NSF RI Small: Swarms that 'hear the shape of the drum'

Final Report

Principal Investigator: Herbert G. Tanner

Department of Mechanical Engineering
University of Delaware
130 Academy Street, 126 Spencer Laboratory
Newark, Delaware 19716-3140

Duration	Funding
3 Years	$450,000

This material is based upon work supported by
the National Science Foundation under Grant No. 0913015.
Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Executive Summary

What & why This project aimed at developing the algorithms and technology that would support the function a coordinated network of mobile sensor nodes tasked with collecting data and autonomously identifying prespecified patterns within this body of data. Potential application domains for such systems were envisioned within surveillance and reconnaissance, search and rescue, underwater acheology, and planetary exploration, among others. The pursuit of the specific research goals identified as necessary to realize small-scale prototypes of such a system was expected to produce new scientific insights and methodological tools that would help us understand better how robotic sensor platform motion, information flow, data management and processing, and pattern recognition interplay with each other in a distributed actuation and computation environment.

The system The envisioned system consisted of a number of small sensor-bearing mobile robots that were able to exchange information with each other through local wireless network connections. These sensor nodes could move and autonomously fall into specific vehicle formations that enabled them to cover fully a given area of interest in a coordinated fashion. Doing so, each one of these nodes would end up collecting a body of sensor data which offers a partial view of the data landscape that the whole sensor ensample was supposed to cover. The idea was that instead of the individual sensors simply dumping their data into a central network node for off-line processing, they would be able to do this processing in a distributed way, based only on their own data and those residing at their closest network neighbors. Thus, without any single network node having to have the whole data picture, they would *collectively* arrive at some conclusion as to whether there is a pattern in the data they have collected. That pattern could stretch over many different individual sensor coverage regions, and thus individual sensors could not by themselves "see" and identify the pattern; they could have captured only bits and pieces of it. It would then be the information exchange between the sensors that would enable them all to "see" through each other's sensor "eyes." The patterns that the swarm of sensors looked for were pre-stored in a form of distributed memory, and the data collected by the whole group would enable the recollection of some particular memory out of this family of pre-stored patterns. Having implemented both information collection, as well as processing, in a distributed manner promised robustness with respect to sensor noise as well as platform failures.

The technical challenges For such a system to be realized, we had to solve a number of interconnected technical problems. First we had to develop algorithms that would guarantee that groups of mobile nodes can fall into tight formations with predefined shapes, and move coherently as a single rigid body without colliding with each other and without leaving gaps between their sensor coverage areas. Succeeding in doing so would ensure that we could collage individual sensor data to bring about the overall picture of the area covered. Forming specific vehicle formation shapes that could float in space along certain paths in a provably convergent manner and without prespecifying individual sensor node reference trajectories, had been an open problem. These sensor nodes would not need that whole picture in order to "see" whatever pattern was hidden in their data. These patterns could be encoded in the eigenvalues of a matrix associated with the discretized version of a certain (Laplacian) partial differential operator. If the eigenvalues of such a matrix, built based on local sensor data, could be evaluated in a distributed way, then they could be matched to some reference values corresponding to memorized patterns. In the above statement are three smaller technical challenges: (a) construct the mathematical representation of the sensor data in the form of this Laplacian matrix without streaming all data to the same network sink node, (b) compute the eigenvalues of this big matrix in a decentralized way so that computation burden and time is distributed over the network of sensor nodes, none of which was expected to have access to significant computational resources, and (c) have the network nodes reach consensus over the global pattern they perceive to have captured, through local information exchange and without any single node holding all the identifying information. In isolation, solutions to these three subproblems have been reported in literature, but it was made clear that these existing methodologies were not directly compatible, and non-trivial technical advancements were needed to bridge them into a functional whole. These technical advancements were realized through this research activity.

How we tested it With the resources provided through this award, a proof-of-concept small-scale physical testbed was developed. The testbed consisted of wheeled mobile robots carrying down-facing cameras, which when moved in a tight formation could take successive pictures of the floor on which they were driving on, and collectivily identify large-scale shapes that were drawn on the floor. The system was shown to be reasonably effective in identifying such shapes in a distributed manner. In an effort to assess the potential of this technology for underwater surveys and archeology, and in collaboration with a colleague specializing in robotic oceanography, we tested the identification and pattern recognition algorithms against very challenging underwater sonar and visual imagery. Transitioning from relatively high-resolution land-based imagery to low-resolution underwater data revealed fragility in the machine learning methodological components that had originally been proposed and integrated, and together with other insights outlined in the following paragraph, motivated a new approach that eventually brought fruit. The new appoach lead to a considerably more robust identification and recognition system which performed admirably against a large body of grainy, noisy, low-resolution underwater imagery recording distributions of north atlantic scallop, collected by field-deployed autonomous underwater vehicles operated by our oceanographer collaborator.

What we learned "Stealing" a quote by Isaac Asimov frequently used by our colleague in oceanography,

> The most exciting phrase to hear in science, the one that heralds new discoveries, is not 'Eureka!', but 'That's funny...'

In this spirit, one of the unanticipated insights we obtained was that, despite all the hype and emphasis on distributed sensing and control, decentralization brings benefits *under the right conditions.* What is sometimes overlooked in the context of distributed control and sensor networks is the cost of passing information around. This cost is not so much in terms of power and bandwidth—which have been at the center of a significant body of recent relevant research; rather it is understood here in terms of time lag and delays which not entirely attributed to bandwidth but also to computation. While we have demonstrated that integrated distributed sensory data processing and decentralized, spatially distributed, collective pattern recognition is feasible, depending on the particular computation/communication budget and risk tolerance, it may not always be preferable to centralized system architectures in terms of performance. Even if performance is not the issue, and robustness takes central role, reliance on nonredundant network links to propagate information around can create multiple possible points of failure that counterbalance the robustness benefits of decentralization. The bottom-line is that at the end of the day, the benefits of decentralization should be weighted against the cost and risk of passing information around. These tradeoffs are expected to change and shift over time with advances in local computation and networking, and thus the right balance has to be re-evaluated and recomputed from time to time. Another realization that came about from our involvement with processing of underwater imagery is that 1. a significant portion of available machine learning and pattern recognition methods rely more heavily than they originally seem on the quality of the data they operate on, and 2. several machine learning techniques which have proven reasonably robust in some contexts, may unexpectedly fail when tried in new domains and a under different set of conditions. This fragility highlights that, besides being cautious even when bringing our best of tools to bear, we still have to understand better the fundamental mechanisms of making decisions in the presense of noise and uncertainty, and that we may have to experiment with new ideas and approaches.

The path forward These basic questions are still open. This research activity delivered on the plan that was set forth in the grant proposal, even via alternative pathways and plan adjustments, but the issues that surfaced in the process of carrying out this research reach beyond the scope of this project and present challenges to be investigated in the future. The very nature of research is such that it is expected to give answers leading to new questions: what exactly are the tradeoffs between decentralization, risk, and performance, under a given set of computational, sensing, and networking parameters? What aspects, features, and attributes should we identify with decision making algorithms that are more likely to perform robustly over a wide range of application domains and based on data of variable degree of fidelity and resolution? In addition to showing us that coordinated swarms of robotic sensor nodes can identify autonomously patterns in their collected data, either by trying to "hear the shape of their drum" or otherwise, this research activity hopefully positions us better to go after some of the fundamental questions identified above.

Acknowledgements

This work has been made possible through an NSF award with number 0822845. The work also leveraged NSF award 00447898. The PI and the students who worked in this project extend special thanks to Art Trembanis and Justin Walker from University of Delaware's Oceanography for close collaboration on topics of underwater imagery and environmental monitoring. Thanks are also extended to Scott Gallager, Amber York, and Dvora Hart of Woods Hole Oceanographic Institution for sharing Habcam scallop image datasets.

1 Project Objectives and Background

1.1 Major goals

The goal of this project is to develop a methodology for automated, collective, and spatially distributed sensing and classification, through the use of networks of interconnected mobile sensor platforms. This robotic sensor network is envisioned to map distributions of spatial physical quantities and discover patterns in the data operating as a distributed associative memory, without any single robot having to collect and process the whole body of data.

The technical hypothesis in this research project was that networked robotic sensors can implement in a distributed fashion a series of algorithms that would enable them to exhibit a new form of collective intelligence. Specifically, the behavior envisioned was the distributed identification of patterns in collectively harvested sensor data, obtained through coordinated motion control and observation strategies. This project therefore aimed at showing that it is possible to realize such systems with current technology, assuming that particular mathematical and computational innovations were introduced on the algorithmic side. The underlying technical assumption was that parellelizing computation, processing, and decision making, and supporting them with new robotic motion control architectures was the key to realization of such systems.

To realize this goal, the project defined three main objectives:

1. Harvest (collect) information in a by means of local measurements, and locally store the data produced.
2. Extract the features of the global sensory data map constructed from the distributed measurements.
3. Match the features produced to a single feature set out of a pre-coded library, and thus identify the pattern in the sensory data collected.

It was recognized that any partial existing solutions to the technical problems underlying these objectives would not readily interface with each other to bring forth the envisioned system. As an example, strategies for collection and storage of sensory data need not necessarily be compatible with processing algorithms, and formation control and area coverage methods have to be coupled and be consistent with any constraints imposed by the need to maintain some communication network topology. In addition, the plan for pattern identification was conceived within the context of neural networks, which had not been realized in any spatially distributed and sparcely networked form. What is more, provably correct formation control and stabilization that did not tie vehicles to specific locations in space or predefined desired trajectories for each one of them, was still elusive.

The research activity thus set out to provide solutions to the particular subproblems associated with the above objectives, and tie these solutions into a technically coherent and practically meaningful ensemble that realizes the key aspects of the envisioned robotic sensor network with distributed decision making capabilities.

In addition to the technical, research-related goals, the project aimed to reach out to K-12 teachers and students, primarily through established programs within the University of Delaware, but also at a national level through widely attended science workshops and festivals.

2 Accomplishments

2.1 Major activities

Research activities involve the development different theoretical, computational, and experimental frameworks for multi-agent motion planning and control, shape identification and classification, and distributed data processing. In parallel to these, simulation and experimental tests are conducted to assess the effectiveness of these methodologies. Outreach activities were related to involvement in a merit badge program with the local (Delmarva) Boy Scout council, organizing field trips for local schools, and participation to national events aimed at attracting wider audiences to subjects related to science and engineering.

2.2 Specific objectives

The research plan outlined three directions, or thrusts, each aligned with an objective identified above:

- describing shapes,
- identifying them, and
- collecting the required data for doing so.

Starting in reverse order, we will next outline what was done and achieved along these directions.

2.3 Significant results

Data collection Collecting sensor data in a way that enables and facilitates subsequent processing, analysis, and decision making, requires that the motion of the mobile platforms bearing the sensors is coordinated. In the case of spatial coverage, for instance, we want the area of interest to be covered quickly, with neither gaps nor excessive redundancy, and obviously in a way that may not compromise the physical integrity of the system. As Figure 1

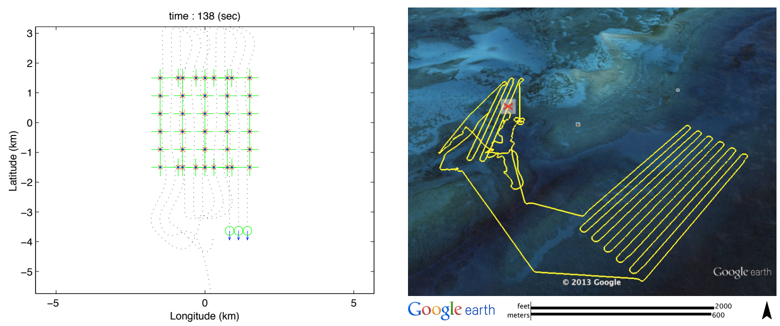

(a) (b)

Figure 1: Left: The concept of covering an area with a flock of robots as presented in the grant proposal. Right: Path, overlayed on a Google map, for a (single) underwater robot during a surveying mission (Courtesy of Art Trembanis).

indicates, the cooperative robot motion planning strategy suggested was quite close to what some of the methods that robot practitioners use in the field for this type of applications. Field implementations, however, typically involve a single vehicle—the reader may recall the efforts for the recovery of the black boxes of the recently lost Malaysian Airlines plane.

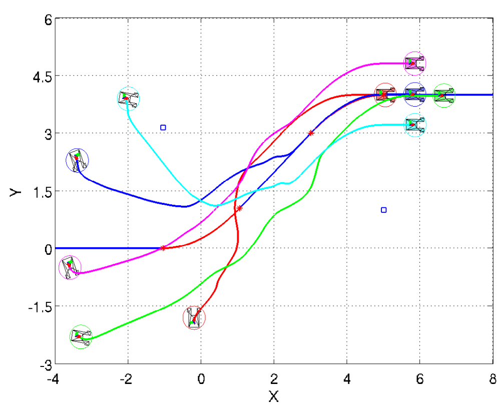

Figure 2: Formation stabilization is combined with flocking and path following. In this control law, the three behaviors coexist and evolve in concert with each other.

In order to scale this idea up from one robot to a group of robots that move like a single, rigid, sensor array and in a provably convergent say, the project had to develop new motion planning and control methodologies for formation control and flocking. To do so, the research team built on the idea of (single-robot) navigation functions and extended them to the multi-agent case while avoiding specifying fixed global destination configurations for each vehicle in the formation—this is one of the features that sets this new methodology apart from its earlier counterparts. Another important characteristic of the developed new methodology is that vehicle controllers, although steering the vehicle group globally from any initial to a specific desired relative configuration, they only have to focus on the intervehicle configuration that presents the most imminent collision risk in order to preserve collision free motion. That feature, in conjunction with the design of the mathematical representations of proximity between vehicles, eventually lead to faster computation, and facilitate decentralization. On the technical side, however, this presented significant new challenges, in the sense that the mathematical proof for collision avoidance and convergence to desired configuration had to employ elements of nonsmooth analysis and stability theory, since the minimum distance function among a set of moving particles changes over time in a non-smooth way. The mathematical framework that was developed to plan and control the motion of the robotic sensor nodes at the kinematic level was reported in [12].

Some challenges, however, remained. Robotic vehicles usually have non-negligible dynamics, and as Figure 1 suggests they should not just float around when they fall into their desired formation but rather follow specific paths. As it turned out, this required a nontrivial technical extension to the previously reported algorithm, in order to bring vehicle dynamics into the picture, and integrate the robot navigation scheme with established flocking control strategies and enable the group, *as a whole* to follow a path like the one shown in Figure 1(b). Practically, a path such as the latter, is prescribed for (the single) underwater vehicle in the form of waypoints; in our case, we assume that we can decompose a desired reference path like this in a sequence of straight line segments and Bézier curves. We therefore extended the multi-agent navigation strategy of [12] to give rise to a combined flocking and formation control behavior on behalf of the vehicle group, which seamlessly transitions to path following while maintaining a rigid formation (Figure 2). These results were reported in [13].

The motion of robotic platforms was thus coordinated in order for consistent data collection. The term consistent here is used with reference to methods that would allow individual data sources to be collated into larger-scale collections. Controlling tightly the relative positions between the robots that were collecting data enabled us to register the data from each robot and create mosaic ensamples [13] (Figure 3) .

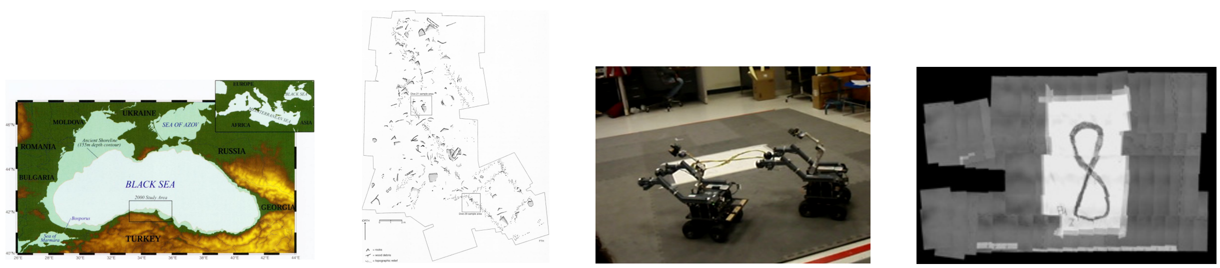

(a) Underwater survey area (b) Sea bed mosaic (c) Robotic testbed (d) Test mosaicking

Figure 3: (a)–(b) (left): suggested applications in the field of undewater archeology, reproduced in the grant proposal under permission from the original authors [1]. (c)–(d) (right): the same mosaicking concept demonstrated via a formation of robots.

Shape identification Once data is collated, the research plan called for the identification of specific patterns within them in a distributed fashion. The idea was to use a spatially distributed neural network, having subsets of its nodes residing on different robotic platforms (Figure 4). The neural network structure originally proposed was that of a Hopfield network, in view of the associative memory behavior of this class of neural networks. The research plan

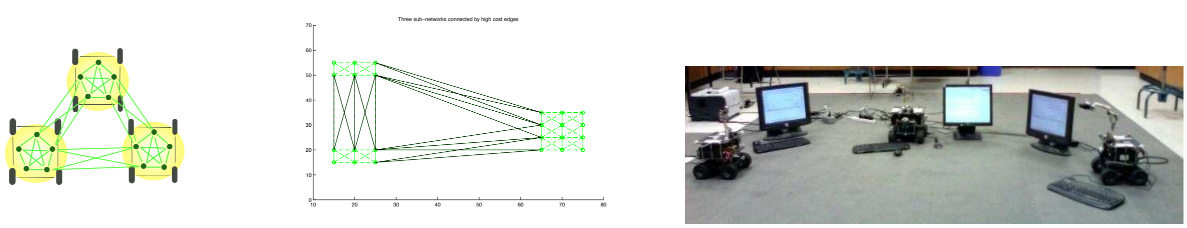

(a) Proposed topology (b) Tested implementation (c) Robot testbed

Figure 4: (a) schematic of an example of the proposed hopfield network topology, drawn from the grant proposal. (b) the topology of a cellular neural network implemented to test the associative memory properties of the developed neural network architecture; the edges color coded black are considered to incur a higher communication cost as they are implemented across physical platforms; (c) the physical testbed of networked robots on which the distributed neural network was realized and tested.

also called for topology optimization on these distributed neural networks, an innovation that aimed at alleviating the cost of communicating information across neurons. As research progressed, we came to the conclusion that a slightly different class of neural networks were more suited to the problem treated, since they did not originally assumed a complete network interconnection topology, but instead were originally designed in a lattice network structure which was relatively sparse and facilitated topology optimization. We thus switched our attention to this class of neural networks and focused our network optimization efforts on the already sparse lattice topology. Along this direction

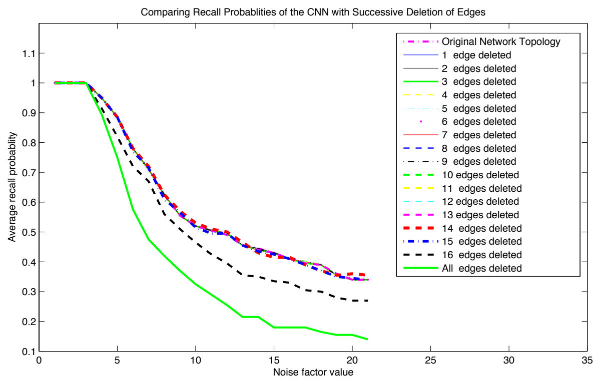

Figure 5: Network performance, parameterized by the number of links removed; the links removed sequentially correspond to the high-cost edges of Figure 4(b).

we were able to show [4, 3] that the performance of the neural network in terms of its ability to recall patterns from memory degrades quite gracefully up to a certain level of edge removal, after which the network stops functioning (Figure 5). The related publications [4, 3] reported for the first time a topology plus weight co-design in this type of neural networks that aims at producing considerably sparser, compared to traditional cellular neural network designs, network topologies with comparable associative memory properties. In fact, in the example depicted in Figure 4(b), the neural network continues functioning without observable performance degradation even when 14 out of its 17 “expensive” long distance links are severed. Figure 6 shows four out of the ten different patterns stored in the small cellular neural network of

Figure 4(b) after it has been optimized by reducing the number of expensive long distance connections. The stability properties of the resulting network can be ensured [3] through an algorithmic process that selects the neural network link weights based on the solution of an LMI that takes into consideration the underlying communication topology as well as the patterns that need to be stored as memories. The network works by accepting a distorted version of some pattern, and then through a sequence of discrete computation steps, it equilibrates to a configuration in its memory that is closest to the one it has been presented with. The network’s inherent stability properties guarantee this convergence, and the weight design algorithm, although it cannot guarantee the complete elimination of spurious local equilibria, ensures that the network’s attractors definitely include the patterns the designer intended to etch into the neural network’s memory.

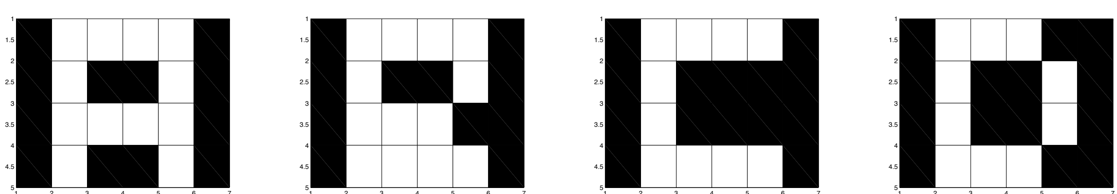

(a) (b) (c) (d)

Figure 6: Some of the patterns used for testing the recalling properties of the sparse neural network

In principle, the distorted input to the neural network could be a quantized version of the raw data that the robotic sensor network collected. In practice, this is not ideal for many reasons. One of them is scalability: the number of nodes in the network must be equal to the number of “pixels” in the raw data. Another very important limitation is that if the pattern in the data is scaled or rotated with respect to the nominal one stored in memory, it will look very different to the network and will not be able to be recalled. For these reasons, the research plan called for the use of particular rotation and scale invariant *shape descriptors*, that is, a fixed number of specific feature representations of (image) data that could describe shapes implicitly. Features that showed promise in terms of serving as such shape descriptors where the eigenvalues of a partial differential operator related to a Dirichlet PDE having as boundary the shape of interest. The following section elaborates on the efforts to use these descriptors for the problem at hand.

Before moving on, it should be noted that although most of the work in this domain of pattern recognition and shape description and identification is in the context of image processing, we need to keep in mind that it is fairly straightforward to represent scalar spatial fields resulting from a wide variety of physical quantities (e.g., pressure, temperature, humidity) in the form of an image.

Shape description The plan proposed for the description of shapes drew from the work on the Dirichlet Laplacian eigenvalue method [9, 14], which promises scale and rotation invariant shape descriptors. Another attractive feature of this methodology is that it does not require (image) segmentation. We thus set off to adapt the Dirichlet Laplacian eigenvalue methodology and apply it to more realistic application scenarios similar to the one in which one would ask to identify shapes within the mosaic of Figure 3(b) that appeared in our grant proposal.

Through our intramural collaborators at Oceanography, we got access to similar data (Figure 7—compare to Figure 3(b)). This is a mosaic of backscatter sonar readings of the ocean floor at Redbird reef off Delaware's Cape Henloppen, where old NYC subway cars were thrown at an attempt to create an artificial habitat for aquatic wildlife.

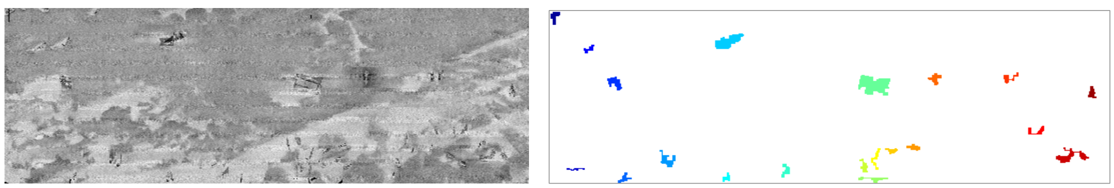

(a) (b)

Figure 7: Backscatter image of Redbird reef ocean floor (left) and a filtered version of it that brings out potential locations of submerged cars

We thus applied the methods with some minor adjustments to data like those shown in Figure 7. Figure 8 shows an example of the outcome of the Dirichlet eigenvalue analysis applied to the the shapes of Figure 7(a) as filtered in Figure 7(b). The example of Figure 8 depicts the best and worst matches of the filtered shapes to a rectangle that would correspond to a subway car outer outline.

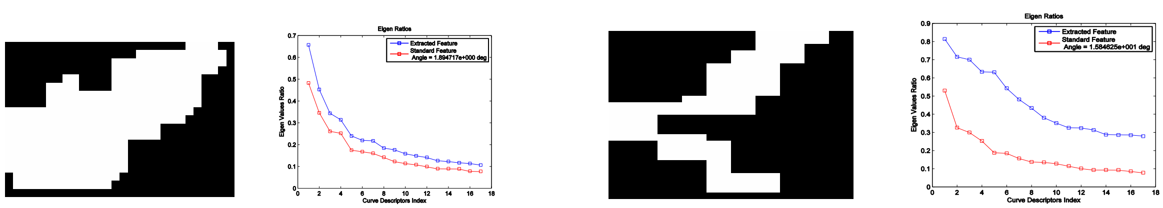

(a) (b) (c) (d)

Figure 8: Sensor data and associated shape descriptors. Parts (a) and (b) show the data and shape descriptor representations, respectively, for the data pattern best matching a rectangle. Parts (c) and (d) do similarly for the worst match. In parts (b) and (d), the red curves show the eigenvalue descriptors for a nominal rectangle of the size and orientation of the data pattern, and the blue curves mark the eigenvalue descriptors for the specific data pattern. Larger deviation of the two curves indicate stronger mismatch.

In principle, the eigenvalue descriptors mapped in blue in Figures 8(b) and 8(d) would be fed to a topology-optimized cellular neural network in order to be associated with the closest pattern in its memory. This way a shape would have a resolution-free numerical signature, one which even a small-scale neural network like the one in Figure 4(b) could process. As we discovered, however, these shape descriptors when applied to this type of data within a discretized mathematical framework, were neither completely scale nor orientation invariant. Actually, hints at this behavior can be found in Figures 8(b), 8(d) where it can be seen that the red curves are not identical; the behavior has been independently observed in more tests and experiments. This fact, in conjunction with the unimpressive matching performance of this method to our underwater sonar data led us eventually to depart from this approach and investigate more conventional methods for large-scale pattern recognition, even if they were not parallelizable in the spirit of the previous sections.

The new approach was multi-layered and presented technical innovations too, only at different domains. In some respects, this approach was more comprehensive than the first attempt, in the sense that it automated the whole process of picking and filtering the potential candidate shapes (done manually in Figure 7(b)), individually segmenting

Figure 9: Low-resolution image taken by an autonomous underwater vehicle conducting marine surveys

them, and then applying template matching, all in one batch. Wanting to push the limits of this more conventional methodological approach, we moved from an application problem of matching fairly easily (to a human) discernable shapes, to that of identifying and counting smaller benthic objects that easily blend to the background: scallops in low-resolution, noisy underwater images (Figure 9). To an untrained eye, identifying these marine animals in the picture of Figure 9 is not immediate. An expert would probably identify a scallop inside the depression in the sand visible on the middle of the figure and slightly to the right. In the related research project our oceanography colleagues were conducting, manual count on a collection of many tens of thousands of such pictures was the only option. Our ambition was that an automated system could do that with reasonable accuracy.

The selection of candidate targets out of a whole data body is thus a critical part of the whole process, if throughput is an issue, because it focuses the algorithm's "eye" to specific regions rather than uniformly scanning the whole image. We attacked this problem by integrating adapted versions of visual attention algorithms into the first stage of the batch identification process outlined in the previous paragraph. Application of visual attention to such applications is very limited (we only found a single other study) and restricted to a very small data set—tens as opposed to several-thousand-sets collected in a single AUV mission. In our application visual attention (see Figure 10) focuses the search on particular regions of the image that "stand out" more; the process is likely to miss some well hidden targets, which an untrained eye would be very likely to miss too, anyway. From that point on, modulo some methodological

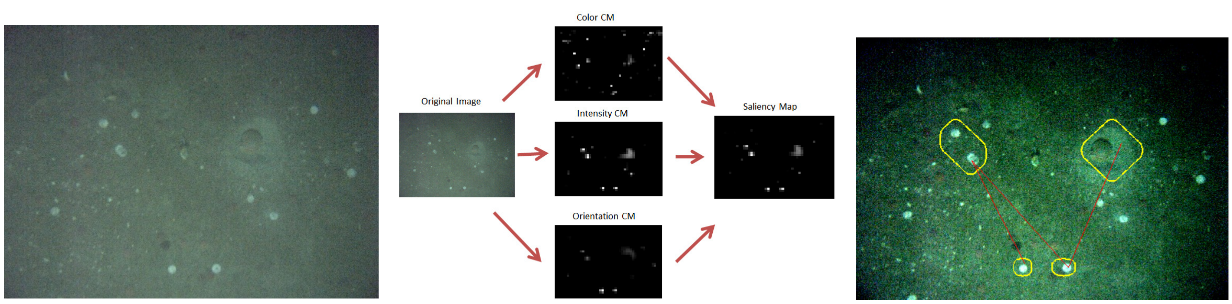

(a) (b) (c)

Figure 10: Application of visual attention on the low-resolution underwater survey pictures

innovations that enable streamlining of this whole process of targeted processing and pattern recognition, the procedure evolves as can be expected. The regions picked up by visual attention are segmented through graph-cut techniques, and the segmented parts are matched to reference templates. In the absense of any distinctive morphological, color, or texture features in the scallops in the images such as the one in Figure 9, we resorted to searching the small shadow (Figure 15(a)) that scallops cast on the sand when illuminated by the strobe light of the AUV, and then using that to fit a circle of dimensions close to the mean size of scallops (Figure 11(e)). The alternative automated shape identification process exemplified in Figure 11 eventually offered reasonably good, relative to the quality of the data, detection performance (Table 1). The algorithm was applied to two large data sets, one containing 1 299 images, and a second with 8 049 images. Table 1 shows the performance results of every stage / layer of the algorithm compared to (manually annotated) ground truth. There are two direct observations: the first relates to the relatively low accuracy (60–70%), and the second to the significant number of false positives. The first observation needs to be put into context; specifically, the quality of the raw data has to be factored in. In comparison to other existing automated scallop identification studies [6], which have been conducted on imagery data taken with different process and equipment (Figure 12), the developed algorithm performs admirably given the difference in the resolution and noise content of the data it had to perform on. The increased success rate of alternative methodologies can be attributed to the bimodal

(a) (b) (c) (d) (e)

Figure 11: (a) Fixation window identified by visual attention; (b) Edge segmented image; (c) Graph-cut segmented image; (d) Region boundaries obtained when the edge segmented image is used as a mask over the graph-cut segmented image boundaries; (e) circle fitted on the extracted region boundaries.

Table 1: Results of multi-layer scallop classification

	Dataset 1	Dataset 2
Number of images	1 299	8 049
Ground Truth Scallops	363	3 698
Valid Ground Truth Scallops	250	2 781
After Visual Attention Layer	231 (92.4%)	2 397 (86.2%)
After Segmentation Layer	185 (74%)	1 807 (64%)
After Classification Layer	183 (73%)	1 759 (63.2%)
False Positives	17 785	52 456

nature of the background histogram, resulting from the absense of speckle noise and color distortion. Details on the development, application, and performance of the algorithm are reported in [8].

2.4 Key Outcomes

Research The research activity in this project produced mathematical and algorithmic tools that enable a group of robots equipped with sensors to coordinate in data collection, and recognize autonomously patterns in the data they have collected. Under certain conditions, and subject to some limitations, this process can be completely decentralized in the sense that no single robot would be required to collect all information from its teammates and perform the bulk of the computation. These ideas were demonstrated first in a proof-of-concept mobile robotic testbed tasked with recognizing large-scale visual patterns on the surface on which they moved, and tested independently against field-collected (underwater sonar and imagery) data. The idea behind the main hypothesis laid out in the grant proposal was therefore verified; at the same time, as some limitations of such a framework became clearer, so did potential paths forward which hold promise in terms of confronting these technical and intellectual challenges.

A first intellectual question that came to the forefront during the execution of the research plan was related to the tradeoff between communication cost and latency when implementing (big) data analysis and processing in a distributed fashion. While decentralization can have clear benefits in certain settings in terms of robustness—specifically, when individual agents within a big group are tasked with performing *similar*, if not identical, tasks—these benefits should be weighted against the cost of passing around the information needed for distributed computation. Information has to be transmitted anyway in any multi-agent sensor and actuator network, irrespectively of where the processing of this information takes place. Yet, distributing the processing, in addition to the data collection, can impose a non-negligible cost in terms of network delay and latency which can impede the performance of the whole system—note that we refer to performance only, since stability and convergence in data collection and inference can be shielded relatively easily in the context of existing mathematical tools. In our testbed implementation, these latencies were noticable but inconsequential. They could primarily be attributed to the parallelized computation of the Dirichlet Laplacian eigenvalues, i.e., they were related to the distributed processing of the information that the network collected, this particular stage in the whole process that comes after data collection and before pattern recognition. In much larger-scale realizations these latencies, as they would scale with the size of the network, could impose a severe time overhead.

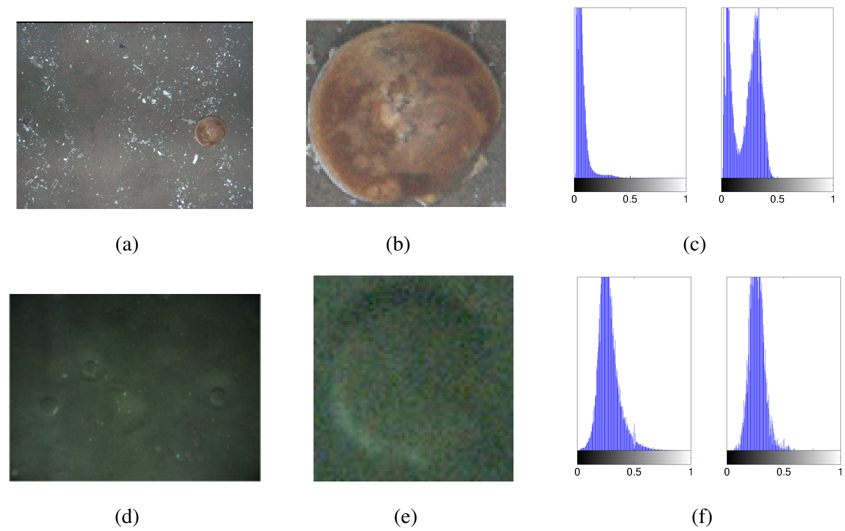

Figure 12: Representative samples of different imagery data on which scallop detection algorithms may be called to operate on. In (a) and (d), we see an image containing a single scallop from the dataset used by [6] (used with permission from the authors) and the datasets used in this paper respectively. A magnified view of a scallop cropped from (a) and (d) can be seen in (b) and (e) respectively. Part (c) gives the saturation histogram of background or the complete image in (a) to left and saturation histogram of (b) to the right. Similarly, part (f) gives the saturation histogram of (d) to the left and saturation histogram of part (e) to the right. The bimodal nature of the scallop histogram in (c) derived from the dataset used in [6], clearly portrays the distinguishing appearance of the scallop pixels from the rest of the image, making it easily identifiable. The datasets we used did not exhibit any such characteristics (as seen in (f)) to aid the identification of scallops.

At a higher level, the technical bottleneck that has been identified is in the coding and compression of the body of information collected, and the lesson learned is that promising approaches may be based on ideas pivoting around local processing of sensor information at the sensor node level, without requiring this node to communicate with its peers specifically for doing this. We have seen some evidence that in some specialized problem domains of distributed detection [10], this can be possible without sacrificing performance. Whether such ideas and methods can be transferred from these specialized domains into the problem area treated here, is an open question.

Another technical issue that became apparent was the possibility for fragility of certain machine learning algorithms, even when applied within a framework which they have been reported to have preformed well in. In this case, the Dirichlet Laplacian eigenvalue methods proved to be somewhat sensitive to rotation and scaling, presumably due to *discretization*. We hypothesize that discretization is at the source of this behavior, as these methods had been reported as able to offer shape descriptors that were rotation and scale invariant. We have found no work that commented on the effect of discretization resolution and quantization on the invariance properties of these descriptors. Having said that, even with the possible effects of quantization factored in, the specific methods still fell short of their reported performance when applied to the body of data resulting from field-surveys.

We found that such fragility is not unique to these eigenvalue methods. A large number of the more conventional machine learning and pattern recognition methods did not always fare well when applied to different domains from the ones been tested and on the type of low quality input that we were experimenting with. Our investigation led us to multi-stage processing and pattern recognition approaches that combined visual attention, image segmentation, and template matching, but also revealed possibilities for image information compression at a local level—facilitated by visual attention. The form of information encoding currently being investigated, in conjunction with the exploitation of sensor motion in the spirit of active sensing, open new exciting possibilities for localized autonomous remote sensing using robotic sensor platforms.

Education The project supported fully or partially four graduate students. Three masters students have already graduated, and one doctoral student is expected to complete all the requirements for his degree within the coming academic year. In addition, elements from the research activity, specifically parts that related to flocking and cooperative behaviors were used as application examples in the junior-level graduate course MEEG620 (Intermediate Dynamics) and the upper-level graduate course MEEG873 (Nonlinear Control) that the PI tought in the fall and spring of 2013, respectively.

(a) (b) (c) (d)

Figure 13: Robotics projects completed by local boy scout teams, part of the requirements for earning their robotics merit badge

Figure 14: Students from first through third grade in a field trip in the PI's lab. The class presented a challenge to the labs robots in the form of a shape (a large hand-drawn number—the class picked a figure 8) that the robots had to identify after scanning the floor, and mosaicking the pictures taken by their cameras—see Figure 3(d).

Outreach There was substantial outreach activity during the duration of and in the context of this project. One of the activities, which by now have become a regular outreach event, is the annual Boy Scout robotics merit badge program, managed through the local Iron Hill Museum and the Delaware Academy of Sciences. The PI serves as a robotics badge councelor, hosts groups of scouts for a two half-day workshop filled with lab tours, educational presentations, team activities, and culminating in a lego robotics challenge (Figure 13) where the scouts present their robotic design, document and explain its development, and answer general questions related to their robotics merit badge certification (e.g., careers in robotics, safety considerations, available competitions, etc.) Some of the experiments and testbed evaluations involved students from a local primary school, giving these youngsters an opportunity to interact with robots, learn more about them, and get to challenge the machines. In fact, the test pattern appearing in Figures 3(c) and 3(d) were created by the class shown in Figure 14 during their field trip to the PI's lab. At the national level, students from the PI's lab participated in the U.S. Science and Engineering Festival held in Washington DC in the spring of 2014. The students worked with UD's College of Engineering Outreach Coordinator Mel Jurist manning the university's booth, demonstrating homegrown robotic devices to children and letting them play around with them through remote controls (Figure 15).

3 Products

3.1 Training and professional development

The primary tangible outcomes of this research activity can be identified in the form of students graduated, research publications, and integration of research elements into the curriculum.

On the education side, the project provided the main source of support for three masters students and one doctoral student, who worked on different aspects of the project.

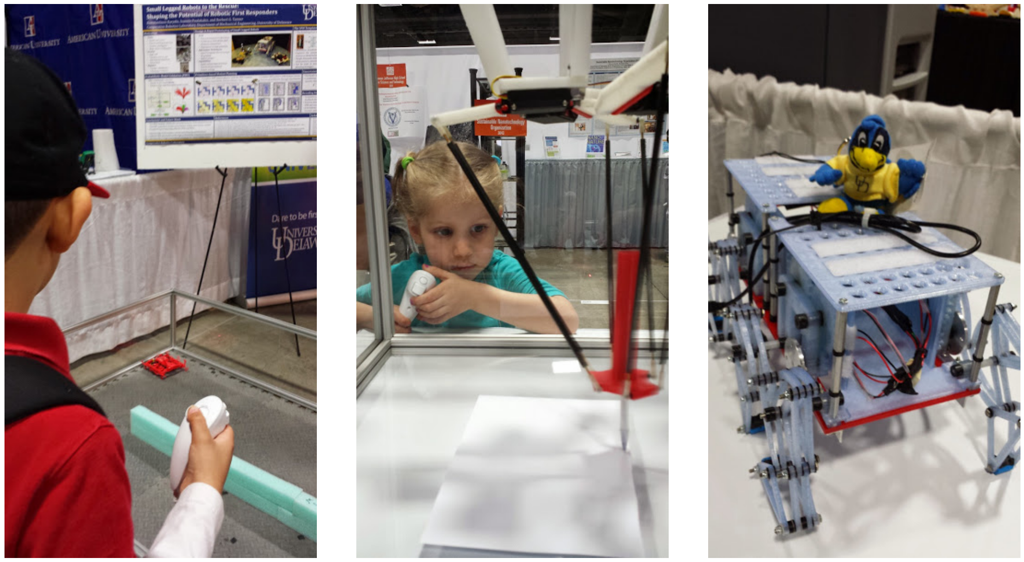

(a) (b) (c)

Figure 15: The PI's students and robotic devices in action, driven by children during the 2014 U.S. Science and Engineering Festival

- Varsha Bhambhani worked on the neural network optimization and its distributed implementation for pattern recognition, successfully defended her masters thesis [2] in the winter of 2012, and now continues her career in industry.
- Adithya Boddu worked on provably correct formation control strategies, defended his masters thesis [5] in the winter of 2012.
- Luis Valbuena R. worked on the hardware implementation of the formation control strategies and their integration with flocking and path following techniques; he defended his masters thesis [11] in the fall of 2012, and went on to continue his graduate studies in the University of New Mexico.
- Prasanna Kannappan is a PhD candidate who still works on the pattern recognition and object identification, while specializing on underwater robots and data collection. He is expected to complete the requirements for his degree in the spring of 2015.

3.2 Dissemination of results

There were a number of journal publications (4) and conference papers and presentations (2) that resulted from this project. There were two published journal papers [3, 12], out of which [3] was selected as the outstanding paper in 2012 for that journal. There are two more journal papers still in review [8, 13], of which [8] documents the direct application of the developed techniques in the domain of marine surveys and environmental monitoring and is coauthored with our colleagues from oceanography. Out of our two conference papers, [4] documents the first resutls on topology optimization in cellular neural networks, and [7] reports on preliminary results on stages of the multi-stage scallop recognition algorithms. In addition, dissemination of the research outcomes in relation to underwater imaging and object recognition was presented in the context of NOAA's January 2014 Undersea Imaging Workshop that was held in Red Bank, NJ. In this forum, our researchers interacted with, and informed, the scalloping community about our recent results on automated scallop detection and counting using autonomous underwater robots.

Further dissemination of the main goals and outcomes of this activity, as well as a copy of this report, is made possible through a dedicated web site for this project:

http://research.me.udel.edu/~btanner/Herbert_G._Tanner/Research/Entries/2009/7/15_Swarms_that_”hear_the_shape_of_the_drum”.html

which will be archived and maintained by the PI for a period of at least 5 more years. Since, in general, project reports are not easily accessible by the public, the PI makes this formated version of the project's final report available through the above website.

4 Participants

Including the PI and author of this report, the following individuals worked on this project for a substantial amount of time and received some form of support from the associated award. All participants were affiliated with the University of Delaware, the awardee institution.

- Herbert Tanner; Principal Investigator
- Prasanna Kannappan; Research assistant—PhD Student
- Luis Valbuena R.; Research assistant—MSc student
- Varsha Bhambhani; Research assistant—MSc student
- Adythya Boddu; Research assistant—MSc student

No other individuals received any form of support from this award. Coleagues who collaborated scientifically or contributed somehow to this project by making available data or advice, are listed in the acknowledgement section of this report.

5 Impacts

On the robotics discipline Component technologies within the whole data collection and autonomous pattern recognition architecture developed in this project, consitute methodological and theoretical advancements with regards to the state of the art in the robotics field. A new class of potential-field-based formation controllers was developed, which can formally guarantee almost global convergence to specific formation shapes, allowing formation shapes to "float in space rather than being pinned to pre-specified positions. The methodology establishes nonsmooth analysis as a mathematical tool for multi-agent robot motion planning and control. Another major development was the recognition that the interconnection topology in cellular neural networks can be a target for optimization aiming at reducing inter-neural communication overhead and the realization that far sparser network topologies are possible without sacrificing performance in terms of the ability of the network to recall stored patterns. Along the lines of this development is also the experimental demonstration that this type of neural networks can be implemented in a spatially distributed way.

On the discipline of marine robotics and environmental monitoring The recent research activity has demonstrated the potential of robotic automation and machine learning in environmental monitoring and oceanographic research. The sea scallop (Placopecten magellanicus) fishery in the US EEZ (Exclusive Economic Zone) of the northwest Atlantic Ocean has been, and still is, one of the most valuable fisheries in the United States. Our work on automated scallop detection and counting—an extension of the shape classification research thrust of this project seems to have potential in applications of environmental monitoring that involve processing of large amounts of data, which currently require many man-hours of visually straining work.

On human resource development The project supported the education of four graduate students, several of who have entered the workforce. The other two continue their graduate study, and will transition to different research projects. In addition, the project provided opportunities for undergraduate research and for two undergraduate students who worked with the PI and his students earning independent study credits toward their bachelor's degree.

6 Changes & Problems

Changes in the originally proposed approach and reasons for the change During the course of research, two realizations required adjustments in the research plans. The first was the emergence of visual detection and classification of objects in unstructured images, obtained in a natural outdoor environment as a major technical challenge on the critical path. Without having particular expertise in the specific area, the research team had the option of circumventing this challenge by artificially simplifying the problem to the point that it was rendered amenable to standard tools. Instead, we chose to address it head on, taking a fresh look at the issue. As a result, we decided that the originally envisioned shape classification method which was based on eigenvalue descriptors is inadequate on its own to solve this problem, and we searched for machine learning, detection, and classification tools with more promise for handling noisy, unstructured, visual scenes from the natural world. Our efforts on the combination of visual attention algorithms with robust image segmentation methods (e.g. graph cut-based) and template matching eventually succeeded, and paved the way for extensions of this idea along a direction that draws from active sensing and exploits sensor motion.

A second realization that moved us away from swarm-based, distributed algorithm implementations, was the tradeoff between communication cost and decentralization of computation tasks. Although some processes can be implemented in a decentralized way, with obvious advantages in terms of robustness and scalability, the necessary communication overhead in terms of bandwidth and time delay that comes with decentralization has to be carefully weighted against these benefits. In conclusion, although a lot of computation processes can be made distributed, only a fraction of those may be particularly suited for such implementation, in the sense that their communication requirements are both small, and scale nicely. Given that the object detection problems we ended up dealing with as a result of our interaction with our colleagues in the College of Earth, Ocean and Environment at the University of Delaware are of particular nature and present specific technical challenges, we have shifted our focus from decentralizing computation to the fundamental limitations of the detection problem itself. Accumulating sufficient insight into the problem, and developing effective centralized solutions is an important first step before decentralizing the computation of these solutions.

These technical approach changes did not impacted expenditures in any significant way.

References

[1] R. D. Ballard, F. T. Hiebert, D. F. Coleman, C. Ward, J. S. Smith, K. Willis, B. Foley, K. Croff, C. Major, and F. Torre. Deepwater archaeology of the black sea: The 2000 season at Sinop, Turkey. *American Journal of Archeology*, 105(4):607–623, October 2001.

[2] V. Bhambhani. Topology optimization in cellular neural networks. Master's thesis, University of Delaware, http://udspace.udel.edu/handle/19716/11197, 2012.

[3] V. Bhambhani, L. A. V. Reyes, and H. G. Tanner. Spatially distributed cellular neural networks. *International Journal of Intelligent Computing and Cybernetics*, 4(4):465–486, 2011.

[4] V. Bhambhani and H. G. Tanner. Topology optimization in cellular neural networks. In *Proceedings of IEEE International Conference on Decision and Control*, pages 3926—3931, 2011.

[5] A. Boddu. Nonsmooth multi-agent navigation functions. Master's thesis, University of Delaware, http://udspace.udel.edu/handle/19716/11190, 2012.

[6] M. Dawkins, C. Stewart, S. Gallager, and A. York. Automatic scallop detection in benthic environments. In *IEEE Workshop on Applications of Computer Vision*, pages 160–167, 2013.

[7] P. Kannappan and H. G. Tanner. Automated detection of scallops in their natural environment. In *IEEE Mediterranean Conference on Control and Automation*, pages 1350–1355, 2013.

[8] P. Kannappan, J. H. Walker, A. Trembanis, and H. G. Tanner. Identifying sea scallops from benthic camera images. *Limnology and Oceanography: Methods*, (submitted).

[9] M. A. Khabou, L. Hermi, and M. B. H. Rhouma. Shape recognition using eigenvalues of the dirichlet laplacian. *Pattern Recognition*, 40:141–153, 2007.

[10] C. D. Pahlajani, I. Poulakakis, and H. G. Tanner. Networked decision making for poisson processes with applications to nuclear detection. *IEEE Transactions on Automatic Control*, 59(1):193–198, 2014.

[11] L. V. R. Flocking with formation control in mobile sensor networks for area search. Master's thesis, University of Delaware, http://udspace.udel.edu/handle/19716/12915, 2012.

[12] H. Tanner and A. Boddu. Multi-agent navigation functions revisited. *IEEE Transactions on Robotics*, 28(6):1346–1359, 2012.

[13] L. Valbuena and H. G. Tanner. Flocking, formation control and path following for a group of mobile robots. *IEEE Transactions on Control Systems Technology*, (submitted) 2013.

[14] M. Zuliani, L. Bertelli, C. Kenney, and B. S. Manjunath. Drums, curve descriptors, and affine invariant region matching. *Image and Vision Computing*, 2007. (in Press).