

Use of Machine Learning for On-line Decision Making in Autonomous Marine Vehicles

Mae L. Seto, Ph.D., P.Eng.
Defence Scientist
Leader / Mine and Harbour Defence Group

IEEE Conference on Intelligent Robots and Systems 2011

Workshop: Current Directions in Marine Vehicle Autonomy
Research

September 26, 2011



Outline

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks

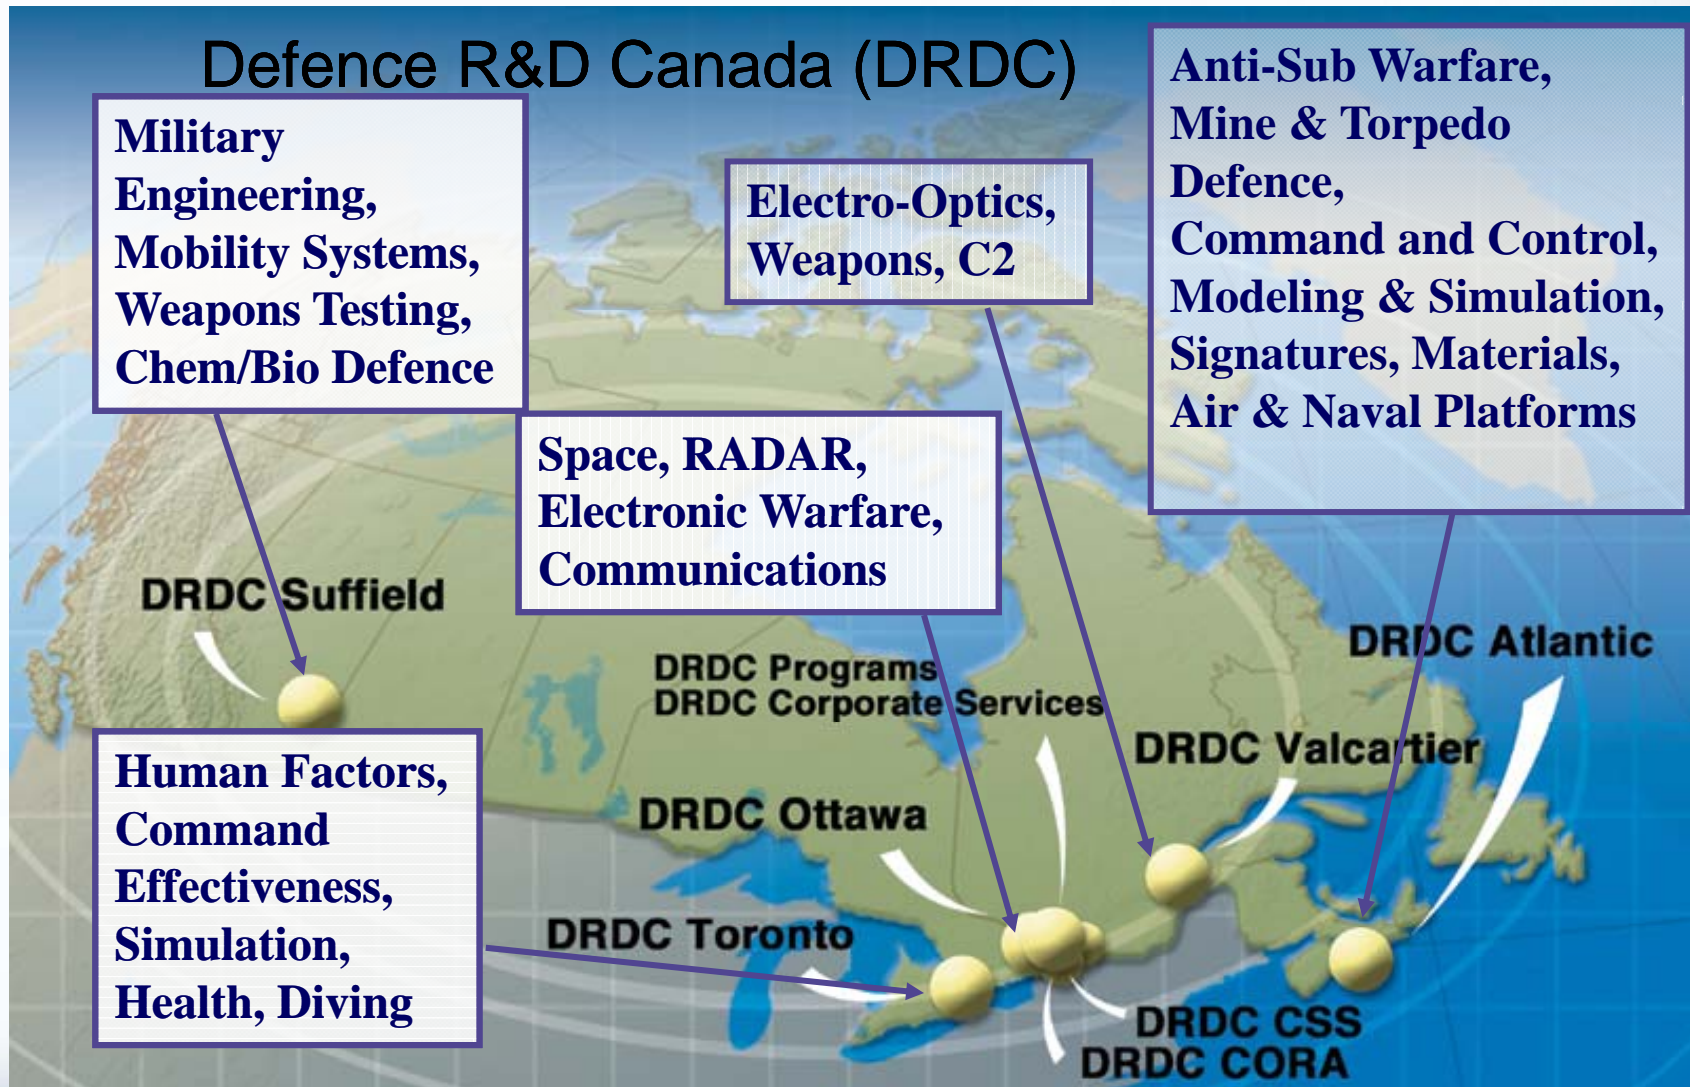
- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks

Defence R&D Canada Labs


- research & development for Canadian Forces



Naval Mine Counter-Measures (NMCM)

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks

- broad objective of mine hunting is to:
detect, classify, identify and neutralize mines in an area up to a given probability, such that the remaining risk is acceptable to follow-on traffic
- traditionally performed with specialized ships with ~80+ personnel – focus shifting to unmanned systems with increase autonomy
- autonomy makes the AUV more operator independent and compensates for poor underwater communications and lack of positioning references
- potential NMCM CONOP with collaborative marine robots for Aug 2012



True Operational Benefit of Autonomous Underwater Vehicles

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



- the ideal is continuous operation over many days (or weeks) with a powerful multi-function sensor suite on-board, and at the same time limit overall size of vehicle to enable simple operation from vessels of opportunity
- limitations that prevent AUVs from realizing their full potential
 - underwater navigation (lack universal positioning system)
 - communications (high attenuation and low bandwidth)
 - on-board energy density and its impact on speed, endurance, sensor choice, and on-board processing
 - on-board intelligence
- this presentation deals with on-board intelligence

Autonomy for Marine Robots - Required Onboard Capabilities

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



- intelligent on-board autonomy allows robots to operate long periods without operator intervention
- require robot have autonomy for decision-making or problem-solving to address *unexpected* robot or mission developments – necessary on long missions as robot and mission evolves
 - Brooks' behavior-based subsumption control architecture (1986) on robots to date – relies on operator scripted missions a priori
 - cannot adapt a mission in-situ based on changes in the ocean environment (as detected fr on-board sensors) or the robot itself
- *mission* autonomy – adapt mission to unanticipated conditions in the environment or in-situ intelligence exploited to better perform the mission, e.g. sonar optimized based on in-situ environment measures
- *robot* autonomy – increases robot fault tolerance to unanticipated self events, for. e.g. a jammed control fin, failed sensor, etc.
- autonomy *paramount* in environment with low bandwidth and high attenuation for acoustic communications



Motivation for On-board Machine Learning

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



- make AUV less dependent on operator
- AUV adapt to unstructured dynamic ocean environment and evolving robot states
 - controllers can adapt on-the-fly to sea state, current, towing effects, etc.
- aids for decision-making and mission re-planning
- on-the-fly learning for image recognition of objects (e.g. mines) and consequent re-tasking of robot based on the image information
- programming the robot easier for less technical users
 - robot learns & generates the optimal robot control policy
- *ideally, define high level task specification and let robot learn implementation*


DRDC Autonomy Testbed IVER2 AUVs

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



3 such vehicles for research in collaborative robotics

- mission planning software interfaces to geo-referenced charts
- WiFi coms to command laptop
- 6 beam Doppler Velocity Log with ADCP
- sensors: Yellowfin Side Scan Sonar (330 / 770 kHz)
MarineSonics SSS (900 / 1800 kHz)
Delta T bathymetric profiling sonar
- B/W and color low light camera and video camera (LED lights)
- 600 Whr energy, 1 – 4 knots speed; 2 knots endurance > 12 hours
- DGPS on surface
- WHOI Micro-modem
- X86 processor with 80 GB drive



Robot and Mission Autonomy with Learning

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



examples described:


- robot fault tolerance
- mission re-planning
- classification of underwater targets for MCM
 - emphasis on marine autonomous systems and on-line learning

Autonomous Fault Tolerance Jammed Control Fin

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks



- AUV fins are used for altitude, depth, yaw, pitch, and roll control
- a jammed plane's impact on a mission can be severe – ranging from vehicle loss to scuttled mission to compromised control
- an AUV fin can be jammed (underactuated) due to:
 - actuator problems
 - powering issues
 - AUV collision resulting in a jammed fin
 - detritus jammed in the actuator
- if AUV can be recovered the plane / actuator is repaired and the AUV is sent off on mission again
- for AUVs on long deployment a support ship may not be nearby for this
- no fail-safes in place to allow the mission to continue in the event of a jammed fin




Jammed Control Fin

Proposed Underway Solution

- Introduction
- **Fault Tolerance**
- Mission Planning
- Target Recognition
- Concluding Remarks



- an on-board re-distribution of control authority can be performed underway to minimize impact on a mission
- requires on-board knowledge about the robot's motion model (AUV dynamics and control)
- also requires a means to optimize the control authority
- in the form of an on-board knowledge-based agent that is invoked when a jammed fin is detected
- optimal gain redistribution, through a genetic algorithm, makes a difference in the AUV's ability to continue a mission or, at least, to transit to a recovery point



Agent Effectiveness for Bow and Stern Fin Jam

- Introduction
- **Fault Tolerance**
- Mission Planning
- Target Recognition
- Concluding Remarks

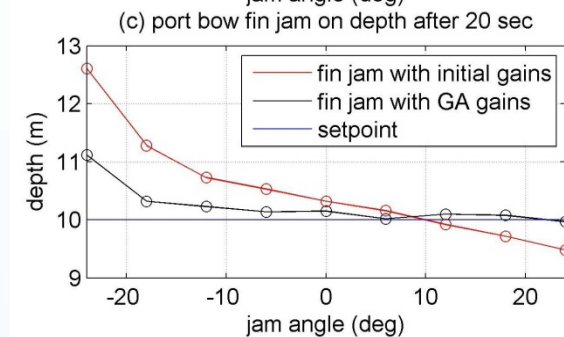

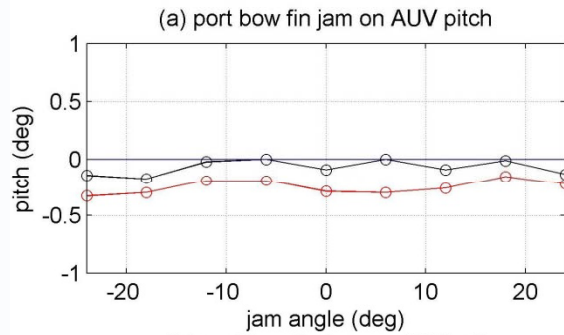
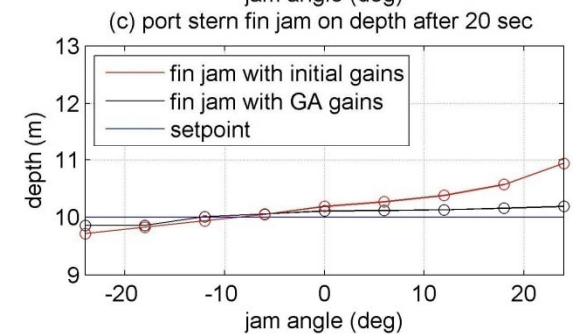
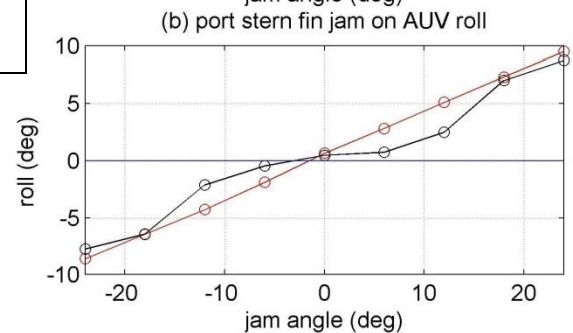
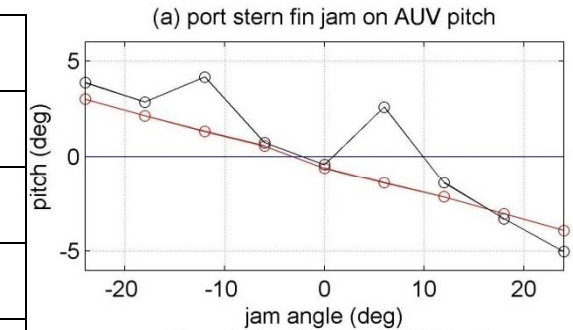


TABLE I: AUV depth rate (m/s) with a bow fin jam

bow fin jam angle	without agent	with agent
-12	0.09	0.003
-18	0.18	0.07
-24	0.41	0.13



Fault Tolerant Control Fins Other Agent Applications

- Introduction
- **Fault Tolerance**
- Mission Planning
- Target Recognition
- Concluding Remarks



- implement agent on-board in-house IVER2 AUVs
- apply to the case of an AUV losing a plane
- apply to AUVs that tow as the tow has an impact on the AUV heading, pitch, roll, and depth, through the tow loads
- adjust to local salinity levels
- improved sea-keeping



Mission Autonomy – Re-plan Mission due to Energy Shortfall

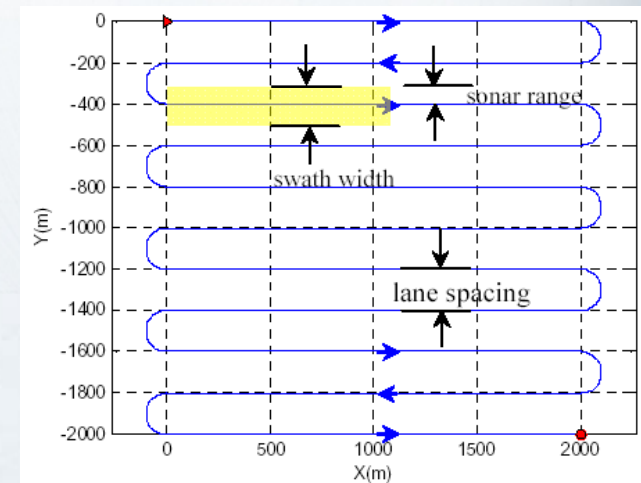
- Introduction
- Fault Tolerance
- **Mission Planning**
- Target Recognition
- Concluding Remarks



- contingency for when on-board energy left is less than anticipated due to stronger currents, higher consumption rate, uncorrected trim, correction for being off-course, etc.

- knowledge-based agent re-plans AUV survey mission given:

1. lower energy budget
2. mission duration left
3. remaining survey area dimensions
4. consumption rate for hotel load
5. projected propulsion energy consumption for AUV



- linear optimizer produces the re-planned mission

Energy Agent Results

- Introduction
- Fault Tolerance
- **Mission Planning**
- Target Recognition
- Concluding Remarks



- optimized mission to survey a 2 km x 2 km survey area with varying available energies and 30,000 seconds left until AUV recovery (swath overlap varies from 20% – 50% or a lane spacing of 80 – 50 m, respectively)

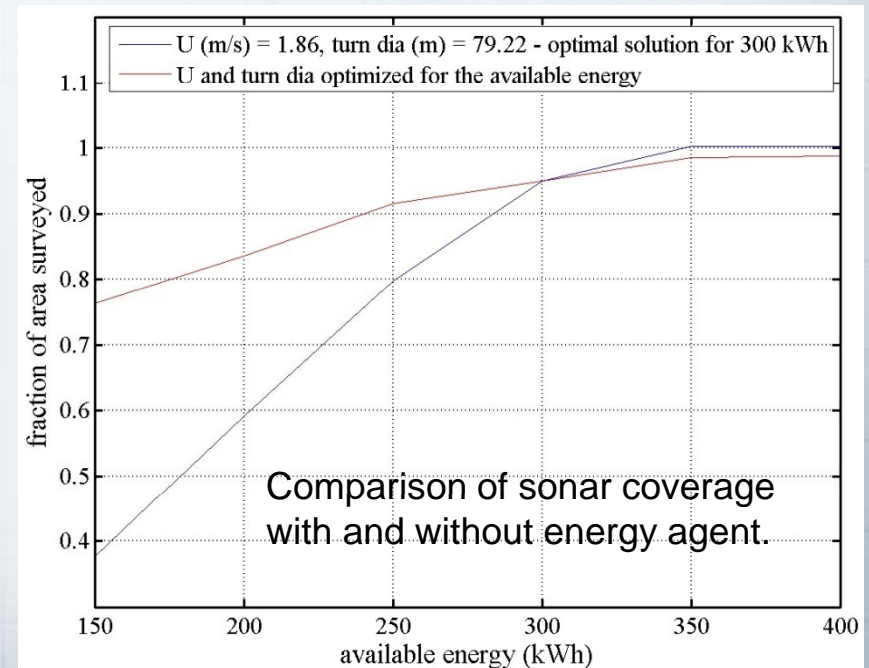
% area surveyed	lane spacing	AUV speed	energy used	energy available	time used (30,000 max)
	m	m/s	kWh	kWh	sec
76.34	78.29	1.46	133.0	150	29,837
83.58	79.83	1.59	179.3	200	28,112
91.61	80.00	1.72	231.0	250	28,600
94.98	79.22	1.86	290.9	300	29,402
98.59	75.91	1.92	335.9	350	29,741
98.76	76.05	1.95	270.3	400	29,652

Energy Agent Effectiveness

- Introduction
- Fault Tolerance
- **Mission Planning**
- Target Recognition
- Concluding Remarks



- evaluated by comparing re-planned missions against missions that are not adaptable
- measured against optimal solution for given energy density
- as available energy drops agent effectiveness apparent
- at the lowest energy considered (factor of 2 less) area coverage increases by a factor of 2

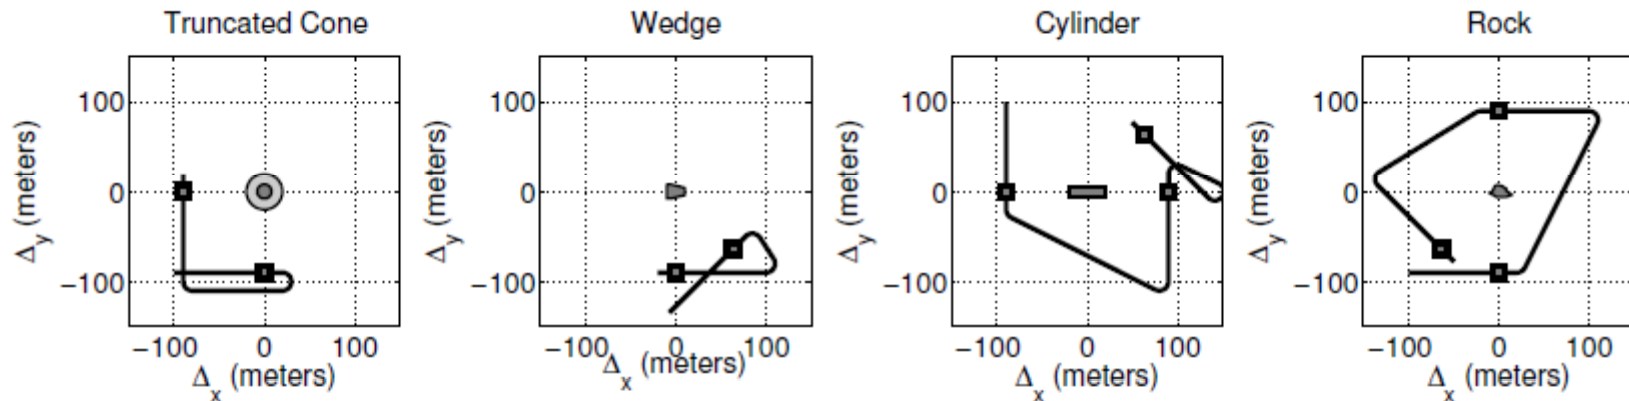
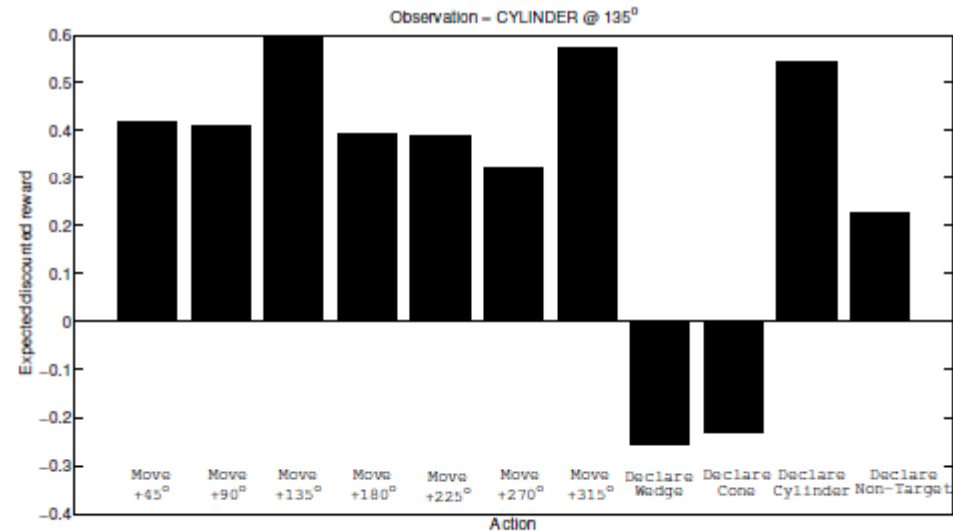


Autonomy + Perception

Adaptive Sensing with POMDP

- Introduction
- Fault Tolerance
- **Mission Planning**
- Target Recognition
- Concluding Remarks

- exploit vehicle manoeuvrability and ATR output to adaptively determine which look to obtain next to reduce uncertainty and improve classification




Sample AUV routes chosen by the POMDP in order to classify different object types


Automated Target Recognition Mine Detection & Classification

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks

- image processing and/or statistical techniques extract features on a sea floor object
- features discriminate between sea floor and clutter
- detection: mine-like objects selected from seabed image (image processing tools, matched filter) and feature extraction
- classification: determines whether object is mine-like



Detection of objects



Template matching based on shadows

Automated Target Recognition Using Learning Methods

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks




- two means of achieving classification:
 1. fit to a multi-class set of potential mines (e.g. Z-test, matched filter, template matching) against a database of known shapes
 - methods sensitive to parameters for sea bottom type and computationally intensive
 2. learning methods that have been trained on examples
 - holds the promise of not having those issues
- Restricted Boltzman Machines (RBM) – generative models that learn to represent distribution of the training data
- earlier work detected the objects through cascaded Restricted Boltzman Machines that learned to extract features from side scan sonar images


Automated Target Recognition Learning Method Effectiveness

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks


- current work has involved developing a lightweight, component driven framework for hosting ATR algorithms.
- this system is a logical system to be embedded in an AUV
- RBM can generate realistic imagery after training



(a) Example cone



(b) CITADEL cone



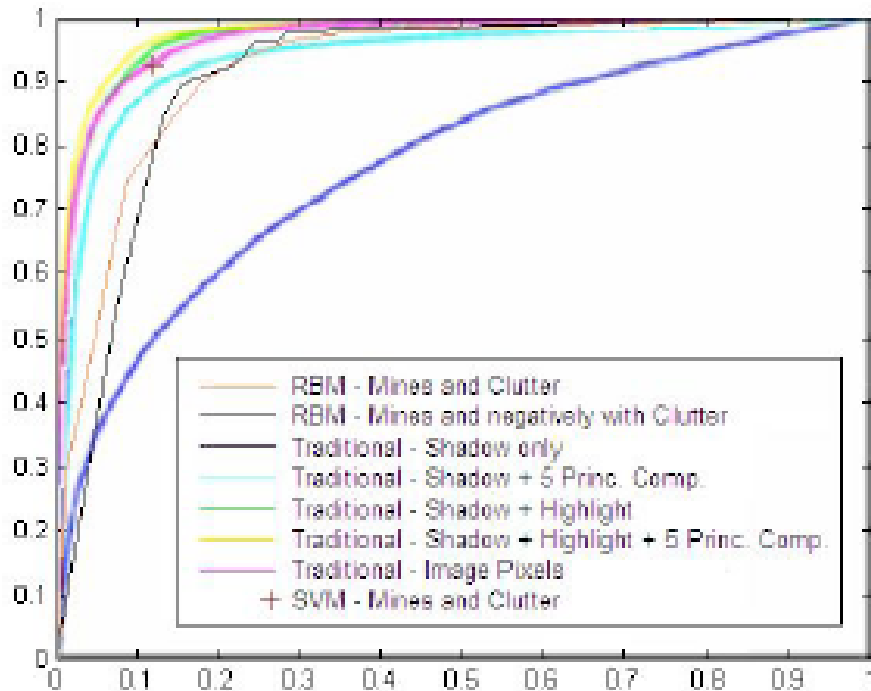
(c) Generated cone

Side scan sonar image of mines showing: (a) typical mine, (b) data from a specific trial, (c) image generated from a trained RBM


Automated Target Recognition Learning Method Effectiveness

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks

- effectiveness comparable to traditional engineered techniques



(a)



(b)

(a) Classification results: probability of detection (y-axis) vs probability of false alarms (x-axis); (b) free energy distribution

- Introduction
- Fault Tolerance
- Mission Planning
- Target Recognition
- Concluding Remarks

Concluding Remarks

- machine learning adds adaptive on-line capability to autonomous marine systems in intelligent fault tolerance, mission-planning, and automated target recognition:
 - re-distribute control authority in an AUV with a jammed control plane allows the AUV to continue mission and avert a potential vehicle loss; improves sea-keeping and performance under tow
 - re-plan mission to adapt for unanticipated energy shortfalls and more efficient missions that are driven by features extracted in the object perception
- challenge: getting the processing power required in a small form factor and without impacting AUV endurance

DEFENCE



DÉFENSE