

Discriminative Random Field Approach to Spatial Outbreak Detection

Kaustav Das, Maheshkumar R. Sabhnani, Eric Xing

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

OBJECTIVE

We propose discriminative random field approach for detecting a disease outbreak. Given observed data on a spatial grid, the goal is to label each node as being under attack and non-attack.

BACKGROUND

Spatial scan finds the most anomalous region that has shown increase in observed counts when compared to the expected baseline. As there can be infinitely many regions to search for, most state-of-the-art algorithms assumes a specific shape of the attack region (circles for Kulldorff and rectangles for Ultra-Fast Spatial Scan Statistics). This assumption might reduce the detection power as real world attacks don't follow standard geometric shapes.

METHODS

We model the entire space as Discriminative Random Field [1] (DRF) where each elemental area is represented by a node. Two nodes in the DRF are connected if they share a common boundary. The observed incidence count at each node is dependent on whether the node is under attack or not. We model it as a normal (or Poisson) distribution with different mean and variance parameters for non-attack and attack conditions.

Given, the observed data y , we would estimate the labels $\{-1, +1\} \equiv \{\text{no attack, attack}\}$ at each node x in the graph. Using the Hammersley Clifford theorem and assuming only up to pair wise clique potentials to be nonzero, the joint distribution over the labels x given the observations y can be written as [1],

$$P(x | y) = \frac{1}{Z} \exp \left(\sum_{i \in S} A_i(x_i, y) + \sum_{i \in S} \sum_{j \in N_i} I_{ij}(x_i, x_j) \right)$$

where, Z is the partition function.

$A(x_i, y)$, the association potential is modeled using a local discriminative model that outputs the association of the site i with class x_i . Using the logistic function as the *link*, the local class posterior can be modeled as,

$$P(x_i = 1 | y) = \frac{1}{1 + \exp(-(\theta_{i0} + \theta_{i1}y))} = \sigma(\theta_{i0} + \theta_{i1}y)$$
$$A(x_i, y) = \log(\sigma(x_i(\theta_{i0} + \theta_{i1}y)))$$

Both $(\theta_{i0}, \theta_{i1})$ will be dependent on the baseline count and variance of the observed node. For $I_{ij}(x_i, x_j)$, the interaction potential, we use a homogeneous Ising model $I = \beta x_i x_j$, which penalizes every dissimilar

pair labels by the cost β . Here β acts as a smoothing parameter in the DRF where larger values encourage more smooth solutions.

To learn the parameters $\lambda = \langle \theta, \beta \rangle$ from the data we alternately calculate the most probable class x^* and the parameter values:

$$x^* = \arg \max_{x \in X} p(x | y, \theta, \beta)$$
$$\theta^*, \beta^* = \arg \max_{\theta, \beta} p(x^* | y, \theta, \beta)$$

The first step is an inference on the DRF, which can be approximated by ICM (Iterative Conditional Mode) algorithm [2]. The second step is a supervised learning of parameters given the labels. We use pseudo-likelihood approximation to approximate the partition function Z .

RESULTS

On synthetic data, our method was able to detect the location and shape of injected outbreaks, with increases as low as 3% of the variance of baseline counts. The running time of the algorithm is linear in the number of nodes on the grid. We achieved 50 – 1000 times speedup over the spatial scan algorithm proposed in [4] with similar detection power. We also ran the algorithm on semi-synthetic data, BARD [3]: simulated outbreaks injected into streams of real emergency department data, with comparable detection times to [4] and much faster execution times.

CONCLUSIONS

This method provides the capability to detect non-uniform irregularly shaped attack regions, and gives a very fast method to detect attack clusters. The proposed method can find clusters in grid sizes of 4 million nodes; hence it is scalable to very high resolution grids.

This work was supported by NSF grant IIS-0325581.

REFERENCES

- [1] Kumar S. & Hebert M. (2003) Discriminative Random Fields: A Discriminative Framework for Contextual Interaction in Classification. IEEE Conference on Computer Vision (ICCV).
- [2] J. Besag. On the statistical analysis of dirty pictures. Journal of Royal Statistical Soc., B-48:259-302, 1986
- [3] Neill DB, Moore AW, Sabhnani M, and Daniel K. Detection of emerging space-time clusters. Proceedings of the 11th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2005, 218-227.
- [4] Neill DB, Moore AW (). Rapid detection of significant spatial clusters. *Proceedings of the 10th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2004, 256-265