Marine Robotics:
Planning, Decision Making and Human-Robot Teaming

Geoff Hollinger

Assistant Professor
Collaborative Robotics and Intelligent Systems (CoRIS) Institute
Oregon State University

June, 2019

Marine Robotics (2019)

Oregon State University

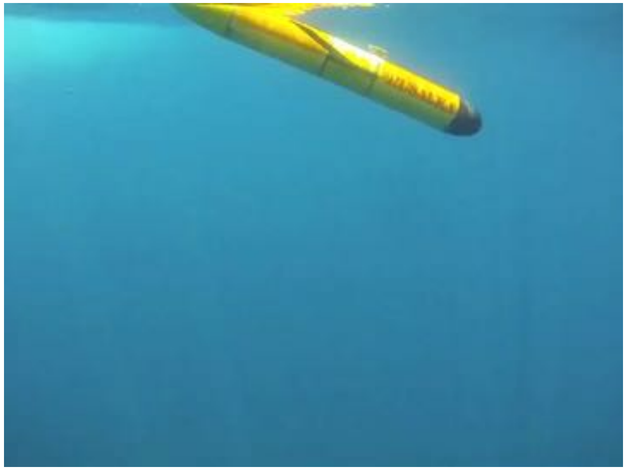

Photo: NASA

Photo: LA Times

Photo: NIWA Science Web

Photo: Bluefin Robotics

Photo: Hydro international

Why use robots?

- ➢ Sampling ocean features is often done from research vessels
  - Costs per day can exceed $30,000 (€26,750)
  - Vessel only samples in a single place at a given time
- ➢ Autonomous vehicles decrease sampling costs while increasing sampling quality

Marine Robotics (2019)

Oregon State University

- State of practice
  - Operator pre-specifies waypoints
  - A team of experts look at the data
  - The team specifies more waypoints
- Goal: marine autonomy
  - In situ decision making
  - Scalability to many vehicles
  - Shared autonomy with technicians

Ocean features

- Many oceanographic features can be described as 'hotspots' in the environment
  - Temperature
  - Bio-Acoustics
  - Chemical Spills

- Planning and monitoring algorithms can improve data collection ability

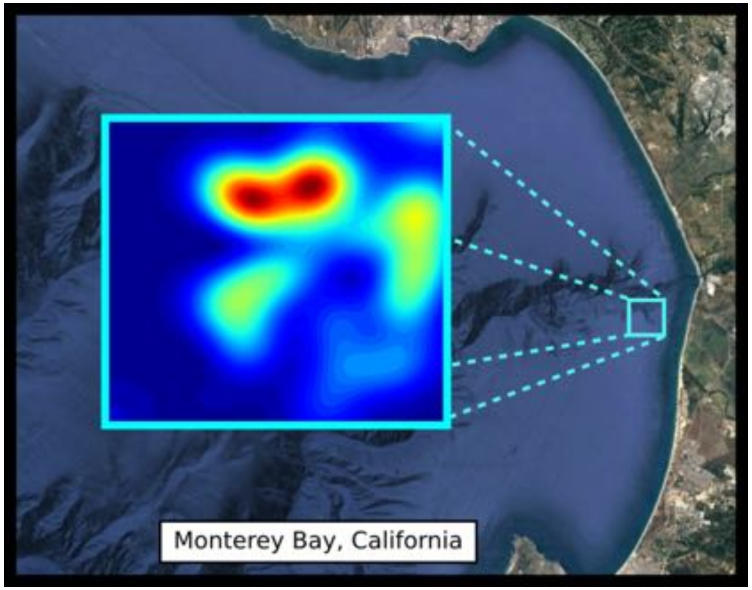

Bio-Acoustic Hotspots in data collected in Monterey, California, May 2017

Adaptive autonomy

- Robot data gathering
  - Mapping from asynchronous observations
  - Topological planning to improve efficiency

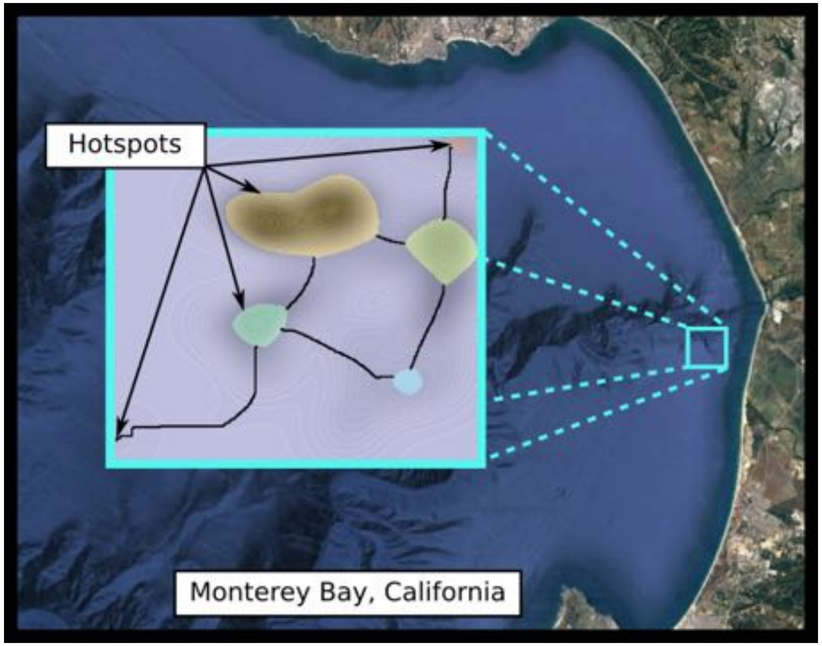

S. McCammon and G. Hollinger "Topological Hotspot Identification for Informative Path Planning with a Marine Robot" ICRA, 2018

Hotspot identification

1. Find local maxima and minima

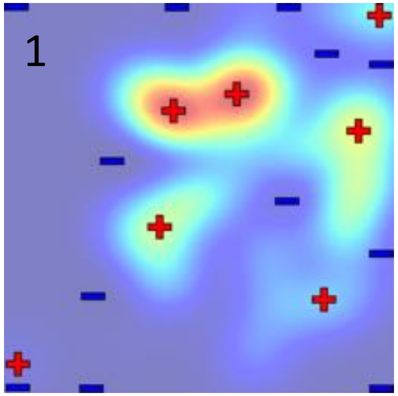

Hotspot identification

1. Find local maxima and minima
2. Expand regions around these points using Fast Marching Method
   - Adapt travel cost function based on which region we are expanding

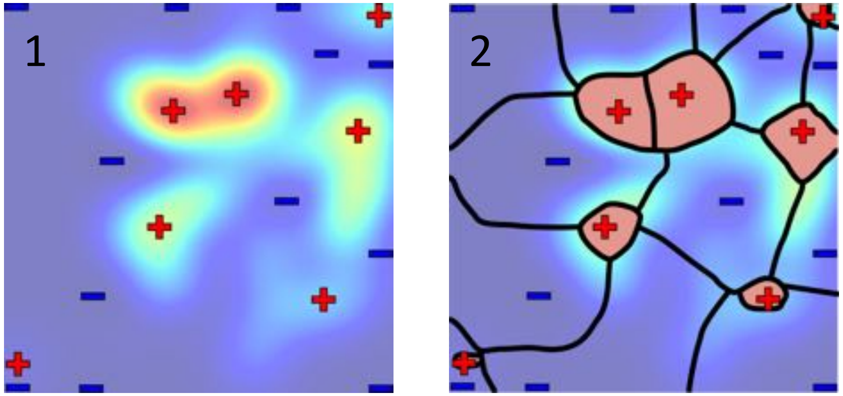

Hotspot identification

1. Find local maxima and minima
2. Expand regions around these points using Fast Marching Method
   - Adapt travel cost function based on which region we are expanding
3. Merge Adjacent regions with the same label
   - Create edges along merged minima locations

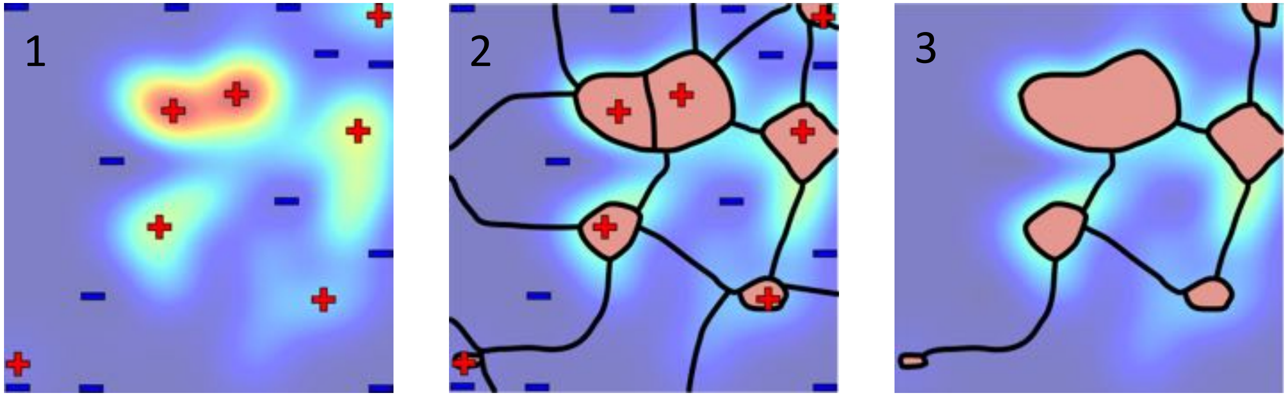

Hotspot identification

Oregon State University

1. Find local maxima and minima
2. Expand regions around these points using Fast Marching Method
   - Adapt travel cost function based on which region we are expanding
3. Merge Adjacent regions with the same label
   - Create edges along merged minima locations
4. Result: Topological Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

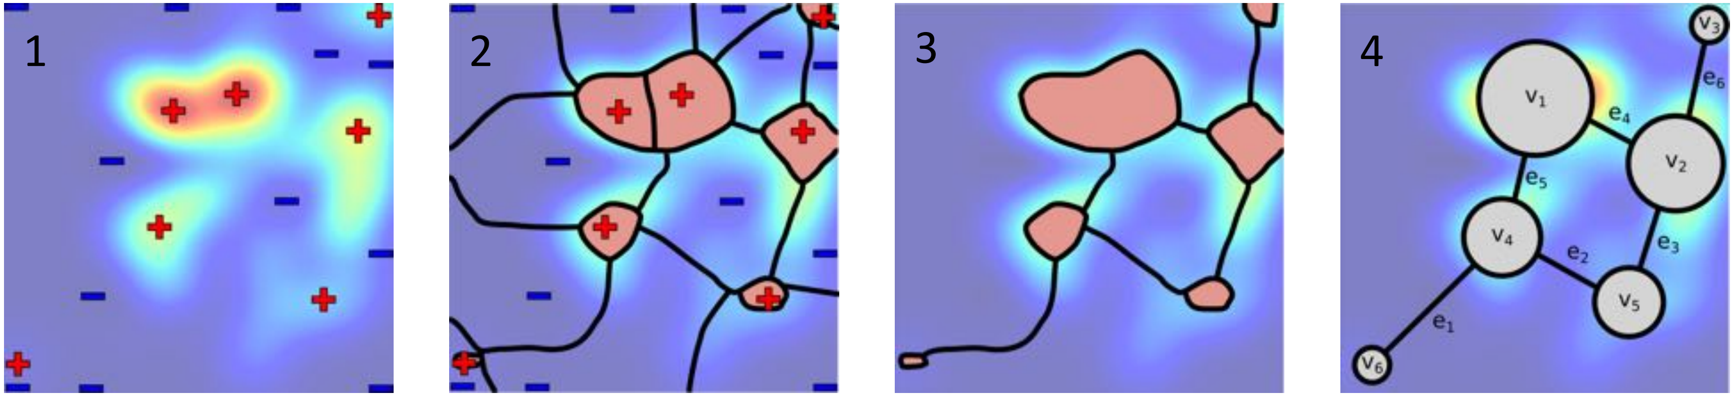

Plan graph

- Goal is to allocate time between hotspot regions to maximize information collected
  - Subject to budget constraints
- Lagrange multiplier method distributes time between hotspots
- Within each hotspot a greedy algorithm used to plan path

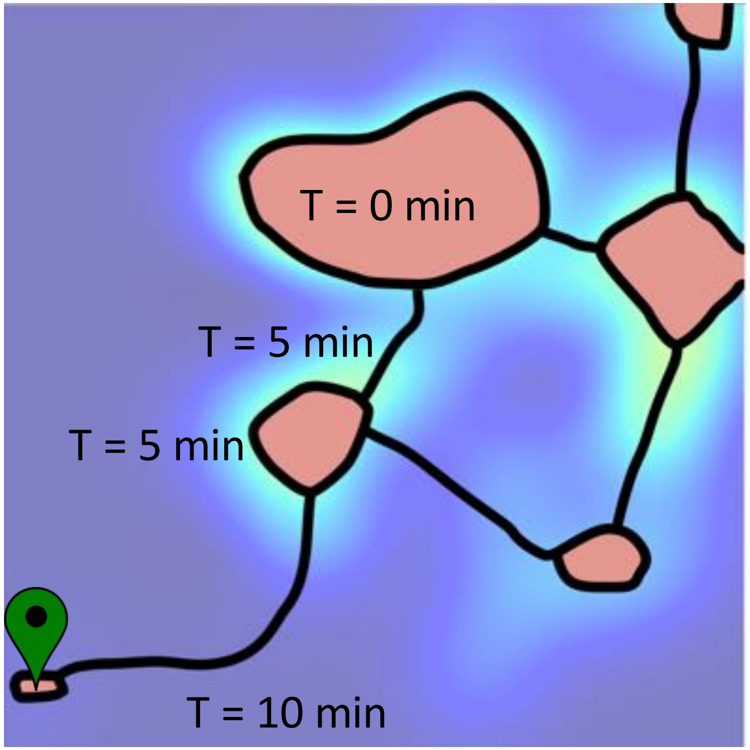

Long-term planning

- ➢ One-shot planning approach assumes a static world
  - Not appropriate for long-term ocean deployments in a dynamic environment
- ➢ Receding Horizon Planning
  - Interleave planning and execution over course of deployment
  - Incorporate new observations during each planning cycle

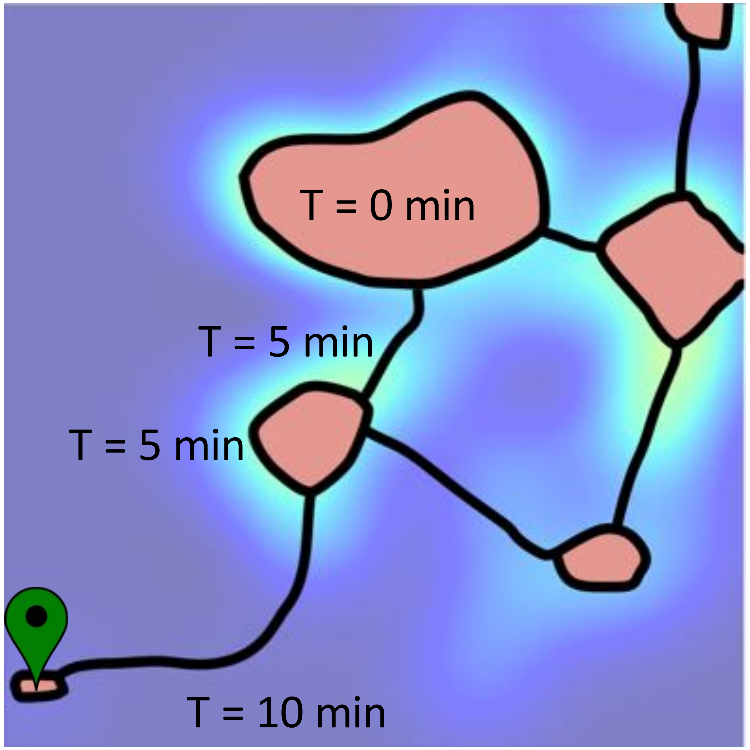

Modeling and estimation

- Gaussian Process (GP) world estimator allows new observations to be incorporated
  - Length-scale parameters control how correlated two observations are
  - Upper Confidence Bound (UCB) encourages exploration

- Add time to GP
  - Additional dimension of prediction
  - Train GP offline using satellite data to determine appropriate length-scale
  - Incorporate increasing uncertainty in past observations

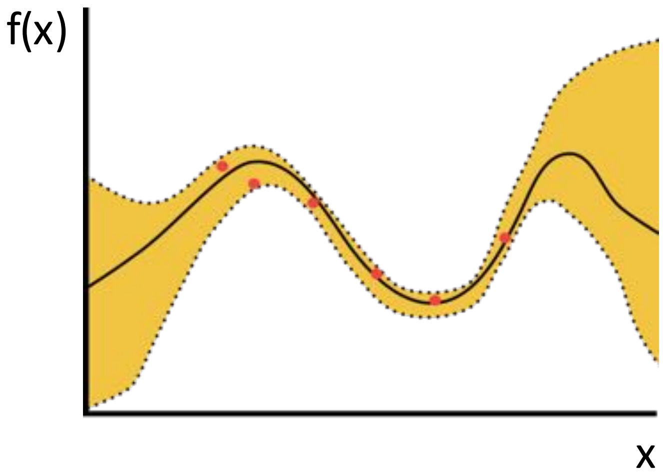

Ocean deployments

➢ Performed two deployments Summer 2018

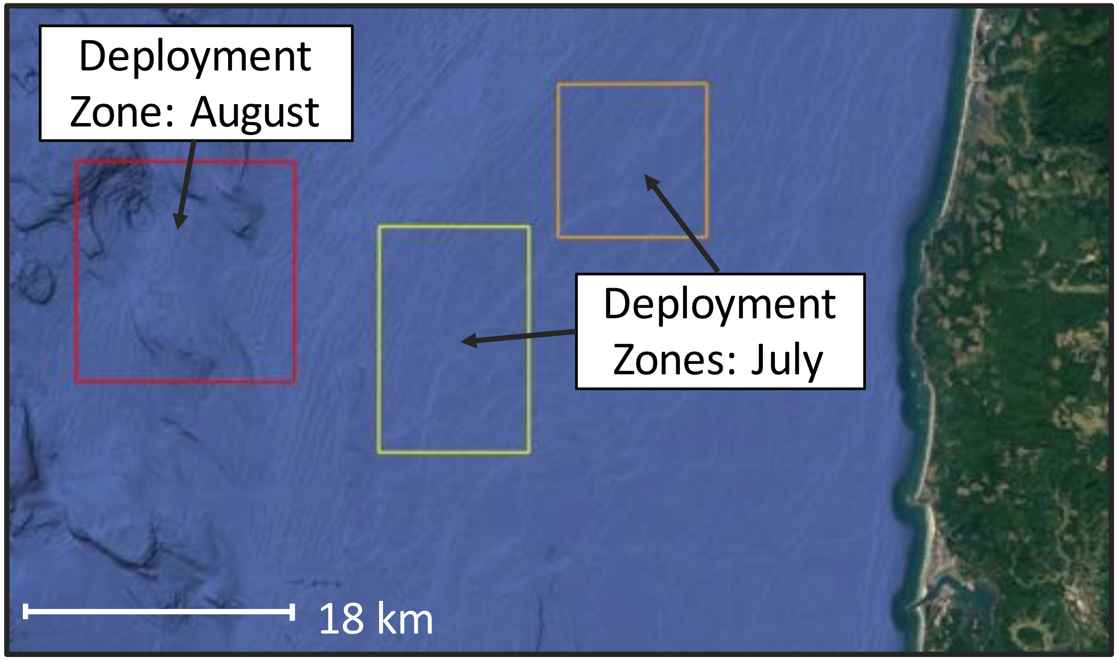

July 2018 deployment

- Original deployment zone in heavy southerly currents
- Limited period of time with active navigation enabled before recovery

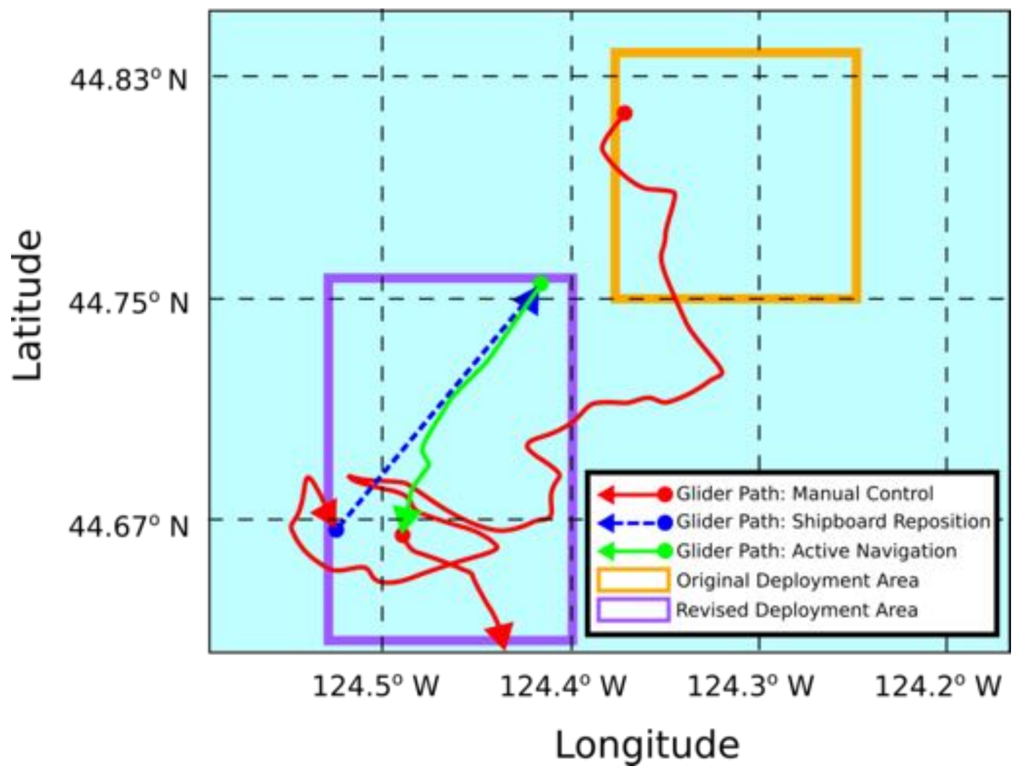

August 2018 deployment

Oregon State University

- New deployment zone farther west to avoid strong currents
- Successfully identified upwelling front in bottom right portion of map

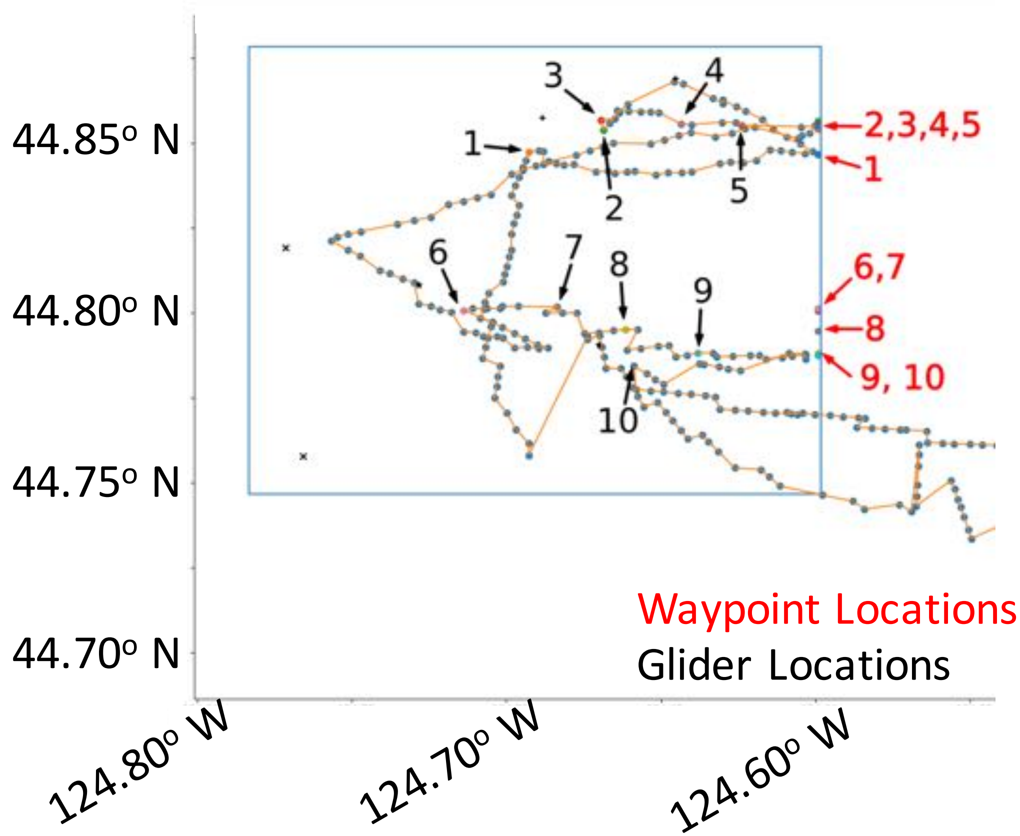

August 2018 deployment

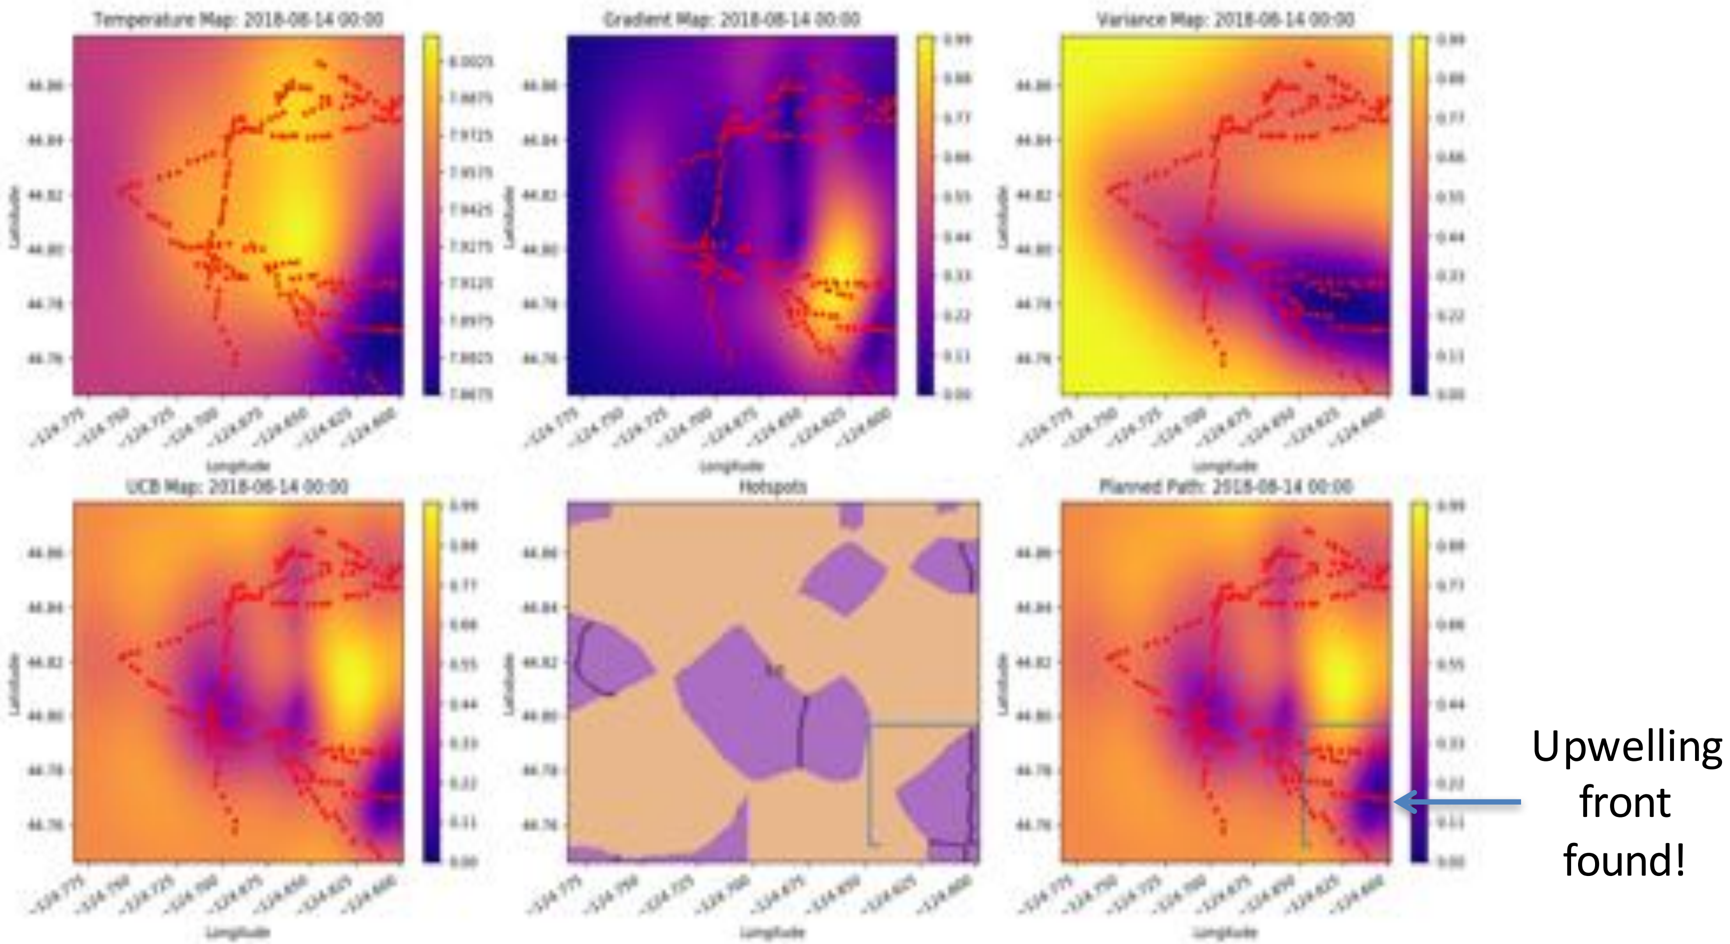

Tracking ocean mixing fronts

Oregon State University

- Ocean fronts occur at the interface between distinct masses of water
  - Warm & Cool
  - Salty & Fresh
- Biological hotspots form at interface between warm water and cool, nutrient-rich water
- Physical processes which drive mixing along fronts are not well understood

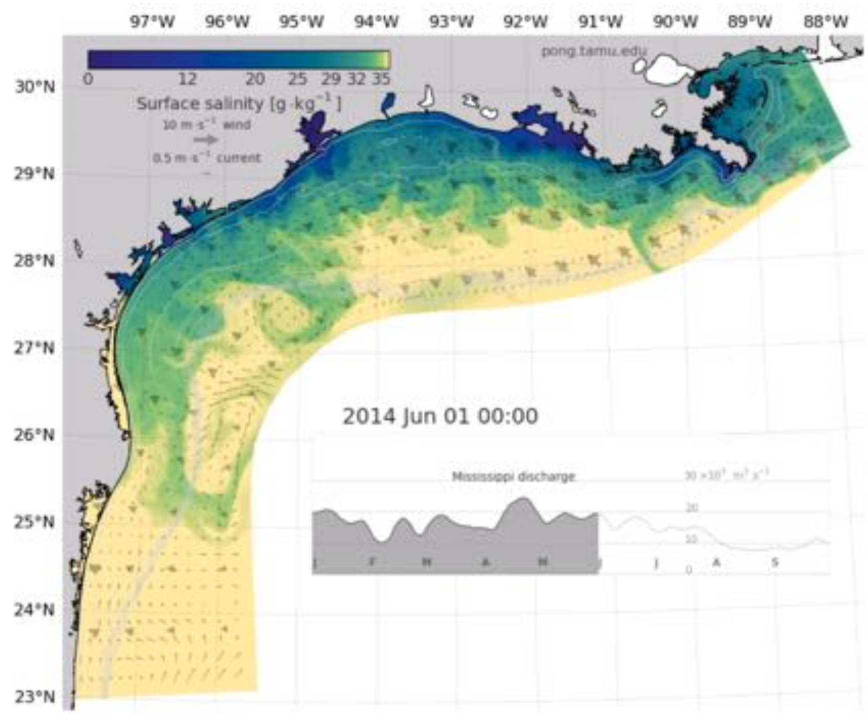

Image Credit: Physical Oceanography Numerical Group (PONG), Texas A&M

Heterogeneous vehicles

- **Robotic Ocean Surface Sampler (ROSS)**
  - Developed at Oregon State University (Dr. Jonathan Nash)
  - Continuous satellite and radio communications
  - .5 - 1 m/s velocity depending on amount of seaweed fouling
  - Two assets

Heterogeneous vehicles

- **Robotic Ocean Surface Sampler (ROSS)**

- **Slocum G3 Glider**
  - Built by Teledyne Industries
  - Infrequent satellite data connection
    - Updates only when surfaced (~2 hours)
  - .4 m/s velocity
  - Four assets

Collaborations with scientists

Oregon State University

Robotics

- Use automation to reduce operator load during sampling tasks
- Planning algorithms coordinate multiple heterogeneous assets

Improve quality of ocean science data through adaptive autonomous sampling

Improve autonomy with knowledge of ocean processes in decision making

Oceanography

- Decades of experience in manually sampling ocean features
- Deep understanding of dynamic physical ocean processes

Methods: Environment estimation

➢ Gaussian Processes (GPs) are a commonly used tool in field robotics
- Provide scalar field and uncertainty estimates
- Several drawbacks
  - Standard RBF Kernel does not extrapolate outside of data points well
  - Scaling issues with large sample data sets

➢ Inform GP with nearest-neighbors prior
- Nearest Neighbors environment estimation provides initial guess for front location
- Allows world model to extrapolate front location beyond data collected

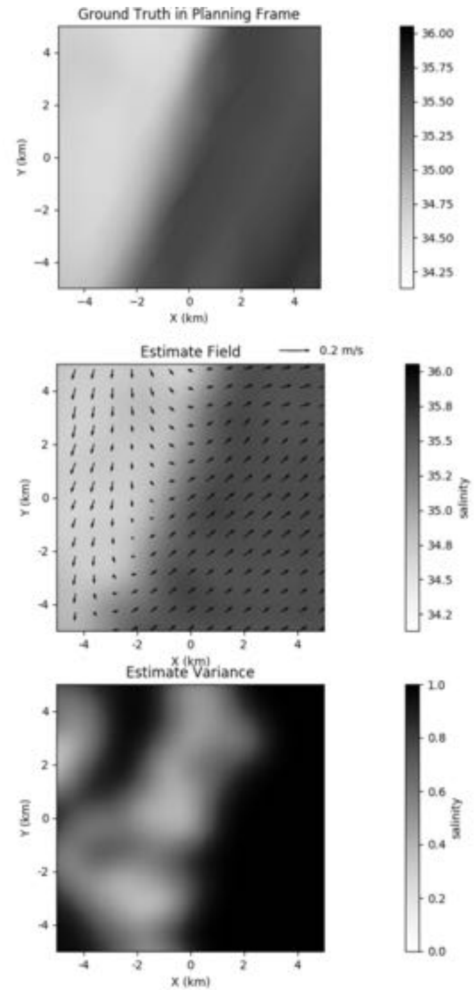

Methods: Environment estimation

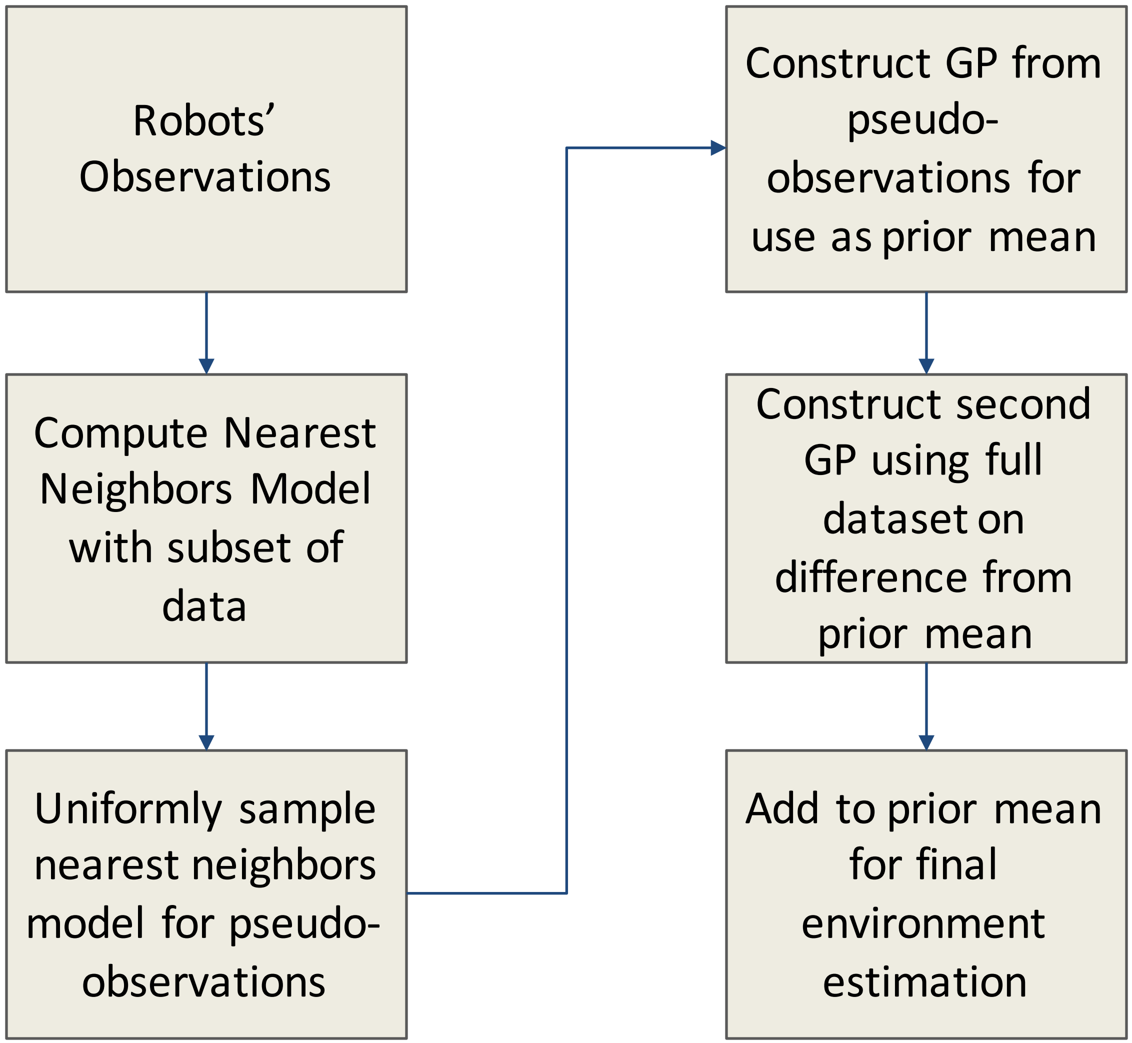

Methods: Adaptive decision making

- **Key challenge:** Adaptive control of 4-6 heterogeneous assets for sampling task
  - *Scalability* - Handle large planning space
  - *Adaptability* - Enable robots to adapt behavior to new information
  - *Flexibility* - Account for the realities of physical robot operation

Methods: Monte Carlo Tree Search (MCTS)

➢ Incremental tree search algorithm
  - Leverages *biased* random sampling
  - Exploits “smoothness” of search space
  - Anytime algorithm
  - Only requires evaluation of *full* paths
  - Can incorporate problem-specific heuristics

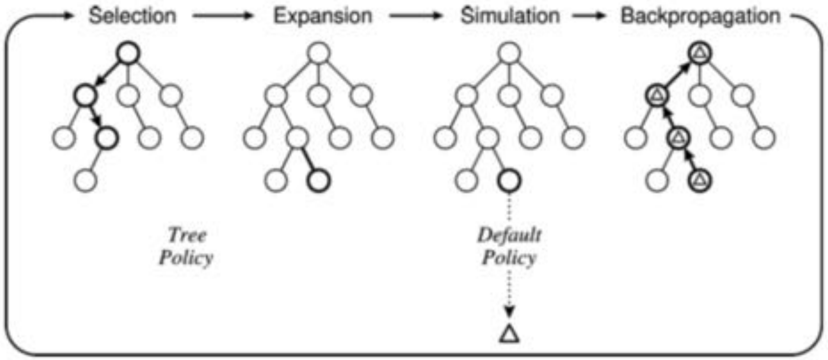

[Browne, 2012]

Methods: Adaptive decision making

Oregon State University

- **Key challenge:** Adaptive control of 4-6 heterogeneous assets for sampling task
  - *Scalability* - Handle large planning space
  - *Adaptability* - Enable robots to adapt behavior to new information
  - *Flexibility* - Account for the realities of physical robot operation

- **Solution:** Iterative optimization using Monte Carlo Tree Search
  - *Scalability* - Iterative optimization sidesteps planning in joint space
  - *Adaptability* - MCTS is proven algorithm for informative path planning tasks
  - *Flexibility* - Simple to add / remove robots from planning step

Methods: Multi-robot coordination

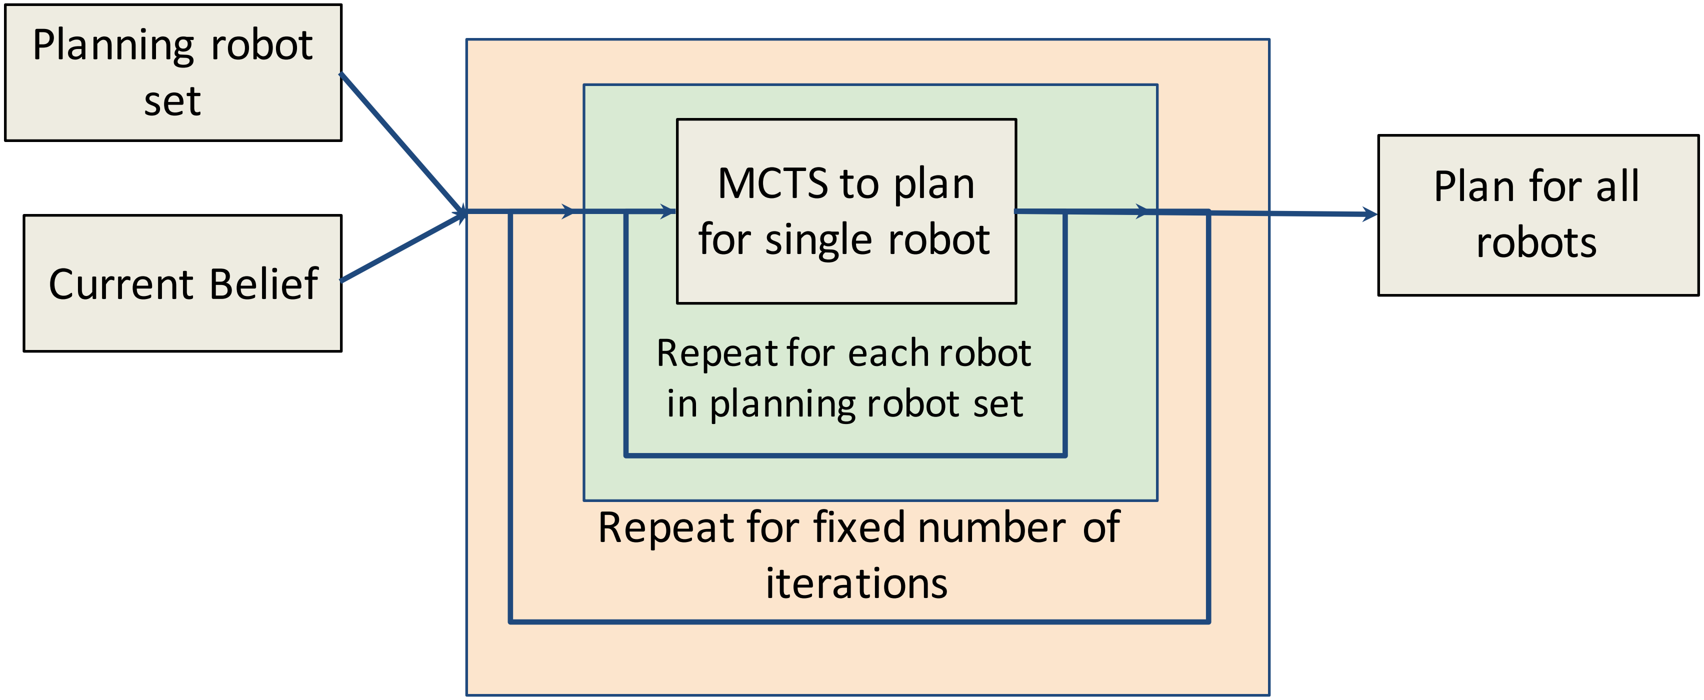

Methods: Objective function

- ➢ Planner’s output is only as good as the objective function
  - Careful tuning of objective function weightings and parameters to achieve good performance in oceanographic task

$$\sum_{\mathcal{E} \in Path} \left[\int_{\mathcal{E}} \left(\alpha_1 * \left| \frac{\delta S}{\delta x} \right| + \alpha_2 * \sigma^2 + \alpha_3 * \left| \frac{\delta S}{\delta x} \right| * \sigma^2 \, dx \right) * T(\mathcal{E}) * Continuity(\mathcal{E}) \right]$$

Methods: Objective function

Oregon State University

- Planner's output is only as good as the objective function
  - Careful tuning of objective function weightings and parameters to achieve good performance in oceanographic task

$$\sum_{\mathcal{E}\in Path}\left[\int_{\mathcal{E}}\left(\underline{\alpha_1 * \left|\frac{\delta S}{\delta x}\right|} + \alpha_2 * \sigma^2 + \alpha_3 * \left|\frac{\delta S}{\delta x}\right| * \sigma^2 \, dx\right) * T(\mathcal{E}) * Continuity(\mathcal{E})\right]$$

Exploitation Term: Magnitude of salinity gradient along edge

Methods: Objective function

Oregon State University

- Planner's output is only as good as the objective function
  - Careful tuning of objective function weightings and parameters to achieve good performance in oceanographic task

$$\sum_{\mathcal{E} \in Path} \left[\int_{\mathcal{E}} \left(\alpha_1 * \left| \frac{\delta S}{\delta x} \right| + \underline{\alpha_2 * \sigma^2} + \alpha_3 * \left| \frac{\delta S}{\delta x} \right| * \sigma^2 \, dx \right) * T(\mathcal{E}) * Continuity(\mathcal{E}) \right]$$

Exploration Term: Amount of edge novelty from GP uncertainty

Methods: Objective function

Oregon State University

- Planner's output is only as good as the objective function
  - Careful tuning of objective function weightings and parameters to achieve good performance in oceanographic task

$$\sum_{\mathcal{E} \in Path} \left[\int_{\mathcal{E}} \left(\alpha_1 * \left| \frac{\delta S}{\delta x} \right| + \alpha_2 * \sigma^2 + \alpha_3 * \left| \frac{\delta S}{\delta x} \right| * \sigma^2 \, dx \right) * T(\mathcal{E}) * Continuity(\mathcal{E}) \right]$$

Combination Term: Encourages exploration in areas likely to hold gradient similar to UCB. Penalizes repeatedly sampling the same gradient

Methods: Objective function

Oregon State University

- Planner's output is only as good as the objective function
  - Careful tuning of objective function weightings and parameters to achieve good performance in oceanographic task

$$\sum_{\mathcal{E} \in Path} \left[\int_{\mathcal{E}} \left(\alpha_1 * \left| \frac{\delta S}{\delta x} \right| + \alpha_2 * \sigma^2 + \alpha_3 * \left| \frac{\delta S}{\delta x} \right| * \sigma^2 \, dx \right) * \underline{T(\mathcal{E})} * Continuity(\mathcal{E}) \right]$$

Temporal Discount: Reduce reward for future edges due to uncertainties

Methods: Objective function

Oregon State University

- Planner's output is only as good as the objective function
  - Careful tuning of objective function weightings and parameters to achieve good performance in oceanographic task

$$\sum_{\mathcal{E} \in Path} \left[\int_{\mathcal{E}} \left(\alpha_1 * \left| \frac{\delta S}{\delta x} \right| + \alpha_2 * \sigma^2 + \alpha_3 * \left| \frac{\delta S}{\delta x} \right| * \sigma^2 \, dx \right) * T(\mathcal{E}) * \underline{Continuity(\mathcal{E})} \right]$$

Continuity Weighting: ADCPs and other sensors cannot collect data when vehicle turns, so encourage straighter paths by penalizing edges that are not collinear with prior and subsequent edges

Example autonomy output

Estimate

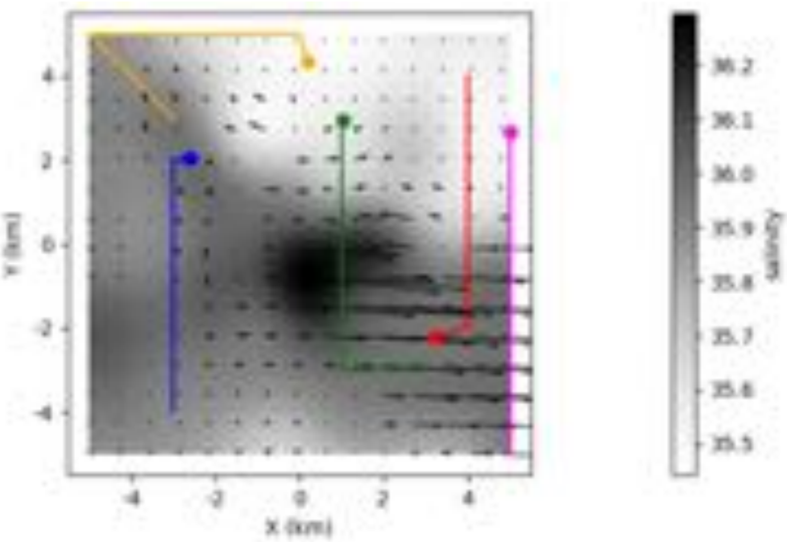

Novelty

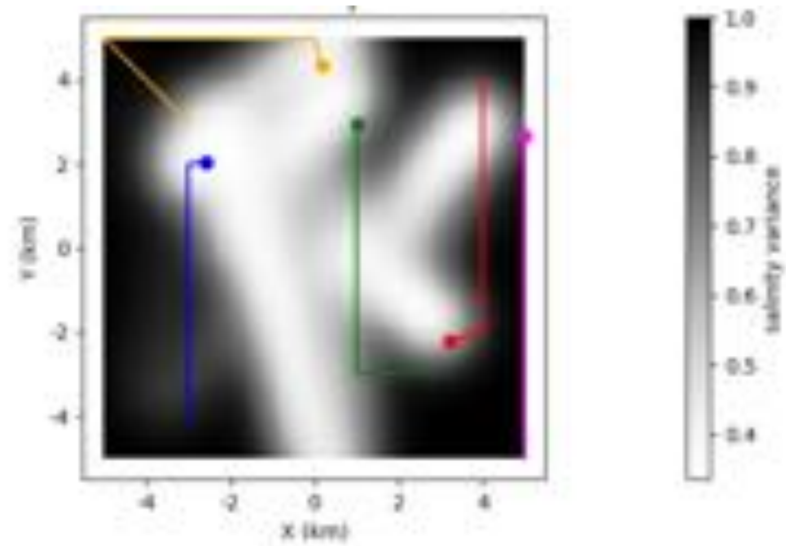

Methods: Lagrangian moving frame

Oregon State University

- Multiple sources of time-varying dynamics of ocean fronts
  - Local forces and mixing change shape of front
  - Front is moved by large global currents

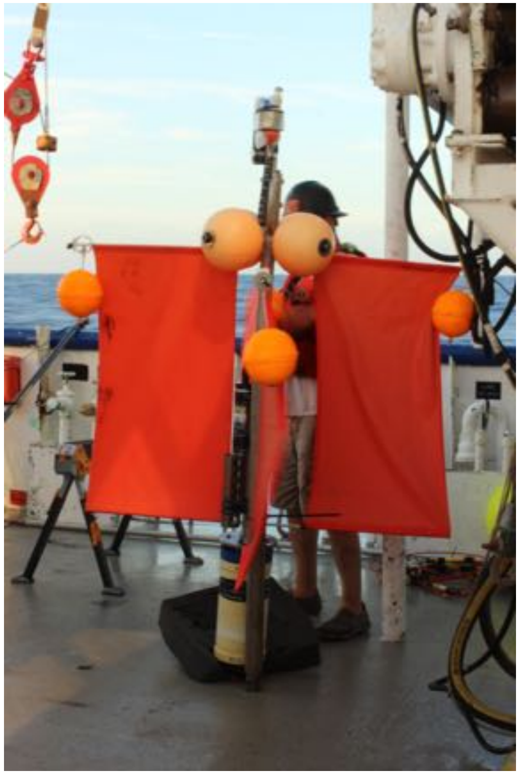

Methods: Lagrangian moving frame

- Tracking Lagrangian “packet” of water increases the time over which sensor observations remain relevant
- Multiple sources of frame motion estimate
  - Physical drifter in water
  - Virtual drifter informed by shipboard Acoustic Doppler Current Profiler (ADCP)
- Integrate data from multiple sources with Kalman Filter
- Project paths into future predicted positions of moving frame using velocity estimates

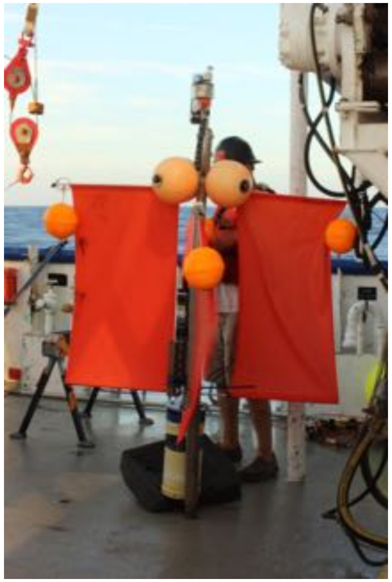

Decision support GUI

- Interface provides state of autonomy system at a glance

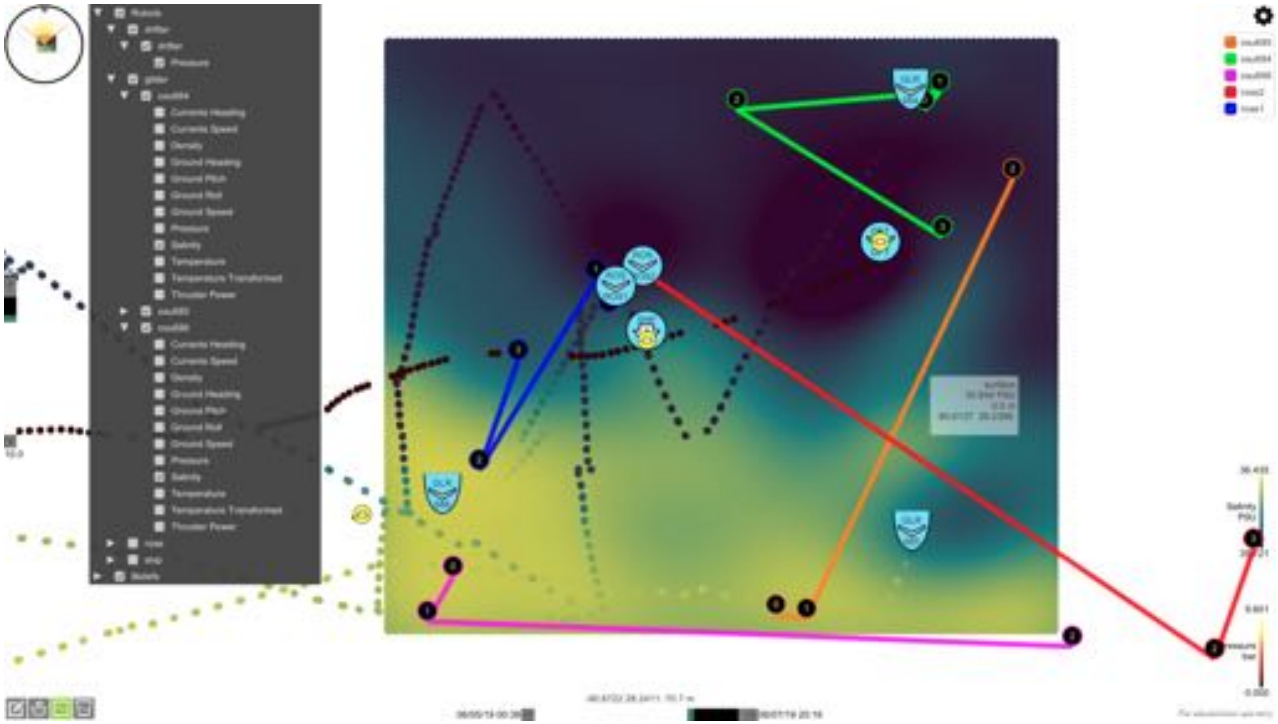

Decision support GUI

- Interface provides state of autonomy system at a glance

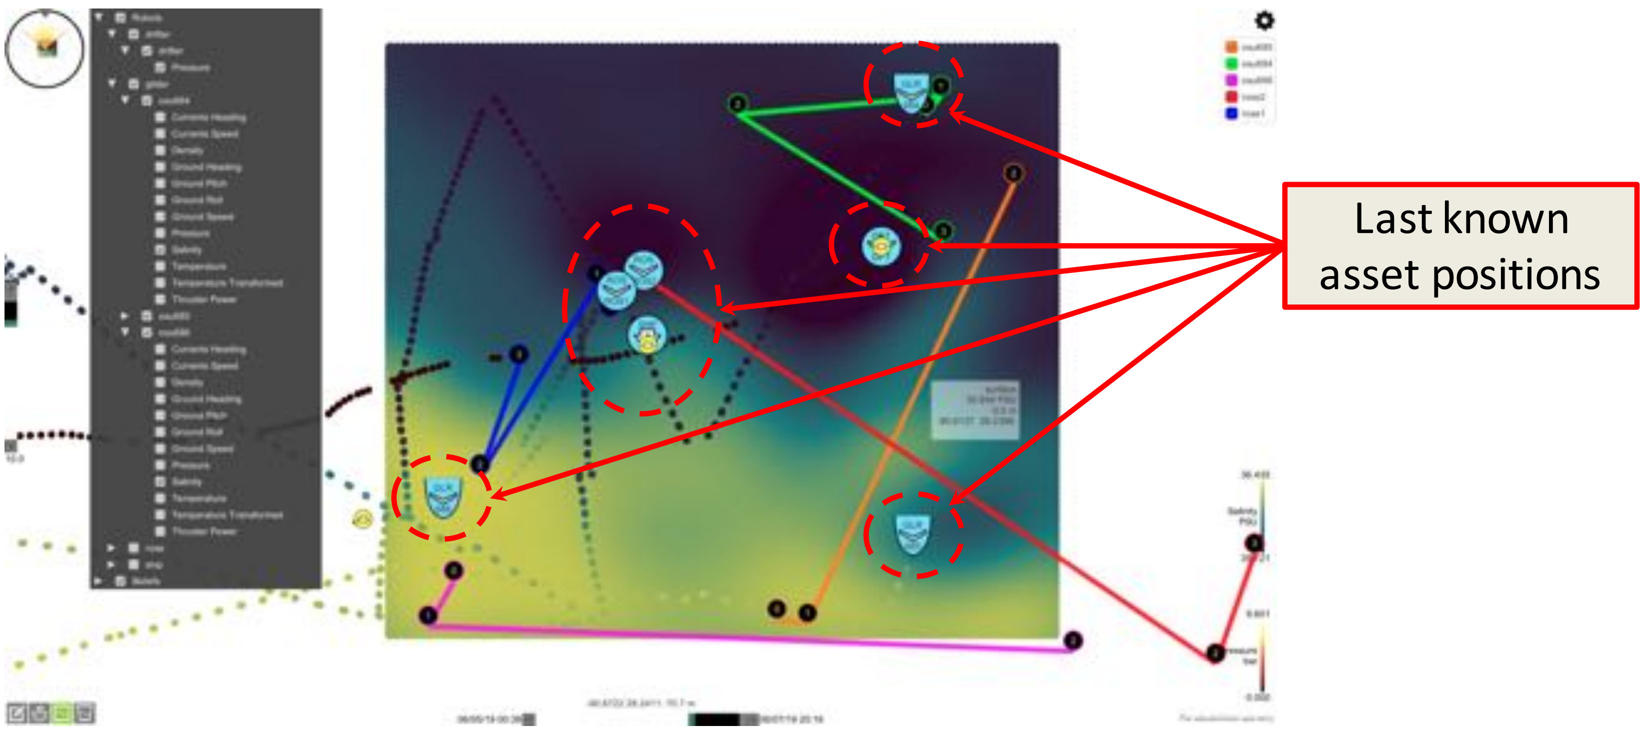

Decision support GUI

- Interface provides state of autonomy system at a glance

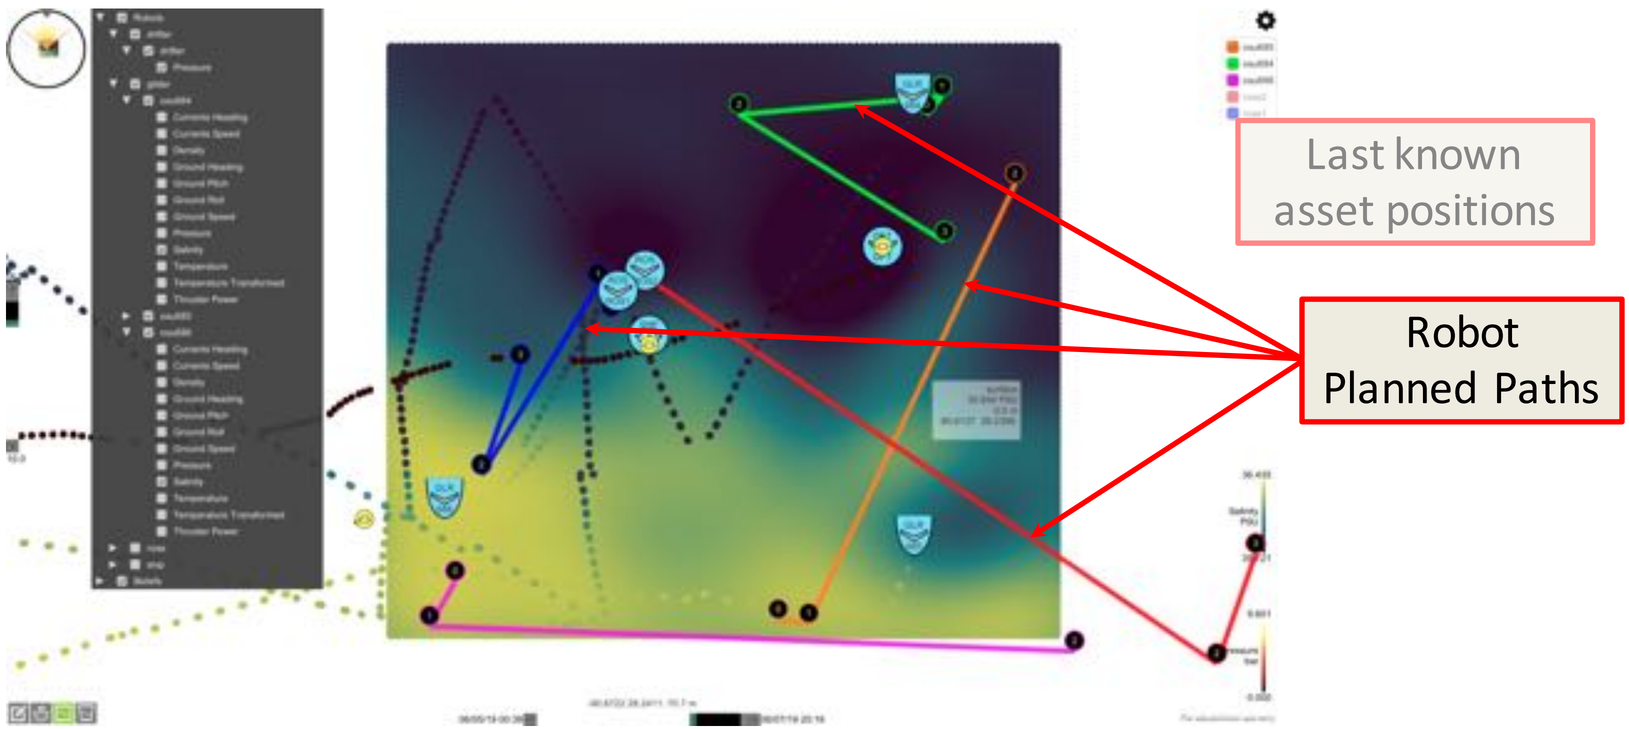

Decision support GUI

- ➢ Interface provides state of autonomy system at a glance

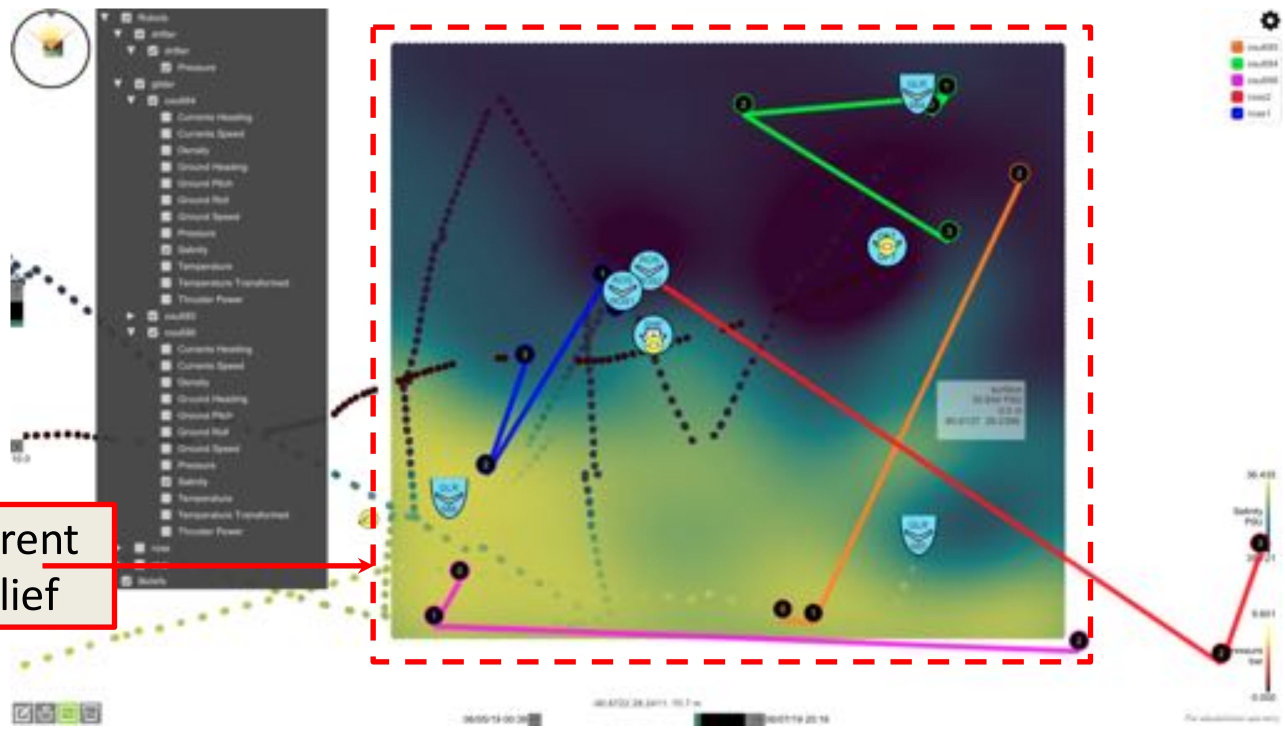

Last known asset positions

Robot Planned Paths

Map of current GP-NN belief

Decision support GUI

Last known obstacle locations

Last known asset positions

Robot Planned Paths

Map of current GP-NN belief

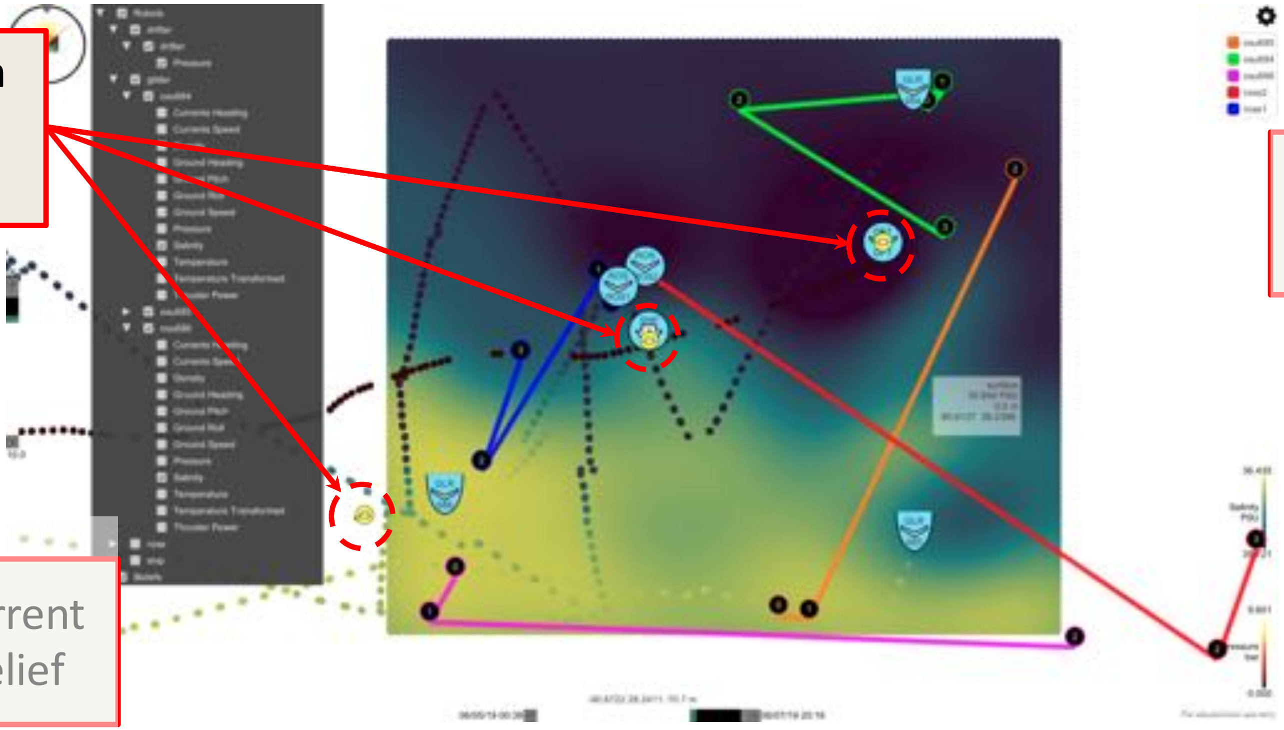

Timeline of experiments

*Timeline approximately to scale

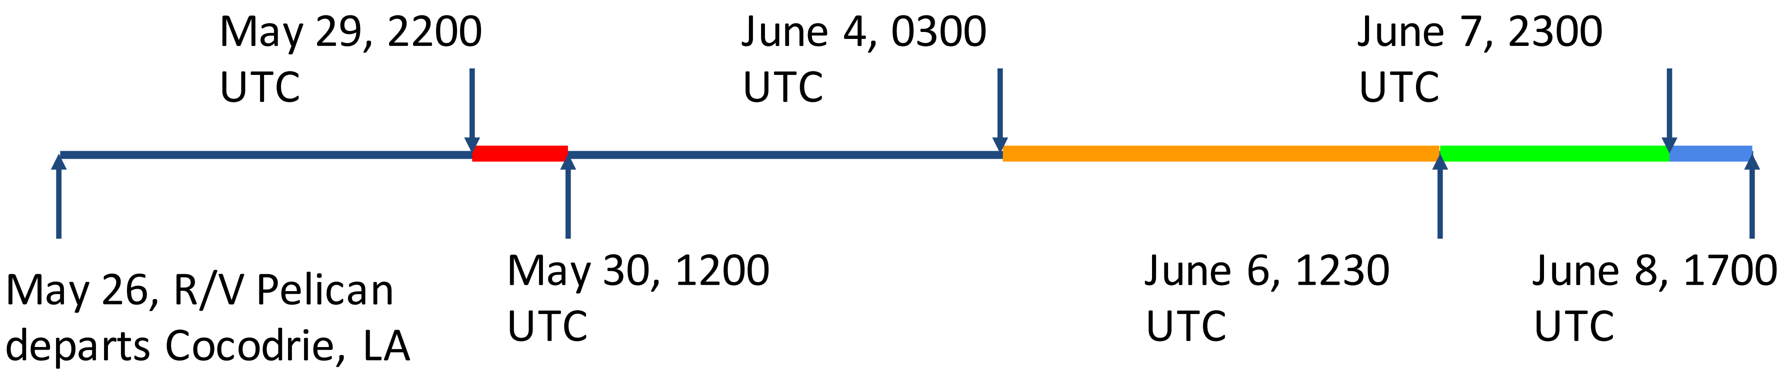

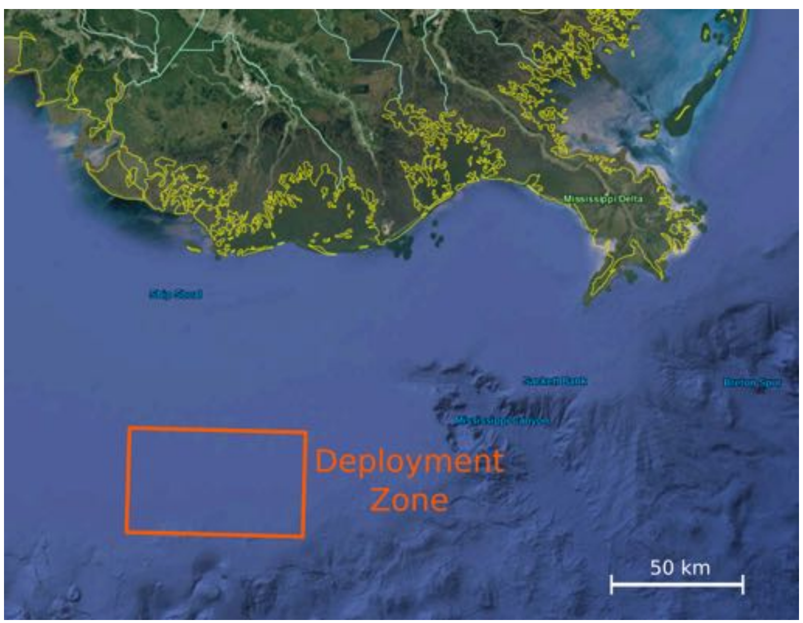

- Robots under manual control
- Autonomy Experiment 1 (~14h)
- Autonomy Experiment 2 (~57h)
- Autonomy Experiment 3 (~35h)
- Autonomy Experiment 4 (~18h)

Experiments in Gulf of Mexico

Oregon State University

- ➢ Autonomy Experiment 1 (14 hours)
  - Initial systems testing and debugging of autonomy system
  - 3 Slocum gliders and R/V Pelican
- ➢ Autonomy Experiment 2 (57 hours)
  - Experimented with different types of planning frames
    - Static, Drifter-based, ADCP-Based
  - 3-4 Slocum Gliders and 0-2 ROSS vehicles
- ➢ Autonomy Experiment 3 (35 hours)
  - Experimented with Combination Term in MCTS objective function
  - 3-4 Slocum Gliders and 0-2 ROSS vehicles
- ➢ Autonomy Experiment 4 (18 hours)
  - Use best parameters discovered in previous 3 experiments
  - 3 Gliders and 0-2 ROSS vehicles

Vehicle up time

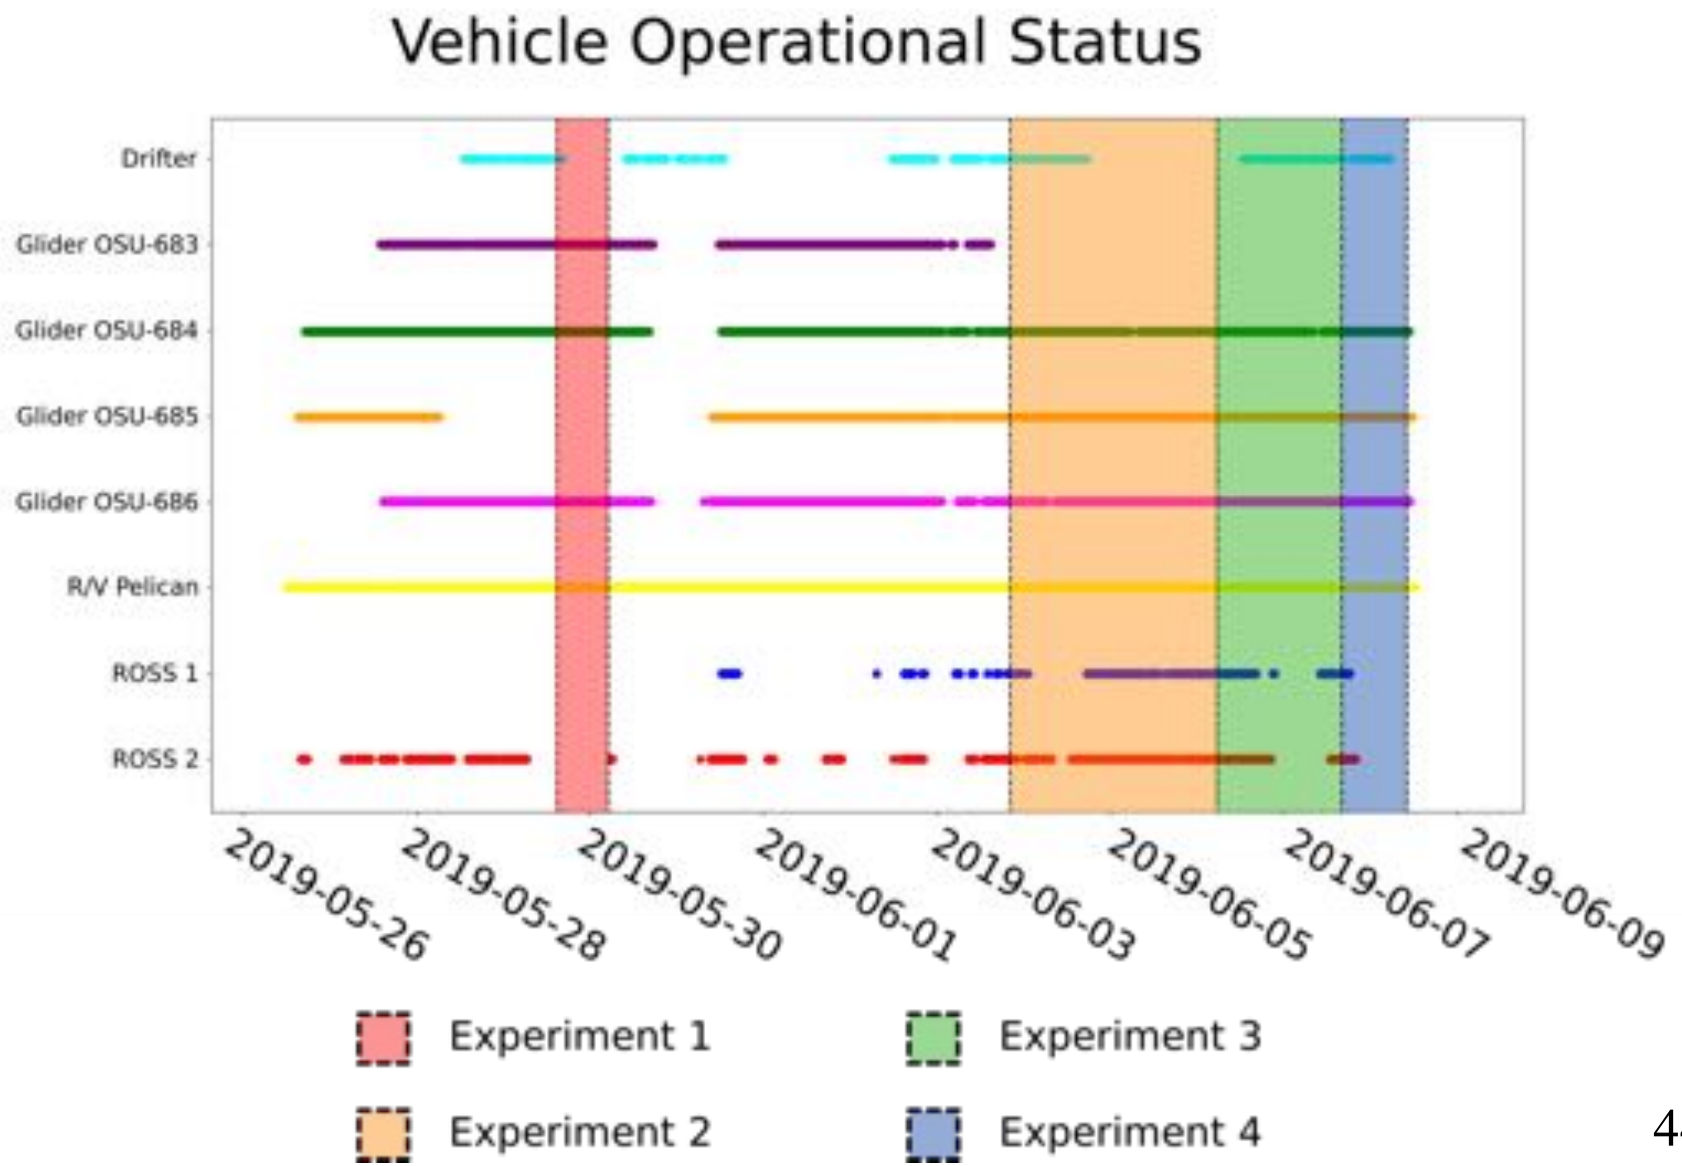

Example front crossings

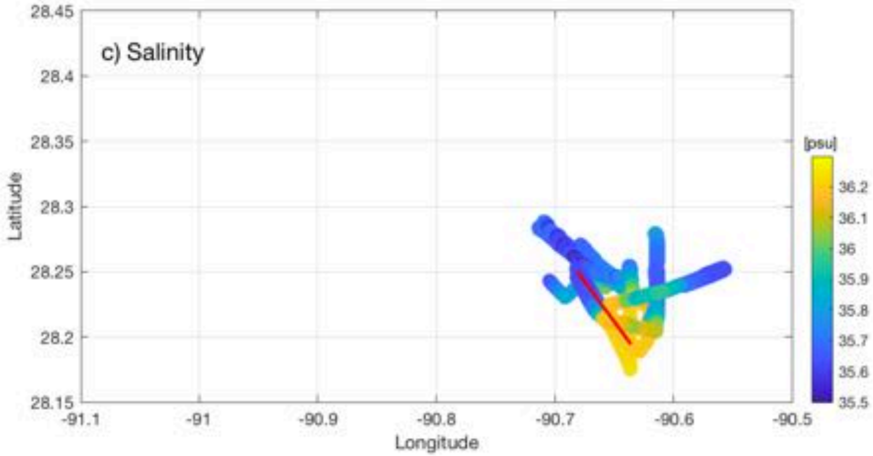

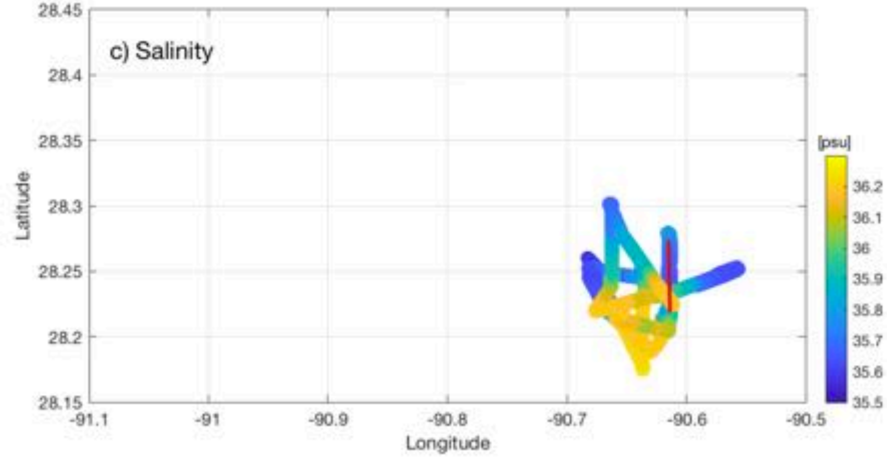

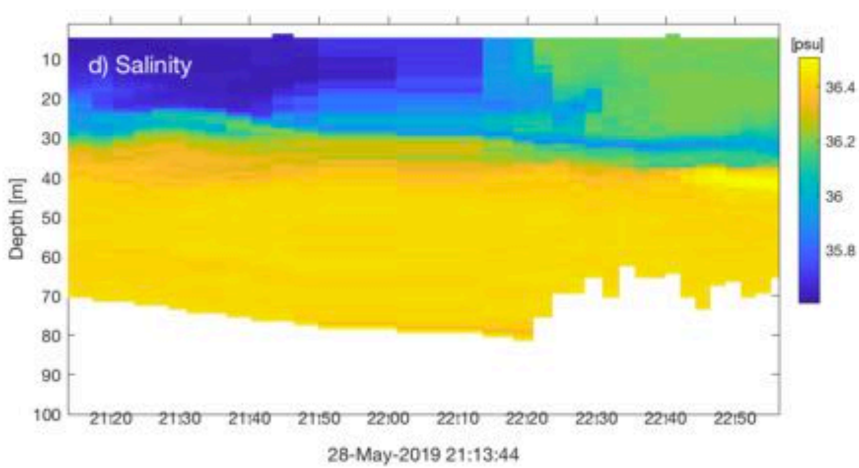

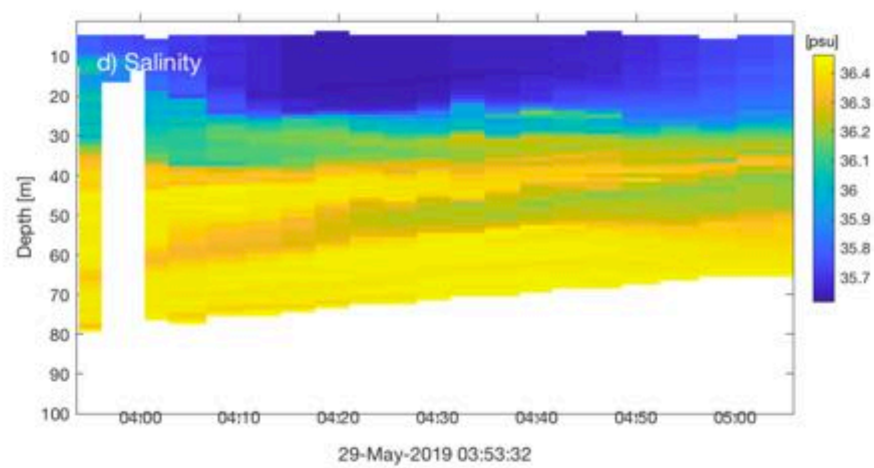

Example front crossings

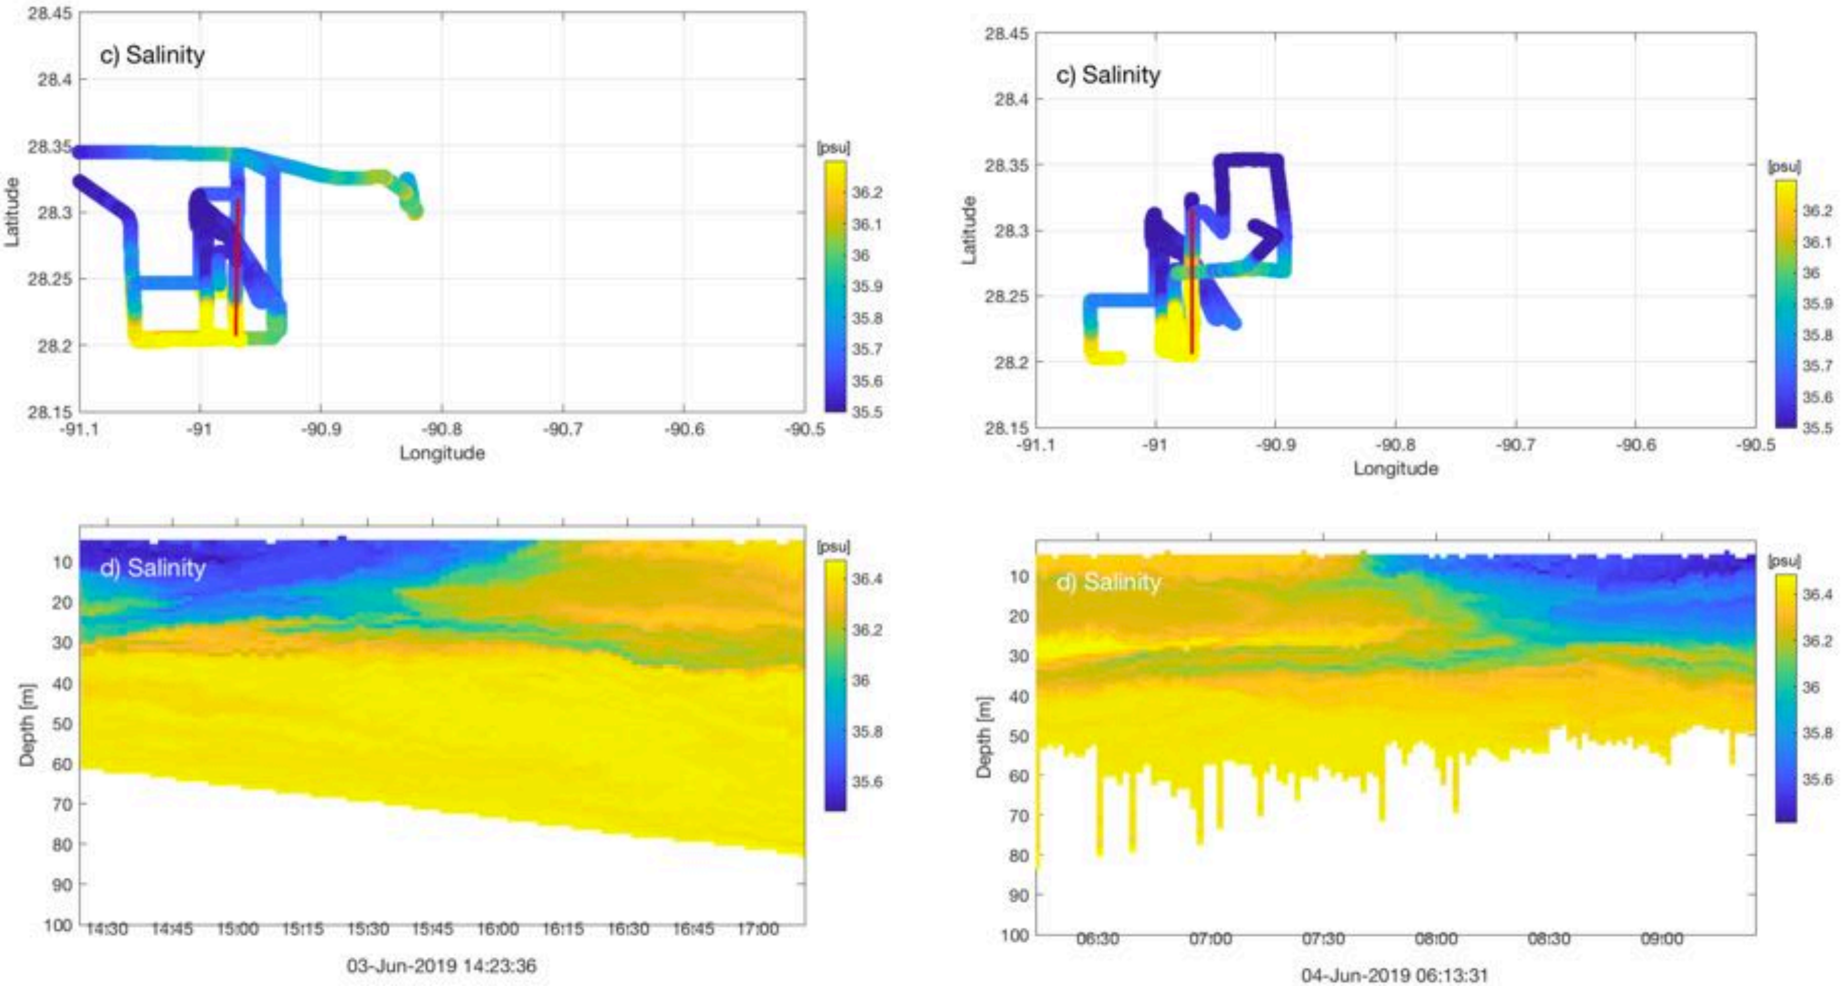

Example front crossings

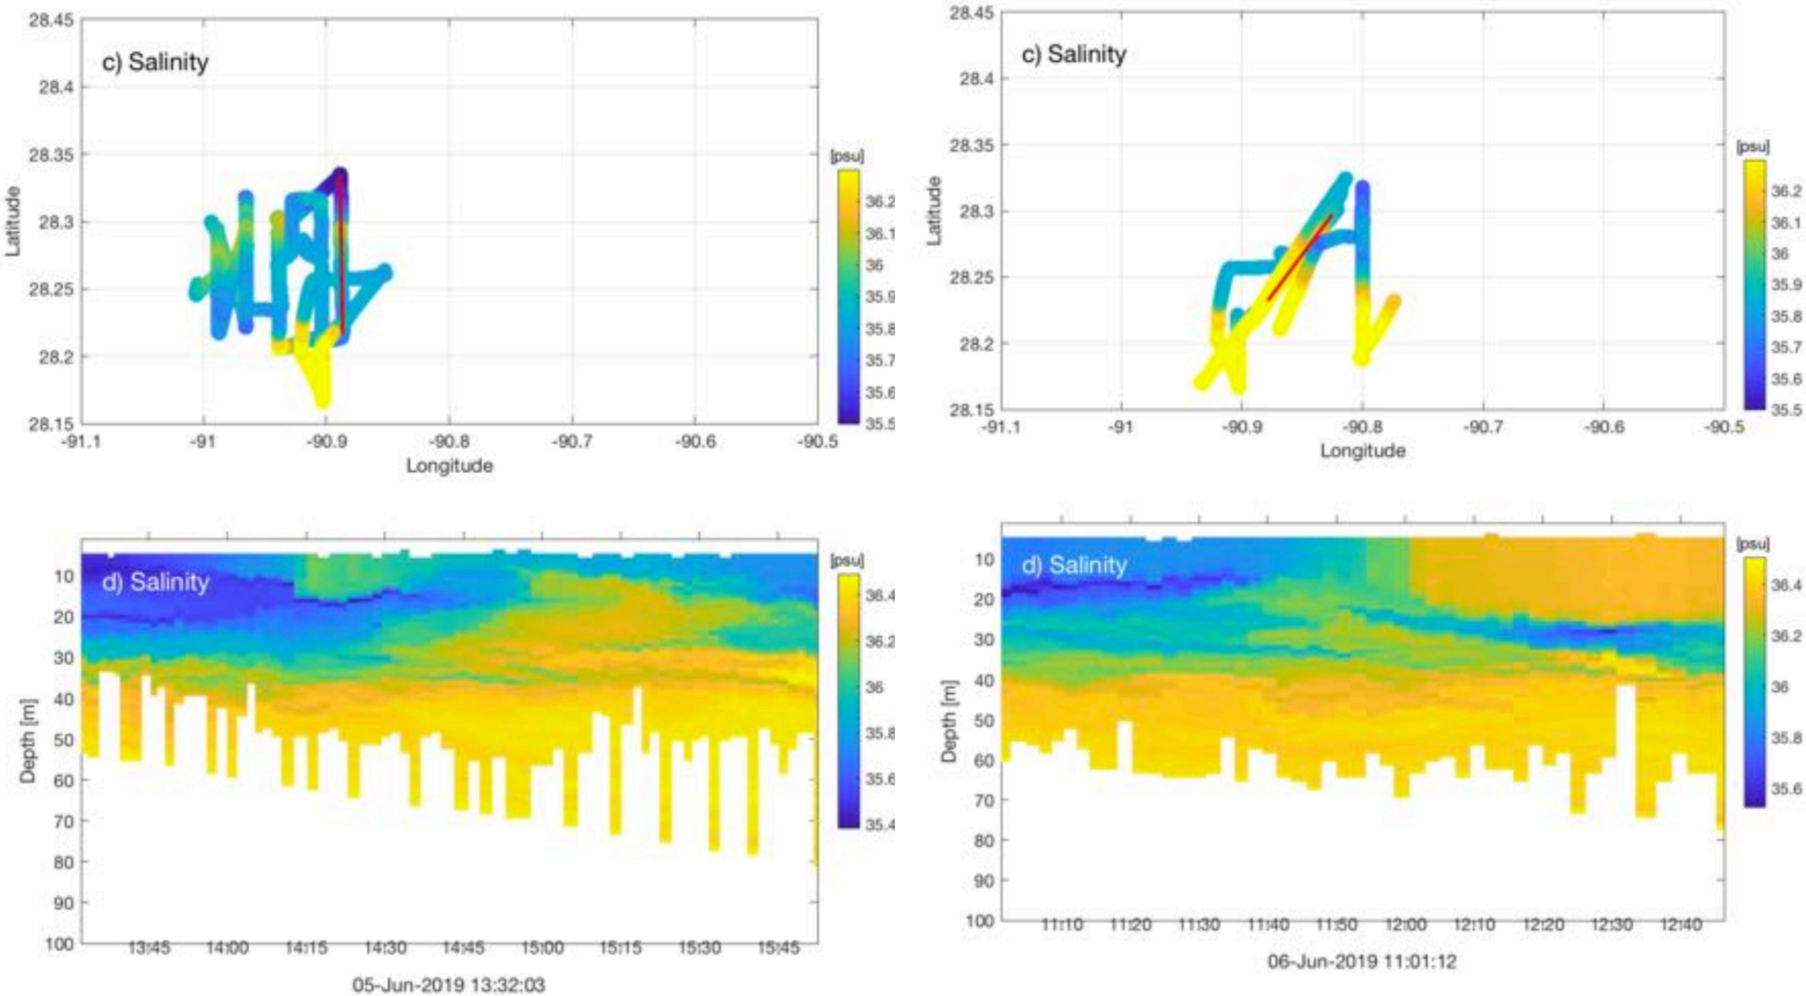

Preliminary autonomy results

Oregon State University

- Autonomy was able to identify a front and track its motion across an ~18 hour period during longest experiment
- More results coming (hot off the presses)!

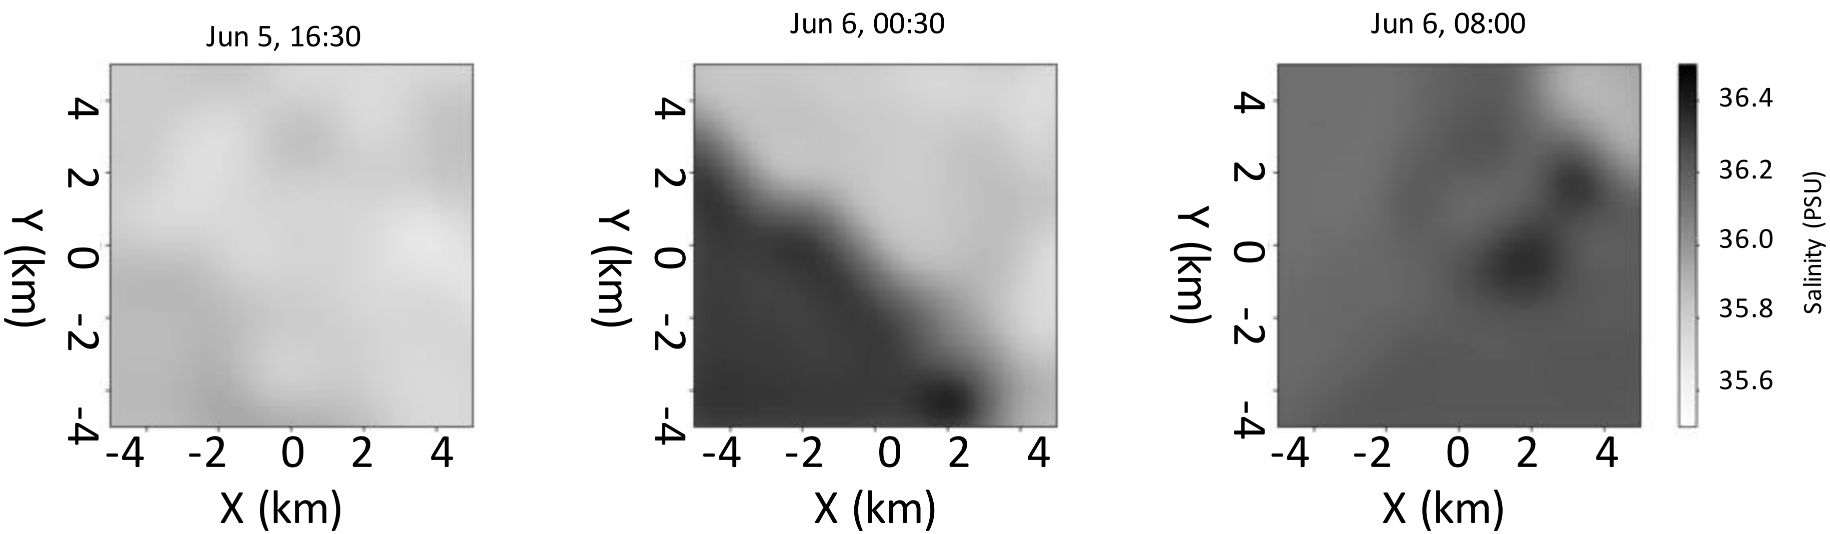

Preliminary GUI results

Oregon State University

- GUI was often used for visualization of vehicle positions and data
  - Both during autonomy and manual experiments (>90% of the time)

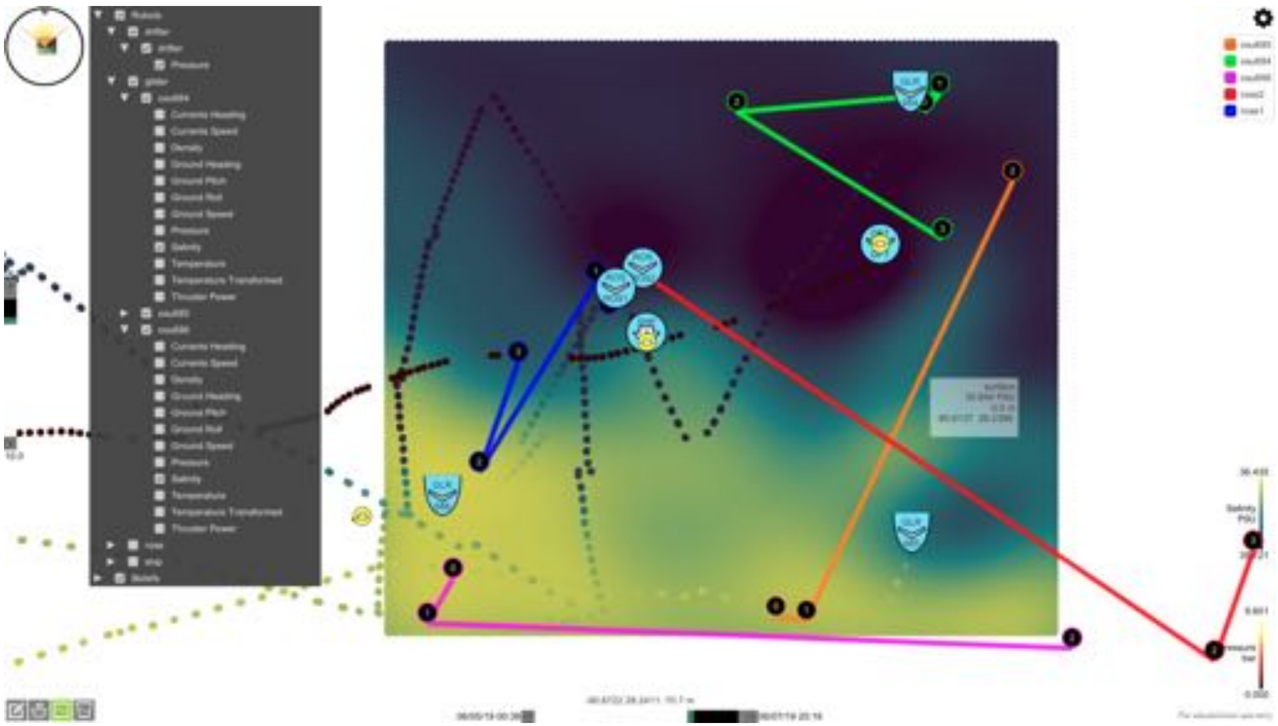

Oregon State University Robotics

Oregon State University

- PhD/MS programs (since 2014)
- 11 core faculty
- ~80 grad students
- 45 affiliated faculty
- Graf hall collaborative space
- Legged, aerial, aquatic surface, underwater, and ground robots

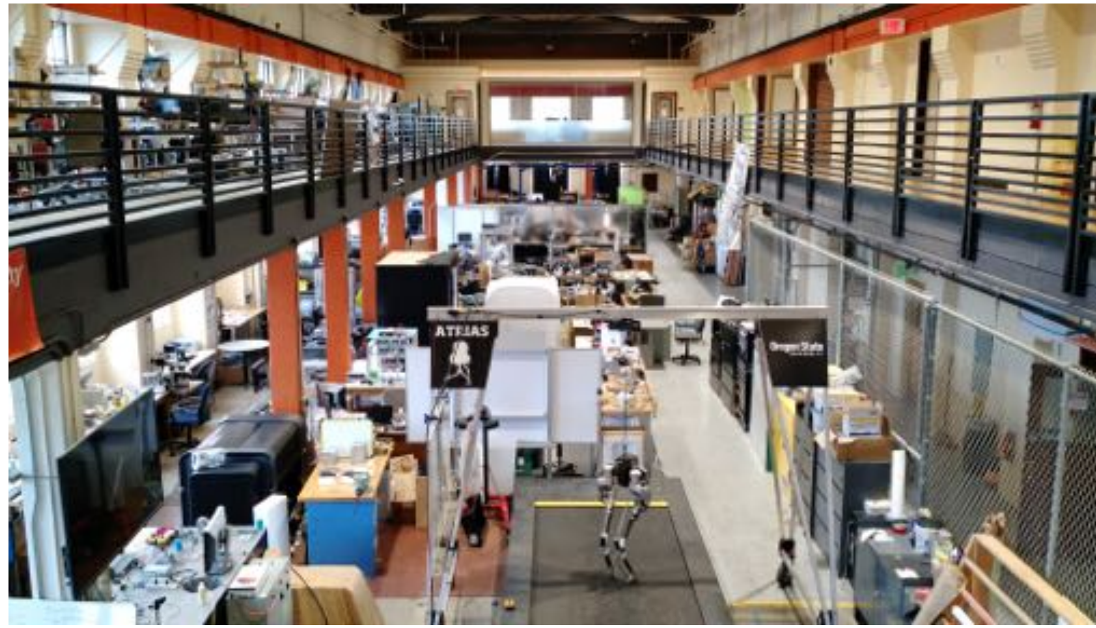

Acknowledgements

Geoffrey A. Hollinger
Robotic Decision Making Laboratory
geoff.hollinger@oregonstate.edu
http://research.engr.oregonstate.edu/rdml/

Funding: ONR N00014-17-1-2581,
NSF IIS-1723924,
and W. M. Keck Foundation

Seth McCammon
(PhD student, OSU)

Graeme Best
(Postdoc, OSU)

Faculty (OSU):
Julie A. Adams
Kipp Shearman
Jonathan Nash
Jack Barth

Gilberto Marcon
(PhD Student, OSU)

Dylan Jones
(PhD student, OSU)

Matt Frantz
(MS student, OSU)