

A Markov Random Field Model of Context for High-Level Information Fusion

Robin Glinton, Joseph Giampapa, Katia Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, U.S.A.
(rglinton, garof, katia)@cs.cmu.edu

Abstract—*This paper presents a method for inferring threat in a military campaign through matching of battlefield entities to a doctrinal template. In this work the set of random variables denoting the possible template matches for the scenario entities is a realization of a Markov Random Field. This approach does not separate low level fusion from high level fusion but optimizes both simultaneously. The result of the added high level context is a method that is robust to false positive and false negative, or missed, sensor readings. Furthermore, the high level context helps to direct the search for the best template match. Empirical results illustrate the efficacy of the method both at identifying threats in the face of false negatives, and at negating false positives, as well as illustrating the reduced computational effort resulting from the incorporation of additional high-level context.*

Keywords: Data association, Markov Random Fields, situation assessment, intent inference.

1 Introduction

In a battlefield environment it is critical that intelligence analysts develop hypotheses about the high-level intent of the enemy. Examples of high-level enemy intent include the intent to retreat or to ambush. Intelligence analysts form such hypotheses by fusing sensor reports with information about the terrain and enemy doctrine. This anticipation of enemy intent allows military planners to allocate resources and to develop counter strategies in order to achieve mission objectives while minimizing losses.

A battlefield environment is high dimensional and noisy. Fusing heterogeneous sources of information to draw useful conclusions about enemy intent in such an environment is extremely difficult even for human beings. Furthermore, sensors are becoming smaller and cheaper. In a network centric battlefield the problem of information overload will become increasingly prevalent and sig-

nificant. The problem of sensor false positives and negatives too will be exacerbated in a sensor rich environment. Long delays in decision cycles result from these problems which in turn have a great cost in terms of human lives. Automating parts of this inference can significantly speed up the decision cycle. An automated analysis can also take advantage of data at a much higher granularity.

Context is critical for information fusion at all levels. A common task in low-level information fusion is determining sensor false positives. Little can be said in terms of discrediting a sensor reading that suggests that a tank is alone in the middle of an open field. However, if that reading suggests that the tank is near other tanks of an appropriate type with a similar heading and velocity this lends credence to it. Other factors like proximity to key terrain features help to disambiguate as well. Context is also invaluable in high-level information fusion. Noting temporal relationships between the activity of spatially separated military units, for example, is an important source of context in determining the intent of the larger force.

We formulate intent inference as a labeling problem. The entities to be labelled are enemy vehicles, friendly vehicles, events, and key terrain features. The labels are drawn from a template that represents a specific enemy plan taken from doctrine. A plan template is modeled as a relational structure [1], a type of graph where the nodes represent roles that enemy vehicles can fulfill as well as key terrain features and significant events. The edges represent contextual relationships between the nodes. For example in an ambush scenario one of the roles might be to be a part of a unit that attacks from the front of an unsuspecting adversary unit while another role could be to be a part of a unit that waits near a suitable piece of terrain in the rear in case the unit under attack attempts a retreat. The contextual relationships in the previous example include relative position and velocity with the respect to other units filling the same role or temporal relationships between significant events. Each entity and event in the

scenario has an associated random variable that takes on the discrete labels from the template. Furthermore, there is a neighborhood system which defines a gating criteria on which entities are related. The set of aforementioned random variable forms a Markov Random Field (MRF). The equivalence between an MRF and a Gibbs distribution [2] allows us to formulate likelihoods and priors in a Bayesian sense in terms of clique potentials which serve to allow the modeling of *a-priori* knowledge about the contextual interactions between labels in the template and hence entities and events in the scenario. We can then find the maximum a posterior (MAP) labeling using combinatorial optimization.

The aforementioned MAP-MRF formulation has proven effective for template matching in high-dimensional, noisy search spaces, particularly in the computer vision and image processing communities [3]. Furthermore, because it is a probabilistic model uncertainty in sensor estimates as well as abstract relationships can be modeled simultaneously. High level context too has been successfully modeled using MRFs for low-level fusion. MRFs, for example, have been used for data association in target tracking [4]. In [5] we developed algorithms for automatic terrain analysis to provide high level terrain context. In [6] we presented algorithms to describe the behavior of battlefield entities in terms of environmental influences. This paper describes a method for uniting these sources of context to provide high-level threat estimates.

The degree of uncertainty in a battlefield environment make statistical models such as Dynamic Bayesian Networks a popular choice for inference on a state space of discrete and continuous variables representing a battlefield. However, as is a general drawback of all statistical methods, in a high dimensional state space the combinatorially large number of possible interactions between variables make it infeasible to calculate conditional distributions. Many approaches resort to particle filtering which approximates the conditional distributions using a weighted sum of state variable samples to reduce the computational complexity of an inference [7],[8]. This approach is does not effectively leverage the strong contextual relationships between the variables that represent a battlefield. By explicitly modeling this context using Markov Random Fields (MRF) the contextual relationships between entities and hence the variables that represent them directly inform which calculations are necessary for a successful inference. This is similar to human reasoning in that instead of sampling a state space based on statistical dependencies between variables as in the case of particle filtering, an MRF approach samples based on configurations of variables that match a high level pre-description of how the system being modeled should work.

That is it concentrates most of the computational effort in the parts of the state space that prescribe logical inferences based on contextual dependencies between variables. This is particularly effective in a battlefield environment because successful military operations require careful organization. Markov Random Fields also make it relatively easy to represent such relationships both temporal and spatial. Military planners use terrain features to organize forces, as reference when interpreting enemy intent. Much of the intent inference literature marginalizes or ignores terrain. Some use a simplified terrain representation. We use a realistic model of terrain and rely on an automated terrain analysis to segment the terrain into high-level features suitable for use in the inference process.

The remainder of the paper is organized as follows. In section 2 the problem is formally defined. In section 3 the formal definition of a Markov Random Field is given. Section 4 illustrates how contextual dependencies both between scenario entities and between roles in the template are modeled as separate relational structures. Section 4.1 gives a derivation of the posterior probability for a given match between the scenario entities and the template in terms of clique potentials derived from the aforementioned unary and binary relationships. Section 4.2 describes how the optimal posterior is found using combinatorial optimization to give the MAP solution for the best match between entities and the template. Finally Section 5 provides empirical results illustrating the efficacy of the model.

2 Problem

The following provides a formal description of the intent inference problem as an instance of template matching. We view a battlespace as a set of entities and events each of which is represented as a vector containing both discrete (e.g. entity id) and continuous (e.g. velocity) variables. The intent inference process is formulated as a labeling problem where labels are taken from a template representing military doctrine. Formally: The significant events in the scenario are indexed by the set $E = \{1, \dots, m\}$. Battlespace entities are indexed by the set $B = \{1, \dots, n\}$. Let $G = E \cup B$. The set of possible labels contained in the template is given by $L = \{0, 1, \dots, M\}$ where the *NULL* match 0 is reserved for sensor false positives and entities not viewed as conforming to the behavioral model represented by the template. Section 4 describes the template representation in detail. An inference of enemy intent is then a labeling or mapping f of the battlespace entities and events in terms of the template labels:

$$f : G \rightarrow L$$

Let V be a real valued functional that assigns to every such mapping an $r \in R$ that indicates the quality of the match. Then an inference requires finding a mapping f' such that:

$$f' = \operatorname{argmin}_f V(f)$$

3 Markov Random Fields

Markov random field (MRF) theory enables the modeling of contextual dependencies between a set of sites S . These sites might be pixels in an image, individuals in a social network, or battlespace entities and events for example. The $i \in S$ are related to each other by a neighborhood system $N = \{N_i \mid \forall i \in S\}$ where N_i is the set of neighboring sites to N . The purpose of N_i is to provide a gating criteria for the set of sites that have a relationship to the site i . In the vision example N_i might be the set of all pixels within a certain euclidean distance of the pixel represented by i . A pair-wise clique C_2 defined on N and S is given by:

$$C_2 = \{\{i, i'\} \mid i' \in N_i, i \in S\}$$

Cliques of arbitrary order can be defined. Pre-defined cliques are effective for modeling the contextual dependencies for certain image properties, for example edges. This rigid pre-definition of a clique is unsuitable for our model so we choose cliques dynamically depending on gating criteria determined by the properties of the entity or event represented by i . A family of random variables $F = \{F_1, \dots, F_m\}$ with a specific configuration $f = \{f_1, \dots, f_m\}$ is said to be Markov random field on S with respect to N subject to the following conditions:

$$P(f) > 0, \forall f \in F$$

$$P(f_i \mid f_{S-\{i\}}) = P(f_i \mid f_{N_i})$$

where f_i denotes that F_i takes on value f_i at site i . The second criteria is the Markov property for a random field and states that the probability of a certain configuration at a site i is statistically independent of the configurations of all other $i \in S$ given the configurations at $i \in N_i$.

Specifying the joint probability $P(f)$ for a Markov random field is in general intractable. However, an equivalence between an MRF and a Gibbs distribution proved by Hammersley and Clifford [2] gives us an alternate means to specify $P(f)$ by reformulating F as a Gibbs distribution.

The possible configurations f of a set of random variables F obey a Gibbs distribution if the joint takes the form:

$$P(f) = Z^{-1} \times e^{\frac{-1}{T}U(f)}$$

where Z is a normalizing constant given by:

$$Z = \sum_{f \in F} e^{\frac{-1}{T}U(f)}$$

and U is given by:

$$U(f) = \sum_{c \in C} V_c(f)$$

U is an energy function that is determined by a sum of clique potentials $V_c(f)$ over all cliques. Clique potentials serve to allow the modeling of *a-priori* knowledge about the contextual interactions between labels at neighboring sites. The clique potentials used in this investigation are developed in Section 4.1. The energy $U(f)$ for a particular configuration of F corresponds to the probability of that configuration. Lower energies correspond to more likely configurations. The parameter T is often referred to as a temperature and controls the sharpness of the distribution. Calculation of the partition function Z is intractable even for relatively small problems. However in the MAP formulation where we seek a configuration f of F that optimizes $P(f)$, calculation of Z is unnecessary.

4 Template Representation

A doctrinal template is a prescription for a military plan used as a reference by Intelligence analysts to interpret incoming battlefield data. For the example of an ambush a doctrinal template might contain the preferred vehicle emplacement, for a unit waiting in ambush, with respect to other units both friendly and adversarial. Such a template might also prescribe the preferred direction of attack and velocity for an attacking unit. We formulate the inference of the intent of a military unit as labeling it relative to a doctrinal template. This might involve finding the best match for the unit to several templates. In this investigation we use only a single template. However the method is easily extended to incorporate several templates.

We model a doctrinal template as a relational structure (RS) [1] as shown in figure 1.

In an RS nodes represent entities which can be both physical entities and events. Each node has an associated vector of unary properties. For the node i the vector of associated unary properties is given by $d_1(i) = [d_1^1(i), \dots, d_1^{k_1}(i)]^T$. For example, if the node represents a tank then one of the elements in its unary properties vector is an *id* for the tank, another might be a confidence




Figure 1: Relational graph showing unary, binary, and triple contextual relationships between nodes.

measure in the identification. Each pair of nodes in the RS has an associated vector of binary relations given by $d_2(i, i') = [d_1^1(i, i'), \dots, d_1^{k_1}(i, i')]^T$. This can be generalized arbitrarily to vectors of n -ary relationships between groups of n nodes. In this investigation we limit our relational structures to contain binary relationships. A neighborhood system N identifies which nodes in the RS are connected by an edge. We represent the battlespace events and vehicles as an RS G with neighborhood system N and unary and binary relations d . In addition we represent the doctrinal template as an RS G' with neighborhood system N' and unary and binary relations D . Then our inference process is reduced to finding the optimal mapping between the battlespace entities and the template:

$$f : G(B, N, d) \rightarrow G'(L, N' D)$$

The following is a description of the nodes as well as unary and binary context that we use in defining a relational structure that represents a doctrinal template. An example of a relational structure representing an ambush using these contextual relationships is given in Figure 2.

In an RS representing a doctrinal template nodes can represent entities or events. The entities can be either key terrain features [6] or vehicles. Events include an abrupt change in velocity of a vehicle, a radio transmission, or a vehicle firing on another. Unary context includes entity ids and related confidence measures as well as symbolic relationships which denote the type of entity that can match a template node. For example we don't expect an event to match a node that represents a terrain feature in the template. Binary contextual relationships include, relative position, velocity and heading between vehicles, temporal distance between events. Modeled too is the goal relationship between a moving vehicle and pos-

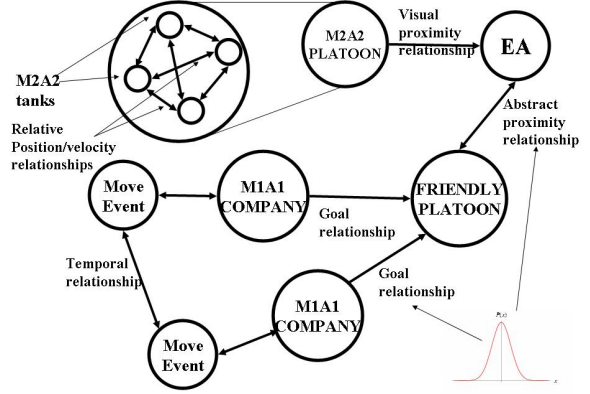


Figure 2: Relational graph representing a doctrinal template of an ambush.

sible target. This relationship is modeled using a potential field technique [6]. In addition abstract relationships like *near* and *soon* can be modeled with an appropriate choice of distributions.

4.1 Derivation of the Posterior Energy

The following summarised from [1] is a derivation of the posterior energy $U(f | d)$ for f a particular labeling of the scenario in terms of the template. We assume that the unary and binary contextual relations observed for and between scenario entities respectively, d_1, d_2 , are a result of the addition of independent zero mean Gaussian noise to the unary and binary relations, D_1, D_2 , given in the template. For a unary relation this gives:

$$d_1(i) = D_1(f_i) + e_1(i)$$

where e is independent zero mean Gaussian noise. Similarly for a binary relation.

$$d_1(i, i') = D_1(f_i, f_{i'}) + e_2(i, i')$$

This assumption is generally invalid but we have found that with an appropriate choice of parameters this assumption produces satisfactory results. The prior energy is given by:

$$U(f) = \sum_{i \in S} V_1(f_i) + \sum_{i \in S} \sum_{i' \in N_i} V_2(f_i, f_{i'})$$

where V_1 and V_2 are constant penalties on the *NULL* match for f_i and $f_{i'}$. This allows for modeling of prior knowledge of the number of false positives expected for a given sensor. The likelihood potential is given by:

$$U(d | f) = \sum_{i \in S, f_i \neq 0} V_1(d_1(i) | f_i) + \sum_{i \in S, f_i \neq 0} \sum_{i' \in S - \{i\}, f_{i'} \neq 0} V_2(d_2(i, i') | f_i, f_{i'}) \quad (1)$$

where the clique potentials are defined as:

$$V_1(d_1(i) | f_i) = \begin{cases} \sum_{k=1}^{K_1} [d_1^{(k)}(i) - D_1^{(k)}(f_i)]^2 / \{2[\sigma_1^{(k)}]^2\} & f_i \neq 0 \\ 0 & f_i = 0 \end{cases} \quad (2)$$

$$V_2(d_2(i, i') | f_i, f_{i'}) = \begin{cases} \sum_{k=1}^{K_2} [d_2^{(k)}(i, i') - D_2^{(k)}(f_i, f_{i'})]^2 / \{2[\sigma_2^{(k)}]^2\} & i \neq i', f_i \neq 0, f_{i'} \neq 0 \\ 0 & f_i = 0 \end{cases} \quad (3)$$

Symbolic unary relations are handled by assigning $\sigma = \infty$. To denote for example that a tank is not an event as opposed to say a confidence measure which can be modeled with a finite σ . The posterior energy is given by:

$$U(f | d) = U(f) + U(d | f)$$

The posterior energy gives a measure of the degree of match to the template for a given labeling of the events and entities in the scenario. This then becomes the objective function in a combinatorial optimization process to find the best labeling in the MAP sense. The following section explains the simulated annealing algorithm that we use to perform the combinatorial optimization.

4.2 Simulated Annealing

Simulated Annealing is a method for combinatorial optimization [9]. The algorithm simulates the physical annealing process of a solid. If we have a posterior energy measure U for a configuration f of a set of discrete random variables F simulated annealing can be used to find an f_i that gives a minima in U . The algorithm starts by randomly sampling the search space of possible configurations. In successive iterations new configurations f_{n+1} are generated by randomly perturbing the previous configuration f_n . The change is accepted with probability $P = e^{-\Delta U/T}$. Where ΔU is the difference in posteriors

evaluated at f_n and f_{n+1} respectively. T is a temperature parameter that starts out high and is lowered according to a schedule that is designed to produce desired results. This is repeated until the value of the function being optimized reaches a steady state. This algorithm is particularly good at avoiding becoming trapped in local minima unlike many gradient descent methods. The cooling schedule we use in this investigation is $T_{n+1} = 0.99T_n$.

4.3 Supporting Algorithms

4.3.1 Data pre-processing

The first step in the template matching algorithm is the pre-processing of sensor data. All of the sensor inputs are processed to detect pre-defined significant events. Possible events are outlined in Section 4. There are three main types of sensor inputs modeled. The first is a simulated Synthetic Aperture Radar (SAR) scan of the battlefield with output processed by a simulated automatic target recognition (ATR) system. The output of the ATR is a confusion matrix with labels of entity types and confidence values which follow a distribution as indicated in the confusion matrix described in [10]. The second sensor input is a simulated Ground Moving Target Indicator (GMTI) which produces tracks for moving entities. Finally detection of radio frequency transmissions is modeled with the start time and source of transmissions input to the template matching algorithm.

Next the raw terrain data elevation, soiltype, etc. is subjected to an automatic terrain analysis, as described in [5], that identifies, segments, and labels militarily significant terrain. This step serves to provide high-level context for a template match based on terrain conditions. For example if key terrain is identified by the terrain analysis as an engagement area then this can be used to match the template for an ambush.

5 Experiments

The following scenario was used for our experiments. Figure 3 shows the scenario as configured in the OTBSaf simulation environment before execution. OTBSaf was chosen as a simulation environment because it models the kind of variability that one would expect in a battlefield environment. For example, during the execution of a scenario military units break formation to respond to terrain features. Furthermore, each individual was represented separately in the doctrinal template to give the algorithm a problem of sufficiently high dimension while keeping the scenario relatively simple to simulate multiple times.

The Area of Operations (AO) is a $50km \times 50km$ area of unpopulated terrain. Two blue force platoons are em-

placed on the map, identified by squares in Figure 3. In




Figure 3: Screenshot of scenario under test as simulated in OTBSaf.

addition, four red force companies of 14 M1A1 tanks each, circled in Figure 3, and a single red force platoon of four M2A2 tanks, signified by a triangle, are conducting maneuvers throughout the AO. Two of the companies are on a reconnaissance mission. The other two companies and the lone platoon are conducting an ambush as prescribed in the doctrinal template given in section 4. In total 56 individual tanks make up the red force. At the start of the scenario the two red force reconnaissance companies begin moving. At ten simulated minutes later, one of the attack companies begins moving along a route towards the blue force platoon being ambushed. Simultaneously the lone red force platoon begins occupying a position with oversight of an engagement area situated along the lone escape route for the blue force platoon under siege. Finally a radio transmission is sent from the first attack company to a second attack company three simulated minutes later. Upon receipt of the transmission a negligible amount of time later, the second attack company begins following a route towards the blue force platoon under threat. The scenario was intentionally designed to be symmetrical to increase its ambiguity. That is, it was designed with two possibilities of friendly platoon under threat and four enemy companies to choose the two attacking companies from.

Two experiments were performed using the data obtained during a fifteen (simulated) minute time window of an execution of the above scenario. The purpose of these experiments was twofold. The first was to show that an MRF approach would produce accurate template matches when confronted with a high degree of sensor false positives and false negatives. The second was to confirm our hypothesis that in the face of a high-dimensional noisy

state space, inter-variable context would help to direct the computational effort. In the experiments the false positive and false negative rates of the sensors were varied and the number of correct template matches collected. The inference required matching the scenario entities to the template given in section 4, that is to correctly identify the friendly platoon under threat and to identify the two attacking enemy companies. In all experiments false negatives were generated by uniformly randomly negating a percentage of the entities sensed. When false positives were generated, the identities, positions, and velocities reported with the false detections were chosen randomly from a uniform distribution. The range of the magnitudes for these velocities were taken between the maximum and minimum for a particular vehicle type according to the random choice of vehicle identity.

In experiment *A* the false negative rate of the sensors was varied between 0 – 90% in increments of ten and the false positive rate was set to zero. For each data-point thirty trials were run. A trial consisted of a single inference. That is, one instance of optimizing the posterior match between scenario entities and the template using simulated annealing. The simulated annealing parameters were kept constant throughout all experiments in an effort to measure the context modeling performance of the MRF template independently. The annealing parameters used were those reported in Section 4.2. For the first experiment the time window was chosen arbitrarily near the midpoint of scenario execution.

In experiment *B* the false negative rate for the sensors was set to zero and the false positive rate was varied between 0 – 50%. The other experimental parameters were the same as in experiment *A*. Experiment *B* was run twice, once with a time window the same as in experiment *A* and once with a time window near the beginning of the scenario so that the added context of the relative start times between the enemy units could be used.

5.1 Model Performance

The results of experimental configuration *A* is shown in Figure 4. This graph shows the percentage of scenario entities correctly assigned to the template vs the percentage of the total scenario entities present that the sensors actually reported. The curve with the dashed lines shows the accuracy of the template match given by the highest probability labeling found in the thirty trials. The solid curve shows the average accuracy among all labelings found in the thirty trials. From the dashed curve it is evident that even with a very high sensor miss rate (with the sensors only reporting 10% of the scenario entities) the percentage of enemy vehicles correctly matched to the template by the highest probability match is still above 80%. This




Figure 4: Graph showing the percentage of scenario entities correctly assigned to the template vs the percentage of scenario entities that the sensors reported. The dashed curve gives the accuracy of the highest probability match found in 30 trials. The solid curve gives the average accuracy of all of the matches found over the 30 trials.

is probably due to the fact that the negated sensor readings were chosen at random. This meant that for most trials the algorithm was presented with at least some vehicles from every unit involved in the template. This likely allowed for strong inter-unit context even when relatively few vehicles were detected. However, one would expect that in this case there would be relatively little intra-unit context in terms of identifying vehicles as belonging to the same unit. However, intra-unit ambiguity is more likely to be caused by false positives which were not considered in this graph. Another possible source of intra-unit ambiguity is when two neighboring units are spatially close enough that it is difficult to resolve them. The main source of intra-unit ambiguity in the scenario under test was due to false positives.

5.2 Influence of Additional Context

In this section we investigate the effect of additional context on the computational effort required to find the MAP-MRF match between the battlefield data and the doctrinal template.

The results of experimental configuration B are given in figures 5 and 6. In Both figures the independent variable is given by the percentage of entities correctly matched to the template, in this case including false positives being labelled with the *NULL* label. The dependent variable is the percentage of entities reported by the sensors that are false positives. For example, a value of 50% for the dependent variable means that the sensors reported a

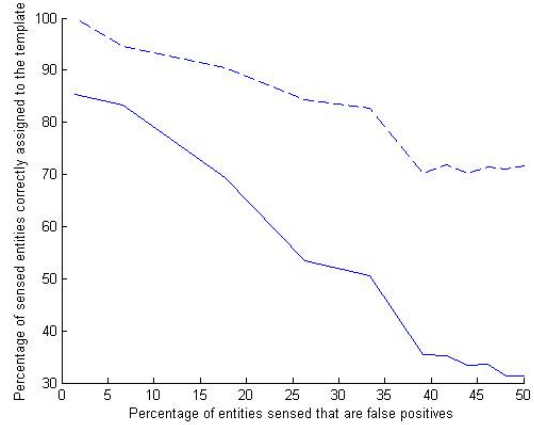


Figure 5: No temporal context - Graph showing the percentage of scenario entities correctly assigned to the template vs the percentage of scenario entities that the sensors reported that were false positives. The dashed curve gives the accuracy of the highest probability match found in 30 trials. The solid curve gives the average accuracy of all of the matches found over the 30 trials.

total of 140 entities, 70 of which do not exist. Figure 5 shows the results of the template matching algorithm when the context of temporal coordination between red force units was not considered, while Figure 6 shows the results with this additional context included. From Figure 5 we can tell that the labeling of the scenario that corresponds with the highest probability labeling found in 30 trials (the dashed curve) is reasonably accurate. However, the average accuracy over the 30 trials (the solid curve) is significantly less. We believe that this illustrates a rough measure of how discriminating the posterior is between configurations. That is, although the posteriors assigned by the model are correct in terms of relative magnitude, assigning lower values to inferior matches, their standard deviation is small making it more difficult for the combinatorial optimization to locate a minima in the posterior energy close to the global minima. Figure 6 shows the same curves except with the context of temporal coordination between units represented as part of the template. This is a powerful context for determining which vehicles are actually taking part in the plan illustrated by the template. Consequently, we see that the average accuracy shown in Figure 6 shows a significant improvement over that in Figure 5. We believe that in this case the posteriors were more discriminating between ambiguous configurations reducing the computational effort necessary to find a good labeling.




Figure 6: Temporal context included-Graph showing the percentage of scenario entities correctly assigned to the template vs the percentage of scenario entities that the sensors reported that were false positives. The dashed curve gives the accuracy of the highest probability match found in 30 trials. The solid curve gives the average accuracy of all of the matches found over the 30 trials.

6 Conclusion

In this paper we present a method for intent inferencing using the MAP-MRF formulation of template matching between battlefield entities and events and a doctrinal template. We feel that the results illustrate that the method is robust to sensor errors. The empirical results show that the template matching algorithm performed well in terms of dealing with sensor false positives and false negatives. We believe that combined with the relative ease of model building, this makes a strong case for incorporating MRF models of context into intent threat inference systems.

References

[1] Stan Z. Li. *Markov random field modeling in image analysis*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2001.

[2] J. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. 1971.

[3] James W. Modestino and Jun Zhang. A markov random field model-based approach to image interpretation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(6):606–615, 1992.

[4] Mukesh A. Zaveri, S. N. Merchant, and Uday B. Desai. Tracking of point targets in ir image sequence using multiple model based particle filtering and mrf based data association. In *ICPR '04: Proceedings of the Pattern Recognition, 17th International Conference on (ICPR'04) Volume 4*, pages 729–732, Washington, DC, USA, 2004. IEEE Computer Society.

[5] Robin Grinton, Sean Owens, Joe Giampapa, Katia Sycara, Charles Grindle, and Mike Lewis. Terrain based information fusion and inference. In *Proceedings of the Seventh International Conference on Information Fusion*, pages 338–345, 2004.

[6] Robin Grinton, Joe Giampapa, Sean Owens, Katia Sycara, Charles Grindle, and Mike Lewis. Intent inferencing using an artificial potential field model of environmental influences. In *Proceedings of the Eighth International Conference on Information Fusion*, 2005.

[7] L. Ronnie M. Johansson and Robert Suzic. Particle filter-based information acquisition for robust plan recognition. In *Proceedings of the Eighth International Conference on Information Fusion*. IEEE, July 2005.

[8] Subrata Das, David Lawless, Brenda Ng, and Avi Pfeffer. Factored particle filtering for data fusion and situation assessment in urban environments. In *Proceedings of the Eighth International Conference on Information Fusion*. IEEE, July 2005.

[9] Rosenbluth A. Rosenbluth M. Teller A. Teller E. Metropolis, N.A. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[10] Bin Yu, Katia Sycara, Joseph Giampapa, and Sean Owens. Uncertain information fusion for force aggregation and classification in airborne sensor networks. In *AAAI Workshop on Sensor Networks*, 2004.