

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

Geoff Hollinger

Assistant Professor

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

May, 2018

Collaborators: Nicholas Lawrance (ETH Zurich),
Thane Somers (Oregon State), Gaurav Sukhatme (USC)

Land and Sea, 2005-2010



2005 DARPA Grand Challenge



DepthX: Deep Phreatic Thermal Explorer

Marine Robotics (2018)

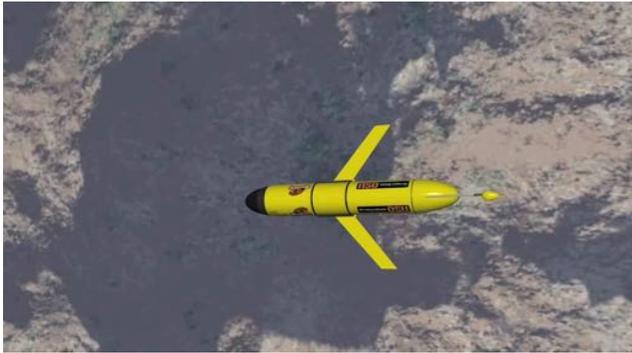


Photo: NASA



Photo: LA Times



Photo: NIWA Science Web



Photo: Bluefin Robotics



Photo: Hydro international

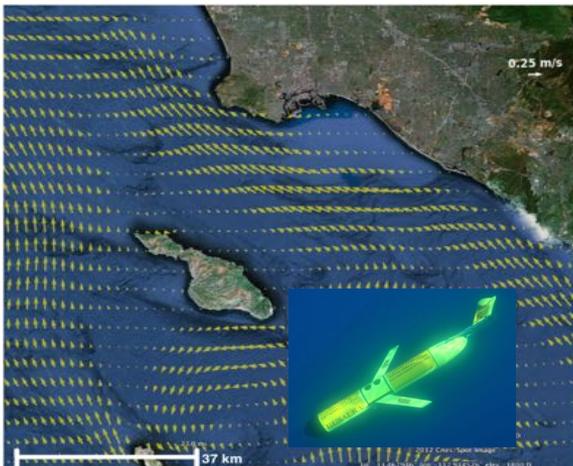
Marine Robotics (2018)

- **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
 - Robustness to errors
 - Shared autonomy with technicians



Marine Robotics

Underwater
Data Collection



(JFR '13, JFR '16, RA-L '17)

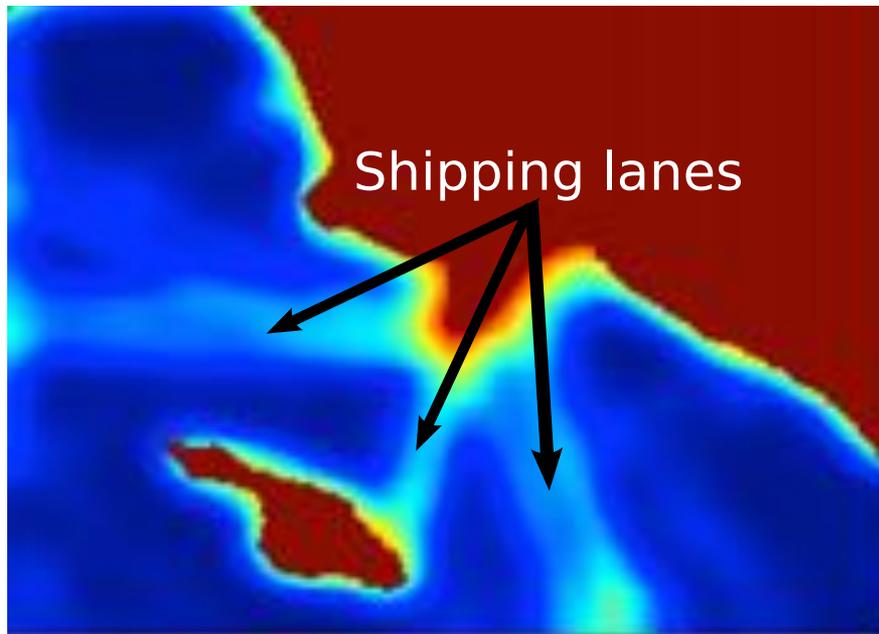
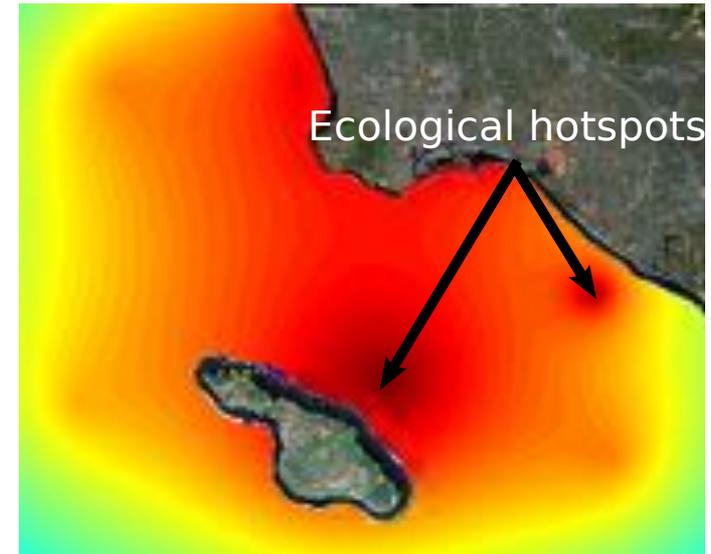
Human-robot
Teaming



(AURO '16)

Ocean data collection

- **Balance:** risk (ships/shore) and information quality (hotspots)
- **Uncertainty:** from ocean currents



Trajectory optimization

Find optimized trajectory: ξ^*

Given scientist's initial trajectory: ξ_0

Risk function: $R(\xi)$

Deviation function: $D(\xi, \xi_0)$

- Minimize:

$$\xi^* = \operatorname{argmin}_{\xi \in \Psi} D(\xi, \xi_0) + \alpha R(\xi)$$



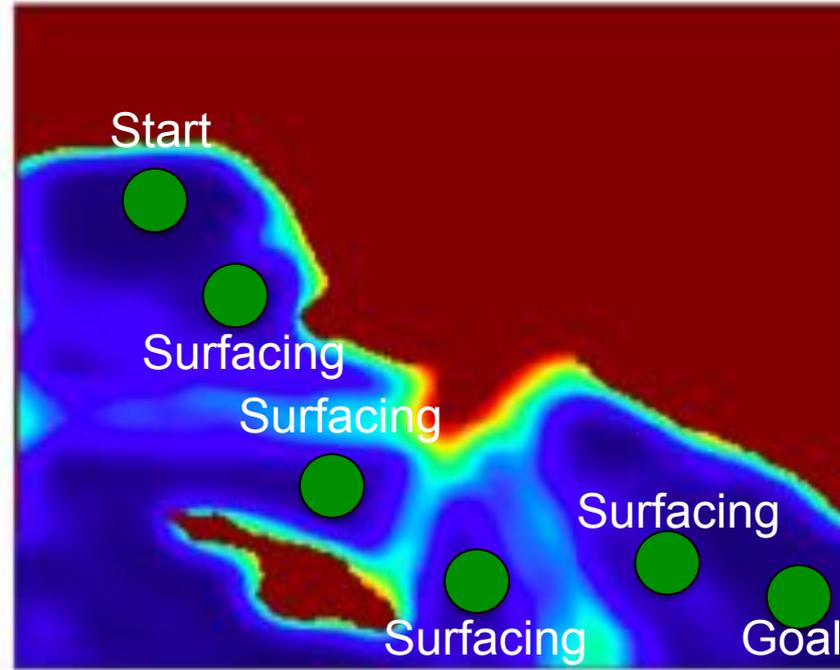
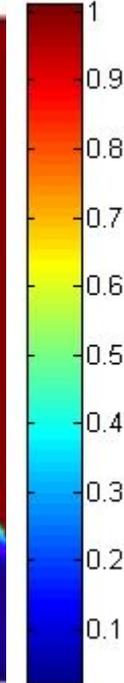
Weighting
parameter

Calculating risk

Sum risk over surfacing locations

Normalized

Risk



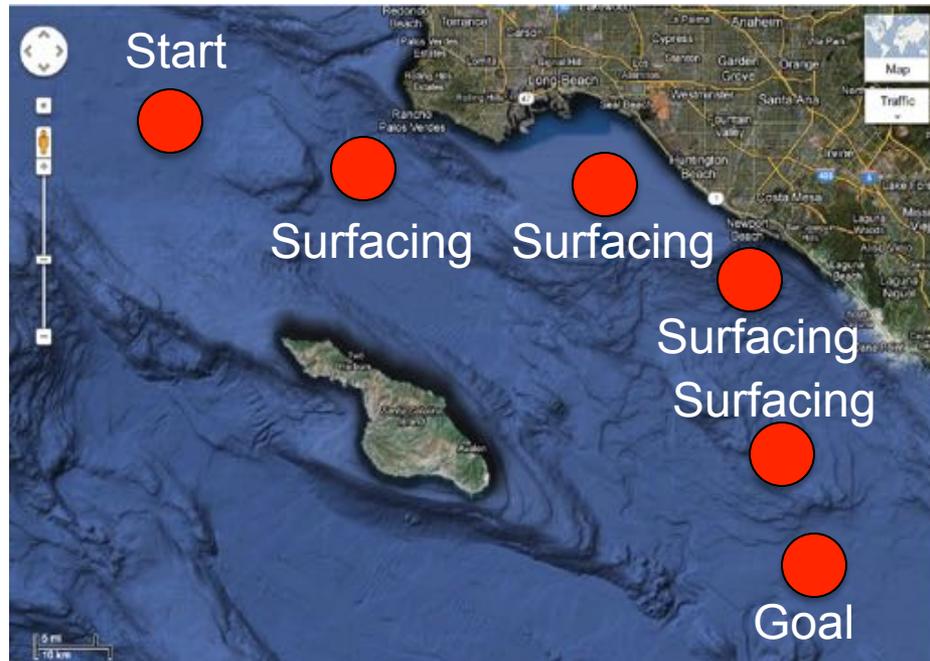
Area of the Southern
California bight

Risk map built from historical
shipping data

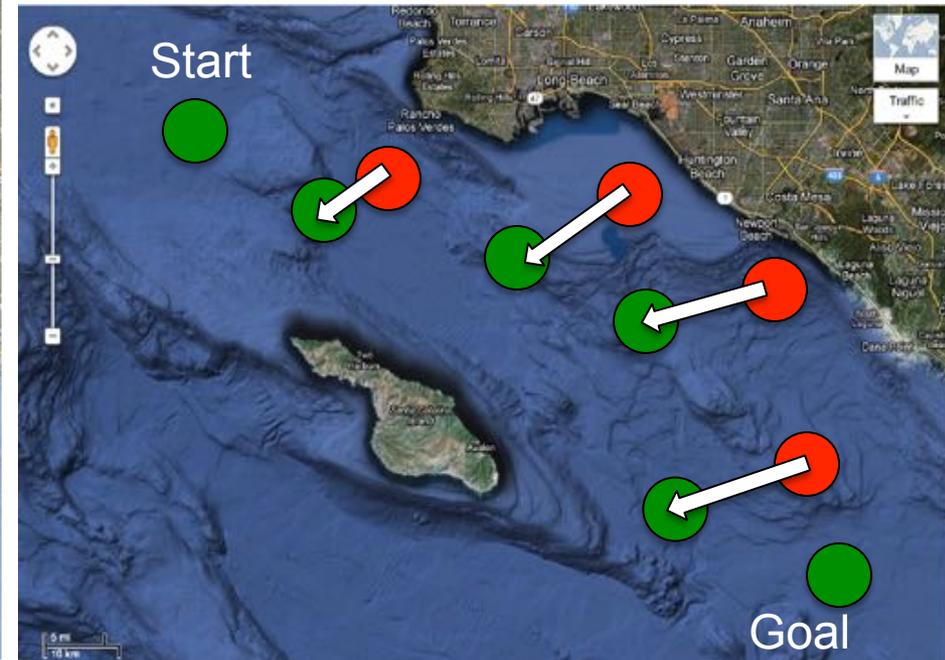
$$\xi^* = \operatorname{argmin}_{\xi \in \Psi} D(\xi, \xi_0) + \alpha R(\xi)$$

Calculating deviation

Sum deviation over surfacing locations



Scientist's trajectory



Optimized trajectory

$$\xi^* = \operatorname{argmin}_{\xi \in \Psi} D(\xi, \xi_0) + \alpha R(\xi)$$

Optimizing the trajectory

- Given a method for calculating deviation and risk (e.g., sum over surfacings)
- Iterative optimization (value iteration)

$$C(s, a) \leftarrow \sum_{s'} T(s' | s, a) (\Delta D(s', \xi_0) + \alpha \Delta R(s'))$$

State/action
pair

s'

State transition
(using uncertainty)

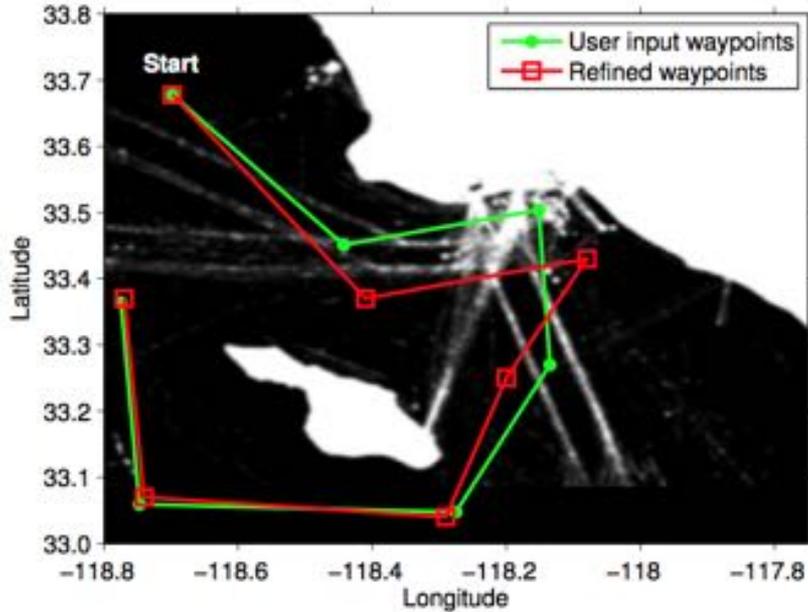
Deviation caused
by state s'

Risk caused by
state s'

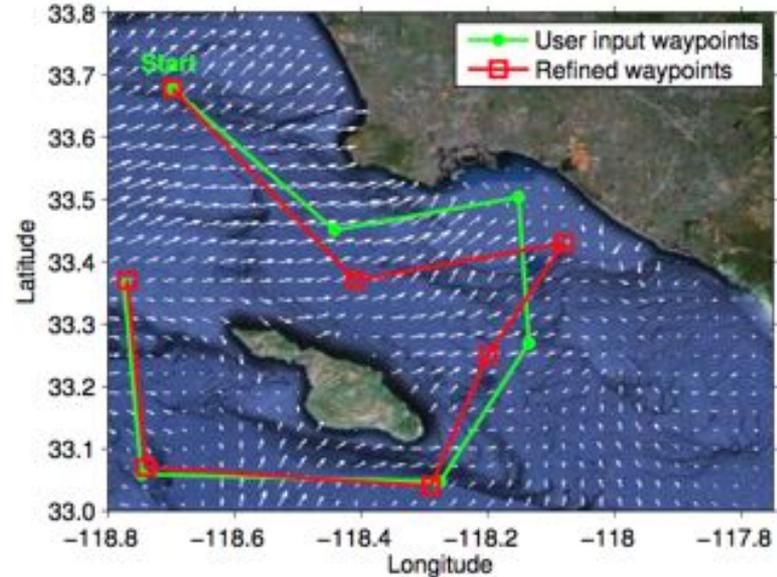
- **Output:** long-term cost for each state/action pair (assuming fixed # of waypoints)
- Optimal solution read from value function

Example trajectory

Risk of collision (shore/ships)

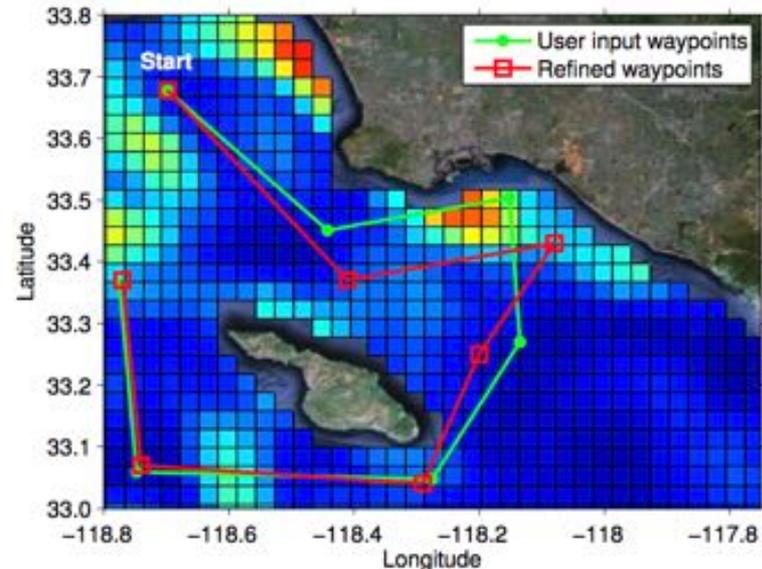


Ocean currents

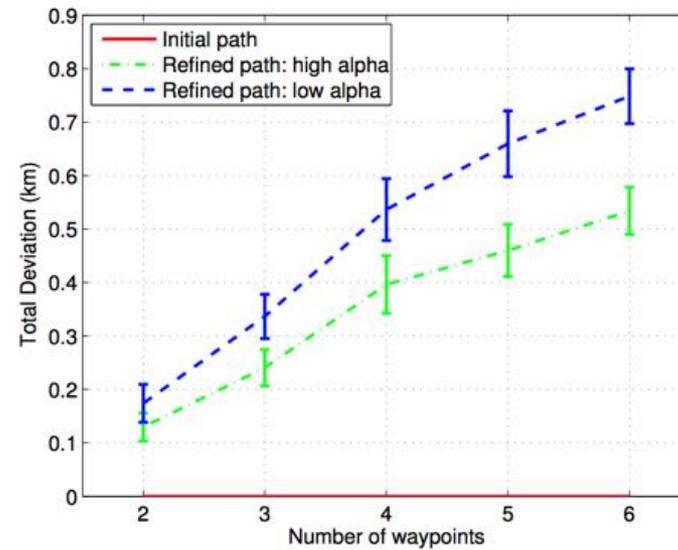
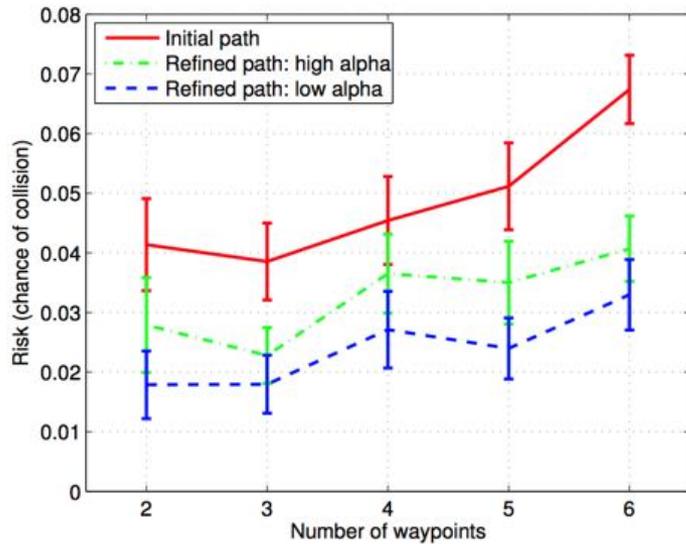


Uncertainty in ocean currents (interpolation variance)

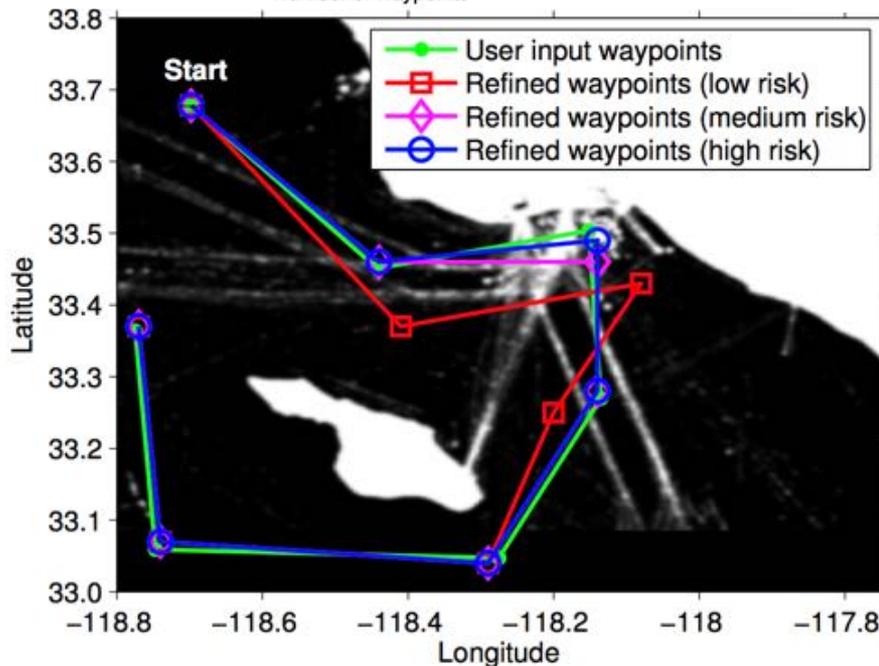
*Running time < 1 second



Incorporating user preference



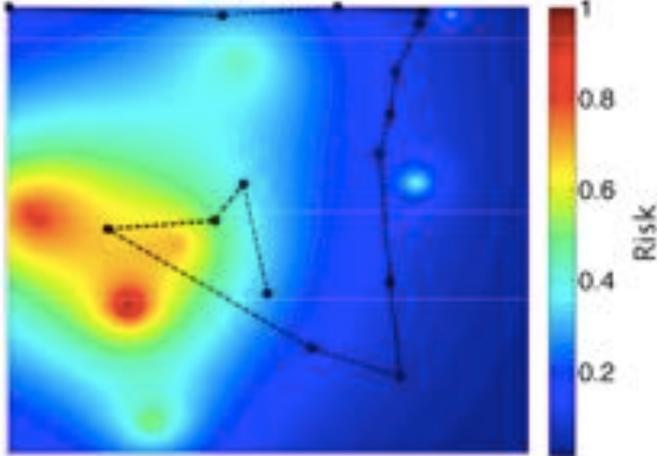
Each point averaged over 20 trajectories



Final path can be modified for low, medium or high risk based on weighting parameter

Balancing risk and reward

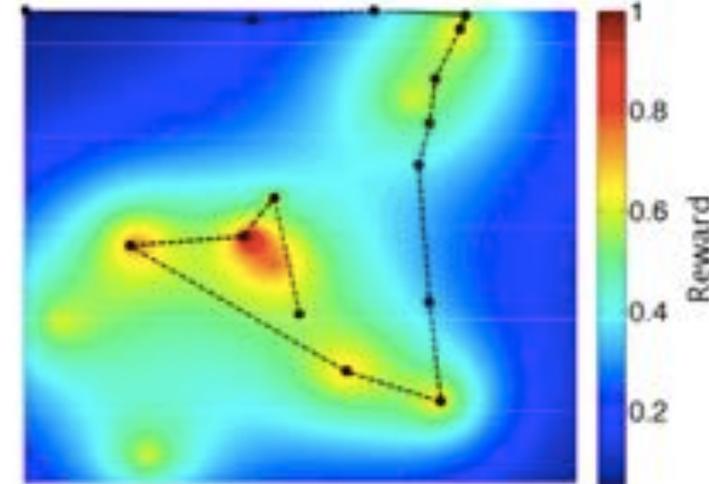
"Risk" Map



Risk of operation,
operators minimize
this implicitly

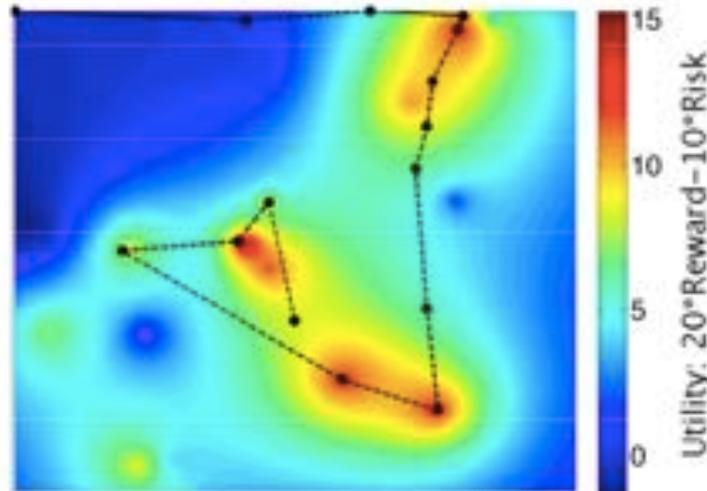
Can risk and
reward be
combined into
"utility" of
operation?

"Reward" Map



Reward of operation,
operators maximize
this implicitly

"Utility" Map



Trajectory optimization

Find optimized trajectory: ξ^*

Risk function: $R(\xi)$

Reward function: $I(\xi)$

- Optimize:

$$\xi^* = \operatorname{argmax}_{\xi \in \Phi} I(\xi) - \alpha R(\xi)$$

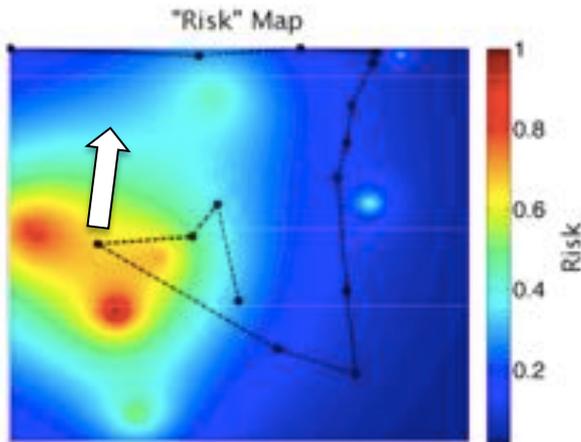


Weighting
parameter

(Somers and Hollinger, AURO '16)

Coactive learning

- **Key idea:** Operator modifies trajectory and algorithm learns human's preferences



Operator moves
point on trajectory
to reduce risk

Operator updates trajectory

$$\vec{\Delta}_t := \vec{\phi}(\langle x_t, y' \rangle) - \vec{\phi}(\langle x_t, y_t \rangle)$$

Algorithm updates weights of risk,
reward, and other factors

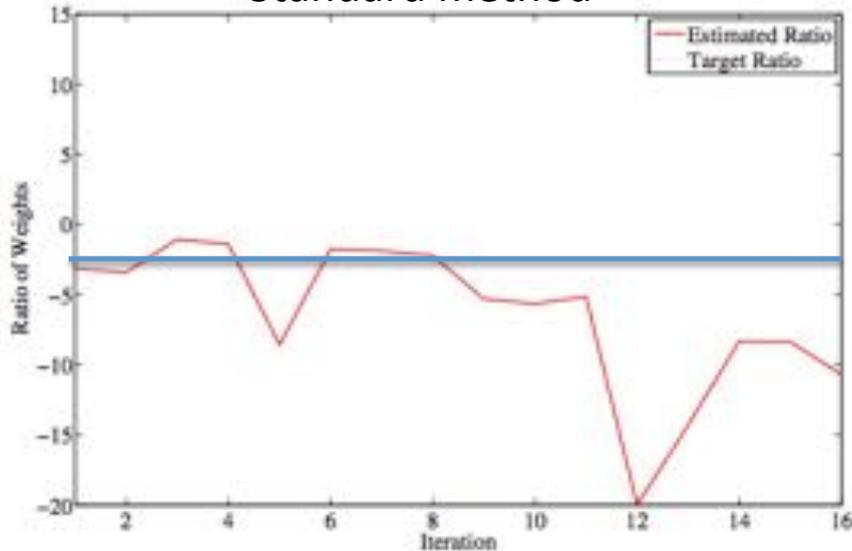
$$\vec{w}_{t+1}^\top = \vec{w}_t^\top + \lambda_t \vec{\Delta}_t$$



Modified
trajectory also
reduces reward

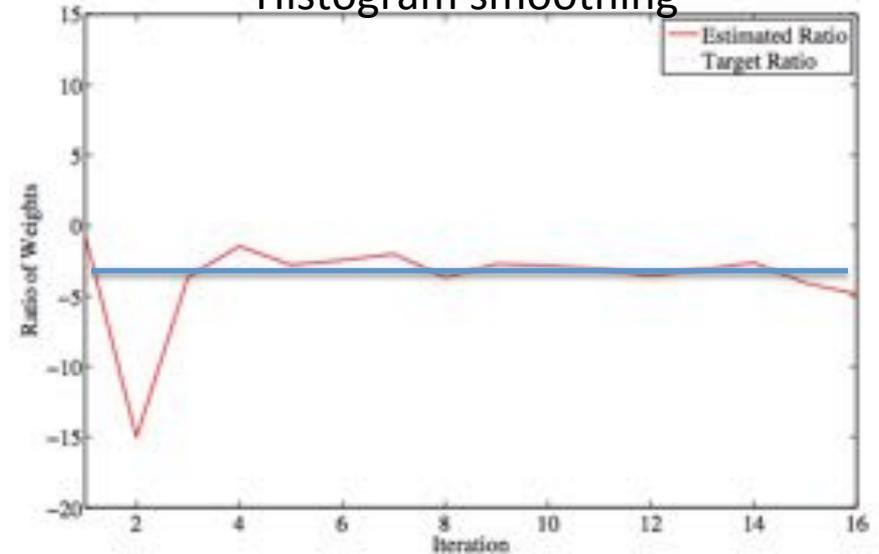
Can we learn the operator's preferences?

Standard method



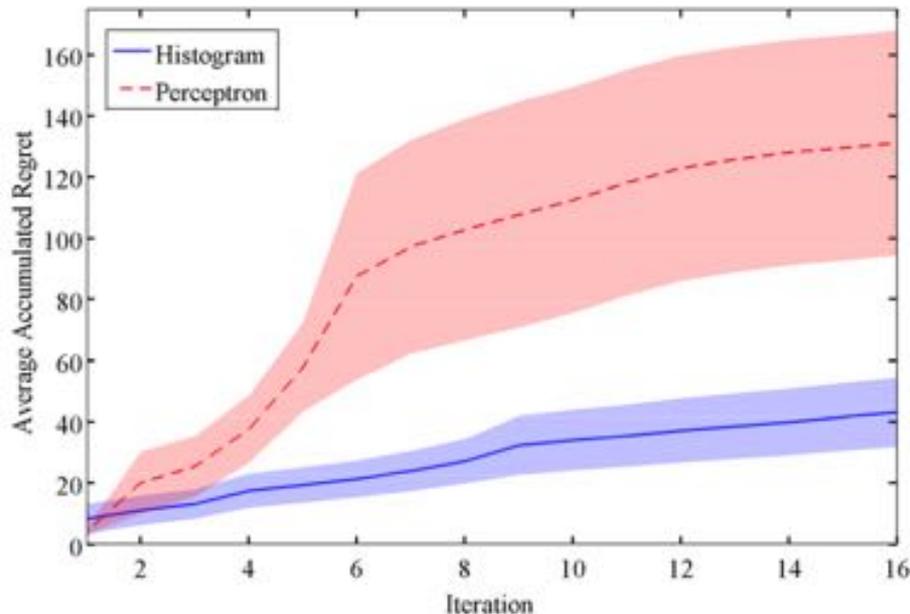
Example ratio of weights for standard
coactive learning method
(Goetschalchx et al., 2014)

Histogram smoothing



Example ratio of weights for histogram
smoothing method
(Somers and Hollinger, 2014)

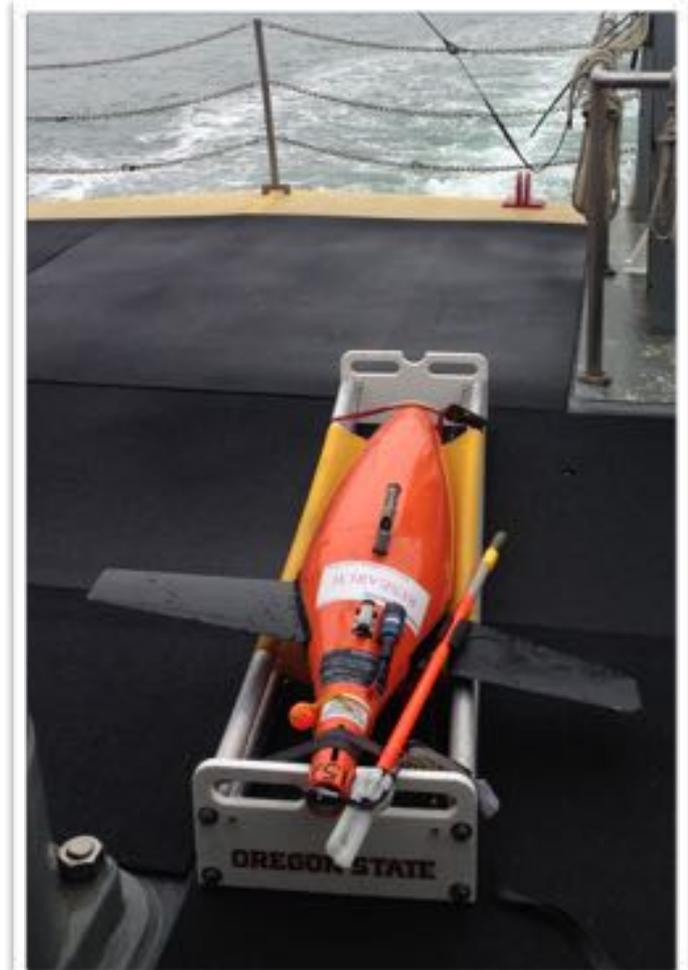
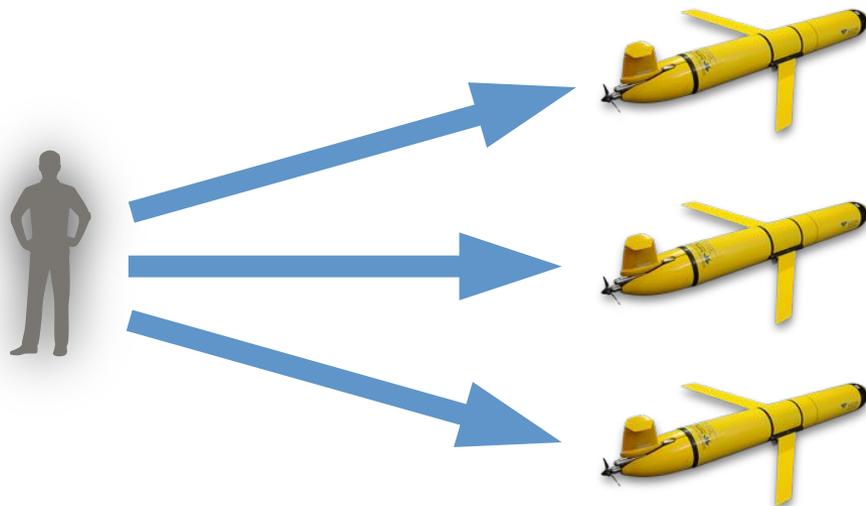
- Histogram smoothing method
 - Store all trajectory modifications made by operator in histogram
 - Infer weighting based on histogram



Regret =
unnecessary risk taken +
reward left on the table

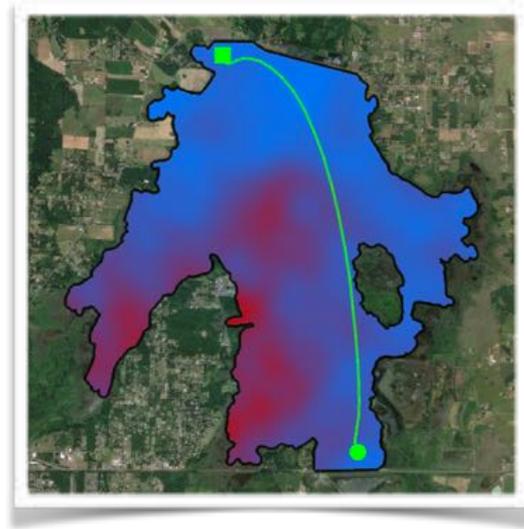
Efficient human-robot learning

- Learn operator's underlying mission utility function
- **Use fast survey-style rating and ranking questions**



Rating and Ranking Queries

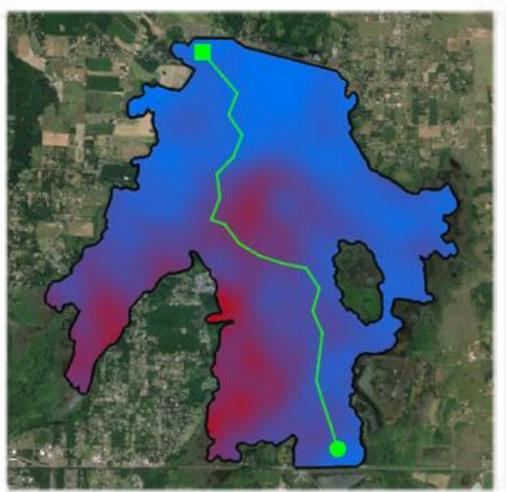
Ratings are specific and localized



Poor	0	0	0	0	0	Excellent
	1	2	3	4	5	

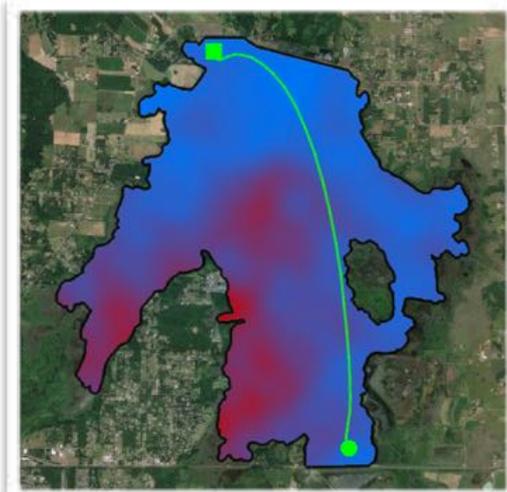
Rating and Ranking Queries

Rankings are non-specific and explore the space



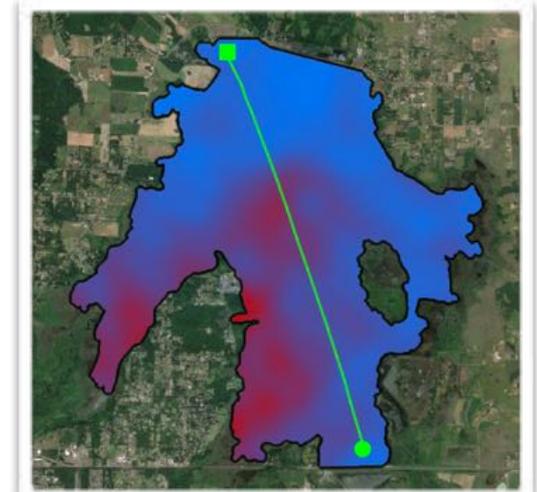
Best?

0



Best?

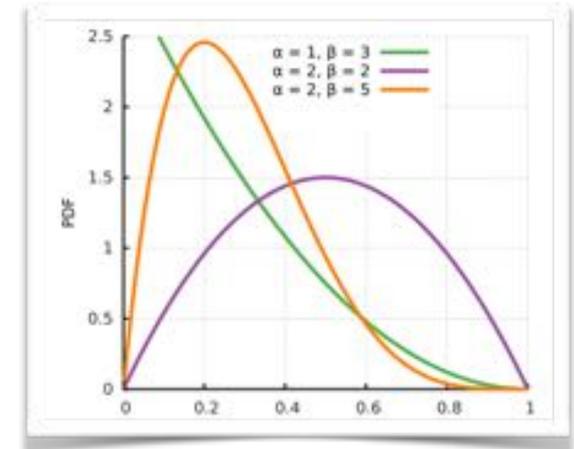
0



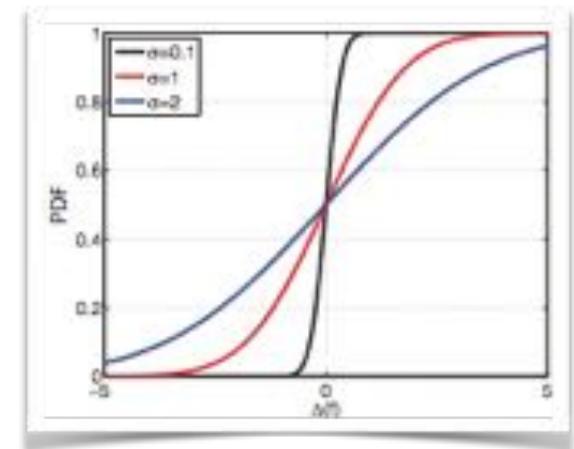
Best?

0

- Single GP models user's underlying value function using:
- **Rating Queries**
 - Bounded beta likelihood
 - Single point, specific value
- **Ranking Queries**
 - Similar to classification likelihood
 - General shape information

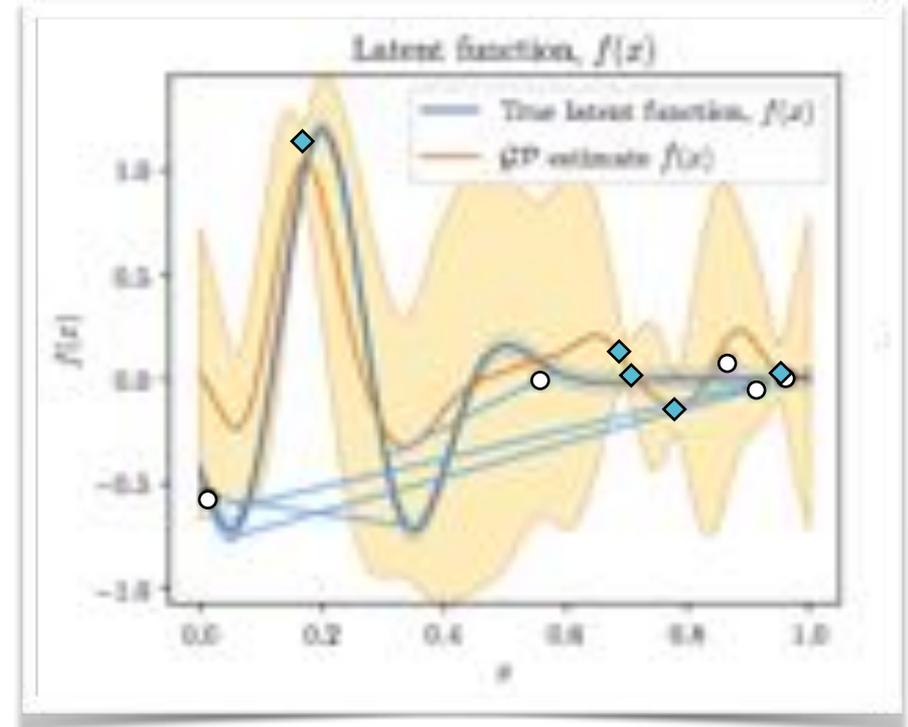
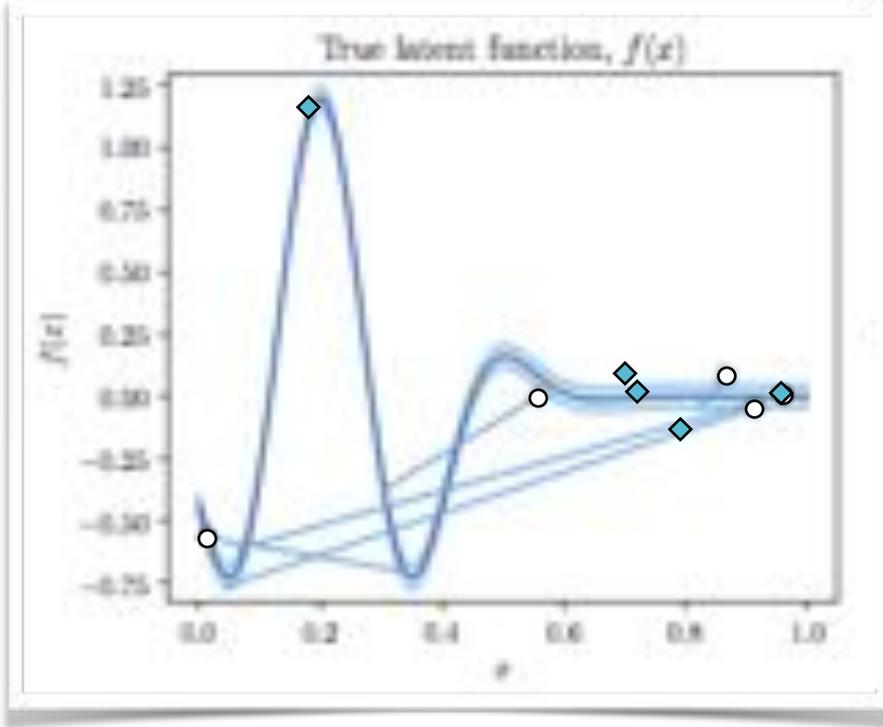


Beta Likelihood



Probit Likelihood

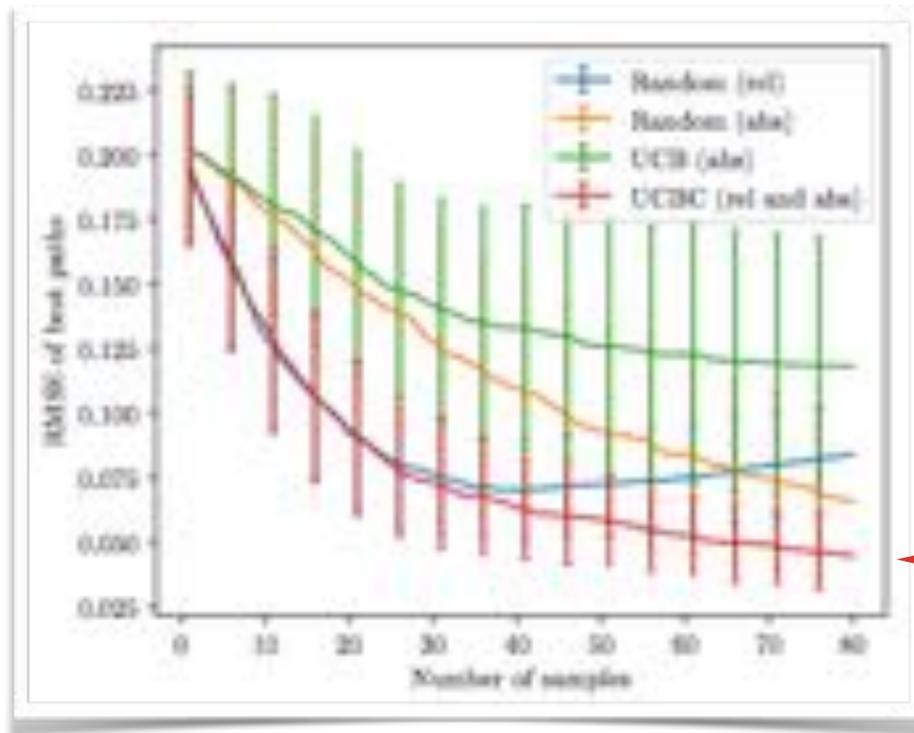
Function Approximation



◆ Rated Point —○ Ranked Points

Active Learning

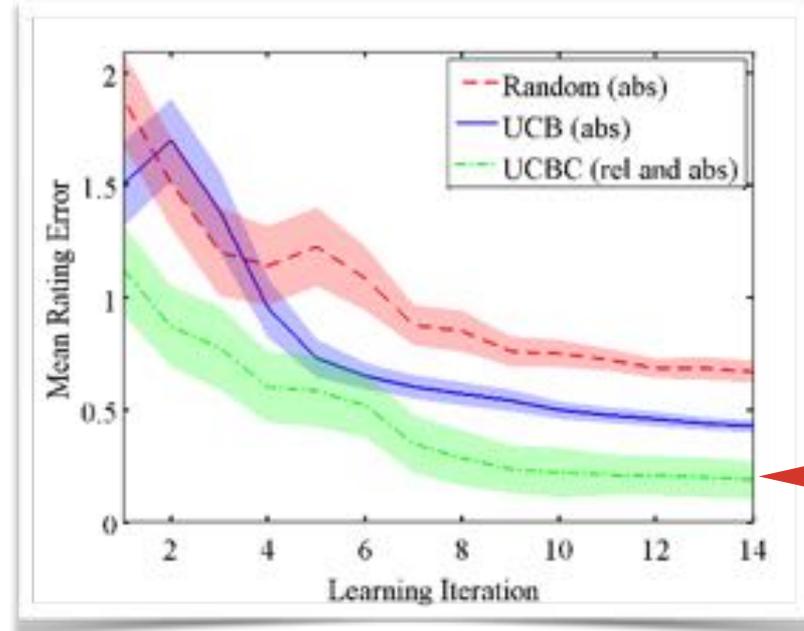
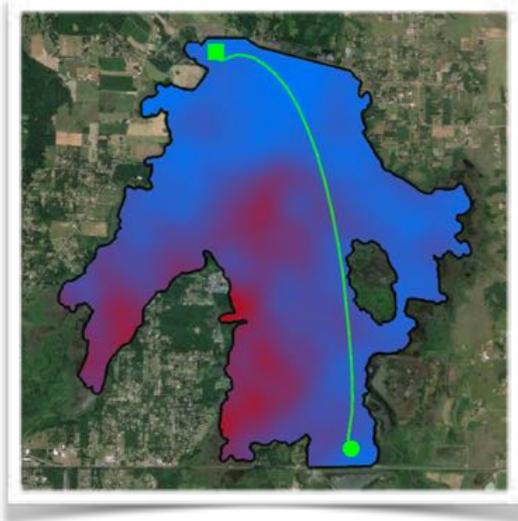
- Use upper confidence bound to select query trajectories
- Transition from ranking queries to rating queries



Upper Confidence Bound Combined (UCBC)

Simulated Lake Trials

- Simulated information gathering mission
- Non-linear objective tradeoffs
- Learns simulated user's preferences in <15 queries



Lake monitoring

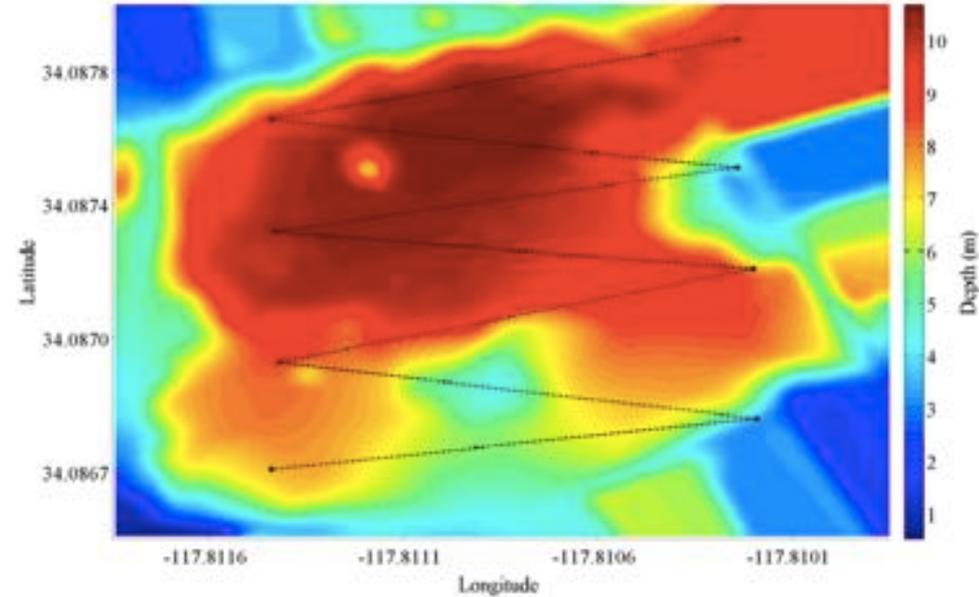


- Puddingstone Reservoir
 - **Goal:** track a given depth/temperature of interest
- Ecomapper Autonomous Underwater Vehicle
 - Propeller-driven up to 5 knots, dives down to 100 meters
 - Navigation: Doppler Velocity Log (DVL) and GPS

Puddingstone reservoir

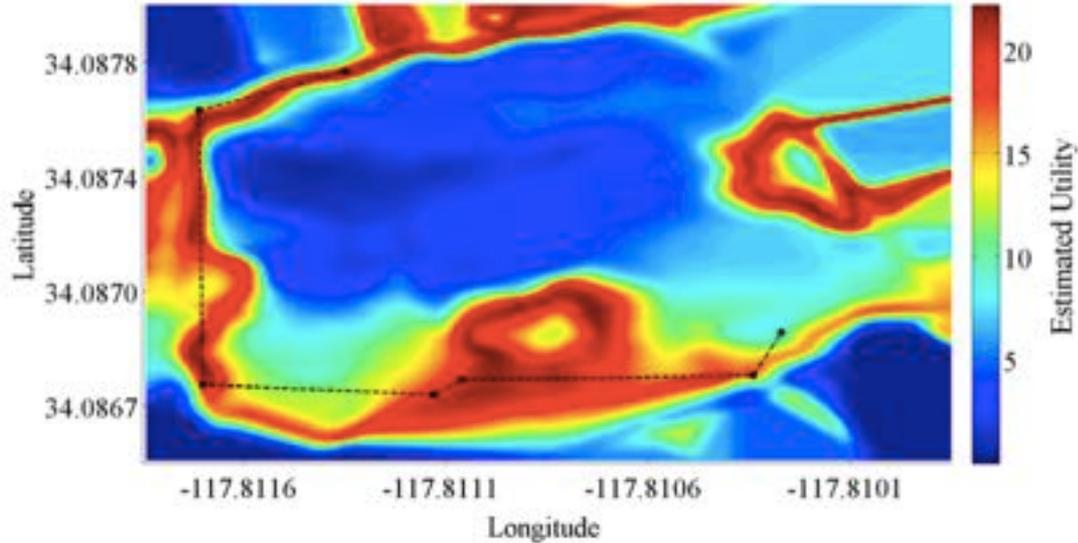


Satellite view of lake



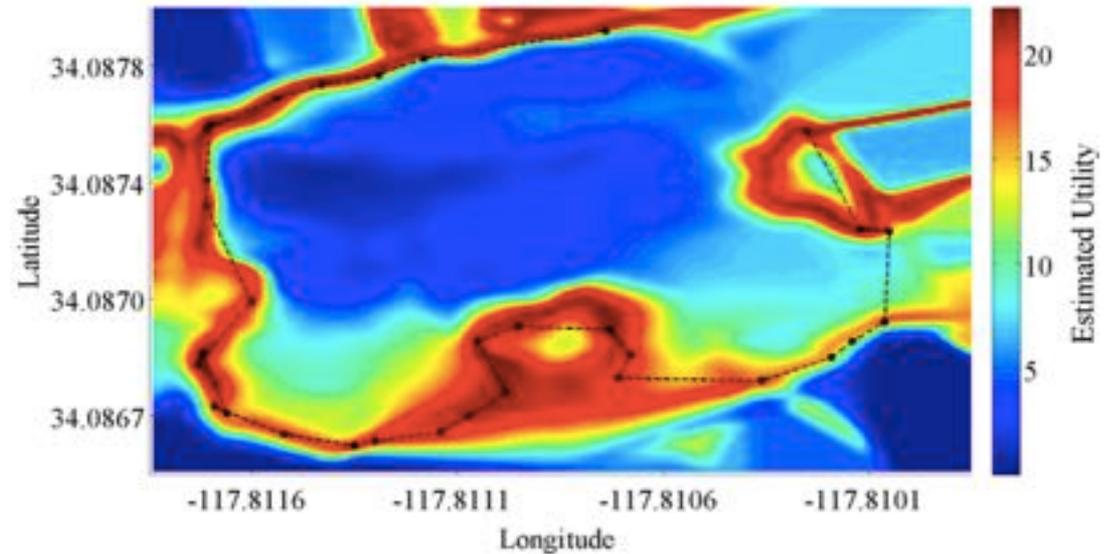
Interpolated depth from standard
lawnmower pattern

Monitoring target depth/temp



Human planned trajectory following a fixed depth and temperature

Computer planned trajectory following a fixed depth and temperature



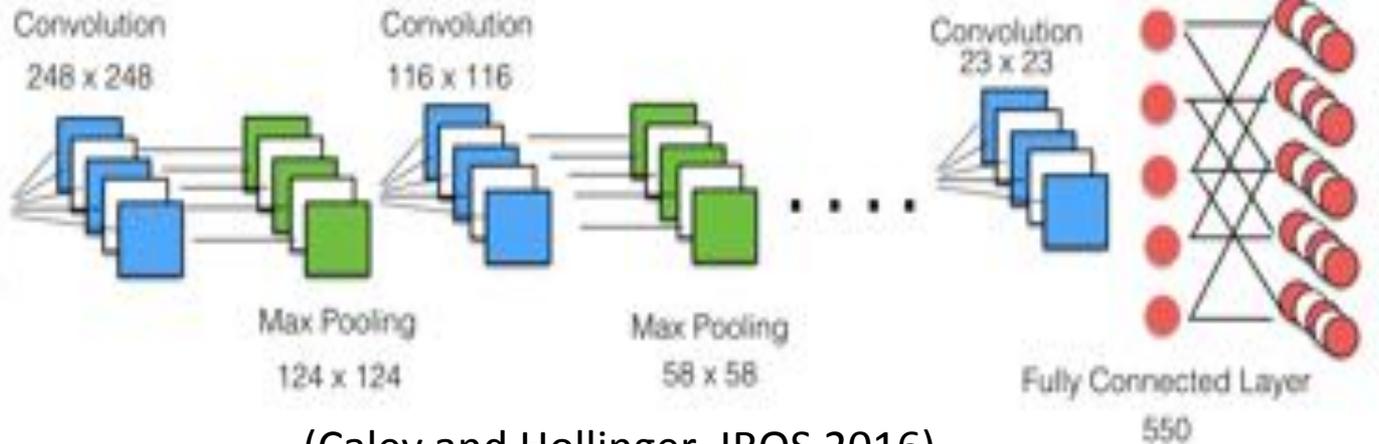
Ongoing work: deep learning for exploration

Input:

map of environment



Input
256 x 256



(Caley and Hollinger, IROS 2016)

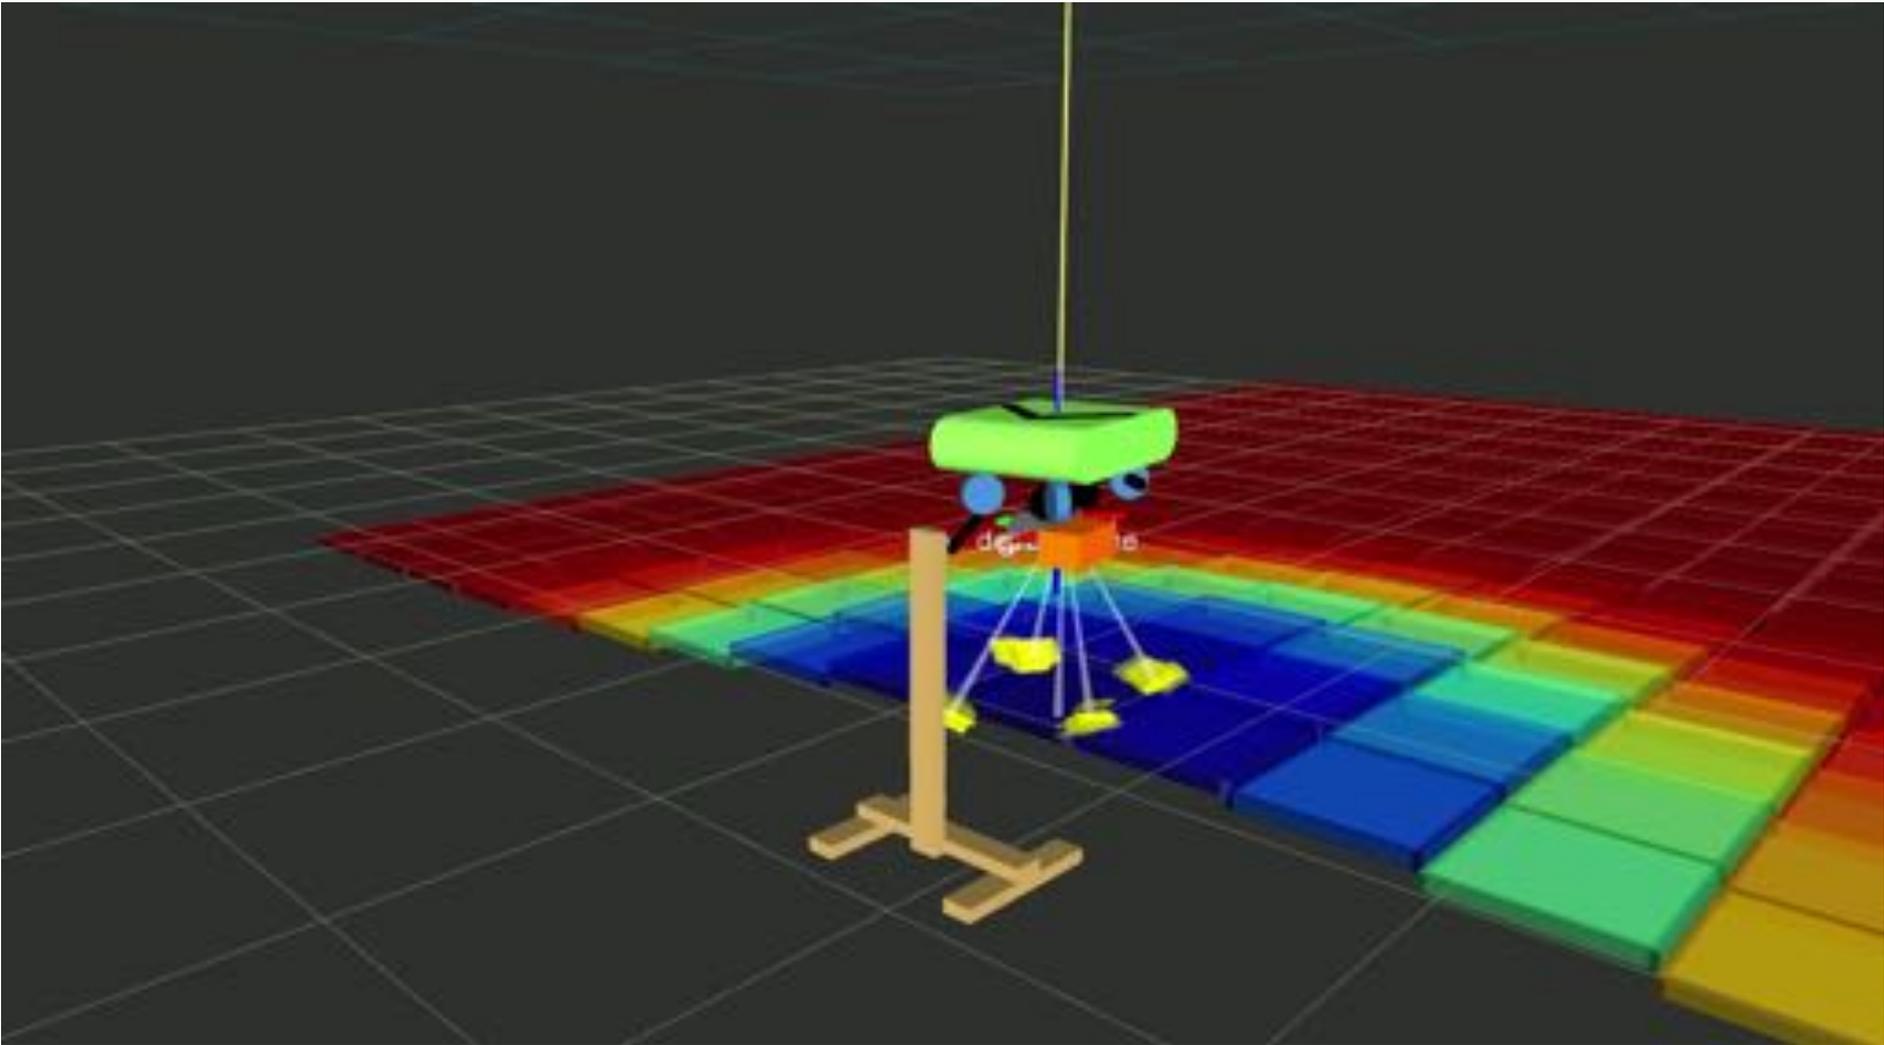
Output:

areas of interest



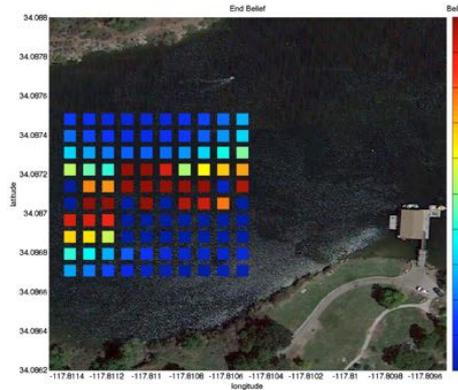
Output
20 x 20

Ongoing work: Semi-Autonomous ROVs

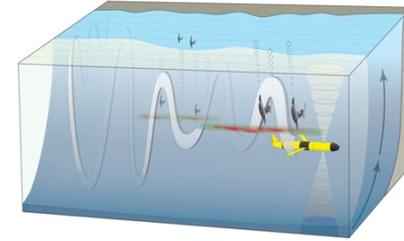


Other contributions

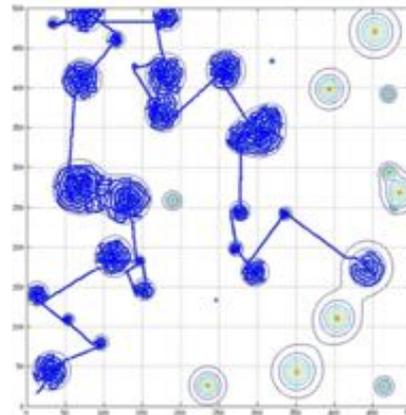
- Rapidly-exploring Adaptive Search and Classification (ReASC)
Hollinger, RSS 2015
IJRR 2014



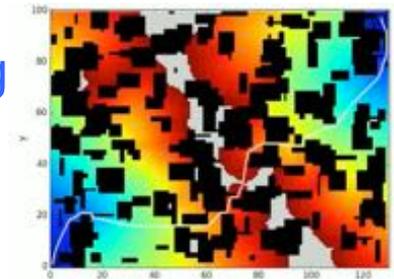
- Case study:
Ocean monitoring
Caley and Hollinger,
OCEANS, 2015
JFR, 2016



- Jump-Horizon (J-Horizon) efficient receding horizon algorithm
Yoo et al., FSR 2015



- Fast-Marching Explorer (FMExplorer) efficient exploration
Lawrance, Chung, Hollinger, RA-L, 2017

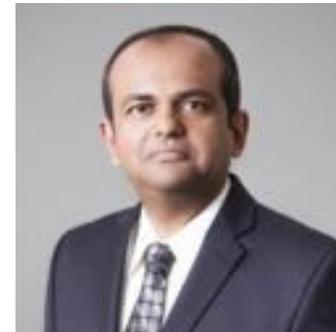


- PhD/MS programs (since 2014)
- 11 core faculty
- ~20 affiliated faculty
- ~70 grad students
- Graf hall collaborative space
- Legged, aerial, aquatic surface, underwater, and ground robots



Acknowledgements

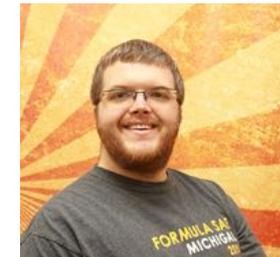
Geoffrey A. Hollinger
Robotic Decision Making Laboratory
geoff.hollinger@oregonstate.edu
<http://research.engr.oregonstate.edu/rdml/>



Gaurav Sukhatme
(Faculty Collaborator, USC)



Nicholas Lawrance
(Postdoc, now at ETH Zurich)



Thane Somers
(PhD student, OSU)

Funding: ONR N00014-14-1-0509,
NSF IIS-1317815, and
DoE DE-EE-0006816