

Multi-robot Markov Random Fields

(Short Paper)

Jesse Butterfield
Dept. of Computer Science
Brown University
Providence, RI 02912-1910
jbutterf@cs.brown.edu

Odest Chadwicke
Jenkins
Dept. of Computer Science
Brown University
Providence, RI 02912-1910
cjenkins@cs.brown.edu

Brian Gerkey
Artificial Intelligence Center
SRI International
Menlo Park, CA 94025-3493
gerkey@ai.sri.com

ABSTRACT

We propose Markov random fields (MRFs) as a probabilistic mathematical model for unifying approaches to multi-robot coordination or, more specifically, distributed action selection. The MRF model is well-suited to domains in which the joint probability over latent (action) and observed (perceived) variables can be factored into pairwise interactions between these variables. Specifically, these interactions occur through functions that evaluate “local evidence” between an observed and latent variable and “compatibility” between a pair of latent variables. For multi-robot coordination, we cast local evidence functions as the computation for an individual robot’s action selection from its local observations and compatibility as the dependence in action selection between a pair of robots. We describe how existing methods for multi-robot coordination (or at least a non-exhaustive subset) fit within an MRF-based model and how they conceptually unify. Further, we offer belief propagation on a multi-robot MRF as a novel approach to distributed robot action selection.

Categories and Subject Descriptors

I.2.9 [Computing Methodologies]: Artificial Intelligence—Robotics

General Terms

Algorithms, Measurement, Performance

Keywords

Multi-robot Task Allocation, Markov Random Field, Belief Propagation

1. INTRODUCTION

A key challenge in the progress of multi-robot systems is **distributed action selection**: a group of robots works collectively as a team, but each must make local decisions, relying on locally available information. Though multi-robot control is a well-studied area with many existing techniques [8,

Cite as: Multi-robot Markov Random Fields (Short Paper), Jesse Butterfield, Odest Chadwicke Jenkins and Brian Gerkey, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1211-1214.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

3, 1], defining comparisons, relationships, and unifying models between these techniques remains an open problem. Tovey et al. [14] have proposed one perspective on a unifying model through a theoretical analysis of multi-robot routing. We propose a novel approach, in which we view the problem of distributed action selection as an instance of **distributed probabilistic inference**. In our formulation, the goal is to construct a joint probability distribution over a set of latent variables (the robots’ actions) from local evidence (a robot’s own sensing) and pairwise interactions between latent variables (inter-robot task dependencies). Action selection at each robot is construed as its **belief** (as a probability distribution) that is conditionally dependent on its own observations and the beliefs of other robots and is a marginal of the overall joint probability.

We model this system of evidence, belief, and conditional dependence using a Markov random field (MRF), which is a pairwise probabilistic graphical model. We claim the MRF is a natural choice for representing multi-robot problems in a general and consistent fashion that allows for simplifying assumptions and analogies (e.g., physics, economics). Further, our MRF model provides a unifying framework in which we can represent, analyze, and compare many existing multi-robot algorithms.

A multi-robot MRF is a mathematical model that can be solved using a variety of inference algorithms. The result of the inference process is a posterior distribution over possible actions for each robot. In this paper, we discuss how existing approaches to multi-robot control can be modeled as MRFs and how their algorithms perform inference in these models. Casting solutions as MRF solving techniques allows us to leverage the large body of existing MRF research to analyze them. In addition, MRF-based models allow for the straightforward application of probabilistic (loopy) belief propagation (BP) [15]. Using BP, each robot locally computes its belief probabilistically over actions as the product of its local evidence and “messages” from other robots. Messages in this case are distributions over a robot’s actions which could be used to generalize robot communication. Although loopy BP provides no direct guarantees for convergence, BP has been shown to have free energy properties that facilitate convergence when all robots can directly communicate (i.e., all-pairs communication) or when some pairs of robots are out of communication range.

2. MULTI-ROBOT MRFS

An MRF [15] is a graphical model that factors a system

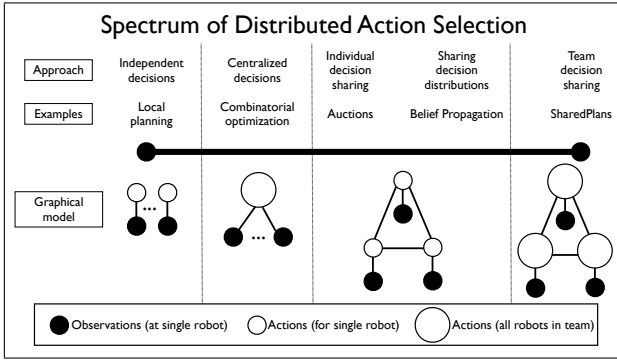


Figure 1: Spectrum of coordination approaches represented by MRF-based graphical models, ranging from acting in isolation (left) to local formation of complete team plans (right).

into a finite set of observed and hidden, or latent, variables with pairwise interactions between them. In a multi-robot MRF, the robots are represented as nodes in the graph, with edges between pairs of robots that are in direct communication range. Each robot i maintains two random variables: y_i , an observed variable representing a robot’s own perception; and x_i , a hidden variable representing the action that it should take. Although these variables can be either discrete or continuous, we assume x_i is a discrete enumeration over the set of possible actions. Further, a latent variable is considered to be a random variable (i.e., a probability distribution) that will assume a single value when the entropy of the distribution is zero.

Consider the example of wireless coverage for a team of robots. In this scenario, robots, each with wireless communication and localization capabilities, work together to provide coverage for a wireless mobile ad-hoc network over a two-dimensional space. In this case, x_i enumerates actions as a discrete set of valid poses and radio transmission power levels for robot i . The action x_i is inferred both from the local observations y_i , containing the current estimated pose of robot i , and in coordination with $x_j | j \neq i$, the action selection of other robots.

Given these variables, a pairwise MRF factors all possible collective team actions x into two functions: **pairwise compatibility** $\psi_{j,i}(x_j, x_i)$ between each robot pair (ij) and **local evidence** $\phi_i(x_i, y_i)$. The joint probability distribution can then be stated as follows:

$$Pr(x) = \frac{1}{Z} \prod_{(ij)} \psi_{j,i}(x_j, x_i) \prod_i \phi_i(x_i, y_i) \quad (1)$$

The normalization constant Z ensures that the distribution sums to 1. The formulation in (1) has two key benefits: we factor the global coordination and local computation into distinct terms; and we can express a spectrum of multi-robot action selection methods by modifying these terms.

The local evidence $\phi(x_i, y_i)$ expresses the likelihood of robot i choosing each of its actions x_i , given its observations y_i . This function is analogous to likelihood models as they are used in Bayes filters [13] for localization and could be instantiated in various forms, from a simple one-step utility estimator to a more sophisticated local planner. The pairwise compatibility $\psi_{j,i}(x_j = a_s, x_i = a_t)$ encodes

the likelihood of robots i and j selecting actions $a_s \in x_i$ and $a_t \in x_j$, respectively, from their combined action space $x_j \times x_i$.

If we ignore pairwise compatibility (i.e., exchange no messages between robots), then each robot will compute its belief (i.e., distribution over actions) and select actions independently, solely based on its own observations. In previous work [2], we have shown that this purely local approach, though naive, can produce good performance in domains similar to wireless coverage. However, it is unclear how these local methods perform when noise and uncertainty are present.

3. UNIFICATION THROUGH MRFS

Multi-robot task assignment Gerkey and Mataric characterized multi-robot task assignment (MRTA) problems as instances of problems in optimization [3]. The simplest version of MRTA, with robots capable of a single task, each task taking only a single robot, and conducting only instantaneous assignment rather than creating time-extended plans (the problem denoted ST-SR-IA in [3]), can be converted into a Markov random field in which the hidden nodes are completely connected (third from the left in Figure 1) and with evidence and compatibility functions defined as follows:

$$\phi_i(x_i, y_i) = \frac{e^{U_i(x_i, y_i)}}{\sum_{x_i} e^{U_i(x_i, y_i)}} \quad (2)$$

$$\psi(x_i, x_j) = \begin{cases} 0 & \text{if } x_i = x_j \\ 1 & \text{otherwise} \end{cases}$$

where $U_i(x_i, y_i)$ is the utility of the robot for a particular task, x_i , based on its local information, y_i . Using the properties of the products of exponents it can be shown that maximizing the joint probability given by (1) in this MRF is equivalent to maximizing the sum of the utilities under the restriction that no two robots may do the same task. That is, the maximum a posteriori assignment that can be computed for this MRF is an optimal solution to the underlying assignment problem. It has been shown [3] how linear programming techniques, such as the Hungarian method [6], compute an optimal solution to this problem in $O(mn^2)$ time, where m is the number of robots and n is the number of tasks. We can thus conclude that for the family of MRFs described by (2), with m nodes and n states, the maximum a posteriori solution can be found in $O(mn^2)$ time, by first converting to the equivalent linear programming representation.

A significant advantage of our MRF model over the optimization view is that the MRF model allows richer, more complex compatibility functions. Importantly, we need not assume that a robot’s utility for a task is independent of which other tasks are being executed. In many real-world problems, especially those involving a component of spatial arrangement, the utility of a set of tasks depends not just on the utilities of the individual tasks, but also on the tasks’ compatibility. For example, in the wireless coverage problem, the compatibility of two robots’ tasks will gradually increase as the robots move away from each other, but then fall off quickly as they move out of communication range. Another advantage of the MRF model is that it allows us to model uncertainty in our utility estimates. For example, by smoothing our local evidence function we can include a

robot’s uncertainty about its current utility estimates into the decision making process.

While the MRF model is best suited for describing the ST-SR-IA variety of MRTA problems, it can be extended to all variations of MRTA. For the multi-task robot and multi-robot task variations, the adaptation can be done by increasing the state space of the nodes. For the multi-task cases the state space of each node becomes the combinations of tasks a robot is capable of, and for the time-extended problems the state space comprises ordered lists of such tasks. The compatibility function is binary: zero when the sets of tasks for two robots overlap and one otherwise. Unfortunately when robots are capable of many tasks or are planning over a long series of tasks, the state space is expanding exponentially. The multi-robot tasks case is difficult to directly cover with our model because compatibility between tasks can no longer be defined as a function of only two nodes. Multi-robot tasks can be described using the SharedPlans framework we discuss later in this section. There are ways to fit it into the MRF model, but they are complicated and require sacrificing some of the advantages of the model.

Auction algorithms Task-swapping auction algorithms [1] can be rephrased into the MRF structure with the functions given in the last section because they are applied to MRTA problems. Tasks are originally assigned in a greedy or a random fashion, but then a subset of the actual joint probabilities are calculated and tasks are exchanged if any superior combinations are found. Calculating all the possible combinations is NP-hard and amounts to the brute force method for solving an MRF. However, by intelligently choosing subsets of the joint probabilities to explore, we can find larger probabilities (better solutions) than by simply using the greedy algorithm alone. The more combinations that are explored, the larger the computational time, but the solution becomes closer to the optimal. For example, the two-party, single task exchange calculates the joint probability of two nodes i and j with tasks a and b after exchanging tasks. If $p(x_i = a, x_j = b) < p(x_i = b, x_j = a)$ then the tasks are exchanged, since the switch will increase the total joint probability. Exchanges are made until no more increases can be found. In auction systems with leaders, there are clusters of nodes for which the leader can evaluate a larger set of the combinations and exchange the values of the nodes to optimize over the whole cluster.

SharedPlans SharedPlans [4] is another way of thinking about planning for multi-robot tasks. Each robot plans for its individual actions and also plans for the actions of other robots that are required to make its individual plan effective. In this case, the local evidence function corresponds to the probability of a particular plan being optimal, where a plan includes a robot’s individual plans and its plans for others. The compatibility function between two plans would then be binary: one if actions in the first robot’s plan for the second robot were included in the second robot’s individual plan and zero otherwise. The joint probabilities will go to zero for any poorly coordinated actions and will be functions of the sum of the utilities when plans are compatible.

Centralized vs. distributed coordination Many of the current algorithms for solving multi-robot coordination problems involve having a central planner or having each robot calculate the total control policy in parallel. These methods correspond to one or each robot building a map of the entire MRF and generating optimal solutions. With al-

gorithms such as belief propagation, marginal probabilities can be calculated locally. Other techniques can be a combination of local and distributed control, such as some auction techniques which require an auctioneer to distribute tasks originally but then allow robots to trade amongst themselves.

Completely connected vs. partially connected In many robot task selection problems, the graph should theoretically be completely connected, because the robots’ task selection depends directly on all other robots. However, in problems where spatial location is relevant, two robots that are unable to communicate are unlikely to be conducting strongly dependent tasks, and so connecting only those robots that are in communication gives a good approximation and allows us to perform inference without using relays to get messages between two unconnected robots. This can also reduce the number of edges in the MRF, accelerating the computation of solutions.

4. MULTI-ROBOT BELIEF PROPAGATION

To perform action selection with a multi-robot MRF, each robot must compute its “piece” of the joint probability distribution (1). That is, for each robot i , we want to compute the marginal probability $p_i(x_i)$, which expresses the likelihood of robot i taking an action, given both its own observations and knowledge of the other robots’ actions. Naively, this inference procedure would require communicating all robots’ observations to a centralized decision-maker. Instead, we will exploit the factored structure of the MRF to apply the BP algorithm [15], which performs inference in a distributed manner.

BP operates by passing “advice messages” between robots, and using these messages, in combination with local observations, to maintain a belief $b_i(x_i)$ for each robot i . When BP converges, the belief $b_i(x_i)$ is exactly equal to the marginal probability $p_i(x_i)$ that we need for coordinated action selection. Robot i exchanges messages only with its neighbors $N(i)$, ensuring that the algorithm can scale to large teams. Robot i ’s belief $b_i(x_i)$ is given by the following product, with normalization constant Z :

$$b_i(x_i) = Z \phi_i(x_i, y_i) \prod_{j \in N(i)} m_{j,i}(x_i) \quad (3)$$

The term $m_{j,i}(x_i)$ is an advice message from robot j to robot i suggesting how robot i should act, given what robot j knows about the world. Robot j computes its message to a neighboring robot i as a product of robot j ’s local evidence, the pairwise compatibility between their actions, and incoming messages from robot j ’s neighborhood (except for those coming from robot i), summed over all possible actions x_j :

$$m_{j,i}(x_i) = \sum_{x_j} \phi_j(x_j, y_j) \psi_{j,i}(x_j, x_i) \prod_{k \in N(j) \setminus i} m_{k,j}(x_j) \quad (4)$$

Multi-robot belief propagation is performed through selective updating of inter-robot messages. The essence of inference with belief propagation lies in the pairwise messages $m_{j,i}(x_i)$. Messages between robots are continually but selectively updated according to (4) and domain-specific compatibility and likelihood functions. Beliefs are computed only when needed to make a decision.

Selecting actions The result of multi-robot belief propagation will be an *action posterior*, a probabilistic belief dis-

tribution over actions for each robot. Whenever an individual robot needs to make a new action decision, it computes its action posterior as a belief from local evidence and incoming messages, according to (3). This belief distribution can then be sampled (e.g., using expectation or MAP estimators) to select a specific action for the robot to execute.

Convergence Because our multi-robot MRFs will most often contain cycles, we must address the issue of convergence. Strictly speaking, the beliefs computed by BP are guaranteed to converge to the true marginal probabilities only when the graphical model is tree-structured [10]. Indeed, pathological graphs can be constructed in which BP fails to converge [7]. However, BP has been successfully applied to a wide array of problems that feature cyclical graphs. Recent analysis has drawn a connection between BP convergence and free energy approximations from physics [15]. While we cannot guarantee convergence to optimal action selection, these approximations provide an analytical foundation showing why convergence of BP in cyclical graphs is typical, even in noisy and ambiguous situations. Recent approaches have demonstrated some success using BP for multi-robot chain-of-sight and group behaviors [11, 12] and for static sensor networks using loopy BP [5] and junction trees [9].

4.1 Wireless Coverage Example

The key to applying our multi-robot belief propagation approach is the appropriate definition of local evidence and joint compatibility functions. In this section we present candidate definitions for the wireless coverage problem introduced earlier.

Evidence functions The coverage evidence function $\phi_i(x_i, y_i)$ could compute the expected contribution to the overall network of each of robot i 's available actions (positions to move to and radio power levels to select). Of course, this function is a domain-dependent policy combination of terms that encode the user requirements, connectivity goals, and energy limits.

One possible local evidence function could be a product of three terms, each of which expresses a preference over robot i 's actions:

$$\phi_i(x_i, y_i) = H(x_i, y_i)C(x_i, y_i)E(x_i, y_i) \quad (5)$$

The term $H(x_i, y_i)$ encodes a preference for actions that provide connectivity (whether single- or multi-hop) to the greatest number of human users or gateways. This term is the foundation for establishing good coverage of users. Similarly, the term $C(x_i, y_i)$ prefers actions that provide connectivity to the other robots. The term $E(x_i, y_i)$ expresses a preference for actions that expend the least energy, incorporating both movement cost and radio transmission cost. This term could be derived from platform- and application-specific characteristics, including motor power requirements, radio power requirements at each transmission level, battery level, and expected lifespan of the network.

Compatibility functions The compatibility $\psi_{i,j}(x_i, x_j)$ will act as a "local coordination" factor, determining the value of joint actions for robots i and j as a function of the resulting pairwise signal strengths between them (recall that BP messages, and thus compatibility functions, are only defined for robots that are in direct communication):

$$\psi_{i,j}(x_i, x_j) = S(x_i, x_j) \quad (6)$$

The term $S(x_i, x_j)$ prefers joint actions that provide "good enough," but not maximal, signal quality between robots i and j . If the robots were to simply maximize their pairwise signal strengths, they would tend to cluster very close together, leading to isolated network components, and providing poor spatial coverage. Such tendencies could be counteracted by defining the term $S(x_i, x_j)$ to prefer signal levels that are high enough to provide the necessary quality of service, but not any higher. Ideally, our planned experimental trials on the wireless coverage problem will demonstrate the effectiveness of using the Markov random fields framework.

5. REFERENCES

- [1] M. B. Dias, R. M. Zlot, N. Kalra, and A. T. Stentz. Market-based multirobot coordination: a survey and analysis. *Proceedings of the IEEE*, 94(7):1257–1270, July 2006.
- [2] B. P. Gerkey, R. Mailler, and B. Morisset. Commbots: Distributed control of mobile communication relays. In *Proc. of the AAAI Workshop on Auction Mechanisms for Robot Coordination*, Boston, Massachusetts, July 2006.
- [3] B. P. Gerkey and M. J. Mataric. A formal analysis and taxonomy of task allocation in multi-robot systems. *The Intl. J. of Robotics Research*, 23(9):939–954, Sept. 2004.
- [4] B. Grosz and S. Kraus. Collaborative plans for complex group action. *AI*, 86(2):269–357, 1996.
- [5] A. T. Ihler, J. W. Fisher III, R. L. Moses, and A. S. Willsky. Nonparametric belief propagation for self-calibration in sensor networks. *IEEE Journal of Selected Areas in Communication*, 23(4):809–819, 2005.
- [6] H. W. Kuhn. The Hungarian Method for the Assignment Problem. *Naval Research Logistics Quarterly*, 2(1):83–97, 1955.
- [7] K. Murphy, Y. Weiss, and M. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proc. of the Conf. on Uncertainty in AI (UAI)*, 1999.
- [8] L. E. Parker. ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, Apr. 1998.
- [9] M. Paskin, C. Guestrin, and J. McFadden. A robust architecture for distributed inference in sensor networks. In *Information Processing in Sensor Networks (IPSN'05)*, 2005.
- [10] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., 1988.
- [11] J. Schwertfeger and O. Jenkins. Multi-robot belief propagation for distributed robot allocation. In *Proc. of the IEEE Intl. Conf. on Development and Learning*, London, England, 2007.
- [12] D. A. Shell and M. J. Mataric. Principled synthesis for large-scale systems: task sequencing. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems*, pages 207–216, Minneapolis/St. Paul, Minnesota, USA, Jul 2006.
- [13] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, September 2005. ISBN 0-262-20162-3.
- [14] C. Tovey, M. Lagoudakis, S. Jain, and S. Koenig. The generation of bidding rules for auction-based robot coordination. In L.E.Parker et al., editors, *Multi-Robot Systems: From Swarms to Intelligent Automata, Volume III*, pages 3–14. Springer, the Netherlands, 2005.
- [15] J. S. Yedidia, W. T. Freeman, and Y. Weiss. *Exploring Artificial Intelligence in the New Millennium*, chapter Understanding Belief Propagation and Its Generalizations. Morgan Kaufmann, 2001.