James D. Motes - Research Statement

Overview: My research is in intelligent multi-robot systems. In both large-scale industrial settings (e.g., manufacturing and logistics) and more intimate settings (healthcare and in-home assistive robotics), we need human-centered robotic systems. While the majority of my work has contributed to **theoretical frameworks for multi-robot planning**, my recent and forthcoming work extends to **making robotic solutions directable by human intelligence**.

Robotic intelligence and the ability to direct it have always developed in tandem. Early robotic systems consisted of little more than motors, so we (humans) developed joysticks to directly control those motors. As robotic platforms increased in intelligence, we developed tablet GUIs to program repeatable action sequences. Current research increasingly uses natural language interfaces, aligned with the VLMs/LLMs that power modern task planning systems. Building on my experience developing theoretical frameworks for multi-robot planning, my research is organized into three thrusts that explore this co-evolution of robotic intelligence and human interaction, each driven by a guiding question:

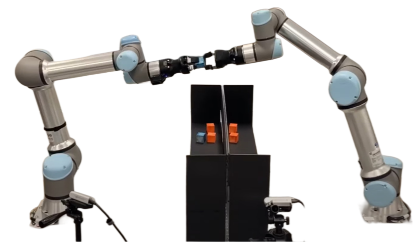

Fig. 1: Robotic manipulator arms performing an object handoff in a shelf stocking scenario [1].

1. **Multi-Robot Planning:** How do we exploit growing general intelligence within the complexity of multi-robot problem spaces?
2. **Human-Robot Co-Planning:** How do we design cooperative planning systems in which humans, robots, and AI services jointly create and refine plans?
3. **Hardware and Software Acceleration:** How do we design planning architectures that fully exploit modern hardware and software acceleration?

I currently lead research teams in each thrust, **mentoring 20+ graduate students** and dozens of undergraduates along the way, **resulting in 14 publications** to date and support from both government (NSF) and industry (Foxconn Interconnect Technologies, IBM, Toyota). I recently served as **primary author on an NSF Foundational Research in Robotics (FRR) proposal** along with several industry partnership proposals (Amazon, Samsung).

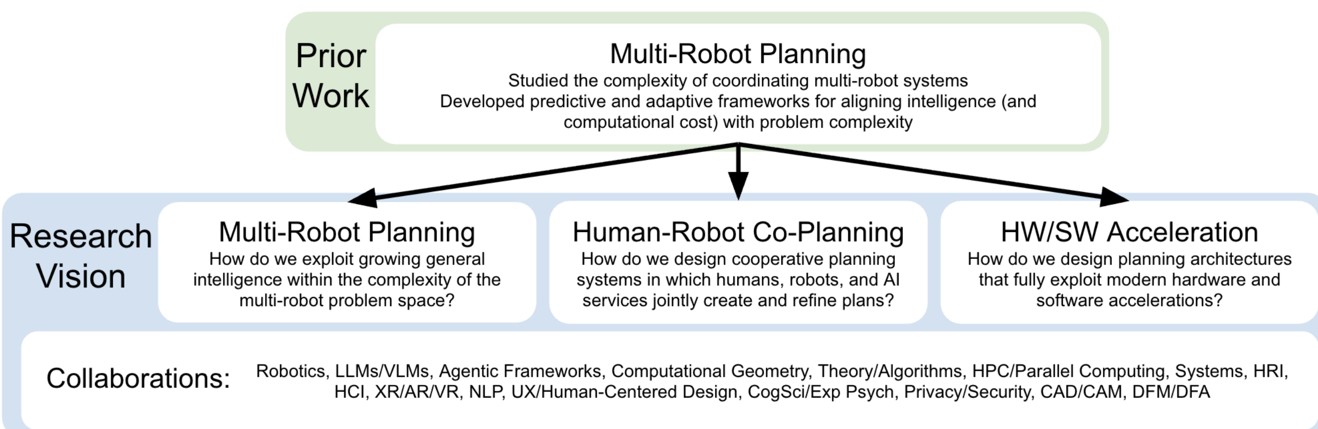

Fig. 2: Schematic of my research: Extending from my prior work into the three thrusts and listing potential collaborations.

Alongside my academic work, I have led DOE SBIR-funded software development for a startup (Optigon Inc.) delivering automated metrology systems, and founded an autonomous manufacturing startup in Houston. Interviewing local manufacturers across sectors made it clear that both current autonomous systems and, critically, the interfaces for instructing them are not yet adequate for large-scale deployment despite strong demand. These experiences, combined with my mentoring-intensive role as a senior student and postdoc in the Parasol Lab at UIUC, have sharpened my research questions. Now, **I aim to develop multi-robot and human-robot co-planning frameworks and interfaces that are theoretically sound, responsive on modern hardware, and grounded in real-world deployment constraints.**

Thrust 1 - Multi-Robot Planning: How do we exploit growing general intelligence within the complexity of multi-robot problem spaces?

The main difficulty of multi-robot planning is exponential growth in the state space as the number of robots or task complexity increases. The simplest version of multi-robot planning, multi-agent pathfinding (MAPF) where agents must find conflict-free paths on a shared graph (typically a grid), is intractable in its general form. Extending to multi-robot motion planning (MRMP) for high degree-of-freedom robots (e.g., robotic arms), where planning a path for a single robot is PSPACE-complete on its own, creates a planning problem of considerable complexity. This is further exacerbated in multi-robot task and motion planning (MR-TAMP) as each action considered by a planning algorithm corresponds to an additional MRMP problem.

Historically, methods either decouple the space, considering robots one at a time, sacrificing completeness and coordination for speed and scalability, or accept the size of the space to meet coordination requirements, providing slower approaches for smaller problems. Instead, my team exploited the insight (illustrated in Fig. 3) that *coordination levels vary during multi-robot problems.* This led to our guiding principle: ***Match the algorithmic intelligence (and computational cost) to the local (sub)problem complexity.*** In other words, de- and re-compose problems into subproblems characterized by consistent internal complexity, so that the cheapest method with sufficient intelligence can be used locally to solve each subproblem. This improves utilization of computational resources on large problems and makes the core challenge of multi-robot planning (sub)problem decomposition and local algorithm matching.

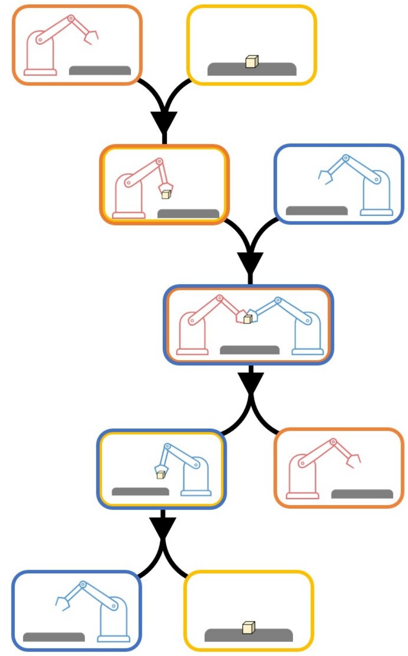

Fig. 3: In a two-robot pick-handoff-place sequence (from Fig. 1), there are varying levels of coordination. Before picking the object, both the blue and red robots can be planned for independently. After performing the pick action, the red robot and the object must be considered jointly. The moment of handoff is the only instance where all elements of the problem need to be coupled. Afterward, the problem can again be decomposed into independent elements. DaSH [2] models these changes in composition explicitly with a *directed hypergraph*.

Prior Results (Theoretical Foundations): My team developed two novel multi-robot planning frameworks, **D**ecomposable St**a**te **S**pace **H**ypergraphs (**DaSH**, *IEEE Transactions on Robotics* [2]) and **A**daptive **R**obot **C**oordination (**ARC**, *IEEE Robotics and Automation Letters* [3]). DaSH [2] is a predictive framework which leverages problem structure to predict where coordination is needed and proactively composes subproblems to isolate the use of more expensive planning techniques. The original method [2] demonstrated **three orders of magnitude speedup in planning times** over state-of-the-art baselines. The most **recent version [1] reaching four orders of magnitude speedup** while planning for up to **16 robots and 60 objects** in object rearrangement problems. In addition to warehouse and factory settings, these results have been demonstrated on biofab labs [1] and mining [4] with the latter reaching **128 robots** in congested MRMP settings.

ARC [3] is a reactive framework which dynamically creates subproblems during planning as coordination requirements are discovered, and the framework dynamically adapts the complexity (and cost) of the planning techniques employed to match the local features. As a result, ARC was the only method that **achieved a 100% success rate for problem sets with large variance in the amount of coordination required**. We have extended the base framework to account for kinodynamic constraints [5] and exploit data from prior solutions for faster planning [6] and have seen additional extensions by other research groups to tackle problems like congested MRMP.

Through my multi-robot research, I have mentored 5 Ph.D. students, 1 M.S. student, and several undergrads. 2 of the Ph.D. students and the M.S. student have graduated with their dissertations and thesis built from this work.

Future Directions (Algorithmic Foundations, VLM/LLM Reasoning, Human-Robot Teams): As a faculty member, I will turn the "match intelligence to local complexity" principle into a general toolbox for multi-robot planning. This includes extending the theoretical foundations to solve new problem classes while maintaining or strengthening theoretical guarantees and creating open source libraries that other labs and industry partners can adopt.

I am also interested in exploring how modern VLM/LLM-based agents can be leveraged in multi-robot planning. The prohibitively large state spaces arising from large numbers of robots and task length and complexity are similar to the challenges of context window limitations in LLM-based multi-agent architectures. The guiding principle and frameworks from my prior work can provide insights into how to efficiently apply VLM/LLM techniques in multi-robot planning, and may also yield further insights for managing context windows within LLM-based multi-agent systems.

I will naturally seek collaborations with colleagues in both academia and industry working on scalable LLM-based multi-agent solutions to better understand and study this overlapping problem.

Finally, I am interested in extending these frameworks to explicitly model humans as agents with their own capabilities, constraints, and preferences. This sets the stage for Thrust 2, studying human-robot interaction (HRI) and cooperative planning, adapting different HRI techniques to match local subproblem coordination requirements.

Potential Funding Sources: NSF Foundational Research in Robotics (FRR), NSF Cyber-Physical Systems (CPS), NSF CISE Robust Intelligence (RI), NSF Advanced Manufacturing programs in CMMI, Amazon Research Awards (ARA), Advanced Robotics for Manufacturing (ARM) Institute

Thrust 2 - Human-Robot Co-Planning: How do we design cooperative planning systems in which humans, robots, and AI services jointly create and refine plans?

As robotic intelligence improves, we are no longer limited to robots that simply execute fixed, precomputed plans. For the first time, it is becoming feasible to deploy teams of capable robots into real homes, labs, and factories, and to ask them to work with people instead of just near them. This presents a new challenge of not only designing interfaces to match the current level of abstraction accessible to robotic intelligence, but of designing systems in which humans, robots, and AI services can **co-plan** over shared representations of tasks, environments, and constraints, matching the level of abstraction with the different elements or participants in the system.

Prior and Ongoing Work (Natural language, AR/VR): I currently lead a graduate research team studying the use of natural language (NL) to parameterize motion planning algorithms. This approach treats NL as a way for (non-expert) users to shape the underlying planning objectives and constraints, not just issue goals, producing customized behaviors in homes and workplaces. As part of this, I was the **primary author on an NSF FRR proposal** in collaboration with natural language processing researchers at UIUC.

Additionally, I have collaborated with Insper and other UIUC researchers to **develop AR/VR interfaces for interacting with motion planning algorithms** [7]. This system provides immediate visual feedback on how changes in the task or environment affect robot behavior.

Finally, as Principal Software Engineer at Optigon, I develop the software and user interfaces that expose control of an automated metrology system and its AI-powered analysis tools in domains where data privacy, reliability, and safety are critical. I plan to carry over the lessons on interviewing users and understanding their thought process and perspective when using automation technologies to my academic research.

Future Directions (Co-Planning Infrastructure, AI-Assisted Design, Privacy and Guarantees): Building from these foundations (NL, AR/VR), I am interested in developing co-planning infrastructure that leverages these interfaces and related modalities as operations on the shared task, environment, and subproblem representations developed in Thrust 1. This focuses not just on one-way task specification, but instead supports iterative dialogue where humans and AI algorithms jointly refine missions and plans. This naturally connects with colleagues in HRI, HCI, human-subjects research, and cognitive science to study what levels of abstraction and types of communication are natural for non-expert users and how they affect mental models and performance.

One promising co-planning application is AI-assisted CAD and design-for-manufacturing/assembly. Human designers and the planning algorithms developed in Thrust 1 share CAD-centric representations of parts, fixtures, and assembly sequences, using multi-robot planning to evaluate metrics such as feasibility, cost, and robustness and to suggest alternative designs which better utilize robot capabilities. I will seek collaborations with CAD, design automation, and manufacturing researchers and companies interested in deploying robotic solutions.

There is also the question of extending co-planning to large multi-robot systems, where humans co-plan high-level policies for dozens or hundreds of robots. The guiding principle in Thrust 1 applies as the system must adopt a hierarchical approach to scale the complexity while staying in the bounds of human mental models.

These directions also raise safety, privacy, and security concerns, from physical safety and privacy in intimate settings, to IP protection in AI-assisted CAD, to cybersecurity for large-scale systems exposed to external co-planning interfaces. I will work with colleagues in these areas to create robust solutions safe to deploy in the real world.

Potential Funding Sources: NSF Human-Centered Computing / Cyber-Human Systems (HCC/CHS), NSF Future of Work at the Human-Technology Frontier (FW-HTF), CPS, FRR, NSF Secure and Trustworthy Cyberspace (SaTC), CMMI, Microsoft Research, ARA

Thrust 3 - Hardware and Software Acceleration: How do we design planning architectures that fully exploit modern hardware and software acceleration?

As robotic intelligence and interactions with it improve, we still face the challenge of making these systems fast and responsive enough for real-world (large-scale) applications. Advances to hardware and software design patterns offer opportunities to accelerate the core components underlying the systems developed in Thrusts 1 and 2.

Prior and Ongoing Work (Computational Geometry, GPU/Serialization, Parallel Computing): I lead a graduate research team (4 Ph.D. students - 1 now graduated, and 1 M.S. student) in collaboration with theory faculty at UIUC to exploit computational geometry to accelerate core motion planning algorithms to respond to physical environment changes (e.g., changes in obstacle positions). Our initial work demonstrated a **60% speedup in basic motion planning settings** [8]. We have four additional manuscripts in preparation for submission extending the basic computational techniques, parallelizing the primitive geometric operations on the GPU, and exploiting the reduction in the cost of motion planning primitives in the frameworks from Thrust 1.

Additionally, as part of our multi-robot research, we **developed a parallelized version of a core multi-agent pathfinding method** underlying many methods developed by both us and dozens of other research labs [9].

Future Directions: I am interested in revisiting planning architectures in the wake of drastic reductions in the cost of basic motion planning primitives both from approximation methods like [8] and modern CPU/GPU acceleration. Many existing approaches defer motion planning to avoid its computational cost, sacrificing feasibility and cost information at higher decision layers to keep runtimes manageable. Now that motion planning is cheaper, these architectures need to be reevaluated to see where the new bottlenecks lie, how the cost-intelligence boundary shifts within frameworks like DaSH and ARC from Thrust 1, and which theoretical guarantees (e.g., optimality or robustness) become practical when motion planning is no longer a prohibitive cost.

Complementary to this, I am interested in leveraging the insights developed in Thrust 2 to inform what information is actually useful when performing co-planning. Co-planning interfaces require fast, informative feedback about feasibility, costs, alternative plans, etc., but this does not mean that all information is necessary or productive when exposed to the human in the loop. This creates an opportunity for collaboration of HCI, cognitive science, systems, and theory researchers to design abstractions and approximations well suited for both co-planning and acceleration.

Finally, I see opportunities to unlock large-scale autonomous robotic systems by co-designing frameworks and representations for hardware acceleration that leverage cloud and edge computing in dynamic settings with hundreds of cooperating robots. This will again require collaborations with privacy and security colleagues along with systems and networking researchers to build safe, robust systems deployable in real-world settings.

Potential Funding Sources: NSF FRR, CPS, RI, CMMI, CISE core systems programs, NVIDIA Academic Programs, Amazon Web Services (AWS) cloud credits, ARA, ARM

References

[1] S. Lee, **James D. Motes**, I. Ngui, M. Morales, and N. M. Amato, "Lazy-dash: Lazy approach for hypergraph-based multi-robot task and motion planning," *Conditionally accepted to IEEE Transactions on Robotics*, 2026.

[2] **James D. Motes**, T. Chen, T. Bretl, M. M. Aguirre, and N. M. Amato, "Hypergraph-based multi-robot task and motion planning," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 4166–4186, 2023.

[3] I. Solis, **James D. Motes**, M. Qin, M. Morales, and N. M. Amato, "Adaptive robot coordination: A subproblem-based approach for hybrid multi-robot motion planning," *IEEE Robotics and Automation Letters*, vol. 9, no. 8, pp. 7238–7245, 2024.

[4] C. McBeth, **James D. Motes**, I. Ngui, M. Morales, and N. M. Amato, "Scalable multi-robot motion planning using guidance-informed hypergraphs," *Conditionally accepted to IEEE Robotics and Automation Letters*, 2026.

[5] M. Qin, I. Solis, **James D. Motes**, M. Morales, and N. M. Amato, "K-arc: Adaptive robot coordination for multi-robot kinodynamic planning," *IEEE Robotics and Automation Letters*, 2025.

[6] I. Solis, **James D. Motes**, M. Qin, M. Morales, and N. M. Amato, "Experience-based subproblem planning for multi-robot motion planning," *arXiv preprint arXiv:2411.08851*, 2024.

[7] I. Ngui, C. McBeth, A. Santos, G. He, K. J. Mimnaugh, **James D. Motes**, L. Soares, M. Morales, and N. M. Amato, "Erupt: An open toolkit for interfacing with robot motion planners in extended reality," *arXiv preprint arXiv:2510.02464*, 2025.

[8] S. Ashur, M. Lusardi, M. Markowicz, **James D. Motes**, M. Morales, S. Har-Peled, and N. M. Amato, "Spite: Simple polyhedral intersection techniques for modified environments," *International Workshop on the Algorithmic Foundations of Robotics*, 2024.

[9] H. Lee, **James D. Motes**, M. Morales, and N. M. Amato, "Parallel hierarchical composition conflict-based search for optimal multi-agent pathfinding," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7001–7008, 2021.