Algorithmic Foundations for Autonomous Environmental Monitoring

Alan Papalia
Northeastern University & University of Michigan

Research Vision: Large-Scale Autonomous Ocean Observation

Major improvements in climate and weather modeling are essential to addressing the changing climate [1]. The ocean – which absorbs 91% of Earth's warming [2] – is particularly central to our climate, yet much of it remains poorly understood due to general difficulty in observing the ocean at scale [3]. This lack of observations, particularly in coastal, polar, and deep-ocean regions, hinders the development of accurate climate models and thus our ability to predict and mitigate the impacts of climate change. In addition to climate-relevance, observation is also critical for the growth of sustainable ocean industries [4], which are projected to reach $3 trillion by 2030 [5]. Altogether, there is mounting societal and economic pressure to improve our ocean observation capabilities.

Large-scale teams of autonomous underwater vehicles (AUVs) have the potential to revolutionize ocean observation and address current data bottlenecks. Such teams can seamlessly collaborate with platforms like buoys, divers, and satellites, enabling high-resolution, adaptive, and reliable ocean observation at unprecedented scales. However, progress is hindered by three key challenges: (i) the high cost of reliable underwater navigation, (ii) a lack of scalable tools for collaborative autonomy, and (iii) the dynamic nature of marine environments [6, 7]. Advancements in these areas are crucial to unlocking the full potential of autonomous ocean observation, allowing diverse teams of AUVs and other platforms to leverage complementary capabilities. *My research enables autonomous ocean observation by developing algorithms for navigation, collaboration, and adaptation in challenging remote environments*. While my research is motivated by ocean observation, the tools developed will impact a wide range of applications in field robotics and environmental monitoring.

Past Research: Collaborative Navigation for Low-Cost Underwater Vehicles

My past research focused on the ability for low-cost autonomous underwater vehicles (AUVs) to reliably navigate underwater. Low-cost navigation is a critical enabling capability for the deployment of AUVs in large-scale scientific endeavors. The current platforms used today are prohibitively expensive, as they rely on high-precision inertial navigation systems, typically costing $100k-$600k per vehicle [8]. These costs prohibit large-scale adoption for even well-funded scientific endeavors (e.g., the Ocean Observatories Initiative).

My work considered the ability to use collaborative navigation – a technique where multiple robots share information to improve each other's navigation estimates – to reduce the cost of navigation equipment. This focused on a team of "minimally instrumented" AUVs which can use (low-cost) imprecise inertial sensors to estimate odometry and simple acoustic equipment for both communication and to estimate the distance between each other (ranging). The underlying principle for this setup is that the whole can be greater than the sum of its parts, with joint information from these low-cost sensors providing a reliable globally consistent estimate of each AUV's position. While my research was inspired by challenges in underwater navigation, this problem is of interest to a wide range of robotics applications, including aerial [9], subterranean [10], and extra-terrestrial [11] navigation.

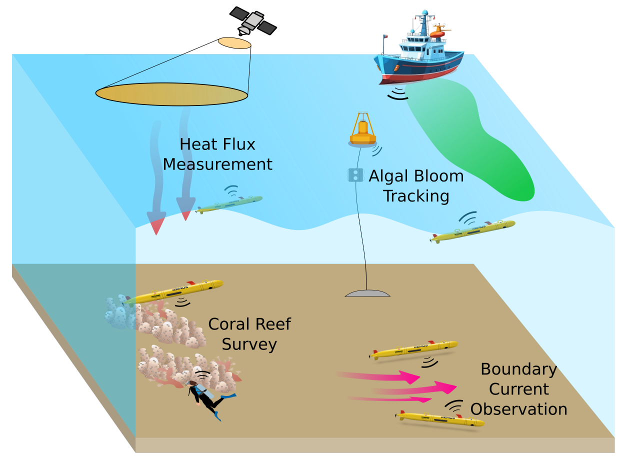

Fig. 1: This figure illustrates several envisioned autonomous ocean monitoring scenarios enabled by my research. My work moves toward future teams of sensing platforms that flexibly adapt and collaborate to observe diverse phenomena.

While this navigation approach has great potential, there remain significant challenges. Range measurement models are nonlinear and are often described as 'underconstrained' (i.e., a single measurement does not uniquely determine the relative position of two robots). Together, these properties manifest as non-convexity and potential ambiguity in the state estimation problem. To date, this has prevented reliable deployments of low-cost collaborative navigation systems.

Specifically, there are two fundamental challenges that appear: (i) the nonlinearity of range measurements means that the optimal state estimate is highly sensitive to the relative position that the measurements are taken from – often referred to as dilution of precision, and (ii) the non-convexity of the state estimation problem means that standard optimization techniques are susceptible to local minima and have no guarantees on the quality of the estimate. In the following sections I outline how I addressed these problems through my work in *planning for improved localization* and *certifiably correct estimation*.

Collaborative Planning to Aid Navigation. Because of the previously mentioned sensitivity of the optimal state estimate to the relative position that range measurements are taken from, the geometry of the relative trajectories of different robots significantly impacts the accuracy of the optimal state estimate. In the worst case, degenerate relative trajectories can lead to full loss of observability and rapid degradation of a robot's estimate of its location. I used tools from information theory (specifically, the *E-optimality* criterion) to connect the trajectories followed by a team of robots to the reliability of the corresponding state estimation problem. I combined these information theoretic properties with graph rigidity theory to develop a collaborative path planning algorithm to rigorously enforce that the collaborative state estimation problem is well-posed and the optimal state estimate is robust to measurement noise [12, 13]. This work established a tight coupling between the team's communication graph and underlying information theoretic properties to enable *lossless simplifications* of the planning task, allowing real-time planning.

This work was expanded in two Master's theses [14, 15]. Furthermore, I similarly leveraged graph and information theory in a collaboration on a convex optimization algorithm for graph pruning [16], a key capability for lifelong autonomy which allows agents to determine what information to 'remember' and 'forget'.

Certifiably Correct Estimation. While performing the work in the previous section I observed that even in best-case estimation problems localization estimates were often highly inaccurate. Even in ideal conditions, existing state estimation algorithms were insufficiently reliable for large-scale autonomous vehicle deployments. The root cause was reliance on local optimization methods like Levenberg-Marquardt [17], which are highly sensitive to initialization and lack any global guarantees on solution quality. These issues were especially problematic for collaborative AUV navigation. The symmetry of range measurements makes it difficult to determine good initializations [18, 19, 20], and incorrect navigation underwater often results in irrecoverable vehicles.

My first work to address these issues developed a novel convex relaxation for the underlying estimation problem [21], providing a principled, deterministic, and efficient way of initializing the state estimation problem. I advanced this line of research and produced a *certifiably correct* state estimation algorithm which is initialization independent and can provide globally optimal estimates with certificates of correctness [22]. This algorithm was the first certifiably correct algorithm to fuse fundamentally different measurement modalities (i.e., range measurements and relative rigid body measurements, such as from odometry). This leveraged semidefinite programming, Riemannian optimization, probabilistic modeling, and graph theory to obtain rigorous performance guarantees along with run times often *faster than state-of-the-art* local-search algorithms. From a theoretical perspective, this work empirically established a regime of this NP-Hard problem for which it can be solved in polynomial time.

The algorithm was validated on two marine surface vehicles using low-cost ($12k per vehicle) acoustic and inertial sensors, achieving high-accuracy localization (0.5% of distance traveled) comparable to $100k+ inertial systems. This work gained traction with researchers at other institutions, leading to a collaboration on a distributed extension [23].

Future Research Agenda

I will build on my past mathematical experience (e.g., graph theory and optimization) to develop *algorithmic tools for autonomous ocean monitoring*. I will pursue three mutually reinforcing directions: (i) low-cost navigation, (ii) heterogeneous teaming, and (iii) environmentally informed autonomy.

Low-Cost Navigation. I will continue my work to drive navigation technologies to cost levels that enable large-scale deployments. I will target two key challenges: outlier-robust navigation (as cheaper sensors increase outlier rates) and inertial navigation. I will explore outlier-robust problem formulations that have the same beneficial structure as my past work [22]. If successful, this would provide the performance guarantees and efficiency lacking in current approaches [24, 25]. For inertial navigation, I will combine learning with geometry and representation theory to construct data-driven inertial corrections as structured Lie group functions — preserving the underlying geometry and computational efficiency [26, 27].

Heterogeneous Teaming. Ocean observation requires diverse sensing modalities and vehicle capabilities, from AUVs and buoys to satellite sensing. Observation strategies must adapt to specific phenomena – e.g., tracking ocean thermal fronts may integrate AUVs with satellite data, while coral reef surveys may rely on AUV-diver collaboration. I will develop optimization-based frameworks to determine ideal team composition and task allocation, maximizing information gain while respecting system constraints. This builds on my past collaboration in diver-AUV teaming [28].

Environmentally Informed Autonomy. Marine environments significantly impact sensing and communication (e.g., sharp temperature gradients induce acoustic refraction), yet current autonomous systems fail to account for these effects, leading to failures in localization and data transmission. I will develop computationally efficient probabilistic environmental models [29] that capture how environmental conditions affect autonomy, enabling environment-aware planning and perception. These models will allow agents to (i) predict how the environment will affect their capabilities and (ii) plan actions to best sample the environment.

Summary. Independently, developments in each of these three areas will advance ocean monitoring as well as the fundamental science of robotics. However, the joint benefits of these advances will be multiplicative. Together, this progress can provide the scale and capability to enable large-scale autonomous ocean observation, unlocking new opportunities in ocean science and conservation while also pushing the frontiers of robotics research across a range of disciplines.

References

[1] Timothy M. Lenton, Johan Rockström, Owen Gaffney, Stefan Rahmstorf, Katherine Richardson, Will Steffen, and Hans Joachim Schellnhuber. Climate tipping points—too risky to bet against. *Nature*, 575(7784):592–595, 2019.

[2] Climate Change: Ocean Heat Content, September 2023. URL https://www.climate.gov/news-features/understanding-climate/climate-change-ocean-heat-content.

[3] NOAA Science Advisory Board. A report on priorities for weather research, 2021. URL https://sab.noaa.gov/wp-content/uploads/2021/12/PWR-Report_Final_12-9-21.pdf.

[4] National Oceanic and Atmospheric Administration. Accelerating growth in an information-based blue economy. URL https://cdn.ioos.noaa.gov/media/2023/05/Accelerating-Growth-In-An-Information-Based-Blue-Economy_1022r.pdf.

[5] OECD. The ocean economy in 2030, 2016. URL http://dx.doi.org/10.1787/9789264251724-en.

[6] Liam Paull, Sajad Saeedi, Mae Seto, and Howard Li. AUV navigation and localization: A review. volume 39, pages 131–149. IEEE, 2013.

[7] Fletcher Thompson and Damien Guihen. Review of mission planning for autonomous marine vehicle fleets. volume 36, pages 333–354. Wiley Online Library, 2019.

[8] Erin Fischell and Peter Traykovski. Doppler shift navigation system and method of using same, July 2022. US Patent App. 17/425,343.

[9] Kexin Guo, Zhirong Qiu, Wei Meng, Lihua Xie, and Rodney Teo. Ultra-wideband based cooperative relative localization algorithm and experiments for multiple unmanned aerial vehicles in GPS denied environments. *International Journal of Micro Air Vehicles*, 9(3):169–186, 2017. ISSN 17568307. doi: 10.1177/1756829317695564.

[10] Nobuhiro Funabiki, Benjamin Morrell, Jeremy Nash, and Ali Akbar Agha-Mohammadi. Range-aided pose-graph-based SLAM: Applications of deployable ranging beacons for unknown environment exploration. *IEEE Robotics and Automation Letters*, 6(1):48–55, 2021. ISSN 23773766. doi: 10.1109/LRA.2020.3026659.

[11] Elizabeth R Boroson, Robert Hewitt, Nora Ayanian, and Jean-pierre De Croix. Inter-robot range measurements in pose graph optimization. *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4806–4813, 2020.

[12] Alan Papalia and John Leonard. Network localization based planning for autonomous underwater vehicles with inter-vehicle ranging. In *IEEE/OES Autonomous Underwater Vehicles Symposium*. IEEE, 2020.

[13] Alan Papalia, Nicole Thumma, and John Leonard. Prioritized planning for cooperative range-only localization in multi-robot networks. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 10753–10759. IEEE, 2022.

[14] Nicole Thumma. Distributed potential field method for cooperative range-only localization in multi-robot networks. MEng. thesis, Massachusetts Institute of Technology, 2022.

[15] Megan C. Flynn. Exploring maneuvering strategies for heterogeneous cooperative navigation in underwater environments. MSc. thesis, Massachusetts Institute of Technology, 2023.

[16] Kevin J. Doherty, Alan Papalia, Yewei Huang, David M. Rosen, Brendan Englot, and John J. Leonard. Mac: Maximizing algebraic connectivity for graph sparsification. 2025. *Under Review*.

[17] Frank Dellaert and GTSAM Contributors. borglab/gtsam, May 2022. URL https://github.com/borglab/gtsam.

[18] George Kantor and Sanjiv Singh. Preliminary results in range-only localization and mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, volume 2, pages 1818–1823. Ieee, 2002.

[19] Joseph Djugash and Sanjiv Singh. A robust method of localization and mapping using only range. In *Experimental Robotics: The Eleventh International Symposium*, pages 341–351. Springer, 2009.

[20] Abhishek Goudar, Frederike Dümbgen, Timothy D Barfoot, and Angela P Schoellig. Optimal initialization strategies for range-only trajectory estimation. 2024.

[21] Alan Papalia, Joseph Morales, Kevin J. Doherty, David M. Rosen, and John J. Leonard. Score: A second-order conic initialization for range-aided slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 10637–10644, 2023. doi: 10.1109/ICRA48891.2023.10160787.

[22] Alan Papalia, Andrew Fishberg, Brendan W. O'Neill, Jonathan P. How, David M. Rosen, and John J. Leonard. Certifiably correct range-aided SLAM. *IEEE Transactions on Robotics*, 40:4265–4283, 2024. doi: 10.1109/TRO.2024.3454430.

[23] Alexander Thoms, Alan Papalia, Jared Velasquez, David M. Rosen, and Sriram Narasimhan. Distributed certifiably correct range-aided slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2025. *To Appear*.

[24] Heng Yang, Jingnan Shi, and Luca Carlone. TEASER: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 37(2):314–333, 2020.

[25] Heng Yang and Luca Carlone. Certifiably optimal outlier-robust geometric perception: Semidefinite relaxations and scalable global optimization. *IEEE Trans. Pattern Anal. Machine Intell.*, 45(3):2816–2834, 2022.

[26] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

[27] Boce Hu, Xupeng Zhu, Dian Wang, Zihao Dong, Haojie Huang, Chenghao Wang, Robin Walters, and Robert Platt. OrbitGrasp: $se(3)$-equivariant grasp learning. In *8th Annual Conference on Robot Learning*, 2024. URL

https://openreview.net/forum?id=clqzoCrulY.

[28] Brendan W. O'Neill, Jesse R. Pelletier, Samuel Calvert, Alan Papalia, John J. Leonard, Lee Freitag, and Eric Gallimore. Loosely-coupled human-robot teams for enhanced undersea operations. In *Proc. of the IEEE/MTS OCEANS Conf. and Exhibition*. IEEE, 2022.

[29] Victoria L Preston, Genevieve E Flaspohler, John W Fisher, Anna PM Michel, and Nicholas Roy. PHORTEX: Physically-informed operational robotic trajectories for scientific expeditions. *IEEE Transactions on Field Robotics*, 2024.