Optimizing sonobuoy placement using multiobjective machine learning

Dr. Christopher Taylor

UNIVERSITY OF LIVERPOOL

Agenda

- Introduction to the sonobuoy placement problem
- Modelling the problem
- The hybrid multiobjective machine learning algorithm
- Results
- Discussion and conclusions
- Q&A

Introduction to the sonobuoy placement problem

- Sonobuoys are portable, expendable sonar systems
- Can be passive or active
- Monostatic, bistatic or multistatic
- In our scenario, we consider fields of passive sonobuoys (e.g. DIFAR)
- Can be placed by an airborne agent – helicopter, aircraft or UAV

https://www.navalnews.com/naval-news/2022/04/ultra-and-sparton-win-u-s-navy-contract-for-new-sonobuoys/

UNIVERSITY OF LIVERPOOL

Introduction to the sonobuoy placement problem

- Sonobuoys are placed sequentially by the agent, typically directed by a surface vessel
- Sonobuoys consist of a flotation device and an array of hydrophones that unfurls on deployment

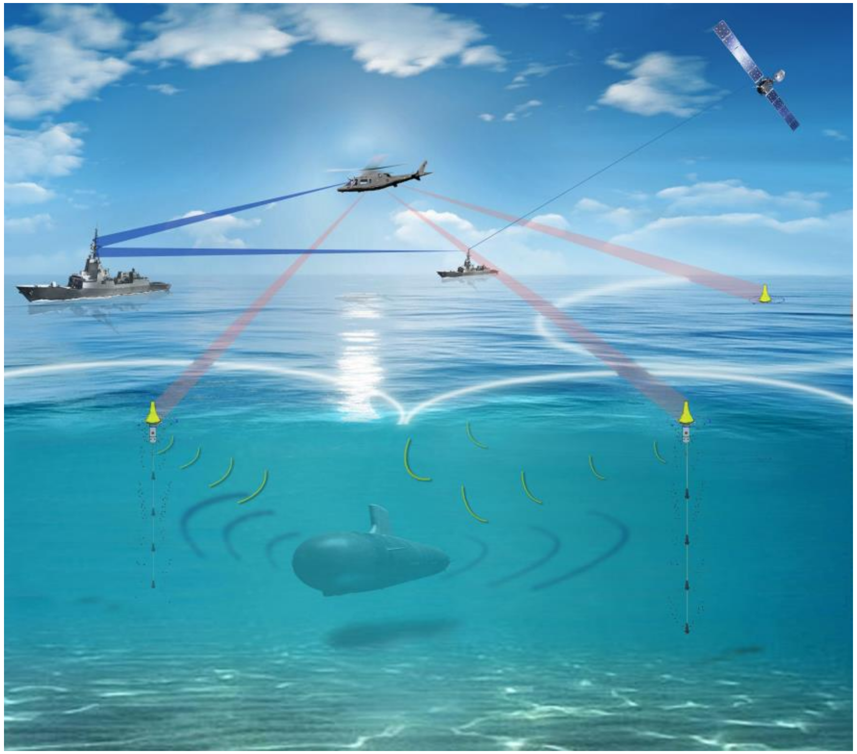

https://www.militaryaerospace.com/sensors/article/14198901/antisubmarine-warfare-asw-sonobuoys-multistatic

Introduction to the sonobuoy placement problem

- Given limited agent payload and the cost of the sonobuoys, we need to use deployment patterns that are the most efficient for the mission objectives
- Grid or lattice patterns are common, but other patterns such as circles/ovals or chevrons might be used for particular mission types
- In our scenario, we seek to minimize localization uncertainty at the point when the entire pattern has been deployed, so that action can be taken
- We also want to minimize the total placement time
- Potentially conflicting objectives

Modelling

- We model sonobuoy placement as a constrained biobjective problem:
  - Minimize total placement time
  - Minimize localization uncertainty immediately after sonobuoy pattern placement is complete
  - Constraints ensure solutions adhere to minimum and maximum distances between placement locations

Modelling

- We assume a single target of interest (TOI)
- Initially, we assume a uniform environment
- Later, we add randomly generated simulations of noise/clutter
- Other simplifying assumptions include:
  - Accurate placement and no drift of sonobuoy positions
  - No failure and sonobuoys always detect the TOI
  - Static noise/clutter

UNIVERSITY OF LIVERPOOL

Modelling

- Environment is discretized to a lattice of hexes
- To simulate uncertainty over the TOI location, measurements are taken from possible contacts in hexes surrounding ground truth
- Transmission loss and noise/clutter calculated for each path from contact to sonobuoy

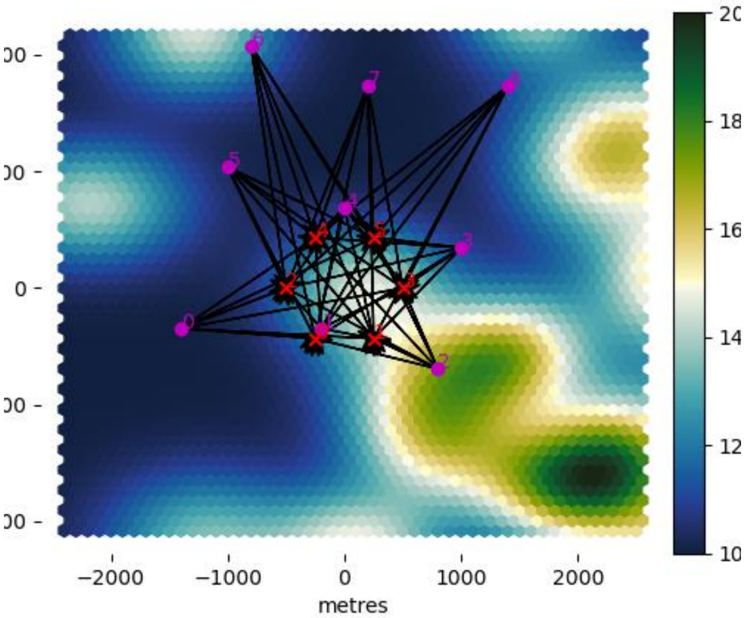

Red x = contact
Purple circle = sonobuoy
Black line=path from contact to sonobuoy
Noise/clutter in dB/km

Modelling

- First optimization objective – total placement time consists of:
  - the time taken for the agent to reach the first placement location plus
  - The time taken for the placement procedure for each sonobuoy plus
  - The time taken to travel between successive placement locations

Modelling

- Second optimization objective – localization uncertainty modelled as a sum of uncertainties over all triangulations between pairs of sonobuoys
  - Transmission loss modelled as spherical for simplicity
  - For a given sonobuoy pair, uncertainty is at a minimum when they are orthogonally placed with respect to a contact
  - Localization uncertainty at a maximum for a given pair when the sonobuoys form a line with the contact

Biobjective machine learning algorithm

- Algorithm has two phases:
  - Phase 1 - offline multiobjective evolutionary algorithm (MOEA)
  - Phase 2 - online multiobjective reinforcement learning (MORL) algorithm
- Multiobjective optimization produces a Pareto front (PF) consisting of a number of nondominated solutions
- The MOEA produces initial solutions in a static environment
- The MORL uses patterns evaluated by the MOEA together with updated information to produce new solutions

Biobjective machine learning algorithm

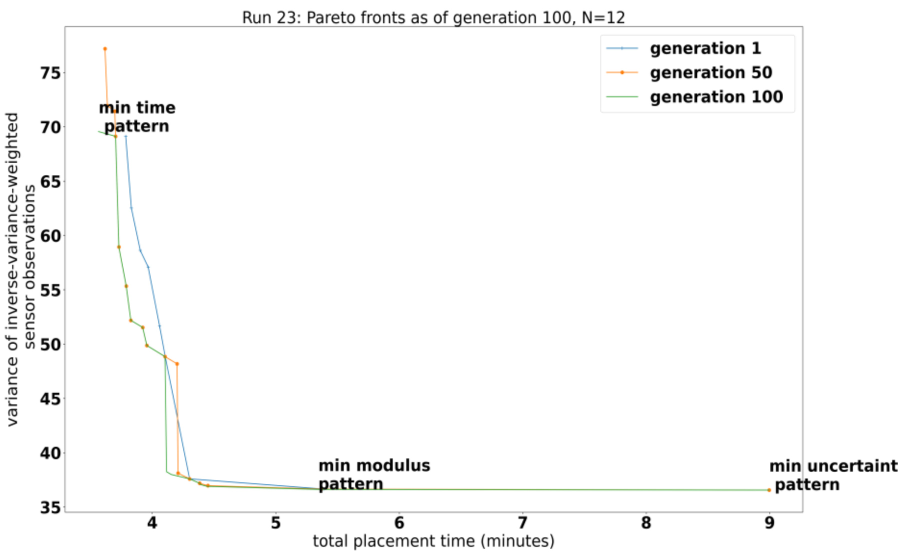

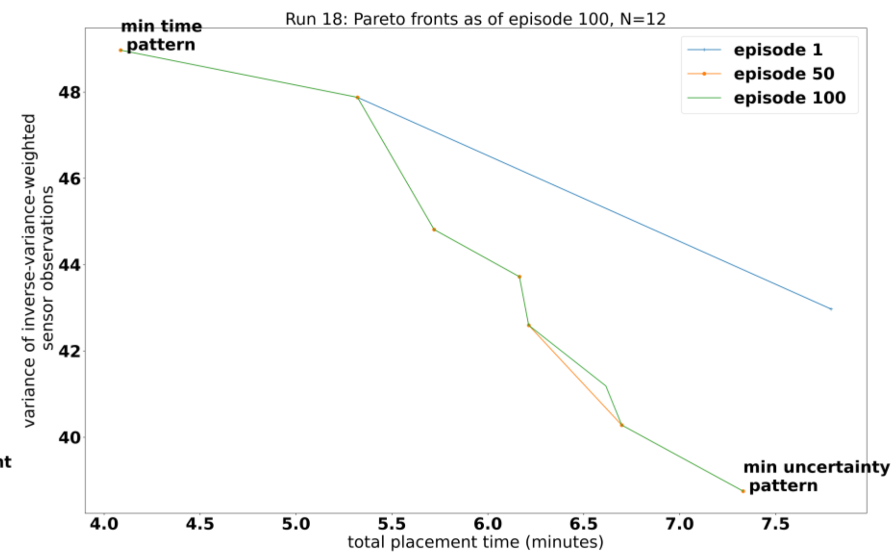

- Pareto front examples:
  - Left hand figure shows MOEA phase progression over 100 generations
  - Right hand figure shows MORL phase progression over 100 episodes

UNIVERSITY OF LIVERPOOL

Biobjective machine learning algorithm

- Why the hybrid approach?

	Evolutionary algorithms	Reinforcement learning
Advantages	• Avoid local minima • Good at multiobjective problems • Generates a diverse population	• Good for online problems • Well suited to learning a sequence of optimal actions
Disadvantages	• Can be slow to converge • Not good at dealing with updated information	• More easily stuck in local minima • Not as well proven in multiobjective setting

Biobjective machine learning algorithm

- MOEA structure:
  1. Initialization – generate all suitable permutations of grid patterns of *N* sonobuoys and add randomly generated patterns
  2. Fitness evaluation – calculate values for each objective and find the PF
  3. Tournament selection – use 2 permutations of the population to generate binary tournaments and pass the winners to genetic operators
  4. Mutation – with probability μ, move one placement location by one hex, chosen at random, subject to constraints
  5. Crossover – with probability 1-μ, use 2 permutations of the new population to generate parent pairs and generate valid children that satisfy constraints

Repeat 2-5 for *G* generations, choose a pattern from the final PF and pass archive of all assessed patterns whose first *k* locations match the chosen pattern to the MORL

Biobjective machine learning algorithm

- This view shows uncertainty plotted against time elapsed for different Pareto-nondominated patterns produced by the MOEA
- An operator could use this view to select which solution to pass to the MORL, depending on mission priorities
- In our experiments, we assume the solution with the minimum modulus (i.e. closest to the origin) is chosen

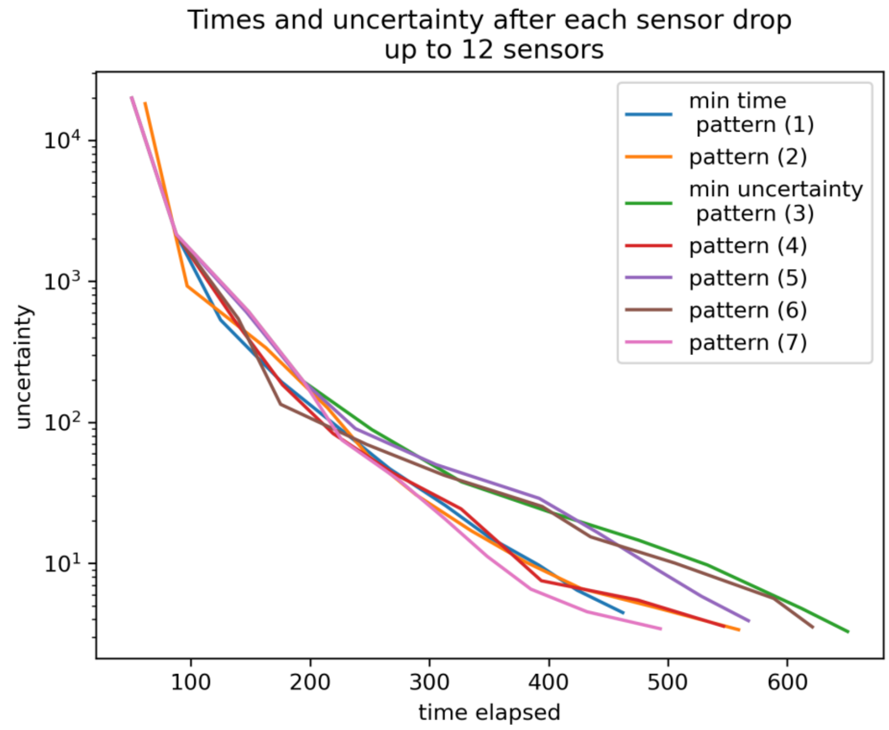

Biobjective machine learning algorithm

- MORL modelling:
- State space is too large for a tabular approach
- Instead, we use an approximation function that takes an average of the objective values for all assessed patterns whose first $n \in [k, N-1]$ placement locations match that of the candidate pattern
- We use complete rollouts and do not discount returns
- Multiobjective approach means there is a set of nondominated *Q*-values for available actions in each state, rather than a single best action
- We use a more complex version of the ε-greedy approach to decide actions

Biobjective machine learning algorithm

- MORL structure:

1. Initialization – import archive of assessed patterns from the MOEA and recalculate fitness values based on updated information about the position of the contacts
2. For each sonobuoy placement, with probability ε choose a next location at random, or else:
    - if any next states exist with zero *Q*-values:
        - with probability ε, select at random from the zero states;
        - or otherwise with probability 1-ε, selects at random from the non-zero states;
    - if no next states exist with zero *Q*-values:
        - select a next placement location at random from patterns on the PF
3. Calculate objective values ,add the new pattern to the archive and go to 2. for required number of episodes

Results – MOEA with uniform environment

Elite patterns for *N*=9

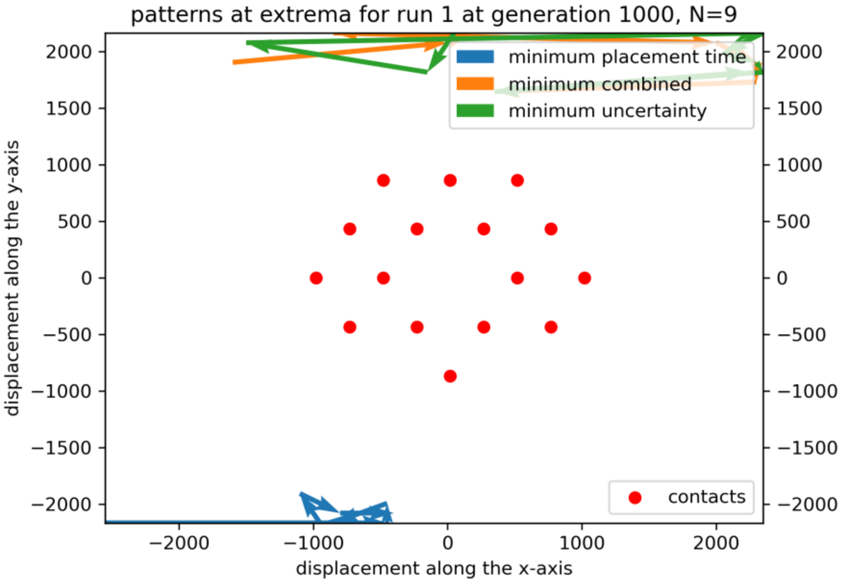

Elite patterns for *N*=12

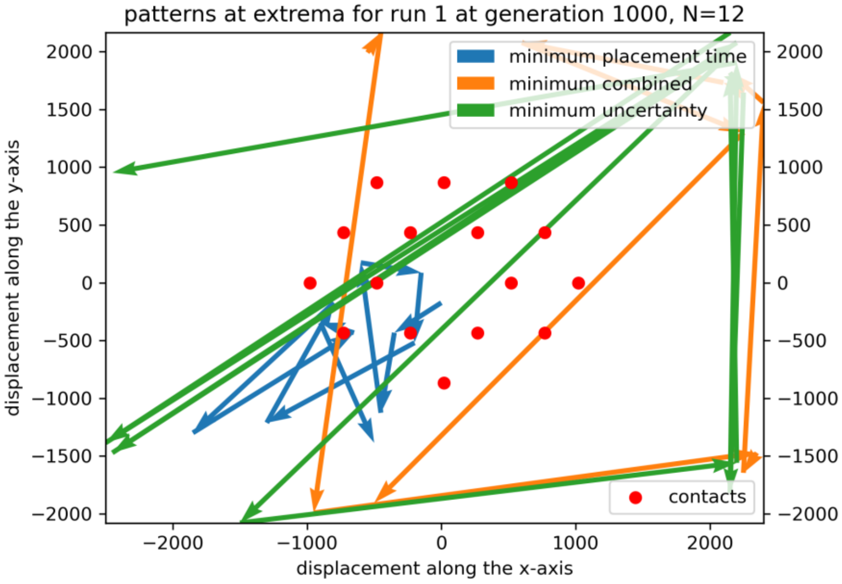

- For lower values of N, patterns tend to be located at one side of the contacts
- As the number of sonobuoys increases, nondominated solutions with :
  - lower placement times tend to be within or close to the zone containing the contacts
  - lower localization uncertainty increasingly surround the contacts

UNIVERSITY OF LIVERPOOL

Results – MOEA+MORL, complex environment

Number of sensors	MOEA mean	MORL mean	MOEA max	MORL max	MOEA %	MORL %
8	3.0%	1.1%	9.4%	11.3%	83.3%	20.0%
9	3.4%	1.3%	9.5%	11.9%	80.0%	30.0%
10	5.1%	0.8%	15.3%	8.8%	90.0%	20.0%
11	5.5%	0.8%	11.5%	6.2%	96.7%	20.0%
12	4.1%	0.3%	7.7%	3.5%	100.0%	16.7%

- Results averaged over 30 runs of 100 MOEA generations + 100 MORL episodes
- Mean figures represent improvement in the modulus of the minimum modulus solution
- Max figures represent the mean maximum improvement
- % figures represent the percentage showing some improvement

Results – MOEA vs. MORL

- With *N*=8 sonobuoys, both MOEA and MORL phases show continuous but improvement
- With *N*=10 sonobuoys, improvement is faster for the MOEA than with 8 sonobuoys, but slower and less continuous for the MORL
- Likely due to higher computational complexity of the MORL – cubic in *N* vs quadratic for the MOEA

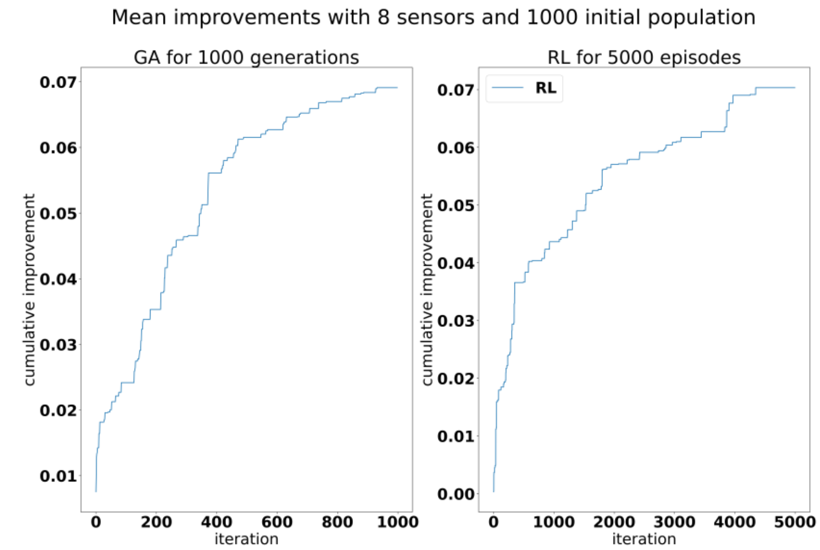

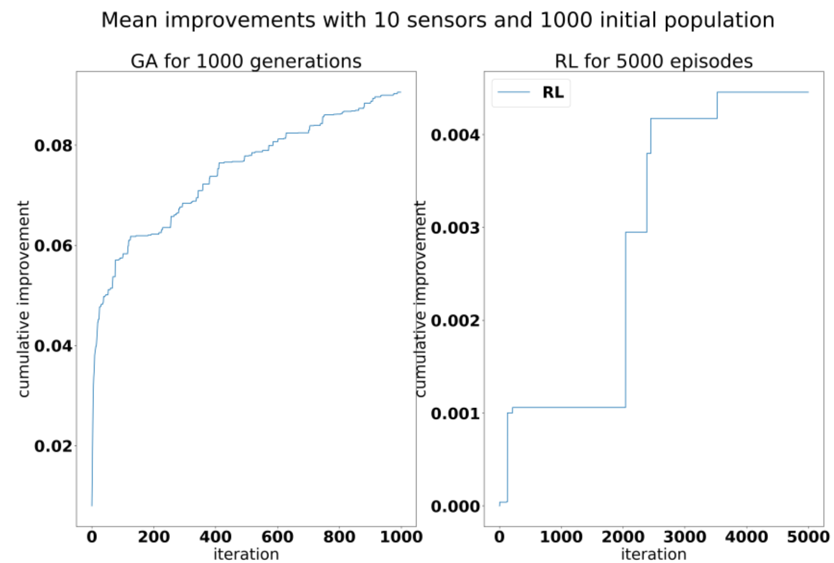

Discussion and conclusions

- MOEA shows good results with the static problem
- MORL has a hard job to do – improving on already optimised results from the MOEA and coping with updated information – but shows good results with smaller numbers of sonobuoys
- MOEA scales well with larger numbers of sonobuoys but MORL shows decreased performance
- Approach gives operator a choice as to how to prioritize localization uncertainty vs. total placement time
- Algorithm may also make it possible to use smaller numbers of sonobuoys, lowering costs

Discussion and conclusions

- Computational complexity is higher for the MORL vs the MOEA
- Many more MORL episodes required to gain good results for higher numbers of sonobuoys
- More sophisticated approximation function may be required
- More advanced parallelization techniques also a topic for future research

Q&A

- Any questions?

Reference: Taylor CM, Maskell S, Ralph JF. Optimizing sonobuoy placement using multiobjective machine learning. In: SSPD 2022; Forthcoming.