RESEARCH STATEMENT

Tim Bretl

My goal is to enable autonomous agents to interact productively with the physical world. To this end, I seek to better understand and control the motion of complex mechanical and biological systems. Recent advances in sensing technology allow the collection of massive amounts of data and the generation of sophisticated world models. My research focuses on interpreting this data to decide how to act.

As a PhD student, I addressed the problem of planning the motion of a multi-limbed robot to “free-climb” vertical rock surfaces. Free-climbing relies on frictional contact with rock features rather than on special fixtures or tools like pitons. It requires strength, but more importantly it requires deliberate reasoning: not only must the robot decide how to adjust its posture to reach the next feature without falling, it must plan an entire sequence of steps, where each one might have future consequences. In my work, I made this process of reasoning manageable by decomposing a free-climbing robot’s configuration space into manifolds associated with each state of contact between the robot and its environment. I created a multi-step planning framework that decides which manifolds to explore by generating a candidate sequence of hand and foot placements first, and a one-step planning algorithm that explores individual manifolds quickly by taking advantage of the interaction between static equilibrium and the topology of closed kinematic chains.

Free-climbing is highly interdisciplinary, requiring advances across the fields of engineering, computer science, and applied mathematics. For example, some of my work focused on foundational geometric algorithms: testing whether a robot can balance on a collection of rock features is equivalent to testing the membership of a point in the projection of a high-dimensional convex set, for which I developed a fast solution method. Likewise, I also dealt with issues of implementation and experimentation: in cooperation with the Mechanical and Robotic Technologies Group at NASA-JPL, I used my planner to enable a four-limbed robot to free-climb an indoor, near-vertical surface covered with artificial rock features. This interplay between abstract theory and a working, physical system is my strongest interest.

As a postdoctoral fellow, I am extending my work to other legged robots that navigate steep or uneven terrain. These robots include a humanoid for industrial and household service and a hexapod for lunar exploration and construction. They have many degrees of freedom and are subject to constraints that require a consideration of dynamics, visual and tactile sensor integration, and the aesthetics of motion. To handle this complexity, I am designing algorithms for the automatic synthesis of motion strategies. My goal is to derive both a method of determining which motions are hard to plan or to execute (thus, are useful to learn) and a method of deciding when a motion strategy should be applied.

The idea of a motion strategy is not restricted to the domain of legged robots. For example, proteins also exhibit characteristic motion. Databases that classify different types of conformational changes are a powerful tool in understanding the behavior of these proteins. However, protein motions are not entirely explained by these databases, just as the actions of legged robots (or of people) do not entirely correspond to a set of learned skills. One challenge when synthesizing and manipulating proteins is to combine case-based reasoning with other methods of planning and control. I intend to address this challenge in future work.

Similar problems arise in other applications, ranging from the very large (intelligent transportation systems, groups of cooperating robots) to the very small (micro- and nano-scale manipulators, continuously deformable robots, wearable electronics and prosthetics). The basic physics of these systems are largely known, but the resulting models are often too complex to be useful. To control their interaction with the physical world, we need ways of conceptualizing and understanding their capabilities. But what characterizes the underlying structure of mechanical and biological motion? Can structural knowledge be incorporated into the design of algorithms for planning, control, and machine learning? Is it possible to take advantage of characteristic uncertainty to control robots that would not otherwise be controllable? How should physical models be used to interpret massive amounts of biometric data? These are some of the questions that drive my research. In each application, I want to bridge the gap between abstract theory, mathematical tools, and physical practice – to understand what we do, why we do it, and how our techniques generalize.