Succinctly defines research vision

My research vision is to *develop planning and task allocation algorithms for cooperative multi-agent marine robotics*.

Motivates the problems they want to address, explains the context of the industry

Autonomous systems have the ability to greatly enhance the capabilities of both crewed and uncrewed vehicles, enabling safer, more sustainable, and efficient maritime operations. Single robots have limited capacity and resources, which can constrain their ability to perform prolonged or resources intensive tasks. Through multiple autonomous underwater vehicles (AUV), robots can safely perform complex operations that are too challenging or dangerous for humans like maritime research, environmental monitoring, offshore industries, and search and rescue operations. However, the underwater environment creates unique planning, navigation and communication challenges for AUVs due to limited visibility, limited bandwidth and signal attenuation, and lack of GPS.

Clear branding on research area

As a professor, my research will focus on the development of learning algorithms and task allocation techniques to enable complex multi-agent missions. My focus will be on development of learning-based control techniques to improve performance of current processes ranging from path planning to coordinated decision making.

Intro

Uses numbering to clearly identify research areas

In pursuit of this, my research has two themes: (1) creating cooperative multi-agent marine robots that are robust to communication losses and (2) developing task allocation algorithms to improve teamwork in complex tasks. Given my research and educational background in autonomy and multi-agent systems, I am uniquely positioned to pursue this work and address this need in marine robotics.

In my prior research, I created an algorithm that leverages information from nearby agents to understand how their behavior would change with better knowledge of the environment. For this work, I used space traffic management as a representative problem, but the techniques are extensible to other nonlinear environments with restrictive communications, like the ocean. In space, each satellite possesses the best knowledge of its physical location and surroundings. If this knowledge was shared, it has the potential to improve the overall efficiency of space operations by enabling collaboration amongst satellites and independent verification of potential collisions. My approach seeks to quantify the value of this information by relying on graph theory to model communication between satellites, and then neural networks to provide an optimized transformation of the significance of each piece of information.

Explains the significance of prior work and how it will inform the field

This approach allows satellites to build an enriched understanding of their local environment, **leading to lower collision rates and improvements in path planning** [1]. The implication of this work is by understanding the impact of improved knowledge on satellite behavior, we can develop techniques that allow satellites to intelligently utilize local information in the planning process.

With my approach, the space environment is converted into a graph, where satellites represent the nodes, and communication pathways represent the edges. When an edge is formed, satellites are communicating information about their exact location or planned future maneuvers. Importantly, this graph structure is dynamic, meaning that edges are changing as the environment changes, thus allowing satellites to form adaptive connections with nearby neighbors. This graph formulation allows us to use graph neural networks (GNNs), a specialized type of neural network that processes graph inputs and learns encodings of key information for each node in the graph. In the case of my formulation, this means that we can get a compact representation of the knowledge a satellite possesses about its environment through information sharing. One key desirable feature of GNNs is their ability to message-hop, meaning that messages can be passed through an intermediary node, allowing nodes that aren't connected to obtain each other's information. This message passing allows satellites to develop an awareness of the location and intentions of other satellites in proximity. The GNN is used in combination with a multi-agent reinforcement learning algorithm, which uses the compact node representations to learn the best course of action to successfully complete tasks.

Bolds key outcomes of previous work and cites own papers

My prior work has shown that this approach **outperforms decentralized standard multi-agent reinforcement learning algorithms**, and can achieve similar performance to centralized methods that have access to global information [2].

My current work focuses on the development of multi-telescope satellite sensor tasking mechanisms for efficient resource allocation. Given that the number of space objects is projected to greatly outpace any increased capacity in sensing capabilities, our ability to gather detailed information on the close encounters will highly depend on the efficient use of these capabilities. I have led an MIT M.Eng. student in the development of a tool that creates a list of viewable satellites given a coordinate location. The output of this tool is being used for training integer programming satellite sensor tasking algorithms to identify synergies in different telescope locations.

Uses bolded headings to separate different sections and improve readability

Future Research

Multi-agent systems require additional work towards the theoretical development of algorithms that enable groups of robotic vehicles in challenging underwater environments. To address this need, my research will follow two streams. The first research stream focuses on creating cooperative multi-agent marine robots through multi-agent reinforcement learning. The second stream focuses on the development of control strategies robust to communication constraints.

This research statement highlights research projects that they would immediately pursue as faculty

I outline below two research questions I will investigate in the next 3-5 years, each in line with the two streams listed above.

1. Graph Neural Networks for Latent Communication: Underwater communications rely on the use of acoustic devices that are suffer from intermittent communication losses, multi-path effects,

reduced bandwidth and low reliability. Through the use of dynamic graphs via graph neural networks, we can capture time-varying communication topologies and water-imposed communication constraints. One promising direction of research I plan to pursue is investigating the use of graph neural networks to model dynamic communication topologies between marine robots. Graph neural networks are able to succinctly learn encodings of graph structures to identify and prioritize different messages passed between nodes on a graph. The outputs of graph neural networks are easily interpretable, as the network seeks to inherently capture relationships and dependencies between interconnected points on a graph, and how information propagates through the network. I plan to develop a graph neural network that is able to learn the dependencies and importance of different messages between marine robots to formalize an adaptive communication network. By encoding the latencies as characteristics within edges on the graph, the graph neural network can learn the prioritization of various time-sensitive parameters. If successful, follow-on work would evaluate the effectiveness of the graph neural network in real-time deployment by incorporate hardware testing and underwater deployment.

2. Efficient Task Allocation for Heterogenous Robots: As multi-agent marine robotics will likely involve groups of heterogenous agents carrying complementary payloads, task allocation will be vital for assigning specific responsibilities to different agents as they complete complex tasks. Leveraging my past experience in telescope task allocation, I am interested in exploring problems related to multi-agent task allocation under limited energy resources using through integer programming and machine learning based techniques. Learning based methods and integer programming are complementary to one another; integer programming provides a rigorous mathematical formulation of the discrete trade-offs between different decision variables, and machine learning allows us to obtain approximations of complex multi-variate data.

Research Collaborations and Funding

More broadly over the course of my career, I would like to establish myself as an expert within the field of multi-agent planning and coordination for marine systems. I believe that this will require an interdisciplinary approach with adjacent fields such as robotics and ocean sciences.

Highlights specific collaborations at the university

As a professor, I am eager to seek collaborations with partnerships with relevant groups like the Graham Sustainability Institute and the University of Michigan Robotics department.

My future research program seeks to advance the current state of the field of marine robotics by improving the state of the art in planning for multi-agent systems. This work is well aligned with existing work in autonomy within the department University of Michigan. I have been fortunate to work with undergraduate, graduate, and faculty researchers from institutions that are part of my educational background (Purdue University, MIT) and others such as University of Washington. As a graduate student, my research collaborations were conducted and published with funding support from my NSF Graduate Research Fellowship.

Directly names funding organizations that they will seek funding from

I will also seek government sources of funding such as NSF, NAVSEA, and NEEC, and FFRDCs (MIT Lincoln Laboratory, The Aerospace Corporation).

References

Directly references prior work, cites it so readers can refer to this work if interested

[1] **Dolan, S.,** Nayak, S., & Balakrishan, H. (2023). Satellite Navigation and Coordination with Limited Information Sharing. *Proceedings of Machine Learning Research*, 211, 1058-1071. https://proceedings.mlr.press/v211/dolan23a/dolan23a.pdf

[2] Nayak, S., Choi, K., Ding, W., **Dolan, S.,** Gopalakrishnan, K., & Balakrishnan, H. (2023). Scalable Multi-Agent Reinforcement Learning through Intelligent Information Aggregation. *Proceedings of the 40th International Conference on Machine Learning* in *Proceedings of Machine Learning Research*. 202:25817-25833 Available from https://proceedings.mlr.press/v202/nayak23a.html