Adaptive Motion Planning for Complex Planning Problems

Lydia Tapia and John Baxter

***Abstract*— Motion planning has been used to solve problems of high complexity in both robotic and biological domains. In robotics, the topology of the planning environment often drives the problem's complexity. Environments can consist of many different regions each of which may be well suited to a specific planning approach. In biological domains, problem complexity is primarily driven by the size of the moveable object. For example, small proteins have hundreds of degrees of freedom, medium proteins have thousands, and two proteins interacting can have even more.**

We present recent intelligent techniques applied to *probabilistic roadmap methods (PRM)* in order to efficiently and automatically solve complex motion planning problems. These techniques use automated methods to learn abut the problem space and then adapt based on the problem's characteristics. We demonstrate the use of these automated methods in both robotics and protein folding applications.

I. INTRODUCTION

In order to address complex planning problems, adaptivity has been proposed as a solution [1]–[6]. While significant improvement has been shown over non-adaptive approaches, these methods have all been seen to have serious drawbacks that limit their usefulness such as requiring significant user intervention (e.g., manual classifications of training instances for supervised machine learning methods, parameter tuning to set learning rates and learning weights) or restricting the types of problems they are able to solve.

II. APPLICATION: ROBOTICS

Although there are many motion planning techniques, there is no method that outperforms all others for all problem instances. Rather, each technique has different strengths and weaknesses which makes it best-suited for certain types of problems. Moreover, since an environment can contain vastly different regions, there may not be a single planner that will perform well in all its regions. Ideally, one would use a suite of planners in concert and would solve the problem by applying the best-suited planner in each region.

In order to take advantage of this existing library of methods, Feature Sensitive Motion Planning has explored using the features of the planning space to help decide where and when to apply particular planners. In preliminary work, a supervised learning method classified features of the space and selected a sampler to apply in a certain region of the space [4], [5]. In recent work, we have used spatially and temporally identified features in order to better decompose the problem and selectively apply planners that adapt over time [9]. This new strategy takes advantage of unsupervised learning methods at all stages of the planning process and produces solutions in complex spaces with little cost and less manual intervention compared to other adaptive methods.

An example is shown in Figure 1 for a maze environment with a movable object. First, features from a small sampling of the space are identified and used to cluster the samples. Each cluster relates to a region of the space (Figure 1(b)). In order to define the optimal number of clusters (n), the elbow criterion is calculated from the variance in the clusters (Figure 1(d)). Intuitively, this criterion selects n such that adding additional clusters does not add sufficient information. Subsequently, an appropriate planner can be selected from a library and applied in each region.

III. APPLICATION: PROTEIN FOLDING

Protein motions play an essential role in many biochemical processes. For example, as proteins fold to their native, functional state, they sometimes undergo critical conformational changes that affect their functionality, e.g., diseases such as Mad Cow or Alzheimer's are associated with protein misfolding and aggregation. Knowledge of the stability, kinetics and detailed mechanics of the folding process may provide insight into how and why the protein misfolds.

In order to computationally study biologically-relevant molecules, we have explored applying PRMs to study molecular motions. For small proteins, both microscopic folding pathways and global folding properties can be studied in a few hours. However, as protein size increases, the computational complexity of the planning problem also increases. Using standard techniques, planning for larger proteins can become infeasible.

In order to efficiently simulate the motions of medium and large proteins, we have explored adaptive planning techniques. One technique, dimensionality reduction, can be applied to landscape models [8]. This computational technique finds the principal features of a high-dimensional space, represented by our motion landscapes, and returns a lower-dimensional representation that still captures the principal features. Dimensionality reduction enables quick and useful global analysis of our landscapes. Through a new use of it as an analysis tool, it can reduce our original model size by almost half, thus facilitating the study of larger proteins.

Our results are quite promising. Our new techniques have been able to capture structural events that have been shown in lab experiments, such as those found for the small proteins

J. Baxter and L. Tapia are with the Department of Computer Science, University of New Mexico, Albuquerque, NM 87131-0001 {tapia, johnbaxter}@cs.unm.edu

Funding for this poster presentation was provided by National Institutes of Health (NIH) Grant P50GM085273 supporting the New Mexico Spatiotemporal Modeling Center and NIH Grant P20RR018754 supporting the Center for Evolutionary and Theoretical Immunology.

(a) (b) (c) (d)

Fig. 1. Automatic region identification in a maze environment. (a) Environment shown with movable body shown above and enlarged. Notice there are three different regions which the robot must traverse: open, constrained, and open. (b) Clustering identifies 3 regions (circled) corresponding to the features of the space. (c) Continued clustering can unnecessarily split the regions further. (d) An automated method, the elbow criterion, determines the best number of regions (red star).

G and its mutants, NuG1 and NuG2 [7]. The application of dimensionality reduction to our roadmaps produced maps that were up to 53% smaller for all proteins studied, yet were still able to capture the experimentally determined folding orders including those for Protein G and its mutants [8]. We also demonstrate its usefulness with proteins of medium to large size.

REFERENCES

[1] J. Berg and M. Overmars. Using workspace information as a guid to non-uniform sampling in probabilistic roadmap planners. *Int. J. Robot. Res.*, 24(12):1055–1072, 2005.

[2] B. Burns and O. Brock. Toward optimal configuration space sampling. In *Proc. Robotics: Sci. Sys. (RSS)*, pages 105–112, 2005.

[3] D. Hsu, G. Sánchez-Ante, and Z. Sun. Hybrid PRM sampling with a cost-sensitive adaptive strategy. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 3885–3891, 2005.

[4] M. Morales, L. Tapia, R. Pearce, S. Rodriguez, and N. M. Amato. A machine learning approach for feature-sensitive motion planning. In *Algorithmic Foundations of Robotics VI*, pages 361–376. Springer, Berlin/Heidelberg, 2005. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), Utrecht/Zeist, The Netherlands, 2004.

[5] M. A. Morales A., L. Tapia, R. Pearce, S. Rodriguez, and N. M. Amato. C-space subdivision and integration in feature-sensitive motion planning. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 3114–3119, April 2005.

[6] S. Rodriguez, S. Thomas, R. Pearce, and N. M. Amato. (RESAMPL): A region-sensitive adaptive motion planner. In *Algorithmic Foundation of Robotics VII*, pages 285–300. Springer, Berlin/Heidelberg, 2008. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), New York City, 2006.

[7] L. Tapia, X. Tang, S. Thomas, and N. M. Amato. Kinetics analysis methods for approximate folding landscapes. *Bioinformatics*, 23(13):539–548, 2007. Special issue of Int. Conf. on Intelligent Systems for Molecular Biology (ISMB) & European Conf. on Computational Biology (ECCB) 2007.

[8] L. Tapia, S. Thomas, and N. M. Amato. Using dimensionality reduction to better capture RNA and protein folding motions. Technical Report TR08-005, Parasol Lab, Dept. of Computer Science, Texas A&M University, Oct 2008.

[9] L. Tapia, S. Thomas, B. Boyd, and N. M. Amato. An unsupervised adaptive strategy for constructing probabilistic roadmaps. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 4037–4044, May 2009.

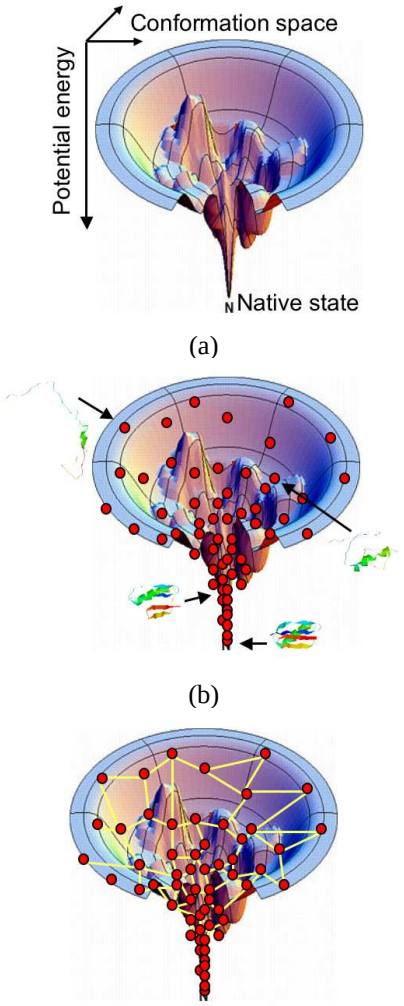

(a) (b) (c)

Fig. 2. (a) The folding energy landscape is the set of all protein conformations and their associated energy. Building an approximate map of the energy landscape consists of two steps: (b) conformation sampling and (c) connecting samples together with feasible transitions.