# Geoffrey Hollinger: Research Statement
# Decision Making & Learning for Robotics

Imagine a world where streams of sensor data are at your fingertips – a world where scientists, first responders, and safety inspectors can retrieve information about any physical location at the push of a button. Far beyond the current boundaries of the internet, such data could be used to assess nuclear and biological accidents, inspect aging facilities, and monitor everything from weather patterns to energy usage. Extrapolating predictions from this data would help to ensure sustainable ecosystems, develop better agricultural models, and potentially save lives through the prediction of disasters. Retrieving this vast quantity of spatially localized information would require a network of mobile sensors capable of operating cooperatively with humans to provide intelligent perception and action on a grand scale. The key problems that must be solved to develop such transformative technology lie at the intersection of physical planning and information optimization, making them problems of robotic decision making.

At the core of robotics is the optimization of *physical* plans for tasks such as grasping and navigation. In contrast, a pervasive notion in engineering and computer science is the idea of *information* optimization. **In my research, I seek to unify information optimization and physical motion planning to bridge the gap between the near-optimal performance possible in theory and the limited performance currently possible in the physical world.**

Data gathering in the physical world is by its nature a *big data* problem, where large streams of information must be acquired and managed. In contrast with many big data applications, such as internet search and database retrieval, information about the physical world must be collected by mobile robots operating over large spatio-temporal scales. In addition, there are constraints on retrieval; it is clearly not possible to access sensor data from the future, nor is it possible to access data on the other side of the world without moving a sensor to that location.

My work treats these challenging robotic information gathering problems within the framework of algorithmic optimization. In the general case, robotic information gathering requires solving the following maximization problem:

$$\mathcal{P}^* = \underset{\mathcal{P} \in \Psi}{\operatorname{argmax}}\ \mathcal{I}(\mathcal{P}) \text{ s.t. } c(\mathcal{P}) \leq B, \tag{1}$$

where $\Psi$ is the space of possible trajectories for a robot or team of robots, B is a budget (e.g., time, fuel, or energy), and $\mathcal{I}(\mathcal{P})$ is a function representing the *information* gathered along the trajectory $\mathcal{P}$. A number of factors make solving this optimization problem particularly difficult: (1) for nearly all interesting objective functions, finding the optimal solution is formally hard (NP-hard or PSPACE-hard), (2) the space of trajectories $\Psi$ grows exponentially in the horizon length and the number of robots, making exhaustive searches intractable, and (3) the exact form of the information function and costs may be unknown a priori, requiring learning and adaptive behavior.

Despite the challenges associated with optimizing motion planning for robotic information gathering, I have shown that formal properties of the objective function (e.g., convexity, submodularity, and monotonicity) along with an informed representation of the trajectory space (e.g., discrete lattices, trajectory trees, and model-predictive controllers) allow for efficient and, in many cases, provably near-optimal approximate solutions to these optimization problems. Such techniques are necessarily *scalable*, *distributed*, and *adaptive* for operation in the physical world.

My research lies at the intersection of algorithm design, machine learning, and motion planning. I have worked to combine methods from disparate areas and to build bridges between the robotics community and these disciplines. In addition, I have forged interdisciplinary alliances with experts in operations research, wireless communication, and environmental science. Throughout my work, I have utilized approaches from multiple communities and unified them into common frameworks.

The breadth of my research agenda spans theoretical analysis and systems development, and my research style is not only to develop decision making techniques and analyze their formal guarantees, but also to implement them on mobile robots operating in the physical world. I strive to develop robotic systems that operate in both natural and built environments, which requires cooperative interaction with humans through supervised autonomy and human-robot collaboration.

The research I have completed to date during my doctoral work at Carnegie Mellon University and my postdoctoral appointment at the University of Southern California has shown that efficient mobile data gathering methods are possible for problems such as autonomous search, underwater data collection, and robotic inspection. Ultimately, my future research agenda provides generalized solutions to high-impact problems (e.g., environmental monitoring, emergency response, and sustainable energy), benefiting society and improving the state of the art in robotics and related fields.

## Research to Date

### Efficient and Guaranteed Search

One focus of my research has been the development of improved techniques for robotic search and pursuit-evasion. A team of autonomous robots enters a real-world environment (e.g., a building, road network, or harbor) where one or more targets may exist. The searchers' goal is to locate the targets or definitively determine that there are no targets in the environment. Prior work typically treats search as either an average-case problem or a worst-case problem. The average-case problem, which I will refer to as *efficient search*, is to locate a non-adversarial target in the minimum expected time. The worst-case problem, or *guaranteed search*, is to find an adversarial target or to clear the environment if no such target exists.

With a focus on scalability and guarantees, my collaborators and I developed the first distributed approximation algorithm for the multi-robot efficient search problem [15]. By utilizing the property of submodularity, the formalization of diminishing returns, I was able to show guarantees on performance. The algorithm generates solutions competitive with general solvers while maintaining linear scalability in the number of robots. I implemented this approach on a state-of-the-art mobile manipulator at Intel Research Pittsburgh [20], and the work was recognized as a finalist for the Best Student Paper Award at the Robotics: Science and Systems Conference in 2008 (conference acceptance rate 24%) [14]. For the case of an adversarial target, I proposed the first "anytime" algorithm in that domain that quickly finds a feasible solution and then continues to generate improved solutions with increasing running time [9]. This technique has become a benchmark in the field, and researchers at Carnegie Mellon University, University of Pittsburgh, and Linköping University are currently extending its capabilities [18].

Figure 1: Pioneer robot searching a building with a team of volunteer firefighters. A graphical representation of the building is shown in the top left. The robot and humans coordinate using a provably near-optimal distributed algorithm that combines submodular optimization and pursuit-evasion [9, 10, 15, 20].

A key component of my work in search has been to unify the adversarial and non-adversarial search problems into a common framework. I derived the first algorithm to generate strategies that perform well under both assumptions, and I provided a successful implementation on mobile robots [10]. As part of my search and pursuit-evasion research thrust, I was the main organizer for a very successful workshop at ICRA 2010, and I served as a guest editor for a special issue of *Autonomous Robots*. Along with my collaborators, I also published the first survey paper that focuses on unifying the once disparate search and pursuit-evasion problems [1].

### Active Perception and Uncertainty Modeling

The search problem described above is a special case of the more general active inspection problem, where a robot plans its views to maximize information gathered about a phenomenon of interest. My work has analyzed the benefit of utilizing new information by quantifying adaptivity gaps when planning to reduce variance in a non-parametric Bayesian framework. I also showed that the property of adaptive submodularity can be applied to these active inspection problems to generate guarantees on performance. I successfully utilized these techniques to improve the efficiency of underwater ship hull inspection [11] and bathymetric mapping [12]. I have extended these techniques to provide novel uncertainty models that improve the safety and reliability of operation of autonomous underwater vehicles operating in ocean currents [5]. A field implementation of this work is now actively in use by the underwater gliders at USC. In addition, I am preparing a paper for the International Conference on Machine Learning (ICML) that applies similar active perception techniques to reduce uncertainty in general Bayesian regression models.

In complementary work, I showed that an extension to Gaussian Process modeling provides a method for estimating the uncertainty of a reconstructed 3D mesh, and this uncertainty can be used to guide inspection planning [7, 8]. A similar uncertainty modeling technique using probabilistic dimensionality reduction and Gaussian Processes can be utilized to reconstruct the locations of an *ad hoc* sensor network for use in localization. I have applied these techniques to a network of ranging radios and was able to reconstruct the locations of all nodes in the network using minimal prior knowledge. The reconstructed network was then used to track a moving agent [6]. This work, along with other work by my collaborators, led to successful spinoff projects with Boeing and Caterpillar on GPS-denied navigation using ranging sensors.

### Communication-Informed Planning

Past robotics research has often examined the problem of communication-constrained planning, where a simple model of communication (e.g., a fixed radius) is assumed, and path planning algorithms or control laws are derived to maintain network connectivity. Such methods are useful, but they do not account for many realistic factors, such as noise, fading, and packet error. My research utilizes models developed alongside communication theorists. To this end, I have proposed distributed data fusion techniques that allow multi-robot underwater search using acoustic communication [16, 17]. This work was honored as a finalist for the KUKA Service Robotics Best Paper Award in 2011, and I was approached by Intelligent Automation, Inc. to develop a field implementation of the work as part of the Navy SBIR program.

I also took the lead on a highly collaborative effort, funded by the Office of Naval Research Science of Autonomy program, that spans four institutions. The goal is to integrate planning and communication for optimizing the performance of robotic sensor networks. My collaborators and I showed that properly utilizing deterministic access and random access scheduling to optimize parameters in a novel path planning algorithm allows for efficient data gathering from underwater sensor networks [4].

In addition, we proposed estimation-theoretic metrics that outperform commonly used metrics such as mutual information and Fisher information [13]. These works were the first to integrate underwater path planning and acoustic communication at a fundamental level. I was heavily involved in the development and writing of a large collaborative grant proposal extending this research, which has been funded by the National Science Foundation and will allow the work to continue through 2015.

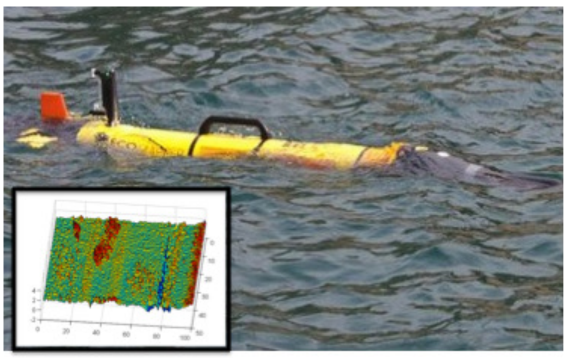

Figure 2: Autonomous underwater vehicle actively planning its dives to provide an accurate bathymetric map of a lakebed. The AUV utilizes a scalable learning and planning framework that generates informed trajectories to minimize uncertainty in the 3D reconstruction (shown in the bottom left) [4, 8, 12].

## Future Directions

I plan to continue research in decision making and learning that bridges the gap between theoretical analysis and real-world applications. By expanding the capabilities of robotic systems using principled methods, we can tame the complexity inherent in optimal planning. I propose to develop scalable and provable techniques that provide solutions to problems previously considered intractable. These techniques will also allow us to understand the best-case and worst-case behavior of systems acting in the physical world. I provide more detail below about two key research threads that are central to achieving this vision.

### Adaptivity and Learning in Robotics

The idea of acting adaptively is a powerful one. The very nature of Darwinian evolution is based on populations adapting for survival. When designing robotic systems, it is only natural to expect that adaptivity and learning can lead to improved performance. Adaptive decision making can take the form of learning a new problem representation *in situ* or utilizing a posterior distribution of a phenomena of interest. Unfortunately, acting adaptively does not come without cost. Not only can errors in adaptive planning potentially lead to system failures, but the necessary computing power for adaptive behavior must be placed on the vehicle or communicated from an external source. This additional complexity, both in computation and in system design, motivates the principled analysis of when adaptivity helps and when it does not.

My prior work has shown that acting adaptively can, in some cases, provide exponential improvements in performance for active perception tasks [11]. However, it is sometimes possible to show a more surprising result that, in other cases, the benefit of adaptivity is provably small or nonexistent. A key idea behind these results is that of a quantifiable *adaptivity gap* [8]. Adaptivity gaps provide a formalization of the benefit of adaptive planning relative to a particular problem domain and objective function.

I plan to utilize adaptivity gaps to develop systems capable of efficiently using limited computational resources to apply adaptive behavior when it is most beneficial. I have previously applied selectively adaptive planning to improve 3D reconstruction [7], and I am currently extending these techniques to mobile healthcare monitoring. In the future, I will develop a unified framework that extends recent advances in adaptivity analysis from machine learning [3] and stochastic optimization [2] to robotic decision making. These fundamental results will provide the necessary tools to utilize adaptive behavior effectively and will lead to adaptivity-aware learning for intelligent systems.

**Mobile Networked Systems**

Networked systems have already changed the world through the development mobile phone networks, wireless computing, and the internet. Such technology has unprecedented power for data collection and assimilation through the use of low-cost sensors included on networked devices. However, a significant limitation of current technology is that the nodes in the network are either stationary (e.g., a wired network in a building) or move independently (e.g., users carrying mobile phones). In contrast, we have barely scratched the surface of the potential for mobile networked systems that are connected to mobile bases (e.g., cars, helicopters, ships, and submarines) and are *actively controlled*. Such networks are multi-robot systems, in that they must perceive with their sensors and act intelligently on this information as part of a human-robot team.

As part of my research agenda, I will study the efficient optimization of actively controlled mobile sensor networks. I will work towards unifying communication theory and robotic motion planning at a fundamental level, which will allow existing theoretical results to be applied in new contexts and will reveal new avenues for theory and applications at the crossover of planning and communication. In my recent work, I have integrated probabilistic planning methods with realistic communication models [4], and I have developed estimation-theoretic performance metrics for robotic sensor networks [5].

I will expand these research threads by analyzing the formal hardness of estimation-theoretic objective functions based on squared error distortion [19], which will facilitate the development of novel motion planning algorithms. In addition, I will work with communication theorists to derive novel communication models that can by readily integrated into multi-robot coordination algorithms (e.g., implicit coordination, market-based optimization, and multi-agent TSP). This research will move towards the realization of active mobile networks that work with human operators to gather data from the physical world.

**Impact and Funding**

Improving decision making and learning for robotic systems has great potential to benefit society and the scientific community. Mobile robotic sensors can assist with tasks ranging from improving the energy efficiency of buildings to monitoring world climate change. Autonomous platforms can provide information about seismic activity, ecological events, and defense operations. Robotic assistance for emergency response and urban search and rescue can even save lives. My research agenda tackles a variety of problems in robotic decision making, maximizing societal impact and creating research opportunities for students at all levels.

My research is in accord with the agendas of many funding agencies, including the Office of Naval Research (ONR), the Defense Advanced Research Projects Agency (DARPA), and the National Science Foundation (NSF). With significant interest in autonomy from the Department of Defense and the development of the multi-agency National Robotics Initiative, there are increasing opportunities to acquire funding in robotics.

My dedication to interdisciplinary research enables my participation in many collaborative programs, such as Multidisciplinary University Research Initiatives (MURI) and Collaborative Technology Alliances (CTA). I have previously taken leadership roles in such projects at both Carnegie Mellon University and the University of Southern California, and I have assisted in the development of successful grant proposals for large multi-investigator projects. In addition to building interdisciplinary academic alliances, I intend to collaborate with industrial research labs, like Boeing and Google, which are becoming increasingly interested in autonomous systems. I am passionate about bringing together the research areas described above and becoming a leader in the development of robotic systems that sense and act intelligently in the physical world.

# References


[1] T. Chung, G. Hollinger, and V. Isler. Search and pursuit-evasion in mobile robotics: A survey. *Autonomous Robots*, 31(4):299–316, 2011.

[2] B. Dean, M. Goemans, and J. Vondrak. Approximating the stochastic knapsack: The benefit of adaptivity. *Mathematics of Operations Research*, 33:945–964, 2008.

[3] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artificial Intelligence Research*, 42:427–486, 2011.

[4] G. Hollinger, S. Choudhary, P. Qarabaqi, C. Murphy, U. Mitra, G. Sukhatme, M. Stojanovic, H. Singh, and F. Hover. Underwater data collection using robotic sensor networks. *IEEE J. Selected Areas in Communications*, 30(5):899–911, 2012.

[5] G. Hollinger, C. Choudhuri, U. Mitra, and G. Sukhatme. Distortion metrics for robotic sensor networks. In *Proc. Int. Conf. Robotics and Automation*, 2013. Under review.

[6] G. Hollinger, J. Djugash, and S. Singh. Target tracking without line of sight using range from radio. *Autonomous Robots*, 32(1):1–14, 2012.

[7] G. Hollinger, B. Englot, F. Hover, U. Mitra, and G. Sukhatme. Uncertainty-driven view planning for underwater inspection. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 4884–4891, 2012.

[8] G. Hollinger, B. Englot, F. Hover, U. Mitra, and G. Sukhatme. Active planning for underwater inspection and the benefit of adaptivity. *Int. J. Robotics Research*, 32(1):3–18, 2013.

[9] G. Hollinger, A. Kehagias, and S. Singh. GSST: Anytime guaranteed search. *Autonomous Robots*, 29(1):99–118, 2010.

[10] G. Hollinger, A. Kehagias, and S. Singh. Improving the efficiency of clearing with multi-agent teams. *Int. J. Robotics Research*, 29(8):1088–1105, 2010.

[11] G. Hollinger, U. Mitra, and G. Sukhatme. Active classification: Theory and application to underwater inspection. In *Proc. Int. Symp. Robotics Research*, 2011.

[12] G. Hollinger, U. Mitra, and G. Sukhatme. Active and adaptive dive planning for dense bathymetric mapping. In *Proc. Int. Symp. Experimental Robotics*, 2012.

[13] G. Hollinger, A. Pereira, and G. Sukhatme. Learning uncertainty models for reliable operation of autonomous underwater vehicles. In *Proc. Int. Conf. Robotics and Automation*, 2013. Under review.

[14] G. Hollinger and S. Singh. Proofs and experiments in scalable, near-optimal search with multiple robots. In *Proc. Robotics: Science and Systems Conf.*, 2008.

[15] G. Hollinger, S. Singh, J. Djugash, and A. Kehagias. Efficient multi-robot search for a moving target. *Int. J. Robotics Research*, 28(2):201–219, 2009.

[16] G. Hollinger, S. Yerramalli, S. Singh, U. Mitra, and G. Sukhatme. Distributed coordination and data fusion for underwater search. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 349–355, 2011.

[17] G. Hollinger, S. Yerramalli, S. Singh, U. Mitra, and G. Sukhatme. Distributed coordination and data fusion for multi-robot search. *IEEE Trans. Robotics*, 2013. Under review.

[18] A. Kleiner, A. Kolling, M. Lewis, and K. Sycara. Hierarchical visibility for guaranteed search in large-scale outdoor terrain. *J. Autonomous Agents and Multi-Agent Systems*, 2011. DOI: 10.1007/s10458-011-9180-7.

[19] N. Merhav and S. Shamai. Information rates subject to state masking. *IEEE Trans. Information Theory*, 53(6):2254–2261, 2007.

[20] S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and J.M. Vande Weghe. HERB: A home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.