VisLab's Experience on Autonomous Driving in Extreme Environments.

Workshop Robotica per esplorazione lunare unmanned
1 Luglio – 2 Luglio
Roma, Italia

Massimo Bertozzi, Luca Bombini, Alberto Broggi, Michele Buzzoni, Andrea Cappalunga, Elena Cardarelli, Stefano Cattani, Pietro Cerri, Mirko Felisa, Rean Isabella Fedriga, Luca Gatti, Paolo Grisleri, Luca Mazzei, Paolo Medici, Pier Paolo Porta, Paolo Zani

VisLab, Dipartimento di Ingegneria dell'Informazione
Università di Parma – Via G.P. Usberti 181/a, 43100 Parma, Italy
www.vislab.it

{bertozzi, broggi, bombini, buzzoni, kappa,cardar, cattani, cerri, felisa,fedriga,lucag, grisleri, mazzei, medici, portap, zani}@ce.unipr.it

INTRODUCTION

The laboratory of Artificial Vision and Intelligent System of Parma University, that recently has started an academy spin-off called VisLab, has been involved in basic and applied research developing machine vision algorithms and intelligent system for the automotive field for more than one decade. The experience gathered through these years covers a wide range of competencies, either under the algorithmic and technological point of view. VisLab developed, tuned and tested several image processing algorithms, such as: obstacle detection, obstacle avoidance, path detection, terrain estimation and classification; often these results have been achieved by fusion between vision and active sensor, for a better 3D reconstruction of the surrounding environment. On the technological side VisLab can rely on a long experience in using several types of vision sensors (FIR, NIR, visible) as well as on laser sensors (both mono and multi plane), multi echo sensors, together with the support of odometric and inertial information.
One of the most distinctive feature that characterize VisLab, that explains its continuous presence in project with manufactures of vehicles of various types, is the very specific experience in developing and setting up out-of-laboratory unmanned prototypes. The application of vision system on board of autonomous vehicles not only requires to fully dominate the cutting edge vision technologies, but also to have a deep experience in dealing with unknown, and often hostile, environments and corresponding challenges: calibration, illumination variation, noise, temperature, installation requirements, etc .
Beside its expertise, the key to VisLab's quick application prototyping is the proprietary software GOLD, that has been developed in the last 10 years and that constitutes the basis of each application developed. GOLD is, at the same time, a powerful development framework for image processing application and a software capable of real-time data acquisition, synchronization, logging, and—ultimately—data processing and visualization, for improved researcher efficiency.

Probably the most comprensive example of the VisLab's competencies in unmanned ground vehicle development, in particular thinking about unknown and unstructured environment, is the TerraMax vehicle. Developed together with Oshkosh Truck Corp and Rockwell Collins, this vehicle was one of the five to reach the finish line at DARPA Grand Challenge 2005, covering with success a 200+ km fully autonomously through the Mojave Desert, staying continuously operative for about 30 hours.

This paper is organized as follow: section "VisLab Milestones" presents some of the most important results achieved by VisLab in the unmanned vehicle field; in "Off-road Sensing" some significant example of off-road sensing systems will be given; in "The VsiLab's Framork: GOLD" will be explained the characteristic of VisLab's proprietary software and its capabilities; the "Conclusions" Paragraph will end the paper.

VISLAB MILESTONES

VisLab has been working on unmanned ground vehicles for more than a decade, both in urban and off-road environment. In this Paragraph some examples on these applications will be presented, starting with the debut on the unmanned vehicle scene with ARGO, and then focusing primarily on the most significant results achieved on off-road/unstructured environment.

The ARGO vehicle.

ARGO [1] 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.
VisLab had been developing, testing, and tuning several solutions for autonomous navigation, particularly for the basic obstacle- and lane detection functionalities. Stereo vision was used to detect and localize 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.

Fig. 1. The ARGO vehicle

The 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.

The RAS project.

With RAS, one of the first efforts toward automating driving functionalities in extreme environments, VisLab tested a snow cat's automatic maneuvering performance during Italian scientific missions to the South Pole.
The project, funded by the national agency for the energy and the environment (ENEA, Ente per le Nuove tecnologie, l'Energia e l'Ambiente), aimed to apply artificial vision to the autonomous driving of a platoon of snow cats (Fig. 2) used to transport people and goods in Antarctica [2]. 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 snowstorms, 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 snow cat'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. 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 (Fig. 3).

Fig. 2. The Snow cat used as test-bed

Fig. 3. Successful track detection in challenging conditions

The DARPA Grand Challenge 2005: TerraMax vehicle

In the recent years 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. 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.

Fig 4. TerraMax concludes the 2005 DARPA Grand Challenge

VisLab, together with Oshkosh Truck Corp. and Rockwell Collins, developed and successfully fielded TerraMax (Fig. 4), an autonomous vehicle that reached the finish line of the 2005 DARPA Grand Challenge[3]. 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.
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.

VisLab's vehicle TerraMax was one of the five to reach the finish line, and the only one relying primarily on vision.

TerraMax at the Pentagon
Due to the outstanding result achieved by TerraMax the DARPA Grand Challenge 2005, the team was invited to present the results of its activities at the Pentagon, Washington DC, in December 2005.

The PLS vehicle.
The TerraMax partners of Oshkosh Truck, Rockwell Collins, and University of Parma have demonstrated the scalability and portability of the autonomous technology by installing and operating the system on an Oshkosh palletized loading system (PLS).
The PLS is a 10x10 vehicle with a gross vehicle weight of 84 000 lb and is capable of delivering a 33 000 lb payload. The unmanned vehicle features an autonomous driving system which is identical to the one that was developed for TerraMax in 2005, and was successfully demonstrated at the Yuma Test Center in January 2006. The project was completed in approximately 75 days. Figure 5 shows the PLS vehicle during an autonomous run.

Fig 5. PLS operating autonomously near Barstow, California

The DARPA Urban Challenge.
After having seen five vehicles reaching the finish line in 2005, the Defense Advanced Research Project Agency (DARPA) moved its third-annual robot race Grand Challenge from the desert into a city environment, calling it Urban Challenge. The Urban Challenge features autonomous ground vehicles maneuvering in urban and sub-urban scenarios, where they had to execute merging into moving traffic, navigate traffic circles, negotiate busy intersections, avoid obstacles, follow lanes and handling parallel parking. Moving traffic was provided by several vehicles driven by professionals, as well as by the robots themselves, so robot-on- robot action was possible. Eighty-nine teams had applied to take part in the competition; DARPA accepted 35 of them, and only 11 were selected for the Final Event.
VisLab was involved in the 2007 DARPA event as a member of Team Oshkosh, developing the artificial vision system for the team's vehicle TerraMax [4]. The vision system is composed of 11 cameras, able to provide the following sensing information: lane/path detection, stop line detection, obstacle detection for straight driving and maneuvering, backward vehicle detection for possible lane changing, oncoming traffic detection when stopped at a stop line. The Vision coverage is shown in Figure 6.

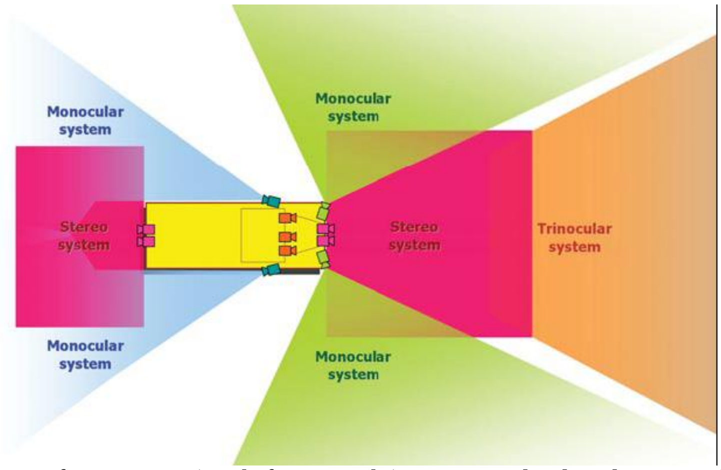

Fig. 6. Vision coverage of TerraMax (truck facing right). Systems displayed: Trinocular (orange) looking forward from 7 to 40 m, stereo front and stereo back (purple) monitoring a 10×10 m area on the front of the truck and 7×5 m in the back, Rearview (blue) monitoring up to 50 m behind the truck, and lateral (green) looking up to 130 m.

The vehicle is also equipped with three LIDARs, capable of providing a 4-planes scan of the surrounding environment. The sensors output both processed data (such as detected obstacles) and raw points: while the former is useful for the high-level path planner, the latter can be exploited to improve the performance of the various vision systems. Vision and LIDAR systems are complementary in many ways: for example, when performing a tridimensional reconstruction of the world LIDAR points offers a very accurate, yet sparse representation of the world, while vision provides dense, but less accurate measurements; some tasks, like lane and path detection, are also better accomplished through the use of vision.

Fig. 7. TerraMax at the Urban Challenge

Integrating the data coming from the two, TerraMax was able to sense its surroundings, understand situations, and react quickly; a service-based system architecture coupled with efficient path planning, trajectory algorithms and behavior modes ensures that sensor data is processed in a timely manner, while robust low level control and X-by-Wire components enable the vehicle to manage situations.

TerraMax, together with 10 other vehicles, after having successfully completed the National Qualification Event, participated to the DARPA Urban Challenge Final Race event in Victorville, California.

BRAiVE and the ERC-founded project.

VisLab is undertaking highly innovative research within its European Research Council founded (ERC) project, whose topic is the development of an open standard for the perception and decision subsystems of intelligent vehicles. Currently, many commercial vehicles include sophisticated control devices like ABS, ESP, and others. These control equipments have been independently developed by car manufacturers and suppliers. Generally, they also act independently, and are singularly tuned. Nevertheless, new methods to improve overall performance are currently under development, exploiting communication and cooperation of these devices: the recently introduced Unified Chassis Control (UCC) is an example. The deployment of the UCC in the mass market, requires to adapt and rethink all control subsystems to provide communication, data fusion, and an overall tuning: namely to integrate all of them together. From the car manufacturers and suppliers point of view, the introduction of the UCC requires the redesign of each single block (ABS, ESP, ...) meaning an additional financial effort, besides the obvious delay in reaching the market. Had a complete UCC architecture been defined well in advance with respect to the development of each single block, its implementation would have been straightforward, less costly, and would have reached the market earlier.

Perception and decision modules are in an earlier development stage than control ones: the advanced driver assistance systems that are currently available on the market are only basic ones (Lane Departure Warning, Blind Spot Monitoring, Automatic Cruise Control, Collision Mitigation), independently developed by different car manufacturers and suppliers. The state of the art of advanced driver assistance systems, in fact, has not yet defined a complete architecture that would allow to fuse together all these basic blocks and benefit from their integration. The availability of such architecture would allow to define a standard module interface so that the following research efforts could be more focused in providing modular systems, already engineered to fit into this architecture.
A new architecture, called Unified Perception & Decision Architecture (UPDA) [5], is the VisLab's answar to this problem. It will take as input all the data coming from the different perception modules, to fuse them, and to build a more sophisticated world model to be used in the following decision level. In doing this, a standard interface will be defined for each perception module, allowing different providers to integrate their own perception system into this standard architecture, thus boosting competition.

Fig. 8. BRAiVE

In order to develop these concepts, VisLab has set up a vehicle prototype (three more are expected for 2010) integrating various sensing technologies, called BRAiVE (BRAin-drIVE): used as the primary test bed for Advanced Driver Assistance Systems products, it is going to be also the development platform for what cancers UPDA and all the fully autonomous driving operations will be carried on for this porpoise. BRAiVE is equipped with 10 cameras, 4 laserscanners, 16 laser beams, GPS, IMU, and full X-by-wire for autonomous driving.
BRAiVE features the following intelligent capabilities:

- Crossing Asstance;
- Overtaking Assistance;
- Obstacle and Pedestrian Detection;
- Lane ans Stop Detection / Lane Keeping
- Parking Assistance;
- Road Sign Detection;
- Stop and Go & Automatic Cruise Control.

Detailed description of these functionality are available at http://braive.vislab.it

BRAiVE has be on display at the last IEEE Intelligent Vehicles Symposium, June 3-5, in Xi'an, China. During this event, the new functionality called 'Stop & Go' has be demonstrated: BRAiVE is able to locate a generic vehicle in front, virtually hook it, and follow it. BRAiVE controls both speed (gas and brake) and steering in autonomous mode. This function is designed to help the driver in queue driving, and brings with it the advangate of improving fuel efficiency as well as providing the driver with a more relaxed driving task.

OFFROAD SENSISNG

In this Section some specific examples of VisLab's vision based sensing system will be presented, again focusing mainly on the off-road application field.

Terrain estimation.

When approaching an unknown environment, one of the most valuable information an unmanned vehicle may gather from its sensing suite is the terrain estimation. It is widely known that with stereo vision it is possible to obtain information about the 3D conformation of the visible scene. Exploiting this feature, VisLab algorithms can provide information about the terrain elevation, slope and drivability.

Stereo Slope Estimation

The first step toward a reliable slope estimation of the terrain in front of the vehicle is, typically, the V-Disparity computation: it consists in calculating the values of similarity (by a correlation measure), for different offset values (disparities), for each pair (left and right) of rows of the stereo images at the same height (*v* coordinate). This operation enables us to produce a 3D graphic structured as it can be seen in Figure 9: the abscissa axis plots the offset for which the correlation has been computed; the ordinates axis plots the image row number; the intensity value is used as a third dimension, and settled proportional to the measured correlation, obtaining an image called V-Disparity image (or correlation image) [6].

Fig. 9. V-Disparity image analysis in a non flat ground scene.

In general this approach requires a terrain with a prevailing slope, and the assumption that ground occupies the largest image part.

Monocular Terrain classification

Another useful information that vision can gather from images taken on front of the vehicle is the classification of the terrain in drivable and not drivable. Most of the algorithms used for path detection look for a single homogeneous road surface in front of the vehicle. To achieve this, the algorithm looks for a variable number of small homogeneous terrain portions, and tries to classify them. Each of these areas 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. It's possible to summarize the 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 usual 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 10 shows the clusters classified as road and off-road [7].

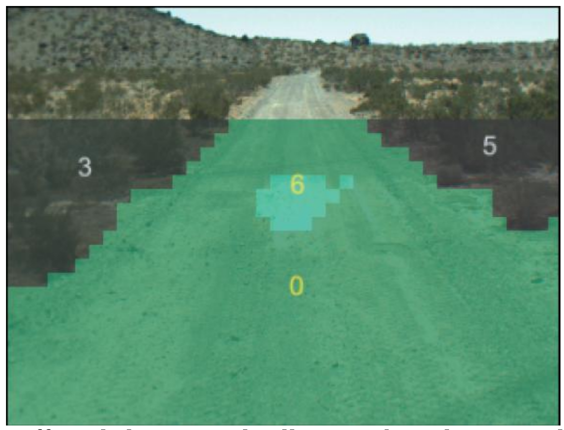

Fig. 10. White numbers denote off-road clusters, and yellow numbers denote road clusters. The interest area is limited to the region within 50 meters from the vehicle. This information is obtained from the preliminary processing of the obstacle-detection module.

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 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.

Stereo 3D Reconstruction

By stereo vision it is also possible to have a global 3D reconstruction of the terrain visible by the cameras. This algorithm is currently underdevelopment and the implementation details are confidential. However it can be disclosed that the system enforces Disparity Space Image and evolutionary computations.
Figure 11 shows an example of off-road terrain estimation.

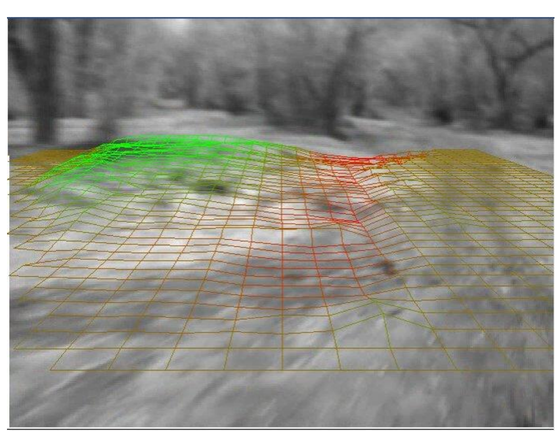

Fig. 11. 3D terrain estimation

Off-road obstacle detection.

Obstacle detection in off-road unstructured environment is particularly challenging since it is not possible to do any assumption on colour, shape, size, of the object to be detected. The only peculiarity of obstacles candidates is their elevation from the ground: any object that emerges significantly from the terrain surface is potentially dangerous for the unmanned vehicle.

VisLab developed several off-road obstacle detector systems, mainly based on 3D stereo reconstruction.

Thin obstacles

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. 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 [8]. We compute the obstacle 3D world coordinates via stereo triangulation. Figure 12 shows an example of this process.

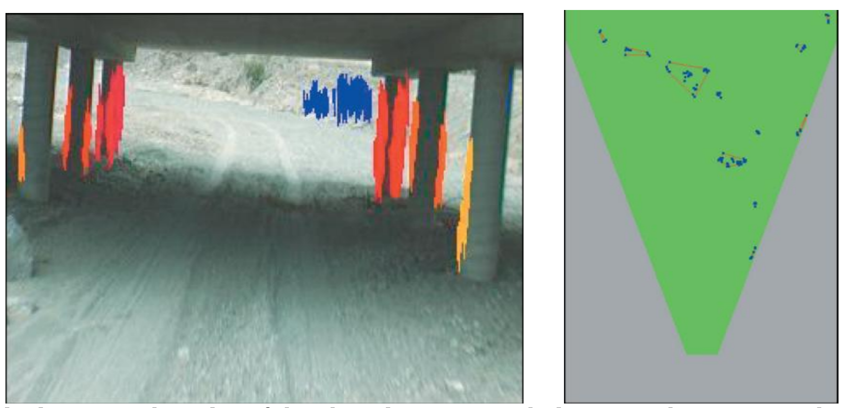

Fig. 12. Obstacle detection: the color of the obstacles varies with distance. The image on the right maps the obstacles in a bird's eye view. The camera's field of view is plotted in green.

THE VISLAB'S FRAMEWORK: GOLD

Effective vision-based systems require developers to consider both hardware and software issues. Hardware issues include the computing engine and the sensing technologies—in vision systems, the cameras. Several camera technologies are available: daylight, far infrared (thermal), near IR, or even range cameras. The choice depends mainly on the specific application and cost-benefit analysis. Software issues include the offline development and debugging software and the final system software. A rapid development tool lets programmers focus on their application's specific vision problem without having to tend to other details.

VisLab has developed GOLD to give programmers such a tool [9]. GOLD includes APIs for common tasks such as I/O operations (acquisition from cameras and other sensors, sensor synchronization, network and file management, data graphical display, and so on) and low-level image-processing functions. A user-friendly GUI supports programming.

GOLD was originally conceived as the main engine to detect lanes and obstacles on board the Argo autonomous vehicle. After 15 years of continuous R&D, GOLD has evolved into a framework aimed at providing a complete set of tools for fast development of computer-vision applications. Moreover, GOLD runs the final onboard system. In spite of this complex functionality, GOLD's user-friendly GUI manages all the system applications. For example, you can visualize the data coming from different sensors, such as a stereo camera system or a laser scanner.

The whole system comprises several subsystem layers, which provide the different functionalities to developers. A cleanly defined interface makes it easy to add or remove subsystems, depending on the target project's specific needs. The framework views all applications as plug-ins that developers can easily remove or connect independently and port to different platforms. The hardware abstraction layers mask the input device complexities from the programmer. Developers can also easily produce highly interactive, consistent user

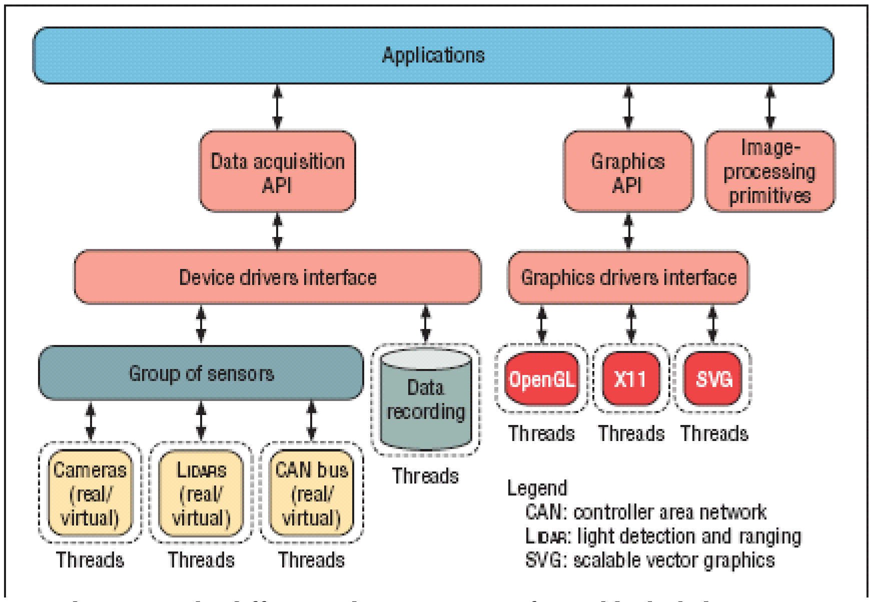

Fig 12. The GOLD architecture. The different colors represent software blocks belonging to same logical level.

interfaces through a set of widgets specifically designed to seamlessly integrate with the processing code. A windowing subsystem provides a simple, powerful set of drawing primitives that developers can use to display intermediate and final results during the debugging phases as well as the final output. Figure 11 gives an overview of the underlying architecture.

Originally developed for Linux/Unix platforms, GOLD has been recently successfully ported on MS Windows and MacOS. The programming techniques applied (OO, open source tools, cross platform libraries) will make easier to port GOLD also to other less general purpoise operating systems and hardware platrforms.

CONCLUSIONS

VisLab has been working on autonomous ground vehicles for more than a decade, being involved in some of the most important and historical events in the field. The experience and knowledge gathered through the years cover basically all the UGV relevant aspect in extreme environment sensing: terrain estimation, terrain classification, path detection, obstacle detection. In particular those experience were made in "out-of -laboratory" condition, where it is necessary to deal with noise, illumination changing, extreme temperatures and where strong real time constrains apply. The more illuminating examples are the 2004 and 2005 DARPA Grand Challenges, as well as the PLS project, that took place in the desert areas of California and Nevada, probably the closest environments to the Moon we have on Heart. This made VisLab one of the most skilled research groups in the world when designing, developing and testing unmanned vehicle with high sensing capabilities and high level of reliability and robustness.

However, this is just the beginning of the VisLab's adventure. Within ERC project VisLab wants to give a strong contribution to the future of unmanned vehicles, proposing new, open and innovative software architecture for better developing and deployment of such systems.

REFERENCES

[1] Alberto Broggi, Massimo Bertozzi, Alessandra Fascioli, and Gianni Conte, *“Automatic Vehicle Guidance: the Experience of the ARGO Vehicle”*. World Scientific, Singapore, April 1999.

[2] Alberto Broggi and Alessandra Fascioli, *Artificial Vision in Extreme Environments for Snowcat Tracks Detection*, *IEEE* Trans. on Intelligent Transportation Systems, 3(3):162-172, September 2002

[3] Deborah Braid, Alberto Broggi, and Gary Schmiedel, *The TerraMax Autonomous Vehicle*, Journal of Field Robotics, 23(9):693-708, Sept 2006

[4] Yi-Liang Chen, Venkataraman Sundareswaran, Craig Anderson, Alberto Broggi, Paolo Grisleri, Pier Paolo Porta, Paolo Zani, and John Beck, *TerraMax: Team Oshkosh Urban Robot*, Journal of Field Robotics, 25(10):841-860, October 2008

[5] Luca Bombini, Stefano Cattani, Pietro Cerri, Rean Isabella Fedriga, Mirko Felisa, and Pier Paolo Porta. *A test bed for Unified Perception & Decision Architecture* In Procs. 13th Int. Forum on Advanced Microsystems for Automotive Applications, Berlin, Germany, May 2009

[6] Alberto Broggi, Claudio Caraffi, Rean Isabella Fedriga, and Paolo Grisleri, *Obstacle Detection with Stereo Vision for Off-Road Vehicle Navigation*, In IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Intl. IEEE Wks. on Machine Vision for Intelligent Vehicles, volume 3, pages 65-72, San Diego, USA, June 2005

[7] Alberto Broggi, Claudio Caraffi, Stefano Cattani, and Rean Isabella Fedriga. *A Decision Network Based Frame-work for Visual Off-Road Path Detection Problem.* In IEEE Intelligent Transportation System Conference (ITSC06), Toronto, Canada, September 2006

[8] Claudio Caraffi, Stefano Cattani, and Paolo Grisleri. *Off-Road Path and Obstacle Detection Using Decision Networks and Stereo Vision.* In IEEE Transactions on Intelligent Transportation Systems, 8(4):607-618, December 2007.

[9] Massimo Bertozzi, Luca Bombini, Alberto Broggi, Pietro Cerri, Paolo Grisleri, and Paolo Zani, *GOLD: A Complete Framework for Developing Artificial Vision Applications for Intelligent Vehicles*, IEEE Intelligent Systems, 23(1):69-71, January-February 2008.