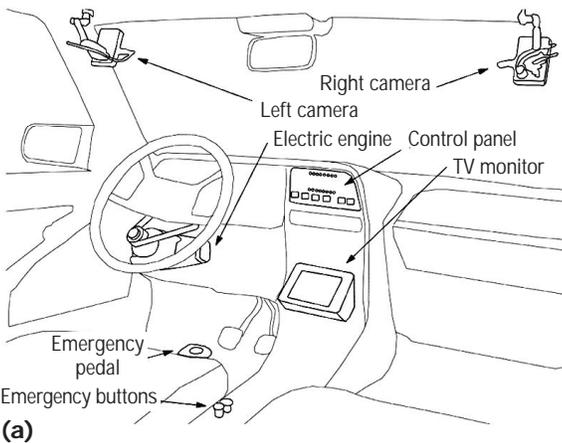


Figure 1. ARGO's equipment: (a) internal ; (b) external.

the image's brightness. For example, in the case of high contrast, the sensor might acquire oversaturated images. We'll describe later how we're dealing with this.

To acquire images, GOLD uses a PCI Matrox board, which can grab three 768×576 -pixel images simultaneously; we use only two of the images. It directly stores the images into the host computer's main memory, thanks to a DMA (direct-memory-access) device. Acquisition can occur in real time, at 25 Hz for full frames or 50 Hz for single-field acquisition.

System calibration. Because the processing is based on stereo vision, camera calibration is key to the approach's success. We painted a grid with a known size on the ground, and we captured two stereo images and used them for the calibration. Thanks to an X-Window-based graphical interface, the user selects the intersections of the grid lines, using a mouse. These intersections represent a small set of points whose world coordinates are known to the system; the system uses this mapping to compute the calibration parameters.

Processing. Vision-based vehicle guidance has two architectural solutions: special-purpose or standard processing systems.

Special-purpose systems offer ad hoc design of the processing paradigm and the overall system architecture. However, they demand complete management of the project, from the hardware level (design of the ASICs—application-specific integrated circuits), to the architecture design, to the programming language and optimizing compiler designs, to the development of applications using the specific computational paradigm.

Although standard processing systems suffer from a less specific instruction set and an architecture less oriented to vehicle guidance, they exploit standard development tools and environments. However, the fast technical improvements tend to reduce the system's lifetime, while the costs of system design and reengineering are justified only for high-volume production.

After considering these advantages and disadvantages, we chose an architectural system for ARGO based on a standard 200-MHz MMX Pentium processor.

MMX technology is an enhancement of

Intel processors that adds instructions, registers, and data types specifically for multimedia data processing. More precisely, it boosts processor performance by exploiting a SIMD (single-instruction, multiple-data) technique. It supports new general-purpose instructions for performing arithmetic and logical operations on multiple data elements packed into 64-bit quantities. These instructions accelerate the performance of applications based on computationally intensive algorithms that perform localized recurring operations on small native data. More specifically, in the processing of gray-scale images, data is represented in 8-bit quantities. So, an MMX instruction can operate on eight pixels simultaneously.

Output. A set of output devices on ARGO feeds the processing's results to the passengers. An acoustical device warns the driver if the system has detected dangerous conditions—for example, when the distance to the leading vehicle is under a safety threshold or when the vehicle's position in the lane is not safe. The system also supplies visual feedback to the driver by displaying the results

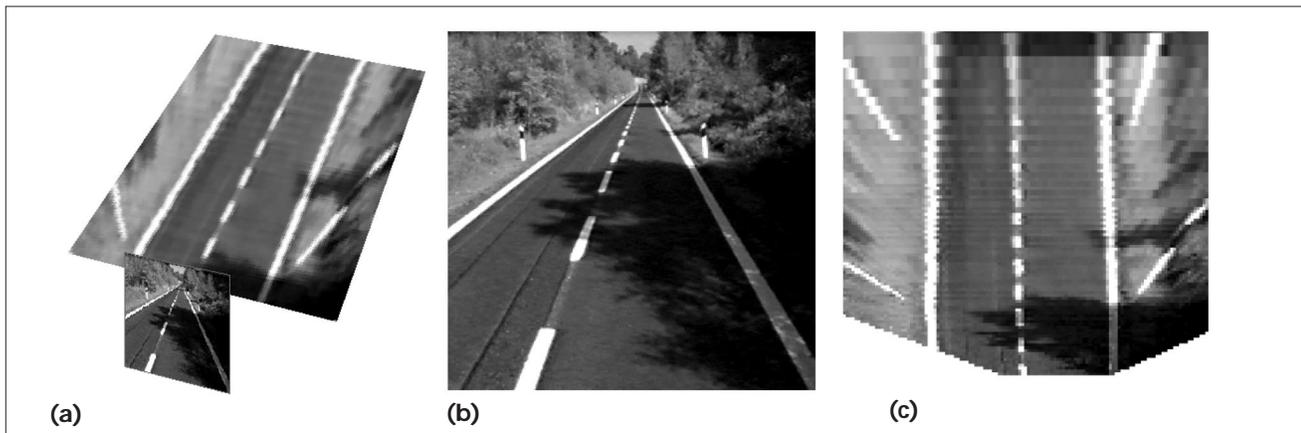


Figure 2. Applying inverse perspective mapping to a road environment: (a) 3D representation of the environment; (b) the acquired image; (c) the remapped image.

on an onboard monitor and on an LED-based control panel. The monitor presents the acquired left image with markers highlighting the lane markings and the position of eventual obstacles. The LEDs encode the vehicle's offset with respect to the road's centerline. Moreover, an actuator on the steering wheel provides autonomous steering; we haven't integrated automatic throttle and brake control yet.

Driving modes. Thanks to the control panel, the driver can adjust the values of various driving parameters and select the system's driving mode:

- *Manual.* The system monitors the driver's activity; in potentially dangerous situations it warns the driver with acoustic and optical signals.
- *Supervised.* In dangerous situations, the system controls the vehicle to keep it safe.
- *Automatic.* The system drives automatically, following the lane and localizing obstacles on the path; it can change lanes.

GOLD

GOLD is basically an automatic lateral-control system, in which lane and obstacle detection provide the necessary input for the automatic steering. (For background on lane detection, see the related sidebar on page 62.)

Vision-based lane-marking detection. The angle of view under which a scene is acquired and the distance of the objects from the camera (namely, the perspective effect) contribute to associate a different information content to each pixel of an image. Image processing must take the perspective effect into account to weigh each pixel according to its informa-

tion content. This differentiated processing makes the use of SIMD machines such as MMX-based computers a knotty problem.

To cope with this situation, we've used a geometrical transform called *inverse perspective mapping*.² IPM lets us remove the perspective effect from the acquired image, remapping it into a new 2D domain (the *remapped domain*) that homogeneously distributes the information among all the pixels. This allows efficient processing on a SIMD computer. Assuming the road in front of the vision system is known, IPM lets us obtain a bird's-eye view of the scene (see Figure 2).

Lane detection relies on the basic assumption that, after the IPM transform, quasi-vertical constant width lines (which are brighter

than their surrounding region) represent road markings. Figure 3 shows the lane-detection steps for the image shown in Figure 2.

So, the first phase of lane detection is based on the search for dark-bright-dark horizontal patterns of a given size. The result encodes the horizontal brightness transitions and the presence of lane markings. Different illumination conditions, such as shadows or sunny blobs, cause road markings to assume different brightness values. Even so, the pixels corresponding to the lane markings maintain a brightness value higher than that of their horizontal neighbors. Taking advantage of the lane markings' vertical correlation, the system enhances this result (see Figure 3a) through a few iterations of *geodesic morphological dilation*.²

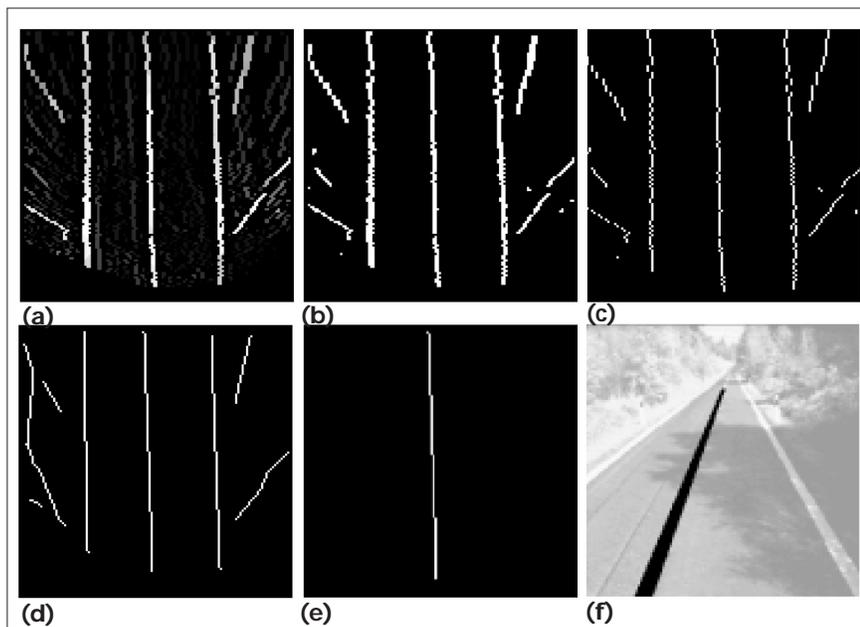


Figure 3. Lane-detection steps: (a) enhanced image; (b) binarized image; (c) concatenation of pixels; (d) segmentation and construction of polylines; (e) identification of the center of the lane; (f) superimposition of the previous result onto a brighter version of the original image for displaying purposes only.



Figure 4. A map of the MilleMiglia in Automatico tour.

Different illumination conditions and the nonuniformity of painted road signs necessitate an adaptive threshold¹ that works on a 3×3 -pixel neighborhood (see Figure 3b). The result is thinned (see Figure 3c) and scanned row by row to build chains of non-zero pixels (see Figure 3d).

Through an iterative process, the system approximates each chain with a polyline made of one or more segments. To get rid of possible occlusions or errors caused by noise, more polylines are joined into longer ones if they satisfy some criteria such as a small distance between the nearest extrema and similar orientation. When more solutions are possible, the system considers all of them.

The system selects the polyline that best approximates the results of the previously processed image; this phase uses no specific model. Initially, the vehicle is assumed to be in a specific position on the road (the lane's center), which is assumed to be almost straight. In this situation, the centerline in the remapped image is a straight vertical line that we expect to find in a circumscribed area of the remapped image. Computed polylines are matched against this model, using several parameters (with specific weights¹) such as distance, parallelism, orientation, and length. The polyline that best fits these parameters is selected (see Figure 3e). It is used to determine a new road model that enables the system to track the road in image sequences and to work on nonstraight roads.

The model assumed for the external environment (a flat road) lets us determine the spatial relationship between image pixels and the 3D world.² So, from the previously processed image, we can derive both the road geometry and the vehicle's lane position.

Automatic steering. GOLD uses the processing results (the position of obstacles and the road geometry) to drive the steering-

wheel actuator. The system combines its outputs (such as ARGO's lateral offset, its yaw relative to the centerline, and the upcoming road curvature) to determine the lane's center at a given distance ahead of ARGO (corresponding to 1.5 seconds). The system turns the steering wheel to guide ARGO toward that point.

GOLD's fast processing rate (that is, the video rate—25 Hz) allows smooth vehicle control.

The MilleMiglia in Automatico tour

The tour consisted of seven stages, each approximately 250 km. Figure 4 and Table 1 outline the tour, which we broadcasted over the Internet (see the "Live broadcasting" sidebar). During the tour, ARGO drove autonomously (see Figure 5) along the Italian highway network, passing through flat areas and hilly regions including viaducts and tunnels. The Italian road network is particularly suited for such an extensive test because it has quickly varying road scenarios that include changing weather conditions and a generally high amount of traffic. The tour took place on highways and freeways, but the system also worked on sufficiently structured rural roads with no intersections.

Performance data. During the journey, besides the normal tasks of data

Table 1. The stages of the MilleMiglia in Automatico tour.

STAGE	DATE	TIME	LEG	DISTANCE		COMMENTS
				(KM)	TERRAIN	
1	1 June	9:00–12:00	Parma–Turin	245	Mainly flat	
2	2 June	9:00–11:00	Turin–Milan–Pavia	175	Flat	Took the Milan bypass
3		16:00–20:00	Pavia–Milan–Ferrara	340	Flat	Took the Milan bypass
4	3 June	12:00–15:30	Ferrara–Bologna–Ancona	260	Flat until Rimini, then hilly crossing the Appennine region toward Ancona	Took the Bologna bypass
5	4 June	8:30–13:00	Ancona–Pescara–Rome	365	Hilly until Rome	
6	5 June	9:00–12:30	Rome–Florence	280	Flat and hilly	Took the Rome bypass
7	6 June	8:30–12:00	Florence–Bologna–Parma	195	Mountainous crossing the Appennine region as far as Bologna, then flat until Parma	

acquisition and processing for automatic driving, the system logged the most significant data, such as speed, steering-wheel position, lane changes, and user interventions and commands. It also dumped its entire status (images included) when it couldn't reliably detect the road lane.

We processed this data offline after the tour to compute the overall system performance, such as the percentage of automatic driving, and to analyze unexpected situations. In particular, for each stage of the tour, we computed ARGO's average and maximum speed. Heavy traffic conditions (especially on Turin's and Milan's bypasses) and toll stations, junctions, and road works, strongly influenced the average speed. The possibility of reprocessing the same images, starting from a given system status, lets us examine the conditions in which a fault was detected, to solve it. At the end of the tour, the system log contained more than 1,200 Mbytes of raw data.

The analysis of the data collected during the tour allowed the computation of a num-



Figure 5. Automatic driving during the MilleMiglia in Automatico tour.

Table 2. System performance during the tour.

STAGE	AVG. SPEED (KPH)	MAX. SPEED (KPH)	AUTOMATIC DRIVING (%)	MAX. DISTANCE IN AUTOMATIC (KM)
1	86.6	109	93.3	23.4
2	80.2	95	85.1	42.2
3	89.8	115	86.4	54.3
4	89.8	111	91.1	15.1
5	88.4	108	91.1	20.8
6	87.5	110	95.4	30.6
7	89.0	123	95.1	25.9

ber of statistics regarding system performance (see Table 2).

The automatic driving percentage and the maximum distance automatically driven show high values despite many tunnels (particularly during the Appennines routes Ancona–Rome and Florence–Bologna), several stretches of road with absent or worn lane markings (Ferrara–Ancona and Ancona–Rome), and stretches with no lane markings at all (Florence–Bologna). (Entering and exiting tunnels can cause sudden, extreme changes in illumination, which can be a problem.) Some stages included passing through toll stations and driving in bypasses with heavy traffic and frequent queues, during which we had to switch off the system.

Figure 6 shows the system's behavior dur-

ing one hour of the tour's third stage. First, at the top, the figure shows the driving mode. During this period, the system was always in automatic mode, except in four situations:

- A: ARGO passed through two tunnels (near Vicenza).
- B: At the end of an overtaking maneuver, the system did not recognize the lane markings because of an occlusion.
- C and D: ARGO passed two working areas with no lane markings.

(In Figure 6, situation E was a junction with a ramp, which we'll discuss in the next section.) In the first two cases, the system switched back to manual driving automatically and warned the driver; in the other two

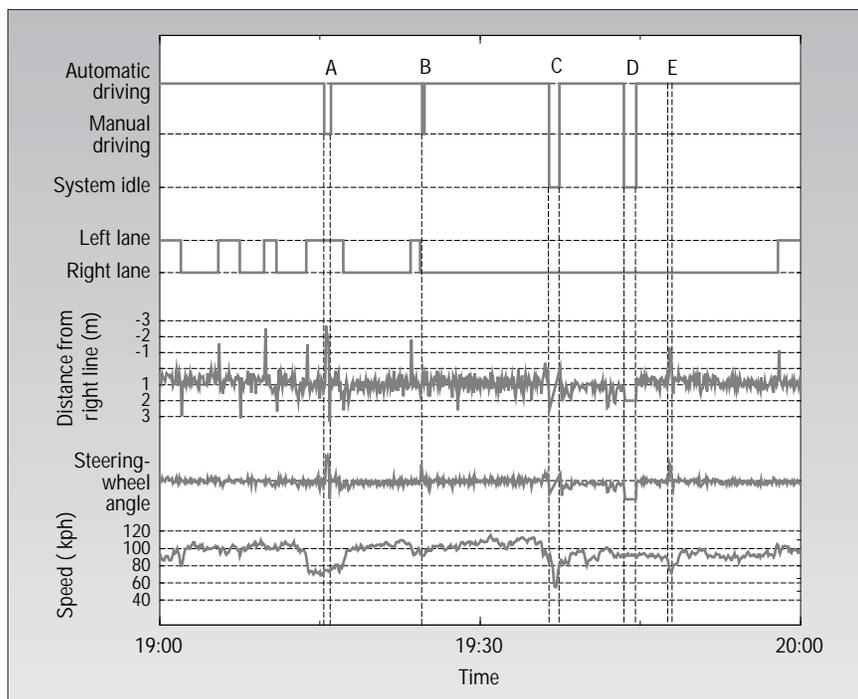


Figure 6. The system's behavior during one hour of automatic driving.

Live broadcasting

During the tour, thanks to the sponsorship of the Telecom Italia Mobile cellular phone company, the ARGO vehicle broadcasted a live video stream on the Internet: Two GSM (global system for mobile communications) cellular modems connected to the Vision Laboratory at the University of Parma's Department of Information Engineering transferred up-to-date news on the tour's progress and camera images of ARGO's interior (see Figure A).

The setup of a permanent data link from a moving platform implies the use of mobile telecommunications facilities, such as GSM modems. The use of GSM modems for live video streaming forces two constraints: a low transmission rate (usually 9,600 bps) and high bandwidth variability during movements.

To increase the link's throughput, we installed two modems—and therefore two channels—that work simultaneously on ARGO. The transmission hardware consists of a PC running Linux with two serial ports and a color camera connected to the parallel port. The network traffic is split across the two serial channels, thanks to the EQL (Equalizer Load-balancer) protocol, which can split the network traffic across multiple links. Moreover, we designed the communication software to dynamically adapt the network traffic to the throughput.

The tour generated high interest with the scientific community, mass media, and general public. The Web site received more than 350,000 contacts during the tour. More than 3,000 Mbytes of information were transferred, with a peak of 16,000 contacts per hour during the tour's first day.



Figure A. The internal camera framing ARGO's interior. A GSM antenna is in the background.



Figure 7. Images dumped in the anomalous situations A, B, C, and D in Figure 6.

cases, the system was switched off manually. For these four situations, the other information in Figure 6 is useless for analyzing system performance. Figure 7 shows the images dumped on disk in these occasions.

Second, Figure 6 shows ARGO's lane position. ARGO made a number of lane changes, moving from the right lane to the central lane and vice versa (it was on a three-lane highway).

Third, the figure shows the distance between the right wheel and the right lane markings (in centimeters). It is possible to detect a lane change from the variation of the distance. When the vehicle reached the middle of the new lane, the system locked onto the new lane, and the distance returned to the nominal value. The selected value of the distance is 80 cm.

Fourth, it shows the steering-wheel angle that the system automatically chose. This

angle is proportional to the distance shown by the previous graph, with the only exception of the peaks caused by lane changes.

Finally, it shows the vehicle's speed. The low speeds in situations C and E are due respectively to road works and driving through a junction (the related steering-wheel movement is also visible in the figure).

Performance analysis. Now let's look at how the vision system, the control system, the user interface, and the environmental conditions affected system performance.

The vision system. Although the system has high sensitivity even in low light (at night, for example), high contrast or a quick change in a scene's illumination degrades the image quality (for example, when entering or leaving a tunnel). In particular, the cameras' slow automatic gain control (designed for applications

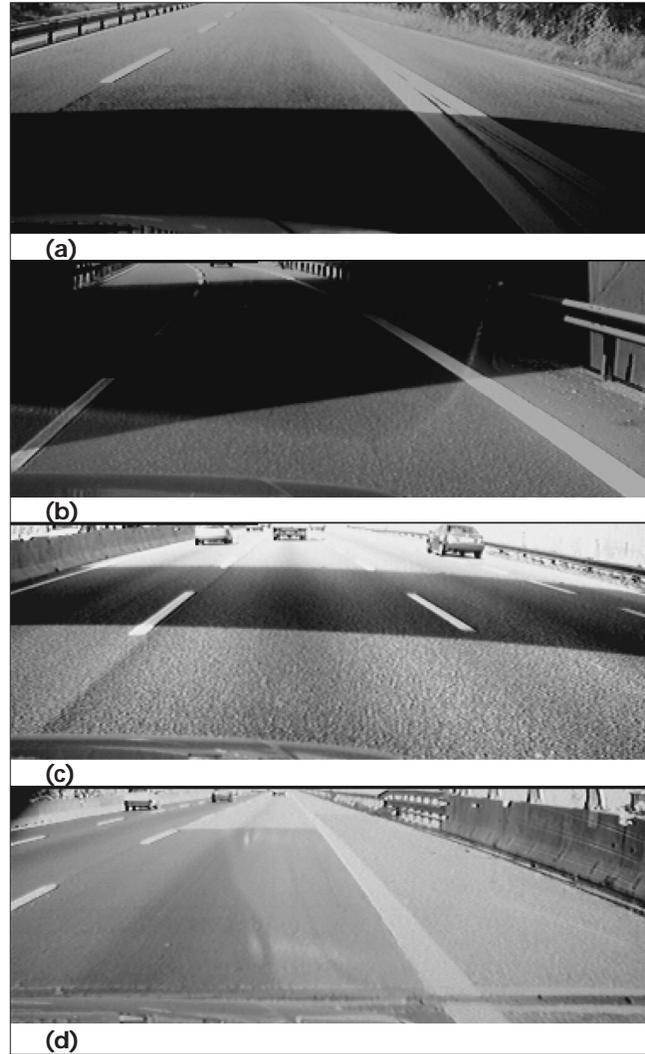


Figure 8. Image saturation caused by high contrast: (a) an exit from a tunnel; (b) a long bridge; (c) a short bridge; (d) two different asphalt patches and a strong reflection on the windshield.

characterized by constant illumination, such as videotelephony) was a problem. During tunnel exits, the gain control caused the acquired images to be completely saturated and impossible to analyze for an interval of 100 to 200 ms. This problem also happened when the vehicle crossed an area that contained patches of new (black) and old (light gray or even white) asphalt. Figure 8 shows examples of high-contrast situations. We are considering several solutions to this problem, which we'll discuss in our conclusion.

The processing system. ARGO's processing system, a commercial PC with a 200-MHz Pentium processor and 32 Mbytes of memory, proved to be powerful enough for automatic driving. Using this system, enhanced by the video frame grabber, GOLD processes up to 25 pairs of stereo frames per second and provides the control signals for au-

tonomous steering every 40 ms. (This is equivalent to one refinement on the steering wheel position for every meter when the vehicle drives at 100 kph.)

The visual processing. During the tour, the system processed about 1.5 million images (each 768×288 pixels), totaling approximately 330 Gbytes of input data. The resolution of the cameras was satisfactory, and the framing of the scene (no horizon to reduce strong light conditions and direct sunlight) was correct for the type of roads we encountered (sloping gently).

The IPM approach proved to be effective for the whole trip. Even if on Italian highways the flat-road assumption (required by the IPM transform) is not always valid, the approximation of the road surface with a planar surface was acceptable. Actually, the vision system's calibration should reflect the modifications of road slope ahead of the vehicle; a wrong calibration generates a lateral offset in the computation of the vehicle's trajectory. Nevertheless, because the highway lanes' width is sufficiently large, this offset has never caused serious problems. However, we are enhancing the IPM transform to eliminate the flat-road assumption and to handle generic roads.³



Figure 9. A vehicle occludes the visibility of the right lane marking near an exit, causing the system to request user intervention.

The IPM transform's only drawback is that vehicle movement (pitch and roll) does not allow reliable detection of obstacles at distances over 50 meters.

Because the system locks onto the right lane marking during navigation, the vehicles in front do not occlude its visibility. This choice, however, is critical during the complex situation of highway exits, where two lines exist—a continuous line for the exit and a dashed line for the lane. Because GOLD cannot automatically choose whether to stay on the highway or to exit, the system requests user intervention. After the user intervenes by pressing a button, the system quickly resumes autonomous driving. (This explains why the percentage of automatic driving was high even when the maximum continuous distance driven in automatic was not that far.) This approach is obviously simpler than approaches based on a fusion of information coming from the detection of both left and right lane markings. We are developing an

extension to handle both markings, to quantitatively compare the two approaches.

Figure 9 shows an exit situation in which the system suffered a major fault: a vehicle occluded the visibility of the right lane marking near a highway exit.

Because the only visible lane marking was the continuous line representing the exit, the system followed that line, therefore requiring user intervention.

Another serious problem happened in a stretch of road just after the conclusion of road works. Lines painted to delimit the diversion were still on the pavement. So, the vehicle followed those lines.

The control system. We developed the control system to be simple. Regarding the mechanical part, an electric stepping motor allows steering-wheel rotation with high resolution and reduced power consumption. The simplicity of the logical part of the control system keeps the entire system robust. The system uses only a temporal average and a simple measurement to compute the steering-wheel angles.

For speeds reaching approximately 90 to 95 kph, the system does not perform noticeably different than a human driver. For higher

The promise and problems of vision-based vehicle guidance

Vision is one of many ways of sensing the surrounding environment. Radio, acoustic, and tactile sensors (such as radar, lasers, sonar, and bumpers) can also be used, but passive sensors such as cameras can offer prominent advantages over other kinds of sensors. One advantage is noninvasive data acquisition, which does not alter the environment. In situations involving massive, widespread use of autonomous sensing agents, such as when many vehicles are moving simultaneously in a road environment, passive sensors are considerably preferable to invasive sensors, which could cause unacceptable environmental pollution.

Vision-based automatic vehicle guidance in outdoor environments faces three key problems:

First, unlike in indoor settings, the sensing in outdoor environments should not rely on ad hoc structured information. This is particularly true in the automotive field, where vehicles can move along roads with different characteristics, and the integration of specific road infrastructures would be extremely expensive and prohibitive. An automatic vehicle should be able to navigate using standard road signs, without requiring further additional infrastructures.

Second, automatic vehicle guidance requires real-time processing because the vehicle's maximum speed is bounded by the processing rate. The main problem, intrinsic to image processing, is the large amount of data—and thus computation—involved.

Third, contrary to indoor applications where robots move in controlled environments, in the automotive field no assumptions can be made about key parameters that are directly measured by the vision sensor. Important examples of such parameters are the illumination or the contrast of the scene. So, the subsequent processing must be robust enough to tolerate changes to these parameters and to adapt them to different road conditions, such as sun (high brightness and contrast caused by shadows), rain (extremely high contrast caused by reflections), and fog (low contrast). The processing must also adapt to dynamic changes in the parameters, such as transitions from sun to shadows or the entrance in a tunnel. Other key problems exist, such as the system's robustness when dealing with camera movement and drifts in its calibration.

However, although computer vision is extremely complex and highly demanding, it is a powerful means for sensing the environment. It has been widely employed to address many complex and challenging automotive tasks,¹ ranging from road following to platooning (when a vehicle automatically follows a manually driven vehicle) and from vehicle overtaking to automatic parking. Of these tasks, we paid the most attention to road following (see the related sidebar) because it comprehends two basic functionalities: obstacle detection and lane detection.

Reference

1. *IEEE Intelligent Systems* (special issue on vision-based vehicle guidance), Vol. 13, No. 6, Nov./Dec. 1998.

speeds (during the test, ARGO reached a peak of 123 kph), the vehicle tends to demonstrate slightly unstable behavior, oscillating inside the lane.

We are studying a more sophisticated control system including a strong road model; this should allow higher speeds with more stability. This new control mechanism is undergoing tests on a system simulator that emulates vehicle behavior on curved roads with more than one lane. We've also used it to test lane-change maneuvers.

The user interface. To manage the automatic driving system, we provided a control panel with six buttons and eight LEDs that indicate the selected driving mode. Because it is still in early development, we can modify the value of a number of driving parameters, which should normally be hard-coded and not adjustable by the final user. Nevertheless, this capability does not increase the interface's complexity. For security, the user must press two buttons simultaneously to execute each command. The system then notifies the user of modifications of its state through vocal messages that confirm the command's

reception and execution. Moreover, for some commands whose execution is not instantaneous (such as the lane-change command), the system also generates an acoustic signal at the maneuver's completion.

The emergency system lets the passenger take control of the vehicle during emergencies by

- tuning driving parameters (through the control-panel buttons),
- temporarily taking over the system (by using a joystick), or
- completely taking over the system (by pressing a pedal close to the clutch, through the control panel buttons, or with an emergency button close to the hand brake).

The user interface's behavior was satisfying. It did not flood the driver with redundant information, and it was simple enough for anyone to use.

Environmental conditions. The tour route included areas with different morphological characteristics: from flat areas (Parma, Piacenza, Turin, Milan, Verona, Padova, Ferrara,

and Bologna), to the sloping territories of the Appennines region (Pescara, L'Aquila, Rome, Florence, and Bologna), and heavy traffic zones (Rome, Turin, and Milan's bypasses). We inevitably encountered stretches of highway with road works, absent or worn horizontal road signs, and diversions.

ARGO did oscillate slightly on sloping stretches of road. However, the system was particularly robust when dealing with horizontal road signs; fog markings; forks, junctions, and exits; heavy traffic conditions; and guardrails. Figure 10 illustrates these situations.

Moreover, the image-processing algorithms were robust in different lighting conditions. The system could always extract the information for navigation even in critical light conditions, with the sun in front of the cameras, with the sun high or low on the horizon, during the evening or the day, and with high or low contrast. At night, the absence of sunlight reflections and shadows improved the system's behavior, while ARGO's headlights constantly illuminated the area of interest. One problem, which we're working on, was the light's reflection on the windshield's

Lane detection

In most prototypes of autonomous vehicles, road following is based on lane detection: First, the guidance system computes the vehicle's position relative to the lane; then it drives actuators to keep the vehicle in a safe position. Generally, such systems reduce lane detection to the localization of specific features such as road markings painted on the road surface. Although this solution can simplify lane detection, it faces two basic problems:

- Shadows (projected by trees, buildings, bridges, or other vehicles) that produce artifacts onto the road surface, thus altering the road texture's appearance;
- Other vehicles on the path that partly occlude the visibility of the road and thus of road markings.

However, some assumptions can aid the detection, speed up the processing, or both:

- Because of physical and continuity constraints, the processing of the whole image can be replaced by the analysis of specific regions of interest (the focus of attention), in which finding the features of interest is more probable.
- The assumption of a fixed or smoothly varying lane width allows enhancement of the search criterion, limiting the search to almost straight lane markings.
- Assumptions of the road geometry's shape can simplify its reconstruction; for example, some implementations use a clothoid (a curve with constantly varying curvature) to model the road's geometry.
- Knowledge of the specific camera calibration and the assumption of a precise road model (for example, a "flat" road without bumps) can be exploited to ease the localization of features. This know-

ledge can also help simplify the mapping between image pixels and their correspondent world coordinates.

Two approaches to lane detection exist: *model-based* and *feature-based*.

The model-based approach consists of a direct match between visual patterns and a model that encodes the road's characteristics. Some implementations iteratively stretch and reshape an initial model according to the content of each frame of the acquired sequence. These deformable templates have been successfully used to detect lane markings, road boundaries, and road signs. Model-based analysis has also been used to process images of intersections in city traffic.

Nevertheless, the model-based approach has several drawbacks. One is the problem of using and maintaining an appropriate geometrical road model. A second drawback is the difficulty of detecting and matching complex road characteristics. A third is the complexity of the computations involved, because the whole processing is performed on large data structures.

Conversely, the feature-based approach limits the computationally intensive task to the extraction of features. It leaves the match with the model to the following medium- and high-level processing steps, thus reducing the amount of data it handles.

Interested readers can find a survey of similar projects in our book,¹ and some examples of solutions to vision-based driving problems in the Nov./Dec. 1998 *Intelligent Systems* (Vol. 13, No. 6).

Reference

1. A. Broggi et al., *Automatic Vehicle Guidance: The Experience of the ARGO Autonomous Vehicle*, World Scientific, Singapore, 1999.

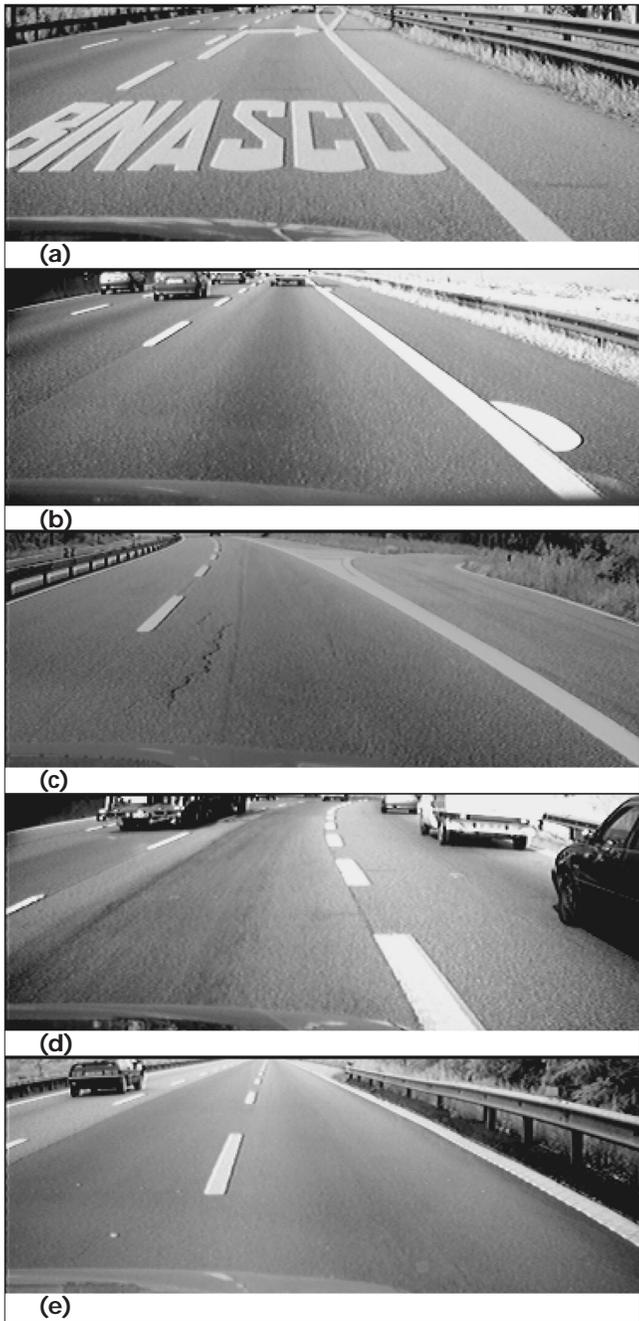


Figure 10. The system demonstrated good reliability when dealing with (a) horizontal road signs, (b) fog markings, (c) exits, (d) heavy traffic, and (e) guard-rails.

internal surface (see Figure 11a). Figures 11b, 11c, and 11d show some typical effects of this reflection.

Finally, the system proved to be surprisingly robust despite the high temperatures during the tour. Even when the external temperature reached 35°C, the system continued to work reliably, even with no air-conditioning.

Minor problems. During the tour, ARGO suffered a few difficulties that did not relate directly to the performance of its systems. These problems included a power failure

caused by an overloading of the power supply, interference of an onboard cellular phone with the cameras and acquisition board, and some human errors in the use of the control panel.

FUTURE DEVELOPMENTS OF ARGO's autonomous control system will deal with enhancing the vision system and controlling the steering wheel. (For the latest informa-



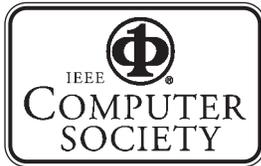
Figure 11. The reflection problem: (a) the light's reflection on the windshield's internal surface causes the images to be oversaturated; (b) an image with no reflection; (c) an image with a strong reflection; (d) an oversaturated image caused by reflection on the windshield.

tion about the ARGO project, visit <http://mille Miglia.ce.unipr.it>.)

In particular, although the lane-detection algorithm is extremely robust, we want to enhance it to look at two lines instead of just one, to remove the errors caused by occlusions or highway exits. We're also improving the computation of an obstacle's distance and will integrate the results with the lane-detection subsystem.

Also, the tour showed that the stereo cameras were incompatible with the application. Because their gain control is too slow and

PURPOSE The IEEE Computer Society is the world's largest association of computing professionals, and is the leading provider of technical information in the field.



MEMBERSHIP Members receive the monthly magazine **COMPUTER**, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

EXECUTIVE COMMITTEE

President: LEONARD L. TRIPP
Boeing Commercial Airplane Group
P.O. Box 3707
M/S 19-RF
Seattle, WA 98124

President-Elect:
GUYLAINE M. POLLOCK *
Past President:
DORIS CARVER *
VP, Press Activities:
CARL K. CHANG †
VP, Educational Activities:
JAMES H. CROSS †
VP, Conferences and Tutorials:
WILLIS KING (2ND VP) *
VP, Chapter Activities:
FRANCIS LAU*
VP, Publications:
BENJAMIN W. WAH (1ST VP)*

STEVEN L. DIAMOND *
VP, Technical Activities: JAMES D. ISAAK *
Secretary:
DEBORAH K. SCHERRER*
Treasurer:
MICHEL ISRAEL*
IEEE Division V Director:
MARIO R. BARBACCI †
IEEE Division VIII Director:
BARRY JOHNSON†
Executive Director:
T. MICHAEL ELLIOTT†

VP, Standards Activities:

*voting member of the Board of Governors †nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 1999: Steven L. Diamond, Richard A. Eckhouse, Gene F. Hoffnagle, Tadao Ichikawa, James D. Isaak, Karl Reed, Deborah K. Scherrer

Term Expiring 2000: Fiorenza C. Albert-Howard, Paul L. Borrill, Carl K. Chang, Deborah M. Cooper, James H. Cross, III, Ming T. Liu, Christina M. Schober

Term Expiring 2001: Kenneth R. Anderson, Wolfgang K. Giloi, Haruhisa Ichikawa, David G. McKendry, Anneliese von Mayrhauser, Thomas W. Williams

Next Board Meeting: 19 Feb. 1999, Houston, Texas

COMPUTER SOCIETY OFFICES

Headquarters Office
1730 Massachusetts Ave. NW,
Washington, DC 20036-1992
Phone: (202) 371-0101
Fax: (202) 728-9614
E-mail: hq.ofc@computer.org

Publications Office
10662 Los Vaqueros Cir.,
PO Box 3014
Los Alamitos, CA 90720-1314
General Information:
Phone: (714) 821-8380
membership@computer.org
Membership and
Publication Orders: (800) 272-6657
Fax: (714) 821-4641
E-mail: cs.books@computer.org

European Office
13, Ave. de L'Aquilon
B-1200 Brussels, Belgium
Phone: 32 (2) 7702198
Fax: 32 (2) 7708505
E-mail: euro.ofc@computer.org

Asia/Pacific Office
Watanabe Building
1-4-2 Minami-Aoyama,
Minato-ku, Tokyo 107-0062, Japan
Phone: 81 (3) 3408-3118
Fax: 81 (3) 3408-3553
E-mail: tokyo.ofc@computer.org

EXECUTIVE STAFF

Executive Director and CEO:
T. MICHAEL ELLIOTT

CFO: VIOLET S. DOAN

Publisher:
MATTHEW S. LOEB

CIO: ROBERT G. CARE

Director, Volunteer Services:
ANNE MARIE KELLY

Manager, Research & Planning:
JOHN C. KEATON

IEEE OFFICERS

President: KENNETH R. LAKER
President-Elect: BRUCE A. EISENSTEIN
Executive Director: DANIEL J. SENESE
Secretary: MAURICE PAPO
Treasurer: DAVID CONNOR
VP, Educational Activities: ARTHUR W. WINSTON
VP, Publications: LLOYD "PETE" MORLEY
VP, Regional Activities: DANIEL R. BENIGNI
VP, Standards Activities: DONALD LOUGHRY
VP, Technical Activities: MICHAEL S. ADLER
President, IEEE-USA: PAUL KOSTEK



their dynamics are not sufficiently high, we will replace them with better-performing cameras. However, we will maintain their position and orientation. Moreover, to get rid of the annoying reflection on the inside of the windshield, we will put shades on the cameras, move them closer to the windshield, and add a light-polarizing filter.

We are also evaluating CMOS-based sensors. Their logarithmic response should allow higher dynamics. Their intrinsic slowness, caused by the specific pixel-by-pixel addressing mode, should pose no additional problems because the early stages of processing (IPM) are based on a strong subsampling of the incoming image. Moreover, their use should speed up image acquisition because a complete scanning and transfer of the whole image is no longer required. ■

Acknowledgments

The Italian National Research Council (CNR) partially supported this project; Telecom Italia Mobile sponsored the MilleMiglia in Automatico tour. We also thank Gianni Conte for his support throughout the project.

References

1. A. Broggi et al., *Automatic Vehicle Guidance: The Experience of the ARGO Autonomous Vehicle*, World Scientific, Singapore, 1999.
2. R. Haralick and L. Shapiro, *Computer and Robot Vision*, Vols. 1 and 2, Addison-Wesley, Reading, Mass., 1992.
3. M. Bertozzi, A. Broggi, and A. Fascioli, "An Extension to the Inverse Perspective Mapping to Handle Non-Flat Roads," *Proc. IEEE Intelligent Vehicles Symp.*, IEEE Press, Piscataway, N.J., 1998.

Alberto Broggi is an associate professor of artificial intelligence at the Department of Computer Science and Systems at the University of Pavia, Italy. He is the ARGO project's coordinator. His research interests include real-time computer-vision approaches for the navigation of unmanned vehicles, and the development of low-cost computer systems used with autonomous agents. He is the newsletter editor for and a member of the Conference and Publication Committees of the new IEEE Intelligent Transportation Systems Council, and will be the program chair of the 2000 IEEE Intelligent Vehicles Symposium, in Detroit. He received his Dr.Eng. (master's degree) in electronic engineering and his PhD in information technology from the University of Parma, Italy. Contact him at the Dipartimento di Informatica e Sistemistica, Università di Pavia, Via Ferrata 1, I-27100 Pavia, Italy; broggi@dis.unipv.it; <http://vision.unipv.it/~broggi>.

Massimo Bertozzi is a researcher at the Department of Information Engineering at the University of Parma, Italy. He received his Dr.Eng. (master's degree) in electronic engineering from the University of Parma, Italy, with a thesis about the simulation of Petri Nets on the CM-2 massively parallel architecture. He received his PhD in information technology, also from the University of Parma. He is member of the IEEE and the International Association for Pattern Recognition. Contact him at the Dipartimento di Ingegneria dell'Informazione, Università di Parma, Parco Area delle Scienze, 181A, I-43100 Parma, Italy; bertozzi@ce.unipr.it; <http://www.ce.unipr.it/people/bertozzi>.

Alessandra Fascioli is a PhD candidate in information technology at the Department of Information Engineering at the University of Parma, Italy. Her research interests focus on real-time computer vision for automatic vehicle guidance and image-processing techniques based on the Mathematical Morphology computational model. She received her Dr.Eng. (master's degree) in electronic engineering from the University of Parma, with a thesis on stereo vision-based obstacle localization in automotive environments. She is member of the IEEE and the International Association for Pattern Recognition. Contact her at the Dipartimento di Ingegneria dell'Informazione, Università di Parma, Parco Area delle Scienze, 181A, I-43100 Parma, Italy; fascal@ce.unipr.it; <http://www.ce.unipr.it/people/fascal>.