

Cooperative Multi-Agent Inference over Grid Structured Markov Random Fields

Ryan K. Williams and Gaurav S. Sukhatme

Abstract—In this work we investigate cooperative inference in multi-agent systems where uncertainty is modeled by the grid structured pairwise Markov random field. A framework is proposed, which we term the *multi-agent Markov random field*, that decomposes the global inference problem into inter-agent belief exchanges over a *hypertree* topology and local intra-agent inference problems. Due to the exponential complexity of exact inference, we propose a *loopy belief propagation* algorithm for approximate inference over appropriately formed local *generalized cluster graphs*. Both synchronous and intelligent message passing are considered and a grid scale-invariant scheme based on the notion of *regions of influence* in a cluster graph is presented. The algorithms are simulated over a grid workspace with a team of virtual Autonomous Surface Vehicles (ASVs), with the goal of spatial plume detection in oceanographic data captured from the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument. We show that while the exact method produces predictably accurate and smooth grid maps, the approximate method competes well in terms of plume detection rate with the region of influence message passing scheme excelling over large tasks due to a lack of dependence on grid size.

I. INTRODUCTION

Probabilistic networks are an elegant framework for modeling and managing uncertainty in a single-agent setting. Well-studied algorithms exist for inference and learning in probabilistic networks and many real-world implementations have been proven highly effective [1]–[3]. Extending probabilistic networks to *multi-agent systems* (MASs) is particularly attractive as such systems can afford significant gains in computational efficiency through concurrency, enhanced scale by spatiotemporal distribution, and robustness to failure via redundancy. The extension of the directed *Bayesian network* (BN) to MASs is most common in the literature, with applications ranging from coordinated tracking and surveillance [4] to distributed fault detection [5].

While BNs share some commonality with undirected *Markov random fields* (MRFs) in the probability distributions that they can represent, there exist certain classes of distributions that cannot be modeled by a BN structure (e.g. distributions that are characterized by cyclical dependencies). For this reason, we study the representation of probabilistic uncertainty and the inference process in MRFs over MASs, which we term *multi-agent Markov random fields* (MAMRFs). In the MAMRF framework, each agent maintains and

exchanges its beliefs over a subset of domain variables with neighboring agents, where the agent *communication topology* is defined by a *hypertree* graph structure over the system. Local and global evidence obtained via inter-agent communication is used to perform intra-agent inference in order to answer queries and perform actions (see [6] for closely related work).

The central example for this paper will be the widely relevant pairwise MRF with grid structure. One of the primary strengths of this formulation of MRF is its simplicity in modeling localized spatial interactions amongst variables in complex domains. In practice, such interactions are prevalent, with applications including image segmentation [7], spatial mapping [8], and object classification [9]. A partitioning of the grid structured MRF over multiple agents will be defined in the context of the MAMRF framework and a presentation of inference therein will be given. The intractability of exact inference over the grid structured MRF motivates an approximate approach based on *loopy belief propagation* (LBP) in *generalized cluster graphs*. To encourage grid scale-invariance, an intelligent message passing scheme based on *regions of influence* (ROIs) in a cluster graph will be contrasted to a synchronous approach in the application of loopy belief propagation.

The proposed inference algorithms are simulated over a spatial grid with the goal of plume detection in oceanographic data by virtual Autonomous Surface Vehicles (ASVs). The test data are captured from the satellite-based Moderate Resolution Imaging Spectroradiometer (MODIS) instrument (see [10] for related work). It is shown that while the exact method produces predictably accurate and smooth grid maps, the method’s exponentially rising cost renders it intractable over large workspaces. The approximate method is shown to compete well in terms of plume detection rate with the ROI message passing scheme enabling increased dimensionality due to its scale-invariance. Additionally, when compared to other plume detection and tracking approaches (e.g. [11]–[13]), our methods generate probabilistic outputs rather than hard decisions or controls, do not require process or environmental gradients, and fully leverage global observations to generate solutions.

The outline of the paper is as follows. The proposed MAMRF framework is presented in Section II. In Section III we analyze the grid structured MAMRF, including network representation, an algorithm for approximate inference, and a discussion of LBP message passing schemes. Simulation results are provided in Section IV. Finally, concluding remarks as well as directions for future work are stated in Section V.

The authors are with the Departments of Electrical Engineering and Computer Science at the University of Southern California, Los Angeles, CA 90089 USA (email: rkwillia@usc.edu; gaurav@usc.edu).

This work was supported in part by the NOAA MERHAB program under grant NA05NOS4781228, a DARPA MURI under grant N00014-08-1-0693, the NSF CPS program under grant CNS-1035866 and a fellowship to R. K. Williams from the USC Viterbi School of Engineering.

II. THE MULTI-AGENT MARKOV RANDOM FIELD

A Markov random field is a triplet $\mathcal{M} = (\mathcal{X}, \mathcal{H}, P)$, where $\mathcal{X} = \{X_1, \dots, X_N\}$ is a set of N domain variables, $\mathcal{H} = (\mathcal{X}, \mathcal{E}, \phi)$ is an undirected graph with nodes labeled by \mathcal{X} , edges \mathcal{E} , and potentials ϕ , and P is a probability distribution over \mathcal{X} . The graph \mathcal{H} encodes a set of independencies present in P over \mathcal{X} . In an MAMRF, a set of K agents $\mathcal{A} = \{A_1, \dots, A_K\}$ each maintains belief over a subset of domain variables $\mathcal{X}_i \subset \mathcal{X}$ represented locally as a *Markov subnet* $\mathcal{M}_i = (\mathcal{X}_i, \mathcal{H}_i, P_i)$. The set of independencies implied by \mathcal{H} are then present as a partitioning $\mathcal{H} = \{\mathcal{H}_1 \cup \mathcal{H}_2 \dots \cup \mathcal{H}_K\}$ across the local graph structures. This partitioning allows each agent to reason over their subset of domain variables and exchange local belief over only a *shared* set of domain variables with neighboring agents. In this way, agents must *cooperate* via structured communication to reason over the global distribution P .

In order to ensure valid probabilistic inference over an MAMRF, the local subnets along with the subnet topology must obey certain conditions. Specifically, let \mathcal{H} be partitioned into subgraphs $\mathcal{H}_i = (\mathcal{X}_i, \mathcal{E}_i, \phi_i)$. Let the subgraphs be organized into a tree $\Psi = (\mathcal{H}, \mathcal{L})$ where each node in Ψ , called a *hypernode*, is labeled by \mathcal{H}_i and each edge in \mathcal{L} , called a *hyperlink*, is labeled with the *interface* $\mathcal{X}_i \cap \mathcal{X}_j$ between \mathcal{H}_i and \mathcal{H}_j . Agent A_j is considered a *neighbor* of agent A_i if $l_{ij} \in \mathcal{L}$, and the set of neighbors of agent A_i is denoted by Nb_i . Assume also that Ψ satisfies the *running intersection property*, that is for each pair of hypernodes \mathcal{H}_i and \mathcal{H}_j , $\mathcal{X}_i \cap \mathcal{X}_j$ is contained in each hypernode on the path from \mathcal{H}_i to \mathcal{H}_j . The tree Ψ is called a *hypertree* over \mathcal{H} [14]. The joint probability distribution

$$P(\mathcal{X}) = \frac{1}{Z} \prod_i \phi_i(\mathcal{X}_i) \quad (1)$$

is then the product of the potentials associated with each \mathcal{H}_i (with duplication ignored), where Z is the standard normalizing function.

The hyperlinks serve as the communication channels between adjacent agents, while the hypertree structure defines the communication topology of the multi-agent system. The running intersection property enforces *d-separation* of hypernodes by the hyperlinks and thus ensures that message-passing operations over the hypertree are probabilistically sound.

III. GRID STRUCTURED MULTI-AGENT MARKOV RANDOM FIELDS

We gain insight into the MAMRF by examining a multi-agent extension of the common and widely applicable pairwise MRF with grid structure. Define the domain $\mathcal{X} = \{X_{ij}\} \cup \{Y_{ij}^k\}$, where $i, j = 1, \dots, N_p$, $k = 1, \dots, N_m$, and $N = N_p^2(N_m + 1)$. Each X_{ij} , which we term a *process variable*, is a random variable with associated node potentials $\phi(X_{ij})$ and edge potentials $\{\phi(X_{ij}, X_{lm}), l, m = 1, \dots, N_p : (X_{ij}, X_{lm}) \in \mathcal{E}\}$, where \mathcal{E} is the edge set of a global graph structure \mathcal{H} . The Y_{ij}^k 's, called *model variables*, are random




Fig. 1. 2×2 grid structured pairwise MAMRF with two subgraphs sharing process variables $\{X_{11}, X_{12}, X_{21}, X_{22}\}$. Each agent maintains a single set of model variables and the hypertree structure is $\mathcal{H}_1 - \mathcal{H}_2$.

variables associated with each X_{ij} , with conditional edge potentials $\phi(X_{ij}, Y_{ij}^k)$.

The graph \mathcal{H} is defined by nodes labeled by each domain variable in \mathcal{X} , with an $N_p \times N_p$ pairwise grid structure over the process variables, and N_m model variables per grid node. We partition \mathcal{H} over \mathcal{A} with local graphs $\{\mathcal{H}_1, \dots, \mathcal{H}_K\}$, and require that \mathcal{H} partitions according to a hypertree structure Ψ . Each agent shares only the set of process variables, i.e. the interfaces are uniformly $\{X_{ij}\}$. The model variables are then partitioned over the agents in some fixed way to meet system requirements (e.g. to satisfy spatial or measurement constraints). We treat the process variables as unobservable; only model variable observations are available for inference over the process variables. Fig. 1 illustrates a trivial 2×2 grid structured MAMRF with two agents and one set of model variables per agent.

A. Inference Methods

Our goal is to infer a set of marginal distributions $\{P(X_{ij} | \mathcal{Z}), i, j = 1, \dots, N_p\}$ over the process variables that represents the collective conditional beliefs of the multi-agent system given a distributed set of observed model variables $\mathcal{Z} = \{Z_1 \cup Z_2 \cup \dots \cup Z_K\}$. To that end, we decompose the global inference problem into a set of inter-agent belief exchanges over the hypertree and a set of local inference problems performed per-agent over each \mathcal{H}_i . Exact inference over the local grid structured MRFs using a clique tree message passing algorithm has a complexity that grows exponentially in N_p [15]. To alleviate this problem for realistic domains we propose an approximate inference algorithm based on the extension of *loopy belief propagation* (LBP) over *generalized cluster graphs* to MASs.

To facilitate the extraction of process marginals, we replace the local graphs \mathcal{H}_i with generalized cluster graphs (denoted \mathcal{C}_i). For each potential in the original network \mathcal{H}_i we introduce a corresponding cluster and connect clusters with overlapping scope [15]. For the simple MAMRF in Fig. 1, the described method generates the cluster graph shown in Fig. 2 for local graph \mathcal{H}_1 . Notice that there now exists univariate clusters from which to extract the set of process marginals without the need for marginalization.

Inference begins with each agent initializing their local cluster graphs by computing initial cluster potentials, denoted ψ , by reducing assigned potentials by the set of local evidence Z_i . The agents then perform message passing over the hypertree to exchange beliefs over the shared set of process variables, induced by local observations on the hidden model variables. We apply a structured collect/distribute message passing scheme for the purposes of clarity (an *asynchronous* scheme is equally appropriate [15]).

During *message collection*, message sets are passed between neighbors upward through the hypertree towards an arbitrarily chosen root node. Each message set, defined by

$$\{\delta_{i \rightarrow j}^k\} = \psi_k \prod_{m \in Nb_i - \{j\}} \delta_{m \rightarrow i}^k, \quad \forall k \in Z_i \quad (2)$$

consists of messages generated by *observed* model clusters (the leaves in \mathcal{C}_i), sent from agent A_i to agent A_j . The domain of each message is the process variable associated with k th observed cluster. The *message distribution* step is analogous to the collection phase, with message sets flowing downward from the root.

At the conclusion of hypertree message passing, each agent has reached *consensus* with respect to the set of model-induced beliefs over the shared process variables. The final phase in the hypertree-based inference algorithm is the *injection step*, where each agent injects the messages received from neighboring agents over the hypertree into a local LBP inference process. Each agent A_i generates an injection message given by

$$\eta_i^j = \prod_{k \in Nb_i} \delta_{k \rightarrow i}^j \quad (3)$$

for each observed model variable (indexed by j) by multiplying the messages received over the hypertree that agree on message scope. The resulting injection messages, whose domains are each in the set of process variables, are assigned to the lightest cluster in the agent's local cluster graph containing η 's domain (i.e. the cluster with the fewest assigned potentials). After injection, each agent then runs an LBP inference process on their injected local cluster graph and extracts a set of marginal distributions over the process variables.

B. Message Passing in Loopy Belief Propagation

The primary drawbacks of the LBP algorithm are non-convergence and message passing complexity. There are various heuristic methods that can be applied to alleviate the issues, several of which are related to *intelligent* message passing approaches [15]. To address computational complexity, we propose an intelligent message passing scheme that exploits the underlying structure of the proposed grid structured MRF to determine regions in the graph where messages would be the most useful, as opposed to the *synchronous* approach. Noting that messages generated by non-observed model clusters are unitary and thus ineffectual in the inference process, we make the fundamental assumption that messages would be most useful in regions surrounding grid




Fig. 2. The generalized cluster graph for approximate inference over \mathcal{H}_1 in Fig. 1 [15].

clusters with *observed* model cluster neighbors. We call these areas of the graph *regions of influence* (ROIs). The ROIs are characterized by a fixed radius R over which standard sum-product messages are recursively passed to neighboring clusters. Message passing over the ROIs eliminates the need to search for regions of potential influence and also reduces significantly the total number of messages passed during an iteration of LBP when the ROIs do not fully overlap. As an example consider a large spatial grid 1000×1000 in size, over which two teams of agents are to operate. Let us assume that the ROI radii are chosen and the teams are arranged such that there exist two effective ROIs of radius 50 over which LBP message passing will occur. In such a scenario, each local cluster graph will contain $3N_p^2 - 2N_p = 2.998 \times 10^6$ non-leaf clusters over which messages are passed in each iteration of synchronous LBP. In comparison the ROI assumption reduces the effective cluster count to 5.96×10^4 , a reduction of about 50-to-1 per iteration. In particular, intelligent ROI message passing is independent of grid size and scales only according to the size of the evidence set (which is proportional to the number of agents) and the radius of the ROIs, enabling mission dependent tuning for achieving feasibility in otherwise infeasible workspaces.

The primary sacrifice that is made by the ROI assumption is that of accuracy. Since the ROIs are of fixed width, regions that do not overlap are rendered independent since messages will not propagate from one region to the other. While this property would certainly be inappropriate in some problem domains, there are others where it could be a benefit, for example in large spatial grids where the independence of distant features is quite a natural assumption.

IV. SIMULATION RESULTS

We present the simulation results for the proposed approximate multi-agent inference algorithm for a dynamic three agent system ($K = 3$) over a 10×10 spatial grid ($N_p = 10$). Agent dynamics are introduced by applying simple waypoint-based locomotion over the grid cells. At a predefined interval, the agents capture local observations and perform cooperative inference to iteratively compute a set of process marginals. The overriding goal of the system is

to probabilistically map a process of interest over the grid. For the purposes of our simulations we chose to consider the problem of mapping oceanographic phenomenon that are typified by plume-like spatial behavior (e.g. harmful algal blooms (HABs) [16]; more generally [10]). Towards this goal, we simulate observations over the workspace by sampling from sea surface temperature (SST, in $^{\circ}\text{C}$) and chlorophyll concentration (CC, in mg/m^3) data taken from the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument aboard NASA satellites Terra and Aqua [17]. In particular we selected data captured from a MODIS sampling taken in 2007 over the coast of California near Santa Cruz. Fig. 3 shows composite images of SST and CC data, and a plume-like region which we select for sampling and simulation. Each agent notionally represents an ASV on the ocean surface.

A single model variable is assigned to each agent per grid cell, where agents A_1 and A_3 are assigned SST and agent A_2 is assigned CC. As we are interested in plume *detection*, we choose binary process variables with potentials

$$\phi(X_{ij}) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \quad \phi(X_{ij}, X_{lm}) = \begin{bmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{bmatrix} \quad (4)$$

The model variables are assumed to be Gaussian distributed with

$$P(Y_{ij}^k | x) \sim \mathcal{N}(\mu_x; \sigma_x^2) \quad (5)$$

where for every $x \in \text{Val}(X_{ij})$ we have an associated mean and variance, μ_x and σ_x^2 . We determine the appropriate conditional Gaussian potentials empirically by selecting a portion of the workspace as being in-plume a priori and calculating mean and variance vectors

$$\begin{aligned} \mu_{sst} &= (14.5 \ 13.75), \quad \sigma_{sst}^2 = (0.5 \ 0.25) \\ \mu_{cc} &= (1.2 \ 1.9), \quad \sigma_{cc}^2 = (0.3 \ 0.2) \end{aligned} \quad (6)$$

The model parameters define the expected in plume and out of plume distribution of SST and CC measurements, while the process potentials reflect both an unbiased predisposition towards cell-wise plume membership and a desire for smoothness of the generated process probabilities. Although we choose a uniform and symmetric set of process potentials for simplicity, varied potentials afford flexibility in modeling process features that are significantly more complex, including spatial biasing and localized asymmetries.

Fig. 4 depicts an example of a 100 time unit simulation of the 3-agent system dynamically performing exact and approximate inference over the grid. The agents were assigned initial positions $x_1 = (0.5, 0.5)$, $x_2 = (9.5, 9.5)$, $x_3 = (0.5, 9.5)$, and a list of waypoints that define *lawnmower* patterned cell traversals (a standard pattern in ocean sampling). Agent measurements were sampled from the MODIS dataset with additive Gaussian noise defined by $\mathcal{N}(0; 0.1)$, to simulate the effects of imperfect sensors.

To assess the relative performance of the proposed inference algorithms, we generated 100 simulation runs with randomized initial positions and sampling patterns over the




Fig. 3. MODIS data used for measurement sampling in agent simulations: (a) sea surface temperature with highlighted region of interest; (b) chlorophyll concentration with highlighted region of interest.

grid. For each simulation we then computed the *Kullback-Leibler divergence* (KL-Div) for each grid cell marginal distribution versus a set of baseline distributions generated by performing exact inference over the grid with full evidence (i.e. all model variables are observed). The error of a grid of inferred process marginals is then given by

$$\epsilon = \sum_{i,j}^{N_p} \sum_x P_{ij}(X_{ij} = x) \ln \left\{ \frac{P_{ij}(X_{ij} = x)}{Q_{ij}(X_{ij} = x)} \right\} \quad (7)$$

where $P_{ij}(\cdot)$ is the baseline distribution and $Q_{ij}(\cdot)$ is the estimated distribution for a cell in row i and column j of the grid. We refer to this error as the cell-wise KL-Div.

The impact of cell-wise error on agent decision making is assessed by calculating the cell-wise classification error by applying a threshold decision rule to each estimated cell and comparing the result to the baseline classification. Fig. 5 shows a comparison averaged over the 100 simulation runs




Fig. 4. Simulation results of a three agent system performing exact and approximate inference (synchronous and intelligent ROI message passing with $R = 12$) over the region shown in Fig. 3: (a)-(c) at $t = 40$; (d)-(f) at $t = 100$. The agents are depicted as white dots, the initial positions are indicated with white x's, the trajectories are the white dashed lines, and the agent direction of travel is indicated by the white arrows. Note that the agents perform lawn-mower patterned cell traversals.

of inference error and classification error (with a threshold of 0.7) as a function of the percentage of samples taken for the exact method, the approximate method with synchronous message passing, and the approximate method with intelligent message passing with ROI radii of 4, 8, and 12.

A. Discussion

We now remark on the efficacy of the inference algorithms based on our simulation results. With only a small number of model variable observations over the grid (20-30%), the agents are able to infer a reasonably accurate map of the underlying process in a cooperative and distributed way. It is clear that the exact approach is superior in terms of accuracy and smoothness, an unsurprising result. The approximation methods require more samples to generate accurate results and the outcome exhibits less smoothness, especially in the case of the localized ROI message passing scheme. Also note that synchronous message passing generates marginally better solutions than the ROI scheme, where the ROI radius acts to modulate the estimate accuracy.

Given a threshold classification rule, the exact and approximate (synchronous and intelligent with ROI radius of 12) inference methods have 10-15% error rates after the first 25-35% of samples, with the exact method outperforming the

approximate methods on average by only 5%. The strength of the approximation methods, especially ROI message passing, lies in computational complexity. In our simulations the computational burden of exact inference became obvious as grids with $N_p \gg 10$ were rendered infeasible by the exponentially rising cost of graph triangulation, cluster size, and message generation. The approximate methods were simulated on grids orders of magnitude greater in size with computation times approaching those required for real-time feasibility. For example, we have achieved LBP computation rates of 5-10 Hz on 100×100 grids in simulation with minimal optimization. We believe that with the recent explosion in computational power and parallelism our approximate algorithms represent a viable option for efficient spatial mapping in realistic multi-agent deployments.

As inter-agent communication is neighbor-wise and not all-to-all or centralized, we expect the proposed algorithms to scale well as the number of agents increases. The number of inter-agent messages is approximately linear in agent count and the computational burden of injecting increased global beliefs by each agent is negligible in comparison to the local LBP process, which scales only according to grid size. The advantages of agent scaling are clear: large workspaces become manageable due to enhanced spatial distribution,




Fig. 5. Average cell-wise KL-Div and plume membership error rate versus percentage of grid samples taken for the 3 agent system described in Section IV. Plots are shown for the exact method, the approximate method with synchronous message passing, and the approximate method with intelligent message passing with ROI radii of 4, 8, and 12.

measurement capabilities and thus process modeling are heightened through flexibility in agent composition, and the impact of agent failures is mitigated by redundancy.

Finally, the proposed MAMRF formulation also has several qualitative benefits compared to other plume detection/tracking implementations (e.g. [11]–[13]). As opposed to gradient based approaches that require local environmental gradients that are difficult to obtain in practice, our method exploits spatial dependence to achieve results based on scalar observations. In addition, our approach leverages distributed global observations in contrast to methods that only incorporate local information to drive control goals. Finally, our algorithm benefits from the flexibility and robustness of probabilistic outputs versus hard decisions or controls.

V. CONCLUSIONS AND FUTURE WORK

In this paper we studied cooperative inference in multi-agent systems where uncertainty is modeled by the grid structured pairwise Markov random field. The *multi-agent Markov random field* framework was proposed, wherein the global inference problem is decomposed into inter-agent belief exchanges and local intra-agent inference problems. An approximate inference algorithm inspired by loopy belief propagation was chosen due to the exponential complexity of exact inference. An intelligent message passing scheme based on the idea of regions of influence in a cluster graph in LBP was discussed. Finally, we presented simulation results for the inference methods over a spatial grid with the goal of identifying plume-like features in oceanographic data using multiple virtual robots. The results demonstrated the accuracy and smoothness of exact inference but also showed its intractability over large grids. The approximate inference

methods were shown to be a reasonable alternative both in terms of computational efficiency and accuracy. In particular, the intelligent ROI message passing scheme is attractive as it exhibits a lack of dependence on grid size.

Directions for future work include identifying less stringent requirements on agent communication, as maintaining a tree structured communication graph may place significant overhead on the system under realistic conditions. Determining methods of approximation that do not exhibit the convergence issues of the LBP algorithm could also be valuable in safety-critical systems. Finally, investigating varying process and model potentials may prove useful in mapping more complex spatial phenomenon.

REFERENCES

- [1] Y. Tamada, S. Imoto, H. Araki, M. Nagasaki, C. Print, D. Charnock-Jones, and S. Miyano, “Estimating Genome-wide Gene Networks Using Nonparametric Bayesian Network Models on Massively Parallel Computers,” *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 2010.
- [2] P. Pinheiro, A. Castro, and M. Pinheiro, “A Multicriteria Model Applied in the Diagnosis of Alzheimer’s Disease: A Bayesian Network,” in *Computational Science and Engineering, 2008. CSE ’08. 11th IEEE International Conference on*, 2008.
- [3] H. Cao and V. Govindaraju, “Preprocessing of Low-Quality Handwritten Documents Using Markov Random Fields,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 2009.
- [4] Y. Xiang and V. Lesser, “Justifying multiply sectioned Bayesian networks,” in *MultiAgent Systems, 2000. Proceedings. Fourth International Conference on*, 2000.
- [5] —, “On the role of multiply sectioned Bayesian networks to cooperative multiagent systems,” *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 2003.
- [6] J. Butterfield, O. C. Jenkins, and B. Gerkey, *Multi-robot Markov random fields*. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [7] X. Li and S. Bian, “Multiscale Image Segmentation Using Markov Random Field and Spatial Fuzzy Clustering in Wavelet Domain,” in *Intelligent Systems and Applications, 2009. ISA 2009. International Workshop on*, 2009.
- [8] Y. Rachlin, J. Dolan, and P. Khosla, “Efficient mapping through exploitation of spatial dependencies,” in *Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on*, 2005.
- [9] J. Kleinberg and E. Tardos, “Approximation algorithms for classification problems with pairwise relationships: metric labeling and Markov random fields,” *J. ACM*, 2002.
- [10] R. N. Smith, Y. Chao, P. P. Li, D. A. Caron, B. H. Jones, and G. S. Sukhatme, “Planning and Implementing Trajectories for Autonomous Underwater Vehicles to Track Evolving Ocean Processes based on Predictions from a Regional Ocean Model,” *International Journal of Robotics Research*, 2010.
- [11] P. Ogren, E. Fiorelli, and N. Leonard, “Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment,” *Automatic Control, IEEE Transactions on*, 2004.
- [12] A. Bertozzi, M. Kemp, and D. Marthaler, “Determining Environmental Boundaries: Asynchronous Communication and Physical Scales,” in *Cooperative Control*, 2005.
- [13] F. Zhang, E. Fiorelli, and N. Leonard, “Exploring scalar fields using multiple sensor platforms: Tracking level curves,” *Decision and Control, 2007 46th IEEE Conference on*, 2007.
- [14] Y. Xiang, *Probabilistic Reasoning in Multi-Agent Systems: A Graphical Models Approach*. Cambridge University Press, 2002.
- [15] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.
- [16] I. Cetinić, “Harmful algal blooms in the urbanized coastal ocean: an application of remote sensing for understanding, characterization and prediction,” Ph.D. dissertation, University of Southern California, 2009.
- [17] B. Maccherone, “MODIS Website,” 2011. [Online]. Available: <http://modis.gsfc.nasa.gov/>