RESEARCH STATEMENT: Joel M. Esposito

OVERVIEW, PHILOSPHY AND VISION

Applicability Recent federal laws require the DoD to make 1/3 of its combat vehicles *unmanned* by 2025; they also forbid the acquisition of new *manned* combat vehicles if it is possible for an unmanned vehicle to meet the mission requirements. I develop algorithms that enable these systems to plan their own motions – faster, safer, more reliably and more collaboratively. A potential application scenario for my work would be teams of unmanned boats cooperatively towing large objects -- imagine towing the USS Cole to a friendly port or dragging mine counter measure arrays through Umm Qasr Port, in Iraq. Because of its DoD relevance, my work has been continuously supported by the Office of Naval Research and the Office of the Secretary of Defense.

Quality When I interviewed at USNA in 2002, the then Associate Dean for Faculty Affairs stated, "While the *quantity* of research done at USNA may not be the same as that of a Research I University, we expect our faculty to produce the same *quality* of work." To me, quality research is deeply rooted in the literature, contains highly novel results, is perceived as important by the research community, and combines a diversity of investigative approaches including theory, and experimental verification. I strive to include all these elements in my work.

Integration of Research and Education Roger Lewin, the science writer, once said: "Too often we give our children answers to remember rather than problems to solve." Undergraduate student research seems to be one of the best remedies for that. From 2004 – 2010 I coordinated the independent research projects (ES49X) within the Systems Engineering department. Personally, I have co-authored 9 conference papers with midshipman and worked with 12 Trident, Honors and independent study students. I have helped create two new courses that exposure students to my research area (ES456 and ES486F); and many of my capstone design teams have supported externally sponsored research project. Every summer I work with students from local high schools through ONR's Science and Engineering Apprenticeship Program.

Leadership Good researchers are intellectual leaders. I have created forums for the exchange of ideas, organizing special session at conferences. I have lead collaborations and mobilized resources for colleagues and students as the principle investigator of [G1 –G10], totaling nearly $1 million. I enjoy mentoring junior faculty getting started in research [J12]. I have been invited to serve on a review panels for NSF, ONR, DoT and NASA.

Vision One of my career goals is to establish USNA as a center of excellence in autonomous vehicle research, recognized by research colleagues, DoD funding sources, and serving as an attractor to students. In FY12 I was awarded $360,000 by the Office of the Secretary of Defense to develop infrastructure to support this goal [G10]. I am in the process of procuring a 24 foot power boat, outfitted with a variety of state of the art sensors which will serve as a test-bed for autonomous vehicle research, motion planning in particular, for years to come.

BACKGROUND: INTRODUCTION TO ROBOT MOTION PLANNING

Consider the maze in Figure 1. Even my 6 year old son is able trace out a solution, but he certainly can't articulate *how* he did it out. This amazing ability comes from specialized "hardware" in our brains dedicated to spatial reasoning (likely the hippocampus), and so far no engineer has been able to replicate it.

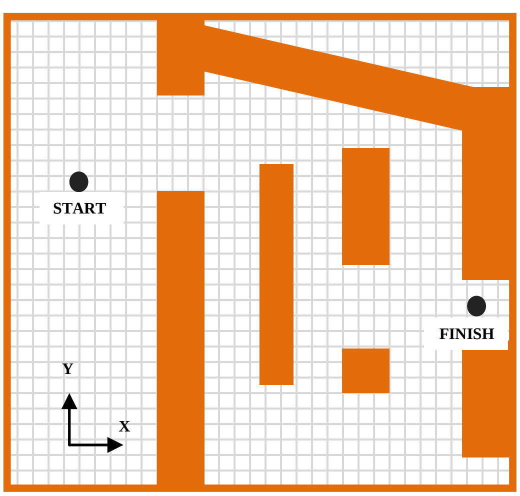

Figure 1: A simple maze.

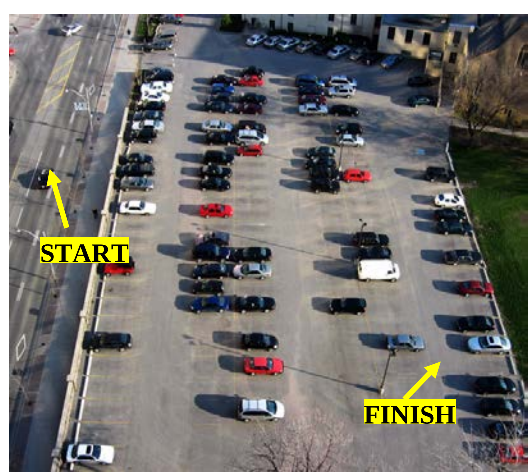

Figure 2: A more realistic version of Figure 1.

Understanding why this is such a challenge for robots requires you to look at the problem as a computer might. First, lacking the ability to interpret Figure 1 directly, it needs a numerical representation of the scene, perhaps something like a table listing the coordinates of the obstacle vertices. Then, it requires a set of algorithmic instructions for using the data to determine a path, without resorting to intuition or the pictorial representation. It becomes clear that the problem is non-trivial; yet it is a capability that is central to autonomy.

Further, consider how grossly abstracted the above example was. What if the problem in Figure 1 was replaced with the scene Figure 2. In addition to reaching the goal (getting to the assigned parking spot), the robot car must account for a variety of secondary objectives such as: not hitting pedestrians, abiding by state traffic laws, and minimizing gas consumption. Also consider that robots that are expected plan their motions quickly – in seconds not hours. The challenge of planning motions fast while accounting for secondary objectives is my core focus.

EARLIER WORK Prior to 2007, my research was split into two related, but distinct, areas. The first concerned numerical methods for simulating robots which are controlled by such computer algorithms. I pioneered a new technique for simulating around discontinuities [B1], later developed further in [J3] and [J7]. This work was part of a larger project, showcased in the prestigious, society-wide journal, *The Proceedings of the IEEE* [J2] and in the case study [J8]. According to Google Scholar these papers have been collectively cited 252 times and I continue to frequently receive inquiries about those works from researchers around the world.

My second focus was robot motion planning. I introduced several algorithms in this area, including: [C6], which was nominated for the best paper award (out of 683 papers) at an international robotics conference; [B4] and [B5] which were published in the two of the most prestigious conferences in robotics (single track, double blind review, with rebuttal phase, acceptance rate under 20%); and [J1] -- cited 83 times.

CURRENT RESEARCH After my promotion to Associate Professor in August 2007 I decide to focus exclusively on robot motion planning. After a year of reading and developing some experimental and computational infrastructure, I began producing conference papers in 2008 [C11-17]; those efforts then matured into [J10-J15].

Cooperative manipulation by large groups of robots

Imagine many robots working collectively to manipulate a larger object, much in the way that army ants cooperate to transport prey (Figure 3). In 2005, I approached ONR with the concept of a team of autonomous tugboats cooperatively towing a barge (Fig. 4). They have supported this from FY05-10 [G2, G3, G5] ($360,000).

Figure 3: Army ants work cooperatively to manipulate and transport large prey.

In 2008 I organized a focus session on cooperative robot manipulation at an international robotics conference, which gave the leading researchers in this field an opportunity to showcase their work and discuss the challenges ahead. At that meeting I presented [C11] which discusses how the robots should select the best contact location around the object to improve the group's maneuverability and mechanical advantage. I developed a mathematical formula that encodes the "quality" of the grasping configuration. I was able to show that this formula has a special structure that enables each tug to compute a solution quickly and independently.

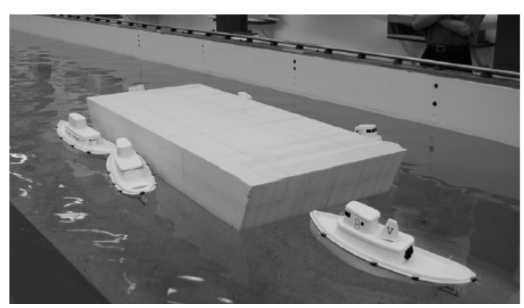

Figure 4: A team of student-built scale tugboats manipulate a small barge.

[C14] addresses the "approach phase" of such tasks, where the robot tugs must navigate to desirable attachment locations around the barge while simultaneously avoiding collisions with other tugs and the barge. I offered a provably correct algorithm to accomplish this; and demonstrated it experimentally on a group of ground robots.

During the "transport phase", the tugboats must coordinate their applied forces to move the barge in the desired fashion. In [C13, J15] , I proposed a series of "pushing" control laws. Surprisingly, when the goal is to move the barge at a constant speed I showed it is possible to use a force control law that does not require any explicit communication between teammates. On the other hand, if the goal is to move the barge along a more complex trajectory we show that each tug must have knowledge of the other tug's attachment locations around the barge. I introduced a distributed scheme which lets the tugs estimate the location of their teammates.

To validate some of these ideas in the field, I partnered with Assoc. Prof Feemster and several midshipman to design and built two experimental test beds: one was a scale model of a YP for use in the tow tank, the other a 10 foot Jon boat for operation in local creeks. Our efforts appeared in [SC7], [C12] and ultimately were featured in special issue of the Journal of Field Robotic entitled "State of the Art in Marine Robotics" [J10]. Note that in 2010 and 2011 this journal had the highest Thomson-Reuters Impact Factor among all robotics publications.

Cooperative manipulation – What can we learning by studying humans?
Consider two humans performing a similar cooperative task – for example moving furniture (Fig. 5). The paradigm by which they interact is quite different from the way most robots coordination protocols are designed. Rather than transmitting their planned accelerations over a WiFi network hundreds of times a second as many robot motion control laws would require, they are much more likely to exchange occasional high level "contextual" information (e.g., "move a little to your right"). And while the resulting performance might defy any mathematical proof of stability or correctness, it is likely to work robustly under a wide range of realistic conditions.

Figure 5: What strategies do humans use when cooperating?

In [C15] and [J14], I worked with Asst. Prof. Parikh and LT Searock to determine why, how and when people prefer to use *verbal* communication to coordinate their efforts. The ultimate goal of this study would be to extract a set of principles that could be used to design better robot-robot or human-robot cooperative manipulation protocols. We found that while people think they perform better and experience less frustration when they are permitted to speak to their teammates, there is no statistically significant improvement in their performance, suggesting that verbal communication fills a primarily social role. We also discovered that the walking trajectories of the people take on a rather predictable pattern, perhaps helping each partner anticipate the actions of the other.

Multi-robot motion planning with wireless network constraints
Another problem I have worked on is *ad hoc* wireless network preservation. Imagine a team of robots cooperatively searching an office building for an intruder, as depicted in Figure 6. Each robot must share information with its teammates about which rooms have been "cleared" and when. In practice, the strength of the wireless links depends in large part on the proximity of the robots and maintaining line-of-sight. These constraints present secondary objectives –in addition to the primary objectives of moving toward a goal and not colliding with each other.

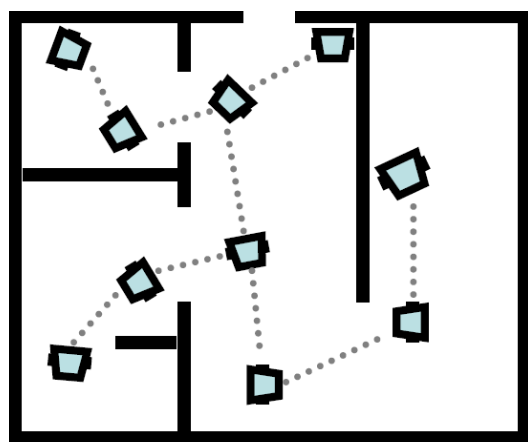

Figure 6: Overhead view of a team of mobile robots exploring an office building while maintaining wireless communication links.

While several other researchers had offered algorithms that addressed the "range constraint", in [J11] I introduced the first algorithm that also accounted for the "line-of-sight constraint" – which turns out to be far more limiting in practice. I offered a topological classification of the obstacle configurations from which a solution is feasible; a mathematical proof of the algorithms performance; and an experimental demonstration. This grew out of earlier effort with my Trident scholar Midn Dunbar [SC1, SC2].

Predicting the Performance of Randomized Algorithms
My last, and most fundamental, set of recent contributions to motion planning provides some theoretical results concerning the speed and reliability of *randomized* motion planning algorithms. By 1995, it had been proven there were some fundamental limitations on how quickly exact motion plans could be computed, so researchers began exploring randomized algorithms – considered the current state of the art. They search for solutions incrementally, by examining thousands of randomly generated sample segments such as the car in Figure 7.

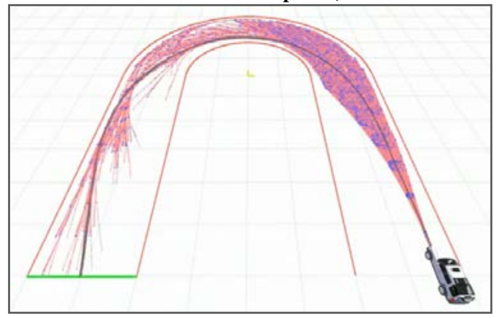

Figure 7: A car uses a randomized motion planning algorithm to negotiate a tight curve.

Despite their practical success and wide spread adoption, the behavior of these algorithms is still not well understood. For example, there was no way to predict how long the algorithms will take to find a solution. In cases where it is not know *a priori* if a solution exists, researchers must guess when to terminate the search.

I have derived some very important predictions on the run times of these algorithms. In [C17] I proposed a simple yet highly accurate statistical model of the average number of samples (proportional to run time) required to solve a given motion planning problem. In 2011 I was invited to present these results at the headline session "50 Years of Robot Motion Planning– Achievements and Emerging Results" at an international meeting of roboticists. In 2012, I modeled the variance as well, and generated experimental evidence that the distribution of the number of samples is Gaussian (a.k.a. bell-curved) [J13].

Together, these results have profound implications for motion planning because, for the first time, researchers can make strong, quantitative statements about a randomized algorithm's performance and progress. Examples include: "If no solution has been found after 10,000 samples have been evaluated, we can be 99.9% confident that this problem is infeasible"; or "There is a 95% chance this problem will require 7,500 (±500) samples to solve". The predictions can also be used by many of the adaptive algorithms proposed in the literature, which require a method to estimate search progress online.

A motion planning algorithm based on these developments was fielded on a midshipman built entry into the Association for Unmanned Vehicle Systems International (AUVSI)'s 2012 and 2013 Intelligent Ground Vehicle Competition, which took 2nd place in the autonomous navigation event, out of a field of 50+ teams from around the world.

FUTURE WORK At present, I am focusing on developing a 24 ft autonomous surface vessel, which will serve as a test bed for to demonstrate high speed, high fidelity motion planning algorithms in the field. Ongoing work in this area is being supported by the Office of Naval Research (ONR) [G8] and Office of the Secretary of Defense [G10]. My goal is to demonstrate a fully autonomous docking scenario by 2018.