Research Statement

Jeremy Gottlieb

1 Overview

My research interest broadly focus on autonomous systems that can successfully complete long-duration missions that require them to work both independently and in concert with human beings. The ability of autonomous systems to successfully operate for long periods of time without human intervention is becoming increasingly important in a wide variety of applications, particularly scientific and military missions. However, robots never operate without some sort of human oversight, even if only occasionally. As autonomous systems become more ubiquitous they will need to have interfaces which allow non-programmers to manipulate their behavior.

A future framework for such systems might be robots working in support of manned Mars missions. In this framework, a group of robots would be sent to Mars with a group of humans. Robots with sufficient autonomy would have a wider range that they could operate in. They would also be able to conduct certain science operations, such as drilling rocks, more easily than spacesuit encumbered humans would. Thus, the robots could perform the bulk of the scientific work while the humans build a semi-permanent base.

At this point, it would be difficult for humans to stay on Mars for an extended period of time. Robots, on the other hand, have successfully operated on Mars for several years. *Opportunity* has been operating on the surface of Mars for almost 11 years, despite originally being designed for only a 90-day mission. Improved long-term autonomy would allow the robots to continue conducting scientific experiments after the humans have left to return to Earth. The robots could also know to return to the base when the next mission arrives to support the next group of humans.

The ability of these robots to successfully operate without human intervention or oversight for extended periods will substantially enhance their ability to conduct scientific operations. The round-trip time for communications with Mars is anywhere from 8 to 48 minutes. There is also a month-long communications blackout every 26 months or so when the Sun gets between Mars and Earth. If we were to place a robot on Pluto, the communications round trip could be up to eight hours.

Even here on Earth there are environments in which improved autonomy would substantially enhance our ability to conduct scientific explorations. Autonomous underwater vehicles (AUVs) are getting cheaper to build and operate. With the right sensor package, research institutions can now launch an AUV from the beach and let it collect oceanographic data for extended periods of time. However, AUVs cannot communicate with scientists on shore unless they surface, which uses expensive satellite communications and exposes them to collisions with surface vessels. They then also have to wait for an oceanographer to evaluate the data and choose the next mission goal.

Given that oceanographic features of interest have variability both spatially and temporally, the time spent surfacing and waiting could mean that the feature of interest moves substantially before the AUV receives further instructions. Recent missions have also sent AUVs to map the underside of Antarctic ice shelves. On these missions, there is no opportunity for humans to communicate with the AUV at all. As with the Mars rovers, the AUVs can be substantially more effective if they are able to make their own decisions about what specific short-term goals fit their long-term mission profiles.

AUVs also highlight the need to provide interfaces that non-technical operators can effectively use. In my research with the autonomy group at the Monterey Bay Aquarium Research Institute (MBARI), one of the senior software engineers was required to accompany their *Dorado* AUV for every deployment to ensure that it booted up properly and correctly received its mission. Since all of the missions ran overnight, that same person frequently stayed up late monitoring the hourly messages from the AUV. If an oceanographer wanted to change what *Dorado* was doing mid-mission, she had to get in touch with the software engineer so that he could send the proper set of commands to the AUV. Clearly a system is needed whereby oceanographers can select new mission goals for the AUV.[1]

[1] We were in the process of designing such a system when MBARI shut down the autonomy group.

2 Past Research

My past research has focused on giving autonomous vehicles the capability to make their own decisions about where to travel in order to find information of scientific interest, and to identify that information in real-time as the vehicle is searching. In particular, I have focused on how vehicles can navigate through environments where the traversability of the terrain and the scientific value of various patches of the environment are both defined stochastically. I have also done work on applying machine learning and data mining techniques to allow an AUV to, in real-time, determine when it has located a particular type of oceanic feature.

In 2012 I was funded by the Intelligent Systems Group at NASA's Ames Research Laboratory. In this work we were exploring how a rover should navigate through an environment with regions of stochastically determined traversability (i.e., there is some probability that a particular region can be driven across). Assume that, for purposes of safety, a rover can only drive through regions that it can eventually determine to be 100% safe. The rover can choose to *a prior*plan a path only through the areas that it knows in advance to be safe.

The first problem with this approach is that the known safe path may be substantially larger than a less certain path. A rover may want to get to the other side of a ravine without knowing for certain whether it can find a safe path into the ravine. On the other hand, going around the ravine could take a substantial amount of time, during which the rover will be performing less effective science. In this case, the cost savings if the rover can get through the ravine are substantial compared with the potential wasted time if it explores the rim of the ravine without finding a safe path through. Should the rover succeed, that cost savings can then be used to conduct more scientific exploration than it might have been able had it gone around.

The second problem is that it would be extremely unlikely for any environment to so well known ahead of time, especially one such as Mars. Thus, a robot needs the capability to evaluate options based on varying levels of uncertainty and choose the one that provides the best trade-off between *a prior*safety and cost savings. Navigation through this environment can then be distilled down to a series of choices. Every time a region of less than 100% safety falls within the rover's sensor horizon, it must decide whether the benefit of continuing to drive towards this region (e.g., reduced transit time or fuel usage) in order to try and gain more information about it outweighs the cost of determining that it cannot be traversed, thus having to divert and take a longer path around the region.

My dissertation research takes this concept and further applies it to robotic foraging. Broadly, the question is how a robot should plan a path that allows it to gather something. In my case the focus has been on gathering scientific information, such as a rover on the Moon or Mars would be doing. The question is difficult because simple maximization is NP-hard, so an optimal solution cannot be tractably found. This is especially problematic on a mobile system, where re-planning needs to take place in real-time.

The algorithms I developed to encompass both of these aspects of the navigation problem, collectively called the Information Foraging Algorithm (IFA), blend techniques from traditional robotic path planning, potential field methods, and optimal foraging theory. In a nutshell, the IFA first evaluates the areas it believes may contain interesting information based on a cost function that takes into account the expected information content as well as the time that will remain for exploration once the information has been harvested. It then uses a greedy algorithm to select the best node. Navigation to a node can utilize any lowest-cost navigation algorithm, such as A*, Djikstra's algorithm, or D*.

In navigating to a particular node, the IFA also takes into account environmental uncertainty, primarily about whether a particular area of the environment can be traversed by the vehicle. It weighs the probability that a region can be traversed and the time savings that would be realized by traversing it with the time that would be lost by exploring the edge of a region that ultimately proves to be untraversable.

In small environments where the optimal path can be tractably computed (generally less than 10 nodes), the IFA typically generates paths that collect better than 75% of the information the optimal path collects, but as a substantial computational savings. This computational savings allows the IFA to be applied to significantly larger domains (for example, 2500 nodes) with paths generated on the order of seconds, as opposed to the intractable amount of computation necessary to generate optimal paths in an environment of that size [**?**].

I have also done From 2011-2013 I was a collaborator with the Autonomy group at the Monterey Bay Aquarium Institute (MBARI). This group was focused on developing techniques for increasing both the autonomy with which their autonomous underwater vehicles (AUVs) can operate, the accuracy with which these vehicles can sample events that MBARI oceanographers are interested in, and allowing oceanographers without a programming background to easily interface with the AUV while it is in the ocean. The bulk of our collaboration focused on developing algorithms that allow the *Dorado* AUV to detect a particular ocean feature called a *front* [?]. In a nutshell, fronts represent boundaries between different types of water, such as the boundary between the Gulf Stream and the Atlantic. Oceanographers are interested in these boundaries because they are hotspots of marine life.

The scientific samples with the most utility to oceanographers are those are within the front, and immediately on either side of it. Many fronts can be identified from satellite imagery of sea surface temperature. However, this is only useful if the area of interest can be clearly imaged on the day of interest. For example, Monterey Bay is frequently covered by clouds and fog, eliminating the possibility of satellite imagery on those occasions. Further, even if conditions are conducive to high quality imagery, by the time the images are received, processed, and analyzed by oceanographers, and the AUV dispatched to the appropriate launch location, several hours have likely passed. Given that fronts have high temporal and spatial variability, there is a very small likelihood that the front is anywhere near where the satellite imagery indicated. Thus, in order to collect high-quality samples, the AUV still needs to have the ability to detect the presence of the front by itself.

3 Current Research

My current research has focused on issues of human-machine interaction where the machines are autonomous systems. As I have noted, developing mechanisms by which non-expert users can still interact with and retarget autonomous systems is only going to be more important as such systems become more ubiquitous. The nature of my current position is such that I am able to devote research energy to this area.

In particular, I am involved in three separate projects, each of which is examining how to improve human-machine interaction so as to supplement and enhance the ability of human-robot teams to successfully perform missions. One project is looking at how to reduce the overhead necessary for the operation of unmanned aerial vehicles (UAVs). Currently the air force requires two people to operate a UAV - a pilot who flies the vehicle and a sensor operator (SO) who controls the camera and watches the sensor feed. We have particularly focused on improving the automatic target recognition (ATR) and tracking system to aid the SO in identifying targets of interest, and especially discriminating among similar targets.

The concept I developed for this involves giving the SO a simple yet flexible interface to an ontology that stores information about prototypical instances of frequent target categories, such as person, boat, car, SUV, etc. The interface then allows the SO to either modify or add parameters that the ATR uses to identify the specific instances the SO is interested in. We also developed a concept for autonomously planning camera motions to scan an area the SO defines according to a pattern s/he defines. This reduces the work the SO has to do actually manipulating the sensor on the UAV.

The second project which I am involved in is developing a mechanism for evaluating the effectiveness of a human-robot team (HRT). Our particular focus is on how to represent a) the structure and information flow of the HRT, b) the tasks the HRT will be expected to perform, and c) the environment in which it will need to perform them. The goal is to give a test and evaluation officer the capability to rapidly assess whether a given autonomous system will provide enough of a benefit for it to be worth purchasing.

The final project is most directly related to human-machine interaction. In this project we are trying to build a system that can, in real-time, estimate an operator's cognitive load and take steps to reduce that load to more manageable levels. My contribution has been to introduce concepts of machine learning and signal processing to enable the real-time component of the system.

4 Future Research

In my future research I hope to get back to working more directly with autonomous systems. In particular, several future improvements to the IFA are planned. The two most significant involve changing how nodes are evaluated to take into account features of the environment near the node, such as information content of adjacent nodes or path-planning for multiple nodes at a time instead of just one. As the IFA is improved, I also want to integrate the event detection techniques developed as part of the my research at MBARI, as well as other information-theoretic methods, to allow robots to better judge how long a robot should forage a particular patch of information before moving on to the next patch.

I also would like to follow up the research I conducted with MBARI more explicitly. MBARI uses an executive system called TREX for controlling the *Dorado* AUV, both in the water and on-shore updating mission plans. One future thread of research is to determine how to use event detection algorithms such as the ones I developed at MBARI to control the behavior of a vehicle to maximize how it samples a given event or patch of interest.

Finally, the research at MBARI evolved into an unrealized project in which we were going to explore how to develop an autonomous system that could utilize multiple, heterogenous agents in a mixed-initiative planning system to target particular scientific goals within a broader research mission specification. In particular, this was going to involve developing a planning system that could construct its own goals to achieve. In effect, it was going to plan about how to plan using a meta-cognitive engine. I believe this capability is going to be at the core of any successful long-term autonomy system, and have a strong interest in pursuing this line of research further.