

Persistent Autonomy

Artificial Intelligence or Biomimesis?


David M Lane
Ocean Systems Laboratory
Heriot-Watt University Edinburgh, Scotland, UK
D.M.Lane@hw.ac.uk

Abstract—We seek to develop autonomous robots that can operate and interact unsupervised for extended lengths of time in unknown environments, adapting their purpose in response to events and goals, learning from successes and failures, recovering from errors in execution whilst monitoring and maintaining self health. Such persistent autonomy is a challenging ambition, and the subject of an increasingly intense research effort. Two broad approaches have evolved, one rooted in artificial intelligence research from the 1970s onward, and the other in studies of animals and even plants that nature has evolved over millennia. Both offer opportunities and challenges in implementation. This paper presents a snapshot of recent and ongoing developments from each approach, and offers some perspectives on the potential that each offers.

Index Terms—autonomous underwater robot, artificial intelligence, biomimesis, bioinspired, architecture, persistent autonomy, world model, ontology, skill learning, adaptive planning, soft robot, fish swimming, flow sense, electric sense, biosonar.

I. INTRODUCTION

Autonomous underwater robots generally succeed best when carrying out well-defined tasks in situations that are well known. Their current applications in various kinds of remote survey exemplify this, requiring them to follow trajectories, or servo around objects. Exceptions, such as obstacles or objects of interest, are handled by switching pre-programmed scripts (behaviours) with appropriate parameters [1,2].



Whilst these work well, there are a new generation of roles envisaged for these service robots that will place them into situations that are different. These new applications and environments may not be precisely specified or known, and situations may not have been previously encountered. Robots must also act autonomously for longer periods of weeks, months or years without failing or becoming confused. This

requires greater flexibility in the way the robot chooses to sense, act and remember in achieving goals, including an ability to identify when it has not been successful executing one or a series of actions, and finding a way to recover within time and energy budgets.

For example, in the oilfield, use cases and business justifications have been mapped that are motivating autonomous inspection vehicles that can not only survey and inspect oilfield infrastructure [3], but eventually will be able to interact unsupervised, to turn valves or replace components. For environmental monitoring, autonomous platforms that can detect interesting events and adapt their mission plans on the fly in scientifically meaningful ways would benefit the scientists' evaluations [4]. And in naval security [5], scenarios involving multiple vehicles that can co-operate and work synergistically, sharing information as communication bandwidth allows, re-planning missions in real time responding to fresh events and instructions are being tested [6].

State of the art autonomous robots typically seek human operator assistance when they recognize they are in situations beyond their logic or their sensing abilities. Such recourse often imposes unwanted operational constraints that limit the use-cases and the business models that motivated their use in the first place. To deal with this we therefore seek to equip our autonomous robots so that they can:

- operate successfully while unsupervised and without recourse to a human operator,
- for extended lengths of time,
- in environments that are partially or completely unknown,
- adapting their purpose in response to events and goals,
- whilst interacting with the environment,
- learning from successes and failures and
- recovering from errors in task execution
- while monitoring and maintaining self health.

We call this *Persistent Autonomy*.

Persistent Autonomy requires developments in several areas, not all mutually exclusive. Smart management of energy consumption and storage, linked to energy harvesting can extend endurance. New forms of acoustic, optical and electric sensing appropriately directed can improve perception. New developments in data compression and acoustic, optical or e/m telemetry and networks can extend communication. But the most challenging area is that of cognition – what approach should be taken to linking sensing and acting so that the

autonomous robot goes beyond switching pre-programmed behaviours designed for specific circumstances? What computational approaches can we take to realize a form of primitive sentence that can reason and adapt in new circumstances, responsive to the local environment, and able to interact with an operator?

The established approach is to use well understood computational architectures, that use traditional Artificial Intelligence based probabilistic, machine learning and related approaches [7]. These can prove to be computationally intensive, and can depend on significant quantities of training data and time for success. Alternatively, and enjoying increased attention, are *Biomimetic* approaches using biological principles from studies of brain, mind and body to copy nature in the design and engineering of materials and machines. Similarly, *Bioinspired* approaches adapt the principles observed in nature into pragmatic designs that are blended with traditional engineering methods. Finally *Biohybrid* systems physically interface biological (i.e. living) and artificially engineered systems [8]. These approaches are motivated by the realization that animals such as ants with as few as 250,000 neurons can achieve impressive feats of chemotaxis, navigation, turbulence and flow sensing/control. It is argued that evolution has captured key neuronal and embodiment structures in the brains and bodies of animals across species, and that foundational research in neuroscience and neuroethology (the study of animal behavior and its underlying mechanistic control by the nervous system) are providing new insights into what these might be [9]

In this paper we will look more closely at some latest developments in these two approaches to cognition, and reflect on the implications for the development of the next generations of autonomous underwater robots.

II. ARTIFICIAL INTELLIGENCE APPROACH

A. Architecture

Figure 1 illustrates a contemporary AI based architecture under development for the study of persistent autonomy for AUV inspection, repair and maintenance (IRM) in the oilfield [10]. As is traditional, the robot's response to change and the unexpected takes place at one or a number of hierarchical levels.




Fig. 1. PANDORA: Computational architecture to develop and study Persistent Autonomy

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 and relative positions. 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 set 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. 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 assess 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.

Perception and Control are much covered subjects in the community, and we will therefore consider only the direction of World Modelling, Task Planning and Skill Learning in a little more detail

B. World Modelling and Status Assessment

World Modelling comprises a geometric representation to model the spatial distribution of objects, and a corresponding semantic representation to store their relationships that are important from the point of view of planning, re-planning to recover from failure, and affordances (i.e. actions that can be applied to the object). Explicit use of such semantics is at the forefront of persistent autonomy robotics research. To represent the semantics of the world and the actions the robot can perform, ontologies are used as a natural development from its use in the semantic web and language processing (fig 2). Ontologies are formal representations of concepts in a domain, and the relationships between those concepts. Examples of concepts could be 'AUV' and 'ForwardLookSonar', and a relation 'hasSensor' could be used to relate the two; 'AUV hasSensor ForwardLookSonar' [11,12]. Their power comes from the ability to store large numbers of such semantic constructs, and infer new information, generally in response to queries.

Such semantic representations are needed as the means to assert world state information required by Task Planning. Recent developments by Tenorth [13] based on long standing efforts of Lenat [14] have fuelled an acceleration of this work applied to robotics [15]. Traditional issues of symbol grounding are addressed by directing sensors as needed to

investigate the world. Ontologies can be made probabilistic for fusion and to represent uncertainty [16], and can evolve automatically. A plan failure is analysed and used to diagnose ontology misalignment and to suggest an ontology repair. [17]

Status assessment uses data from the various strategic, tactical and execution layers of the architecture's Acting hierarchy to determine if a task has failed in execution. If so, this is propagated through the World Model to the Task Planner to realise actions to recover from the failure. This may involve diagnosis to understand why failure occurred and to inform future planning. The detail of detecting task failure is very system specific, involving identifying the out of bounds parameters (e.g. controller errors, probabilities in execution of learned tasks).




Fig. 2. Ontology classes for describing spatial, temporal and mathematical things (from [13], [14])

Whilst there is a substantial literature on the subject of diagnosis in general, in practice it reduces to the specification of an ontology of fault conditions that can be reasoned with as status data changes during task execution. Previous work inferring facts about state and mode, from noisy systems based on data has been carried out in model-based reasoning community [18]

C. Task Planning

Task Planning sequences actions on the vehicle through skill learning or directly to other lower level behaviours. It adds a level of sophistication to the simple switching of behaviours in a traditional autonomy architecture. Actions can be fine grained or coarse, and are made runnable according to assertions in the world model ontology. Actions can fail, new information can be asserted through the world model, and so planning must be adaptive i.e. plans can be reformulated on the fly. This is distinct from path planning [19], which is waypoint based, and for which numerous techniques exist. They become computationally intensive when based on probabilistic methods in 6 degrees of freedom [20].

Whilst the international community has been very active in the development of classical planning algorithms and models, there has been little development of such adaptive planning approaches. A planning system is adaptive if it can respond to execution conditions to re-plan when failures or opportunities

occur [21-23] Within a forward search framework hindsight optimisation can be used to select the most likely best next action, and re-planning to respond to unexpected failures in execution. Adaptive planning must:

- allow the removal and insertion of plan fragments to modify a plan under execution. Removal of fragments will occur when failure of an action renders parts of the plan useless. This will free-up resource which can be re-used
- support the insertion of plan fragments to use up released resources in ways that allow the system to take advantage of “opportunities”. Plan fragment insertion and removal represent a form of re-planning that exploits existing plan structure and reduces the waste incurred when a plan has to be completely replaced
- modify the plan while preserving high confidence that the plan will achieve its highest priority goals.

The computational mechanisms involved in achieving robust plan modifications are constraint reasoning and convolution of the distributions representing probable resource demands of actions (this can be done analytically in simple cases, or by Monte Carlo methods for more complex distributions).

D. Skill Learning

Skill learning provides greater flexibility and robustness in the execution of actions, as well as a means to teach the robot what an action is. Rather than represent the action as series of explicit moves or waypoints, probabilistic encoding is used. Probabilistic parameters can be learned when an operator manually teaches the robot through repeated demonstration. This has the advantage that the robot can subsequently adapt to deal with variations in the action caused by noise or perturbation. Such a skill reduces the likelihood of task failure due to small errors in positioning or contact, due to calibration errors, sensor noise or transient responses in low level control.

Hidden Markov Models (HMMs) are a well-studied tool with the abilities to do both generation and recognition of continuous motion sequences [24]. It allows the consideration of both spatial and temporal variabilities across various demonstrations. A multitude of variants of HMM exist, which is appropriate for the various problems that a network of dissimilar robots needs to face at the subtask level in a noisy real-world environment. As the approach is probabilistic and relies on well-founded learning algorithms, it offers numerous possibilities of adaptation and combination with other machine learning techniques.

Reinforcement Learning (RL) is the general designation for methods and algorithms that learn Markov Decision Process (MDP) solutions, i.e., state-action maps (policies) that maximize the infinite horizon discounted cumulative rewards. Recent challenges for reinforcement learning concern handling large state spaces and their inherent complexity [25]. However, RL has come a long way from the initial small discrete toy problems to complex continuous state-action space real-world problems. Recent advances in Expectation-Maximization based RL (EM-RL) have shown rapid learning on high-dimensional problems [26]. Unfortunately, these algorithms are not

guaranteed to find the global optimum and may end up in some sub-optimal local optimum, or alternatively, the optimum they find might be quite unstable and therefore sensitive to noise.

E. Persistent Autonomy Through AI

Architectures such as Fig 1 offer the prospect of realizing the cognition desired for persistent autonomy. Control stabilizes, linearizes and tunes the robot motion and contact conditions during task execution. Skill learning provides the means to respond to disturbances and errors within the behaviour, without recourse to planning. Task planning provides the means to adapt behaviour in response to execution failure at a lower level, new observations in the environment from sensors or health issues on the vehicle (component failure, energy usage). World modeling processes sensor data to maintain a logical description of the world, while directing sensors according to the task, asserting the logical primitives and probabilities that are the preconditions to trigger actions in the Task Planner. Status Assessment monitors the performance of action execution as well as live sensor data and the state of the world to detect and diagnose failures of the task and on the vehicle. Implementing this functionality presents several practical challenges:

- A large amount of software is involved, generally from different teams. This requires a significant skilled manpower to implement, expert in different branches of sensor processing, control, computer science and so on.
- More code increases the likelihood of bugs, and errors interpreting interface specifications if not well defined. This requires careful coding and testing during development and system integration, followed by rigorous testing in the target environment
- Training and tuning are required where probabilistic methods are involved. This requires both accesses to data and time in the field for training to converge. Such training delays kill operational effectiveness, and lead to the system being unused unless essential
- Special purpose computing resources may be required (e.g. access to a GPU board) adversely affecting heat generation, space use and power consumption.

III. BIOMIMETIC APPROACH

A. Architecture

Large-scale efforts by the neuroscience and modelling research communities have resolved the broad connectivity and functioning of significant parts of the brain to determine an understanding of its architecture as a series of nested sensorimotor loops [27] (figure 3)

Beyond the classical work of Brooks Subsumption Architecture [28], efforts to build robotic models of brain-like structures abstracted from the architecture of neurons and synapses are typified by the more sophisticated Distributed Adaptive Control Architecture (DAC) [29]. DAC consists of three, tightly coupled, layers for behavioral control; the reactive, adaptive and contextual layers (Fig 4), each

performing increasingly more memory dependent mappings from sensory states to action.




Fig. 3. Global architecture (simplified) of the mammalian brain showing key nuclei and connections for sensing and control. It can be viewed as a series of nested sensorimotor loops. From [27].




Fig. 4. DAC implementation of a synthetic foraging ant. From [23]

The reactive layer (bottom) performs reflex actions including exploration, collision avoidance, chemical search, homing, etc. using a range of sensory modalities and predefined sensori-motor mappings. During foraging, the adaptive layer performs landmark recognition, feature extraction, heading direction computation (HDA) and constructs memory segments for each encountered landmark comprising heading direction and landmark information. These memory segments are sequenced in the short-term memory (STM) of the contextual layer until a goal state occurs, e.g.

feeder detection. At this time the contents of STM are retained into long-term-memory (LTM). During a recall phase (homing, landmark navigation, exploration), the segments of LTM are matched against current sensory events and an optimal trajectory is computed from the heading directions suggested by the recalled LTM segments.

This architecture has successfully been implemented on modest general purpose computing resources using the large-scale neural simulator IQR [30] for mapless navigation and foraging exhibited by ants and moths. Using the definition of section I, it implements a form of Persistent Autonomy. The difference with the oilfield IRM application of section IIA is one of task complexity. Tasks such as water jet cleaning, grasping and turning valves, connector insertions and assembly require more than landmark recognition, storage and tracking. However, coupled to appropriate low level robot control, the principles of training to pre-record and store successful action segment sequences in a contextual layer LTM may inspire workable solutions using sensing primitives beyond landmarks. The same mechanisms used online may assist error recovery caused by unexpected events or situations during execution.

B. Embodiment

A central tenet of the biomimetic approach is a move away from the notion of a central controlling system upon which all sensors, actuators and limbs depend. So called *morphological computation* outsources computation to the body and the environment, leading to the notion of *soft robotics*. Here, the material properties of the robot body effectively carry out computation through their properties of elasticity, viscosity and compliance.

Nowhere is this truer than in the aquatic environment. The existence proof of the swimming modes [31] of fish, cetacians (dolphins, whales), pinnipeds (sea lions, walruses, seals) and cephalopods (squid, octopus) show us that agile, efficient and low drag propulsion can be obtained from relatively simple patterns of actuator motion played out through variable stiffness end effectors such as or muscular hydrostats (tentacles) or caudal fins (fish tail) with variable stiffness [32] (Fig 5).




Fig. 5. Variation of caudal fin stiffness with length (from [32])

Analysis and experiment in various outputs from [33] and [34] have demonstrated through simulation and experiment that fish use flow as a source of energy, particularly when swimming in a karmen vortex sheet (KVS). This is a particular

form of turbulence found behind bluff bodies or when swimming in diamond formation (figure 6).




Fig. 6. Karmen vortex street formation and hydrodynamic benefit derived from swimming in a diamond formation (from [34])

The variable stiffness of the fish body linked to the period and amplitude of its swimming gait can extract energy from the vortices as they pass alternately down the left and right sides of the fish body. Tuning of the gait can affect energy efficiency by up to 25%. In the limit, the fish can become more than 100% efficient. It has been demonstrated by experiment [35] and in simulation [34] that a dead fish can passively swim upstream in a KVS as a result of it's soft embodiment.




Fig. 7. Snapshots of a dead fish swimming passively in karmen vortex street (from [34])

For persistent autonomy therefore, embodiment and soft actuation is an area where our growing understanding and our ability to fabricate soft autonomous vehicles can lead to pay offs in energy harvesting, simplified control, and reduced energy consumption for motion, extending the robot's mission endurance.

C. Sensing

Three areas of biomimetic and bioinspired underwater sensing are providing insights that can open new capabilities for persistently autonomous robots.

1) Flow

In the same way that flow can be a source of energy for an underwater robot, in can also be a source of information.



Fig. 8. Flow as a source of information (courtesy [33])

Fish have sophisticated lateral line sensors that enable them to detect small changes in pressure along the length of their body, and therefore indirectly flow.




Fig. 9. Artificial fish lateral line measures pressure and flow (from [33])

Because they are moving in an incompressible fluid, they use this to detect the presence of adjacent objects (e.g. neighbouring fish in a shoal). Fish are also known for their exceptional odour tracking capabilities. The very few existing studies suggest that this is partially achieved through multimodal sensing, including flow directionality. It has been hypothesized that as the Reynolds number grows, the animals rely more on flow information that gives them information about the varying directionality of the odour trail and fills in the gaps of odourless regions between the eddy trails. Some recent studies on shark odour detection confirm this hypothesis [36]. If flow features are stable therefore, they can be used as landmarks for navigation and path following in the same way optical features are used by ants. New developments in MEMS lateral line sensors [33] are enabling greater experimentation with sensory-motor co-ordination, and can provide the basis to validate through experiment the ways in which flow can assist situational awareness underwater.

2) Electric Sense

Several species of fish as well as sharks have been demonstrated to use electric sense over both short (cm) and long (m) ranges according to their environment. Experiments have demonstrated that fish use different types of electric organs and electroreceptors not only for detection of adjacent objects and fish, but also for communication in both passive and active modes. By adjusting their body shape they adjust the field so as to image objects, and their communication strategies work so as to avoid jamming when swimming in shoals.




Fig. 10. The African Mormyrid fish *Gnathonemus petersii*, and a top view of the fish basal electric field. From [37]

Recent work [34], [38] has developed multi-physics simulators and robotic testbeds to enable study the morphology of the sensor and the physics of electric interactions for navigation in

confined environments. Taking inspiration from passive electro-location in real fish, a set of reactive control laws encoding simple behaviors such as avoiding any electrically contrasted object, or seeking a set of objects while avoiding others according to their electric properties, is developed. In water prototypes of sensors and swimming robots have been demonstrated, including behaviours such as wall following, servoing, docking and obstacle avoidance.

Electric sense has advantage over optical sensing, as it is not affected by poor visibility or disturbed sediment. Compared to acoustic methods it is not subject to diffraction or shadowing and has significantly faster propagation times and bandwidth potential. Fish can therefore ‘see’ round corners and make rapid friend, foe or food decisions. One can therefore imagine ways electric sense could be used on the oilfield, for example to characterize the state of sacrificial anodes or to determine operational state of a previously mapped wellhead, manifold or other electrical device without contact.




Fig. 11. Multi-robot docking using electric sense and permanent magnets to form an anguilliform swimmer (from [34])




Fig. 12. Electric sense: sensor testbed & wall following behaviour (from [34])

Design challenges remain in current prototypes to achieve operational ranges of 10s or 100s of metres. Key is the design of sensors and receiver electronics that have sufficient sensitivity to detect weak fields from a distance, with shielding to remove effects of electric fields naturally occurring on the robot, particularly from any electric motors.

3) Biosonar

Cetaceans such as Dolphins have developed sonars able to recognize targets at 100s of metres in noisy and reverberant environments. Despite having mediocre equipment (from the sonar engineering perspective) and working at relatively low frequencies lacking resolution, they are able to make fine discrimination of wall thickness, the difference between water-filled cylinders, material differences in metallic plates, and to discriminate and recognise species of fish food. The high temporal resolution of the biosonar signals along with the high dynamic range of its auditory system are critical factors for target discrimination.

Recently, [39], advanced time-frequency signal analysis techniques have been used to generate new models for bio-inspired sonar signals. The inspiration comes from the analysis of bottlenose dolphin clicks. These pulses are very short duration, between 50 and 80 μ s, where a double down-chirp structure has been delineated using fractional Fourier methods.

The majority of clicks have energy distributed between two main frequency bands with the higher frequencies delayed in time by 5–20 μ s. Signal syntheses using a multiple chirp model based on these observations are able to reproduce much of the spectral variation seen in earlier studies on natural dolphin echolocation pulses. Analyses of the detailed echo structure for these pulses ensonifying two solid copper spherical targets indicate differences in discriminatory potential between the signals. It is suggested that target discrimination could be improved through the transmission of a signal packet in which the chirp structure is varied between pulses. Evidence that dolphins may use such a strategy themselves comes from observations of variations in the transmissions of dolphins carrying out target detection and identification tasks.




Fig. 13. Dolphin multi-chirp click variation in time-frequency from one animal analysed using fractional Fourier method [39]. Clicks contain a double down chirp in the 30-130KHz range, exhibiting fairly constant low frequency but variable high frequency components.

More recent efforts [40] have developed commercial wideband sonar prototypes and evaluated them for tasks including AUV based cable detection and tracking, and classification of man made objects such as mines. Modelling and experimental studies have demonstrated the way these signals can penetrate objects and excite characteristic resonances that form the basis for their identification. Their classification performance is significantly superior to the normal highlight and shadow methods characteristic of sidescan and synthetic aperture imaging sonars. Bioinspired sonar designs therefore can be expected to have an important operational role in security and offshore energy applications in the future.

D. Persistent Autonomy Through Biomimesis

The biomimetic and bioinspired methods described above are showing great promise in developing new forms of sensing, embodiment, control and architecture that can have a lasting impact on the way we engineer our autonomous robots for the ocean. The promise of simpler software architectures, efficient propulsion, energy harvesting and higher fidelity sensing all of which use the water as a medium to the advantage of the robot

are offering new horizons for the development of smaller, smarter and more persistent systems.

IV. CONCLUSIONS

Technology readiness levels (TRL) describe a scale from 1 to 10 that defines the journey of a technology from an initial idea (0) to an operational system in use in the field (10). AI approaches to autonomy have been in development for decades, and it is not surprising therefore, that currently fielded systems with high TRL are largely based on simplifications of the AI approach – tri-level software architecture, pre-defined behaviours, ATR for landmark detection in SLAM navigation and obstacle avoidance, conventional control systems design through linear model approximation and controller response shaping in time or frequency. Consequently, the AI approach to persistent autonomy of section II is also at a higher TRL, offering an easier development path to implementation. On the other hand, the biomimetic and bioinspired approaches of section III are yielding some fundamental insights that could lead to more profound developments – soft swimming robots, energy harvesting, navigation based on electric sense and flow. Although these are currently low TRL, this will change in time as research and eventually application follows. Since natural selection will operate on the methods becoming those of choice in the future, the most likely outcome will be artificial intelligence AND biomimesis in realizing a more effective persistent autonomy.

ACKNOWLEDGMENT

This work is supported by the European Union Seventh Framework Programme FP7/20072013 Challenge 2 Cognitive Systems, Interaction, Robotics under grant agreement No 288273 PANDORA. Thanks to colleagues and associates in the FILOSE, ANGELS, BIOTACT and SYNTHETIC FORAGER projects for discussions and materials. In particular Maarja Kruusmaa, Paul Verschure, Tony Prescott, Frederic Boyer, Maria Fox and Petar Kormushev offered valuable insights.

REFERENCES

- [1] P. Newman "MOOS: Mission Oriented Operating Suite"
- [2] J.Evans, P.Patrón, B.Smith, D.M. Lane. "Design and evaluation of a reactive and deliberative collision avoidance and escape architecture for autonomous robots" *Autonomous Robots* Volume 24, Number 3 (2008), 247-266,
- [3] J.A. Mair, J. Jamieson, I. Tena, J. Evans "Autonomous Vehicle Qualification Demonstrates Potential for a Game Change", *Offshore Technology Conference OTC 2010*, paper 20456 Houston, Texas, 6 May 2010
- [4] C McGann, F. Py, K. Rajan, P.J. Ryan, R. Henthorn,. "Adaptive Control for Autonomous Underwater Vehicles". In *Proc. of 23rd AAAI Conference on Artificial Intelligence*, (AAAI), 1319–1324. 2008
- [5] A. Cormack, D.M. Lane, J. Wood "Operational Experiences for Maritime Homeland Security Operations" *IEEE International Conference Oceans 2010*, Sydney 24-27 May 2010
- [6] F. Maurelli, P.Patron, J.Cartwright, J.Sawas, Y.Petillot and D.M. Lane "Integrated MCM missions using heterogeneous

- fleets of AUVs" Proc of IEEE-MTS Oceans'12, Yesou, South Korea
- [7] S. Thrun, W. Burgard, D. Fox "Probabilistic Robotics" MIT Press 2006
 - [8] T.J. Prescott, N.F. Leopora, A. Mura, P.F.M.J. Verschure (Eds) "Biomimetic and Biohybrid Systems" Proceedings First International Conference, Living Machines 2012, Barcelona, Spain July 2012. Springer. Lecture Notes in Artificial Intelligence 7375
 - [9] M. W. Kirschner, J.C. Gerhart, "The plausibility of life: Resolving Darwin's dilemma" Yale University Press 2006
 - [10] D.M. Lane, F.Maurelli, T.Larkworthy, D.Caldwell, J. Salvi, M. Fox, K.Kyriakopoulos "PANDORA: Persistent Autonomy through Learning, Adaptation, Observation and Re-planning" Proceedings of IFAC NGCUV'12. Portugal
 - [11] K. Hamilton, D.M. Lane, K Brown, N. Taylor "An Integrated Diagnostic Architecture for Autonomous Underwater Vehicles" Journal of Field Robotics. Vol. 24, No 6, pp 497-526 (2007)
 - [12] E. Miguelañez, P. Patrón, K.E. Brown, Y.R.Petillot, D.M. Lane, "Semantic knowledge-based framework to improve the situation awareness of autonomous underwater vehicles" IEEE Transactions on Knowledge and Data Engineering, v23 , No 5 pp759-773
 - [13] M. M. Tenorth "Knowledge Processing for Autonomous Robots" PhD Thesis, Technical University of Munich 2011
 - [14] D. Lenat "CYC: A large-scale investment in knowledge infrastructure" Communications of the ACM, 38(11):33–38. 1995
 - [15] M. Tenorth, M. Beetz, "KnowRob — Knowledge Processing for Autonomous Personal Robots." In IEEE/RSJ International Conference on Intelligent RObots and Systems., pages 4261–4266. 2009
 - [16] K.B. Laskey, P.C.G. Costa, T. Janssen "Probabilistic ontologies for knowledge fusion". In Proceedings of the 11th IEEE/ISIF International Conference on Information Fusion 2008
 - [17] F. McNeill, A. Bundy "Dynamic, automatic, first-order ontology repair by diagnosis of failed plan execution", IJSWIS, Vol 3, no 3, pp 1-35, IGI Publishing, 2007
 - [18] T. Mikaelian, B.C. Williams, M. Sachenbacher: "Model-Based Monitoring and Diagnosis of Systems with Software-Extended Behaviour". AAAI 2005: 327-333
 - [19] C. Petres, Y Petillot, D.M.Lane "Path Planning for Autonomous Underwater Vehicles" IEEE Transaction on Robotics Vol. 23, No 3. (2007)
 - [20] Z. A. Saigol "Automated planning for hydrothermal vent prospecting using AUVs" PhD Thesis. University of Birmingham, UK. 2011
 - [21] M. Fox, A. Gerevini, D. Long, I.Serina "Plan stability: Replanning versus plan repair" 16th International Conference on Automated Planning and Scheduling 2006.
 - [22] P. Patron "Semantic Based Adaptive Mission Planning for Unmanned Underwater Vehicles" PhD Thesis, Heriot-Watt University, Edinburgh, 2010
 - [23] C. C. Sotzing, D. M. Lane "Improving the Co-ordination Efficiency of Limited Communication Multi-AUV Operations Using a Multi- Agent Architecture" Journal of Field Robotics Vol 27 No 4 July 2010 pp412-429
 - [24] Z. Wang, A. Peer, M. Buss "An HMM approach to realistic haptic human-robot interaction" In Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, pages 374–379 2009.
 - [25] J. Peters and J. Kober, "Using reward-weighted imitation for robot reinforcement learning," in Proc. IEEE Intl Symp. on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL), March 2009, pp. 226–232.
 - [26] P. Kormushev, S. Calinon, D.G. Caldwell, "Robot motor skill coordination with EM-based reinforcement learning". In Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS), pages 3232–3237, Taipei, Taiwan 2010
 - [27] D. Kleinfeld, E. Ahissar, M.E. Diamond, "Active sensation: insights from the rodent vibrissa sensorimotor system" Curr Opin Neurobiol, 16(4), 435-444 2006
 - [28] R. Brooks "A robust layered control system for a mobile robot". IEEE Journal of Robotics and Automation, 2 (1): 14–23 1986
 - [29] Z.Mathews. S Bermúdez i Badia, P.F.M.J. Verschure, "An Insect-Based Method for Learning Landmark Reliability Using Expectation Reinforcement in Dynamic Environments" IEEE International Conference on Robotics and Automation 2010
 - [30] U. Bernardet, M. Blanchard, and P. F. M. J. Verschure, "Iqr: a distributed system for real-time real-world neuronal simulation," Neu-rocomputing, vol. 44-46, pp. 1043–1048, 2002.
 - [31] M. Sfakiotakis, D.M. Lane, J.B.C. Davies "A Review of Fish Swimming Modes for Aquatic Locomotion" IEEE Journal of Oceanic Engineering, April 1999 Vol. 24 (2).
 - [32] T. Salumäe, M. Kruusmaa, "A flexible fin with bio-inspired stiffness profile and geometry", Journal of Bionic Engineering 8.4, Elsevier, 2011, pp. 418-428
 - [33] FILOSE Project Final Report and Deliverables <http://www.filose.eu>
 - [34] ANGELS Project Final Report and Deliverables <http://www.theangelsproject.eu/>
 - [35] J. C. Liao, D.N. Beal G.V. Lauder, M.S. Triantafyllou "The Karman gait: novel body kinematics of rainbow trout swimming in a vortex street. Journal of Experimental Biology," vol. 206, 1059 - 1073, 2003
 - [36] J.M. Gardiner, J. Atema, "Sharks need the lateral line to locate odor sources: rheotaxis and eddychemotaxis" Journal of Experimental Biology, 210(11), 1925-1934, 2007
 - [37] G. von der Emde, S. Schwarz, L. Gomez, R. Budelli, and K. Grant. "Electric fish measure distance in the dark" Letters to Nature, Nature, 395:890–894, 1998.
 - [38] F. Boyer, V. Lebastard, C. Chevallereau, N. Servagent "Underwater reflex navigation in confined environment based on electric sense" IEEE Trans Robotics and Automation, In Press.
 - [39] C. Capus, Y. Pailhas, K. Brown, D. M. Lane, P.W. Moore, and D. Houser. "Bio-inspired wideband sonar signals based on observations of the bottlenose dolphin (*Tursiops truncatus*)". J. Acoust. Soc. Am., 121(1): pp. 594–604, January 2007.
 - [40] C. Capus, Y.Pailhas, K.E Brown, D.M. Lane "Detection of Buried and Partially Buried Objects Using A Bio-inspired Wideband Sonar" Proc. IEEE Int. Conf. OCEANS 2010, Sydney, May 2010