

Spatial Event Prediction by Combining Value Function Approximation and Case-Based Reasoning

Hua Li¹, Héctor Muñoz-Avila², Diane Bramsen¹, Chad Hogg², Rafael Alonso¹

¹ SET Corporation, 1005 N. Glebe Rd.,
Suite 400, Arlington, VA 22201
{hli,dbramsenralonso,}@setcorp.com

² Department of Computer Science and Engineering, 19 Memorial Drive West,
Lehigh University, Bethlehem, PA 18015
{hem4,cmh204}@lehigh.edu

Abstract. This paper presents a new approach for spatial event prediction that combines a value function approximation algorithm and case-based reasoning predictors. Each of these predictors makes unique contributions to the overall spatial event prediction. The function value approximation prediction is particularly suitable to reasoning with geographical features such as the (x,y) coordinates of an event. The case-based prediction is particularly well suited to deal with non-geographical features such as the time of the event or income level of the population. We claim that the combination of these two predictors results in a significant improvement of the accuracy in the spatial event prediction compared to pure geographically-based predictions. We support our claim by reporting on an ablation study for the prediction of improvised explosive device (IED) attacks.

Keywords: spatial prediction, case-based prediction, function value approximation.

1 Introduction

Spatial event prediction is a problem for which the input is a series of events e_1, e_2, \dots, e_n and their location in a map [1,2,3]. These events have time stamps associated with them, in addition to the locations in the map where they occur and some additional information (e.g., type of event). Based on these locations, regions or influence zones are found. Within an influence zone, cells may have different *influence values*, which are weights associated with cells reflecting a prediction about the potential locations of future events.

Figure 1 presents an example of an influence map generated by the PITS++ system, our function value estimation predictor, for improvised explosive device (IED) attacks in an urban location. PITS++ uses a function value approximation mechanism to update the influence values each time a new IED event is entered into the system. IED attacks are a type of attack where groups of insurgents place an explosive device that

is triggered to explode when a target moves close by. These kinds of attacks have become very common in Iraq and elsewhere and are frequently discussed by the news media. The colors are not visible in black and white printout but, basically, we use cyan colored numbers to indicate the locations of IED attacks and the colored-areas indicate the likelihood of attacks. Each cell is colored white (zero likelihood), green (very unlikely), yellow (somewhat likely), orange (likely), or red (very likely).

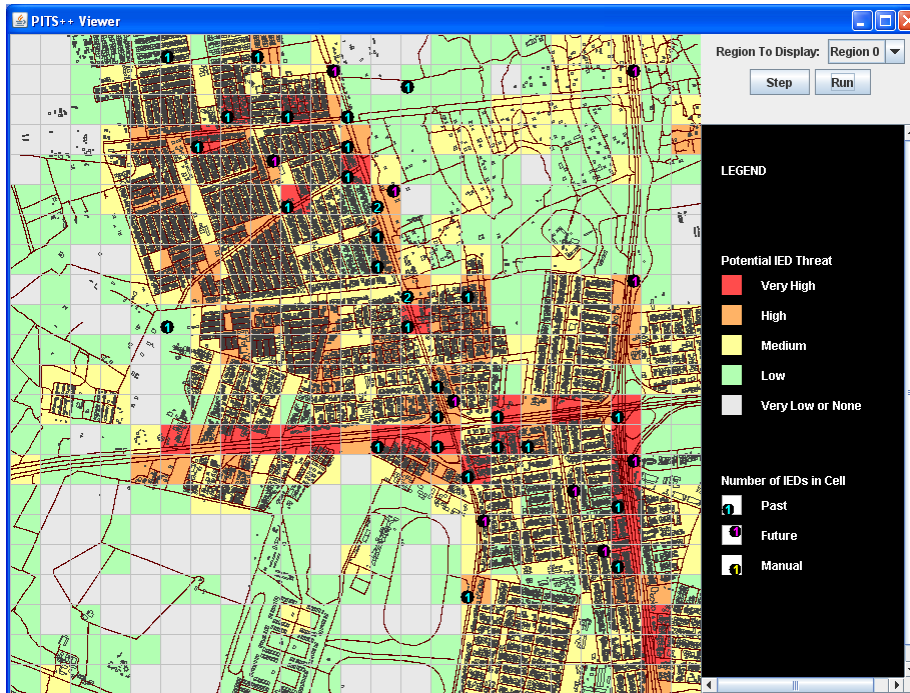


Figure 1: PITS++ viewer showing the training (cyan) and test (magenta) events

The PITS++ system is based on its predecessor PITS system [4], which computes these influence maps based on purely geographical features such as the (x,y) location of the map. Our goal is to enhance the influence map by adding non-geographical features such as time or income level in the area of the attacks. To accomplish this we added case-based reasoning capabilities to PITS++ to directly modify or “retouch” the influence values to take into account the contribution of the non-geographical features. The case-based reasoning module, DINCAT (Domain Independent Case-base Assistant) stores copies of the events originally used to train the PITS++ system but also annotated with the non-geographical features associated with the event such as the time of the attack. Then, for each cell in the influence map it retrieves all cases whose similarity to the features of the case is greater than a certain threshold. These cases are then used to retouch the influence values by taking into account the following factors:

- The time stamps from the retrieved cases

- The similarities of the retrieved cases
- The number of cases retrieved

In this paper we will discuss how these factors were combined into a retouching formula and show in an ablation study on synthetic data that, CBR substantially improves the accuracy of the prediction of the PITS++ system. To the best of our knowledge, this is the first time that case-based reasoning approaches have been combined with function value estimation predictors for the task of spatial event prediction. Our results demonstrate the significant impact that CBR can have for this task.

The paper continues as follows: the next section describes the PITS++ system function value estimation predictor; Section 3, the main section of the paper, describes in detail the CBR techniques used to enhance the spatial prediction process; Section 4 discusses the results of our empirical evaluation; Section 5 discusses related work; finally, we make concluding remarks.

2 PITS++: Function Value Estimation Prediction

SET Corporation's PITS++ tool dynamically assesses the potential of IED threat, i.e., the likelihood of insurgents emplacing IEDs in a geographic area, by making an estimation of the prediction function value based on a collection of input events.

The PITS++ tool is built on our previous work on PITS [4]. Note that the main difference between the two systems is that PITS++ incorporates a CBR mechanism into the original PITS system. Figure 2 provides an overview of the original PITS system function value estimation mechanism [4]. The inputs to PITS include terrain, IED events, and a history of friendly (blue) and opponent (red) force activity. In PITS the region of interest is a rectangle bounded by geographic coordinates and divided into cells of configurable dimension. PITS extracts IED-relevant features from an input message stream and populates each cell with the terrain and history data relevant to that cell. PITS computes over these IED relevant features to determine the influence value, which we call the PIT value, for each cell. These features (e.g., intersections and corners) are systematically determined using behavioral heuristics as well as knowledge from subject matter experts (SMEs). Each feature has a weight associated with it that indicates the opponent's preference for a feature in the context of IED emplacement activities. The feature weights are dynamically adapted with the latest IED events using function value estimation algorithms [5,6]. The cells are grouped into IED influence regions based on a cell's location and PIT value.

Prediction of IED emplacements is captured by the IED map, which is a grouping of all the IED attractiveness regions in the terrain at a given point in time. Each cell in the grid is thermally colored according to its potential IED threat level. In Figure 1, past, future (during evaluation phases), and manually input (current) IED events are indicated by numbers displayed in the lower left, upper right, and lower right corners

of the cells, respectively. As a temporally ordered list of events are entered into the system, the corresponding PIT values are adjusted based on a scalar function on the preferences elicited so far and the features. In Figure 2, The Feature Map lists all the cells that contain non-zero values for each feature. The BattlefieldAOP class is a representation of the area of interest as a grid of cells. It is responsible for populating each cell with the terrain and history information relevant to the cell.

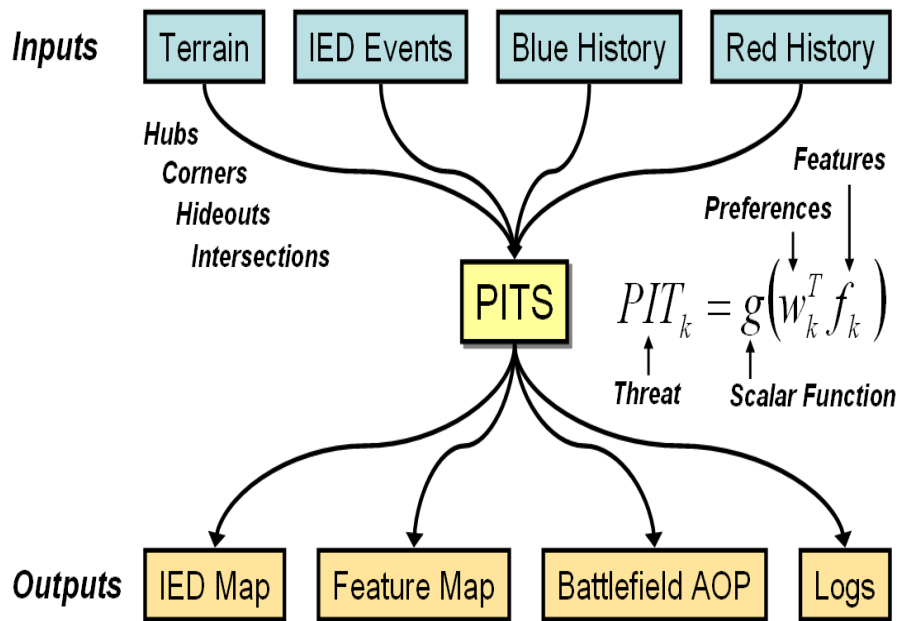


Figure 2: Function value estimation mechanism of the Potential IED Threat System (PITS).

The following are the geographical features computed in PITS (their values are normalized so they are always between 0 and 1). These features were obtained from interviews with subject matter experts:

- **Roads:** We calculate this value by summing up the number of roads in the cell.
- **Corners/Intersections:** Because of the data, we do not distinguish corners from intersections. To calculate this value, we simply look to see if there is at least one corner or intersection in the cell. If so, the cell gets a value of .5

for this feature¹. If the cell does not contain a corner or intersection, then the value is 0. Multiple corners/intersections have no additional impact on the feature value.

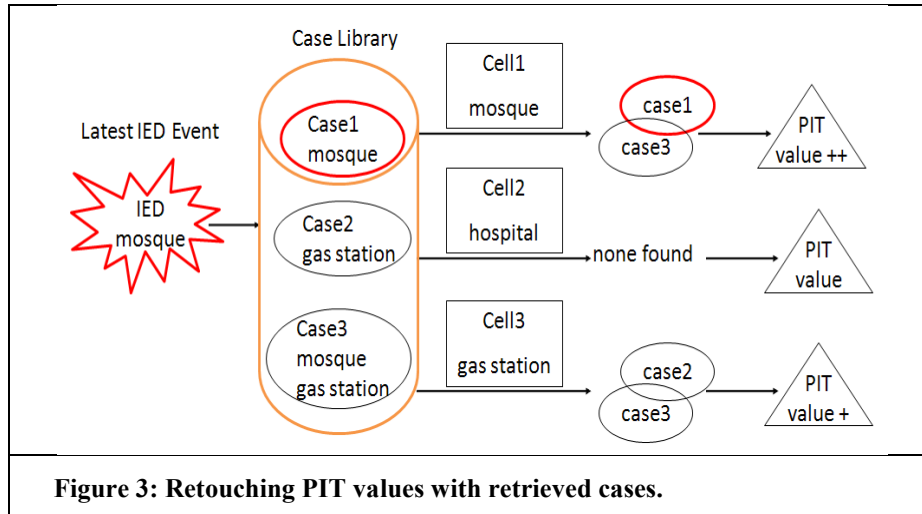


Figure 3: Retouching PIT values with retrieved cases.

- **Buildings:** This feature is meant to identify dense areas of the city, so we are looking to see if the cell contains at least 5 buildings. If so, we give it a value of 1 for this feature, and 0 otherwise.
- **Prior IEDs:** If an IED has gone off in the cell, the cell will have a value of 1 for this feature, and 0 otherwise.

3 Integrated Prediction with CBR

The basic premise is to update the PIT value (or influence values) by a “retouching” process based on the cases stored in the case base. Retouching works as follows. Suppose the latest IED attack occurred near a mosque. It will be saved in the case library after being processed by the CBR module. At the time of prediction, for each cell in the battlefield grid IED map, a query case will be created using all features associated with this cell. The query case will be dispatched to the CBR module, which will retrieve a list of similar cases. In Figure 3, the case library contains three cases where case1 is more recent than case2 and case3. Cell1 gets two similar cases (case1 and case3) because they all share the fact that they are near a mosque. Cell3 gets two similar cases (case2 and case3) because they are all linked to a gas station. Cell2, on the other hand, failed to retrieve any similar cases because none of the cases in the

¹ For this feature the values are either 0 or 0.5. We did not assign a max value of 1 when corners or intersections were present because this feature was deemed less significant as those with max value of 1, e.g. Roads.

library is related to a hospital. Both cell1 and cell3 will have their PIT value bumped up because they found similar cases whereas cell2 will not. In addition, cell1 will have a larger increase than cell2 because the former contains a more recent case case1.

For the purposes of using CBR in the context of IED attack prediction, cases represent IED events. Formally, we define a case to be:

$$\text{Case} = (\text{feature}_1, \dots, \text{feature}_n), \quad (1)$$

where feature_i includes both geographical features and non-geographical features. Geographical features, such as if the cell contains a major road intersection, are represented in the original PITS system. Non-geographical features are divided into human terrain (e.g., religion) and attack specific features (e.g., the type of explosive used) [7]. Note that non-geographical features are not represented in the original PITS. So a case can be seen as representing a possible correlation between the geographical and non-geographical features.

3.1 Retouching formula

The correlations between geographical and non-geographical features stored in the cases are used to determine how the PIT value is retouched. Specifically, the increment in PIT value is a function of the following factors:

- Date stamp of the cases. Prediction should be influenced by the date when an event took place. An event that occurred one year ago should carry less weight than a week-old event.
- Similarity of the features of the event and the retrieved cases. Prediction should be influenced by the similarity between the cell in consideration and where an event took place. Closer events should carry more weight than those farther away.
- Number of cases retrieved. The more cases are retrieved, the larger the change in the PIT value.

The old PIT value is updated by a factor of the summation of the retrieved cases, factoring in their similarity and their time stamps. We developed the following formula which commits to these three constraints:

$$PITS_{NEW} = PITS_{OLD} + PITS_{MIN,MAX} * \sum_C (SIM(C)/(SIM_{MIN,MAX} * TIME_{NOW,MIN}(C))) \quad (2)$$

Where:

- C is a variable iterating over all retrieved cases
- $PITS_{OLD}$ is the current PIT value for the cell
- $PITS_{NEW}$ is the value we are trying to compute

- $PITS_{MIN,MAX}$ is a scaling factor that determines the relative significance of the original PIT value and the cases. It is currently defined as a factor of a simple linear interpolation of the possible PIT values, $(PITS_{MAX} - PITS_{MIN})$.
- $SIM(C)$ is the similarity between the case and the PITS++ system cell whose value is being retouched
- $SIM_{MIN,MAX}$ is a factor based on the minimum similarity and maximum similarity values of the cases. We currently set it to 1.
- $TIME_{NOW,MIN}(C)$ is a factor based on how close is the case's time stamp, $TIME(C)$, to the date when the retouch is done (NOW) and the earliest date (MIN) for which we consider data useful. The closer the time stamp of C to NOW, the smaller the value of $TIME_{NOW,MIN}(C)$, which in turn makes the fraction larger. Conversely, the closer it is to MIN the larger the value of $TIME_{NOW,MIN}(C)$, which in turn makes the fraction smaller. It is currently defined as a simple linear interpolation: $(TIME(C) - MIN) / (NOW - MIN)$

3.2 Similarity metric

The similarity metric in DINCAT aggregates local similarities. The local similarities measure how close are two values of the same feature. For example, if a feature represents the (x,y) location in a map, the similarity between two locations can be defined as a function of the inverse of the distance between the two locations. Local similarity $sim_i()$ for a feature is defined such that it returns a value between 0 (non similar) and 1 (most similar). We define three forms of local similarities depending on the type of feature:

- **Symbolic.** For symbolic features we assign 1 if they are the same and 0 if they are different.
- **Numeric.** For numeric features we assume that the minimum (min) and maximum (max) values are given, and we define the similarity between two values X and Y as the inverse of the ratio of the distance between them and the largest possible distance: $1 - (|X - Y| / (max - min))$.
- **Date.** For date values, we convert them into absolute times measured in hours relative to a fixed date in time. We assign min and max to be the absolute time for the range dates for the events and use the same formula as with the numeric features.

With these local similarities we compare two vectors of features $\langle X \rangle$ and $\langle Y \rangle$ by computing the aggregated similarity metric of the local similarities, $SIM_{GLOBAL}()$, defined as:

$$SIM_{GLOBAL}(X_{1..n}, Y_{1..n}) = \alpha_1 sim_1(X_1, Y_1) + \dots + \alpha_n sim_n(X_n, Y_n) , \quad (3)$$

where the values of the vector weights, $\alpha_1 + \dots + \alpha_n$, sum to 1. As a result, $SIM_{GLOBAL}()$ also returns a value between 0 and 1 (1 been most similar). For our current implementation we set each α_i to $1/n$.

Table 1: Non-geographical features currently implemented.

Name	Type	Parameters	Description
TimeIEDAttack	numeric	0, 24	The time of the IED attack
DateIEDAttack	date	MM/DD/YYYY	The date of the IED attack
DeliveryMode	symbolic	boat-borne, animal-borne, collar-bombs, suicide-bombers, platter-chargers, explosively-formed-penetrators, improvised-rocket-assisted-munitions	Classification of the IED by the delivery mechanism as per JCS Pub 1-02
Target	symbolic	US-Armored-Vehicle, Iraqi-Police-Vehicle, US-Contractor-Vehicle, US-Foot-Patrol, Iraqi-Foot-Patrol, Civilian-vehicle, Civilian-foot	Type of target of the IED attack
TriggerMechanism	symbolic	infrared-light-beam, radio-signal, hard-wire, contact	Trigger mechanism used in the IED
Academic	symbolic	academic, non-academic	Indicates if the area of the IED attack is close to an university
Income	symbolic	1, 2, ..., 9	Indicates the income level in the area of the IED attack (1 is lowest; 9 is highest)
Tribe	symbolic	AlDulaim, BaniTamim, Shamar, AlJanabi, Aljubour, Alazza	Tribe in the area of the IED attack. There are more than 100 tribes in Iraq. Current values reflect the fact that in any one area only a few tribes are present.
Religion	symbolic	Shia, Sunni, Christian	Predominant religion in the area of the IED attack. There are more than 10 religions in Iraq. Current values reflect the fact that in any one area only a few religions are present.

3.3 Non-geographical features

A non-geographical feature may take a date, a numeric value (between a minimum value and a maximum value) or a symbolic value (from a predefined set). For each feature (see Table 1) we identify the following elements:

- **Name:** indicates the name of the feature
- **Type:** indicates the type of the feature; this can be symbolic, numeric, or date.
- **Parameters:** for numeric features this will indicate the minimum and maximum value and for symbolic features this will indicate the set of possible values.
- **Description:** a description of the feature

4 Empirical Evaluation

The purpose of the experiment is to evaluate the contributions, if any, of the CBR approach to the event spatial prediction made by PITS. For this purpose we performed an ablation study where we compared the results of PITS versus PITS++ (PITS augmented with the CBR retouch mechanism) on the same data. As mentioned before, cases can be seen as storing information of previous events co-relating geographical and non-geographical features. Therefore, it is conceivable that using our CBR retouching approach will result in improvements in the prediction when, as a whole, correlations exist between the geographical and non-geographical features. However, it might be detrimental to use the CBR approach when no such co-relations exist. Therefore, we created 3 data sets to observe the performance in 3 scenarios:

- **Correlated.** The data set consisted of 48 events in which 3 kinds of correlations exist between the geographical and non-geographical features. Every event in the data set commits to one of these correlations.
- **Partially correlated.** The data set consisted of 44 events in which 5 kinds of correlations exist between the geographical and non-geographical features. Roughly $\frac{1}{4}$ of the events have no correlation whatsoever.
- **Chaotic.** The data set consists of 42 events. No co-relations exist between the geographical and non geographical features.

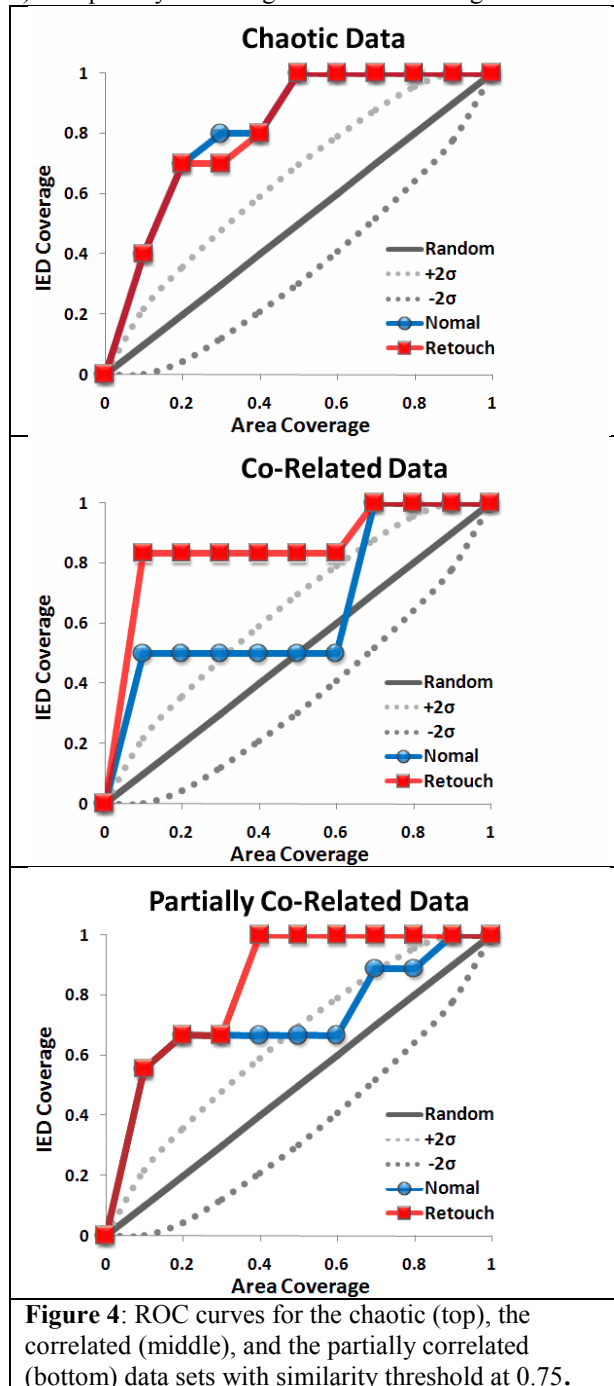
Each data set was divided into a training set consisting of $\frac{3}{4}$ of the data and a testing set consisting of the remaining $\frac{1}{4}$ of the data. The retrieval threshold was set to 0.75.

The pseudo-ROC (Receiver Operating Characteristics) curve was used to evaluate the PITS++ system's performance as an IED event predictor (Figure 4). It is defined as the IED coverage, which is the percentage of future IED events covered by threat regions, plotted as a function of the area coverage, which is the percentage of the

playbox (i.e., area of prediction) occupied by threat regions. The threat regions are determined by the PITS++ IED threat values. An area coverage of 10% looks at the top 10% of cells that have the highest IED attractiveness value. In general, the more rapidly the pseudo-ROC curve rises, the better the predictor.

A random predictor has a curve close to the diagonal, which is drawn in the plots as a solid black line. We also plot the curves that mark 2-standard deviations from the random predictor as dotted lines around the diagonal (Figure 4). For a predictor to be statistically better than the random predictor, its curve should be above the upper dotted line.

Figure 4 shows the resulting ROC curves for the chaotic, partially correlated, and correlated data sets. For each of these, two figures are drawn: one indicating the results with retouch and for PITS without any retouch. In addition we present the curve $Y = X$, which indicates a random prediction. The X-axis denotes the area coverage for the prediction. So 0.1 represents a prediction within 10% of the map whereas 1 represents all area. The Y-axis denotes the accuracy of the



prediction. So, for example, the random predictor achieves 100% prediction only when 100% of the area is covered. Therefore, the results when comparing the two curves are particularly interesting for lower values of X or at the very least to the first X-point where Y=1 is achieved. Overall a good comparison of the performance is the area under their curves. We compare the ratio: $\text{area(CBR)}/\text{area(non-CBR)}$.

The results are as follows: for the correlated dataset the curve for the CBR is always above the non-CBR approach until $X = 0.7$. The ratio of the CBR over the non-CBR is 77.7% if we consider only the areas until both curves reaches $Y = 1$ (i.e., $X = 0.7$). The ratio reduces to 66.6% if we count all area between $X = 0$ and $X = 1$. For the partially correlated, the performance is the same for both CBR and non-CBR until $X = 0.3$. Between 0.3 and 0.9 the CBR performance improves over the non-CBR and they are tied again between 0.9 and 1. The ratio of the CBR over the non-CBR is 86.2% and it augments to 82% if we consider only the areas until both curves reaches $Y = 1$ (i.e., $X = 0.9$). Although it is noteworthy that with CBR it reaches 1 at $X = 0.4$ whereas non-CBR reaches 1 only at $X = 0.9$. Finally for chaotic both curves are the same and only in the interval $[0.2, 0.4]$ the non-CBR does slightly better. However, the ratio of the non-CBR over the CBR is only 98.8% and it augments to only 97.2% if we consider only the areas until both curves reaches $Y = 1$ (i.e., $X = 0.5$). We made only one run with each data set because the system behaves deterministically: with the same input sequence of events, PITS and the CBR retouching algorithm produces the same values. The random predictor may produce multiple values, but on average it produces the $y=x$ curve shown in our analysis.

We ran student t-tests on the results obtained from the experiments comparing the data points for the normal versus the retouched results. The difference in score for the chaotic data set is not significant (t-test score: 87%), for the correlated data set is significant (t-test score: 99.7%), and for the partially correlated data set is significant (t-test score: 98.6%). In conclusion, if there is a co-relation between the geographical and the non-geographical features, -even if only partial- the CBR retouch will improve the performance of the prediction. If there is no co-relation the CBR retouch would have a negligible negative impact on the prediction.

5 Related Work

Related approaches can be divided into three kinds: time prediction from time series, prediction from influence maps, and case-based prediction. We briefly discuss each insofar as to contrast with our approach. Each of these has been the subject of extensive and well-established research. The problem of prediction from time series in its most simple form can be defined as to obtain a time range prediction $[t, t']$ for an event from a history of event in time t_1, t_2, \dots, t_n [8]. This is a well-founded field. Some methods assume an implicit model for time series while others assume an explicit model. The problem of IED time prediction is dependent on geographical,

human terrain, and attack-specific features, for which, to the best of our knowledge, no time prediction model exists capable of incorporating all of these kinds of features.

Influence maps is a method for spatial analysis which receives as input a series of events e_1, e_2, \dots, e_n and their location in a map [1,2]. These events do not necessarily have time stamps; just the locations in the map where they occur and some additional information (e.g., type of event). Based on these locations, regions or influence zones are found. Within an influence zone, cells may have different weights reflecting more or less influence from the events in the cell. These weights are typically represented with colors for visualization purposes. In influence maps the geographical location (e.g., the “x,y” coordinates) play a significant role in how the regions are determined. Whereas indeed the graphical locations are important, our work aims to find common geographical features between attacks rather than just the “x,y” location. Moreover, we want to incorporate non-geographical features (i.e., human terrain and attack-specific) into the process of determining these regions. Again, to the best of our knowledge no work exists accomplishing this.

Case-based prediction refers to the use of case-based reasoning (CBR) as the prediction technique [9,10,11,12,13,14]. Predictions in CBR include time prediction as well as class prediction (e.g., predict the kind of object based on partial observations), and strategy prediction (e.g., predicting the next movement from an opponent) among others. The difficulty of using CBR for IED prediction is that events might be related at different and contrasting levels (e.g., they might be geographically close but rather different from the perspective of the human terrain). Therefore, instead of tackling the whole problem with CBR we aim at using CBR for making predictions with the non-geographical features and combine this prediction with the one from the PITS system which models the geographical features. Analogously, adding all features to the PITS model, which is in essence the approach taken in Liu and Brown [3], would introduce the curse of dimensionality. This is a well documented limitation of value function approximation algorithms [15].

Case-based reasoning has been combined with value function approximation algorithms such as reinforcement learning [16,17,18,19] and neural networks [20,21]. The particular value function approximation algorithm developed in PITS++ has been shown to be particularly useful for spatial event prediction and, therefore, a suitable base line to measure performance gains by using case-based reasoning techniques.

6 Conclusions

One major aspect of defeating the use of IEDs is defeating the device itself. There has been intensive effort devoted to combining sensor data of various types to identify those who planted the IEDs. There are several limitations in those approaches. First, it typically tries to identify the emplacers after the attack. In other words, the current work is post-mortem and reactive in nature rather than predictive and proactive. Second, this work has been heavily relying on a human expert to perform the

historical pattern analysis and recognition. There is an outcry for predictive algorithms that operate in an autonomous or semi-autonomous manner. Third, this work has limited ability to fuse all potential data sources (geospatial and temporal event information, social and cultural data, coalition traffic patterns, multi-spectral sources of sensor data, etc). Existing methods could be used by adding all necessary features into a given value function approximation algorithm. However, this will incur the curse of dimensionality.

We presented an alternative approach in PITS++ for spatial event prediction that combines function value estimation and a case-based prediction. PITS, the function value estimation predictor, contributes with its capability to reason with geographical features such as the (x,y) coordinates of an event. We enhance this capability by using case-based reasoning techniques to model non-geographical features such as the time of the event or trigger mechanism of the IED. Cases capture, in essence, instances of correlations between geographical and non-geographical features. This correlation is exploited by our system to update the geographical prediction of PITS. We observed that this update results in a significant improvement of the accuracy in the empirical prediction compared to pure geographically-based predictions when there is some correlation between the geographical and the non-geographical features in the input event traces. In the worst scenario, when no such a co-relation exists, CBR does not help but it is not detrimental either. We support these claims with an experiment on an ablation study for the prediction of improvised explosive device (IED) attacks.

There are several directions that we want to explore in the future. First, there was little tuning in the CBR component. In particular, each local similarity metric was assigned the same weight. It is conceivable that further improvements in the prediction performance can be achieved if weights are tuned by, for example, performing statistical analysis of the data. Second, the integration between the value function approximation and the CBR prediction is one way; the CBR retouching process does not permanently change the PIT values. Instead, it is currently designed to perform the CBR retouch on the PIT values every time a new prediction is needed. We need to investigate what would be the implication of keeping the retouched PIT values; particularly what does this mean for the value function approximation process. Third, we would like to perform an evaluation with real data. For the current evaluation we used simulated data because of the unavailability of adequate real data at the time of the study. Although public data exists about IED attacks, this contains mostly geographical features. Fourth, all features currently used in the study were manually given based on interviews with Subject Matter Experts. A potential research direction is to learn such features by extracting them from raw IED data (e.g., from a repository of field reports). Fifth, the values of parameters in the retouch formula (Formula 2) are currently set manually. A potential research direction is to use Bayesian or other methods to learn these parameters. Sixth, we would like to investigate time prediction of IED events, building on existing work for explicit representations and reasoning methods for temporal events.

Acknowledgements

This research was supported by grants from the Air Force Research Laboratory and the National Science Foundation Grant No. NSF 0642882. The opinions stated in this paper are those of the authors and not of the funding agencies supporting this work.

References

1. Sweetser, P.: Strategic Decision-Making with Neural Networks and Influence Maps. AI Programming Wisdom 2, Charles River Media (2004)
2. Tozour, P.: Influence Mapping. Game Programming Gems 2, Charles River Media Preece, Eds. Lecture Notes In Computer Science, vol. 2416. Springer-Verlag, London, 174-188 (2001)
3. Liu, H., Brown, D.E.: Spatial-Temporal Event Prediction: A New Model, In: Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, October 1998, San Diego, California (1998)
4. Li, H., Bramsen, D., Alonso, R.: Potential IED Threat System (PITS). In: Proceedings of 2009 IEEE International Conference on Technologies for Homeland Security (HST 2009). IEEE Xplore. (2009) (To appear)
5. Alonso, R., Bloom J.A., Li, H., Basu C.: An Adaptive Nearest Neighbor Search for a Parts Acquisition ePortal, In: Proceedings of the Ninth ACM International Conference on Knowledge Discovery and Data Mining, SIGKDD'03 (2003)
6. Alonso, R., Li, H.: Model-Guided Information Discovery for Intelligence Analysis. In: Proceedings of CIKM '05, Bremen, Germany (2005)
7. Kipp, J., Grau, L., Prinslow, K., Smith, D.: The Human Terrain System: A CORDS for the 21st Century. Military Review. Sept-Oct (2006)
8. Hamilton, J.D.: Time Series Analysis. Princeton University Press, Princeton, NJ (1994)
9. Faltings, B.: Probabilistic Indexing for Case-Based Prediction. In: Proceedings of the 2nd International Conference on Case-Based Reasoning, pages 611–622. Springer-Verlag (1997)
10. Hansen, B.K.: Weather prediction using case-based reasoning and fuzzy set theory. M.S. thesis, Dept. of Computer Science, Technical University of Nova Scotia (2002)

11. Jære, M. D., Aamodt, A., Skalle, P.: Representing Temporal Knowledge for Case-Based Prediction. In: Proceedings of the 6th European Conference on Case-Based Reasoning (ECCBR-02). Springer (2002)
12. Zehraoui, F., Kanawati, R., Salotti, S.: Case base maintenance for improving prediction quality. In: Proceedings of the 5th International Conference on case-based reasoning (ICCR-2003). Springer (2003)
13. Redmond, M., Line, C.B.: Empirical Analysis of Case-Based Reasoning and Other Prediction Methods in a Social Science Domain: Repeat Criminal Victimization. In: Proceedings of the 5th International Conference on case-based reasoning (ICCR-2003). Springer (2003)
14. Weber, R. Evanco, W., Waller, M., Verner, J.: Identifying critical factors in case-based prediction. In: Proceedings of the 18th FLAIRS Conference. AAAI Press (2004)
15. Bellman, R.E.: Adaptive Control Processes. Princeton University Press, Princeton, NJ (1961)
16. Bridge, D.: The virtue of reward: Performance, reinforcement and discovery in case-based reasoning. Invited Talk at the 6th International Conference on Case-Based Reasoning (ICCBR-05) (2005)
17. Gabel, T., Riedmiller, M.: Multi-agent Case-Based Reasoning for Cooperative Reinforcement Learners. Proceedings of the 8th European Conference on Case-Based Reasoning (ECCBR-02). Springer (2006)
18. Sharma, M., Holmes, M., Santamara, J.C., Irani, A., Jr., C.L.I., Ram, A.: Transfer learning in real-time strategy games using hybrid CBR/RL. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07). AAAI Press (2007)
19. Auslander, B., Lee-Urban, S., Hogg, C., Muñoz-Avila, H.: Recognizing The Enemy: Combining Reinforcement Learning with Strategy Selection using Case-Based Reasoning. In: Proceedings of the 9th European Conference on Case-Based Reasoning (ECCBR-08). Springer (2008)
20. Chen, D., Burrell, P.: Case-based reasoning system and artificial neural networks: A review. Neural Computing and Applications (2001)
21. Fdez-Riverola, F., Corchado, J.M., Torres, J.M.: An Automated Hybrid CBR System for Forecasting. In: Proceedings of the 6th European Conference on Case-Based Reasoning (ECCBR-02). Springer (2002)