Persistent autonomy: the challenges of the PANDORA project ⋆

David M. Lane [*] **Francesco Maurelli** [*] **Petar Kormushev** [**]
Marc Carreras [***] **Maria Fox** [****]
Konstantinos Kyriakopoulos [†]

[*] *Ocean Systems Lab, Heriot-Watt University, Edinburgh, UK, contact: d.m.lane@hw.ac.uk*
[**] *Advanced Robotics Department, Italian Institute of Technology (IIT), Genova, Italy*
[***] *Depart. d'Arquitectura i Tecnologia de Computadors, Campus Montilivi, Universitat de Girona, Spain*
[****] *Department of Informatics, King's College London, UK*
[†] *Control Systems Laboratory, National Technical University of Athens (NTUA), Greece*

Abstract: PANDORA is a three year project that is developing new computational methods to make underwater robots Persistently Autonomous, significantly reducing the frequency of assistance requests. The aim of the project is to extend the range of tasks that can be carried on autonomously and increase their complexity while reducing the need for operator assistances. Dynamic adaptation to the change of conditions is very important while addressing autonomy in the real world and not just in well-known situation. The key of Pandora is the ability to recognise failure and respond to it, at all levels of abstraction. Under the guidance of major industrial players, validation tasks of inspection, cleaning and valve turning will be trialled with partners' AUVs in Scotland and Spain.

1. INTRODUCTION

Whilst humans and animals perform effortlessly doing complicated tasks in unknown environments, our human-built robots are not very good at being similarly independent. Operating in real environments, they easily get stuck, often ask for help, and generally succeed only when attempting simple tasks in well-known situations. We want autonomous robots to be much better at being autonomous for a long time (persistent autonomy), and to be able to carry out more complicated tasks without getting stuck, lost or confused. Following the Deep Water Horizon disaster in the BP Macondo oilfield in the Gulf of Mexico in 2010, Oil Companies are developing improved ways to cost effectively and safely carry out more frequent inspection, repair and maintenance tasks on their subsea infrastructure. This is particularly challenging in deep water. To date, Autonomous Underwater Vehicles (AUVs) have been deployed very successfully for various forms of seabed and water column transit survey. First commercial units will soon be applied to simple hovering inspection tasks, with future units expected to address much harder intervention where contact is made to turn a valve or replace a component. Because these vehicles reduce or remove the need for expensive ships, their adoption is expected to grow over the next 5 to 10 years.

⋆ The research leading to these results has received funding from the European Union Seventh Framework Programme FP7/20072013 Challenge 2 Cognitive Systems, Interaction, Robotics under grant agreement No 288273 PANDORA

To be successful commercially, these hovering AUVs must operate for extended periods (12-48 hours +) without the continual presence of a surface vessel. They must therefore demonstrate persistent autonomy in a challenging environment. We therefore choose this application focus to evaluate the projects research, with guidance from BP, Subsea7 and SeeByte Ltd. on the project's Industrial Advisory Group. Three essential areas have been identified:

- Describing the World
- Directing and Adapting Intentions
- Acting Robustly

We believe that they are core research areas in which significant advancements is pivotal for Persistent Autonomy. This paper is structured as follow: section II briefly describes the system architecture and the relations between different core fields; section III presents the scenario tasks Pandora is working on; section IV presents the preliminary results of the network; section V presents the validation metrics developed to evaluate the work of the network.

2. ARCHITECTURE

Figure 1 outlines the computational architecture designed for development and study of persistent autonomy. Key is the notion that the robots response to change and the unexpected takes place at one or a number of hierarchical levels. At an Operational level, sensor data is processed in Perception to remove noise, extract and track features, localise using SLAM, in turn providing measurement values for Robust Control of body axes, contact forces/torques

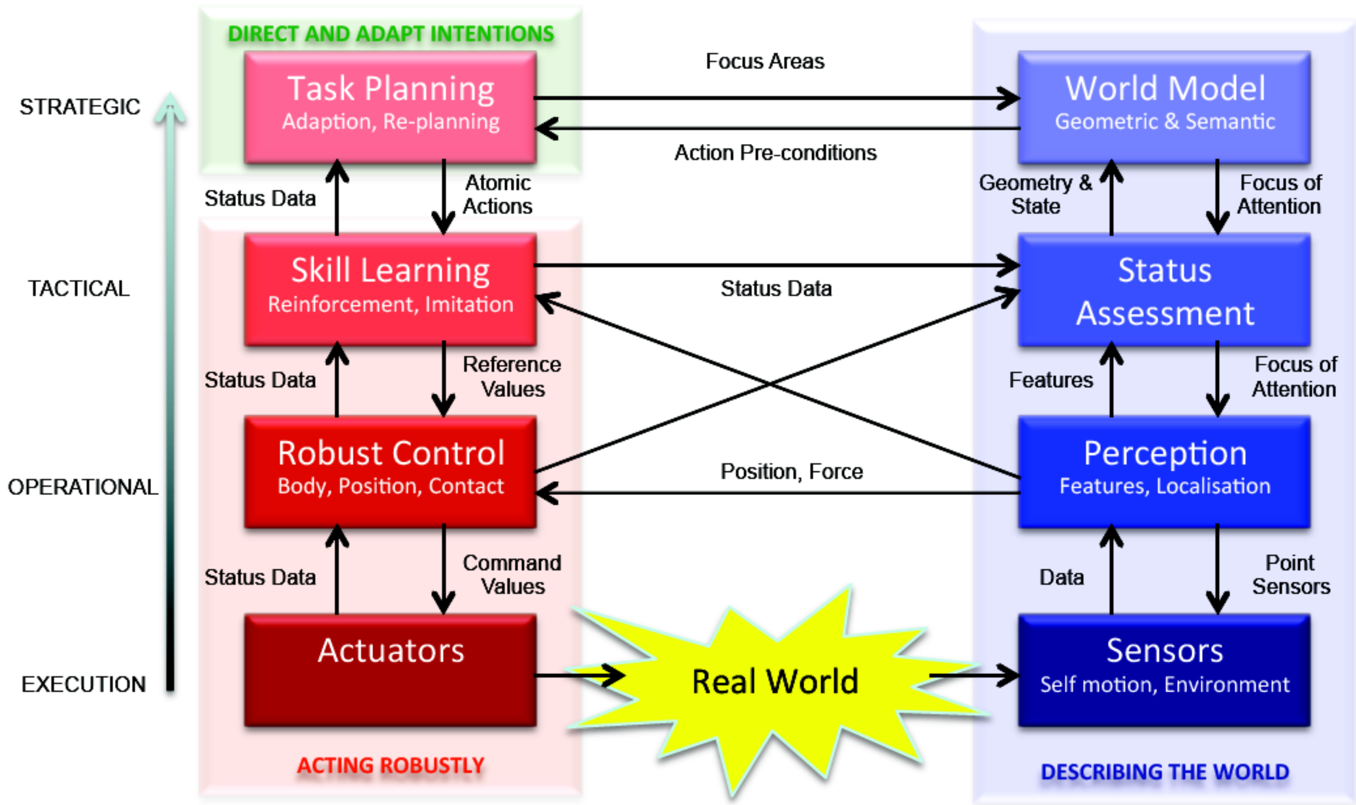

Fig. 1. PANDORA: Computational architecture to develop and study persistent autonomy

and relative positions. One of the goals is to further explore some of the current approaches (Aulinas et al. [2011], Lee et al. [2012]) and integrate them on a real vehicle. In the cases where a map is given, localisation techniques will be used (Petillot et al. [2010]), with a specific attention to active localisation (Maurelli et al. [2010]). Relevant work on robust control can be found in Panagou and Kyriakopoulos [2011], Karras et al. [2011]. At a Tactical Level, Status Assessment uses status information from around the robot in combination with expectations of planned actions, world model and observed features to determine if actions are proceeding satisfactorily, or have failed. Alongside this, reinforcement and imitation learning techniques are used to train the robot to execute pre-determined tasks, providing reference values to controllers. Fed by measurement values from Perception, they update controller reference values when disturbance or poor control causes action failure. The learning block will be lead by IIT, with relevant expertise in the field (Kormushev et al. [2011], Calinon et al. [2010]) Finally at a Strategic level, sensor features and state information are matched with geometric data about the environment to update a geometric world model. These updates include making semantic assertions about the task, and the world geometry, and using reasoning to propagate the implications of these through the world description. Task Planning uses both semantic and geometric information as pre-conditions on possible actions or action sequences that can be executed. When Status Assessment detects failure of an action, Task Planning instigates a plan repair to asses best response, if any. Where there is insufficient data to repair, Task Planning specifies Focus Areas where it would like further sensor attention directed. These are recorded in the World Model and propagated through Status Assessment as Focus of Attention to direct the relevant sensors to make further measurements. Relevant work on Planning has been performed by Fox et al. [2011, 2012].

3. TEST SCENARIOS

The goal of the Pandora project is to use the architecture defined in section 2 for the following three tasks:

3.1 Task A: Autonomous inspection of a submerged structure e.g. a ship hull (FPSO) or manifold (Fig. 2)

A hover capable autonomous underwater vehicle is equipped with a forward looking sonar, a video camera and dead reckoning navigation system. The structure is partially known, but there are inconsistencies between it and the geometric world model. The vehicles high-level goal is to autonomously inspect the entire structure with no data holidays, and bring back a complete data set of video and sonar for mosaicking and post processing. There may be a current running, and the optical visibility may be very poor. In some cases, the sonar inspection sensors must be kept at a constant grazing angle relative to the structure, for best performance. In the absence of a pan and tilt unit, the vehicle must dynamically pitch, yaw and roll to maintain this orientation.

3.2 Task B: Autonomous location, cleaning and inspection of an anchor chain (Fig. 3)

A hover capable autonomous underwater vehicle is equipped as above, but in addition carries a high-pressure water jet.

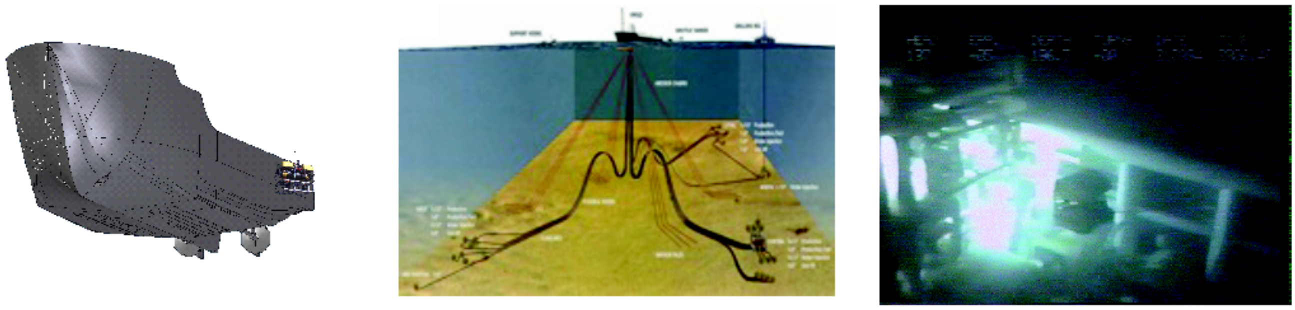

Fig. 2. Task A: Autonomous Inspection of Ship Hull or Subsea Structure

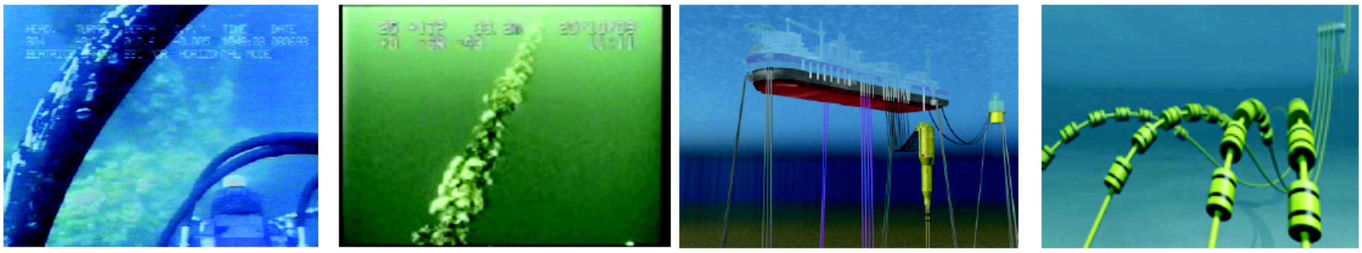

Fig. 3. Task B: (a) Marine growth (b) Anchor Chain (c) FPSO, Anchors and Risers

Fig. 4. Task C: Valve Turning: (a) Docked ROV (b) Hover-capable Inspection AUV Prior to Launch

Its goal is to locate the correct anchor chain of an FPSO and traverse it to remove the marine growth on all sides using the water jet. Thereafter it revisits the chain and brings back complete video inspection data for subsequent post processing. The reaction forces from the water jet introduce significant forces and moments onto the vehicle, and also disturb the anchor chain. Both are therefore in constant disturbed motion. The optical visibility drops to zero during jetting as the marine growth floats in the water. There may be sea currents moving over the anchor, creating minor turbulence downwind of the chain. The chain is located adjacent to flexible risers of slightly larger dimension bringing oil to the surface.

3.3 Task C: Autonomous grasping and turning of a valve from a swimming, undocked vehicle (Fig. 4)

A hover capable autonomous underwater is equipped as in Task A, with a simple robot arm at the front. Its goal is to locate the correct valve panel of a subsea manifold and open the correct valve. On each panel a selection of valve heads are exposed, each with a T bar attached for grasping. The vehicle must identify the state of the valves (open, close, in-between) from the T bar orientations, and if appropriate, use the robot arm to grasp the correct valve and open it. The vehicle does not dock, because there are no docking bars on the panel. It must therefore hover by swimming, counteracting any reaction forces from the turning. It must also ensure that the gripper position and orientation of the gripper after grasping does not cause significant shear forces in the T bar, and break it off. The visibility is generally good, but there may be sea currents running and minor turbulence down current from the manifold.

4. PRELIMINARY RESULTS

4.1 Work Package 1: Describing the World

Efforts have been made mainly on those areas:

(i) **Ontologies**: Ontology representation represents a very good way to represent the reality and the relations among different entities. The work is carried on building on previous work from the OSL lab, like Miguelanez et al. [2011]. Software tools like Protege have been used to represent ontologies and to reason on the world representation. The main challenges encountered are the computational load and therefore the scalability of growing ontologies. Regarding this area, the work of Beetz's group (Tenorth and Beetz [2009]) is certainly very relevant. It looks a very good tool to link sensor data into symbol, and to perform computational methods associated to ontologies.

(ii) **Feature Detection and World Model Update**: sensing the environment and analysing sensor data is crucial for this work package, as the world representation is intimately linked with the reality. Work towards 3D SLAM has shown preliminary results in simulation, with integration with ROS. Some optimisations are required for the integration in the AUVs, but the results are very promising. Regarding SLAM and World Modeling, the issue of updating the World Model according to the sensor data has been raised and discussed. Key factors are related to the decision of updating existing information because inconsistent with the sensor data, and when to take that decision. It is to be noticed that the inconsistency not necessarily means that the information in the World Model is wrong, but it may mean an inconsistency in the vehicle belief.

Additionally, the group has worked on an interface between the world modeling and the planning system (Work Package 2).

4.2 Work Package 2: Directing and Adapting Intentions

The activities in WP2 mainly focused on the following tasks:

(i) **Definition of the link between planning and control layers**: in the initial phase of the project, WP2 focused on the inspection task. In this domain, the high-level component is represented by the planner which has the goal of defining the sequence of waypoints to visit in order to inspect all the surface, while minimising the time and the energy consumption. The low-level component is represented by the controller that is responsible for deciding the physical actions required to move the AUV between different waypoints. This interaction raises an important issue, which is the need for a communication protocol between the planner and the controller. This protocol needs to be defined both from a modelling point of view (defining a model for describing the tasks and the different actions), and from an architectural point of view (implementing the integration of the planner within the ROS architecture). This issue is an important challenge that will be addressed in the next phase of the project.

(ii) **Definition of the models required to describe the different tasks**: PDDL (Planning Domain Definition Language) is the standard language for describing planning domains. As a first result in WP2, it has been agreed that PDDL can be used to model both the missions and the actions of the AUV. In the initial phase of the project, WP2 focused on the inspection task. In particular, the inspection task will be performed in two phases. In the first phase, a low-level inspection of the structure provides the set of relevant waypoints covering the structure. Then, in the second phase, the planner is executed to decide the optimal sequence of waypoints to visit and observations to perform. This domain has been modelled in PDDL. Some features of this domain are the set of waypoints resulting from the low-level inspection and the cost of moving between different waypoints, as well as the observability degree of each waypoint (i.e., which portion of the structure can be observed from each waypoint). The PDDL domain also includes the different actions the AUV can perform. Some examples are: moving between two waypoints (which can be done using a fast navigation or hovering), observing, relocating, turning, etc. Part of this activity has been used to solve a feature-tracking problem, as described in the ICAPS12 paper "Plan-Based Policy-Learning for Autonomous Feature Tracking".

(iii) **Integration of the planning technology within the ROS architecture**: the architecture used in the project is ROS, where all the AUV and sensor simulators are already implemented. WP2 focused on integrating a planner (namely POPF) into ROS. This integration required the implementation of POPF as a node into ROS, and the definition of a set of messages to allow the communication between the planner and the simulators. In particular, an executor has been implemented as well as a well defined protocol for passing as parameters to the planner actions the data received by the sensors. Currently, a preliminary demo is available, where the planner receives in input the domain description (i.e., a set of waypoints, observability degree etc.), and decides the sequence of actions to be performed. Then these actions are executed by the AUV simulator.

Next steps include the improvement of the simulator integration as well as discussions about how ontology can be used to improve the domain description and the communication between planner and controller.

4.3 Work Package 3: Skill Learning for Persistent Autonomy

The ability to learn and adapt in dynamic underwater environment is crucial for achieving persistent autonomy of the AUVs. All three scenarios in the PANDORA project - structure inspection, chain cleaning and valve turning - can benefit from applying machine learning algorithms to increase the AUV adaptability. So far, the research efforts in skill learning have been mainly focused around the third scenario, which consists of autonomous grasping and turning of a valve with a free-floating AUV. In this task, the AUV uses a manipulator to grasp the correct valve on a panel and open or close it. Since the vehicle does not dock, it needs to hover by swimming while counteracting reaction forces from the turning and from the sea currents and even minor turbulence. Also it must ensure that the gripper position and orientation of the gripper after grasping does not cause significant shear forces in the valve handle (T bar shape), and break it off. Initial experiments are being conducted with a commercial manipulator (KUKA Light Weight Arm) in lab conditions, not underwater. In the current experiments, the valve position is being estimated from the sensory input with an Extended Kalman Filter. Imitation learning approach is used in order to learn the trajectory to follow with the robotic arm. To perform the task safely, a fuzzy system is developed which generates appropriate decisions for the arm movement in real-time according to the instant dynamics of the valve. The ongoing research in skill learning focuses on achieving multiple important desired properties of movement planning, which are: ease of representing and learning, compactness of the representation, robustness against perturbations and changes in a dynamic environment, ease of reuse for related tasks and easy modification for new tasks, and ease of categorization for movement recognition. Regarding the first scenario in PANDORA, where an AUV has to inspect a submerged structure under potentially unknown environment conditions, the controller can be adapted in real-time using machine learning methods. This task can be considered as a problem of searching in policy space for episodic reinforcement learning. The ongoing research in this direction focuses on developing derivative-free optimization methods for this scenario.

4.4 Work Package 4: Robust Control Strategies for Efficient Positioning and Interaction

A state feedback control approach for the motion of 3D under-actuated UUVS was developed. The objective is to steer the vehicle to a desired configuration w.r.t. a target of interest. A vision-based sensor system is considered, thus imposing configuration constraints due to the limited field-of-view (FOV) of the onboard camera. The control design is based on dipolar reference vector fields and tools from viability theory, and guarantees the convergence of the vehicle in a neighborhood of the desired configuration, without violating the FOV constraints. Moreover, the non-actuated DOFs are shown to be input-to-state stable. An alternative approach to the aforementioned problem was proposed based on Model Predictive Control (MPC). The approach combines the notion of dipolar vector fields along with a constrained nonlinear MPC formulation, which incorporates the visibility constraints. The proposed control strategy falls into the class of dual-mode MPC schemes, i.e. the system trajectories are forced by the model predictive controller into a suitably defined terminal region that contains the goal configuration. Therefore, convergence of the system trajectories to the goal configuration is guaranteed by switching to the dipolar vector field based controller once in the terminal region. The problem of autonomous inspection of an anchor chain has also been dealt for under-actuated underwater vehicles. The dynamics and the kinematics of the UUV as well as exogenous disturbances representing ocean currents and waves were included in the analysis. Employing the Prescribed Performance Control technique, a robust time-varying control scheme is designed incorporating visibility constraints as well. Moreover, the non-actuated DOF (sway) is shown to remain bounded. Finally, the off-line identification problem of the dynamic parameters of AUVs was approached. Relevant information regarding the sensor suite of *Girona 500 AUV* has been gathered and various technical issues have been tackled. Subsequently, a way-point tracking controller for *Girona 500 AUV* has been designed. It utilizes the estimated parameters from the off-line identification procedure as well as a parameter tuning algorithm to increase the closed loop performance.

4.5 Work Package 5: Testbeds Integration and Experimental Validation

The experimental validation has started by developing an augmented reality framework that will be used to test the systems in simulation before using the real test-beds. Simulator and experimental testbeds are available with the consortium, like for example Davis and Lane [2010]. *Girona 500 AUV* from UdG and *Nessie VI AUV* from HWU will be used in the 3 scenarios of the project: a) structure inspection, b) chain cleaning and c) valve turning. The simulator has been already set up and allows the use of both AUVs in a first simplified version of the scenarios. All the software is based on the ROS framework, and this allows each partner to test and integrate their systems in a cooperative and continuous way. In the following months, the 3 scenarios will be partially solved with the simulator. As far as real test-bed experimentation, the work has consisted till now in upgrading the vehicles with new equipment. Both vehicles are right now completely operative for pre-programmed autonomous missions and, therefore, once the new equipment is integrated, the three tasks will be able to be tested with the new methods developed in the project. *Nessie VI AUV* will be used in the structure inspection task. A multibeam imaging Forward Looking Sonar (FLS) will be integrated for detecting the structure and performing the correct inspection movements accordingly. *Girona 500 AUV* will be used in the chain cleaning and valve turning tasks. For the first task, also a FLS will be integrated for detecting the chain with zero visibility and generating the trajectories for cleaning the surface with a high pressure water jet, that will also be integrated. The valve turning tasks will require the integration of a new underwater manipulator that will allow the operation of the valves from the AUV in free-floating mode. All the equipment has already been chosen and it is planned to have it completely integrated by mid 2013, following the plan of the project. Some of the demonstration platforms are showed in Fig. 5.

5. VALIDATION

A comprehensive list of validation metrics have been developed to test the ability of each system to meet its criteria. Ultimately the best all-encompassing metrics of persistent autonomy is quantitative analysis of robot engineer interactions, which will be greatly reduced by PANDORA's success. So overall, we measure our success by the reduction in the number of times the operator is called to assist a stuck robot during execution of tasks and sequences with noisy sensor data. We also count the number of successful automatic recoveries the robot achieves from an execution failure. These require world modelling to detect the failure, task planning to indicate corrective action, and reinforcement learning/adaptive control to successfully execute once more.

Beyond this, the overarching performance indicators within each core theme are:

- Describing the World:
    - · Trends in numerical errors of vehicle location and object location/geometries (or other indirect measures such as residuals, covariance, OSPA

Fig. 5. Demonstration Platforms: (a) top left: HWU Nessie VI AUV; (b) top right: UdG Girona 500 AUV; (c) HWU Augmented Reality Testbed

error, Hellinger distance), and probabilities in semantic relationships (ontologies).

  · Number of correct and incorrect diagnoses of task execution failure.

- Directing and Adapting Intentions:
  Comparison of the number of failing actions in an automated planning/re-planning system with the number of failures produced by hand-built plans, assuming given world information (efficiency and resource utilisation are therefore implicit)
- Acting Robustly:
  Position and orientation error norm per unit distance, and the average wrench (force and torque) error norm.

6. CONCLUSION

This paper has presented the challenges which the FP7 Project PANDORA is currently addressing, focusing on persistent autonomy. Current existing autonomous systems require frequent operator intervention. The focus of Pandora is to enhance the long-term autonomy of AUV missions, through increased cognition, at all the levels of abstraction. An agile development approach is used to allow frequent measurements of development metrics to rapidly optimize the system on the three main experimental scenarios. The preliminary results after only a few months from the start of the project are very encouraging.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Union Seventh Framework Programme FP7/20072013 Challenge 2 Cognitive Systems, Interaction, Robotics under grant agreement No 288273 PANDORA.

REFERENCES

Josep Aulinas, Marc Carreras, Yvan R. Petillot, Joaquim Salvi, Xavier Llado, Rafael Garcia, and Ricard Prados. Feature extraction for underwater visual slam. In *IEEE/OES OCEANS 2011*, 2011.

S. Calinon, F. D'halluin, E. L. Sauser, D. G. Caldwell, and A. G. Billard. Learning and reproduction of gestures by imitation: An approach based on hidden Markov model and Gaussian mixture regression. *IEEE Robotics and Automation Magazine*, 17(2):44–54, 2010.

B. C. Davis and D. M. Lane. Guided construction of testing scenarios for autonomous underwater vehicles using the augmented-reality framework and javabeans. *Journal of Engineering for the Maritime Environment: IMechE Vol. 224 Part M*, 2010.

Maria Fox, Derek Long, and Daniele Magazzeni. Automatic construction of efficient multiple battery usage policies. In *Proceedings of the 21st International Conference on Automated Planning and Scheduling, ICAPS 2011*, 2011.

Maria Fox, Derek Long, and Daniele Magazzeni. Plan-based policy learning for autonomous feature tracking. In *Proceedings of the 22nd International Conference on Automated Planning and Scheduling, ICAPS 2012 (to appear)*, 2012.

George C. Karras, Savvas G. Loizou, and Kostas J. Kyriakopoulos. Towards semi-autonomous operation of under-actuated underwater vehicles: sensor fusion, on-line identification and visual servo control. *Auton. Robots*, 31(1):67–86, 2011.

Petar Kormushev, S. Calinon, and D. G. Caldwell. Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input. *Advanced Robotics*, 25(5):581–603, 2011.

Chee Sing Lee, Daniel E. Clark, and Joaquim Salvi. Slam with single cluster phd filters. In *Proceedings of the 21st International Conference on Robotics and Automation, ICRA 2012*, 2012.

F. Maurelli, A. Mallios, S. Krupinski, Y. Petillot, and P. Ridao. Speeding-up particle convergence with probabilistic active localisation for auv. In *IFAC IAV*, 2010.

Emilio Miguelanez, Pedro Patron, Keith E. Brown, Yvan R. Petillot, and David M. Lane. Semantic knowledge-based framework to improve the situation awareness of autonomous underwater vehicles. *IEEE Transactions on Knowledge and Data Engineering*, 23:759–773, 2011. ISSN 1041-4347. doi: http://doi.ieeecomputersociety.org/10.1109/TKDE.2010.46.

Dimitra Panagou and Kostas J. Kyriakopoulos. Control of underactuated systems with viability constraints. In *Proc. of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 5497–5502, Orlando, Florida, December 2011.

Y. Petillot, F. Maurelli, N. Valeyrie, A. Mallios, P. Ridao, J. Aulinas, and J. Salvi. Acoustic-based techniques for auv localisation. *Journal of Engineering for Maritime Environment*, 224(4):293–307, 2010.

Moritz Tenorth and Michael Beetz. KnowRob – Knowledge Processing for Autonomous Personal Robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4261–4266, 2009.