Research Statement

Diane Uwacu
https://parasollab.web.illinois.edu/people/duwacu

Overview

Motion Planning involves finding a valid path for a moving object in a given environment. Motion planning algorithms are applied to several disciplines, including robotics, automated design, computational geometry, and biology.

Due to their complexity, motion planning problems are generally solved using sampling-based algorithms that rely on an abstraction of the problem called a configuration space ($\mathcal{C}_{space}$) and good computational geometry tools.

In my research, I have proposed algorithms that take advantage of workspace guidance to discover important planning regions in the environment efficiently. This has entailed creating a general representation of skeleton annotations to guide traditional planners. I am interested in applying motion planning algorithms to robotics and computational biology applications and studying accessibility. These algorithms are included in the recently open-sourced Parasol Planning Library (PPL). I want to highlight the main ideas in my research and discuss my interests in advancing the area.

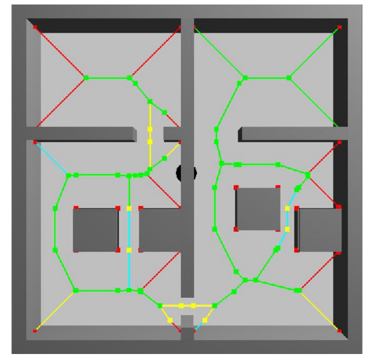

Figure 1: Four connected rooms and their clearance annotated workspace skeleton color-coded from red (lowest clearance) to green (highest clearance).

Algorithms

Workspace Skeleton-guided Dynamic Regions Roadmaps

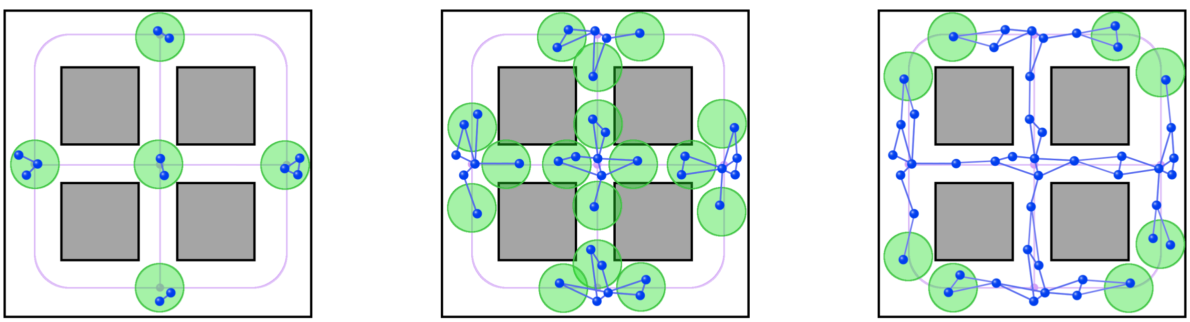

(a) (b) (c)

Figure 2: An illustration of Dynamic Regions Roadmaps. Obstacles are shown in gray. The *workspace skeleton* is shown in purple. (a) The algorithm samples initial connected components (blue) in regions (green) around each skeleton vertex. (b) Sampling regions expand outward along the skeleton edges. (c) The components in the middle tunnels successfully connect to form bridges. The outer passages are still expanding.

Workspace guidance involves using the workspace structure to direct how sampling-based planners explore the environment. A *workspace skeleton* is a minimal representation of the workspace that contains information on the connectivity of the workspace. Some examples include the medial axis skeleton in 2D environments and the mean curvature skeleton in 3D environments.

I have collaborated with a fellow graduate student in the Parasol lab to develop the Dynamic Regions Roadmaps method [5] illustrated in Figure 2. The algorithm was tested on environments with long narrow passages that are difficult to plan in with regular methods and yielded well-connected roadmaps in record time.

Annotated Skeleton Guided Planning

The success of the skeleton-guided planning strategies depends on how good the guidance of the skeleton is for the $\mathcal{C}_{space}$. However, additional features in the workspace often provide more insight to the planner. I

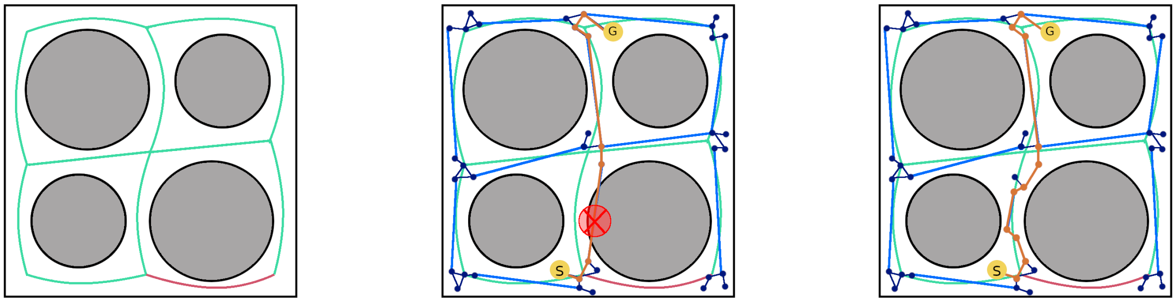

(a) (b) (c)

Figure 3: An illustration of the Hierarchical Annotated Skeleton Planner. Obstacles are shown in gray. (a) The skeleton is annotated with obstacle clearance. (b) Sampling is done inside regions centered at the skeleton vertices, and the local connected components are connected lazily by assuming C-space connectivity. An initial path is queried. (c) The path is validated and fixed accordingly.

worked with several undergraduate students to develop mechanisms to annotate the skeleton with workspace information relevant to the problem. For example, a skeleton annotated with obstacle clearance, like the one shown in Figure 1, would provide better guidance to find safe paths for a moving robot. The guidance of annotated workspace skeletons makes it possible to use basic sampling-based planning algorithms to achieve specific path requirements.

More recently, I have worked with an undergraduate research student to develop the Hierarchical Annotated Skeleton-guided Planning Strategy [3] illustrated in Figure 3. The strategy starts with the skeleton-based path, then fixes it hierarchically based on the solution's acceptance criteria. It improves its reliance on the skeleton as planning evolves and considers additional information to bias its path-finding efficiency.

Currently, I am working with three undergraduate students to extend this algorithm to tree-based algorithms. The algorithm will handle goal-oriented motion planning problems with kinodynamic constraints and be applied to computational biology.

Applications

Robotics

We have applied our techniques to solve motion planning in robot navigation problems. Empirically, we have shown that our strategies are among the best to tackle the narrow passage problem, which is a major contribution to motion planning research.

Figure 4 shows examples of robotics problems on which we have tested our algorithms, including but not limited to long narrow passages, heterogeneous environments, 3D environments, and many obstacles. With the guidance of the workspace skeleton, the planner easily finds topological intersections, discovers the connectivity of the workspace, and solves the motion planning queries fast, as illustrated by the experiment done in a garage environment in Figure 4a. Annotating the skeleton with additional information about the environment biases the planner to prioritize exploring regions containing the desirable solution. This is useful for robotic navigation problems in crowded spaces like a shopping store shown in Figure 4b, for which the goal is to find safe paths.

Computational Biology

A protein is a large molecular structure involved in several essential reactions. It reacts with a drug molecule called a ligand to activate or inhibit its functions. Protein-ligand binding research focuses on one of two challenges: binding site identification or binding site accessibility.

I applied the annotated-skeleton-guided motion planning algorithm to identify binding site tunnels that con-

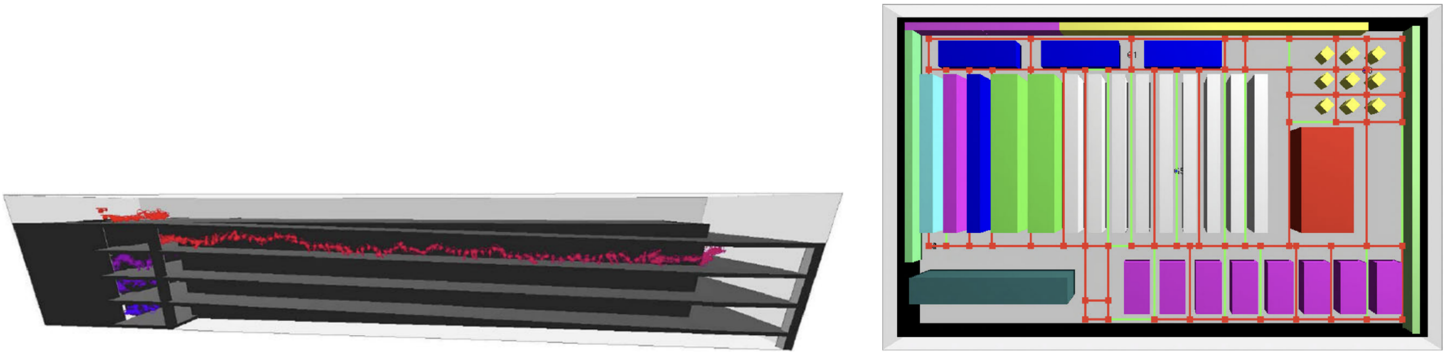

(a) (b)

Figure 4: Examples of robotics motion planning environments. (a) The Garage problem presents a series of ten queries scattered across the levels of the structure. The primary sources of difficulty are thin walls and large scale. (b) The grocery store environment with a three DOF nonholonomic robot. The aisles in the store environment have different lengths, and the environment has fifty obstacles that are not easily scalable.

nect the ligand to the protein's binding site, as illustrated in Figure 5a. The workspace skeleton, annotated with protein-ligand interaction energy, indicates geometric tunnels and biases the motion planning algorithm to first explore tunnels with favorable energy that are likely to lead the ligand to the binding site. Due to the protein's geometric shape, there are many misleading dead-end voids that the workspace skeleton saves the planner from exploring.

With this algorithm, I worked with an undergraduate summer student to study several problems related to protein-ligand binding [4]. In one example, we studied the haloalkane dehalogenase enzyme, *DhaA*, shown in Figure 5b. DhaA is essential in the biodegradation of soil pollutants like 1,2,3-trichloropropane (TCP), an environmental pollutant with no known natural biodegradation pathway. DhaA's activity is limited by its low stability and vulnerability to the harsh conditions of its reactive environment. Mutants have been engineered in laboratories to enhance DhaA's stability and/or activity towards TCP. In all the proposed mutants, our algorithm correctly identified the active tunnels and simulated the path of the ligand to reach the binding site.

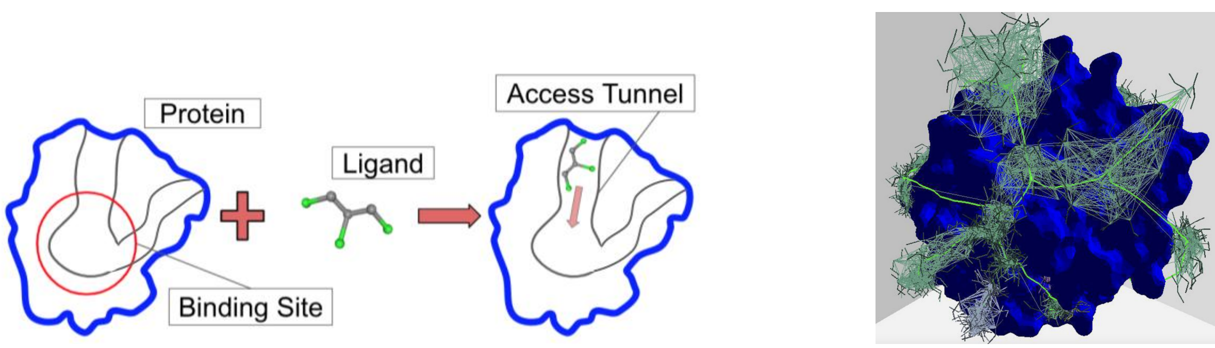

(a) (b)

Figure 5: Binding Site Accessibility: (a) The ligand can reach the binding site through an access tunnel. (b) Roadmap of TCP-DhaA accessibility created with skeleton guidance.

Future Work

I am interested in discovering fast computational solutions to geometrically complex planning environments. More recently, I collaborated with a junior Ph.D. student in the lab to expand the skeleton-guided motion planning solutions to the composite space of multi-robot systems. Our work [1] is in the pipeline for publication. We have started defining and evaluating guidance for motion planning algorithms in another project.

Our preliminary work [2] was presented at the IROS workshop evaluating motion planning performance.

The algorithms used to model protein-ligand interaction are not robust yet to have a more general application in molecular pharmacology and drug selection. I want to expand these algorithms to improve their accuracy and make them easily accessible to experts. I plan to work with undergraduate students in this domain to find testing data, design and implement algorithms, and contextualize our findings through technical writing. In the near future, I would like to explore the protein-ligand binding problem further as I develop multi-disciplinary algorithms that can be applied to robotics and computational biology.

The research that I did in my dissertation research is proof that these general algorithms can be applied to interesting multi-disciplinary problems. In addition, I have worked with many undergraduate research students, which has helped me think of ways to divide the research into modularized projects that can be completed in a semester or two and could fit capstone project requirements.

To help students get comfortable with the foundations of research in planning, I plan to develop a crash course similar to the one I described in my teaching statement and adapt it to the school's curriculum. This will help students get a guided introduction to research through technical reading, reflection questions, and coding activities. Here are examples of capstone projects that I could guide an undergraduate student to work on:

- **Real-time adaptive guided planning with machine learning**: Applying reinforcement learning to the annotated skeleton guidance, the algorithm can adapt to changes based on the real-time feedback received. It would be interesting to explore this idea with examples where planning performance relies on local and global adaptability. At the end of this study, we will be able to assess the impact of learning on guided algorithms.
- **Factors of protein-ligand affinity**: Understanding how different protein and ligand properties affect binding affinity would help us make ligand recommendations and validate structures.
- **Parallel Guided Planning Strategies**: Using computing resources to parallelize skeleton-guided planning strategies would boost performance in harder problems.

References

1. C. McBeth, J. Motes, **D. Uwacu**, M. Morales, N. M. Amato, "Scalable Multi-robot Motion Planning for Congested Environments Using Topological Guidance," *arXiv:2210.07141*, Oct 2022
2. A. Attali, S. Ashur, I. B. Love, C. McBeth, J. Motes, **D. Uwacu**, M. Morales, N. M. Amato, "Evaluating Guiding Spaces for Motion Planning," *Evaluating Motion Planning Performance Workshop (IROS-EMPP)*, Kyoto, Japan, *arXiv:2210.08640*, Oct 2022
3. **D. Uwacu**, A. Yammanuru, M. Morales and N. M. Amato, "Hierarchical Planning With Annotated Skeleton Guidance," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11055-11061, Oct. 2022, doi: 10.1109/LRA.2022.3196885. Presented at IEEE/RSJ Int. Conf. on Intel. Robots and Systems (IROS 2022), Kyoto, Japan
4. **D. Uwacu**, A. Ren, S. Thomas, N. M. Amato, "Using Guided Motion Planning to Study Binding Site Accessibility," *Proceedings of the 11th ACM Int. Conf. on Bioinformatics, Computational Biology and Health Informatics*, Issue: 109, (Virtual) New York, USA, Sep 2020
5. R. Sandström, **D. Uwacu**, J. Denny and N. M. Amato, "Topology-Guided Roadmap Construction With Dynamic Region Sampling," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6161-6168, Oct. 2020. Presented at IEEE/RSJ Int. Conf. on Intel. Robots and Systems (IROS 2020), Las Vegas, USA