# Research Statement

**Volkan Isler** isler@cs.umn.edu

My research field is robotics. Most of my group's efforts focus on problems that arise when using robots as mobile sensors charged with observing phenomena. Such problems appear in numerous settings ranging from surveillance and search-and-rescue to environmental monitoring and agricultural automation. In terms of style, my contributions can be grouped into two major categories. (1) *Algorithmic foundations:* Most of these results are in the form of papers containing a theorem and its proof. (2) *Field systems:* These results demonstrate a capability in the field. Examples include building a solar map using robots equipped with solar panels and the semantic reconstruction of a farm with cameras mounted on aerial vehicles. A distinct feature of my research lab is that we are comfortable in both domains and have validated foundational results in field experiments numerous times.

I will start with an overview of the field systems we have been developing as part of ongoing projects in my lab. Afterward, I will focus on two sets of foundational problems which arise in these applications[1].

## 1 Field Systems

In this section, I describe our work on environmental monitoring followed by our more recent work on agricultural robotics.

### 1.1 Monitoring Radio-Tagged Carp with Autonomous Surface Vehicles

The biggest project in my lab is on finding and tracking radio-tagged invasive fish (specifically carp) in inland lakes. Since 2011, my group has been working on building autonomous vehicles to automate this environmental monitoring task. We have three focus areas.

(1) Search: The range of the radio-tags (which are surgically implanted into the fish) can be rather short. Therefore, the first task is to establish contact with the radio signal. In one extreme, we can ignore the fish motion and assume that they are mostly stationary. The resulting search problem can be reduced to coverage [26]. Of course, the target motion cannot be ignored in most search applications. A simple model for the motion of the fish is the random-walk. Surprisingly, this fundamental problem (how to find a random-walker) is open in very basic settings. We solved the problem for the one-dimensional case [13] (imagine a river or shore-line) and have been working on the two-dimensional case [20]. In the other extreme, we can model fish motion as an adversarial process actively trying to avoid capture. I have been working in this domain for more than a decade and will revisit it in Section 2.1.

(2) Active localization: Once we hear the signal, we can rotate the directional antenna on board to obtain a coarse bearing measurement. Multiple measurements can be merged to accurately localize the target. Imagine that a single boat obtains the first measurement. Where should it take the next set of measurements so as to accurately and quickly localize the target [29, 33]? What if there are multiple robots and communication constraints [34]? I will further discuss our work on some of these active localization problems in Section 2.2.

(3) Energy: If a robotic system will collect data for long periods of time, its energy limitations must be addressed. We took a two-pronged approach. At a basic level, we studied energy-efficient navigation as a subroutine for the two tasks above [28]. However, we quickly realized that energy savings simply by optimizing velocity profiles will not be enough. Therefore, we studied solar energy harvesting as an alternative, and developed methods to build spatio-temporal solar maps and ways of incorporating harvesting into navigation [16–18].

In addition to these three sub-problems, we studied problems such as online navigation with a sonar [19] and initializing multi-target tracking [32] that came up during our field work. As we work on these problems, we developed considerable expertise in field robotics. We have now started doing preliminary work on the next generation of this project where we will find and record the behavior of (tagged and untagged) animals using aerial vehicles [6]. The main goal is to first perform high-level surveillance with a fixed wing airplane to obtain whereabouts of moose, bears and similar animals. Then we would like to obtain close-up video footage using multi-rotor UAVs.

[1] Intermediate sections are mostly self-contained and can be read (or skipped :) in any order.

### 1.2 Agricultural Robotics

These days, most of our effort is dedicated to agricultural robotics projects. The main project in my lab in this domain is on surveying specialty farms such as apple orchards. Specifically, we work on collecting yield-related parameters such as the total number of apples, number of apples per each tree and apple diameters [22] using footage from aerial and/or ground vehicles. We now have promising results which we will evaluate this summer with systematic field trials.

As part of this work, we have had to revisit fundamental vision algorithms since many assumptions they rely on are invalid in the agriculture domain: Orchard footage is filled with occlusions and specularities. Standard point features such as SIFT are not always reliable in orchard settings because they come from occlusion boundaries (e.g. one leaf over another) and look drastically different across views. Without matching features, standard motion estimation algorithms fail which makes it hard to merge views or maintain accurate counts. To solve this problem, we developed a technique where we use higher-level objects directly as features to compute the camera motion [21]. Specifically, we used apple contours as image features and developed a bundle-adjustment technique to compute the camera motion to align them. We are also spending considerable effort to speed up our algorithms using GPU to achieve real-time performance. In the next phase of this project, we will work on developing *active sensing* algorithms where we actively control the camera viewpoint to circumvent image quality issues such as specularities.

As part of a smaller project, we have been working on collaboration mechanisms between an aerial and a ground system for precision agriculture applications in maize fields. In particular, we have been looking into developing ways for the aerial robot to help ground robot navigation by identifying rows from aerial footage [8]. We also developed an adaptive sampling mechanism which collects aerial and ground-based measurements. A novel aspect of the system is the capability of the ground system to mule the aerial system so as to maximize its battery life [30]: if two locations to be sampled are far apart, the aerial vehicle can simply land on the ground robot and travel for free!

I am excited about agriculture as an application domain and will be talking more about it in Section 3.

## 2 Algorithmic Foundations

In this section, I will present an overview of our results on two sets of fundamental robotics problems.

### 2.1 Search and Pursuit-Evasion

The lion and man game is a classical game in which a lion tries to capture a man who tries to avoid being captured. The players are assumed to have the same speed. The original version of this game, that takes place in a circular arena, was studied as a recreational math problem in the 1950s. Many robotics applications such as target tracking, surveillance, search can be formulated variants of the lion and man game. Of course, the original formulation is not general enough to capture real-life scenarios. I have been working on solving robotics related versions of the lion and man game.

These contributions, which I summarize below, establish a pretty comprehensive theory of the lion and man game for robotics.

*General environments:* What happens if the environment is more complex than a circular arena? Our contributions include showing that a single lion can capture the man in any simply-connected polygon [7] and that three lions can capture the man in any polygon [4]. More recently, we have been working on solving the game in higher-dimensions and showed that a single lion can win the game on a convex terrain [12]. Three lions can win the game on any three-dimensional surface homeomorphic to a disk (possibly with holes) [11].

*Sensing constraints:* In the original version of the game the players can see each other at all times. What if the lion has sensing limitations? For example, what if it can see the man only if he is in his line-of-sight? We have shown that a single lion can still capture the man but this might take exponential time (in the number of vertices). Two lions can capture the man quickly. A lion with visibility constraints can win the game in a monotone polygon [10] and in strictly sweepable polygons [2]. In these formulations, the lion can observe the exact position of the man when he is in its line of sight. However, most sensors such as cameras can only give the bearing (direction). Recently, we showed that [31] if there is uncertainty in bearing measurements, there are simple circular environments in which the man can escape forever!

*Mobility constraints:* Most mobile robot systems are differential drives which can not move sideways. Solving the lion and man game under such mobility constraints is challenging. These problems are usually formulated as differential games and solved using dynamic-programming like numerical techniques for a given environment. We have made recent progress and obtained a much general result: a differential drive lion can get within a step-size of the man in convex environments [23].

In order to broadly disseminate these results, I have also contributed to a survey paper [5] and more recently, a "toolkit" paper which gives an overview of tools and techniques for solving variants of the lion and man game [9].

*Non-adversarial target models:* There are applications in which it makes sense to use a non-adversarial evader model. For example, in tracking animals, it might make more sense to model the targets as a stochastic process. In search and rescue, both parties can actively try to meet. I have also contributed to to solving such games. In [3], we present what is currently the most efficient strategy for rendezvous search on the line. In [14], we solve the rendezvous problem in two-dimensional settings. As detailed in Section 1.1, we have also made significant progress on the problem of finding a random-walker [13, 20].

## 2.2 Sensor Planning

Most sensors used in robotics do not directly measure the position of a target. Cameras, microphone-arrays, directional antennas all give an estimate of the target's bearing. To estimate the target's position, at least two distinct bearing measurements must be obtained. The uncertainty in the estimation is a function of the target-sensor geometry: suppose the two measurements were obtained at locations $s_1$ and $s_2$ and the true location of the target is $x$. The uncertainty is proportional to the distances $d_1 = d(s_1, x)$ and $d_2 = d(s_2, x)$. But it is also a function of the angle $\alpha = \angle s_1 x s_2$. In general, $\pi/2$ is a good angle whereas 0 and $\pi$ result in large uncertainty. A commonly used metric to express the uncertainty is $d_1 d_2 / \sin\alpha$. The dependency of the estimate on the angle makes associated sensor planning problems very challenging, and a rich source of interesting sensor planning problems.

First, consider the following basic sensor placement problem: we are given a set of candidate target locations. We would like to place a minimum number of sensors so as to guarantee that for every possible target location, there are at least two "good" sensors to observe it. That is, the uncertainty given by $d_1 d_2 / \sin\alpha$ is below a given threshold. Recently, we studied a similar problem for a robotic sensor. The goal is to compute a tour and measurement locations along the tour to ensure simultaneously that (i) the total time to traverse the tour and collect measurements is minimized, and (ii) the uncertainty objective is achieved (for every candidate target location, we have taken at least two good measurements). For both problems, we have obtained polynomial-time constant-factor approximation algorithms. The mobile version has been implemented and tested in the context of the "carp tracking" application described above [1, 25]. Earlier in [24], we studied a similar data gathering problem for downloading data from stationary sensors. The novelty of the formulation is a "two ring" communication model which captures the stochastic nature of download time as a function.

We have also worked on online versions of these problems where measurements obtained during data gathering are incorporated in trajectory planning. These solutions might be more effective in field settings. However, obtaining theoretical guarantees is harder. Nevertheless, we were able to obtain algorithms for the problem of localizing a stationary target with one or more robots collecting bearing measurements. The algorithms come with provable performance guarantees and have been field-tested. For multiple robots, we also incorporate communication constraints [33, 34].

*These days, I am interested in solving sensor planning problems for aerial vehicles.* We have studied the problem of tracking multiple targets with multiple aerial robots [27]. My group is now working on localizing a radio source with an aerial vehicle for habitat monitoring applications [6]. Another result, which we are preparing for submission is on how to quickly visit a given set of cones with an aerial vehicle. The angle of the cones correspond to the field of view of a camera on the vehicle. The height is associated with the desired level of resolution.

# 3 Looking Ahead

Many of us in the field of robotics work toward taking robots out of structured environments such as factory floors. We would like them to work effectively in dynamic, unstructured, large and complex environments. In my opinion, agricultural automation is an ideal next frontier toward this goal. Orchards do have some structure: there are rows of trees, trees have branches, leaves and fruit. However, operating in these environments is challenging: we need robust systems that can navigate on rough terrain. Yet, they need to be precise enough to observe diseases, spray pesticides, pick weeds and harvest fruit. Farms can be hard to access for a human operator (the width of a corn row can be less than 20 inches whereas the plans can get up to 12 feet tall.) The amount of data to be collected can be enormous with no communication infrastructure to quickly upload the data to a server. The scale of these environments necessitate efficient algorithms.

One of my goals for the next couple of years is to develop a system which can semantically reconstruct a farm both spatially and temporally [15]. Semantics would include object level information (e.g. apples). At a finer level, diseases, bruises or plant parts such tillers and spikes can be mapped and tracked over time.

More broadly, environmental monitoring with robots is an exciting application domain with potential for huge positive societal impact. In order to monitor endangered species and their habitats we need fundamental new knowledge addressing energy and communication limitations. In the next chapter of my career, I plan to maintain focus on these application domains and address the following technical questions:

(1) Active Perception: Most of my work revolves around controlling robots so as to achieve a sensing objective. As described above, we have developed a pretty comprehensive set of results for fundamental robot capabilities such as search and coverage. However, the sensing models used in these results do not always capture some of the challenges in more sophisticated sensing tasks. For example, can we plan a UAV trajectory to ensure that all apples have been covered up to a desired resolution and without specularities? I would like to close this gap between sensing and planning and develop realistic active perception algorithms with provable performance guarantees.

(2) Real-time Big Data: The amount of data collected from sensors such as cameras and LIDARs can be overwhelming. For many robotics applications, the idea of shipping all of this data to a server for processing is infeasible. There many not be a sufficient communication infrastructure or the latency requirements of the application can be very strict. I believe we need to perform in-network processing of the sensor data to circumvent these challenges. In addition to processing sensor data, the algorithms must be informed by communication constraints.

(3) Heterogeneous systems: These days, we are fortunate to have access to numerous commercial-grade robotics platforms. The capabilities and variety of these systems are rapidly improving. Still, there is no "universal robot" which is ideal for all tasks. Multi-rotor aerial vehicles can navigate in tight spaces but they have battery limitations. Fixed-wing aerial vehicles can travel long distances but they are designed primarily for open spaces. These systems have severe payload limitations whereas ground robots can carry large sensors and manipulators. I am interested in exploiting the synergies in a heterogeneous system of robots. We have been exploring air and ground collaboration in the agriculture domain. We are also starting a new collaboration with UMN's Large Lakes Observatory in Duluth to track plumes using surface and underwater vehicles. These applications only scratch the surface of the potential of heterogeneous systems. I plan to place increasing emphasis on this domain.

This is truly an exciting time for robotics research. We have the privilege of working on intellectually stimulating problems whose solutions can be used in a wide range of applications. Devices which can sense and interact with the environment will soon drastically change our lives. As part of this revolution, my focus will be on developing robotic systems which can perpetually collect data and make sense of it in applications such as agricultural and environmental monitoring as well as search and rescue and disaster response. I will continue to push the frontiers of theoretical knowledge and system capabilities toward this goal.

## References

[1] H. Bayram, J. Vander Hook, and V. Isler. Gathering bearing data for target localization. *IEEE Robotics and Automation Letters*, 1(1):369–374, 2016.

[2] L. Berry, A. Beveridge, J. Butterfield, V. Isler, Z. Keller, A. Shine, and J. Wang. Line-of-sight pursuit in strictly sweepable polygons. *arXiv preprint arXiv:1508.07603*, 2015.

[3] A. Beveridge, V. Isler, and D. Ozsoyeller. Symmetric rendezvous search on the line with an unknown initial distance. *IEEE Transactions on Robotics*, 29(6), 2013.

[4] D. Bhadauria, K. Klein, V. Isler, and S. Suri. Capturing an evader in polygonal environments with obstacles: The full visibility case. *The International Journal of Robotics Research*, page 0278364912452894, 2012.

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

[6] N. S. H. Bayram, K. Doddapaneni and V. Isler. Active localization of vhf collared animals with aerial robots. In *12th Conference on Automation Science and Engineering (CASE)*, 2016. to appear.

[7] V. Isler, S. Kannan, and S. Khanna. Randomized pursuit-evasion in a polygonal environment. *Robotics, IEEE Transactions on*, 21(5):875–884, 2005.

[8] Z. Li and V. Isler. Large scale image mosaic construction for agricultural applications. *IEEE Robotics and Automation Letters*, 1(1):295–302, 2016.

[9] N. Noori, A. Beveridge, and V. Isler. A pursuit-evasion toolkit. Technical report, Technical Report TR-15-013, University of Minnesota, Computer Science & Engineering Department, 2015. To appear in IEEE Robotics and Automation Magazine.

[10] N. Noori and V. Isler. Lion and man with visibility in monotone polygons. *The International Journal of Robotics Research*, page 0278364913498291, 2013.

[11] N. Noori and V. Isler. The lion and man game on polyhedral surfaces with boundary. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 1769–1774. IEEE, 2014.

[12] N. Noori and V. Isler. The lion and man game on convex terrains. In *Algorithmic Foundations of Robotics XI*, pages 443–460. Springer International Publishing, 2015.

[13] N. Noori, A. Renzaglia, J. Vander Hook, and V. Isler. Constrained probabilistic search for a one-dimensional random walker. *IEEE Transactions on Robotics (T-RO)*, 2015.

[14] D. Ozsoyeller, V. Isler, and A. Beveridge. Symmetric rendezvous in planar environments with and without obstacles. In *AAAI*, 2012.

[15] C. Peng, P. Roy, J. Luby, and V. Isler. Semantic mapping of orchards. In *5th IFAC Conference on Sensing, Control and Automation for Agriculture*, 2016. to appear.

[16] P. A. Plonski and V. Isler. A competitive online algorithm for exploring a solar map. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 5766–5771. IEEE, 2014.

[17] P. A. Plonski, P. Tokekar, and V. Isler. Energy-efficient path planning for solar-powered mobile robots*. *Journal of Field Robotics*, 30(4):583–601, 2013.

[18] P. A. Plonski, J. Vander Hook, and V. Isler. Environment and solar map construction for solar-powered mobile systems. *IEEE Transactions on Robotics*, 32(1):70–82, 2016.

[19] P. A. Plonski, J. Vander Hook, C. Peng, N. Noori, and V. Isler. Navigation around an unknown obstacle for autonomous surface vehicles using a forward-facing sonar. In *IEEE Search, Safety and Rescue Robotics*, 2015.

[20] A. Renzaglia, N. Noori, and V. Isler. The role of target modeling in designing search strategies. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 4260–4265. IEEE, 2014.

[21] P. Roy and V. Isler. Surveying apple orchards with a monocular vision system. In *12th Conference on Automation Science and Engineering (CASE)*, 2016. to appear.

[22] P. Roy and V. Isler. Vision-based apple counting and yield estimation. In *International Symposium on Experimental Robotics (ISER)*, 2016. to appear.

[23] U. Ruiz and V. Isler. Capturing an omnidirectional evader in convex environments using a differential drive robot. *IEEE Robotics and Automation Letters*, 2016.

[24] O. Tekdas, D. Bhadauria, and V. Isler. Efficient data collection from wireless nodes under the two-ring communication model. *The International Journal of Robotics Research*, page 0278364912439429, 2012.

[25] O. Tekdas and V. Isler. Sensor placement for triangulation-based localization. *Automation Science and Engineering, IEEE Transactions on*, 7(3):681–685, 2010.

[26] P. Tokekar, E. Branson, J. Vander Hook, and V. Isler. Tracking aquatic invaders: Autonomous robots for monitoring invasive fish. *Robotics & Automation Magazine, IEEE*, 20(3):33–41, 2013.

[27] P. Tokekar, V. Isler, and A. Franchi. Multi-target visual tracking with aerial robots. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 3067–3072. IEEE, 2014.

[28] P. Tokekar, N. Karnad, and V. Isler. Energy-optimal trajectory planning for car-like robots. *Autonomous Robots*, 37(3):279–300, 2014.

[29] P. Tokekar, J. Vander Hook, and V. Isler. Active target localization for bearing based robotic telemetry. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 488–493. IEEE, 2011.

[30] P. Tokekar, J. Vander Hook, D. Mulla, and V. Isler. Sensor planning for a symbiotic uav and ugv system for precision agriculture. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 5321–5326. IEEE, 2013.

[31] J. Vander Hook and V. Isler. Pursuit and evasion with uncertain bearing measurements. In *CCCG*, 2014.

[32] J. Vander Hook, P. Tokekar, E. Branson, P. G. Bajer, P. W. Sorensen, and V. Isler. Local-search strategy for active localization of multiple invasive fish. In *Experimental Robotics*, pages 859–873. Springer International Publishing, 2013.

[33] J. Vander Hook, P. Tokekar, and V. Isler. Cautious greedy strategy for bearing-only active localization: Analysis and field experiments. *Journal of Field Robotics*, 31(2):296–318, 2014.

[34] J. Vander Hook, P. Tokekar, and V. Isler. Algorithms for cooperative active localization of static targets with mobile bearing sensors under communication constraints. *Robotics, IEEE Transactions on*, 31(4):864–876, 2015.